UNIVERSITY OF TWENTE.

de heus

FUTURES HEDGING OF COMMODITY RISK

Designing Hedging Strategies with a Focus on the Study of the Optimal Hedge Ratio

by
Tobias Kocks

MARCH 25, 2015

Document Information

Document type:	Master Thesis
Author:	Tobias Kocks
Date:	25th of March, 2015
Institute:	University of Twente
Faculty:	School of Management and Governance Drienerlolaan 5 7522 NB Enschede
Program:	Industrial Engineering & Management (IEM)
Track:	Financial Engineering & Management (FEM)
Company:	De Heus Voeders B.V. Rubensstraat 175 6717 VE Ede

Graduation Committee

First examiner:	Dr. Berend Roorda University of Twente
Second examiner:	Dr. Reinoud Joosten University of Twente
External supervisor:	Ben Tacken De Heus Voeders B.V.

Management summary

The goal of this research project was to develop hedging strategies for the three commodities corn, wheat and soy using futures securities. Under a futures hedge, the variability of the return of a spot commodity is sought to be alleviated by taking an offsetting position in a correlated futures position. The prime purpose of a hedging strategy then is to find that ratio of the futures position relative to the spot position, such that it yields on optimal outcome. This is what is termed the optimal hedge ratio – the strategic variable of this research.

The first step in this endeavor concerns finding futures products that best mimic the spot return movement. While most of the music in the agricultural commodity market plays in the US, for corn and wheat it was found that local spot price movement was best matched by the corresponding Matif exchange products. CBOT soybean meal contracts were found to be best correlated with local soy exposure. While the term 'local' is somewhat generic, it has been used intentionally, for the different spot exposures are represented by two different datasets per commodity – the internal replacement values on the one hand, and external (i.e., Oil World and Reuters) price series from related products on the other hand.

The second step is concerned with settling on an objective that is to be achieved by the hedge. The two targets that are considered in this research are reducing the variance of the hedge portfolio to the maximum extent on the one hand, and optimizing both variance and average return on the other hand, giving rise to the minimum variance –, and mean-variance objective function, respectively. For both functions, the variance and correlation of the spot and futures returns are the paramount determinants of the optimal hedge ratio. Therefore the major task of the internship was concerned with setting up databases for the internal and external price series and programming a model that would, amongst other things, flexibly calculate and plot the volatility and correlation of any two variables to get an idea of the development of the correlation and volatility dynamics over time.

Having established that correlation and volatility are the crucial ingredients in determining the hedge ratio under both objective functions, the third step is then about applying models to estimate the volatility and correlation parameters. Four industry standards are applied in this paper – the OLS, the EWMA, the DCC, and the diagonal BEKK model. The former is a static model – it simply calculates the sample average and standard deviation and applies the resulting hedge ratio throughout the entire life of the hedge. The latter three represent econometric, conditional models that incorporate new information into the model as it arrives, to come up with a dynamic, time-varying estimate. Each of these models can be used under both objective functions.

In line with more recent findings in the literature, the time-varying models seem to add little benefit in the context of the minimum-variance objective function. That is, this research indicates that it is beneficial to stick to the straightforward, static OLS regression method. Under the mean-variance objection function, however, time-varying models do add value if the average rate of returns of the underlying spot and futures securities are sufficiently big, which has proven to be the case for the soybean meal commodity case.

Preface

There we are. After months of work I can finally present my master Financial Engineering & Management graduation thesis. As a former bachelor Business Administration graduate I was not quite sure whether switching studies was the right thing to do. After all, since FEM is a lot more technical, there was a risk component attached to it. However, with the overwhelming experience of the Financial Engineering minor in mind, my gut instinct indicated that this was the right way to go.

Indeed, in hindsight, this switch was the best decision of my life so far. Not only could I finally study the precious Finance discipline with the technical twist that I missed so much in my bachelors, it also, in an early stage, put me in touch with my passion for programming that I was only latently aware of before. Fortunately, as the master FEM program is twice as big as a master BA program, I could still devote a considerable amount of elective courses to further develop my programming proficiency.
Moreover, the size of the program allowed me to take a couple of my elective courses abroad in Lisbon. What an amazing experience that was. I had the pleasure to make a lot of friends and meet inspirational colleagues, teachers, and other persons. It made me realize that collaboration with different cultures is something that I definitely do not want to miss out on in my future employment career.

Typing those final sentences of my very last report as a student, I would like to thank a couple of people explicitly that helped me get this far. First of all, I would like to thank Wouter Slot, a peer of mine. We have taken a lot of courses together and in some ways complemented one another. Without him, studying would not have been as enriching and joyful as it was. Moreover, I believe that my curriculum would look different to date if it was not for his encouragement.
Next, I want to thank Dustin Schilling, my roommate throughout the bachelors and part of my masters study. He significantly shaped my study experience in a positive way, especially in the bachelor stage. Of course, I would also like to thank all my other colleagues, friends and family for their ongoing support.

As for the thesis in particular I first and foremost owe gratitude to the company De Heus Voeders B.V, which gave me the opportunity to conduct this thrilling project that combines elements of both of my favorite disciplines, Finance and IT. I want to thank them for the intensive introduction phase, the numerous meetings where they have let me been a part of, the company events that they invited me to, but most of all I want to thank the staff of the Purchase & Trade division, especially Ben Tacken, my external supervisor, for taking the time to answer my questions. I gained a lot of insights from the agricultural commodity market and their IT-landscape. I hope that I can, at least in part, return the favor with this thesis and the programmed application.

Last but absolutely not least, I am very thankful for the valuable feedback of my first supervisor, Berend Roorda, and my second supervisor, Reinoud Joosten. Knowing that your time is scarce, I am very grateful for the time and effort you took to provide input to further improve my thesis.
Apart from that, I would also like to thank both of you for the numerous lectures that you provided and especially Berend for being confident about the study switch and his input for tailoring my curriculum.

Table of Contents

Management summary ... IV
Preface ... V
Chapter 1 - Introduction ... 1
1.1 Company and Background Information ... 1
1.2 Problem Statement ... 3
1.3 Research Questions ... 7
1.4 Thesis Structure ... 8
Chapter 2 – Literature Review ... 10
2.1 Review of Alternative Objective Functions ... 11
2.1.1 Minimum Variance (MV) Objective Function ... 11
2.1.2 Mean - Variance Objective Function ... 11
2.1.3 Maximum Expected Utility Objective Function ... 12
2.1.4 Conclusion on Objective Functions ... 14
2.2 Review of MV OHR Estimation Methods ... 15
2.2.1 Ordinary Least Squares (OLS) Regression ... 16
2.2.2 OHR and Time Variation ... 17
2.2.3 Exponentially Weighted Moving Average (EWMA) ... 18
2.2.4 GARCH Class OHRs ... 19
2.2.5 Diagonal VECH Parameterization ... 21
2.2.6 (Diagonal) BEKK Parameterization ... 21
2.2.7 Constant Conditional Correlation (CCC) Parameterization ... 22
2.2.8 Dynamic Conditional Correlation (DCC) Parameterization ... 22
2.2.9 A Note on GARCH Models and Symmetry ... 23
2.2.10 Conclusion on MV OHR Estimation Methods ... 23
2.3 Review of Hedging Performance Parameters ... 24
2.4 Conclusion on Literature Review ... 25
Chapter 3 – Methodology ... 27
3.1 The Data ... 27
3.2 The Tools ... 29
3.3 Minimum Variance Hedge Effectiveness Measurement ... 32
Chapter 4 – Analysis ... 33
4.1 Correlation Study – Identification of the Hedge Securities ... 33
4.2 Discussion of Descriptive Statistics ... 37
4.3 Parameterization of Hedging Strategies ... 40
4.3.1 Corn Parameterization ... 41
4.3.2 Wheat Parameterization ... 46

4.3.3 Soy Parameterization 47
4.4 Application of Minimum Variance Hedging Strategies 48
4.4.1 Case Corn 49
4.4.2 Case Wheat 55
4.4.3 Case Soy 59
4.5 Sensitivity Analysis 63
4.6 Application of Mean-Variance Hedging Strategies 66
4.6.1 Case Corn 67
4.6.2 Case Soy 69
4.6.3 Case Wheat 71
4.7 Discussion and Implications 73
4.7.1 Minimum-Variance 73
4.7.2 Mean-Variance 75
Chapter 5 – Conclusion 77
Chapter 6 – Recommendations 79
Chapter 7 – Limitations & Future Research 81
References 83
APPENDICES 87
APPENDIX A – Objective functions 87
APPENDIX B - OHR and cointegration 89
APPENDIX C – Meta-analysis of empirical studies: OLS vs. time-varying models 90
APPENDIX D – Constant Conditional Correlations 92
APPENDIX E - Custom program interface 94
APPENDIX F - Histograms and descriptive statistics 95
APPENDIX G - Correlograms: Autocorrelation in squared returns 97
APPENDIX H - EViews Wheat Parameterization Output 99
APPENDIX I – EViews Soy Parameterization Output 101
APPENDIX J – EViews OLS Regression Monthly returns 103
APPENDIX K – EViews BEKK estimation Monthly returns 104

Chapter 1 - Introduction

This chapter commences with an introduction of De Heus Voeders B.V., the problem owner of this project, passes over to the problem identification and the corresponding research questions, and concludes with an overview of the structure of the remainder of this thesis.

1.1 Company and Background Information

De Heus Holding

De Heus Voeders B.V. is a global animal nutrition producer founded in 1911. While its headquarter is located in Ede, the Netherlands, which is also the center of their group activity, De Heus is also commercially active in more than 50 countries, including its foreign affiliates in Poland, Czech Republic, Russia, Vietnam, South Africa, Egypt, and Ethiopia, its joint venture with Wellhope in China and its partnership with Nutrifarms in Brazil.

With its more than 3000 employees, De Heus accomplishes an annual production output of more than 4 million tons of compound feed, with about 50% of the output being contributed by the 8 production plants in the Netherlands, 25% originating from Poland, and the remaining 25% being scattered across the remaining countries. There are numerous product segments, which can roughly be divided into five groups – cattle, pigs, layers, broilers, and miscellaneous (including horses, goats, sheep, and fish).

While the headquarter in Ede features departments that are present in any corporate institution (e.g., Finance & Legal, HR, Marketing, Sales, etc.) it is worth to highlight their Formulation department for a moment. This department is closely integrated with the Purchase & Trade department, the problem owner of this thesis. It optimizes the feed recipes for all production plants, local as well as foreign ones. The tool used for optimization is based on linear programming, and calculates the optimal feed recipes by taking into account output -, quality -, nutritional -, and storage constraints as well as raw material prices. Given that prices fluctuate over time, so does the composition of the recipes.

Purchase and Trade Department

The Purchase & Trade department can further be subdivided into two teams – one responsible for buying macro ingredients and one for purchasing micro ingredients. As the names suggest, the former term denotes the raw materials that make up the bigger portion of the feed in terms of weight and includes grains, proteins, fats & liquids, and byproducts. Though there is a great variety of macros that are frequently used in production, the ones with the biggest share in terms of weight are *corn*, *wheat* and *soy*. Those (macro) commodities also form the focus of our research. The term micro ingredients accommodates all the additives that are mixed into the compound feed. Think of *vitamins*, *minerals*, *enzymes*, *trace elements*, *amino acids*, *aromas*, etc. While their share in weight in the compound feed is much lower, the price per unit of weight is much higher.

But it is not just the characteristics of the two raw material groups that differ. There are also differences in terms of market structure. Whilst macro ingredients are produced and traded by numerous market participants throughout the world, the market structure of the micro ingredients market is much more oligopolistic. In this case, it is the production capacity of a handful of suppliers that primarily drives market prices. Moreover, micro ingredients lack a clear, universally accepted reference market price.

The Commodity Market

A number of macro ingredients, on the other hand, are traded on most futures exchanges around the world. Think of cereals like corn, wheat, barley, oats or oilseeds like soybeans, rapeseeds, or palm. The word futures is key here, and indicates that agents are able to agree on a transaction in the future, based on a price agreed upon today. On a futures exchange, each product has its own delivery calendar with mostly around four to eight delivery dates per year, where the most immediate delivery date is referred to as the lead, front, or spot contract month. While each futures contract specifies a set of quality standards, unit volume, and a pool of destination ports and thus makes a physical delivery possible, contracts are usually either closed out or rolled forward (say, at November 10, selling 2000 long November contracts and purchasing 2000 January contracts) prior to the settlement date (e.g., November 14), which renders them an efficient hedging security.

The notion of forward trading is not only key for futures commodity trading, but is also common practice in the physical cash market, where it is not unusual to book a delivery and thus fix a price for a period of, say, 12 months. The physical commodity market is a global one with international import and export flows. The US is by far the most prominent exporter for corn (50.71 million tons in 2013/14 vs. 129.62 million tons worldwide), most of which is imported by Asia (Rabobank Food & Agribusiness Research, 2007). Likewise, it is also a leading exporter of soybeans together with Brazil (US: 44.82 million tons, Brazil: 46.83 million tons vs. 112.83 million tons worldwide in 2013/14) (USDA, 2015).

America's role as the paramount producer and exporter for the most important grain and oilseed partly explains why its futures exchange (the Chicago Board of Trade, henceforth CBOT) is also leading – leading in terms of trade volume, but also in the sense that it is *the* exchange that is used as reference for pricing cash market contracts. The Matif (a division of Euronext Paris, formerly known as Paris Bourse before the merger with Euronext NV in 2000) is the biggest futures exchange for agricultural products in Europe. However, its trade volume represents only but a small fraction of what the CBOT turns over. Using futures exchanges as a reference to price cash market contracts is common practice, where the unit price is the futures exchange price plus a premium or discount depending on, amongst other things, the supply & demand and logistics of the products' country of origin, as well as the quality and the destination of the product.

Purchase Policy, Hedging and Coordination

In deciding whether or not to trigger a long or short position, market fundamentals (supply and demand mostly) are the prime source of information. Insights from statistical and numerical analyses (e.g., volatility, moving averages of correlations and prices, curve shifts, etc.) are regarded as an additional, second opinion, piece of the puzzle.

As a processor of agricultural commodities, De Heus builds up length in a certain commodity if they are under the impression that prices for that commodity are at the verge of increasing. If, on the other hand, a bearish sentiment prevails, the quantity bought of that commodity is such that it is just enough to keep the factories running. Even though the different subsidiaries are responsible for their own procurement, their purchase positions and performances are evaluated at corporate level.

Apart from purchasing raw materials from local traders for spot or future, physical delivery, De Heus also takes cash settled positions at futures exchanges (CBOT mostly). Occasionally, this is done to hedge price fluctuations in the cash market. From the point in time that a commodity enters the books of De Heus, up to the time it is consumed in the factories, De Heus is *long* the commodity. Subsequently, a physical position is hedged with a *short* position in a futures contract.
Other times, however, taking futures positions is done on speculative grounds in order to exploit a certain vision on the market. That is, they would go long (short) if commodity prices are expected to increase (decrease) over time or use a combination of both (i.e. spreads) for limited exposure.

Replacement Values
Each week, the purchasers at De Heus estimate *replacement values* for an array of commodities. The replacement value reflects the purchaser's estimate of what it would cost per unit of measurement to buy a certain commodity in the local cash market. This is done for the most immediate, spot delivery but also for delivery in i months, for i = 1, 2, …, 12. The motivation behind this is that, in any point in time, there are likely to be book exposures spanning a delivery period of more than just a single month. The replacement values thus allow them to continuously monitor the economic value of their purchase books. In the context of replacement values, the term 'spot' as used in our thesis explicitly refers to this most immediate delivery period.

1.2 Problem Statement

The goal of our research is to

> *"Contribute to the risk management of De Heus by designing a hedging strategy for their most impactful commodity items."*

To understand why it is this issue that is the most relevant to tackle, we have to turn our attention to the interplay of a twofold of recent developments.

Structural Changes and Risk Appetite
First of all, De Heus has been growing tremendously over the past decade and has the ambition to *double* its output scale in the near future. While the company has grown in terms of people and output, processes are somewhat lagging behind, calling for more centralization. A couple of trajectories have already been put into place to tackle this problem. Their corporate IT system is being upgraded, trying to integrate, among other things, the contract -, price -, and stock data of all holding countries. Moreover, a global procurement and supply management system is in the making, to bundle corporate purchases and facilitate intra-holding transfer and distribution. However, there is still a variety of information that is not currently utilized.

Second, the purchase and trade department has a risk taking attitude. Even though most of the competitors cover their demand for raw materials mostly by buying spot in the market, De Heus is more pronounced in taking positions up to e.g., 18 months ahead in the future, thus locking in a presumed low price. This attitude is based on the philosophy that their professional insights into the (demands of the) feed industry as well as market fundamentals enable them to make informed decisions and take riskier (i.e., longer term) positions. Obviously, while this can save you a lot of money, there is also a risk of incurring extra costs if the price development takes on an adverse path, where the degree of risk advances with the distance of the delivery period.

Market Risk Management

Apart from the big risk appetite, we are also dealing with a market that has become more and more difficult to read over the past years. While market fundamentals are still the driving force in dictating price movement, the increased activity of institutional investors (mutual -, hedge -, pension funds, etc.) has lead volatility to increase. Figure 1.1 illustrates this by depicting the 40 days moving average volatility of the front month corn and milling wheat Matif futures contracts over the past 14 years. It shows how a) volatility swings have become bigger over those 14 years, where especially big swings can be observed in the past 4 years, b) how the base volatility level has increased in the period after 2006 and c) how the volatility of the most recent months has sharply risen. Note, that the CBOT futures contracts show similar movements in volatility.

The combination of those three factors (increased volume, volatility and risk appetite) translates into a considerable market risk exposure, which should be managed accordingly.

Market risk has four components – commodity risk, currency risk, interest rate risk, and equity risk. The latter component is, apart from a few minority interests, not applicable to our case.

Interest rate risk, though applicable, only has a marginal impact compared to currency – and commodity risk. While currency risk certainly is a hot item within the corporation, commodity risk – the risk stemming from an adverse development of the price or volatility of a commodity – has the biggest impact on the financial results of the company and is thus the focus of this research.

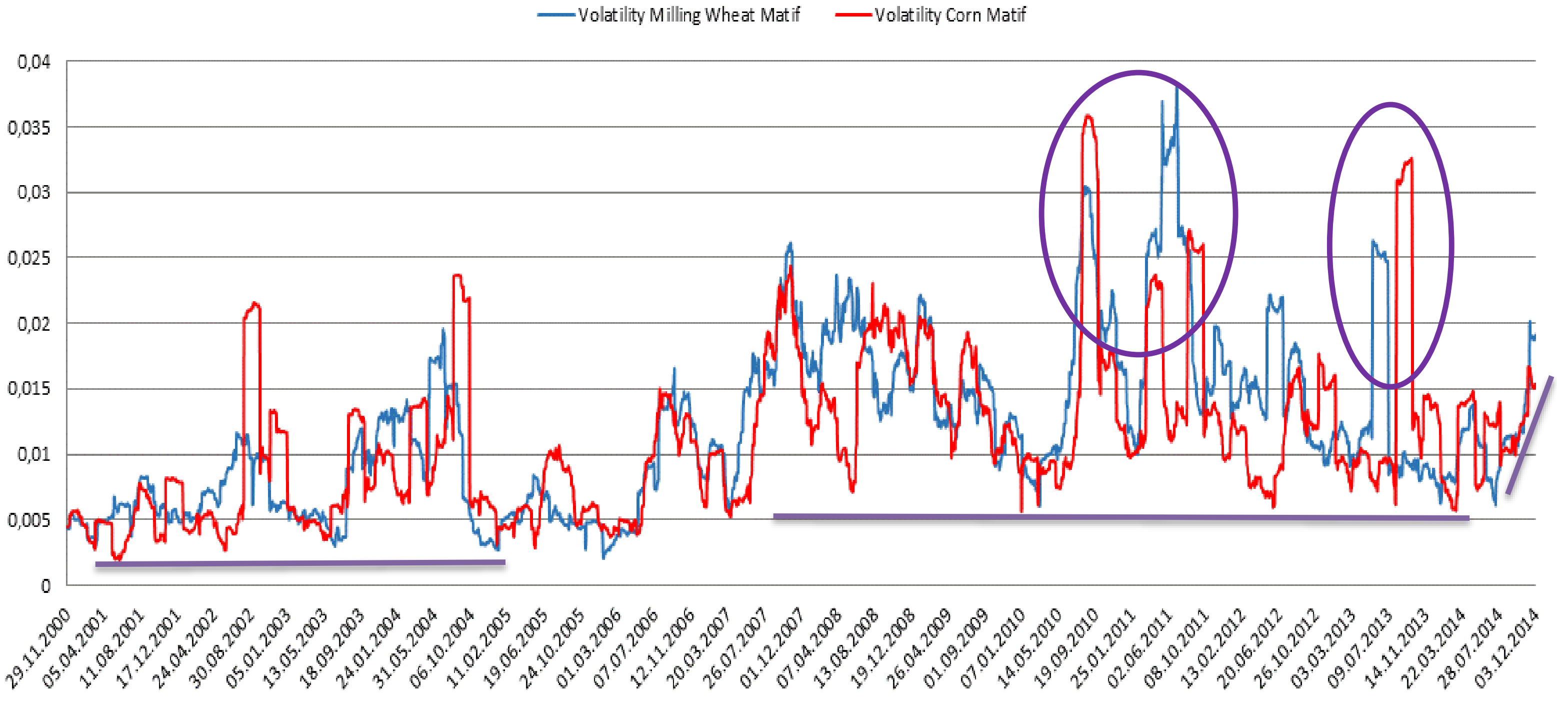

Figure 1.1 – *40 days moving average volatility of the EMAc1 (corn) and BL2c1 (MW) contracts.*

As has already been established earlier, the commodity market is mainly a futures market. Even though commodity options are available on some exchange markets, De Heus thinks that the complexity of those derivatives presumably does not match their current state of development. Hence, it is self-evident to resort to futures as a financial security in order to hedge market risk exposure.

Basis Risk

The intent of a futures hedge is to take an 'opposite' position in a futures contract (e.g., long physical and short futures) so that any potential loss in the physical position can be offset by a gain in the futures position. The problem with that undertaking is that the returns of a commodity eventually paid for in the cash market, are *less than perfectly correlated* with the returns of any of the available futures contracts (irrespective of the futures exchange). Less than perfect correlation translates into what is termed as **basis risk**. Consider the case where the cash commodity price decreases by Δ. Assume that the futures price of that commodity, however, only decreases by .7 Δ. This leaves an extra loss of .3Δ, which is not offset by the gain in the short futures position.

Note that the term basis is not applied consistently in the literature. Usually it is defined as the excess of the cash market asset price over the futures security price, but sometimes it is also used the other way around. In this research we stick to the common notation, i.e., basis = cash asset price – futures price.

Figure 1.2 gives an example of the movement of the basis (green line) of a) the cash market price of one of our corn products (blue line) and b) the futures corn product traded on the Matif (red line) over the past two and half years. As we can see, even though both products refer to the same commodity, there is still considerable fluctuation in the basis, moving within the range of
€ -24 to € 46 per metric ton.

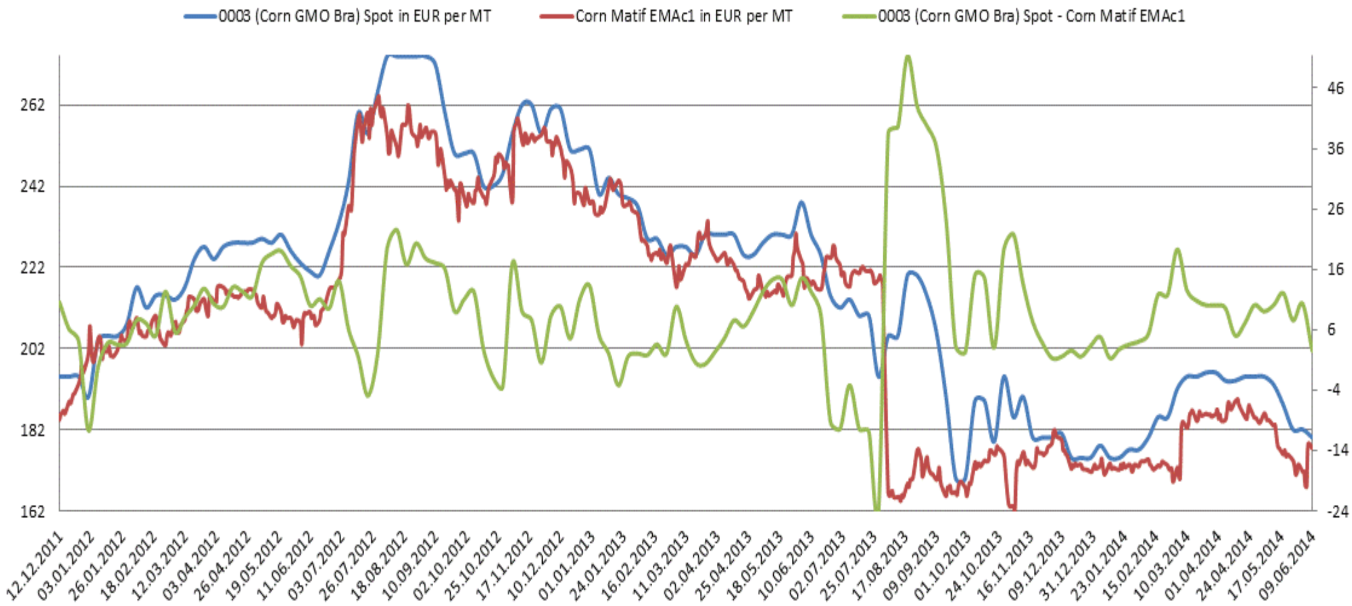

Figure 1.2 – *Basis (green line) movement of corn cash (blue) and corn futures (red) over time.*

Optimal Hedge Ratio

The presence of basis risk is relevant in terms of hedging strategies, because it affects the optimal number of contracts that is needed to hedge a certain cash commodity position. If the correlation between the two variables in a hedge were perfect (i.e., 1), we would simply buy an amount of futures contracts such that the total value (price times quantity) of all those contracts matches the total value of the outstanding exposure (*bona fide* hedge). However, as stated earlier, correlation is imperfect and also fluctuating over time. The body of literature governing the **optimal hedge ratio**, i.e., the value of the futures position relative to the value of the cash commodity position, is concerned with this issue.

To see how the *hedge ratio* at time *t-1*, h_{t-1}, the strategic variable in this research, *determines* the *outcome* of a *hedge*, consider the following notation of $R_{h,t}$, the return of the hedged portfolio at time *t* as given by Chen *et al*. (2003):

$$R_{h,t} = \frac{C_{s,t-1}S_{t-1}R_{s,t} - C_{f,t-1}F_{t-1}R_{f,t}}{C_{s,t-1}S_{t-1}} = R_{s,t} - h_{t-1}R_{f,t} \tag{1a}$$

where $C_{s,t-1}$ denotes the number of long spot units, $C_{f,t-1}$ denotes the number of short futures units at *t-1*, S_{t-1} and F_{t-1} represent the spot and futures prices at time *t-1*, respectively, and $R_{s,t}$ stands for the return of the spot security, $R_{s,t} = (S_t - S_{t-1})/S_{t-1}$. Likewise, we have $R_{f,t} = (F_t - F_{t-1})/F_{t-1}$.[1] Note, that in the context of this research, the spot prices, S_t, are e.g., the replacement values of De Heus. It follows from the above that the hedge ratio is given by

$$h_{t-1} = \frac{C_{f,t-1}F_{t-1}}{C_{s,t-1}S_{t-1}} \tag{1b}$$

which is indeed the total value of the futures position over the total value of the spot position at time *t-1*. The notations of (1a) and (1b) stress that the optimal hedge ratio is set *prior* to period *t* (i.e., at the end of period *t-1*) in order to optimize the expected outcome *of* period *t*.

Note, that in this research, we consider each commodity hedge separately. That is, per spot commodity that we analyze, the hedge portfolio in (1a) consists of one type of spot security and one type of futures security, only. This restriction greatly simplifies the analysis and is also persistently applied in literature. Figure 1.3 visualizes the factors and their interactions leading to the central research question.

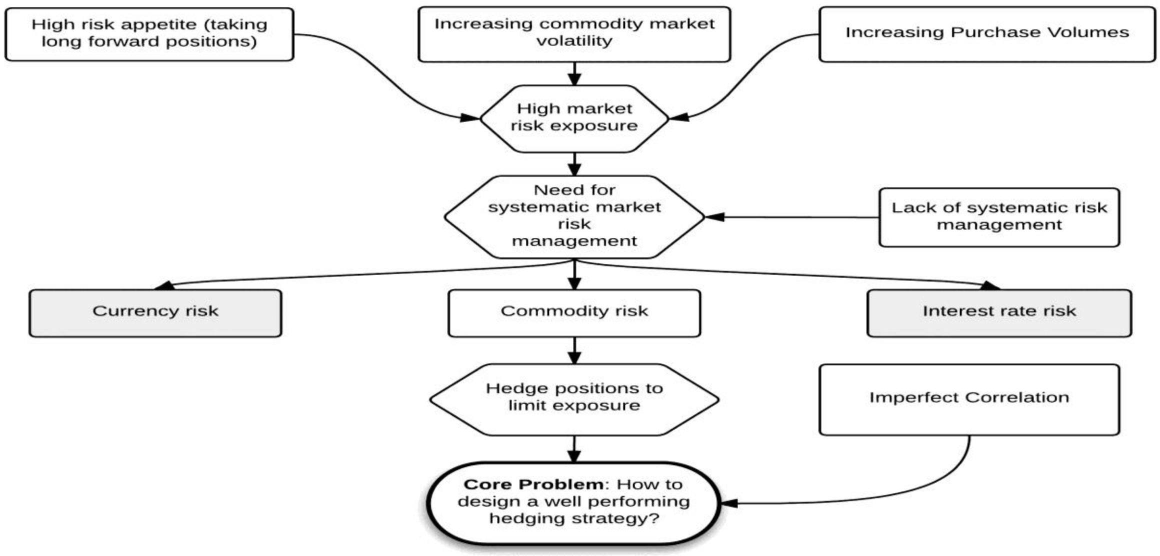

Figure 1.3 – *Problem Cluster.*

[1] In practice futures contracts do not require an initial outlay – the cash flows stem from the daily mark-to-market valuation. Still, it is conventional to base hedging strategy calculations on returns rather than plain changes. Amongst other reasons, this is because the correlations would be distorted in a cross hedge scenario, where different securities may be quoted in different units of measurements and/or currencies. In that regard, returns provide a standardized unit of measurement.

1.3 Research Questions

Central Research Question
The *central research question* of our thesis is therefore the following:

> *Central Research Question: "How to design a hedging strategy at De Heus for the most important commodities with special attention to modeling the optimal hedge ratio?"*

De Heus purchases and processes more than 300 ingredients. However, as the research question suggests, we limit ourselves to study a handful of commodities only, due to limitations of scope. The ingredients we do analyze, are *corn*, *wheat*, and *soy*. Those are the most liquid commodities and together they make up about 50% of the weight of the compound feed.

In order to answer the main research question, we set up a number of research sub questions. By answering those sub questions, we systematically aim to give an answer to the main research question.

Research Sub Questions
The most basic action in setting up a hedging strategy is picking a futures contract that best mimics the price movement of the cash market variable that we want to hedge. Since we have different continuations per futures contract, different commodity exchanges that trade futures for the same commodity, and since there is a certain degree of substitutability between the (agricultural) commodities, we have various futures products that qualify as suitable hedging vehicle per cash commodity we want to hedge.

More specifically, we are going to inquire basic *basis spreads* (e.g., long February cash corn and short March CBOT corn, both initiated in December), *time spreads* (e.g., long February cash corn and short May CBOT corn, both initiated in December) and *cross spreads* (e.g., long February cash corn and short March CBOT wheat, both initiated in December) as hedging setups.

Though this might be another point of discussion, in the spirit of Hull (2012), the suitability of a futures security as a hedging vehicle will be judged solely on the grounds of its correlation with the cash commodity that we want to hedge. Moreover, due to the earlier mentioned preeminence of the CBOT and Matif in Europe, we confine our pool of candidate futures contracts to all the agricultural commodities that are traded on those two exchanges. The first research sub question therefore is:

> *Sub Question 1: "Of all the agricultural futures products traded on the CBOT and Matif, based on the degree of correlation, per cash commodity, which futures contract qualifies as the most suitable hedging security?"*

The most important parameter in the context of a futures hedge is the optimal hedge ratio (OHR), i.e., the proportion of the total value of a futures position relative to the total value of the original exposure to be hedged. There are numerous alternatives that attempt to model this ratio such that the hedge performs effectively and/or efficiently. Each of these models makes its own assumptions about the stochastic behavior of the variables at issue in the hedge. Correspondingly, our second research sub question is the following:

Sub Question 2: "Which models are available to model the OHR and what restrictions do they pose on the variables involved in the hedge?"

The third research sub question directly follows from Sub Question 2. Given the set of models that the literature research yields, and given the assumptions those models make, it would be interesting to see in how far, according to basic descriptive statistics, our spot and futures time series conform to those assumptions. Therefore we have:

Sub Question 3: "Given the historic data of the cash and futures commodities, to what extent do those time series comply with the assumptions posed by the models identified in Sub Question 2 and what would be the implication of a potential violation of those assumptions?"

Once we have identified a set of methods to model the OHR, we can apply those models to our historic data and can compare their performance for the different hedging setups identified in Sub Question 1. While the term optimal hedge ratio is casually used in the finance community, what is really optimal certainly is a point of dispute. Therefore, we need to establish the grounds on which we compare the performance of the different OHR models, i.e., we have to ask ourselves:

Sub Question 4: "On the basis of which indicators do we measure the performance of a hedge?"

Once we have distilled a set of key indicators we can rank the performance of the different hedging strategies. However, it is likely that the answer will not be terminal. Model output is a function of model input and if the input parameters vary, so may the outcome and the corresponding conclusions. A sensitivity analysis should therefore yield insight into the robustness of the strategies' performances:

Sub Question 5: "How volatile are the hedging strategy outcomes with respect to the choice of the base parameters?"

1.4 Thesis Structure

This section gives an outline of how the remaining chapters will be devoted to answering the research sub questions identified in the preceding section.

In Chapter 2, the *literature research*, we will shed light into *Sub Questions 2* and *4*. That is, first, we will scan the literature for the different OHR models and their assumptions.
Moreover, we will investigate the indicators used to determine the performance of a hedging strategy as well as the factors that we need to take into account when evaluating a strategy's performance, thus answering *Sub Question 4*.

In Chapter 3, the *methodology* chapter, a more detailed description of the different databases for the cash prices is given. Moreover, the significant part of this research project was concerned with structuring currently underutilized data. We have set up structured databases for the different (internal and external) data sources and have developed an application that allows for the query and plotting of plain price -, spread-, volatility-, and correlation – series, as well as marginal and bivariate distributions or scatterplots of any of the available price series. This program is also introduced more thoroughly in Chapter 3.

The *analysis* of the data is executed in Chapter 4. First, we kick off by answering *Sub Question 1*. We use our programmed model to review the possible hedging setups to filter a number of combinations that are worth further pursuing in the remainder of the analysis.
Additionally, we analyze the stochastic characteristics of those series and check in how far they might violate the assumptions of the models identified in Chapter 2, i.e., we approach *Sub Question 3*.
Subsequently, we apply our data to the chosen models and analyze the performance of the different models based on the criteria that we found in *Sub* Question 4. Finally, a sensitivity analysis is performed to test the robustness of the proposed solutions (*Sub Question 5*).

The *conclusions* of this research are drawn in Chapter 5. Here, we get back to the central research question and answer it in light of the empirical findings of Chapter 4. It summarizes the takeaways of this project and how they have been derived.

In Chapter 6, we discuss the *recommendations* that follow from our research. That is, it describes the insights from the research that are relevant for an actual implementation of a hedging strategy.

The final chapter highlights the *limitations* of this research that have to be taken into account when interpreting the results of this thesis. In the same vein, it suggests issues for further research that could not be touched upon in the limited scope of this research project.

Chapter 2 – Literature Review

Due to the attractiveness of futures contracts as a hedging vehicle amongst practitioners, futures hedging has also become the focus of a lot of research. Being one of the most central questions in futures hedging, the study of the optimal hedge ratio, i.e., the total value of the futures position relative to that of the spot position, has furthermore attracted a great deal of attention within that field of research.

Recall from (1a) and (1b) in the introduction that the return of the hedged portfolio at time t, $R_{h,t}$ and the optimal hedge ratio, *set at* time *t-1*, h_{t-1}, are respectively given as

$$R_{h,t} = R_{s,t} - h_{t-1}R_{f,t} \qquad and \qquad h_{t-1} = \frac{C_{f,t-1}F_{t-1}}{C_{s,t-1}S_{t-1}}$$

Once more, bear in mind that we set the optimal hedge ratio *for* period *t at* the end of period *t-1*. *From this point onwards*, to stress the fact that we choose the hedge ratio to optimize the outcome *for* period t, we will refer to this variable as h_t, rather than h_{t-1}.
Now, the very first step in modeling the OHR concerns the choice of the objective function that ought to be minimized or maximized, i.e., the function that we apply on the variable $R_{h,t}$ and optimize with respect to the parameter h_t in order to yield the OHR, h_t^*.

The objective function is the function that describes how the return of the portfolio, $R_{h,t}$, is transformed into the investor's utility. Obviously, it would be theoretically most rigorous to directly optimize the investor's utility function (Cecchetti *et al*., 1988). In practice, however, one usually applies a set of assumptions concerning the shape of the investor's utility function so as to make the calculations traceable.

Having settled on an objective function, the expression of h_t^* is mostly a function of the moments of the distribution of $R_{h,t}$. The second step is then concerned with the various econometric methods that are available to estimate those parameters.

The remaining chapter is structured into four parts. The *first part* provides a review of the different objective functions that have been postulated and applied in literature in order to derive an expression for the OHR, i.e., it deals with step 1 in modeling the OHR. It highlights their requisites and assumptions as well as their major advantages and shortcomings and eventually settles on an objective function that is going to be used in this research.

Part two of this chapter provides a review of the different econometric techniques that are applied in literature in order to statistically estimate the OHR. It is thus concerned with the second stage in modeling the OHR. Together, part one and part two are targeted to tackle *Sub Question 2*.

The *third part* of this chapter concludes with an overview and selection of the available hedging performance indicators. After all, we can only judge the performance of the models once we have established on what grounds they will be compared. That is, it tries to resolve *Sub Question 4*.

2.1 Review of Alternative Objective Functions

Even though one could think of any objective function, there are six such function groups that are frequently referred to in literature. Those are the the *minimum-variance* -, the *mean-variance* -, the *expected utility*, -the *Sharpe ratio* -, the *minimum mean-extended Gini*-, and the *minimum semi-variance* objective functions. This section provides a brief description of the former three objective functions. The others are outlined in Appendix A.

2.1.1 Minimum Variance (MV) Objective Function

When researchers review the evolution of OHR estimation, they usually start with Johnson, who was the first person to analytically derive the OHR on the basis of the minimum variance criterion in 1960 (Alexander & Barbosa, 2007).

As the name implies, the objective function F(.) in this case is simply the variance of the portfolio return. Sticking to the definition of the portfolio return in (1a) without time subscript, we get that:

$$\begin{aligned} F(R_h) &= Var[R_h] = Var[R_s - hR_f] = Var[R_s] + h^2 Var[R_f] - 2h\, Cov[R_s, R_f] \\ &= \sigma_{R_s}^2 + h^2\sigma_{R_f}^2 - 2h\,[\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}] \end{aligned} \tag{2a}$$

Minimizing the variance of the portfolio with respect to the hedge ratio, we first calculate the corresponding derivative:

$$\frac{\partial F(R_h)}{\partial h} = 2h\, Var[R_f] - 2\,[\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}]$$

Setting the derivative equal to zero, the OHR is given by:

$$\begin{aligned} &2h\, Var[R_f] - 2\,[\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}] = 0 \\ &h^* = \frac{\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}}{\sigma_{R_f}^2} = \frac{\sigma_{R_sR_f}}{\sigma_{R_f}^2} = \rho_{Rs,Rf}\,\frac{\sigma_{Rs}}{\sigma_{Rf}} \end{aligned} \tag{2b}$$

Where $\rho_{Rs,Rf}$ is the correlation between the futures and spot returns, and σ_{Rf} and σ_{Rs} are their corresponding standard deviations.
While the minimum variance approach is intuitive and easy to understand, and allows for an analytical derivation, a shortcoming of this approach is that it ignores the expected return of the portfolio.

2.1.2 Mean - Variance Objective Function

An approach that attempts to eradicate the absence of the mean in the utility function is the mean-variance approach. Hedging is costly. If you hedge away part of the variance of your portfolio, you will inevitably also hedge away part of your expected return. It would thus be reasonable to optimize $R_h$with respect to h such that it takes account of the effect on both risk and return.

In this context, Chen *et al*. (2003) propose the following objective function that has been empirically applied by Hsin *et al*. (1994):

$$F(R_h, A) = E[R_h] - 0.5A\,Var[R_h] \tag{3a}$$

where A is the risk aversion parameter, which is divided by 2 for computational convenience. Plugging in expressions (1a) and (2a) into (3a) we get:

$$\begin{aligned} F(R_h, A) &= E\left[R_s - hR_f\right] - 0.5A\left[\sigma_{R_s}^2 + h^2\sigma_{R_f}^2 - 2h\left[\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}\right]\right] \\ &= \mu_{R_s} - h\mu_{R_f} - 0.5A\sigma_{R_s}^2 - 0.5Ah^2\sigma_{R_f}^2 + Ah\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf} \end{aligned}$$

Again, optimizing (now maximizing) the objective function with respect to the hedge ratio, the first derivative is given by:

$$\frac{\partial F(R_h)}{\partial h} = -\mu_{R_f} - Ah\sigma_{R_f}^2 + A\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}$$

Equating the derivative with 0, the following expression for the OHR unfolds:

$$\begin{aligned} &-\mu_{R_f} - Ah\sigma_{R_f}^2 + A\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf} = 0 \\ &-h\sigma_{R_f}^2 = \frac{\mu_{R_f}}{A} - \rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf} \\ &h^* = \rho_{Rs,Rf}\frac{\sigma_{Rs}}{\sigma_{Rf}} - \frac{\mu_{R_f}}{A\sigma_{R_f}^2} \end{aligned} \tag{3b}$$

The resemblance of expression (3b) with (2b) is obvious. In fact, if we drop the second term of (3b) it is similar to expression (2b). This happens if either the numerator (μ_{R_f}) is zero, or the denominator ($A\sigma_{R_f}^2$) tends to infinity. This implies that the mean-variance optimization approach and the minimum variance approach yield the same outcome under the assumption that either the investor is infinitely risk averse (A → ∞), and/or that the futures products' price series follows a martingale process (i.e., $\mu_{R_f} = 0$) (Chen *et al*., 2003). [2]

While the mean-variance approach is advantageous in that it incorporates the expected return of the hedged portfolio, it introduces the problem that we now need the risk aversion parameter in order to solve for the OHR.

2.1.3 Maximum Expected Utility Objective Function

As mentioned earlier, the expected utility objective function is the most theoretically rigorous approach as it optimizes the utility function with respect to the OHR parameter, *without prescribing* the *shape* of the *utility function*. It is the overarching approach that embeds all objective functions.

[2] Under a martingale stochastic process, the expected value of random variable S at time t+1, S_{t+1}, is equal to the prior realized observation, of S, S_t – regardless of the set of values observed up to time t. That is, $E[S_{t+1}\,|S_1 \ldots S_t] = S_t$. It follows that in that case $E_t[(S_{t+1} - S_t)/S_t] = (S_t - S_t)/S_t = 0$.

Cecchetti *et al*. (1988) provide a graphical representation (see Figure 2.1) of this risk-return trade-off that is inherent in futures hedging and which has been introduced in the prior section. Here, R_h refers to the return of the hedged portfolio and is equal to the notation in (1a).

As the investor varies the hedge ratio h, the ratio of expected return to standard deviation also varies. This is indicated by the risk-return frontier, which is derived from the joint distribution of the spot and futures returns.

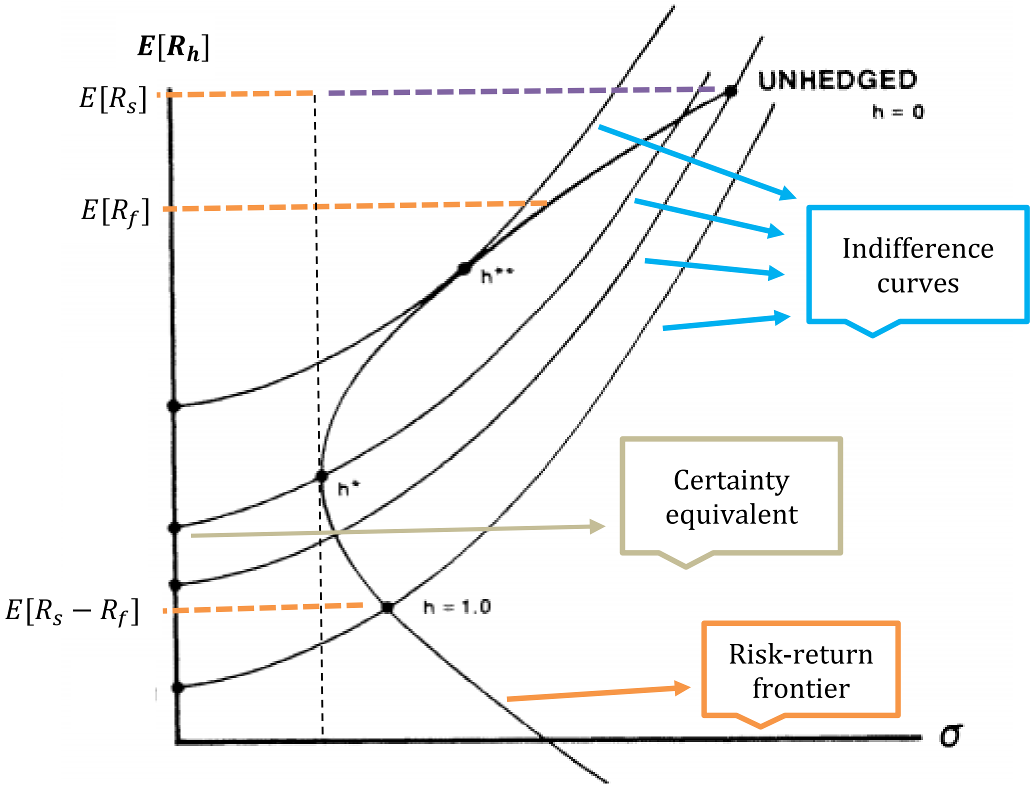

Figure 2.1 – *Risk return frontier (Cecchetti et al., 1988) and hedging strategies.*

In this figure, the hedge ratio increases as we move from the top to the bottom of the figure. Being completely unhedged, the expected return of the portfolio is the highest at $h = 0$ as there is no correlated short position that eats away part of the spot positions' expected return. In the same vein, however, there is not covariance term to reduce the overall risk of the spot commodity. The figure highlights how we sacrifice one (expected return) for the other (less risk). This trade-off culminates in h^*, the point of minimum variance discussed in Section 2.1.1. This is the leftmost point on the risk-return frontier.

As we further increase the hedge ratio, we enter the set of inefficient hedge ratios. Consider for example $h = 1$. Assuming the presence of basis risk, this constitutes a scenario in which we are overhedged. By definition, the variance at this point is bigger than at h^* since $h \neq h^*$ and h^* is the variance minimizing hedge ratio.[3] Moreover, since $h > h^*$ the proportion of the expected spot return that is erased by the expected futures return, is also bigger. Note, that in most cases, the variance minimizing hedge ratios are lower than 1.

[3] In case of constant hedge ratios, there is only one hedge ratio, h^*, that leads to the corresponding minimum variance point in Figure 2.1

Notice, that when the futures price series would actually follow a martingale process, in which case $E[R_f] = 0$, then the expected value of the hedged portfolio return, $E[R_h]$, would be immune to changes in the hedge ratio. In this event, the set of feasible hedge ratios would be given by the purple, horizontal line in Figure 2.1 (for $0 \leq h \leq 1$). This constitutes a very attractive scenario for we can reduce the variance without sacrificing part of the expected value. Ironically, however, if the spot exposure earns a risk premium it will have a positive beta. But for a futures security to be a suitable hedge instrument, it must be correlated with the spot asset, in which case it will have a positive beta as well and thus also earn a risk premium, i.e., have an expected value bigger than zero.

Cecchetti *et al*. (1988) now argue that the optimal futures hedging ratio is the one that maximizes the investors utility. To compute this ratio, we first of all need the utility function of the investor. For any given function, we can then derive a set of indifference curves (see Figure 2.1), where, as the name implies, each curve represents a set of points along which the investor, in terms of utility, is indifferent as to the ratio between risk and expected return. The OHR is then the point at which the slope of the indifference curve is equivalent to the slope of the risk-return frontier. This point is indicated as h^{**}in Figure 2.1 (Cecchetti *et al*., 1988).

While the expected utility approach is theoretically appealing, it requires us to know the analytic functions of the investor's utility and the joint distribution of the spot and futures returns, which makes it cumbersome to implement in practice.

Note, that both the minimum variance and the mean-variance objective functions are quadratic utility functions. Therefore, both of these functions are not generally consistent with the expected utility paradigm unless, of course, the investor happens to make investment decisions based on quadratic utility functions. Moreover, the three approaches are also identical to each other in case the spot and futures returns are bivariate normally distributed (Lien & Tse, 2002).

2.1.4 Conclusion on Objective Functions

Section 2.1 and Appendix A have given a brief review of the most common theories that are applied in literature in order to model the optimal hedge ratio. A number of important points crystallize alongside the comparison of the different theories.

First and foremost, all of the rather advanced methods analytically resemble the basic *minimum variance* approach under a few assumptions. If we assume that the futures price series, F_t, follows a martingale process then both the *mean-variance* and the *Sharpe-ratio* approach equal the *minimum-variance* function. The same is also true for the *expected utility*-, the *minimum MEG*-, and the *minimum GSV* approach if we impose the stronger assumption that the spot and futures returns, R_s and R_f, are jointly normally distributed (Chen *et al*., 2008). Dropping the assumption of joint normality, Baillie & Myers (1991, p. 117) less restrictively claim that

> *"However, it can be shown that, provided expected returns of holding futures are zero, then the minimum variance hedging rule is also generally the expected utility-maximizing hedging rule."*

Lien & Tse (2002) agree with the proposition of Baillie & Myers (1991) under the condition that hedgers are not allowed to borrow or lend and that there are no transaction costs involved in the hedge.

Since the martingale and joint normality assumption are so often referred to in literature, Chen *et al*. (2008) set out to investigate in how far those two assumptions actually empirically apply for a set of 25 commodities (amongst which soybean, soy meal, corn and wheat). The tests were conducted on daily, as well as weekly, monthly, 2 month, and quarterly returns. They found that the pure martingale hypothesis held for all commodities, whereas the joint normality assumption was rejected in all but one case. But if the expected value of a commodity return was zero, then this would imply that its beta would be zero. In fact, Dusak & Young applied a CAPM analysis to corn, wheat and soybean and found neither systematic risk nor evidence of risk-adjusted return premium (Bjornson & Carter, 1997).

Second, we see that the rather advanced theories (i.e., the *expected utility* -, the *minimum MEG* - and the *minimum GSV* approach) all represent parametric approaches i.e., they require an analytic expression of the joint distribution of the spot and futures return series. This makes their application infeasible, as, in case of our spot replacement series, we only have around 150 data points to begin with, which renders an analytical derivation of the joint distribution function troublesome.

Third, it appears empirically that the added value of more sophisticated approaches to model the optimal hedge ratio is rather small (Lence, 1995; Lien *et al*., 2002). Generally, Chen *et al*. (2003, p. 433) claim that, while different objective functions will yield different optimal hedge ratios,

> *"[…] there is no single optimal hedge ratio that is distinctly superior to the remaining ones."*

It is probably a mixture of these three reasons that explains why the overwhelming proportion of scientific articles on futures hedging almost exclusively centers on OHR estimation in the context of the minimum variance approach (Harris & Shen, 2003). And it is also why we will focus on the *minimum variance* objective function in an attempt to model the optimal hedge ratio. Using this objective function to show the maximum possible variance reduction, we will, as a secondary measure, also apply the *mean-variance* hedging strategy to inspect whether, and if so, how the series of hedge ratios changes, if we also take account of the futures rate of return.

2.2 Review of MV OHR Estimation Methods

Having settled on the minimum-, and mean-variance approaches as our objective functions in Section 2.1, we can now proceed to reviewing the techniques that have empirically been applied in literature in order to estimate the input parameters for both approaches. As has been noted before, the issue of modeling the OHR has been of resounding interest in literature, featuring a proliferating number of articles, especially in the minimum variance context.

As it is virtually impossible to spotlight the whole array of models with all their adaptions and extensions within the scope of our research, we confine our review to the more recurring models that are frequently used as a benchmark in order to judge the performance of new models and which, as we will see, usually perform just as well, if not even better.

Recall from Equations (2b) and 3(b) that, apart from the constant risk aversion coefficient and the expected futures return, which is simply estimated as the sample average, the involved input parameters are identical. Therefore, any method that models those input parameters will suit both objective functions. To avoid unnecessary duplication, this section only describes the minimum variance case. According to (2b) the variance minimizing optimal hedge ratio is given as:

$$h^* = \frac{\rho_{Rs,Rf}\sigma_{Rs}\sigma_{Rf}}{\sigma_{R_f}^2} = \rho_{Rs,Rf}\frac{\sigma_{Rs}}{\sigma_{Rf}}$$

where σ_{Rs} and σ_{Rf} are the standard deviation of the spot and futures return series, respectively, and $\rho_{Rs,Rf}$ is the correlation between those two return series. Since the variables in (2b) are all unknown population parameters, we need a method to estimate those parameters – and this is what all the upcoming models in Section 2.2 are about. While the methods of estimating those variables may differ, the way that those estimates are used to assemble the minimum-variance OHR is always the one given in (2b), or (3b) in case of the mean-variance OHR.

2.2.1 Ordinary Least Squares (OLS) Regression

The most basic, yet still frequently applied approach to model the OHR is a simple OLS regression.

An OLS regression fits a straight line to the independent and dependent variable observations (in our case the spot and futures returns) in a scatter plot such that the total variance (i.e., the sum of the squared vertical distances between the observations and the straight line) is minimized (Poortema, 2011). As the joint variance minimization is indeed what we are after, the suitability of the OLS method in our context is self-evident.

More concretely, we regress the realized, historic spot rates of return, $R_{s,t}$, against the realized, historic futures rates of return, $R_{f,t}$ That is, we have:

$$R_{s,t} = \alpha + \beta R_{f,t} + \varepsilon_t \tag{4}$$

The model is thus set up such that the realized spot rate of return in period t, $R_{s,t}$, is explained by a constant term α, the impact of the independent variable, i.e., the realized futures rate of return in period t, R_{ft}, the degree of impact of which is captured by β, and an error term, ε_t (Bos & Gould, 2007).

Notice, that we estimate the β in (4) with $\hat{\beta} = \frac{\hat{\sigma}_{R_sR_f}}{\hat{\sigma}_{R_f}^2} = \hat{\rho}_{R_sR_f}\frac{\hat{\sigma}_{R_s}}{\hat{\sigma}_{R_f}}$.

As one can see, the expression for $\hat{\beta}$ is the same as in (2b) except for that the population parameters have been substituted by their corresponding sample parameters. Therefore, according to the OLS method, the optimal hedge ratio is given by the slope coefficient in (4), i.e., $h^* = \beta$.

According to most authors in literature, cointegration poses a problem for the OLS method. Specifically, they argue that if the spot and futures return series were cointegrated then the OLS regression equation as given in (4) would be miss-specified (Chen *et al*., 2003; Hatemi-J & Roca, 2006; Lahiani & Guesmi, 2014). The argument is that in the presence of cointegration, the OLS regression would over-difference the data and cloud the long-run relationship between the two series (Kroner & Sultan, 1993). A better estimate of the OHR could then be attained by incorporating error correction terms in the regression analysis. Hence, the beta from (4) would not be the best linear unbiased estimator ('BLUE') anymore.

However, due to limitations in scope and the controversies surrounding the topic of cointegration (see Appendix B for more information), this is something that is left for future research.

2.2.2 OHR and Time Variation

A much more serious problem with the *OLS* method than a potentially spurious definition of the regression equation due to the omission of a cointegration relationship, is the fact that it is *time-invariant*. As we have seen in Equation (4) the optimal hedge ratio is based on *unconditional* sample moments.

Suppose that the OLS optimal hedge ratio estimation is applied on a window of *n* weekly observations. Of course one will get new, varying estimates of the OHR as the OLS method is applied across time. However, per application, it implicitly assumes that the moments of the bivariate distribution of spot and futures returns do not change *within* the period of time spanning the *n* observations. Put differently, any variation of the OHR observed over time is solely due to sampling error – there is no feature in the unconditional OLS method that otherwise would explain variation over time (Alexander, 2011).

The paramount issue with unconditional modeling approaches is that they are at odds with the *heteroscedasticity* that is oftentimes observed in financial time series (Lien & Tse, 2002). Heteroscedasticity refers to the observation that different degrees of *volatility* usually come in *clusters*. That is, there are relatively tranquil periods (or clusters) where there is little divergence from the mean, and there are periods where returns are rather spiky. Volatility clustering is also termed *autoregressive conditional heteroscedasticity* which stresses the fact that there is usually autocorrelation (cross correlation of a series with its own, lagged values) in the squared returns and that the second moment of the distribution is not constant throughout time. Analog to that, we would also expect the correlation between two assets to vary over time as well (Bos & Gould, 2007).

It is for this reason that there is strong support in literature for *time-varying* approaches that model the moments of the bivariate distribution *conditional* on new price information (Cecchetti *et al*., 1988; Baillie & Myers, 1991; Kroner & Sultan, 1993; Lien & Tse, 2002; Engle, 2002; Chen *et al*., 2003; Bauwens *et al*., 2006; Hatemi-J & Roca, 20006; Bos & Gould, 2007; Hsu *et al*., 2008; Ramlall, 2009; Chang *et al*., 2011). However, while the majority of the researchers advocates the use of time-varying models from a theoretical perspective, empirical validation of the hedging performance improvement of those models over time-invariant models is a much more controversial matter.

Note, that when literature compares the hedging performance of different models, they usually compare what are termed the *in-sample* and *out-of-sample* periods. The in-sample period is that period of the whole sample period from which we draw the data to estimate the parameters of a model. The out-of-sample period then tests the performance of the models on the remaining part of the whole sample. Suppose that we have a total sample of 800 observations and split it into 500 in-sample and 300 out-of-sample observations. The models are parameterized on the basis of the first 500 observations. Of course, we are generally less interested in the in-sample performance of the models, as they have been parameterized in order to increase the likelihood of the in-sample data in the first place. What we are more interested in, is how flexibly and well a model copes with new, out-of-sample data. And it is this period, the out-of-sample period, where time-varying models usually prevail (Kroner & Sultan, 1993; Choudhry, 2003; Ramlall, 2009; Alexander, 2011).

To disclose the digressing nature of the findings in literature on the performance differential between static and dynamic hedge ratio models, Appendix C provides an overview of the settings and outcomes of a number of articles that have empirically tested the relative performance of time-varying over time-invariant models. While some authors document a strict dominance of dynamic models (e.g., Baillie & Myers, 1991; Kroner & Sultan, 1993; Sephton, 1993, Choudhry, 2003; Lien & Yang, 2008), others table the exact opposite (e.g., Lien & Wilson, 2001; Lien *et al*., 2002; Bystrom, 2003, Alexander & Barbosa, 2007, Kenourgios, 2008). This ambiguity is moreover not restricted to the meta-level. We also find differences on the level of individual articles. In the case of Bera *et al*. (1997), Choudhry (2004), Bos & Gould (2007), Hsu *et al*. (2008), Park & Jei (2010), and Kostika & Markellos (2013) for example, the ranking order of static and dynamic models is either interspersed and/or contingent on the sample context (i.e., the rankings of the in-sample and out-of-sample contexts differ).

Methodologically, the variability in markets, time, and applied parameterization approaches make it very difficult to compare the outcomes of the various models to come to a conclusive answer – even after three decades of research (Kostika & Markellos, 2013).

The idea that the performance of time-varying OHR models is contingent on the market environment to which it is applied to, is also supported by Alexander (2007) who approves the usefulness of more sophisticated models in the commodity environment where the basis can be highly volatile due to transportation costs, storage constraints, and logistic problems (Alexander, 2011), but highly questions their use for hedging for example stock indices.

2.2.3 Exponentially Weighted Moving Average (EWMA)

Another shortcoming with respect to the OLS regression method is that it attributes equal weight to all observations in the sample. This makes it stiff and less reactive to more recent information in the sample. Owing to this inflexibility, it is also less useful for providing short term estimates for the moments of a (bivariate) distribution (Alexander, 2011).

The EWMA method seeks to overcome this shortcoming by introducing the exponential smoother, λ, a constant that is between 0 and 1. Assuming that $E[R_s] = E\left[R_f\right] = 0$, the *conditional* variance and covariance of the spot and futures return series are modeled as (Alexander, 2011):

$$\sigma_{R_{f,t}}^2 = (1-\lambda)\sum_{i=1}^{n}\lambda^{i-1}\ R_{f,t-i}{}^2 = (1-\lambda)R_{f,t-1}{}^2 + \lambda\sigma_{R_{f,t-1}}^2 \tag{5a}$$

$$\sigma_{R_sR_{f,t}} = (1-\lambda)\sum_{i=1}^{n}\lambda^{i-1}\ R_{s,t-i}\ R_{f,t-i} = (1-\lambda)R_{s,t-1}\ R_{f,t-1} + \lambda\sigma_{R_sR_{f,t-1}} \tag{5b}$$

As can be seen, the approach models the (co-) variance by both, a component that recursively contains past information of the (co-) variance parameter, as well as a component that captures the latest shock observed in the market. Again, it follows form (2b) that:

$$h_t^* = \frac{\sigma_{R_sR_{f,t}}}{\sigma_{R_{f,t}}^2} \tag{5c}$$

The subscript of the OHR in (5c) signifies that we just entered the realm of time-varying optimal hedge ratios. Notice also how the smoothing constant in (5a) and (5b) indeed provides for the effect that the latest observation is always attributed the biggest weight, as the impact of an observation that is *i* periods back in the past is scaled down by a factor of λ^{i-1}.

For modeling volatility in financial markets, the smoothing parameter λ usually takes on values in the range between 0.75 and 0.98 (Alexander, 2011). As we can see from Equations (5a) and (5b), the higher λ, the more weight is placed on the past volatility parameter, which recursively contains all the past *n* volatility observations and therefore, the *smoother* the estimated volatility series. Lower values for λ on the other hand, ensure that the volatility series reacts more drastically to new shocks in the market.

2.2.4 GARCH Class OHRs

A class of models that perfectly lends itself and is extensively used in literature and practice in order to model the conditional covariance and variance of (two) financial time series is the GARCH class. In fact, the EWMA model that has just been introduced, is a specific case of this class.

GARCH (generalized autoregressive heteroscedasticity) models are a generalizations of ARCH(*p*) models. Under an ARCH(*p*) model, the volatility for period *t* is represented as a weighted average of the past *p* innovations:

$$\sigma_t^2 = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \cdots + \alpha_p\varepsilon_{t-p}^2 \tag{6a}$$

for $\alpha_0 > 0$, $\alpha_1, \ldots, \alpha_p \geq 0$ and where the innovations are residuals from a mean equation, usually of the most simple form:

$$R_t = c + \varepsilon_t \tag{6b}$$

The GARCH(*p*, *q*) model extends Equation (6a) by adding q autoregressive terms:

$$\sigma_t^2 = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \cdots + \alpha_p\varepsilon_{t-p}^2 + \beta_1\sigma_{t-1}^2 + \cdots + \beta_q\sigma_{t-q}^2 \tag{6c}$$

for $\alpha_0 > 0$, $\alpha_1, \ldots, \alpha_p, \beta_1, \ldots, \beta_q \geq 0$. However, usually it suffices to take $p = q = 1$, yielding the vanilla GARCH method that has found overwhelming applicability in financial institutions:

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{6d}$$

For $\omega > 0$ and $\alpha, \beta \geq 0$. It is important to understand that the GARCH model consists of both, the mean equation such as in (6b) and the variance equation such as in (6d). Note also, that for the underlying data generating process to be stationary, we require that $\alpha + \beta < 1$ (Alexander, 2011).[4]

The EWMA approach resembles the integrated GARCH (I-GARCH) model where we have that $\alpha + \beta = 1$. Setting $\beta = \lambda$, which implies that $\alpha = 1 - \lambda$ and dropping the constant term, i.e., $\omega = 0$ then Equation (6d) becomes:

$$\sigma_t^2 = (1 - \lambda)\varepsilon_{t-1}^2 + \lambda\sigma_{t-1}^2$$

which is indeed equal to (5a) as the EWMA assumes zero mean and thus causes the constant term in (6b) to drop and therefore $r_t = \varepsilon_t$.

The generic GARCH(1,1) model is both more efficient and more effective than the ARCH(p) model as the former needs to estimate less parameters, yet has better forecasting qualities (Alexander, 2011). It follows from (6d) that the higher β, the persistency parameter, the longer it takes for volatility shocks to fade away. The reactivity parameter, α, determines in how far volatility is affected by squared errors.

In what follows we provide a review of *bivariate* GARCH models which are just bivariate generalizations of the univariate GARCH(1,1) model and thus look very familiar to expression (6d). All of those bivariate models assume that the *conditional joint distribution* of the *returns* is *normal* with the variance and covariance between the two variables being governed by H_t, the variance-covariance matrix. In our case:

$$H_t = \begin{bmatrix} \sigma_{R_{s,t}}^2 & \sigma_{R_{s,t}R_{f,t}} \\ \sigma_{R_{f,t}R_{s,t}} & \sigma_{R_{f,t}}^2 \end{bmatrix} \tag{7}$$

where, of course, $\sigma_{R_{s,t}R_{f,t}} = \sigma_{R_{f,t}R_{s,t}}$. To distinguish time-varying OHRs, based on conditional second moment estimation methods, from the time-invariant OHRs, we denote them with $h_t^*|\Omega_{t-1}$ where Ω_{t-1} represents the information set up to and including t-1.

The sole goal of all of the upcoming approaches is to model H_t. In fashion of expression (2b), they all construct the optimal hedge ratio as:

$$h_t^*|\Omega_{t-1} = \frac{[H_t]_{1,2}}{[H_t]_{2,2}} = \frac{\sigma_{R_{f,t}R_{s,t}}}{\sigma_{R_{f,t}}^2} \tag{8}$$

[4] See appendix B for an explanation of what stationarity is and what it implies.

2.2.5 Diagonal VECH Parameterization

The VECH GARCH model is the most basic multivariate GARCH model and has originally been proposed by Bollerslev *et al.* (1988). In its more parsimonious, diagonal representation, the covariance matrix, H_t, is given as (Alexander, 2011):

$$vech(H_t) = C + Avech(\xi_{t-1}\xi'_{t-1}) + Bvech(H_{t-1}) \tag{9}$$

where *vech* is the operation that stacks the lower diagonal values of a matrix into a column vector, A and B are diagonal coefficient matrices and C is a coefficient column vector. That is, $C = (\omega_1, \omega_2, \omega_3)'$, $A = diag\{\alpha_1, \alpha_2, \alpha_3\}$, $B = diag\{\beta_1, \beta_2, \beta_3\}$, and $\xi_t = (\varepsilon_{r_{s,t}}\varepsilon_{r_{f,t}})'$.

By restricting the parameter matrices A and B to be diagonal we effectively make sure that each element in the covariance matrix, $H_t, H_{ij,t}$, is only a function of its own lagged value $H_{ij,t-1}$ and $\varepsilon_{i,t-1}\varepsilon_{j,t-1}$ for $i, j = R_s, R_f$. In the bivariate case, the proposed model involves the estimation of 9 parameters (Choudhry, 2004).

A serious shortcoming of the (diagonal) VECH parameterization is that it does not ensure positive definiteness of the covariance matrix (Baillie & Myers, 1991; Engle, 2002; Lien *et al.*, 2002).

2.2.6 (Diagonal) BEKK Parameterization

There are various restrictions that can be imposed in order to ensure positive definiteness of the covariance matrix, H_t. A popular model in that regard is the BEKK model, which builds on the work of Baba, Engle, Kraft and Kroner. It is also the GARCH model that has been used in the seminal paper on time-varying GARCH optimal hedge ratios by Baillie & Myers (1991). The plain version with a single ARCH and GARCH term is:

$$H_t = C'C + A'\xi_{t-1}\xi'_{t-1}A + B'H_{t-1}B \tag{10a}$$

where C, A, and B are 2x2 parameter matrices. While C is upper triangular, A and B are unrestricted. In our bivariate case, (10a) involves the estimation of $N(5N + 1)/2 = 11$ parameters (Chang *et al.*, 2011).

There are a couple of restrictions that can further be imposed on the parameter matrices to reduce the number of parameters to be estimated. The two most cited approaches in this respect are the *scalar* and *diagonal BEKK* parameterization. Since the performance of the latter form of parameterization is often found to be superior (Alexander, 2011) – also in the context of futures hedging (Chang *et al.*, 2011) – it is also the one that we will further consider in our review.

The diagonal BEKK bivariate GARCH model requires the parameter matrices A and B in (10a) to be diagonal resulting in 7 parameters to be estimated (Bauwens *et al.*, 2006):

$$H_t = \begin{bmatrix} c_1 & 0 \\ c_{12} & c_2 \end{bmatrix}\begin{bmatrix} c_1 & c_{12} \\ 0 & c_2 \end{bmatrix} + \begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix}\begin{bmatrix} \varepsilon_{1,t-1} \\ \varepsilon_{2,t-1} \end{bmatrix}\begin{bmatrix} \varepsilon_{1,t-1} & \varepsilon_{2,t-1} \end{bmatrix}\begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix} + \begin{bmatrix} b_1 & 0 \\ 0 & b_2 \end{bmatrix} H_{t-1} \begin{bmatrix} b_1 & 0 \\ 0 & b_2 \end{bmatrix} \tag{10b}$$

In the diagonal BEKK version we can also derive a stationarity condition solely on the basis of the elements of the diagonal matrices. In particular, for an n-dimensional diagonal BEKK model with just one lag, we require that $(a_{ii}^2 + b_{ii}^2) < 1$ for all i (Engle & Kroner, 1995).

2.2.7 Constant Conditional Correlation (CCC) Parameterization

Since parameter estimation can become quite a problematic venture in the context of higher dimensions, there are a number of simplifications and approximations that can be applied to further ease the estimation process. Recall how the covariance matrix in the VECH and BEKK approaches were modeled by *bivariate* GARCH models.

The constant conditional correlation (CCC) proposed by Bollerslev in 1990 simplifies the estimation process by modeling the elements of the covariance matrix H_t by means of *univariate* GARCH processes such as in (6d) (Engle, 2002). The model can be written as:

$$H_t = D_t R D_t, \qquad \text{where } D_t = diag\{\sqrt{h_{i,t}}\} \tag{11a}$$

where the $h_{i,t}$'s are the univariate GARCH processes, and R is the constant correlation matrix, containing the conditional correlations (see Appendix D for a proof). In practice, one frequently uses the unconditional sample coefficient of correlation of the standardized residuals from the involved univariate GARCH processes to model the time-invariant correlation matrix, which is termed the *variance targeting* approach.

2.2.8 Dynamic Conditional Correlation (DCC) Parameterization

Constant correlation is of course, a strong assumption to make. Moreover, variance targeting comprises the disadvantage that, should the sample correlation between the standardized residuals be nonnegative, then it follows from (11a) that there is no way for the conditional covariance to temporarily take on negative values. A negative hedge ratio is therefore impossible. To circumvent this problem, the correlation matrix can be modeled to be time-varying as well. More specifically, expression (11a) is adapted to become (Engle, 2002):

$$H_t = D_t R_t D_t \tag{11b}$$

A simple way to model R_t is to deploy the exponential smoothing method, such that:

$$[R_t]_{i,j} = \rho_{i,j,t} = \frac{q_{i,j,t}}{\sqrt{q_{ii,t} q_{jj,t}}} \qquad \text{where} \tag{11c}$$

$$q_{i,j,t} = (1-\lambda)\left(\varepsilon_{i,t-1}\varepsilon_{j,t-1}\right) + \lambda(q_{i,j,t-1})$$

and where $\varepsilon_{i,t-1}$ is the standardized residual of security i at time t-1. Note that while modeling R_t using expression (11c) indeed gives rise to a dynamic conditional correlation model, it is not the DCC model proposed by Engle (2002) that is frequently referred to in literature when using the term DCC.

2.2.9 A Note on GARCH Models and Symmetry

The GARCH models presented so for all rest on the assumption of symmetry. That is, they assume that bad market news have the same impact on the volatility of an asset as good news. Empirical research, however, suggests that pessimistic information might have a more profound impact on the volatility series of an asset (Brooks *et al.*, 2002) due to leverage effects (Park & Jei, 2010).
In the same vein, Lien & Yang (2008) argue that a positive basis (defined as spot – futures) has a bigger influence on the variance-covariance structure between the spot and futures time series than a negative basis.

However, despite the effort of those authors to develop GARCH models that account for asymmetry, the performance improvement over conventional GARCH models was only marginal at best. In the case of Brooks *et al*. (2002), the out-of-sample hedging performance was even lower than that of the conventional models. Park & Jei (2010) even found profound in-sample deficiencies of some asymmetric models.

Given that the added value of asymmetric models is disputable, especially in our small sample context, and since accounting for asymmetry considerably complicates the model estimation process (Alexander, 2011), the investigation of asymmetry is something that we leave for future research.

2.2.10 Conclusion on MV OHR Estimation Methods

While countless methods have been developed in literature to model the optimal hedge ratio under the minimum variance objective function, Section 2.2 provided a review of the most commonly applied models. Basically, those models can be sub-divided into time-invariant (OLS) and time-varying (EWMA, VECH, BEKK, CCC and DCC) models. Since the former group assumes that the moments of the distribution do not change over time, it fails to account for the heteroscedasticity that is often observed in the volatility of financial time series.

Still, there is some dispute of whether the added value of more advanced time-varying models justifies their increased complexity in the modeling process. Hence, it will be interesting to see how they contest in our context.

As for other restrictions, we have seen that the EWMA method assumes that the expected value of both the spot and futures series returns are zero. While this assumption is not made in the context of the (VECH, BEKK, CCC, and DCC) GARCH models, they do, however, in most cases assume that the residuals are conditionally (jointly) normally distributed as they rely on maximum likelihood estimation in order to extract the model parameters, in which case the assumption of normality simplifies the computations. However, it is shown that even when the assumption of normality of the error terms is violated, then maximizing the Gaussian log-likelihood still provides for quasi-maximum-likelihood estimates, which are normally and asymptotically distributed should the GARCH functions (6b) and (6d) be correctly specified (Bollerslev & Wooldridge, 1992).

Of all the considered models, we will only discard the bivariate VECH GARCH and the CCC model. The former does not ensure positive definiteness and the latter relies on the strong assumption of constant correlation, effectively denying the OHR to take on negative values. We therefore resort to its dynamic extension instead.

2.3 Review of Hedging Performance Parameters

Up to now, we have been talking about hedging effectiveness without having specified exactly how we would measure it. Section 2.3 features an overview of the various performance measures that are proposed and used in literature.

Since we have already settled on our *objective* functions, the pool of performance measures is already refined as it only makes sense to choose from indicators that measure the performance with regard to the *objective* that is sought to be optimized in that function. This insight is backed by Alexander (2007) who states that the question of how to estimate the OHR and how to measure its effectiveness are integrally related.

Within the space of risk-minimizing measures, Ederington's measure of hedging effectiveness proposed in 1979 is the one that has predominantly been applied in literature to date and is given as

$$E = 1 - \frac{\sigma^2_{R_h}}{\sigma^2_{R_s}} = \frac{\sigma^2_{R_s} - \sigma^2_{R_h}}{\sigma^2_{R_s}} \tag{12a}$$

where, as in (1a), R_h is the return of the hedged portfolio, and R_s is the return of the sole spot position, i.e., the unhedged portfolio. Equation (12a) is therefore just the proportion of the variance reduction induced by the futures hedge relative to the proportion of the variance of the unhedged portfolio.

A problem with Ederington's measure is that it is based on *unconditional* sample moments. However, as we have seen in Section 2.2, the time-varying approaches seek to minimize *conditional* variance. It is for this reason that Lien (2005) disapproves the use of Ederington's measure of hedging effectiveness in any but the OLS setting. He consequently contests all research in which alternative estimation approaches are compared to the OLS method based on Ederington's measure.

The reason is simply that, since the OLS method is designed specifically such that it minimizes the unconditional variance, there is no way that any method other than the OLS method beats the OLS method on the grounds of Ederington's measure, since unconditional variance is exactly what it measures. This is strictly true in the in-sample context, but also extends to the out-of-sample context in case both the in – and out-of-sample sizes are sufficiently large and if there is no structural change in the two sample periods (Lien, 2005).

While Lien explicitly claims that the prevalence of OLS methods in terms of unconditional minimum variance performance measures holds in comparison with any static or dynamic OHR estimation method, Kroner & Sultan (1993) say that this only applies to the static realm. In fact, there are numerous papers that empirically observe a bigger unconditional variance reduction of dynamic estimation methods over the OLS method – even in the in-sample period (e.g., Baillie & Myers, 1991; Kroner & Sultan, 1993; Bera *et al.*, 1997).

Moreover, Lien (2006) also argues that Ederington's measure yields downwards biased estimates of the hedging effectiveness and thus undermines the true value of futures contracts as hedging vehicles. It follows that we have to be cautious when using Ederington's measure to compare the performances of the different models.

To account for the uncondtionality shortcoming, Alexander (2007) simply swaps the unconditional moments in (12a) for the conditional moments, which of course yields

$$E_t = 1 - \frac{\sigma^2_{R_{h,t}}}{\sigma^2_{R_{s,t}}} = \frac{\sigma^2_{R_{s,t}} - \sigma^2_{R_{h,t}}}{\sigma^2_{R_{s,t}}} \quad (12b)$$

As for the mean-variance objective function, the effectiveness of a hedging strategy will be based on its location on the risk-return plane introduced in Figure 2.1

Apart from focusing on the effectiveness of a hedging strategy, i.e., the degree to which it is capable of reducing the variance of the unhedged portfolio, it is reasonable to also shed light onto the efficiency of the model. In case of the time-varying models it can be expected that the corresponding OHRs are more volatile, inducing the need to more frequently rebalance the hedged portfolio, which is costly as each rebalancing act requires a variable amount of transaction costs (Sultan & Kroner, 1993; Harris & Shen, 2003; Hatemi-J & Roca, 2006; Alexander, 2011).

2.4 Conclusion on Literature Review

Being armed with the insights acquired in the literature review, we are now able to answer Sub Questions 2 and 4.

> *Sub Question 2: "Which models are available to model the OHR and what restrictions do they pose on the variables involved in the hedge?"*

In order to specify a set of OHR estimation models, we first of all had to settle on an objective function – a function that translates risk and return into the investor's utility. From the set of six functions that have been reviewed, we have chosen *variance minimization* as our primary objective function as this is what all the other functions reduce to if we assume a martingale process of the futures price series (*mean-variance* approach, *Sharpe-ratio* approach), or if we assume joint normality of the spot and futures return series (also including the *expected utility*-, *minimum MEG* -, and the *minimum GSV* approach). Moreover, it is nonparametric - an important feature given our limited number of observations (Chen *et al*., 2003). As a secondary measure, we will also deploy the *mean-variance* objective function to check whether the hedge ratios indeed remain unchanged upon introducing the average futures return in the optimization problem, as assumed by the minimum-variance approach.

Suitable for both objective functions, we identified four models that we will apply in the analysis, namely the OLS regression, the EWMA, the Dynamic Conditional Correlation, and the BEKK GARCH model. The OLS model is a time-invariant model and hence does not account for the heteroscedasticity that is usually found in the volatility of financial return series. That is, it assumes homoscedasticity. While the dynamic EWMA does not make this assumption, it does, however, assume that the expected value of the return series is zero.
The time-varying GARCH models do neither assume homoscedasticity nor zero mean. They do on the other hand, usually but not necessarily, assume that the residuals of both series be conditionally

(jointly) normally distributed. Violation of this assumption, however, does not immediately invalidate the use of those models on a theoretical level.
Moreover, should the underlying data generating process be stationary, then the sum of the (squares of the) ARCH and GARCH parameters in the (diagonal BEKK) GARCH model ought to be smaller than unity.

Sub Question 4: "On the basis of which indicators do we measure the performance of a hedge?"

In order to measure the hedging performance in the minimum variance context, literature almost exclusively applies Ederington's measure of hedging performance, which simply divides the reduction in the hedged portfolio (including the futures hedge security) by the variance of the unhedged portfolio (i.e., the sole spot security). The way that this measure was originally proposed, it favors the OLS method as it measures the effectiveness in terms of the unconditional sample moments, which is exactly what is optimized by the OLS method. To account for this shortcoming, Alexander came up with a measure that was similar to that of Ederington except for that she swapped the unconditional for the conditional sample moments.

Moreover, since the EWMA, BEKK GARCH and DCC GARCH will vary over time and thus require the investor to rebalance the portfolio over time, it is reasonable to also take into account the impact of transaction costs in judging the performance of the four models. As for the mean-variance context, the different hedging strategies will be judged by their location in the risk-return plane.

Chapter 3 – Methodology

In Chapter 2 we have described the different estimation approaches that will be applied in order to model the OHR. This chapter describes exactly on *what data* those models are applied (Section 3.1) and the *tools* that are used in order to derive and implement those models (Section 3.2). Finally, we will describe exactly how the unconditional and conditional minimum variance hedging effectiveness is measured (Section 3.3).

3.1 The Data

The primary objective of De Heus is of course to hedge spot market price fluctuations. As has been mentioned in the introduction, the purchasers of De Heus make weekly price estimates for a number of commodities, which indicate what it would cost to buy a certain unit of a commodity in the spot market, for spot delivery or delivery up to 12 months in the future. This is what they call the *replacement values*.

As was also mentioned in Chapter 1, this research focuses on the products *corn*, *wheat*, and *soy*. There are however, different kinds of corn, wheat, and soy (by-) products that are processed at the plants of De Heus, which differ in terms of form, quality and origin and which have their own, distinct product codes. We therefore clarify that we represent *corn* by the product code *0003*, *wheat* by *0041* and *soy* by *0300*. All replacement products are denoted in € per metric ton (MT).
A primary issue concerned with the internal replacement values is that their weekly records just started in December 2011, giving us, at the time of writing, a sample size of 158 observations per product. Keeping in mind that we preferably would want to split the total sample into an in-sample and out-of-sample period (see Section 2.2.3), the sample size is rather thin.

It is for this reason, that we also apply our analysis to weekly spot data from external parties (Oil World, and Reuters). While hedging the replacement values is still the prime target, it will be interesting to see how the models behave in the context of bigger sample sizes.
To represent *corn*, we chose the product "*French Corn FOB Rhine*" from the Reuters database. *Soy* is sought to be mirrored by "*Soya pellets 48% Brazil, cif Rotterdam*", which is also in the high protein range. Both price series already include the costs to get it to ports or on waterways in the Netherlands and can thus both be considered prices for the local spot market. This is, unfortunately, not true for the product "*Wheat, U.S., No.2, SRW, fob Gulf*" to which we had to resort to in order to represent *wheat*.

Product category	Product name	Currency	Unit	Interval	n
Corn (Replacement)	0003 (GMO corn)	EUR	MT	Weekly	158
Corn (Reuters)	French Corn FOB Rhine	EUR	MT	Weekly	677
Wheat (Replacement)	0041 (Wheat boat)	EUR	MT	Weekly	158
Wheat (Oil World)	Wheat, U.S., No.2, SRW, fob Gulf	USD	MT	Weekly	597
Soy (Replacement)	0300 (Soya Hipro ADM GMO)	EUR	MT	Weekly	158
Soy (Oil World)	Soya pellets 48% Brazil, CIF Rott.	USD	MT	Weekly	670

Table 3.1 – *Overview of spot sample data.*

The futures data will briefly be introduced once we have decided on our futures contract based on the correlation study in Section 4.1.

Hedging Horizon

As is noted by numerous authors, the length of the interval that is used to calculate the returns of the price data and on which the models are applied, should ideally be of the same length as the time interval for which you want to hedge the commodity (Lien & Tse, 2002; Chen, 2003; Hull, 2012). That is, if you want to calculate the optimal hedge ratio to hedge a spot exposure for the coming 4 weeks, then you should also estimate your model based on monthly return data.

The problem with that undertaking is that the return intervals in the sample are not allowed to overlap, which causes a serious reduction of the sample size if we were to expand the horizon. For example, if we were to estimate a model for a hedging horizon for, say, 2 months, then our replacement sample size would further shrink down to 20, which impedes any meaningful analysis. If, on the other hand, the return periods did overlap, then the observations would be autocorrelated and the model estimates would be biased (Chen *et al*., 2003).

Still, insights from hedging in the monthly horizon setting are very much appreciated as long cash positions at De Heus are usually maintained for a longer period of time. Therefore, apart from conducting the analysis on the basis of weekly returns, we will also consider monthly returns. However, given the aforementioned concerns for sample size reductions, the latter analysis will only be conducted for the external Reuters and Oil World series.

Possessing both weekly and monthly return data, we can then investigate the effect of different hedge horizons on the effectiveness of the hedges. Note that Lien & Tse (2000) have shown that under the *stability under aggregation* property (which is satisfied by the constant OLS regression method) the optimal hedge ratio *theoretically* holds for whatever hedging horizon.

Sample Splitting

In order to test the performance of our models, we first of all split each sample into an in-sample and an out-of-sample part, where the former is used to parameterize the model. Once the model is estimated, it is then applied on the complete sample.

Finding a useful point to cut the full sample into two parts asks for finding a reasonable balance between having as much observations as possible to estimate a stable model on the one hand, and yet having enough observations to test the models in an out-of-sample setting on the other hand.

For the weekly return samples, we believe that in case of the replacement values, this balance is found at n = 130. For the remaining spot price series that cut is made at n = 500. As for the monthly return samples, there will not be any out-of-sample period at all.

Futures Contract Rolling

As the trading activity in the front month contract usually decreases considerably as it approaches its maturity date, a number of researchers consider it good practice to roll it over to the next maturity contract one or two weeks before the front month contract expires (Kroner & Sultan, 1993; Harris & Shen, 2003; Lien *et al*., 2002). This is also done in this research, because for various products there are indeed a lot of instances in the database that feature dips in the prices of the front month contract around maturity.

Table 3.2, showing an extract from the Matif corn database, illustrates a case in point of the. If we did not roll the front month contract (EMAc1 in this case) prior to maturity (5th of March), then the automatic contract rolling would cause a jump from 170 to 182.75 in the lead month series. If however, we manually rolled the contract from EMAc1 to EMAc2 two weeks before, i.e., on the 19th of February, the price jump would only be 6 (from 173.25 to 179.25). The custom tool has been programmed such that it makes the contract switch 14 days before expiry. It is found that for quite a few front month contracts, the correlation with local spot commodity series improved around 5 to 10 per cent.

1	Date	EMAc1	EMAc2
3681	18-02-2014	173	178.25
3682	19-02-2014	173.25	179.25
3683	20-02-2014	174.25	179.75
3684	21-02-2014	174	180
3685	24-02-2014	173.25	178.5
3686	25-02-2014	172.25	178.25
3687	26-02-2014	172.5	178.25
3688	27-02-2014	172.5	178
3689	28-02-2014	169.25	177
3690	03-03-2014	172.75	181.75
3691	04-03-2014	170.75	181.75
3692	05-03-2014	170	183
3693	06-03-2014	182.75	187.25
3694	07-03-2014	184	189.25

Manual contract rolling

Automatic contract rolling

Table 3.2 – *Dip in Matif corn front month series (maturity date 5th of March).*

Another alteration that has been made to the original series is the exclusion of return outliers that are *caused* by automatic *contract rolling*. To clarify what we mean by that consider for example Figure 3.1, which shows the corn Matif prices for each of the six continuations. Suppose that there is a market consensus that due to fundamental changes, there will be plenty of corn supply in the market starting in period t_F, one year from the present.

Suppose also, that period t_F is not currently included in the contract lives of any of the traded maturities. As time passes, however, there will be a rollover date at which the new contract initiation of the most distant maturity contract (in this case EMAc6) will include period t_F. At that rollover date, we would observe a drastic price adjustment in the contract. And as time further passes, that price adjustment cascades through to the following contracts at each rollover date (EMAc6 becomes EMAc5, EMAc5 becomes EMAc4 and so forth) until it reaches the front month contract. As we can see in Figure 3.1 there are no such adjustments for the second half of 2013 and for the entire year 2014, which suggests that those adjustments are indeed irregular events.

The point is that those rollover cascading effects are predictable. It thus makes sense to exclude them from our analysis as they would otherwise seriously downward bias our hedging effectiveness, irrespective of the applied method.

3.2 The Tools

We apply two tools in our research. The first is EViews – an econometric software that will be applied in order to derive the (OLS, BEKK & DCC) model parameters based on the in-sample data. When estimating a (bivariate BEKK) GARCH model, the software provides for the option to choose between a multivariate normal and a multivariate Student's t-distribution of the error terms. We are therefore not confined by the assumption of joint normality.

EViews

The way that EViews –or any other econometric software – estimates the parameters of the GARCH models is by deploying numerical optimization algorithms (in case of EViews the Marquardt or Berndt-Hall-Hall-Hausman algorithm) to maximize some likelihood function. More specifically, the likelihood function takes as input the functions given in (10b) and (6b) & (6d) for BEKK and DCC, respectively as well as the in-sample observations. The algorithm then fits the model by virtually attempting to maximize the likelihood of the data observed in the sample by numerically optimizing the likelihood function with respect to the model parameters.

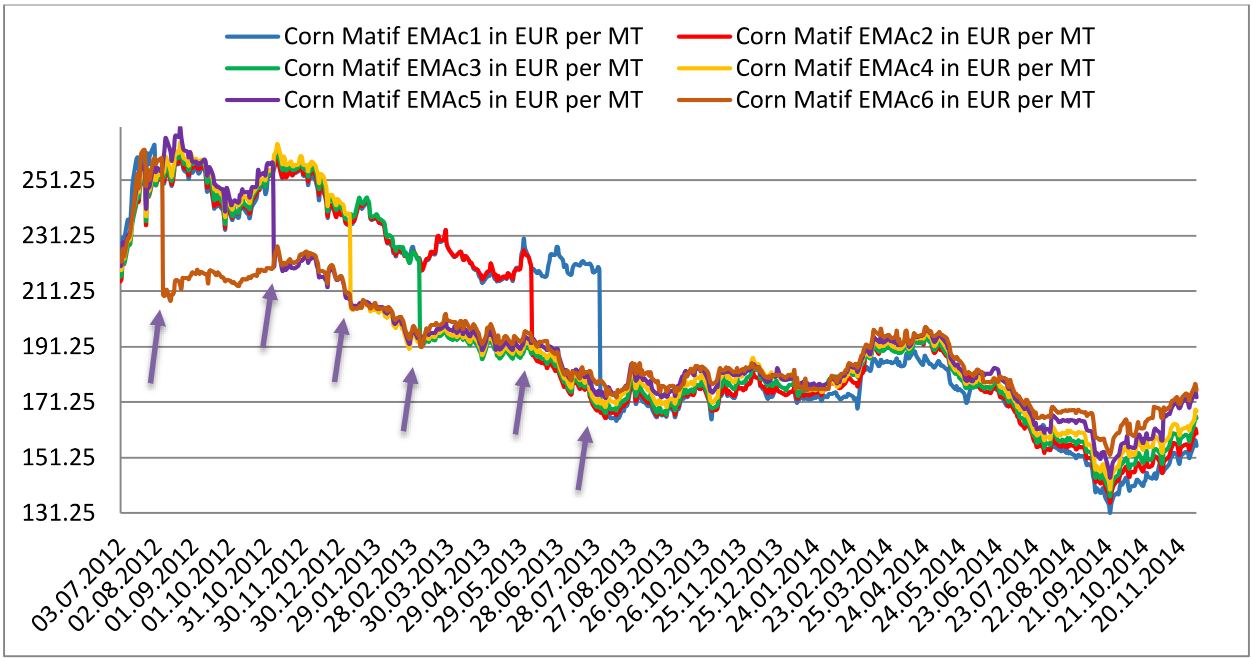

Figure 3.1 – *Matif Corn and rollover cascading.*

As has been explained in Section 2.2.9 conditional normality is often assumed to make those optimization calculations traceable. However, it will be established in Section 4.2 that the assumption of normality is significantly rejected for all series – at least for the unconditional case. Still, as is also shown in Section 2.2.9, even under violation of the assumption of conditional normality, optimization of the Gaussian log-likelihood function still yields quasi-maximum likelihood estimates (QMLE), which are consistent and asymptotically normal when the GARCH mean and variance functions are correctly specified.

It is found in the analysis that enforcing the assumption of a conditional t-distribution in some cases results in unstable parameters estimates – especially in the small sample context. For example, it will be shown in Section 4.2 that all return distributions are stationary. Yet, in some cases, the sum of the estimated ARCH and GARCH parameters is found to be bigger than unity, in which case the assumption of a mean-reverting volatility process would not hold.[5]

Moreover, in cases where estimation under the t-distribution does result in consistent estimates, they are found not to differ substantially from the QMLEs. Consequently, unless stated otherwise, all estimates will be based on conditional normality of the errors.

[5] A mean reverting volatility process is a process, where, over time, the volatility tends to its average level and does not follow any trend. It is thus a direct implication of a stationary process, where the moment statistics do not vary across time.

To ensure quasi-maximum likelihood estimation, we check the option 'Heteroskedasticity consistent covariance (Bollerslev-Wooldridge)' for the coefficient covariance.

Custom Tool

The second tool is the custom application that we have programmed in VBA Excel. For a picture of the interface assembled from screenshots, please refer to Appendix E. The program can best be described as a *time series manager*, which lets you query price information from databases, construct new time series from those information and plot them on a graph. We have set up structured databases for the internal replacement values and external sources including but not limited to the already mentioned Oil World and Reuters, the latter of which is also our supplier of real time FX and (commodity) exchange market price quotes.

While there are more features and functionalities, the program can roughly be split into three parts. In the first part, you choose your base products – i.e., the basic price series of any internal or external product, including FX series. You can further edit a given selection by changing its maturity, or converting its unit or currency.

Once a pool of base products has been selected, one may create linear combinations (weighted subtractions, additions, divisions, or multiplications) of those base series to create new series. Somewhat hidden in the interface, the application also lets you create and plot a volatility series from the returns of the base and newly constructed series based on variable rolling window input.

Those new series may further be combined with the base series of part one to create correlation series in part three. Any combination is possible – this is true for the spread as well as for the correlation part. Part three moreover lets you specify the time interval that ought to be applied in order to calculate the returns as well as the rolling window size based on which the correlations are calculated per point in time. It is important to stress that while the graphs only plot the *level* prices of the series of part one and two, the volatility and correlation computations are based on the *returns* of those series.

Research Execution

The two models thus perfectly complement each other. While EViews provides the parameter estimates (i.e., the regression coefficient β in (4), the BEKK coefficient matrices A, B, and C in (10a), and the univariate GARCH estimates of the DCC model in (11b)) for the different models, the custom application takes those parameters as input to calculate the return series of the hedged portfolio.

In case of the OLS approach, this series can be calculated right away using the built-in spread calculation feature. In the context of the time-varying models, we first of all have to combine the model parameters with the observations of the base series in order to calculate a set of series containing the second moment (i.e., (co-) variance) forecasts per point in time and per series so as to consequently construct the series of time-varying optimal hedge ratios.

In the case where the hedged portfolio includes non EUR - denoted prices, the built-in FX converter of the custom model is applied to translate those prices back to EUR. There are two possibilities to approach the FX problem in the hedging context. We could either *a)* conduct all the OHR calculations on the basis of the series' base (i.e., foreign) currencies to receive some value h_t' and then translate the return of the hedged portfolio back to the domestic currency, or we could *b)* first convert the return series to the domestic currency and then base the OHR calculation on the translated series to generate some value h_t'' for the OHR.

Since the (co-) variance structure will be slightly altered after currency translation we have that, in general, $h'_t \neq h''_t$. However, if we translate the currency *ex post* then the FX component causes variability in hedge portfolio return, which is not accounted for in the models. Since variance minimization is what we want to achieve in the first place, it makes more sense to choose for alternative b) – and thus optimize the portfolio on the basis of the translated series, which is indeed what is also done in our thesis.

Note, that the Oil World wheat commodity spot exposure is not translated to EUR, as the price quotes refer to the US cash market. The series is rather used in order to compare the hedging effectiveness in different countries.
Since the application conveniently prints any time series (thus also the hedge portfolio return series) we can then carefully analyze the statistics of the new portfolio and judge both the conditional as well as the unconditional performance of the different hedging strategies.

EWMA and GARCH Model Initialization
Since all the time-varying methods work recursively, we somehow have to manually initialize the starting value –after all, there is no *t-1* at t=0. If we have not yet collected any price data, then there is no past information on which to apply the econometric models. To make sure that the initialization phase is not characterized by excessively volatile OHR estimates, the moment estimates for t=0 are simply opted to be the sample moments.

3.3 Minimum Variance Hedge Effectiveness Measurement

The application of each of the four hedging methods will result in distinct streams of futures returns and therefore different outcomes of the hedged portfolio return per point in time, $R_{h,t}$. Per commodity and per sample, the unconditional hedge effectiveness is simply the variance reduction measured by the unconditional sample variance of the realized spot and hedge portfolio returns:

$$\hat{E} = \frac{\hat{\sigma}^2_{R_s} - \hat{\sigma}^2_{R_h}}{\hat{\sigma}^2_{R_{s,t}}} \tag{13a}$$

To measure the conditional variances at point *t*, we apply the EWMA method with a smoothing constant of 0.94 on the realized spot and hedge portfolio returns- similar to Alexander (2007). This yields *n* values for the hedging performance for each point in time, i.e.:

$$\hat{E}_t = \frac{\hat{\sigma}^2_{R_{s,t}} - \hat{\sigma}^2_{R_{h,t}}}{\hat{\sigma}^2_{R_{s,t}}} \tag{13b}$$

The overall conditional hedge effectiveness is then measured as an equally weighted average:

$$\bar{E}_t = \sum_{i=1}^{n} \frac{1}{n} \hat{E}_t \tag{13c}$$

Chapter 4 – Analysis

Having described the data and models in Chapter 3, this chapter then empirically applies the models identified in Chapter 2 on those data in order to answer the central research question – that is, to identify the optimal hedging strategy per type of commodity. The analysis is structured into 4 parts. It *starts off* with the correlation study, which aims at identifying the most appropriate hedging security per type of product with the aid of our custom tool. It thus tackles *Sub Question 1*.
Section 4.2 follows with a brief review of the statistical characteristics of the different futures and spot series and discusses their implications regarding the application of the different models, i.e., it is concerned with *Sub Question 3*.
Armed with those findings, we then proceed to estimate the different model parameters in *Section 4.3*.
In *Section 4.4* we finally apply the different models to the hedge portfolios identified in Section 4.1, based on the parameters estimated in Section 4.3. We discuss the performance of those models both in terms of conditional as well as unconditional effectiveness measures.
Since the outcome of the different models is contingent on the applied input parameters it would be interesting to see how the results vary with those input parameters. The robustness of the proposed strategies is investigated in *Section 4.5* and thus answers *Sub Question 5*.
Do the futures commodities really all follow a martingale process? In light of the findings from Section 4.2, *Section 4.6* pursuits to find out in how far consideration for the average portfolio return might alter the strategies proposed under minimum-variance.
The analysis is concluded by *Section 4.7*, which discusses the main findings of the analysis in the light of the central research question.

4.1 Correlation Study – Identification of the Hedge Securities

This section is concerned with identifying which futures contracts best mimic the various spot commodities. Recall that we want to hedge the spot exposure of three types of products, namely *corn*, *wheat*, and *soy* – and that per commodity we have two different samples to represent the spot positions, which are the replacement values of De Heus on the one hand, and the series from external data vendors on the other hand. An overview of all the available agricultural futures commodities in our database is given in Table 4.1. Note that the database also includes the Matif Malting Barley contract as well as commodities from the Dalian commodity exchange. However, the correlations of all of those series with local spot exposures are not mention worthy. Moreover, there was no trading activity in the Matif Malting Barley contract at the time of writing.

As mentioned earlier, we consider both basic –, but also cross –, and inter- time hedges. Using the custom application, we therefore correlate the weekly returns of each spot commodity series with each maturity contract of each futures product, yielding 408 different combinations as we have 68 different futures maturity contracts in total and 3 different spot products, which are represented by two different samples. In the correlation matrix we therefore only show the correlation coefficient of the maturity contract with the highest correlation per spot – futures pair.

Product	Exchange	Contracts	Currency	Unit	Size
Corn	CBOT	Cc1 to Cc6	USc	Bsh	5,000
Corn	Matif	EMAc1 to EMAc6	EUR	MT	50
Wheat	CBOT	Wc1 to Wc6	USc	Bsh	5,000
Milling Wheat	Matif	BL2c1 to BL2c5	EUR	MT	50
Soybean	CBOT	Sc1 up to Sc8	USc	Bsh	5,000
Soybean Meal	CBOT	SMC1 up to SMC9	USD	STn	100
Soybean Oil	CBOT	BOc1 to BOc9	USc	Lbs	60,000
Rapeseed	Matif	COMc1 to COMc4	EUR	MT	50
Oats	CBOT	Oc1 to Oc6	USc	Bsh	5,000
Crude Palm Oil	CBOT	POc1 to POc9	USD	MT	25

Table 4.1 – *Futures contracts overview.*

Note that in case the correlation of a more distant maturity contract is only marginally higher than that of a front month or nearby contract, then the preference is given to the closer maturity contracts as they, in general, feature a higher trading liquidity, which is of course also a major concern when it comes to actually implementing the hedging strategies in practice. In general, we have only considered those maturities that were also actively traded at the time of writing.[6] The results of the correlations are summarized in Table 4.2.[7]

Note, that the numbers in brackets indicate the correlation of the monthly, non-overlapping, return periods. For each column (i.e., spot exposure) the contract in bold refers to the futures contract, which has the highest correlation with that spot exposure.

A number of things can be observed upon studying the correlation matrix. First of all, there are no big surprises as for the spot-futures combinations that yield the highest correlations. A corn spot position is best mirrored by a futures corn product. The same holds for the other two product categories. Thus, no cross-hedges were found to outperform the rather obvious, basic combinations.

Second, domestic spot commodities are best mimicked by derivatives traded on domestic futures exchanges. To a certain extent, this was also to be expected. While the market for the agricultural commodities considered in this analysis is a global one, transportation costs, levies as well as domestic supply and demand factors still enact distinct price dynamics. Yet, weekly correlations as low as .37 (wheat CBOT and wheat replacement), .35 (corn CBOT and corn replacement) or even .22 (corn CBOT and French corn FOB Rhine) are remarking.

Third, in most cases, the monthly correlations are stronger than the weekly correlations, which also makes sense. Financial variables carry both a systemic and an idiosyncratic risk component. While variables may diverge in the short term, systemic market forces are likely to make them move in the same direction as we stretch the time horizon. This reasoning applies to our case since all the variables considered represent agricultural commodities.

[6] For all the Matif contracts, contract continuities up to + 4 have been traded. As for the CBOT products, corn, soybean, and wheat traded actively for continuities up to + 5, soybean oil and soybean meal traded up to + 7, and oats was reasonably liquid only for the first two contract maturities.
[7] All displayed correlations are based on the base series (i.e., without FX conversion) of the products.

Futures	Corn		Wheat		Soy	
	Rep.	Reuters	Rep.	Oil World	Rep.	Oil World
Corn CBOT	.35 (.74) Cc5	.22 (.44) Cc1	.34 (.60) Cc2	.51 (.51) Cc1	.30 (.32) Cc3	.34 (.49) Cc1
Corn Matif	.50 (.71) **EMAc2**	.67 (.79) **EMAc2**	.46 (.65) EMAc2	.52 (.61) EMAc2	.34 (.39) EMAc1	.28 (.49) EMAc1
Wheat CBOT	.29 (.61) Wc4	.31 (.51) Wc2	.37 (.58) Wc1	.79 (.92) **Wc1**	.16 (.40) Wc1	.25 (.46) Wc1
MW Matif	.32 (.52) BL2c2	.41 (.62) BL2c1	.58 (.77) **BL2c1**	.62(.75) BL2c1	.28 (.48) BL2c1	.25 (.45) BL2c1
Soybean CBOT	.21 (.60) Sc2	.23 (.34) Sc1	.35 (.57) Sc2	.39 (.41) Sc3	.48 (.39) Sc1	.62 (.71) Sc1
SBM CBOT	.23 (.53) SMc2	.21 (.34) SMc1	.28 (.46) SMc2	.33 (.37) SMc3	.52 (.49) **SMc1**	.65 (.78) **SMc1**
SBO CBOT	.14 (.43) BOc1	.21 (.26) BOc1	.26 (.43) BOc1	.32 (.35) BOc1	.12 (.11) BOc1	.34 (.46) BOc1
Rapeseed Matif	.29 (.60) COMc1	.30 (.42) COMc1	.41 (.72) COMc2	.39 (.45) COMc1	.28 (.51) COMc2	.34 (.46) COMc2
Oats CBOT	.14 (.40) Oc2	.15 (.27) Oc2	.14 (.25) Oc2	.31 (.32) Oc2	.23 (.34) Oc2	.18 (.38) Oc2
CPO CBOT	.16 (.37) CPOc1	.26 (.30) CPOc1	.20 (.33) CPOc1	.33 (.37) CPOc1	.07 (.14) CPOc1	.22 (.29) CPOc1

Table 4.2 – *Spot Futures Correlation Matrix.*

Fourth and most importantly, the *weekly* correlations of the internal replacement spot series with respect to their futures pendants are somewhat restrained. While they all hover around 50% (50% for corn, 58% for wheat and 52% for soy), and thus provide something to work with, they fall about 15% short of the correlations inherent in the spot series given by the external vendors (i.e., 67% for *French corn FOB Rhine*, and 65 % for the *soya pellets 48% Brazil, cif Rotterdam*).

Why is that? If we take a closer look at the 16 weeks moving average correlation between the wheat replacement and the BL2c1 futures contract in Figure 4.1, then we see that the correlation moves within its high – low range with the plain price levels matching each other quite well for most of the time. There are, however, three dips (i.e., diverging price movements) that distort the correlation relationship. Consider for example the period around end of May 2013, where a sharp and immediate price correction takes place in the futures contract. The wheat replacement series, however, takes a lot more time to adjust. A similar instance is witnessed in the period between the start of June and start of August 2014. Here, the futures prices slides down in a steady slope. The replacement series, however, even slightly increases in the first four weeks, but then sharply drops in the four weeks to follow. In the period around September 20012, the overall price pattern of the two series matches. However, the Matif contract is excessively volatile. This degree of volatility is not matched by the replacement series, causing the rolling average correlation to drop to a level as low as 21%.

If those periods were filtered out from the analysis, then the unconditional correlation coefficient would increase considerably.

Getting back to the correlation matrix, we observe that in general, the correlations of US spot commodities are greater than those of the EU spot commodities. The weekly US wheat FOB Gulf return series for example features a 79% (92% even on a monthly basis) correlation with the CBOT Wc1 futures return series. While not displayed in this table, the same results are also found to hold for other US spot commodities such as corn and soybeans. Possibly, this can be ascribed to either the role of the CBOT as the leading market place for agricultural commodity futures, providing a degree of liquidity that the Matif cannot live up to, and/or the fact that the US is the leading producer of corn, soybeans and other agricultural commodities.

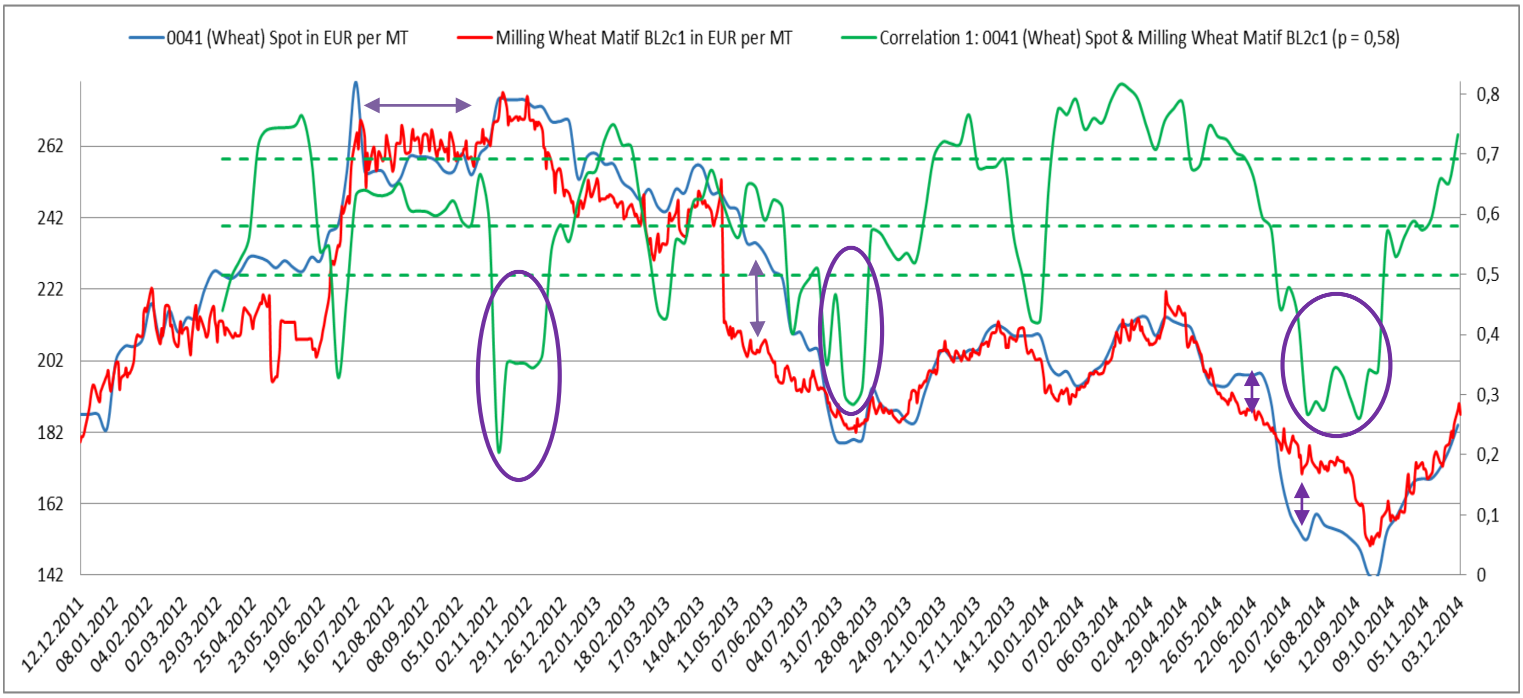

Figure 4.1 – *Illustration of correlation dips for the series wheat replacement and wheat Matif.*

Having elaborated on the correlation matrix, we conclude this section by answering Sub Question 1:

> *Sub Question 1: "Of all the agricultural futures products traded on the CBOT and Matif, based on the degree of correlation, per cash commodity, which futures product qualifies as the most suitable hedging security ?"*

Judged by the unconditional coefficient of correlation, we match both spot corn commodity series with the EMAc2 corn Matif contract. In case of soy, we trace both series by the SMc1 CBOT soybean meal contract. Finally, the replacement wheat and the Oil World wheat spot series are mimicked by the BL2c1, and Wc1, respectively. Take note, that the 0300 soy and SMc1 hedge portfolio is the only case in which the spot and futures contracts are denoted in different currencies. If we translate the SMc1 returns from USD to EUR, then the correlation between the two variables increases to .55 (.51) in the case of weekly (monthly) returns. Likewise, the soy Oil World and SMC1 correlation increases to .66 (.79) after translating both series from USD to EUR.

4.2 Discussion of Descriptive Statistics

In this section, we will discuss the basic statistical properties of the *univariate* spot and futures return series such as their mean, variance, skewness, and kurtosis. Moreover, as we have seen in the model descriptions in Chapter 2, the OLS method assumes that the variance process is homoscedastic. We will therefore also investigate the degree of autocorrelation in the different series.

In terms of GARCH model estimation, it has been pointed out that normality of the error distribution is usually incorporated in the GARCH parameterization process. We thus also check in how far this assumption holds. Finally, we have seen that additional restrictions on the parameters in the GARCH model are applied if the underlying data generating process is deemed stationary. Whether or not this is actually the case is also being statistically investigated.

The sample statistics provided in Table 4.3 are all based on non-overlapping, weekly return intervals. As for the sample start and end dates, all of the futures return series start on 03.01.2000 and end on 03.12.2014 (resulting in 753 observations for the EMAc2 & BL2c1 series and 706 observations for the SMc1 and Wc1 contract series). The replacement return series all start on 19.12.2011 and terminate on the 22nd of December, 2014, yielding 158 observations per return series). While the Oil World soy and wheat contracts have similar start and end dates (25.10.2001 and 18.12.2014, respectively) the wheat series has fewer observations (599) than the soy contract (673) as the former has gaps in its database.

To keep the discussion concise and to preserve space, we will not discuss each series individually, but rather describe the EMAc2 series as a case in point. The descriptive statistics, histograms, and correlograms (with lags from 1 to 5) of the other series can be found in Appendices F and G, respectively. An overview of the descriptive statistics is given in Table 4.3.

Series	Mean	Std. Dev.	Skewness	Kurtosis	Jarque-Bera	Q-Stat	ADF
EMAc2	0.000640	0.0277	-0.5610	8.5046	990.15*	11.006*	-23.83*
BL2c1	0.001113	0.0336	-0.3025	7.6052	676.90*	20.038*	-26.46*
SMc1	0.002507	0.0473	-0.4662	5.5477	216.51*	25.086*	-28.13*
Wc1	0.002318	0.046	0.5564	4.3389	89.16*	2.2041	-26.15*
0003	-0.001441	0.0292	-0.2646	5.5749	45.49*	7.2182*	-11.80*
0041	0.000233	0.0285	0.1509	6.0013	59.90*	10.312*	-10.80*
0300	0.002791	0.0324	0.9805	6.7955	120.15*	0.1879	-11.51*
Corn Reuters	0.000730	0.0330	-0.0833	8.3061	814.94*	0.7164	-24.80*
Soy OW	0.001764	0.0346	-0.1303	4.4127	57.87*	2.3550	-25.01*
Wheat OW	0.002915	0.0546	1.3660	14.256	3348.62*	13.218 *	-27.24*

Table 4.3 – *Descriptive statistics.*[8]

[8] A star (*) indicates that the null-hypothesis is rejected at the 1% significance level.
Jarque-Bera test null-hypothesis: Data is normally distributed.
Ljung-Box Q test (k) null-hypothesis: Data is independently distributed, i.e., features no autocorrelation at lag k.
Augmented Ducky Fuller test null-hypothesis: Data contains a unit root, i.e., is I(1) and thus not stationary.

Now, if we take a first look at the histogram of the weekly EMAc2 return series in Figure 4.2, we immediately recognize that the returns are far from being normally distributed. While the *unconditional* distribution is almost symmetric and only slightly skewed to the left (which is captured by the skewness values of -.560981) it has a lot of probability weight in the tails of the distribution.

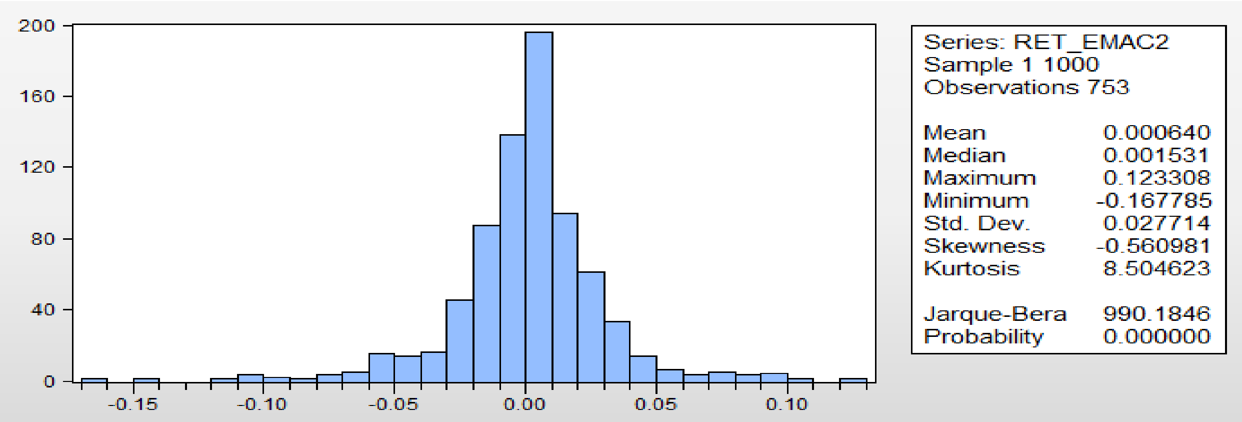

Figure 4.2 – *Statistics of the EMAc2 contract (H_0: Data is normally distributed).*

The fatness of the tails of a distribution is measured by the kurtosis statistic. For a normal distribution the reference value of this statistic is 3. The EMAc2 series therefore features a great deal of excess kurtosis (5.504), which is oftentimes observed in financial data. Based on the Jarque-Bera sample test statistic, which measures the degree of deviation from a normal distribution, we consequently strongly *reject* the null hypothesis that the unconditional distribution is normal at the 1% significance level. As we can see from Table 4.3, those two conclusions (symmetry and excess kurtosis) also apply to all the other futures and spot return series in our analysis.

On a second glance, we see that the average weekly rate of return of the EMAc2 series is 0.00064 with a standard deviation of 0.0277 (or 3.34% and 19.97% annualized). As for the other European futures product in our list, BL2c1, the average is 0.00111 with a standard deviation of 0.0336 (5.95% and 24.23% annualized)**.** Those averages are significantly higher than the averages that Chen *et al.* (2008) reported (e.g., 0.000094 for soy meal) and which lead them to fail to reject the martingale hypotheses. Those discrepancies are likely due to structural changes in the moments between the two sampling periods. In fact, if we apply the same sample period that Chen *et al.* (2008) used (1980 to 1997) we also receive significantly lower mean values than the ones reported in Table 4.3 and which are based on the 2000-2014 sampling period. Furthermore, the annualized weekly averages and standard deviations of the DAX and AEX over the same sample period were 6.79% & 24.06% and 5.77% & 21.40%, respectively. By definition, those indices have a beta of (close to) 1. If there would still be no risk premium in the EMAc2 and BL2c1 contracts, we would expect to observe averages that are significantly lower than the market equity indices, which we do not. The same results apply upon comparing the SMc1 and Wc1 contracts to e.g., the S&P500 index.

To diagnose the potential for autocorrelation in the series we eyeball the return series to get a first impression. Indeed, Figure 4.3 seems to suggest that volatile periods come in clusters, where returns

The Q-statistics refer to the Ljung-Box Q test for the first lag.

of high magnitude are likely to entail consecutive observations of rather extreme returns of either sign. We can also see that those clusters take turns with more tranquil periods of return.

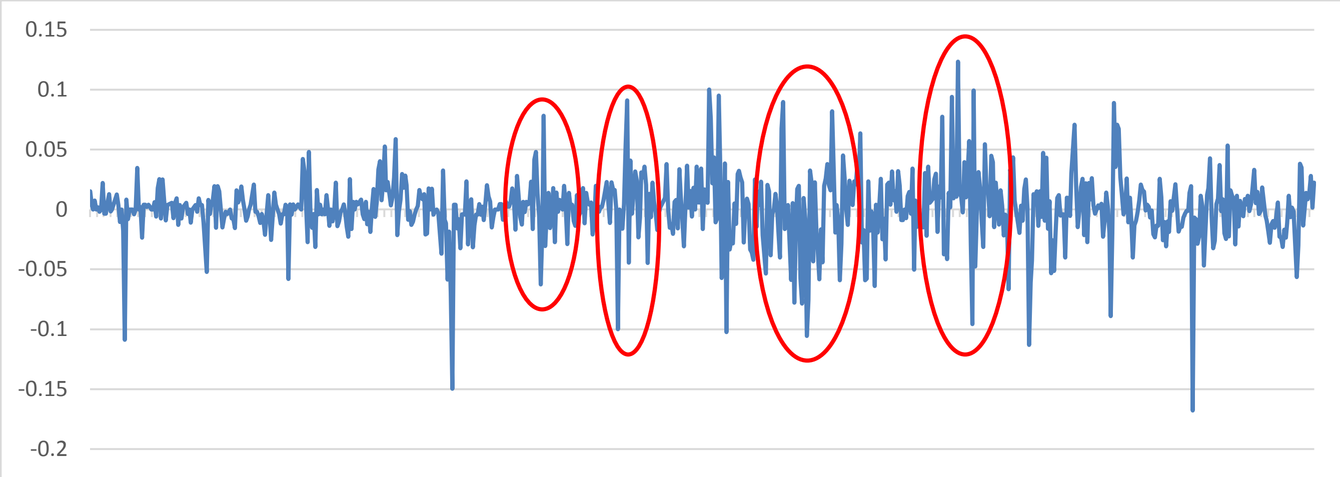

Figure 4.3 – *Volatility clustering of the EMAc2 return series.*

To test the potential for autocorrelation with more scrutiny, we conduct the Ljung – Box Q test on the squared return series (after all, we want to test for heteroscedasticity in the second moments). The correlogram in Table 4.4 reports significant degrees of autocorrelation for the first five lags. In fact, the series exhibits significant autocorrelation all the way up to lag 22 at the 5 % level. The diagnostic test therefore provides strong evidence that extreme returns have an echoing effect on successive observations and that recent information is viable in forecasting the conditional variance of the series.

While most of the futures and spot series in our analysis show signs of autocorrelation starting from the first lag at the 1% significance level (see Table 4.4), some series do not. However, even those series that fail to comply with this strict significance level, still feature reasonable degrees of autocorrelation. The Oil World soy spot series for example successively increases its degree of autocorrelation as we increase the lag size (see Appendix G). While it is almost autocorrelated at lag 1 at the 10 % significance level, it is significantly autocorrelated at the 5% level for lag 2 and 3. For lags 4 through to 36 the p-value of no autocorrelation even approaches zero.

Correlogram of RET2_EMAC2

Date: 01/24/15 Time: 16:35
Sample: 1 1000
Included observations: 753

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.121	0.121	11.006	0.001
		2	0.045	0.030	12.509	0.002
		3	0.057	0.049	14.992	0.002
		4	0.053	0.040	17.143	0.002
		5	0.056	0.042	19.486	0.002

Table 4.4 – *EMAc2 Correlogram (H_0: No autocorrelation at lag k).*

Moreover, once we translate the series to EUR, which is the series that will also be used in the analysis, then it is autocorrelated at each lag at the 1% significance level, except for lag 2, where the

degree of autocorrelation is significant at the 5% level. The situation for the Wc1 wheat futures series is very similar.
The only series that seem to entail no autocorrelation at all are the Reuters corn and the replacement soy spot series.

Another aspect that we can extract from Figure 4.3 upon further inspection is that the volatilities seem to *mean-revert*. This concept is related to stationarity. We statistically test for stationarity by means of the Augmented Dickey-Fuller (ADF) test. As we can see from Table 4.3, all return series are indeed stationary, as the null-hypothesis of a unit root has been rejected in all cases with p-values approaching zero. This implies that we do not need to bother with cointegration and that the OLS regression analysis can be conducted without modifying it with error correction terms.

Having reviewed the unconditional sample statistics of our time series, we conclude this section by answering Sub Question 3:

> *Sub Question 3: "Given the historic data of the cash and futures commodities, to what extent do those time series comply with the assumptions posed by the models identified in Sub Question 2 and what would be the implication of a potential violation of those assumptions?*

The strongest assumption that has been made in Chapter 2 is the assumption of homoscedasticity of the returns proposed by the OLS regression. Clearly, the autocorrelation and heteroscedasticity findings in this section contest this claim and rather provide evidence in favor of the hypothesis that the conditional variance of the series changes throughout time. At first, the use of the more advanced time-varying models therefore seems justified.

Moreover, it has been crystallized in the literature review, that under the assumption of a martingale process of the futures prices (i.e., $E[R_{f,t}] = 0$), the variance minimizing function would yield the same outcome as the mean-variance -, and Sharpe ratio objective function. The EWMA method also relies on the zero mean assumption for both the futures and spot return series.
In this section we have shown that the sample average rates of return deviate from the low values observed by (Chen *et al*., 2008) and Choudhry (2009). While focusing on the minimum variance hedge ratios, it therefore seems reasonable to also take a glimpse at the mean-variance hedge ratios and outcomes.

Last but not least, we have shown that the *unconditional* return distributions of all spot and futures return series feature fat tails. When estimating the GARCH models it would therefore seem more reasonable to assume that the errors are conditionally distributed according to a t-distribution rather than a normal distribution.

4.3 Parameterization of Hedging Strategies

Recall that the number of in-sample observations to estimate the models is n = 130 and n = 500 for the replacement and external cash commodity series.
For the EWMA method, we start with the parameter value λ = .94. In order to save space, the EViews parameterization output is only provided for the first case (corn) – for the other output the interested reader may consult Appendices H and I for the wheat and soy case, respectively.

4.3.1 Corn Parameterization

In this section we consider the parameter estimation for the hedging strategies for the spot commodities *0003 corn (replacement)* and *French Corn FOB Rhine (Reuters)*, which are sought to be mimicked by the futures contract EMAc2 of the Matif commodity exchange. To avoid repetition, we denote the former spot as v_1 and the latter spot as v_2. See Table 4.14 for a summary of the parameters under the different samples and methods.

OLS Regression

To get the OLS regression coefficients, we run Regression (4) twice. Once with v_1 as the dependent, and once with v_2 as the dependent variable and with EMAc2 as the independent variable in both cases. This yields the following outputs. See Tables 4.5 and 4.6, respectively.

Not surprisingly, the effect of the futures returns on both spot series is statistically significant. The magnitude of those effects is captured by β_1 = 0.662873 and β_2 = 0.747109. Both are a bit higher than their corresponding coefficients of correlation as the spot volatilities are somewhat higher than the respective futures volatilities. In this case, Equation (4) therefore becomes:

$$R_{1,t} = -0.001285 + 0.662873\, R_{ft} + \varepsilon_{1,t} \quad \text{and} \tag{14a}$$
$$R_{2,t} = -0.000375 + 0.747109\, R_{ft} + \varepsilon_{2,t} \tag{14b}$$

Dependent Variable: RET_REP_CORN
Method: Least Squares
Date: 03/21/15 Time: 15:08
Sample: 1 130
Included observations: 130
RET_REP_CORN = C(1) + C(2)*RET_EMAC2_REP

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	-0.001285	0.002268	-0.566511	0.5720
C(2)	0.662873	0.105566	6.279216	0.0000

Table 4.5 – *EViews OLS regression output Corn replacement* (H_0: $\beta = 0$).

Dependent Variable: RET_REUT_CORN
Method: Least Squares
Date: 03/21/15 Time: 15:14
Sample: 1 500
Included observations: 500
RET_REUT_CORN = C(1) + C(2)*RET_EMAC2_REUT

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	-0.000375	0.001128	-0.332218	0.7399
C(2)	0.747109	0.042235	17.68943	0.0000

Table 4.6 - *EViews OLS regression output Corn Reuters* (H_0: $\beta = 0$).

Dynamic Conditional Correlation

As has been described in Equation (11a) in Section 2.2.7, the elements of the diagonal matrix, D_t, can be modeled as *univariate* GARCH processes. In the light of the motivation in Section 3.2 we estimate those models under the assumption of conditional normality of the errors and hit the option for quasi-maximum likelihood estimates.

Resorting to the frequently applied vanilla GARCH(1,1) model we then get the following estimates for the variables v_1 and v_2, see Tables 4.7 and 4.8. For the EMAc2 futures contract, we also get two GARCH(1,1) models, depending on which sample period (that of v_1 or v_2) we base the estimation on. The outputs of those two models are presented in Tables 4.9 and 4.10, respectively. For estimation of the small sample EMAc2 series, we applied the I-GARCH restriction as in all other setups, EViews could not achieve convergence, which resulted in low, and insignificant ARCH and GARCH terms.[9]

Dependent Variable: RET_REP_CORN
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:17
Sample: 1 130
Included observations: 130
Convergence achieved after 500 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000833	0.001727	0.482498	0.6295
		Variance Equation		
C	8.23E-05	4.04E-05	2.036773	0.0417
RESID(-1)^2	0.265879	0.112308	2.367405	0.0179
GARCH(-1)	0.638667	0.116382	5.487679	0.0000

Table 4.7 – *Corn Replacement GARCH(1,1) estimation output* (H_0: $coeff. = 0$).

Upon comparing Table 4.7 through to 4.10, we learn that, apart from the forced EMAc2 small sample I-GARCH model, the ARCH and GACH terms in all the models are statistically significant at the 5% level according to the z-test, which, for each coefficient, tests the null-hypothesis that the coefficient is zero. The ARCH and GARCH terms are therefore momentous in modeling the conditional volatility. Again, this suggests that in general, heteroscedasticity is present in the considered time series, which confirms our first impression of Section 4.2.

Dependent Variable: RET_REUT_CORN
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:20
Sample: 1 500
Included observations: 500
Convergence achieved after 17 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	-0.000224	0.001112	-0.201595	0.8402
		Variance Equation		
C	5.20E-05	2.33E-05	2.233173	0.0255
RESID(-1)^2	0.270921	0.066287	4.087106	0.0000
GARCH(-1)	0.738139	0.044956	16.41932	0.0000

Table 4.8 – *Corn Reuters GARCH(1,1) estimation output* (H_0: $coeff. = 0$).

To formally test whether those GARCH models offer a good fit with the underlying data, we take a quick glance at the autocorrelation of the standardized residuals. If the models were specified correctly, then autocorrelation should be absent at each lag. The results for each model up to lag 10 are summarized in Table 4.11.

[9] Under the I-GARCH (integrated GARCH) model, the ARCH and GARCH parameters, α and β, sum to 1. In EViews, the constant term, ω, moreover disappears. Note that, as mentioned in Section 2.2.4, this would imply that the underlying returns process is *not* stationary. Clearly, this is at odds with the observation made in Section 4.1 that all return series are found to be stationary. Therefore, the I-GARCH restriction is only applied as a last resort.

Dependent Variable: RET_EMAC2_REP
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:22
Sample: 1 130
Included observations: 130
Convergence achieved after 24 iterations
Bollerslev-Wooldridge robust standard errors & covariance
Presample variance: backcast (parameter = 0.7)
GARCH = C(2)*RESID(-1)^2 + (1 - C(2))*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.001155	0.001475	0.783177	0.4335
Variance Equation				
RESID(-1)^2	0.045642	0.028985	1.574682	0.1153
GARCH(-1)	0.954358	0.028985	32.92604	0.0000

Table 4.9 – *EMAc2 GARCH(1,1) estimation output under replacement sample* (H_0: $coeff. = 0$).

Dependent Variable: RET_EMAC2_REUT
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:24
Sample: 1 500
Included observations: 500
Convergence achieved after 15 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.001510	0.000745	2.027386	0.0426
Variance Equation				
C	2.35E-05	1.26E-05	1.867703	0.0618
RESID(-1)^2	0.272499	0.071518	3.810217	0.0001
GARCH(-1)	0.740540	0.053749	13.77764	0.0000

Table 4.10 – *EMAc2 GARCH(1,1) estimation output under Reuters sample* (H_0: $coeff. = 0$).

Model	k =1	k =2	k =3	k =4	k =5	k =6	k =7	k =8
GARCH Replacement	1.66 (.20)	2.38 (.30)	2.44 (.49)	2.62 (.62)	3.00 (.70)	3.05 (.80)	4.48 (.72)	4.68 (.79)
GARCH Reuters	0.83 (.36)	1.43 (.49)	3.35 (.34)	5.58 (.23)	6.58 (.25)	6.69 (.35)	7.50 (.38)	7.50 (.48)
GARCH EMAc2 (Rep. sample)	8.37 (.00)	10.30 (.01)	10.30 (.02)	10.87 (.03)	10.97 (.05)	11.95 (.06)	11.96 (.10)	13.84 (.09)
GARCH EMAc2 (Reut. sample)	2.60 (.11)	2.98 (.22)	4.30 (.23)	6.33 (.18)	7.11 (.21)	7.12 (.31)	16.64 (.02)	17.91 (.02)

Table 4.11 – *Autocorrelation of standardized residuals. Ljung Box Q-Stat per lag k (*H_0*: No autocorrelation at lag k).*

As we can see, the GARCH models offer a good fit for both spot series. As for the EMAc2 futures series, however, this is not true. Especially in the replacement sample context, where we had to enforce the I-GARCH restriction, there is still statistically significant autocorrelation in the first eight lags at the 10% significance level. In the Reuters sample context, the situation looks a lot better. Yet, the seventh and eighth lags still feature statistically significant autocorrelation at the 5% level.

Inserting the estimated values into Equations (6b) and (6d) we get the following mean and variance equations for v_1 and the EMAc2 series, v_f:

$$R_{1,t} = 0.000833 + \varepsilon_{1,t} \tag{15a}$$
$$\sigma^2_{1,t} = 0.0000823 + 0.265879\,\varepsilon^2_{1,t-1} + 0.638667\,\sigma^2_{1,t-1} \tag{15b}$$
$$R_{f,t} = 0.001155 + \varepsilon_{f,t} \tag{16a}$$
$$\sigma^2_{f,t} = 0.045642\,\varepsilon^2_{f,t-1} + 0.954358\,\sigma^2_{f,t-1} \tag{16b}$$

The DCC system is then specified according to (11b) as

$$H_t = \begin{bmatrix} \sigma_{1,t}^2 & \sigma_{1,f,t} \\ \sigma_{f,1,t} & \sigma_{f,t}^2 \end{bmatrix} = \begin{bmatrix} \sqrt{\sigma_{1,t-1}^2} & 0 \\ 0 & \sqrt{\sigma_{f,t-1}^2} \end{bmatrix} \begin{bmatrix} 1 & \rho_{1,f,t} \\ \rho_{f,1,t} & 1 \end{bmatrix} \begin{bmatrix} \sqrt{\sigma_{1,t-1}^2} & 0 \\ 0 & \sqrt{\sigma_{f,t-1}^2} \end{bmatrix} \tag{17}$$

where $\rho_{1,f,t}$ is obtained by applying the EWMA method according to (11c) to the standardized residuals (i.e., the residual values we get after subtracting from each return observation the constant in mean equations (15a) and (16a), divided by their respective standard deviations) of R_1 and R_f.

Applying the same approach to v_2 and with EMAc2 now being based on the Corn Reuters sample period we have:

$$R_{2,t} = -0.000224 + \varepsilon_{2,t} \tag{18a}$$
$$\sigma_{2,t}^2 = 0.0000520 + 0.270921\, \varepsilon_{2,t-1}^2 + 0.738139\, \sigma_{2,t-1}^2 \tag{18b}$$

$$R_{f,t} = 0.001510 + \varepsilon_{f,t} \tag{19a}$$
$$\sigma_{f,t}^2 = 0.0000235 + 0.272499\, \varepsilon_{f,t-1}^2 + 0.740540\, \sigma_{f,t-1}^2 \tag{19b}$$

Similar to (17), expression (11b) then becomes:

$$H_t = \begin{bmatrix} \sigma_{2,t}^2 & \sigma_{2,f,t} \\ \sigma_{f,2,t} & \sigma_{f,t}^2 \end{bmatrix} = \begin{bmatrix} \sqrt{\sigma_{2,t-1}^2} & 0 \\ 0 & \sqrt{\sigma_{f,t-1}^2} \end{bmatrix} \begin{bmatrix} 1 & \rho_{2,f,t} \\ \rho_{f,2,t} & 1 \end{bmatrix} \begin{bmatrix} \sqrt{\sigma_{2,t-1}^2} & 0 \\ 0 & \sqrt{\sigma_{f,t-1}^2} \end{bmatrix} \tag{20}$$

BEKK Parameterization

In estimating the BEKK models, we assume the errors to be conditionally normally distributed. As for the constant coefficient matrix C'C in (10a), EViews provides the possibility to ensure various restrictions such as confining C'C to be scalar (i,e. *b* times a matrix of ones), diagonal, rank 1, or full rank.

Whenever model fitting was done for the larger sample sizes, choosing C'C to be indefinite resulted in the same output as restricting C'C to be full rank. That is, both rows (or columns) of C'C turned out to be linearly independent without the need to apply any further restrictions to the elements of C in deriving the BEKK model. As the output is usually similar under both absence or presence of this restriction, all large sample models have been estimated applying the full rank restriction, as this ensures that the covariance matrix is positive semi definite.

In the replacement sample context, however, the full rank restriction sometimes led to negative ARCH/GARCH parameter estimates, which is why the rank 1 restriction has been applied for all replacement samples. Under this restriction, the elements of the second column in C are zero. The transformed constant coefficient matrix C'C is still triangular, however.

As can be seen from the parameter estimates in Tables 4.12 and 4.13, $(a_{ii}^2 + b_{ii}^2) < 1$ holds for both v_1 and v_2. The displayed parameters correspond to simple mean equations such as in (6b). For v_1 those constants were 0.001628 and -0.002495 for the spot and futures series, respectively. In case of v_2 we have 0.0000140 and 0.001496. In the fashion of (10b) we then have:

$$C'C_{v_1} = \begin{bmatrix} 0.0000505 & 0.000115 \\ 0 & 0.000260 \end{bmatrix},\ A_{v_1} = \begin{bmatrix} 0.458564 & 0 \\ 0 & 0.161442 \end{bmatrix}, B_{v_1} = \begin{bmatrix} 0.850681 & 0 \\ 0 & 0.634052 \end{bmatrix}$$

$$C'C_{v_2} = \begin{bmatrix} 0.0000879 & 0.0000302 \\ 0 & 0.0000361 \end{bmatrix},\ A_{v_2} = \begin{bmatrix} 0.426650 & 0 \\ 0 & 0.545985 \end{bmatrix}, B_{v_2} = \begin{bmatrix} 0.868627 & 0 \\ 0 & 0.840175 \end{bmatrix}$$

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a rank one matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	5.05E-05	3.75E-05	1.346311	0.1782
M(1,2)	0.000115	7.96E-05	1.439611	0.1500
M(2,2)	0.000260	0.000186	1.399881	0.1615
A1(1,1)	0.458564	0.084444	5.430364	0.0000
A1(2,2)	0.161442	0.184393	0.875531	0.3813
B1(1,1)	0.850681	0.055339	15.37230	0.0000
B1(2,2)	0.634052	0.326633	1.941178	0.0522

Table 4.12 – *Corn Replacement and EMAc2 Diagonal BEKK output* ($H_0: coeff. = 0$).

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a full rank matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	8.79E-05	2.92E-05	3.008144	0.0026
M(1,2)	3.02E-05	9.68E-06	3.119104	0.0018
M(2,2)	3.61E-05	1.19E-05	3.025659	0.0025
A1(1,1)	0.426650	0.065550	6.508737	0.0000
A1(2,2)	0.545985	0.072475	7.533442	0.0000
B1(1,1)	0.868627	0.031801	27.31436	0.0000
B1(2,2)	0.840175	0.034583	24.29425	0.0000

Table 4.13 – *Corn Reuters and EMAc2 Diagonal BEKK output* ($H_0: coeff. = 0$).

	OLS		DCC		BEKK	
Parameter	Rep	Reut	Rep	Reut	Rep	Reut
β	0.662873*	0.747109*	-	-	-	-
μ_1	-	-	0.000833	-0.000224	0.001628	0.0000140
μ_2	-	-	0.001155	0.001510 **	0.002495	0.001496
C_{11}	-	-	0.0000823**	0.0000520**	0.0000505	0.0000879**
C_{12}	-	-	-	-	0.000115	0.0000302**
C_{22}	-	-	0	0.0000235**	0.000260*	0.0000361**
A_{11}	-	-	0.265879**	0.270921*	0.458564*	0.426650*
A_{22}	-	-	0.045642	0.272499*	0.161442**	0.545985*
B_{11}	-	-	0.638667*	0.738139*	0.850681*	0.868627*
B_{22}	-	-	0.954358*	0.740540*	0.634052*	0.840175*

Table 4.14 – Corn *Parameter summary* ($H_0: coeff. = 0$).[10]

[10] * = significance at the 1% level, ** = significance at the 5% level, *** = significance at the 10% level,

4.3.2 Wheat Parameterization

For a thorough discussion on parameterization refer to the corn parameterization case. In this section we will only briefly provide and discuss the wheat parameter outcomes. All the EViews output can be found in Appendix H. In this section v_1 represents *Wheat (0041) replacement* and v_2 stands for *Wheat, U.S., No.2, SRW, fob Gulf* from the Oil World data source. The former is hedged with the *BL2c1* Matif futures contract and the latter is hedged via the *Wc1* CBOT contract. A summary of the coefficients is provided in Table 4.16 at the end of this section.

OLS Parameterization

Running the regression with v_1 as dependent and *BL2c1* variable yields β_1 = 0.639467. Choosing to regress *Wc1* on v_2, we get β_1 = 0.842738.

DCC Parameterization

Upon fitting the various univariate GARCH(1,1) models, we have made a number of alterations in the software settings. First, in the v_1 context, the parameters of the estimated GARCH models are not significant for both the v_1 as well as for the BL2c1 series. To enforce meaningful parameters, the I-GARCH restriction has been applied. As can be seen in Appendix H, the ARCH and GARCH parameters for both GARCH models sum to 1.[11]

Second, as already mentioned in the methodology chapter, there are gaps in the sample of the v_2 series. Those are most pronounced in the range from observation 400 to 500. This results in meaningless GARCH parameters. After all, the frequency of the dates corresponding to the observations must be consistent throughout the sample in order to study autocorrelation.

Contrary to the other external data series, we therefore based the estimation of the GARCH model on the first 400 observations. The autocorrelation in the standardized residuals for the GARCH models of the different series is presented in Table 4.15. Overall, it seems that the model fitting went better than in the corn parameterization case. It is only the first two lags of the standardized residuals of v_2 and BL2c1 (v_1 context) that show significant autocorrelation at the 10% level.

Model	k =1	k =2	k =3	k =4	k =5	k =6	k =7	k =8
GARCH Replacement	0.85 (.35)	2.72 (.26)	4.58 (.21)	4.75 (.31)	4.46 (.45)	7.69 (.26)	7.70 (.36)	8.21 (.41)
GARCH Oil World	4.89 (.03)	4.91 (.09)	5.12 (.16)	5.31 (.26)	6.50 (.26)	7.57 (.27)	8.27 (.31)	8.43 (.39)
GARCH BL2c1 (Rep sample)	2.87 (.09)	5.28 (.07)	6.29 (.10)	7.69 (.10)	7.82 (0.17)	7.87 (.25)	8.05 (.33)	8.57 (.38)
GARCH Wc1 (OW sample)	0.64 (.42)	0.88 (.64)	1.50 (.68)	5.16 (.27)	8.81 (.12)	9.10 (.17)	9.11 (.24)	9.13 (.33)

Table 4.15 – *Autocorrelation in standardized residuals (H_0: No autocorrelation at lag k).*

BEKK Parameterization

Restricting C'C to be full rank in the Reuters sample context and rank 1 in the replacement sample (full rank restriction caused negative entries for the GARCH parameter matrix) context has yielded

[11] See footnote 9.

parameter estimates that resulted in a positive semi definite covariance matrices. Both BEKK models satisfy the stationarity condition.

	OLS		DCC		BEKK	
Parameter	Rep	OW	Rep	OW	Rep	OW
β	0.639467*	0.842738*	-	-	-	-
μ_1	-	-	-0.000139	0.002757	0.000484	0.003372
μ_2	-	-	0.000258	0.003022	0.001054	0.003271
C_{11}	-	-	0	0.0000324	0.0000721	0.000223*
C_{12}	-	-	-	-	0.000115***	0.000194*
C_{22}	-	-	0	0.0000385	0.000183***	0.000168*
A_{11}	-	-	0.045586**	0.069250**	0.023135	0.315406*
A_{22}	-	-	0.085688*	0.022427	0.380860*	0.103625**
B_{11}	-	-	0.954414*	0.925548*	0.940794*	0.905451*
B_{22}	-	-	0.914312*	0.964046*	0.702232*	0.958960*

Table 4.16 – Wheat *Parameter summary* $(H_0: coeff. = 0)$.

4.3.3 Soy Parameterization

Again, we only briefly discuss the results of the estimation outcomes that are to be found in Appendix I. In this parameterization section, v_1 = *Soya Hipro ADM GMO (0300)* and v_2= *Soya pellets 48% Brazil, CIF Rotterdam*. Both are hedged with the *SMc1* contract from the CBOT. A summary of the parameter estimates is appended at the end of this section. See Table 4.18.

OLS Regression
Running the OLS regression with v_1 as dependent and SMC1 as independent variable yields $\beta_1 = 0.360171$. Swapping v_1 for v_2 we get $\beta_2 = 0.535843$. One immediately notices how β_1 and β_2 are both substantially lower than the coefficient of correlation between the variables in the corresponding OLS regression. It follows from (4) and (2b) that this is due to the fact that the standard deviation of the returns of the futures series, SMc1 is considerably larger than the standard deviation of the respective spot series. In the regression involving v_1 one might suspect that this is because the SMc1 series also includes the variability from the FX conversion, whereas the spot series is already denoted in the domestic currency. However, having double checked the regression with the USD-denoted series, this hypothesis can be rejected, as the regression output was similar.

DCC Parameterization
Just as in the OLS case, the univariate GARCH estimations of v_2 and SMc1 have also been tested with their base currency series. No deviations have been observed in the parameter estimation output.
As for the GARCH estimation of v_1, we again had to apply the I-GARCH restriction such that $\alpha + \beta = 1$.[12] For the small sample GARCH estimation of the SMc1 series, neither the ARCH nor the GARCH term is statistically significant. This is also reflected in Table 4.17, which shows the autocorrelation in the standardized residuals. There, we can see that there is still statistically significant autocorrelation at the first lag at the 10 % level. Overall, though, the fitted GARCH models seem to do well in filtering the autocorrelation.

[12] See footnote 9.

Model	k =1	k =2	k =3	k =4	k =5	k =6	k =7	k =8
GARCH Replacement	1.52 (.22)	1.52 (.47)	3.90 (.27)	3.92 (.42)	4.11 (.53)	5.51 (.48)	6.00 (.54)	6.03 (.64)
GARCH Oil World	0.15 (.70)	0.16 (.92)	3.40 (.33)	3.77 (.44)	4.35 (.50)	4.78 (.57)	5.31 (.62)	5.33 (.72)
GARCH SMc1 (Rep sample)	3.54 (.06)	3.58 (.17)	3.69 (.30)	3.74 (.44)	7.63 (.18)	7.93 (.24)	8.32 (.30)	8.81 (.36)
GARCH SMc1 (OW sample)	0.03 (.87)	0.03 (.99)	1.10 (.78)	3.11 (.54)	5.91 (.32)	10.45 (.11)	10.55 (.16)	11.10 (.20)

Table 4.17 – *Autocorrelation in standardized residuals (H_0: No autocorrelation at lag k).*

BEKK Parameterization

As explained in the corn BEKK parameterization, the system of $v_1$is based on the assumption that C'C is rank 1, while in the context of v_2 it is rank 2. The parameterizations of the BEKK models are stationarity consistent for both v_1 and v_2. Refer to Table 4.18 for the parameterization summary.

	OLS		DCC		BEKK	
Parameter	Rep	OW	Rep	OW	Rep	OW
β	0.360171*	0.535843*	-	-	-	-
μ_1	-	-	0.002206	0.001298	0.002919	0.001667
μ_2	-	-	0.004543	0.002440	0.003864	0.002786
C_{11}	-	-	0	0.0000497***	0.0000777***	0.0000332*
C_{12}	-	-	-	-	0.000234**	0.0000617*
C_{22}	-	-	0.000548	0.000323**	0.000703**	0.000154*
A_{11}	-	-	0.050926**	0.099218*	0.108238	0.253757*
A_{22}	-	-	0.246582	0.203071*	0.455762*	0.351549*
B_{11}	-	-	0.949074*	0.865535*	0.945226*	0.955032*
B_{22}	-	-	0.481810	0.640135*	0.660496*	0.894837*

Table 4.18 – Soy *Parameter summary (H_0: $coeff. = 0$).*

4.4 Application of Minimum Variance Hedging Strategies

Per commodity, we review the outcome of the hedging effectiveness of each method in terms of conditional and unconditional hedging effectiveness, both in – and out-of- sample. As the OLS method is perceived as the baseline method in literature, we discuss the relative performance of each model with respect to the straightforward OLS method and shed light into the question of how the OHR estimates under the different models vary through time. Taking a broader perspective, we also judge the suitability of the hedge setup in general.

As there are likely to be discrepancies between the hedging strategies of the different commodities, we consider each individually, starting with *corn*, followed by *wheat* and finishing with *soy*. Per commodity, we will of course inspect the hedging strategies for both, the replacement and the external vendor series (e.g., Reuters and Oil World).

4.4.1 Case Corn

We first reveal the output for the small (replacement) sample and then table the results for the Reuters sample.

Corn Replacement (0003)
As a reference point, Tables 4.19 and 4.20 present the results of the hedge outcome under each of the four methods both in– and out-of-sample. Note that $\overline{Var_t}(R_{s,t})$ and $\overline{Var_t}(R_{h,t})$ denote the average of the conditional variances of the spot (or unhedged) and the hedged portfolio, respectively. Note that in general, $(\overline{Var_t}(R_{s,t})$ - $\overline{Var_t}(R_{h,t}))/\ \overline{Var_t}(R_{s,t}) \neq \bar{E}_t$as denoted in (12c). That is, the difference of the average of the conditional variances is not the same as the average of the difference of the conditional variances.

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0.000859596	0.000859596	0.000859596	0.000859596
$Var(R_{h,t})$	0,000657166	0,000692255	0,000712092	0,000684584
$\hat{E}$	23.55%	19.47%	17.17%	20.36%
$\overline{Var_t}(R_{s,t})$	0,000893622	0,000893622	0,000893622	0,000893622
$\overline{Var_t}(R_{h,t})$	0,000685523	0,000721173	0,000740599	0,000714103
$\bar{E}_t$	25.73%	21.93%	20.17%	22.73%

Table 4.19 – *In-sample Hedging Effectiveness (Corn Replacement).*

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,000927608	0,000927608	0,000927608	0,000927608
$Var(R_{h,t})$	0,000635513	0,00067238	0,000625247	0,000605969
$\hat{E}$	31.49%	27.51%	32.26%	34.67%
$\overline{Var_t}(R_{s,t})$	0,00076743	0,00076743	0,00076743	0,00076743
$\overline{Var_t}(R_{h,t})$	0,000567221	0,000591321	0,000585012	0,000563196
$\bar{E}_t$	24.67%	20.15%	22.23%	25.03%

Table 4.20 – *Out-of-sample Hedging Effectiveness (Corn Replacement).*

There are two conclusions that can immediately be drawn from the first look.
First, the straightforward, static OLS method significantly outperforms all time-varying models in the in-sample context. In the out–of-sample context, however, it is slightly outperformed by the BEKK model, both conditionally (25.03% vs. 24.67%) and unconditionally (34.67% vs. 31.49%).
In terms of unconditional hedge effectiveness, the DCC also scores better than the OLS method, but the degree of performance improvement is less striking (32.26% vs. 31.49%).
Even though we have to be careful in interpreting the results from the out-of-sample context, given the rather sparse sample size, the results seem to be consistent with findings in literature that the performance of time-varying methods relative to the static OLS model improves in the out-of-sample context.

The source of the in-sample dominance of the OLS method is, in part, likely to be located in the poor model specification of the time-varying models. As we have seen in Section 4.1, it was rather challenging to fit proper univariate GARCH models to the EMAc2 futures series. After model fitting, there was still a significant degree of autocorrelation in the standardized residuals, especially in the replacement sample context, which might explain the DCC's performance dip of 17.17%. Apart from the residual autocorrelation of the EMAc2 series in the replacement sample setting, we have also seen that the GARCH parameter of the replacement spot series took on a low value of around 0.64. This of course makes the covariance estimates of the DCC model less persistent and more reactive.

The OHR series of the DCC model therefore intensively reacts to shocks in the market, which explains its spiky shape in Figure 4.4. A particularly harsh reaction of this model occurs in the period around October 2013. In general, the time-varying models follow a similar pattern.

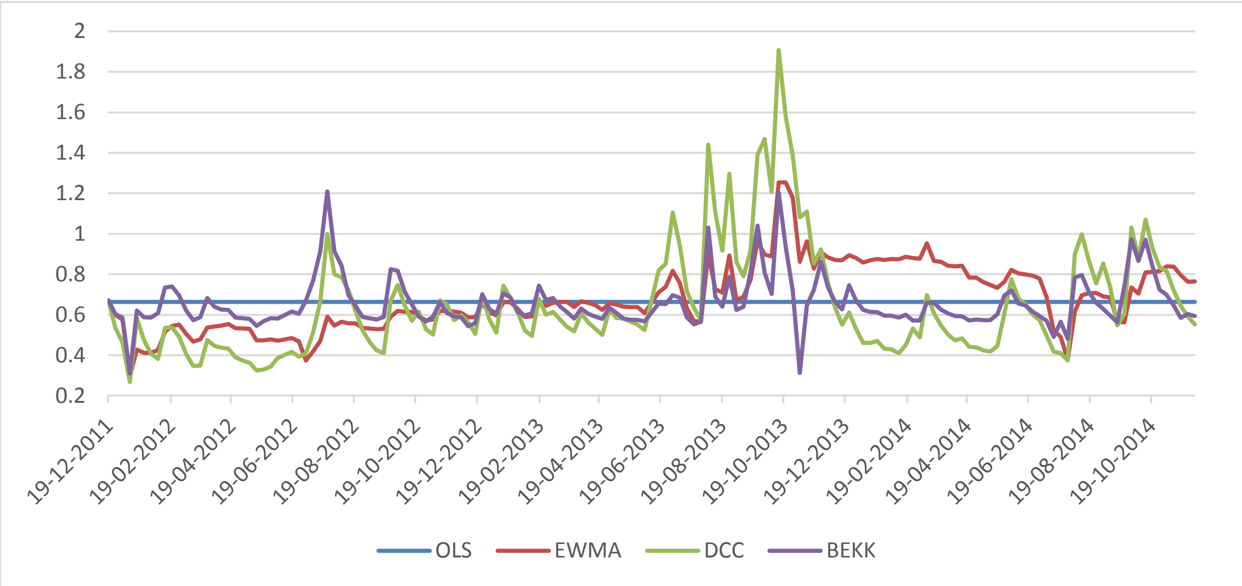

Figure 4.4 – *Corn Replacement OHR series.*

But what exactly is it that the time-varying hedge ratios so unanimously, albeit to a different extent, seem to respond to? Recall that the hedge ratio can be defined as the correlation times the standard deviation of the spot divided by the standard deviation of the futures. With that in mind, the answer to the observed leverage is given in the following Figure (4.5), which plots the volatility of both securities and their correlation.

While both, the correlation and the futures volatility parameters do not fundamentally change, the volatility of the corn replacement value series increases rapidly. It follows from the formula of the OHR that the time-varying hedge ratios will respond with a higher hedge ratio. This behavior makes sense. After all, potentially bigger, more volatile returns of one direction have to be offset by a multitude of potentially lower, less volatile returns of the opposite direction.

However, it is ironically exactly this virtue, which costs the time-varying hedge strategies their score in performance. If we take a look at the conditional hedge performances over time (see Figure 4.6), it is exactly this period where the performance gap among the different methods widens the most. The problem is that in this period, correlation hovers around only 30%, which is relatively low, compared to the rest of the sample period. The low degree of correlation causes the high degree of leverage in the futures contracts to occasionally backfire. In fact, the hedging performance of the DCC methodology for example, which features the highest degree of leverage, drops slightly below zero, which implies that the variance of the hedged portfolio during that time was indeed higher than what it would have been without a futures hedge altogether.

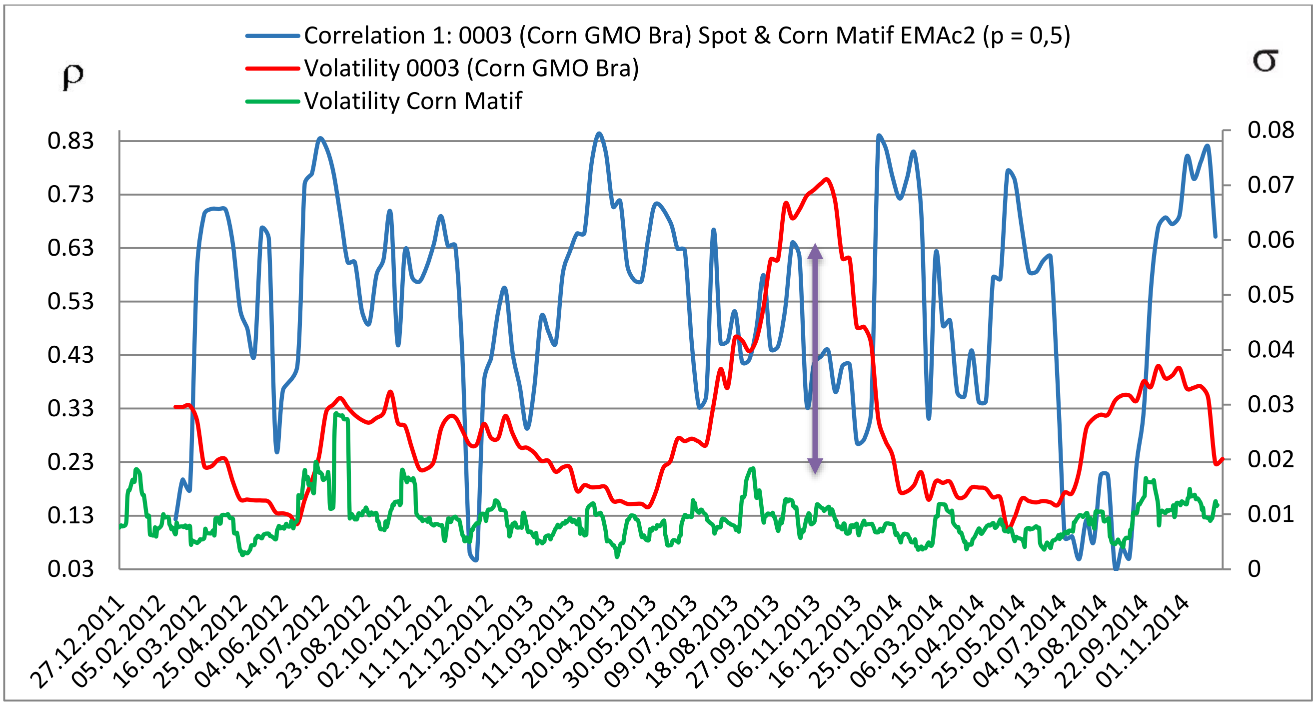

Figure 4.5– *Volatility and Correlation of 0003 and EMAc2 (10 weeks rolling average).*

The second major conclusion is that, regardless of the relative performance between the different models, an overall hedging effectiveness of around 25% is rather disappointing. To see how the hedging effectiveness unfolds in a bigger sample context, we now turn to the *French Corn FOB Rhine* cash commodity of the Reuters database.

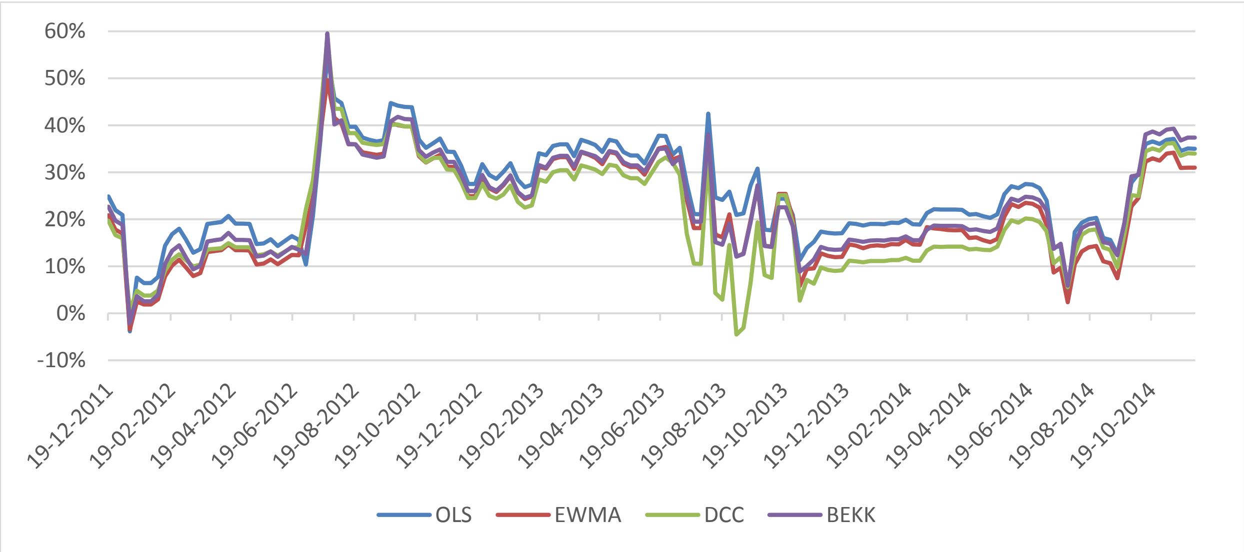

Figure 4.6 – *Codintional Hedging Effectiveness over time (corn replacement).*

French Corn FOB Rhine (Reuters)

Among comparison of the different hedging strategies in Table 4.21 and 4.22 it once again seems that the static hedge ratio outperforms the time-varying strategies – but now to a much stricter degree. Another interesting point to notice upon comparison is that the unconditional hedging effectiveness in the out-of-sample period is about twice as much as it is in the in-sample period. This is true for each strategy. For the conditional hedging effectiveness the divide is even higher.

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,00103007	0,00103007	0,00103007	0,00103007
$Var(R_{h,t})$	0,000632587	0,000683634	0,000713197	0,00069271
$\hat{E}$	38.59%	33.63%	30.76%	32.75%
$\overline{Var_t}(R_{s,t})$	0,000986014	0,000986014	0,000986014	0,000986014
$\overline{Var_t}(R_{h,t})$	0,000630797	0,000679892	0,00070325	0,000678682
$\bar{E}_t$	28.20%	24.35%	23.07%	25.67%

Table 4.21 – *In-sample Hedging Effectiveness (Corn Reuters).*

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,00088588	0,00088588	0,00088588	0,00088588
$Var(R_{h,t})$	0,000310624	0,000318586	0,000344264	0,000353537
$\hat{E}$	64.94%	64.04%	61.14%	60.09%
$\overline{Var_t}(R_{s,t})$	0,000986139	0,000986139	0,000986139	0,000986139
$\overline{Var_t}(R_{h,t})$	0,000305648	0,000319117	0,000367589	0,000379728
$\bar{E}_t$	66.45%	65.13%	59.18%	60.23%

Table 4.22 – *Out-of-sample Hedging Effectiveness (Corn Reuters).*

This impression is further visualized in Figure 4.7, which plots the evolution of the hedging effectiveness of the four strategies. With the out-of-sample period starting in May 2010, the hedging effectiveness of all the hedging strategies has moved around 50% ever since.

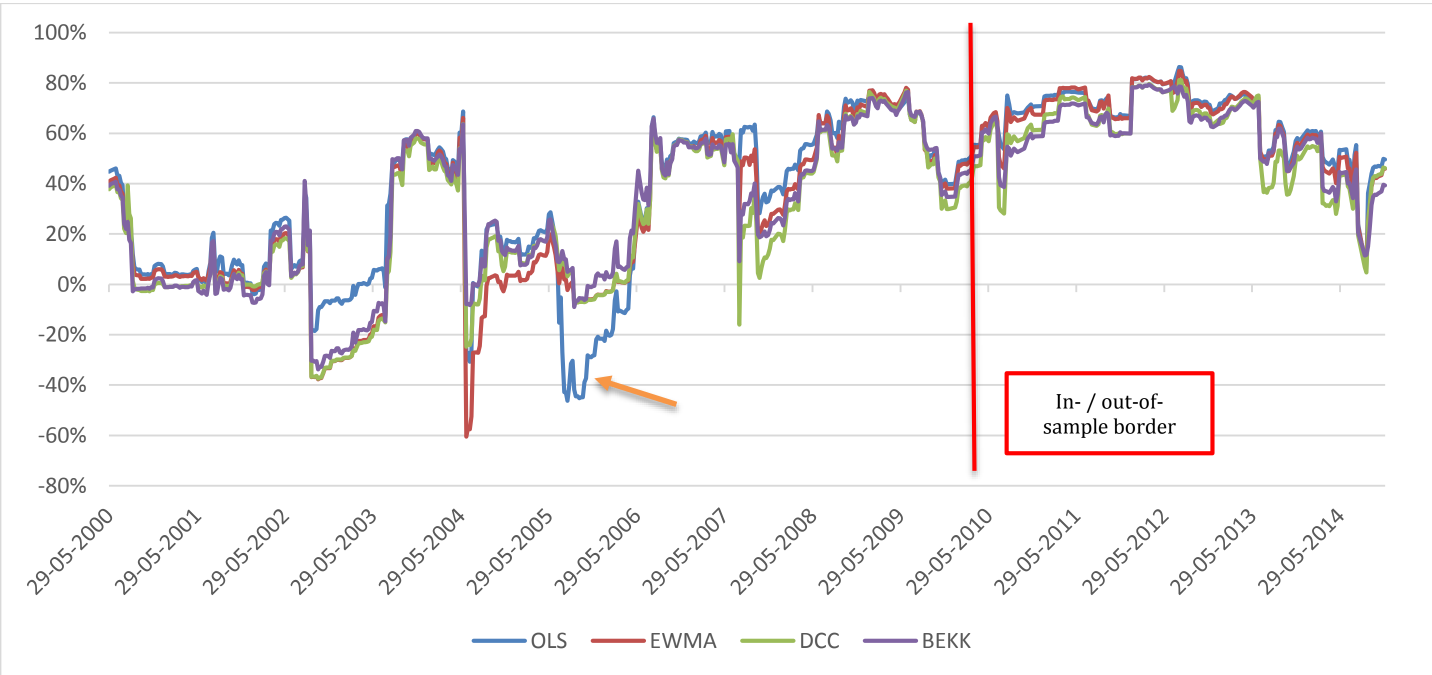

Figure 4.7 – *Conditional Hedging Effectiveness (Reuters Corn).*

Another, very interesting point that unfolds from this graph is that now there is a period in which the time-varying models significantly outperform the static OLS method. While the conditional hedging performance of the OLS method around September/October 2005 drops to as low as -45%, the time-varying methods oscillate around 0%. To see how this can be the case, we again take a look at the volatility - correlation plot for the period 2005. See Figure 4.8.

As we can see, the September/October period constitutes one of the rare instances throughout the lifetime of the short-term correlation series, in which the correlation actually drops to a negative level – up to –77% in this case, based on an 8 weeks rolling average. The static hedge ratio, of course, does not take any account of this and continues to maintain a hedge ratio of 0.747109, which in the scenario of a negative correlation leads to an outright long rather than a hedged position, thus considerably increasing the variance of the "hedged" portfolio.

The time-varying models, on the other hand, recognize how the conditional correlation (a weighted average of past information) slides to zero over time, and thus close their short futures positions (see Figure 4.9). The BEKK method, being the most reactive model in terms of parameterization, even calculates a negative covariance between the futures and the spot and therefore comes up with a negative hedge ratio, resulting in a long futures position, which, given the negative correlation, cancels the movements of the long spot position and hence results in a lower variance of the hedged portfolio.

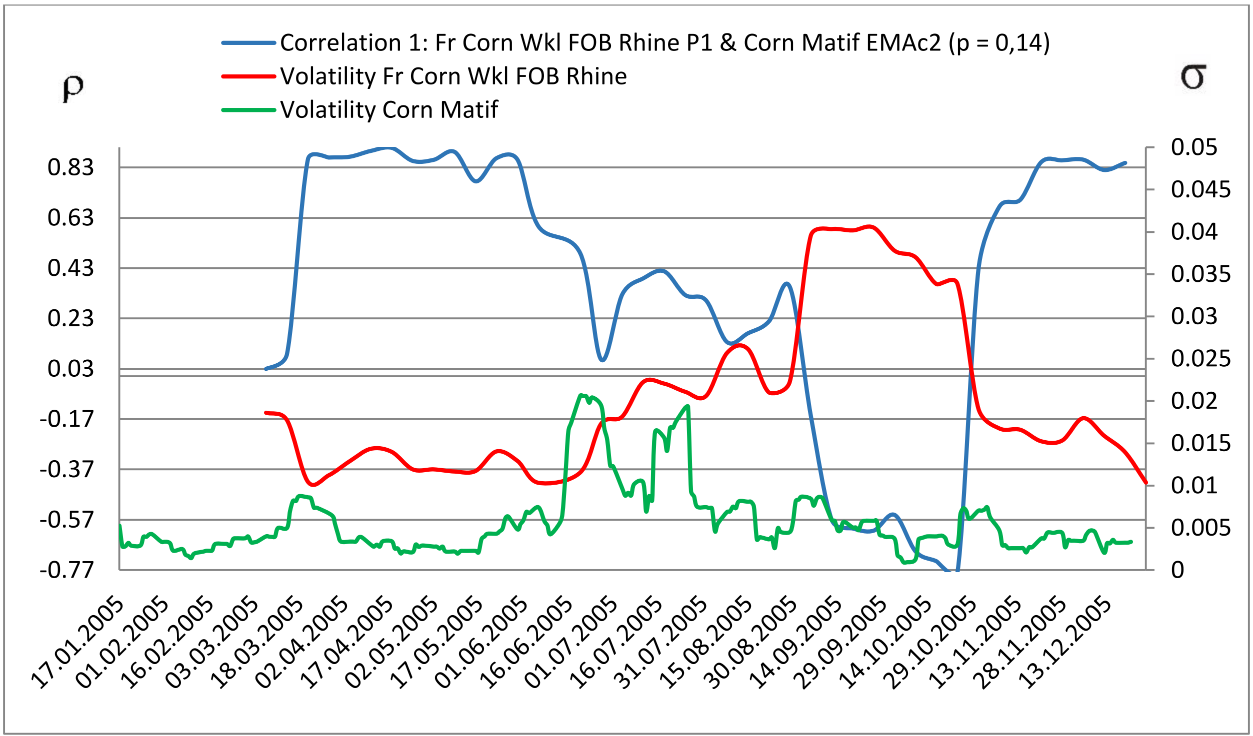

Figure 4.8 - *Volatility and Correlation of Corn Reuters and EMAc2 (8 weeks rolling average).*

Judging the hedging effectiveness of the Reuters corn sample in general, there is a 20% boost in hedging effectiveness compared to the replacement sample case (the unconditional, full sample hedge effectiveness is 24.86% and 44.7% for the replacement and Reuters sample, respectively).

While the volatility of the spot and futures also play a part, the key parameter to a successful hedging strategy is predominantly the degree of correlation. When we compare Figure 4.10 with Figure 4.7, we see that the periods of hedging ineffectiveness and correlation dips clearly coincide.

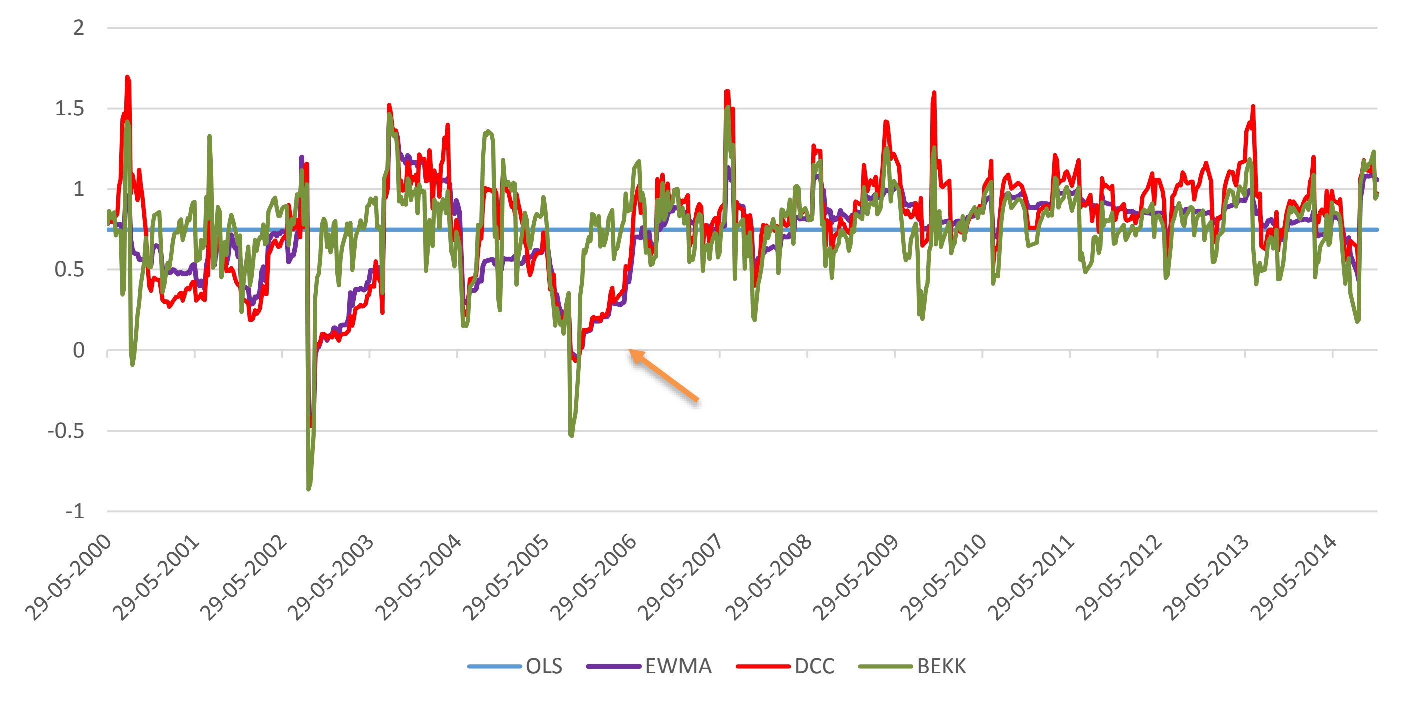

Figure 4.9 – *Corn Reuters OHR series.*

So, while 44.7% may not seem too much either, we have to bear in mind that this refers to the unconditional hedging effectiveness, which attributes equal weight to all observations in the sample. We can, however, see from Figure 4.10, how the degree of correlation has improved and stabilized since the beginning of 2007. The contemporary hedging effectiveness of the EMAc2 futures contract is therefore likely to be higher, which is also supported by the remarkably higher degree of out-of-sample hedging effectiveness reported in Table 4.22.

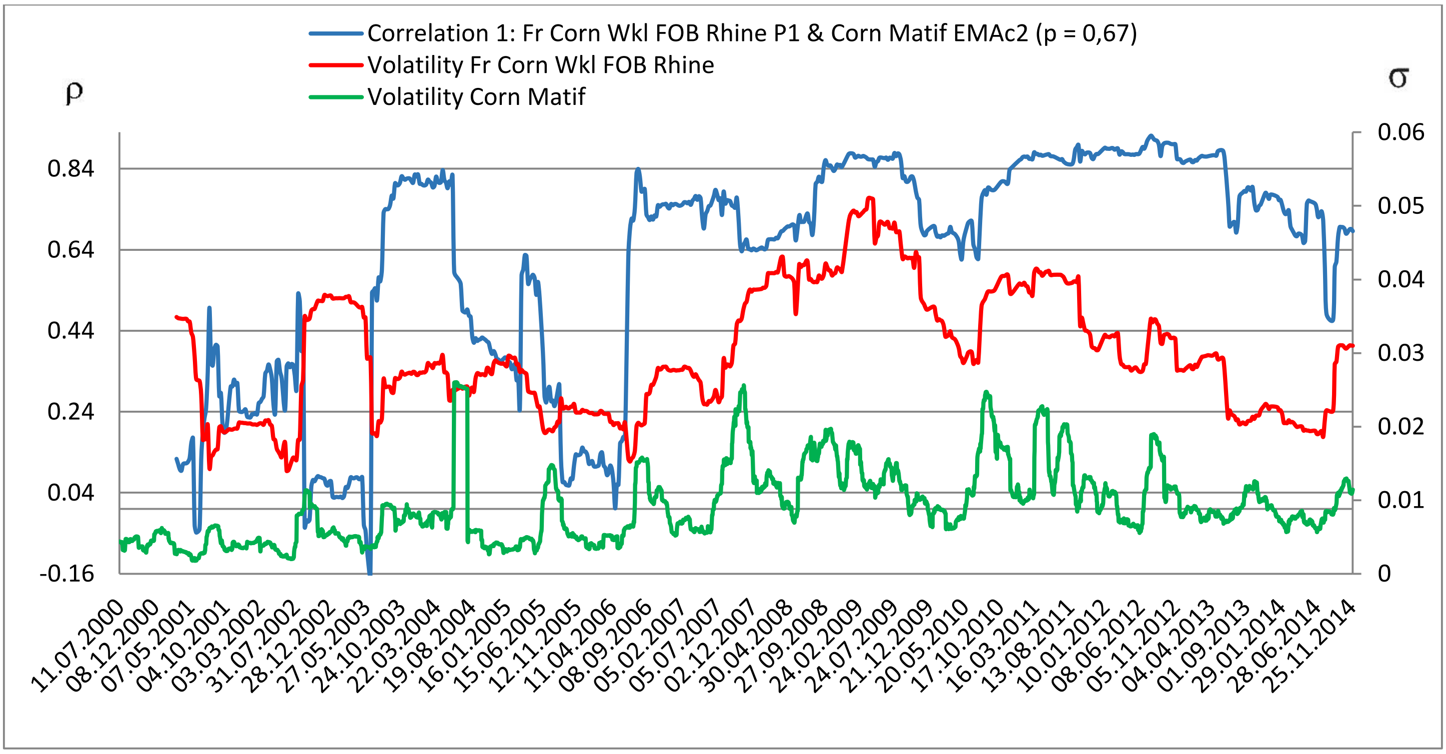

Figure 4.10 - *Volatility and Correlation of Corn Reuters and EMAc2 (40 weeks rolling average).*

4.4.2 Case Wheat

Again, we first present the findings for the replacement sample, followed by the results for the external (Oil World) sample.

Wheat Replacement (0041)

Owing to the higher coefficient of correlation, we learn from Tables 4.23 and 4.24 that the overall wheat replacement hedging effectiveness is about 6% to 10% higher than the corn replacement hedging effectiveness, whatever model or sample we compare. Once more, we also observe a superior performance of the OLS method for both the in – and out-of-sample period. However, just as in the corn replacement context, the performance gap is less pronounced in the out-of-sample period.

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,000711284	0,000711284	0,000711284	0,000711284
$Var(R_{h,t})$	0,00047742	0,000496758	0,000499031	0,00051172
$\hat{E}$	32.88%	30.16%	29.84%	28.05%
$\overline{Var_t}(R_{s,t})$	0,000771547	0,000771547	0,000771547	0,000771547
$\overline{Var_t}(R_{h,t})$	0,000520696	0,000541069	0,000542789	0,00055753
$\bar{E}_t$	33.55%	30.85%	30.43%	29.47%

Table 4.23 – *In-sample Hedging Effectiveness (Wheat Replacement).*

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,001658295	0,001658295	0,001658295	0,001658295
$Var(R_{h,t})$	0,001063388	0,001119915	0,001112036	0,001094289
$\hat{E}$	35.88%	32.46%	32.94%	34.01%
$\overline{Var_t}(R_{s,t})$	0,001034388	0,001034388	0,001034388	0,001034388
$\overline{Var_t}(R_{h,t})$	0,000710177	0,00073522	0,000733411	0,000718214
$\bar{E}_t$	31.82%	29.25%	29.40%	31.00%

Table 4.24 – *Out-of-sample Hedging Effectiveness (Wheat Replacement).*

While the face statistics document a strict outperformance, it is worthwhile to take a look at how the methods compare over time. Figure 4.11 plots the conditional hedge effectiveness over time. Upon studying this figure, we realize that the time-varying BEKK method actually beats the OLS method for quite some time during the period from end of June 2013 until the end of September 2014. However, prior to that period, the converse is true. Here, the BEKK method seriously falls short performance wise, also compared to its time–varying counterparts.

Again, this is related to leverage and can be explained by taking a look at the volatility and correlation structure throughout the lifetime of the series (see Figure 4.12). If we look closely at Figure 4.11, then we see that the period of severe underperformance stretches from the 23rd of July 2012 until the 5th of November 2012. As we can see in Figure 4.12, at the start of this period, both the correlation and the spot volatility are at periodic highs. The BEKK method being by far the most reactive to changes in the (co-) variance structure (see wheat parameterization section) it is also the method that responds the quickest by increasing the short position in the futures contracts, assuming a hedge ratio of about 80% (see Figure 4.13).

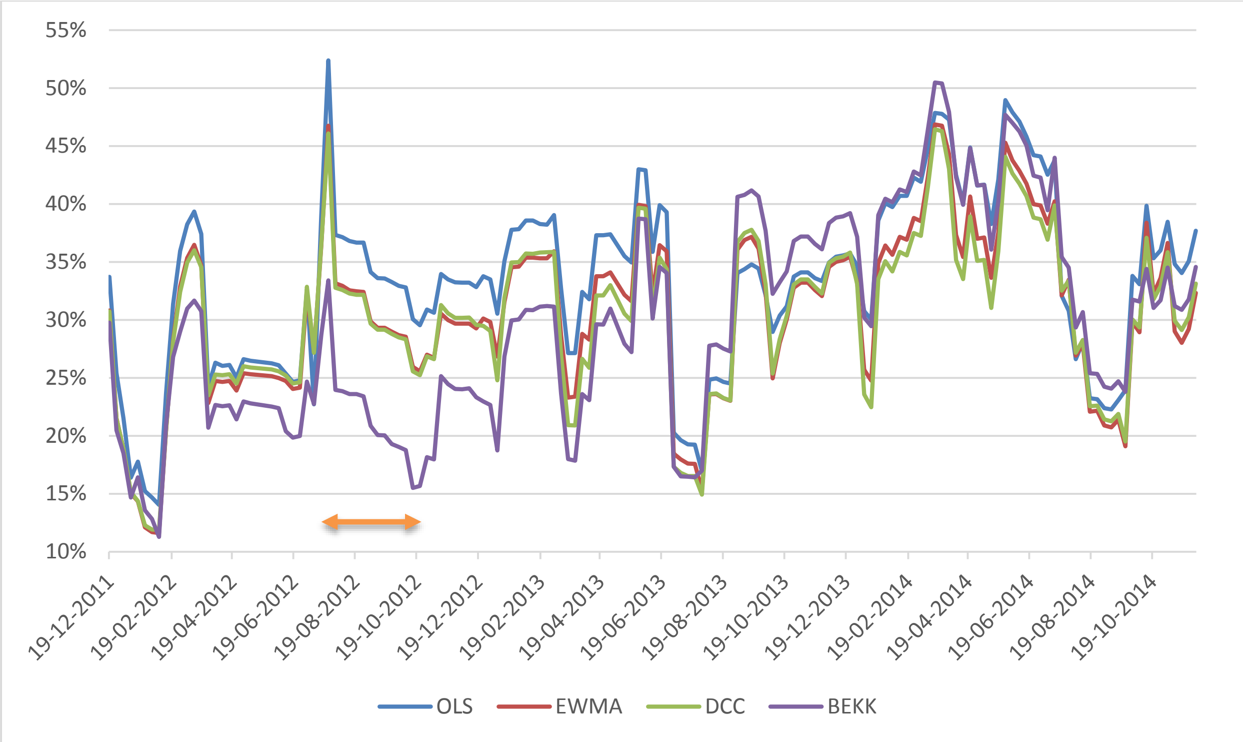

Figure 4.11 - *Conditional Hedging Effectiveness (Wheat Replacement).*

However, as quickly as the correlation and spot volatility rallied up, they collapse again at the end of September. The BEKK method still having the biggest hedge ratio, it suffers the most from the sudden decrease in the correlation and spot volatility, which explains the performance gap in Figure 4.12. It is only when the correlation stabilizes again that it can catch up with the other methods again.

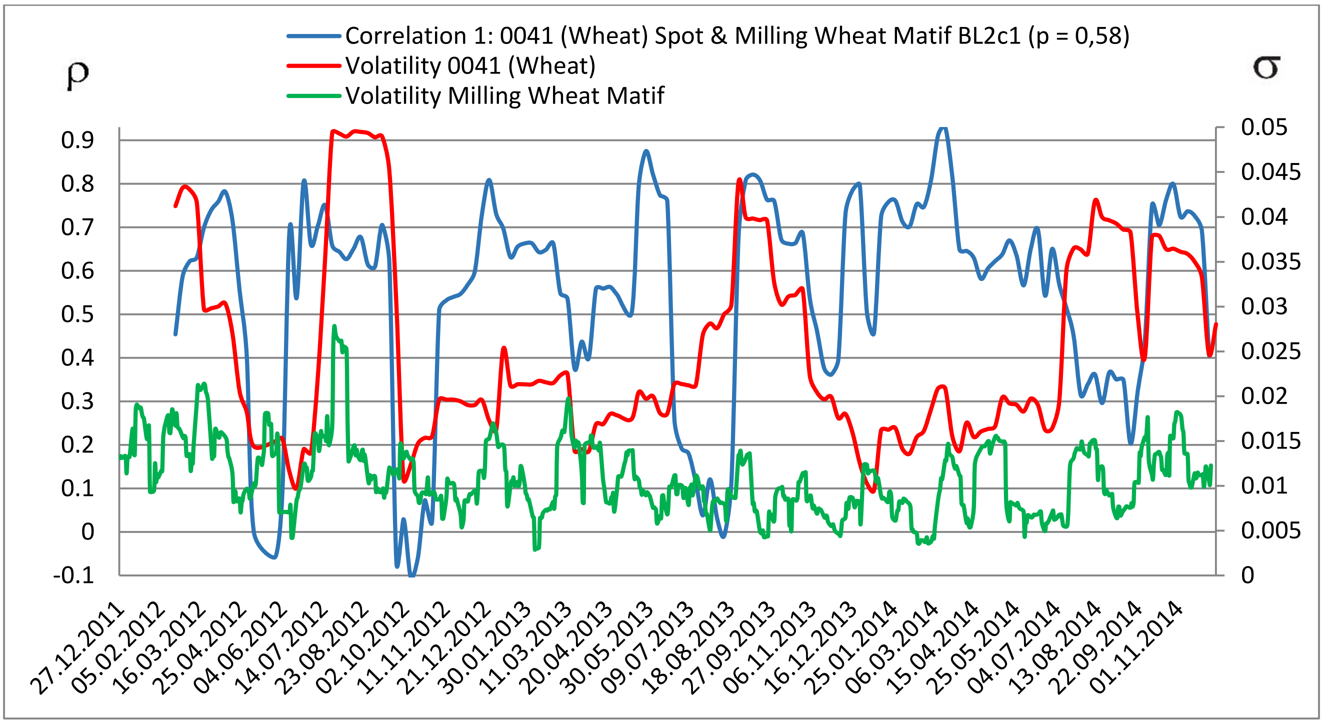

Figure 4.12 - *Volatility and Correlation of Wheat replacement and BL2c1 (10 weeks rolling average).*

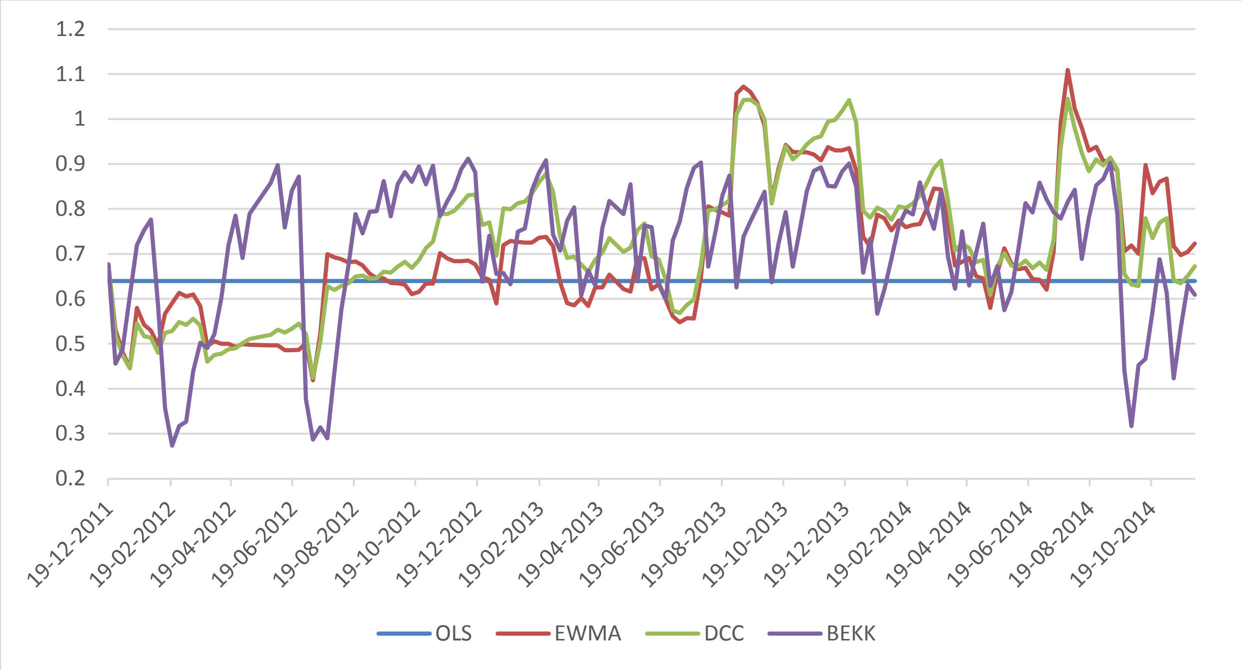

Figure 4.13 – *Wheat replacement OHR series.*

Wheat, U.S., No.2, SRW, fob Gulf (Reuters)

While this case does not represent a local cash market commodity exposure, it serves as a welcomed reference point for studying the behavior of the OHR models under a more stable correlation structure. Tables 4.25 and 4.26 show that the straightforward OLS method is still the favorable approach under most criteria. It is only in case of the in-sample average conditional hedging effectiveness that the BEKK method outperforms the OLS model by about 0.07%.

In general, the tables show that the performance race is now much tighter than what it used to be in the panels before.

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,00294242	0,00294242	0,00294242	0,00294242
$Var(R_{h,t})$	0,001314975	0,001367311	0,001419725	0,001378438
$\hat{E}$	55.31%	53.53%	51.75%	53.15%
$\overline{Var_t}(R_{s,t})$	0,002923135	0,002923135	0,002923135	0,002923135
$\overline{Var_t}(R_{h,t})$	0,001308529	0,001356309	0,001402696	0,001362541
$\bar{E}_t$	59.87%	58.97%	58.80%	59.94%

Table 4.25 – *In-sample Hedging Effectiveness (Wheat Oil World).*

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,003065436	0,003065436	0,003065436	0,003065436
$Var(R_{h,t})$	0,000668159	0,000686193	0,00076789	0,000774643
$\hat{E}$	78.20%	77.61%	74.95%	74.73%
$\overline{Var_t}(R_{s,t})$	0,003205866	0,003205866	0,003205866	0,003205866
$\overline{Var_t}(R_{h,t})$	0,00071971	0,000749434	0,000851123	0,000853348
$\bar{E}_t$	71.57%	70.89%	69.32%	69.76%

Table 4.26 – *Out-of-sample Hedging Effectiveness (Wheat Oil World).*

This is also reflected in the evolution of the hedging effectiveness in Figure 4.14. Performance gaps open up only occasionally, and if they do, they are of a much smaller magnitude compared to what they used to be in the prior cases.

It follows from this observation that the applied hedge ratios under the different models should not drift too far apart either. In fact, if we take a look at Figure 4.15 we see that the various methods move within the same range for most of the time. It furthermore highlights the only OHR outlier, which also explains the big performance deviation highlighted in Figure 4.14.

As the correlation and spot volatility peaked at the end of August 2010 (see figure 4.16), the reactive BEKK method is again the first approach to promptly increase the short futures position. However, as the spot volatility collapses in no time, the then overhedged portfolio exhibits excessive return variance compared to the other models.

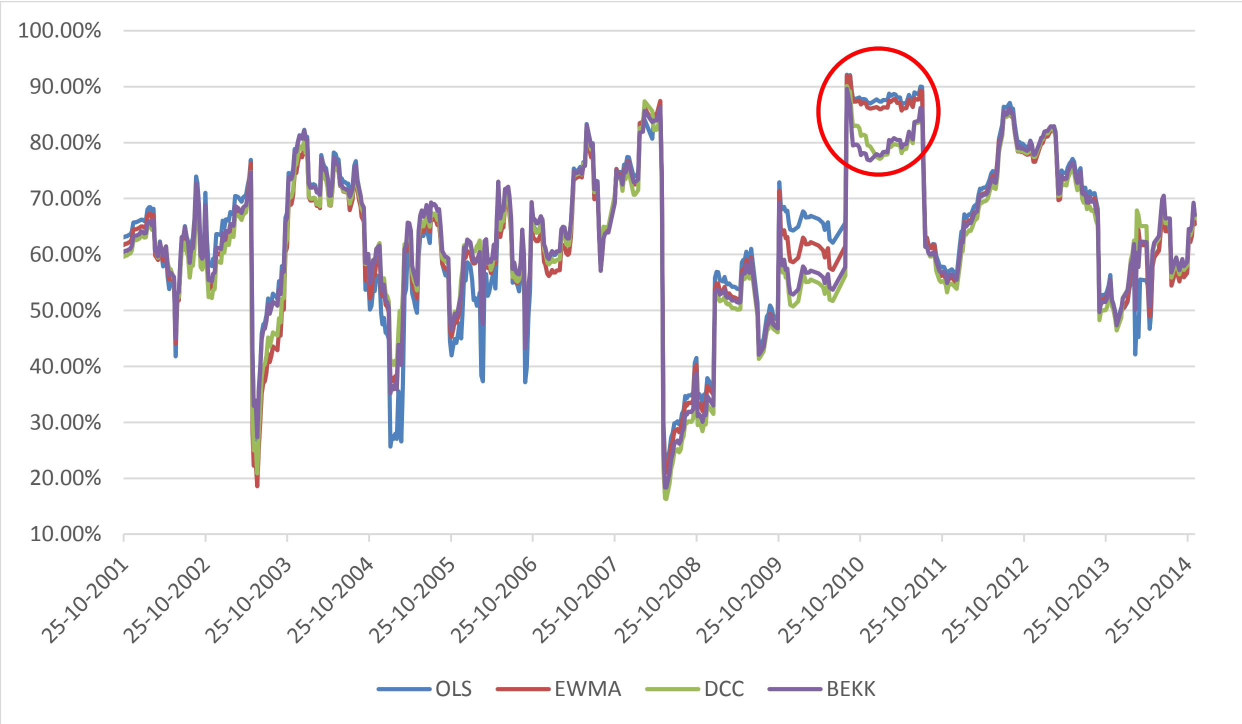

Figure 4.14 - *Conditional Hedging Effectiveness (Wheat Oil World).*

Overall, the effectiveness of hedging a U.S. cash commodity position with a CBOT futures contract seems much higher than what we have seen in the panels considered before. Especially during the out-of-sample period, which started in May 2010, the unconditional hedging effectiveness of the OLS and EWMA method almost hit 80%.

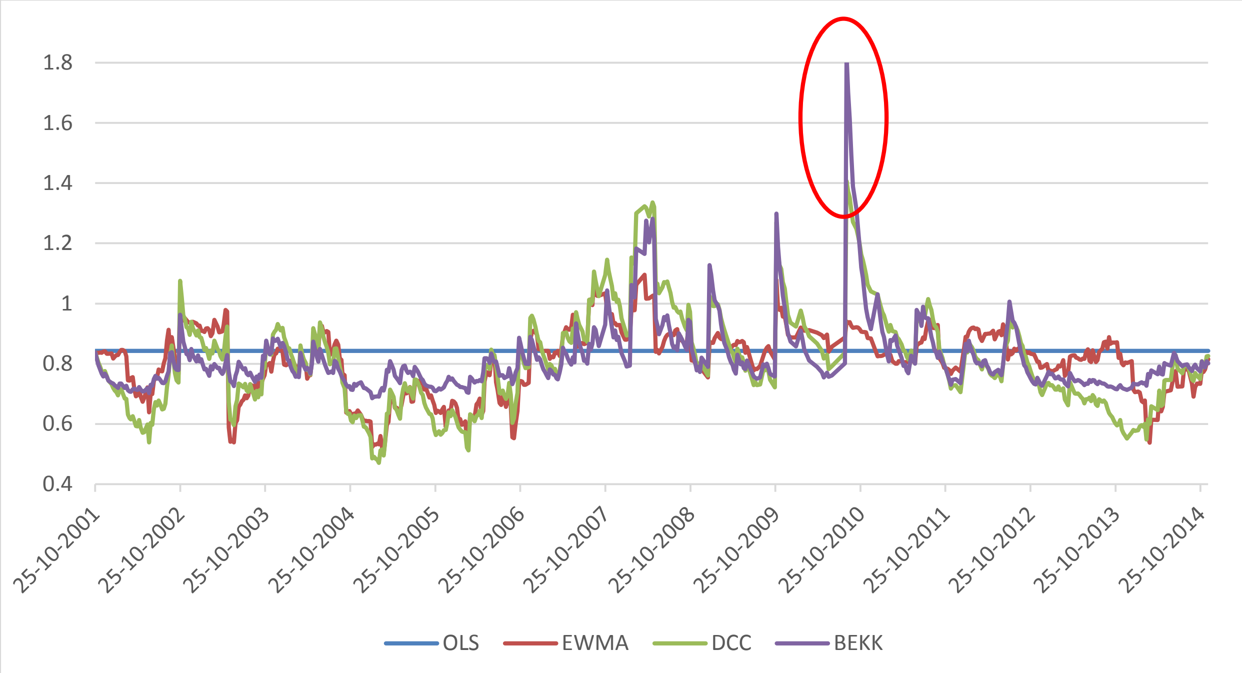

Figure 4.15 - *Wheat Oil World OHR series.*

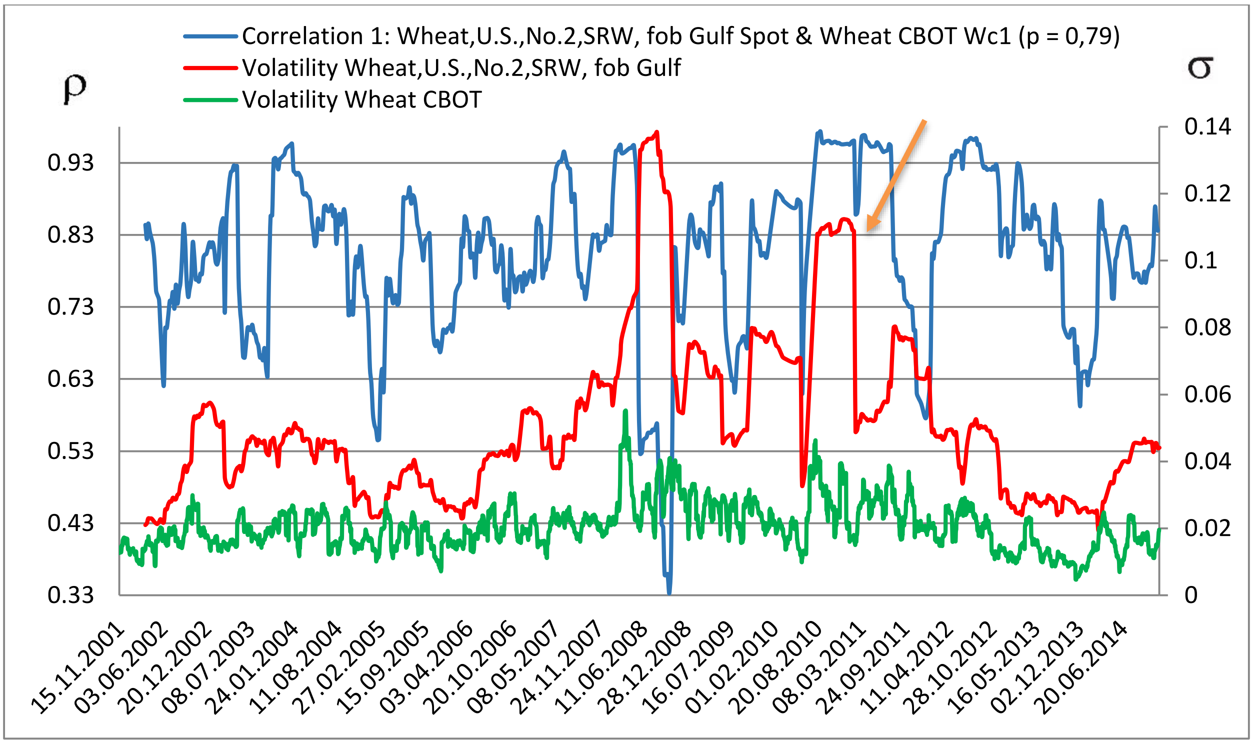

Figure 4.16 - *Volatility and Correlation of Wheat Oil World and Wc1 (10 weeks rolling average).*

4.4.3 Case Soy

In the final commodity group we first present the results for the replacement product *0300 (Soya Hipro ADM GMO)* and then head to the findings for the Oil World product *Soya pellets 48% Brazil, CIF Rotterdam.*

Soy Replacement (0300)

The soy replacement out-of-sample case represents the second instance in which the OLS method is outperformed by a time-varying method (BEKK) in terms of both conditional, and unconditional hedge effectiveness. However, this is not true for the in–sample context. And if we consider the total sample as a whole, then the OLS method performs slightly better (30.65% vs. 29.34% and 29.65% vs. 27.66% for the unconditional vs. average conditional hedge effectiveness).

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,001034044	0,001034044	0,001034044	0,001034044
$Var(R_{h,t})$	0,000753622	0,00077668	0,000780107	0,000785165
$\hat{E}$	27.12%	24.89%	24.56%	24.07%
$\overline{Var_t}(R_{s,t})$	0,001111164	0,001111164	0,001111164	0,001111164
$\overline{Var_t}(R_{h,t})$	0,000798204	0,000825088	0,000827693	0,0008282
$\bar{E}_t$	28.53%	26.23%	25.94%	25.76%

Table 4.27 – *In-sample Hedging Effectiveness (Soy Replacement).*

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,00119212	0,00119212	0,00119212	0,00119212
$Var(R_{h,t})$	0,000653354	0,000804962	0,000681221	0,000561777
$\hat{E}$	45.19%	32.48%	42.86%	52.87%
$\overline{Var_t}(R_{s,t})$	0,000807971	0,000807971	0,000807971	0,000807971
$\overline{Var_t}(R_{h,t})$	0,000505621	0,000562502	0,000534532	0,000485664
$\bar{E}_t$	35.49%	29.48%	31.77%	37.54%

Table 4.28 – *Out-of-sample Hedging Effectiveness (Soy Replacement).*

Apart from the difference in the out-of sample hedging performance, which is also visualized in Figure 4.17, there are, conform to the information in Table 4.27, no notable differences among the time-varying approaches within the in-sample period.

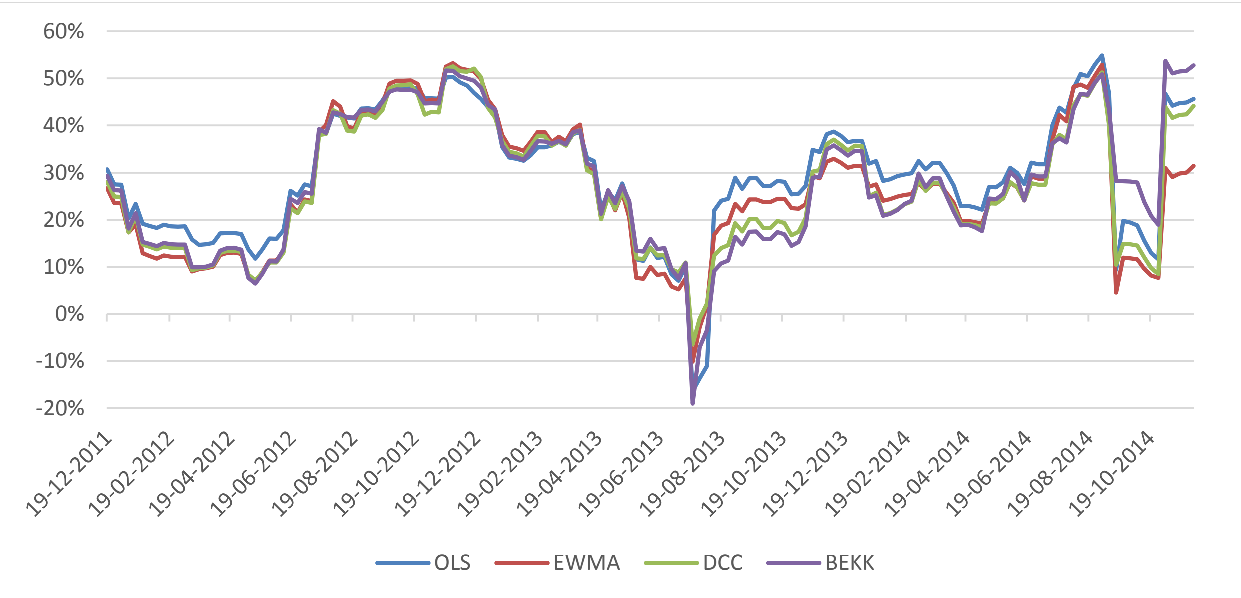

Figure 4.17 - *Conditional Hedging Effectiveness (Soy Replacement).*

One peculiar characteristic of the replacement soy panel is the low hedge ratio. As we already noted in the parameterization section, the volatility of the SMc1 futures series is relatively high compared to that of the spot soy replacement series. We therefore theoretically only need a relatively light futures position to cancel out swings in the cash market. Consequently, Figure 4.18 displays hedge ratios that are mostly below the level of 0.5.

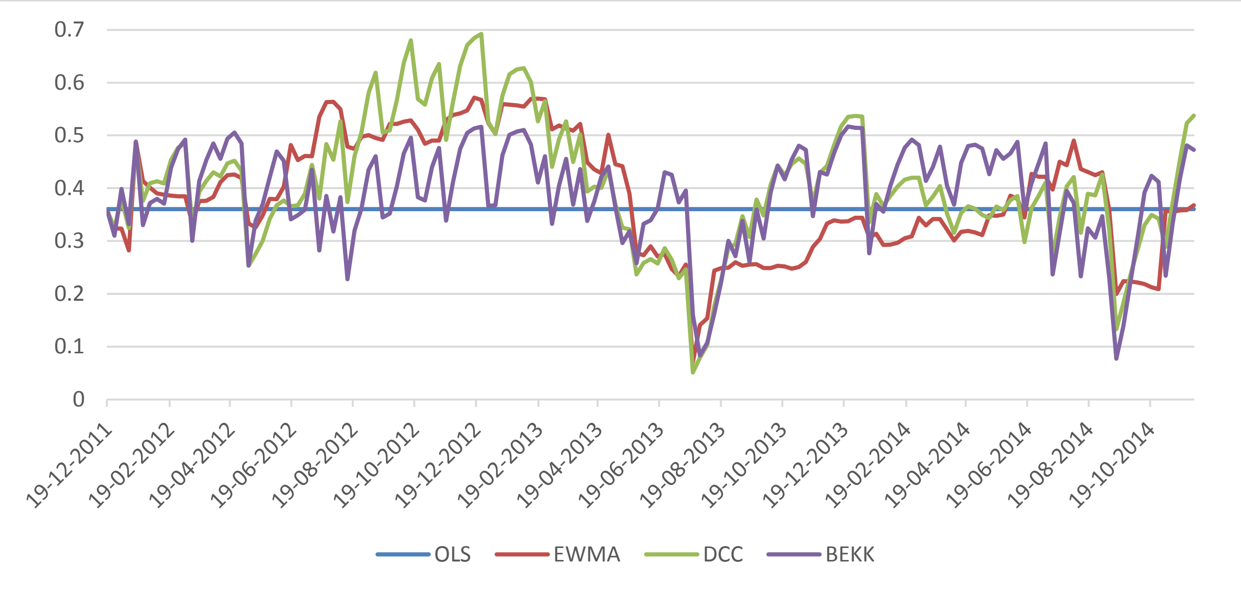

Figure 4.18 - *Wheat Oil World OHR series.*

Soya pellets 48% Brazil, CIF Rotterdam (Oil World)

The big soy sample presents yet another interesting case – the first in which the in-sample OLS hedging performance is found to be dominated by a time-varying alternative, BEKK. While, oddly enough, the OLS method wins the race in the out-of-sample period - when considering the whole sample, the BEKK method is still 1.51% (1.97%) better off in terms of (un-) conditional hedging effectiveness. To get an idea of why that is, we move on to consider Figure 4.19.

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,001320084	0,001320084	0,001320084	0,001320084
$Var(R_{h,t})$	0,000753307	0,000731973	0,000729776	0,000712647
$\hat{E}$	42.93%	44.55%	44.72%	46.01%
$\overline{Var_t}(R_{s,t})$	0,001337301	0,001337301	0,001337301	0,001337301
$\overline{Var_t}(R_{h,t})$	0,000764785	0,000742272	0,000740143	0,000723176
$\bar{E}_t$	40.99%	42.07%	42.09%	43.25%

Table 4.29 – *In-sample Hedging Effectiveness (Soy Replacement).*

Measure	OLS	EWMA (0.94)	DCC	BEKK
$Var(R_{s,t})$	0,001141313	0,001141313	0,001141313	0,001141313
$Var(R_{h,t})$	0,000485704	0,000502112	0,000538679	0,000518287
$\hat{E}$	57.44%	56.00%	52.80%	54.59%
$\overline{Var_t}(R_{s,t})$	0,001091041	0,001091041	0,001091041	0,001091041
$\overline{Var_t}(R_{h,t})$	0,000431087	0,000444036	0,000476515	0,000457185
$\bar{E}_t$	60.90%	60.26%	57.79%	59.59%

Table 4.30 – *Out-of-sample Hedging Effectiveness (Soy Replacement).*

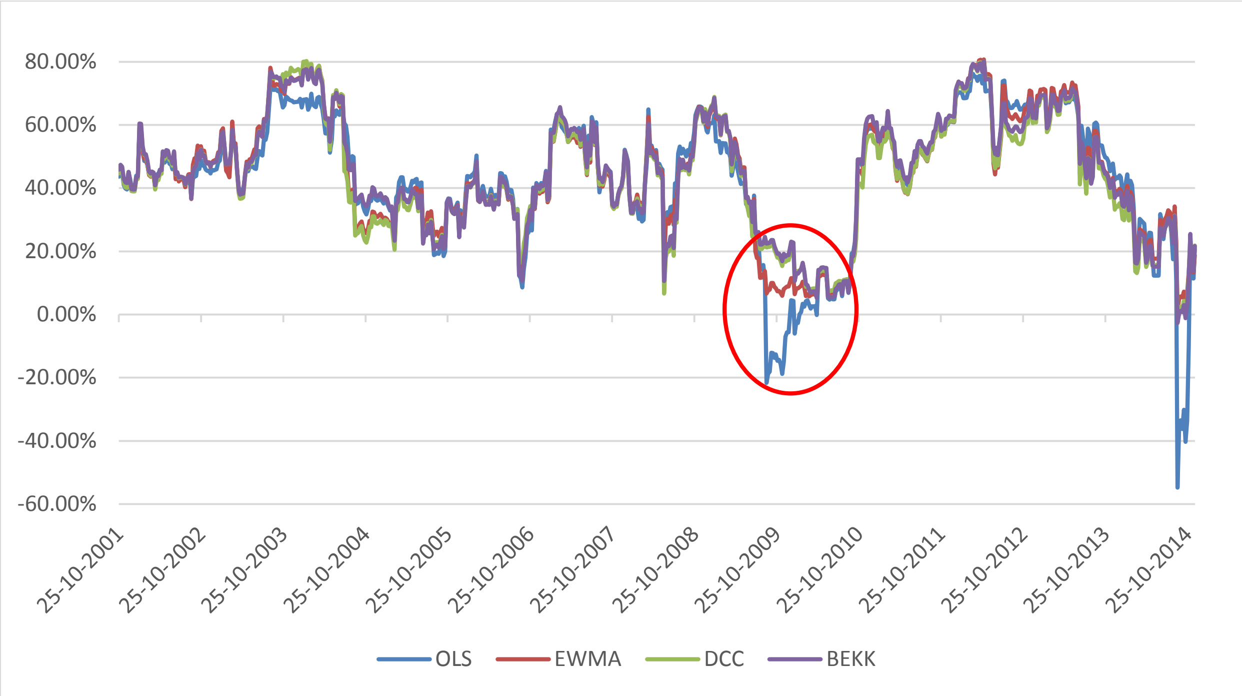

Figure 4.19 - *Conditional Hedging Effectiveness (Soy Oil World).*

As we see, there are two moments where the OLS conditional hedging effectiveness falls short of meeting the performance of the time-varying models. For the rest of the whole sample period, there are no noticeable differences among the time-invariant and the time-varying methods.

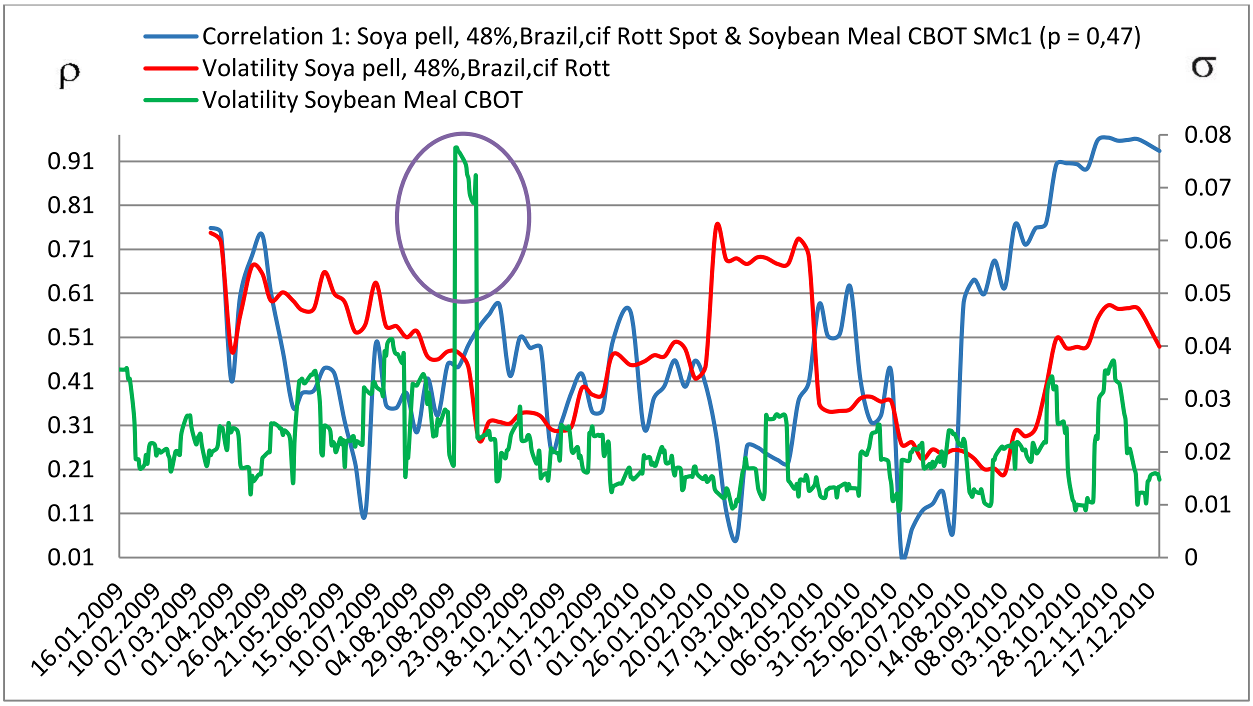

Figure 4.20 – *Oil World Soy & SMc1 correlation and volatility (10 weeks rolling average).*

Figure 4.20 provides an overview of what was going on during the time of the first performance gap, which peaked at the 24th of September, 2009. At that time there was a surge in the volatility of the futures series. Whereas one would usually respond with lowering the hedge ratio, the static OLS method sticks to its OHR of 0.535843, thus being seriously overhedged. The time-varying methods on the other hand utilize the conditional information to lower their positions in the futures contracts. Note, that the reasoning behind the second performance shortfall of the OLS method in Figure 4.19 is similar.

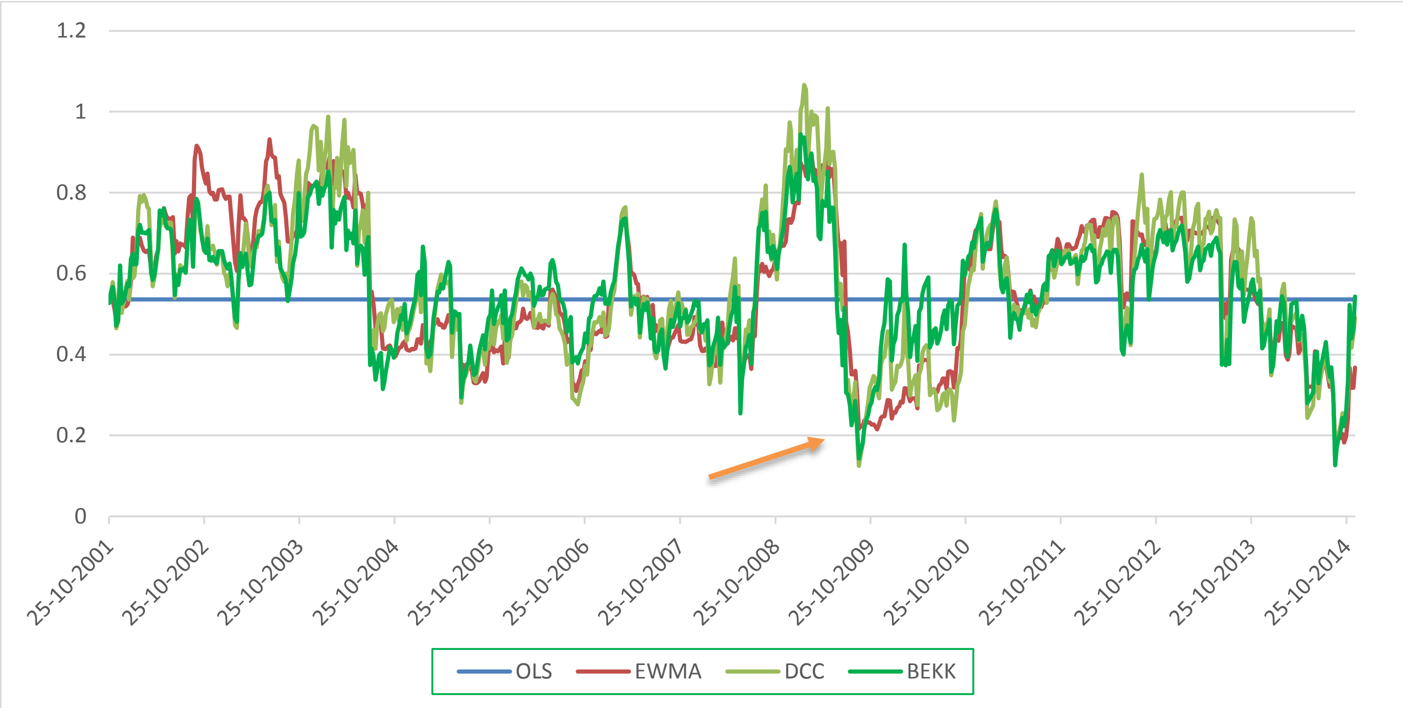

Figure 4.21 - *Soy Oil World OHR series.*

4.5 Sensitivity Analysis

All outcomes presented in Section 4.4 are based on parameter choices. As we vary those parameters, so does the output. In this section we will therefore investigate what happens as we change the smoothing parameter λ of the EWMA method or the β parameter of the OLS method in order to answer Sub Question 5:

> *Sub Question 5: "How volatile are the hedging strategy outcomes with respect to the choice of the base parameters?"*

More interestingly, we have already discussed that the degree of correlation between two return variables is likely to increase as we stretch the horizon over which we calculate those returns. Since hedges at De Heus are more likely to reside in the sphere of months rather than weeks, we will also study what happens as we extent the hedging horizon from 1 to 4 weeks.

To keep the discussion concise, we only provide the output of the λ sensitivity analysis for the full sample context. The results for various levels (.85, .90, .94, .97) of the smoothing parameter are provided in Table 4.31. As we can see from this table, both the unconditional as well as the conditional hedging effectiveness linearly increase as we increase λ. The higher the smoothing parameter, the more it approaches the OLS hedging performance. In this regard, the Oil World soy panel constitutes the only exception, where the hedging effectiveness slightly decreases as we increase λ from 0.94 to 0.97.

Lambda	HE	Corn Repl.	Corn Reut.	Wheat Repl.	Wheat OW	Soy Repl.	Soy OW
λ = 0.85	E	15.64%	37.51%	27.66%	59.87%	19.23%	45.56%
	$\bar{E}_t$	16.59%	32.26%	27.40%	61.03%	19.35%	44.49%
λ = 0.90	E	18.45%	39.22%	29.22%	60.82%	23.50%	46.63%
	$\bar{E}_t$	19.43%	33.83%	29.03%	62.05%	23.75%	45.65%
λ = 0.94	E	20.81%	40.72%	30.80%	61.68%	26.65%	46.72%
	$\bar{E}_t$	21.67%	35.01%	30.62%	62.90%	26.79%	45.92%
λ = 0.97	E	22.80%	42.09%	32.22%	62.38%	28.42%	45.99%
	$\bar{E}_t$	23.52%	35.94%	31.97%	63.57%	28.09%	45.40%
OLS	E	24.86%	44.72%	33.70%	63.06%	30.65%	45.66%
	$\bar{E}_t$	25.60%	38.20%	33.30%	63.73%	29.65%	45.20%

Table 4.31 – *Full sample OLS and EWMA hedging effectiveness (HE) for various λ.*

Overall, this observation makes sense. We have already seen in Section 4.3.2, that the OLS method yields the best outcome in most of the cases. Since the EWMA initializes the first (co-)variance estimate as the sample (co-) variance, bigger values for λ cause a higher persistency of the initial sample estimates and thus result in a closer resemblance with the OLS method, which is solely based on the sample estimates. In fact, if λ = 1, then the EWMA method equals the OLS approach.

We have just witnessed in Table 4.31 that, depending on the sample context, decreasing the EWMA parameter λ by about 0.05 can cause a hedging effectiveness reduction of about 2% to 3%. In how far then is the static OLS method susceptible to changes in beta? To answer this question, refer to Figure 4.22, which shows the distribution of the Oil World soy OLS hedging effectiveness as a function of its input parameter, beta. Three points can be distilled from this figure. First, the figure obviously peaks at its optimal value for beta, in this case $h^* = 0.535843$. Second, the graph is axially symmetric around h^* and intersects the x-axis at *0* and $2h^*$. This is logical. Apparently, if $h = 0$, then the hedged portfolio is simply the lone spot position, in which case there is no variance reduction, hence zero hedging effectiveness. If on the other hand, $h = 2h^*$ then we have

$$E = \frac{\sigma_{R_s}^2 - \sigma_{R_h}^2}{\sigma_{R_s}^2} = \frac{\sigma_{R_s}^2 - Var[R_s - 2h^* R_f]}{\sigma_{R_s}^2} = \frac{\sigma_{R_s}^2 - [\sigma_{R_s}^2 + 4(h^*)^2\sigma_{R_f}^2 - 4h^*\sigma_{R_sR_f}]}{\sigma_{R_s}^2}$$
$$= \frac{4h^*\sigma_{R_sR_f} - 4(h^*)^2\sigma_{R_f}^2}{\sigma_{R_s}^2}$$

Since $h^* = \rho\frac{\sigma_{R_s}}{\sigma_{R_f}}$, substituting this and working it out algebraically we get

$$E = \frac{4\left(\rho\frac{\sigma_{R_s}}{\sigma_{R_f}}\right)\sigma_{R_sR_f} - 4(\rho\frac{\sigma_{R_s}}{\sigma_{R_f}})^2\sigma_{R_f}^2}{\sigma_{R_s}^2} = \frac{4\left(\rho\frac{\sigma_{R_s}}{\sigma_{R_f}}\right)\rho\sigma_{R_s}\sigma_{R_f} - 4\rho^2\sigma_{R_s}^2}{\sigma_{R_s}^2} = \frac{4\rho^2\sigma_{R_s}^2 - 4\rho^2\sigma_{R_s}^2}{\sigma_{R_s}^2} = 0$$

Third and most importantly, we can see that the slope of the graph gets steeper as we move to the tails of the distribution. If we for example decrease beta from 0.1 to 0, then this is associated with a 15% drop in hedging effectiveness. If, on the other hand, we shift beta by the same magnitude from 0.5 to 0.4, then the associated loss is "only" 2%. Note, that the shapes for the other samples look very similar. This implies that the OLS solution is relatively robust.

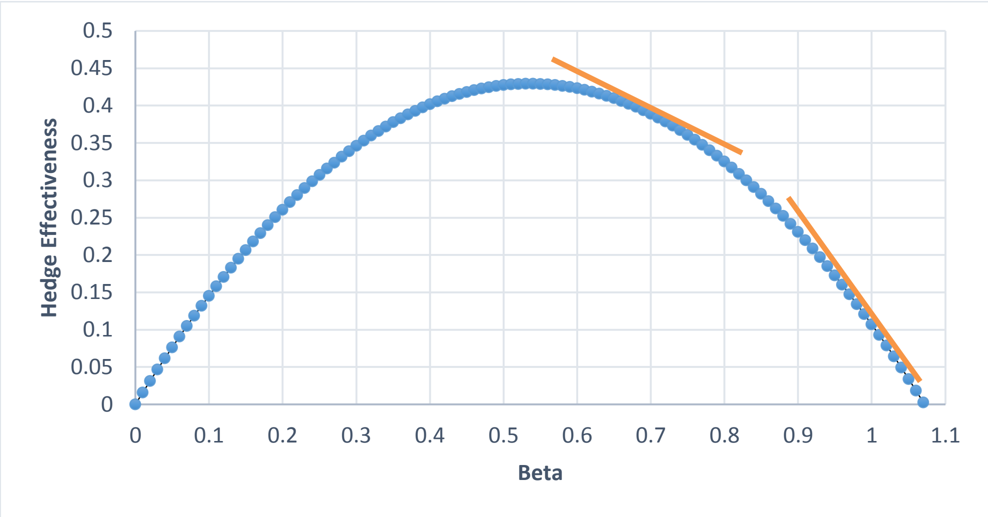

Figure 4.22 – *Oil World soy OLS distribution for varying betas* ($h^* = 0.535843$).

To get an impression of how the hedging effectiveness develops as we extend the hedging horizon, we matched the price series such that we get non-overlapping, monthly returns. As we have not noticed much performance differences under the time-varying methods in Section 4.3.2, we will only discuss the OLS, EWMA and diagonal BEKK methods. Moreover, for reasons discussed in Chapter 3, we will only provide the output for the full sample context of the external vendor panels.

Parameters	Corn Reuters	Wheat OW	Soy OW
β	0.929397*	0.882999*	0.648100*
μ_1	-0.001594	0.010522***	0.003394
μ_2	0.004705	0.011812***	0.005610
C_{11}	0.001275***	0.0000271	0.002188*
C_{12}	0.000653***	0	0.001566*
C_{22}	0.000334***	0.0000122	0.001121*
A_{11}	0.062935	0.257541*	0.471200*
A_{22}	0.069913	0.244583*	0.342050*
B_{11}	0.839981*	0.966901*	0.608452*
B_{22}	0.947272*	0.971411*	0.859135*

Table 4.32 – *OLS and BEKK parameterization output.*

The OLS and BEKK parameterization output is given in Table 4.32. As for the BEKK parameter estimation, the coefficient matrix in the Oil World wheat sample had to be restricted to be diagonal in order to ensure that the covariance matrix is positive definite. For the remaining two samples, the coefficient matrix is rank 1. Refer to Appendices J and K for the EViews OLS and BEKK estimation output, respectively.

Entering the parameters of Table 4.32 into our custom model, this then yields the following outcome. See Table 4.33. Note, that the values in brackets refer to the full sample performance score of the weekly return series. The table shows that by enlarging the hedging horizon, the effectiveness increases substantially. For the Reuters corn and Oil World wheat sample, the hedging effectiveness improves, on average, by 20%, whereas the Oil World soy sample is about 15% better off.

Method	HE	Corn Reuters	Wheat OW	Soy OW
OLS	$\hat{E}$	**62.28%** (44.72%)	**83.99%** (63.06%)	**62.63%** (45.66%)
	$\bar{E}_t$	**58.71%** (38.20%)	**84.02%** (63.73%)	**59.87%** (45.20%)
EWMA (.94)	$\hat{E}$	**60.07%** (40.72%)	**83.66%** (61.68%)	**61.92%** (46.72%)
	$\bar{E}_t$	**55.09%** (35.01%)	**83.69%** (62.90%)	**59.36** (45.92%)
BEKK	$\hat{E}$	**62.63%** (39.11%)	**83.59%** (60.45%)	**62.78%** (47.62%)
	$\bar{E}_t$	**58.50%** (34.70%)	**83.71%** (63.18%)	**60.27%** (46.70%)

Table 4.33 – *Full sample hedging effectiveness based on monthly return series.*

Moreover, the performances of the OLS and BEKK methods seem to be at par in the monthly interval panel. Here, the BEKK method slightly outperforms the OLS method in the Oil World soy sample, both conditionally and unconditionally (62.78% & 60.27% vs. 62.63% & 59.87%). However, the converse is true in the Oil World wheat sample (83.59% & 83.71% vs. 83.99% & 84.02%). As for the Reuters corn sample, the BEKK method scores higher judged by the unconditional hedge effectiveness criterion (62.63% vs. 62.28%), but finds itself slightly outstripped based on conditional hedge effectiveness (58.50% vs 58.71%).

In general, those results were to be expected and are in line with findings in literature (Ederington, 1979). We have argued both theoretically and empirically that correlation is a key ingredient to a successful hedging strategy. We have also argued and shown that the degree of correlation increases as we stretch the return intervals. It follows from the combination of both insights that a higher hedge horizon ought to lead to higher hedging effectiveness.

4.6 Application of Mean-Variance Hedging Strategies

While the minimum variance approach is the most effective one in terms of variance reduction, it does so without concern for the average rate of return of the futures products. However, we have revealed in Section 4.2 that the rate of return of the futures are not negligible. In this section we therefore categorize the impact of the reduction in the average rate of return of the spot, due to variance minimization. Moreover, we apply the mean-variance OLS, EWMA, DCC, and BEKK hedging strategies to see how the hedging strategies would change if we were to take account of the mean return.

Computationally, there does not change much with respect to the minimum-variance strategy. According to (3b), the only alteration that we need to make for the mean-variance hedges is to subtract from the minimum-variance hedge ratio the futures expected rate of return divided by the product of its variance times the risk aversion parameter (A). To avoid redundancy, we therefore also abstain from a detailed explanation of how the OHR varies over time according to the different mean-variance models, but rather only provide occasional examples.

For the expected value we simply choose the sample average based on the in-sample observations. The conditional variance is drawn from the conditional covariance matrix. To keep the discussion concise, we only study the complete sample period. As for the risk-aversion parameters, we consider the values 1 through to 10. The risk-return planes used to illustrate the inherent trade-off show annualized values, where

$$R_a = (1 + R_w)^{52} - 1 \quad \text{and} \qquad \sigma_a = \sigma_w \sqrt{52}$$

where R_a (r_w) and σ_a (σ_w) are the annualized (weekly) average and standard deviation, respectively.

4.6.1 Case Corn

Figure 4.23 provides an overview of the annualized risk return plane of the corn replacement hedge portfolio. It accommodates the realized outcomes of all mean-variance and minimum-variance hedging strategies, as well as the unhedged outcome and the static frontier. As for the mean-variance strategies we have sets of outcomes as each risk aversion parameter yields a unique outcome.

The 'static frontier' represents the risk-return frontier of Figure 2.1. That is, it shows the set of realized average rates of return and standard deviations of the hedged portfolios for $0 \leq h \leq 1$. The outcome of any time-invariant hedging strategy will have to lay on this line for values of h between 0 and 1 or on the extrapolated line for values above 1 or below 0. Hence, the unhedged and the OLS strategy outcomes are all elements of this solution set. Outcomes of time-varying hedging strategies are, of course, free to take on risk and return values offside of this line.

Just as in Figure 2.1 the hedge ratio of the static frontier increases as we move down the vertical axis. The unhedged portfolio is located at the top of the frontier, where *h = 0*. The point with the greatest curvature on that line is given by the static (OLS) minimum-variance. We can therefore distill from Figure 2.1 that *mean-variance strategies* generally feature *lower hedging ratios* than minimum-variance hedging strategies as all mean-variance OLS outcomes are above the minimum-variance OLS point. This result is to be expected and directly follows from formula (3b). As the expected futures return is subtracted from the expected spot return, concern for the expected portfolio rate of return will lead to lower offsetting futures positions than would otherwise be the case.

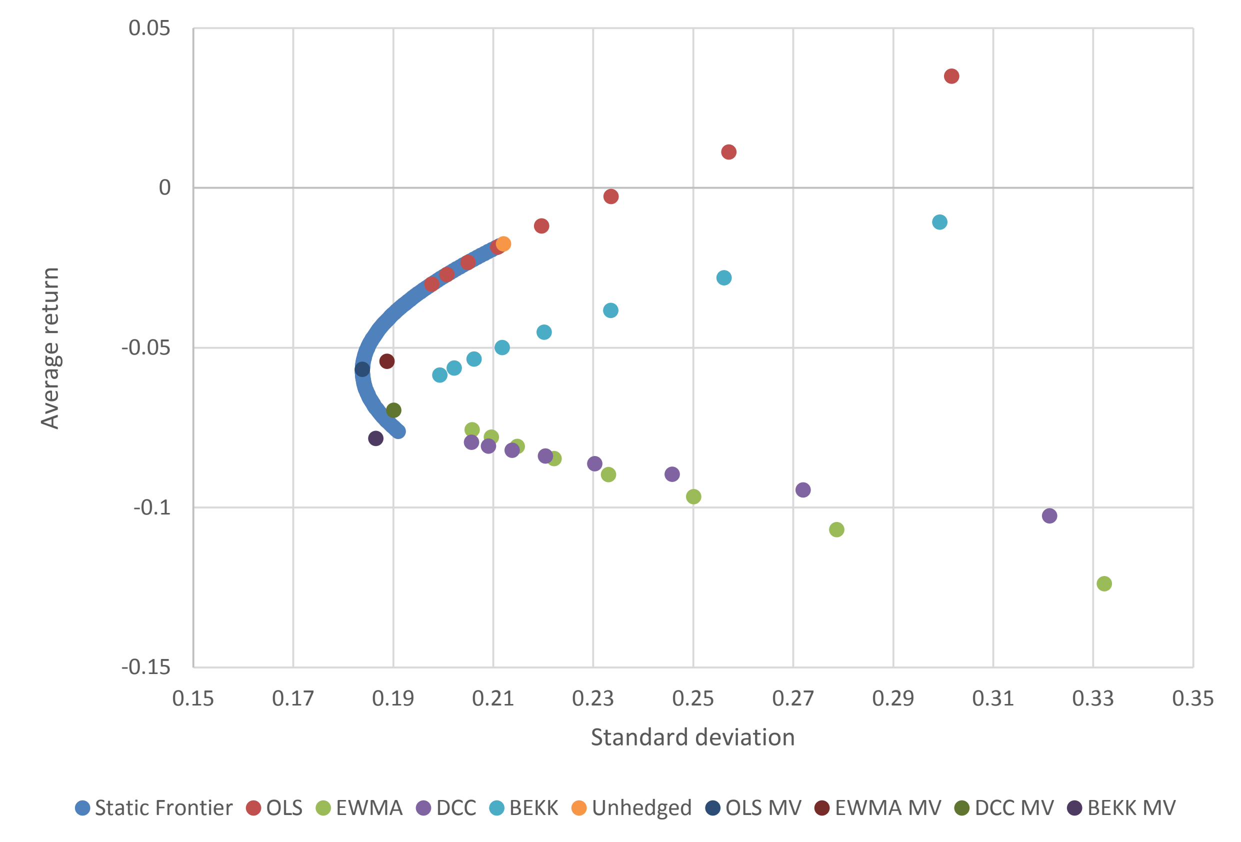

Figure 4.23 – *Corn Replacement annualized risk-return plane.*

The next interesting point that follows from Figure 4.23 is that not only were the time-varying methods less successful in minimizing the portfolio variance, most of them (i.e., DCC & BEKK) also reduced the average return to a greater degree. The time-varying mean-variance strategies all produce outcomes that are far off the efficient risk-return frontier, with the EWMA and the BEKK strategy featuring the worst results.

As discussed in Section 2.1.2 the mean-variance strategy approaches the minimum-variance strategy as we increase the risk aversion parameter. It follows that points more to the right refer to outcomes of lower risk aversion parameters, while higher coefficients make them resemble their minimum variance pendants.

When considering the Reuters sample we observe somewhat identical results. Just as in the corn replacement sample, the time-varying mean-variance strategies consumed a bigger portion of the average return, while also providing for a lower degree of variance reduction. Likewise, most of the solutions of the high risk aversion time-varying mean-variance strategies are below the efficient frontier. Only for lower risk aversion coefficients ($A \leq 3$) do those strategies start to pay off.

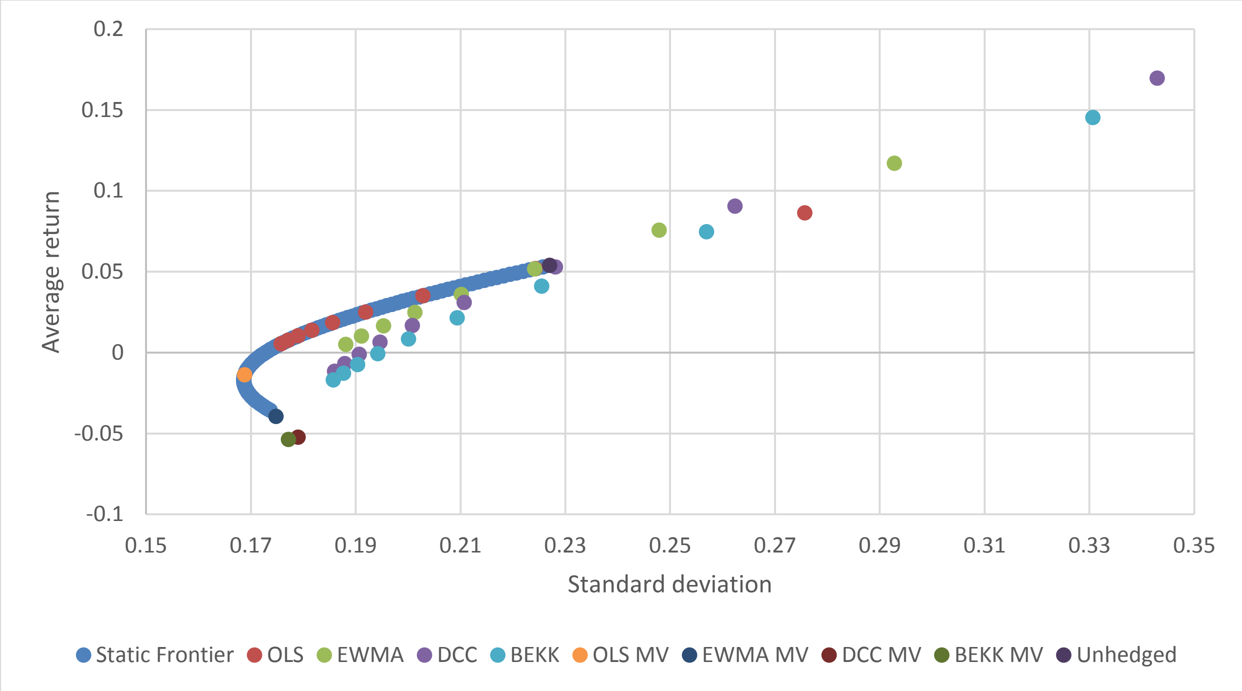

Figure 4.24 – *Corn Reuters annualized risk-return plane.*

4.6.2 Case Soy

In the case of soy, the results look quite different in two regards. First, the location of the risk-return outcomes of the various minimum-variance hedging strategies, both time-invariant and time-varying, is much denser. Only the BEKK MV strategy falls short of delivering the average rate of return of the other strategies.

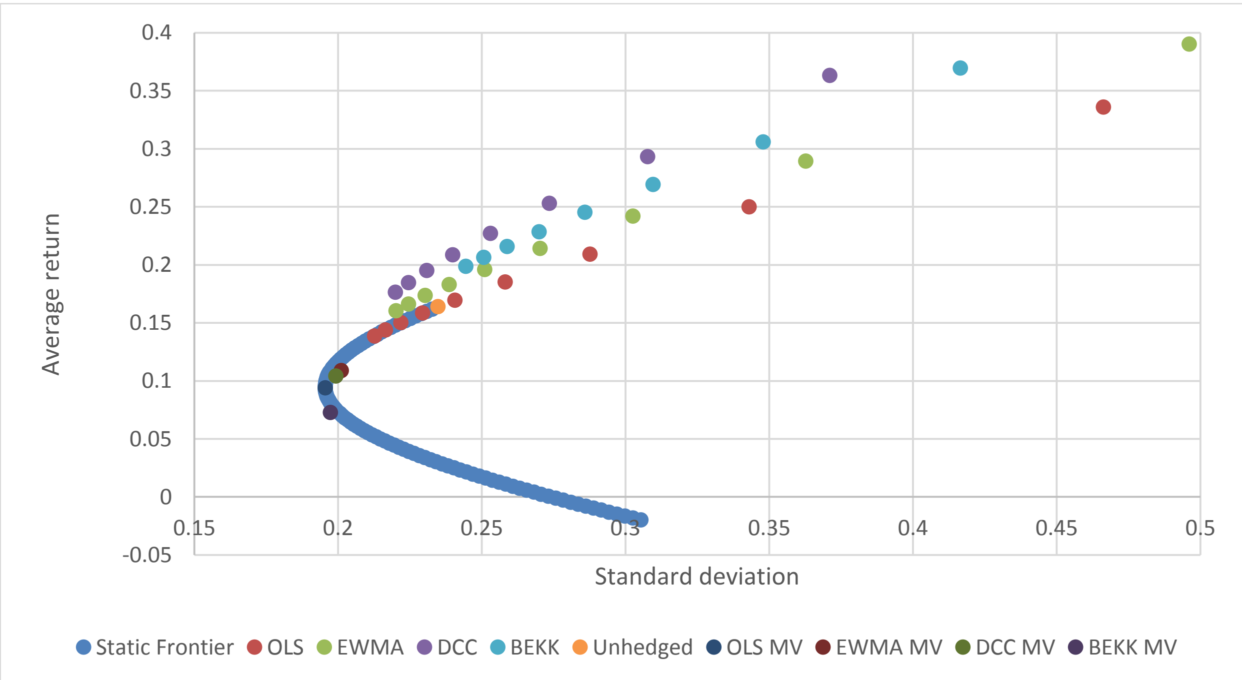

Figure 4.25 – *Soy Replacement annualized risk-return plane.*

Second, and more interestingly, all the time-varying mean-variance strategies provide outcomes that are more efficient than any static hedging strategy with the DCC method performing best. This is true for any risk aversion coefficients in the range from 1 to 10. For each degree of standard deviation in the feasible set, there is a time-varying (i.e., EWMA, DCC, and BEKK) hedging mean-variance strategy that achieves a higher average rate of return than that of a static mean-variance strategy.

Notice, how most of the time-varying strategies propose solutions based on optimal hedge ratios that are, in general, lower than zero (since they are located at the top right from the unhedged solution). This might seem as a counter-intuitive outcome at first as we are, on average, effectively increasing the outright exposure in a hedging strategy. Keep in mind, however that we are now optimizing both, risk and return. Given the low static minimum-variance ratio, due to the high volatility of the soy futures series (see Section 4.4.3), contemporaneously optimizing for the mean leads to hedging strategies that frequently happen to involve negative hedge ratios.

But how is it even possible that such positive performance differences establish? Realize, that since we are naturally short the futures contract in a long spot hedge, we would ideally want to observe an inverse relationship between the rate of return and the applied hedge ratios in order to maximize the return. That is, if the futures return is positive, we would want to go long in the futures contract, which is achieved by a negative hedge ratio and vice versa for negative futures returns.
As we can see in Figure 4.26 the DCC mean-variance hedging strategy (here plotted with a risk aversion coefficient of A = 7) neatly implements the inverse-relationship principle. For most of the time, whenever the average return (plotted as a 10 weeks rolling average), is positive, the DCC(7) hedge ratio is negative, and vice versa. Static hedge ratios naturally do not possess this flexibility - hence the average portfolio return difference. In order to visualize in how far a dynamic mean-variance hedging strategy differs from its minimum-variance pendant, Figure 4.26 also includes the minimum variance DCC strategy.

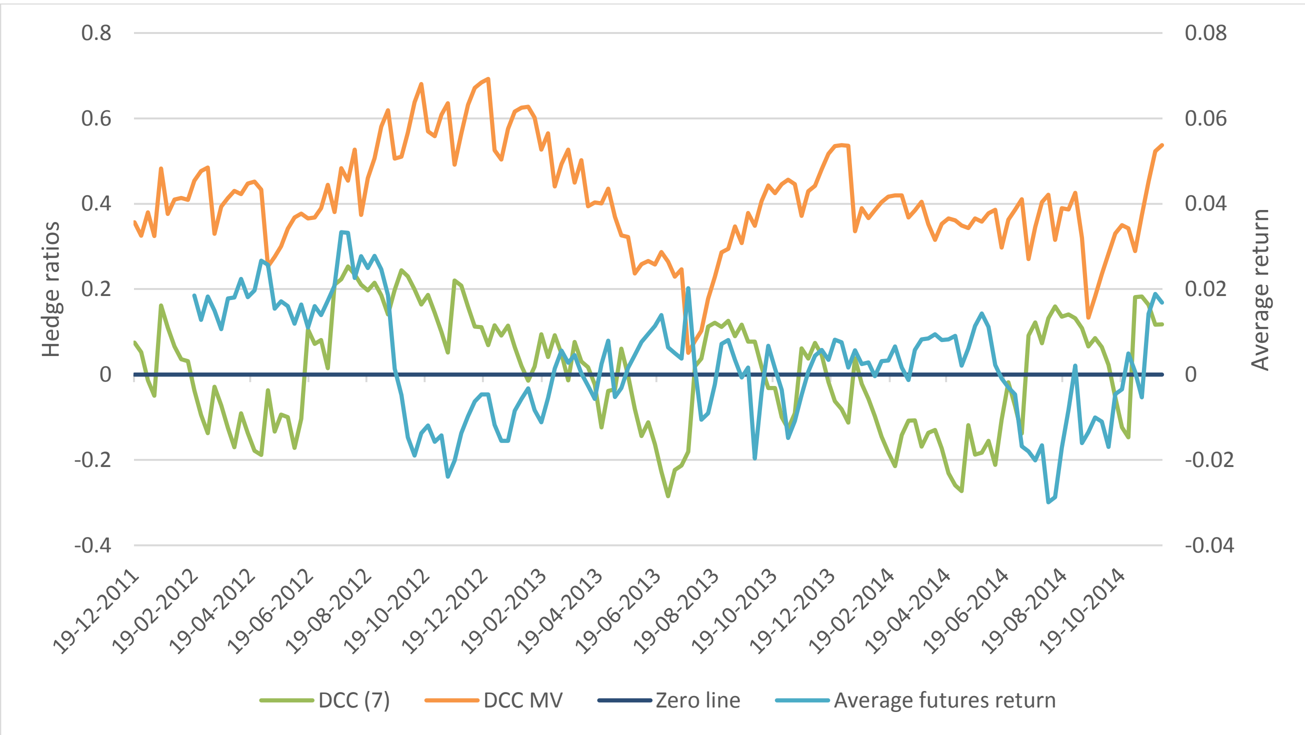

Figure 4.26 - *Evolution of the DCC(7), DCC MV OHR and the average futures rate of return.*

As can be seen, the overall patterns of the two approaches are quite similar. Moreover, as described in the paragraph above, the minimum-variance hedge ratios are significantly higher. While this contributes to a greater portfolio variance reduction, the inverse relationship now only holds in relative terms (i.e., when the average return decreases, the hedge ratio increases and vice versa) but not anymore in complete terms (i.e., when the average return is negative, the hedge ratio is positive and vice versa). We thus suffer a reduction in the average rate of return.

The results described above also hold for the Oil World soy case, even though the performance differential is a bit less pronounced in this case and the proposed strategies tend to rely on higher hedge ratios. See Figure 4.27.

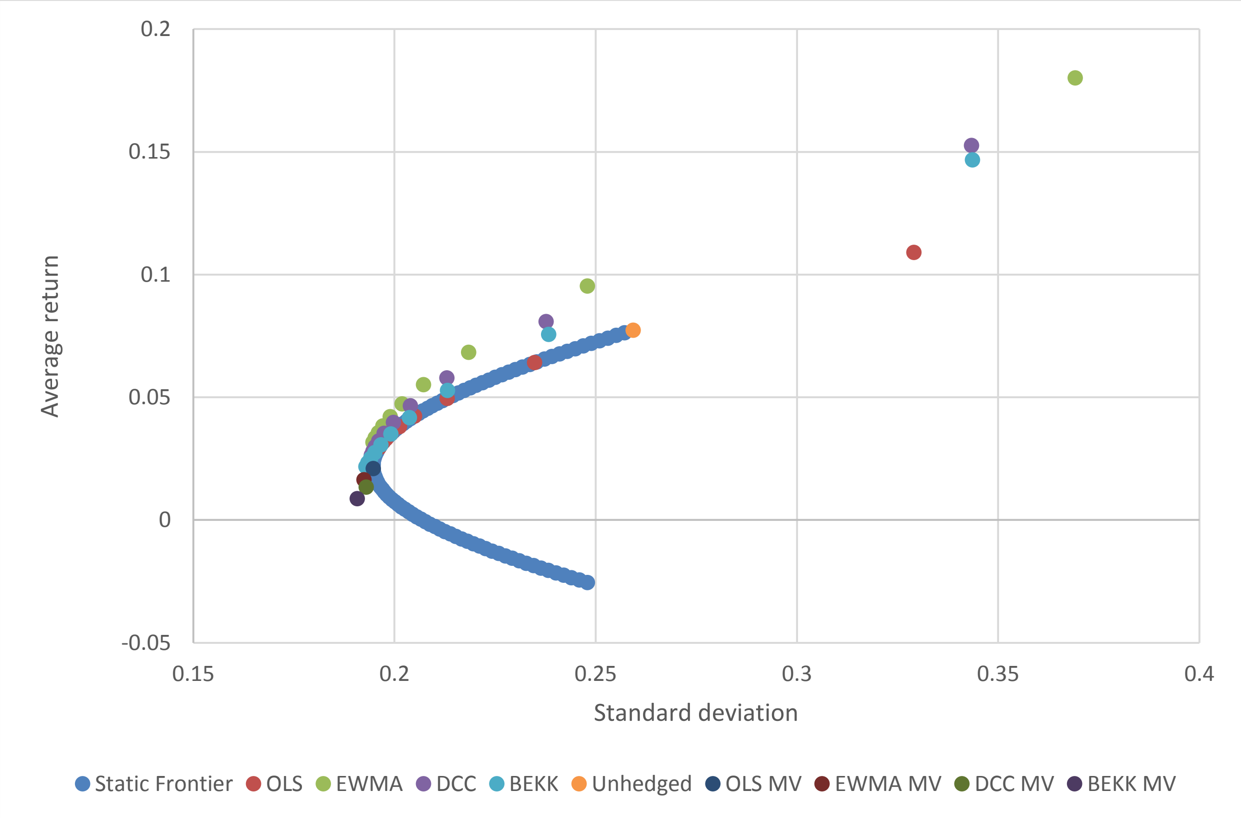

Figure 4.27 – *Soy Oil World annualized risk-return plane.*

4.6.3 Case Wheat

Just as in the corn replacement case, the wheat replacement sample presents a case where the EWMA and DCC mean-variance hedging strategies malfunction. Both feature realized risk-return solutions that are way off the static frontier, irrespective of the risk aversion coefficient under consideration.

To investigate the source of those performance discrepancies, we again take a look at the hedge ratio evolution. Figure 4.29 shows how the hedge ratios of the DCC (A=6) and BEKK (A=10) mean-variance strategies change over time. For the same time frame it depicts how the average return (measured as a 10 weeks rolling average) develops. Clearly, Figure 4.29 suggests how the inverse relationship holds more frequently for the BEKK strategy than it does for the DCC strategy, which exhibits a pattern very similar to that of the EWMA (A=6) strategy.

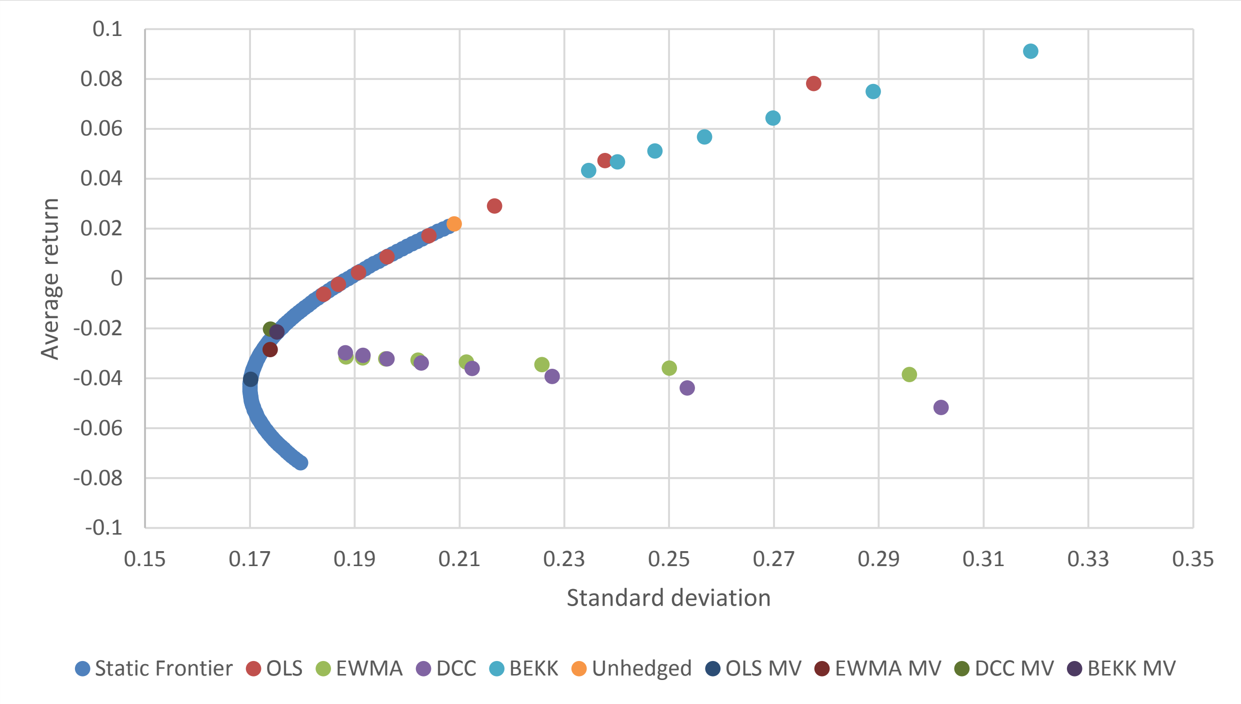

Figure 4.28 – *Wheat Replacement annualized risk-return plane.*

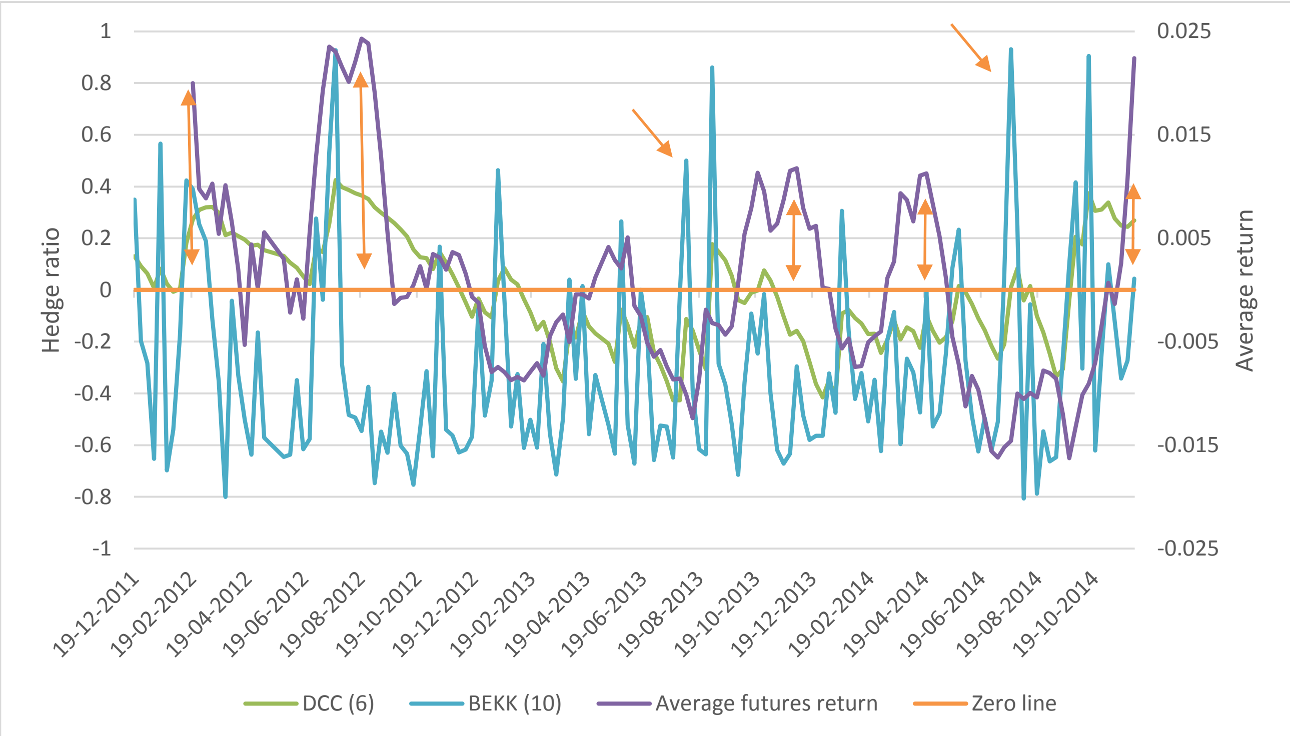

Figure 4.29 – *Evolution of the DCC (6), BEKK (10) OHR and the average futures rate of return.*

Consider for example the period from July 2012 to September 2012 where the futures reaches its sample average high. While both strategies propose positive hedge ratios at the start of this time frame, the BEKK quickly turns the short futures position in a long position, whereas the DCC strategy continues to hold its short futures position. Likewise, the BEKK method proposes a short position when the SMc1 series hits its 2013 average temporal low in mid August, while the DCC method continues to apply negative hedge ratios.

Turning to the wheat Oil World sample, we, similar to both soy sample cases, again observe a scenario where the time-varying mean-variance hedge strategies produce risk-return results that are superior to those of static strategies. In this case, however, the performances differences also unfold for higher risk aversion coefficients and in the realm of higher hedge ratios. The annualized standard deviation of the mean-variance OLS (A=3) strategy for example is 27.41% and offers an average rate of return of 8.52%. The EWMA (A=4) achieves a slightly lower standard deviation of 27.19%, yet also provides an annualized rate of return of 10.81%.

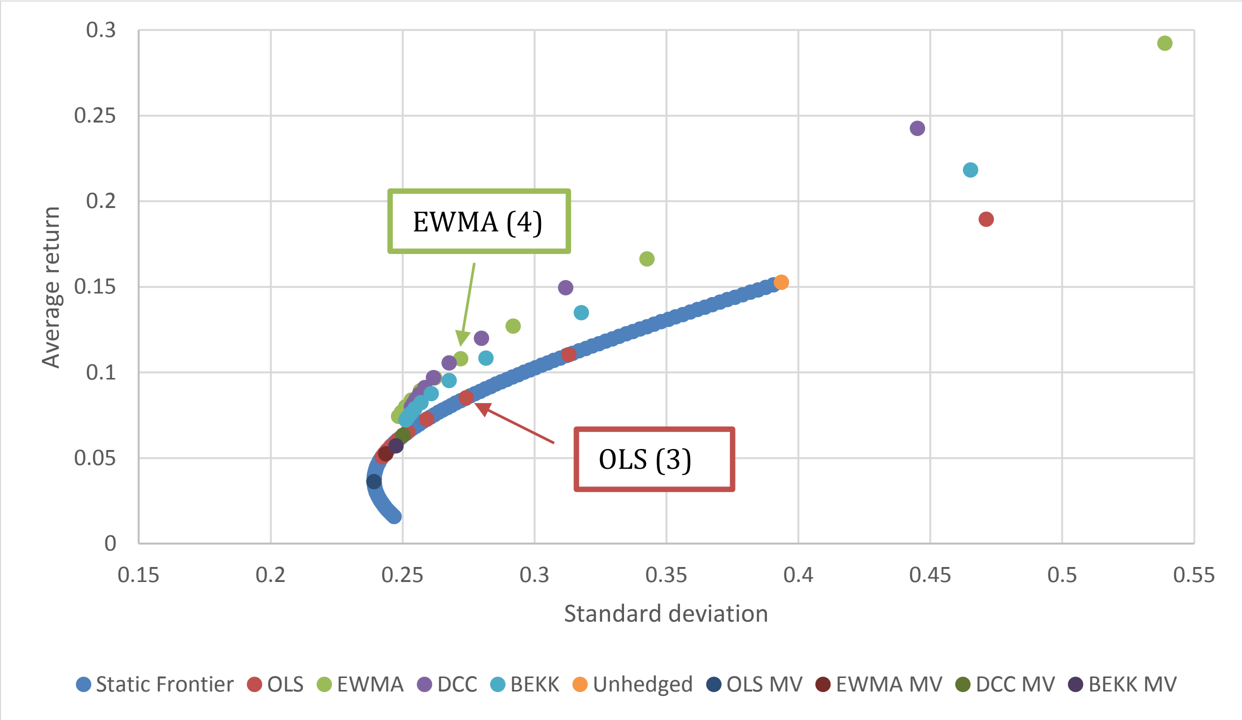

Figure 4.30 – *Wheat Oil World annualized risk-return plane.*

4.7 Discussion and Implications

This section discusses the main findings from the analysis. Section 4.7.1 examines the results from the minimum-variance hedging strategies. Supplemental insights from the mean-variance hedging strategies are dissected in Section 4.7.2.

4.7.1 Minimum-Variance

The central theme of our research has been whether or not time-varying hedge ratios would contribute to a higher hedging effectiveness. Indeed, we have seen that those time-varying methods (especially the vanilla diagonal BEKK GARCH model) were capable of utilizing recent market information to periodically outperform the standard time-invariant OLS regression method in each of the six samples. Moreover, it is well established that the maximum benefit of the time-varying GARCH (i.e., DCC and BEKK) models can only be obtained if the parameters of the models are kept up to date. This has not been the case in this research, where the parameters were kept constant throughout the out-of-sample period. So there is some extra potential in the models, which has not been tapped in this paper. Furthermore the literature review has brought forward the point that the unconditional hedging effectiveness measure underestimates the true performance of the conditional models.

	Corn Rep.		Corn Reut.		Wheat Rep.		Wheat OW		Soy Rep.		Soy OW	
	In	Out	In	out	In	Out	In	Out	In	Out	In	Out
OLS	✓		✓	✓	✓	✓	~	✓	✓			✓
TV		✓								✓	✓	

Table 4.34 – *Meta performance comparison of time-varying (TV) models and OLS.*

Nonetheless, the preference clearly goes out to the time-invariant OLS method. It is much more efficient in the sense that it can easily be estimated. Moreover, it seemed to be more effective in most of the sample settings, too - even in terms of the unconditional measure. See Table 4.34 for a global overview of which model had the lead performance wise, per sample. While the literature research already indicated that this would likely be the case for the in-sample setting, the analysis has shown that in four out of six cases, the OLS method also outperforms the time-varying methods in the out-of-sample context. Another factor that supports the preference for the time-invariant model is *transaction costs*. To capture the true benefit of the time-varying models in practice, the futures position in the hedged portfolio would frequently have to be rebalanced to account for changes in the conditional covariance structure. As we learned from the figures that depict how the OHRs of the different models evolve over time, there is considerable volatility in the optimal ratio of futures contracts relative to the spot position. This, of course, translates into substantial transaction costs. Thus, if we net the potential for additional performance gain due to refreshing parameter estimation with the potential for transaction costs, then the preference of the static OLS method over the dynamic GARCH models becomes even stronger.

Given the observed dominance in terms of variance reduction of the OLS method – even in the out-of-sample periods – one itching question remains. If the time-varying models are built to exploit shifts in the variance and covariance structure of a set of variables and if we indeed observe heteroscedasticity in our return series (which has been shown in both 4.2 as well as 4.3), then why do those time-varying methods fail to outperform the static OLS method?

According to Harris *et al*. (2009) this may be due to a fourfold set of factors, amongst which the following two. First, while the true, unobservable hedge ratio (that is, the series of hedge ratios based on the true conditional population parameters) may be time-varying, it may not be time-varying to the degree that it would be financially beneficial to estimate deviations from the average OHR. Second, the problem could reside in the model itself. It could be that either the model does not contain enough information about the moment dynamics of the hedged and hedging variable, or it could be that the model incorrectly uses this information, i.e., that it is miss-specified.

While we can of course not investigate whether the first factor plays a role, the second explanation may arguably be applicable in our case. In literature, model estimation is usually done on the basis of samples that contain a number of observations in the quadruple digit sphere. In this research, however, especially the replacement sample size was severely restrained. In Section 4.3.1 we have seen that it was exactly in this small sample size context, where model estimation was problematic.

The claim that sample size could be an issue is further substantiated by the observation that it were the big, external sample contexts in which the performance gap between time-varying and time-invariant models was the closest. For the Oil World soy and wheat series, the in-sample hedging effectiveness of the BEKK model was even superior to that of the OLS method.

Irrespective of the methodologies to estimate the optimal hedge ratios, do the available futures contracts in general provide a useful instrument to hedge local cash commodity exposure? The answer to this question is contingent on what we define as the local cash market variable. If we understand this to be the internal replacement values, then the outlook is rather daunting. With unconditional weekly return sample correlations of 0.50, 0.58, and 0.55 for corn, wheat, and soy, respectively, the potential for variance reduction was shown to be only within the range of 25% to 35%. While this relates to weekly exposures, the effectiveness in the monthly return context is not likely to exceed the 40% range either. In this context, however, we should not forget that the replacement sample size is limited. Single correlation dips therefore have a material impact on the measured hedging performance.

If, on the other hand, we measure cash market exposure with data from external parties, then the outlook for hedging potential is more promising. Those series exhibited a much higher degree of correlation with their corresponding products on commodity exchange markets and featured monthly variance reduction in the 60% dimension.

4.7.2 Mean-Variance

In Section 4.2 we have established that there is reason to believe that the zero mean assumption would not hold for our futures return series. Upon locating the minimum-variance strategy outcomes in the risk-return plane, this impression has further been supported. In fact, going from a completely unhedged portfolio to an OLS minimum-variance hedge portfolio is accompanied by a drop of about 5% in the average, annualized rate of return for all cases except the wheat Oil World sample where the cut is even about twice as high. Therefore, unless the investor is highly risk averse, one should carefully consider in how far one wants to sacrifice parts of the expected return for a greater variance reduction. In case of wheat- and soy replacement for example, one would have to deploy a risk aversion coefficient of 100 or 40, respectively, in the OLS mean-variance strategy, in order to approach the OLS minimum variance outcome. Those values are remarkably high and it is questionable whether decision makers at De Heus embody coefficients in that sphere.

Furthermore, insights from the risk-return analysis endorse the conclusion of the OLS method being the superior minimum-variance hedging strategy. While in some cases (e.g., the wheat replacement sample) time-varying minimum-variance hedging strategies feature solutions that have a higher average rate of return, those solutions are not in general more efficient (i.e., they also feature a higher volatility).

The unambiguous superiority faints, however, as we introduce the mean rate of return into the optimization problem. More specifically, the time-varying mean-variance hedging strategies have been shown to outperform the static mean-variance hedging strategy in 3 out of 6 cases. That is, they were able to achieve levels of return for a given standard deviation that a static hedge portfolio could not achieve, whatever value we choose for the static hedge ratio.

Why is it then, that the mean-variance strategies perform so well in some cases (soy replacement, soy Oil World, and wheat Oil World), yet perform so poorly in the other cases (corn replacement, corn Reuters, wheat replacement)? Looking at Figure 4.31, we see that this is likely due to the average spot and futures rates of return. The worse the average rate of return of the spot and futures in a hedge, the worse the outcome of the dynamic mean-variance hedging strategies relative to the static hedging strategies.

This suggests that incorporating mean optimization in a hedging strategy is beneficial only, if the averages exceed a certain threshold level.

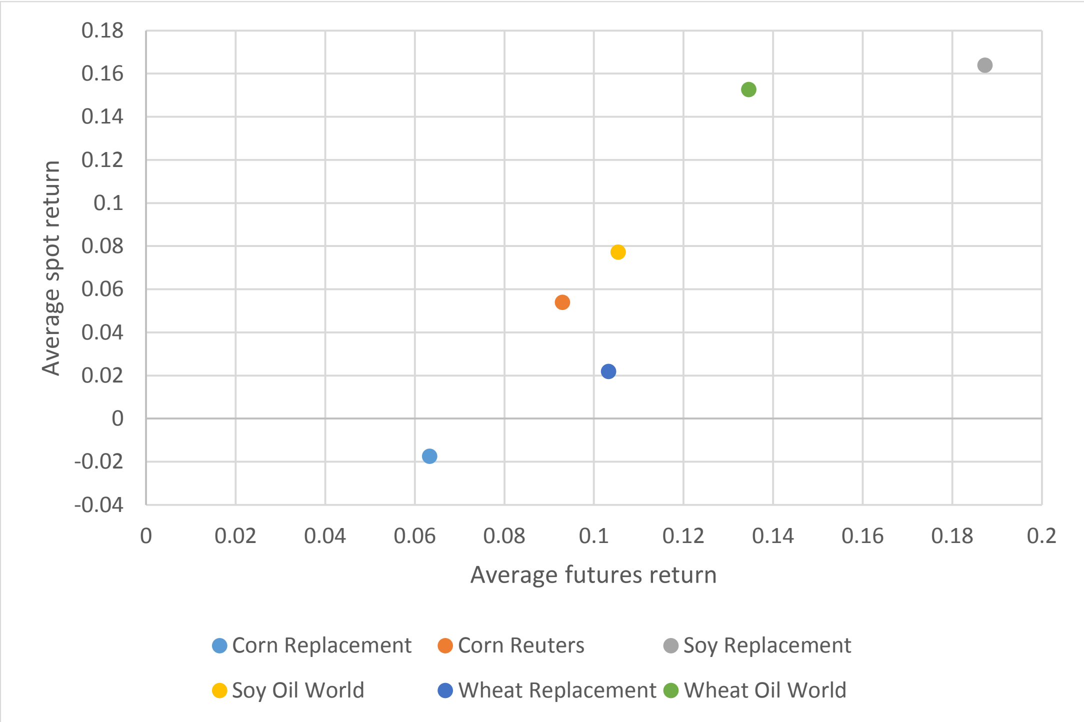

Figure 4.31 – *Spot and futures sample, annualized, average return per hedge portfolio.*

Chapter 5 – Conclusion

The purpose of this project was to facilitate the current risk management at De Heus by investigating opportunities for hedging strategies for the three commodities *corn*, *wheat*, and *soy*. The central research question was:

> "*How to design a hedging strategy at De Heus for the most important commodities with special attention to modeling the optimal hedge ratio?*"

To this end, we have set up structured databases including both cash market and commodity exchange securities and programmed a model using VBA Excel that, amongst other things, allows the user to visually investigate the dynamics of volatility and correlation changes over time for any security by plotting rolling averages of the corresponding series.

Understanding the dynamics of the volatility and correlation of the return series of the spot and futures variables is crucial in setting up a proper hedge. In fact, they constitute the sole trinity of factors that are used in order to set up a proper hedge. More specifically, it has been shown in the literature review that under the minimum variance object function, the optimal hedge ratio, i.e., the position in futures relative to the spot position is given as the coefficient of correlation times the spot volatility over the futures volatility.
While minimum variance hedging strategies offer the limit of what is possible in terms of risk reduction, they do so without concern for the return of the hedge portfolio. This is appropriate only if either the investor is highly risk averse, or if the expected return of the futures security is zero. Since both assumptions are found to be unrealistic, we have also calculated and studied mean-variance hedge ratios, which are similar to the minimum-variance hedge ratios, with the only alteration that they also include a subtrahend written as the expected futures return divided by the risk aversion coefficient times the variance of the futures.

In general, the choice of the futures products constitutes the first step in setting up a hedging strategy. A sixfold of spot series have been sought to be hedged, two per commodity type. Apart from the internal replacement values, external price series from Oil World and Reuters have been incorporated in our analysis as they featured bigger sample sizes. The considered spot series then are *0003 (GMO corn)* & *French Corn FOB Rhine* for *corn*, *0041 (Wheat boat)* & *Wheat, U.S., No.2, SRW, fob Gulf* for *wheat*, and *0300 (Soya Hipro ADM GMO)* & *Soya pellets 48% Brazil, CIF Rotterdam* for *soy*. An in-depth correlation study revealed that the returns of the corn spot commodities could best be mimicked by the EMAc2 corn contract of the Matif commodity exchange, while the soy spot exposures were found to best be mirrored by the SMc1 soybean meal contract from the CBOT. As for wheat, the replacement series was matched most appropriately by the BL2c1 Milling Wheat contract (Matif), while the US wheat spot exposure was best handled by the Wc1 CBOT contract.

With the optimal hedge ratio being a function of the correlation and volatility of the spot and futures return series, the central theme of this research was whether or not the hedge ratio should be allowed to vary over time. Undertaking this quest is motivated by the concept termed *heteroscedasticity*, which refers to the tendency of a financial securities' volatility to usually come in clusters.

As modeling the conditional covariance matrix for a portfolio of securities is of great interest in the risk management discipline, there are numerous econometric models available to approach this endeavor. Three conditional, time-varying, industry-standard models have been applied in our research and integrated into the VBA program, namely the EWMA, DCC, and the bivariate BEKK model. The effectiveness of those models has been compared against the time-invariant OLS regression method, where the optimal hedge ratio is simply obtained by using the unconditional sample estimates.

While the minimum-variance analysis has shown that the time-varying models were periodically successful at deploying conditional market information by adequately adapting the hedge ratio to changes in the volatility and correlation structure, the standard regression method was overall found to be more effective in terms of variance reduction. Given those dynamic shifts in the correlation and volatility structure, the static character of the OLS method causes the spot position to be occasionally overhedged or underhedged. On average, however, it outperformed the other hedging strategies.

In the mean-variance analysis, however, the superiority of the standard regression has vanished. Here, it has been shown that the OLS method was outperformed by time-varying models in 3 out of 6 cases (soy replacement, soy Oil World, wheat Oil World) where the realized solutions of the dynamic strategies had more efficient risk-return ratios than any of the OLS strategies, whatever the risk aversion coefficient. The other cases exhibited inferior spot and futures returns, which was found to inhibit dynamic mean-variance optimization.

However, a profound shortcoming of the time-varying models is their need to regularly rebalance the hedge portfolio in order to exploit their conditional benefits. This gives rise to transaction costs, which have not even been considered in this analysis.

Overall, the minimum variance hedging effectiveness for the replacement exposures was on the low end. The variance reduction of the weekly spot returns for corn, wheat, and soy was only within the range of 25% to 35%. Upon applying the analysis on the external Reuters and Oil World price series, however, the hedging effectiveness increased to values in the range of 45% and 65%.
It has moreover been shown that the hedging effectiveness increases as we widen the hedging horizon. In fact, having also conducted the external series analysis on the basis of monthly return intervals, the hedging effectiveness could further be improved to a range of 60% to 80%.

Chapter 6 – Recommendations

If you actually decide to set up a hedge – what then are the implications of our research that should be taken into account?

The very first step constitutes the clear definition of the hedging objectives. If the target is to reduce the variability of the spot return to a maximum degree, then we clearly advocates to use the minimum-variance OLS regression method in an attempt to find the optimal hedge ratio. Not only has the paper shown that the rather complicated methods add little benefit in terms of variance reduction, the OLS method is also a lot easier to estimate, executing only a few commands Excel.
If however, one wants to jointly optimize the average rate of return and the variance, then one should opt for mean-variance hedging strategies. Here, the preference for either a static or a dynamic model depends on the return rates of the spot and futures series. If they are sufficiently high, it seems that a dynamic model is more appropriate. As for the dataset underlying this study, this has particularly been the case for the soy commodity.

However, what is as important as the choice of the estimation model, is the choice of the dataset upon which the estimations are to be based. We have seen profound differences in the outcomes based on whether we chose to represent a local spot exposure with its corresponding internal replacement product or with the series provided by external data vendors.
Equally eminent is the question of the data window that is chosen to study correlation and volatility. What is correlation after all? There is no such thing as a true correlation when we talk about pairs of financial variables. The parameter can only be interpreted in the light of the data that have been used in order to come up with an estimate. Figure 4.12, which shows the corn Reuters and EMAc2 correlation, provides an excellent case in point. The correlation structure remarkably improved in the period after August 2006. A correlation estimated today based on the complete sample starting in May 2000 will be significantly lower (.67) than the one we receive if we only consider the period after August 2006 (.80). Likewise, the optimal hedge ratio is destined to change considerably as well. We thus have to be very careful in choosing a time window that we believe is representative for the future.

Another issue concerned with data is the frequency of the dataset. It has been argued that ideally, the interval used to calculate the returns ought to be the same as the length of the period for which you want to set up the hedge. The sensitivity analysis has shown that the input parameters, the optimal hedge ratio and the hedging effectiveness change considerably as we change the return interval. As the programmed model is capable of handling flexible return intervals, this is something that should be taken into account when studying opportunities for a hedge.

Finally, an important implication for deriving the optimal hedge ratio that has been touched upon is the need to study not just the correlation between the variables in the hedge, but also their respective volatilities. We have observed a number of instances where the conditional hedge effectiveness decreased not because the correlation changed, but since the spot volatility exploded temporarily. If, in this instance, the futures volatility remains unchanged, or even decreases, then we do of course need more futures contracts to match the spot return swings. An important advice that follows from our thesis is therefore to use the custom model in order to study the past behavior of the correlation and volatilities to get an impression of how they might develop over the time for which we want to construct the hedge.

Having derived an assumed optimal hedge ratio, how ought it to be used in practice? The hedge ratio is simply a scalar that we use to multiply the proportion of the total value of the spot exposure relative to the total value per futures contract in order to calculate the optimal number of futures contracts. Suppose that we have $C_{s,t} = 40{,}000$ tons of corn in our books, which currently sell for $S_t =$ € 140/MT and of which we want to hedge the price fluctuations. Assume further that we decide to use a corn Matif contract, currently valued at F_t = € 136/MT, to hedge the exposure. It follows from Table 3.1 that each contract is concerned with a volume of $Q_f = 50$ tons. Suppose also that our analysis yields an optimal hedge ratio of $h^*_{t+1} = 0.67$. According to (1b) then, the optimal number of futures contracts, N_t, that should be shorted, is given as $N_t = h^*_{t+1} \frac{C_{s,t} S_t}{Q_f F_t} = 0.67 \frac{40000*140}{50*136} = 551.76$, or 552 if we round up. Note that in the context of (1b), $C_{f,t} = Q_f N_t$.

Notice that, as mentioned in Section 1.1, De Heus makes weekly estimates for replacement values of each of the upcoming 12 months. Thus, while our research has only focused on the most immediate spot month exposure, you may actually also have booked a delivery 4-5 months ahead.
If you want to hedge the full product, then you would of course have to hedge each forward month separately for which you have booked delivery. To do this, we simply run the complete analysis as pointed out in Chapter 4 for each month exposure on the basis of the corresponding replacement series. Depending on the resulting OHR, you then short a number of contracts related to the total value that is booked for that particular month.

Chapter 7 – Limitations & Future Research

Our research is based on an array of assumptions, and choices made in order to derive the results presented in the analysis. The results should not be accepted at face value but rather be interpreted in the light of the assumptions underlying them. A couple of those assumptions have already been mentioned in the course of our thesis.

First, the performance of time-varying models is dependent upon the way they are parameterized. If, for example, one had chosen a different in-sample interval, then model estimation would have yielded different parameters and therefore the optimal hedge ratio series and hedge effectiveness would have changed. While the minimum variance sensitivity analysis has shown that the OLS hedge performance is relatively robust with respect to changes in beta, such insights do not exist for the minimum variance DCC and BEKK model and demand further research.
It has been argued that the relative performance of the time-varying models rather shines in the out-of-sample period. In fact, some researchers in literature went so far as to reserve 80% of the sample space for out-of-sample performance testing and thus only use the first 20% for model specification. In this research, the proportion is rather the other way around. The base sample sizes have simply not been large enough to allow for such an allocation.

Second, it has further been argued that continuous model parameter updating would further benefit the usefulness of the DCC and BEKK models, whereas consideration for transaction costs due to rebalancing would deteriorate their performance.
The above mentioned points then imply that that our reported performances, especially in the out-of-sample context, and thus also the usefulness of the time-varying models in general, have to be interpreted with care and that further research is welcomed to get a more colored picture of the true performance difference, especially in the mean-variance optimization context.

Third, the models have been applied in their vanilla versions. That is, in case of the DCC and BEKK model, only the most recent ARCH and GARCH terms have been included for their autoregressive representation. While it is unlikely that higher terms would have led to significant performance improvements, we cannot reject this hypothesis empirically. The same holds for asymmetric GARCH models.
In general, we often made inference about the performance of time-varying models, while only having applied a threefold of the numerous models available in estimating the second moment of a distribution. The same is also true for time-invariant models. While we only used the commonly applied OLS method, there are other alternatives that could have been applied. The GLS (generalized least squares) method is a good example, as it is a linear regression model that can also cope with heteroscedasticity in the error terms. Further research would therefore be appropriate to test how other time-invariant and time-varying models would score.

Fourth, the mean value in the mean-variance hedging strategies has been modeled as a static sample average. One could, however, also allow the mean to vary over time, just as we allowed the covariance to vary over time. It would be interesting to see whether this might further improve performance of the mean-variance hedging strategies.

Fifth, futures contracts are marked to market on a daily basis, creating a stream of unpredictable cash flows during the course of the hedging horizon. This gives rise to what are termed multi-period hedging models. Like most models applied in literature, the models used in this research only consider a single-period setting.

Sixth, another problem concerned with multidimensionality is the fact that we have only considered bivariate hedges, involving only one spot exposure and one futures contract. If we theoretically were to revoke the diagonality restriction pertaining to the (ARCH and) GARCH parameter matrices in the DCC and BEKK model, then adding new securities to the hedge portfolio would introduce new layers of complexity as we would have to take into account and model the covariance of any resulting combination of variables. Thus, if we would set up simultaneous hedges for both corn and soy, then we would not only have to be cautious about the correlation between spot corn and futures corn, but also between spot corn & spot soy, and spot corn & futures soy. The same is true for the other variables. Further research is therefore necessary to investigate the nature and impact of multivariate hedge constructs.

Finally, it has been shown in the literature research that cointegration is an important concept when it comes to model specification. More specifically, numerous authors argue that if the spot and futures series in a hedge are cointegrated then the OLS and GARCH methods should incorporate error correction terms. Due to limitations in scope, cointegration has not been paid any attention to and therefore also presents fruitful potential for further research.

References

Alexander, C. (2011). *Market Models*. Chichester, West Sussex: John Wiley & Sons Ltd.

Alexander, C. & Barbosa, A. (2007). Effectiveness of minimum-variance hedging. *The Journal of Portfolio Management*, 33(2), 46-59.

Baillie, R. T. & Myers, R. J. (1991). Bivariate GARCH estimation of the optimal commodity futures hedge. *Journal of Applied Econometrics*, 6(2), 109-124.

Bauwens, L., Laurent, S. & Rombouts, J. V. K. (2006). Multivariate GARCH models: A survey. *Journal of Applied Econometrics*, 21, 79-103.

Bera, A. K., Garcia, P. & Roh, J. S. (1997). Estimation of time-varying hedge ratios for corn and soybeans: BGARCH and random coefficient approaches. Office for Futures and Options Research, University of Illinois. Retrieved on December 5, 2014, from http://econpapers.repec.org/paper/wpawuwpfi/9712007.htm

Benartzi, S. & Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110, 75-92.

Bjornson, B & Carter, C. A. (1997). New evidence on agricultural commodity return performance under time-varying risk. *American Journal of Agricultural Economics*, 79(3), 918-930.

Bos, C. & Gould, P. (2007). Dynamic correlations and optimal hedge ratios. Tinbergen Institute, Vrije Universiteit Amsterdam. Retrieved on January 4, 2015, from http://hdl.handle.net/10419 /86538

Brooks, C., Henry, O.T. & Persand, G. (2002). The effect of asymmetries on optimal hedge ratios. *Journal of Business*, 75(2), 333-352.

Bystrom, H.N.E. (2003). The hedging performance of electricity futures on the Nordic Power Exchange. *Applied Economics*, 1, 1-11.

Bollerslev, T. & Wooldridge, T. (1992). Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews*, 11, 143-172

Chang, C. L., McAleer, M. & Tansuchat, R. (2011). Crude oil hedging strategies using dynamic multivariate GARCH. *Energy Economics* 33, 912-923.

Cecchetti, S. G., Cumby, R. E., & Figlewski, S. (1988). Estimation of the optimal futures hedge. *The Review of Economics and Statistics*, 70, 623-630.

Chen, S. S., Lee, C. F., & Shrestha, K. (2003). Futures hedge ratios: a review. *The Quarterly Review of Economics and Finance*, 43, 433-465.

Chen, S. S., Lee, C. F., & Shrestha, K. (2008). Do the pure martingale and joint normality hypotheses hold for futures contract? Implications for the optimal hedge ratio. *The Quarterly Review of Economics and Finance*, 48, 153-174.

Cheung, C. S., Kwan, C. C. Y. & Yip, P. C. Y. (1990). The hedging effectiveness of Options and Futures: A Mean-Gini approach. *Journal of Futures Markets*, 10, 61-73.

Choudhry, T. (2003). Short-run deviations and optimal hedge ratio: Evidence from stock futures. *Journal of Multinational Financial Management*, 13, 171-192.

Choudhry, T. (2004). The hedging effectiveness of constant and time-varying hedge ratios using three

Pacific Basin stock futures. *International Review of Economics & Finance*, 13, 371-385.

Choudhry, T. (2009). Short-run deviations and time-varying hedge ratios: Evidence from agricultural futures markets. *International Review of Financial Analysis*, 18, 58-65

Ederington, L. H. (1979). The hedging performance of the new futures markets. *Journal of Finance*, 34, 157-170

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

Engle, R. F. & Kroner, K. F. (1995). Multivariate Simultaneous Generalized ARCH. *Econometric Theory* 11, 122-150.

Harris, R. D. F. & Shen, J. (2003). Robust estimation of the optimal hedge ratio. *The Journal of Futures Markets*, 23(8), 799-816.

Hatemi-J, A. & Roca, E. (2006). Calculating the optimal hedge ratio: Constant, time varying and the Kalman Filter approach. *Applied Economics Letters*, 13(5), 293-299.

Hsin, C. W., Kuo, J. & Lee, C. F. (1994). A new measure to compare the hedging effectiveness of foreign currency futures versus options. *The Journal of Futures Markets*, 14(6), 685 - 707

Hsu, C. C., Tseng, C. P. & Wang, Y. H. (2008). Dynamic hedging with futures: A copula-based GARCH model. *Journal of Futures Markets*, 28(11), 1095-1116.

Hull, J. C. (2012). *Options, futures, and other derivatives*. Harlow, Essex: Pearson Education Limited.

Kenourgios, D., Samitas, A. & Drosos, P. (2008). Hedge ratio estimation and hedging effectiveness: the case of the S&P 500 stock index futures contract. *International Journal of Risk Assessment and Management* 9, 121-134

Kostika, E. & Markellos, R. N. (2013). Optimal hedge ratio estimation and effectiveness using ARCD. *Journal of Forecasting*, 32, 41-50

Kroner, K. F. & Sultan, J. (1993). Time-varying distributions and dynamic hedging with foreign currency futures. *Journal of Financial and Quantitative Analysis*, 28(4), 535-551.

Ku, Y. H. H., Chen, H. C. & Chen, K. H. (2007). On the application of the dynamic conditional correlation model in estimating optimal time-varying hedge ratios. *Applied Economics Letters*, 14, 503-509

Lahiani, A. & Guesmi, K. (2014). Commodity price correlation and time varying hedge ratios. *IPAG Business School*. Retrieved on December 15, 2014, from http://www.ipag.fr/wp-content/uploads/recherche/WP/IPAG_WP_2014_142.pdf

Lence, S. H. (1995). On the optimal hedge under unbiased futures prices. *Economics Letters*, 47, 385-388

Lien, D. (2005). The use and abuse of the hedging effectiveness measure. *International Review* of Financial Analysis, 14, 277-282

Lien, D. (2006). Estimation bias of futures hedging performance: A note. *The Journal of Futures Markets*, 26(8), 835-841.

Lien, D. & Tse, Y. K. (2000). A note on the length effect of futures hedging. *Advances in Investment Analysis and Portfolio Management*, 7, 131-143.

Lien, D. & Tse, Y. K. (2002). Some recent developments in futures hedging. *Journal of Economic Surveys*, 16(3), 357-396.

Lien, D., Tse, Y. K. & Tsui, A. K. C. (2002). Evaluating the hedging performance of the constant-correlation GARCH model. *Applied Financial Economics*, 12, 791-798.

Lien, D. & Yang, L. (2008). Asymmetric effect of basis on dynamic futures hedging. Empirical evidence from commodity markets. *Journal of Banking & Finance*, 32, 187-198.

Park, S. Y. & Jei, S. Y. (2010). Estimation and hedging effectiveness of time-varying hedge ratio: Flexible bivariate GARCH approaches*. The Journal of Futures Markets*, 30(1), 71-99.

Poortema, K. (2011*). Statistische Technieken*. University of Twente.

Ramlall, I. (2009). Striving for the optimal hedge ratio. University of Mauritius. Retrieved on December 17, 2014, from http://ssrn.com/abstract=1490570

Rabobank Food & Agribusiness Research (2007). World Map Grains & Oilseeds. Utrecht.

Sephton, P. S. (1993). Optimal hedge ratios at the Winnipeg commodity exchange. The Canadian *Journal of Economics*, 26(1), 175-193.

Shalit, H. (1995). Mean-Gini hedging in futures markets. *Journal of Futures Markets*, 15, 617-635.

Shalit, H. & Yitzhaki, S. (1984). Mean-Gini, portfolio theory and the pricing of risky assets. *Journal of Finance*, 38, 1449-1468.

USDA (2015). Production, Supply and Distribution Online. Retrieved on December 20, 20014, from http://apps.fas.usda.gov/psdonline/psdHome.aspx

APPENDICES

APPENDIX A – Objective functions

Sharpe ratio objective function

Another objective function that also takes into consideration the average return of the hedged portfolio is the Sharpe ratio objective function, which is given by

$$F(R_h, R_F) = \frac{E[R_h] - R_F}{\sigma_h} \tag{21a}$$

where R_F is the risk free rate of return. Hence, this method maximizes the excess of the portfolio return over the risk free rate per unit of standard deviation of the portfolio. Setting the derivative of (21a) with respect to h equal to 0 and solving for h, we get the following OHR (see Chen *et al.*, 2003):

$$h^* = \frac{(\sigma_s/\sigma_f)[(\sigma_s/\sigma_f)(\mu_{R_f}/\mu_{R_S} - R_F)) - \rho_{Rs,Rf}]}{[1 - (\sigma_s/\sigma_f)(\mu_{R_f}\rho_{Rs,Rf}/(\mu_{R_S} - R_F))]} \tag{21b}$$

Note that, again, this method is similar to the variance minimization approach under the assumption that futures price series follows a martingale process (Chen *et al.*, 2003), i.e., if $\mu_{R_f} = 0$ then (21b) reduces to (2b).

While the Sharpe ratio approach succeeds in involving the expected portfolio rate of return without introducing the risk aversion parameter, it complicates the optimization procedure in that (21a) is a non-linear function of the hedge ratio (Ramlall, 2009). Moreover, like the mean-variance approach, it is consistent to the expected utility paradigm only if either the utility function is quadratic and/or the returns are jointly normally distributed (Chen *et al.*, 2003).

Minimum MEG objective function

The minimum MEG objective function makes use of the MEG (mean extended-Gini) coefficient – a measure of statistical dispersion that also takes account of the average.

A more pragmatic expression for the MEG coefficient has been put forward by Shalit & Yitzhaki (1984). Under this expression, the objective function is given by

$$F(R_h, v) = -vCov(R_h, (1 - G(R_h))^{v-1}) \tag{22}$$

where v is the risk aversion parameter, and G(.) is the cumulative distribution function of R_h. The OHR is then found by minimizing (22) with respect to h.

An advantage of the minimum MEG approach is that it is consistent with the expected utility paradigm. However, it is quite difficult to implement in practice (Lien & Tse, 2002). Moreover, having compared the effectiveness of futures and options as hedging vehicles, Cheung *et al.* (1990)

concluded that futures favor the application of the minimum variance objective function, while the MEG approach lends itself more to modeling the OHR in an options context.

If the futures and spot returns were jointly normally distributed, then the minimum MEG method would yield the same OHR as the minimum variance approach (Shalit, 1995).

Minimum GSV objective function

The final approach accounts for the well-established fact that investors in general are more concerned with the risk of a downside movement of an asset than with a potential upside movement (Benartzi & Thaler, 1995).

The generalized semi-variance approach integrates this notion of asymmetric risk perception by introducing two new parameters, δ and α, where the former stands for the target rate of return, and the latter represents the perceived damage from failing to reach that target rate of return (Lien & Tse, 2002). The objective function is given by

$$F(R_h, \delta, \alpha) = \int_{-\infty}^{\delta} (\delta - R_h)^{\alpha} dG(R_h), \qquad \alpha > 0 \tag{23}$$

as we can see from the boundaries of the integrals in expression (23), only returns lower than the target rate of return are taken into consideration when evaluating the objective function. The OHR is derived by minimizing (23) with respect to h.

The two paramount advantages of this approach are the capability to adjust for the asymmetric risk perception of managers and the consistency with the stochastic dominance concept. On the flipside, however, the approach requires us to know the risk aversion parameter, α, as well as the analytic expression for the distribution function of R_h (Chen *et al*., 2003).

As was the case for the minimum MEG objective function, if spot and futures returns are both jointly normally distributed, then the minimum GSV objective function yields an OHR that is similar to that under the minimum variance approach (Chen *et al*., 2003).

Note on the MEG and GSV objective functions

The objective functions of the minimum MEG and minimum GSV methods can easily be augmented to also incorporate the expected return of the hedged portfolio. In this case we would have

$$F(R_h) = E[R_h] - \Gamma(R_h)$$

where $\Gamma(R_h)$ is either (22) or (23). Unlike in the mean-variance approach we do not multiply $\Gamma(R_h)$ by a risk aversion parameter anymore, as (22) and (23) already include a risk aversion parameter themselves. For both approaches, the OHR is found by maximizing $F(R_h)$ with respect to h. The resulting OHRs are termed *optimum mean-MEG* and *optimum mean-generalized semi-variance*, respectively.

APPENDIX B - OHR and cointegration

The idea of cointegration builds on the concept of (non-) stationarity. A time series is said to be stationary (i.e., I(0) in the sense of the Engle & Granger, 1987) if the joint probability distribution of the stochastic process underlying the time series is constant throughout time. That is, for a stochastic process, X_t, a constant k, and a time frame of length m

$$P(X_{t1} \leq x_{t1}, X_{t2} \leq x_{t2}, \dots, X_{tm} \leq x_{tm}) = F(X_{t1}, X_{t2}, \dots, X_{tm}) = F(X_{t1+k}, X_{t2+k}, \dots, X_{tm+k}) \qquad (24)$$

It follows from this expression that distribution parameters such as the mean or the variance do not change as we wander along the time line. According to Engle & Granger (1987) if X_t is stationary and has zero mean, then the variance of X_t is finite and innovations in the series only have temporary effects. The series thus never drifts too far away from its equilibrium.

If, however, the time series contains a unit root (i.e., I(1)) then the time series is said to be nonstationary and (24) does not hold anymore. Moreover, the variance of X_t goes to infinity as t tends to infinity and innovations have a permanent effect on the time series (Engle & Granger, 1987). Reversion to historic values thus happens much less frequently, if ever.

The concept of cointegration now states that if there are two time series that are I(1) and there exists a liner combination of those two time series such that the linear combination is I(0), i.e., stationary, then the two series are said to be cointegrated. If this were the case then, as mentioned in Section 2.1.1 the OLS regression in (4) would be miss-specified.

The presence of cointegration in spot and futures time series is, however, a very controversial issue in literature with divergent contributions on a theoretical (Brooks *et al*., 2002), as well as on an empirical level (Lien & Wilson, 2001). The latter is true even within the realm of agricultural commodities. Baillie & Myers (1991) for example found no cointegration among the spot and futures series for each of their tested commodities, amongst which corn and soybeans. The opposite is true in case of (Choudhry, 2009) who included corn, wheat, and soybeans in the list of tested commodities.

The controversies also extend to the question in how far the omission of the cointegration relationship affects the hedging performance due to diverging estimates of the OHR parameter.
According to Ghosh (1995), Ghosh & Clayton (1996), and Kroner & Sultan (1993) the hedging performance suffers considerably if cointegration is left out of consideration (Choudhry, 2009).
Lien (2004) however states that the loss in hedging effectiveness is likely to be minimal.

APPENDIX C – Meta-analysis of empirical studies: OLS vs. time-varying models[13]

Article	Commodity	Tested models	Outcome
Baillie & Myers (1991)	Agricultural commodities: Beef, coffee, corn, cotton, gold, soybeans	Diagonal Vech GARCH, OLS	In & out: GARCH>OLS
Kroner & Sultan (1993)	FX: British Pound, German Mark, Japanese Yen, Swiss Franc	GARCH-X, OLS, naive	In & out: GARCH-X>OLS>naive
Sephton (1993)	Agricultural commodities: Feed wheat, feed barley, canola	Vech GARCH, CCC, OLS	Complete sample: GARCH>CCC>OLS
Bera *et al.* (1997)	Agricultural commodities: Corn and soybean	DVECH GARCH, BEKK, CCC, OLS	In & out: GARCH>OLS>CCC >BEKK
Lien & Wilson (2001)	Crude Oil	SV, EGARCH, OLS, naive	In & out: OLS>naive>EGARCH>SV
Lien *et al.* (2002)	Mixed: FX(BP, DM, JY), agri (SBO, wheat, crude oil, corn, cotton), and equity indices (NYSE composite, S&P 500)	CCC, OLS	Out: OLS>CCC for each considered market
Bystrom (2003)	Electricity spot and futures	CCC, OGARCH, OLS	Out (uncond.): Naïve>OLS>CCC>OGARCH
Choudhry (2003)	Equity indices from Australia, Germany, Hong Kong, Japan, South Africa & the UK	GARCH-X, GARCH, OLS, naive	Complete & out: GARCH>GARCH-X>OLS>naive
Choudhry (2004)	Equity indices from Australia, Hong Kong & Japan	GARCH (DVECH), OLS, naive	Out 1: OLS>GARCH>naive Out 2: GARCH>OLS>naive
Chen *et al.* (2006)	Mixed: S&P500, swiss franc, gold, coffee, corn, wheat, cotton, sugar, soybean oil, soybean	DCC, CCC, OLS, naive	Out: DCC>CCC>OLS>naive
Alexander & Barbosa (2007)	Equity indices: Nasdaq 100, S&P 500, FTSE 100, CAC 40, Hang Seng Composite, Kospi 200, Ibovespa	OLS, ECM, EWMA, GARCH	Out 1 & out 2: OLS>ECM>EWMA>GARCH. Difference in HE highly marginal, though
Bos & Gould (2007)	S&P 500, FTSE 100	OLS, BEKK, DCC	In: BEKK>OLS>DCC Out: OLS>DCC>BEKK
Ku *et al.* (2007)	FX: British pounds and Japanese yen	DCC, CCC, ECM, OLS	In & out: DCC>OLS>ECM>CCC

[13] A few of the articles also test models not included in this list. The overview, however, is mainly restricted to the models that are also discussed in our thesis, i.e., VECH, BEKK, CCC, DCC, EWMA, OLS and the naïve hedge.
"In" and "out" refer to in-sample and out-of-sample performance, respectively.
Usually there are differences in performance with respect to the different tested commodities. The ranking given under "outcome" refers to a joint impression, when considering all commodities together.

Lien (2007)	Mixed: 24 commodities – FX, equity indices as well as agricultural commodities	ECM, OLS, naive hedge	In: OLS>ECM>naive Out: ECM>OLS>naive
Hsu *et al*. (2008)	Equity indices: S&P 500, FTSE 100, MSCI-SWI	OLS, CCC, DCC	In: DCC>OLS>CCC Out: DCC>CCC>OLS
Kenourgios (2008)	S&P 500	OLS, ECM, GARCH, EGARCH	In: ECM>OLS>EGARCH>GARCH
Lien & Yang (2008)	Commodities (corn, soybeans, cotton, coffee, pork belly, lean hog, heating oil, crude oil, copper, silver)	GARCH, OLS, naive	In: GARCH>OLS>naive Out: GARCH>OLS>naive
Lai *et al*. (2009)	Equity indices from Hong Kong, Japan, Korea, Singapore and Taiwan	OLS, DCC, naive	In & out: DCC>OLS (in 3/5 cases)>naive
Park & Jei (2010)	Agricultural: Corn & soybeans	CCC, DCC, OLS, naive	In (corn): OLS>CCC>DCC>naive In (soybeans): CCC>DCC>OLS>naive Out (both): OLS>CCC>DCC>naive
Kostika & Markellos (2013)	Equity indices: DJI, FTSE, and DAX	CCC, BEKK (diagonal), EWMA (.94), ECM, OLS, naive	In: CCC>OLS>ECM>BEKK>EWMA>naive Out: CCC>ECM>OLS>BEKK>EWMA>naive

APPENDIX D – Constant Conditional Correlations

As mentioned in section 2.2.7, the CCC model can be written as:

$$H_t = D_t R D_t, \qquad \text{where } D_t = diag\{\sqrt{h_{i,t}}\} \tag{25a}$$

Assuming zero mean of the returns, the variance is simply the expected value of the squared returns, i.e., $h_{i,t} = E_{t-1}(r_{i,t}^2)$. Defining the returns as a product of the conditional standard deviation times the standardized disturbance

$$r_{i,t} = \sqrt{h_{i,t}}\varepsilon_{i,t} \tag{25b}$$

we can show that even though R is constant, it contains the conditional correlations. Note that the conditional correlation between $r_{1,t}$ and $r_{2,t}$ is given by:

$$\rho_{12,t} = \frac{E_{t-1}[r_{1,t}r_{2,t}]}{\sqrt{E_{t-1}[r_{1,t}^2]E_{t-1}[r_{2,t}^2]}} \tag{25c}$$

Substituting (25b) into (25c) and simplifying algebraically we get

$$\rho_{12,t} = \frac{E_{t-1}[r_{1,t}r_{2,t}]}{\sqrt{E_{t-1}[r_{1,t}^2]E_{t-1}[r_{2,t}^2]}} = \frac{E_{t-1}[\sqrt{h_{1,t}}\varepsilon_{1,t}\sqrt{h_{2,t}}\varepsilon_{2,t}]}{\sqrt{E_{t-1}[(\sqrt{h_{1,t}}\varepsilon_{1,t})^2]E_{t-1}[(\sqrt{h_{2,t}}\varepsilon_{2,t})^2]}} =$$

$$\frac{\sqrt{h_{1,t}}*\sqrt{h_{2,t}}*E_{t-1}[\varepsilon_{1,t}\varepsilon_{2,t}]}{\sqrt{h_{1,t}}*\sqrt{h_{2,t}}*\sqrt{E_{t-1}[(\varepsilon_{1,t})^2]E_{t-1}[(\varepsilon_{2,t})^2]}} = \frac{E_{t-1}[\varepsilon_{1,t}\varepsilon_{2,t}]}{\sqrt{E_{t-1}[\varepsilon_{1,t}^2]E_{t-1}[\varepsilon_{2,t}^2]}} = E_{t-1}[\varepsilon_{1.t}\varepsilon_{2,t}] \tag{25d}$$

which shows that the conditional correlation between $r_{1,t}$ and $r_{2,t}$ is the same as the conditional covariance of the standardized disturbances. Moreover, it follows from (25b) that

$$\varepsilon_{i,t} = \frac{r_{i,t}}{\sqrt{h_{i,t}}} \xrightarrow[notation]{matrix} \varepsilon_t = D_t^{-1} r_t = \begin{bmatrix} \frac{1}{\sigma_{1,t}} & 0 \\ 0 & \frac{1}{\sigma_{2,t}} \end{bmatrix} \begin{bmatrix} r_{1,t} \\ r_{2,t} \end{bmatrix} = \begin{bmatrix} \frac{r_{1,t}}{\sigma_{1,t}} \\ \frac{r_{2,t}}{\sigma_{2,t}} \end{bmatrix} \tag{25e}$$

Writing the conditional correlations in (25d) in matrix form and plugging (25e) into the matrix form of (25d) we get that

$$E_{t-1}[\varepsilon_t\varepsilon_t'] = E_{t-1}\left[\begin{bmatrix} \frac{r_{1,t}}{\sigma_{1,t}} \\ \frac{r_{2,t}}{\sigma_{2,t}} \end{bmatrix}\begin{bmatrix} \frac{r_{1,t}}{\sigma_{1,t}} & \frac{r_{2,t}}{\sigma_{2,t}} \end{bmatrix}\right] = E_{t-1}\begin{bmatrix} \frac{r_{1,t}^2}{\sigma_{1,t}^2} & \frac{r_{1,t}r_{2,t}}{\sigma_{1,t}\sigma_{2,t}} \\ \frac{r_{2,t}r_{1,t}}{\sigma_{2,t}\sigma_{1,t}} & \frac{r_{2,t}^2}{\sigma_{2,t}^2} \end{bmatrix} = \begin{bmatrix} 1 & \frac{\sigma_{12,t}}{\sigma_{1,t}\sigma_{2,t}} \\ \frac{\sigma_{21,t}}{\sigma_{2,t}\sigma_{1,t}} & 1 \end{bmatrix} \tag{25f}$$

Likewise, we get the same result as in (25f) when we combine the diagonal and the variance matrix:

$$D_t^{-1}H_tD_t^{-1} = \begin{bmatrix} \frac{1}{\sigma_{1,t}} & 0 \\ 0 & \frac{1}{\sigma_{2,t}} \end{bmatrix}\begin{bmatrix} \sigma_{1,t}^2 & \sigma_{12,t} \\ \sigma_{21,t} & \sigma_{2,t}^2 \end{bmatrix}\begin{bmatrix} \frac{1}{\sigma_{1,t}} & 0 \\ 0 & \frac{1}{\sigma_{2,t}} \end{bmatrix} = \begin{bmatrix} 1 & \frac{\sigma_{12,t}}{\sigma_{1,t}\sigma_{2,t}} \\ \frac{\sigma_{21,t}}{\sigma_{2,t}\sigma_{1,t}} & 1 \end{bmatrix} \tag{25g}$$

Finally, it follows from (25a) that $D_t^{-1} H_t D_t^{-1} = R$. But if R = (25g) = (25f) = (25c) then this means that R indeed contains the *conditional correlations*. While the term "*constant conditional*" seems as an oxymoron at first, expression (25a) sheds light into what is meant by that. The constant element of the model resides in the correlation matrix, R, which is time-invariant. The covariance matrix is, however, allowed to vary over time by modeling the individual variance elements with univariate GARCH processes, thus providing for the conditional part of the expression in (25a).

APPENDIX E - Custom program interface

APPENDIX F - Histograms and descriptive statistics

The following figures and tables show the histograms and descriptive statistics of the various weekly spot and futures series, with exception of the EMAc2 Corn Matif series. The reported p-values refer to the Jarque-Bera test with the null-hypothesis that the data is normally distributed.

Return series Milling Wheat Matif, continuation BL2c1:

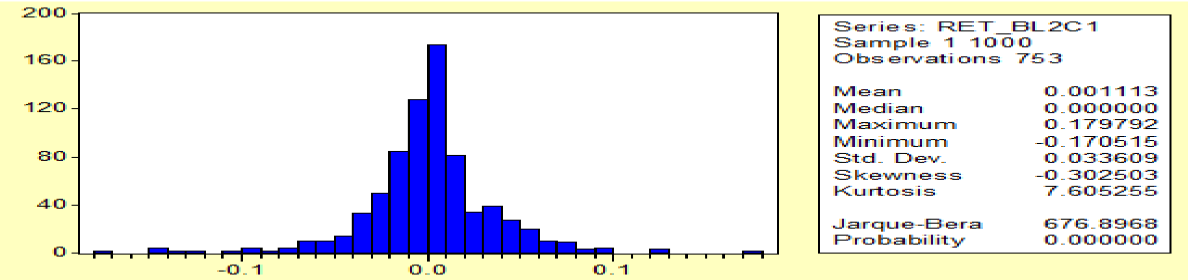

Return series Soybean Meal CBOT, continuation SMc1:

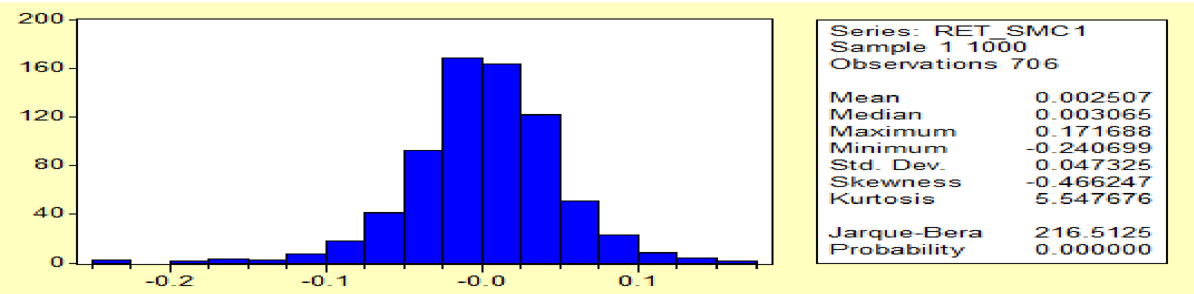

Return series Wheat CBOT, continuation Wc1:

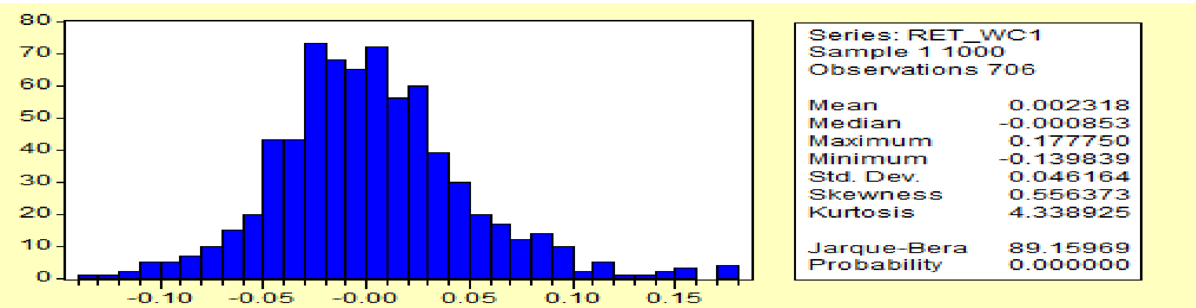

Return series Corn Replacement (0003), spot continuation:

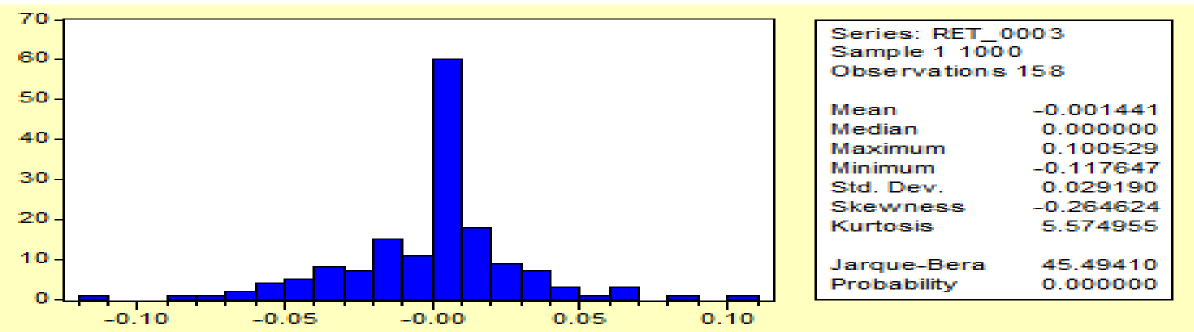

Return series Wheat Replacement (0041), spot continuation:

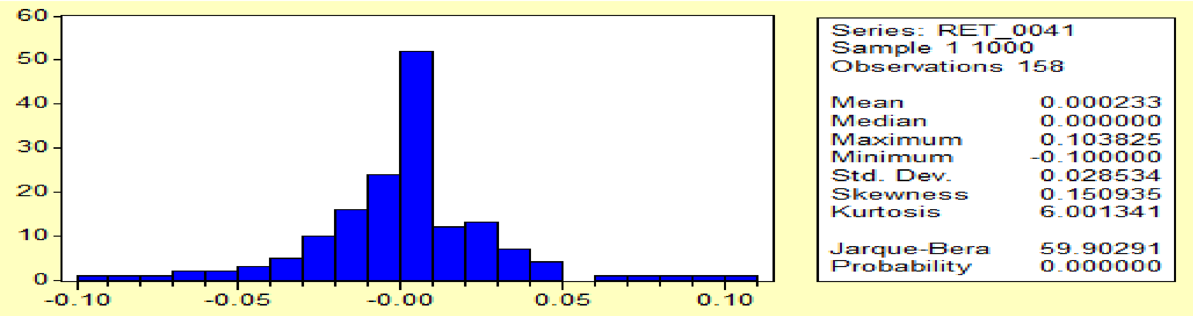

Return series Soy Replacement (0300), spot continuation:

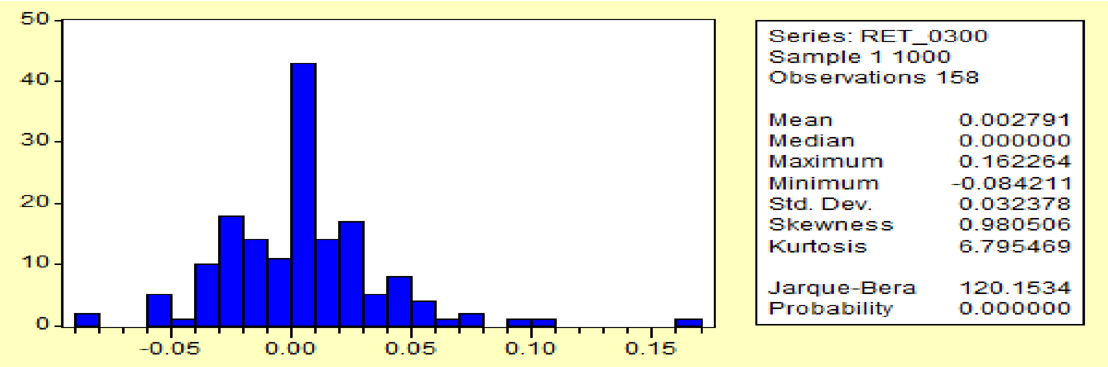

Return series Corn Reuters, spot continuation:

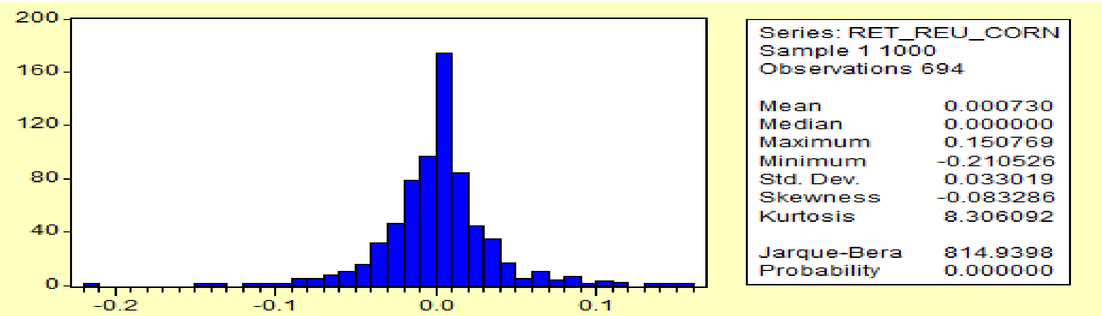

Return series soy Oil World, spot continuation:

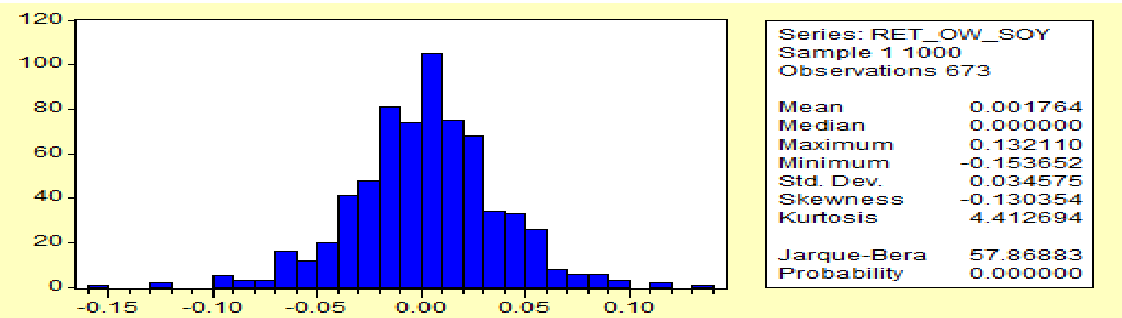

Return series wheat Oil World, spot continuation:

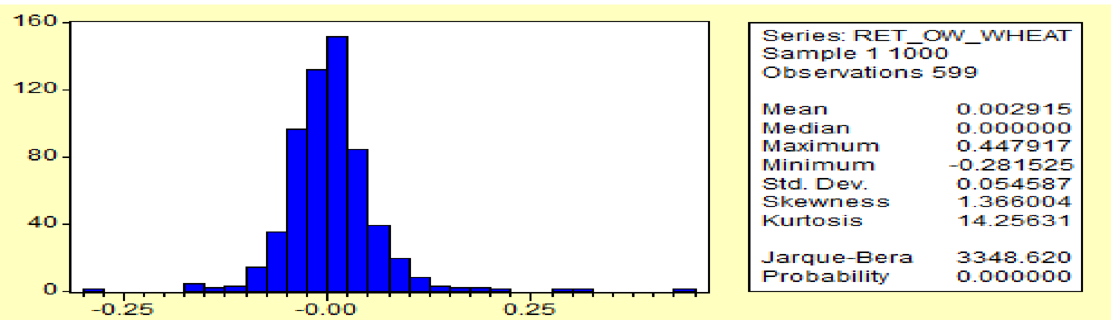

APPENDIX G - Correlograms: Autocorrelation in squared returns

The following tables show the correlograms of the various weekly, squared spot and futures series, with exception of the EMAc2 Corn Matif series. The reported p-values at each lag k refer to the Ljung Box Q test, which tests the null-hypothesis that the data features no autocorrelation at lag k.

Squared return series Milling Wheat Matif, continuation BL2c1:

Correlogram of RET2_BL2C1

Date: 01/24/15 Time: 16:14
Sample: 1 1000
Included observations: 753

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.163	0.163	20.038	0.000
		2	0.029	0.003	20.685	0.000
		3	0.125	0.123	32.486	0.000
		4	0.142	0.106	47.703	0.000
		5	0.089	0.052	53.749	0.000

Squared return series Soybean Meal CBOT, continuation SMc1:

Correlogram of RET2_SMC1

Date: 01/24/15 Time: 16:18
Sample: 1 1000
Included observations: 706

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.188	0.188	25.086	0.000
		2	0.096	0.063	31.697	0.000
		3	0.094	0.068	37.966	0.000
		4	0.125	0.096	49.121	0.000
		5	0.104	0.059	56.866	0.000

Squared return series Wheat CBOT, continuation Wc1:

Correlogram of RET2_WC1

Date: 01/24/15 Time: 16:19
Sample: 1 1000
Included observations: 706

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.056	0.056	2.2041	0.138
		2	0.004	0.001	2.2142	0.331
		3	0.139	0.140	16.015	0.001
		4	0.091	0.078	21.953	0.000
		5	0.011	0.003	22.044	0.001

Squared return series Corn Replacement, continuation spot:

Correlogram of RET2_0003

Date: 01/24/15 Time: 16:20
Sample: 1 1000
Included observations: 158

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.212	0.212	7.2182	0.007
		2	0.312	0.280	22.980	0.000
		3	0.260	0.174	33.967	0.000
		4	0.146	0.004	37.465	0.000
		5	0.242	0.125	47.125	0.000

Squared return series Wheat Replacement, continuation spot:

Correlogram of RET2_0041

Date: 01/24/15 Time: 16:21
Sample: 1 1000
Included observations: 158

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.253	0.253	10.312	0.001
		2	-0.001	-0.069	10.312	0.006
		3	-0.096	-0.084	11.822	0.008
		4	0.051	0.106	12.257	0.016
		5	0.036	-0.010	12.475	0.029

Squared return series Soy Replacement, continuation spot:

Correlogram of RET2_0300

Date: 01/24/15 Time: 16:22
Sample: 1 1000
Included observations: 158

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	-0.034	-0.034	0.1879	0.665
		2	-0.076	-0.078	1.1306	0.568
		3	-0.056	-0.061	1.6330	0.652
		4	0.007	-0.004	1.6407	0.801
		5	-0.056	-0.066	2.1628	0.826

Squared return series Corn Reuters, continuation spot:

Correlogram of RET2_REU_CORN

Date: 01/24/15 Time: 16:23
Sample: 1 1000
Included observations: 694

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.032	0.032	0.7164	0.397
		2	0.084	0.083	5.6366	0.060
		3	0.012	0.007	5.7339	0.125
		4	0.039	0.031	6.7800	0.148
		5	-0.009	-0.013	6.8423	0.233

Squared return series Soy Oil World (USD), continuation spot:

Correlogram of RET2_OW_SOY

Date: 01/24/15 Time: 16:24
Sample: 1 1000
Included observations: 673

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.059	0.059	2.3550	0.125
		2	0.087	0.084	7.4655	0.024
		3	0.070	0.061	10.818	0.013
		4	0.128	0.115	21.861	0.000
		5	0.054	0.033	23.874	0.000

Squared return series Soy Oil World (EUR), continuation spot:

Correlogram of RET2_OW_SOY_EUR

Date: 01/26/15 Time: 11:35
Sample: 1 1000
Included observations: 670

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.101	0.101	6.8091	0.009
		2	0.038	0.028	7.7593	0.021
		3	0.078	0.072	11.826	0.008
		4	0.085	0.071	16.741	0.002
		5	0.019	0.001	16.995	0.005

Squared return series Wheat Oil World, continuation spot:

Correlogram of RET2_OW_WHEAT

Date: 01/24/15 Time: 16:25
Sample: 1 1000
Included observations: 599

Autocorrelation	Partial Correlation		AC	PAC	Q-Stat	Prob
		1	0.148	0.148	13.218	0.000
		2	0.097	0.076	18.850	0.000
		3	0.047	0.024	20.207	0.000
		4	0.048	0.032	21.627	0.000
		5	0.049	0.034	23.098	0.000

APPENDIX H - EViews Wheat Parameterization Output

The following tables show the EViews (OLS, univariate/bivariate GARCH) estimation output for the various hedge setups in the wheat panel, based on weekly returns. The reported p-values refer to the null-hypothesis that the coefficient in question is zero (i.e., redundant in the regression equation).

OLS: Wheat Replacement (0041), continuation spot &
Milling Wheat Matif, continuation BL2c1:

Dependent Variable: RET_REP_WHEAT
Method: Least Squares
Date: 03/21/15 Time: 15:35
Sample: 1 130
Included observations: 130
RET_REP_WHEAT = C(1) + C(2)*RET_BL2C1_REP

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	-0.000385	0.001931	-0.199296	0.8423
C(2)	0.639467	0.080757	7.918385	0.0000

OLS: Wheat Oil World, continuation spot & Wheat CBOT, continuation Wc1:

Dependent Variable: RET_OW_WHEAT
Method: Least Squares
Date: 03/21/15 Time: 15:37
Sample: 1 500
Included observations: 500
RET_OW_WHEAT = C(1) + C(2) * RET_WC1_OW

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.000677	0.001579	0.428529	0.6685
C(2)	0.842738	0.028953	29.10735	0.0000

Univariate GARCH (1,1): Wheat Replacement (0041), continuation spot:

Dependent Variable: RET_REP_WHEAT
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:39
Sample: 1 130
Included observations: 130
Convergence achieved after 16 iterations
Bollerslev-Wooldridge robust standard errors & covariance
Presample variance: backcast (parameter = 0.7)
GARCH = C(2)*RESID(-1)^2 + (1 - C(2))*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	-0.000139	0.002059	-0.067547	0.9461
Variance Equation				
RESID(-1)^2	0.045586	0.022941	1.987096	0.0469
GARCH(-1)	0.954414	0.022941	41.60287	0.0000

Univariate GARCH (1,1): Wheat Oil World, continuation spot:

Dependent Variable: RET_OW_WHEAT
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:41
Sample: 1 400
Included observations: 400
Convergence achieved after 17 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.002757	0.002104	1.310682	0.1900
Variance Equation				
C	3.24E-05	2.55E-05	1.274341	0.2025
RESID(-1)^2	0.069250	0.034429	2.011391	0.0443
GARCH(-1)	0.925548	0.039596	23.37478	0.0000

Univariate GARCH (1,1): Milling Wheat Matif, continuation BL2c1 (replacement sample):

Dependent Variable: RET_BL2C1_REP
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:43
Sample: 1 130
Included observations: 130
Convergence achieved after 9 iterations
Bollerslev-Wooldridge robust standard errors & covariance
Presample variance: backcast (parameter = 0.7)
GARCH = C(2)*RESID(-1)^2 + (1 - C(2))*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.000258	0.001925	0.133985	0.8934
Variance Equation				
RESID(-1)^2	0.085688	0.027989	3.061506	0.0022
GARCH(-1)	0.914312	0.027989	32.66697	0.0000

Univariate GARCH (1,1): Wheat CBOT, continuation Wc1 (Oil World sample):

Dependent Variable: RET_WC1_OW
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:44
Sample: 1 400
Included observations: 400
Convergence achieved after 14 iterations
Bollerslev-Wooldridge robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.003022	0.002270	1.331327	0.1831
Variance Equation				
C	3.85E-05	2.53E-05	1.519257	0.1287
RESID(-1)^2	0.022427	0.015054	1.489777	0.1363
GARCH(-1)	0.964046	0.021935	43.94986	0.0000

B-GARCH (Diagonal BEKK): Wheat Replacement (0041), continuation spot & Milling Wheat Matif, continuation BL2c1:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a rank one matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	7.21E-05	4.46E-05	1.618779	0.1055
M(1,2)	0.000115	6.06E-05	1.897386	0.0578
M(2,2)	0.000183	9.60E-05	1.906844	0.0565
A1(1,1)	0.023135	0.197587	0.117088	0.9068
A1(2,2)	0.380860	0.138645	2.747007	0.0060
B1(1,1)	0.940794	0.032022	29.38006	0.0000
B1(2,2)	0.702232	0.152974	4.590535	0.0000

B-GARCH (Diagonal BEKK): Wheat Oil World, continuation spot & Wheat CBOT, continuation Wc1:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a full rank matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	0.000223	6.51E-05	3.425528	0.0006
M(1,2)	0.000194	6.44E-05	3.010250	0.0026
M(2,2)	0.000168	5.89E-05	2.859649	0.0042
A1(1,1)	0.315406	0.057788	5.458009	0.0000
A1(2,2)	0.103625	0.040828	2.538091	0.0111
B1(1,1)	0.905451	0.028020	32.31490	0.0000
B1(2,2)	0.958960	0.011364	84.38788	0.0000

APPENDIX I – EViews Soy Parameterization Output

The following tables show the EViews (OLS, univariate/bivariate GARCH) estimation output for the various hedge setups in the soy panel, based on weekly returns. The reported p-values refer to the null-hypothesis that the coefficient in question is zero (i.e., redundant in the regression equation).

OLS: Soy Replacement (0300), continuation spot & Soybean Meal CBOT, continuation SMc1:

Dependent Variable: RET_REP_SOY
Method: Least Squares
Date: 03/21/15   Time: 15:48
Sample: 1 130
Included observations: 130
RET_REP_SOY = C(1) + C(2) * RET_SMC1_REP

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.002364	0.002431	0.972204	0.3328
C(2)	0.360171	0.052188	6.901360	0.0000

OLS: Soy Oil World, spot continutation & Soybean Meal CBOT, continuation SMc1:

Dependent Variable: RET_OW_SOY
Method: Least Squares
Date: 03/21/15   Time: 15:50
Sample: 1 500
Included observations: 500
RET_OW_SOY = C(1) + C(2) * RET_SMC1_OW

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.000260	0.001230	0.211372	0.8327
C(2)	0.535843	0.027682	19.35684	0.0000

Univariate GARCH (1,1): Soy Replacement (0300), continuation spot:

Dependent Variable: RET_REP_SOY
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15   Time: 15:53
Sample: 1 130
Included observations: 130
Convergence achieved after 12 iterations
Bollerslev-Wooldridge robust standard errors & covariance
Presample variance: backcast (parameter = 0.7)
GARCH = C(2)*RESID(-1)^2 + (1 - C(2))*GARCH(-1)

Variable	Coefficient	Std. Error	z-Statistic	Prob.
C	0.002206	0.002575	0.856687	0.3916
		Variance Equation		
RESID(-1)^2	0.050926	0.024607	2.069625	0.0385
GARCH(-1)	0.949074	0.024607	38.56994	0.0000

Univariate GARCH (1,1): Soy Oil World, spot continutation:

Dependent Variable: RET_OW_SOY
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15   Time: 15:54
Sample: 1 500
Included observations: 500
Convergence achieved after 11 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.001298	0.001542	0.841585	0.4000
		Variance Equation		
C	4.97E-05	2.72E-05	1.828872	0.0674
RESID(-1)^2	0.099218	0.031487	3.151098	0.0016
GARCH(-1)	0.865535	0.040147	21.55904	0.0000

Univariate GARCH (1,1): Soybean Meal CBOT, continuation SMc1 (replacement sample):

Dependent Variable: RET_SMC1_REP
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:56
Sample: 1 130
Included observations: 130
Convergence achieved after 32 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.004543	0.003831	1.185852	0.2357
Variance Equation				
C	0.000548	0.000459	1.193656	0.2326
RESID(-1)^2	0.246582	0.262189	0.940475	0.3470
GARCH(-1)	0.481810	0.379824	1.268508	0.2046

Univariate GARCH: Soybean Meal CBOT, continuation SMc1 (Oil World sample):

Dependent Variable: RET_SMC1_OW
Method: ML - ARCH (Marquardt) - Normal distribution
Date: 03/21/15 Time: 15:57
Sample: 1 500
Included observations: 500
Convergence achieved after 17 iterations
Bollerslev-Wooldrige robust standard errors & covariance
Variance backcast: ON
GARCH = C(2) + C(3)*RESID(-1)^2 + C(4)*GARCH(-1)

	Coefficient	Std. Error	z-Statistic	Prob.
C	0.002440	0.001744	1.399084	0.1618
Variance Equation				
C	0.000323	0.000144	2.245155	0.0248
RESID(-1)^2	0.203071	0.061389	3.307968	0.0009
GARCH(-1)	0.640135	0.105187	6.085711	0.0000

B-GARCH (Diagonal BEKK): Soy Replacement (0300), continuation spot
& Soybean Meal CBOT, continuation SMc1:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a rank one matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients				
	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	7.77E-05	4.35E-05	1.784999	0.0743
M(1,2)	0.000234	0.000114	2.044242	0.0409
M(2,2)	0.000703	0.000353	1.991982	0.0464
A1(1,1)	0.108238	0.083051	1.303268	0.1925
A1(2,2)	0.455762	0.171286	2.660830	0.0078
B1(1,1)	0.945226	0.028955	32.64490	0.0000
B1(2,2)	0.660496	0.175550	3.762432	0.0002

B-GARCH (Diagonal BEKK): Soy Oil World, spot continutation
& Soybean Meal CBOT, continuation SMc1:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a full rank matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients				
	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	3.32E-05	1.38E-05	2.404622	0.0162
M(1,2)	6.17E-05	2.06E-05	2.995056	0.0027
M(2,2)	0.000154	5.05E-05	3.046011	0.0023
A1(1,1)	0.253757	0.035348	7.178866	0.0000
A1(2,2)	0.351549	0.049114	7.157824	0.0000
B1(1,1)	0.955032	0.011263	84.79439	0.0000
B1(2,2)	0.894837	0.022998	38.90894	0.0000

APPENDIX J – EViews OLS Regression Monthly returns

The following tables show the EViews OLS estimation output for the three hedge setups that include external spot series, based on monthly returns. The reported p-values refer to the null-hypothesis that beta is zero (i.e., redundant in the regression equation).

OLS: Corn Reuters, spot continuation & Corn Matif, continuation EMAc2:

Dependent Variable: RET_REUT_CORN
Method: Least Squares
Date: 03/21/15 Time: 16:01
Sample: 1 150
Included observations: 150
RET_REUT_CORN = C(1) + C(2) * RET_EMAC2_REUT

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.001458	0.003680	0.396219	0.6925
C(2)	0.841471	0.057210	14.70836	0.0000

OLS: Wheat Oil World, spot continuation & Wheat CBOT, continuation Wc1:

Dependent Variable: RET_OW_WHEAT
Method: Least Squares
Date: 03/21/15 Time: 16:03
Sample: 1 150
Included observations: 150
RET_OW_WHEAT = C(1) + C(2) * RET_WC1_OW

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.001003	0.002996	0.334748	0.7383
C(2)	0.896901	0.031361	28.59882	0.0000

OLS: Soy Oil World, spot continuation & Soybean Meal CBOT, continuation SMc1:

Dependent Variable: RET_OW_SOY
Method: Least Squares
Date: 03/21/15 Time: 16:05
Sample: 1 150
Included observations: 150
RET_OW_SOY = C(1) + C(2) * RET_SMC1_OW

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	0.000908	0.003643	0.249280	0.8035
C(2)	0.658917	0.042339	15.56279	0.0000

APPENDIX K – EViews BEKK estimation Monthly returns

The following tables show the EViews Diagonal BEKK B-GARCH estimation output for the three hedge setups that include external spot series, based on monthly returns. The reported p-values refer to the null-hypothesis that the coefficient in question is zero (i.e., redundant in the regression equation).

B-GARCH (Diagonal BEKK): Corn Reuters, continuation spot
& Corn Matif, continuation EMAc2:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a rank one matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	0.001275	0.000713	1.787843	0.0738
M(1,2)	0.000653	0.000366	1.786365	0.0740
M(2,2)	0.000334	0.000197	1.696646	0.0898
A1(1,1)	0.062935	0.144894	0.434354	0.6640
A1(2,2)	0.069913	0.087918	0.795208	0.4265
B1(1,1)	0.839981	0.092613	9.069844	0.0000
B1(2,2)	0.947272	0.027432	34.53190	0.0000

B-GARCH (Diagonal BEKK): Wheat Oil World, spot continuation
& Wheat CBOT, continuation Wc1:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a diagonal matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	2.71E-05	3.63E-05	0.747090	0.4550
M(2,2)	1.22E-05	3.89E-05	0.312895	0.7544
A1(1,1)	0.257541	0.038995	6.604390	0.0000
A1(2,2)	0.244583	0.046600	5.248512	0.0000
B1(1,1)	0.966901	0.012155	79.54999	0.0000
B1(2,2)	0.971411	0.014198	68.42068	0.0000

B-GARCH (Diagonal BEKK): Soy Oil World, continuation spot
& Soybean Meal CBOT, continuation SMc1:

Covariance specification: Diagonal BEKK
GARCH = M + A1*RESID(-1)*RESID(-1)'*A1 + B1*GARCH(-1)*B1
M is a rank one matrix
A1 is a diagonal matrix
B1 is a diagonal matrix

Transformed Variance Coefficients

	Coefficient	Std. Error	z-Statistic	Prob.
M(1,1)	0.002188	0.000537	4.073945	0.0000
M(1,2)	0.001566	0.000402	3.896750	0.0001
M(2,2)	0.001121	0.000336	3.336000	0.0008
A1(1,1)	0.471200	0.079125	5.955119	0.0000
A1(2,2)	0.342050	0.061050	5.602782	0.0000
B1(1,1)	0.608452	0.097013	6.271891	0.0000
B1(2,2)	0.859135	0.033195	25.88138	0.0000