The Carbon Market, Oil, and the Macroeconomy

Inauguraldissertation

zur Erlangung des akademischen Grades
Doctor Rerum Politicarum

an der
Fakultät für Wirtschafts- und Sozialwissenschaften
der Ruprecht-Karls-Universität Heidelberg

vorgelegt von

Daniel Rittler

Heidelberg, September 2012

Acknowledgments

First of all, I would like to thank my supervisor Christian Conrad. I am deeply indebted to him for guiding me through the process of writing this thesis. Without his ongoing advice and support the completion of the thesis would not have been possible. I greatly benefitted from his stimulating and valuable suggestions, comments, and feedback on my research during the last years.

I thank Andreas Löschel for his immediate readiness to be a member of the thesis committee.

I am indebted to my co-authors Christian Conrad, Karin Loch, and Waldemar Rotfuß for the fruitful collaboration. Their intensive effort considerably contributed to the quality of this thesis.

The thesis benefitted from suggestions of participants of numerous international academic conferences and the Departmental Workshop at the Alfred-Weber-Institute. In particular, I would like to thank Maria Mansanet-Bataller, Julien Chevallier, Thomas Eife, Zeno Enders, Marc Gronwald, Menelaos Karanasos, Nicolas Koch, Bruce Mizrach, Benoit Sevi, and Sandra Schmidt.

I greatly enjoyed the free and stimulating working atmosphere at the Chair of Empirical Economics at the Alfred-Weber-Institute and the collaboration with my colleagues Balazs Cserna, Karin Loch, Yuhong Mangels, Alexander Rohlf, Jan-Philip Schmidt, and Peter Schmidt. I also would like to thank Hartmut Kogelschatz of the former Chair of Statistics.

Finally, I am very grateful for the mental support of my wife Catherine during the last years. I also would like to thank my family who supported me throughout my academic years. I owe them a lot.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 General remarks

With the construction of the European market for greenhouse gas emission permits, the European Union has created a new energy market which is also referred to as the European carbon market. This thesis predominantly deals with the econometric analysis of this new market. Section 1.1.1 describes the general framework of the carbon market, provides an overview on the carbon price development, and motivates the central research questions related to this market. In addition to the analysis of the European carbon market, the thesis is concerned with the econometric investigation of the relationship between the oil price and U.S. stock market performance in consideration of developments in the macroeconomic environment. Section 1.1.2 motivates and summarizes the corresponding research questions.

1.1.1 The carbon market

Given the concerns of adverse effects of anthropogenic climate change on ecological systems and mankind caused by a considerable extension of the concentration of greenhouse gases in the atmosphere (see United Nations Framework Convention on Climate Change (1992)), in 2003 the member states of the European Union agreed to establish an economically efficient market for tradable greenhouse gas emission permits. This market is referred to as the European Union Emissions Trading Scheme (EU-ETS) and entered into force in January 2005. According to Directive 2003/87/EC, the EU-ETS is supposed to assist the EU member states to cost-effectively comply with their target to reduce aggregate greenhouse gas emissions in the period 2008 to 2012 relative to the level of 1990 by 8%, as defined under the Kyoto Protocol. While multilateral emissions trading between countries is defined under Article 17 of the Kyoto Protocol, the EU-ETS is designed as a trading platform for greenhouse gas emissions at the company level that according to Chevallier (2012) covers roughly 50% of the EU-wide greenhouse gas emissions. Companies that own production units in energy-intensive sectors, including the sectors combustion, pulp and paper, iron and steel, and cement and lime, do have to underlay their greenhouse gas emissions by European Union Allowances (EUAs).[1] Each EUA warrants the right to emit one tonne of CO_2-equivalent during a specified commitment period such that with the entry into force of the EU-ETS companies do have to pay a price for the emission of greenhouse gases which in turn implies that EUAs constitute an input factor in the companies' production process.[2] The first commitment period covered the years 2005 to 2007 and was considered a test period. The second period coincides with the Kyoto period and lasts from 2008 to 2012. Finally, the post-Kyoto period covers the years 2013 to 2020. The three commitment periods are referred to as Phases I, II, and III.

The EU-ETS is designed as a cap-and-trade system in which the cap defines the number of EUAs available on the market, and hence, determines EUA supply. According to Article 9 of Directive 2003/87/EC each member state has to set up a National Allocation Plan (NAP) for Phases I and II. The NAPs contain the number of EUAs the member states intend to assign to the production units covered by the EU-ETS and have to be approved by the European Commission. Chevallier (2012) points out that the number of annually issued EUAs has been reduced from 2.2 billion in Phase I to 2.08 billion in Phase II. During Phase III the number of issued EUAs is linearly reduced by 1.74% each year such that 2.039 billion EUAs will be issued in 2013 (see also Resolution 2010/634/EU). A second conceptual change concerns the allocation mechanism. While in Phases I and II at least 90% of the issued EUAs were allocated free of charge, the portion of freely distributed EUAs will be reduced from 80% in 2013 to 30% in 2020 (see Directive 2009/EC/29). The second component of the cap-and-trade system is the opportunity to freely trade issued EUAs on the carbon market.[3]

In accordance with Article 12 of Directive 2003/87/EC companies covered by the EU-ETS have to verify their emissions for each compliance year and surrender the appropriate number of EUAs to the national authorities by April 30. Companies that lack EUAs to fulfill regulatory requirements can purchase EUAs on the markets, while superfluous EUAs can be sold on the markets. Within a given commitment period, companies can also transfer

[1] According to directive 2009/29/EC the European Commission has decided to also cover the aviation sector's emissions by the EU-ETS from 2012 on.

[2] Besides carbon dioxide, the EU-ETS accounts for methane, nitrous oxide, hydro fluorocarbons, perfluorocarbons, and sulphur hexafluoride.

[3] These markets are: ECX (London), NordPool (Oslo), EEX (Leipzig), Eurex (Stuttgart), BlueNext (Paris), EXAA (Vienna) and Climex (Utrecht).

EUAs of a compliance year to the consecutive one, or employ EUAs of the next compliance year to fulfill regulatory requirements of the current year. The transfer of EUAs between Phases I and II is prohibited.

Figure 1.1 shows the development of EUA spot and futures prices during the period April 25, 2005 to April 25, 2011. The bold line refers to the spot price of the EUA contract traded at BlueNext in Paris. The dashed line shows the price development of the rolled-over December EUA futures contract with maturity in Phase II traded at ECX in London.[4,5] Apparently, Figure 1.1 reveals substantial differences in EUA spot and futures

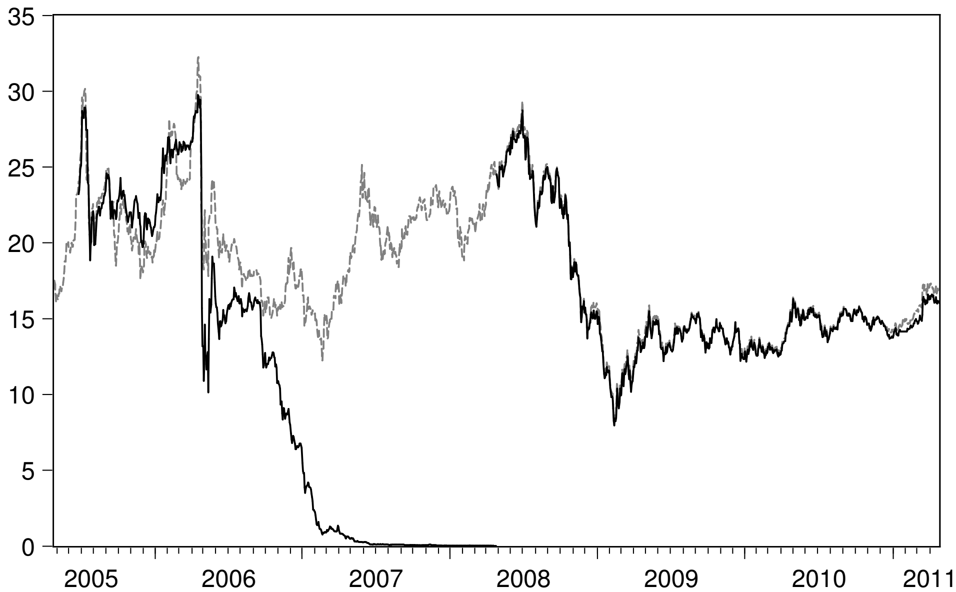

Figure 1.1: Carbon price dynamics from April 2005 to April 2011

price dynamics. These differences were particularly pronounced in Phase I. The most significant break in both markets' price development occurred in the last week of April 2006, when the first EU member states announced that their verified emissions were considerably below the number of allocated allowances, which in turn pointed to a substantial overallocation with EUAs in the first commitment period (see also Alberola et al. (2008)). While futures prices have stabilized immediately after the breakdown and even increased after 2006, spot prices converged to zero towards the end of Phase I. The development of spot prices can be traced back to the prohibition of the utilization of Phase I allowances to

[4]According to Chevallier (2012) BlueNext (ECX) is the most liquid EUA spot (futures) market attracting 72% (96%) of the total EUA spot (futures) market turnover. In contrast to futures market trading at ECX, BlueNext spot market trading started on June 24, 2005.

[5]For a detailed description of the construction of the futures price series see Sections 2.4.1, 3.4.1, and 4.4.1.

satisfy Phase II regulatory requirements. However, even for the period before April 2006, Figure 1.1 exhibits substantial deviations of spot prices from futures prices. Moreover, the prices of both markets prove to be more volatile in the first commitment period than in the second one. With the convergence of spot prices towards zero, trading activity in this market entirely collapsed, while trading activity in the futures market has steadily increased over both commitment periods. In strong contrast to the significant differences in both markets' price dynamics in Phase I, from the start of Phase II on spot and futures prices have obviously evolved very similar throughout the whole compliance period. EUA prices declined strongly during the financial crisis of 2008/2009 and increased slightly during the economic recovery after the crisis.

After the EU-ETS entered into force a new area of economic research called "carbon finance" has emerged. Studies of this research area are concerned with the empirical analysis of EUA price dynamics. While Paolella and Taschini (2008), Benz and Trück (2009), and Daskalakis et al. (2009) focus on the stochastic properties of EUA spot or futures returns, Benz and Hengelbrock (2008) analyze the joint price development of EUA futures contracts traded at different venues. Other studies such as Mansanet-Bataller et al. (2007) and Alberola et al. (2009a) identify fundamental determinants of EUA prices, including energy prices and weather. The analysis of the relevance of regulatory conditions on the EUA price development is provided by Alberola et al. (2008) and Mansanet-Bataller and Pardo (2009), while Oberndorfer (2009) and Veith et al. (2009) analyze the link between EUA prices and the stock market performance of electricity companies. In contrast to the above-mentioned studies which refer to the first commitment period, this thesis predominantly focuses on the econometric analysis of EUA prices during the second period.

The first part of the thesis is concerned with the incorporation and transmission of information in EUA spot and futures prices. Figure 1.1 suggests that regulatory conditions (e.g. the market collapse as a consequence of the overallocation with allowances) as well as general macroeconomic conditions (e.g. the financial crisis and the economic recovery) affect EUA price dynamics substantially. This observation motivates the first research question of whether EUA prices systematically respond to the release of new information concerning such regulatory and macroeconomic conditions. Interesting releases include announcements on the number of EUAs distributed to the EU member states for the second commitment period as well as announcements on key macroeconomic figures reflecting current and future economic activity.

The second research question is directed towards the transmission of information between EUA spot and futures markets. Commencing with the spot market crash in April 2006, Figure 1.1 clearly reveals the decoupling of spot and futures prices that is most pronounced in the end of Phase I. For Phase II the figure detects remarkably similar dynamics of spot and futures prices which points to the existence of a stable long-term relationship between both prices. This in turn raises the second research question of whether one market systematically adjusts prices faster to new information and such sets a price signal to which the other market responds. Put differently, the question is whether one market can be identified as the price-leading market.

In contrast to the first and the second research questions that are related to the investigation of the EUA spot and futures price development, the third research question is directed towards the link between EUA prices and the stock market performance of companies covered by the EU-ETS. Since EUAs constitute an input factor in their production process, EUA price fluctuations could affect the profits of these companies which in turn could have an effect on their stock price. This research question aims at investigating whether companies of different sectors and countries are heterogeneously affected by changing EUA prices and whether the link between EUA prices and the stock market performance is stable over Phases I and II.

While the implementation of the EU-ETS can be rationalized by economic theory, the trading system has been attracting substantial attention in the current political debate, even seven years after its entry into force. In this context, questions related to the appropriateness of the system to assist EU member states in achieving their emission reduction targets, the efficiency of the carbon market, and the overall stringency of emissions caps play a major role in the discussion. This thesis can be considered a contribution to the clarification of these questions.

1.1.2 Oil, stock prices, and the macroeconomy

According to the annual energy review 2010 of the Energy Information Association (EIA), U.S. energy expenditure as share of U.S. GDP has been in the range of 5.9% to 13.7% over the period from 1970 to 2009, in which a considerable fraction of total energy expenditure went back to expenditure on crude oil. Given the high relevance of crude oil for the U.S. economy as implied by these numbers substantial empirical economic research on the link between the oil price and macroeconomic performance as well as stock market performance has been provided. Studies such as Hamilton (1983, 1985, 2003) find a negative effect of oil supply shocks on macroeconomic growth which suggested the conjecture that recessions were caused by oil supply shocks induced by exogenous events in the Middle East. More recently, Kilian (2009) argues that changes in global aggregate demand and changes in oil market specific demand rather than changes in oil supply are the main determinants of the oil price development. In the methodological framework

of Kilian (2009), Kilian and Park (2009) show that an increase in oil prices caused by global aggregate demand induces a positive stock price response, while a negative stock price reaction is caused by an increase in oil prices that can be traced back to oil market specific demand. Consequently, the relation between oil and stock prices may change over time and depend on whether the oil price is predominantly driven by oil market specific demand or by global aggregate demand. Further evidence for time-variation in the oil-stock relation is provided by Miller and Ratti (2009) and Filis et al. (2011).
The evidence of a time-varying relationship between oil and stock prices combined with the relevance of aggregate demand in the determination of the oil price imposes the research question of whether the oil-stock relation is affected by changes in the macroeconomic environment. In this context, the thesis aims at analyzing whether macro variables reflecting current and future economic activity as well as inflation measures can be used to anticipate changes in the oil-stock relation. Moreover, the question of how external events like the invasion of Kuwait and the first Gulf War in 1990/1991 affect the relation between stock and oil prices is addressed in this thesis.
While the previous research question considers the link between oil prices and aggregate stock index prices the second research question is directed towards the effects of the oil price on industry-specific stock returns. Previous studies such as Nandha and Faff (2008) and Narayan and Sharma (2011) find that the oil price affects stock returns across different industries heterogeneously. While increasing oil prices are accompanied by increasing stock prices of oil and gas exploring companies, companies of other industries rather exhibit negative stock price responses to increasing oil prices. Against this background, the thesis is concerned with the question of whether investors holding these stocks require premia for risk associated with the oil price and whether the required risk premia vary over time and across industries.

1.2 Outline of the thesis

The thesis consists of five self-contained research articles and is organized in two parts. Part I contains Chapters 2 and 3; Part II consists of Chapters 4, 5 and 6. Part I assesses the incorporation of information in the European markets for emission allowances at high-frequency. While Chapter 2 investigates the effects of the announcement of macroeconomic and market-specific news on EUA futures returns during the first and the second commitment period of the EU-ETS, Chapter 3 analyzes the transmission of information between the spot and the futures market in the second commitment period. Part II analyzes the link between stock market performance and energy prices. Chapter 4 concentrates on the relationship between stock returns of companies covered by the EU-ETS and the EUA price, while the relation between the U.S. stock market and the oil price is investigated in Chapters 5 and 6.

Chapter 2 is based on the article "Modeling and explaining the dynamics of European Union Allowance prices at high-frequency" which is joint work with my supervisor Prof. Dr. Christian Conrad and my colleague Dr. Waldemar Rotfuß. This article is published in *Energy Economics*.[6] The basis of Chapter 3 is the single-author paper "Price discovery and volatility spillovers in the European Union Emissions Trading Scheme: A high-frequency analysis" published in the *Journal of Banking and Finance*. Chapter 4 is based on the single-author paper "Carbon and the stock market: A policy evaluation of the EU-ETS". The paper is invited for resubmission to *Ecological Economics*. Chapter 5 is joint work with Prof. Dr. Christian Conrad and my colleague Karin Loch entitled "On the macroeconomic determinants of the long-term oil-stock correlations". Chapter 6 is based on the single-author paper "Pricing the risk of oil in the Intertemporal CAPM: An industry-level study".

Explaining EUA prices at high-frequency

Chapter 2 models EUA price returns at high frequency. While previous studies such as Benz and Trück (2009) or Paolella and Taschini (2008) focus on the statistical properties of daily EUA returns, we provide a thorough investigation of such properties at the intraday horizon. At high-frequency we find a seasonality pattern in the second conditional moment of the returns that is very similar to the one of stock returns and exchange rates described in Andersen and Bollerslev (1997, 1998). In a first step, we remove this deterministic pattern from the high-frequency data. Then, we model the conditional mean and the conditional variance equation jointly making use of various GARCH-type models that account for asymmetric effects, power transformations, and the long memory property. We show that the Fractionally Integrated Asymmetric Power GARCH (FIAPGARCH) model of Tse (1998) outperforms more parsimonious specifications obtained by certain parameter restrictions imposed on the FIAPGARCH model in terms of in-sample fit, while restricted specifications that exclusively account for the long memory property (and asymmetric effects) yield the best out of sample volatility forecast accuracy. Second, we investigate the EUA price response to the release of macroeconomic as well as market-specific news. While previous studies predominantly employ simple regressions to capture

[6] Parts of this paper are published in the Ph.D. dissertation "On the High-Frequency Price Reactions of European Union Allowances to News" of Dr. Waldemar Rotfuß accepted by the Rechts- und Wirtschaftswissenschaftliche Fakultät at the Friedrich-Alexander-Universität Erlangen-Nürnberg.

the impact of energy prices, economic activity, and weather conditions on the EUA price at daily frequency (see Mansanet-Bataller et al. (2007) and Alberola et al. (2009a)), we investigate the high-frequency response of EUA returns to the surprise component of price-relevant announcements. Following Andersen et al. (2003) and Conrad and Lamla (2010), we employ the market's expectations obtained from survey data which are subtracted from the realized value of the according macroeconomic variables to construct the surprise components. In addition to macro variables, we investigate the relevance of announcements on the total number of allowances allocated to the EU member states as determined by the European Commission and published in the National Allocation Plans (NAPs). In contrast to survey expectations, we use the expectation formation model of Rotfuß et al. (2009) to construct the surprise component. The empirical analysis reveals that EUA returns instantaneously react to the surprise component in the NAPs. In accordance with economic theory, tighter than expected emissions caps induce an EUA price increase. Moreover, we find that the surprise component of announcements that reflect current and future economic activity in Germany and the U.S. generate instantaneous price reactions which are also consistent with the implications of economic theory. In particular, EUA price increases are induced by higher than expected macroeconomic activity reflecting higher than expected demand for EUAs. This result challenges the finding of the loose connection between macroeconomic variables and the carbon market as pointed out in Chevallier (2009). The results of Chapter 2 imply that EUA prices incorporate new information similar to asset prices of more mature markets.

Price discovery and volatility spillovers

Chapter 3 turns to the analysis of price discovery and volatility spillovers in EUA spot and futures markets during Phase II of the EU-ETS, and hence, investigates the transmission of information in the most liquid spot and futures markets after the spot market re-entered into operation. In contrast to Chapter 2, the analysis of Chapter 3 has to be restricted to Phase II due to the collapse of the spot market in Phase I as a consequence of the overallocation with allowances in this phase associated with the prohibition of the utilization of Phase I allowances to fulfill Phase II regulatory requirements. For Phase I, Benz and Hengelbrock (2008) find a stable long-term relation between the prices of EUA futures contracts traded at different venues and identify the European Climate Exchange in London as the price leading market. Given the cost-of-carry model that relates spot and futures prices of a commodity good by directly assessing the costs and benefits of physically holding the commodity good also implies the existence of such a stable long-term relation between EUA spot and futures prices. Nevertheless, the results of the previous literature are mixed. Studies such as Uhrig-Homburg and Wagner (2009), Milunovic and Joyeux (2010), and Chevallier (2010) analyze the link between EUA spot and futures prices at daily frequency. In contrast to these studies, we assess both markets' contribution to the price discovery process at the intraday frequencies of 10 and 30 minutes making use of the common factor measures based on vector error correction models (VECM) as proposed by Schwarz and Szakmary (1994), Gonzalo and Granger (1995), and Hasbrouck (1995). We clearly reveal the existence of a stable long-term relation between EUA spot and futures prices. In line with Hasbrouck (1995) and Theissen (2002) we show that the analysis of the price discovery process should take place at highest frequencies due to increasing contemporaneous correlation between both markets' residuals at lower frequencies which induces an identification problem. For the early stage of the second commitment period we find that about 70% of the price discovery process can be attributed to the futures market. During the more mature stage of Phase II the relevance of the futures market's informational role even increases. For the second step, we remove the intraday seasonality pattern from the residuals of the VECM (see also Chapter 2). In addition to the investigation of information transmission in EUA spot and futures prices, we analyze the structure of volatility spillovers. Put differently, we assess the transmission of information in the second conditional moment. For this we estimate a dynamic version of the unrestricted extended CCC-GARCH model introduced by Conrad and Karanasos (2010). We reveal close links between the conditional volatilities of both markets which clearly contradicts the Phase I results of Milunovich and Joyeux (2010) who observe a weak link between both markets' uncertainties employing daily data. In particular, we find unidirectional volatility spillovers from the futures to the spot market, while both markets' uncertainties are affected by lagged shocks in the respective other market. In conclusion, our results clearly point to the existence of a close link between the spot and the futures market. Even for the period directly after the spot market re-entered into operation the data provide evidence for a structure of informational spillovers between the spot and futures market that is similar to the one observed for mature markets as reported by Tse (1999) among others.

Carbon and the stock market

In Chapter 4 we empirically analyze the impact of the EU-ETS on the stock performance of companies covered by the trading system. Oberndorfer (2009) and Veith et al. (2009) reveal a positive link between the carbon price and the stock returns of European electricity companies during the first commitment period. Both studies focus on electricity companies since more than 60% of the allowances available on the market have been allocated to the electricity sector. Because of this, the authors argue that the restriction to electricity companies should maintain relevance for the whole EU-ETS. In contrast to these studies, we conjecture that the generalization of the results on the relation between electricity stocks and the carbon price to companies of other sectors could induce spurious or even misleading conclusions. Kettner et al. (2008), Ellerman and Joskow (2008), and Convery et al. (2008) reveal considerable differences in sector-specific net-compliance positions. On average, electricity companies are net-short in allowances, whereas non-electricity companies have received more allowances than needed to fulfill regulatory requirements. In general, companies with net-long positions should rather benefit from increasing prices since superfluous allowances can be sold on the market at higher prices implying the realization of higher regulatory profits. On the other hand, net-short companies have to purchase additional allowances on the market which negatively affects their profitability as far as they are not able to pass through the full costs to consumers. Given these findings, we also consider companies operating in the sectors iron and steel, pulp and paper, chemicals, and cement and lime. In contrast to Oberndorfer (2009) and Veith et al. (2009) who focus on Phase I, we explicitly investigate whether the link between the carbon and the stock market has changed over Phases I and II. Further, given the findings of considerable heterogeneity in country-specific net-compliance positions pointed out in Kettner et al. (2008), we also analyze whether the effect of the carbon price on stock performance is different across individual member states. In particular, the results of Chapter 4 allow us to draw inference on the transmission of regulatory burden to the shareholders of companies under the EU-ETS. We reveal a rather loose relationship between the carbon and the stock market during Phase I, where only electricity stocks are affected by changes in the carbon price. Consistent with Oberndorfer (2009) and Veith et al. (2009), we find that increasing carbon prices are accompanied by electricity stock price increases. The main contribution of this chapter to the existing literature is the identification of a close relationship between the carbon and the stock market during Phase II, where the structure of the relationship across the individual sectors tends to reflect the sector- and country-specific net compliance positions. Stock returns of electricity companies located in countries with more restrictive emissions caps are negatively affected by increasing carbon prices. In strong contrast, for non-electricity companies we reveal that stock price increases are accompanied by increasing carbon prices. The effects are stronger in sectors characterized by more generous allowance allocation. We also find that the effect tends to be more pronounced in countries with less restrictive emissions caps. In conclusion, the results are consistent with the view that investors consider the carbon price as a relevant pricing factor during the second commitment period.

Long-term oil-stock correlations

In Chapter 5 we analyze the relationship between the U.S. stock market and the oil price. Hamilton (1983, 1985, 2003) argues that during the 1970s and 80s the oil price has strongly been affected by oil supply shocks which preceded 9 out of 10 recessions in the U.S. after World War II. According to Hamilton (2003) such supply shocks were caused by exogenous events in the Middle East including the oil crises of 1973 and 1979 as well as the Iraq-Iran War (1980-1988) and the Gulf War (1990-1991). In contrast to these studies, Kilian (2009) argues that even during the 1970s the oil price has mainly been driven by oil market-specific demand shocks, rather than by supply shocks. In line with Harris et al. (2009), Kilian (2009) finds that shocks in the global demand for industrial goods predominantly have determined the oil price development since the second half of the 1990s. Therefore, Kilian and Park (2009) argue that explaining the impact of the oil price on stock prices by regressing stock returns on oil returns could induce spurious conclusions because of reverse causality since the global demand for industrial goods is also a driving factor of U.S. stock prices. Moreover, the authors argue that not only the magnitude but even the sign of the stock price response to an oil price shock depends on the type of the specific underlying shock. While oil price shocks caused by shocks in the global demand for industrial goods induce positive stock price reactions, a negative stock price response is observed in case that the oil shocks can be traced back to oil market-specific demand shocks. To reconsider the oil-stock relation, we develop an econometric specification in the general framework of the Dynamic Conditional Correlation MIxed DAta Sampling (DCC-MIDAS) approach of Colacito et al. (2011) that allows us to distinguish between daily fluctuating short-term correlations and slowly moving long-term correlations which in turn are explicitly linked to the macroeconomic environment. We endogenize the long-term correlation between crude oil and stock returns with respect to economic activity. We consider macro variables that reflect the current stance of the economy, the future economic outlook, and the inflation dynamics. In addition to the long-term correlation analysis, we employ the MIDAS approach of Engle et al. (2009) to reveal the link between long-term oil market volatility and the macro environment including the same variables used in the correlation analysis. We reveal that the long-term oil market volatility is counter-cyclically linked to the macroeconomic activity and can be well anticipated by key macroeconomic figures. We also show that such figures affect oil market and stock market long-term volatility similarly. Current or expected economic expansions (contractions) predict decreases (increases) in both markets' volatilities. The main contribution of Chapter 5 is the identification of a counter-cyclical relationship between the long-term oil-stock correlation and macroeconomic activity that is driven by those macro variables that also anticipate oil and stock market long-term volatility. During recessions and the early phase of the economic recovery the long-term correlation is positive, while it is negative in periods of economic expansions. We argue that during recessions declined profit expectations induce the depreciation of stock prices while the contraction in aggregate demand generates negative oil price movements. This in turn explains the positive correlation in this phase. With the economic recovery stock prices increase due to improved profit expectations, while stimulated aggregate demand induces the oil price to rise such that the long-term correlation remains positive. During economic expansions characterized by strong growth above trend increases in the oil price induce higher production costs which negatively affect the profit expectations. This negative effect overcompensates the positive effect of simultaneously increasing oil and stock prices observed for the early phase of the expansion and such forces the oil-stock correlation to turn negative.

Pricing the risk of oil

Motivated by the findings of Chapter 5, Chapter 6 investigates the relationship between the U.S. stock market and the oil price from an empirical finance viewpoint. In contrast to previous studies such as Jones and Kaul (1996) or Driesprong et. (2008) that use ad hoc specifications to reveal the effects of the oil price on the stock market, we take the Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973) as the theoretical fundament of our empirical analysis. The ICAPM models the expected excess return on a risky asset as a linear function of the covariance between the excess returns on the risky asset and the market portfolio and the covariance between the excess return on the asset and changes in state variables that are related to changes in the investor's future consumption. Merton (1973) shows that the sign of the impact of the covariance with the state variable is determined by the elasticity of the investor's marginal utility of wealth with respect to the state variable. For the oil price this sign is unclear ex ante. However, the empirical literature implies two competing views. Using the structural decomposition of the global oil price suggested in Kilian (2009), Kilian and Park (2009) show that the stock price response to oil price shocks depends on the type of the specific oil price shocks. Aggregate demand shocks lead to increasing stock prices, while shocks related to oil market-specific (precautionary) demand induce stock price contractions. Hence, the sensitivity of the investor's marginal utility of wealth with respect to the oil price depends on the dominance of the shocks within the period investigated. Our empirical analysis reveals that the marginal utility of wealth and the oil price are inversely related which confirms the finding of Kilian and Park (2009) that mainly aggregate demand shocks have driven the price of oil. Moreover, the results reveal that the investor sacrifices some expected return for holding stocks that tend to pay off when the oil price is low and the marginal utility of wealth is high. We identify that such stocks belong to the industries consumer staples, consumer discretionary, health care, industrials, and financials. However, we also reveal that the picture changes dramatically after the peak of the financial crisis in September 2008. The correlations between oil and stock price changes turn positive across all industries. This generates positive expected risk premia for stocks of all industries which is consistent with the view that none of the stocks can be used to hedge against decreasing oil prices after September 2008. Besides the specific patterns of the expected risk premia associated with the oil price, we investigate the relationship between expected return and market risk. Our results suggest a positive risk-return relation. We show that the dynamic evolution of the market risk premia is pretty similar across the individual industries, while their magnitude differs considerably. Finally, we reveal that the conditional covariation between stock returns and changes in macroeconomic risk factors induce statistically significant risk premia, where in particular, the signs of the premia are consistent with those implied by economic theory.

Part I

High-frequency analysis of the European carbon market

Chapter 2

Explaining EUA prices at high-frequency

2.1 Introduction

In this article we analyze the high-frequency price dynamics of European Union Allowances (EUAs) traded on the European Union Emissions Trading Scheme (EU ETS). Our contribution to the literature on the modeling of the EUA price dynamics is twofold. First, we establish that EUA high-frequency return data are characterized by a distinguished intraday seasonality pattern in their second conditional moment. Such intraday seasonality has been proven to exist in stock returns and exchange rates (see Andersen and Bollerslev (1997) and Andersen and Bollerslev (1998)), but the finding of such a pattern which is linked to the intensity of the intraday market activity is novel for EUA returns. Hence, a meaningful econometric analysis of the data requires a preliminary step in which the returns are filtered in order to remove the seasonality. The autocorrelation function of the absolute values of the appropriately filtered return series is then shown to display a very slow decay behavior, which is typical for time series obeying long memory in their conditional second moment. In addition, there is clear evidence for heteroskedasticity and asymmetric responses to positive and negative shocks in the conditional variance. We find that a fractionally integrated asymmetric power GARCH (FIAPGARCH) model is best suited to capture all the stylized facts in the high-frequency EUA returns. This model has been suggested by Tse (1998) and combines the long memory property of the fractionally integrated GARCH (FIGARCH) specification of Baillie et al. (1996) with the asymmetric power GARCH (APGARCH) model of Ding et al. (1993). While previous research on EUA prices had already established the conditional heteroskedasticity in daily returns (e.g. Benz and Trück (2009), Paolella and Taschini (2008) and Chevallier (2009)), the finding of long memory in intraday returns is novel to this article. Since the FIAPGARCH model nests several other GARCH specifications under certain parameter constraints, we can use standard information criteria and likelihood ratio tests in order to rank the competing models. We clearly show that the long memory FIAPGARCH specification outperforms the short memory GARCH models in terms of in-sample modeling performance. In addition, the superiority of the long memory models is also confirmed out-of-sample by means of a volatility forecast comparison. Since accurate volatility forecasts are crucial in many financial applications such as option-pricing or risk management, the explicit modeling of the long-memory property is of fundamental economic importance.

Second, we provide a detailed analysis of the real-time response of EUA prices to the releases of major macroeconomic announcements. The previous literature has identified political and institutional decisions on the overall cap intensity, economic activity, energy prices and temperature as the main EUA price drivers. For a comprehensive survey on current research we refer to the overview article by Zhang and Wei (2010). Our contribution has many distinguishing features. First of all, the previous literature was entirely based on the analysis of daily data. However, since the response to news usually occurs very quickly in financial markets, our high-frequency perspective appears to be more appropriate. Further, for measuring the strength of the response of EUA prices to the release of new information, we construct surprise variables which are based on the difference of the actual figures and the market's expectations. The expectations data are obtained from surveys among market participants. This approach to measure announcement effects is commonly used in e.g. exchange rate markets (see Andersen et al. (2003) and Conrad and Lamla (2010)). In contrast, previous articles on the link between EUA prices and economic fundamentals have either exclusively made use of the actual figures, but did not take into account expectations, or simply employed dummy variables which indicated the occurrence of certain news events, but did not control for the specific content of the news. In the empirical analysis we focus on the releases of *i)* the European Commission's (EC's) decisions on second National Allocation Plans (NAPs), *ii)* macroeconomic indicators about the future economic outlook and *iii)* figures about the current economic stance. We show that EUA prices react most strongly to the EC's decisions on second NAPs. The price adjustment occurs immediately after the release of the new information and is highly significant. The direction of the price adjustment is in line with what economic theory would suggest: a higher than expected allocation of emission rights leads to a fall in the EUA price. This finding extends and complements the results in Mansanet-Bataller and Pardo (2009) and Rotfuß et al. (2009).

In addition, we find that EUA prices react to the releases of macroeconomic figures.

Announcements related to the current economic activity as well as the future economic development in Germany and the U.S. induce significant and immediate EUA price reactions. Positive news which indicate higher than expected economic activity lead to the expectation of increasing demand for emission allowances and, hence, an increase in the EUA price. This finding is of central importance because previous studies based on daily data, such as Chevallier (2009), have concluded that the carbon market is only remotely connected to macroeconomic variables. In contrast, our results based on a high-frequency framework provide strong evidence for the existence of a link between EUA prices and variables related to macroeconomic performance. In response to surprise macroeconomic announcements, EUA prices typically react on impact, i.e. within a few minutes after the release. The finding that mainly German announcements lead to significant price reactions can be explained by the observation that the German economy is the largest within Europe, highly industrialized and, thus, highly emissions-driven. Finally, the significance of the U.S. announcements can be rationalized by the fact that financial market participants preferably use U.S. figures to gauge the future development of the European economy and, thereby, the future demand for European emissions allowances.
The remainder of this article is organized as follows. Section 2.2 discusses the main features of the EU ETS. Related literature is reviewed in Section 2.3. Section 2.4 describes the data and the econometric models. Section 2.5 presents the empirical results, while Section 2.6 concludes.

2.2 The European Union Emissions Trading Scheme

In 2003 the European Union (EU) established a scheme for greenhouse gas emission allowance trading. The scheme is substantially larger and far more complex than the pioneering U.S. system for sulfur dioxide. It is based on the Directive 2003/87/EC and formally entered into operation in January 2005; ten years after the U.S. predecessor began operating. The purpose of the European trading scheme is to promote reductions of greenhouse gas emissions in a cost-effective and economically efficient manner. It aims to assist EU Member States (member states in the following) in meeting their commitments under the Kyoto Protocol at minimum costs and has been called the “New Grand Policy Experiment” of market-based policies in environmental regulation (see Kruger and Pizer (2004), for more details). The scheme requires selected industrial units to participate in the trading of emission allowances. The program covers carbon dioxide emissions from four broad sectors: energy, production and processing of ferrous metals, minerals, and other energy-intensive activities (in particular production of pulp and paper). One emission allowance grants the participating installation (or some other holder of it) the right to emit one metric tonne of carbon dioxide equivalent (tCO2e) during a specified commitment phase. For a legal description of the EU ETS, see European Parliament and Council (2003).

The EU ETS is divided into three commitment phases (Phase I: 2005-2007, Phase II 2008-2012, Phase III: 2013-2020) and runs on the basis of a "cap-and-trade" system. The EU ETS emission cap is defined for each commitment phase by the so called "National Allocation Plans". We term these various plans as first, second, or third NAPs according to the commitment phases. The NAPs are defined by each member state and contain both the national total of allowances as well as a rule for distributing the allowances to the participating installations. The EC approves each NAP and thereby sets the EU ETS emission cap. In total there are 27 NAPs and 27 decisions of the EC on the first and second NAPs, respectively.[1] The allowances are grandfathered or auctioned, with grandfathering having been the most common allocation rule in the first two phases. According to European Parliament and Council (2009) auctioning should be the basic principle for allocation from 2013 onwards. The allowances are freely tradable after they have been allocated to the participating installations.

The participating installations are required to verify their emissions and to surrender the equivalent number of EUAs or other eligible instruments to a national competent authority on an annual basis. Installations which have spare allowances can sell them on the market. Inversely, any installation which lacks allowances has to purchase them from other installations or market participants.

Trading in emission rights takes place on organized markets and over-the-counter (OTC). The trading in EUAs is not specifically regulated or supervised by the EC, although it sets the framework. Trading is regulated by the member states and their national regulating authorities. The most liquid EUA spot market is BlueNext in Paris, which attracts approximately 70 percent of the total daily turnover of the whole organized spot market. Besides an active EUA spot market there is also a vital derivatives market, where futures, options, and other derivatives on EUAs are traded. The most liquid futures market is ICE Futures in London, which absorbs circa 90 percent of the daily turnover in EUA futures. The trading rules on all organized EUA spot and futures markets are largely identical.

[1]Note that the national totals in Phase III will decrease linearly from the average national quantity of allowances in Phase II without any additional EC approval. For more details, see European Parliament and Council (2009).

2.3 Related literature

This article can be considered part of a relatively new research area called “carbon finance”. The recent papers by Paolella and Taschini (2008), Benz and Trück (2009) and Daskalakis et al. (2009) primarily focus on stochastic properties of EUA prices at a daily frequency. Among other things, these authors provide evidence for conditional heteroskedasticity in daily EUA returns. In another stream of articles the authors try to establish which fundamentals are the main price drivers in the EUA market. The role of regulatory issues has been considered by Mansanet-Bataller and Pardo (2009) and Alberola et al. (2008) who find that the approval of the overall cap or the verification of actual emissions significantly affect the EUA price. Other fundamental factors such as energy prices, weather or the overall economic activity are analyzed, among others, in Mansanet-Bataller et al. (2007) and Alberola et al. (2009a). Their results suggest that EUA prices are closely connected to energy markets, in particular to electricity, gas and crude oil prices. In addition, Alberola et al. (2009b) conclude that sectoral production also significantly determines the EUA price. On the contrary, it has been argued by Chevallier (2009) that macroeconomic risk factors, such as the default spread, short-term interest rates or selected market portfolios are only loosely related to the EUA price.
Despite the growing interest in “carbon finance”, very few studies focus on the relation between the EUA price and its fundamentals at high-frequency. The first work in this direction has been done by Benz and Hengelbrock (2008) and Rotfuß (2009), where the former analyze the joint development of two different exchange-based EUA price series and the latter provides selected features of the intraday price formation and volatility in the EU ETS. More recent studies, for example Rittler (2012), focus on the relation between EUA spot and futures prices at high-frequency. The present article builds on the study of Rotfuß et al. (2009) in which the relation between EUA prices and one fundamental factor, namely the determination of the overall supply of EUAs in the second phase is analyzed.

2.4 Data and methodology

2.4.1 Data

Price data

In the empirical analysis we employ high-frequency price data for second-phase EUAs which were obtained from the ICE Futures/European Climate Exchange (ECX), the leading exchange for trading in EUA futures. We focus on price series of the EUA futures contracts maturing in December 2008 (from 01/11/2006 to 15/12/2008), in December 2009 (from 16/12/2008 to 15/12/2009), and in December 2010 (from 16/12/2008 to 09/07/2010). Data prior to November 2006 is not considered due to low liquidity in these instruments. Each EUA futures contract has a maturity of more than three years. They were launched in 2005 and matured on 15/12/2008, 14/12/2009, and 15/12/2010, respectively. In total, we consider 931 trading days, whereby we restrict the analysis to on-exchange transactions. Trading in ICE Futures takes place every working day between 7:00 and 17:00 GMT. The raw data files contain a total of 1,012,646 irregularly spaced transaction records. Each transaction record consists of the transaction price, the transaction volume and the corresponding time stamp (measured up to the second) in GMT. In order to explore the intraday price dynamics, we transform the irregularly spaced transaction prices to equidistant price series. Figure 2.1 displays the equidistant 10-minute EUA prices for the period under consideration. The frequency of the series used in our analysis is chosen to be $h = 10, 30$ and 60 minutes, with 7:00 GMT being the first equidistant point in time. At each equidistant point in time the corresponding price is calculated as the mean of the preceding and the immediately following price, unless there is a transaction at the equidistant point itself. If there is no transaction at 17:00 GMT, the last equidistant price equals the last recorded transaction price. To avoid overnight effects, we do not take the mean of transaction prices of two consecutive trading days.

Figure 2.1 about here.

Equidistant returns are constructed from the price series as follows

$$R_{t,k}(h) = 100 \times (\log(P_{t,k}(h)) - \log(P_{t,k-1}(h))), \quad t = 1, ..., T \quad \text{and} \quad k = 1, ..., K(h),$$

where $P_{t,k}(h)$ represents the equidistant EUA price at the end of the k-th interval at day t given frequency h. T is the total number of trading days and $K(h)$ the number of equidistant intervals per trading day. $P_{t,0}(h)$ is defined as the last equidistant price on the preceding trading day $t-1$, unless there is a transaction exactly at 7:00 GMT.[2]

At all three frequencies, the descriptive statistics (not reported) reveal that the EUA returns have a mean which is not significantly different from zero, are slightly skewed to the right and have a kurtosis that is considerably greater than three. The Jarque-Bera statistics reject normality and the outcome of the Engle LM tests suggest that there is conditional heteroscedasticity. As often reported for high-frequency data, there is some

[2]In addition to returns we also calculate the transaction volumes at the three intraday frequencies. In the empirical analysis the transaction volumes will be used as a predictor for the market volatility.

evidence for serial correlation at low lags in the high-frequency returns, possibly due to microstructure effects. In sharp contrast, Figure 2.2 shows that the absolute returns are highly correlated even for long lags.

Figure 2.2 about here.

The figure depicts exemplarily the sample autocorrelation function of the 10-minute absolute EUA returns for five consecutive trading days. The intraday periodic pattern is clearly observable. Due to the 10-minute frequency and a trading session of ten hours, we have 60 equidistant intervals per trading day. The figure reveals a peak at the first lag and a fast decay of the sample autocorrelation function within the first half of the trading day (up to lag 30). After lag 30, the autocorrelation begins to increase towards a second peak at the beginning of the next trading day (lag 60). The same seasonality pattern is observed for the following days, whereas the amplitude of the subsequent peaks is slowly decreasing. The pattern results from the intraday seasonality due to the time-varying intensity of market activity which is illustrated in Figure 2.3.

Figure 2.3 about here.

The figure shows the average absolute returns, $f_k(h) = \frac{1}{T}\sum_{t=1}^{T}|R_{t,k}(h)|$, for each interval $k = 1, ..., 60$. The average absolute returns are high at the beginning of the trading session and then decrease until mid-day. After mid-day, the average absolute returns again increase slightly. Although less pronounced, this pattern of the absolute EUA returns resembles the typical intraday U-shaped pattern observed in other financial markets (see, e.g., Andersen and Bollerslev (1997)).

There are several ways to deal with the intraday seasonality (see, e.g., Martens et al. (2002)). A simple but very effective method is to standardize $R_{t,k}(h)$ according to the following rule:

$$r_{t,k}(h) = \frac{R_{t,k}(h)}{f_k(h)}.$$

The standardization simply scales each return $R_{t,k}(h)$ by the average absolute return of the interval k. Figure 2.4 illustrates the effect of the filtering by displaying the sample autocorrelation function of $|r_{t,k}(h)|$ for five consecutive trading days, where again we fix the frequency at $h = 10$. As evident from the figure, the sample autocorrelation function does not exhibit any remaining seasonality. In fact, it decreases smoothly with increasing lags. However, while the autocorrelations initially decay fast they are characterized by an extremely slow rate of decay thereafter. Such behavior is typical for long memory processes and suggests that the volatility of EUA returns should be modeled as a fractionally integrated process (see Andersen and Bollerslev (1997)). In order to formally test the hypothesis of long memory and fractional integration we apply the R/S statistic suggested by Lo (1991) as well as the Geweke (1983) estimator of the fractional differencing parameter d to the squared and absolute filtered return data. The results presented in Table 2.1 show that according to the R/S statistic we can reject the short memory null hypothesis against the long memory alternative for all series. Similarly, the estimated fractional differencing parameters d are significantly greater than zero in all cases.[3]

Figure 2.4 about here.

Table 2.1 about here.

Announcement data

As mentioned before, previous research has identified political and institutional decisions on the overall cap stringency, energy prices, temperature events and economic activity as the main EUA price drivers. In our high-frequency analysis of announcement effects we will focus on regulatory issues and measures of economic activity. Further, among the measures of economic activity we distinguish between those which capture the future economic outlook and those which capture the contemporaneous macroeconomic situation. More specifically, we consider the releases of *i)* the EC's decisions on second NAPs, *ii)* leading economic indicators that are supposed to reflect the views of the market participants on the future economic development, and *iii)* macroeconomic figures that are supposed to capture the real economic activity in the European Union and its biggest members states (Germany, Great Britain, and France). In addition, we also make use of data from those U.S. announcements which have been shown to affect other European assets prices in previous studies (see, e.g., Andersson et al. (2009)). The econometric analysis of the announcement effects is then based on the differences of *realized* and *expected* figures.

Table 2.2 about here.

In total we make use of $i = 13$ announcement series which span the period 11/2006 to 07/2010. The individual observations for each of the announcement series consist of the time stamp of the announcement release (t, k), the realized value $y^i_{t,k}$, and the median

[3]Interestingly, the long memory property has recently been established in the volatility of other energy price series, e.g. crude oil prices (see Kang et al. (2009) and Wei et al. (2010)). Hence, our findings for EUA returns complement the results from this strand of literature.

$\widehat{y}^i_{t,k}$ of the corresponding expectations of the market participants. With the exception of the decisions on second NAPs, all realized and expected data were obtained from the Forex Factory (*forexfactory.com*, FF in the following) database. FF compiles the expectations either from the Bloomberg or from the Reuters press releases, which are publicly available shortly after the announcements. The expectation data are consensus forecasts, which Bloomberg and Reuters obtain by means of surveys few days prior to the announcements. Since there are no expectations data available regarding the EC's decisions on second NAPs, we construct the expectation variable by assuming that market participants anticipate the EC's decisions by using the national total in Phase I as a reference point. The expectation for a certain member state is set equal to the number of EUAs submitted to the EC if this number is below a lump sum cut amount of the national total approved in Phase I and equal to the lump sum cut amount otherwise. Assuming that expectations are formed in such a way is reasonable since it was clear to all market participants that the EC will only allow tight caps in the second commitment period. For the subsequent analysis we have chosen the lump sum cut to be 7.5 percent.[4]

An overview of the announcement data along with a test for the unbiasedness of the expectations is provided in Table 2.2. The standard procedure to test for the unbiasedness of expectations is to run a linear regression of the realized value on an intercept and the expected value. The expectations can be assumed to be unbiased if the estimates for the intercept and the slope coefficient are not jointly significantly different from zero and unity. As can be seen from Table 2.2, in all regressions the R^2 is relatively high and the null hypothesis of unbiased expectations, $H_0 : \beta_1 = 0$ and $\beta_2 = 1$, can only be rejected in one out of the thirteen cases at the 5% level. Thus, we conclude that the expectations data can be considered as being of good quality.

Since the units of measurement differ across the various announcement variables, we follow Balduzzi et al. (2001) and standardize the surprise variables, i.e. the difference between the realized and expected values, as:

$$S^i_{t,k} = \frac{y^i_{t,k} - \widehat{y}^i_{t,k}}{sd^i},$$

where sd^i is the sample standard deviation of the forecast error for announcement i. We refer to the resulting variables $S^i_{t,k}$ as the standardized surprises.

Next, we briefly discuss the expected effects of surprises in the three groups of announcement variables. Since a positive surprise in the NAP represents an unexpected increase

[4]This value was suggested in Rotfuß et al. (2009) and is in line with unbiased expectations (see Table 2.2). Moreover, we have checked the robustness of our results with respect to reasonable variations in the lump sum cut ranging from 5 to 10 percent.

of the national total in Phase II of a member state, we should expect a price decline simply because of increasing supply. Similarly, better than expected figures on the actual economic activity or the future economic development should signal a higher (future) demand for emission allowances and, hence, induce an increase in EUA prices.[5]

2.4.2 Methodology

In order to capture the news effects on EUA prices, we model the continuously compounded EUA returns as a function of their own first lag and the contemporaneous and P lagged values of the standardized surprise variables. To simplify the notation, we now change the index of the return and surprise variables from (t, k) to n, where $n = 1, \ldots, TK$.[6] The mean equation is then given by

$$r_n = \mu + \theta r_{n-1} + \sum_{i=1}^{13} \sum_{p=0}^{P} \theta_{i,p} S^i_{n-p} + \varepsilon_n, \tag{2.1}$$

where the error term ε_n is given by $\varepsilon_n = \sigma_n Z_n$ with $\{Z_n\}$ being a sequence of independent and identically distributed random variables with $\mathbf{E}(Z_n) = 0$, $\mathbf{E}(Z_n^2) = 1$ and σ_n^2 being the conditional variance. Since the empirical autocorrelation function of the absolute filtered returns revealed a clear pattern of long memory and persistence, we follow Tse (1998) and model the conditional variance as a fractionally integrated asymmetric power GARCH (FIAPGARCH$(1, d, 1)$) process given by

$$\begin{aligned}(1 - \beta L)\sigma_n^\delta &= \omega + \sum_{j=1}^{J} \sum_{q=1}^{Q} \omega_{j,q} W^j_{n-q} \\ &\quad + \left((1 - \beta L) - (1 - \phi L)(1 - L)^d\right)(|\varepsilon_n| - \gamma\varepsilon_n)^\delta, \end{aligned} \tag{2.2}$$

where L denotes the lag operator and β/ϕ are the autoregressive/moving average parameters of the variance equation. The fractional differencing parameter $0 \leq d \leq 1$ captures the long memory in the volatility and $\delta > 0$ denotes the optimal power transformation. In addition, the asymmetry term $|\gamma| < 1$ ensures that positive and negative innovations of the same size can have asymmetric effects on the conditional variance. Note that for the conditional variance to be positive almost surely for all n, the parameter combination

[5] Alberola et al. (2009b) argue that the response to surprises in the actual activity may depend on the potential long/short compliance positions held by the industrial sectors. Their results suggest that important industrial sectors had net long compliance positions in the period 2005-2006. In such a situation even a positive surprise on the actual economic activity can lead to declining EUA prices.

[6] Note that we also suppress the reference to the frequency h by simply writing r_n instead of $r_n(h)$.

(β, d, ϕ) has to satisfy the inequality constraints derived in Conrad (2010) and Conrad and Haag (2006). The explanatory variables W_{n-q}^j are the (lagged) values of the filtered transaction volume and dummy variables which indicate whether an announcement takes place or not.[7]

The flexible FIAPGARCH specification nests several standard GARCH models. For $\delta = 2$ it reduces to an asymmetric FIGARCH (FIAGARCH) specification and under the additional constraint that $\gamma = 0$ to the symmetric FIGARCH one (see Baillie et al. (1996)). On the other hand, for $d = 0$ the model reduces to a short memory asymmetric power GARCH (APGARCH) specification and under the additional constraint that $\delta = 2$ to the asymmetric GARCH (AGARCH) model (see Ding et al. (1993)). In this last case, the conditional variance takes the familiar form $\sigma_n^2 = \omega + \alpha(|\varepsilon_{n-1}| - \gamma\varepsilon_{n-1})^2 + \beta\sigma_{n-1}^2$ with $\alpha = \phi - \beta$.

All models are estimated by using the quasi-maximum-likelihood method as implemented by Laurent and Peters (2003) in the G@RCH package for Ox, which allows us to draw robust inference even if the return data are non-Gaussian. Finally, we can use standard information criteria and likelihood ratio tests to discriminate between the most general FIAPGARCH model and the nested GARCH specifications (see also Conrad et al. 2011)).[8]

2.5 Empirical results

2.5.1 Baseline specifications

Estimation at intraday frequencies

We begin the empirical analysis by first estimating the four GARCH specifications at the 10-, 30- and 60-minute frequency without including any explanatory variables. By doing so, we will identify the GARCH specification which is best suited for modeling the high-frequency EUA return data. The results are presented in Tables 2.3, 2.4 and 2.5. At the 10-minute frequency, the autoregressive and moving average parameters are highly significant in all four GARCH specifications. In addition, the asymmetry term is significantly positive in all four models, which suggests that the conditional variance of

[7]As for the return series, we filtered the transaction volume by first eliminating a time trend and then removing the intraday seasonality pattern.

[8]Because of the excessive number of zero returns in the U.S. SO_2 market in the period 1999-2006, Paolella and Taschini (2008) suggest to use mixed normal GARCH specifications to analyze emissions allowances. However, in our data set for Phase II of the EU ETS the liquidity in the market for CO_2 allowances is sufficiently high such that the zero returns problem is not evident (in our sample, the average time between two consecutive trades is 56 seconds). Instead the predominant feature of the return series is the long memory property in the second conditional moment.

EUA returns increases considerably more in response to negative innovations than to positive ones of the same size. The fractional differencing parameter in the FIAGARCH and FIAPGARCH models take values around 0.25 which are significantly above zero. Such a degree of volatility persistence is well in line with the ones observed in other financial markets, such as exchange rates (see, e.g., Conrad and Lamla (2010)). Clearly, the evidence for long memory which was already evident from Figure 2.4 and Table 2.1 is reinforced by our coefficient estimates.[9] The estimated FIAGARCH and FIAPGARCH parameter combinations (β, d, ϕ) satisfy the inequality constraints derived in Conrad (2010) and, thus, guarantee the non-negativity of the conditional variances. Further, the optimal power transformation in the FIAPGARCH model is estimated to be significantly greater than two and, hence, restricting it to two may lead to suboptimal modeling and forecasting performance (see Brooks et al. (2000)). Comparing the Akaike and Schwartz information criteria (AIC and SIC) of the FIAPGARCH model with those from the models which impose the restrictions $\delta = 2$ (FIAGARCH), $d = 0$ (APGARCH) or $\delta = 2$ and $d = 0$ (AGARCH) clearly leads to the conclusion that the most general FIAPGARCH specification is the preferred one. This conclusion is reinforced by the results of the likelihood ratio tests, which clearly reject the restricted models in favor of the most general specification in all cases. Also, the Ljung-Box statistics show that there is no remaining serial correlation in the squared standardized residuals from the FIAPGARCH specification. In sharp contrast, the hypothesis of uncorrelated squared standardized residuals is strongly rejected for the short memory APGARCH and AGARCH models. Finally, in the mean equation the estimated constants are insignificant for all four models, while the first order autoregressive coefficient is significant for the FIAGARCH and FIAPGARCH at the 15% level only.

Table 2.3 about here.

As can be seen from Table 2.4, the empirical results at the 30-minute frequency are quite similar to the ones at the 10-minute frequency. The information criteria and the likelihood ratio tests unanimously favor the FIAPGARCH specification as the best one. At the 60-minute frequency, the FIAPGARCH model still clearly dominates the short memory specifications, but we obtain identical information criteria for the FIAGARCH and the FIAPGARCH model. As can be seen from Table 2.5 the power term δ in the FIAPGARCH model is not significantly different from two and, consequently, the likelihood ratio test does no longer reject the nested FIAGARCH against the FIAPGARCH.

[9] Also note that the sum of the α and β parameters in the two short memory GARCH specifications is close to and not significantly different from one. This feature is often called the 'integrated GARCH (IGARCH) effect' and – as argued in Baillie et al. (1996) – is due to misspecification of the conditional variance equation because of the neglected long memory property.

Table 2.4 about here.

Table 2.5 about here.

In summary, based on the results for the four GARCH specifications the most general FIAPGARCH model can be clearly identified as producing the best fit to the data at the 10- and 30-minute frequencies. At the lower 60-minute frequency the FIAGARCH and the FIAPGARCH model produce comparable fits. In the subsequent analysis of the announcement effects we will employ the FIAPGARCH model as our baseline specification.

On the importance of long memory

The analysis in the previous subsection has clearly shown that the *in-sample* fit of the long memory GARCH specifications is superior to the one of the corresponding short memory specifications. Because volatility forecasts are a crucial input into option pricing formulas and portfolio selection models, we analyze the economic significance of the long memory property by comparing the success of the different GARCH specifications in forecasting volatility *out-of-sample*. The usual applications such as Value-at-Risk computations typically require volatility forecasts at a daily frequency. Hence, we first construct a series of 931 daily EUA returns from our full sample. Then each specification is re-estimated for the period 01/11/2006 to 31/12/2009 (i.e. over the first 800 observations) and out-of-sample volatility forecasts are constructed for the period 01/01/2010 to 09/07/2010. In order to be able to evaluate the forecast performance of the competing specifications at different forecast horizons, we produce 5-, 10- and 15-step-ahead predictions in addition to one-step-ahead forecasts. For the forecast evaluation we construct the daily realized volatilities as the sum of the squared 10-minute intraday returns and then calculate the mean square error (MSE) and the mean absolute error (MAE) statistics over the relevant forecast horizons.

We only briefly comment on the in-sample estimation results for the daily data, because the findings were quite similar to the ones using the intraday data. In particular, we again find strong evidence in favor of long memory, i.e. the fractional differencing parameter is significantly different from zero as well as one in all long memory specifications (see Table 2.6, first column). As at the 60-minute frequency, the estimated power terms are not found to be significantly different from two and, in addition, the significance of the asymmetry term becomes weaker. Because of these two findings we additionally estimated the simple GARCH and FIGARCH models.

The results of the forecast comparison are summarized in Table 2.6. According to the

MSE and MAE statistics, the long memory specifications generate considerably more accurate volatility forecasts than their short memory counterparts. That is, at each forecast horizon the FI(A)GARCH specifications clearly provide the best volatility forecasts, as indicated by both evaluation criteria. More specifically, the FIGARCH model dominates at the 1- and 5-step-ahead horizons and the FIAGARCH one at longer forecast horizons. This finding is in line with the recent evidence in Engle (2010) who argues that modeling asymmetry is particularly important for producing accurate long-run volatility forecasts. The forecast comparison clearly shows that neglecting the long memory property leads to inferior volatility predictions.[10] Hence, the analysis of daily data reinforces the evidence of long memory in the volatility of EUA returns and highlights the importance of this finding for financial applications.[11]

2.5.2 Measuring announcement effects

Announcement effects over different time horizons

We now augment the FIAPGARCH specification from above by the contemporaneous standardized surprise variables from the 13 macroeconomic announcement series in the conditional mean and by additional control variables in the conditional variance. The corresponding results for the three intraday high-frequency series are presented in Table 2.7. We first focus on the releases of the EC's decisions on second NAP's. The reaction to announcements of the EC's decisions is significant at the 1% level at the 10-minute frequency and at the 5% level at the 60-minute frequency. In line with our discussion in Section 2.4.1, a higher than expected allocation of emission rights leads to an immediate decline in the EUA price. The price drop is strongest after ten minutes with a value of -3.1364 and takes a value of -0.6703 after 60 minutes.

Table 2.7 about here.

Among the variables which capture the tendency of the future economic development, the U.S. ISM manufacturer index clearly evokes the strongest price reaction. At the 10- and 30-minute frequencies the effect is positive and highly significant (at least at the 5% level). Similarly, the coefficients for surprises in the DE ifo index and DE new orders are positive and significant (at the 5% or 10% level) at various intraday frequencies. The fact that the estimated coefficients are always positive is in line with the argument that the

[10]Interestingly, the APGARCH model has the worst performance at all forecast horizons. Also the FIAPGARCH model does not perform as well as the long memory models which impose the restriction $\delta = 2$. Thus, for forecasting daily EUA volatility the optimal power transformation appears to be two.

[11]Similarly, Kang et al. (2009) and Wei et al. (2010) provide convincing evidence for the economic implications of long memory in crude oil price volatility.

prospect of stronger than expected economic growth and, hence, rising emissions will lead to an increase in the demand for allowances.[12]

Looking at the variables which reflect the current stance of the economy, we find a similar picture as for the indictors on the future economic development. That is, surprise announcements on German and U.S. variables lead to immediate and significant market reactions. More specifically, the U.S. nonfarm payrolls evoke the strongest price reaction which is highly significant (at least at the 5% level) at all three frequencies. Similarly, the surprises in the DE industrial production induce significant price reactions at the 10- and 30-minute frequency (at the 5% and 15% level, respectively). In all cases, higher than expected production levels lead to increasing EUA prices.

Note that we augmented the conditional variance equation by the first lag of the filtered trading volume and a binary variable (EU NAP dummy) which indicates whether a NAP announcement takes place or not. At the intraday frequencies the lagged trading volume has a positive and highly significant (at the 1% level) effect on the conditional variance, i.e. higher trading volume leads to more volatility. This finding is in line with the sequential information arrival hypothesis as first developed by Copeland (1976) and empirically supported by Darrat et al. (2003) among others. On the contrary, the binary NAP variable is weakly significant (at the 10% level) at the 10-minute frequency only. This is a remarkable finding because it suggests that the NAP announcements significantly affect the level of the EUA prices with almost no effect on volatility.[13] Further, note that the estimated values of the structural coefficients in the conditional mean and variance equations of the three high-frequency specifications are almost identical to the corresponding ones that we obtained for the baseline specifications. Also, based on the Ljung-Box statistics there is no evidence for misspecification in the squared standardized residuals.

In summary, we find that the NAP approvals have an immediate, long lasting and by far the strongest effect on EUA prices. In addition, German and U.S. figures on the future economic development as well as on the contemporaneous economic activity have immediate – but compared to the NAP announcements – rather short-lived effects on the EUA price process. In line with the argument that higher than expected contemporaneous/future economic activity will increase the demand for EUA allowances, the estimated reaction coefficients for those variables are all positive. Interestingly, we find no significant

[12] The only exception is the DE ZEW index for which the results are mixed: a negative coefficient at the 60-minute frequency and a positive coefficient at the 10-minute frequency (significant at the 10% and 15% level, respectively).

[13] We also experimented with dummy variables for the other announcements. However, none of these were found to be significant.

effects in response to the EU, French or British announcements. The fact that German and U.S. announcements are predominant may be explained as follows: i) Germany is the largest European economy, heavily industrialized and, thus, highly emissions-driven and ii) financial markets typically use information about the U.S. economy as a predictor for the future development of the European economy.
Finally, we would like to note that at a daily frequency none of the macroeconomic indicators appear to evoke significant price reactions (results not reported). This observation squares with the findings in Chevallier (2009) and at the same time highlights the importance of taking a high-frequency perspective. In contrast, even based on daily data the negative effect of the NAP announcements is still observable and significant at the 5% level.

Lagged announcement effects

In the previous subsection we investigated contemporaneous announcement effects but did not take into account the complicated lead lag structure which is often observed in financial markets in response to the release of macroeconomic news. Hence, in a second step we rerun the regression at the 10-minute frequency, but now include several lagged values of the surprise variables. The results for this extended model are presented in Table 2.8. We include up to 3 lags of the standardized surprise variables such that a 40-minute period is covered.[14]

Table 2.8 about here.

As before, the overall reaction to the NAP announcements is negative. However, it is interesting to note that after the strong price decline within the first ten minutes, there is a significant price reversal which is then followed by a final price decline. This interesting observation cannot be made by simply looking at the reaction at different frequencies (as is done in Table 2.7), but requires the inclusion of lagged surprises.
The results for the forward looking indicators confirm our previous findings. Again, positive surprises in DE new orders and the U.S. ISM manufacturer index lead to increasing EUA prices on impact. Both effects are highly significant. In line with our previous results, the coefficient estimates for the DE Ifo index suggest that the reaction to this indicator is somewhat delayed. Apart from the second lag of the DE ZEW index, all significant coefficients are of the expected sign and, hence, imply positive price reactions to positive surprises.

[14]Including additional lagged values did not change our conclusions in a significant way. Of course, the results are available from the authors upon request.

Concerning the variables representing the current economic stance, the results are also in line with those of Table 2.7. Again, macroeconomic announcements related to the economies of Germany and the U.S. affect EUA prices the strongest. Positive surprises in DE industrial production and the U.S. nonfarm payrolls lead to positive EUA price reactions which are immediate and highly significant. In addition, the second lag of surprises in EU industrial production affects EUA prices positively and highly significantly (at the 1% level). No significant effects can be observed in response to surprises in the industrial production of the U.K. and France.

Definitely, the analysis with lagged standardized surprise variables allows us to gain some additional insights into the way the EUA prices adjust to new information. First, EUA prices react most strongly to NAP announcements. However, while the overall reaction is negative there is some price reversal effect at the second lag. This effect might also explain why we do not find a significant reaction to the NAP announcements in Table 2.7 at the 30-minute frequency. Second, immediate reactions are observed in response to German and U.S. macroeconomic announcements on the future economic development as well as the current economic stance. In most cases the new information is priced within the first ten minutes after the release of the announcement. Lagged surprise variables are significant in a few cases only. Hence, for the German and U.S. announcements our results are in line with the findings for other financial markets in which the adjustment to new information typically takes place within a few minutes or even seconds (see, e.g., Andersen et al. (2003)). This is remarkable, since the EU ETS is still a relatively new market. On the contrary, the finding that there are almost no price reactions in response to EU, French and British announcements is rather surprising.

2.6 Conclusions

This article contributes to the steadily expanding literature on the modeling and explaining of the movements in EUA prices. The distinguishing feature of our contribution lies in investigating the price dynamics from a high-frequency perspective. We show that the price dynamics of the EUA futures contracts maturing in December 2008, 2009, and 2010 are very well captured by a fractionally integrated asymmetric power GARCH specification. Thus, we establish that high-frequency EUA returns do not only obey conditional heteroscedasticity, but are also characterized by long memory, power effects and asymmetry in their second conditional moments. The finding of long memory in the volatility of EUA returns complements the recent evidence by Kang et al. (2009) and Wei et al. (2010) of long memory in oil price volatility. Additionally, we have shown that the long memory specifications clearly outperform their short memory competitors in terms of out-of-sample forecast performance. Since volatility forecasts are crucial for risk management as well as optimal portfolio allocations, the long memory property in the volatility of EUA returns is of considerable economic relevance.

Moreover, we investigate how EUA prices respond to the release of the EC's decisions on second NAPs, macroeconomic announcements on the future economic development and the actual economic activity. We find that the reaction to the EC's decisions on second NAPs is immediate and of the expected sign. Similarly, the empirical evidence suggests that German as well as U.S. leading economic indicators, which point towards higher than expected growth in the future, induce an immediate increase in EUA prices. The same results are found in response to German and U.S. announcements on the actual economic activity, i.e. positive news regarding the current stance of the economy generate positive and significant price reactions within the first ten minutes after the release. However, we were not able to establish a link between EUA prices and announcements in other major European countries. This result is surprising and opens an interesting avenue for future research.

Finally, from the size and the strength of the response of EUA prices to the various data releases, we conclude that NAP announcements which directly affect the supply of emission allowances are by far the most important driving force of EUA prices, while on the demand side U.S. nonfarm payrolls, the U.S. ISM manufacturer index, German industrial production as well as German new orders generate the strongest effects on EUA prices.

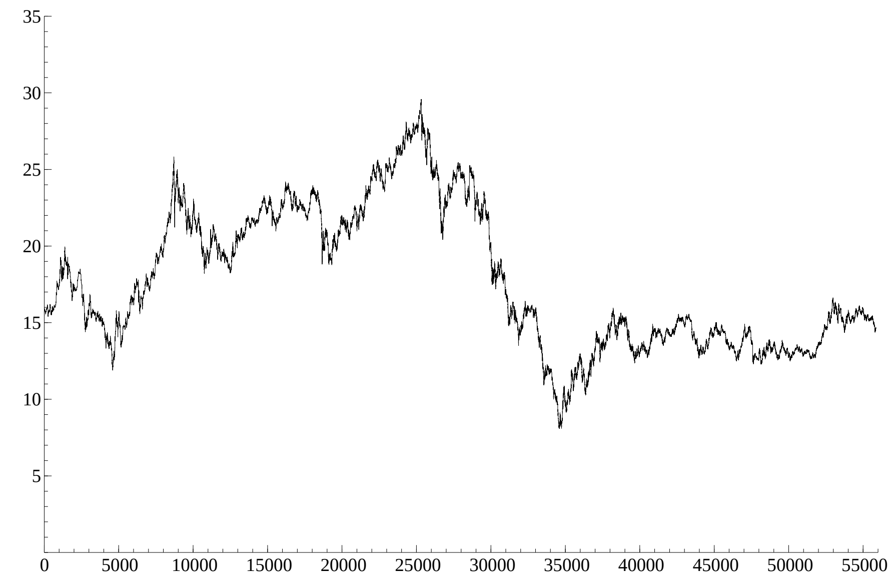

Figure 2.1: Equidistant 10-minute prices of the EUA futures

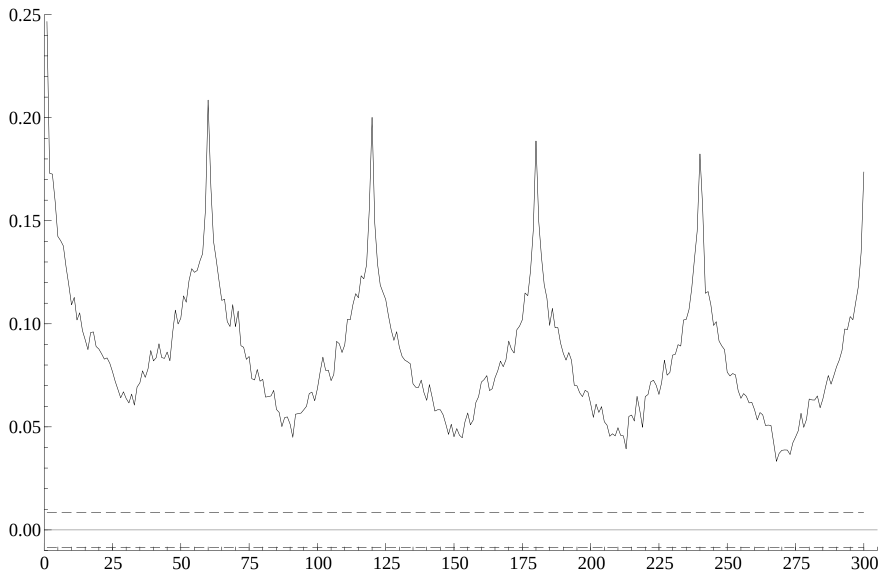

Figure 2.2: Autocorrelation function of absolute 10-minute EUA returns for five consecutive trading days

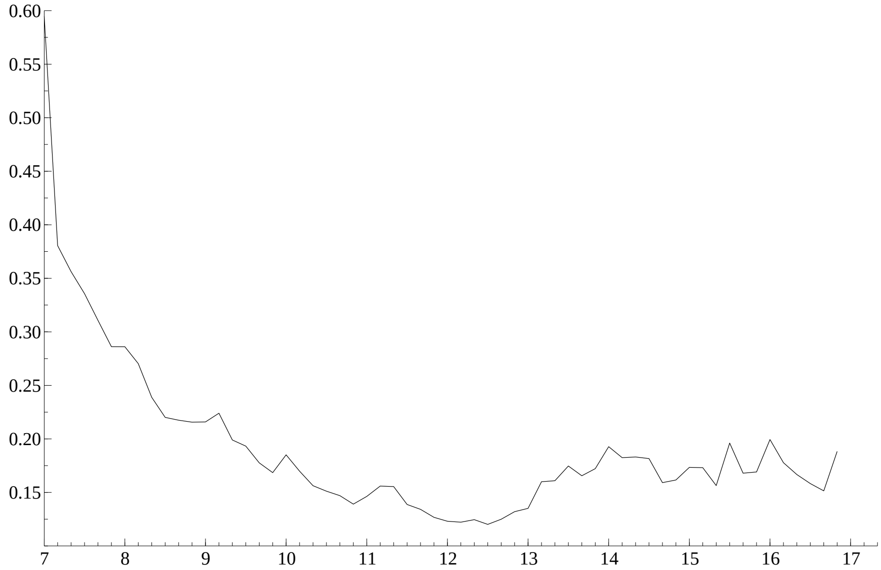

Figure 2.3: Average absolute 10-minute EUA returns for each 10-minute interval during a trading day

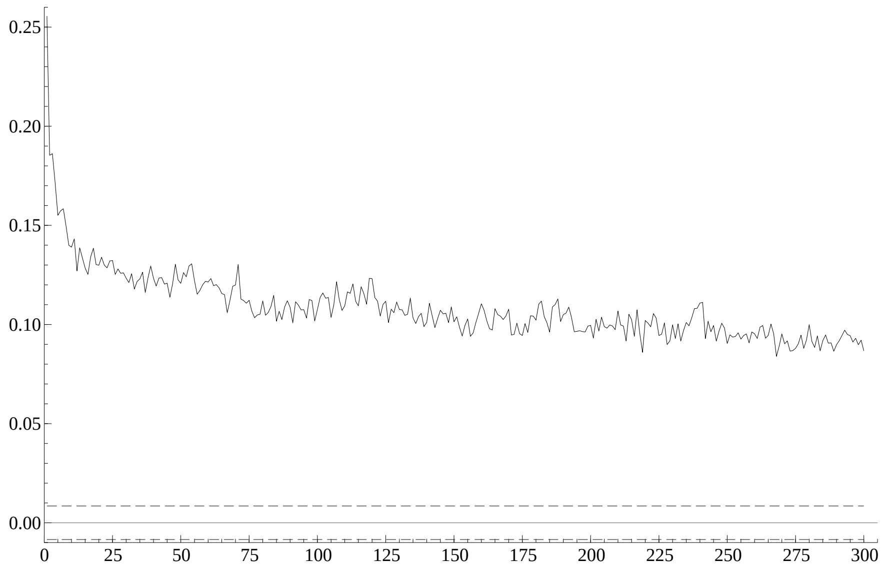

Figure 2.4: Autocorrelation function of filtered 10-minute absolute EUA returns for five consecutive trading days

Table 2.1: Tests for Long Memory in Absolute/Squared Filtered Returns

	$\lvert r_{tk}(10)\rvert$	$r^2_{tk}(10)$	$\lvert r_{tk}(30)\rvert$	$r^2_{tk}(30)$	$\lvert r_{tk}(60)\rvert$	$r^2_{tk}(60)$
R/S	5.3912	5.0109	4.1346	3.0370	3.4485	2.2169
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
d	0.1309	0.3074	0.1647	0.3814	0.1909	0.3127
	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]

Notes: The entries in the first line are the values of Lo's (1991) R/S statistic. The entries in the second line are the Geweke (1983) estimated fractional differencing parameters. The numbers in brackets are p-values.

Table 2.2: Announcement Data and Tests of Unbiasedness of Expectations

Announcement	# obs.	β_1	β_2	R^2	Wald test
Policy Variable					
EU NAP	27	0.434	0.982	0.997	1.75
		(1.397)	(0.011)		[0.194]
Future Economic Outlook					
DE Ifo index	44	0.241	0.998	0.972	0.03
		(2.547)	(0.026)		[0.970]
DE ZEW index	44	0.771	0.964	0.955	0.82
		(1.275)	(0.032)		[0.447]
EU Consumer confidence	44	−0.325	0.992	0.969	0.38
		(0.490)	(0.027)		[0.688]
DE New orders	45	−0.003	0.637	0.169	1.59
		(0.005)	(0.216)		[0.216]
EU New orders	44	0.002	0.987	0.605	0.42
		(0.002)	(0.123)		[0.662]
U.S. ISM manufacturer index	45	1.607	0.971	0.914	0.39
		(2.267)	(0.045)		[0.683]
U.S. Uni-Michigan index adv.	44	3.167	0.946	0.888	1.34
		(3.849)	(0.052)		[0.274]
Current Economic Activity					
DE Industrial production	45	−0.003	1.316	0.484	1.98
		(0.002)	(0.207)		[0.150]
EU Industrial production	44	−0.001	0.944	0.735	0.75
		(0.001)	(0.088)		[0.481]
FR Industrial production	45	−0.002	1.231	0.268	0.71
		(0.002)	(0.310)		[0.499]
GB Industrial production	45	−0.002	1.464	0.503	4.81
		(0.001)	(0.222)		[0.013]
U.S. Nonfarm payrolls	45	−11.916	0.979	0.933	0.66
		(10.916)	(0.040)		[0.520]

Notes: The first two columns contain the name and the number of observations of the announcement series. The third to fifth columns contain the estimates from the regression $y^i_{t,k} = \beta_1 + \beta_2 \cdot \widehat{y}^i_{t,k} + \eta^i_{t,k}$ and the corresponding R^2. The last column shows the results of a Wald test of the joint hypothesis $H_0 : \beta_1 = 0$ and $\beta_2 = 1$. Numbers in parentheses are standard errors, numbers in brackets are p-values.

Table 2.3: GARCH Models at 10-Minutes Frequency

	AGARCH	APGARCH	FIAGARCH	FIAPGARCH
		Mean Equation		
μ	-0.0072	-0.0067	-0.0069	-0.0066
	(0.0055)	(0.0054)	(0.0052)	(0.0052)
θ	-0.0029	-0.0027	0.0086^{+}	0.0082^{+}
	(0.0058)	(0.0057)	(0.0054)	(0.0053)
		Variance Equation		
ω	0.0138^{**}	0.0138^{**}	0.1057^{***}	0.0862^{***}
	(0.0062)	(0.0063)	(0.0153)	(0.0161)
α	0.0317^{***}	0.0304^{***}		
	(0.0089)	(0.0092)		
ϕ			0.4625^{***}	0.4659^{***}
			(0.0395)	(0.0467)
β	0.9628^{***}	0.9625^{***}	0.6427^{***}	0.6102^{***}
	(0.0111)	(0.0111)	(0.0348)	(0.0437)
γ	0.0986^{***}	0.0944^{***}	0.1082^{***}	0.0972^{***}
	(0.0235)	(0.0222)	(0.0222)	(0.0197)
δ		2.1056^{***}		2.1774^{***}
		(0.1055)		(0.0653)
d			0.2923^{***}	0.2515^{***}
			(0.0176)	(0.0236)
AIC	3.565	3.565	3.546	*3.544*
SIC	3.566	3.566	3.547	*3.546*
$Q^2(20)$	163.028	153.724	23.787	20.373
	[0.000]	[0.000]	[0.162]	[0.312]
LR	1143.558	1133.066	68.494	
	[0.000]	[0.000]	[0.000]	

Notes: The numbers in parentheses are Bollerslev-Wooldridge robust standard errors. ***, **, *, $^{+}$ indicate significance at the 1 %, 5 %, 10 % and 15 % level. AIC and SIC are the Akaike and Schwartz information criteria. The numbers in italic letters indicate the model with the smallest value of the information criteria. $Q^2(20)$ is the Ljung-Box statistic for the squared standardized residuals at lag 20. LR is the likelihood ratio test $LR = 2[L_{UR} - L_R]$, where L_{UR} is the likelihood of the unrestricted FIAPGARCH specification and L_R the likelihood of the restricted model. The numbers in brackets are p-values.

Table 2.4: GARCH Models at 30-Minutes Frequency

	AGARCH	APGARCH	FIAGARCH	FIAPGARCH
		Mean Equation		
μ	-0.0081	-0.0076	-0.0054	-0.0053
	(0.0086)	(0.0085)	(0.0084)	(0.0084)
θ	-0.0019	-0.0014	-0.0023	-0.0022
	(0.0088)	(0.0087)	(0.0086)	(0.0084)
		Variance Equation		
ω	0.0133^{***}	0.0137^{***}	0.0757^{***}	0.0591^{**}
	(0.0049)	(0.0050)	(0.0205)	(0.0235)
α	0.0444^{***}	0.0419^{***}		
	(0.0088)	(0.0085)		
ϕ			0.4158^{***}	0.4009^{***}
			(0.0730)	(0.1156)
β	0.9494^{***}	0.9481^{***}	0.6080^{***}	0.5362^{***}
	(0.0108)	(0.0108)	(0.0743)	(0.1252)
γ	0.1067^{***}	0.0985^{***}	0.1421^{***}	0.1279^{***}
	(0.0302)	(0.0299)	(0.0315)	(0.0279)
δ		2.1825^{***}		2.2467^{***}
		(0.1594)		(0.0279)
d			0.3034^{***}	0.2389^{***}
			(0.0225)	(0.0340)
AIC	3.374	3.378	3.364	*3.363*
SIC	3.381	3.380	3.368	*3.366*
$Q^2(20)$	92.137	81.875	24.494	21.797
	[0.000]	[0.000]	[0.139]	[0.241]
LR	280.344	274.336	26.444	
	[0.000]	[0.000]	[0.000]	

Notes: As in Table 2.3.

Table 2.5: GARCH Models at 60-Minutes Frequency

	AGARCH	APGARCH	FIAGARCH	FIAPGARCH
		Mean Equation		
μ	-0.0092 (0.0121)	-0.0103 (0.0124)	-0.0086 (0.0118)	-0.0086 (0.0118)
θ	-0.0027 (0.0119)	-0.0026 (0.0120)	-0.0032 (0.0126)	-0.0032 (0.0126)
		Variance Equation		
ω	0.0138*** (0.0051)	0.0133*** (0.0047)	0.0629*** (0.0207)	0.0621*** (0.0221)
α	0.0450*** (0.0093)	0.0516*** (0.0104)		
ϕ			0.4497*** (0.0874)	0.4504*** (0.0888)
β	0.9480*** (0.0111)	0.9482*** (0.0105)	0.6436*** (0.0767)	0.6415*** (0.0832)
γ	0.1301*** (0.0402)	0.1603*** (0.0531)	0.1638*** (0.0475)	0.1627*** (0.0510)
δ		1.6245*** (0.1670)		2.0116*** (0.1459)
d			0.3177*** (0.0340)	0.3145*** (0.0537)
AIC	3.349	3.348	*3.339*	*3.339*
SIC	3.355	3.353	*3.346*	*3.346*
$Q^2(20)$	41.060 [0.000]	48.864 [0.000]	19.488 [0.362]	19.410 [0.367]
LR	94.836 [0.000]	81.316 [0.000]	0.220 [0.639]	

Notes: As in Table 2.3.

Table 2.6: Forecast Evaluation

	1-step-ahead		5-step-ahead		10-step-ahead		15-step-ahead	
Model	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
FIAPGARCH $d = \underset{(0.1564)}{0.5101}$	6.435	2.099	9.835	2.734	14.430	3.449	17.670	3.879
FIAGARCH $d = \underset{(0.1620)}{0.3954}$	5.137	**1.887**	8.487	2.558	**12.520**	**3.205**	**15.700**	**3.621**
FIGARCH $d = \underset{(0.1725)}{0.4669}$	**5.093**	1.894	**8.423**	**2.531**	13.030	3.289	16.990	3.792
APGARCH	*7.117*	*2.249*	*11.000*	*2.894*	*16.590*	*3.722*	*20.700*	*4.246*
AGARCH	5.611	2.007	8.922	2.616	13.590	3.360	17.650	3.878
GARCH	5.539	2.012	8.827	2.600	13.960	3.382	18.140	3.937

Notes: MSE and MAE are the mean square and mean absolute forecast error, respectively. The bold/italic numbers indicate the model with the lowest/highest value of the MSE/MAE.

Table 2.7: FIAPGARCH Model with Contemporaneous Surprises

	10-min	30-min	60-min
Mean Equation			
μ	−0.0059 (0.0052)	−0.0052 (0.0084)	−0.0073 (0.0117)
θ	0.0101* (0.0053)	−0.0025 (0.0085)	−0.0017 (0.0127)
EU NAP	−3.1364*** (0.8419)	−0.7492 (0.5584)	−0.6703** (0.3486)
DE Ifo index	0.2244 (0.1798)	0.2993** (0.1505)	0.0685 (0.1329)
DE ZEW index	0.2716+ (0.1876)	−0.1211 (0.1615)	−0.2964* (0.1740)
EU consumer confidence	−0.0209 (0.2351)	−0.1701 (0.2209)	−0.1826 (0.2170)
DE new orders	0.3313** (0.1636)	0.1894 (0.1972)	0.0401 (0.1961)
EU new orders	−0.0484 (0.1873)	−0.0347 (0.2020)	−0.1599 (0.1674)
U.S. ISM manufacturer index	0.5119*** (0.1966)	0.3609** (0.1906)	−0.0697 (0.1979)
U.S. Uni-Michigan index adv.	0.0726 (0.1425)	0.2256 (0.1622)	0.1173 (0.1520)
DE industrial production	0.7018** (0.3080)	0.2616+ (0.1746)	0.1903 (0.1793)
EU industrial production	−0.0207 (0.1484)	0.1370 (0.1108)	0.0986 (0.1269)
FR industrial production	0.1576 (0.2038)	−0.0645 (0.1750)	0.1147 (0.1662)
GB industrial production	0.1792 (0.1322)	−0.0878 (0.2200)	−0.0119 (0.1374)
U.S. nonfarm payrolls	1.2040*** (0.2635)	0.5520** (0.2559)	0.5357*** (0.2451)
Variance Equation			
ω	0.1025*** (0.0183)	0.0667*** (0.0239)	0.0621*** (0.0206)
ϕ	0.4608*** (0.0502)	0.4184*** (0.1253)	0.5044*** (0.0904)
β	0.6034*** (0.0483)	0.5516*** (0.1313)	0.6830*** (0.0723)
γ	0.1146*** (0.0207)	0.1348*** (0.0285)	0.1919*** (0.0573)
δ	2.1619*** (0.0667)	2.2300*** (0.1014)	1.9928*** (0.1449)
d	0.2416*** (0.0230)	0.2357*** (0.0329)	0.3027*** (0.0497)
EU NAP dummy	9.1105* (4.9633)	3.3068 (2.5301)	0.5638 (0.8472)
Lagged volume	0.0553*** (0.0146)	0.0385*** (0.0148)	0.0395*** (0.0151)
$Q^2(20)$	20.124 [0.326]	20.764 [0.291]	20.961 [0.281]

Notes: The numbers in parentheses are Bollerslev-Wooldridge robust standard errors. ***, **, *, + indicate significance at the 1 %, 5 %, 10 % and 15% level. The numbers in brackets are p-values.

Table 2.8: FIAPGARCH Model with Lagged Surprises at 10-Minutes Frequency

Mean Equation				
	Lag			
	0	1	2	3
μ	-0.0060 (0.0053)			
θ		$0.0101^{\star}$ (0.0053)		
EU NAP	$-3.3259^{\star\star\star}$ (0.8308)	$2.3826^{\star\star\star}$ (0.4588)	$-0.8866^{\star\star}$ (0.4227)	0.0415 (0.2696)
DE Ifo index	0.2239 (0.1805)	0.0235 (0.1218)	0.2819^{+} (0.1678)	-0.2093 (0.1877)
DE ZEW index	0.2766^{+} (0.1842)	-0.2963 (0.2304)	$-0.3685^{\star\star}$ (0.1497)	-0.0617 (0.1349)
EU consumer confidence	-0.0209 (0.2347)	0.0086 (0.2484)	-0.2239 (0.2299)	-0.0780 (0.2714)
DE new orders	$0.3261^{\star\star}$ (0.1646)	0.1265 (0.3433)	-0.0740 (0.2032)	0.2184 (0.2700)
EU new orders	-0.0476 (0.1851)	0.0543 (0.1925)	-0.1292 (0.2298)	-0.2012 (0.1790)
U.S. ISM manufacturer index	$0.5106^{\star\star\star}$ (0.1953)	-0.1085 (0.2104)	0.1873 (0.1882)	-0.1251 (0.2759)
U.S. Uni-Michigan index adv.	0.0768 (0.1417)	-0.1585 (0.1771)	0.1572 (0.2101)	-0.0127 (0.1785)
DE industrial production	$0.7003^{\star\star}$ (0.3072)	0.0083 (0.1843)	-0.1956 (0.2580)	0.0729 (0.2347)
EU industrial production	-0.0224 (0.1439)	-0.1133 (0.0874)	$0.3640^{\star\star\star}$ (0.0994)	-0.2010 (0.1620)
FR industrial production	0.1645 (0.2044)	-0.2112 (0.1590)	0.0603 (0.1391)	0834 (0.3029)
GB industrial production	0.1963 (0.1411)	-0.3778 (0.3480)	0.1851 (0.2141)	0.0581 (0.2565)
U.S. nonfarm payrolls	$1.2115^{\star\star\star}$ (0.2635)	-0.1853 (0.2398)	-0.0845 (0.2177)	0.2053 (0.2046)

Variance Equation								
ω	ϕ_1	β_1	γ	δ	d	EU NAP dum.	Lagged Vol.	$Q^2(20)$
$0.1026^{\star\star\star}$ (0.0184)	$0.4560^{\star\star\star}$ (0.0501)	$0.5988^{\star\star\star}$ (0.0487)	$0.1138^{\star\star\star}$ (0.0207)	$2.1664^{\star\star\star}$ (0.0665)	$0.2410^{\star\star\star}$ (0.0231)	7.5040^{+} (4.8805)	$0.0559^{\star\star\star}$ (0.0148)	19.875 [0.339]

Notes: As in Table 2.7.

Chapter 3

Price discovery and volatility spillovers

3.1 Introduction

Since the implementation of the European Union emissions trading scheme (EU-ETS) in January 2005, trading activity within the futures markets for European Union Allowances (EUAs) has steadily expanded over the first two commitment periods. However, as a consequence of the overallocation with allowances in Phase I in association with the prohibition of the utilization of Phase I allowances to fulfill regulatory requirements in Phase II, spot market trading activity broke down and prices converged to zero within this period. With the start of Phase II, spot market trading activity strongly rose and was even higher compared to the period prior to the spot market collapse.

The main objective of this study is to analyze the price discovery process in the most liquid EUA spot and futures markets in Phase II of the EU-ETS, that is, after the spot market re-entered into operation. In addition, we investigate the joint volatility dynamics in both markets. Consequently, the paper directly attempts to assess the structure of information transmission in the EU-ETS. Contrary to previous studies such as Uhrig-Homburg and Wagner (2009), Milunovich and Joyeux (2010), and Chevallier (2010), we make use of daily as well as intraday data at the frequencies of 10 and 30 minutes. We conduct the investigation of the price transmission between both markets on the basis of vector error correction models. Besides the analysis of common factor measures as suggested by Schwarz and Szakmary (1994), Gonzalo and Granger (1995), and Hasbrouck (1995) to reveal the long-run price discovery process, we also investigate the short-run causality structure by means of Granger-causality tests. In order to assess the transmission of information in the second conditional moment, we estimate a dynamic version of the unrestricted extended CCC-GARCH model as developed by Conrad and Karanasos (2010), where each market's conditional volatility is determined by lagged volatilities and lagged shocks of both markets. This model is flexible enough to capture negative volatility spillovers, leverage effects and dynamic conditional correlations.

The first result is the absence of a cointegration relationship in daily spot and futures prices. Hence, at this frequency we cannot identify any market to be the price leading market. This result is in line with the findings of Milunovich and Joyeux (2010) and Chevallier (2010). However, extending the data from daily to intraday frequency, the analysis reveals a completely different picture. Based on high-frequency data, the results strongly support the existence of a cointegration relationship, and hence underpin the close link between both markets. Moreover, we show that drawing meaningful economic inference on each market's contribution to the price discovery process requires to conduct the analysis on the basis of data at the highest frequency of 10 minutes. The reason for this is an increasing correlation between the innovations of the two markets at lower frequencies which induces an identification problem. Most importantly, we find that the futures market incorporates information first and then transfers it to the spot market. While at the early stage of Phase II the futures market attracts 70 percent of the price discovery process, this portion even increases over time. Consequently, our results considerably extend the findings of Uhrig-Homburg and Wagner (2009) and Chevallier (2010) as they show the close relationship between both markets and the futures market's informational role.

Second, concerning the short-run causality structure, we find unidirectional Granger-causality from the futures to the spot market in daily data. However, the investigation of high-frequency data reveals a bidirectional causality structure between both markets. This result is robust with respect to the choice of the intraday frequency.

Third, in the volatility analysis we observe a similar pattern as in the price discovery analysis. In the early stage of Phase II we find unidirectional spillovers from the futures market volatility and from shocks in the futures market to the spot market's volatility. There is no such impact into the opposite direction. Contrary, in the more mature stage only lagged spot market shocks but not lagged spot market volatility affect futures market volatility. In addition, the impact of lagged futures market volatility on current spot market volatility considerably increases over time. Consequently, the results of the volatility analysis confirm the existence of the close link between both markets, which we also find in the price discovery analysis. Further, these results contradict the findings of Milunovich and Joyeux (2010) who observe a weak link between both markets' uncertainties in Phase I making use of daily data. Finally, the investigation of the DCC-structure indicates

that the dynamic conditional correlation between spot and futures returns increases from about 0.1 at the start of Phase II to approximately 0.6 at the end of the sample period.
In summary, we find strong evidence for a close relationship between the price and volatility dynamics in both markets that even intensifies over time. Further, making use of high-frequency data, we identify the futures market to be the price leading market. This result is consistent with previous findings for mature financial markets, as Tse (1999) among others reports.
We organize the remainder of the paper as follows. In Section 3.2, we give an overview on the related literature while Section 3.3 summarizes the key elements of the EU-ETS. Section 3.4 describes the data and gives an overview on the relationship between commodity spot and futures prices in general. Section 3.5 outlines the methodology we use in the empirical analysis, while Section 3.6 summarizes the estimation results and provides an interpretation of the empirical findings. Finally, Section 3.7 concludes.

3.2 Related literature

With improved data availability since the introduction of the EU-ETS, a fast growing number of empirical studies related to this market has been conducted. Besides the analysis of the impact of market fundamentals and regulatory aspects on the allowance price dynamics (see Mansanet-Bataller et al. (2007), Alberola et al. (2008) or Mansanet-Bataller and Pardo (2009) among others) and the relationship between macroeconomic performance and allowance prices (see Chevallier (2009) or Conrad et al. (2012) among others), the investigation of statistical price properties is in the focus of this field of research. While Paolella and Taschini (2008), Daskalakis et al. (2009) or Chevallier and Sevi (2011) investigate individual volatility dynamics in the spot or the futures market, other studies explicitly assess the relationship of the joint dynamics of spot and futures prices.
Uhrig-Homburg and Wagner (2009) investigate the joint development of spot and futures EUA prices in Phase I in the framework of a cost-of-carry relationship. The authors argue that for companies under the EU-ETS there is no benefit of holding EUAs in terms of meeting unexpected demand to keep the production process going since these companies need EUAs only once a year to fulfill regulatory requirements. In their empirical analysis, Uhrig-Homburg and Wagner (2009) find a cointegration relationship between observed futures prices and theoretical futures prices which they derive in the cost-of-carry model. They find that the futures contract leads the long-run price discovery process. Contrary, Milunovich and Joyeux (2010) find mixed evidence on the existence of such a relationship,

and doubt the validity of the cost-of-carry relation. More recently, making use of vector error correction models and controlling for structural breaks, Chevallier (2010) confirms the results of Milunovich and Joyeux (2010) for Phase II.

As a whole, previous studies assessing the relationship between spot and futures prices yield mixed evidence. Yet, apart from Chevallier (2010) all studies refer to Phase I, and moreover, Uhrig-Homburg and Wagner (2009), Milunovich and Joyeux (2010), and Chevallier (2010) conduct their analysis on the basis of daily data, which the authors justify by the low spot market liquidity in Phase I. However, Hasbrouck (1995) and Tse (1999) among others show that the usage of intraday data leads to more informative results compared to daily data. For the EU-ETS, Benz and Hengelbrock (2008) provide a first high-frequency price discovery analysis for Phase I. They study the joint price dynamics of futures contracts traded at the ECX and at NordPool, respectively. The authors find strong evidence for the existence of a cointegration relationship and the price leadership of the futures contract traded at the ECX. However, they critically mention the low trading activity at NordPool.

Concerning the transmission of information in the second conditional moment, that is the analysis of volatility spillovers, empirical evidence is rare. Only Milunovich and Joyeux (2010) address this topic in the framework of the GARCH-BEKK model. The authors conclude that there seems to be minor relevance of informational spillovers in the volatility of spot and futures prices. However, the study again refers to Phase I and is based on daily data.

3.3 The European Union emissions trading scheme

In this section we briefly summarize the most important features of the EU-ETS according to Directive 2003/87/EC.[1] Within the framework of the Kyoto Protocol the European Union has established the EU-ETS with the main objective to reduce greenhouse gas emissions in a cost efficient way.[2] To fulfill their commitments defined under the Kyoto Protocol, the European Community and its Member States agreed to construct an efficient European market for greenhouse gas emission allowances. Companies operating in the sectors steel and iron, pulp and paper, minerals, and energy do have to under-

[1] For a detailed formal description of the EU-ETS we refer to the Directives 2003/87/EC and 2009/29/EC.

[2] In contrast to the multilateral trade of emissions between national states as defined in Article 17 of the Kyoto Protocol, the EU-ETS is designed as a trading platform for greenhouse gas emissions on the firm level.

lay their greenhouse gas emissions with European Union Allowances (EUAs) where one EUA warrants the right to emit one tonne of CO_2-equivalent during one of three specified commitment phases.[3,4] The commitment phases cover the periods 2005-2007 (Phase I), 2008-2012 (Phase II), and 2013-2020 (Phase III).

For Phases I and II the total number of EUAs available in the market is determined by National Allocation Plans (NAPs). In accordance with Article 9 of Directive 2003/87/EC each Member State has to set up a NAP for each period containing the number of EUAs the state intends to assign to participating installations. These plans have to be approved by the European Commission. With the start of Phase III the number of issued allowances is linearly reduced yearly by 1.74% from the average annual quantity of allowances issued in the Phase II NAPs.[5] In Phases I and II, at least 95% and 90% of the issued EUAs have been allocated free of charge, while only a small portion has been auctioned. According to Directive 2009/29/EC, from 2013 onwards the portion of EUAs allocated free of charge is reduced from 80% to 30% in 2020. After their allocation, allowances can be traded over-the-counter and on several organized markets.[6] In addition to spot market trading, some of the exchanges also offer derivative products.

By April 30, each participating installation has to surrender the adequate number of EUAs required to underlay its greenhouse gas emissions of the previous year to the national authorities. Installations that do not have enough EUAs to fulfill the regulatory requirements can purchase additional allowances on the markets, while on the other hand superfluous allowances can be sold on the markets. Within a given commitment period, installations can also transfer superfluous allowances of one compliance year to the consecutive one, or employ EUAs of the next compliance year to fulfill regulatory requirements in the current year. However, the transfer of allowances between Phases I and II is not permitted. Installations that fail to meet the requirements, do have to pay a penalty of 40 (100) Euros for each excess tonne of CO_2-equivalent in Phase I (Phase II).

[3] According to Directive 2009/29/EC the European Commission plans to also cover the aviation sector's emissions by the EU-ETS from 2012 on.

[4] Besides carbon dioxide, the EU-ETS accounts for methane, nitrous oxide, hydro fluorocarbons, perfluorocarbons and sulphur hexafluoride.

[5] According to Resolution 2010/634/EU the total number of EUAs that will be allocated in 2013 equals 2,039 millions.

[6] These markets are: ECX (London), NordPool (Oslo), EEX (Leipzig), Eurex (Stuttgart), BlueNext (Paris), EXAA (Vienna) and Climex (Utrecht). For a comprehensive description of the trading frameworks at the various markets see Mansanet Bataller and Pardo (2008).

3.4 Data

3.4.1 Spot and futures markets

To analyze price discovery, causality and volatility spillovers, we construct time series based on spot market and futures market trading. For the spot market series, we use tick-by-tick data provided by the BlueNext, which attracts about 70 percent of the total daily spot market transaction volume (Rotfuß (2009)). For the futures market series we use tick-by-tick data provided by the ECX, which attracts about 90 percent of the total daily futures market transaction volume. Hence, our analysis covers a substantial portion of exchange-based trading. Due to the lack of spot market transactions in Phase I of the EU-ETS, we only consider transactions within the period 01/05/2008 to 15/12/2009 where we concentrate on the futures contracts with maturity in December 2008 and in December 2009, respectively.[7] Before March 2009, spot market trading takes place from 07:00 to 15:00 GMT. From March 2009 onwards, spot market trading time is extended to 06:00 to 15:30 GMT. Trading in the futures market is feasible from 06:00 to 16:00 GMT. To exclusively consider the trading period where spot as well as futures market trading is possible, we restrict the daily series within the empirical analysis to 07:00 to 15:00 GMT. We transform the irregular price data to equidistant intraday log prices at frequencies $h = 10$ and 30 minutes.[8] Taking the immediately preceding and following quote at the end of each h-minute interval, we compute the mean to obtain the log price at the h-minute mark. If the observed time stamp of the transaction equals the h-minute mark, we use the corresponding price as the equidistant intraday price at frequency h. If there is no transaction at the first h-minute mark at 07:00 the first intraday price equals the last price of the preceding trading day. To avoid overnight effects, we do not take the mean of transaction prices of two different days. The price of the last h-minute mark of the trading day at 15:00 equals the price of the last observed transaction before 15:00.

3.4.2 Relating EUA spot and futures prices

A considerable part of the commodity pricing literature addresses the relationship between spot and futures prices of particular commodity goods. According to Fama and French (1987), valuation of futures contracts consists of two approaches. The first approach

[7]Since the deadline for submitting allowances for the preceding year's emissions is on April 30 of the consecutive year, the trading period corresponding to Phase I does not end on 31/12/2007 but on 30/04/2008. Hence, we do not consider transactions before 01/05/2008.

[8]Besides the frequencies of 10 and 30 minutes we analyze other intraday frequencies as well. The results are similar to those Section 6 reports.

uses a risk premium to model the relationship between spot and futures prices. The second approach directly investigates the cost and benefit of holding a commodity good and models the relationship between this good's spot and futures price in the cost-of-carry framework such that the no-arbitrage condition holds (see also Borak et al. (2006)). Generally, the cost-of-carry relationship states

$$p_t^F(T) = e^{(r_t+u_t-\delta_t)(T-t)} p_t^S, \tag{3.1}$$

where $p_t^F(T)$ denotes the observed futures price in t of a contract with maturity in T, p_t^S denotes the spot price in t, r_t states the risk-free interest rate in t, u_t are the storage costs in t, and, following Brennan (1991) and Uhrig-Homburg and Wagner (2009), δ_t is the convenience yield in t. On the one hand, there is a negative effect of holding a commodity good in the form of forgone interest yields (r_t) and storage costs (u_t). On the other hand, there is a positive effect (δ_t) of holding the commodity good due to uncertainty caused by fluctuations in supply and demand, where the opportunity of meeting unexpected demand in the production process immediately justifies this benefit.

The cost-of-carry relationship as stated above holds for a range of commodity goods and has to be evaluated in case of the European Union allowance market.[9] Uhrig-Homburg and Wagner (2009) and Borak et al. (2006) point out that contrary to other production factors like raw materials or energy, companies under the EU-ETS need EUAs only once a year for compliance requirements. Moreover, the authors argue that storage costs are negligible. This implies that there is no economic rationale for the existence of convenience yields in the European Union allowance market. Hence, the relationship between spot and futures prices given by Equation (3.1) reduces to

$$p_t^{TF}(T) = e^{r_t(T-t)} p_t^S. \tag{3.2}$$

Insert Figure 3.1 about here.

We now compare the theoretical futures price $p_t^{TF}(T) = e^{r_t(T-t)} p_t^S$ with the observed futures price $p_t^F(T)$ which should be identical if the cost-of-carry relationship without convenience yields holds.[10] The risk-free interest rate r_t we use in the empirical analysis is the monthly EURIBOR on a daily basis. Observing theoretical futures prices lying above

[9]The cost-of-carry relationship holds within markets for intertemporally storable commodity goods like gold or oil. Caution is advised when modeling the relation of intertemporally non-storable commodity goods within the cost-of-carry framework.

[10]Note that other authors as Joyeux and Milunovich (2010) or Chevallier (2010) model the relationship between spot and futures prices by directly comparing the spot and the futures prices. However, applying this methodology does not alter our empirical findings.

observed futures prices could be evidence for the existence of convenience yields. Using the aforementioned emissions market data, we illustrate the price dynamics of the observed futures price of the contract with maturity in December 2009 and the theoretical futures price derived within the cost-of-carry-model neglecting convenience yields in Figure 3.1. Besides the observed and theoretical futures prices, Figure 3.1 also presents the spot price. The figure shows that within the first two months of Phase II the theoretically derived futures price lies above the observed futures price. From 01/07/08 (observation 40) on, the relationship between the theoretical and the observed futures price postulated by the cost-of-carry model virtually seems to hold.[11] Within the whole sample period, the spot price clearly lies below the observed futures price, where the difference decreases as the time to maturity decreases.

Insert Figure 3.2 about here.

Figure 3.2 shows the log-returns in the spot and the futures market at the frequency of 30 minutes. The graphs clearly exhibit volatility clustering. Furthermore, a first visual inspection implies that the evolution of volatility in both markets is closely linked. Periods of high (low) volatility in the futures market accompany periods of high (low) volatility in the spot market.

3.5 Methodology

In this section, we describe the models we employ to investigate price discovery, causality and volatility transmission in the spot and futures markets. We make use of a two step sequential estimation procedure. In the first step, we estimate a vector error correction model within the Engle and Granger (1987) framework, using the series of the theoretically derived futures prices and the series of the empirically observed futures prices. Following Theissen (2002) and Benz and Hengelbrock (2008), we afterwards compute price discovery measures based on common factors as introduced by Schwarz and Szakmary (1994) and by Hasbrouck (1995). We perform Granger-causality tests, and finally, use the residuals of the first estimation step to estimate a bivariate asymmetric unrestricted extended DCC-GARCH model, as introduced by Conrad and Karanasos (2010) to investigate the transmission of volatility between both markets.

[11] Besides the illustrated contract with maturity in December 2009, we also observe the circumstance that the observed futures price of the contract with maturity in December 2008 is smaller than the theoretical futures price within the period from 02/05/08 to 30/06/08. Again, from 01/07/08 the observed futures price and the theoretical futures price are virtually identical.

3.5.1 Estimating the vector error correction model

Let $\mathbf{p}_t = (p_t^{TF}, p_t^F)'$ be the two-dimensional price vector containing the theoretical futures price p_t^{TF} and the observed futures price p_t^F, where we assume both individual series to be $I(1)$.[12] Following Engle and Granger (1987), the series p_t^{TF} and p_t^F are cointegrated if both components of $\mathbf{p}_t$ are $I(1)$ and a (2×1) vector $\boldsymbol{\beta} \neq \mathbf{0}$ exists, such that $z_t = \boldsymbol{\beta}'\mathbf{p}_t \sim I(0)$. Consider the vector autoregressive (VAR) model in levels of order k, given by

$$\mathbf{p}_t = \boldsymbol{\mu} + \sum_{i=1}^{k} \mathbf{A}_i \mathbf{p}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{3.3}$$

where the matrices $\{\mathbf{A}_i\}_{i=1}^k$ are the (2×2) coefficient matrices of the lagged endogenous variables and $\boldsymbol{\mu}$ is a (2×1) vector of constants. We assume the (2×1) error vector $\boldsymbol{\varepsilon}_t$ to be uncorrelated over time with mean vector $\mathbf{0}$ and covariance matrix $\boldsymbol{\Omega}$. If p_t^{TF} and p_t^F are cointegrated, then following the Granger Representation Theorem the series have a vector error correction model (VECM) representation of order $(k-1)$, given by

$$\Delta\mathbf{p}_t = \boldsymbol{\mu} + \boldsymbol{\alpha}\boldsymbol{\beta}'\mathbf{p}_{t-1} + \sum_{i=1}^{k-1} \boldsymbol{\Gamma}_i \Delta\mathbf{p}_{t-i} + \boldsymbol{\varepsilon}_t, \tag{3.4}$$

where $\boldsymbol{\alpha}\boldsymbol{\beta}' = \boldsymbol{\Pi} = \sum_{i=1}^k \mathbf{A}_i - \mathbf{I}$ and $\boldsymbol{\Gamma}_i = -\sum_{j=i+1}^k \mathbf{A}_j$. The vector $\boldsymbol{\Pi}\mathbf{p}_{t-1} \sim I(0)$ states the long-run relationship, while the matrices $\{\boldsymbol{\Gamma}_i\}_{i=1}^{k-1}$ contain the short-run dependencies. If p_t^{TF} and p_t^F are cointegrated then $\boldsymbol{\Pi}$ is singular with $rk(\boldsymbol{\Pi}) = 1$ and $rk(\boldsymbol{\alpha}) = rk(\boldsymbol{\beta}) = 1$. Moreover, $\boldsymbol{\beta} = (\beta^{TF},\ \beta^F)'$ is the cointegrating vector and $\boldsymbol{\alpha} = (\alpha^{TF},\ \alpha^F)'$ is the loading vector that controls the speed of adjustment to the long-run equilibrium path. According to the efficient-market hypothesis of Fama (1970), both markets' price processes should incorporate new information immediately. Moreover, price differences should exist only temporarily due to arbitrage reasons such that both price processes should incorporate the new information equally. Therefore, we impose the restriction $\boldsymbol{\beta} = (1\ \ -1)'$ on the cointegrating vector. Since $\Delta\mathbf{p}_t$ is a stationary process it has a Wold representation and can be expressed as a vector moving average (VMA) process of infinite order

$$\Delta\mathbf{p}_t = \tilde{\boldsymbol{\mu}} + \sum_{s=0}^{\infty} \boldsymbol{\Psi}_s L^s \boldsymbol{\varepsilon}_t = \tilde{\boldsymbol{\mu}} + \boldsymbol{\Psi}(L)\boldsymbol{\varepsilon}_t, \tag{3.5}$$

[12]Making use of the ADF and the KPSS tests, we emirically justify the assumption of the $I(1)$ hypothesis (not reported).

where $\mathbf{\Psi}_0 = \mathbf{I}$, $\tilde{\boldsymbol{\mu}} = \mathbf{\Psi}(1)\boldsymbol{\mu}$, and $\mathbf{\Psi}(1) = \sum_{s=0}^{\infty} \mathbf{\Psi}_s$. Following Hayashi (2000) $\mathbf{\Psi}(z)$ is of full rank everywhere on $|z| \leq 1$. Moreover, $rk(\mathbf{\Psi}(1)) = 1$, $\boldsymbol{\beta}'\mathbf{\Psi}(1) = 0$ and $\mathbf{\Psi}(1)\boldsymbol{\alpha} = 0$. With $\mathbf{\Psi}(L) = \mathbf{\Psi}(1) + (1-L)\mathbf{\Psi}^*(L)$, where $\mathbf{\Psi}_s^* = -\sum_{j=s+1}^{\infty} \mathbf{\Psi}_j$, iterating backwards and summing up yields the relationship in levels

$$\mathbf{p}_t = \tilde{\boldsymbol{\delta}}_0 + \tilde{\boldsymbol{\mu}}t + \mathbf{\Psi}(1)\sum_{j=1}^{t} \boldsymbol{\varepsilon}_j + \mathbf{\Psi}^*(L)\boldsymbol{\varepsilon}_t. \tag{3.6}$$

The elements of the (2×2) moving average impact matrix $\mathbf{\Psi}(1)$ are the cumulative VMA coefficients (see also Yan and Zivot (2010)). $\mathbf{\Psi}(1)\sum_{j=1}^{t} \boldsymbol{\varepsilon}_j \sim I(1)$ measures the long-run impact of the innovations at time t on each of the prices. Hence, $\mathbf{\Psi}(1)\sum_{j=1}^{t} \boldsymbol{\varepsilon}_j$ describes the stochastic trend. The term $\mathbf{\Psi}^*(L)\boldsymbol{\varepsilon}_t \sim I(0)$ represents the transitory portion of the price change. Further, $\tilde{\boldsymbol{\delta}}_0$ is a (2×1) vector of constants depending on $\mathbf{p}_0$ and $\boldsymbol{\varepsilon}_0$.
Due to the orthogonality of $\boldsymbol{\beta} = (1 \;\; -1)'$ and $\mathbf{\Psi}(1)$, the moving average impact matrix contains identical rows, that is, the long-run impacts of an innovation on the prices in both markets are identical. Defining $\boldsymbol{\psi} = (\psi_1, \psi_2)$ as the common row vector of $\mathbf{\Psi}(1)$, we can rewrite the VMA in levels as

$$\mathbf{p}_t = \tilde{\boldsymbol{\delta}}_0 + \tilde{\boldsymbol{\mu}}t + \begin{pmatrix} 1 \\ 1 \end{pmatrix} \boldsymbol{\psi} \sum_{j=1}^{t} \boldsymbol{\varepsilon}_j + \mathbf{\Psi}^*(L)\boldsymbol{\varepsilon}_t. \tag{3.7}$$

Following Hasbrouck (1995), we define $(1\ 1)'\boldsymbol{\psi}\sum_{j=1}^{t} \boldsymbol{\varepsilon}_j \sim I(1)$ as the common trend that describes the common efficient price in the two markets. Finally, Baillie et al. (2002) show that for $\boldsymbol{\beta} = (1,\ -1)'$ the impact matrix can be expressed as

$$\mathbf{\Psi}(1) = \boldsymbol{\beta}_\perp \pi \boldsymbol{\alpha}_\perp' = \begin{pmatrix} \psi_1 & \psi_2 \\ \psi_1 & \psi_2 \end{pmatrix} = \tilde{\pi} \begin{pmatrix} -\alpha^F & \alpha^{TF} \\ -\alpha^F & \alpha^{TF} \end{pmatrix}, \tag{3.8}$$

where the vectors $\boldsymbol{\alpha}_\perp$ and $\boldsymbol{\beta}_\perp$ are the orthogonal complements to the vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ such that $\boldsymbol{\alpha}'\boldsymbol{\alpha}_\perp = 0$ and $\boldsymbol{\beta}'\boldsymbol{\beta}_\perp = 0$ and π as well as $\tilde{\pi}$ are scaling factors.

3.5.2 Price discovery measures and causality analysis

The main objective of price discovery analysis is to identify the process of the incorporation of permanent changes in the price of equal or closely linked assets traded on more than one market (see Hasbrouck (1995) and Harris et al. (2002)). Put differently, the central question is whether one market adjusts prices faster as a response to new information than the other market which would lead to different prices across both markets after the incorporation of the information in the faster reacting market. Economic theory, however, suggests that according to the law of one price and due to arbitrage considerations price differences can only exist temporarily. Hence, theory predicts that the price reaction in the first market generates a price reaction in the second market such that arbitrage opportunities do not remain permanently.

In his seminal paper, Hasbrouck (1995) develops a general framework for analyzing a market's contribution to the price discovery process. The starting point is the assumption of a common implicit efficient equilibrium price in all the markets the asset is traded on. In particular, the measure of a market's contribution to the price discovery process is called the market's information share, which is defined as the portion of the variance of an innovation to the common efficient price that can be attributed to this market relative to the total variance of such an innovation. As an alternative price discovery measure, Schwarz and Szakmary (1994) propose the common factor weights. In the framework of an error correction model a market's common factor weight equals the relative magnitude of the coefficients of the adjustment vector. Gonzalo and Granger (1995) show that the common factor weights can be derived from a common factor model. Moreover, Baillie et al. (2002) show that both measures provide similar results as far as the correlation between the error correction model's residuals is weak or moderate.

Besides the analysis of price discovery in markets for identical assets (see Hasbrouck (1995) or Theissen (2002) among others) the investigation of price discovery in spot and futures markets has gained interest in the literature. In general, for mature markets the futures usually is the price leading market. Tse (1999) explains this finding by relatively low transaction costs and inherent leverage in the futures market. In case of the EU-ETS only Uhrig-Homburg and Wagner (2009) and Benz and Hengelbrock (2008) provide empirical studies. While the former report the predominant informational role of the futures relative to the spot market based on daily data making use of common factor weights, the latter also compute information shares to investigate the price discovery process of EUA futures traded on different markets. Based on intraday data, Benz and Hengelbrock (2008) highlight the relevance of the ECX. However, both studies refer to the first commitment period.

In the following, we present the common factor weights and the information shares. Moreover, we introduce a methodology for assessing the short-run causality structure.

Consider the loading matrix $\boldsymbol{\alpha}$ of the VECM described by Equation (3.4). As can be seen from Equation (3.8) the coefficients α^{TF} and α^{F} represent the permanent effect that a shock has on the system. In particular, the equation shows that $\alpha^{TF} - \alpha^{F}$ constitutes the system's total adjustment to a shock in the markets. An intuitive measure of a market's contribution to the price discovery process in this bivariate system constitutes the portion of the total adjustment that can be attributed to the respective other market (see also Theissen (2002)). Hence, following Schwarz and Szakmary (1994) the common factor weights of the futures and the spot market are given by

$$CFW^F = \frac{\alpha^{TF}}{\alpha^{TF} - \alpha^F} \quad \text{and} \quad CFW^{TF} = \frac{-\alpha^F}{\alpha^{TF} - \alpha^F}. \tag{3.9}$$

If the price discovery process exclusively takes place in the futures market, then only the spot market reacts to deviations from the equilibrium path which implies $\alpha^{TF} \neq 0$, while $\alpha^F = 0$ yielding $CFW^F = 1$ and $CFW^{TF} = 0$. If each of the two markets equally contributes to the price discovery process, then the markets' common factor weights are identical.

Besides the common factor weights, Hasbrouck (1995) develops the information share concept to assess a market's contribution to the price discovery process. Consider the moving average impact matrix $\mathbf{\Psi}(1)$ of the VMA representation derived by applying the Beveridge-Nelson decomposition to the VECM. Hasbrouck (1995) defines a market's information share as the market's contribution to the total variance $\boldsymbol{\psi}\mathbf{\Omega}\boldsymbol{\psi}'$ of an innovation to the permanent portion of the price change relative to $\boldsymbol{\psi}\mathbf{\Omega}\boldsymbol{\psi}'$. Hence, if the VECM errors are uncorrelated, that means $\mathbf{\Omega}$ is diagonal, the information share of market i is given by

$$IS_i = \frac{\psi_i^2 \sigma_{ii}}{\boldsymbol{\psi}\mathbf{\Omega}\boldsymbol{\psi}'}, \tag{3.10}$$

where ψ_i is the ith element of $\boldsymbol{\psi}$ and σ_{ii} is the ith diagonal element of $\mathbf{\Omega}$. If the residuals are contemporaneously correlated, the off-diagonal elements of $\mathbf{\Omega}$ representing the covariance between both markets' residuals have to be accounted for in the construction of the price discovery measure. In order to minimize and bound the impact of the correlation, Hasbrouck (1995) suggests to conduct the analysis on the basis of data at highest frequencies and to make use of the Cholesky decomposition $\mathbf{\Omega} = \mathbf{F}\mathbf{F}'$, where $\mathbf{F}$ is a lower triangular matrix, leading to the information share

$$IS_i = \frac{([\boldsymbol{\psi}\mathbf{F}]_i)^2}{\boldsymbol{\psi}\mathbf{\Omega}\boldsymbol{\psi}'} \tag{3.11}$$

of market i, where $[\boldsymbol{\psi}\mathbf{F}]_i$ is the ith element of the row vector $\boldsymbol{\psi}\mathbf{F}$. Baillie et al. (2002) point out that by substituting the expressions of Equation (3.8) into Equation (3.11), the ratio of the information shares is given by $IS_1/IS_2 = (-\frac{\alpha^F}{\alpha^{TF}}F_{11} + F_{21})^2/F_{22}^2$, and since

$IS_1 + IS_2 = 1$ the information shares can be expressed as

$$IS_1 = \frac{(-\frac{\alpha^F}{\alpha^{TF}} F_{11} + F_{21})^2}{(-\frac{\alpha^F}{\alpha^{TF}} F_{11} + F_{21})^2 + F_{22}^2} \quad \text{and} \quad IS_2 = \frac{F_{22}^2}{(-\frac{\alpha^F}{\alpha^{TF}} F_{11} + F_{21})^2 + F_{22}^2}. \tag{3.12}$$

However, the computed value of a market's information share is not invariant to the ordering of the series in the VECM as we attribute the contemporaneous correlation to the market represented by the VECM's first equation. Nevertheless, we derive an upper (lower) bound for a market's information share by ordering the price series of this market first (second). Having constructed the upper and lower bounds of the information shares, we compute the mean and each information share's range.

Finally, in order to analyze the short-run causality structure, we apply Granger-causality tests as proposed by Granger (1969). Within the framework of the VECM, we test whether the lagged log-returns of one market are jointly significant in the equation of the other market. Rejecting the null hypothesis indicates that the corresponding market does Granger-cause the other market.

3.5.3 Volatility spillovers

In addition to the transmission of information in the spot and futures returns, we analyze the volatility transmission between the two markets. According to Andersen (1996), the latent underlying flow of information to the market drives an asset's volatility process. Consequently, the price leadership of one market could also induce the existence of a causal relationship in both markets' volatility dynamics (see also Chan et al. (1991)).

In order to capture the joint volatility dynamics of both markets, as a baseline specification, we employ the unrestricted extended CCC-GARCH (UECCC-GARCH) model of Conrad and Karanasos (2010) that allows for volatility spillovers of either sign.[13] Moreover, in order to allow the conditional correlations between spot and futures returns to vary over time, we generalize the UECCC-GARCH model to the asymmetric unrestricted extended DCC-GARCH (AUEDCC-GARCH) specification by adopting the dynamic conditional correlation (DCC) structure of Engle (2002).[14]

We now set up a specification along the lines of the aforementioned models which we estimate on the basis of the fitted residuals of the spot and the futures market obtained from the VECM. However, before estimating the bivariate volatility specification, we first

[13] For a comprehensive overview on multivariate GARCH models and appropriate estimation techniques see Bauwens et al. (2006).

[14] We leave the application of fractionally integrated specifications allowing for long memory in the variance equation (see Conrad et al. (2011)) for future research.

have to transform the residual vector. Conrad et al. (2012) point out that the absolute EUA returns exhibit a strong intraday pattern which induces deterministic seasonality in the autocorrelation function of the absolute returns and reflects the time-varying intraday trading activity. This pattern carries over to the fitted residuals of the error correction model as Figure 3.3 illustrates in case of the absolute spot market residuals at 10 minutes frequency.

Insert Figure 3.3 about here.

To extract the deterministic pattern, we scale each fitted residual of an intraday interval by the average absolute residual of the corresponding interval over all trading days (see also Tse (1999) and Conrad et al. (2012)).
Let $\mathbf{e}_t = (e_{1t}, e_{2t})'$ be the (2×1) vector of the filtered residuals with filtration $\mathcal{F}_{t-1}$ generated by the information up through time $t-1$. For the volatility specification, we define $\mathbf{e}_t = \mathbf{z}_t \odot \mathbf{h}_t^{\wedge 1/2} = (z_{1t}\sqrt{h_{11t}}, z_{2t}\sqrt{h_{22t}})'$, where $\odot$ denotes the Hadamard product and $\wedge$ indicates elementwise exponentiation. We assume the (2×1) vector $\mathbf{z}_t = (z_{1t}, z_{2t})'$ to be independent and identically distributed with mean zero, finite second moments, and dynamic correlation matrix

$$\mathbf{R}_t = \left[\boldsymbol{\rho}_{ij,t}\right]_{i,j=1,2} = \left[\frac{q_{ij,t}}{\sqrt{q_{ii,t}q_{jj,t}}}\right]_{i,j=1,2}, \tag{3.13}$$

where

$$\mathbf{Q}_t = [q_{ij,t}]_{i,j=1,2} = (1 - \alpha^{DCC} - \beta^{DCC})\overline{\mathbf{Q}} + \alpha^{DCC}\mathbf{z}_{t-1}\mathbf{z}_{t-1}' + \beta^{DCC}\mathbf{Q}_{t-1}. \tag{3.14}$$

Furthermore, $\overline{\mathbf{Q}}$ is the unconditional covariance matrix of $\mathbf{z}_t$ and α^{DCC} and β^{DCC} are non-negative scalars satisfying $\alpha^{DCC} + \beta^{DCC} < 1$. Consequently, the DCC-specification models the covariance matrix of $\mathbf{z}_t$ as a GARCH-type equation itself.
Finally, the structure we impose on the conditional variances is given by

$$\mathbf{h}_t = \boldsymbol{\omega} + \mathbf{A}\mathbf{e}_{t-1}^{\wedge 2} + \mathbf{B}\mathbf{h}_{t-1} + \boldsymbol{\Gamma}\left(\mathbf{e}_{t-1}^{\wedge 2} \odot \mathbf{1}_{\mathbf{e}_{t-1}<0}\right), \tag{3.15}$$

where $\boldsymbol{\omega} = (\omega_1, \omega_2)'$ are constants, $\mathbf{A} = [a_{ij}]_{i,j=1,2}$, $\mathbf{B} = [b_{ij}]_{i,j=1,2}$, $\boldsymbol{\Gamma} = [\gamma_{ij}]_{i,j=1,2}$ and $\mathbf{1}_{\mathbf{e}_{t-1}<0} = (\mathbf{1}_{e_{1,t-1}<0}, \mathbf{1}_{e_{2,t-1}<0})'$ is a bivariate vector of indicator functions where $\mathbf{1}_{e_{i,t-1}<0}, i = 1, 2$, is equal to one if $e_{i,t-1} < 0$ and zero elsewise. For $[\gamma_{ij}]_{i,j=1,2} = \mathbf{0}$ and $\alpha^{DCC} = \beta^{DCC} = 0$ the model equals the UECCC-GARCH(1,1) specification of Conrad and Karanasos (2010). The authors derive necessary and sufficient conditions for $\mathbf{h}_t \geq \mathbf{0}$ that do not place any restrictions on the signs of the coefficients in the $\mathbf{B}$ matrix, such that negative volatility spillovers are explicitly allowed for. The coefficient a_{12} (a_{21}) indicates to which extend the squared lagged innovation of the futures market (spot market) affects the conditional variance of the spot market (futures market). Hence, we refer to this effect as ARCH spillover effect. Further, b_{12} (b_{21}) indicates the extend to which the lagged conditional variance of the futures market (spot market) determines the conditional variance of the spot market (futures market) within the current period and is called GARCH spillover effect. Furthermore, the structure we impose on the conditional variances allows both market's positive and negative shocks to asymmetrically affect both markets variances.

We estimate all specifications making use of quasi maximum likelihood estimation (QMLE) following Bollesrlev and Wooldridge (1992).

3.6 Empirical results

Section 3.4 provides a first graphical illustration of spot and futures price dynamics.[15] Our full sample covers data from 02/05/2008 to 15/12/09, where we conduct the investigation on the basis of the most actively traded futures contract.[16] Table 3.1 presents the descriptive statistics of daily and high-frequency returns.

Insert Table 3.1 about here.

As a whole, the sample covers 394 trading days yielding 6304 and 18912 equidistant high-frequency observations at the 30 and 10 minutes frequency, respectively. The means of the returns are negative but very close to zero at each frequency, while the standard deviations increase with decreasing frequencies. Moreover, the summary statistics indicate that the return distributions are slightly right-skewed at most of the analyzed frequencies. Positive skewness in combination with strong excess kurtosis leads to the rejection of the null-hypothesis of normally-distributed returns as the Jarque-Bera-statistics indicate. In order to analyze the evolution of both markets' informational roles over time, we split up the full sample into two subsamples. The first subsample refers to the pre 15/12/2008 period, and hence, studies the December 2008 contract, while the second subsample covers the post 15/12/2008 period, and consequently studies the December 2009 contract. Descriptive

[15]Note that in order to save space we denote the theoretically derived futures price as spot price and the empirically observed futures price as futures price.

[16]Within the period 01/05/2008 to 15/12/2008 the futures contract with maturity in December 2008 is the most actively traded futures contract, whereas in the post 15/12/2008 period the contract with maturity in December 2009 attracts the highest trading volume. Note that the contract with the highest trading volume does not have to be the contract with the closest time to maturity, since, futures with maturity in March of each year are traded as well.

statistics of both subsamples are similar to those Table 3.1 presents (not reported).
In a next step, we check the integration order of the series. For this, we apply the ADF test and the KPSS test. Concerning the series in levels, we account for a time trend and a constant, while in case of the series in first differences, we do not account for a time trend. The Schwarz information criterion chooses the number of lags in the tests. The results (not reported) clearly indicate the existence of a unit root in the levels series and the absence of a unit root in the series in first differences. Consequently, we treat all series as integrated of order one.
Based on these results, the appropriate specification to investigate the price transmission mechanisms is a vector error correction model. Again the Schwarz information criterion determines the optimal lag order.

3.6.1 The vector error correction model

Table 3.2 shows the long-run relationship between the spot and the futures series as implied by the estimated VECM with restricted cointegration vector $\boldsymbol{\beta} = (1,\ -1)'$.[17] Within the VECM, we check whether the error correction term $p^{TF}_{t-1} - p^{F}_{t-1}$ of period $t-1$ significantly enters the equation of the spot and/or futures series. While Panel A shows the results of the estimated error correction models on the basis of the spot and the futures market concerning the full sample, Panels B and C refer to the first and the second subsample, respectively. The results of Panel A indicate that the error correction term enters the spot equation statistically significant at least at the five percent level at each of the analyzed intraday frequencies. Further, the error correction term enters the futures equation statistically significant at the five percent level at the frequency of 10 minutes, while it does not enter the futures equation at the 30 minutes frequency statistically significant. Spot prices exceeding futures prices in period $t-1$ lead to positive price reactions in the futures market and to negative price movements in the spot market in the consecutive period. This is consistent with what one would expect in case of deviations from the equilibrium relation with $p^{TF}_{t-1} > p^{F}_{t-1}$. Finally, the price reaction in the spot market is about 2.5 times as strong as the price reaction in the futures market. Moreover, both reactions decrease with increasing frequencies.
The results of Panels B and C refer to the futures contracts with maturity in December 2008 and 2009, respectively. The estimated adjustment coefficients ($\hat{\alpha}^{TF}$ and $\hat{\alpha}^{F}$) shown in Panel B imply that both markets react statistically significant to deviations from the long-

[17]We also estimate the models with unrestricted cointegration vector $\boldsymbol{\beta}$. The results are not qualitatively different from those we report in this section. Moreover, in almost each case, we cannot reject the hypothesis of the cointegrating vector being equal to $\boldsymbol{\beta} = (1, -1)'$ at least at the five percent level.

run equilibrium. In each of the estimated error correction models the term $p_{t-1}^{TF} - p_{t-1}^{F}$ enters the spot and the futures equation at the one percent level at the frequencies of 10 and 30 minutes. At each intraday frequency the sign of the estimated adjustment coefficients is the expected one. Contrary to the full period, price reactions in the spot market induced by deviations from the long-run equilibrium are only about 1.5 times as strong as price reactions in the futures market at 10 minutes frequency, while at the lower frequency of 30 minutes the effect in the spot market exceeds the one in the futures market slightly. The results provided by Panel C, however, reveal a different situation in the second subsample.

Insert Table 3.2 about here.

While the error correction term enters the spot market equation at least at the 10 percent level at both intraday frequencies, the error term does not significantly enter the futures market equation at any intraday frequency. This indicates that in the second subsample exclusively the spot price adjusts to the long-run price equilibrium in case of deviations which implies that there is no transmission of information from the spot to the futures market anymore. The speed of adjustment in the spot market is similar in both subsamples.

In order to compare the results of our study with those of previous investigations such as Uhrig-Homburg and Wagner (2009), Milunovich and Joyeux (2010), and Chevallier (2010), we also report the results of the analysis on the basis of daily data. Neglecting the information contained in the high-frequency data, however, leads to the spurious conclusion of no cointegration relationship between spot and futures prices since the error correction term does neither enter the spot nor the futures equation significantly in the full sample and the second subsample. This in turn is in strong contrast to the findings of the high-frequency analysis.[18] Moreover, the absence of a cointegration relationship in daily data is in line with Chevallier (2010) who does not find such a relationship for Phase II after controlling for structural breaks, whereas the sample he analyzes is slightly shorter than our sample. In addition, the result is mainly consistent with the one of Milunovich and Joyeux (2010) who consider Phase I. Finally, the result contradicts the one of Uhrig-Homburg and Wagner (2009) who also for Phase I find evidence for the existence of a long-run relationship in daily data.

In conclusion, there is strong evidence for a cointegration relationship at each of the intraday frequencies. This finding implies the importance of expanding the analysis to a high-frequency level and extends the results of Uhrig-Homburg and Wagner (2009),

[18] We also confirm this result by the Johansen cointegration test (not reported).

Milunovich and Joyeux (2010), and Chevallier (2010). In particular, the intradaily findings underpin the close relationship of the joint dynamics of spot and futures prices which we do not detect in daily data. Moreover, the results are consistent with findings of high-frequency studies for mature financial markets, as e.g. Tse (1999) reports for the DJIA-index and the corresponding index-futures.

3.6.2 Price discovery measures and causality analysis

Based on the estimates Table 3.2 presents, we now investigate both markets' contributions to the price discovery process. Since the error correction term does neither enter the spot nor the futures market equation at the daily frequency, we abdicate to report the price discovery measures in these cases. Table 3.3 presents the common factor weights for both markets as well as the mean and the range of the futures market's information share. In the last column, we report the estimated residual correlation $\rho^{TF,F}$ between both series. Panel A shows the estimates concerning the complete period. According to the common factor weights, the portion the futures market contributes to the price discovery process is about 70 percent at both frequencies. This is confirmed by the futures market's information share of 0.687 at 10 minutes frequency, while the range of the information share equals 0.521. Hence, the ranges of the spot and futures markets' information shares slightly overlap. At the lower frequency of 30 minutes the futures market's information share reduces to 57.9 percent while the range increases to 0.794. The extension of the range is consistent with Theissen (2002) who observes that an information share's range increases considerably by switching the analyzed frequency from 1 to 5 minutes intervals in the analysis of German stocks. The reason for this is that according to Equation (3.11), we have to attribute the residual correlation that increases from 0.587 to 0.827 with decreasing frequency to one of both markets. Consequently, the range of the information shares at the lower frequency of 30 minutes is more than 1.5 times as large as the range of the information shares at the 10 minutes frequency. This is in line with Baillie et al. (2002) who show that high residual correlation strongly affects the upper and the lower bounds of a market's information share. In conclusion, the results of Panel A indicate the futures market's predominant informational role in the price discovery process which is consistent with the findings for more mature markets as e.g. Tse (1999) reports. Panels B and C shed light on the evolution of both markets' contribution to the price discovery process. While in the first subsample the common factor weights of the futures market are less than the ones we report for the full period, the futures market's factor weights in the second subsample are clearly higher at both frequencies. Both measures confirm this finding, whereas the increase is less pronounced in case of the information share which we trace back to an increase of the contemporaneous correlation that induces an increase of the range of the information shares. Consequently, at the lower frequency of 30 minutes the range of both information shares are strongly overlapping which leads to the dilution of any conclusion, such that we have to draw economically meaningful inference on the basis of the data at the highest frequency. Again, this is consistent with the findings of Baillie et al. (2002) and Theissen (2002). As a whole, the results indicate the futures market's predominant informational role in the price discovery process which even increases over time.

Insert Table 3.3 about here.

Consequently, the results clearly extend those of Uhrig-Homburg and Wagner (2009) who find the price leadership of the futures contract in Phase I. Moreover, the results underpin the predominant role of the ECX as already pointed out by Benz and Hengelbrock (2008). Table 3.4 addresses the question of the short-run causality dynamics. We perform Granger-causality tests to check whether one market's lagged log-returns affect the other market's current log-returns. The χ^2-distributed test statistics imply that at the 10 percent level, we have to reject the hypothesis that the futures market does not Granger-cause the spot market in the short-run context on the basis of daily data, while we cannot reject the hypothesis that the spot market does not Granger-cause the futures market. Hence, the results imply unidirectional Granger-causality from the futures to the spot market in the full sample. Concerning the high-frequency analysis, we reject the hypothesis that the spot market does not Granger-cause the futures market at least at the five percent level for each of the analyzed samples and frequencies. Similarly, we reject the opposite direction hypothesis that the futures market does not Granger-cause the spot market at the one percent level for each of the analyzed samples and frequencies. In contrast to the long-run price discovery analysis in which we identify the futures market as predominant contributor to the price discovery process, Table 3.4 implies bidirectional feedback in the returns. In conclusion, the results of the daily analysis are in line with the findings of Chevallier (2010). However, considering more informative intraday data again reveals a different situation.

Insert Table 3.4 about here.

3.6.3 Volatility spillovers

In the previous subsections, we provide an analysis of the price transmission mechanisms between spot and futures markets. In addition, we now study the informational spillovers in the volatility of both processes.

Insert Table 3.5 about here.

For this, we only make use of 10 minutes frequency data since in the previous subsection, we show that meaningful inference should be based on data at the highest frequency. Table 3.5 presents the results, where Panel A refers to the full sample, Panels B and C to the two subsamples, respectively. For each sample, we estimate three specifications. In the first row of each panel, we present the general specification as described by Equation (3.15). The second row refers to the model that does only allow for ARCH spillovers, that is, we impose the restriction $b_{12} = b_{21} = 0$. Finally, in each panel's last row, we additionally restrict a_{12} and a_{21} to zero and do not allow for asymmetric effects. Consequently, the last rows represent the DCC-model of Engle (2002). We compare the two restricted specifications with the general specification making use of the χ^2_k-distributed likelihood-ratio statistic $LR = 2 \cdot (L_{UR} - L_R)$, where k is the number of restricted parameters, L_{UR} is the likelihood of the unrestricted asymmetric UEDCC-GARCH specification and L_R is the likelihood of the restricted model. Each specification satisfies the stationary conditions of Engle (2002) and the non-negativity conditions of Conrad and Karanasos (2010). In the following, we discuss the estimation results.

First, we provide a summary of the results of the general bivariate asymmetric UEDCC-GARCH specification. For each sample the GARCH spillover coefficient from the spot to the futures market, b_{21}, is insignificant. Consequently, we set this coefficient equal to zero in each of the estimated specifications. In the full sample, the volatilities of both markets significantly react to lagged shocks in the other market, whereas the response of spot market volatility to a shock in the futures market is about six times higher than the response of the futures market volatility to a shock in the spot market.[19] Moreover, the current spot volatility significantly reacts to the lagged volatilities of both markets, whereas the reaction to its own past is slightly stronger than the reaction to lagged futures volatility. Finally, both markets asymmetrically react to positive and negative shocks of the same size. Bad news increase volatility stronger than good news. Investigating the subsamples reveals further interesting details of the evolution of both markets' conditional volatilities. In the first subsample both markets' lagged volatilities affect the spot market volatility. Moreover, we observe ARCH spillovers only from the futures to the spot market. Furthermore, the effect of lagged futures market volatility on current spot volatility is slightly weaker than in the full sample. Consequently, in the early stage of Phase I, there is evidence for unidirectional transmission of information from the futures to the spot market. The situation is very different for the second subsample. While

[19]Even though, we do not standardize the series, the magnitude of the unconditional volatilities of both markets are very similar, as Table 3.1 shows.

lagged shocks of both markets significantly affect both markets' volatilities, the impact of lagged futures market volatility on current spot market volatility is considerably stronger than in the first subsample. Most importantly, lagged spot market volatility does not affect the current spot market volatility anymore, while the spillover effect $b_{12} = 0.8285$ is very large and highly significant. As a whole, the findings strongly confirm the futures market's informational role that we observe in the price discovery analysis. Furthermore, the results contradict the findings of Milunovich and Joyeux (2010) who only find a weak relation between the volatilities of both series. Equally important, our results are in line with those for mature markets as e.g. Tse (1999) reports.

Second, we compare the results of the general asymmetric UEDCC-GARCH specifications with those of the restricted models. For each sample and for each restricted specification, the likelihood ratio test statistics unambiguously indicate that restricting coefficients to zero significantly reduces the likelihood of the model. Additionally restricting b_{21} to zero induces severe biases in the coefficient estimates. Compared to the general specification, the restricted models strongly overestimate the effect of the lagged spot volatility. This overestimation is even more pronounced in the second subsample. Contrary, setting b_{12} equal to zero does not influence the estimated coefficients in the futures market equation. The same is the case for additionally restricting a_{12} and a_{21} to zero and not allowing for asymmetric effects.

Finally, we analyze the evolution of the dynamic conditional correlation between both series. Figure 3.4 illustrates the evolution of the dynamic correlation and provides a very interesting picture.

Insert Figure 3.4 about here.

At the beginning of the sample period the correlation between both series is only slightly larger than zero. Moreover, in the early stage of the sample the dynamic correlation is very volatile compared to the later stage of the period. The figure shows that over time the correlation between both series strongly increases reaching values of approximately 0.65. These results are in line with the residual correlations of the full sample and both subsamples as Table 3.3 reports. Moreover, Table 3.5 shows that these findings are also consistent with the estimated coefficient β^{DCC} which is higher in the second subsample than in the first one.

3.7 Conclusion

This paper addresses the question of information transmission in European spot and futures markets for emission allowances in the second commitment period of the EU-ETS making use of high-frequency data. Previous studies based on daily data such as Uhrig-Homburg and Wagner (2009) and Milunovich and Joyeux (2010) find mixed evidence for the relationship between spot and futures prices in the first commitment period. Moreover, based on low-frequency data and after controlling for structural breaks, Chevallier (2010) does not find a cointegration relationship between spot and futures prices in the early stage of the second commitment period.

We show that this is due to the focus on the analysis of daily data. Based on high-frequency data, however, we find unambiguous evidence for the existence of a cointegration relationship between spot and futures prices. Moreover, we clearly identify the futures market as the price leading market. According to the price discovery measures of Schwarz and Szakmary (1994), Gonzalo and Granger (1995), and Hasbrouck (1995), 70 percent of the price discovery process take place in the futures market. Further, the informational content of futures trading increases over time. The results of the conditional volatility analysis further confirm the findings of the price discovery analysis. Contrary to Milunovich and Joyeux (2010) who find evidence for a loose dynamic relationship between spot and futures volatility in the first commitment period, we find strong evidence for volatility spillovers from the futures to the spot market but not into the opposite direction. Moreover, the analysis reveals that both markets' lagged shocks affect the volatilities in the other market. Finally, the results of the DCC-structure analysis indicate that the link between both markets considerably intensifies over time as the dynamic conditional correlation increases from 0.1 at the early stage of Phase II to about 0.65 at the end of the sample period.

In conclusion, our results help to understand the mixed evidence previous studies report and highlight the informational role of the futures market. Moreover, they indicate that the price discovery process in the European allowance markets is similar to the one in more mature markets as Tse (1999) reports. This finding is remarkable due to the immature character of the EU-ETS.

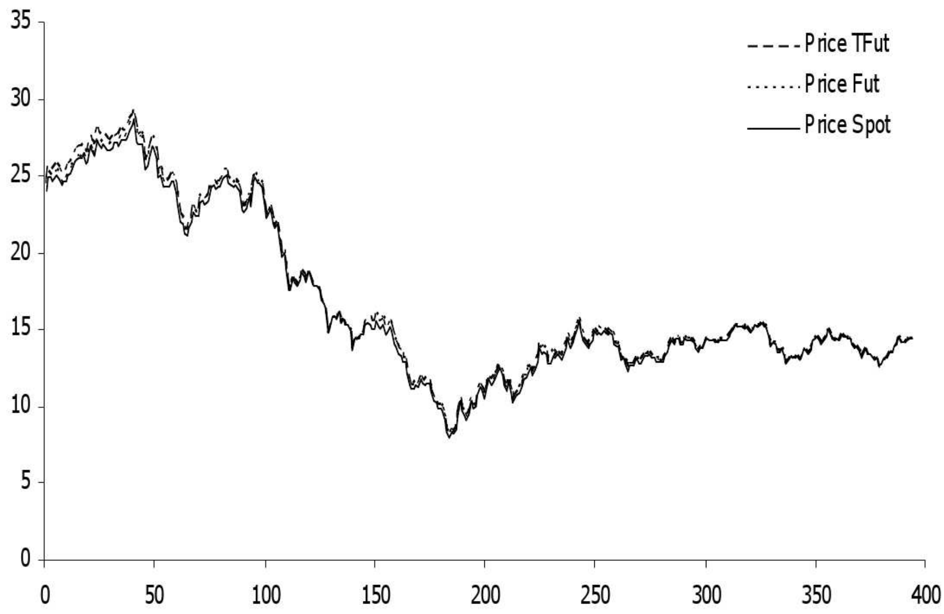

Figure 3.1: Cost-of-Carry Relationship

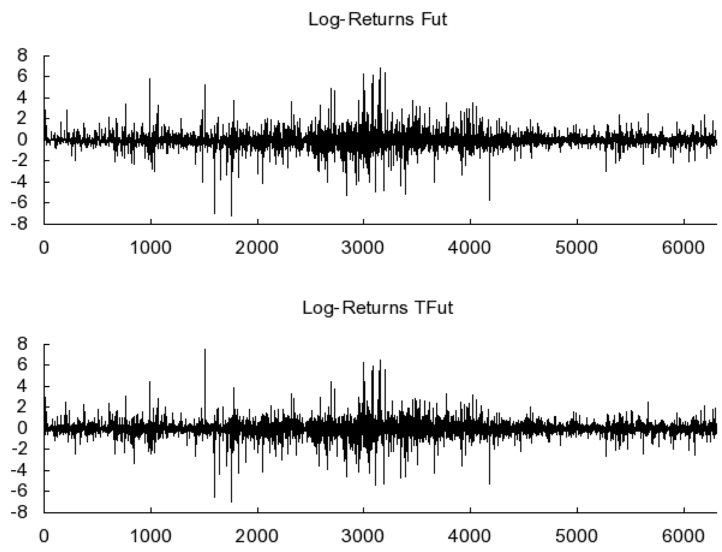

Figure 3.2: Log-returns in the spot and futures market at the frequency of 30 minutes.

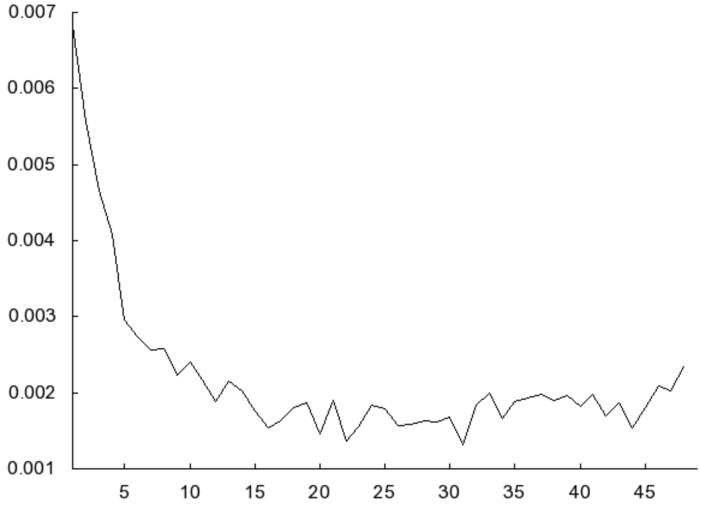

Figure 3.3: Intraday pattern of the absolute spot residuals at 10-minute frequency.

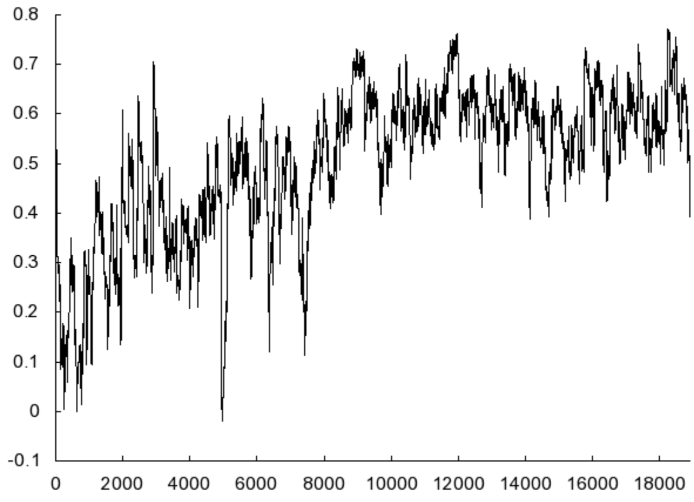

Figure 3.4: Evolution of dynamic correlation between both series at 10-minutes frequency.

Table 3.1: Descriptive Statistics

Full Sample						
Series	# obs.	Mean	Stand. Dev.	Skewness	Kurtosis	Jarque-Bera
$\Delta P_t^{TF}(10)$	18912	$-2.82 \cdot 10^{-5}$	0.0039	0.8735	76.6531	> 1000 [0.000]
$\Delta P_t^{F}(10)$	18912	$-2.52 \cdot 10^{-5}$	0.0042	0.4806	58.2754	> 1000 [0.000]
$\Delta P_t^{TF}(30)$	6304	$-8.46 \cdot 10^{-5}$	0.0070	0.2478	20.672	> 1000 [0.000]
$\Delta P_t^{F}(30)$	6304	$-7.52 \cdot 10^{-5}$	0.0075	0.0092	18.8685	> 1000 [0.000]
$\Delta P_t^{TF}(d)$	394	-0.0014	0.0285	-0.0807	4.3949	32.369 [0.000]
$\Delta P_t^{F}(d)$	394	-0.0013	0.0283	0.0316	4.3406	29.568 [0.000]

Notes: p-values in brackets.

Table 3.2: Long-run Relationship I - VECM

	adjustment vector	
frequency	$\hat{\alpha}^{TF}$	$\hat{\alpha}^{F}$
Panel A - Full Sample		
10	$-0.0176^{\star\star\star}$ (0.0029)	$0.0069^{\star\star}$ (0.0032)
30	$-0.0213^{\star\star}$ (0.0094)	0.0079 (0.0096)
daily	-0.0347 (0.1444)	0.0244 (0.1439)
Panel B - Contract with maturity in December 2008		
10	$-0.0427^{\star\star\star}$ (0.0049)	$0.0302^{\star\star\star}$ (0.0057)
30	$-0.0540^{\star\star\star}$ (0.0162)	$0.0607^{\star\star\star}$ (0.0178)
daily	0.4241 (0.2964)	$0.5993^{\star\star}$ (0.2935)
Panel C - Contract with maturity in December 2009		
10	$-0.0297^{\star\star\star}$ (0.0055)	0.0073 (0.0059)
30	$-0.0339^{\star}$ (0.0177)	0.0008 (0.0175)
daily	-0.0979 (0.2572)	-0.0092 (0.2564)

Notes: The table shows the estimation results of the VECM with restricted cointegration vector $\boldsymbol{\beta} = (1,\ -1)'$. The numbers in parentheses are robust standard errors. $\star\star\star$, $\star\star$, $\star$ indicate significance at the 1 %, 5 % and 10 % level.

Table 3.3: Long-run Relationship II - Price Discovery

frequency	CFW		IS for futures		$\rho^{TF,F}$
	futures	spot	mean	range	
			Panel A - Full Sample		
10	0.7193	0.2807	0.6866 (0.4260;0.9472)	0.5212	0.5874
30	0.7293	0.2707	0.5793 (0.1825;0.9761)	0.7936	0.8272
			Panel B - Contract with maturity in December 2008		
10	0.5857	0.4143	0.6473 (0.4675;0.8272)	0.3597	0.3793
30	0.4709	0.5291	0.4973 (0.1498;0.8448)	0.6950	0.6949
			Panel C - Contract with maturity in December 2009		
10	0.8021	0.1979	0.6904 (0.4016;0.9792)	0.5776	0.6743
30	0.9762	0.0238	0.6071 (0.2144;0.9998)	0.7804	0.8810

Notes: Price discovery measures are only shown if the error term at least enters one of the equations at the corresponding frequency and both coefficients obey the expected sign. Numbers in parentheses are lower and upper bounds of the futures market's information share.

Table 3.4: Short-run Dynamics - Granger-causality

frequency	futures does not cause spot	spot does not cause futures
	Panel A - Full Sample	
10	1379.01 [0.000]	241.60 [0.000]
30	315.87 [0.000]	102.03 [0.000]
daily	6.85 [0.077]	5.76 [0.124]
	Panel B - Contract with maturity in December 2008	
10	522.98 [0.000]	18.03 [0.001]
30	215.52 [0.000]	9.81 [0.020]
	Panel C - Contract with maturity in December 2009	
10	662.19 [0.000]	265.21 [0.000]
30	83.76 [0.000]	110.04 [0.000]

Notes: The shows the χ^2-distributed test-statistic on the Granger-causality tests in the bivariate specifications.

Table 3.5: Coefficient Estimates AUEDCC-GARCH Specification

α^{DCC}	β^{DCC}	ω_1	ω_2	a_{11}	a_{22}	b_{11}	b_{22}	γ_{11}	γ_{22}	a_{12}	a_{21}	b_{12}	b_{21}	LR
Panel A - Full Sample														
0.0159***	0.9801***	0.0071	0.0057***	0.0616***	0.0226***	0.4975***	0.9563***	0.0356***	0.0153***	0.0764***	0.0127***	0.3379***	–	–
(0.0020)	(0.0029)	(0.0100)	(0.0008)	(0.0079)	(0.0021)	(0.0527)	(0.0027)	(0.0109)	(0.0024)	(0.0081)	(0.0014)	(0.0458)		
0.0212***	0.9722***	0.0135**	0.0095***	0.0293**	0.0431***	0.9246***	0.9381***	0.0153	0.0119**	0.0351***	0.0116***	–	–	154.06
(0.0028)	(0.0125)	(0.0059)	(0.0033)	(0.0117)	(0.0086)	(0.0213)	(0.0120)	(0.0109)	(0.0056)	(0.0115)	(0.0041)			[0.000]
0.0122**	0.9853***	0.0089**	0.0062***	0.0342***	0.0396***	0.9634***	0.9585***	–	–	–	–	–	–	705.24
(0.0049)	(0.0068)	(0.0040)	(0.0020)	(0.0102)	(0.0064)	(0.0107)	(0.0068)							[0.000]
Panel B - Contract with maturity in December 2008														
0.0392***	0.8810***	0.1899**	0.0148**	0.0287	0.0425**	0.4674***	0.9390***	−0.0021	0.0177	0.1881***	0.0058	0.2699**	–	–
(0.0123)	(0.0422)	(0.0835)	(0.0068)	(0.0268)	(0.0139)	(0.1427)	(0.0166)	(0.0415)	(0.0113)	(0.0574)	(0.0049)	(0.1096)		
0.0124	0.9858***	0.0496***	0.0118	0.0005	0.361**	0.8949***	0.9406***	0.0181	0.0271***	0.0877***	0.0078	–	–	139.2
(0.0101)	(0.0128)	(0.0190)	(0.0074)	(0.0048)	(0.0157)	(0.0289)	(0.0197)	(0.0148)	(0.0096)	(0.0299)	(0.0054)			[0.000]
0.0283***	0.9052***	0.0514	0.0192**	0.0436*	0.0544***	0.9419***	0.9387***	–	–	–	–	–	–	699.42
(0.0109)	(0.0450)	(0.0333)	(0.0082)	(0.0230)	(0.0145)	(0.0287)	(0.0169)							[0.000]
Panel C - Contract with maturity in December 2009														
0.0155***	0.9796***	−0.0931*	0.0040**	0.0377**	0.0145***	0.0424	0.9664***	0.0391	0.0081**	0.0628***	0.0146***	0.8285***	–	–
(0.0052)	(0.0100)	(0.0523)	(0.0018)	(0.0162)	(0.0038)	(0.3276)	(0.0027)	(0.0248)	(0.0041)	(0.0167)	(0.0050)	(0.3243)		
0.0161***	0.9756***	0.0017	0.0046**	0.0201***	0.0300***	0.9671***	0.9591***	–	–	0.0108***	0.0092**	–	–	23.8
(0.0052)	(0.0103)	(0.0014)	(0.0020)	(0.0056)	(0.0052)	(0.0081)	(0.0085)			(0.0034)	(0.0039)			[0.000]
0.0105***	0.9851***	0.0021**	0.0025**	0.0205***	0.0267***	0.9787***	0.9724***	–	–	–	–	–	–	168.56
(0.0025)	(0.0044)	(0.0010)	(0.0011)	(0.0048)	(0.0049)	(0.0050)	(0.0052)							[0.000]

Notes: The table shows the estimation results of the volatility transmission structure. The first equation represents the spot market, while the second equation refers to the futures market. Panels A, B, and C refer to the full sample, the first subsample and the second subsample. In each Panel the first row represents the general bivariate asymmetric UEDCC-GARCH specification with ARCH and GARCH spillover effects. The specifications represented by the second row only allow for asymmetric effects and ARCH volatility effects. The specifications represented by the third row do neither allow for spillover nor for asymmetric effects. The entries in the last column are likelihood ratio statistics LR, with $LR = 2 \cdot (L_{UR} - L_R)$, where L_{UR} is the likelihood of the unrestricted asymmetric UEDCC-GARCH specification, while L_R is the likelihood of the restricted models. Numbers in parentheses are Bollerslev-Wooldridge robust standard errors, numbers in brackets are p-values. $\star\star\star$, $\star\star$, $\star$ indicate significance at the 1 %, 5 % and 10 % level.

Part II

Energy markets, stock markets, and the macroeconomy

Chapter 4

Carbon and the stock market

4.1 Introduction

With the implementation of the EU-ETS companies covered by the system do have to pay a price for the emission of greenhouse gases.[1] Hence, emission allowances which certify the right to emit a specific amount of CO_2-equivalent can be considered an input factor in the production process of these companies. This in turn generates additional costs which constitute the regulatory burden of the system. From the viewpoint of a company covered by the EU-ETS allowances are assets in the company's balance sheet, while the number of allowances the company receipts is determined by decisions of national authorities approved by the European Commission. If the number of allowances a company holds exceeds the number of allowances it needs to fulfill regulatory requirements the company is characterized by a net-long compliance position and can sell superfluous allowances on the market and thus realize additional revenues.[2] Contrary, a company that is net-short in allowances has to purchase extra allowances on the market. Consequently, profits of companies with different net-compliance positions are likely to be affected heterogeneously by carbon price variations. While companies with net-long positions should rather benefit from increasing carbon prices, the opposite is the case for companies with net-short positions. Kettner et al. (2008), Ellerman and Joskow (2008), and Convery et al. (2008) describe strong heterogeneity in the generosity of allowance allocation across the sectors covered by the EU-ETS which has led to considerable discrepancies in sector-specific net-compliance positions. While electricity companies on average hold net-short positions, non-electricity companies are characterized by net-long positions. Ellerman et al. (2007)

[1] For a detailed description of the trading scheme see e. g. Ellerman et al. (2010) or Chevallier (2012).

[2] Profits generated by the disposal of superfluous allowances should not be confused with windfall profits arising from the incorporation of the market value of allowances allocated free of charge in the electricity price (see Woerdman et al. (2009)).

trace back the sector-specific differences in emissions caps to the policy makers' concerns of negative impacts of the EU-ETS on the competitiveness of non-electricity companies. While electricity companies are assumed to be able to pass through additional costs to the consumer, non-electricity companies face international competition that prevents them from passing through the full costs. This in turn implies that even a net-short electricity company could benefit from increasing carbon prices by incorporating the market price of carbon in the electricity price as far as allowances are (at least partly) allocated free of charge. Besides sector-specific differences in the emissions caps, Kettner et al. (2008) and Ellerman and Joskow (2008) also highlight essential heterogeneity in the stringency of emissions caps across the member states of the European Union.

The main objective of this study is to put the considerations of Kettner et al. (2008), Ellerman and Joskow (2008), and Convery et al. (2008) one step further and analyze whether shareholders of companies that belong to individual sectors covered by the EU-ETS do have to carry regulatory burden of the system or whether they even benefit from the regulation. For this, we analyze the effect of variations in the carbon price on the market value of these companies. Hence, the study investigates the relevance investors attach to the input factor carbon in the valuation of company stock returns and enables us to assess the effects of the carbon price on net profits of investors. We also aim at analyzing changes in the relationship between the carbon and the stock market over the first and the second commitment period.

Our econometric specification is based on multifactor panel regression models (see e.g. Fama and French (1992)) motivated by arbitrage pricing theory. Excess returns on company stocks are explained by excess returns on the market portfolio and energy price excess returns as well as macroeconomic risk factors. In order to account for the specific error structure of large finance panels, we adopt the estimation method of Thompson (2011) which yields consistent estimates of the covariance matrices even if the errors are correlated over time and across companies.

The first result of the paper is the finding of a loose relationship between the carbon market and the stock market in the first commitment period. Allowing carbon price changes to affect stock prices heterogeneously across individual sectors but restricting the effects to be identical across countries reveals that carbon price changes affect stock returns in the electricity sector only.[3] For the electricity sector this confirms the results of Oberndorfer (2009) and Veith et al. (2009) who find a significantly positive effect of carbon price changes on stock returns adopting an estimation strategy that does not explicitly correct covariance matrices for correlation across time and companies. In the more de-

[3] Sectors included in the study are (i) electricity, (ii) iron and steel, (iii) chemicals, and (iv) cement.

tailed country- and sector-specific analysis, we find a positive effect on electricity stock returns of German&Austrian, Scandinavian, and UK companies which is again consistent with the findings of Oberndorfer (2009). For non-electricity companies there is only very weak evidence for such an effect. In contrast to our study neither Oberndorfer (2009) nor Veith et al. (2009) analyze the impact of the carbon price on non-electricity stocks. Both studies justify the restriction to electricity companies by the major portion of allowances allocated to electricity companies.[4] In contrast to these studies we argue that it is essential to not only focus on electricity stocks but also on stocks of companies which belong to other sectors under the EU-ETS. The main argument behind the inclusion of such companies is the considerable heterogeneity in the restrictiveness of sector-specific emissions caps which have led to different net compliance positions across the sectors. Our results imply that policy decisions on the sector- and country-specific stringency of allowance allocation are not reflected in the relationship between carbon and stock returns.

Second, for Phase II our study is the first one to analyze the link between the carbon and the stock market. In contrast to Phase I, we observe a completely different picture. First, the results clearly reveal close links between carbon price changes and stock returns. For electricity companies located in countries with more restrictive emissions caps, we find a negative impact of increasing carbon prices on stock returns. Shareholders of these companies are hit by the regulatory burden of the system. For non-electricity companies we find positive and significant effects, which are stronger in countries with more generous allowance allocation. Hence, in contrast to the concerns of negative impacts of the EU-ETS on the competitiveness of internationally active non-electricity companies (Ellerman et al. (2007)), shareholders of these companies do not bear regulatory burden of the system but benefit from the regulation. The results reflect the sector- and country-specific stringency of emissions caps. Finally, we control for asymmetric price effects. In contrast to the Phase I results of Oberndorfer (2009), for Phase II we find that carbon price increases (decreases) are accompanied by stronger price responses of electricity (non-electricity) stocks than price decreases (increases) of the same size. This result is consistent with the view that with increasing carbon prices investors expect that electricity companies have to purchase additional allowances at advanced prices which in turn negatively affects these companies' profitability.

In conclusion, this study can be interpreted as an evaluation of policy decisions of national authorities and the European Commission. Hence, our findings are relevant not only for market participants in terms of risk management and investment strategy purposes but

[4]Oberndorfer (2009) and Veith et al. (2009) argue that such installations receive almost two-thirds of the total number of allowances. Veith et al. (2009, p. 607) even argue that "*a restriction to the power industry maintains relevance for the trading system as a whole.*".

especially for regulatory authorities themselves.

We organize the remainder of the paper as follows. In Section 4.2, we summarize the related literature, while Section 4.3 outlines the methodology used in the empirical analysis. Moreover, we construct hypotheses concerning the relation between sector-specific returns and the carbon price factor. Section 4.4 describes the data. Section 4.5 summarizes the estimation results and provides an interpretation of the empirical findings. Finally, we provide a policy discussion and the conclusion in Section 4.6.

4.2 Related literature

Recent empirical work in the field of carbon finance often addresses the modeling of the carbon price and price volatility dynamics in one or more markets (see Paolella and Taschini (2008) or Benz and Trück (2009) for analyses of Phase I and Rittler (2012) for Phase II). These studies estimate various GARCH-type specifications to capture specific statistical properties of the data. Moreover, the investigation of fundamental factors driving the carbon price has gained interest in the carbon finance literature. Mansanet-Bataller et al. (2007) and Alberola et al. (2009b) identify energy prices and weather as relevant determinants of the carbon price development during Phase I. This is confirmed by Bredin and Muckley (2011) and Mansanet-Bataller et al. (2011) for Phase II. Moreover, Albreola et al. (2008, 2009) show that the sector production in the sectors electricity generation, paper, and iron affects the carbon price, while the authors reveal heterogeneity in the effects across countries included in the studies. In contrast to the sector-specific analysis of Albreola et al. (2008, 2009), Chevallier (2011) aims at analyzing the impact of aggregate economic activity in the EU on the carbon price. Using a Markov-switching methodology, Chevallier (2011) shows that the effect of the EU industrial production on the carbon price depends on the state of the economy. Finally, in their high-frequency analysis Conrad et al. (2012) show that the carbon price immediately adjusts to the surprise component in macroeconomic and market-specific announcements.

Even though the afore-mentioned studies imply that there is a close relationship between the carbon market and the real economy which could also imply a link between the carbon and the stock market, empirical evidence on the impact of the EU-ETS on the stock performance of companies covered by the system is rare. Given the concerns of negative impacts on the competitiveness of companies under the EU-ETS relative to competitors from outside the EU (see Ellerman et al. (2007) among others), some studies analyze the economic effects of the introduction of the system making use of simulation studies. In the framework of a European energy market model Lise et al. (2010) conclude that electricity

companies can pass through 70%-90% of the carbon price to consumers which in case of free allocation implies the realization of considerable windfall profits. Demailly and Quirion (2008) and Smale et al. (2006) find similar results for the steel sector and Smale et al. (2006) for the cement sector, whereas in these sectors the portion of the carbon price producers can pass through to consumers is smaller than in the electricity sector. In contrast to these simulation studies, Oberndorfer (2009) and Veith et al. (2009) directly assess the impact of changes in the carbon price on the net profits of European electricity companies during Phase I. While Oberndorfer (2009) analyzes spot market data of the European Energy Exchange (Leipzig), Veith et al. (2009) also include futures market data in their analysis.[5] Both studies adopt the methodological framework of multifactor modeling to assess the link between the carbon and the stock market. Excess returns on electricity stocks are explained by excess returns on the market portfolio and excess returns on energy factors including carbon. Oberndorfer (2009) as well as Veith et al. (2009) reveal a weak but significant positive effect of the carbon price on electricity stock returns. In addition, Oberndorfer (2009) does not find evidence for asymmetric effects of positive and negative carbon price changes. Veith et al. (2009) show that the carbon price does not affect stock returns of electricity companies which generate electricity based on green technologies that avoid the emission of greenhouse gases.
In a related strand of the literature Convery et al. (2008), Kettner et al. (2008), and Ellerman and Buchner (2008) reveal considerable heterogeneity in the net-compliance positions across the sectors covered by the EU-ETS. Figure 4.1 shows these specific net compliance positions as percentage of allocated allowances for selected sectors. The numbers are taken from Kettner et al. (2008, p. 52). Apparently, the sectors pulp and paper, iron and steel, and cement and lime are characterized by net-long positions, while the figure presents a net-short position for the power and heat generating sector. Ellerman et al. (2007) point out that the sector-specific differences reflect the concerns of policy makers that in contrast to electricity companies, non-electricity companies cannot pass through the full carbon price to consumers because of international competition.[6]

Insert Figure 4.1 about here.

Besides varying allocation patterns across individual sectors, Kettner et al. (2008) and Ellerman and Buchner (2008) show that there is also strong evidence for heterogeneity

[5]Note that Paolella and Taschini (2008) hint at the problem of extremely low trading activity in the carbon market before 2006. Moreover, Chevallier (2011) points out that in Phase I spot prices are contaminated by the carbon market crash in April and May 2006.

[6]This argument is consistent with the regulatory framework of the EU-ETS. According to Directive 2009/29/EC, the allocation of allowances to electricity companies is more restrictive compared to companies of other sectors.

in the stringency of country-specific emissions caps. Figure 4.2 summarizes the net compliance positions of selected member states. The numbers are taken from Kettner et al. (2008, pp. 46).

Insert Figure 4.2 about here.

For Germany, the Scandinavian countries, and countries of Western Europe, the figure presents net-long positions while the opposite is the case for Italy, Spain, Ireland and the UK. Finally, Ellerman and Joskow (2008) explain that national authorities and the European Commission further tightened the Phase II emissions caps of Italy, Spain and the UK compared to 2005 verified emissions of these countries.

More recently, Bushnell et al. (2011) adopt the framework of an event study analysis to capture the impact of the carbon price on abnormal returns of companies that belong to different sectors covered by the EU-ETS during the carbon market breakdown in April/May 2006. The authors find that net compliance positions indeed affect the stock price response during this specific period.

4.3 Methodology and hypothesis construction

4.3.1 Estimation strategy

The factor model we estimate to capture the link between the carbon and the stock market is given by the panel regression

$$r_{it} = \mathbf{x}_{it}'\boldsymbol{\beta} + \varepsilon_{it}, \tag{4.1}$$

with $i = 1, ..., n$ companies and $t = 1, ..., T$ time periods. r_{it} are the company stock returns, $\mathbf{x}_{it}$ is a vector of explanatory variables and $\boldsymbol{\beta}$ contains the corresponding factor loadings. ε_{it} is a potentially heteroskedastic error term with $\mathbf{E}(\varepsilon_{it}|\mathbf{x}_{it}) = 0$. We assume that the error can be correlated over time for a particular company (time effect), that is $\mathbf{E}(\varepsilon_{it}\varepsilon_{ik}|\mathbf{x}_{it}, \mathbf{x}_{ik}) \neq 0$. Moreover, we assume that the error can also be correlated across companies at a given period of time (company effect), that is $\mathbf{E}(\varepsilon_{it}\varepsilon_{jt}|\mathbf{x}_{it}, \mathbf{x}_{jt}) \neq 0$.

Under these assumptions, we have to apply a specific estimation strategy in order to obtain unbiased standard errors. According to Petersen (2009), former approaches like Fama and MacBeth (1973), Huber (1967), or Rogers (1983) assess the problem of error correlation only in one dimension (either correlations over time or correlations across companies). Thompson (2011) suggests an alternative procedure that generates consistent standard errors in the presence of correlation simultaneously over time and across companies. In the panel regression described by Equation (4.1) the OLS-estimator of the coefficient vector can be written as

$$\hat{\boldsymbol{\beta}} = \mathbf{H}^{-1}\sum_{i,t}\mathbf{x}_{it}y_{it} = \boldsymbol{\beta} + \mathbf{H}^{-1}\sum_{i,t}\mathbf{x}_{it}\varepsilon_{it}, \tag{4.2}$$

with $\mathbf{H} = \sum_{i,t}\mathbf{x}_{it}\mathbf{x}'_{it}$. The variance of $\hat{\boldsymbol{\beta}}$ can be approximated by $\mathbf{H}^{-1}\mathbf{G}\mathbf{H}^{-1}$ in large samples, where $\mathbf{G}$ is given by $\mathbf{G} = \mathbf{Var}\left(\sum_{i,t}\mathbf{x}_{it}\varepsilon_{it}\right) = \sum_{i,j,t,k}\mathbf{E}\left(\mathbf{x}_{it}\varepsilon_{it}\varepsilon_{jk}\mathbf{x}'_{jk}\right)$. Thompson (2011) shows that under the given error assumptions, the covariance matrix that is robust with respect to correlation over time and across companies is given by

$$\mathbf{Var}(\hat{\boldsymbol{\beta}}) = \underbrace{\mathbf{H}^{-1}\sum_{i}\mathbf{c}_i\mathbf{c}'_i\mathbf{H}^{-1}}_{\mathbf{V}_{company}} + \underbrace{\mathbf{H}^{-1}\sum_{t}\mathbf{s}_t\mathbf{s}'_t\mathbf{H}^{-1}}_{\mathbf{V}_{time}} - \underbrace{\mathbf{H}^{-1}\sum_{i}\sum_{t}\mathbf{x}_{it}\varepsilon^2_{it}\mathbf{x}'_{it}\mathbf{H}^{-1}}_{\mathbf{V}_{white}}, \tag{4.3}$$

where $\mathbf{c}_i = \sum_t \mathbf{x}_{it}\varepsilon_{it}$ is the sum over all observations for company i and $\mathbf{s}_t = \sum_i \mathbf{x}_{it}\varepsilon_{it}$ is the sum over all observations for time t. Replacing $\mathbf{c}_i$ and $\mathbf{s}_t$ by their estimates $\hat{\mathbf{c}}_i = \sum_t \mathbf{x}_{it}\hat{\varepsilon}_{it}$ and $\hat{\mathbf{s}}_t = \sum_i \mathbf{x}_{it}\hat{\varepsilon}_{it}$ with $\hat{\varepsilon}_{it} = r_{it} - \mathbf{x}'_{it}\hat{\boldsymbol{\beta}}$ leads to

$$\widehat{\mathbf{Var}}(\hat{\boldsymbol{\beta}}) = \hat{\mathbf{V}}_{company} + \hat{\mathbf{V}}_{time} - \hat{\mathbf{V}}_{white}, \tag{4.4}$$

where $\hat{\mathbf{V}}_{company} = \mathbf{H}^{-1}\sum_i \hat{\mathbf{c}}_i\hat{\mathbf{c}}'_i\mathbf{H}^{-1}$ is the formula for the estimation of standard errors clustered by companies, $\hat{\mathbf{V}}_{time} = \mathbf{H}^{-1}\sum_t \hat{\mathbf{s}}_t\hat{\mathbf{s}}'_t\mathbf{H}^{-1}$ is the formula for the estimation of standard errors clustered by time, and $\hat{\mathbf{V}}_{white} = \mathbf{H}^{-1}\sum_i\sum_t \hat{\varepsilon}^2_{it}\mathbf{x}_{it}\mathbf{x}'_{it}\mathbf{H}^{-1}$ is the covariance matrix of White (1980) which is robust to heteroskedasticity.

4.3.2 Empirical specifications

In this section we describe the vector of explanatory variables $\mathbf{x}_{it}$ and the corresponding factor loadings $\boldsymbol{\beta}$ of the general model defined in the previous section. In our baseline specification $\mathbf{x}'_{it}\boldsymbol{\beta}$ is given by

$$\beta_0 + \beta_m r_{m,t} + \beta_{br} r_{br,t} + \beta_g r_{g,t} + \beta_c r_{c,t} + \beta_{ts}\Delta ts_t + \beta_{ds}\Delta ds_t. \tag{4.5}$$

In this model, we control for the excess returns on the market portfolio, $r_{m,t}$ and for the excess returns on the energy commodities oil, $r_{br,t}$, gas, $r_{g,t}$, and carbon, $r_{c,t}$. In addition, we augment the specification by two macroeconomic risk factors, the term spread, ts_t, and the default spread, ds_t. In case of both variables we construct first differences to obtain stationary series. This specification does not account for sector-specific effects. Rather,

the effect of the carbon factor is restricted to be identical across all sectors.
The second specification explicitly allows for sector-specific effects, in order to ensure that the carbon price factor as well as the other energy price factors can affect stock returns of companies operating in different sectors heterogeneously. In this model $\mathbf{x}'_{it}\boldsymbol{\beta}$ is given by

$$\begin{aligned}&\beta_0 + \sum_{s\in S}\beta_{0,s}D_s + \beta_m r_{m,t} + \beta_{br} r_{br,t} + \beta_g r_{g,t} + \beta_c r_{c,t} + \beta_{ts}\Delta ts_t + \beta_{ds}\Delta ds_t \\ &+ \sum_{s\in S}\beta_{m,s}D_s r_{m,t} + \sum_{s\in S}\beta_{br,s}D_s r_{br,t} + \sum_{s\in S}\beta_{g,s}D_s r_{g,t} + \sum_{s\in S}\beta_{c,s}D_s r_{c,t},\end{aligned} \tag{4.6}$$

where $S = \{steel,\ chemicals,\ cement\}$ and D_s is a dummy variable that equals one if the corresponding observation belongs to sector s and zero elsewise. β_c describes the effect of the carbon price factor on electricity stock returns, $\beta_c + \beta_{c,s}$ captures the effect on stock returns of companies operating in sector s. Analogously, we construct the sector-specific effects of the oil and the gas price factors. Sector-specific constants are given by β_0 for electricity companies and by $\beta_0 + \beta_{0,s}$ for sector s. To check whether the carbon price significantly affects electricity stock returns we test the hypothesis $H_0 : \beta_c = 0$ against the alternative $H_1 : \beta_c \neq 0$. To check whether the carbon price significantly affects stock returns of sector s, $s \in \{steel,\ chemicals,\ cement\}$, we test the hypothesis $H_0 : \beta_c + \beta_{c,s} = 0$ against the alternative $H_1 : \beta_c + \beta_{c,s} \neq 0$. Accordingly, we perform similar tests to assess the sector-specific impact of the oil and gas price. The F-distributed test statistics are computed on the basis of the covariance matrix given by Equation (4.4).
In order to capture individual effects of the carbon price on country- and sector-specific company stock returns, we estimate a further specification in which $\mathbf{x}'_{it}\boldsymbol{\beta}$ is given by

$$\begin{aligned}&\beta_0 + \sum_{ind\in I}\sum_{r\in R}\beta_{0,ind,r}D_{ind}D_r + \beta_m r_{m,t} + \beta_{br} r_{br,t} + \beta_g r_{g,t} + \beta_c r_{c,t} + \beta_{ts}\Delta ts_t + \beta_{ds}\Delta ds_t \\ &+ \sum_{ind\in I}\beta_{m,ind}D_{ind} r_{m,t} + \sum_{ind\in I}\beta_{br,ind}D_{ind} r_{br,t} + \sum_{ind\in I}\beta_{g,ind}D_{ind} r_{g,t} + \sum_{ind\in I}\beta_{c,ind}D_{ind} r_{c,t} \\ &+ \sum_{r\in R}\beta_{c,ind,r}D_{ind}D_r r_{c,t}.\end{aligned} \tag{4.7}$$

D_{ind} is a dummy variable that equals one if the corresponding observation belongs to sector ind, $ind \in \{electricity,\ steel,\ chemicals,\ cement\}$, and zero elsewise, while D_r equals one if the corresponding observation belongs to region r, $r \in R$ with $R = \{northern\ europe,$ $western\ europe,\ southern\ europe,\ united\ kingdom\}$ and zero elsewise.[7] We individually

[7] Countries assigned to Northern Europe are Denmark, Finland, Norway, and Sweden. Western Europe consists of France, Belgium, The Netherland, and Luxembourg. Southern Europe contains Portugal, Spain, Italy and Greece. Great Britain and Ireland are assigned to the region United Kingdom.

estimate Equation (4.7) for each sector. The set I contains all sectors apart from the sector for which we estimate the model. To assess the impact of the carbon price factor on electricity stock returns, we estimate Equation (4.7) for the electricity sector, such that $I = \{steel,\ chemicals,\ cement\}$. Hence, β_c represents the effect of the carbon price factor on German&Austrian electricity stock returns, while $\beta_c + \beta_{c,electricity,r}$ captures the effect on stock returns of electricity companies located in region r. Accordingly, we compute the country-specific effects for the other sectors. To check whether the carbon price factor significantly affects German&Austrian electricity stock returns we test the hypothesis $H_0 : \beta_c = 0$ against the alternative $H_1 : \beta_c \neq 0$. To check whether the carbon price factor affects electricity stock returns in region r we test the hypothesis $H_0 : \beta_c + \beta_{c,electricity,r} = 0$ against the alternative $H_1 : \beta_c + \beta_{c,electricity,r} \neq 0$.
Following Zachmann and von Hirschhausen (2008) and Oberndorfer (2009), we investigate the impact of asymmetric effects. That is, we analyze whether there are different responses to increasing and decreasing carbon prices in sector-specific stock returns. In the corresponding specification $\mathbf{x}'_{it}\boldsymbol{\beta}$ is given by

$$\begin{aligned}
&\beta_0 + \sum_{ind \in I} \sum_{r \in R} \beta_{0,ind,r} D_{ind} D_r + \beta_m r_{m,t} + \beta_{br} r_{br,t} + \beta_g r_{g,t} + \beta_c r_{c,t} + \beta_{ts} \Delta ts_t + \beta_{ds} \Delta ds_t \\
&+ \sum_{ind \in I} \beta_{m,ind} D_{ind} r_{m,t} + \sum_{ind \in I} \beta_{br,ind} D_{ind} r_{br,t} + \sum_{ind \in I} \beta_{g,ind} D_{ind} r_{g,t} + \sum_{ind \in I} \beta_{c,ind} D_{ind} r_{c,t} \\
&+ \sum_{r \in R} \beta_{c,ind,r} D_{ind} D_r r_{c,t} + \beta_{pos} D_{pos} r_{c,t} + \sum_{ind \in I} \beta_{pos,ind} D_{pos} D_{ind} r_{c,t} \\
&+ \beta_{c,ind,pos,g\&a} D_{pos} D_{g\&a} D_{ind} r_{c,t} + \sum_{r \in R} \beta_{c,ind,pos,r} D_{ind} D_r D_{pos} r_{c,t}. \qquad (4.8)
\end{aligned}$$

We introduce the dummy variable D_{pos} that equals one for positive carbon price changes and zero elsewise. As Equation (4.7), we estimate Equation (4.8) individually for each sector. To capture asymmetric effects of the carbon price on electricity stock returns we estimate Equation (4.8) for the electricity sector. Hence, $\beta_c + \beta_{c,pos} + \beta_{c,elec,pos,g\&a}$ (β_c) describes the effect of increasing (decreasing) carbon prices on the stock returns of German&Austrian electricity companies. $\beta_c + \beta_{c,elec,r} + \beta_{pos} + \beta_{c,elec,pos,r}$ ($\beta_c + \beta_{c,elec,r}$) represents the effect of increasing (decreasing) carbon prices on stock returns of electricity companies located in region r.

4.3.3 Hypothesis construction

From the viewpoint of companies under the EU-ETS, allowances constitute an input factor in the production process which is comparable to other production factors like energy commodities (see also Benz and Trück (2009) and Borak et al. (2006)). Formally, emission allowances are assets in a company's balance sheet such that changes in the carbon price also affect the market value of the company which could influence the company's stock price. Based on these considerations we construct the following hypotheses.

Hypothesis H1: Carbon price changes do not affect company stock returns.

Alberola et al. (2008, 2009b) and Chevallier (2011) show that sectoral and aggregate industrial production affect the carbon price. Moreover, for Phase I Oberndorfer (2009) and Veith et al. (2009) find that the carbon price positively influences electricity stock returns. Thus, we expect that hypothesis *H1* has to be rejected.

Hypothesis H2: The effect of the carbon price on stock returns is stable over both phases.

First, Borak et al. (2006) point out different statistical properties of the carbon price dynamics (see also Figure 1). Second, according to Ellerman and Buchner (2008) and Ellerman and Joskow (2008) there is a fundamental difference in the net-compliance positions of the whole system. While due to the overallocation with emission allowances in Phase I the system on average is long, there is a short position of the whole system in Phase II. Thus, we expect the impact of the carbon price on company stock returns to vary over both commitment periods so that hypothesis *H2* is expected to be rejected.

Hypothesis H3: The effect of the carbon price on stock returns is stable across sectors.

As aforementioned, main determinants of the carbon price impact on company stock returns are the company's net-compliance position and its ability to pass through the carbon price to consumers. Figure 4.1 shows considerable differences in the net compliance positions across the sectors covered by the EU-ETS. In Phase I the electricity sector is the only sector characterized by a net-short compliance position, while all other sectors are net-long. Ellerman and Joskow (2008) describe a similar picture for the second commitment period. Moreover, the portion of regulatory costs companies can pass through to consumers also seems to vary across sectors (see Smale et al. (2006) and Lise et al. (2010)). Hence, we expect stock returns of companies of different sectors to be affected heterogeneously by carbon price changes such that hypothesis *H3* has to be rejected.

Hypothesis H4: The effect of the carbon price on company stock returns is stable across participating member states.

Besides sector-specific differences in the stringency of emissions caps, Ellerman and Joskow (2008) also highlight heterogeneity in the generosity of allowance allocation across member states (see also Figure 4.2). We expect stock returns of companies located in countries with more restrictive emissions caps to be more (less) sensitive to increasing carbon prices in sectors characterized by net-short (net-long) compliance positions. On the other hand, we expect stock returns of companies located in countries with less restrictive emissions caps to be less (more) sensitive to increasing carbon prices in sectors characterized by net-long (net-short) compliance positions. Consequently, we conjecture that hypothesis *H4* has to be rejected.

Hypothesis H5: The asymmetric effects of the carbon price on company stock returns is identical across sectors.

Given the existence of sector-specific effects of the carbon price on stock returns, decreasing and increasing carbon price changes could induce heterogenous stock price responses which also differ across the individual sectors. While Oberndorfer (2009) does not find asymmetric effects on electricity stock returns during Phase I, Zachmann and von Hirschhausen (2008) reveal that electricity prices are asymmetrically affected by positive and negative carbon price changes.

4.4 Data

4.4.1 Carbon market data

Our sample covers the period from January 2006 to June 2010. We make use of carbon price data of the European Climate Exchange (London) as this is the most liquid market for EUAs. We drop all bank holidays and all trading days with non-regular trading as e.g. December 31. Based on daily closing prices of futures contracts with maturities in December 2008, 2009, and 2010[8], we construct daily log-returns according to

$$r_{c,t} = 100 \cdot \left[ln(p_{c,t}) - ln(p_{c,t-1}) \right], \tag{4.9}$$

and compute excess returns by subtracting the risk free interest rate from the daily returns. For the risk free returns we use the 3-month EURIBOR on a daily basis.
Our full sample covers a substantial part of the whole development of the price dynamics

[8]For the period January 1, 2006 to December 15, 2008 we use the contract with maturity in December 2008. For the period December 16, 2008 to December 14, 2009 we adopt the December 2009 contract and for the post December 14, 2009 period we consider the December 2010 futures contract. This implies that we use the contract with the highest trading activity as measured by transaction volume for each period.

in the EU-ETS since its formal implementation. The full sample consists of the two subsamples that constitute the first and the second commitment period. The first subsample covers data of the period January 2006 to April 2008, the second subsample spans the period May 2008 to June 2010. This translates to 1139 observations in the full period and to 592 (547) observations in the first (second) subsample, respectively. Figure 4.3 shows the carbon price development over the full sample.

Insert Figure 4.3 about here.

Apparently, the carbon price is more volatile in the beginning of the sample. Moreover, for the early stage of the EU-ETS the figure shows the carbon market breakdown in April and May 2006, where the carbon price decreases by about 50 percent within only two weeks. After a further decrease to 12 Euros, a strong upward trend follows until July 2008. Then, accompanied by the economic and financial crisis the carbon price decreases to less than 10 Euros. Towards the end of the sample period the carbon price slightly increases and realizes values of about 15 Euros.[9] Table 4.1 summarizes the descriptive statistics of the carbon returns.

Insert Table 4.1 about here.

4.4.2 Stock market data

Motivated by the sector-specific compliance positions (see Figures 4.1 and 4.2), besides electricity companies we include companies that belong to the sectors pulp and paper, iron and steel, and cement and lime. Furthermore, we also consider companies of the chemical sector since these companies also own installations covered by the EU-ETS. According to the industry and sector classification of the EURO STOXX 600, we include each company listed in the index over the full period.[10] As a whole, the corresponding sample consists of 16 electricity, 11 steel, 16 chemical, and 7 cement companies.[11] Table 4.2 summarizes these four sectors and the classification criteria according to the EURO STOXX 600. Table 4.3 provides detailed descriptions of all companies.

Insert Tables 4.2 and 4.3 about here.

[9] For a more detailed description of the price development in the EU-ETS we refer to Chevallier (2011) and Chevallier and Sevi (2011) and the references therein.

[10] All company data is taken from DataStream.

[11] The EURO STOXX 600 does also contain stocks of companies located in Switzerland. However, the EU-ETS does not cover Switzerland. Consequently, we dropped these observations to avoid dilution of the estimation results.

For electricity stocks, we exclusively consider companies that belong to the subsectors conventional electricity or multiutilities. We do not include companies of the subsector alternative electricity. According to the supersector classification of the EURO STOXX 600, we assign the single company of the pulp and paper sector to the group of steel stocks. Following Equation (4.9), we construct company stock returns based on daily closing prices, where we first transform all prices into Euros making use of the corresponding exchange rates. Then, we compute excess returns by subtracting the risk free rate from the daily returns (see Table 4.1 for the descriptive statistics of the company stock returns).

4.4.3 Control variables

In the previous asset pricing literature, Fama and French (1992) among many others identify variables which are relevant pricing factors for company stock returns. This set contains an overall market factor, firm-related factors as well as macroeconomic risk factors. More recently, Bali and Engle (2010) confirm the relevance of these factors for company stock valuation in the framework of the Intertemporal CAPM. Moreover, they argue that each variable that influences a company's market value potentially constitutes a pricing factor. Following Obernorfer (2009) and Veith et al. (2009) we also include oil and gas price changes in addition to the carbon price factor.

The EURO STOXX 600 index serves as a proxy for the market portfolio as this index is the broadest European index that covers the 600 largest listed European companies. For the oil price we take the one month ahead crude oil brent futures contract. Finally, the price of the one year ahead gas futures contract traded at the Anglo-Dutch energy exchange (APX-ENDEX) is used for the gas price. Prices of contracts in foreign currencies are transformed into Euros. Then, we compute excess returns of the market portfolio, the oil price, and the gas price on the basis of daily closing prices according to Equation (4.9).[12] Table 4.1 presents the descriptive statistics. We also control for the default spread and the term spread. We define the default spread as the difference between daily yields on BAA-rated and AAA-rated corporate bonds. Finally, the term spread equals the difference between daily yields on 10-year and 3-month bonds of the ECB. In order to have stationary series, we take the first difference of each macroeconomic risk factor. Again, Table 4.1 shows the descriptive statistics.

[12] Linking the gas price series of contracts with different maturities induces extreme jumps. In order to make sure that these jumps do not contaminate our estimation results, we replace the corresponding returns by the mean of the return of the last day before and the return of the first day after the linking of both series.

4.5 Empirical results

4.5.1 Baseline specification

We first show the results of the initial specification. Equation (4.5) describes the model which explains company stock returns by returns on the market portfolio and by energy returns as well as by macroeconomic risk factors. We estimate Equation (4.5) based on data of all companies. That is, we do not account for potential heterogeneity across companies of different sectors. Table 4.4 presents the estimation results, where the columns refer to the full sample as well as to the first and to the second subsample, respectively. The subsample analysis allows us to capture changes in investors' expectations concerning the impact of the carbon price on stock returns.

Insert Table 4.4 about here.

First, Table 4.4 shows that the market beta is highly significant for each sample and takes on values about one. Compared to Oberndorfer (2009) and Veith et al. (2009) the market beta is slightly larger which we attribute to the inclusion of a broader set of companies in our study. Second, the table shows that the estimated coefficients on the energy price factors considerably vary over the two subsamples. While the oil price factor is positive and significant in the first subsample, the gas price factor is only marginally significant in this period. Moreover, the table shows that there is also a difference concerning the impact of the carbon price. Only in the second subsample this factor affects company stock returns significantly. This in turn, can be interpreted as first evidence that in Phase II shareholders on average benefit from the regulation of the EU-ETS. Most importantly, the table shows that there seems to be a change in the investors' valuation of the carbon price factor's relevance over the two phases.

Finally, we refer to the results concerning the macroeconomic risk factors.[13] In the second subsample, the default spread is negatively related to company stock returns at the 1% level. The impact of this factor in the second subsample has more than doubled compared to Phase I. Reasons for this finding could be the recession and the financial crisis that mainly occur in this period. In contrast, we do not find a significant impact of the term spread on company stock returns for any subsample.

In conclusion, the initial analysis indicates that the effects of the pricing factors on stock returns vary over Phase I and II, which requires to perform subsample analysis. Consequently, hypotheses *H1* and *H2* as constructed in Section 4.3.3 have to be rejected.

[13]Note that Chevallier (2009) finds a loose relationship between macroeconomic risk factors and the carbon price.

Moreover, the results imply that we should control for the impact of the macroeconomic risk factors in the subsequent analysis since at least the default spread affects stock returns.

4.5.2 Sector-specific analysis

As pointed out in Sections 4.3.3 and 4.4.2 heterogeneity in sector-specific net compliance positions could affect the impact of the carbon price factor on company stock returns. In the previous regressions we do not explicitly allow for this, since the effect of the carbon price on stock returns is restricted to be identical across all sectors. Now, we augment the initial specification by interaction terms which control for sector-specific effects. Equation (4.6) describes the model and Table 4.5 presents the estimation results.

Insert Table 4.5 about here.

The table shows that the interaction terms between the market returns and the sector-specific dummy variables are highly significant for each period and sector which implies that the stock price response to price changes in the market portfolio differs across the individual sectors. Further, the table confirms the estimation results on the macroeconomic risk factors obtained in the previous section. For the gas price factor we do not find any significant interaction terms at all, while there is weak evidence for heterogeneity in the effects of oil price changes on sector-specific stock returns. Most importantly, for the carbon price factor we find that 5 out of 9 interaction terms are significantly different from zero, while the magnitude of the estimated coefficients on the interaction terms is higher in Phase II. These findings reveal that the carbon price factor affects stock returns across individual sectors differently. In order to formally capture the sector-specific effects, we perform F-tests to check the joint significance of the corresponding variables as described in Section 4.3. Table 4.6 shows the estimated effects and the corresponding values of the F-distributed test statistics in brackets.

Insert Table 4.6 about here.

First, we summarize the results concerning the sector-specific impact of price changes in the market portfolio. For each sample period the market beta for electricity stocks is considerably smaller than 1 which implies that electricity stocks tend to pay off less than the market when the market return is high. In strong contrast, steel and cement stocks rather tend to pay off more than the market, when the market return is high as indicated by a market beta that is considerably greater than 1 for such stocks. Finally, given a market beta close to 1, chemical stocks perform similar to the market portfolio. We now turn to the sector-specific effects of the oil and gas price factors. For the electricity sector we find that neither the oil nor the gas price factor significantly affects stock returns. For Phase I, the findings concerning the gas price factor are predominantly in line with Veith et al. (2009) and Oberndorfer (2009), while depending on their econometric specification these authors find positive or negative effects of the oil price factor. In the steel sector increasing oil prices are accompanied by increasing stock prices (at least in the full sample and the first subsample), while there is no effect of gas price changes on stock returns. For the chemical sector we do not find any significant effects of the gas and oil price factors. Finally, the results for the cement sector show a negative and highly significant effect of the oil price factor and no significant effect of the gas price factor in Phase II and the full sample, while we do not find any significant effects of these two factors during Phase I.

Next, we investigate the relation between sector-specific returns and the carbon price factor. For the full sample, we do not find a significant effect of this factor on electricity stock returns. For the first subsample, the estimated coefficient is positive and significant at the 1% level. This finding confirms the Phase I results of Veith et al. (2009) and Oberndorfer (2009) and implies that increasing carbon prices are accompanied by increasing electricity stock prices, even though the electricity sector is net-short in this period (see Kettner et al. (2008)). This in turn could reflect the realization of windfall profits caused by the free allocation of allowances in Phase I. Consequently, on average owners of electricity companies benefit from the regulation under the EU-ETS in the pilot period. In strong contrast, for Phase II we find a completely different picture for the electricity sector. Now, the effect of the carbon price factor on electricity stock returns is negative and the magnitude of the effect is similar to the one of Phase I. However, the effect is not statistically significant. Despite the net-short position of the electricity sector in Phase II (see in particular Ellerman and Joskow (2008)), the results are consistent with the view that investors expect changing carbon prices to not affect the market value of electricity companies on average. Hence, the result does not point to the shift of regulatory burden to the owners of electricity companies.

For the steel sector, we observe a completely different picture. Even though the system as a whole and the steel sector in particular are characterized by net-long compliance positions in Phase I, we do not find a significant effect of the carbon price factor on steel stock returns for this period. Consequently, investors do not consider the carbon price factor to be relevant in the valuation of steel stocks. This also implies that shareholders of steel companies do not carry regulatory burden of the system in Phase I. Again, looking at the results of Phase II reveals a considerably different picture. For this period, we find a positive and highly significant effect of the carbon price factor on steel stock returns, which implies the high relevance of the factor in the valuation of steel company stocks. A 1% carbon price increase is accompanied by a 0.06% stock price increase. This result points to the realization of regulatory profits of owners of steel companies. With increasing carbon prices, steel companies can sell superfluous allowances at advanced prices which finally leads to positive stock price reactions. In addition, our findings clearly document that the magnitude of the effect of the carbon price factor on steel stock returns is considerably higher compared to all other sectors. This reflects the steel sector's most significant net-long position among all sectors covered by the system as described by Ellerman and Joskow (2008) and illustrated by Figure 4.1, such that the number of superfluous allowances that steel companies can sell at advanced prices is higher than in other sectors.

The findings for the chemical sector are pretty similar to those of the steel sector. For Phase I, we do not find a significant effect of the carbon price on stock returns. For Phase II, we reveal a positive significant effect that is considerably stronger than the one of Phase I but smaller than the Phase II effects in the steel sector and the cement sector. In particular, a 1% carbon price increase is accompanied by a 0.04% stock price increase. In conclusion, the sector-specific analysis reveals that on average shareholders of chemical companies do neither benefit from the EU-ETS nor do they carry regulatory burden of the system in Phase I, while in the Phase II such shareholders benefit from increasing carbon prices.

Finally, for the cement sector, the results are consistent with those of the steel sector and the chemical sector. This holds for both commitment periods. For Phase I, we find a negative but insignificant effect of the carbon price on cement stock returns indicating that shareholders of cement companies are not hit by the regulatory burden of the system nor that they benefit from the regulation. In contrast, the carbon price factor positively affects cement stock returns in Phase II indicating the factor's relevance in the valuation of cement company stocks, where a 1% carbon price increase is accompanied by a 0.04% stock price increase. The results for Phase II imply that owners of cement companies benefit from the regulation of the EU-ETS. However, the effect is smaller than the one in the steel sector.

In conclusion, the results of the sector-specific analysis point to a considerable change in the investors' perception of the relevance of the carbon price factor in the stock valuation of companies covered by the EU-ETS such that hypothesis *H3* has to be rejected. Generally, the link between the carbon and the stock market is rather loose in Phase I. Positive effects are only found for electricity companies implying the realization of regulatory profits of the corresponding companies' owners in this period. For Phase II, there is evidence for a closer link between the carbon and the stock market. While owners of non-electricity companies benefit from the EU-ETS and realize regulatory profits, shareholders of electricity companies do neither carry regulatory burden of the system, nor do they benefit.

4.5.3 Country and sector-specific analysis

By additionally controlling for country-specific effects, we now analyze whether the impact of the carbon price on sector- and country-specific company stock returns also reflects the heterogeneity in the stringency of country-specific emissions caps. This allows us to analyze whether regulatory burden of the system is heterogeneously shifted to shareholders of companies located in different regions. We estimate the specification given by Equation (4.7) individually for each sector for Phases I and II.[14] In contrast to the analysis conducted in the previous section, we have to compute the F-statistics on the basis of the White covariance matrix or the covariance matrix clustered by time. The reason for this is the small number of companies per country and sector which can induce misleading results by computing standard errors clustered by firms. Table 4.7 presents the estimation results, where F-statistics are computed on the basis of the covariance matrix clustered by time.[15]

Insert Table 4.7 about here.

The table shows that pooling the effects of the carbon price factor on company stock returns over all countries leads to a number of spurious conclusions. First, we summarize the results for Phase I. We find positive and significant effects of the carbon price factor on electricity stock returns in Germany&Austria, Northern Europe, and the UK. This is mainly in line with Oberndorfer (2009) and with our sector-specific results. Hence, our estimation results imply that owners of electricity companies located in these regions benefit from the regulation under the system. For the remaining regions Southern Europe and Western Europe, we do not find any significant effect. Interestingly, countries that belong to regions in which stock returns are significantly affected by the carbon price exhibit strong heterogeneity in the stringency of allowance allocation. The same is the case for countries that belong to regions in which the carbon price factor does not affect stock returns. Hence, the results do not reflect the country-specific net compliance positions illustrated in Figure 4.2. For the steel sector the findings of Section 4.5.2 are broadly

[14]Due to space considerations we do not report the coefficient estimates of the respective models. Results are available upon request from the author.

[15]The F-statistics and the corresponding p-values obtained from the specification with White standard errors are very similar.

confirmed. Only in case of Western Europe the carbon price factor significantly affects steel stock returns at the 15% level, while we do not find significant effects in any of the other regions. For the sectors chemicals and cement we do not find any links to the price of carbon at the 10% significance level. Hence, for these sectors the results support the findings of the sector-specific analysis.

Second, we refer to the results of Phase II. Taking country-specific effects into account reveals some further interesting relationships between the carbon price factor and company stock returns during this commitment period. Despite the insignificant effect of the carbon price factor on electricity stock returns in the sector-specific analysis (see Table 4.6), Table 4.7 shows that the stock returns of electricity companies located in Southern Europe and the UK are negatively affected by increasing carbon prices (at least at the 5% level). In both regions a 1% carbon price increase is accompanied by a stock price decrease of 0.055%. The negative effects closely reflect the net-short position of the electricity sector in combination with less generous allocation of allowances in countries covered by these regions as pointed out by Kettner et al. (2008) and Ellerman and Joskow (2008) and shown in Figure 4.1. In order to fulfill regulatory requirements, these companies do have to purchase additional allowances on the market at advanced prices due to the net-short position which in turn induces higher production costs and reduces their profitability such that finally stock prices decrease. In strong contrast, for Northern Europe we even find a positive and significant effect which is consistent with the least restrictive emissions caps in the countries of this region. For the other regions, we do not find a significant impact of the carbon price factor on electricity stock returns. The results imply that in electricity stock valuation investors seem to take sector- and country-specific compliance positions into account. This in turn points to a considerable change in the investors' perception of the relevance of the allowance price in electricity stock valuation over both commitment periods. Most importantly, the analysis clearly shows that regulatory burden that is shifted to the owners of electricity companies is carried by owners of companies located in countries with more restrictive emissions caps. In contrast, owners of companies located in regions with more generous allowance allocation are not hit by the regulatory burden. Restricting the effects to be identical across all regions as in the sector-specific analysis leads to the spurious conclusion of no significant link between the price of carbon and the performance of electricity stocks. Moreover, the result contradicts the conjecture that electricity companies can pass-through the full carbon price to consumers as put forward in simulation studies (see e.g. Lise et al. (2010)).

For steel companies the findings also extend the results of the sector-specific analysis. We find the strongest effects for Western Europe and Germany&Austria where a 1% carbon price increase is accompanied by a stock price increase of 0.22% and 0.09%. This implies that the effect is more pronounced for companies located in countries with less restrictive emissions caps. In line with more restrictive caps in Southern Europe, we do not observe a significant effect of the carbon price factor on steel stock returns. Surprisingly, we also fail to observe a significant impact on the performance of Northern European steel stocks. For Western Europe and Germany&Austria investors consider the carbon price factor to be highly relevant for steel stock valuation. Moreover, the results are consistent with the view that with increasing carbon prices investors expect companies in these regions to realize additional profits since such companies can sell superfluous allowances at advanced prices. Consequently and in strong contrast to owners of electricity companies, shareholders of steel companies are not hit by the regulatory burden of the system but benefit from the regulation.

For chemical companies, the country-specific analysis also reveals some further interesting details on the heterogeneity in the impact of the carbon price factor on stock returns. While we observe positive and highly significant effects for the regions Northern Europe and Germany&Austria, we do not find any significant effect for companies of the remaining regions. In Northern Europe (Germany&Austria) a 1% carbon price increase is accompanied by a 0.09% (0.06%) stock price increase. Again, these results reflect the heterogeneity in the stringency of the country-specific emissions caps. The results imply that in the identified countries owners of chemical companies benefit from the regulation under the EU-ETS. However, compared to the steel sector the magnitude of these profits tends to be smaller.

Finally, for cement companies, we find positive and significant effects for the regions Germany&Austria and Western Europe which also confirms the findings of the sector-specific analysis. In Western Europe (Germany&Austria) a 1% carbon price increase is accompanied by a 0.06% (0.08%) stock price increase. Again, consistent with the stricter UK emissions cap, owners of UK cement companies do not benefit from increasing carbon prices as implied by a negative but insignificant coefficient. As in the other non-electricity sectors the results point to the existence of regulatory profits of shareholders which tend to be more pronounced in countries with less restrictive emissions caps.

In conclusion, our results extend those of the previous analysis. We find considerable differences in the investors' perception of the relevance of the carbon factor in stock valuation. While in Phase I the link between the carbon and the stock market is rather loose, the link has considerably intensified in Phase II. For Phase I, we find no evidence for the shift of regulatory burden to investors. In Phase II regulatory burden is shifted to owners of electricity companies, while shareholders of other sectors' companies benefit from the regulation. As a consequence, hypothesis *H4* as constructed in Section 4.3.3 has to be rejected.

4.5.4 On the role of asymmetric price effects

Zachmann and von Hirschhausen (2008) find that carbon price changes affect electricity prices asymmetrically. In particular, increasing carbon prices induce stronger electricity price reactions. In contrast, Oberndorfer (2009) does not reveal asymmetric effects of the carbon price on electricity stock returns in Phase I. To analyze whether the impact of increasing and decreasing carbon prices on stock returns differs across the individual sectors in Phase II, we estimate Equation (4.8). The results are summarized in Table 4.8.

Insert Table 4.8 about here.

Basically the results support the findings of the previous analysis. In contrast to Oberndorfer (2009) who does not find asymmetric effects of carbon price changes on electricity stock returns during Phase I, the results of our analysis reveal further details when controlling for positive and negative carbon price changes which lead to the rejection of hypothesis *H5*. Stock returns of electricity companies located in the UK and Southern Europe are affected stronger by increasing prices which is consistent with the view that with increasing carbon prices investors expect that such companies have to buy additional allowances at advanced prices which in turn negatively affects these companies' profitability. Contrary, negative carbon price changes tend to be accompanied by more pronounced stock price responses of non-electricity companies.

4.6 Conclusion and policy discussion

We empirically investigate the impact of changes in the price of European Union emission allowances on the market value of companies covered by the EU-ETS. In particular, the study constitutes a comprehensive policy evaluation of the EU-ETS since it investigates the relevance of price variations in the input factor emission allowances in the stock valuation of companies covered by the system from the view point of investors. The study allows us to capture the role of policy decisions concerning the stringency of emissions caps on the link between the carbon and the stock market. Hence, the study enables us to detect whether shareholders of companies under the EU-ETS carry regulatory burden of the system or even benefit from the environmental regulation. In contrast to previous studies, we make use of a comprehensive data set containing companies of all sectors under the EU-ETS, which we motivate by the heterogeneity in sector-specific net-compliance positions found in the previous literature (see for example Kettner et al. (2008) or Ellerman and Joskow (2008)).

To analyze the effect of the carbon price on company stock returns we adopt the methodology of multifactor panel regression models making use of the estimation strategy of Thompson (2011) in order to obtain consistent standard errors. We explicitly investigate the impact of sector- and country-specific effects on the link between the carbon and the stock market. Even though we have to interpret the estimation results with caution since we only consider companies listed in the EURO STOXX 600, we find that the investors' perception of the system's relevance in stock valuation has considerably changed over the first and the second commitment period.

In Phase I the links between the carbon market and the stock markets are rather weak. For this period the results imply that the carbon price almost exclusively affects the market value of electricity companies in Germany&Austria, Northern Europe, and the UK. There are no or at most very weak effects on the market value of non-electricity companies. Despite strong heterogeneity in sector-specific compliance positions the hypothesis of identical effects of the carbon price on the market value of companies across different non-electricity sectors cannot be rejected. A possible explanation following Zachmann and von Hirschhausen (2008) and Oberndorfer (2009) could be the inexperience of investors during this commitment period. Moreover, this result points to the realization of regulatory profits in the electricity sector in Phase I. Further, the results imply that shareholders of non-electricity companies neither bear regulatory burden, nor benefit from the regulation.

For Phase II the results point to a significant change in the investor's perception of the system's relevance in stock valuation compared to Phase I. In particular, we detect a close link between the carbon market and financial markets on which stocks of the corresponding companies are traded. First, increases in the carbon price are accompanied by decreasing stock prices of electricity companies located in countries with more restrictive emissions caps. Second and in strong contrast, the stock prices of non-electricity companies are positively related to increasing carbon prices. The effects are stronger in countries with more generous emissions caps. Especially, the direction and the magnitude of the links reflect the sector- and country-specific net compliance positions which in turn are determined by policy decisions of national authorities approved by the European Commission. The results of our study are consistent with the view that restrictive allowance allocation to electricity companies leads to the shift of regulatory burden of the system to owners of such companies. On the other hand, our results unambiguously point to the realization of regulatory profits of non-electricity company owners which reflects the generous allocation of allowances to these companies. In contrast to the concern of negative consequences on the competitiveness of internationally active non-electricity companies accompanied by the shift of regulatory burden of the system to such companies' owners as argued by Ellerman et al. (2007), investors rather seem to perceive the EU-ETS as an advantage for these companies. Superfluous allowances that are not needed to fulfill regulatory requirements can be sold on the market which generates additional revenues that shareholders would not have realized without the implementation of the EU-ETS.

Finally, we also control for asymmetric price effects. Generally, our results support the findings of the sector- and country-specific analysis. The results provide weak evidence that sector-specific stock returns asymmetrically respond to positive and negative carbon price changes. Positive price changes induce stronger electricity stock price reactions in countries with more restrictive emissions caps, while negative price changes tend to induce stronger price reactions of non-electricity stocks.

In conclusion, we find that policy decisions of national authorities and the European Commission concerning the stringency of allowance allocation in individual sectors and member states are reflected in the relationship between the carbon price and the market value of companies covered by the EU-ETS.

Figures and tables

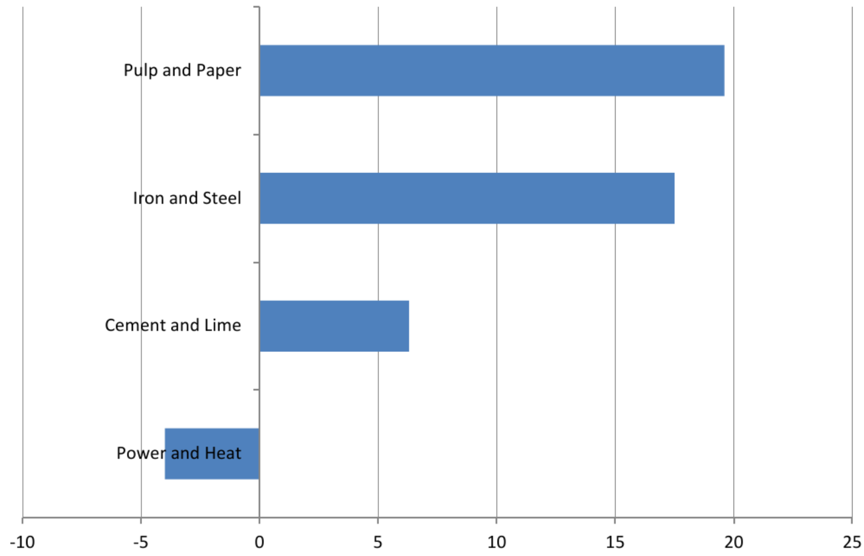

Figure 4.1: Sector-specific compliance positions in Phase I according to Kettner et al. (2008)

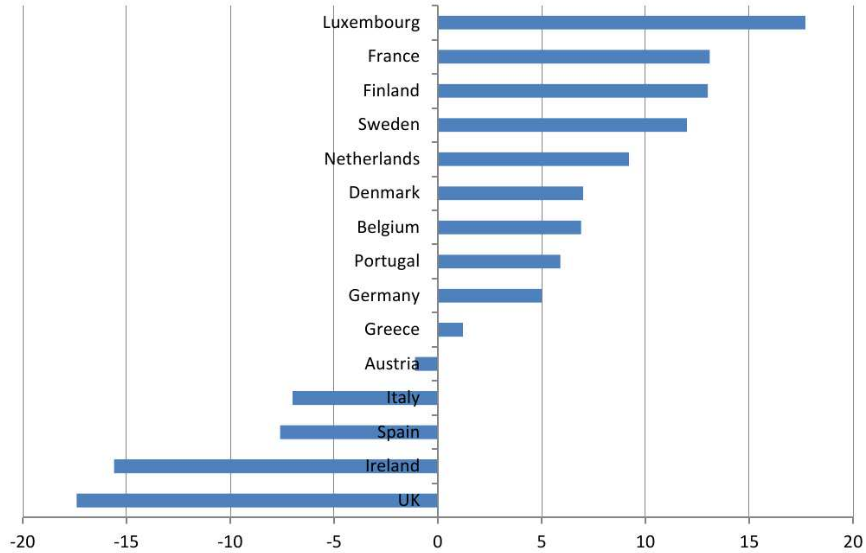

Figure 4.2: Country-specific compliance positions in Phase I according to Kettner et al. (2008)

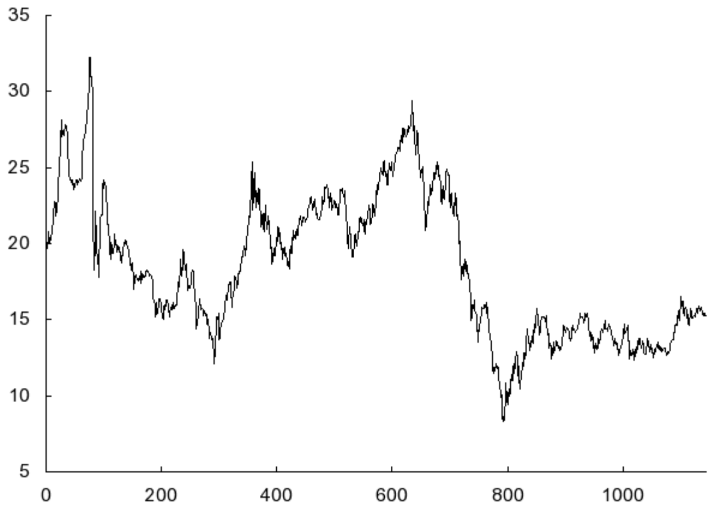

Figure 4.3: Carbon price dynamics from January 2006 to June 2010

Table 4.1: Descriptive statistics

series	mean	sd	kurtosis	skewness	JB-statistic	series	mean	sd	kurtosis	skewness	JB-statistic
r_c	-0.021	2.895	11.245	-0.878	3372.703 [0]	M:EGR	-0.057	2.664	2.999	0.291	16.096 [0]
r_m	-0.021	1.515	6.036	-0.035	437.77 [0]	I:TEN	0.037	2.976	5.586	-0.457	357.112 [0]
r_o	0.018	2.203	4.017	0.162	54.068 [0]	M:UPM	-0.037	2.445	2.546	0.069	10.685 [0.005]
r_g	-0.069	1.651	3.166	0.445	38.835 [0]	O:VAS	0.005	3.371	4.738	-0.16	148.176 [0]
$\Delta TERM$	0.001	0.05	3.428	0.562	68.655 [0]	D:TKA	0.013	2.75	4.26	0.072	76.305 [0]
ΔDEF	0	0.029	31.878	2.9	41174.337 [0]	H:AKZO	0.008	2.216	8.627	0.517	1553.076 [0]
I:A2A	-0.031	1.968	7.877	0.091	1130.279 [0]	D:BAS	0.029	2.092	7.156	-0.102	821.596 [0]
DRXG	-0.039	2.198	2.813	-0.085	3.014 [0.222]	D:BAYN	0.023	1.973	4.178	-0.124	68.797 [0]
EON	-0.024	2.091	8.775	0.061	1583.616 [0]	CRDA	0.053	2.32	3.139	-0.084	2.25 [0.325]
F:EDF	-0.002	2.121	5.573	-0.039	314.471 [0]	JMAT	-0.01	2.435	4.152	-0.038	63.281 [0]
E:ELE	0	1.84	6.08	0.049	450.619 [0]	D:SDF	0.1	3.215	3.442	-0.333	30.31 [0]
I:ENEL	-0.045	1.817	12.555	-0.115	4335.105 [0]	M:KEMR	0.008	2.59	12.503	-0.218	4294.706 [0]
M:FORT	0.012	2.155	5.338	-0.115	261.856 [0]	D:DSMX	0.001	2.217	7.782	-0.177	1090.987 [0]
F:GSZ	-0.005	2.218	12.999	0.691	4835.635 [0]	D:LXS	0.023	2.874	2.744	0.178	9.14 [0.01]
E:IBE	-0.016	2.264	10.792	0.536	2935.793 [0]	D:LIN	0.028	1.942	5.367	0.154	270.347 [0]
NG.	-0.019	1.79	12.355	0.051	4154.281 [0]	F:RHA	-0.04	3.588	3.028	-0.056	0.639 [0.727]
IPR	0.004	2.33	3.875	-0.243	47.506 [0]	B:SOL	-0.025	1.838	4.334	0.143	88.302 [0]
G:PPC	-0.039	2.721	8.242	-0.309	1322.193 [0]	B:UM	0.016	2.906	2.857	-0.207	9.104 [0.011]
E:REE	0.01	1.691	7.3	-0.017	877.402 [0]	VCTA	0.028	2.503	36.922	-2.37	55677.258 [0]
D:RWE	-0.013	1.791	11.092	0.08	3108.688 [0]	N:YARA	0.055	3.356	3.908	-0.396	68.854 [0]
SSE	-0.007	1.791	9.149	-0.11	1796.645 [0]	W:ASSB	0.019	2.525	3.109	0.194	7.725 [0.021]
I:TERN	0.03	1.358	5.745	0.027	357.854 [0]	CRH	-0.024	2.767	2.811	-0.074	2.737 [0.254]
E:ACX	0.004	2.037	3.01	0.095	1.732 [0.421]	DK:FLB	0.066	2.877	2.922	-0.154	4.767 [0.092]
H:MT	0	3.587	4.93	-0.304	194.235 [0]	D:HEI	-0.05	3.057	9.614	0.131	2079.053 [0]
M:OUTO	-0.001	3.195	5.242	-0.289	254.503 [0]	F:LFG	-0.034	2.578	4.643	0.035	128.347 [0]
M:RRUK	-0.047	3.055	2.648	-0.278	20.54 [0]	F:GOB	-0.034	2.901	5.532	0.203	312.049 [0]
D:SZG	0.007	3.432	5.929	-0.309	425.368 [0]	O:WNBA	-0.097	3.181	5.141	-0.483	261.841 [0]
W:SSAA	0.015	3.399	2.917	-0.22	9.513 [0.009]	F:AIR	0.025	1.767	3.962	0.021	44.006 [0]

Notes: The table presents the company covered in the study and the classification according to industry, supersector, sector, and subsector as applied by the EURO STOXX 600. The column Code contains the Data Stream Codes. The last column shows the countries in which the head quarters of the companies are located.

Table 4.2: Company classification

Group	Industries	Supersectors	Sectors	Subsectors
Electricity	Utilities	Utilities	Electricity; Gas, Water and Multiutilities	Conventional Electricity; Multiutilities
Steel	Basic Materials	Basic Resources	Industrial Metals and Mining; Forestry and Paper	Iron and Steel; Paper
Chemicals	Basic Materials	Chemicals	Chemicals	Commodity Chemicals; Special Chemicals
Cement	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures

Notes: The table presents the classification structure according to the EURO STOXX 600. We set up the groups energy, steel, chemicals, and cement. Each EURO STOXX 600 company that belongs to the respective industry, supersector, sector and subsector is assigned to the corresponding group.

Table 4.3: Company classification

Company	Industry	Supersector	Sector	Subsector	Code	Country
A2A	Utilities	Utilities	Electricity	Conventional Electricity	I:A2A	IT
ACERINOX	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	E:ACX	ES
AIR LIQUIDE	Basic Materials	Chemicals	Chemicals	Commodity Chemicals	F:AIR	FR
AKZO NOBEL	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	H:AKZO	NL
ARCELORMITTAL	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	H:MT	LU
ASSA ABLOY	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	W:ASSA	SE
BASF	Basic Materials	Chemicals	Chemicals	Commodity Chemicals	D:BAS	DE
BAYER	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	D:BAYN	DE
CRH	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	CRH	IE
CRODA INTERNATIONAL	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	CRDA	GB
DRAX GRP	Utilities	Utilities	Electricity	Conventional Electricity	DRXG	GB
E.ON	Utilities	Utilities	Gas, Water and Multiutilities	Multiutilities	D.EON	DE
EDF	Utilities	Utilities	Electricity	Conventional Electricity	F:EDF	FR
ENDESA	Utilities	Utilities	Electricity	Conventional Electricity	E:ELE	ES
ENEL	Utilities	Utilities	Electricity	Conventional Electricity	I:ENEL	IT
FLSMIDTH & COMPANY	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	DK:FLB	DK
FORTUM	Utilities	Utilities	Electricity	Conventional Electricity	M:FORT	FI
GDF SUEZ	Utilities	Utilities	Gas, Water and Multiutilities	Multiutilities	F:GSZ	FR
HEIDELBERGCEMENT	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	D:HEI	DE
IBERDROLA	Utilities	Utilities	Electricity	Conventional Electricity	E:IBE	ES
INTERNATIONAL POWER	Utilities	Utilities	Electricity	Conventional Electricity	IPR	GB
JOHNSON MATTHEY	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	JMAT	GB
K + S	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	D:SDF	DE
KEMIRA	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	M:KEMR	FI
KONINKLIJKE DSM	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	D:DSMX	NL
LAFARGE	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	F:LFG	FR
LANXESS	Basic Materials	Chemicals	Chemicals	Commodity Chemicals	D:LXS	DE
LINDE	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	D:LIN	DE
NATIONAL GRID	Utilities	Utilities	Gas, Water and Multiutilities	Multiutilities	NG.	GB
OUTOKUMPU	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	M:OUTO	FI
PUBLIC POWER CORPORATION	Utilities	Utilities	Electricity	Conventional Electricity	G:PPC	GR
RAUTARUUKKI K	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	M:RRUK	FI
RED ELECTRICA CORPORATION	Utilities	Utilities	Electricity	Conventional Electricity	E:REE	ES
RHODIA	Basic Materials	Chemicals	Chemicals	Specialilty Chemicals	F:RHA	FR
RWE	Utilities	Utilities	Gas, Water and Multiutilities	Multiutilities	D:RWE	DE
SAINT GOBAIN	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	F:GOB	FR
SALZGITTER	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	D:SZG	DE
SCOTTISH & SOUTHERN ENERGY	Utilities	Utilities	Electricity	Conventional Electricity	SSE	GB
SOLVAY	Basic Materials	Chemicals	Chemicals	Speciality Chemicals	B:SOL	BE
SSAB A	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	W:SSAA	SE
STORA ENSO R	Basic Materials	Basic Recources	Forestry and Paper	Paper	M:EGR	FI
TENARIS	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	I:TEN	IT
TERNA	Utilities	Utilities	Electricity	Conventional Electricity	I:TERN	IT
THYSSEN KRUPP	Industrials	Industrial Goods and Sevices	General Industrials	Diversified Industrials	D:TKA	DE
UMICORE	Basic Materials	Chemicals	Chemicals	Speciality Chemicals	B:UM	BE
UPM KYMMENE	Basic Materials	Basic Recources	Forestry and Paper	Paper	M:UPM	FI
VICTREX	Basic Materials	Chemicals	Chemicals	Speciality Chemicals	VCTA	GB
VOESTALPINE	Basic Materials	Basic Recources	Industrial Metals and Mining	Iron and Steel	O:VAS	AT
WIENERBERGER	Industrials	Construction and Materials	Construction and Materials	Building Materials and Fixtures	O:WNBA	AT
YARA	Basic Materials	Chemicals	Chemicals	Speciality Chemicals	N:YARA	NO

Notes: The table presents the company covered in the study and the classification according to industry, supersector, sector, and subsector as applied by the EURO STOXX 600. The column Code contains the Data Stream Codes. The last column shows the countries in which the head quarters of the companies are located.

Table 4.4: Energy prices and macroeconomic risk factors

	Full sample	First subsample	Second subsample
market portfolio	$1.0224^{\star\star\star}$ (0.0448)	$0.9956^{\star\star\star}$ (0.0507)	$1.0330^{\star\star\star}$ (0.0470)
oil price	0.0159 (0.0116)	$0.0338^{\star\star}$ (0.0134)	0.0017 (0.0155)
gas price	0.0197^{+} (0.0122)	0.0155 (0.0119)	0.0206 (0.0176)
carbon price	$0.0131^{\star\star}$ (0.0065)	0.0054 (0.0073)	$0.0229^{\star}$ (0.0137)
term spread	0.0100 (0.0220)	-0.0005 (0.0245)	0.0174 (0.0281)
default spread	$-0.0826^{\star\star\star}$ (0.0270)	-0.0418 (0.0328)	$-0.0907^{\star\star\star}$ (0.0306)
constant	0.0217^{+} (0.0144)	$0.0509^{\star\star\star}$ (0.0168)	-0.0118 (0.0257)

Notes: The table presents the results of the panel regression of the baseline specification augmented by macroeconomic risk factors without sector and country-specific effects. Numbers in parentheses are Thompson (2011) robust standard errors. $\star\star\star, \star\star, \star, +$ indicate significance at the 1%, 5%, 10%, and 15% level.

Table 4.5: Sector specific effects

	Full sample	First subsample	Second subsample
market portfolio	0.7515*** (0.0541)	0.6622*** (0.0537)	0.7962*** (0.0628)
market portfolio×steel	0.5004*** (0.0848)	0.6783*** (0.1074)	0.5107*** (0.0921)
market portfolio×chem	0.2442*** (0.0869)	0.3462*** (0.0888)	0.1881** (0.0971)
market portfolio×cement	0.4872*** (0.0953)	0.5247*** (0.0903)	0.4591*** (0.1206)
oil price	0.0091 (0.0167)	0.0050 (0.0166)	0.0061 (0.0242)
oil×steel	0.0435* (0.0347)	0.0974*** (0.0348)	0.0168 (0.0435)
oil×chem	0.0086 (0.0241)	0.0149 (0.0214)	0.0081 (0.0343)
oil×cement	0.0184 (0.0276)	0.0184* (0.0276)	−0.0395 (0.0293)
gas price	0.0165 (0.0184)	0.0056 (0.0178)	0.0222 (0.0258)
gas×steel	0.0171 (0.0351)	0.0332 (0.0401)	0.0064 (0.0460)
gas×chem	0.0036 (0.0244)	0.0166 (0.0253)	−0.0055 (0.0343)
gas×cement	−0.0123 (0.0320)	−0.0191 (0.0341)	−0.0091 (0.0445)
carbon price	0.0068 (0.0101)	0.0248*** (0.0086)	−0.0251 (0.0205)
carbon×steel	0.0165 (0.0183)	−0.0263 (0.0208)	0.0856** (0.0355)
carbon×chem	0.0093 (0.0139)	−0.0214* (0.0125)	0.0607** (0.0280)
carbon×cement	−0.0021 (0.0225)	−0.0481** (0.0204)	0.0699* (0.0382)
term spread	0.0100 (0.0220)	−0.0005 (0.0246)	0.0174 (0.0282)
default spread	−0.0826*** (0.0270)	−0.0418 (0.0330)	−0.0907*** (0.0306)
steel	0.0176 (0.0382)	0.0248 (0.0456)	0.0070 (0.0656)
chem	0.0360 (0.0282)	−0.0132 (0.0344)	0.0895* (0.0503)
cement	0.0000 (0.0392)	−0.0216 (0.0431)	0.0246 (0.0702)
constant	0.0063 (0.0195)	0.0527** (0.0218)	−0.0455 (0.0331)

Notes: The table presents the results of the panel regression of the baseline specification augmented by macroeconomic risk factors and sector-specific effects but no country-specific effects. Numbers in parentheses are Thompson (2011) robust standard errors. ⋆⋆⋆, ⋆⋆, ⋆, + indicate significance at the 1%, 5%, 10%, and 15% level.

Table 4.6: Sector specific effects II

	Electr.	Steel	Chemicals	Cement
		Full sample		
market	0.751⋆⋆⋆	1.318⋆⋆⋆	0.996⋆⋆⋆	1.239⋆⋆⋆
	[192.79]	[390.38]	[231.01]	[311.18]
oil	0.009	0.053⋆	0.018	−0.030+
	[0.30]	[3.28]	[1.38]	[2.51]
gas	0.017	0.034	0.020	0.004
	[0.80]	[1.39]	[1.30]	[0.04]
carbon	0.007	0.023⋆	0.016⋆	0.005
	[0.45]	[2.92]	[3.73]	[0.08]
		First subsample		
market	0.662⋆⋆⋆	1.340⋆⋆⋆	1.008⋆⋆⋆	1.187⋆⋆⋆
	[152.18]	[234.94]	[185.73]	[375.03]
oil	0.005	0.102⋆⋆⋆	0.020	0.023
	[0.09]	[10.67]	[1.76]	[1.30]
gas	0.006	0.039	0.022	−0.014
	[0.10]	[1.51]	[1.26]	[0.35]
carbon	0.025⋆⋆⋆	−0.002	0.003	−0.023
	[8.23]	[0.01]	[0.15]	[0.18]
		Second subsample		
market	0.796⋆⋆⋆	1.307⋆⋆⋆	0.984⋆⋆⋆	1.125⋆⋆⋆
	[160.51]	[321.24]	[198.55]	[191.01]
oil	0.006	0.023	0.014	−0.070⋆⋆⋆
	[0.01]	[0.43]	[0.46]	[7.47]
gas	0.022	0.029	0.017	0.013
	[0.74]	[0.54]	[0.45]	[0.21]
carbon	−0.025	0.061⋆⋆	0.036⋆⋆	0.045+
	[1.50]	[5.82]	[3.63]	[2.56]

Notes: The table presents the results of the panel regression of the baseline specification augmented by macroeconomic risk factors and sector-specific effects but no country-specific effects. Numbers in brackets are F-statistics. ⋆⋆⋆, ⋆⋆, ⋆, + indicate significance at the 1%, 5%, 10%, and 15% level.

Table 4.7: Sector and country specific effects

	Electr.	Steel	Chemicals	Cement
First subsample				
Northern Europe	0.085*** [16.57]	−0.002 [0.01]	0.003 [0.03]	−0.036 [0.38]
Southern Europe	0.007 [0.61]	−0.008 [0.25]	–	–
Germany&Austria	0.035* [3.50]	−0.013 [0.17]	−0.008 [0.29]	−0.012 [0.33]
Western Europe	0.015 [1.08]	0.045^{+} [2.47]	0.014^{+} [2.43]	−0.021^{+} [2.43]
United Kingdom	0.041*** [15.21]	–	0.000 [0.00]	−0.024 [1.63]
Second subsample				
	Electr.	Steel	Chemicals	Cement
Northern Europe	0.061* [2.96]	0.023 [0.55]	0.090*** [7.02]	0.037 [0.03]
Southern Europe	−0.056*** [7.84]	0.024 [0.59]	–	–
Germany&Austria	0.038 [1.46]	0.092** [6.00]	0.056** [4.93]	0.077* [2.72]
Western Europe	0.034 [1.21]	0.225*** [16.48]	0.015 [0.42]	0.065* [3.73]
United Kingdom	−0.054** [4.12]	–	0.008 [0.11]	−0.045 [0.96]

Notes: The table presents the results of the panel regression of the baseline specification augmented by macroeconomic risk factors, sector-specific effects, and country-specific effects. Numbers in brackets are F-statistics. $\star\star\star, \star\star, \star, +$ indicate significance at the 1%, 5%, 10%, and 15% level.

Table 4.8: Asymmetric effects

	Electr.	Steel	Chemicals	Cement
Positive carbon price changes				
Northern Europe	0.013 [0.02]	−0.059 [0.47]	0.057 [0.41]	−0.061 [0.56]
Southern Europe	−0.124* [2.87]	−0.043 [0.27]	–	–
Germany&Austria	−0.014 [0.03]	0.045 [0.24]	0.007 [0.01]	−0.112 [1.14]
Western Europe	−0.040 [0.23]	0.190^{+} [2.66]	−0.051 [0.44]	−0.023 [0.07]
United Kingdom	−0.142* [3.04]	–	−0.048 [0.36]	−0.125 [1.54]
Negative carbon price changes				
	Electr.	Steel	Chemicals	Cement
Northern Europe	0.063 [1.25]	0.056 [1.52]	0.080^{+} [2.29]	0.085*** [3.25]
Southern Europe	−0.035 [1.05]	0.044 [0.63]	–	–
Germany&Austria	0.044 [0.67]	0.094^{+} [2.08]	0.059 [1.89]	0.205*** [8.77]
Western Europe	0.044 [1.23]	0.216** [5.61]	0.033 [0.42]	0.103** [4.17]
United Kingdom	−0.016 [0.17]	–	0.018 [0.16]	−0.014 [0.04]

Notes: The table presents the results of the panel regression of the baseline specification augmented by macroeconomic risk factors, sector-specific effects, and country-specific effects. Numbers in brackets are F-statistics. $\star\star\star, \star\star, \star, +$ indicate significance at the 1%, 5%, 10%, and 15% level.

Chapter 5

Long-term oil-stock correlations

5.1 Introduction

Given the empirical evidence in, e.g., Hamilton (1983, 1985, 2003) on the negative impact of oil price shocks on economic activity, it does not seem surprising that studies such as Jones and Kaul (1996) also find a negative relationship between oil prices and stock returns. In this article, we revisit the oil-stock market relationship by analyzing the dynamic correlations between crude oil prices and U.S. stock market returns during the period 1993–2011. The rolling window of yearly realized correlations in Figure 5.1 clearly reveals that there is considerable time-variation in the correlation between the two return series with extended periods of positive correlations. Using a two-component dynamic correlation model, we aim at explaining these variations by changes in the U.S. macroeconomic environment. Our specification allows us to separate day-to-day fluctuations (the dashed line in Figure 5.1) from gradual long-term movements (the bold line) which are related to the stance of the economy. The dynamic correlations plotted in Figure 5.1 are obtained from a specification which explains the long-term component by variations in the Chicago Fed national activity index (NAI). Figure 5.1 clearly shows the close link between the oil-stock correlation and the business cycle. In particular, note the positive oil-stock correlation during recessions and thereafter.

Figure 5.1 about here

Our econometric specification is based on the Dynamic Conditional Correlation - MIxed Data Sampling (DCC-MIDAS) model proposed in Colacito et al. (2011). The DCC-MIDAS model combines the Engle (2002) DCC specification with the GARCH-MIDAS framework of Engle et al. (2009). The GARCH-MIDAS framework extends the simple GARCH specification by modeling volatility as consisting of a short-term and a long-term component. Most importantly, the long-term component is specified as a function of the macroeconomic environment. In the original DCC specification with correlation targeting each quasi-correlation follows a 'GARCH type' process which is mean-reverting to the unconditional correlation of the volatility-adjusted residuals. The basic idea of Colacito et al. (2011) is to replace this unconditional correlation with a slowly time-varying long-term component which is driven by lagged realized correlations. The quasi-correlation then fluctuates around this long-run trend. Hence, the new specification can be considered as a two component model for the dynamic correlations. In the spirit of Engle et al. (2009) the short-term component fluctuates at a daily frequency while the long-term component adjusts at the lower monthly frequency. Colacito et al. (2011) assume that the long-term component can be expressed as a weighted sum of the lagged monthly realized correlations between the volatility-adjusted residuals.

Using the GARCH-MIDAS framework, we first analyze whether the long-term oil market volatility is related to the U.S. macroeconomy and whether oil and stock volatility respond to the same macroeconomic information. Next, we extend the DCC-MIDAS model by directly incorporating information on the macroeconomic development in the long-term correlation component, i.e. we replace the realized correlations by monthly macroeconomic variables. Since the macroeconomic variables – unlike the realized correlations – are not restricted to the minus one to plus one interval, we suggest a new specification for the long-term component. Similar to Christodouklakis and Satchell (2002), we assume that the Fisher-z transformation of the long-term component can be written as a linear function of the weighted lagged macroeconomic variables. The weights are again determined using the MIDAS approach. We refer to this new specification which includes a macroeconomic explanatory variable as the DCC-MIDAS-X model.

In broad terms, our results can be summarized as follows. First, we find that the movements in long-term oil market volatility can be well predicted by various measures of U.S. macroeconomic activity. Our empirical results provide convincing evidence for a counter cyclical relationship between measures which either describe the current stance of the economy, e.g. industrial production, or provide forward looking information about the future state of the economy, e.g. the leading index for the U.S., and oil market volatility. Current and expected increases (decreases) in economic activity clearly anticipate downswings (upswings) in long-term oil volatility. This result strengthens the argument of Barsky and Kilian (2004) and Harris et al. (2009) that the oil price development is very much synchronized with the business cycle and that there is indeed reverse causality from macroeconomic variables to the oil price. Interestingly, we also find that long-term oil and stock market volatility are determined by the same macroeconomic factors, while

Kilian and Vega (2011) report that oil price returns in contrast to stock returns do not respond instantaneously to macroeconomic news.

Second, our empirical results show that changes in the long-term oil-stock correlation can be anticipated by the same macroeconomic factors which also affect the long-term volatilities. We provide strong evidence for a counter cyclical behavior of the long-term oil-stock correlation. The economic rationale behind is best explained by again looking at Figure 5.1 which exemplarily relates the oil-stock correlation to changes in the NAI. The phases with positive long-term oil-stock correlations correspond to values of the NAI which either indicate recessions or the beginning of expansions with growth still below or at trend. On the other hand, a negative long-run correlation emerges when the NAI signals strong growth above trend. Clearly, the positive correlation during recessions is driven by the simultaneous drop in oil and stock prices. The economic recovery during the early phase of an expansion then leads to increasing oil prices due to higher demand as well as to rising stock prices because of the improved outlook for corporate cash flows. The combination of these two effects causes the long-run oil-stock correlation to remain positive. This interpretation squares with the findings in Kilian and Park (2009) regarding the positive short-run effect of an unexpected increase in global demand on oil and stock prices. Finally, during boom phases with strong growth above trend both the further increasing oil prices as well as the expectation of rising interest rates have a depressing effect on the stock market. Hence, for these periods our model predicts a negative long-term correlation.

Third, the long-term correlation component can be interpreted as the predicted or expected correlation given a certain state of the economy. Since the macroeconomic variables which drive the long-term component represent aggregate demand, the deviations of the short-term from the long-term component should be driven by other factors related to the stock and/or the oil market. Typical examples would be either oil specific, i.e. precautionary, demand shocks or supply shocks. The fact that various measures of macroeconomic activity lead to a convincing and coherent fit of the long-term correlation suggests that aggregate demand is the most important factor for the oil-stock relationship. Our results can thus be understood as further empirical evidence for Kilian's (2009, p.1068) claim that "models of endogenous oil prices should focus on the aggregate demand side of the oil market".

Fourth, the fact that the sign of the oil-stock correlation critically depends on the state of the economy reinforces the argument by Kilian and Park (2009) that simple regressions of stock returns on oil price changes can be very misleading. This point may well explain the conflicting empirical evidence on the oil-stock relationship in Jones and Kaul (1996), Wei (2003) and others.

Fifth, as shown in Colacito et al. (2011) the explicit modeling of the long-term correlation component can be very beneficial when it comes to portfolio choice, hedging decisions or risk management. In the oil-stock context we expect the potential efficiency gains to be highly relevant, since the time-varying correlations are relatively large and – in contrast to backward looking models – the DCC-MIDAS-X specification allows us to anticipate changes in correlations.

Finally, several remarks are in order. The DCC-MIDAS-X specification remains a reduced form model. Hence, while we find that measures of economic activity are helpful predictors for the long-term oil-stock correlation, our estimates do not necessarily have a causal interpretation. Further, the model does not explicitly distinguish between different types of shocks to oil prices as in, e.g., Kilian and Park (2009) or Kilian (2009). However, we can interpret our long-term correlation component as the correlation that would be prevalent if the aggregate demand side dominates. Oil specific shocks due to precautionary demand or supply shocks are rather reflected by the short-term component and can be considered as a reason why the short-term component can deviate from the long-run trend. The behavior of the short-term component during the invasion of Kuwait in August 1990 and the second Gulf War in 2003 are in line with this interpretation. Finally, we focus on economic activity measures for the U.S. only, while the oil price is driven by global demand. Nevertheless, we believe that our U.S. activity measures are likely to be highly correlated with global demand for most of the time.

The remainder of the article is organized as follows. Section 5.2 reviews the related literature while Section 5.3 discusses the GARCH-MIDAS and DCC-MIDAS models. The data and empirical results are presented in Sections 5.4 and 5.5. Section 5.6 provides some robustness analysis and Section 5.7 concludes the article.

5.2 Related literature

Our analysis is based on two strands of literature. The first one is concerned with the modeling of long-term movements in volatilities and correlations, the second one with the relationship between oil, the macroeconomy and stock prices.

The idea of having short- and long-term component models of volatilities dates back to Ding and Granger (1996) and Engle and Lee (1999). In their specifications, both components simply follow 'GARCH-type' processes but with different degrees of persistence. Similarly, Davidson (2004) proposed the HYGARCH specification which can be considered as a two component model with the short-term component being a GARCH process while the long-term component follows a FIGARCH process (see also Conrad, 2010). While these specifications allow to separate the two volatility components, both components are assumed to be driven by the same shocks. In addition, the unconditional variance is still constant over time. Engle and Rangel (2008) and Engle et al. (2009) relax this assumption and propose specifications in which the long-term component can be considered a time-varying unconditional variance. While in the Engle and Rangel (2008) Spline-GARCH model both components fluctuate at the same frequency, Engle et al. (2009) assume that the long-term component evolves at a lower frequency than the short-term component. Using the MIDAS framework of Ghysels et al. (2005, 2007), they directly relate the long-term component to the evolution of macroeconomic time series such as industrial production or inflation. In line with the earlier findings in Schwert (1989), the GARCH-MIDAS model provides strong evidence for a counter cyclical behavior of financial volatility. Recently, Conrad and Loch (2011) extend the analysis of Engle et al. (2009) by using a broader set of macroeconomic variables including leading indicators and expectations data from the Survey of Professional Forecasters. The DCC-MIDAS model proposed in Colacito et al. (2011) simply extends the two component idea from volatilities to correlations. However, instead of relating the long-term correlation directly to its potential macroeconomic sources, Colacito et al. (2011) only consider lagged realized correlations as explanatory variables.

Since the seminal articles of Hamilton (1983, 1985, 2003) exogenous oil supply shocks were suspected to be causal for recessions and periods of low economic growth. Based on this presumption, several empirical studies have analyzed the relationship between oil prices and stock market returns. While Jones and Kaul (1996) or Nandha and Faff (2008) indeed find that oil price increases negatively affect stock prices, Huang et al. (1996) or Wei (2003) cannot establish a significant relationship. Recently, Miller and Ratti (2009) provide evidence for a time-varying relationship. For the period after 1999 they even report a positive connection. Hence, the empirical evidence is far from being uncontroversial. Kilian and Park (2009) provide two explanations for the conflicting results. First, although the oil price is often assumed to be exogenous with respect to the U.S. economy, there may be reverse causality at work (see also Barsky and Kilian, 2004). Similarly, Harris et al. (2009) argue that – in contrast to the 1970s when supply shocks were likely to be predominant – oil prices are mainly driven by high global aggregate demand since the mid-1990s. Thus, stock and oil price changes may be induced by the same macroeconomic factors and, hence, regressions of stock returns on oil price changes may be misleading due to endogeneity. The empirical results in Ewing and Thompson (2007) confirm the procyclical behavior of oil prices and specifically indicate that crude oil prices lag industrial production. Second, Kilian and Park (2009) argue that the sign of the effect of an oil price increase on the stock market depends on the type of the underlying shock and, hence, may change over time. While shocks due to an unanticipated economic expansion may have a positive impact, shocks related to precautionary demand are likely to have a negative impact. For several oil-importing and oil-exporting countries Filis et al. (2011) show that the oil-stock correlation is indeed time-varying. Although they informally relate phases of positive or negative correlations to demand and supply shocks, their simple DCC-GARCH model does not explicitly incorporate information on the state of the economy. In particular, their model does not allow to forecast changes in correlations in response to changes in the macro environment.

5.3 The DCC-MIDAS model

In this section, we develop the econometric framework to analyze the impact of macroeconomic variables on long-term volatility and correlations. We consider the bivariate vector of asset returns $\mathbf{r}_t = (r_{1,t}, r_{2,t})'$, where $r_{1,t}$ refers to the stock and $r_{2,t}$ to the oil returns, and denote by $\mathcal{F}_{t-1} = \sigma(\mathbf{r}_{t-1}, \mathbf{r}_{t-2}, \ldots)$ the σ-field generated by the information available through time $t-1$. Let $\mathbf{E}[\mathbf{r}_t|\mathcal{F}_{t-1}] = \boldsymbol{\mu}_t = (\mu_{1,t}, \mu_{2,t})'$ and define the vector of residuals $\mathbf{r}_t - \boldsymbol{\mu}_t = \boldsymbol{\varepsilon}_t = (\varepsilon_{1,t}, \varepsilon_{2,t})'$. We assume that conditional on $\mathcal{F}_{t-1}$ the residuals are normally distributed with $\mathbf{Var}[\boldsymbol{\varepsilon}_t|\mathcal{F}_{t-1}] = \mathbf{H}_t$, i.e. $\boldsymbol{\varepsilon}_t|\mathcal{F}_{t-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{H}_t)$. Following Engle (2002), we decompose the conditional covariance matrix into $\mathbf{H}_t = \mathbf{D}_t\mathbf{R}_t\mathbf{D}_t$ where

$$\mathbf{R}_t = \begin{pmatrix} 1 & \rho_{12,t} \\ \rho_{12,t} & 1 \end{pmatrix} \qquad \text{and} \qquad \mathbf{D}_t = \begin{pmatrix} h_{1,t}^{1/2} & 0 \\ 0 & h_{2,t}^{1/2} \end{pmatrix}. \tag{5.1}$$

Finally, we define the standardized residuals $\boldsymbol{\eta}_t = (\eta_{1,t}, \eta_{2,t})'$ as $\boldsymbol{\eta}_t = \mathbf{D}_t^{-1}\boldsymbol{\varepsilon}_t$. Note that $\mathbf{Var}[\boldsymbol{\eta}_t|\mathcal{F}_{t-1}] = \mathbf{R}_t$. The DCC framework allows us to separately model the conditional variances and the conditional correlations.

5.3.1 Conditional variances

To capture the impact of macroeconomic variables on return volatility, we adopt the GARCH-MIDAS framework of Engle et al. (2009). We assume a multiplicative component model for each conditional variance, i.e. we specify $h_{i,t} = g_{i,t}m_{i,\tau}$, where $g_{i,t}$ is the short-run and $m_{i,\tau}$ the long-run component. While the transitory volatility component changes at the daily frequency t, the long-run component changes at the monthly frequency τ only. We denote by $N^{(\tau)}$ the number of days within month τ. Specifically, we assume that the short-run volatility component follows a mean-reverting unit GARCH(1,1) process

$$g_{i,t} = (1 - \alpha_i - \beta_i) + \alpha_i \frac{(r_{i,t-1} - \mu_{i,t-1})^2}{m_{i,\tau}} + \beta_i g_{i,t-1}, \tag{5.2}$$

with $\alpha_i > 0$, $\beta_i \geq 0$, and $\alpha_i + \beta_i < 1$. The long-term component is modeled as a slowly varying function of exogenous variables X_τ using the MIDAS specification

$$\log(m_{i,\tau}) = m_i + \theta_i \sum_{k=1}^{K_v} \varphi_k(\omega_i) X_{\tau-k}, \tag{5.3}$$

where the log transformation guarantees the non-negativity of the conditional variances when the exogenous variables can take negative values. The X_τ will be monthly macroeconomic variables. For the weighting scheme, we follow Engle et al. (2009) and adopt a restricted beta weighting scheme where the weights are computed according to

$$\varphi_k(\omega_i) = \frac{(1 - k/K_v)^{\omega_i - 1}}{\sum_{l=1}^{K_v}(1 - l/K_v)^{\omega_i - 1}}, \qquad k = 1, ..., K_v. \tag{5.4}$$

For all $\omega_i > 1$, the weighting scheme guarantees a decaying pattern, where the rate of decay is determined by ω_i. Large (small) values of ω_i generate a rapidly (slowly) decaying pattern. Given a maximum lag order K_v, the weighting scheme entails a data driven lag-length selection, depending on the scale of ω_i. K_v itself can be determined by the Akaike information criterion (AIC).

In the following, we will refer to the component model with explanatory variables as GARCH-MIDAS-X. Finally, note that when $\theta_i = 0$ the long-run component is simply a constant and $h_{i,t}$ follows a GARCH(1,1) process with unconditional variance $\sigma_i^2 = \exp(m_i)$.

5.3.2 Conditional correlations

The DCC-MIDAS specification proposed by Colacito et al. (2011) provides a natural extension of the GARCH-MIDAS model to dynamic correlations. We first decompose the conditional correlation matrix as $\mathbf{R}_t = diag\{\mathbf{Q}_t\}^{-1/2}\mathbf{Q}_t diag\{\mathbf{Q}_t\}^{-1/2}$, with $\mathbf{Q}_t = [q_{ij,t}]_{i,j=1,2}$, and specify the quasi-correlations as

$$\mathbf{Q}_t = (1 - a - b)\bar{\mathbf{R}}_t + a\boldsymbol{\eta}_{t-1}\boldsymbol{\eta}'_{t-1} + b\mathbf{Q}_{t-1}, \tag{5.5}$$

with $a > 0$, $b \geq 0$, and $a + b < 1$. In the Engle (2002) DCC model with correlation targeting the matrix $\bar{\mathbf{R}}_t$ does not depend on time and equals the empirical correlation matrix of $\boldsymbol{\eta}_t$, i.e. has ones on the main diagonal while the off-diagonal elements are $\bar{\rho}_{12} = T^{-1}\sum_{t=1}^{T}\eta_{1,t}\eta_{2,t}$. In contrast, in the DCC-MIDAS framework introduced in Colacito et al. (2011) the off-diagonal elements are the long-term correlations $\bar{\rho}_{12,\tau}$. As in the GARCH-MIDAS equation the long-term correlation component does not vary at the daily frequency t but at the lower frequency τ. That is, the short-run quasi-correlations fluctuate around the time-varying long-run correlations:

$$q_{12,t} = \bar{\rho}_{12,\tau} + a(\eta_{1,t-1}\eta_{2,t-1} - \bar{\rho}_{12,\tau}) + b(q_{12,t-1} - \bar{\rho}_{12,\tau}). \tag{5.6}$$

Colacito et al. (2011) assume that $\bar{\rho}_{12,\tau}$ can be expressed as a weighted average of the K_c past realized correlations RC_τ:

$$\bar{\rho}_{12,\tau} = \sum_{k=1}^{K_c} \varphi_k(\omega_{12}) RC_{\tau-k}, \tag{5.7}$$

with

$$RC_\tau = \frac{\sum_{t=N_{\tau-1}+1}^{N_\tau} \eta_{1,t}\eta_{2,t}}{\sum_{t=N_{\tau-1}+1}^{N_\tau} \eta_{1,t}^2 \sum_{t=N_{\tau-1}+1}^{N_\tau} \eta_{2,t}^2}, \tag{5.8}$$

where $N_\tau = \sum_{i=1}^{\tau} N^{(i)}$ and $N_0 = 0$. The weights are again given by equation (5.4) with ω_i and K_v replaced by ω_{12} and K_c, respectively. Since the weights $\varphi_k(\omega_{12})$ sum up to one and the RC_τ are correlations, the long-run correlation will itself lie within the $[-1, +1]$ interval.

We extend the DCC-MIDAS model by directly incorporating information on the macroeconomic development in the long-run component. Similarly as in the GARCH MIDAS setting – where the specification for $m_{i,\tau}$ has to ensure the non-negativity of the long-term volatility – our specification has to ensure that the long-run correlation lies within the $[-1, +1]$ interval although the macroeconomic explanatory variables do not. We follow Christodoulakis and Satchell (2002) and use the Fisher-z transformation of the correlation coefficient, i.e. we assume that

$$\bar{\rho}_{12,\tau} = \frac{\exp(2z_{12,\tau}) - 1}{\exp(2z_{12,\tau}) + 1}, \tag{5.9}$$

with

$$z_{12,\tau} = m_{12} + \theta_{12}\sum_{k=1}^{K_c} \varphi_k(\omega_{12}) X_{\tau-k}, \tag{5.10}$$

where X_τ denotes either a macroeconomic explanatory variable or the realized correlations. Note that in our non-linear specification, from θ we can only infer the sign but not directly the marginal effect of a macroeconomic variable on the long-term correlation. Finally, in the DCC-MIDAS model - as in the standard DCC model - the short-run correlations are obtained by rescaling, i.e. $\rho_{12,t} = q_{12,t}/\sqrt{q_{11,t}q_{22,t}}$. In the subsequent analysis we refer to the specifications with either macroeconomic explanatory variables or the realized correlations as DCC-MIDAS-X or DCC-MIDAS-RC models, respectively.

5.3.3 Estimation

Following Engle (2002) and Colacito et al. (2011) the model parameters can be estimated using a two-step procedure. This is feasible because the log likelihood function to be maximized

$$\mathcal{L} = -\sum_{t=1}^{T}\left(2\log(2\pi) + 2\log(|\mathbf{D}_t|) + \boldsymbol{\varepsilon}_t'\mathbf{D}_t^{-2}\boldsymbol{\varepsilon}_t\right) - \sum_{t=1}^{T}\left(\log(|\mathbf{R}_t|) + \boldsymbol{\eta}_t'\mathbf{R}_t^{-1}\boldsymbol{\eta}_t - \boldsymbol{\eta}_t'\boldsymbol{\eta}_t\right) \tag{5.11}$$

can be separated into two parts. The first sum in equation (5.11) contains the data and the variance parameters while the second sum depends on the volatility-adjusted residuals and the correlation parameters. Hence, in the first step we estimate the GARCH-MIDAS parameters individually for each return series and use the estimated volatility-adjusted residuals in the second step to obtain the correlation parameters.

5.4 Data

Since we apply the MIDAS approach, our data consists of observations at the daily as well as the monthly frequency. We combine daily stock market and crude oil price data with monthly observations on the macroeconomic variables. While the stock series was obtained from the Kenneth R. French data library, the oil prices and the macroeconomic data are taken from the FRED database at the Federal Reserve Bank of St. Louis. Our data covers the period from January 1993 to November 2011.

5.4.1 Oil and stock market data

For the stock series, we employ the daily returns on the CRSP value-weighted portfolio, which is based on all NYSE, AMEX and NASDAQ stocks and can be considered the best available proxy for 'the stock market'. In addition, the CRSP data facilitates comparison of our results with those of Engle et al. (2009) and Conrad and Loch (2011). As in Kilian and Vega (2011), the oil price returns are constructed from the West Texas Intermediate (WTI) crude oil spot price. Panel A of Table 5.1 provides summary statistics for the two return series. While the sample mean of the returns is positive for both markets, the table provides first evidence for stronger price fluctuations in the oil than in the stock market. The annualized unconditional standard deviation of the oil price returns is 39.18% and, hence, considerably higher than the 19.41% of the CRSP returns. Finally, the unconditional correlation between oil and stock returns is 0.15.

Table 5.1 about here

5.4.2 Macroeconomic data

We divide the monthly macroeconomic data into three categories. Those which measure the current stance of the economy, forward looking indicators and measures of inflation. The first category contains the following variables: industrial production (IP), nonfarm payrolls (NFP), and the unemployment rate (UR). The forward looking indicators are the national activity index (NAI)[1] and the leading index (LI)[2] for the U.S. They are supposed to reflect the role of market participants' expectations concerning the future economic development. The final category consists of the producer price index (PPI) and the consumer price index (CPI) and captures inflation dynamics.

For the variables IP, PPI, and CPI we compute month-to-month growth rates according to $100 \cdot [\ln(X_\tau) - \ln(X_{\tau-1})]$, while in case of UR and NFP we use month-to-month changes. Finally, NAI and LI are included in levels. Panel B of Table 5.1 provides the summary statistics for the macroeconomic data. Figure 5.2 shows the dynamics of the macroeconomic variables for the period from January 1993 to November 2011.

Figure 5.2 about here

[1]The NAI is standardized weighted average of 85 monthly indicators of national economic activity including figures that represent (i) production and income, (ii) employment, unemployment, and hours, (iii) personal consumption and housing, and (iv) sales, orders and inventories. The NAI is computed and published by the Federal Reserve Bank of Chicago. Positive realizations indicate growth above trend, while negative realizations indicate growth below trend. The variables IP, NFP, and UR are among the indicators used for the computation of the NAI.

[2]The LI predicts the six-month growth rate of the US coincident index based on variables that lead the economy including housing permits, unemployment insurance claims, delivery times from the ISM manufacturing survey, and the term spread. The LI is published by the Federal Reserve Bank of Philadelphia.

5.5 Empirical results

We first present the estimation results for the GARCH-MIDAS models which relate the long-term volatilities to the macroeconomic environment. Thereafter, the DCC-MIDAS specifications which focus on the long-run correlations are discussed.

5.5.1 Determinants of long-term volatilities

Tables 5.2 and 5.3 present the estimates for the various stock and oil GARCH-MIDAS models. In addition to the models which include the macroeconomic explanatory variables, we also consider the restricted version of equation (5.3) with $\theta_i = 0$. Recall, that in this benchmark specification the GARCH-MIDAS model reduces to a GARCH(1,1) with constant unconditional variance. Since this model is nested within the class of GARCH-MIDAS models, we can use likelihood-ratio tests (LRT) and the AIC to compare the fit of the models which are augmented by macroeconomic variables with the benchmark specification. Further, since the serial correlation in daily stock and oil returns is negligible, we choose $\mu_{i,t} = \mu_i$ in both conditional means. Based on the AIC we choose $K_v = 36$ for both markets, i.e. our specifications cover three MIDAS lag years. However, all results are robust to moderate changes in K_v. The constant μ_i is significant in all stock return models, but insignificant in the oil return specifications. In all cases the estimated α_i and β_i parameters are highly significant. Interestingly, while the α_i (β_i) parameters are estimated to be slightly higher (lower) in the stock than in the oil market, the sum $\alpha_i + \beta_i$ is almost identical in both markets and always less than one. That is, in all specifications the short-run volatility component is mean-reverting to the long-run volatility trend. Next, we discuss the estimated θ_i and ω_i parameters individually for the two markets.

Tables 5.2 and 5.3 about here

Table 5.2 shows that each variable in the two categories current stance of the economy and future economic outlook has a significant effect on long-term stock market volatility. For IP, NFP, NAI, and LI the estimated coefficient $\hat{\theta}_1$ is negative and highly significant, while it is positive and highly significant in case of UR. Since the sign of θ_1 measures whether an increase of the respective variable leads to an upswing or downswing in the long-run volatility, the estimates imply that higher (lower) levels of economic activity lead to a reduction (rise) in long-term stock market volatility. In stark contrast, both inflation measures do not significantly affect long-term stock market volatility. The LRT which compare the GARCH-MIDAS-X models with the restricted benchmark specification imply that we can reject the hypothesis that $\theta_1 = 0$ for all specifications with significant macroeconomic variables. This result is also confirmed by the AIC. Finally, the loglikelihood function, the LRT, and the AIC unambiguously identify the model including the unemployment rate as the one with the best fit.

Our results are consistent with the findings in Engle et al. (2009) and Conrad and Loch (2011). Engle et al. (2009) consider industrial production and producer price inflation as explanatory variables and report that industrial production strongly influences long-term U.S. stock market volatility. In line with our results, they find significant effects of inflation when it was high and volatile in the 1970s, but insignificant ones during the post-1985 period of the Great Moderation.[3] Conrad and Loch (2011) consider various other measures of economic activity including several leading indicators and find that variables which can predict the future state of the economy have explanatory power for long-run volatilities. This squares with our highly significant θ_1 estimates for LI and NAI. These variables are likely candidates to affect uncertainty concerning future cash flows and risk premia. In summary, our findings deliver further support for the view that long-term stock market volatility behaves counter cyclical.

In Table 5.3 we now turn to the analysis of the macroeconomic determinants of long-term oil market volatility. As in case of the stock market, the estimates for θ_2 suggest that long-term oil price volatility is closely linked to each of the macroeconomic variables describing the current stance of the economy as well as the future economic outlook. In particular, the results imply that downturns in U.S. economic activity, i.e. decreases in IP, NFP, NAI, and LI and increases in UR lead to higher levels of long-term oil market volatility. The empirical evidence implies that changes in variables which measure economic activity do precede changes in long-term oil market volatility. Although this result does not necessarily invalidate the assumption that "oil price changes cannot be predicted from earlier movements in macro variables" (see Hamilton, 2008), it challenges the view that oil price movements are exogenous with respect to the U.S. economy. We will return to this issue in the next subsection. The fact that measures of economic activity help to anticipate changes in oil price volatility also supports the argument of Barsky and Kilian (2004), Kilian (2009), and Harris et al. (2009) that oil prices are mainly driven by aggregate demand and to a much lesser extend by oil supply shocks. Hence, an economic downturn can be viewed as a negative aggregate demand shock which increases long-term oil price volatility.[4] Similarly as for the stock returns, neither PPI nor CPI significantly affect oil price uncertainty. This, in turn, is consistent with the argument in Harris et

[3] In addition to the levels, Engle et al. (2009) also investigate whether the uncertainties about IP and PPI affect stock volatility.

[4] The finding is analogous to the leverage effect in the stock market. A positive (negative) demand shock leads to increasing (decreasing) oil prices and thereby decreases (increases) oil market uncertainty.

al. (2009) that in contrast to the 1970s, the relationship between inflation and oil prices is muted during the 2000s. Similarly, Ewing und Thompson (2007) have shown that oil prices lag industrial production but lead consumer prices in the period 1982-2005. Lastly, the LRT and the AIC in Table 5.3 reveal that all GARCH-MIDAS-X specifications with significant macroeconomic variables achieve a better fit than the restricted GARCH(1,1). While the information criteria of the various GARCH-MIDAS-X specifications are pretty similar, it is interesting that UR and LI achieve the best fit which is in line with the stock market results.

Figure 5.3 shows the GARCH-MIDAS-UR estimates of the annualized long-term volatility components for the two markets. While the level of oil price volatility is about twice as high as the one of the stock prices, the evolution of the two components is very similar across markets. The observation that the macroeconomic environment affects long-term oil and stock volatility in almost the same way is very interesting. Recently, Kilian and Vega (2011) investigated whether oil prices can be viewed as asset prices. By regressing daily oil price changes on macroeconomic news they find that oil prices do not react to U.S. macroeconomic aggregates and, hence, conclude that oil prices behave very differently from asset prices. However, our results suggest that at least the second moments of oil and stock returns respond in a comparable fashion to macroeconomic news.

Figure 5.3 about here

Based on the estimates $\hat{\omega}_i$ and $\hat{\theta}_i$, we now quantitatively compare the persistence and the magnitude of the effect of changes in the macro variables on long-term volatility in both markets. As can be seen from Tables 5.2 and 5.3, for each macroeconomic variable with significant θ_i, the corresponding estimate $\hat{\omega}_i$ is considerably larger in the oil than in the stock market. Hence, the effect of changes in macro variables on long-term volatility is less persistent in the oil market than in the stock market. Following Engle et al. (2009), we can compute the magnitude of an effect of a one percent (unit) change in the macroeconomic variable X_τ on the long-term volatility in month $\tau+1$ according to $exp(\hat{\theta}_i \cdot \varphi_1(\hat{\omega}_i))-1$. Even though we observe differences in the persistence of the effects across markets, we find that the magnitude of the effects in $\tau+1$ is pretty similar. In case of the GARCH-MIDAS-UR, a one percentage point increase in this variable is accompanied by an increase in long-term stock market volatility of 0.799%, while oil price volatility increases by 0.712%.

5.5.2 Determinants of long-term correlations

Next, we analyze the macroeconomic determinants of the long-term oil-stock correlation. Now we consider two benchmark specifications. The first natural benchmark is the

DCC-GARCH model which is obtained from the DCC-MIDAS specification by replacing $\bar{\rho}_{12,\tau}$ with the unconditional correlation of the volatility-adjusted GARCH residuals. The second benchmark specification follows Colacito et al. (2011) and uses backward-looking monthly realized correlations as explanatory variables. We estimate two versions: one where m_{12} und θ_{12} vary freely (DCC-MIDAS-RC) and one where we restrict these parameters to $m_{12} = 0$ and $\theta_{12} = 1$ (DCC-MIDAS-RC restr). In the DCC-MIDAS-X specifications we replace the realized correlations with key macroeconomic figures. In order to facilitate comparison between the various DCC, DCC-MIDAS-RC and DCC-MIDAS-X models, the first step volatility-adjusted residuals for all models are taken from the benchmark GARCH(1,1) specification. As in case of the long-term volatilities, we find that the optimal lag length is equal to three MIDAS lag years, i.e. we choose $K_c = 36$.
Table 5.4 presents the estimation results. Clearly, in all specifications the estimated parameters a and b are highly significant and sum up to a value of less than one. That is, the quasi-correlations are mean-reverting either to the unconditional correlation in the DCC-GARCH case or to the long-term correlation in the various DCC-MIDAS-X specifications. The estimates of θ_{12} indicate that all variables which represent the current stance of the economy or the future economic outlook significantly affect the long-run oil-stock correlation. In line with our analysis in Section 5.5.1, we find negative θ_{12} coefficients on IP, NFP, NAI, and LI, while the coefficient on UR is positive. The estimates imply that a contraction of macroeconomic activity leads to an increase of the long-term correlation. Moreover, none of the two inflation measures can explain the long-run co-movements in stock and oil prices which again reinforces our findings from the long-term volatility analysis.

Table 5.4 about here

According to the LRT, all DCC-MIDAS-X specifications with significant θ_{12} estimates as well as the restricted DCC-MIDAS-RC model are preferred to the nested DCC-GARCH. Hence, there is convincing evidence in favor of the component models which allow for a time-varying long-term correlation. In addition, a comparison of the information criteria also confirms the superiority of the DCC-MIDAS-X models relative to the DCC-MIDAS-RC benchmark specification. Finally, according to the AIC the DCC-MIDAS-NAI model achieves the best fit among all specifications. We explain below that the forward looking properties of the NAI which gauges future economic activity as well as inflationary pressures (and thereby future monetary policy) are particularly relevant for anticipating changes in the oil-stock correlation. Note that the model which includes UR still performs second best.
Figure 5.1 shows the estimated dynamics of the short- and long-run correlations based

on the DCC-MIDAS-NAI specification together with a rolling-window of yearly realized correlations. First, although the unconditional correlation between stock and oil returns was found to be 0.15, the figure shows that there is substantial time-variation in the realized correlations with prolonged periods of positive or negative correlations. While the short-run component closely follows the behavior of the realized correlations, the long-run correlation evolves much more smoothly. Both the realized correlations as well as the short-run correlations follow this long-run trend component. Figure 5.1 reveals a very interesting cyclical pattern in the evolution of the long-run correlation. At the beginning of the sample period in 1993 the correlation takes a value of 0.14 and then starts to decrease until it reaches a minimum of -0.12 in 1994. It stays in the negative territory until mid-2000. From mid-2000 onwards the long-term correlation starts to increase and turns positive before the recession of 2001 (first shaded area). The correlation further increases until it reaches a peak of 0.25 at the end of the recession. The figure shows that the long-term correlation remains above 0.2 for the two subsequent years, which are followed by a smooth decrease and a period of negative correlations during the years 2005 to 2006. Again, the long-term correlation starts to increase almost two years before the recession of 2007-2009 and becomes positive clearly before the beginning of the recession (second shaded area). At the end of this recession we observe a peak at 0.60. Finally, the correlation starts to decrease smoothly.

To provide an economic interpretation of the correlation dynamics we refer to Figure 5.4 which depicts the long-term correlation along with the NAI. First, the figure clearly shows the inverse relationship between the NAI and the long-term oil-stock correlation which was already evident from the negative θ_{12} estimate in Table 5.4. On average, the oil-stock correlation is positive (negative) when the NAI takes negative (positive) values, i.e. when the economy is expanding below (above) trend growth. Interestingly, when the NAI turns negative before and during the 2001 and 2007-2009 recessions the long-term correlation steeply increases, while it decreases more gradually when the NAI stays in the negative territory in the aftermath of the recessions. On the other hand, the long period of growth above trend from 1994 to 1999 is accompanied by a period of negative oil-stock correlations.[5] Our empirical evidence for a counter cyclical oil-stock correlation is again perfectly in line with the recent evidence in Harris et al. (2009) and Kilian (2009) in favor of a positive oil-growth relation. Similarly, the results in Section 5.5.1 support the view that good news on the macroeconomy are also good news for the oil price, i.e. reduce oil volatility. Increasing economic activity leads to higher oil demand and, consequently,

[5] Only during 1995 when the NAI takes a few negative values the long-term correlation temporarily increases but remains negative.

higher oil prices. Further, Kilian and Park (2009) argue that in an early phase of an expansion increasing oil prices may not have negative effects on the stock market. This is because in the short-run the positive effect of higher economic activity on expected future cash flows dominates and, hence, the oil-stock correlation will be positive. However, in the long-run the negative effect of increasing oil prices on corporate cash flows will dominate and turn the oil-stock correlation negative.

Figure 5.4 about here

The long-term correlation in Figure 5.4 very much supports these views. Before and during both recessions bad news on the NAI lead to sharply decreasing stock and oil prices and, therefore, to a positive oil-stock correlation. The fact that the correlation turns positive well before both recessions is remarkable and suggests that the long-term oil-stock correlation may itself be used as an early recession indicator. During the recovery phases in 2002-2003 and 2010-2011 the improvement in the NAI leads to increasing oil prices and at the same time to upward revisions concerning firms' expected dividends and cash flows. In these periods the oil-stock correlation remains positive, but smoothly decreases. The same rationale also applies to the first year of our sample, which falls into the recovery period after the recession of 1990/91 (see Section 5.6). Finally, during the years 1994-1999 and 2005-2006 the NAI grows above trend for a protracted period which again should positively affect oil prices. However, the (expected) oil price increases now dampen the outlook for future corporate cash flows, i.e. during these periods the good news on the macroeconomy – through the indirect effect via increasing oil prices – turn into bad news for the stock market. Alternatively, the negative effect might also work via interest rates. When the economy is already close to full employment, good news on the NAI should signal higher future interest rates and, hence, be bad news for the stock market. During these strong boom phases the negative effect dominates and leads to a negative long-run oil-stock correlation.

Since the evolution of the long-term correlation is purely driven by variables which represent U.S. aggregate demand, deviations of the short-term component from the long-run trend must be related to other factors which either affect stock and/or oil returns. Typical oil related factors would be oil supply shocks or oil specific, i.e. precautionary, demand. Specifically, the temporary deviation in 2003 may be due to precautionary demand provoked by the second Iraq war (see Figure 5.1). Another example would be the positive correlation signaled by the short-term component as well as the realized correlations around 1998/99. Following the Asian and Russian financial crises, this positive short-term correlation can be explained by simultaneously falling oil and stock prices. Nevertheless, the fact that these deviations occur only for relatively short periods suggests that the

oil-stock correlation can be well explained by U.S. economic activity for most of the time. This result is very much in line with Kilian (2009, p.1068) who reasons that "models of endogenous oil prices should focus on the demand side of the oil market".
A particulary interesting conclusion that can be drawn from the time-varying oil-stock correlation is that regressions of stock returns on oil price changes are likely to be misleading, since the result will depend on the state of the economy. This insight may explain the controversial empirical findings on the oil-stock relationship and squares with the arguments put forward in Kilian and Park (2009).
Next, we discuss the MIDAS lag structure and its implications more closely. Recall that the higher ω_{12} the more weight will be given to the more recent observations of the macro variable and, hence, the faster the weights will decline to zero. Table 5.4 reveals that the lowest ω_{12} is estimated for IP and the highest for NFP. Since the DCC-MIDAS-NAI model produced the best fit for the correlations, we plot in Figure 5.5 the corresponding weighting function. For comparison, we also plot the weighting functions for the GARCH-MIDAS-NAI models for the stock and oil market. The figure shows that the weighting function of the correlation model is nearly linear while the weighting functions of the volatility specifications are rapidly declining. This in turn implies that changes in the NAI have a much more persistent effect on the long-run correlation than on the long-run volatilities. We obtain similar results for each of the other significant macroeconomic variables.

Figure 5.5 about here

In the previous considerations we mainly focused on the DCC-MIDAS-NAI specification to explain the dynamic behavior of the slowly-moving long-run correlation component. However, Table 5.4 clearly reveals that the fit of the DCC-MIDAS-X specifications with IP, NFP, UR, and LI are only slightly inferior. Figure 5.6 displays the estimated long-run correlations from the corresponding specifications. The figure nicely illustrates that the long-term components of all specifications follow the same pattern and, hence, further support our argument that the long-term oil-stock correlation is counter cyclical. Note that the exceptional deviation in the long-term correlation component predicted by IP for October 2005 can be traced back to a significant contraction in industrial production one month earlier which is not reflected to such a strong intend in the other macroeconomic figures (compare Figure 5.2).

Figure 5.6 about here

5.6 Robustness

In this section we present evidence on the robustness of our results by considering alternative measures of the stock market and extending our sample period. We first make use of two alternative return series representing the stock market: the S&P 500 and the DJIA index (both were obtained from the FRED database). Second, we extend the initial sample to the period 1986-2011 and thus include the first Gulf War from 1990 to 1991.

5.6.1 S&P 500 and DJIA

Tables 5.5 and 5.6 refer to the specifications including the S&P 500 and the DJIA index, respectively. The coefficient estimates in both tables broadly confirm the results presented in Table 5.4. Again, all variables on the current economic stance as well as the future economic outlook significantly affect the long-term correlation between stock and oil prices. As for the CRSP, the DCC-MIDAS-NAI specification achieves the best fit in both cases.

Tables 5.5 and 5.6 about here

5.6.2 Extended sample

The parameter estimates for the extended sample are qualitatively identical to the ones for the original sample and, hence, strongly confirm our previous interpretations.[6] Nevertheless, the extended sample allows for some further insights into the behavior of the long- and short-term correlation components. Both components are plotted in Figure 5.7 for the DCC-MIDAS-NAI model. While the behavior of the long-term correlation component during the recession of 1990/91 exhibits the same pattern as described above, the short-term correlation component sharply declines from 0.15 to -0.30 with the invasion of Kuwait on August 2. In line with Kilian (2009), we view the evolution of the short-term component as mainly triggered by precautionary demand. On January 18, the short-term component realized an all-time minimum of -0.47 as a consequence of the 40% oil price drop accompanied by a stock market recovery of more than 3%. This was caused by the decision of the Bush administration to compensate for shortfalls in oil supply by releasing the strategic crude oil reserves. Finally, at the beginning of 1993 the short-term correlation reverts to the long-term component.

Figure 5.7 about here

[6]The estimates are not reported but are of course available upon request to the authors.

5.7 Conclusion

We investigate the effect of changes in the U.S. macroeconomic environment on the long-term co-movements between crude oil and stock price returns. For this, we extend the two-component DCC-MIDAS model of Colacito et al. (2011) by allowing the slowly-moving long-term correlation component to be determined endogenously by the variation of key macroeconomic figures. We show that changes in macroeconomic variables which reflect the current stance of the economy as well as the future economic outlook can anticipate counter cyclical fluctuations in the long-term correlation. More specifically, our model predicts a negative correlation during prolonged periods of strong economic expansions, while a positive correlation is observed during recessions and recoveries. The correlation pattern suggests that during recessions (expansions with growth below or at trend), bad (good) news on the macroeconomy are bad (good) news for the stock as well as for the oil market. However, during periods with strong growth above trend, good news on the macroeconomy are still good news for the oil market but become bad news for the stock market. This is because both the further increasing oil prices as well as the expectation of rising interest rates have a depressing effect on the stock market.

Our results provide further evidence for the argument put forward in Barsky and Kilian (2004) and Kilian (2009) that oil price changes should not be considered exogenous with respect to the U.S. economy. The counter cyclical behavior of the long-term oil-stock correlation squares with the recent evidence in Harris et al. (2009) and Kilian and Park (2009) that oil price developments have been synchronized with the business cycle. Moreover, the finding that the sign of the oil-stock correlation varies with the state of the economy, may explain the conflicting empirical evidence in previous studies on the oil-stock relationship when simple regressions of stock returns on oil price changes are employed.

Finally, we also assess the impact of macroeconomic developments on the long-term volatilities of crude oil and stock price returns. Our results show that the long-term volatilities in both markets are driven by the same macroeconomic factors. Hence, while Kilian and Vega (2011) report that oil prices in contrast to asset prices do not respond to U.S. macroeconomic news, at the least the second moments of oil price returns behave very much like those of asset prices.

Tables

Table 5.1: Descriptive Statistics (January 1993 - November 2011)

Variable	Obs	Min	Max	Mean	Std. Dev.*	Skewness	Kurtosis
Panel A (Daily return data)							
Oil	4744	-17.09	16.41	0.0332	39.18	-0.19	7.73
CRSP	4744	-9.00	11.52	0.0365	19.41	-0.15	11.07
Panel B (Monthly macro data)							
Current stance of the economy							
IP	227	-4.23	2.15	0.17	0.69	-1.73	11.30
NFP	227	-820	508	98.21	234	-1.46	5.93
UR	227	-0.50	0.60	0.01	0.18	0.62	4.37
Future economic outlook							
NAI	227	-4.46	1.55	-0.14	0.87	-1.92	8.71
LI	227	-3.82	2.84	1.17	1.20	-1.83	7.25
Inflation rates							
PPI	227	-5.48	2.94	0.24	1.11	-1.30	8.97
CPI	227	-1.83	1.37	0.21	0.28	-1.79	16.56

Notes: *The standard deviations are annualized for the daily return series.

Table 5.2: GARCH-MIDAS parameter estimates: CRSP

Variable	μ_1	α_1	β_1	m_1	θ_1	ω_1	LLF	LRT	AIC
Current stance of the economy									
IP	0.0678***	0.0843***	0.9056***	0.3873*	−0.9893**	2.7737**	-6520.74	4.48	2.7487
	(0.0119)	(0.0123)	(0.0137)	(0.2339)	(0.4961)	(1.2659)		[0.1065]	
NFP	0.0680***	0.0873***	0.8998***	0.3714*	−0.0019***	8.8520	-6518.19	9.58	2.7476
	(0.0119)	(0.0128)	(0.0147)	(0.2038)	(0.0005)	(5.6496)		[0.0083]	
UR	0.0688***	0.0863***	0.9010***	0.1403	3.9929***	5.5678*	-6515.07	15.82	**2.7463**
	(0.0117)	(0.0121)	(0.0138)	(0.2015)	(0.7600)	(2.9746)		[0.0004]	
Future economic outlook									
NAI	0.0682***	0.0864***	0.9010***	0.1036	−0.6120***	6.0568	-6517.63	10.7	2.7474
	(0.0118)	(0.0126)	(0.0144)	(0.2077)	(0.1382)	(4.1461)		[0.0047]	
LI	0.0681***	0.0867***	0.8999***	0.6557***	−0.4055***	5.8714	-6515.83	14.3	2.7466
	(0.0118)	(0.0125)	(0.0142)	(0.2117)	(0.0814)	(4.5227)		[0.0008]	
Inflation rates									
PPI	0.0666***	0.0823***	0.9098***	0.2506	−0.0895	15.8835	-6522.44	1.08	2.7494
	(0.0119)	(0.0120)	(0.0131)	(0.2658)	(0.1545)	(27.6558)		[0.5827]	
CPI	0.0664***	0.0815***	0.9106***	0.1806	0.1907	532.2708***	-6521.69	2.58	2.7491
	(0.0119)	(0.0117)	(0.0127)	(0.2566)	(0.1486)	(0.0188)		[0.2753]	
Benchmark model									
GARCH(1,1)	0.0665***	0.0822***	0.9099***	0.0112***	-	-	-6522.98	-	2.7488
	(0.0119)	(0.0119)	(0.0129)	(0.2281)					

Notes: The numbers in parentheses are Bollerslev-Wooldridge robust standard errors. ***, **, * indicate significance at the 1 %, 5 %, and 10 % level. LLF is the value of the maximized likelihood function and AIC is the Akaike information criterion. The numbers in bold letters indicate the model with the smallest value of the information criterion. LRT is the likelihood ratio test $LR = 2[L_{UR} - L_R]$, where L_{UR} is the likelihood of the unrestricted GARCH-MIDAS-X specification and L_R is the likelihood of the restricted benchmark model. The numbers in brackets are p-values.

Table 5.3: GARCH-MIDAS parameter estimates: Oil market

Variable	μ_2	α_2	β_2	m_2	θ_2	ω_2	LLF	LRT	AIC
Current stance of the economy									
IP	0.0486	0.0582***	0.9230***	1.8422***	−0.4641**	6.9842***	-10589.1	5.2	4.4620
	(0.0331)	(0.0157)	(0.0221)	(0.1389)	(0.2353)	(2.2342)		[0.0743]	
NFP	0.0475	0.0597***	0.9225***	1.8526***	−0.0008**	16.5369***	-10589.2	5.0	4.4621
	(0.0331)	(0.0158)	(0.0212)	(0.1485)	(0.0004)	(6.1789)		[0.0821]	
UR	0.0497	0.0579***	0.9230***	1.7406***	1.6849***	13.4396**	-10586.3	10.8	**4.4609**
	(0.0327)	(0.0146)	(0.0202)	(0.1310)	(0.6116)	(6.1886)		[0.0045]	
Future economic outlook									
NAI	0.0484	0.0571***	0.9251***	1.7239***	−0.2890**	15.1178**	-10588	7.4	4.4616
	(0.0331)	(0.0156)	(0.0215)	(0.1378)	(0.1148)	(6.8522)		[0.0247]	
LI	0.0481	0.0573***	0.9239***	2.0059***	−0.2084***	18.8496***	-10586.5	10.4	**4.4609**
	(0.0329)	(0.0154)	(0.0212)	(0.1542)	(0.0609)	(7.0180)		[0.0055]	
Inflation rates									
PPI	0.0461	0.0582***	0.9264***	1.8007***	−0.0881	16.0532*	-10591.1	1.2	4.4629
	(0.0333)	(0.0164)	(0.0209)	(0.1549)	(0.1196)	(9.3534)		[0.5488]	
CPI	0.0470	0.0599***	0.9248***	1.7799***	0.0291	5.9839	-10591.7	0.0	4.4631
	(0.0331)	(0.0156)	(0.0197)	(0.2252)	(0.7773)	(14.9089)		[1.0000]	
Benchmark model									
GARCH(1,1)	0.0470	0.0599***	0.9248***	1.7858***	-	-	-10591.7	-	4.4623
	(0.0332)	(0.0156)	(0.0196)	(0.1586)					

Notes: See Notes of Table 5.2.

Table 5.4: DCC-MIDAS parameter estimates: CRSP and oil market

Variable	a	b	m_{12}	θ_{12}	ω_{12}	LLF	LRT	AIC
Current stance of the economy								
IP	0.0189***	0.9713***	0.2143***	−0.6888***	1.5996*	-4665.09	13.1	1.9668
	(0.0063)	(0.0107)	(0.0594)	(0.1936)	(0.8845)		[0.0044]	
NFP	0.0190***	0.9706***	0.1982***	−0.0010***	3.8517	-4664.52	14.24	1.9665
	(0.0064)	(0.0118)	(0.0545)	(0.0003)	(3.3013)		[0.0026]	
UR	0.0204***	0.9636***	0.0582	2.6018***	1.7203**	-4663.00	17.28	1.9659
	(0.0058)	(0.0112)	(0.0360)	(0.6054)	(0.7918)		[0.0006]	
Future economic outlook								
NAI	0.0192***	0.9659***	0.0462	−0.3502***	2.0487**	-4662.10	19.08	**1.9655**
	(0.0057)	(0.0108)	(0.0368)	(0.0771)	(1.1143)		[0.0003]	
LI	0.0192***	0.9684***	0.3450***	−0.2142***	2.2960	-4663.71	15.86	1.9662
	(0.0059)	(0.0110)	(0.0800)	(0.0544)	(1.6828)		[0.0012]	
Inflation rates								
PPI	0.0190**	0.9774***	0.1669	−0.1590	7.7937	-4671.24	0.8	1.9694
	(0.0076)	(0.0105)	(0.1055)	(0.1990)	(10.2830)		[0.8495]	
CPI	0.0201***	0.9744***	0.4247*	−1.5224	3.7484*	-4670.59	2.1	1.9691
	(0.0074)	(0.0113)	(0.2352)	(1.1307)	(2.1652)		[0.5519]	
Benchmark models								
DCC-RC	0.0225***	0.9582***	0.0324	0.8704***	5.3086*	-4668.95	5.38	1.9684
	(0.0058)	(0.0108)	(0.0366)	(0.3202)	(2.7355)		[0.1460]	
DCC-RC restr	0.0228***	0.9574***	-	-	4.7764**	-4669.56	8.98	1.9678
	(0.0060)	(0.0101)			(2.3004)		[0.0414]	
DCC	0.0191***	0.9775***	-	-	-	-4671.64	-	1.9683
	(0.0035)	(0.0046)						

Notes: See Notes of Table 5.2. The LRT compares the unrestricted DCC-MIDAS-X models with the DCC-GARCH specification.

Table 5.5: DCC-MIDAS parameter estimates: S&P 500 and oil market

Variable	a	b	m_{12}	θ_{12}	ω_{12}	LLF	LRT	AIC
Current stance of the economy								
IP	0.0212***	0.9680***	0.1888***	−0.6734***	1.6217*	-4662.74	12.96	1.9658
	(0.0065)	(0.0109)	(0.0585)	(0.1877)	(0.8294)		[0.0047]	
NFP	0.0214***	0.9668***	0.1737***	−0.0010***	4.0441	-4661.84	14.76	1.9654
	(0.0065)	(0.01162)	(0.0528)	(0.0002)	(3.1617)		[0.0020]	
UR	0.0229***	0.9595***	0.0354	2.5942***	1.7862**	-4659.96	18.52	1.9646
	(0.0060)	(0.0114)	(0.0354)	(0.5858)	(0.7948)		[0.0003]	
Future economic outlook								
NAI	0.0215***	0.9619***	0.0233	−0.3485***	2.1222*	-4659.20	20.04	**1.9643**
	(0.0059)	(0.0110)	(0.0362)	(0.0744)	(1.0954)		[0.0002]	
LI	0.0215***	0.9648***	0.3189***	−0.2114***	2.4474	-4660.93	16.58	1.9650
	(0.0061)	(0.0111)	(0.0769)	(0.0516)	(1.6603)		[0.0009]	
Inflation rates								
PPI	0.0212**	0.9745***	0.1416	−0.1549	7.3416	-4668.84	0.76	1.9683
	(0.0079)	(0.0111)	(0.0962)	(0.2013)	(9.3929)		[0.8590]	
CPI	0.0227***	0.9701***	0.4476*	−1.7569	3.2581	-4667.84	2.76	1.9679
	(0.0074)	(0.0114)	(0.2585)	(1.2303)	(2.3148)		[0.4301]	
Benchmark models								
DCC-RC	0.0251***	0.9543***	0.2739	0.8521***	5.5409**	-4666.33	5.78	1.9673
	(0.0059)	(0.0116)	(0.0346)	(0.3037)	(2.6691)		[0.1228]	
DCC-RC restr	0.0254***	0.9534***	-	-	4.8424**	-4666.92	4.60	1.9667
	(0.0061)	(0.0109)			(2.2679)		[0.0320]	
DCC	0.0211***	0.9749***	-	-	-	-4669.22	-	1.9672
	(0.0074)	(0.0102)						

Notes: See Notes of Table 5.4.

Table 5.6: DCC-MIDAS parameter estimates: DJIA and oil market

Variable	a	b	m_{12}	θ_{12}	ω_{12}	LLF	LRT	AIC
Current stance of the economy								
IP	0.0248***	0.9628***	0.1692***	-0.6670^{***}	1.3994**	-4642.36	12.7	1.9572
	(0.0068)	(0.0115)	(0.0558)	(0.1806)	(0.5774)		[0.0053]	
NFP	0.0252***	0.9607***	0.1552***	-0.0010^{***}	3.0553	-4641.41	14.6	1.9568
	(0.0065)	(0.0116)	(0.0515)	(0.0003)	(2.2618)		[0.0022]	
UR	0.0264***	0.9533***	0.0153	2.6614***	1.5454***	-4638.81	19.8	1.9557
	(0.0060)	(0.0116)	(0.0345)	(0.5519)	(0.5393)		[0.0002]	
Future economic outlook								
NAI	0.0252***	0.9552***	0.0034	-0.3542^{***}	1.8093**	-4638.19	21.04	**1.9554**
	(0.0060)	(0.0115)	(0.0350)	(0.0714)	(0.7539)		[0.0001]	
LI	0.0252***	0.9591***	0.3020***	-0.2126^{***}	2.0039*	-4640.40	16.62	1.9564
	(0.0063)	(0.0114)	(0.0741)	(0.0503)	(1.0740)		[0.0008]	
Inflation rates								
PPI	0.0239**	0.9707***	0.1216	-0.1418	6.4377	-4648.21	1.00	1.9597
	(0.0087)	(0.0127)	(0.0868)	(0.2063)	(9.4750)		[0.8013]	
CPI	0.0267***	0.9630***	0.5362*	-2.2800	2.3720	-4646.32	4.78	1.9589
	(0.0072)	(0.0117)	(0.3168)	(1.4860)	(2.1948)		[0.1886]	
Benchmark models								
DCC-RC	0.0288***	0.9460***	0.0248	0.9111***	4.8276***	-4643.71	10.00	1.9578
	(0.0059)	(0.0123)	(0.0321)	(0.2532)	(1.7707)		[0.0186]	
DCC-RC restr	0.0289***	0.9456***	-	-	4.5605***	-4644.13	9.16	1.9571
	(0.0060)	(0.0116)			(1.6072)		[0.0025]	
DCC	0.0237***	0.9713***	-	-	-	-4648.71	-	1.9586
	(0.0081)	(0.0115)						

Notes: See Notes of Table 5.4.

Figures

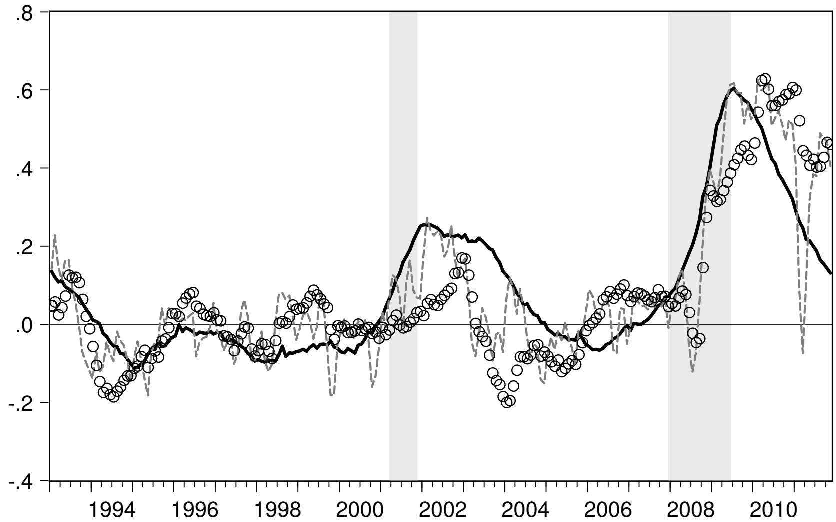

Figure 5.1: The figure shows the DCC-MIDAS-NAI estimates of the short-term (dashed line) and long-term (bold black line) oil-stock correlation. The circles correspond to one-year rolling window realized correlations. Each series is shown at a monthly frequency. Monthly realizations of the daily short-term and realized correlations are obtained by computing monthly averages. Shaded areas represent NBER recession periods.

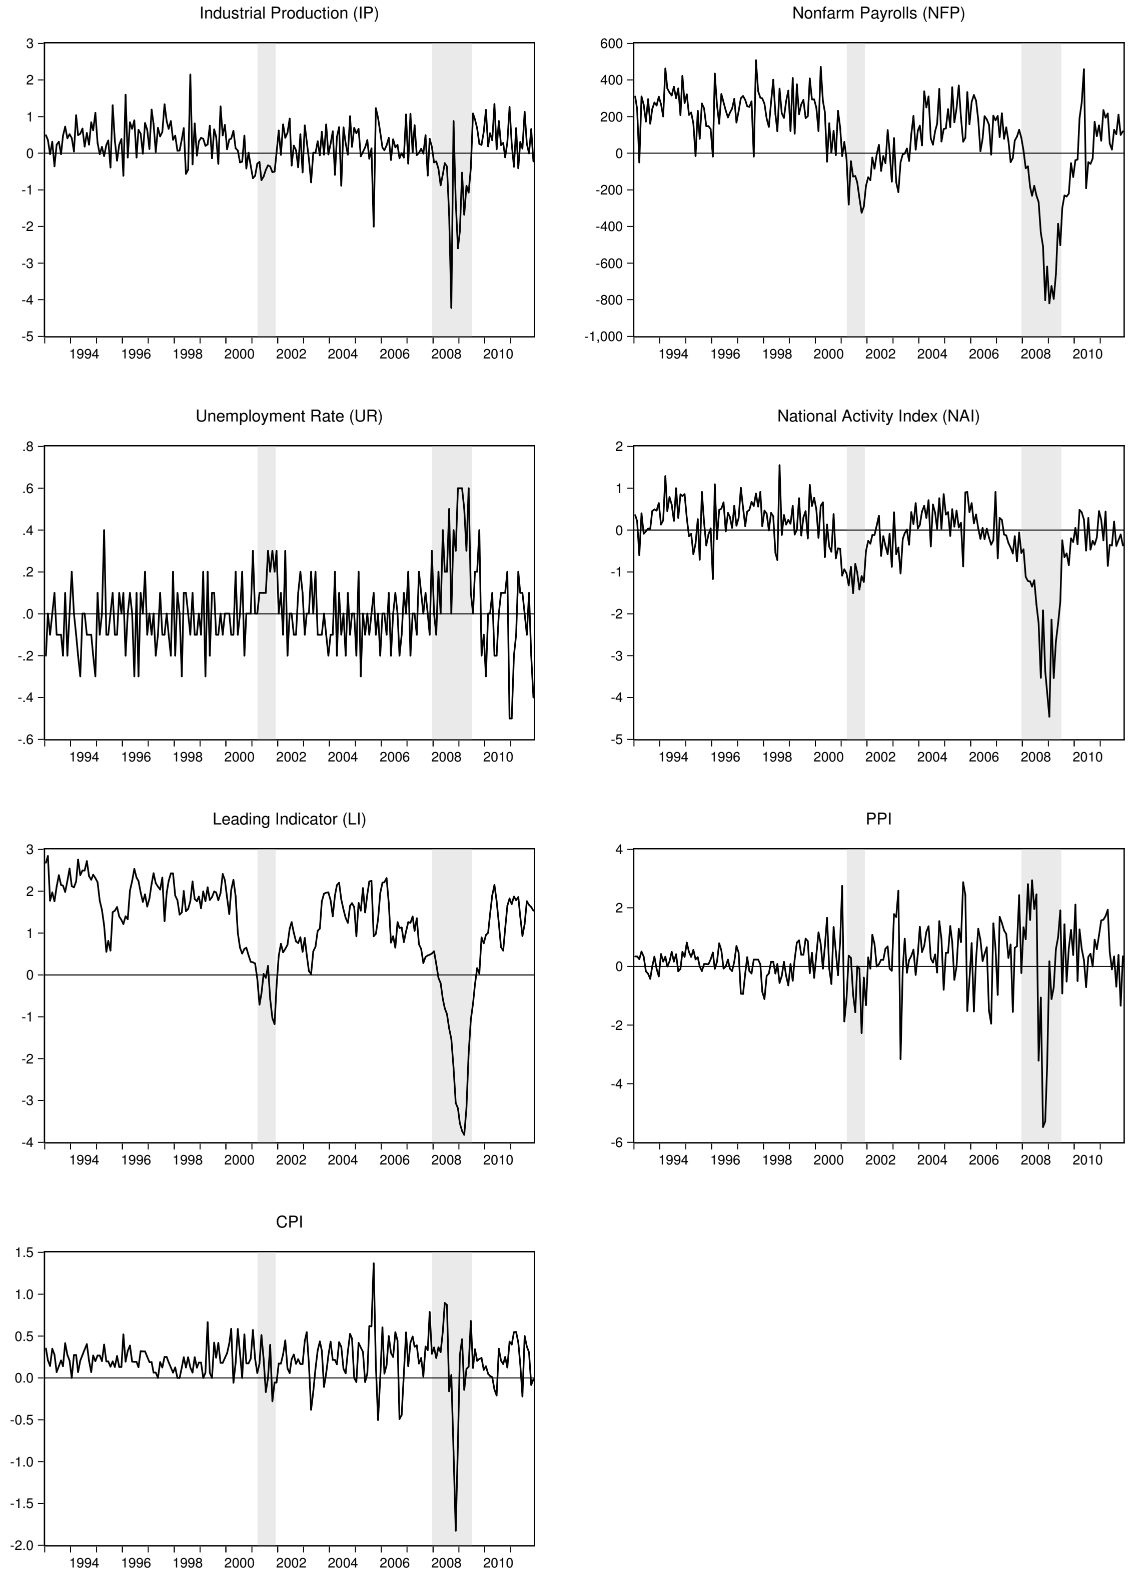

Figure 5.2: The figure shows the development of the macroeconomic explanatory variables. Shaded areas represent NBER recession periods.

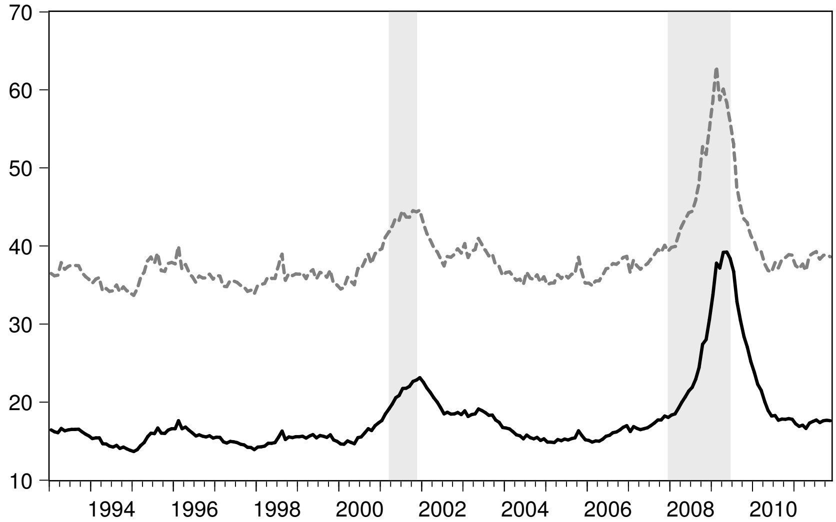

Figure 5.3: The figure shows the annualized long-term volatility components (standard deviations) obtained from the GARCH-MIDAS-UR specification. The bold line refers to the stock market, the dashed line to the oil market. Shaded areas represent NBER recession periods.

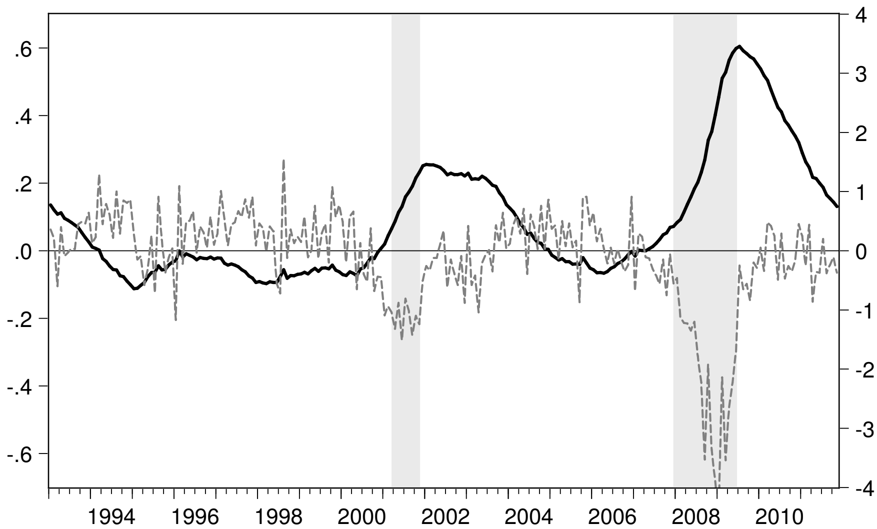

Figure 5.4: The bold black line (left scale) represents the DCC-MIDAS-NAI estimate of the long-term oil-stock correlation. The dashed line (right scale) corresponds to the NAI. Shaded areas represent NBER recession periods.

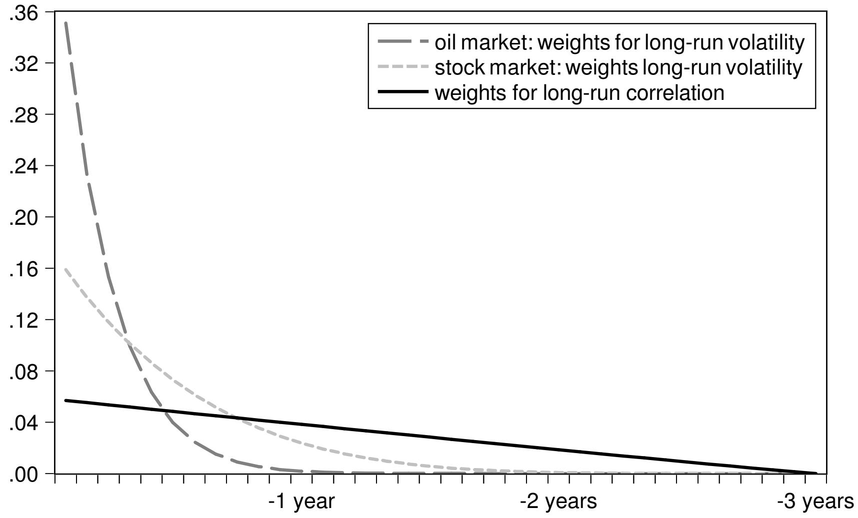

Figure 5.5: The figure shows the estimated weighting functions for the long-term volatilities based on the GARCH-MIDAS-NAI and for the long-term correlation based on the DCC-MIDAS-NAI. While the bold black line refers to the long-term correlation, the light-gray and the dark-gray dashed lines refer to the long-term volatilities of CRSP and of oil price returns, respectively.

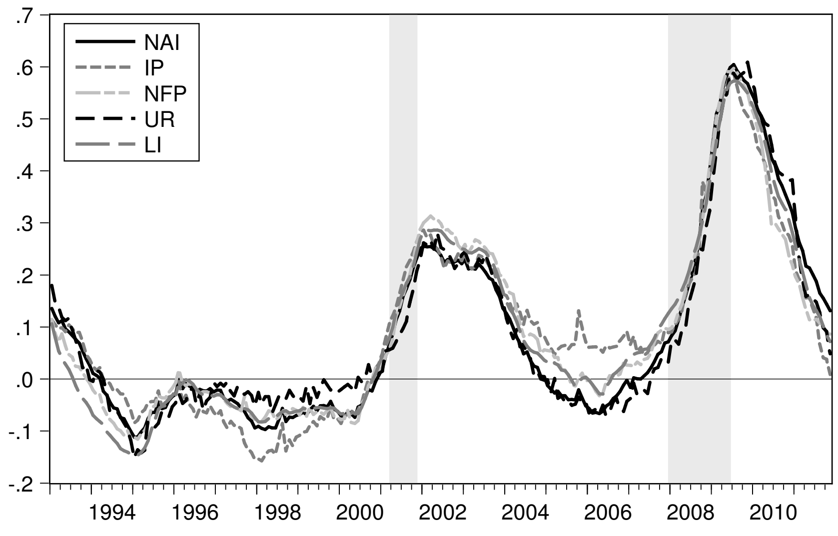

Figure 5.6: The figure shows the DCC-MIDAS-X estimates of the long-term oil-stock correlations for all significant macroeconomic variables. Shaded areas represent NBER recession periods.

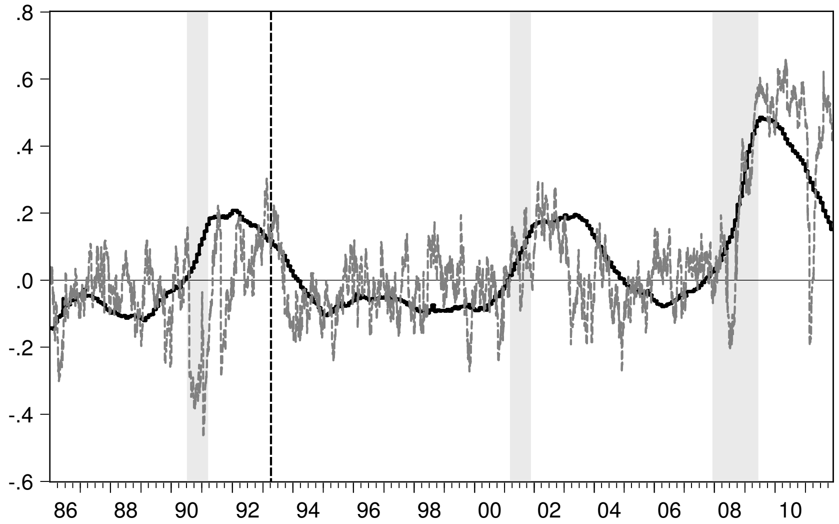

Figure 5.7: The figure shows the DCC-MIDAS-NAI estimates of the short-term (dashed dark-gray line) and long-term (bold black line) correlation for the extended sample (1986 - 2011). Each series is shown at a daily frequency. The vertical dashed line indicates the beginning of the shorter sample. Shaded areas represent NBER recession periods.

Chapter 6

Pricing the risk of oil

6.1 Introduction

In this paper, we empirically investigate the relationship between the U.S. stock market and the oil price during the 10-year period from June 2001 to June 2011 in the framework of Merton's (1973) Intertemporal Capital Asset Pricing Model (ICAPM).

Even though there has been substantial empirical research on the impact of the oil price on stock performance, previous studies usually employ rather ad hoc specifications to assess industry-specific oil-stock relations. Using a market model augmented by oil returns, Faff and Brailsford (1999) and Nandha and Faff (2008) analyze the impact of the oil price on industry-specific stock returns for the Australian stock market and for global industry indices, respectively. Both studies conclude that the oil and gas industry benefits from increasing oil prices, while for other industries the opposite is the case. Recently, Narayan and Sharma (2011) broadly confirm these results for individual U.S. stocks. Other studies such as Sadorsky (2001) or Boyer and Filion (2007) estimate panel specifications to reveal whether the oil price affects stock returns. Studies on the aggregate level are provided by Jones and Kaul (1996) and Driesprong et al. (2008).

Besides the lack of theoretical motivation, Kilian and Park (2009) argue that the specifications used in the previous literature are generally inappropriate to capture the impact of the oil price on the stock market. First, the authors hint at the problem of reverse causality as oil and stock prices have been driven by the same forces since the 1970s. Second, making use of the structural oil price decomposition introduced in Kilian (2009), Kilian and Park (2009) reveal that the stock price reaction to an oil shock depends on the specific type of the shock. Moreover, the authors argue that the relative importance of the shocks could change over time. Finally, Conrad et al. (2012) show that the conditional correlation between stock index and oil returns is closely linked to the business cycle and considerably fluctuates over time taking on positive as well as negative values which confirms the argument of Kilian and Park (2009) that regressing stock returns on oil price changes could induce spurious conclusions.

In line with Cifarelli and Paladino (2010, 2012), we suggest to take Merton's (1973) ICAPM as the theoretical basis for our empirical analysis as opposed to the ad hoc regressions estimated in the previous studies. In equilibrium, the ICAPM predicts that the representative risk-averse investor charges a market risk premium and a premium for bearing the risk of adverse shifts in the investment opportunity set that negatively affect the investor's future wealth. While the market risk of an individual stock is measured by the covariance between the returns on the stock and the market portfolio, the risk of adverse shifts in the investment opportunity set is represented by the covariance between the return on the stock and changes in state variables that are related to the investor's intertemporal consumption. Merton (1973) shows that the model implies a positive relationship between market risk and expected return. The sign on the relationship between the expected return and the risk associated with changes in the state variables is determined by the sensitivity of the investor's marginal utility of wealth with respect to the state variables. Put differently, a state variable is priced in the ICAPM as far as the investor's marginal utility of wealth depends on the level of the state variable. In general, the investor demands stocks that tend to pay off in periods of a high marginal utility of wealth.

In this paper we aim at analyzing whether the price of oil can be considered a state variable in the ICAPM. In particular, we are interested in whether the marginal utility of wealth is positively or negatively related to the oil price. Based on the sign of the sensitivity of the investor's marginal utility of wealth with respect to the oil price we can re-evaluate the findings of the previous literature. According to the structural oil price decomposition of Kilian (2009) used in Kilian and Park (2009), the dominance of aggregate demand shocks would imply that high oil prices are accompanied by strong economic performance. Hence, the investor is well off when the oil price is high such that the marginal utility of wealth and the oil price are inversely related. In contrast, the dominance of supply-side shocks and/or precautionary demand shocks would imply that periods of high oil prices are accompanied by weak stock market performance. In this case the marginal utility of wealth would be high when the oil price is high.

Using an extension of the bivariate CCC-GARCH-in-mean specification of Scruggs (1998), Cifarelli and Paladino (2010, 2012) reveal a positive relation between expected weekly DJIA index returns and market risk, though, the estimated market risk coefficient is considerably smaller than the reasonalbe estimates of Scruggs (1998) or Bali and Engle (2010). Furthermore, the authors find that DJIA returns are positively affected by the conditional covariation between stock and oil returns during the period 2000-2009. Scruggs and Glabadanidis (2003), however, point out that models that impose a constant correlation structure are inferior to models that allow the correlation to vary over time. Given the strong fluctuations in the oil-stock correlations shown in Conrad et al. (2012), we conjecture that this is of crucial importance for our study. Following Bali and Engle (2010) our identification strategy is a three-step estimation approach that allows to include a large number of individual stocks. First, we estimate conditional variances of and dynamic conditional correlations between excess returns on each available S&P 500 stock, the market portfolio, and changes in the state variables. Based on the estimated variances and correlations we construct conditional covariances. For the last step, Bali and Engle (2010) propose to estimate panel regressions of the individual stock returns on the conditional covariances employing Zellner's seemingly unrelated regression approach. Given this empirical framework, we investigate the risk-return relation with a special focus on whether the price of oil constitutes a priced state variable in the ICAPM. Furthermore, we analyze the dynamic structure of the market and the oil risk premia across the individual industries contained in the S&P 500 index.

The main findings of the paper can be summarized as follows. Our empirical results clearly confirm the existence of a positive relationship between market risk and expected return which advocates the findings of previous studies such as Scruggs (1998), Guo and Withelaw (2006), or Bali and Engle (2010) among many others. We find that investors require significantly positive risk premia for holding stocks that positively covary with the market, and hence, tend to pay off when market performance is strong. Our model predicts statistically significant and economically reasonable market risk premia across individual industries. While the evolution of the industry-specific premia is pretty similar, their magnitude differs strongly. Highest expected risk premia are observed for financials and energy stocks. Investors require relatively low premia for consumer staples, health care, and utility stocks which implies that these stocks are used to hedge against the market. For each industry the model predicts exploding required risk premia after the bankruptcy filing of Lehman Brothers in September 2008 where again financials and energy stocks exhibit the highest premia. Interestingly, we find that the risk premia for such stocks are still considerably higher after the crisis compared to the pre-crisis levels.

The inclusion of the conditional covariance between stock and oil returns in the baseline specification reveals that the relative risk aversion coefficient estimated in the baseline specification is too large which generates too high required industry-specific market risk premia. Nevertheless, controlling for intertemporal hedging demand induced by the conditional covariation with changes in the oil price confirms the existence of a positive risk-return relationship. Most importantly, our results show that the price of oil can be considered a priced state variable in the ICAPM. In particular, we find that the investor's marginal utility of wealth is inversely related to the oil price. In periods of high oil prices, the investor values an additional unit of wealth less than in periods of low oil prices since the investor is already financially better off. This result is consistent with the argument in Barsky and Kilian (2004) and Kilian and Park (2009) that the oil price is predominantly driven by aggregate demand shocks, and hence, by the same forces as economic activity and stock prices. The finding contradicts the view that within our sample period supply-side shocks have been responsible for the oil price development.

The finding that the investor's marginal utility of wealth is negatively related to the oil price implies that decreasing oil prices reflect an adverse shift in the investment opportunity set. To hedge against this shift the investor increases his demand for stocks that are negatively correlated with the oil price since these stocks tend to pay off when the oil price is low. Hence, the investor accepts negative risk premia for such stocks, while he charges positive risk premia for stocks that have high association with the oil price and tend to pay off when the oil price is high. We show that even the sign of the correlation between oil returns and most industry-specific stock returns changes over time such that the usage of models that are not able to capture dynamic correlations (as in Scruggs (1998) and Cifarelli and Paladino (2010, 2012)) are inappropriate to model the dynamic risk premia. Taking the dynamic correlations into account reveals significant differences in the estimated risk premia across the industries. During the expansion period between the recessions 2000/2001 and 2007-2009 investors require positive oil risk premia for stocks of the industries energy, materials, and utilities since the returns on these stocks positively covary with the oil price. For stocks of the remaining industries the investor accepts lower expected returns due to increased hedging demand. The picture changes dramatically in the aftermath of the bankruptcy filing of Lehman Brothers in September 2008. Our model predicts extreme upswings in the required oil risk premia for stocks of each industry that can be traced back to simultaneously collapsing stock and oil prices leading to high positive oil-stock correlations in combination with high conditional stock and oil price volatilities. The extreme risk premia reflect that none of the stocks serves as a hedging vehicle against decreasing oil prices anymore due to the high positive covariation, and that investors massively reduce their demand for risky stocks because of extreme uncertainty. Towards the end of the recession 2007-2009 and during the expansion period after this recession the estimated oil risk premia steadily decrease. However, for none of the industries the pre-crisis level is reached until the end of our sample period.

Motivated by the findings in the empirical asset pricing literature and the literature on the oil-stock relation, we include additional factors in our model to evaluate the robustness of the relation between risk and expected return. In particular, we control for the impact of the Fama-French factors, macroeconomic risk factors as well as lagged oil and market returns. Finally, we control for additional risk premia induced by the conditional covariation with financial and macro factors. The empirical results imply that our findings on the risk-return relation are robust with respect to all additional factors.
We organize the remainder of the paper as follows. In Section 6.2, we briefly summarize the theoretical model and set up the econometric framework for our empirical analysis. Section 6.3 describes the data. Section 6.4 presents the estimation results and provides an interpretation of our findings. Finally, Section 6.5 concludes.

6.2 Methodology

6.2.1 Theoretical model

The theoretical framework for our empirical specification is given by the Intertemporal CAPM of Merton (1973). In this section we briefly summarize the ICAPM and discuss the role of the oil price in the model.

The Intertemporal CAPM

In the ICAPM the representative investor's indirect utility $U(W_t, \mathbf{S}_t, t)$ of period t depends on his wealth W_t and k state variables $S_{1,t}, \ldots, S_{k,t}$ which are related to the investor's investment opportunity set. The state variables are collected in the vector $\mathbf{S}_t$. In accordance with Scruggs (1998), for $k = 1$ the equilibrium relation between risk and return for asset i is given by

$$\mathbf{E}\left[r_{i,t}|\mathcal{F}_{t-1}\right] = \left[\frac{-U_{WW}W}{U_W}\right] \mathbf{Cov}\left[r_{i,t}, r_{m,t}|\mathcal{F}_{t-1}\right] + \left[\frac{-U_{WS_1}}{U_W}\right] \mathbf{Cov}\left[r_{i,t}, \Delta S_{1,t}|\mathcal{F}_{t-1}\right], \quad (6.1)$$

where $\mathbf{E}\left[r_{i,t}|\mathcal{F}_{t-1}\right]$ is the expected excess return on asset i given the information set $\mathcal{F}_{t-1}$, $\mathbf{Cov}\left[r_{i,t}, r_{m,t}|\mathcal{F}_{t-1}\right]$ is the expected conditional covariance between the excess return on asset i and the market portfolio, and $\mathbf{Cov}\left[r_{i,t}, \Delta S_{1,t}|\mathcal{F}_{t-1}\right]$ is the expected conditional covariance of the excess return on asset i with the changes in the state variable. Consequently, in equilibrium the investor is compensated for the market risk given by $\mathbf{Cov}\left[r_{i,t}, r_{m,t}|\mathcal{F}_{t-1}\right]$ and for the risk of adverse changes in the investment opportunity set represented by $\mathbf{Cov}\left[r_{i,t}, \Delta S_{1,t}|\mathcal{F}_{t-1}\right]$.

U_W, U_{WW}, and U_{WS_1} denote the first and the second partial derivatives of the utility function with respect to wealth W and the state variable S_1. In particular, U_W measures the investor's marginal utility of wealth. The risk aversion assumption implies that $U_W > 0$ and $U_{WW} < 0$. Hence, the expression $\lambda_0 = [-U_{WW}W/U_W]$, which represents the investor's relative risk-aversion, is greater than zero such that the ICAPM predicts a positive risk-return relation. The elasticity of the marginal utility of wealth with respect to the state variable, U_{WS_1}, can either be positive or negative. This implies that the ratio $\lambda_1 = [-U_{WS_1}/U_W]$ can be positive or negative. If an increase in the state variable affects the optimal consumption negatively, this increase leads to an unfavorable shift in the investment opportunity set.[1] In periods accompanied by high realizations of the state variable the marginal utility of wealth is high. An additional unit of wealth increases the investor's utility relatively strong because in such periods the investor is financially less well off. This in turn implies $U_{WS_1} > 0$. In periods of small realizations of the state variable the marginal utility of wealth is low, since the investor is financially already better off. In contrast, if an increase in the state variable affects the optimal consumption positively then the marginal utility of wealth is high (low) when the state variable takes on small (high) values.

To smooth intertemporal consumption the investor wants to hedge against unfavorable shifts in the investment opportunity set. In general, the investor demands assets that tend to pay off when the marginal utility of wealth is high. If increases in the state variable negatively affect the optimal consumption then the investor demands assets that have high positive covariance with the state variable since these assets tend to pay off when the marginal utility of wealth is high. This in turn implies that the investor sacrifices some expected return for holding assets that have higher association with the state variable. If the optimal consumption is positively affected by increases in the state variable then the investor requires higher expected returns for assets that have higher covariance with the state variable since such assets pay off when the marginal utility of wealth is low.

The role of oil in the ICAPM

In general, we aim at investigating whether the oil price can be considered a state variable that affects the investor's optimal consumption which implies that the investor's marginal utility of wealth depends on the level of the oil price such that $U_{WS_1} \neq 0$. Even if the marginal utility of wealth indeed depends on the oil price, it is unclear ex ante whether it is positively or negatively related to the oil price. If $U_{WS_1} < 0$ ($U_{WS_1} > 0$), then the

[1] An unfavorable shift in the investment opportunity set implies that the risk-return trade-offs of each attainable portfolio changes adversely.

investor is financially well off (less well off) when the oil price is high. While a priori the sign of U_{WS_1} is not determined, there are two competing views that can be derived from the empirical literature.

Numerous studies focus on the impact of oil price shocks on the U.S. economy and stock market. Hamilton (1983, 1985, 2003) conjectures that exogenous oil supply shocks were responsible for economic contractions during the 1970s. In his seminal paper, Kilian (2009) introduces a structural VAR model that facilitates the decomposition of the global oil price into three components. In particular, Kilian (2009) makes a distinction between shocks related to the global supply of crude oil, shocks related to the global demand for industrial goods and shocks related to precautionary demand. In contrast to Hamilton (1983, 1985, 2003), Kilian (2009) argues that even during the 1970s economic contractions were caused by precautionary demand rather than supply shocks, while thereafter the oil price has predominantly been driven by global demand. For the period starting in the mid-1990s, Harris et al. (2009) confirm the finding of Kilian (2009), which is further supported by Hamilton (2008) for the 2000s. Given the methodological framework of Kilian (2009), Kilian and Park (2009) show that the stock market response to oil price shocks strongly depends on the type of the underlying shock. While increasing oil prices caused by precautionary demand induce stock price depreciations, demand-side-related shocks lead to increasing stock prices.[2] Consequently, the elasticity of the investor's marginal utility of wealth with respect to the oil price depends on the type of the shocks that predominantly drive the oil price. In case of the dominance of demand-side-related shocks, the optimal consumption is positively linked to the oil price, which implies that $U_{WS_1} < 0$ and $\lambda_1 > 0$. The investor is well off in periods of high oil prices such that an additional unit of wealth generates a relatively small utility increment. Decreasing oil prices would imply an adverse shift in the investment opportunity set. To hedge against such shifts, the investor increases his demand for stocks that tend to pay off when the oil price is low. In equilibrium, negative oil risk premia should be observed for stocks that negatively covary with the oil price. In contrast to such stocks, the investor requires positive oil risk premia for stocks that have high association with the oil price, since they tend to pay off in periods the investor is already well off. On the other hand, the dominance of precautionary demand shocks and/or supply shocks would imply an inverse relation of the optimal consumption and the oil price. High oil prices would be accompanied by weak stock market performance such that increasing oil prices would imply an unfavorable shift in the investment opportunity set. In particular, the investor would be well off in periods

[2] In particular, Kilian and Park (2009) find that shocks in the aggregate demand induce decreasing stock prices in the short-run, while in the long-run the effect turns positive.

of low oil prices such that $U_{WS_1} > 0$ and $\lambda_1 < 0$.
Given the empirical evidence that the oil price has mainly been determined by aggregate demand-side-related shocks as argued in Kilian (2009) and Kilian and Park (2009), we conjecture that the investor's marginal utility of wealth is decreasing in the oil price such that $\lambda_1 > 0$.

6.2.2 Empirical specification

Following Bali and Engle (2010), we implement a three-step estimation strategy that allows to precisely capture the development of the time-varying variances and dynamic correlations based on the DCC-GARCH model of Engle (2002) and the cross-sectional variation among stock returns in the framework of Zellner's seemingly unrelated regression approach.
Let $r_{m,t}$ be the excess return on the market portfolio, $\mathbf{r}_t = (r_{1,t}, \ldots, r_{n,t})'$ the n-dimensional vector of stock excess returns, and $\mathbf{S}_t = (S_{1,t}, \ldots, S_{k,t})'$ the k-dimensional vector of state variables. We define $\mathbf{y}_t = (\mathbf{r}_t', r_{m,t}, \mathbf{S}_t')'$ with $\mathbf{E}[\mathbf{y}_t|\mathcal{F}_{t-1}] = \boldsymbol{\mu}_t$, where $\mathcal{F}_{t-1} = \sigma(\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \ldots)$ is the filtration generated by the information available through time $t-1$. Given $\mathcal{F}_{t-1}$, the residual vector $\boldsymbol{\varepsilon}_t = \mathbf{y}_t - \boldsymbol{\mu}_t$ is assumed to be normally distributed with $\mathbf{Var} = [\boldsymbol{\varepsilon}_t|\mathcal{F}_{t-1}] = \mathbf{H}_t$, where $\mathbf{H}_t$ can be decomposed into $\mathbf{H}_t = \mathbf{D}_t\mathbf{R}_t\mathbf{D}_t$. $\mathbf{R}_t$ is the correlation matrix of $\boldsymbol{\varepsilon}_t$, while $\mathbf{D}_t$ contains the standard deviations $h_{i,t}^{1/2}$ of the individual residuals $\varepsilon_{i,t}$ on the main diagonal. All off-diagonal elements of $\mathbf{D}_t$ equal zero. Finally, the vector of standardized residuals is given by $\mathbf{z}_t = \mathbf{D}_t^{-1}\boldsymbol{\varepsilon}_t$. We implement the following three-step estimation strategy as suggested in Bali and Engle (2010).
In the first step, for each component of $\mathbf{y}_t$ we individually model the conditional mean and the conditional variance as a univariate AR(1)-GARCH(1,1) process. The conditional mean is given by

$$y_{i,t} = \kappa_i + \theta_i y_{i,t-1} + \varepsilon_{i,t}, \tag{6.2}$$

where κ_i is the intercept and θ_i the autoregressive parameter. The equation of the conditional variance is given by

$$h_{i,t} = \omega_i + \alpha_i \varepsilon_{i,t-1}^2 + \beta_i h_{i,t-1}, \tag{6.3}$$

where ω_i, α_i, and β_i are constants with $\omega_i, \alpha_i > 0$, $\beta_i \geq 0$, and $\alpha_i + \beta_i < 1$. The models are estimated using the quasi-maximum-likelihood method of Bollerslev and Wooldridge (1992). For each element of $\mathbf{y}_t$ we construct the conditional variances $h_{i,t}$ and the standardized residuals $z_{i,t} = \varepsilon_{i,t}/h_{i,t}^{1/2}$.

In the second step, we estimate the elements of the conditional correlation matrix $\mathbf{R}_t$ based on the first-step standardized residuals. Following Bali and Engle (2010), we estimate bivariate specifications which yield the correlation dynamics between each stock and the market portfolio, the oil price, and the state variables. As in Engle (2002), $\mathbf{R}_t$ is decomposed as $\mathbf{R}_t = diag\{\mathbf{Q}_t\}^{-1/2}\mathbf{Q}_t diag\{\mathbf{Q}_t\}^{-1/2}$, where

$$\mathbf{Q}_t = (1 - a - b)\mathbf{R}_0 + a\mathbf{z}_{t-1}\mathbf{z}'_{t-1} + b\mathbf{Q}_{t-1}, \tag{6.4}$$

with $a > 0$, $b \geq 0$, and $a + b < 1$. $\mathbf{R}_0 = T^{-1}\sum_{t=1}^{T} \mathbf{z}_t\mathbf{z}'_t$ is the empirical covariance matrix of $\mathbf{z}_t$. The DCC-parameters a and b are allowed to take on different values in each of the bivariate specifications. Given this estimation strategy, Bali and Engle (2010) point out that the resulting correlation matrix for the full system is not enforced to be positive definite, however, each correlation is bounded to the interval $[-1, 1]$. Based on the results of the first and the second step we compute the conditional covariance between the returns on stock i and the market portfolio, $h_{i,m,t}$, and the conditional covariance between the return on stock i and changes in the state variable $S_{j,t}$, $h_{i,\Delta S_j,t}$.

In the third estimation step, we use the estimated conditional covariances to explain variations in stock excess returns. Following Bali and Engle (2010), the empirical formulation of Equation (6.1) with a single state variable $S_{1,t}$ can be written as

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,\Delta S_1,t} + u_{i,t}, \tag{6.5}$$

where $r_{i,t}$ is the excess return on stock i at day t and $h_{i,m,t}$ is the time-t conditional covariance between the return on stock i and the return on the market portfolio obtained from the second-step estimation. $h_{i,\Delta S_1,t}$ is the conditional covariance between the return on stock i and the change in state variable S_1. c_i are constant intercepts that are allowed to vary across each stock. In Equation (6.5), expected returns are replaced by ex-post returns, while expected covariances are replaced by covariances estimated on the basis of the information available through time $t - 1$. According to Equation (6.1), the stock-specific intercepts equal zero, while λ_0 is given by $[-U_{WW}W/U_W]$ and represents the investor's relative risk aversion coefficient. As implied by theory, we restrict λ_0 to be identical across all stocks included in the panel. Similarly, Equation (6.1) requires that $\lambda_1 = [-U_{WS_1}/U_W]$ is identical across all stocks. Following Bali and Engle (2010) the panel is estimated using Zellner's seemingly unrelated regression approach that yields estimated standard errors which are robust to heteroskedasticity, autocorrelation, and contemporaneous correlation in the residuals.

6.3 Data

Our sample covers the 10-year period from June 15, 2001 to June 14, 2011, yielding a total of 2502 observations at the daily frequency.

6.3.1 Stock and oil data

While Bali and Engle (2010) predominantly focus their analysis on DJIA 30 companies, we extend the number of companies contained in our analysis considerably. Based on the composition of the S&P 500 index of June 14, 2011, we include all companies for which data are available for the whole sample period, yielding a total of 417 companies. According to the Global Industry Classification Standard (GICS) the companies belong to the 10 industries: energy, materials, industrials, consumer discretionary, consumer staples, health care, financials, information technology, telecommunication services, and utilities. All company data are taken from DataStream. Table 6.1 lists the industries and the number of companies that belong to the respective industries. Given the large number of companies, we do not report descriptive statistics due to space considerations. We compute continuously compounded stock returns and subtract the risk free rate given by the one-month Treasury bill rate transformed to daily frequency as obtained from the data library of Kenneth R. French to construct excess returns.

Insert Table 6.1 about here

For the market portfolio we use the CRSP value-weighted NYSE/AMEX/NASDAQ index which is considered the best proxy for the "market" in the previous literature. Returns are obtained from the data library of Kenneth R. French. Following Kilian and Vega (2011) and Conrad et al. (2012) the West Texas Intermediate (WTI) crude oil spot price is used to construct continuously compounded oil returns. The data are taken from the FRED database at the Federal Reserve Bank of St. Louis.

6.3.2 Macroeconomic risk and financial factors

Given the daily frequency in our empirical analysis we have to restrict macroeconomic factors to variables available at that frequency. Following Bali and Engle (2010), we opt for the default spread computed as the difference between the BAA-rated and AAA-rated corporate bond yields. In addition, we include the term spread computed as the difference between the yields on the 10-year Treasury bond and the 3-month Treasury bill. All macroeconomic variables are taken from the FRED database.

Besides macroeconomic risk factors, we include the Fama-French size (SMB) and book-to-market (HML) factors. While the SMB factor is computed as the difference between the average return on three small portfolios and the average return on three big portfolios, the HML factor is given by the difference between the average return on two value portfolios and the average return on two growth portfolios (see also Fama and French (1992, 1993)). The financial factors are obtained from the data library of Kenneth R. French.

6.4 Empirical results

In this section we summarize the empirical results. First, we present the industry-specific results without accounting for intertemporal hedging demand. Second, we analyze whether investors require a risk premium for the conditional covariation with the oil price. Further, we check whether the results are robust with respect to additionally controlling for macroeconomic risk factors, financial factors, and lagged returns. Finally, we analyze whether conditional covariation with financial and macro factors induce additional risk premia and whether the inclusion of these factors affects the results on the risk-return relation.

6.4.1 Baseline specification

In the reference specification, stock returns are exclusively explained by the conditional covariance between stock and market returns. That is, we do not control for any other factor nor do we account for intertemporal hedging demand. We estimate conditional covariances by Equations (6.2)-(6.4). Each stock that violates the conditions $\alpha_i > 0$, $\beta_i \geq 0$, and $\alpha_i + \beta_i < 1$ in the GARCH(1,1)-specification is removed from the sample. The numbers of stocks of the corresponding industries included in the regression are summarized in Table 6.1. For each pair of variables that does not satisfy the conditions $a > 0$, $b \geq 0$, and $a + b < 1$ in the DCC-specification, we estimate the conditional correlation making use of the CCC-GARCH model of Bollerslev (1990). The panel regression of the third estimation step is given by

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + u_{i,t}, \tag{6.6}$$

where $h_{i,m,t}$ is the conditional covariance between the excess returns on stock i and the market portfolio. The estimation results are summarized in Table 6.2 under Model 1.

Insert Table 6.2 about here

The table shows that the estimated relative risk aversion coefficient equals 2.388 and is significant at the 1% level. The Wald-statistic referring to the null hypothesis that all intercepts are zero cannot be rejected at any reasonable level of significance which in combination with the positive estimate of λ_0 implies the validity of Merton's (1973) ICAPM for our sample period. The result confirms a positive relation between expected returns and market risk. The investor charges a positive risk premium for stocks that have higher association with the market. He is willing to sacrifice some expected return for holding stocks which are negatively correlated with the market and tend to pay off when the market performance is weak. The estimate of the relative risk aversion coefficient is economically reasonable and confirms the previous findings for earlier periods as reported in Scruggs (1998), Guo and Whitelaw (2006), and Bali and Engle (2010) among many others.

To assess the economic significance of our estimation results we compute the expected dynamic market risk premium. Implied by Equations (6.1) and (6.6), this premium is given by the product of the relative risk aversion coefficient λ_0 and the time-varying covariance between the returns on the stocks and on the market, $h_{i,m,t}$. Figure 6.1 presents the expected daily annualized market risk premia for each industry computed as the daily average of the estimated risk premia over all stocks that belong to the respective industry.[3]

Insert Figure 6.1 about here

First, the graphs reveal that the expected risk premia across the different industries evolve very similarly. During the recession 2000/2001 (first shaded area) the risk premia realize a first peak in each of the non-energy industries. It is not surprising that the highest expected premium is observed for the information technology industry as this recession is closely linked to the collapse of the dot-com bubble. After significantly decreasing required market risk premia during the economic recovery of 2002, the stock market correction of July 2002 has led to higher expected risk premia, whereas the period 2003-2007 is characterized by extremely low expected risk premia in each industry. With the beginning of the recession 2007-2009 (second shaded area) we only observe a modest upswing in the premia. According to the logic of the ICAPM, the sharp price declines across all industries on September 15, 2008 induced by the bankruptcy filing of Lehman Brothers can be rationalized by exploding market risk premia. This in turn reflects the investor's reluctance to hold risky assets during this period of high uncertainty at all. The highest required premia are observed for financial and energy companies. Towards the end of the recession 2007-2009 the expected risk premia strongly decrease. However, even for

[3]The annualized risk premia are obtained by multiplying the product of the risk aversion coefficient and the conditional covariance by 252.

the phase of economic recovery after the recession the graphs reveal higher required risk premia compared to the period 2003-2007. Second, despite the similar evolution of the expected premia, the graphs reveal significant differences in the magnitude of the premia across individual industries. Table 6.3 shows the average and the standard deviation of the expected industry-specific risk premia for the periods identified above.

Insert Table 6.3 about here

In the period between the two recessions (referred to as Expansion period I in the table) investors require the highest risk premium for holding stocks of the industries information technology (8.69%), materials (6.55%), and consumer discretionary (6.31%) implying that during the rather calm period these stocks exhibit the highest association with the market. Lower risk is observed for stocks of the industries consumer staples (3.65%), utilities (4.21%), and health care (4.85%). For the recession period 2007-2009 the increase in the premia predicted by the model is most pronounced in the industries financials and energy. Interestingly, in the period after the recession the expected risk premia for consumer staples (4.33%), utilities (4.52%), and health care (5.32%) have almost reached the pre-recession levels, while the required premia for financial and energy stocks still are considerably higher than before the recession. In conclusion, investors expect that consumer staples, utility, and health care stocks tend to (also) pay off when market performance is weak and such serve as a hedge against the market.

6.4.2 Oil price risk premium

We now extend the baseline specification by the conditional covariation between stock and oil returns. Again, the conditional covariances are estimated in the first two steps described by Equations (6.2)-(6.4). The third-step panel of Equation (6.5) is given by

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,o,t} + u_{i,t}, \tag{6.7}$$

where λ_1 is the common slope coefficient on the conditional covariance between stock and oil returns, $h_{i,o,t}$. The estimation results are summarized in Table 6.2 under Model 2. The estimated coefficient λ_0 equals 1.932 and is significant at the 1% level. This result confirms the finding of a positive relation between expected return and market risk found in the previous section. However, including the conditional covariance between oil and stock returns as an additional explanatory variable reveals that the specification without intertemporal hedging demand suffers from omitted variable bias. The relative risk aversion coefficient is considerably overestimated which in turn generates too high expected market risk premia. Nevertheless, in line with our previous results the Wald-statistic implies that the null hypothesis that all intercepts equal zero cannot be rejected, which leads to the conclusion that the ICAPM holds given the positive estimate of the relative market risk aversion coefficient.

The estimate of λ_1 is 1.398 and significant at the 1% level such that the oil price can be considered a priced state variable in the ICAPM of Merton (1973) which is consistent with the results in Cifarelli and Paladino (2010, 2012). According to Equation (6.1), $\lambda_1 > 0$ implies that the marginal utility of wealth and the oil price are inversely related, that is, $U_{WS_1} < 0$. In periods of high oil prices the investor values an additional unit of wealth lesser than in periods of low oil prices which reflects that high oil prices are realized in periods in which the investor is financially well off, while he is less well off during periods of low oil prices. Put differently, the oil price is positively related to optimal consumption opportunities which indicates that decreasing oil prices reflect an adverse shift in the investment opportunity set. To hedge against such adverse shifts the investor expands his demand for stocks that tend to pay off when the oil price is low. In general, stocks that are negatively correlated with the oil price tend to pay off in such periods. Hence, the investor sacrifices some expected return for holding such stocks. On the other hand, stocks that have high association with the oil price tend to pay off when the oil price is high and the marginal utility of wealth is low such that the investor requires a positive risk premium for holding stocks that positively covary with the oil price. In equilibrium, such stocks generate higher expected returns.

The estimated elasticity of the marginal utility of wealth with respect to the oil price, $U_{WS_1} < 0$, can be interpreted as further evidence in favor of Barsky and Kilian (2004) and Kilian and Park (2009). The authors argue that the oil price is mainly driven by aggregate international demand, and hence, by the same forces as stock prices. According to this argument periods of high (low) oil prices are related to strong (weak) stock market performance reflecting low (high) marginal utility of wealth which in turn implies $\lambda_1 > 0$. In contrast, $U_{WS_1} < 0$ clearly contradicts the view that positive oil price shocks and/or precautionary demand shocks have predominantly been the drivers of the oil price within our sample period. In this case, high oil prices should be accompanied by weak stock market performance which would imply that with increasing oil prices the investor expands his demand for stocks that positively covary with the oil price.

We now turn to the estimated industry-specific oil risk premia. To provide a first impression of the evolution of the conditional correlation between oil price changes and industry-specific returns we refer to Figure 6.2.

Insert Figure 6.2 about here

The figure shows the dynamic conditional correlations between oil and stock returns averaged over all stocks that belong to the respective industries. The graphs reveal extreme variability in the dynamic correlations in each industry. In the energy industry the correlation is positive throughout the whole sample, while in the industries materials and utilities positive correlations are observed for the majority of the sample period. In strong contrast, for the remaining seven industries we observe prolonged periods of positive and negative correlations which is in line with the finding on the dynamic correlation between returns on the CRSP portfolio and oil returns as reported by Conrad et al. (2012). Consequently, restricting the dynamic correlations to a constant as suggested in Scruggs (1998) or Cifarelli and Paladino (2010, 2012) could generate covariances that exhibit the wrong sign during certain periods, and hence, induce misleading estimates of the risk premia. In particular, each specification with constant conditional correlation is able to produce either negative or positive risk premia depending on the unconditional correlation between the corresponding series. The heavily fluctuating dynamic correlations shown in Figure 6.2 clearly advocate the usage of models that explicitly capture the dynamic structure in the correlations. Figure 6.3 presents the daily annualized oil risk premia for each industry computed as the daily average of the estimated risk premia over all stocks that belong to the respective industry.

Insert Figure 6.3 about here

The graphs clearly exhibit considerable differences in the estimated oil risk premia across the individual industries. Implied by negative dynamic correlations in the non-energy industries in combination with $\lambda_1 > 0$, the estimated oil risk premia in the expansion period between the recessions of 2000/2001 and 2007-2009 is negative for the majority of these industries which reflects the hedging demand for such stocks. This in turn shows that investors sacrifice some expected return for holding stocks that tend to pay off in periods of low oil prices and high marginal utility of wealth. On the other hand, positive correlation especially in the energy industry, but also in the industries materials and utilities, reduces the hedging demand for such stocks which leads to positive required oil risk premia that generate higher expected returns. This mechanism explains the more pronounced negative risk premia in the non-energy industries during the first half of the recession 2007-2009. After a first increase in the recession period the oil price starts to decline considerably. The investor wants to hedge against such an unfavorable shift and increases his hedging demand for stocks that exhibit negative correlation with the oil price which induces the downswing in the estimated oil risk premia before the bankruptcy filing of Lehman Brothers in September 2008. However, in the aftermath of the bankruptcy filing the picture changes dramatically. The price of all stocks as well as the oil price begin to decrease massively, which induces the conditional oil-stock correlations across all industries to turn positive. In particular, this implies that stocks which the investor has used to hedge against adverse shifts in the investment opportunity set now exhibit high association with the oil price such that the stocks do not serve as a hedging vehicle anymore. As a consequence the ICAPM implies that the investor requires high positive risk premia to hold any risky asset. This effect is even magnified by extremely high conditional variances reflecting the uncertainty in the oil and the stock market during this period. Towards the end of the recession 2007-2009 we observe strong downswings in the required oil risk premia across all industries which can mainly be traced back to the considerable contraction in the conditional volatilities. During the economic recovery in the second half of 2009 the correlations between oil and industry-specific stock returns remain high due to simultaneously raising oil and stock prices. Hence, even during this period of economic expansion the investor requires positive risk premia for non-energy stocks.

Finally, we provide summary statistics of the industry-specific oil risk premia for the periods defined above. The corresponding results are reported in Table 6.4.

Insert Table 6.4 about here

As in Table 6.3, Panel A shows the respective market risk premia across the individual industries based on the estimates of Equation (6.7). The magnitude of these premia is about 81% of those obtained in the previous section which has to be traced back to the biased estimate of λ_0 because of the omission of $h_{i,o,t}$ in the baseline specification. However, the interpretation of the previous section is still valid and our model produces economically reasonable and statistically significant market risk premia. Panel B summarizes the expected industry-specific oil risk premia. We identify the highest required risk premia for oil companies for each sample, while for the full sample we observe positive oil risk premia for each industry caused by the high positive correlation during the recession 2007-2009. Most importantly, the table shows that for the industries industrials, consumer staples, consumer discretionary, health care, financials, and information technology the negative average of the estimated risk premia during the first expansion period supports our previous arguments. Note that the significant downswings in the oil risk premia of non-energy stocks in the early phase of the recession 2007-2009 are not accounted for in the numbers referring to Expansion period I. Extending this period until the bankruptcy filing of Lehman Brothers, however, results in more pronounced negative oil risk premia for all industries apart from energy, materials, and utilities. Finally, a direct comparison of the market and oil risk premia across the industries shows that the investor requires similar oil and market risk premia for energy stocks, while for non-energy stocks the market risk premia considerably exceed the oil risk premia. Again, this finding has to be traced back to the high correlation between oil and energy stock returns. In conclusion, not only the industry-specific market risk premia but also the oil risk premia is economically reasonable.

6.4.3 Robustness

In this section we control for factors that the previous literature has identified to affect company stock returns.

Financial factors

In our first robustness check, we extend Equation (6.7) by the size (SMB) and the book-to-market factor (HML) introduced by Fama and French (1992, 1993). Based on the conditional covariances obtained from the first two steps, we estimate the following system of equations

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,o,t} + \lambda_2 \cdot SMB_{t-1} + \lambda_3 \cdot HML_{t-1} + u_{i,t}. \tag{6.8}$$

The estimation results are summarized in Table 6.2 under Model 3. The table shows that both factors negatively enter the specification significant at the 5% level. Most importantly, Table 6.2 shows that the coefficient estimates of λ_0 and λ_1 are robust with respect to the inclusion of the Fama-French factors which confirms our findings on the relation between market and oil risk and expected returns. Moreover, the outcome of the Wald-statistic is in favor of the null hypothesis that all intercepts equal zero.

Macroeconomic risk factors

Following Engle and Bali (2010) we directly include unexpected news in macroeconomic risk factors into Equation (6.7). The panel regression of the third estimation step is given by

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,o,t} + \lambda_2 \cdot \Delta DEF_{t-1} + \lambda_3 \cdot \Delta TERM_{t-1} + u_{i,t}, \tag{6.9}$$

where ΔDEF and $\Delta TERM$ denote the change in the default spread and the change in the term spread, respectively. Model 4 reported in Table 6.2 shows the estimation results. Interestingly, neither unexpected news in the term spread nor unexpected news in the default spread affect the stock returns. As in the previous robustness check the estimates and significance levels of λ_0 and λ_1 are consistent with those obtained from Equation (6.7).

The Wald-statistic implies that the null hypothesis that all intercepts equal zero cannot be rejected at any level of significance.

Lagged market and oil returns

Recently, Narayan and Sharma (2011) find that stock returns of specific industries are negatively affected by lagged oil price changes in the post-2000 period. To control for this effect in the ICAPM framework, we directly include lagged oil returns into Equation (6.7). In addition, we also include lagged market returns. For both variables the slope coefficient is restricted to be identical across all stocks that belong to the same industry. The estimated panel model is given by

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,o,t} + \lambda_{2,j} \cdot r_{m,t-1} + \lambda_{3,j} \cdot r_{o,t-1} + u_{i,t}, \tag{6.10}$$

where $\lambda_{2,j}$ and $\lambda_{3,j}$ are the industry-specific slope coefficients on the lagged market and oil returns, respectively. The results are presented in Table 6.2 under Model 5. The table shows that the lagged market return negatively affects current stock returns across all industries highly significant. We reveal a negative impact of lagged oil returns on expected stock returns of the industries industrials, consumer discretionary, consumer staples, and health care, which is broadly in line with Narayan and Sharma (2011). The table confirms the magnitude of the estimated coefficients λ_0 and λ_1. Both coefficients are significant at the 1% level.

Hedging demand induced by financial factors

In the first two robustness checks we have investigated whether the estimates on λ_0 and/or λ_1 are affected by the incorporation of financial factors or macroeconomic risk factors. Following Bali and Engle (2010), we now reveal whether conditional covariation with these factors induce additional risk premia and whether our previous results are affected. Our third-step panel regression is given by

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,o,t} + \lambda_2 \cdot h_{i,SMB,t} + \lambda_3 \cdot h_{i,HML,t} + u_{i,t}, \tag{6.11}$$

where $h_{i,SMB,t}$ ($h_{i,HML,t}$) is the time $t-1$-expected conditional covariance between the returns on stock i and the size (book-to-market) factor. The estimation results are summarized in Table 6.2 under Model 6. First, the estimates of λ_0 and λ_1 are in line with those of the previous section and the Wald-statistic is in favor of the null-hypothesis that all intercepts equal zero. Hence, our conclusions on the relation between risk and expected return are robust with respect to the inclusion of risk premia induced by conditional covariation with the Fama-French factors. Second, we observe a positive and significant (insignificant) coefficient on the covariance between stock returns and the SMB (HML) factor. The result concerning the SMB factor is economically plausible. SMB is defined as the difference between the average return on small portfolios and big portfolios. On average small companies are hit stronger by economic crises than larger companies which can be traced back to the lower degree of diversification of small companies. Hence, the SMB factor should tend to have smaller values during crises such that optimal consumption and the SMB factor are positively linked. This implies that a positive change in the SMB factor indicates a favorable shift in the investment opportunity set. The investor reduces his hedging demand for stocks that positively covary with the SMB factor which in turn induces the expected excess return to increase.

Hedging demand induced by macroeconomic risk factors

In line with the last robustness check we now augment Equation (6.7) by the conditional covariance between stock returns and macroeconomic risk factors. The corresponding panel regression is given by

$$r_{i,t} = c_i + \lambda_0 \cdot h_{i,m,t} + \lambda_1 \cdot h_{i,o,t} + \lambda_2 \cdot h_{i,\Delta DEF,t} + \lambda_3 \cdot h_{i,\Delta TERM,t} + u_{i,t}, \tag{6.12}$$

where $h_{i,\Delta DEF,t}$ ($h_{i,\Delta TERM,t}$) is the time $t-1$-expected conditional covariance between the excess return on stock i and the change in the default (term) spread. The estimation results are shown in the last column of Table 6.2. The table shows that in line with our previous robustness checks the findings on the risk-return relation discussed above are confirmed. The coefficient estimates and significance levels of λ_0 and λ_1 are very similar to those of Section 6.4.2. We observe a negative (positive) and significant coefficient on the covariance between stock returns and changes in the default (term) spread. Both results are economically reasonable. The default spread measures the difference between corporate bond yields on BAA- and AAA-rated stocks. This difference increases during recessions such that the investor's marginal utility of wealth is increasing in the default spread, that is, $U_{WS_2} > 0$ and $\lambda_2 < 0$. A positive change in the default spread implies an adverse shift in the investment opportunity set. In equilibrium, the investor is willing to sacrifice some expected return for holding stocks that tend to pay off when the default spread is large and the marginal utility of wealth is high. In contrast to the default spread, the term spread is most pronounced during economic expansions and starts to decline prior to recessions. Consequently, $U_{WS_3} < 0$ and $\lambda_3 > 0$. A decrease in the term spread hints at a deterioration of economic performance, and hence, implies an unfavorable shift in the investment opportunity set which leads to increasing expected returns of stocks that have high association with the term spread.

6.5 Conclusion

In this paper we use Merton's (1973) ICAPM to empirically assess the link between the U.S. stock market and the oil price. In equilibrium, the investor is compensated for bearing the market risk and for the risk of adverse changes in state variables that are related to the investor's future consumption. Following Bali and Engle (2010) we implement a three-step estimation strategy. First, using the DCC-GARCH model of Engle (2002) we estimate conditional variances of and covariances between each stock of the S&P 500 index, the market portfolio, and a set of potential state variables including the oil price, financial factors, and macro factors. Then, we explain stock returns by the estimated conditional covariances in a large panel making use of Zellner's seemingly unrelated regression approach. We focus on the 10-year period June 15, 2001 to June 14, 2011 and provide a thorough investigation of the dynamics of the estimated industry-specific market and oil risk premia.

Our study reveals that the ICAPM holds for the analyzed period and confirms the existence of a positive relation between risk and return which is in line with previous studies such as Scruggs (1998) or Bali and Engle (2010). In particular, our findings confirm Merton's (1973) prediction that the investor accepts lower returns for stocks that less strong covary with the market and tend to (also) pay off when market performance is weak. The evolution of the expected market risk premia across the individual industries is pretty similar over time, whereas the magnitude of the premia differs strongly. Investors require the highest risk premia for financials and energy stocks. Interestingly, the risk premia for these stocks are still considerably higher after the recession 2007-2009 compared to the period of economic expansion during the recessions 2000/2001 and 2007-2009, while the premia for consumer staples or health care stocks have almost reached their low pre-crisis levels.

In contrast to the relation between market risk and expected stock returns, the link between stock returns and the oil price is unclear ex ante. Our findings indicate that the investor's marginal utility of wealth is negatively related to the oil price. Put differently, if the level of the oil price is high the marginal utility of wealth is low which implies that the investor is financially well off in periods of high oil prices. According to Kilian (2009) and Kilian and Park (2009), the result provides further evidence for the view of Barsky and

Kilian (2004) that the oil price is predominantly driven by shocks induced by unexpected changes in aggregate demand, rather than by precautionary demand shocks or shocks related to the supply-side in our sample period. We find that the risk premia associated with the oil price fluctuate considerably over time and across industries. During the expansion period 2002-2007 investors charge positive risk premia for energy, materials, and utility stocks since the returns of these stocks are positively correlated with the oil price and tend to pay off when the oil price is high. In contrast, during this period investors accept negative risk premia for stocks that are inversely related to the oil price which implies that these stocks are used as a hedge against decreasing oil prices. We show that the picture changes dramatically in the months after the bankruptcy filing of Lehman Brothers in September 2008. Now, extreme upswings in the correlations lead to positive required risk premia for stocks of each industry. Hence, none of the stocks serves as a hedging vehicle in this period anymore.

Finally, we show that our results are robust with respect to the inclusion of hedging demand induced by macroeconomic risk factors and by the financial factors of Fama and French (1992).

Table 6.1: Industries

Industry	number of companies	companies included
S&P 500	417	388
Energy	35	33
Materials	26	25
Industrials	55	51
Consumer Discretionary	64	63
Consumer Staples	37	37
Health Care	42	39
Financials	64	51
Information Technology	58	58
Telecommunication Services	6	5
Utilities	30	26

Notes: The table summarizes the industries included. The second column presents the number of companies that belong to the respective industry in the initial sample. The last column presents the number of companies included in the empirical analysis. Stock return series that do not satisfy the conditions $\alpha > 0$, $\beta \geq 0$, and $\alpha + \beta < 1$ in the first estimation step are removed from the sample.

Table 6.2: Estimation results company-level

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
$h_{i,m,t}$	$2.388^{\star\star\star}$ (0.270)	$1.932^{\star\star\star}$ (0.282)	$1.935^{\star\star\star}$ (0.282)	$1.836^{\star\star\star}$ (0.282)	$1.836^{\star\star\star}$ (0.282)	$2.034^{\star\star\star}$ (0.332)	$1.683^{\star\star\star}$ (0.287)
$h_{i,o,t}$		$1.398^{\star\star\star}$ (0.313)	$1.364^{\star\star\star}$ (0.313)	$1.498^{\star\star\star}$ (0.314)	$1.466^{\star\star\star}$ (0.313)	$1.242^{\star\star\star}$ (0.315)	$1.149^{\star\star\star}$ (0.314)
SMB_{t-1}			$0.044^{\star\star}$ (0.017)				
HML_{t-1}			$-0.041^{\star\star}$ (0.017)				
ΔDEF_{t-1}				-0.002 (0.004)			
$\Delta TERM_{t-1}$				0.000 (0.001)			
$h_{i,SMB,t}$						$3.469^{\star\star\star}$ (1.192)	
$h_{i,HML,t}$						1.395 (0.887)	
$h_{i,\Delta DEF,t}$							$-0.837^{\star\star\star}$ (0.309)
$h_{i,\Delta TERM,t}$							$0.456^{\star\star\star}$ (0.081)
Wald-statistic	218.62 [0.999]	218.58 [0.999]	217.59 [0.999]	217.87 [0.999]	219.92 [0.999]	215.49 [0.999]	216.16 [0.999]

Model 5

	EN	MA	IND	CD	CS	HC	FI	IT	$TELE$	UT
$r_{m,t-1}$	$-0.042^{\star\star\star}$ (0.013)	$-0.057^{\star\star\star}$ (0.010)	$-0.046^{\star\star\star}$ (0.009)	$-0.036^{\star\star\star}$ (0.009)	$-0.050^{\star\star\star}$ (0.008)	$-0.025^{\star\star\star}$ (0.010)	$-0.141^{\star\star\star}$ (0.010)	$-0.077^{\star\star\star}$ (0.010)	$-0.026^{\star\star}$ (0.014)	$-0.046^{\star\star\star}$ (0.011)
$r_{o,t-1}$	0.002 (0.007)	-0.006 (0.005)	$-0.010^{\star\star}$ (0.004)	$-0.011^{\star\star}$ (0.005)	$-0.008^{\star\star}$ (0.004)	$-0.008^{\star}$ (0.005)	-0.004 (0.005)	0.002 (0.005)	-0.007 (0.007)	0.009 (0.006)

Notes: The first row entries of the table show the estimated relative risk aversion coefficients. Model 1 refers to the baseline specification including the conditional covariance between the excess returns on the individual stocks and the market. Column 2 shows the results on the baseline specification augmented by the conditional covariance between the excess return on stock i and changes in the oil price. Models 3 to 7 constitute robustness checks. The Wald statistics correspond to the null-hypothesis that the stock-specific intercepts $c_1, \ldots, c_{388}$ are jointly equal to zero. Numbers in parentheses are standard errors. $\star\star\star, \star\star, \star$ indicate significance at the 1%, 5%, 10%. Numbers in brackets are p-values.

Table 6.3: Industry-specific market risk premia

	Full sample		Expansion period I		Recession period		Expansion period II	
Industry	Mean	Stand. Dev.	Mean	Stand. Dev.	Mean	Stand. Dev.	Mean	Stand. Dev.
Energy	12.30	22.91	6.04	5.01	40.56	47.10	9.86	5.38
Materials	11.96	18.35	6.55	4.39	35.90	36.34	9.35	4.56
Industrials	10.31	13.74	6.14	4.35	28.66	26.10	8.15	4.22
Consumer Discretionary	10.64	14.47	6.31	4.76	30.04	27.32	7.87	4.03
Consumer Staples	5.78	7.84	3.65	2.55	15.71	15.47	4.33	2.00
Health Care	7.11	9.54	4.85	3.53	17.93	19.44	5.32	2.33
Financials	13.29	24.67	5.54	4.48	48.92	46.93	9.22	5.24
Information Technology	11.98	13.06	8.69	6.75	27.51	24.01	8.50	3.66
Telecommunication	8.83	13.99	6.00	6.36	25.07	27.52	4.61	2.20
Utilities	6.66	11.74	4.21	4.69	19.15	24.31	4.52	2.69

Notes: The table shows the predicted annualized market risk premium averaged over all stocks of the individual industries obtained from Model 1. Annualized risk premia are obtained by multiplying the product of the conditional covariances and the relative risk aversion coefficient by 252. The full sample contains all observations from June 15, 2001 to June 14, 2011. Expansion period I covers the phase between the NBER recessions of 2001 and 2007-2009. Hence, the period contains the observations from December 1, 2001 to November 30, 2007. The recession periods are June 15, 2001 to November 30, 2001 and December 1, 2007 to June 30, 2009. Expansion period II coincides with the period after the recession of 2007-2009.

Table 6.4: Industry-specific market and oil risk premia

	Full sample		Expansion period I		Recession period		Expansion period II	
Industry	Mean	Stand. Dev.	Mean	Stand. Dev.	Mean	Stand. Dev.	Mean	Stand. Dev.
Energy	9.95	18.54	4.89	4.06	32.81	38.10	7.98	4.35
Materials	9.68	14.84	5.30	3.55	29.05	29.40	7.56	3.69
Industrials	8.34	11.12	4.97	3.52	23.19	21.12	6.60	3.42
Consumer Discretionary	8.61	11.71	5.11	3.85	24.30	22.10	6.37	3.26
Consumer Staples	4.67	6.35	2.96	2.06	12.71	12.51	3.50	1.61
Health Care	5.76	7.72	3.92	2.86	14.51	15.73	4.30	1.89
Financials	10.75	19.96	4.48	3.63	39.58	37.97	7.46	4.24
Information Technology	9.69	10.57	7.03	5.46	22.26	19.42	6.88	2.96
Telecommunication	7.14	11.31	4.85	5.15	20.28	22.27	3.73	1.78
Utilities	5.39	9.49	3.41	3.79	15.49	19.67	3.66	2.17

Panel B: Oil risk premia in individual industries

	Full sample		Expansion period I		Recession periods		Expansion period II	
Industry	Mean	Stand. Dev.	Mean	Stand. Dev.	Mean	Stand. Dev.	Mean	Stand. Dev.
Energy	9.19	10.45	5.70	4.22	23.51	20.07	8.01	2.91
Materials	3.28	5.77	0.66	3.50	10.68	10.50	5.57	2.55
Industrials	1.57	4.18	-0.20	3.43	5.50	7.97	3.94	2.53
Consumer Discretionary	0.71	3.85	-0.48	3.63	2.82	8.16	2.89	2.18
Consumer Staples	0.28	1.78	-0.48	1.97	1.52	3.49	1.63	0.84
Health Care	0.68	1.92	-0.03	2.72	1.94	3.84	1.97	1.01
Financials	1.17	4.97	-0.43	3.75	3.88	10.70	4.14	2.66
Information Technology	1.49	3.73	-0.03	5.10	4.53	6.86	3.75	2.07
Telecommunication	0.98	3.00	0.03	4.94	3.69	5.98	1.90	1.31
Utilities	1.48	2.57	0.51	3.87	4.04	5.08	2.32	1.32

Notes: Panel A (B) shows the predicted annualized market (oil) risk premium averaged over all stocks of the individual industries obtained from Model 2. Annualized market (oil) risk premia are obtained by multiplying the product of the respective conditional covariances and λ_0 (λ_1) by 252. The full sample contains all observations from June 15, 2001 to June 14, 2011. Expansion period I covers the phase between the NBER recession of 2001 and 2007-2009. Hence, the period contains the observations from December 1, 2001 to November 30, 2007. The recession periods are June 15, 2001 to November 30, 2001 and December 1, 2007 to June 30, 2009. Expansion period II coincides with the period after the recession of 2007-2009.

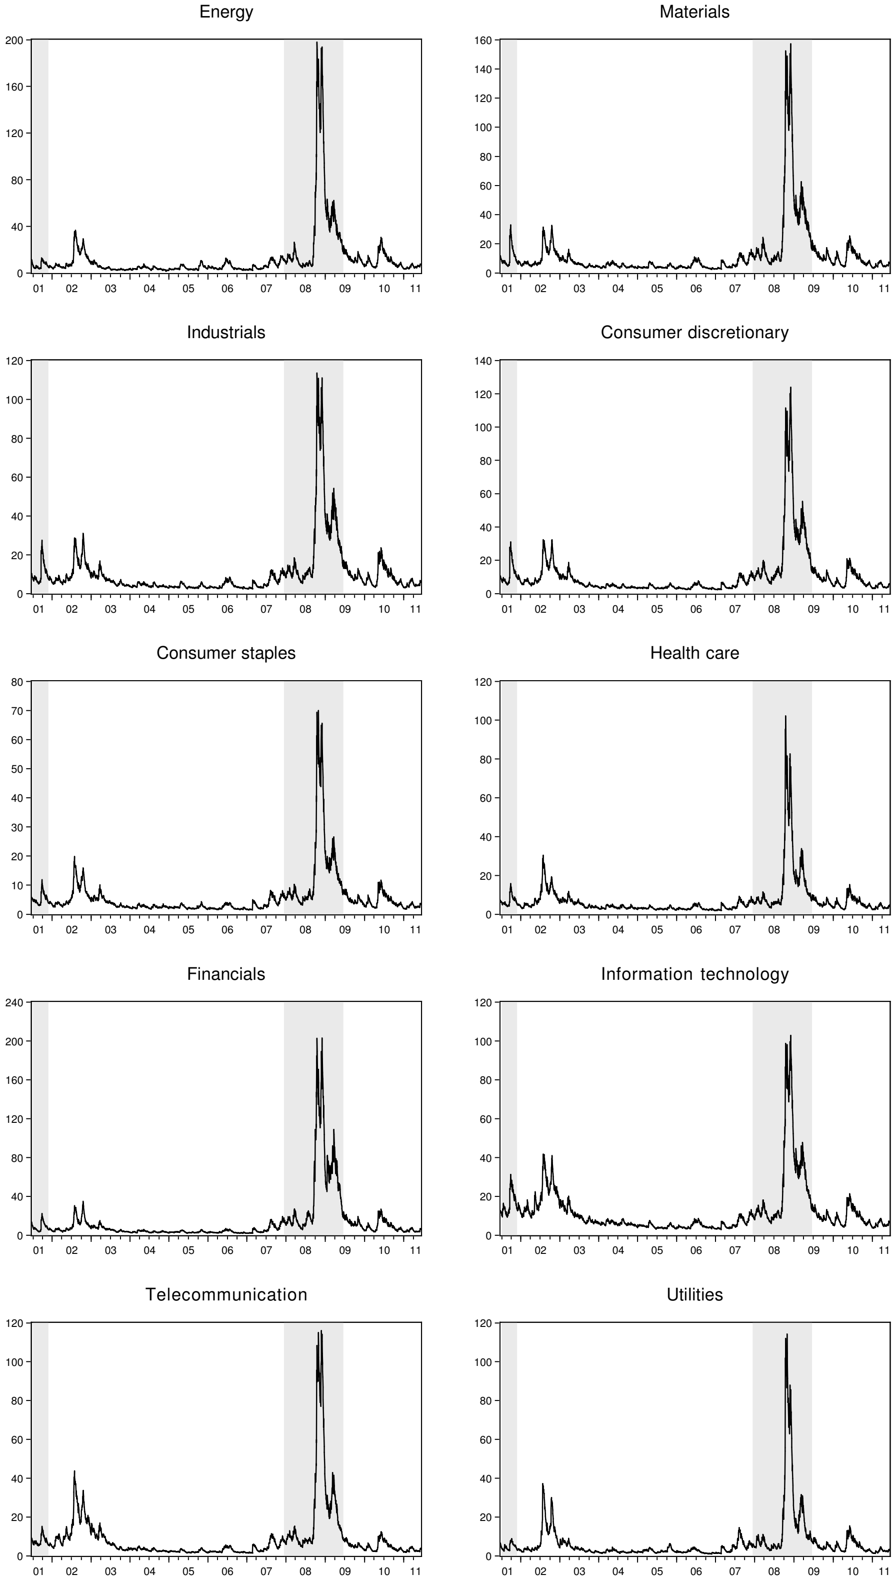

Figure 6.1: Annualized average daily market risk premium for the individual industries.

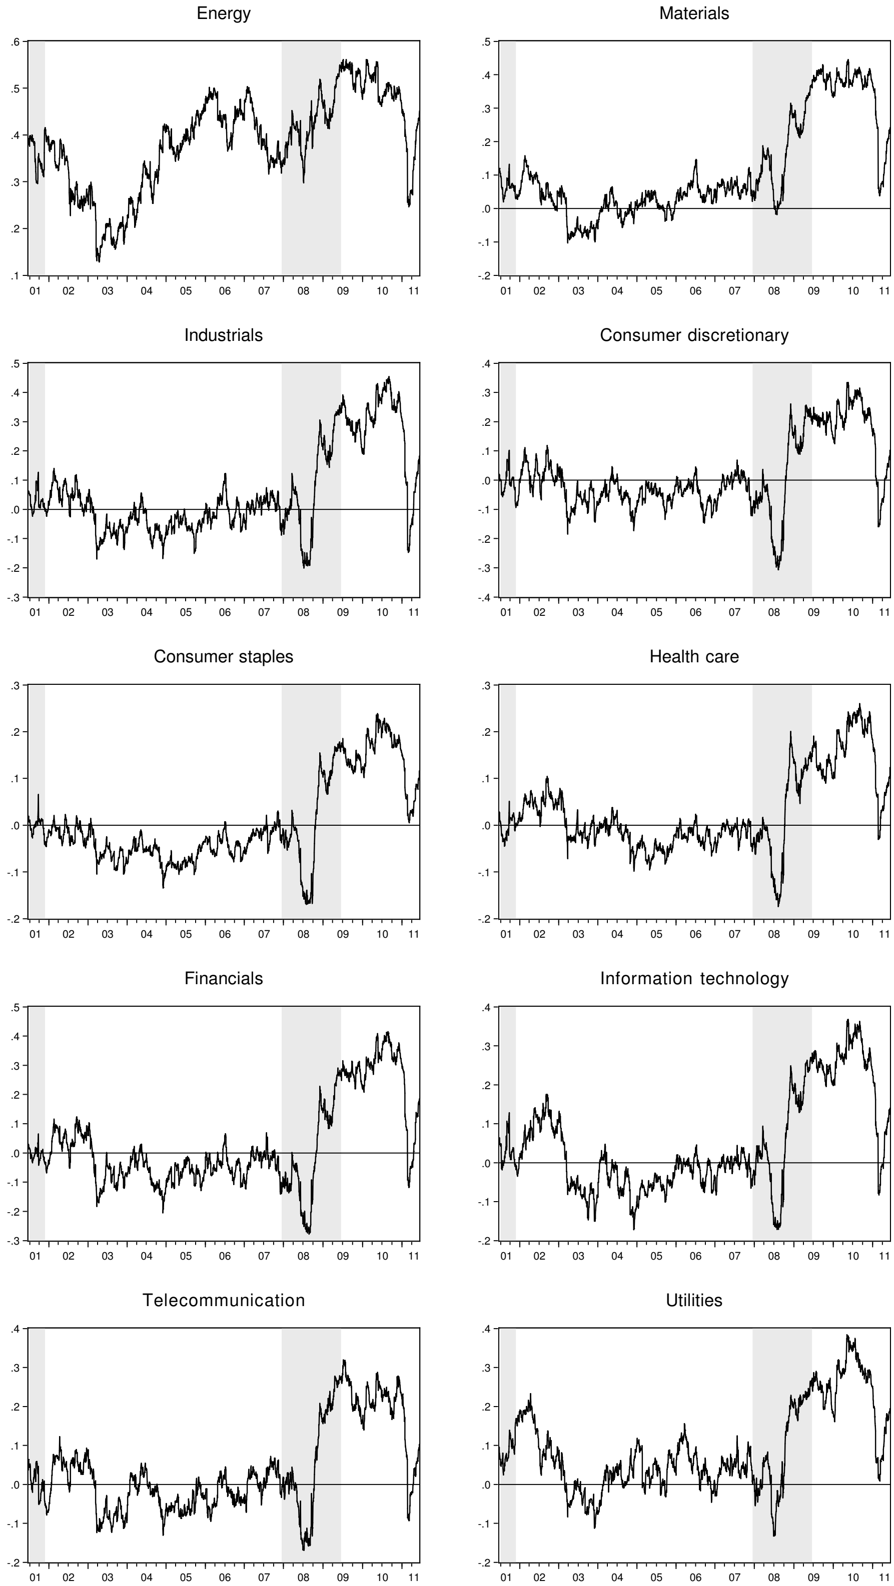

Figure 6.2: Dynamic Conditional Correlations between oil return and industry returns.

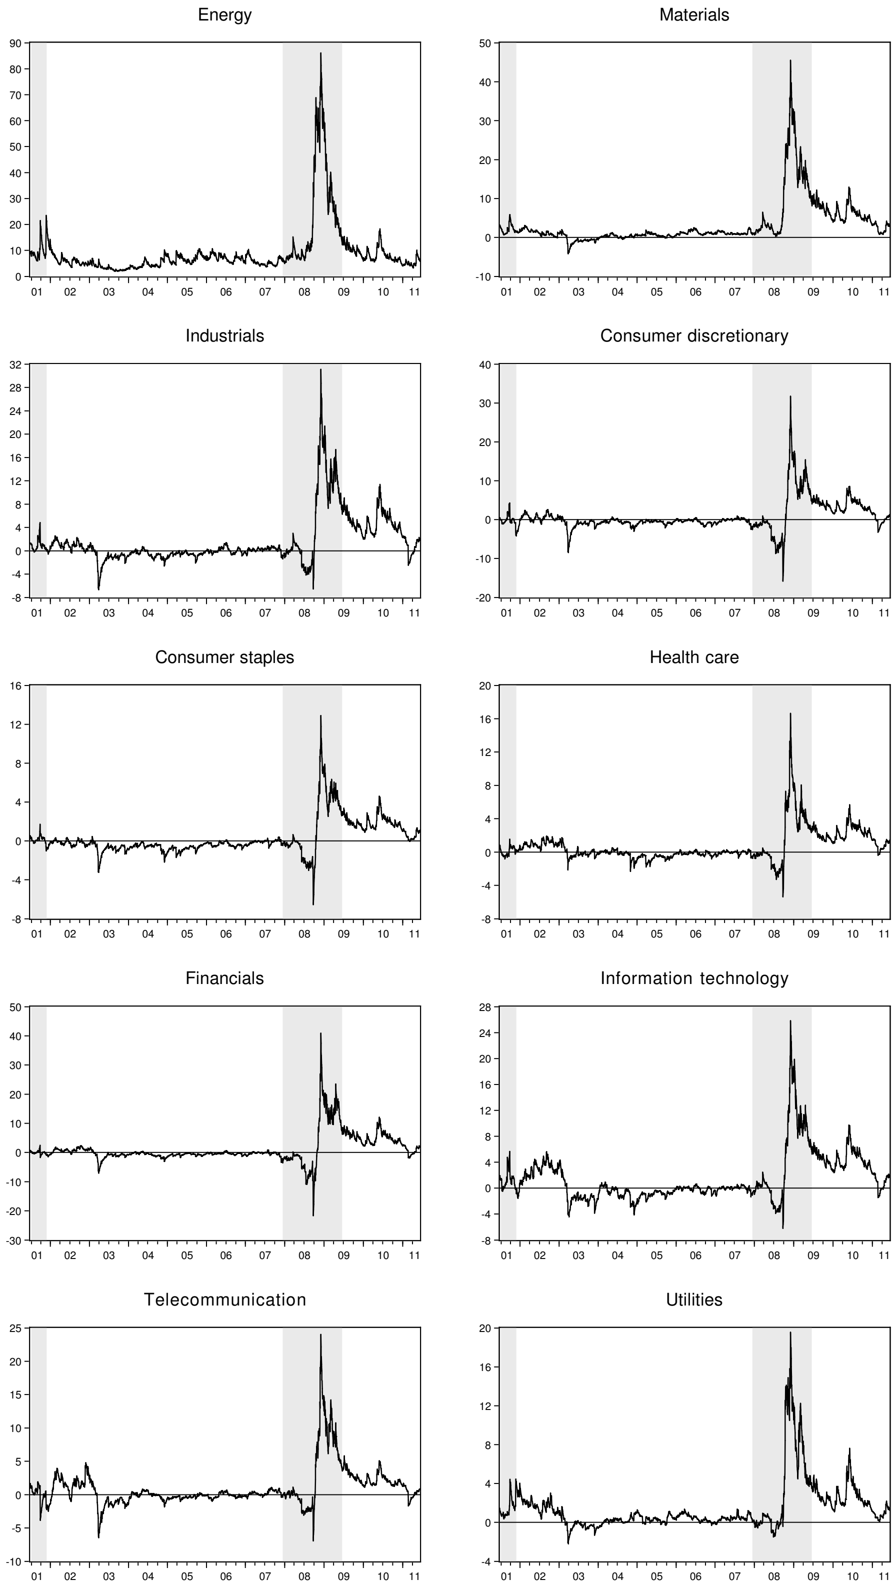

Figure 6.3: Annualized average daily oil risk premium for the individual industries.

Bibliography

[1] Alberola, E., Chevallier, J., Cheze, B., 2008. Price drivers and structural breaks in European carbon prices 2005-07. Energy Policy 36, 787-797.

[2] Alberola, E., Chevallier, J., Cheze, B., 2009a. The EU emissions trading scheme: The effects of industrial production and CO2 emissions on European carbon prices. International Economics 116, 95-128.

[3] Alberola, E., Chevallier, J., Cheze, B., 2009b. Emissions compliances and carbon prices under the EU ETS: A country specific analysis of industrial sectors. Journal of Policy Modeling 31, 446-462.

[4] Andersen, T., 1996. Return volatility and trading volume: An information flow interpretation of stochastic volatility. Journal of Finance 51, 169-204.

[5] Andersen, T., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. Journal of Empirical Finance 4, 115-158.

[6] Andersen, T., Bollerslev, T., 1998. Deutsche mark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. Journal of Finance 53, 219-265.

[7] Andersen, T., Bollerslev, T., Diebold, F., Vega, C., 2003. Micro effects of macro announcements: Real-time price discovery in foreign exchange. American Economic Review 93, 38-62.

[8] Andersson, M., Overby, L.J., Sebestyen, S., 2009. Which news moves the euro area bond market? German Economic Review 10, 1-31.

[9] Baillie, R., Bollerslev, T., Mikkelsen, H., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 74, 3-30.

[10] Baillie, R., Booth, G., Tse, Y., Zabotina, T., 2002. Price discovery and common factor models. Journal of Financial Markets 5, 309-321.

[11] Balduzzi, P., Elton, E., Green, C., 2001. Economic news and bond prices: Evidence from the U.S. treasury market. Journal of Financial and Quantative Analysis 36, 523-543.

[12] Bali, T., Engle, R., 2010. The intertemporal capital asset pricing model with dynamic conditional correlations. Journal of Monetary Economics 57, 377-390.

[13] Barsky, R., Kilian, L., 2004. Oil and the macroeconomy since the 1970s. Journal of Economic Perspectives 18, 115-134.

[14] Bauwens, L., Laurent, S., Rombouts, J., 2006. Multivariate GARCH models: A survey. Journal of Applied Econometrics 21, 79-109.

[15] Benz, E., Hengelbrock, J., 2008. Price discovery and liquidity in the European CO2 futures market. Working Paper. Bonn Graduate School of Economics.

[16] Benz, E., Trück, S. 2009. Modeling the price dynamics of CO_2 emission allowances. Energy Economics 31, 4-15.

[17] Bollerslev, T., 1990. Modeling the coherence in short-run exchange rates: A multivariate generalized ARCH model. Review of Economics and Statistics 72, 498-505.

[18] Bollerslev, T., Wooldridge, J., 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. Econometric Reviews 11, 143-172.

[19] Borak, S., Härdle, W., Trück, S., Weron, R., 2006. Convenience yields for CO_2 emission allowance futures contracts. Discussion paper, Humbolt University of Berlin.

[20] Boyer, M., Filion, D., 2007. Common and fundamental factors in stock returns of Canadian oil and gas companies. Energy Economics 29, 428-453.

[21] Bredin, D., Muckley, C., 2011. An emerging equilibrium in the EU emissions trading scheme. Energy Economics 33, 353-362.

[22] Brennan, M., 1991. The price of convenience and the valuation of commodity contingent claims. In: Lund, D., Oksendahl, B. (Eds), Stochastic Models and Option Values, Applications to Resources, Environment and Investment Problems. North-Holland, Amsterdam.

[23] Brooks, R., Faff, R., McKenzie, M., 2000. A multi-country study of power ARCH models and national stock market returns. Journal of International Money and Finance 19, 377-397.

[24] Bushnell, J., Chong, H., and Mansur, E., 2011. Profiting from regulation: An event study of the European carbon market. Working paper, Iowa State University.

[25] Chan, K., Chan, K., Karolyi, A., 1991. Intraday volatility in the stock index and stock index futures markets. Review of Financial Studies 4, 657-684.

[26] Chevallier, J., 2009. Carbon futures and macroeconomic risk factors: A view from the EU ETS. Energy Economics 31, 614-625.

[27] Chevallier, J., 2010. A note on cointegrating and vector autoregressive relationships between CO_2 allowances spot and futures prices. Economics Bulletin 30, 1564-1584.

[28] Chevallier, J., 2011. A model of carbon price interactions with macroeconomic and energy dynamics. Energy Economics 33, 1295-1312.

[29] Chevallier, J., 2012. Econometric Analysis of Carbon Markets. Springer, New York.

[30] Chevallier, J., Sevi, B., 2011. On the realized volatility of the ECX CO2 emissions 2008 futures contract: Distribution, dynamics and forecasting. Annals of Finance 7, 1-29.

[31] Christodoulakis, G., Satchell, S., 2002. Correlated ARCH (CorrARCH): Modelling the time-varying conditional correlation between financial asset returns. European Journal of Operational Research 139, 351-370.

[32] Cifarelli, G., Paladino, G., 2010. Oil price dynamics and speculation. Energy Economics 32, 363-372.

[33] Cifarelli, G., Paladino, G., 2012. Can oil diversify away the unpriced risk of a portfolio? International Journal of Finance and Economics 17, 73-88.

[34] Colacito, R., Engle, R., Ghysels, E., 2011. A component model for dynamic correlations. Journal of Econometrics 164, 45-59.

[35] Conrad, C., 2010. Non-negativity conditions for the hyperbolic GARCH model. Journal of Econometrics 157, 441-457.

[36] Conrad, C., Haag, B., 2006. Inequality constraints in the fractionally integrated GARCH model. Journal of Financial Econometrics 4, 413-449.

[37] Conrad, C., Karanasos, M., 2010. Negative volatility spillovers in the unrestricted ECCC-GARCH model. Econometric Theory 26, 838-862.

[38] Conrad, C., Karanasos, M., Zeng, N., 2011. Multivariate fractionally integrated APARCH modelling of stock market volatility: A multi-country study. Journal of Empirical Finance 18, 147-159.

[39] Conrad, C., Lamla, M., 2010. The high-frequency response of the EUR-US dollar exchange rate to ECB monetary policy announcements. Journal of Money, Credit and Banking 42, 1391-1417.

[40] Conrad, C., Loch, K., 2011. Anticipating long-term stock market volatility. Working paper. Heidelberg of University.

[41] Conrad, C., Loch, K., Rittler, D. 2012. On the macroeconomic determinants of the long-term oil-stock correlations. Discussion paper, University of Heidelberg.

[42] Conrad, C., Rittler, D., Rotfuß, W., 2012. Modeling and explaining the dynamics of European Union allowance prices at high-frequency. Energy Economics, 34, 316-324.

[43] Convery, F., Ellerman, D., De Perthuis, C., 2008. The European carbon market in action: Lessons from the first trading period. Report 158. MIT Joint Program on the Science and Policy of Global Change.

[44] Copeland, T., 1976. A model of asset trading under the assumption of sequential information arrival. Journal of Finance 31, 1149-1168.

[45] Darrat, A., Rahman, S., Zhong, M., 2003. Intraday trading volume and return volatility of DJIA stocks: A note. Journal of Banking and Finance 27, 2035-2043.

[46] Daskalakis, G., Psychoyios, D., Markellos, R., 2009. Modeling CO2 emission allowance prices and derivatives: Evidence from the European trading scheme. Journal of Banking and Finance 33, 1230-1241.

[47] Davidson, J., 2004. Moment and memory properties of linear conditional heteroscedasticity models, and a new model. Journal of Business and Economic Statistics 22, 16-19.

[48] Demailly, D., Quirion, P. 2008. European emissions tading scheme and competitiveness: A case study on the iron and steel industry. Energy Economics 30, 2009-2027.

[49] Ding, Z., Engle, R., Granger, C., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83-106.

[50] Ding, Z., Granger, C., 1996. Modeling volatility persistence of speculative returns: A new approach. Journal of Econometrics 73, 185-215.

[51] Driesprong, G., Jacobsen, B., Maat, B., 2008. Striking oil: Another puzzle? Journal of Financial Economics 89, 307-327.

[52] Ellerman, D., Buchner, B., 2008. Over-allocation or abatement? A preliminary analysis of the EU-ETS based on the 2005-06 emissions data. Environmental and Resource Economics 41, 267-287.

[53] Ellerman, D., Buchner, B., Carraro, C., 2007. Allocaton in the European Union Emissions Trading Scheme: Rights, Rents and Fairness. Cambrigde University Press, Cambrigde.

[54] Ellerman, D., Convery, F., de Perthuis, C. 2010. Pricing Carbon: The European Union Emissions Tading Scheme. Cambridge University Press, Cambrigde.

[55] Ellerman, D., Joskow, P., 2008. The European Union's emissions trading system in perspective. Pew Center for Global Change.

[56] Engle, R., 2002. Dynamic conditional correlation: A simple class of multivariate GARCH models. Journal of Business and Economic Statistics 20, 339-350.

[57] Engle, R., 2010. Long term skewness and systemic risk. New York University, Volatility Institute, Discussion Paper.

[58] Engle, R., Ghysels, E., and Sohn, B., 2009. Stock market volatility and macroeconomic fundamentals. (Former title: On the economic sources of stock market volatility). Working Paper. New York University.

[59] Engle, R., Granger, C., 1987. Co-integration and error correction: Representation, estimation and testing. Econometrica 55, 251-276.

[60] Engle, R., Lee, G., 1999. A permanent and transitory component model of stock return volatility. In: Engle, R., White, H. (Eds.). Cointegration, Causality and Forecasting: A Festschrift in Honor of Clive W.J. Granger. Oxford University Press.

[61] Engle, R., Rangel, J., 2008. The spline GARCH model for unconditional volatility and its global macroeconomic causes. Review of Financial Studies 21, 1187-1222.

[62] European Parliament and Council, 2003. Directive 2003/87/EC of the European Parliament and of the Council of 13 October 2003. Official Journal of the European Union L 275, 32-46.

[63] European Parliament and Council, 2009. Directive 2009/29/EC of the European Parliament and of the Council of 23 April 2009. Official Journal of the European Union L 140, 63-87.

[64] European Parliament and Council, 2010. Resolution 2010/634/EC of the European Parliament and of the Council of 23 October 2010. Official Journal of the European Union L 279, 34-35.

[65] Ewing, B., and Thompson, M., 2007. Dynamic cyclical comovements of oil prices with industrial production, consumer prices, unemployment, and stock prices. Energy Policy 35, 5535-5540.

[66] Faff, R., Brailsford, T., 1999. Oil price risk and the Australian stock market. Journal of Energy Finance and Development 4, 69-87.

[67] Fama, E., 1970. Efficient capital markets: A review of theory and empirical work. Journal of Finance 25, 383-417.

[68] Fama, E., French, K., 1987. Commodity futures prices: Some evidence on forecast power. Journal of Business 60, 55-73.

[69] Fama, E., French, K., 1992. Cross-section of expected stock returns. Journal of Finance 47, 427-465.

[70] Fama, E., French, K., 1993. Common risk factors in the returns of stocks and bonds. Journal of Financial Economics 33, 3-56.

[71] Fama, E., MacBeth, J., 1973. Risk, return and equilibrium. Journal of Political Economy 81, 607-636.

[72] Filis, G., Degiannakis, S., Floros, C., 2011. Dynamic correlation between stock market and oil prices: The case of oil-importing and oil-exporting countries. International Review of Financial Analysis 20, 152-164.

[73] Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long memory time series models. Journal of Time Series Analysis 4, 221-238.

[74] Ghysels, E., Santa-Clara, P., Valkanov, R., 2005. There is a risk-return trade-off after all. Journal of Financial Economics 76, 509-548.

[75] Ghysels, E., Sinko, A., Valkanov, R., 2007. MIDAS regressions: Further results and new directions. Econometric Reviews 26, 53-90.

[76] Gonzalo, J., Granger, C., 1995. Estimation of common long memory components in cointegrated systems. Journal of Busines and Economic Statistics 13, 27-35.

[77] Granger, C., 1969. Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37, 424-438.

[78] Guo, H., Whitelaw, R., 2006. Uncovering the risk-return relation in the stock market. Journal of Finance 61, 1433-1463.

[79] Hamilton, J., 1983. Oil and the macroeconomy since World War II. Journal of Political Economy 91, 228-248.

[80] Hamilton, J., 1985. Historical causes of postwar oil shocks and recessions. Energy Journal 6, 97-116.

[81] Hamilton, J., 2003. What is an oil shock? Journal of Econometrics 113, 363-398.

[82] Hamilton, J., 2008. Oil and the macroeconomy. In: Durlauf, S., and Blume, L. (Eds.). New Palgrave Dictionary of Economics. 2nd edition. Palgrave McMillan Ltd.

[83] Hamilton, J., 2009. Understanding crude oil prices. Energy Journal 39, 179-206.

[84] Harris, E., Kasman, B., Shapiro, M. and West, K., 2009. Oil and the macroeconomy: Lessons for monetary policy. US Monetary Policy Forum Conference Paper.

[85] Harris, F., McInish, T., Wood, R., 2002. Security price adjustments across exchanges: An investigation of common factor components for Dow stocks. Journal of Financial Markets 5, 277-308.

[86] Hasbrouck, J., 1995. One security, many markets: Determining the contributions to price discovery. Journal of Finance 50, 1175-1199.

[87] Hayashi, F., 2000. Econometrics. Princton University Press, New Jersey.

[88] Huang, R., Masulis, R., Stoll, H., 1996. Energy shocks and financial markets. Journal of Futures Markets 16, 1-27.

[89] Huber, P., 1967. The behavior of the maximum likelihood estimates under nonstandard conditions. Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability 1, 221-233.

[90] Jones, C., Kaul, G., 1996. Oil and the stock markets. Journal of Finance 51, 463-491.

[91] Kang, S., Kang, S., Yoon, S., 2009. Forecasting volatility of crude oil markets. Energy Economics 31, 119-125.

[92] Kettner, C., Koppl, A., Schleicher, S., Thenius G., 2008. Stringency and distribution in the EU Emissions Trading Scheme: First evidence. Climate Policy 8, 41-61.

[93] Kilian, L., 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. American Economic Review 99, 1053-1069.

[94] Kilian, L., Park, C., 2009. The impact of oil price shocks on the U.S. stock market. International Economic Review 50, 1267-1287.

[95] Kilian, L., Vega, C., 2011. Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices. Review of Economics and Statistics 93, 660-671.

[96] Kruger, J., Pizer, W., 2004. Greenhouse gas trading in Europe: The new grand policy experiment. Environment 46, 8-23.

[97] Laurent, S., Peters, J., 2002. G@RCH 2.2: An Ox package for estimating and forecasting various ARCH models. Journal of Economic Surveys 16, 447-485.

[98] Lise, W., Sijm, J., Hobbs, B., 2010. The impact of the EU ETS on prices, profits and emissions in the power sector: Simulation results with the COMPETES EU20 model. Environmental and Resource Economics 47, 23-44.

[99] Lo, A.W., 1991. Long-term memory in stock market prices. Econometrica 59, 1279-1313.

[100] Mansanet-Bataller, M., Chevallier, J., Herve-Mignucci, M., Alberola, E., 2011. EUA and sCER Phase II price drivers: Unveiling the reasons for the existence of the EUA-sCER spread. Energy Policy 39, 1056-1069.

[101] Mansanet-Bataller, M., Pardo, T., 2008. What you should know to trade in CO2 markets. Energies 1, 120-153.

[102] Mansanet-Bataller, M., Pardo, T., 2009. Impact of regulatory announcements on CO2 prices. Journal of Energy Markets 2, 1-33.

[103] Mansanet-Bataller, M., Pardo T., Valor, E., 2007. CO2 prices, energy and weather. Energy Journal 28, 73-92.

[104] Martens, M., Chang, Y., Taylor, S., 2002. A comparison of seasonal adjustment methods when forecasting intraday volatility. Journal of Financial Research 25, 283-299.

[105] Merton, R., 1973. An intertemporal asset pricing model. Econometrica 41, 867-887.

[106] Miller, J., Ratti, R., 2009. Crude oil and stock markets: Stability, instability, and bubbles. Energy Economics 31, 559-568.

[107] Milunovich, G., Joyeux, R., 2010. Market efficiency and price discovery in the EU carbon futures. Applied Financial Economics 20, 803-809.

[108] Nandha, M., Faff, R., 2008. Does oil move equity prices? A global view. Energy Economics 30, 986-997.

[109] Narayan, P., Sharma, S., 2011. New evidence on oil price and firm returns. Journal of Banking & Finance 35, 3253-3262.

[110] Oberndorfer, U., 2009. EU emission allowances and the stock market: Evidence from the electricity industry. Ecological Economics 68, 1116-1126.

[111] Paolella, M., Taschini, L., 2008. An econometric analysis of emission allowances prices. Journal of Banking and Finance 32, 2022-2032.

[112] Petersen, M., 2009. Estimating standard errors in finance panel data: Comparing approaches. Review of Financial Studies 22, 435-480.

[113] PointCarbon, 2008. Carbon 2008, post-2012 is now. Point Carbon's 5th annual conference: Carbon market insights 2008, Copenhagen.

[114] Rittler, D., 2012. Price discovery and volatility spillovers in the European Union emissions trading scheme: A high frequency analysis. Journal of Banking and Finance 36, 774-785.

[115] Rogers, W., 1983. Analyzing Complex Survey Data. Rand Corporation Memorandum, Santa Monica.

[116] Rotfuß, W., 2009. Intraday price formation and volatility in the European emissions trading scheme: A first analysis. Working Paper, ZEW Mannheim.

[117] Rotfuß, W., Conrad, C., Rittler, D., 2009. The European commission and EUA prices: A high-frequency analysis of the EC's decision on second NAPs. ZEW Discussion Paper, Mannheim.

[118] Sadorsky, P., 2001. Risk factors in stock returns of Canadian oil and gas companies. Energy Economics 23, 17-27.

[119] Schwarz, T., Szakmary, A., 1994. Price discovery in petroleum markets: Arbitrage, cointegration and the time interval of analysis. Journal of Futures Markets 14, 147-167.

[120] Schwert, W., 1989. Why does stock market volatility change over time? Journal of Finance 44, 1115-1153.

[121] Scruggs, J., 1998. Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach. Journal of Finance 53, 575-603.

[122] Scruggs, J., Glabadanidis, P., 2003. Risk premia and dynamic covariance between stock and bond returns. Journal of Financial and Quantitative Analysis 38, 295-316.

[123] Smale, R. Hartley, M., Hepburn, C., Ward, J., Grubb, M. 2006. The impact of the CO_2 emissions trading on firm profits and market prices. Climate Policy 6, 29-46.

[124] Theissen, E., 2002. Price discovery in floor and screen trading systems. Journal of Empirical Finance 9, 455-474.

[125] Thompson, S., 2011. Simple formulas for standard errors that cluster by both firm and time. Journal of Financial Economics 99, 1-10.

[126] Tse, Y., 1998. The conditional heteroscedasticity of the Yen-Dollar exchange rate. Journal of Applied Econometrics 13, 49-55.

[127] Tse, Y., 1999. Price discovery and volatility spillovers in the DJIA index and futures markets. Journal of Futures Markets 19, 911-930.

[128] United Nations Framework Convention on Climate Change, 1992. Text of the Convention. URL: unfccc.int/key_documents/the_convention/items/2853.php.

[129] United States Energy Information Administration, 2011. Annual energy review 2010. URL: www.eia.gov/totalenergy/data/annual/pdf/aer.pdf.

[130] Uhrig-Homburg, M., Wagner, M., 2009. Futures price dynamics of CO_2 emission allowances: An empirical analysis of the trial period. Journal of Derivatives 17, 73-88.

[131] Veith, S., Werner, J., Zimmermann, J., 2009. Economic consequences of emission trading schemes: Evidence from the European power sector. Energy Economics 31, 605-613.

[132] Wei, C., 2003. Energy, the stock market, and the Putty-Clay investment model. American Economic Review 93, 311-323.

[133] Wei, Y., Wang, Y., Huang, D., 2010. Forecasting GARCH-crude oil market volatility: Further evidence using class models. Energy Economics 32, 1477-1484.

[134] White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817-838.

[135] Woerdman, E., Couwenberg, O., Nentjes, A. 2009. Energy prices and emissions trading: Windfall profits from grandfathering? European Journal of Law and Economics 28, 185-202.

[136] Yan, B., Zivot, E., 2010. A structural analysis of price discovery measures. Journal of Financial Markets 13, 1-19.

[137] Zachmann, G., von Hirschhausen, C., 2008. First evidence of asymmetric cost pass-through of EU emissions allowances: Examining wholesale electricity prices in Germany. Economics Letters 99, 465-469.

[138] Zhang, Y., Wei, Y., 2010. An overview of current research on EU ETS: Evidence from its operating mechanism and economic effect. Applied Energy 87, 1804-1814.