



LUND UNIVERSITY

Multivariate Modelling of Energy Markets

Lunina, Veronika

2016

Document Version:

Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (APA):

Lunina, V. (2016). *Multivariate Modelling of Energy Markets* (Lund Economic Series no. 198 ed.). [Doctoral Thesis (compilation), Department of Economics]. Department of Economics, Lund University.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Multivariate Modelling of Energy Markets

Veronika Lunina



LUND
UNIVERSITY

DOCTORAL DISSERTATION

by due permission of the School of Economics and Management,
Lund University, Sweden.

To be defended at Holger Crafoords EC3:210
on February 2, 2017 at 10:15.

Faculty opponent

Professor Fred Espen Benth,
Department of Mathematics, University of Oslo

Organization Lund University Department of Economics P.O. Box 7082 S-220 07 Lund, Sweden	Document name Doctoral dissertation
	Date of issue
Author Veronika Lunina	Sponsoring organization
Title and subtitle Multivariate Modelling of Energy Markets	
Abstract <p>This thesis contributes to the empirical energy finance literature and consists of three research papers. The common denominator for all papers is the multivariate modelling approach, placing the electricity market at the core and delving into its interdependencies on fundamentally related markets and factors.</p> <p>In the first paper, <i>Modelling Cross-Commodity Interdependencies in Volatility: A Case Study of the German Energy Markets</i>, we focus on multivariate modelling of the return series of electrical power, natural gas, coal and carbon emission allowances in the German market. We pay special attention to selecting an appropriate volatility model allowing for cross-commodity effects, coupled with a flexible skew-Student distribution for the error terms. We discuss the relationship between the discovered volatility spillover effects and the fundamental developments in the energy markets.</p> <p>The second paper, <i>Volatility Transmission in the German Energy Markets: A Variance Impulse Response Analysis</i>, develops a comprehensive analysis of the transmission of independent shocks from the gas, coal and carbon markets to the power market, building on the estimation results from the first paper and employing the novel Volatility Impulse Response Function (VIRF) methodology. We find that spillover effects show significant time variation and are substantial in size.</p> <p>The third paper, <i>Joint Modelling of Power Price, Temperature, and Hydrological Balance with a View towards Scenario Analysis</i>, presents a model for the joint dynamics of the Nordic system spot power price together with its major demand-side factor, the outdoor temperature, and its major supply-side factor, the hydrological balance. We demonstrate how the model can be used in meteorological scenario analysis.</p>	
Keywords Energy markets, electricity, volatility spillovers, skew-Student asymmetric BEKK, volatility impulse response function, temperature, hydrobalance	
Classification system and/or index terms (if any) JEL Classification: C32, C58, G17, Q41	
Supplementary bibliographical information	Language English
ISSN and key title 0460-0029 Lund Economic Studies no. 198	ISBN 978-91-7753-056-5 (print) 978-91-7753-057-2 (pdf)
Recipient's notes	Number of pages 123
	Price Security classification

I, the undersigned, being the copyright owner of the abstract of the above-mentioned dissertation, hereby grant to all reference sources permission to publish and disseminate the abstract of the above-mentioned dissertation.

Signature _____ Date _____

Multivariate Modelling of Energy Markets

Veronika Lunina



LUND
UNIVERSITY

LUND ECONOMIC STUDIES NUMBER 198

Copyright © Veronika Lunina 2016

Distributed by the Department of Economics

Lund University

P.O. Box 7082

S-220 07 Lund

SWEDEN

ISBN 978-91-7753-056-5 (print)

ISBN 978-91-7753-057-2 (pdf)

ISSN 0460-0029

Printed in Sweden by Media-Tryck, Lund University

Lund 2016

Contents

Acknowledgement	i
Chapter 1: Introduction	1
1 Winds of change: The evolution of the European energy markets	2
2 Approaches to electricity price modelling	5
3 Summary of the thesis	7
Chapter 2: Modelling Cross-Commodity Interdependencies in Volatility: A Case Study of the German Energy Markets	15
1 Introduction	16
2 The data	19
3 Model framework	22
3.1 Conditional mean	23
3.2 Conditional covariance	23
3.3 Distributional assumptions	26
4 Estimation	29
4.1 Methodology	30
5 Results	33
5.1 Model specification	33
5.2 Parameter estimates and volatility spillovers	36
5.3 Conditional correlations and volatilities	42
6 Summary and conclusions	45
References	46
Chapter 3: Volatility Transmission in the German Energy Markets: A Variance Impulse Response Analysis	51
1 Introduction	52
2 Volatility impulse response function	55
2.1 Background	55
2.2 Definition of VIRF and its properties	57
2.3 Construction of independent shocks	58
2.3.1 An illustration	60
2.4 VIRF computation via Monte Carlo integration	62
3 Results	64

3.1	Volatility transmission patterns: The yearly perspective	65
3.2	Day-ahead variance responses: The daily perspective	68
3.3	Sector-wide news and critical news	71
4	Summary and conclusions	74
	References	75

Chapter 4: Joint Modelling of Power Price, Temperature, and Hydrological Balance with a View towards Scenario Analysis		79
1	Introduction	80
2	Data and preliminary analysis	82
2.1	The dataset	82
2.2	Data analysis	84
2.3	Relationship between power price, temperature, and hydrobalance	91
3	The model	93
3.1	General framework	93
3.1.1	Conditional mean	93
3.1.2	Conditional covariance	96
3.1.3	Distributional assumptions	98
3.2	Model identification and estimation procedure	100
4	Results	103
4.1	Model identification results	104
4.2	Estimation results	106
5	Application: Scenario analysis	113
6	Summary and conclusions	119
	References	119
	Appendix A	123

Acknowledgement

There is no better way of starting this thesis than by thanking the person without whose advice I would probably not enter the PhD programme. Thank you, Björn, for encouraging me to apply, for useful tips, and for helping me to figure out my research interests. Björn Hansson was also the one who recommended me to take courses at the Department of Mathematical Statistics, where I first heard about someone from my own department, whose name is Rikard, and who is doing research in mathematical finance, which might be of interest.

I properly met Rikard Green for the first time after the departmental Christmas dinner in 2012. I remember walking down the long set of stairs at Grand Hotel simultaneously with Rikard. Someone spoke to me in Swedish, and Rikard asked me if I managed to understand, and then introduced himself. I was planning to head home, but then I thought to myself “ah, this must be the Rikard” and decided to join him together with other colleagues for a beer. At that moment I was struggling to find a research topic. I remember Rikard first telling me about modelling of electricity markets. He spoke about it with such a rare, genuine passion that I was truly inspired. Later I came to Rikard and asked him if he can be my supervisor. Thank you, Rikard, for opening energy finance to me, for introducing me to people both in industry and in academia, thank you for inspiring me! Further, thank you for helping to bring two other great supervisors on board.

I was quite nervous before my first conversation with Karl Larsson about potential research topics. I knew that Karl is a very bright mathematician, and I was a little intimidated. What I did not know, but quickly understood, is that Karl is also an extremely humble, nice, and open-minded person. Thank you, Kalle, for always being helpful and considerate, for all the time you invested together with Rikard,

teaching me energy markets finance, for kindly agreeing to schedule meetings after lunch (although I know you prefer mornings), thank you for being a fun co-author to work with!

I would like to say a huge thank you to my main supervisor Birger Nilsson. Birger is one of the most talented and also easy-going people I have ever met. He is extremely rigorous, but at the same time always maintains a relaxed attitude and is not overly obsessive about minor things. Birger's contribution to this thesis is hard to overestimate. Thank you, Birger, for never exposing me to stress, for bearing with my impatience, for listening to my endless talking, for being kind and thoughtful, thank you for being my mentor and my friend!

I am grateful to Frederik Lundtofte and the Knut Wicksell Centre for Financial Studies for the generous financial support. Thank you to everyone who has ever been my teacher, and to all members of the finance group both for research advice and for many joyful memories.

Thank you, Valeriia, for being the best office mate, for supporting all of my ideas, for always saying that I am doing great no matter what, thank you for all the laughs. Thank you, Bujar, for sharing my slightly twisted sense of humour, for all the talks and fun that we had while organizing the Arne Ryde workshops. Thank you, Hans, for supporting me in your own peculiar way and being honest. Thank you, Patrik, for encouraging me to apply for internships in London, for your numerous career tips, thank you for motivating me. Thank you, Yana, for keeping me company in SingStar and being a good friend. Thank you, Dominika, Hassan, and Caglar for being ready to help when I need it.

Thank you, Tania, for being my best friend in the whole world for the past ten years, for supporting me in all of my endeavours, for always being on my side. Thank you, Alex, for being a loving and caring husband, thank you for always choosing me first, I am truly lucky to have you.

Above all, I wish to thank my parents for raising me to be the person I am today, for loving me unconditionally, for giving valuable advice and yet letting me make my own choices. Thank you!

Veronika

Lund, October 2016

CHAPTER 1

Introduction

This thesis consists of three research papers that contribute to the empirical energy finance literature. The first paper investigates the interdependencies between the electrical power, natural gas, coal, and carbon emission allowances in the German energy market, with a particular focus on volatility spillovers. Building on the results from the first paper, the second paper develops a comprehensive analysis of the transmission of independent shocks from the gas, coal, and carbon markets to the power market. The third paper explores the joint evolution of power price, outdoor temperature, and hydrological balance in the Nordic energy market. The common denominator for all papers is the multivariate modelling approach, placing the electricity market at the core and delving into its interdependencies on fundamentally related markets.

The thesis contains empirical results of interest to energy market participants and which are discussed with reference to the developments in the underlying markets. However, the thesis is as related to multivariate modelling as it is to energy markets. The first paper features an extensive model selection procedure based on six conditional covariance model specifications and seven distributional assumptions for the error terms. In the second paper, we employ the recently developed concept of a volatility impulse response function, and suggest our own method of normalizing it, which facilitates the interpretation and comparability of the results. In the third paper, we also address simulations from the suggested model and illustrate a potential application in meteorological scenario analysis. Further, in all of the papers, we deviate from the traditional normality assumption for the error terms and utilize a flexible skew-Student distribution, as Bauwens and Laurent (2002, 2005) propose. A simulation exercise in the third paper demonstrates the benefits of using a distribution that

can accommodate the non-Gaussian properties of both energy prices and meteorological data series.

The rest of the introductory chapter is organized as follows. Section 1 discusses the recent trends in the European energy markets, revealing the importance of the multivariate modelling framework. Section 2 provides an overview of the existing approaches to modelling the price of electricity, which is the core energy commodity in this thesis. Section 3 summarizes the research papers and outlines the main results.

1 Winds of change: The evolution of the European energy markets

Over the past two decades, we have witnessed the European electricity and natural gas markets transform from the paradigm of vertical integration into liberalized competitive structures. This led to the establishment of energy exchanges, such as the Nord Pool, the European Energy Exchange, the Amsterdam Power Exchange, and so on, where electrical power and other energy commodities are traded, both for physical delivery and for financial settlement. While the process of deregulation and the subsequent formation of competitive energy markets is over in most European countries, new trends have emerged, which yet again challenge the established market model and are reshaping the energy landscape.

The most profound trend is the growing reliance on renewable power generation sources, such as hydro, wind, solar, and biomass. Renewable generation accounted for almost 30% of electricity demand in Europe in 2015, up from 17% in 2008. Increases in the installed renewable capacity, offering power generation at the lowest marginal cost and subsidized by governments, undermined the competitiveness of thermal generation and eroded the credit quality of many utilities.

Many fossil-fuel power plants have become unprofitable over the recent years due to low loads. In addition to decreasing the overall mean level of power prices, renewables also significantly affected intraday profiles, leading to lower peak prices and higher volatility. Despite their major benefit of environmental friendliness, most renewable energy sources are highly intermittent in nature, which creates a lot of uncertainty in production planning. Therefore, many European countries are considering the introduction of capacity payments to incentivize thermal generators to remain online and secure the power supply when the sun is not shining or the wind is not blowing. The increased reliance on intermittent generation capacity will inevitably lead to a greater risk of supply disruptions and more frequent occurrences of extreme price movements. However, whether we like it or not, renewables are here to stay and we will see further changes in the coming years as the markets adapt to the new paradigm and learn how to make different generation technologies co-exist rather than compete.

One way to adapt is to increase the interconnections between national electricity markets, leading to more efficient power plant dispatch and supply stability. The Nordic electricity market, encompassing Denmark, Norway, Sweden, Finland, Latvia, Lithuania, and Estonia, is already highly interconnected, with a common system power price and actual area prices coinciding most of the time. Denmark, for instance, has significant wind capacity, and benefits from the opportunity to export the surplus wind power on windy days, and import power during calm periods. Norway, on the other hand, has plenty of hydrological generation capacity, which can be partially stored in water reservoirs. Due to the common market, other Nordic countries can take advantage of Norway's extra hydro capacity. There is currently potential for increasing the interconnections between Italy and the rest of the Europe, and between the United Kingdom and continen-

tal Europe, where the largest power price differentials are observed. In addition, the Nordic market, with the lowest prices, can become more connected to neighbouring markets such as Germany. Despite the fact that connection presents significant regulatory, political, and technological challenges, the European Commission is promoting the creation of a fully integrated European power network.

Apart from interconnectivity between regional markets, we also observe the growing interdependencies between types of energy. Until the past decade, there was little substitution between different energy sources. For example, oil has traditionally dominated as a transportation fuel. Nowadays, however, natural gas is increasingly considered as an alternative, and electric cars are becoming more accessible. Energy markets are less segregated, resulting in developments in one specific market (e.g., price shocks, regulatory changes, technological advances) that have far-reaching implications for other markets.

Finally, technological progress is reshaping both supply and demand in the energy markets. On the one hand, we see the wide spread of energy efficient technologies, ensuring that the growing demand can be met at a lower energy expense rate. On the supply side, power storage solutions are being developed. The exploration of unconventional sources of oil and gas (such as tight and shale) and cost-minimizing innovations in the extraction process prevent the prices of conventional oil and gas from rallying as a result of scarcity. While the exploration of shale oil and gas is more relevant for the U.S. market, it took a toll on the shift to greater power generation using coal in Europe, where the surplus U.S. coal has been exported in recent years.

To thrive in an environment of interconnected markets, energy market participants have to understand their exposure to a higher number of risk factors and potential scenarios than ever. This highlights the relevance of the multivariate modelling approach. As the interconnectivity between the energy markets strengthens, consider-

ing them in a system is becoming a necessity rather than an advantage in many practical applications, from production planning and capital allocation decisions to risk management.

2 Approaches to electricity price modelling

The literature on electricity price modelling generally follows one of two approaches: *reduced-form* modelling and *structural* modelling.

Reduced-form models are concerned with capturing essential empirical properties while keeping the underlying risk factors unidentified. These can be further categorized into continuous time models and discrete time models. Continuous time models are mathematically elegant, and in many cases offer closed-form solutions for pricing and hedging derivative assets. The classical continuous time model for the spot dynamics of commodity prices is the one-factor Schwartz model (see Schwartz, 1997). This model, based on the exponential of the Ornstein-Uhlenbeck process, was extended to a two-factor model that Lucia and Schwartz (2002) applied to electricity prices. In addition to the short-run mean-reverting component, the two-factor model contains a long-run non-stationary component. Originally, the driving noise for both components was assumed to be a Brownian motion. Further, multi-factor jump-diffusion models were developed based on a more general family of stochastic processes with stationary independent increments, called Levy processes (see, e.g., Benth et al., 2007; Meyer-Brandis and Tankov, 2007).¹ Continuous time models are well suited for analytical derivations and derivatives pricing. However, they require discretization to be estimated since the data are always observed at discrete time points. Therefore, unless the primary application of a model is derivatives pricing, it often makes sense to

¹See Benth, Šaltytė-Benth and Koekebakker (2008) for a rigorous introduction to the reduced-form stochastic modelling of electricity and related markets.

specify the dynamics in discrete time to begin with. Discrete time models aim to capture the stylized properties of electricity prices by fitting their conditional mean and conditional variance processes. The former is typically based on the autoregressive integrated moving average methodology of Box and Jenkins (1970), while the latter builds on the seminal papers by Engle (1982) and Bollerslev (1986). While being less elegant, the time-series models are easier to augment with additional features, more straightforward to estimate, and useful in many practical applications, such as value-at-risk estimations, short-term forecasting, and scenario analysis.

Whether continuous or discrete, reduced-form models strive for mathematical tractability and practical convenience, and are not concerned with identifying the fundamental sources of randomness underlying the stochastic processes. However, electricity is a physical commodity fundamentally linked via production and consumption to a set of known and observed factors. The wide availability of supply and demand data that provides insights into the movements of power prices led to the development of the structural modelling approach. Structural models are concerned with the fundamentals of supply and demand, with prices resulting from equilibrium considerations. The first structural model for spot power prices was suggested in Barlow (2002). It featured a vertical demand curve, representing price inelasticity, and a supply curve given by a non-linear function of a mean-reverting process, which captured the evolution of demand over time. This simple approach later accounted for price relationship among multiple fuels, available generation capacity, and other fundamental factors.² Structural models adapt easily to changing market conditions and appeal by their realism, while still allowing for closed-form forward prices in some cases. On the other hand, they tend

²See Carmona and Coulon (2013) for a detailed survey of structural models for power prices.

to become complicated faster than reduced-form models once additional factors are considered, and typically contain a large number of parameters.

This thesis adopts the reduced-form discrete time modelling approach. However, the multivariate setting enables us to utilize the fundamental information in the spirit of structural modelling. Reduced-form or structural, continuous time or discrete, all models, are by definition imperfect simulations of reality. There is always a trade-off between complexity and realism, between bias and efficiency, between generality and specificity. When it comes to research, there is no one-size-fits-all answer. The only universal truth is that every model has its limitations, and good researchers understand the limitations of their models. To quote Box and Draper (1987) in their discussion on empirical model building, “*Essentially, all models are wrong, but some are useful*”.

3 Summary of the thesis

The rest of the thesis is organized into three chapters that present the research papers.

In Chapter 2, *Modelling Cross-Commodity Interdependencies in Volatility: A Case Study of the German Energy Markets*, we focus on multivariate modelling of the return series on electrical power, natural gas, coal, and carbon emission allowances. Our sample comprises the nearest delivery yearly futures contracts and covers the beginning of 2008 to the beginning of 2014. In order to perform the analysis in a reliable and statistically robust way, we employ a vector autoregressive (VAR) system coupled with time-varying volatilities and correlations, as captured by a relatively general Baba-Engle-Kraft-Kroner (BEKK) specification. In addition, we make use of the flexible skew-Student distribution proposed in Bauwens and Laurent (2002, 2005), moti-

vated by the distinct non-Gaussian properties of the energy return series. Within this model framework, we perform an extensive model selection procedure. The likelihood ratio tests indicate that accounting for cross-commodity effects in second moments, as well as the asymmetric effects of positive and negative shocks on future volatility, leads to a statistically significant improvement in the likelihood value. Our preferred distributional specification features individual skewness parameters for the four commodities, and a common degrees of freedom parameter. We confirm the existence of volatility spillover effects to the power market from all other markets, with the most economically significant effect coming from the coal market. This result is consistent with coal's larger share in the German power generation mix during the sample period and its higher profitability compared to gas. Overall, coal is the driving commodity in terms of volatility in our system because we find plenty of spillover effects channelling from the coal market and none channelling to it. Analysis of the estimated conditional correlations reveals a decrease in co-movement between power and fossil fuels, which can be related to the on-going structural shift from carbon-intensive to renewable generation technologies.

In Chapter 3, *Volatility Transmission in the German Energy Markets: A Variance Impulse Response Analysis*, we investigate the transmission of volatility on a deeper level, building on the results of the previous study. We analyse the impact of large exogenous shocks in gas, coal, and carbon on the expected variance of power at different horizons using the volatility impulse response function methodology introduced in Hafner and Herwartz (2006). The results indicate that volatility spillover effects show large variations across commodities and over time. Spillovers from coal are substantial throughout the sample period; however, with significant variation on a daily basis. Spillovers from gas are generally weaker, although more persistent, and decrease during the latter part of our sample. Spillovers from

carbon show the fastest decay and are economically non-significant until around 2011. We find that positive news (price increases) in gas and coal lead to much larger responses in the expected variance of power than negative news (price decreases). Distinguishing between the effects of positive and negative news is less important for the carbon market. Finally, benchmarking the magnitude of the responses in power variance against the responses in own-market variance reveals the non-trivial size of cross-market effects.

In Chapter 4, *Joint Modelling of Power Price, Temperature, and Hydrological Balance with a View towards Scenario Analysis*, we propose a model for the joint evolution of spot power price, outdoor temperature, and hydrological balance. Temperature is a major demand-side factor affecting power prices, while hydrobalance is a major supply-side factor in energy markets with a dominant share of hydrological power generation, such as the Nordic market. Our time series modelling approach coupled with the skew-Student distribution allows for interrelations in both mean and volatility, and accommodates most of the discovered empirical features, such as periodic patterns and long memory. We find that in the Nordic market, the relationship between temperature and power price is driven by the demand for heating, while the cooling effect during summer months does not exist, likely due to mild climate conditions. Hydrobalance, on the other hand, negatively affects power prices throughout the year, since in dry hydrological conditions a larger share of demand has to be covered by higher marginal cost generation sources. Further, we confirm the existence of volatility spillover effects from temperature and hydrobalance to power. We illustrate how the proposed model can be used to generate a variety of meteorological scenarios and analyse the implications for power prices.

References

- [1] Barlow, M., 2002. A diffusion model for electricity prices. *Mathematical Finance* 12, 287–298.
- [2] Bauwens, L., Laurent, S., 2002. A new class of multivariate skew densities, with applications to GARCH models. CORE Discussion Paper 20.
- [3] Bauwens, L., Laurent, S., 2005. A new class of multivariate skew densities, with application to Generalized Autoregressive Conditional Heteroscedasticity models. *Journal of Business and Economic Statistics* 23, 346–354.
- [4] Benth, F.E., Kallsen, J., Meyer-Brandis, T., 2007. A non-Gaussian Ornstein-Uhlenbeck process for electricity spot price modelling and derivatives pricing. *Applied Mathematical Finance* 14, 153–169.
- [5] Benth, F.E., Šaltytė-Benth, J., Koekebakker, S., 2008. Stochastic modelling of electricity and related markets. *Advanced series on statistical science and applied probability*; v. 11. World Scientific, Singapore.
- [6] Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- [7] Box, G.E.P., Draper, N.R., 1987. *Empirical model-building and response surfaces*. Wiley, New Jersey.
- [8] Box, G.E.P., Jenkins, J.M., 1970. *Time series analysis forecasting and control*. Holden-Day, San Francisco.
- [9] Carmona, R., Coulon, M., 2013. A survey of commodity markets and structural models for electricity prices. In: Benth, F.E., Kholodnyi, V., Laurence, P. (Eds.). *Quantitative energy finance: modeling, pricing, and hedging in energy and commodity mar-*

- kets. Springer-Verlag, New York, 41–83.
- [10] Engle, R.F., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987–1007.
 - [11] Hafner, C.M., Herwartz, H., 2006. Volatility impulse responses for multivariate GARCH models: An exchange rate illustration. *Journal of International Money and Finance* 25, 719–740.
 - [12] Lucia, J., Schwartz, E.S., 2002. Electricity prices and power derivatives: evidence from the Nordic power exchange. *Review of Derivatives Research* 5, 5–50.
 - [13] Meyer-Brandis, T., Tankov, P., 2007. Multi factor jump-diffusion models of electricity prices. Working paper, Universite Paris VII, France.
 - [14] Schwartz, E.S., 1997. The stochastic behaviour of commodity prices: Implications for valuation and hedging. *Journal of Finance* 52, 923–973.

CHAPTER 2

Modelling Cross-Commodity Interdependencies in Volatility: A Case Study of the German Energy Markets

Abstract

This study investigates the interdependencies in volatilities of the returns on electrical power, natural gas, coal, and carbon emission allowances in the German energy markets. These commodities are physically linked through the electricity production process, and are natural to consider as a system. We pay special attention to selecting an appropriate econometric volatility model within the VAR-BEKK framework, coupled with a flexible skew-Student distribution for the error terms. The results indicate the existence of volatility spillover effects to the power market from all other markets, with the largest in magnitude effect coming from the coal market. In addition, we observe a decreasing trend in the correlations between power and fossil fuels. We interpret our results in terms of the fundamental developments in energy markets during the sample period from 2008 to 2013, particularly the changes in spark and dark spreads and in the actual generation mix.

Keywords: energy forward markets, volatility spillovers, skew-Student asymmetric BEKK

JEL Classification: C32, C58, G1, Q41

1 Introduction

In this study, we examine a multivariate model of the return series of electrical power, natural gas, coal, and carbon (CO_2) in the German markets. These commodities are connected through the production process. The prices of input fuels, namely gas and coal, constitute the main portion of the variable costs of producing electricity. Because using gas or coal to produce electricity is associated with carbon emissions, the price of these emissions also enters the cost side of electricity generation. The spread between the power price and the generation cost defines the producer's gross margin. For more than a decade, the spark and dark spreads have stood out as dominant spreads in the European power markets because they correspond to the payoffs of standard gas- and coal-fired production units, respectively. Indeed, they drive power plant profitability, and further serve as indicators that provide incentives for agents in the energy sector to invest in future production capacity. For these reasons, it is of fundamental importance that energy companies and policymakers understand how the cost side of price spreads impacts electricity prices. Here, we focus on the German power market because it is an excellent example of a liquid and increasingly transparent energy market. Germany is Europe's largest power market, and relevant price data for power and fuels (for German production units) are reliable and publicly available from the European Energy Exchange (EEX).

The interrelations between a set of commodities can manifest in first-order moments (expected returns), and in second-order moments, affecting both volatilities and correlations. In this study, we employ a multivariate cross-commodity model for the returns on power, gas, coal, and carbon.

There is a growing body of literature on the interrelations between different energy commodities. A common approach in previous

research has been to employ methods of co-integration in order to investigate first-order interrelations between different prices. For example, Gjolberg and Johnsen (1999) study crude oil prices. Casassus, Liu, and Tang (2013) provide empirical evidence of co-integration between several petroleum-related markets. A comparable study of electricity markets is that of De Vany and Walls (1999), who test for co-integration in 11 regional power spot prices in the U.S. market.

All studies mentioned so far are concerned with first-moment interrelations. While second-order interrelations have been extensively studied in the context of equity markets, there is relatively little research related to energy markets. Lin and Tamvakis (2001), Efimova and Serletis (2014), and Karali and Ramirez (2014) examine second-order interrelations in different segments of the U.S. oil and gas markets. Interrelations between crude oil and various other commodities in the European markets are investigated in Reboredo (2014) and Liu and Chen (2013). Koenig (2011) studies the time variation in the correlations among power, fuels, and carbon in the U.K. market.

This paper presents a comprehensive study uncovering the interdependencies in volatilities between power, fossil fuels, and carbon in the German energy market. We focus on the German market for two reasons. First, it represents the largest European power market, and exhibits a growing degree of transparency and openness towards surrounding markets. Second, and perhaps more important, the German market is currently undergoing a structural transition of its energy portfolio in order to reduce its dependence on fossil fuels, moving towards a larger proportion of renewable energy sources. For this reason, Germany has attracted much attention within the international energy arena. We analyse the correlations and volatility spillover effects between power, fuels, and carbon in light of the on-going structural changes in this energy market. Since fossil-fired power plants can be regarded as real options on the spreads between

the power price and the production cost, we implicitly concentrate on aspects related to the extrinsic value of power plants in this study, which are of fundamental interest for energy companies and policy-makers.

In order to perform a reliable and statistically robust analysis, we employ a vector autoregressive (VAR) system coupled with time-varying volatilities and correlations, as captured by a general Baba-Engle-Kraft-Kroner (BEKK) specification. The VAR part of the model allows for commodity interrelations in expected returns, while the BEKK part of the model allows for interrelations and spillover effects in volatilities. We estimate the model with six different covariance specifications, each under seven different distributional assumptions, by making use of the flexible skew-Student distribution proposed in Bauwens and Laurent (2002, 2005). The flexible distributional assumption allows each return series to have individual statistical properties. Within this model framework, we perform an extensive analysis of the model specification, with particular focus on the conditional second moments and distributional assumptions. According to the likelihood tests, the preferred covariance matrix specification is the most general, allowing for asymmetric effects in volatility and spillovers across a number of different channels. The tests rejected models that do not allow for spillover effects. The preferred distributional specification, however, is not the most general. While excess kurtosis is an important feature, we find that allowing for individual tail properties does not provide a statistically significant improvement in the likelihood value.

We continue the paper with an economic analysis of the estimation results from the best model based on our specification tests. We find statistically significant spillover effects to the power market from all other markets, with the spillover from coal standing out in magnitude. Coal appears to be the driving commodity in our system, as

we report plenty of spillover effects channelling from the coal market and none channelling to it. These results are consistent with developments in the underlying markets. For example, we observe that gas plays a less important role in the generation mix compared to coal, and that spark spreads become negative toward the second half of our sample period. In contrast, coal remains in-the-money, with dark spreads staying positive throughout the sample period. Further, inspecting the estimated conditional correlations reveals the weakening link between power and fossil fuels. We argue that with the growing share of renewables in the German power generation mix, there is less co-movement between the prices of power and thermal generation sources.

The remainder of this chapter is organized in five further sections. Our data are presented in Section 2. Sections 3 and 4 describe the model framework and the estimation procedures, respectively. The estimation results and analysis are discussed in Section 5, while Section 6 contains a summary and concluding remarks.

2 The data

Our data set comprises the daily closing prices of the following futures contracts:

1. Gas TTF¹, traded on the APX-ENDEX exchange in EUR/MWh.
2. German base load power, traded on the EEX in EUR/MWh.
Base load profile refers to the delivery of power as a constant flow during the delivery period.

¹The Title Transfer Facility (TTF) is a virtual market place for natural gas in the Netherlands. It is commonly used as a price reference for gas contracts in both the Netherlands and Germany.

3. Coal API2², traded on the EEX in USD/t.
4. CO₂ EU Allowances, traded on the EEX in EUR/t. One EUA permits the emission of one ton of carbon dioxide. The futures contract size is 1000 EUAs.

We perform the estimations on a sample that includes the front-year contracts on gas, power, and coal, together with the futures contract on carbon EUAs for delivery at the end of the current year. The futures prices are sampled daily, and are organized as rolling contracts. The front-year futures contracts are traded until the last trading day of a year for the delivery of the underlying over the next calendar year. The EUA futures contracts are settled in December of the specified year. Our sample starts on January 3, 2008, and ends on January 15, 2014. The choice of the starting point is related to the specifics of the carbon EUA market. The European Emissions Trading Scheme (EU ETS) was introduced in 2005 and was planned to be implemented in three phases, or three trading periods. The first phase (2005 – 2007) was highly volatile and, during this period, prices could triple or collapse by half over a one-week period. In 2007, carbon prices fell to almost zero, compared to a peak level of around 30 EUR/t, when it became known that the aggregate emissions were in fact lower than the number of allowances issued. The carbon derivatives market was highly illiquid until the beginning of the second phase in 2008. The daily futures settlement prices for the first phase are available, but there were no actual trades on most of these days. Therefore, we chose to start our sample in 2008. The carbon market is still extremely volatile, and experienced a number of sharp rises and falls, and not only in the early stages. For exam-

²API2 is a price index calculated as the average of the Argus cif (cost, insurance, and freight), ARA (Amsterdam, Rotterdam, and Antwerpen) assessment, and McCloskey's northwest European steam coal marker, and is the primary price reference for coal contracts in North Western Europe.

ple, on April 16, 2013, the price of a yearly carbon futures contract dropped by 42%, from 4.97 EUR/t to 3.25 EUR/t. This happened after the European Parliament rejected a proposal to delay the sales of 900 million EUAs as a supply restriction measure intended to artificially raise the price during the period of economic slowdown and the drop in power production. Extreme price movements and high volatility were characteristic features of the carbon market, and will likely remain so unless there are significant regulatory changes, such as, for example, price caps and floors. Therefore, we keep the extreme observations in our sample, and take their presence into account in the estimations.

Table 2.1 presents the descriptive statistics for the daily log-returns on our futures contracts. The returns resulting from rolling to a new contract are deleted from the sample.

Table 2.1: Summary statistics for log-returns.

	Gas	Power	Coal	Carbon
Mean (%)	-0.036	-0.048	-0.033	-0.099
Median (%)	-0.047	-0.101	-0.017	0.000
Maximum (%)	9.223	6.508	8.622	22.369
Minimum (%)	-7.406	-5.908	-9.820	-42.476
Std. Dev. (%)	1.378	1.087	1.462	3.357
Skewness	0.360	0.222	-0.334	-1.013
Kurtosis	7.212	7.509	8.943	23.108

Table 2.1 shows that all series are leptokurtic, with the carbon returns being the most extreme case. It is worth noting that we have both positively and negatively skewed series in our sample; the gas and power returns are positively skewed, while coal and carbon returns feature a negative skew. This motivates us to not only introduce asymmetry and heavy tails when making distributional assumptions, but also allow for different properties in the individual series.

Figure 2.1 indicates that all return series show time-varying volatility and volatility clustering. The energy markets were affected by the financial crisis and, as can be seen in Figure 2.1, the gas, power, and coal markets experienced a period of particularly high volatility from the middle of 2008 to the middle of 2009. The carbon market features several periods of pronounced volatility. Perhaps the most striking price changes occurred during the first half of 2013, which includes the turmoil caused by the European Parliament's decision not to delay EUA sales.




Figure 2.1: Time series of log-returns.

3 Model framework

This section describes the econometric specification that we use to analyse the volatility dynamics of the energy forward markets. It

consists of three building blocks: the conditional mean model, the conditional covariance model, and the choice of the distribution of innovations.

A general model within our multivariate framework can be formulated as follows:

$$r_t = \mu_t + \varepsilon_t, \quad (2.1)$$

where r_t is a $k \times 1$ vector of log-returns for k different assets, μ_t is a $k \times 1$ vector, and ε_t is a $k \times 1$ vector of zero-mean error terms with conditional covariance matrix H_t . The process μ_t is the conditional mean. Below, we discuss each component of the model in more detail.

3.1 Conditional mean

The conditional mean vector is modelled within the vector autoregression (VAR) framework. That is, each return series is assumed to be a linear function of its own past lags and the past lags of the other return series. An unrestricted VAR(p) model (lag order p) can be written as follows:

$$\mu_t = \eta + \Phi_1 r_{t-1} + \dots + \Phi_p r_{t-p}, \quad (2.2)$$

where Φ_j , for $j = 1, \dots, p$, are $k \times k$ matrices and η is a $k \times 1$ vector of constants.

3.2 Conditional covariance

We assume that the conditional covariance matrix H_t follows a multivariate Generalized Autoregressive Conditional Heteroskedasticity (GARCH) process of the Baba-Engle-Kraft-Kroner (BEKK) type developed by Engle and Kroner (1995) and Kroner and Ng (1998):

$$H_t = C'C + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + B'H_{t-1}B + D'\zeta_{t-1}\zeta'_{t-1}D, \quad (2.3)$$

where A , B , C , and D are $k \times k$ matrices, ε_{t-1} is the $k \times 1$ vector of error terms in Eq. (2.1), and ζ_{t-1} is a $k \times 1$ vector of asymmetric error terms. Each element in $\zeta_{t-1} = (\zeta_{1,t-1}, \dots, \zeta_{k,t-1})$ is defined either as:

$$\zeta_{i,t-1}^+ \equiv \max(\varepsilon_{i,t-1}, 0) \quad \text{or} \quad \zeta_{i,t-1}^- \equiv \min(\varepsilon_{i,t-1}, 0), \quad (2.4)$$

depending on whether the conditional variance is higher following a positive or a negative shock. We determine the specification of ζ by estimating univariate GARCH models with asymmetric residuals on each individual time series. This approach enforces consistency between the individual series in both the multivariate and univariate models.

The specification of the $k \times k$ lower triangular parameter matrix C is such that $C'C$ is guaranteed to be positive semi-definite, while A , B , and D are, apart from identifiability conditions, unrestricted $k \times k$ parameter matrices.³

We estimate six versions, M1–M6, of the BEKK model, which are summarized in Table 2.2. The models differ in terms of the specification of the parameter matrices A , B , and D , and whether or not the asymmetric term (D) is included. The least complex parameterization is M1, in which the parameter matrices A and B are diagonal matrices, and the asymmetric term is not included. The most complex specification is M6, in which the parameter matrices A , B , and D are non-diagonal and non-symmetric (i.e., with no restrictions on the elements), and where we include the asymmetric term. By definition, diagonal specifications allow for own-market influences on conditional volatility only, while non-diagonal specifications also allow for cross-market influences. If the parameter matrices are symmetric, spillovers between two markets are automatically the same in both directions,

³A sufficient condition to eliminate observationally equivalent structures is to fix the sign of one of the diagonal parameters in A , B , and D (see Engle and Kroner, 1995; Kroner and Ng, 1998).

while non-symmetric parameter matrices remove this restriction.

Table 2.2: Covariance specifications.

Model	Parameter matrices $A, B,$ and D	Asymmetric BEKK term
M1	diagonal	no
M2	diagonal	yes
M3	non-diagonal symmetric	no
M4	non-diagonal symmetric	yes
M5	non-diagonal non-symmetric	no
M6	non-diagonal non-symmetric	yes

An important feature of the BEKK parameterization is that it models the full covariance matrix directly. Within the most general model, a typical diagonal element $h_{ii,t}$, representing the variance of commodity i at time t , can be explicitly written as a function of the past variances, covariances, and shocks, as follows:

$$\begin{aligned}
 h_{ii,t} = & \sum_{j=i}^k c_{ji}^2 + \sum_{j=1}^k \left(a_{ji} \varepsilon_{j,t-1} \sum_{j=1}^k a_{ji} \varepsilon_{j,t-1} \right) + \\
 & + \sum_{j=1}^k \left(b_{ji} \sum_{l=1}^k b_{li} h_{lj,t-1} \right) + \sum_{j=1}^k \left(d_{ji} \zeta_{j,t-1} \sum_{j=1}^k d_{ji} \zeta_{j,t-1} \right).
 \end{aligned} \tag{2.5}$$

Consequently, off-diagonal elements in the parameter matrices A , B , and D have immediate interpretations in terms of cross-market volatility spillover effects. In particular, the parameters a_{ji} , b_{ji} , and d_{ji} , for $j \neq i$, control volatility spillovers from commodity j to commodity i . The *signs* of the off-diagonal parameters do not have a straightforward interpretation because these parameters appear in several non-linear terms determining each element of the H -matrix at each point in time. Thus, the total effect of a shock in one market on the volatility in another market is a non-linear function of the

shocks to all variables in the system. However, if, for example, parameter a_{ji} is nonsignificant, then there is automatically no effect of $\varepsilon_{j,t-1}$ on $h_{ii,t}$. This holds regardless of whether the other off-diagonal elements are significant. In contrast, the presence of a significant off-diagonal parameter in any of the A -, B -, or D -matrices allows us to conclude that spillovers exist, and also to determine their direction. In addition, we can compare the size of the spillover effects based on the magnitude of these parameters. This motivates our choice of framework. An alternative model framework might have been the dynamic conditional correlation (DCC) models proposed in Engle (2000) and Engle and Sheppard (2001), which specify the evolution of the conditional correlation matrix instead. However, as discussed above, we prefer to work with BEKK-type models because these models directly specify the evolution of the full covariance matrix, in which we can interpret the relevant parameters straightforwardly in terms of second-moment spillovers.⁴

3.3 Distributional assumptions

We complete the model framework with a specification of the joint distribution for the vector of innovations ε in Eq. (2.1). A common approach in the literature assumes the multivariate normal distribution and argues that, even if the true conditional distribution of the innovations is not normal, the Quasi-Maximum Likelihood (QML) estimator is consistent and asymptotically normal, provided that the conditional mean and conditional variance equations are correctly specified (see Bollerslev and Wooldridge, 1992). However, Engle and Gonzales-Rivera (1991) show that the QML estimator is inefficient and, furthermore, that its inefficiency increases with the degree of departure from normality. This point is particularly important for finan-

⁴See Caporin and McAleer (2012) for an interesting discussion of the similarities and differences between the BEKK and DCC models.

cial assets, for which the returns are generally skewed and leptokurtic. Furthermore, in many practical applications that involve estimating tail quantiles, distributions that incorporate non-zero skewness and excess kurtosis are highly relevant; for example, in parametric value-at-risk estimations (see Giot and Laurent, 2003; Hung, Lee, and Liu, 2008; and Cheng and Hung, 2011). Therefore, while the normal distribution may serve as a benchmark case, we believe that more flexible distributions are an important building block when modelling energy-related asset returns. In this study, we choose to deviate from the normality assumption. In particular, we implement the VAR-BEKK process in conjunction with the multivariate skew-Student density of Bauwens and Laurent (2002, 2005).

The most general version of their multivariate skew-Student distribution is constructed such that the univariate marginal distributions can have individual skewness coefficients and tail properties. Given the nature of our data, which we summarize in Table 2.1, with the carbon return series being considerably more leptokurtic, we want to relax the restriction of equal degrees of freedom implied by the standard multivariate Student distribution. We also want to allow for different skewness coefficients for the individual series, especially since we have both positively and negatively skewed variables in our sample. The multivariate skew-Student density with the independent components of Bauwens and Laurent (2002, 2005) introduces such flexibility at a reasonable computational cost. In addition, this skew-Student distribution is relatively straightforward to augment with GARCH-type second-moment dynamics. Naturally, the skewness coefficients, as well as the degrees of freedom, can be restricted to a single value, creating different types of nested distributions, the relevance of which can be statistically contrasted using standard likelihood ratio tests.

Following Bauwens and Laurent (2002), a $k \times 1$ random vector z_t is standard multivariate skew-Student distributed with independent

components if its probability density function is given by:

$$f(z_t) = \left(\frac{2}{\sqrt{\pi}} \right)^k \left[\prod_{i=1}^k \frac{\xi_i s_i}{1 + \xi_i^2} \frac{\Gamma\left(\frac{v_i+1}{2}\right)}{\Gamma\left(\frac{v_i}{2}\right) \sqrt{v_i-2}} \left(1 + \frac{\kappa_{i,t}^2}{v_i-2} \right)^{-\frac{1+v_i}{2}} \right], \quad (2.6)$$

where

$$\kappa_{i,t} = (s_i z_{i,t} + m_i) \xi_i^{-I_{i,t}}, \quad (2.7)$$

and

$$I_{i,t} = \begin{cases} 1 & \text{if } z_{i,t} \geq -\frac{m_i}{s_i} \\ -1 & \text{if } z_{i,t} < -\frac{m_i}{s_i} \end{cases}, \quad (2.8)$$

with skewness parameters $\xi = (\xi_1, \dots, \xi_k)$ and degrees of freedom parameters $v = (v_1, \dots, v_k)$ for $v_i > 2$. We let $\Gamma(x)$ denote the Gamma function evaluated at $x > 0$. We obtain the density function $f(z_t)$ by taking the product of k independent skew-Student components, thereby allowing each marginal distribution to have a different tail behaviour. In the present setting, we define the multivariate skew-Student distribution for the vector of standardized residuals z_t as follows:

$$z_t = H_t^{-1/2} \varepsilon_t, \quad (2.9)$$

where ε_t is the vector of actual residuals from the model in Eq. (2.1) and H_t is the BEKK covariance matrix in Eq. (2.3). The constants $m_i = m_i(\xi_i, v_i)$ and $s_i = s_i(\xi_i, v_i)$ are the means and standard deviations of the non-standardized univariate skew-Student density as in Fernandez and Steel (1998), respectively, defined by:

$$m_i(\xi_i, v_i) = \frac{\Gamma\left(\frac{v_i-1}{2}\right) \sqrt{v_i-2}}{\sqrt{\pi} \Gamma\left(\frac{v_i}{2}\right)} \left(\xi_i - \frac{1}{\xi_i} \right), \quad (2.10)$$

$$s_i^2(\xi_i, v_i) = \left(\xi_i^2 + \frac{1}{\xi_i^2} - 1 \right) - m_i^2. \quad (2.11)$$

The parameter ξ_i^2 is the ratio of probability masses above and below the mode, and can be interpreted directly as a measure of skewness. In the case where $\xi_i < 1$, the data are negatively skewed, and $\xi_i > 1$ indicates positive skewness. The symmetric case corresponds to $\xi_i = 1$, which implies that $m_i = 0$ and $s_i = 1$. If we restrict all ξ_i to be equal to 1, and all v_i to be the same, Eq. (2.6) reduces to a distribution similar to the textbook multivariate Student density.

4 Estimation

In our final four-asset model framework, we organize the vector of futures log-returns as $r_t = (r_{1,t}, r_{2,t}, r_{3,t}, r_{4,t})'$, where $r_{1,t}$ denotes the return on natural gas, $r_{2,t}$ denotes the return on power, $r_{3,t}$ denotes the return on coal, and $r_{4,t}$ denotes the return on carbon.

We start by determining the appropriate lag order p for the VAR part governing the mean equation. To do so, we employ a number of criteria. Individual correlograms of return series indicate that autocorrelation is present at the first lag, and in some cases, at the second lag as well. Next, we compare the VAR models of up to the fifth order based on the information criteria (AIC, SIC, and HQ), the sequential likelihood ratio (LR) test statistics, and the final prediction error.⁵ We select two lags by three out of these five criteria, and none favour more than two lags. The final check is a test for any remaining serial correlation in the residuals. According to the multivariate LM test, we can reject the null hypothesis of no autocorrelation up to the second order in residuals of VAR(1) at any conventional significance level. In contrast, for the VAR(2) model, we cannot reject the hypothesis of no autocorrelation up to the second order at the 5% significance level. Therefore, we choose a VAR(2) specification for the conditional mean process.

⁵We estimate these VAR models assuming normally distributed innovations.

To determine an appropriate definition for the asymmetric error term of each asset, ζ_i , we estimate univariate GARCH models of the type proposed in Glosten, Jagannathan, and Runkle (1993) for each series. Then, we pick the best specification based on the likelihood value. This analysis leads us to specify the vector ζ as:

$$\zeta_t = \left(\zeta_{1,t}^+, \zeta_{2,t}^-, \zeta_{3,t}^+, \zeta_{4,t}^- \right)',$$

where $\zeta_{i,t}^+$ and $\zeta_{i,t}^-$ are defined in Eq. (2.4). This specification is consistent with conditional variance being higher after a negative shock for power and carbon (the leverage effect), but higher after a positive shock for gas and coal (the inverse leverage effect).

4.1 Methodology

We estimate all parameters in our models simultaneously using full information maximum likelihood (ML). Our estimation methodology proceeds in three steps. First, we use the OLS method to estimate the parameters in the mean equations, ignoring the GARCH error structure. Then, we estimate the GARCH parameters by QML, assuming normality and conditional on the given VAR parameters. These two steps yield consistent estimates of all mean and covariance parameters (Bollerslev and Wooldridge, 1992). However, to obtain efficient estimates, we require a joint estimation of all parameters. This motivates our final step, in which we re-estimate all parameters using the parameter estimates from the two initial steps as starting values only.⁶ We implement this procedure for all six covariance specifications in Table 2.2.

⁶As starting values for the distributional parameters in the skew-Student distributions, we use $\xi_i = 1$ for all skewness parameters. For the degrees of freedom v_i , we use either the value corresponding to the average kurtosis of the data series, or the values corresponding to the individual kurtosis of the data series, depending on the specification (see Table 2.3).

Let θ denote the parameter vector for the full model. Then, the log-likelihood function is given by:

$$\ln L(\theta) = \sum_{t=3}^T \left\{ \ln f(z_t) - \frac{1}{2} \ln |H_t| \right\}, \quad (2.12)$$

where $f(z_t)$ is the probability density function in Eq. (2.6), T is the number of time series observations, and $|H_t|$ denotes the determinant of H_t . Note that the summation starts from $t = 3$ because the estimation is conditional on the first two time series observations owing to the VAR(2) specification of the mean equation. The second term in the sum in Eq. (2.12) is the Jacobian correction term arising in the transformation from z to ε . To evaluate the likelihood function, we calculate the inverse of the square root matrix $H_t^{-1/2}$ in Eq. (2.9) at each time point using a standard spectral decomposition. We set the initial H_t equal to the sample covariance matrix and the initial values of the residuals are set equal to zero.

We estimate the six BEKK specifications described in Table 2.2 by maximizing the log-likelihood function in Eq. (2.12). In addition, we estimate each model under the assumption of the six types of multivariate Student distributions summarized in Table 2.3.

We also estimate all BEKK specifications for the benchmark case of the normal distribution. In the case of normally distributed residuals, we replace the density function in Eq. (2.12) by the standardized normal density obtained as the limiting distribution of $f(z_t)$ when $\xi_i = 1$, as $v_i \rightarrow \infty$.

Table 2.3: Types of Student distributions.

Parameter	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6
v_1	v	v_1	v	v_1	v	v_1
v_2	v	v_2	v	v_2	v	v_2
v_3	v	v_3	v	v_3	v	v_3
v_4	v	v_4	v	v_4	v	v_4
ξ_1	1	1	ξ	ξ	ξ_1	ξ_1
ξ_2	1	1	ξ	ξ	ξ_2	ξ_2
ξ_3	1	1	ξ	ξ	ξ_3	ξ_3
ξ_4	1	1	ξ	ξ	ξ_4	ξ_4

Note: Types 1 and 2 correspond to symmetric distributions. Types 3 and 4 are asymmetric with a common value of the skewness parameter, while Types 5 and 6 allow the variables to have different skewness properties. With respect to kurtosis, Types 1, 3, and 5 restrict degrees of freedom parameters to a common value, while Types 2, 4, and 6 allow them to vary.

In total, we estimate 42 different model specifications. The number of conditional mean parameters is always 36. The total number of conditional covariance parameters together with the skewness and degrees of freedom parameters ranges from 18 in the simplest specification (M1/Normal) to 66 in the most complex specification (M6/Type6). The log-likelihood function is maximized by simulated annealing, which is a derivative-free stochastic search algorithm. The fundamental property of simulated annealing is that it is allowed to accept worse intermediate solutions (downhill moves) while searching for the optimum, which leads to a more extensive exploration of the parameter space and prevents the algorithm from becoming stuck in local optima. In theory, this property also makes the algorithm insensitive to starting values. However, to further increase the chance of identifying the global optimum, we implement the sequential strategy described above, which involves using consistent QML estimates as starting values. Our particular implementation of the algorithm follows the approach in Goffe, Ferrier, and Rogers (1994) closely. The

advantages of simulated annealing come at the cost of a higher execution time compared to conventional algorithms. Thus, we execute all optimizations on a high-performance computer cluster.⁷ We calculate the standard errors for individual parameters by estimating the outer product of the gradients matrix using numerical first derivatives.⁸

5 Results

This section presents the estimation results. We first describe the likelihood ratio tests for the preferred model specification. Next, we discuss the estimated parameters from the best model, with a particular focus on volatility spillovers. Finally, we examine the estimated conditional correlations and volatilities.

5.1 Model specification

The mean equation in our model is a preselected VAR(2) specification. Therefore, the choice of the preferred model involves two parts: the covariance specification and the distributional assumption.

Table 2.4 reports the results of the LR tests of the six BEKK specifications. We find that, regardless of the additional assumptions, the models with an asymmetric term and with non-diagonal, non-symmetric parameter matrices are superior. Because specifications with diagonal parameter matrices are rejected against their non-diagonal counterparts, we conclude that volatility spillovers in the energy forward markets exist. We also infer that these spillovers are not the same in both directions because we reject the models with

⁷The computations were performed on resources provided by the Swedish National Infrastructure for Computing (SNIC) at LUNARC, Lund University.

⁸An alternative is to calculate the standard errors based on the inverse of the Hessian. However, implementing stable and reliable numerical second derivatives is a challenge, even in less complex settings than ours. Thus, we leave this topic for future research.

non-diagonal, symmetric parameter matrices against the models with non-diagonal and non-symmetric parameter matrices.

Table 2.4: Likelihood ratio tests of covariance specifications.

Test	5% cr.v.	1% cr.v.	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6
1 M1 vs. M2 LR (4)	9.49	13.28	50.39	48.35	49.89	47.66	50.86	49.29
2 M3 vs. M4 LR (10)	18.31	23.21	71.14	70.63	70.94	70.02	70.75	70.46
3 M5 vs. M6 LR (16)	26.30	32.00	74.30	75.27	76.13	77.18	76.57	76.93
4 M1 vs. M3 LR (12)	21.03	26.22	36.20	36.25	35.02	35.24	32.00	32.36
5 M2 vs. M4 LR (18)	28.87	34.81	56.95	58.53	56.06	57.60	51.89	53.53
6 M1 vs. M5 LR (24)	36.42	42.98	65.27	64.07	62.48	60.96	57.28	57.15
7 M2 vs. M6 LR (36)	51.00	58.62	89.18	90.99	88.72	90.47	82.99	84.80
8 M3 vs. M5 LR (12)	21.03	26.22	29.07	27.82	27.46	25.72	25.28	24.80
9 M4 vs. M6 LR (18)	28.87	34.81	32.23	32.46	32.66	32.87	31.10	31.27

Note: Tests 1–3 are for the asymmetric BEKK term ($H_0: D = 0$ versus $H_1: D \neq 0$). Tests 4–5 are for H_0 : diagonal matrices versus H_1 : non-diagonal, symmetric matrices. Tests 6–7 are for H_0 : diagonal matrices versus H_1 : non-diagonal, non-symmetric matrices. Tests 8–9 are for H_0 : non-diagonal symmetric matrices versus H_1 : non-diagonal, non-symmetric matrices.

The remaining combinations of models are not nested. The degrees of freedom are reported in parentheses. Columns 3 and 4 report the upper-tail critical values of χ^2 -distribution with the corresponding degrees of freedom.

In summary, based on the LR tests, we prefer the most general covariance model (M6), with non-diagonal, non-symmetric parameter matrices and an asymmetric term.

Irrespective of the covariance model, the log-likelihood value increases significantly when switching from the normal distribution to

the simplest Student distribution (Type 1). This suggests that allowing for excess kurtosis is also highly important in a conditional setting, although we cannot do an LR test for non-nested distributional specifications. The test results for the nested Student specifications are summarized in Table 2.5. These results hold regardless of the covariance model specification, and we report the test statistics and the p-values for the M6 covariance model only.

Table 2.5: Likelihood ratio tests of distributional specifications for the M6 model.

LR	Type 1	Type 2	Type 3	Type 4	Type 5	Type 6
Type 1		5.81 (0.1212)	7.25 (0.0071)	13.44 (0.0093)	25.20 (0.0000)	31.27 (0.0001)
Type 2			–	7.63 (0.0057)	–	25.46 (0.0000)
Type 3				6.19 (0.1027)	17.95 (0.0005)	24.02 (0.0005)
Type 4					–	17.83 (0.0005)
Type 5						6.07 (0.1083)
Type 6						

Note: ‘–’ indicates that the distributional types are not nested. The χ^2 p-values for upper-tail one-sided tests with the corresponding degrees of freedom are reported in parentheses.

First, we can reject Type 1 against Type 3, and Type 2 against Type 4, which means that introducing non-zero skewness leads to a statistically significant increase in the log-likelihood value. Moreover, allowing for individual skewness coefficients is also statistically important because we can strongly reject Type 3 against Type 5, and Type 4 against Type 6. However, as indicated by the p-values exceeding 0.10, we cannot strongly reject Type 1 against Type 2, Type 3 against Type 4, and Type 5 against Type 6. This indicates that allowing for

individual excess kurtosis is not as important as allowing for individual skewness. Based on these tests, our preferred distribution is Type 5, which is the multivariate skew-Student distribution with a common degrees of freedom parameter and with individual skewness parameters.

We base the further analysis in this section on the results obtained from the preferred model: M6/Type5.

5.2 Parameter estimates and volatility spillovers

Tables 2.6 and 2.7 present the estimation results for the M6/Type5 model. We are particularly interested in the off-diagonal elements of the A , B , and D parameter matrices because they control the cross-market effects in the second-moment dynamics. However, we start with a general overview of the estimation results.

The diagonal elements of the B -matrix indicate a high level of persistence in the volatility of energy log-returns. Among the diagonal elements of the D -matrix, we find a significant d_{33} for coal and d_{44} for carbon. This confirms the presence of asymmetric effects in the conditional volatility of these assets. The inverse leverage effect that we find in the coal series is consistent with the theory of storage in commodity markets, arguing that high prices occur at periods of low inventory levels and are associated with high volatility (see Deaton and Laroque, 1992).

Based on the 95% confidence intervals, the skewness parameters ξ_1 of gas and ξ_2 of power are statistically significantly larger than 1, which is consistent with the positive sign on the unconditional skewness in the data. The skewness parameter estimates of coal and carbon are in the negative region, slightly below 1. However, since the upper bound of the confidence intervals is above 1 for these parameters, we can draw no conclusions. Further, we estimate the common

degrees of freedom parameter ν of 7.6792.⁹

Table 2.6: Conditional mean parameter estimates for the M6/Type5 model.

η_1	0.0080	η_2	-0.0425**	η_3	-0.0277	η_4	-0.0674
	(0.0207)		(0.0186)		(0.0250)		(0.0603)
$\phi_{11}^{(1)}$	0.0433	$\phi_{12}^{(1)}$	0.0150	$\phi_{13}^{(1)}$	0.0064	$\phi_{14}^{(1)}$	0.0007
	(0.0300)		(0.0410)		(0.0250)		(0.0063)
$\phi_{21}^{(1)}$	0.0214	$\phi_{22}^{(1)}$	-0.0467	$\phi_{23}^{(1)}$	0.0213	$\phi_{24}^{(1)}$	0.0283***
	(0.0222)		(0.0349)		(0.0206)		(0.0059)
$\phi_{31}^{(1)}$	0.0673**	$\phi_{32}^{(1)}$	0.1099**	$\phi_{33}^{(1)}$	-0.0191	$\phi_{34}^{(1)}$	-0.0006
	(0.0297)		(0.0448)		(0.0291)		(0.0078)
$\phi_{41}^{(1)}$	-0.0779	$\phi_{42}^{(1)}$	0.0047	$\phi_{43}^{(1)}$	-0.1056**	$\phi_{44}^{(1)}$	0.0346
	(0.0499)		(0.0849)		(0.0503)		(0.0294)
$\phi_{11}^{(2)}$	0.0407	$\phi_{12}^{(2)}$	-0.0163	$\phi_{13}^{(2)}$	0.0017	$\phi_{14}^{(2)}$	-0.0022
	(0.0294)		(0.0394)		(0.0250)		(0.0066)
$\phi_{21}^{(2)}$	0.0140	$\phi_{22}^{(2)}$	-0.0356	$\phi_{23}^{(2)}$	0.0160	$\phi_{24}^{(2)}$	-0.0076
	(0.0220)		(0.0343)		(0.0212)		(0.0065)
$\phi_{31}^{(2)}$	-0.0179	$\phi_{32}^{(2)}$	0.0166	$\phi_{33}^{(2)}$	-0.0153	$\phi_{34}^{(2)}$	-0.0026
	(0.0290)		(0.0436)		(0.0287)		(0.0075)
$\phi_{41}^{(2)}$	0.0446	$\phi_{42}^{(2)}$	0.0005	$\phi_{43}^{(2)}$	-0.0406	$\phi_{44}^{(2)}$	-0.0602**
	(0.0532)		(0.0851)		(0.0539)		(0.0295)

Note: The conditional mean is given by Eq. (2.2), where η is a 4×1 vector of constants, and Φ_1 and Φ_2 are 4×4 VAR parameter matrices with elements denoted by $\phi_{ij}^{(1)}$ and $\phi_{ij}^{(2)}$, for $i, j = 1$ (gas), 2 (power), 3 (coal), 4 (carbon), respectively. $\phi_{ij}^{(p)}$ represents the effect of commodity j on commodity i in lag p . Standard errors are reported in parentheses. Superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

⁹We estimate ν^{-1} instead of ν for numerical reasons.

Table 2.7: Conditional covariance and distributional parameter estimates for the M6/Type5 model.

a_{11}	0.1881*** (0.0216)	a_{12}	0.0322* (0.0175)	a_{13}	0.0275 (0.0184)	a_{14}	-0.0182 (0.0456)
a_{21}	0.0279 (0.0358)	a_{22}	0.1695** (0.0316)	a_{23}	0.0481 (0.0338)	a_{24}	-0.0572 (0.0907)
a_{31}	-0.0061 (0.0201)	a_{32}	0.0281 (0.0189)	a_{33}	0.1413*** (0.0217)	a_{34}	0.0185 (0.0554)
a_{41}	-0.0021 (0.0051)	a_{42}	0.0106* (0.0054)	a_{43}	-0.0003 (0.0054)	a_{44}	0.2727*** (0.0329)
b_{11}	0.9806*** (0.0054)	b_{12}	0.0053 (0.0050)	b_{13}	-0.0076 (0.0048)	b_{14}	0.0090 (0.0133)
b_{21}	-0.0212** (0.0107)	b_{22}	0.9606*** (0.0096)	b_{23}	-0.0145 (0.0097)	b_{24}	0.0572* (0.0295)
b_{31}	0.0073* (0.0038)	b_{32}	-0.0033 (0.0042)	b_{33}	0.9895*** (0.0035)	b_{34}	-0.0206* (0.0120)
b_{41}	0.0022 (0.0016)	b_{42}	-0.0024 (0.0019)	b_{43}	-0.0003 (0.0018)	b_{44}	0.9211*** (0.0093)
c_{11}	0.0017 (0.0533)						
c_{21}	-0.0302 (0.0286)	c_{22}	-0.0434 (0.0330)				
c_{31}	0.0035 (0.0244)	c_{32}	0.0715* (0.0272)	c_{33}	0.0556** (0.0243)		
c_{41}	-0.0124 (0.0275)	c_{42}	0.0250 (0.0268)	c_{43}	0.0137 (0.0326)	c_{44}	0.3263*** (0.0637)
d_{11}	0.0602 (0.0391)	d_{12}	-0.0869*** (0.0273)	d_{13}	-0.0128 (0.0261)	d_{14}	-0.0845 (0.0593)
d_{21}	-0.0055 (0.0456)	d_{22}	0.0523 (0.0467)	d_{23}	-0.0360 (0.0479)	d_{24}	0.1047 (0.0976)
d_{31}	0.0808*** (0.0306)	d_{32}	0.1378*** (0.0280)	d_{33}	0.0987*** (0.0367)	d_{34}	0.1525** (0.0772)
d_{41}	0.0037 (0.0069)	d_{42}	-0.0169** (0.0083)	d_{43}	0.0028 (0.0083)	d_{44}	-0.3608*** (0.0448)
ξ_1	1.1454*** (0.0455)	ξ_2	1.1187*** (0.0455)	ξ_3	0.9893*** (0.0405)	ξ_4	0.9641*** (0.0388)
ν^{-1}	0.1302*** (0.0131)						

Note: The conditional covariance matrix is given in Eq. (2.3) and specified by the 4×4 matrices A , B , C , and D , the elements of which are denoted by a_{ij} , b_{ij} , c_{ij} , and d_{ij} , for $i, j = 1$ (gas), 2 (power), 3 (coal), and 4 (carbon), respectively. The C -matrix is lower triangular and, therefore, there are no estimates for the entries above the diagonal. The distributional parameters are ν^{-1} , ξ_1 , ξ_2 , ξ_3 , and ξ_4 . Standard errors are reported in parentheses. Superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Of the nine parameters that control volatility spillovers to the power market (a_{12} , a_{32} , a_{42} , b_{12} , b_{32} , b_{42} , d_{12} , d_{32} , d_{42}), we find that a_{12} , a_{42} , d_{12} , d_{32} , and d_{42} are statistically significant. These spillovers originate from all other markets, and are channelled mostly through the asymmetric BEKK term, with all three of the d -parameters in question being significant.

The results presented in Table 2.7 indicate that the most sizeable and statistically significant volatility spillover effect to the power market comes from the coal market (as captured by d_{32}). We also find statistically significant spillover effects to the power market from the gas and carbon markets, but these effects are of lower magnitude. This result is consistent with the stronger interrelations in volatility between electrical power and the fuel type that is currently in-the-money. Figure 2.2 plots the yearly clean spark and dark spreads during our sample period, which indicate the profit margins of gas- and coal-fired power plants, respectively.¹⁰

Figure 2.2 shows that the dark spread is mostly higher than the spark spread, which means that power generation from coal was more profitable than that from gas during almost the entire sample period. Moreover, coal never falls out-of-the-money, unlike gas, for which the spark spread becomes negative from the beginning of 2013 onwards.

¹⁰We use the following typical definitions of the clean spark spread (CSS) and clean dark spread (CDS):

$$\text{CSS} = \text{Power} - 2 \times \text{Gas} - 0.4 \times \text{Carbon},$$

and

$$\text{CDS} = \text{Power} - 0.4 \times \text{Coal} - 0.9 \times \text{Carbon}.$$

We use the peak load power price in all calculations.

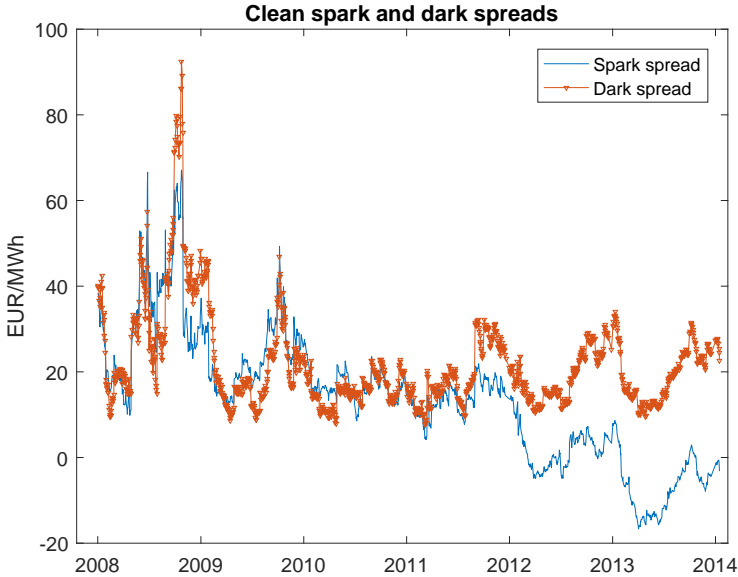


Figure 2.2: Time series of clean spark and clean dark spreads.

To gain additional insights, we examine the data on actual electricity production in Germany from different technologies.¹¹ Figure 2.3 plots the weekly volumes of electricity (in TW) generated in Germany using coal and gas by a number of electricity producers.

Much more electricity is generated from coal almost every week during our sample period. Apart from winter–summer seasonality, we can clearly observe that generation from gas is steadily falling, owing to decreasing spark spreads. Additionally, the difference between the electricity volumes produced from the two fuels grows steadily as the dark spreads wander further up from the spark spreads. This trend is further supported by the recent low carbon price environment. We believe that these fundamental relationships in the supply stack are

¹¹Source: <http://www.transparency.eex.com>. This information is reported under voluntary commitments by around 40% of market participants registered on the EEX transparency platform.

the reason why the spillovers from gas to power are smaller than from coal to power.




Figure 2.3: Time series of weekly electricity generation from coal and gas in Germany.

In addition to the reported volatility spillovers from the fuel components to the power price, we find statistically significant results in the opposite direction, namely from power to gas and carbon (the significant parameters are b_{21} and b_{24}). There is no clear fundamental reason why volatility spillovers would occur in this direction. Our interpretation is that a given increase in the forward power volatility creates uncertainty about the future power price level, which, in turn, leads to additional uncertainty about the future gas and carbon volumes planned for power production. Finally, this uncertainty is transmitted to the gas and carbon prices. We see similar results for the coal and gas volatilities, where we report a statistically significant

spillover effect from the coal market to the gas market (parameters b_{31} and d_{31}). In general, coal appears to be the driving commodity in terms of volatility in our system since we find many significant spillovers channelled from the coal market, and none channelled to it. We believe that an increase in coal volatility impacts the uncertainty of the future production mix (coal/gas) for power production, which manifests as gas volatility. Moreover, we report a fairly sizeable and statistically significant volatility spillover effect from the coal market to the carbon market (parameters b_{34} and d_{34}). We note that the coal and carbon markets are clearly linked via the power market because coal power plants are major CO₂ emitters. In terms of electrical energy, coal emits more than twice the CO₂ of gas. This is verified by the carbon coefficients in the clean spark and dark spread definitions (see footnote 10), which represent the number of carbon credits necessary to cover the respective power production. We believe that the connection between the coal and carbon markets, along with coal having been a profitable technology (in-the-money) for power production during the full sample period explain the significant volatility spillover effects from coal to power. An increase in coal volatility likely creates uncertainty about the future production mix (coal/gas), and hence uncertainty about future emitted CO₂ volumes, which finally transmits to the carbon prices.

5.3 Conditional correlations and volatilities

Next, we use the estimated parameters to calculate the conditional second moments. Figure 2.4 displays the conditional volatilities of the log-returns on our four commodities as implied by the preferred model.

We can see that the estimated processes are in line with the time

series of log-returns in Figure 2.1.¹² Carbon has the highest volatility on average, characterized by a dramatic increase towards the end of the sample period. This period of elevated volatility was mostly related to the uncertainty about the EUA supply in the beginning of 2013, and the subsequent policy decisions mentioned in Section 2. Note also the contemporaneous spike in power volatility, while gas and coal markets retained tranquillity. This might be an example of a volatility spillover effect, since to our knowledge, nothing idiosyncratic happened at the same time in the German power market.




Figure 2.4: Conditional volatilities of commodity log-returns implied by the M6/Type5 model.

We proceed our analysis by investigating conditional correlations. Correlations measure the degree of co-movement between energy com-

¹²In addition, we find our GARCH sample averages of volatilities to be 1.373 (gas), 1.061 (power), 1.408 (coal), and 3.569 (carbon), which is in line with the sample standard deviations of log-returns in Table 2.1.

modities, which is an essential determinant of power producers' hedging strategies. Figure 2.5 presents the power/gas and power/coal correlations, as implied by the preferred model.

We find a clear decreasing trend in both series during the sample period. The average power/gas correlation coefficient falls by half, from 0.72 in 2008 to 0.36 in 2013, while the average power/coal correlation coefficient decreases from 0.67 to 0.46.




Figure 2.5: Conditional correlations between power and fuels implied by the M6/Type5 model.

The correlations between power and gas exhibit a particularly sharp decline from the beginning of 2013, which coincides with the spark spreads turning negative. Less co-movement between power and fossil fuels can be attributed to the rise in renewables, which offer power generation at the lowest marginal cost, decreasing the long-run mean power price level. In Germany, the share of renewables

in the power generation mix grew from 16% in 2008 to 25% in 2013. Furthermore, the German government aims to raise it to 35% by 2020, and to 50% by 2030.¹³ The build-up of renewable generation capacity also leads to lower loads on thermal power plants. These structural changes in the supply stack tend to loosen the relationship between power and fossil fuels, which is reflected in the decreasing correlations.

6 Summary and conclusions

In this study, we investigate modelling of electrical power, gas, coal, and carbon in the multivariate setting. These commodities are fundamentally linked through the electricity generation process, which gives rise to a number of interdependencies. The literature traditionally focused on the interdependencies between the prices or returns, i.e., in the first moments. This study finds that non-trivial cross-market effects also exist in the second moments, and, in explaining the changes in volatility in the power market, for instance, one should take into account volatility in the related markets. We estimate a large number of model specifications within the VAR-BEKK framework. The model selection results indicate that accounting for cross-market effects and asymmetric effects in volatility leads to a statistically significant improvement in the likelihood value. Our preferred skew-Student distributional specification features individual skewness parameters for the four commodities, and a common degrees of freedom parameter to determine the tail properties. An analysis of the estimated conditional correlations reveals a decrease in co-movement between power and fossil fuels during 2008 – 2013, which can be related to the growing share of renewable generation sources in Germany. We find that the highest magnitude volatility spillover effect to the power market comes from the coal market. This result is consistent with coal's

¹³Source: BDEW German Association of Energy and Water Industries (2014).

higher profitability and larger share in the generation mix compared to gas. While the estimated in this study coefficients provide an idea of the average spillover effects during the sample period, a natural extension is to adopt a dynamic perspective, revealing how volatility transmission has evolved over time.

References

- [1] Bauwens, L., Laurent, S., 2002. A new class of multivariate skew densities, with applications to GARCH models. CORE Discussion Paper 20.
- [2] Bauwens, L., Laurent, S., 2005. A new class of multivariate skew densities, with application to Generalized Autoregressive Conditional Heteroscedasticity models. *Journal of Business and Economic Statistics* 23, 346–354.
- [3] Bollerslev, T., Wooldridge, J. M., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews* 11, 143–172.
- [4] Caporin, M., McAleer, M., 2012. Do we really need both BEKK and DCC? A tale of two multivariate GARCH models. *Journal of Economic Surveys* 26, 736–751.
- [5] Casassus, J., Liu, P., Tang, K., 2013. Economic linkages, relative scarcity, and commodity futures returns. *Review of Financial Studies* 26, 1324–1362.
- [6] Deaton, A., Laroque, G., 1992. On the behaviour of commodity prices. *Review of Economic Studies* 59, 1–23.
- [7] Cheng, W.-H., Hung, J.-C., 2011. Skewness and leptokurtosis in GARCH-typed VaR estimation of petroleum and metal asset returns. *Journal of Empirical Finance* 18, 160–173.

- [8] De Vany, A.S., Walls, W.D., 1999. Cointegration analysis of spot electricity prices: insights on transmission efficiency in the western US. *Energy Economics* 21, 435–448.
- [9] Efimova, O., Serletis, A., 2014. Energy markets volatility modelling using GARCH. *Energy Economics* 43, 264–273.
- [10] Engle, R., 2000. Dynamic conditional correlation – a simple class of multivariate GARCH models. *Journal of Business and Economic Statistics* 20, 339–350.
- [11] Engle, R., Sheppard, K., 2001. Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. NBER Working Paper Series No. 8554.
- [12] Engle, R., Gonzalez-Rivera, G., 1991. Semiparametric ARCH Model. *Journal of Business and Economic Statistics* 9, 345–360.
- [13] Engle, R., Kroner, K., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory* 11, 122–150.
- [14] Fernandez, C., Steel, M.F.J., 1998. On Bayesian modelling of fat tails and skewness. *Journal of the American Statistical Association* 93, 359–371.
- [15] Giot, P., Laurent, S., 2003. Value-at-Risk for long and short positions. *Journal of Applied Econometrics* 18, 641–664.
- [16] Gjolberg, O., Johnsen, T., 1999. Risk management in the oil industry: can information on long-run equilibrium prices be utilized? *Energy Economics* 21, 517–527.
- [17] Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779–1801.
- [18] Goffe, W., Ferrier, G., Rogers, J., 1994. Global optimization of statistical functions with simulated annealing. *Journal of Econo-*

metrics 60, 65–99.

- [19] Hung, J.C., Lee, M.C., Liu, H.C., 2008. Estimation of Value-at-Risk for energy commodities via fat-tailed GARCH models. *Energy Economics* 30, 1173–1191.
- [20] Karali, B., Ramirez, O.A., 2014. Macro determinants of volatility and volatility spillover in energy markets. *Energy Economics* 46, 413–421.
- [21] Koenig, P., 2011. Modelling correlation in carbon and energy markets. *Cambridge Working Paper in Economics* 1123.
- [22] Kroner, K., Ng, V., 1998. Modeling asymmetric comovements of asset returns. *Review of Financial Studies* 11, 817–844.
- [23] Lin, S.X., Tamvakis, M.V., 2001. Spillover effects in energy futures markets. *Energy Economics* 23, 43–56.
- [24] Liu, H.-H., Chen, Y.-C., 2013. A study on the volatility spillovers, long memory effects and interactions between carbon and energy markets: The impacts of extreme weather. *Economic Modelling* 35, 840–855.
- [25] Reboredo, J.C., 2014. Volatility spillovers between the oil market and the European Union carbon emission market. *Economic Modelling* 36, 229–234.

CHAPTER 3

Volatility Transmission in the German Energy Markets: A Variance Impulse Response Analysis

with Rikard Green, Karl Larsson, and Birger Nilsson

Abstract

This study investigates the transmission of volatility between electrical power, natural gas, coal, and carbon emission allowances in the German energy market during 2008 – 2013. We focus on the impact of large exogenous shocks in gas, coal, and carbon on the expected variance of power. We base our methodology on the concept of the Volatility Impulse Response Function (VIRF) introduced in Hafner and Herwartz (2006). We suggest a method to normalize the VIRF such that the results are straightforward to interpret and comparable over time and across different markets. The results indicate that positive news (i.e., price increases) in both gas and coal markets leads to a much larger variance response in the power market than negative news (i.e., price decreases). The impact of the gas market news is weaker on average compared to the coal market news, however it takes longer to die out. Spillovers from the carbon market show the fastest decay and are nonsignificant until 2011. Benchmarking the magnitude of the responses in the variance of power against the own-market variance responses reveals a non-trivial size of the cross-market effects, and suggests the relevance of taking them into account in practical applications.

Keywords: volatility impulse response function, time-varying volatility spillovers, German energy markets, asymmetric BEKK

JEL Classification: C32, C58, G1, Q41

1 Introduction

In Chapter 2, we uncovered interdependencies in volatility between electrical power, natural gas, coal, and carbon emission allowances in the German energy forward markets. Although we estimated a model with time-varying conditional second moments, our interpretation of the volatility spillover effects is rather static and limited to constant spillover coefficients. The aim of this study is to provide a rich dynamic perspective on the volatility transmission, and to quantify the impact of the shocks in a way that is straightforward to interpret and ensures comparability both over time and across different markets.

To achieve this goal, we apply the volatility impulse response analysis, which is a powerful tool for understanding the mechanics of shock transmission. Our implementation is based largely on the methodologies in Koop et al. (1996) and Hafner and Herwartz (2006). Koop et al. (1996) develop the Generalized Impulse Response Function (GI), which deals with the problems of history, shock, and compositional dependence of traditional impulse response functions when applied to multivariate non-linear systems. This methodology was traditionally applied to conditional mean systems. Hafner and Herwartz (2006) extend it to multivariate Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models and introduce the concept of the Volatility Impulse Response Function (VIRF). The VIRF measures the difference in expected (co)variance following an independent shock, and the expected ‘baseline’ (co)variance without the information about this shock. They derive analytical expressions for the VIRF in a general symmetric multivariate GARCH setting. In this study, however, we work with an asymmetric BEKK model selected based on the analysis in Chapter 2, and therefore, we do not use the VIRF as defined in Hafner and Herwartz (2006). Instead, we compute the VIRF numerically via Monte Carlo integration tech-

niques outlined for the GI in the earlier work of Koop et al. (1996). Further, we suggest an effective way of normalizing the responses that facilitates the interpretation of the results.

Volatility impulse response analysis is a relatively recent development. The VIRF methodology of Hafner and Herwartz (2006) was applied to equity and commodity markets in a number of very recent studies (see Olson et al., 2014; Hasanov et al., 2016; Jin and An, 2016). Le Pen and Sevi (2010) investigate volatility transmission between a number of European electricity markets, also with the help of the analytical VIRF expressions in the symmetric BEKK setting. To our knowledge, this is the first comprehensive study of volatility transmission between the different energy markets.

We investigate how the exogenous shocks in the gas, coal, and carbon markets affected the expected power market volatility at different times throughout the sample period, from the beginning of 2008 until the end of 2013. We draw these exogenous shocks (*news*) from the underlying skew-Student return distributions. We consider news of the magnitude that occur two or three trading days per year. Working with an asymmetric BEKK model provides us with an opportunity to distinguish between the effects of positive and negative news, which we find highly relevant. We develop the VIRF analysis along several dimensions. First, we compute the average responses in power variance for each year. Second, we examine the responses for different horizons, ranging from one day ahead to three calendar months ahead. This allows us to measure not only the strength of the spillover effects, but also the speed of their decay. Lastly, we compare the size of the cross-market effects to the size of the own-market variance responses.

The results show that positive news in gas and coal (i.e., news corresponding to an unexpected price increase) tends to generate significant volatility spillovers to power. On the other hand, negative

news (i.e., a price decrease) has only weak effects. For carbon, the story is different. Here, neither positive nor negative news generated any spillovers to the variance of power during the first three years of our sample period, 2008 – 2010. However, during 2011 – 2013, we find that both large upward and downward moves in the carbon market lead to significant responses in power market volatility.

Overall, there is a lot of variation in the strength of the spillover effects, both over time and across markets. We find a day-ahead increase in the expected power variance following the positive news in coal to range between 12% and 25% on average for different years, with a few extremes of up to 50%. The response to the positive news in gas is generally weaker, ranging from 3% in 2012 to 16% in 2010 on average, except for 2013, when we do not find any positive spillover effect. These results are consistent with the developments in the underlying markets. For example, we observe that the role of gas in the German power generation mix declines, and that spark spreads indicating the profitability of gas-fired power plants enter the negative region from the beginning of 2013. In contrast, coal remains an important fuel for power generation, and dark spreads stay positive throughout the sample period. Despite lower magnitude spillovers from gas, we find that it takes longer on average for the effect of news in gas to die out compared to the effect of news in coal.

Comparing the responses that exogenous news generates in the power variance to the responses in the own-market variance reveals that the cross-market effects are far from trivial, and taking them into account in practical applications can clarify a bigger picture. Distinguishing between the effects of price increases and price decreases in the asymmetric setting provides further valuable insights.

The remainder of this chapter is organized in three sections. Section 2 presents the VIRF methodology. We address the properties of the function in detail, discuss its interpretation, and explain how we

compute it. Section 3 contains the empirical results, while Section 4 concludes. The analysis in this study is based on the VAR-BEKK model estimated previously using the sample of daily returns on the front-year futures contracts and presented in Chapter 2.

2 Volatility impulse response function

This section presents the methodology that we use to analyse volatility transmission in the German energy markets. We start by providing a general background of the impulse response analysis, followed by the definition of the VIRF. Next, we describe how we construct the independent shocks and how they propagate through the system. We address the properties of the VIRF, which are important for interpreting the results. Finally, we explain the numerical computation procedure.

2.1 Background

In the general context, an impulse response function measures the time profile of the effect of a shock on the behaviour of a series. Sims (1980) introduced the concept of an impulse response for linear systems, such as vector autoregressive (VAR). The impulse response function was designed to provide an answer to the question of “What is the effect of a shock of size δ hitting the system at time t on the state of the system at time $t + n$, given the information available at $t - 1$?” It was traditionally assumed that no other shocks occurred between t and $t + n$, i.e., the future shocks were ‘switched off’. As impulse response functions were applied to non-linear systems (see, e.g., Gallant et al., 1993), the traditional approach of choosing a zero-shock baseline led to several conceptual problems, even in univariate settings. Potter (1994) highlights that non-linear models produce impulse responses that are history- and shock-dependent, which means

that the traditional impulse response function depends on t and δ , in addition to n . A further problem of composition dependence was discovered in multivariate systems, where it was impossible to disentangle the effect of a shock to a single variable from the effects of the contemporaneous shocks to other variables.

To address these limitations, Koop et al. (1996) develop the concept of a Generalized Impulse Response Function (GI). The GI is treated as a random variable defined on the underlying probability space of the time series under consideration. For the case of an arbitrary current shock ε_t and history ω_{t-1} , the GI for each horizon $n = 0, 1, \dots$ is given by the following difference:

$$GI(n, \varepsilon_t, \omega_{t-1}) = E(Y_{t+n} | \varepsilon_t, \omega_{t-1}) - E(Y_{t+n} | \omega_{t-1}), \quad (3.1)$$

where Y is the $k \times 1$ conditional mean vector to forecast. Taking the expectations solves the problem of treating future shocks since their effects are averaged out. The first expectation represents the state of the system conditioning on the history and the current shock, while the second expectation conditions on the history only and serves as the baseline. Both the current shock and history are treated as realizations from the stochastic data generation process (DGP) of Y , which means that both expectations in Eq. (3.1) are realizations of random variables themselves.

Hafner and Herwartz (2006) extend Koop et al.'s (1996) methodology to multivariate GARCH models. Rather than looking at the effect of a shock on the conditional mean process, they consider the effect on the conditional covariance, and put forth the concept of a Volatility Impulse Response Function (VIRF).¹ VIRF is used to

¹Lin (1997) provides an alternative methodology that applies Gallant et al.'s (1993) approach to multivariate GARCH models. The impulse response function suggested in Lin (1997) is defined as the impact of a small perturbation in a historic innovation on the future predicted volatility (i.e., a derivative w.r.t.

quantify the impact of independent shocks on volatility.

2.2 Definition of VIRF and its properties

In this study, we define a VIRF at time t and horizon n following Hafner and Herwartz (2006):

$$VIRF(t, n, \omega_{t-1}, \varepsilon_t) = E[vech(H_{t+n}) | \omega_{t-1}, \varepsilon_t] - E[vech(H_{t+n}) | \omega_{t-1}], \quad (3.2)$$

where H_{t+n} is the $k \times k$ conditional covariance matrix at time $t+n$, and $vech()$ is the operator that stacks the lower triangular fraction of a $k \times k$ matrix into a $k^* = k(k+1)/2$ dimensional vector. Similarly to the GI in Eq. (3.1), the first expectation in Eq. (3.2) is conditioned on the history ω_{t-1} and the current shock to the mean system ε_t ,² while the second expectation is conditioned only on the history and provides the baseline. The natural choice of the history is $\omega_{t-1} = \{\varepsilon_{t-1}, H_{t-1}\}$. Since H_{t-1} depends on all previous innovations, the history becomes the set of all past innovations: $\omega_{t-1} = \{\varepsilon_{t-1}, \varepsilon_{t-2}, \dots, \varepsilon_1\}$. Thus, the VIRF measures the response to the shock ε_t in variances and covariances by comparing their expected values at time $t+n$, with and without specific information about the shock. This response is the average of what might happen given the past and the present. In our case of four commodities, the VIRF as given by Eq. (3.2) is a 10×1 vector, with four entries representing the responses in variances and six entries representing the responses in covariances.

When interpreting the results of the VIRF analysis, it is important to understand the following properties, which distinguish VIRF from the traditional impulse response functions in linear systems. First, in symmetric GARCH models, VIRF is an even function, i.e., the effect

historic innovation), and is fundamentally different from the approach in Hafner and Herwartz (2006).

²See Eq. (2.1) and Eq. (2.2) for the conditional mean system.

of a positive shock is the same as the effect of a negative shock of the same size. In this study, however, we apply the VIRF analysis to an asymmetric BEKK model, and thus are able to distinguish between the effects of positive and negative shocks. Secondly, VIRF is not a homogeneous function of any degree. However, it is a monotonic function: a larger shock generates a larger response, all else being equal. Finally, VIRF depends on the history of the process through the state of the covariance matrix when an independent shock occurs, which is explained in detail in the following subsection.

2.3 Construction of independent shocks

Naturally, we can choose both ε_t and ω_{t-1} from the estimated conditional covariance model, leading to an analysis of the impact of a historical shock given the observed volatilities and correlations on the date of the shock event. However, we can obtain a bigger picture by constructing the random shocks from the distribution of the underlying time series, while still conditioning on the observed covariance matrix. This allows us to find the potential effect of a large exogenous shock on every day during the sample period, rather than only looking at those few days when some abnormal events actually occurred.

The next question is how to produce the independent shocks that hit the system and how to perturb it such that it is possible to quantify the impact of a particular variable-specific shock. Here, we use the spectral decomposition of the conditional covariance matrix. First, we assume that one of the commodity markets is hit by an exogenous shock, which we refer to as ‘news’. Accordingly, the news vector z

has a non-zero element in one position and zero elements elsewhere:

$$\begin{aligned} z &= (z_1, 0, 0, 0)' \\ z &= (0, z_2, 0, 0)' \\ z &= (0, 0, z_3, 0)' \\ z &= (0, 0, 0, z_4)', \end{aligned}$$

where z_i denotes news in the return of commodity $i = 1, \dots, 4$; gas, power, coal, and carbon. In the spirit of Koop et al. (1996) and Hafner and Herwartz (2006), we treat z as a random variable, and its realizations are drawn from the DGP of the underlying time series. In Chapter 2, we define the multivariate skew-Student distribution with independent components of Bauwens and Laurent (2002, 2005) for the vector of innovations. Therefore, we can draw the news z_i from the marginal densities using the estimated skewness and degrees of freedom parameters.³ We consider the 99% and the 1% quantiles in the news distributions, which we refer to as positive news and negative news, respectively.⁴ Therefore, the markets, on average, experience news of this magnitude two or three trading days per year, both for positive and negative news. Positive news corresponds to unexpected price increases, while negative news can be thought of as unexpected price decreases. Note that the news vectors are time invariant because the marginal densities are standardized w.r.t. the mean and variance, and have constant higher moments.

³See Eq. (2.6) for the joint density function, obtained as the product of independent marginal densities.

⁴The model-implied 99% quantiles (positive news) are $z_1 = 3.191$, $z_2 = 3.169$, $z_3 = 2.968$, and $z_4 = 2.908$. The 1% quantiles (negative news) are $z_1 = -2.638$, $z_2 = -2.706$, $z_3 = -3.013$, and $z_4 = -3.062$. The 99% quantile for a given commodity is numerically different from the 1% quantile due to the estimated non-zero skewness. Gas and power are right-skewed and, therefore, the 99% quantile is larger in absolute terms than the 1% quantile. For coal and carbon, the opposite is true.

The news originates in *one* of the commodities, and then propagates through the non-linear BEKK system to generate a system-wide shock ε_t :

$$\varepsilon_t = H_t^{1/2} z = (\varepsilon_{1t}, \varepsilon_{2t}, \varepsilon_{3t}, \varepsilon_{4t})', \quad (3.3)$$

where $H_t^{1/2}$ is the ‘square root’ of the covariance matrix, calculated using the spectral decomposition. In contrast to the news, which is time invariant, the system-wide shock induced by the news depends on the state of the system at a particular point in time. In general, all elements in ε_t are time-varying and non-zero due to the dependence structure imposed by the time-varying and non-diagonal covariance matrix H_t .

2.3.1 An illustration

Consider the following simplified example providing the intuition behind our news propagation mechanism, and the implications for the VIRF analysis. Let us assume we have only two commodities: gas (1) and power (2). The gas market is hit by an exogenous shock, leading to an unexpected increase in the gas price. Thus, the news vector z is $(3.19, 0)'$, where the first element is the 99% quantile from the assumed distribution for the gas return series. Next, assume that at a given point in time, the conditional covariance matrix is

$$H = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix} = \begin{pmatrix} 4 & 1 \\ 1 & 4 \end{pmatrix}.$$

This corresponds to the correlation of 0.25. Applying the spectral decomposition of H and using Eq. (3.3) gives the corresponding system-wide shock $\varepsilon_1 = (6.33, 0.80)'$.

Consider next that at some other point in time, gas volatility is higher, while the correlation between power and gas remains the same,

and the conditional covariance matrix is

$$H = \begin{pmatrix} 9 & 1.5 \\ 1.5 & 4 \end{pmatrix}.$$

In this case, the corresponding shock is $\varepsilon_2 = (9.52, 0.96)'$.

Finally, consider the third case where both power and gas volatility remain the same, but the correlation between the two series increased from 0.25 to 0.50:

$$H = \begin{pmatrix} 4 & 2 \\ 2 & 4 \end{pmatrix}.$$

This yields $\varepsilon_3 = (6.16, 1.65)'$. Because the covariance matrix is different in the three cases above, the expected covariance matrix at the next point in time is also different, both with the information about the news z and without it. Recall that in the asymmetric BEKK(1,1) model with k variables, the conditional covariance matrix at time t is defined as follows:

$$H_t = C'C + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + B'H_{t-1}B + D'\zeta_{t-1}\zeta'_{t-1}D,$$

where A , B , C , and D are $k \times k$ matrices, ε_{t-1} is the $k \times 1$ vector of error terms, and ζ_{t-1} is a $k \times 1$ vector of asymmetric error terms.

To investigate the behaviour of the VIRF in this illustration, we use the following parameters:⁵

$$A = \begin{pmatrix} 0.19 & 0.03 \\ 0.03 & 0.17 \end{pmatrix}; B = \begin{pmatrix} 0.98 & 0.01 \\ -0.02 & 0.96 \end{pmatrix}; D = \begin{pmatrix} 0.06 & -0.09 \\ -0.01 & 0.05 \end{pmatrix}.$$

We can now compute both expectations in Eq. (3.2) and evaluate

⁵Here we are using a subset of the actual estimated parameters presented in Table 2.7. C -matrix is set to a zero-matrix for simplicity.

the VIRF numerically. In this example, we consider the day-ahead VIRF only, that is, the horizon n is equal to 1. The exact procedure is described in the next subsection, and the result for each of the three cases is a 3×1 vector with elements measuring the impact of the independent news z on gas variance, covariance between gas and power, and power variance. We investigate the impact on power variance only, and the results are as follows:

$$\begin{aligned} E(\sigma_2^2 | \varepsilon_1, \omega_1) - E(\sigma_2^2 | \omega_1) &= 4.14 - 3.91 = 0.23, \\ E(\sigma_2^2 | \varepsilon_2, \omega_2) - E(\sigma_2^2 | \omega_2) &= 4.65 - 3.93 = 0.72, \\ E(\sigma_2^2 | \varepsilon_3, \omega_3) - E(\sigma_2^2 | \omega_3) &= 4.25 - 3.94 = 0.31. \end{aligned}$$

To facilitate comparability, we suggest normalizing the obtained responses with respect to the baseline expectations, thus obtaining the relative measures. In particular, we find that in the first case, the expected power variance is 5.88% larger conditioning on the shock ε_1 and the history ω_1 than conditioning on the history only. The percentage responses for the second and third cases are 18.32% and 7.87%, respectively. We can see that both higher gas variance (the second case), and higher gas/power correlation (the third case), amplify the effect of the news in gas z on the response in power variance. These results are in line with the intuition that larger volatility spillover effects occur at more turbulent times and between highly correlated markets.

2.4 VIRF computation via Monte Carlo integration

Hafner and Herwartz (2006) derive analytical VIRF expressions for the general vec representation of symmetric multivariate GARCH models. However, we work with an asymmetric BEKK model that lacks analytical expressions for the expectations in Eq. (3.2). Therefore, we perform the integrations numerically using the Monte Carlo

techniques outlined in Koop et al. (1996).

To calculate the VIRFs, we proceed in two steps:

1. Construct shocks ε_t for each trading day t corresponding to positive and negative news z_i for each commodity. This gives two time series of shocks for each commodity (in total, eight time series of shocks).

2. Calculate a VIRF for each trading day t for each time series of shocks and for horizons from $n = 1$ to $n = 63$ (i.e., approximately three calendar months ahead).

In step two, we need to integrate out the effects of the future shocks from $t + 1$ to $t + n$. To perform the numerical integration, we simulate future random realizations of the vector of innovations. In our case, each ‘future’ consists of 63 realizations. We need 62 random shocks in order to iterate the BEKK process forward up to $t + 63$ in the first expectation (for a given t , ω_{t-1} , and ε_t), and an extra random shock for the second expectation, where we condition on the history only. By using common random shocks to evaluate both expectations, we ensure that the difference between them in a single future scenario is driven by the different conditions at time t only. Further, using common random numbers is a variance reduction technique in stochastic simulation.

We draw the random shocks from the underlying standardized marginal skew-Student densities using the following analytical quantile function given in Laurent (2002):

$$F^{-1}(p | \xi) = \begin{cases} \frac{\xi^{-1} G^{-1}(\frac{p}{s}(1+\xi^2)) - m}{s} & \text{if } p < \frac{1}{1+\xi^2} \\ \frac{-\xi G^{-1}(\frac{1-p}{s}(1+\xi^{-2})) - m}{s} & \text{if } p \geq \frac{1}{1+\xi^2} \end{cases}, \quad (3.4)$$

where ξ is the skewness parameter, G^{-1} is the quantile function of the original symmetric Student density (G being the cumulative distribution function, accordingly), and m and s are the mean and the stan-

dard deviation of the non-standardized skew-Student density given by Eq. (2.10) and Eq. (2.11), respectively.

For each of the eight time series of shocks constructed in the first step, we create 20,000 possible ‘futures’, over which we take the averages.⁶ By the Law of Large Numbers, these averages converge to the conditional expectations required to calculate the VIRF. Each of the 20,000 realizations of the right-hand side of Eq. (3.2) is conditioned on the same shock ε_t and the same history ω_{t-1} , but corresponds to a different future.

Finally, we normalize the variance responses for each commodity with the corresponding initial baseline expected variance (i.e., the baseline expectation at $t + 1$). This normalization creates a relative response in the variances comparable over time and across commodities.

3 Results

We focus our empirical analysis on the transmission of news in gas, coal, and carbon to the variance of power, and therefore what we refer in this section as a VIRF is actually a single element in the VIRF vector measuring the response in the variance of power.

We present two sets of complementary results based on the calculations outlined above. First, we show the average variance response in power for each year, 2008 to 2013, for all horizons. We calculate the yearly VIRFs by taking the average of the daily VIRFs over all trading days each year. This yearly analysis gives an overall dynamic picture of the strength of volatility spillovers between the different

⁶We calculate a VIRF for each trading day from January 7, 2008, to October 4, 2013. We cannot iterate the BEKK model outside the sample because we need the actual covariance matrix to construct the shocks (futures) given the simulated news. Therefore, for 2013, we calculate VIRFs up to and including October 4, when there are 63 trading days remaining in the sample.

commodities. It also provides us with insights into the speed of decay in the responses when the horizon increases. Second, we show the day-ahead responses in power variance for each trading day in the sample. This daily time series analysis gives a different perspective of the time variation in the news transmission, which allows us to trace changes in the volatility spillovers over time in more detail, and to connect these changes to specific market events.

3.1 Volatility transmission patterns: The yearly perspective

We plot the yearly VIRFs in Figures 3.1 – 3.3. These plots provide us with an insight into the nature of the spillover effects from three perspectives. First, we can track changes over time. Second, since the normalized VIRFs are expressed as percentages, we can compare the strength of the news transmission between different commodities in a straightforward way. It is also possible to determine the half-life of a shock and assess how fast the system recovers after an abnormal event. Finally, we differentiate between unexpected price increases (positive news) and unexpected price decreases (negative news).

Consistent with the estimation results from Chapter 2, Figures 3.1 – 3.3 suggest that coal has the largest effect on power volatility throughout the sample period. For example, in 2008, positive news in the coal market resulted in a 25% larger than otherwise expected day-ahead variance on average in the power market compared to the corresponding effect of less than 5% for positive news in the gas market. It is worth noting that the estimated volatility of gas and coal is very similar overall during the sample period (see Figure 2.4).




Figure 3.1: Response in power variance to news in the gas market.




Figure 3.2: Response in power variance to news in the coal market.




Figure 3.3: Response in power variance to news in the carbon market.

Further, we can see in Figure 3.1 that positive news in gas leads to much lower responses in the power variance towards the end of the sample period, with the 2013 VIRF having a strikingly different pattern. Technically, a negative response means that the expected power variance, given the news in gas, is lower than what we would otherwise expect it to be. Economically, this means that an abnormal event in the gas market does not lead us to revise our expectations about future volatility in the power market upwards. We discussed previously that spark spreads were below zero from the beginning of 2013, and that the correlation between gas and power declined sharply (see Figures 2.2 and 2.5). A negative gas-to-power VIRF in 2013 is another manifestation of the weakening link between the two markets.

We find news in gas to have a longer-lasting effect, on average, than news in coal. It takes approximately 14 trading days for a power VIRF to decrease by half after news in gas, compared with nine trad-

ing days after news in coal. News in carbon has the shortest effect on expected power variance, with an average half-life of four trading days.

3.2 Day-ahead variance responses: The daily perspective

Figure 3.4 displays the evolution of the day-ahead responses in the power variance over the sample period. We can see that the spillovers to the power variance resulting from news in all other commodities show considerable variation over time.




Figure 3.4: Time series of day-ahead variance responses in the power market.

Econometrically, these varying volatility responses are a consequence of the complex dynamic interrelations between the different

commodities in the BEKK system. As illustrated earlier on a simplified example, one particular aspect of this is the time-varying correlations. Low correlations lead to weaker shock transmission, all else being equal.

Interestingly, the response in power variance is much larger after positive news than after negative news in both gas and coal throughout the sample period. This result makes sense because increases in fuel prices have a negative effect on producers' profit margins. For carbon, the story is different. There are no large differences in the responses to positive and negative news.⁷ We see no significant difference during 2008 – 2011, while during 2012 – 2013, negative news has a slight tendency to produce higher VIRFs. However, it is only during the latter part of the sample period that the carbon news leads to spillovers of any significant magnitude with a large daily variation. This increase in spillovers from carbon roughly coincides with a simultaneous, and significant, increase in the volatility of carbon (see Figure 2.4). Recall from an illustration in Subsection 2.3.1 that this phenomenon is consistent with our model and how we specify the covariance matrix. Thus, while the increase in carbon volatility may be the cause of the increase in VIRFs seen in 2012 – 2013, the underlying economic reasons for the patterns we see in response to positive versus negative carbon news are much harder to identify.

As the lower-right panel in Figure 3.4 indicates, the day-ahead response in power variance to own shocks is, on average, around 9% for negative news and around 14% for positive news on average. We can see that responses in power variance to news in other energy markets are more volatile than responses to own news, yet at certain times much stronger.

⁷Note also the lower absolute value of the spillover coefficient $d_{42} = -0.0169$, representing the asymmetric spillover effect from carbon to power compared to $d_{12} = -0.0869$ for gas and $d_{32} = 0.1378$ for coal (see Table 2.7).

Investigating how gas, coal, and carbon markets react to own shocks can provide an additional feel for the strength of the cross-market effects. Figure 3.5 displays the evolution of the day-ahead responses to own-market shocks. We can see that positive news originating in the coal market leads, on average, to a 20% higher than otherwise expected coal variance. Although the corresponding response in power variance (the upper-right panel in Figure 3.4) shows high variation, it is above 20% during many periods. This reveals that volatility spillover effects in the energy markets are not trivial. We believe that the benefits in terms of the explanatory and predictive power of taking them into account in practical applications outweigh the costs of increased model complexity.




Figure 3.5: Time series of own-market day-ahead responses in variance.

3.3 Sector-wide news and critical news

We assumed previously that one of the markets is hit by exogenous price news, and investigated the spillover effects to the power market. While the main focus of this study is on comparing the strength and dynamics of idiosyncratic shock transmission from fuels and carbon to power, we also consider the case when all four commodities experience a tail event simultaneously. This case, which we refer to as ‘sector-wide news’, corresponds to the news vector z having all elements equal to the 99% or the 1% quantiles in the corresponding news distributions. We think of it as price news to the entire energy sector, as opposed to news specific to a certain market, such as an effect of a policy announcement in the carbon market, for instance.

Figure 3.6 displays the yearly power VIRFs resulting from concurrent positive and negative news in all commodities.




Figure 3.6: Response in power variance to sector-wide news.

We can see that these VIRFs have less variation on a yearly basis, especially at the horizon of up to five days, compared to the ones in Figures 3.1 – 3.2, although we still observe weaker responses in 2013. Sector-wide price increases have a greater impact on expected power variance than sector-wide price decreases. If a 99% quantile news event occurs simultaneously in all the four markets, we can expect the day-ahead variance in the power market to increase by 180% – 210%, on average for different years. We plot the corresponding daily VIRFs in Figure 3.7.




Figure 3.7: Time series of day-ahead responses in power variance to sector-wide news.

As Figure 3.7 indicates, the day-ahead variance responses are generally weaker in the second half of the sample period. This result is partly due to decrease in volatilities in all markets except carbon (see Figure 2.4), and partly due to decrease in correlations, as discussed

previously.

We can gain another perspective by calculating the ‘critical news’ at time t required to generate a power variance response of zero at time $t+1$. How large does news in gas, coal, and carbon have to be to obtain a VIRF of zero? Recall that a VIRF of zero means a variance response in power that is equal to the baseline expected power variance. These critical news levels have to be calculated numerically, and we do so for both positive (right-tail) and negative (left-tail) news. We report the results as critical quantiles rather than actual news values because this gives a more direct sense of the magnitude of news required to generate a VIRF of zero. These quantile values can also be related to the 99% and 1% quantiles we used in our previous VIRF calculations. We report the results in Table 3.1, where we list the positive and negative quantiles (as percentages) for all commodities and for each year in our sample.

Table 3.1: Critical quantiles (%) in news distribution.

	Positive News			Negative News		
	Gas	Coal	Carbon	Gas	Coal	Carbon
2008	98.12	92.49	99.42	0.56	1.54	0.72
2009	96.05	95.24	99.42	1.17	0.91	0.77
2010	94.73	94.90	99.57	1.35	0.47	0.56
2011	97.43	95.63	98.36	0.56	0.36	2.46
2012	98.32	94.40	98.15	0.25	0.38	3.17
2013	99.78	95.35	96.22	0.01	0.28	6.32

Note: The table shows critical quantiles in the news distribution for different markets. News of a lower magnitude does not lead to a higher than otherwise expected day-ahead variance in the power market.

We observe that the news that makes the VIRF at $t + 1$ equal to zero varies across commodities and over time. For example, gas in 2013 required large positive and negative news (at the 99.78% quantile in the right tail and the 0.01% quantile in the left tail). This is a

manifestation of a very weak link between the power and gas markets at that time. On the other hand, for carbon in 2013, we need only news at the 95.35% level in the right tail and 6.32% in the left tail.

4 Summary and conclusions

In this study, we investigate the transmission of news and volatility spillovers between electrical power, gas, coal, and carbon in the German energy market during 2008 – 2013. We base our methodology on the VIRF concept introduced in Hafner and Herwartz (2006). However, the lack of analytical expressions for asymmetric GARCH models that we work with requires us to compute VIRFs numerically, following the earlier work of Koop et al. (1996) on impulse response functions in the conditional mean systems. The VIRF measures the impact of an independent shock on the conditional covariance matrix given the history.

We focus our analysis on the impact of exogenous large news in the gas, coal, and carbon markets on the expected volatility in the power market. We distinguish between the effects of positive news, corresponding to price increases, and negative news, representing price decreases. Further, we calculate the VIRF for horizons from one day to three calendar months following the news event. The results indicate that spillover effects show large variation across commodities and over time. The spillovers from coal are substantial throughout our sample period, but with significant time variation on a daily basis. The spillovers from gas are generally weaker and are declining toward the end of the sample period, which is consistent with gas becoming an unprofitable generation technology and playing a less important role in the German power generation mix. Despite the lower magnitude spillovers from gas, we find that it takes longer on average for the effect of news in gas to die out compared to the effect of news in

coal. Spillovers from carbon have the fastest decay and were more or less nonsignificant until 2011, when they started to show more variation and a larger impact. We find that positive news in gas and coal markets leads to a much larger response in the variance of power compared to negative news, generating only small or no spillover effects. Distinguishing between the positive and negative news appears to be less important for the carbon market.

The overall implication is that modelling the volatility in the power market, whether for explanatory or forecasting purposes, gains substantially from incorporating the information from the fundamentally related fossil fuel and carbon markets.

References

- [1] Bauwens, L., Laurent, S., 2002. A new class of multivariate skew densities, with applications to GARCH models. CORE Discussion Paper 20.
- [2] Bauwens, L., Laurent, S., 2005. A new class of multivariate skew densities, with application to Generalized Autoregressive Conditional Heteroscedasticity models. *Journal of Business and Economic Statistics* 23, 346–354.
- [3] Gallant, A.R., Rossi, P.E., Tauchen, G., 1993. Nonlinear dynamic structures. *Econometrica* 61, 871–907.
- [4] Hafner, C.M., Herwartz, H., 2006. Volatility impulse responses for multivariate GARCH models: An exchange rate illustration. *Journal of International Money and Finance* 25, 719–740.
- [5] Hasanov, A.S., Do, H.X., Shaiban, M.S., 2016. Fossil fuel price uncertainty and feedstock edible oil prices: Evidence from MGARCH-M and VIRF analysis. *Energy Economics* 57, 16–27.

- [6] Jin, X., An, X., 2016. Global financial crisis and emerging stock market contagion: A volatility impulse response function approach. *Research in International Business and Finance* 36, 179–195.
- [7] Koop, G.M., Pesaran, H.M., Potter, S.M., 1996. Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74, 119–147.
- [8] Laurent, S., 2002. Asymmetry and fat-tails in financial time series. Doctoral dissertation, Maastricht University.
- [9] Le Pen, Y., Sevi, B., 2010. Volatility transmissions and volatility impulse response functions in European energy forward markets. *Energy Economics* 32, 758–770.
- [10] Lin, W.-L., 1997. Impulse response function for conditional volatility in GARCH models. *Journal of Business & Economic Statistics* 15, 15–25.
- [11] Olson, E., Vivian, A.J., Wohar, M.E., 2014. The relationship between energy and equity markets: Evidence from volatility impulse response functions. *Energy Economics* 43, 297–305.
- [12] Potter, S., 1994. Nonlinear impulse response functions. Department of Economics Working Paper. University of California, Los Angeles, CA.
- [13] Sims, C., 1980. Macroeconomics and reality. *Econometrica* 48, 1–48.

CHAPTER 4

Joint Modelling of Power Price, Temperature, and Hydrological Balance with a View towards Scenario Analysis

Abstract

This study presents a model for the joint dynamics of power price, temperature, and hydrological balance, with a view towards scenario analysis. Temperature is a major demand-side factor affecting power prices, while hydrobalance is a major supply-side factor in power markets dominated by hydrological generation, such as the Nordic market. Our time series modelling approach coupled with the skew-Student distribution allows for interrelations in both mean and volatility, and accommodates most of the discovered empirical features, such as periodic patterns and long memory. We find that in the Nordic market, the relationship between temperature and power price is driven by the demand for heating, while the cooling effect during summer months does not exist. Hydrobalance, on the other hand, negatively affects power prices throughout the year. We demonstrate how the proposed model can be used to generate a variety of joint temperature/hydrobalance scenarios and analyse the implications for power price.

Keywords: spot power price, temperature, hydrological scenarios, VARFIMA-BEKK, skew-Student

JEL Classification: C32, G17, Q41

I wish to thank Stefan Schneider, Pierre Arzounian, and Carsten Teller at Uniper Global Commodities (Düsseldorf) for providing the data and valuable feedback on this project.

1 Introduction

In this study, we develop a model for the joint evolution of the spot electrical power price, outdoor temperature, and hydrological balance. The model is relevant for power markets with a large share of hydrological generation, such as the Nordic market, and offers a wide range of opportunities for scenario analysis.

Consider the following example illustrating the fundamental relationship between power price, temperature, and hydrobalance. Energy producers and retailers plan their business activities based on the estimated demand for power (load) for a certain time horizon. During the heating seasons, the load is driven to a large extent by temperature. If a particular season is actually colder than expected, more power than planned will be consumed for heating purposes. Temperature affects demand, but to understand the implications for prices, we also need to consider the supply side of price formation. We define hydrobalance as the measure of the potential capacity of a hydrological power generation system. If a year has been relatively wet, with lots of precipitation, i.e., hydrobalance is high, this excess demand may be covered at a low cost without moving the price. On the contrary, a combination of low temperature and low hydrobalance is a major source of price risk in power markets. Therefore, it is natural to model these three variables as a system.

Both temperature modelling and hydrological modelling are large research areas on their own. Among the literature related to power markets, Halldin (2005) studies modelling of the time series of water inflows and stochastic optimization of a hydro-power system. Green (2015) shows that the intra-daily profiles of the Nord Pool system price are affected by hydrological balance and develops an hourly forward curve model with hydrological dependence. Bivariate power-temperature models have been developed in Benth et al. (2012) and

Caporin et al. (2012), for the purpose of pricing an exotic type of weather derivatives called energy quanto options. This study adopts an econometric modelling approach, similar in certain respects to that of Caporin et al. (2012), introducing the third dimension of hydrobalance into the system.

We analyse both the univariate properties of our three data series, and the dependencies between them in detail. The model is identified within the Vector Autoregressive Fractionally Integrated Moving Average (VARFIMA) framework, coupled with a time-varying covariance process of Baba-Engle-Kraft-Kroner (BEKK) type. The need for fractional integration is motivated by long memory, observed in all series. In addition to the stochastic part, our model contains a deterministic component capturing the yearly periodic patterns in power prices and temperatures. Due to the highly pronounced non-normal statistical properties of our data, we apply a flexible multivariate skew-Student distribution proposed in Bauwens and Laurent (2002, 2005), while treating the normal distribution as the benchmark for comparison. The model allows for interrelations, both in means and volatilities, restricted such that power price can be affected by temperature and hydrobalance, but not the other way around.

We find that in the Nordic market, the relationship between temperature and power price is driven by the demand for heating, while the cooling effect during summer months does not exist, likely due to mild temperature conditions. Hydrobalance, on the other hand, has a significant inverse effect on power prices throughout the year. Further, estimation results indicate the existence of volatility spillover effects from hydrobalance and temperature to power. Correlations between power and temperature show seasonal patterns, ranging from -0.5 during winter periods to 0 during summer periods. Correlations between power and hydrobalance oscillate around an average level of -0.25 . Finally, the simulation exercise reveals the benefits of

skew-Student distribution in reproducing the distinct non-Gaussian properties of both power price and meteorological series.

The remainder of the chapter is organized as follows. Section 2 describes the data and the results of preliminary data analysis. Section 3 presents the modelling framework, as well as the identification and estimation procedure. The empirical results are discussed in Section 4. Section 5 addresses the simulation from the model and provides a scenario analysis example, which demonstrates how our model can be utilized to generate a variety of joint temperature/hydrobalance scenarios and analysing the implications for power prices. Section 6 contains a summary and concluding remarks.

2 Data and preliminary analysis

This section describes the data and the results of preliminary data analysis, which lays the foundation for our choice of modelling framework. We first investigate the properties of the univariate series that we want the model to reproduce, and then discuss the desired dependence structure.

2.1 The dataset

Our dataset consists of daily observations of the power price series, the temperature series, and the hydrobalance series. Since the model is designed specifically for power markets with a large share of hydrological generation, we use Nordic market data.¹ The sample period spans from January 1, 2008, to February 21, 2016.

We obtain the power price data from the Nordic power exchange,

¹Hydro power is the largest generation source in terms of installed capacity in the Nordic power market. According to the Nord Pool, in a year with normal precipitation, hydro power accounts for half of Nordic countries' demand (98% in Norway, in particular).

the Nord Pool. The Nord Pool Elspot is the spot market where agents trade power on an auction basis for physical delivery during each hour of the following day, which is why it is often referred to as the day-ahead market. The daily spot power price is the arithmetic average across 24 hourly prices, and is quoted in EUR/MWh. Being the common marketplace for the Nordic (Denmark, Finland, Norway, and Sweden) and the Baltic (Estonia, Latvia, and Lithuania) countries, Nord Pool is divided into a number of bidding areas, which can have different prices in the presence of transmission constraints. In addition to these area prices, all participating countries share a common *system spot price*, calculated under the assumption of unconstrained transmission capacity. In this study, we consider the Nord Pool system spot power price since it is the reference price for trading and clearing the majority of financial contracts.

The temperature and hydrobalance data are obtained from Thomson Reuters. We use the daily average temperature (DAT) in Sweden, which is the population-weighted average across a basket of several cities. Alternatively, we could consider the average temperature across all Nord Pool area countries, but since the Baltic countries joined the market during 2010 – 2013, it is more straightforward to use a single country as a proxy for the whole region. This does not lead to any loss of generality because temperature series in the individual Nordic and Baltic countries are highly correlated, and Sweden would have had the largest population weight in the index anyway. The daily average temperature is the average of the minimum and the maximum temperature during a given day measured in degrees of Celsius.

Finally, the hydrobalance series represents the deviations of the total hydrological resources from the seasonal normal level measured in terms of energy capacity (TWh). The total hydrological resources are defined as the sum of the water reservoir content, the snow pack,

and soil water, and reflect the available capacity of the hydro-power system. The seasonal normal levels of the hydrological resources for each day of a year are computed by Thomson Reuters based on the 1981 – 2005 period, and thus account for recent weather trends. Consequently, a positive (negative) hydrobalance value on a given day indicates that the hydrological condition is wetter (drier) than it has been on average for the same day of a year during 1981 – 2005. Most hydrological data are usually available at a weekly granularity, but Thomson Reuters provides daily Nordic hydrobalance series starting from 2008. This series contains more fundamental information than if we were to interpolate between the weekly observations, which motivates our choice of the sample period start.

2.2 Data analysis

Figure 4.1 plots the power price series along with the fitted seasonal mean function given by Eq. (4.1) below.




Figure 4.1: Time series of Nord Pool power system spot prices.

We can clearly see a decreasing trend, as an increasing amount of low marginal cost wind, solar, and biomass generation displaces coal and gas technologies. Another important feature is yearly seasonal patterns with higher prices during winter periods related to demand for heating and shorter day-light length. Seasonal patterns can be captured by a periodic function, such as sine or cosine. In addition, day-of-the-week effects in power prices have been widely documented, with lower prices observed on weekends and holidays due to limited business activity (see, e.g., Lucia and Schwartz, 2002). Therefore, we choose the following specification for power seasonal mean at time t , denoted $\Lambda_{1,t}$:

$$\Lambda_{1,t} = \lambda_{1,1} + \lambda_{1,2} \cos\left(\frac{2\pi}{365}(t - \lambda_{1,3})\right) + \lambda_{1,4}t + \lambda_{1,5}D_t, \quad (4.1)$$

where $\lambda_{1,i}$, $i = 1, \dots, 5$, are the parameters to estimate, t is time measured in days, and D_t is a dummy variable taking a value of 1 if day t is a non-business day (i.e., a weekend or a holiday), and 0 otherwise. $\lambda_{1,1}$ represents the overall (non-seasonal) average price level, $\lambda_{1,2}$ is the amplitude of the mean price, and $\lambda_{1,3}$ is the phase angle. The amplitude of a cosine wave reflects how large the distance between peaks and troughs is. A phase angle shifts the time to adjust for the fact that yearly maximum and minimum mean prices do not necessarily have to occur on January 1 and July 1, respectively. Note that the period of oscillation is equal to one year, or 365 days, ignoring leap years.

A closer look at Figure 4.1 reveals that power prices can have large upward spikes, followed by fast mean reversion. Spikes typically occur during winter seasons if high demand for power coincides with an unexpected supply-side shock, like an outage at a major power plant. It is worth noting that spikes are relatively less dramatic in power markets with a large share of hydro generation, such as the

Nordic market, since water reservoirs can serve as a safety buffer against unforeseen imbalances between supply and demand. This, of course, heavily depends on hydrological conditions and whether there is enough excess capacity in the hydro-power system to provide this sort of safety cushion.

The temperature series is plotted in Figure 4.2, along with the fitted seasonal mean function given by Eq. (4.2) below.




Figure 4.2: Time series of daily average temperature in Sweden.

We adopt the following specification for temperature seasonal mean at time t , denoted $\Lambda_{2,t}$:

$$\Lambda_{2,t} = \lambda_{2,1} + \lambda_{2,2} \cos\left(\frac{2\pi}{365}(t - \lambda_{2,3})\right), \quad (4.2)$$

where $\lambda_{2,1}$ is the overall (non-seasonal) average level of the series, $\lambda_{2,2}$ is the amplitude of the mean temperature, and $\lambda_{2,3}$ is the phase

angle. Alaton et al. (2002) document a small yet statistically significant increasing linear trend in Stockholm temperatures during 1957 – 1997, which is attributed to global warming and urbanization. We do not find such an effect in our sample, probably since we are looking at a much shorter time span. Since our model is not meant for multi-year forecasts, we omitted the linear trend term. Benth and Šaltytė-Benth (2005) reach the same conclusion while examining the Norwegian temperature data during 1990 – 2003.

We form de-seasonalised power price and de-seasonalised temperature series by subtracting from the original series Λ_1 and Λ_2 , respectively. Further analysis in this section is concerned with the properties of de-seasonalised series.

Finally, we examine hydrobalance. Note that hydrobalance is a series of deviations from the normal state, and therefore should have a long-run mean level of zero, assuming that the chosen normal state is representative of recent dynamics, or stable over long periods of time. However, when looking at shorter horizons, the sample mean can move away from zero, as is the case in our sample. As Figure 4.3 shows, dry periods are more common than wet periods during the sample.

We can also see how extremely persistent hydrological condition is: once a trend is established, it might take months for hydrobalance to revert back to zero. It is also fairly uncommon for the series to change signs within a single year; therefore, it makes sense to classify the whole years as ‘wet’ or ‘dry’. Nevertheless, it is by no means binary, and a wide range of possible scenarios are likely. The hydrobalance series in itself is not seasonal, although the total hydrological resources do exhibit strong yearly seasonal patterns, with water reservoirs being gradually filled after the spring flood and melting of the snow pack.

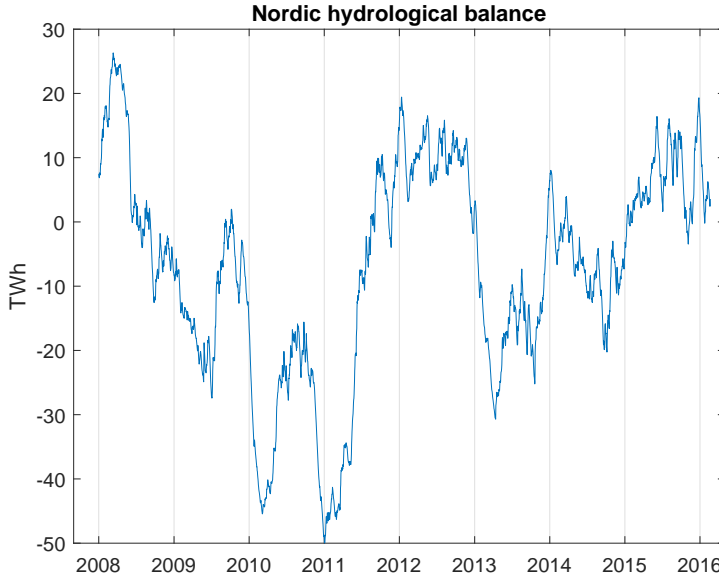


Figure 4.3: Time series of daily Nordic hydrological balance.

Table 4.1 contains the results of stationarity tests.

Table 4.1: Stationarity tests.

	Power price	Power log-price	Temperature	Hydrobalance
ADF	-5.33***	-6.48***	-14.35***	-2.27
KPSS	0.39*	0.33	0.36*	0.81***
Sample size	2974	2974	2974	2974

Note: The table reports the results of stationarity tests. ADF refers to the Augmented Dickey-Fuller test with the null hypothesis of a unit root. KPSS refers to Kwiatkowski et al.'s (1992) test with the null of a stationary $I(0)$ process. Stationarity tests for power price, log-price, and temperature were applied to de-seasonalised series. The number of lags in the ADF tests were selected based on Schwartz Information Criterion. KPSS tests used the Bartlett kernel with Newey-West automatic bandwidth selection. Superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

The presence of a unit root can be strongly rejected based on the Augmented Dickey-Fuller test for all series except hydrobalance.

Interestingly, the KPSS test with the opposite null hypothesis of stationarity can also be rejected for power price and temperature series, although at the 10% significance level only. Haldrup and Nielsen (2006) report similar results in conducting a wide range of stationarity tests on Nordic power prices, showing that neither $I(0)$ nor $I(1)$ processes seem to be appropriate. Let us get further insights into the time series properties of our data by examining Figure 4.4, which plots the sample autocorrelation functions (ACFs) for de-seasonalised power price, de-seasonalised temperature, and hydrobalance.




Figure 4.4: Sample autocorrelation functions.

We can see slight peaks in power price ACF at lags 7 and 14, which indicate the presence of weekly seasonal effects not captured by Λ_1 . Hydrobalance shows an extremely high degree of persistence in the autocorrelation function, typical for $I(1)$ processes. Both power price and temperature show a slow (hyperbolic) decay in autocorrela-

tions that cannot be captured by traditional Autoregressive Moving Average (ARMA) models. This type of behaviour is known as long memory, or long-range dependence. Saying that a given process has long memory means that the effect of a single shock is extremely persistent. However, unlike in the case of a random walk, persistence can be combined with mean-reversion in long-memory models. There is no reason to assume that any of the series in question could wander arbitrarily away from their fundamental mean levels, and therefore, a model that enforces mean-reversion while allowing high shock persistence appears to be a good choice.

Figure 4.5 illustrates the daily changes in de-seasonalised power log-price (i.e., log-returns), de-seasonalised temperature and hydrobalance.




Figure 4.5: Time series of daily changes.

We can see pronounced clustering effects in power volatility, and seasonal patterns in temperature volatility with peaks during winter

periods. Daily hydrobalance fluctuations have notable positive skewness, which suggests the need for a non-symmetric distribution. In addition, the graph reveals a frequent occurrence of extreme observations in power log-returns.

2.3 Relationship between power price, temperature, and hydrobalance

In the previous subsection, we investigated the properties of power prices, temperature, and hydrobalance separately. We now discuss how they are related, and the kind of dependence structure we want our model to impose.

Temperature is the main demand-side stochastic factor affecting power price. Temperature determines demand for heating during the winter periods and demand for cooling during the summer periods. Ebbeler, Benth and Kiesel (2014) find that the correlation between de-seasonalised temperature and German power spot price is negative in the winter months, and positive, though lower in magnitude, in the summer months. Seasonal effects of this kind can be accommodated in the linear model, either by allowing the temperature coefficient in the power mean equation to take different values during cooling and heating seasons, or by having a single temperature coefficient for the entire sample and an additional coefficient for heating seasons only. We test both specifications by regressing de-seasonalised power log-returns on de-seasonalised temperature and find that there is no extra heating season effect. In fact, temperature has no effect on power evolution during summer months in our sample, so one could argue that the entire effect is driven by demand for heating. When considering October – March periods only, the estimated coefficient is negative and very close to the coefficient for the whole sample period. Thus, we conclude that Nordic summer temperatures are too mild to generate any significant cooling demand effect in the power market,

and therefore, a single coefficient for each temperature lag in the power mean equation is sufficient.

Halldin (2005) discussed the inverse relationship between the Nordic power price and the hydro reservoir level in the context of the stochastic optimization of a hydro-thermal power system. We now compare the Nord Pool system prices under different hydrological conditions, but the same demand conditions. Here, we use the Nord Pool consumption data in addition to the data described previously. Figure 4.6 presents a scatterplot of power prices against power consumption, where we group observation pairs depending on the hydrobalance level. The red dots represent the lower quartile of hydrobalance, i.e., the driest 25% of days during the sample period. The blue crosses, on the other hand, mark the upper quartile of hydrobalance, i.e., the wettest 25% of days.




Figure 4.6: Power price versus power consumption scatterplot.

It is evident that given the same consumption, power price is gen-

erally higher under dry hydrological conditions. This is explained by the fundamentals of power price formation: as less hydro power is available, more generation technologies with higher marginal costs are utilised. In addition to a higher average price level, dry hydrological conditions lead to an increased probability of extreme prices. Additional regression analysis reveals that the relationship between power price and hydrobalance can be approximated sufficiently well by a single linear term, and there seems to be no motivation for including any non-linear effects.

To summarize, we would like the model to allow temperature and hydrobalance to influence the power price dynamics, but not the other way round. We expect a reasonable model to produce negative correlations between power and hydrobalance. Further, we expect the correlations between power and temperature to be negative during the heating seasons and close to zero otherwise.

3 The model

This section presents the model for the joint evolution of power price, temperature, and hydrobalance. We first describe the model in its general form, and motivate how it is expected to capture the numerous features discovered in the preliminary analysis. Further, we discuss the model identification and estimation methodology.

3.1 General framework

3.1.1 Conditional mean

Considering all of the theoretical and empirical arguments outlined in Section 2, we suggest a model of the mean evolution of the system within the Vector Autoregressive Fractionally Integrated Moving Average (VARFIMA) framework. Univariate ARFIMA processes are a

well-known class of long-memory models, introduced by Granger and Joyeux (1980) and Hosking (1981) as a generalization of the traditional ARIMA (p, d, q) model, which allows the difference parameter d to take fractional values. Beran (1994) and Palma (2007) discuss the statistical properties and inference for long-memory processes in detail. Prior to becoming a subject of interest for econometricians, the fractional behaviour of certain time series was extensively studied in hydrology and climatology. The first published papers describing how to test and model the long-range dependence date back to the 1950s, and are concerned with modelling the inflows of the river Nile (see Hurst, 1951).² Brody et al. (2002) document fractional behaviour in the English temperature series. Haldrup and Nielsen (2006) explain that $I(d)$ processes with fractional d fit the power price data well in the presence of long memory effects. Similar to this study, Caporin et al. (2012) apply the VARFIMA framework to a joint model of power prices and temperatures.

We denote the power log-price series by $x_{1,t}$, the temperature series by $x_{2,t}$, and the hydrobalance series by $x_{3,t}$. Their joint evolution is governed by the following dynamic system:

$$\Phi(L)\Delta(L) \begin{pmatrix} x_{1,t} - \Lambda_{1,t} \\ x_{2,t} - \Lambda_{2,t} \\ x_{3,t} \end{pmatrix} = \Theta(L) \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \varepsilon_{3,t} \end{pmatrix}, \quad (4.3)$$

$$\Phi(L) = I - \sum_{j=1}^p \begin{pmatrix} \phi_{11}^j & \phi_{12}^j & \phi_{13}^j \\ 0 & \phi_{22}^j & 0 \\ 0 & 0 & \phi_{33}^j \end{pmatrix} L^j, \quad \Theta(L) = I + \sum_{k=1}^q \begin{pmatrix} \theta_{11}^k & \theta_{12}^k & \theta_{13}^k \\ 0 & \theta_{22}^k & 0 \\ 0 & 0 & \theta_{33}^k \end{pmatrix} L^k,$$

²Hurst was an English civil servant sent to Egypt as a hydrological consultant to predict how much the Nile floods from year to year. He developed rescaled range statistics, which became known later as the Hurst's exponent (H), and is related to the fractional difference parameter d .

where $\Phi(L)$ is a restricted vector autoregressive (VAR) polynomial of order p , with L denoting the lag operator ($L^j x_t = x_{t-j}$), and I denoting a 3×3 identity matrix. $\Delta(L)$ is a diagonal long-memory matrix with a typical diagonal element $(1-L)^{d_i}$, d_i is a fractional difference parameter for variable i ; $\Theta(L)$ is a restricted vector moving average (VMA) polynomial of order q ; $(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t})'$ is the vector of innovations; and finally, $\Lambda_{1,t}$ and $\Lambda_{2,t}$ are the seasonal mean functions given by Eq. (4.1) and Eq. (4.2), respectively. Recall that hydrobalance is de-seasonalised a priori, representing the deviations of the total hydrological resources from the seasonal mean level. $\Phi(L)$ and $\Theta(L)$ are restricted such that temperature and hydrobalance can affect power prices, while the opposite is not possible.³ Note that in the process of lag order selection, we will restrict $\Phi(L)$ and $\Theta(L)$ even further to achieve the highest possible sparsity while retaining the essential effects. The parameters d_i determine the long-range behaviour of the series, while the parameters in $\Phi(L)$ and $\Theta(L)$, together with the lag order p and q , determine the short-range properties.

The difference operator $(1-L)^d$, for any real d , is an infinite linear filter given by the following binomial expansion:

$$(1-L)^d = \sum_{k=0}^{\infty} \binom{d}{k} (-1)^k L^k, \quad (4.4)$$

with the binomial coefficients $\binom{d}{k} = \frac{d!}{k!(d-k)!} = \frac{\Gamma(d+1)}{\Gamma(k+1)\Gamma(d-k+1)}$, where $\Gamma(\cdot)$ denotes the Gamma function. Hosking (1981) shows that under certain assumptions ensuring stationarity and invertibility, a fractional process has infinite moving average and autoregressive representations with coefficients based on the binomial expansion of the difference operator. In practice, the truncated versions of these rep-

³It is reasonable to allow temperature to influence the evolution of hydrobalance in the mean equation. However, since the primary focus of this study is power price dynamics, we do not explore the effect of temperature on hydrobalance.

representations are often estimated with approximate maximum likelihood methods. Not all fractionally integrated processes have long memory. An $I(d)$ process shows the long-memory property in the form of hyperbolic autocorrelation decay rate only if $d > 0$. Further, $d < \frac{1}{2}$ corresponds to a stationary process with a finite variance and an integrable spectral density. For the range $\frac{1}{2} \leq d < 1$, an $I(d)$ process is not stationary, but still mean reverting, and we can define its spectral density in a more general form, although not integrable (see Beran, 1994, for details). Finally, $d > 1$ leads to a non-stationary and non-mean-reverting case with long memory. Note that we can reduce the case of $d > 1$ to one of the cases mentioned above by taking the appropriate number of integer differences (e.g., if x_t is $I(1.2)$, then $(1 - L)^1 x_t = x_t - x_{t-1}$ is $I(0.2)$).

3.1.2 Conditional covariance

The innovation process in our model follows a conditional distribution with zero mean and time-varying covariance matrix H_t :

$$(\varepsilon_{1,t}, \varepsilon_{2,t}, \varepsilon_{3,t})' \mid \omega_{t-1} \sim D(0, H_t), \quad (4.5)$$

where ω_{t-1} denotes the information set at $t - 1$, which constitutes all past observations. The importance of modelling the time-varying volatility both in financial and meteorological data is an established fact. There is far less consensus, however, on what the best way to do this is, and the choice of model is often driven by the specifics of the dataset and the application in mind. Previous studies on temperature modelling suggest that temperature volatility has yearly cycles, similar to the mean. Benth and Šaltytė-Benth (2005, 2007) calibrate the truncated Fourier series to the daily temperature residuals. Campbell and Diebold (2005) propose conditional volatility dynamics for temperature that combines a seasonal component captured by Fourier

series and a cyclical component captured by a Generalized Autoregressive Conditional Heteroskedasticity (GARCH) process. Including a periodic component in the volatility process seems to be relevant for meteorological series. However, a very limited number of multivariate volatility models allow for inclusion of exogenous variables and deterministic terms, mostly due to the excessive parameter restrictions required to guarantee positive semi-definiteness of the covariance matrix. In the setup most similar to ours, Caporin et al. (2012) model conditional variances of power and temperature by long-memory log-GARCH processes with deterministic components, and model conditional correlations separately. They mention that their approach to modelling correlations cannot be generalized to systems of a dimension higher than two.

In this study, we specify the evolution of the full covariance matrix directly using the multivariate GARCH framework of the Baba-Engle-Kraft-Kroner (BEKK) type defined in Engle and Kroner (1995). Following the multivariate GARCH literature, the vector of model error terms from Eq. (4.3) is written as:

$$\varepsilon_t = H_t^{1/2} z_t, \quad (4.6)$$

where z_t is a 3×1 vector of independently identically distributed (i.i.d.) innovations with zero mean and unit variance, and $H_t^{1/2}$ is the 3×3 square root of the conditional covariance matrix, which imposes the desired dependence structure.

We assume that the conditional covariance matrix H_t follows a BEKK(1,1) process:

$$H_t = C'C + A'\varepsilon_{t-1}\varepsilon'_{t-1}A + B'H_{t-1}B, \quad (4.7)$$

where A , B , and C are 3×3 parameter matrices, C is lower triangular, and ε_{t-1} is the 3×1 vector of innovations in Eq. (4.3).

An important advantage of the BEKK model is that H_t is positive semi-definite by construction. In addition, the off-diagonal elements in the A and B matrices have immediate interpretations in terms of the cross-variable volatility spillover effects.⁴ Due to the nature of our series, we restrict some of these off-diagonal elements to zero. In particular, we rule out any cross-effects to the temperature series and allow temperature, but not power, to affect hydrobalance.

3.1.3 Distributional assumptions

We complete the model framework with a specification of the joint distribution of the i.i.d. innovation vector z_t in Eq. (4.6).

Despite the wide acknowledgement that financial data series exhibit heavy tails and skewness, the normal distribution is still dominant in the modelling literature for several reasons. First, it is convenient to resort to the asymptotic properties of the Quasi-Maximum Likelihood (QML) estimator, which is consistent even if the true conditional distribution of innovations is not normal, provided that the conditional mean and variance models are correctly specified. Second, the normal distribution often allows for closed-form pricing and hedging of derivative assets while introducing any non-normal dynamics requires computationally intensive numerical methods to price even standard derivatives in most cases.

Temperature series, on the other hand, can be much better approximated by the normal distribution than any price series. Most of the temperature modelling papers we referred to (Alaton, 2002; Brody et al., 2002; Campbell and Diebold, 2005) rely on the normal distribution assumption for the residuals. However, Benth and Šaltyté-Benth (2005) show that normality is rejected for some of the Norwegian temperature data and propose to apply the generalized hy-

⁴See Chapter 2 Section 3.2 for a more detailed discussion of the BEKK model.

perbolic distribution family. A closer look at Figure 4.2, which plots the Swedish temperature series, reveals that extreme deviations from the seasonal mean are quite common, especially in the winter periods. Recall that it is those extreme values, and not the average dynamics, that give rise to excess power demand and are of primary interest for any risk management or production planning purposes.

Finally, hydrological time series are known to be significantly positively skewed (see Helsen and Hirsch, 2002), which is also the case with our hydrobalance data, as confirmed by Figure 4.5. Considering all of the arguments above, we suggest using a flexible heavy-tailed and skewed distribution while keeping the normal distribution as the benchmark for comparison. The multivariate skew-Student distribution with independent components of Bauwens and Laurent (2002, 2005) appears to be an excellent choice, since it allows the univariate marginal distributions to have individual skewness and tail properties. Furthermore, this distribution is relatively straightforward to augment with GARCH-type dynamics, as discussed in Chapter 2.

In this study, we specify the multivariate skew-Student distribution for the vector of standardized innovations z_t in Eq. (4.6). Following Bauwens and Laurent (2002), a $k \times 1$ random vector z_t is standard multivariate skew-Student distributed with independent components if its probability density function is given by:

$$f(z_t) = \left(\frac{2}{\sqrt{\pi}} \right)^k \left[\prod_{i=1}^k \frac{\xi_i s_i}{1 + \xi_i^2} \frac{\Gamma\left(\frac{v_i+1}{2}\right)}{\Gamma\left(\frac{v_i}{2}\right) \sqrt{v_i - 2}} \left(1 + \frac{\kappa_{i,t}^2}{v_i - 2} \right)^{-\frac{1+v_i}{2}} \right], \quad (4.8)$$

where

$$\kappa_{i,t} = (s_i z_{i,t} + m_i) \xi_i^{-I_{i,t}}, \quad (4.9)$$

and

$$I_{i,t} = \begin{cases} 1 & \text{if } z_{i,t} \geq -\frac{m_i}{s_i} \\ -1 & \text{if } z_{i,t} < -\frac{m_i}{s_i} \end{cases}, \quad (4.10)$$

with skewness parameters $\xi = (\xi_1, \dots, \xi_k)$ and degrees of freedom parameters $v = (v_1, \dots, v_k)$ for $v_i > 2$, and $\Gamma(\cdot)$ denoting the Gamma function. The constants $m_i = m_i(\xi_i, v_i)$ and $s_i = s_i(\xi_i, v_i)$ are the means and standard deviations of the non-standardized skew-Student density, respectively, defined by:

$$m_i(\xi_i, v_i) = \frac{\Gamma\left(\frac{v_i-1}{2}\right) \sqrt{v_i-2}}{\sqrt{\pi} \Gamma\left(\frac{v_i}{2}\right)} \left(\xi_i - \frac{1}{\xi_i} \right), \quad (4.11)$$

$$s_i^2(\xi_i, v_i) = \left(\xi_i^2 + \frac{1}{\xi_i^2} - 1 \right) - m_i^2. \quad (4.12)$$

$\xi_i = 1$ corresponds to the symmetric density, while the thickness of the tails is decreasing in v_i . Note that the standardized multivariate normal density is the limiting distribution of $f(z_t)$ in Eq. (4.8), when $\xi_i = 1$ and $v_i \rightarrow \infty$.

3.2 Model identification and estimation procedure

Long-memory model estimation is a well-addressed area, and many estimation methods have been proposed in the literature. Most are based on either a time domain or frequency domain representation of the density function. The time domain procedures include various implementations of exact maximum likelihood, such as the Durbin-Levinson algorithm and state space methods, as well as a number of approximate likelihood methods based on truncated versions of autoregressive and moving average representations of long-memory processes (see, e.g., Hasslett and Raftery, 1989). The frequency domain procedures include Whittle estimators and various semiparametric methods, and are based in one way or another on the calculation of the

periodogram of the series using Fast Fourier Transform (FFT). These methods offer significant computational advantage over the time domain methods, but at the cost of lower precision of estimates. In this study, we estimate all model parameters jointly in the time domain using exact maximum likelihood. However, model complexity requires us to proceed in several steps.

First, we estimate the parameters in Λ_1 and Λ_2 using the least squares method. Subtracting these functions from the original power log-price and temperature observations yields the de-seasonalised series, which are the input to the next step.

The second step is to get the initial estimates of the fractional difference parameters. We would like to get a consistent estimate of the degree of fractional integration in the series without making any prior assumptions about the short-range properties. This can be achieved by using a semiparametric estimation method, which does not require specification of the parametric model and relies only on the assumption about the shape of the spectral density of the time series. The most common semiparametric methods to estimate long-memory parameters are local Whittle (see Künsch, 1987, and Robinson, 1995a), and log-periodogram regression (see Geweke and Porter-Hudak, 1983, and Robinson, 1995b). However, as Shimotsu and Phillips (2005) point out, these estimators are inconsistent for $d > 1$, and discontinuous at several points in the non-stationary region, leading to non-normal limit theory. Instead, they suggest a general purpose semiparametric estimator called the exact local Whittle estimator with well-behaved asymptotic properties in the wide range of stationary and non-stationary values. We use the exact local Whittle estimator of Shimotsu and Phillips (2005) to obtain the initial estimates of the d -parameters.⁵

⁵MATLAB code for exact local Whittle estimation is available at Katsumi Shimotsu's personal website.

The third step is to apply the fractional difference filter to the series and identify the short-range part of the model, that is, the structure of $\Phi(L)$ and $\Theta(L)$. Calculating fractional differences is in itself a non-trivial task. Standard implementations of fractional differencing based on the binomial expansion of the difference operator have $\mathcal{O}(n^2)$ time complexity, which means that the number of operations performed to compute the differenced series is a quadratic function of the input size. This is acceptable if differences are only to be computed once, but makes the joint estimation of all parameters in a trivariate model with several thousand time series points practically infeasible. However, Jensen and Nielsen (2013) suggest a fast fractional difference algorithm that takes advantage of a frequency-domain transform of the series. Their algorithm is of $\mathcal{O}(n \log n)$ time complexity and offers substantial computational advantages. We identify the short-range dynamics by inspecting the ACFs and PACFs of the differenced series following the standard practice. In addition, we conduct a number of univariate estimations assuming constant variance and compare them based on the information criteria.

The final step is the joint estimation of all model parameters by exact maximum likelihood using the parameter estimates from the previous steps as starting values only.⁶ We implement this procedure for the cases of normal distribution and skew-Student distribution, separately. In the case of normally distributed residuals, the log-likelihood function is given by the log of the multivariate normal density function. In the case of the skew-Student distributed residuals,

⁶Parameters in the seasonal mean functions are not re-estimated to decrease computational time. As starting values for the shape parameters in the skew-Student distributions, we use $\xi_i = 1$ for all skewness parameters and $\nu_i = 100$ for all kurtosis parameters ($\nu_i \rightarrow \infty$ corresponds to normality).

the log-likelihood function is given by:

$$\ln L(\theta) = \sum_{t=\max(p,q)+1}^T \left\{ \ln f(z_t) - \frac{1}{2} \ln |H_t| \right\}, \quad (4.13)$$

where θ is the parameter vector for the full model, $f(z_t)$ is the probability density function in Eq. (4.8), T is the number of time series observations, and $|H_t|$ denotes the determinant of H_t . Note that the summation is conditional on the first p or q observations, whichever is larger, owing to the lag order of $\Phi(L)$ and $\Theta(L)$ in the mean equations. The second term in the sum in Eq. (4.13) is the Jacobian correction term arising in the transformation from z to ε . We calculate the square root matrix $H_t^{1/2}$, which is required to obtain the vector of standardized residuals z_t as given by Eq. (4.6) at each time point using a standard spectral decomposition. The initial H_t is set to the sample covariance matrix of the fractionally differenced data and the initial values of the residuals are set to zero.

The log-likelihood function is maximized by simulated annealing, following Goffe, Ferrier and Rogers (1994). To further increase the chance of identifying the global optimum, we use consistent QML estimates as starting values for the model with skew-Student distributed innovations. Finally, we calculate the standard errors of the parameters using the outer product gradient method with numerical first derivatives.

4 Results

This section presents the results of model identification and estimation. We start by discussing how we identified the conditional mean system within our general framework. We then discuss the estimated parameters. Finally, we examine the model implied second moments

in light of the findings from our preliminary analysis.

4.1 Model identification results

The long-memory parameters are estimated by the exact local Whittle method as follows: 0.6747 for power, 0.2970 for temperature, and 1.1115 for hydrobalance. These results are consistent with the autocorrelation functions of the non-differenced data in Figure 4.4. Specifically, hydrobalance has the slowest ACF decay, reflected in the highest d -parameter, while temperature shows the fastest (yet still hyperbolic) ACF decay, with the lowest d -parameter. Further, the d -parameter of temperature lies in the stationary region, while the other two take values in the non-stationary region. It is worth noting that the d -parameter of hydrobalance is above one, which means that mean reversion property is lost. The implications of this result on the model simulation will be discussed further.

We proceed to identifying the short-range properties of the conditional mean system. Figure 4.7 displays the autocorrelation and partial autocorrelation functions of the fractionally differenced series. The 95% white noise confidence bounds are given by the horizontal blue lines.

We can see that power has a slight spike in both functions at the first lag, and weekly periodic patterns. One alternative to capture weekly periodicities is to take seasonal differences. However, due to the presence of non-seasonal fractional differencing in our model, we prefer to include seasonal autoregressive lags instead. Further investigation in the univariate framework reveals that two weekly seasonal terms, in addition to the non-seasonal AR(1) term, are sufficient to whiten the residuals. Moreover, a parsimonious specification with a single temperature term and a single hydrobalance term in the power mean equation is preferred based on the information criteria.




Figure 4.7: Autocorrelation and partial autocorrelation functions after fractional differencing.

Both temperature and hydrobalance show an AR signature, with ACF decaying gradually and PACF truncated at a certain lag. The PACF of temperature features three significant lags, while the PACF of hydrobalance is truncated at the first lag. Overall, the interpretation of the functions after appropriate fractional differencing is straightforward, suggesting to model temperature as an AR(3) process and hydrobalance as an AR(1) process. We argue that there is no need to include any moving average terms, thus reducing the general VARFIMA framework presented in the previous section to its special case of fractionally integrated vector autoregression. This leads us to

identify the following conditional mean equation system:

$$\begin{aligned}
 y_{1,t} - \phi_{11}^1 y_{1,t-1} - \phi_{11}^7 y_{1,t-7} - \phi_{11}^{14} y_{1,t-14} - \phi_{12}^1 y_{2,t-1} - \phi_{13}^1 y_{3,t-1} &= \varepsilon_{1,t} \\
 y_{2,t} - \phi_{22}^1 y_{2,t-1} - \phi_{22}^2 y_{2,t-2} - \phi_{22}^3 y_{2,t-3} &= \varepsilon_{2,t} \\
 y_{3,t} - \phi_{33}^1 y_{3,t-1} &= \varepsilon_{3,t}
 \end{aligned}
 \tag{4.14}$$

where y_1 , y_2 , and y_3 denote de-seasonalised fractionally differenced power log-price, de-seasonalised fractionally differenced temperature, and fractionally differenced hydrobalance, respectively.

4.2 Estimation results

We first comment on the seasonal mean parameter estimates presented in Table 4.2.

Table 4.2: Seasonal mean parameter estimates.

Power		Temperature	
$\lambda_{1,1}$: constant	3.9575 (0.0113)	$\lambda_{2,1}$: constant	6.9372 (0.0614)
$\lambda_{1,2}$: amplitude	0.1606 (0.0085)	$\lambda_{2,2}$: amplitude	-9.7656 (0.0861)
$\lambda_{1,3}$: phase angle	7.9960 (3.1498)	$\lambda_{2,3}$: phase angle	21.728 (0.5202)
$\lambda_{1,4}$: linear trend	-0.0003 (0.0000)		
$\lambda_{1,5}$: non-business days	-0.1205 (0.0132)		

Note: The table reports the estimated coefficients and their standard errors (in parentheses). The seasonal mean functions are given by Eq. (4.1) and Eq. (4.2). All coefficients are statistically significant at the 1% level. The R^2 values are 0.400 for power and 0.846 for temperature.

We can interpret the parameters $\lambda_{1,1}$ and $\lambda_{2,1}$ as constant average levels of power log-price and temperature, respectively. The

estimated average log-price corresponds to the average price of 52.32 EUR/MWh. The parameters $\lambda_{1,2}$ and $\lambda_{2,2}$ represent half of the distance between the peaks and the troughs of the yearly seasonal functions. Thus, the difference between the temperature highs and lows is around 19 degrees Celsius, while the average power price difference between warm and cold seasons is 16.88 EUR/MWh.⁷ Parameters $\lambda_{1,3}$ and $\lambda_{2,3}$ act as time shifts, placing the cosine waves in the correct phase of power price and temperature yearly cycles. Further, we report a significant decreasing linear time trend in the power price series. As mentioned in Section 2, this phenomenon is related to the changes in the Nordic power generation mix, with a growing share of load covered by low marginal cost renewable generation sources. Finally, note the non-trivial effect of non-business days, comparable in magnitude to the amplitude of the yearly cycle.

We estimated the remaining parameters under two distributional assumptions: the normal distribution, serving as the benchmark case, and the more flexible skew-Student distribution. Since the two model specifications are not nested, it is not possible to formally test them against each other using the likelihood ratio test. However, we can still get an idea of the gain from departing from normality by comparing the starting log-likelihood value of the skew-Student specification, which is -4559.23 , with the final log-likelihood value of -3696.94 . Recall that the optimal parameter values from the normal specification were the starting values for the skew-Student specification, while the starting values of the ξ - and ν -parameters were set to roughly correspond to normality.

Table 4.3 presents the estimated values of the parameters from the stochastic component of the conditional mean system.

⁷ $e^{(3.9575+0.1606)} - e^{(3.9575-0.1606)}$

Table 4.3: Stochastic conditional mean parameter estimates.

		Normal			Skew-Student						
Power		Temperature	Hydrobalance		Power	Temperature	Hydrobalance				
d_1	0.6963	d_2	0.2660	d_3	1.2103	d_1	0.6445	d_2	0.2957	d_3	1.1478
	(0.0231)		(0.0371)		(0.0227)		(0.0214)		(0.0360)		(0.0145)
ϕ_{11}^1	<i>0.0130</i>	ϕ_{22}^1	0.8506	ϕ_{33}^1	0.1351	ϕ_{11}^1	0.0755	ϕ_{22}^1	0.8283	ϕ_{33}^1	0.1596
	(0.0290)		(0.0411)		(0.0280)		(0.0278)		(0.0394)		(0.019)
ϕ_{11}^7	0.2128	ϕ_{22}^2	-0.2561			ϕ_{11}^7	0.2058	ϕ_{22}^2	-0.2367		
	(0.0155)		(0.0259)				(0.0143)		(0.0244)		
ϕ_{11}^{14}	0.1289	ϕ_{22}^3	0.0629			ϕ_{11}^{14}	0.1170	ϕ_{22}^3	0.0445		
	(0.0124)		(0.0193)				(0.0120)		(0.0183)		
ϕ_{12}^1	-0.0037					ϕ_{12}^1	-0.0035				
	(0.0005)						(0.0005)				
ϕ_{13}^1	-0.0153					ϕ_{13}^1	-0.0126				
	(0.0014)						(0.0015)				

Note: The table reports the estimated coefficients and their standard errors (in parentheses). The conditional mean system is given by Eq. (4.3) in the general form for the original series, and Eq. (4.14) in the restricted form for the de-seasonalised and differenced series. Non-significant coefficients are reported in italics, while the remaining coefficients are statistically significant at the 1% level.

The estimates of the memory parameters are very close to the univariate exact local Whittle estimates discussed earlier. Haldrup and Nielsen (2006) find that the Nordic zonal hourly spot price series show long memory with d ranging between 0.31 and 0.52. Caporin et al. (2012) report d -parameters of 0.39 and 0.19 for the Oslo area daily power and temperature series. We find a higher degree of long memory for both series, with d_1 of power being well in the non-stationary region.

In general, there are no extreme differences between the parameter estimates from the two distributional specifications. We can see the natural trade-off between the degree of long memory and the magnitude of the non-seasonal autoregressive coefficients in all the three series. In particular, a higher estimate of d_1 under the normal dis-

tribution is coupled with a lower and nonsignificant ϕ_{11}^1 . The autoregressive coefficients from the temperature mean equation are slightly lower in magnitude than those reported in Benth, Šaltytė-Benth and Koekebakker (2008) for the Stockholm temperature series, due to the presence of the long-memory component in our specification. Further, we report significant ϕ_{11}^7 and ϕ_{11}^{14} coefficients capturing the first- and the second-order weekly seasonal autoregressive patterns in the power price. Finally, we find significant negative temperature and hydrobalance effects in the power mean equation, which is in line with our preliminary analysis.

Table 4.4 presents the estimated values of the conditional covariance and distributional parameters.

We first analyse the diagonal coefficients of A and B -matrices. Caporin et al. (2012) find that temperature has a lower degree of persistence in volatility than power price. According to our results, temperature and hydrobalance volatility have weaker ARCH-effects (as measured by a_{22} and a_{33}) compared to power price volatility. However, the GARCH-coefficient b_{22} for temperature suggests a higher degree of persistence in volatility than the corresponding coefficient b_{11} for power. This might be related to the fact that we do not explicitly model the seasonality in temperature volatility, so part of it is accommodated by b_{22} . Interestingly, the power GARCH-coefficient b_{11} appears to be lower than the corresponding estimate for the German power futures contract reported in Chapter 2.

The majority of the off-diagonal coefficients both in A and B -matrices are statistically significant, which confirms the existence of volatility spillover effects. There are differences, however, in the estimates of these effects between our two distributional specifications. The normal specification features significant spillovers from temperature to power volatility, as measured by a_{12} and b_{12} , and smaller magnitude a_{13} coefficient, representing the spillover from hydrobalance to

power volatility. The skew-Student specification, on the other hand, results in nonsignificant temperature-to-power effects, but highly statistically significant and sizeable hydrobalance-to-power effects, as measured by both a_{13} and b_{13} .

Table 4.4: Conditional covariance and distributional parameter estimates.

		Normal			Skew-Student						
c_{11}	0.0102 (0.0134)				c_{11}	0.0153*** (0.0017)					
c_{21}	-0.0189*** (0.0072)	c_{22}	0.1169*** (0.0377)		c_{21}	-0.0069*** (0.0020)	c_{22}	0.1904*** (0.0274)			
c_{31}	-0.0076*** (0.0020)	c_{32}	-0.1312*** (0.0244)	c_{33}	0.0791*** (0.0044)	c_{31}	-0.0088*** (0.0016)	c_{32}	-0.0792*** (0.0221)	c_{33}	0.0752*** (0.0084)
a_{11}	0.5435*** (0.0142)	a_{12}	-0.5599** (0.2255)	a_{13}	-0.1963** (0.0956)	a_{11}	0.4414*** (0.0223)	a_{12}	-0.1430 (0.1733)	a_{13}	-0.2178*** (0.0691)
		a_{22}	0.1830*** (0.0117)				a_{22}	0.1615*** (0.0137)			
		a_{32}	-0.0609** (0.0274)	a_{33}	0.1350*** (0.0067)		a_{32}	-0.0196 (0.0258)	a_{33}	0.1956*** (0.0126)	
b_{11}	0.8311*** (0.0086)	b_{12}	0.2075* (0.1071)	b_{13}	0.0561 (0.0400)	b_{11}	0.8786*** (0.0094)	b_{12}	-0.0272 (0.0707)	b_{13}	0.1554*** (0.0288)
		b_{22}	0.9734*** (0.0032)				b_{22}	0.9764*** (0.0038)			
		b_{32}	0.0471*** (0.0078)	b_{33}	0.9829*** (0.0016)		b_{32}	0.0224*** (0.0074)	b_{33}	0.9764*** (0.0028)	
						ξ_1	0.9843*** (0.0212)	ξ_2	1.0288*** (0.0270)	ξ_3	1.8734*** (0.0574)
						v_1^{-1}	0.2598*** (0.0168)	v_2^{-1}	0.1351*** (0.0218)	v_3^{-1}	0.2569*** (0.0198)

Note: The table reports the estimated coefficients and their standard errors (in parentheses). The conditional covariance model is given by Eq. (4.7) and is parameterized by the 3×3 matrices C , A and B , with typical elements c_{ij} , a_{ij} , and b_{ij} , for $i, j = 1$ (power), 2 (temperature), and 3 (hydrobalance), respectively. The C -matrix is lower triangular, while the A and B matrices are restricted such that temperature volatility dynamics is exogenous, hydrobalance volatility is allowed to be affected by temperature but not by power, while all variables in the system can affect power volatility. Skew-Student distributional parameters are reported in the last two rows. Superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

We also find the b_{32} -coefficient on temperature-to-hydrobalance volatility spillover effect to be statistically significant at the 1% level

in both specifications. However, its magnitude is much lower than of the coefficients capturing spillovers to power. Note that the signs of the off-diagonal parameters do not have a straightforward interpretation because these parameters appear in several non-linear terms determining each element of the conditional covariance matrix at each time point.

The last two rows in Table 4.4 report the estimates of the skewness and the inverses of the degrees of freedom parameters.⁸ Although the ξ_1 of power is below 1, and the ξ_2 of temperature is above 1, the 95% confidence intervals for these parameters leave the question of asymmetry open, with the lower bound in the negative region and the upper bound in the positive region. The story is different with the ξ_3 of hydrobalance, which is well in the positive asymmetry region, as expected. Further, we find that power and hydrobalance have very similar tail properties with ν -parameters close to 4. Temperature shows less heavy tails with the ν_2 estimate of 7.4, although it still implies a relatively fat-tailed distribution.⁹ Bauwens and Laurent (2005) report similar degrees of freedom parameter values for the exchange rate series and several U.S. stocks. Taking another look at Figure 4.2, we expect that this estimate is mostly driven by the extreme temperature occurrences in winter periods, and this has to be taken into account in a simulation from the model.

Diagnostic checks of the residuals from both model specifications reveal that we are left with zero-mean uncorrelated noise. Examining the estimated volatilities and correlations, plotted in Figures 4.8 and 4.9, allows us to further assess the in-sample model performance.

We can see the resemblance between the estimated volatility processes and the daily changes series in Figure 4.5. Temperature volatil-

⁸We estimate ν^{-1} instead of ν itself for numerical reasons.

⁹In Chapter 2, we report a common ν estimate of 7.7 for power, gas, coal, and carbon emission allowances front-year futures log-return series. It is natural that futures returns, especially in the long-end, show less kurtosis than spot returns.

ity displays peaks during the winter periods and troughs during the summer periods. Hydrobalance volatility starts the yearly cycle at a relatively low level, reaches the minimum around April, and then takes on an upward trend with a peak in August – September, followed by a sharp decline.




Figure 4.8: Model implied volatilities.

The estimated conditional correlations are well in line with our expectations. The correlation between power and temperature reaches the minimum of -0.5 to -0.4 during the winter months, and is roughly zero during the summer months. The correlation between power and hydrobalance does not have a pronounced seasonal shape, and mostly stays in the negative region between -0.5 and 0 , oscillating around the average level of -0.25 , with a few extremes.




Figure 4.9: Model implied correlations.

Finally, there seem to be very minor differences between the estimated second moments from the two distributional specifications. Hydrobalance volatility and the power/hydrobalance correlations display more variation under the skew-Student specification. On the rest of the plots, the red lines and the blue lines coincide almost perfectly.

5 Application: Scenario analysis

This section illustrates how our model can be used to generate a number of power price scenarios under different hydrological and temperature conditions. We address the issues related to simulation from the model and present an overview of the simulation results from the skew-Student specification.

We start the simulation on February 22, 2016 (the day after the

sample period ends), and finish on February 28, 2017, yielding a simulation length of 373 days. We use the last sample values of power price, temperature, and hydrobalance as the starting points for all simulated paths. In addition, the estimated conditional covariance matrix on the last day of the sample period (H_T) and the last values of the residuals are used to iterate the BEKK process forward. For each day of the simulation period, we draw the random shocks z_t from the underlying skew-Student distributions using the analytical quantile function.¹⁰ Appendix A presents the kernel density estimates of the random samples drawn from the univariate skew-Student distributions with skewness and degrees of freedom parameters equal to our estimates. We construct the error terms ε_t using Eq. (4.6). Next, the long memory is created by applying numerical fractional integration of the error term series.¹¹ Further, we follow Eq. (4.14) to generate the stochastic mean component. Finally, we add the predicted seasonal mean component for power and temperature, and transform the log-price back to the natural units.

Figures 4.10 – 4.12 show the historical data series starting from April 27, 2015, followed by ten simulated paths.

A first glance at the simulated paths suggests that the model does a fairly good job of capturing the empirical properties of the modelled series. It is worth noting that since the extreme power price and temperature observations are typically observed during the winter months, we can adjust the random shock generation to reflect this.

¹⁰See Eq. (3.4) for the quantile function of the skew-Student distribution.

¹¹To generate a fractionally integrated process, we approximate the binomial expansion of $(1 - L)^{-d}$ by truncating at 100 terms. The first 100 simulated values use the actual model residuals.




Figure 4.10: Historical power prices and simulated paths.




Figure 4.11: Historical temperature and simulated paths.

We can see in Figure 4.11 that only one temperature path out of ten produced an outlier during the winter period. However, we can easily adjust the simulation procedure to increase the likelihood of outliers occurring in winter, and if necessary, decrease the likelihood of outliers in other periods.¹² Overall, the simulation stage reveals the true benefits of using the skew-Student distribution, since the normal distribution cannot generate large enough moves frequently enough to produce realistic behaviour.




Figure 4.12: Historical hydrobalance and simulated paths.

As Figure 4.12 illustrates, we can generate a wide variety of hydrological scenarios. Recall from the discussion of the difference parameter properties that a value of d above 1 corresponds to a non-stationary case without mean reversion property. This means that

¹²For instance, this can be done by drawing the quantiles from a distribution other than uniform while generating the random values from the skew-Student distribution.

some of the hydrobalance paths produced by our model will be unreasonable, and have to be discarded, which can also be automated in the simulation by introducing the bounds on how far hydrobalance can wander away from zero.

A limitation of the model that is more fundamental is that although extreme power price occurrences are likely, there is nothing in the model enforcing fast mean reversion and enabling spikes to appear once in a while. On the contrary, the otherwise desirable long-memory property makes it impossible for a single extreme value to occur. However, for an application such as meteorological scenario analysis, this limitation is of minor importance.

Our model creates plenty of interesting opportunities in scenario analysis. For instance, an average power price can be calculated for a range of temperature and hydrobalance combinations. This might be of interest for production planning in power markets with heavy reliance on hydrological generation, such as the Nordic market. As an illustration, we simulate 1000 scenarios from the model and calculate the average power price, temperature, and hydrobalance values for November 1, 2016 – February 28, 2017, for each scenario. Figure 4.13 plots the average power price and temperature observation pairs grouped by hydrological conditions. The red dots mark the driest quarter of scenarios, while the blue crosses mark the wettest quarter. We omit the observations in the middle of the range.

We can see that, with rare exceptions, the average power price is higher in dry scenarios, given the same average temperature. Further, extremely high average power prices tend to occur under a combination of cold and dry conditions. A similar analysis can be done for statistics other than the average.




Figure 4.13: A scenario analysis example.

Another potential application of the model is Monte Carlo pricing of weather derivatives, such as energy quanto contracts. Energy quanto contracts have a payoff that depends on the product of two indices: an energy index (e.g., an average power price during the delivery period) and a temperature index (e.g., Heating Degree Days, or HDD). A wide variety of payoff structures is possible for quanto contracts, such as swap, put/call, collar, and so on. Caporin et al. (2012) illustrate how their bivariate power-temperature model can be utilised to price such contracts. Our model can be applied in a similar fashion, with the additional flexibility of computing ‘hydrological bounds’ on contract prices.

6 Summary and conclusions

In this study we propose a model for the joint evolution of spot power price, temperature, and hydrobalance. Our model successfully captures most of the discovered empirical features, such as long memory and heavy tails in all series, yearly seasonal patterns in power price and temperature, weekly periodic patterns in power price, pronounced positive skewness in hydrobalance, and time-varying conditional second moments. We find that in the Nordic market, power price is inversely related to temperature throughout the year, except for summer months, when the effect is nonsignificant. Hydrobalance, on the other hand, negatively affects power price in all periods, since in dry hydrological conditions, higher marginal cost generation sources set the price. Further, we confirm the existence of volatility spillover effects from temperature and hydrobalance to power. We illustrate how our model can be used to generate a variety of weather scenarios and to analyse the implications for power prices. The model is relevant for power markets with a dominant share of hydrological generation and provides a wide scope of opportunities for scenario analysis with relatively little meteorological input.

References

- [1] Alaton, P., Djehiche, B., Stillberger, D., 2002. On modelling and pricing weather derivatives. *Applied Mathematical Finance* 9, 1–20.
- [2] Bauwens, L., Laurent, S., 2002. A new class of multivariate skew densities, with applications to GARCH models. CORE Discussion Paper 20.
- [3] Bauwens, L., Laurent, S., 2005. A new class of multivariate skew densities, with application to Generalized Autoregressive Condi-


- tional Heteroscedasticity models. *Journal of Business and Economic Statistics* 23, 346–354.
- [4] Benth, F.E., Šaltytė-Benth, J., 2005. Stochastic modelling of temperature variations with a view towards weather derivatives. *Applied Mathematical Finance* 12, 53–85.
- [5] Benth, F.E., Šaltytė-Benth, J., 2007. The volatility of temperature and pricing of weather derivatives. *Quantitative Finance* 7, 553–561.
- [6] Benth, F.E., Šaltytė-Benth, J., Koekebakker, S., 2008. Stochastic modelling of electricity and related markets. *Advanced series on statistical science and applied probability*; v. 11. World Scientific, Singapore.
- [7] Benth, F.E., Lange, N., Myklebust, T.Å., 2012. Pricing and hedging quanto options in energy markets. SSRN Working Paper.
- [8] Beran, J., 1994. *Statistics for long-memory processes*. Chapman and Hall, London.
- [9] Brody, D.C., Syroka, J., Zervos, M., 2002. Dynamical pricing of weather derivatives. *Quantitative Finance* 3, 189–198.
- [10] Campbell, S.D., Diebold, F.X., 2005. Weather forecasting for weather derivatives. *Journal of the American Statistical Association* 100, 6–16.
- [11] Caporin, M., Pres, J., Torro, H., 2012. Model based Monte Carlo pricing of energy and temperature Quanto options. *Energy Economics* 34, 1700–1712.
- [12] Ebbeler, S., Benth, F.E., Kiesel, R., 2014. Indifference pricing of weather derivatives based on electricity futures. In: Prokopczuk, M. (Ed.). *Energy pricing models: recent advances, methods, and tools*. Palgrave Macmillan, New York.

- [13] Engle, R., Kroner, K., 1995. Multivariate simultaneous generalized ARCH. *Econometric Theory* 11, 122–150.
- [14] Geweke, J., Porter-Hudak, H., 1983. The estimation and application of long memory time series models. *Journal of Time Series Analysis* 4, 221–238.
- [15] Goffe, W., Ferrier, G., Rogers, J., 1994. Global optimization of statistical functions with simulated annealing. *Journal of Econometrics* 60, 65–99.
- [16] Granger, C.W.J., Joyeux, R., 1980. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 15–29.
- [17] Green, R., 2015. A power market forward curve with hydrology dependence - an approach based on Artificial Neural Networks. Knut Wicksell Working Paper Series 2015/1, Knut Wicksell Centre for Financial Studies, Lund University.
- [18] Haldrup, N., Nielsen, M.Ø., 2006. A regime switching long memory model for electricity prices. *Journal of Econometrics* 135, 349–376.
- [19] Halldin, R., 2005. Stochastic modeling and optimization under uncertainty of a hydro power system. Doctoral dissertation, Centre for Mathematical Sciences, Lund Institute of Technology.
- [20] Hasslett, J., Raftery, A.E., 1989. Space-time modelling with long-memory dependence: Assessing Ireland’s wind power resource. *Journal of Applied Statistics* 38, 1–50.
- [21] Helsel, D.R., Hirsch, R.M., 2002. Statistical methods in water resources. *Techniques of water resources investigations*, Book 4, chapter A3. U.S. Geological Survey.
- [22] Hosking, J.R.M., 1981. Fractional differencing. *Biometrika* 68,

165–176.

- [23] Hurst, H.E., 1951. Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers* 116, 770–799.
- [24] Jensen, A.N., Nielsen, M.Ø., 2013. A fast fractional difference algorithm. QED Working Paper 1307, Queen’s University.
- [25] Künsch, H., 1987. Statistical aspects of self-similar processes. In: Prokhorov, Y., Sazanov, V. (Ed.): *Proc. First World Congress of the Bernoulli Society* 1, 67–74. VNU Science Press, Utrecht.
- [26] Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 54, 159–178.
- [27] Lucia, J., Schwartz, E., 2002. Electricity prices and power derivatives: evidence from the Nordic power exchange. *Review of Derivatives Research* 5, 5–50.
- [28] Palma, W., 2007. *Long-memory time series models: theory and methods*. Wiley Series in Probability and Statistics. John Wiley and Sons, New Jersey.
- [29] Robinson, P.M., 1995a. Gaussian semiparametric estimation of long range dependence. *Annals of Statistics* 23, 1630–1661.
- [30] Robinson, P.M., 1995b. Log-periodogram regression of time series with long-range dependence. *Annals of Statistics* 23, 1048–1072.
- [31] Shimotsu, K., Phillips, P.C.B., 2005. Exact local Whittle estimation of fractional integration. *Annals of Statistics* 33, 1890–1933.

Appendix A



Note: The figure shows the smooth kernel density estimates of the random samples drawn from the standard normal distribution and the standardized skew-Student distribution with ξ and v parameters equal to our estimates. Each random sample consists of 1000 values.