The Hedging Performance of Electricity Futures on the Nordic Power Exchange Nord Pool

Hans NE Byström
Department of Economics
Lund University
P.O. Box 7082
S-220 07 Lund
Sweden
phone: +46-46-222 79 09
E-mail: hans.bystrom@nek.lu.se

November 2, 2000

Abstract

The Nordic Power Exchange (Nord Pool), the first multinational exchange for electricity trading, has existed since January 1996. Spot and futures contracts are traded on this exchange and its typical characteristics are very high volatility as well as non-normally distributed returns. In this paper I look at electricity futures and how they can be used for short term hedging in the spot market. I study the minimum variance hedge ratio and how it can be estimated in different ways. The traditional naive hedge and the OLS hedge are compared out-of-sample to more elaborate moving average and GARCH hedges. The empirical results indicate some gains from hedging with futures despite the lack of straight-forward arbitrage possibilities in the electricity market. Furthermore, I find that the relative performance of the different variance minimizing hedges depends on whether unconditional or conditional variances are studied.

JEL Classification Codes: C22, C53, G13, Q49.
Keywords: electricity prices, hedging.

1 Introduction

One of the latest markets to become standardized and organized on an exchange is the electric power market. Allover the world, there are only a handful of operating electricity exchanges, with the Nordic Power Exchange, "Nord Pool", being the only multinational one. On this exchange, there are actors from Norway, Sweden, Denmark, Finland, and England. The early deregulation of their electricity markets is an important factor behind the development of an exchange for trading in electricity in the Nordic countries; in particular the Norwegian and the Swedish markets have become fairly competitive in the last few years.

With the setup of an organized exchange, the situation for electricity producers and distributors has changed; from a situation where a reliable supply of energy was most important, the focus has partly shifted to obtaining optimal financial performance and efficient risk management. On the Nordic Power Exchange, electricity is traded both on a spot market and a futures market. The main reason for trading in futures is for actors to monitor the volatility of their power portfolios and to minimize the negative effect of adverse fluctuations in electricity prices.

This paper investigates the statistical and distributional properties of spot and futures prices on Nord Pool, as well as the short-term hedging performance of these futures. An evaluation of the hedging performance is principally of interest in the electricity market, due to problems with the storage of electricity when arbitrage arguments are used for the pricing of futures. The high volatility in this market, many times as large as in traditional financial markets, also contributes to make hedging important.

When hedging *price risk*, the optimal proportion of the future contract that should be held to offset the cash position is called the optimal *hedge ratio*. This ratio is traditionally estimated by examining the ratio between the *unconditional* covariance between cash and futures prices and the *unconditional* variance of the price of futures. This method can be criticized on a number of grounds. First, the traditional optimal hedge ratio is only utility maximizing under certain assumptions on the futures returns, otherwise it is merely variance minimizing as shown by Myers (1991). Second, since almost all financial assets and commodities have time varying second moments, the hedge ratio will be time varying and possibly best modelled in a dynamic framework as a function of *conditional* covariances and variances (Baillie and Myers (1991)). These conditional covariances and variances are often modelled with the conditionally heteroscedastic ARCH- and GARCH-type models developed by Engle (1982) and Bollerslev (1986). Baillie and Myers (1991), for instance, employ a bivariate GARCH model to estimate hedge ratios for commodities and find that dynamic GARCH-based hedge ratios out-perform those ratios coming from the traditional unconditional approach.

In this paper, I consider time varying variances and covariances of Nordic electricity price returns over the period January 1996-October 1999 and investigate how the time variation affects the hedging performance out-of-sample on the Nordic Power Exchange. In addition to the traditional unconditional hedges, I apply different conditional hedges. On the one hand, I apply continuously updated 50-day moving averages of the second moments, and on the other, I apply two different GARCH models; first, the constant conditional correlation bivariate GARCH model proposed by Bollerslev (1990), and second, a multivariate GARCH model called Orthogonal GARCH where a "diagonalization" of the bivariate problem simplifies the multidimensional GARCH estimation (Ding (1994), Alexander and Chibumba (1998), Byström (1999)).

I find that short-term hedging of electricity spot prices with electricity futures, using different estimates of the optimal (or actually minimum variance) hedge ratio, significantly reduces the variability of the portfolio returns. My empirical findings also confirm that variances and hedge ratios vary significantly over time, but that the two GARCH models only slightly improve the hedging performance out-of-sample compared to the unconditional hedges. The traditional simpler hedging models perform as well as the more elaborate conditional models, if the performance of the hedges are evaluated on the basis of their ability to reduce the unconditional (sample) portfolio variance. If we look at the conditional variance of the hedging portfolios instead, then the GARCH based hedges outperform the other hedges.

The paper is organized as follows: chapter 2 describes the data and the statistical particularities of electricity prices as well as the general features of the Nordic electricity market and Nord Pool, chapter 3 deals with hedging and hedge ratios, and chapter 4 presents the hedging results. Chapter 5 concludes the paper.

2 Spot and Futures on The Nordic Power Exchange (Nord Pool)

In January 1996, the Swedish electricity market was deregulated and integrated with the previously deregulated Norwegian electricity market. At the same time, the first multinational electricity exchange, Nord Pool, was created. The exchange has participants from Sweden, Norway, Denmark, Finland, and England (January 1998), but only the Norwegian and the Swedish markets are fully integrated. Of all trade in electric power in these two countries, around 25% (January 1998) is managed by Nord Pool; the main part of the electricity trade is still organized as bilateral contracts between producers and consumers.

In this paper, I will deal with the two major markets on Nord Pool, the spot market and the futures market. The spot market is a market for physical delivery of electricity, while the futures market is organized as a purely financial market without physical delivery. In reality, the spot market is also a short-term (one-day) futures market, however. Each day at noon, spot prices and volumes for each hour the *following* day are determined at an auction; what is called a futures contract in the electricity market is actually a futures contract with a *future* as the underlying asset[1]. Otherwise, the futures contracts are highly standardized and defined in terms of a given number of MWhs of electricity for (hypothetical) delivery during a given future week. The contracts are designed to reduce risk and make it possible to secure electricity prices up to three years in advance. It is possible to go short (and long) in the futures market, but the physical properties of the commodity in question, electricity, makes it difficult (or impossible) to go both short and long in the spot market. This complicates the theoretical treatment of these markets as regards pricing and hedging.

2.1 Data

Due to the deregulated electricity market being a fairly new phenomenon, data of any useful size and quality has only recently become available. In this paper, I use daily (trading days) spot and futures prices (Fig.1) from Nord Pool for almost four years, January 2, 1996 to October 21, 1999[2]. From these daily prices (quoted in Norwegian kroner (NOK/MWh)) I calculate 946 daily returns (log-differences).

Each day a number of futures contracts with different maturities can be chosen as a hedge instrument and a decision of which contract to use must be taken. My intention is to study the short term hedging performance (one-week holding period), and since short term future contracts are more liquid as well as more correlated with the underlying spot prices than the longer term contracts, futures with three weeks left to maturity are chosen for the hedging investigations[3]. In order to avoid thin market and expiration effects, however, I roll over to the next three-week maturity contract one week prior to the expiration of the current contract. With this roll-over procedure, a time-series is created for the full time period 1996-1999[4]. A drawback from closing out the futures positions before expiration is that it introduces some basis risk to the hedge,

[1] Throughout this paper the "spot" price is treated as if it were a true spot price.

[2] The daily spot prices are so called "dygnspriser". These daily prices are the average prices of the 24 hourly quoted prices each day. The futures contracts are so called "veckokontrakt" for (hypothetical) delivery of a certain average amount of power over one week (vecka).

[3] In addition to these futures, I have also hedged spot movements with longer maturity futures (still with one-week holding period). In these cases, the hedging performance deteriorates compared to the shorter maturity futures case.

[4] In this way, the futures in my hedges always have a remaining lifetime of between 5 and 15 days.

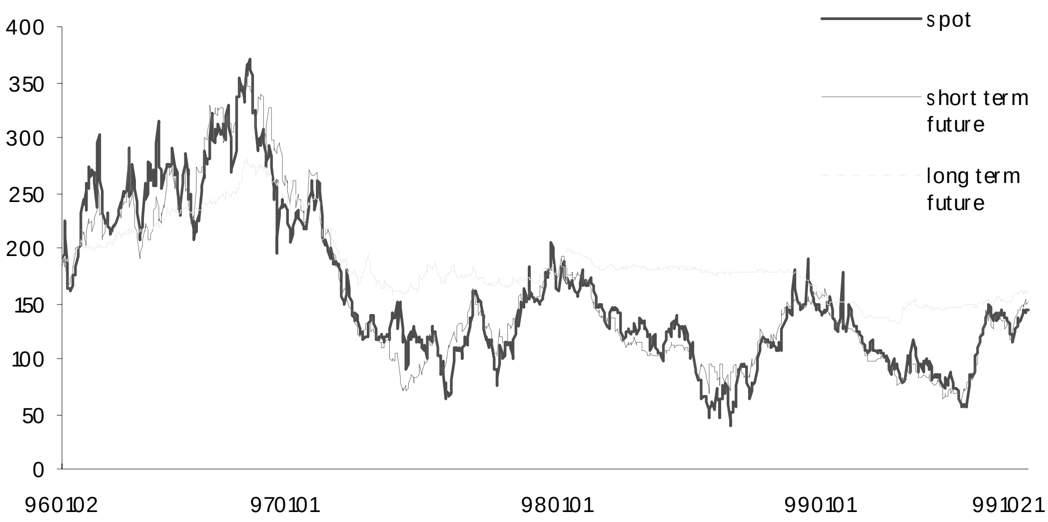

Figure 1: Electricity Prices at the Nordic Electricity Exchange (Nord Pool).

since the future price is not directly tied to the underlying spot price prior to the maturity date[5]. This basis risk is particularly serious in the electricity market, where large temporary deviations between spot and future prices appear, due to non-straightforward arbitrage possibilities.

In Fig.2, I have plotted daily spot and futures return volatilities over the second half of the data set; GARCH volatilities (from the constant correlation bivariate GARCH model) as well as sample volatilities (annualized standard deviations). The volatilities are apparently very high compared to ordinary asset and currency markets. It can also be seen how the GARCH model successfully captures the swings in volatilities.

The volatility of both spot as well as futures returns clearly varies over time and the assumption of identically and independently normally distributed returns seems unrealistic. Further evidence of this is given by the investigation of the return distributions in Fig.3 and Fig.4; the return distributions are skewed and fat-tailed compared to the normal distribution. A time varying conditional variance might be one reason for the fat-tailed unconditional distributions.

According to Table 1, both the spot and the futures price series (logarithms) in this sample are stationary already on the level. The assumption of stationary electricity prices might seem

[5] I already have basis risk due to the complexity associated with the closing of the futures; the futures are price contracts for an entire future week, not a single day. The basis risk due to this phenomenon should be partly eliminated by the roll over procedure though; no actual expiration is ever allowed.

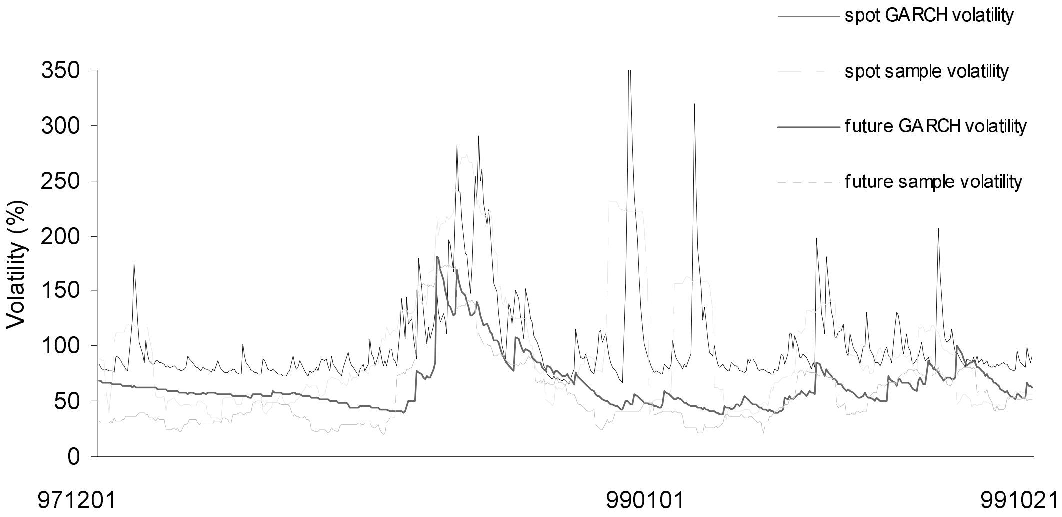

Figure 2: Spot Return Volatility and Futures Return Volatility. The GARCH volatility comes from the constant correlation bivariate GARCH model. The sample volatility is annualized in percent (20-day window).

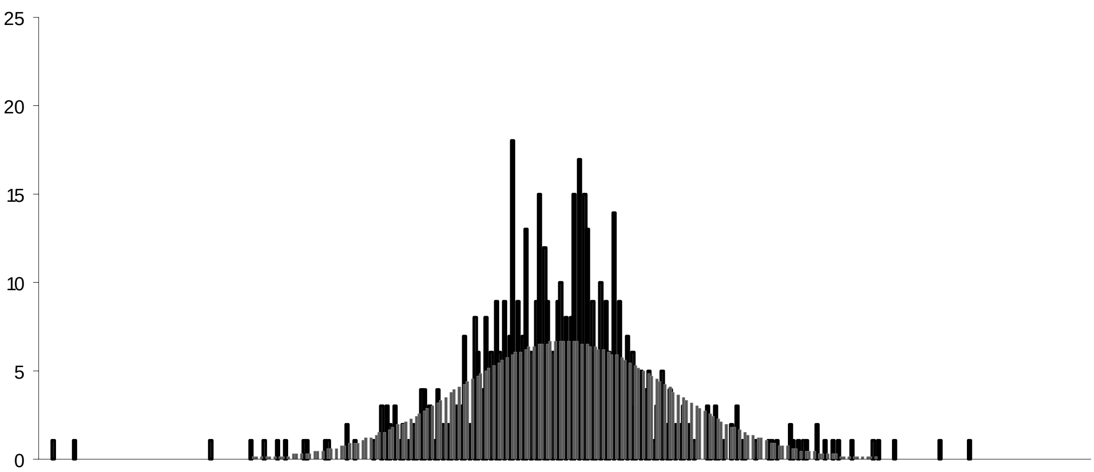

Figure 3: Spot Return Distribution vs. normal Distribution.

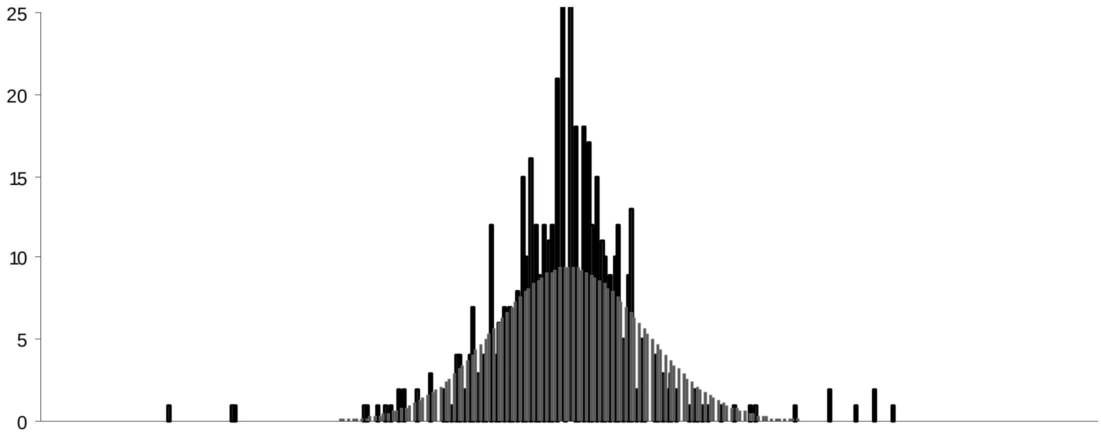

Figure 4: Future Return Distribution vs. normal Distribution.

Table 1: Stationarity of logprices over the sample period January 2, 1996 to October 21, 1999.

	Phillips-Perron (no trend)	Phillips-Perron (with trend)
Spot prices	-38.97	-84.16
Future prices	-24.71	-42.12

The 99 percent critical values for the Phillips-Perron test with and without trend are -3.96 and -3.43.

unrealistic in the long run, but for the sample at hand, it cannot be rejected. Error correction due to cointegration is therefore not expected to improve the behavior of the GARCH models.

Table 2 reports some statistics on the return series. Both the futures returns and the spot returns have means not significantly different from zero. The unconditional distributions of spot returns and, in particular, futures returns are non-normal, as evidenced by skewness, high excess kurtosis, and highly significant Bera-Jarque statistics[6]. Tests for autocorrelation, using different Q-statistics, indicate that no autocorrelation is present in the spot market while some correlation exists at long lags in the futures market. Finally, a test for ARCH, using the Q^2-statistic, finds significant ARCH effects at all lags in the spot market and at long lags in the futures market.

[6] The Bera-Jarque (B-J) statistic is χ^2_2-distributed under the null of normality. The statistic is $n \cdot [\frac{skewness^2}{6} + \frac{excess\ kurtosis^2}{24}]$,where n is the sample size.

Table 2: Return statistics January 2, 1996 to October 21, 1999.

	mean	variance	excess kurtosis	skewness	B-J	Q(6)	Q(18)	$Q^2(6)$	$Q^2(18)$
Spot	-0.000273 0.066	0.00429	6.698	-0.055	1768.83	6.92	20.56	146.16	175.20
Future	-0.000210 0.042	0.00175	9.211	1.004	3503.14	4.57	52.72	8.74	175.40

Small figures denote standard errors. B-J is the Bera-Jarque test for non-normality and Q(.) is the Ljung-Box test for autocorrelation. 99 percent critical value for Bera-Jarque is 9.21 and 99 percent critical values for Ljung-Box are 16.8, and 34.8.

3 Hedging Strategies

In this chapter, the minimum variance hedge ratio is estimated[7]. This hedge ratio determines how many futures contracts should be bought or sold for each spot contract for an investor to minimize the variance of his portfolio returns. There are many alternative ways of estimating this ratio, and this chapter looks at both unconditional and conditional estimates.

3.1 The Minimum Variance Hedge Ratio

To secure positions in a spot market, traders use futures as hedging instruments. For each spot contract, the hedge ratio tells us how many futures contracts should be purchased or sold. Let s_{t+1} and f_{t+1} be the changes in spot and futures prices, respectively, between time t and $t+1$, respectively, and let h_t be the hedge ratio at time t. Then,

$$x_{t+1} = s_{t+1} - h_t f_{t+1} \tag{1}$$

is the return to a trader going long in the spot market and going short in the futures market at time t. The variance of this return portfolio is

$$var_t(x_{t+1}) = var_t(s_{t+1}) + h_t^2 \cdot var_t(f_{t+1}) - 2 \cdot h_t \cdot cov_t(s_{t+1}, f_{t+1}), \tag{2}$$

and the minimum variance ratio, $h_{t,\min.var.}$, can then be derived by simply minimizing this variance with respect to h_t. We end up with the following expression for $h_{t,\min.var.}$:

$$h_{t,\min.var.} = \frac{cov_t(s_{t+1}, f_{t+1})}{var_t(f_{t+1})}. \tag{3}$$

[7]The minimum variance hedge ratio is also the optimal one (maximizing the mean-variance expected utility) if the futures prices either behave as martingales or the investors have infinite risk aversion. Neither of these conditions are expected to hold in the electricity market, however.

In the particular case of electricity I invent a scenario where an electricity producer/distributor knows he is to sell electricity on the spot market on a particular day in the near future and wants to hedge with electricity futures. Since the aim is study the short term hedging performance, a one-week investment horizon is assumed. For each spot position, the producer goes short in $h_{t,\min.var.}$ contracts in the futures market, where $h_{t,\min.var.}$ comes from (3). This operation is then repeated each day in the (out-of-sample) *test period*. As described below, $h_{t,\min.var.}$ is either modelled as constant throughout the test period or as time varying and updated on a daily basis.

3.2 Estimating the Hedge Ratio

The hedge ratio is estimated in five different ways; the naive one-to-one hedge ratio, where each spot contract is offset by exactly one futures contract, the OLS-hedge ratio where a regression of the spot returns on futures returns gives the hedge ratio expressed as in (3), but in an invariant unconditional version[8], a dynamic hedge ratio calculated by using continuously updated moving averages (50 days back in time) of variances and covariances and finally, two different dynamic hedge ratios based on bivariate GARCH modelling of the two return series (estimated by using daily data).

Several time series, including our spot and futures return series, exhibit periods of unusually large volatility followed by periods of tranquility. Under such circumstances, the assumption of a constant variance is obviously not appropriate. In order to capture the varying variance, the conditional variance can be modelled as a function of past errors as well as its own lags. This is done in GARCH models, and the first GARCH model to be estimated is the bivariate constant conditional correlation GARCH model introduced by Bollerslev (1990). The second choice is a multivariate GARCH model called Orthogonal GARCH, where the use of principal components analysis "diagonalizes" and simplifies the problem (Ding (1994), Alexander and Chibumba (1998) and Byström (1999)).

In the first GARCH model, spot returns and futures returns (daily returns) are modelled within the bivariate constant conditional correlation framework of Bollerslev (1990). The mean equations are specified as AR(2) processes and the conditional variance equations as GARCH(1,1):

$$\begin{aligned} y_{s,t} &= \alpha_{s,1} + \alpha_{s,2} y_{s,t-1} + \alpha_{s,3} y_{s,t-2} + \varepsilon_{s,t} \\ y_{f,t} &= \alpha_{f,1} + \alpha_{f,2} y_{f,t-1} + \alpha_{f,3} y_{f,t-2} + \varepsilon_{f,t} \end{aligned}$$

[8] Regression of futures returns on spot returns over the estimation period gives an OLS estimate of the optimal hedge ratio $h_{opt} = 0.408$, with a $standarderror = 0.049$. Daily returns, instead of weekly, are used since this is expected to improve the estimate (Duffie (1989)).

$$\begin{aligned}\sigma^2_{s,t} &= \phi_{s,1} + \phi_{s,2}\varepsilon^2_{s,t-1} + \phi_{s,3}\sigma^2_{s,t-1} \\ \sigma^2_{f,t} &= \phi_{f,1} + \phi_{f,2}\varepsilon^2_{f,t-1} + \phi_{f,3}\sigma^2_{f,t-1} \\ \sigma_{sf,t} &= \rho\sigma_{s,t}\sigma_{f,t},\end{aligned} \tag{4}$$

where $\sigma^2_{s,t}$ and $\sigma^2_{f,t}$ are the conditional variances of $\varepsilon_{s,t}$ and $\varepsilon_{f,t}$, $\sigma_{sf,t}$ is the conditional covariance between $\varepsilon_{s,t}$ and $\varepsilon_{f,t}$, and $\varepsilon_t = \sigma_t u_t$, where $u_t \sim N(0,1)$[9]. The time-varying conditional covariance between spot and futures returns are parametrized to be proportional to the product of the corresponding conditional standard deviations. This assumption greatly simplifies the computational burden in the estimation compared to more elaborate multivariate models. On the left-hand side of Table 3, estimation results from the Maximum Likelihood estimation (using the BHHH algorithm implemented in GAUSS) of the bivariate GARCH model are presented. It is shown how all GARCH parameters as well as the correlation coefficient are significantly different from zero. According to the statistics in the lower part of Table 3, both the spot and futures residuals distribution approximate a standard normal distributions fairly well, even though some autocorrelation remains for the futures residuals. Some skewness and kurtosis remain as well.

In the second model, the Ding (1994) and Alexander and Chibumba (1998) Orthogonal GARCH model is applied by using orthogonal factors. The main idea is to use principal components analysis to generate a number of orthogonal factors which can each be treated in a simple univariate GARCH framework.

There are *two* daily return series, spot and futures, with t observations represented as a $t \times 2$ matrix, $\mathbf{Y}_t$. The $t \times 2$ matrix, $\mathbf{P}_t$, of principal components is then defined as

$$\mathbf{P}_t = \mathbf{Y}_t\mathbf{W}_t \tag{5}$$

where $\mathbf{W}_t$ is the orthogonal 2×2 matrix of eigenvectors of $\mathbf{Y}_t^T\mathbf{Y}_t$ ordered according to the size of the corresponding eigenvalue. Notice that $\mathbf{P}_t$, just as $\mathbf{W}_t$, is now an orthogonal matrix. By inverting (5), one gets the principal components representation of the system

$$\mathbf{Y}_t = \mathbf{P}_t\mathbf{W}_t^T.$$

Ω_t, the variance of $\mathbf{Y}_t$ at time t, can now be calculated as

$$\Omega_t = var(\mathbf{Y}_t) = var(\mathbf{P}_t\mathbf{W}_t^T) = \mathbf{W}_t\mathbf{D}_t\mathbf{W}_t^T, \tag{6}$$

[9] This paper is limited to the standard GARCH model without asymmetrical extensions or dummies, and it is usually sufficient to limit the order of the GARCH model to (1,1). The same model has also been estimated with the conditional distribution for the error term modelled as the Student's t-distribution. No significant change in either the time-varying hedge ratio or the hedgeportfolio variance has been observed.

Table 3: Typical GARCH parameter estimates. Maximum Likelihood estimates as well as standardized residual statistics.

	Bivariate GARCH		Orthogonal GARCH	
	Spot	*Futures*	*First Principal Component*	*Second Principal Component*
α_1	0.000464 0.00185	-0.000128 0.00129	-0.000254 0.00181	0.000588 0.00127
α_2	0.0750 0.0392	-0.0661 0.0470	0.104 0.0411	-0.0846 0.0456
α_3	-0.0121 0.0421	-0.0420 0.0471	-0.00146 0.0440	-0.0802 0.0461
ϕ_1	0.000299 0.000173	0.0000570 0.0000341	0.000442 0.000269	0.0000461 0.0000278
ϕ_2	0.199 0.0305	0.0758 0.0110	0.229 0.0388	0.0637 0.0105
ϕ_3	0.736 0.0284	0.898 0.0146	0.681 0.0418	0.910 0.0158
ρ	0.290 0.0270			
Mean	-0.0189	-0.0179	0.0159	-0.0020
Standard Deviation	1.003	1.003	1.001	0.999
Skewness	-0.024	0.326	-0.127	0.050
Excess Kurtosis	3.005	6.822	2.982	7.361
$Q(6)$	4.49	12.73	8.05	18.98

The first 722 daily observations are used. Small figures are standard errors and Q(6) denotes the Ljung-Box test for the first six lags.

where $\mathbf{D}_t$ is a diagonal matrix of principal component variances at t and where W_t is assumed to be known at time t[10]. This also enables us to calculate the forecasted covariance matrix $\Omega_{t+1} \mid \Psi_t$ by univariate methods, where Ψ_t is the information set at t; for each principal component, the conditional variance of the principal component i, $var_{t+1}(P_i \mid \Psi_t)$, can easily be forecasted by, for instance, any univariate GARCH model. This gives us the Orthogonal GARCH specification. This particular conditional covariance matrix also has the advantage of always being positive definite, since D_t is diagonal with positive elements along its diagonal.

On the right-hand side of Table 3, estimation results from the univariate GARCH estimates of the principal components are presented by using the previous AR(2)-GARCH(1,1) model. As for the constant correlation bivariate model, all GARCH parameters are significant. As for the constant correlation bivariate GARCH model, the Orthogonal GARCH residuals are skewed and peaked, indicating (slightly) less than a perfect fit.

[10] W_t does not change very much from day to day and W_t can be approximated with W_{t-1}, without introducing large errors in the calculation of the covariance matrix. This is particularly the case when calculating the *forecasted* covariance matrix. In this case, we forecast at time $t-1$, only using information up to $t-1$ (including W_{t-1}).

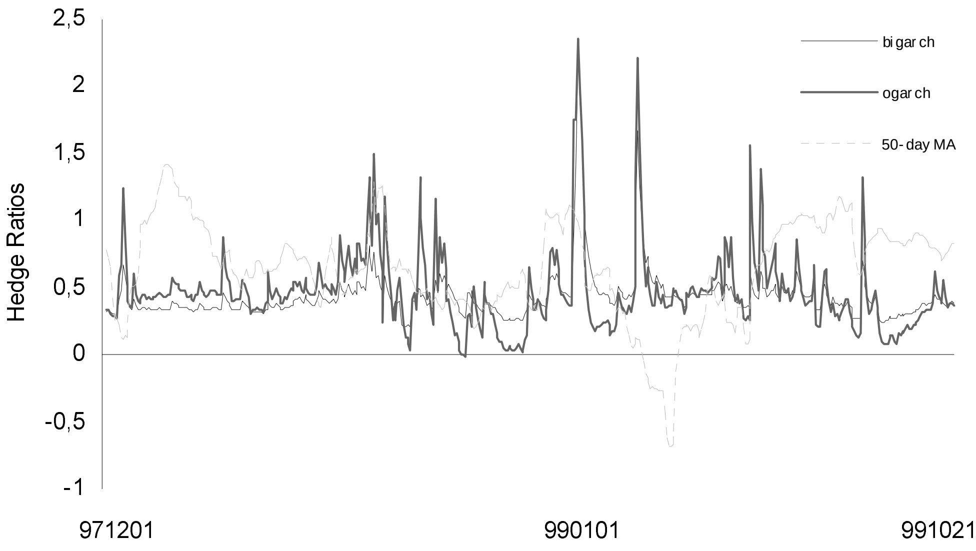

Figure 5: Conditional hedge ratios over the test period.

It is important to remember that Orthogonal GARCH, just like most other multivariate GARCH specifications, is based on certain assumptions. When assuming the conditional covariance matrix of the principal components to be diagonal, we foresee the fact that only the *unconditional* covariance matrix is diagonal. The *conditional* covariances need not be zero (Alexander and Chibumba (1998)), which could create some problems for the orthogonal GARCH methodology.

4 Hedging Performance

The data set is divided into an *estimation period* of 473 days, January 2, 1996 to November 30, 1997, and an equally long 473 day *test period,* December 1, 1997 to October 21, 1999. Calculating weekly returns from daily data over the test period, test series with 469 *overlapping* weekly spot- and futures returns are constructed. From each of these series, I then construct five different test series with 93 *non-overlapping* weekly spot- and futures returns[11]. To further assess the

[11] The series with 469 overlapping observations contain Monday to Monday weekly returns, Tuesday to Tuesday weekly returns, etc. Each of the five series with 93 non-overlapping returns contains returns from only one of the five working days, for instance Monday to Monday weekly returns. In this way, almost all autocorrelation found in the overlapping weekly return series is removed.

stability of the results the test period is divided into three equal long subperiods; December 1, 1997 to July 23, 1998, July 24, 1998 to March 8, 1999, and March 9, 1999 to October 21, 1999. For the dynamic models, the (one-week) hedge ratio is updated each day in the *test period;* in the two GARCH cases, GARCH *forecasts* of weekly (5-days) covariances and variances by iteration of the variance equation are used (using an expanding sample when estimating the GARCH parameters). In this way, some of the hedge ratio variability is captured and we get the three series of varying hedge ratios pictured in Fig.5. These ratios are all stationary and the means of the different hedge ratios all close to the constant OLS hedge ratio[12].

4.1 Different Evaluation Approaches

The purpose of the entire hedging exercise has been to minimize the variance (uncertainty) of the hedge portfolio (which, in the case of an infinitely risk avert trader or Martingale futures, is equal to utility maximizing). To evaluate the hedging performance out-of-sample, I look at the *test period* and the test series defined above, covering approximately the months December 1997 to October 1999. Throughout the paper, I evaluate out-of-sample and all hedge ratios are *predicted* hedge ratios using *predicted* variances and covariances.

I have chosen to evaluate to what extent the hedges minimize both the unconditional variance and the conditional variances over the test period. The unconditional variance is simply the sample variance of the (one-week) spot returns and hedge portfolio returns over the test period. In calculating the conditional variance I assume either the bivariate GARCH model or the Orthogonal GARCH model to be the underlying "data generating process", and each day I calculate (one-week) conditional spot and hedge portfolio variances that can be compared. The average conditional variances over the test period are then calculated and compared. The close fit of both the GARCH models to data, as demonstrated in Table 3 and in Fig.2, makes the assumption of data generated by either of these GARCH models plausible. The interesting information from this conditional evaluation is to what additional extent GARCH hedges reduce the conditional portfolio variance, in the presence of heteroscedasticity, compared to hedges that do not incorporate GARCH effects.

Both the unconditional variance approach, Kroner and Sultan (1993), and Park and Switzer (1995), as well as the conditional variance approach, Baillie and Myers (1991), Sephton (1993), and Bera, Garcia and Roh (1997), can be found in the literature.

---

[12]The Phillip Perron values with and without trend are -71.7 and -71.9 for the BIGARCH hedge ratio, -110.4 and -111.5 for the OGARCH hedge ratio, and -16.13 and -16.09 for the 50-days MA hedge ratio. The mean hedge ratios are 0.43, 0.47, and 0.64, respectively.

Table 4: Unconditional portfolio return variance and percentage variance reduction out-of-sample compared to the non-hedged spot position. Overlapping weekly portfolio returns.

	Spot	Naive	OLS	BIGARCH	OGARCH	50-days
Variance	0.0227	0.0193	0.0195	0.0201	0.0208	0.0200
% Variance Reduction	–	17.79	16.44	12.87	8.97	13.46

4.2 The Unconditional Variance

This subsection studies the unconditional variance of the spot and portfolio returns by simply calculating the sample variance of the different return series over the test period as in Kroner and Sultan (1993) and Park and Switzer (1995). In Table 4, the variances and variance reductions (out-of-sample) of the different hedges compared to the open spot position are presented for the test series with overlapping returns. All hedges reduce the variance. Somewhat surprisingly, the best hedges are the unconditional naive hedge and the OLS hedge. All the dynamic hedge ratios perform worse than the static ones and there do not seem to be any major gains from modelling spot and futures returns on Nord Pool with time-varying volatilities[13]. The finding that the naive hedge performs equally well as (or even slightly better than) the OLS hedge was also found in the US stock index market by Park and Switzer (1995), which is an example of how simpler models sometimes work better than more elaborate ones.

In order to assess the significance of the results in a statistical sense, I turn to "bootstrap" techniques to get the distributions of the portfolio variance estimates. Bootstrapping return series with overlapping returns is possible by systematically picking (with replacement) groups of five (the dependences reach over five days) returns from the series until a new bootstrapped series of equal length as the original series is constructed (Shau and Tu (1995)). This procedure is repeated 1000 times and for each bootstrapped series, a new estimate of the unconditional return variance is found. In Table 5, means and standard deviations from these 1000 variance estimates are presented. It can be noticed how, as expected, the mean variances correspond closely to the actual sample variances in Table 4. From the size of the standard deviations, it can be concluded that even if the hedge portfolios have a smaller variance for our sample than the spot position, no hedge differs in a statistically significant way (at traditional significance levels) from the unhedged spot position. At the same time, no hedge significantly differs from any other hedge in its hedging performance.

[13] In the case of orthogonal GARCH, the weak performance could at least partly be due to a break down in the orthogonalization methodology

Table 5: Means and standard deviations of the bootstrapped (1000x) unconditional variance distributions. Overlapping weekly portfolio returns.

	Spot	Naive	OLS	BIGARCH	OGARCH	50-days
Mean Variance	0.0230	0.0193	0.0195	0.0201	0.0209	0.0203
Standard Deviation	0.0033	0.0035	0.0033	0.0034	0.0039	0.0034

Table 6: Unconditional portfolio return variance and percentage variance reduction out-of-sample compared to the non-hedged spot position. Non-overlapping weekly portfolio returns.

		Spot	Naive	OLS	BIGARCH	OGARCH	50-days
	series 1	0.0169	0.0145	0.0140	0.0144	0.0149	0.0147
	series 2	0.0271	0.0191	0.0213	0.0218	0.0221	0.0210
Variance	series 3	0.0253	0.0227	0.0226	0.0229	0.0241	0.0229
	series 4	0.0270	0.0263	0.0252	0.0263	0.0274	0.0267
	series 5	0.0191	0.0154	0.0159	0.0163	0.0170	0.0163
	series 1	–	14.19	17.03	14.50	11.61	12.65
	series 2	–	29.27	21.15	19.40	18.49	22.50
% Variance Reduction	series 3	–	10.30	10.63	9.58	4.82	9.41
	series 4	–	2.74	6.68	2.47	−1.64	0.95
	series 5	–	19.50	16.62	14.59	11.18	14.60

Series 1 represents Monday to Monday returns, series 2 represents Tuesday to Tuesday returns, etc. In this way, I work with non-overlapping returns with almost all autocorrelation removed.

The overlapping nature of the returns (in the test period) obviously induces dependences in the return series above. To study the effect of this on the relative hedge performance, Table 6 shows the same entries as Table 4, but now calculated for non-overlapping return series where almost all autocorrelation in the returns is removed. From the five different test series it can be seen how the conclusions drawn from Table 4 remain valid. All the different hedges systematically reduce the portfolio variance compared to the spot variance even though the reduction is much smaller for some series (for instance series 4) than for others[14]. Once more, even for the non-overlapping return series, the naive hedge and the OLS hedge perform better than the conditional hedges.

As a final variation of this theme, the 469 observation long test period is split into the three

[14]The Orthogonal GARCH hedge actually increases the variance for the fourth non-overlapping series.

Table 7: Unconditional portfolio return variance and percentage variance reduction out-of-sample compared to the non-hedged spot position. The three subperiods.

		Spot	Naive	OLS	BIGARCH	OGARCH	50-days
	subperiod 1	0.0118	0.0101	0.0102	0.0100	0.0102	0.0114
Variance	subperiod 2	0.0376	0.0372	0.0354	0.0370	0.0388	0.0374
	subperiod 3	0.0176	0.0106	0.0123	0.0125	0.0130	0.0113
	subperiod 1	–	17.25	15.58	18.03	15.98	3.89
% Variance Reduction	subperiod 2	–	0.99	6.04	1.52	−3.09	0.43
	subperiod 3	–	66.91	43.08	40.60	36.01	55.73

The three subperiods correspond to the three time periods of equal length 1997-12-01 to 1998-07-23, 1998-07-24 to 1999-03-08, and 1999-03-09 to 1999-10-21.

subperiods of equal length. As can be seen in Table 7, the hedging performance differs somewhat between the three subperiods but for all hedges and time periods, the hedged positions vary less than the unhedged spot position. In particular, for the special case of the Orthogonal GARCH hedge, it can clearly be seen how the performance deteriorates in period two; probably due to the high return variance in this period. The high volatility spills over to both hedge ratios and portfolio returns and, as mentioned above, the performance of the Orthogonal GARCH technique is particularly sensitive to highly volatile time periods.

It has been shown how the non-hedged spot position is, overall, more volatile than the hedge portfolios, indicating how hedging spot positions with futures contracts on the Nordic Power Exchange can be profitable for a variance-minimizing trader. To further evaluate the performance of the different hedges, I have chosen to look at another performance measure; how often the weekly portfolio return is actually smaller (in an absolute sense) than the weekly spot return. In Table 8, I count the number of times (out of 469) the absolute weekly return of the hedged positions is smaller than the absolute weekly spot return[15]. While the naive hedge now performs relatively worse, the two GARCH hedges perform equally well as the OLS hedge. Even though the moving average hedge is just slightly better than the naive hedge, it is clear how the conditional hedges tend to reduce the variance *as often* as the unconditional hedges, even if they, on average, do not remove *as much* of the variability as the unconditional hedges. It should also be noticed that the portfolio returns are actually larger than the spot returns in

[15]This is an alternative measure of the variability of the different portfolios. The large variance of a particular portfolio might be due to a small number of very large returns, even though most of the returns are smaller than the other portfolios. Such a portfolio should perform relatively better with this alternative variability measure.

Table 8: The number of times (out of the 469 weekly returns) that the hedge portfolios absolute weekly return is smaller (in an absolute sense) than the absolute weekly spot return.

	Naive	OLS	BIGARCH	OGARCH	50-days
number of times	280	306	306	305	281
% of the full sample	59.70	65.25	65.25	65.03	59.91

about 35 to 40% of the weeks.

So far, we can conclude that it has been possible to use futures hedges on Nord Pool to reduce the variance in the Nordic electricity market in the last two years, while the performance of the different hedge portfolios does not generally differ significantly from each other and also depends on the choice of evaluation measure as well as the time period. The fairly weak performance of the more elaborate GARCH hedges compared to the simple OLS hedge might be explained by estimation problems or the fact that GARCH models do not always perform as well out-of-sample as in-sample. There are no indications of severe estimation problems, however, and the problems associated with forecasts from GARCH models estimated on daily data should be quite limited at such short horizons as one week. When it comes to Orthogonal GARCH, which has primarily been developed for large highly correlated multidimensional problems, the weak performance could partly be explained by a breakdown of the orthogonalization technique in periods of high volatility and asymmetrically behaving weakly correlated assets.

4.3 The Conditional Variance

Studying the reduction of the conditional instead of the unconditional variance gives somewhat different results. Assuming that the true return processes, and the conditional covariance matrix, are generated by either of the two GARCH models gives us the possibility to compare the relative performance of the data-generating GARCH model and the other models in minimizing the conditional variance (Baillie and Myers (1991), Sephton (1993) and Bera, Garcia and Roh (1997)). Hedging the spot position each day in the test period and taking the average value of the conditional variance of the spot and hedge portfolio returns, provides a measure for each portfolio that should be as small as possible.

In Table 9, the (average) conditional variances and variance reductions (out-of-sample) of the different hedges compared to the open spot position are presented for two different choices of underlying covariance matrix[16]. The constant correlation GARCH model is (not unexpectedly)

[16] In Tables 9, 10, and 11 the percentage reduction of the average (over the test periods) conditional variances

Table 9: Portfolio return conditional variance (average value over the test period) and percentage reduction of this average conditional variance out-of-sample compared to the non-hedged spot position. Overlapping conditional weekly portfolio variances.

	Spot	Naive	OLS	BIGARCH	OGARCH	50-days
Conditional Variance, BIGARCH Cov. Matrix	0.0250	0.0268	0.0235	0.0232	0.0237	0.0250
Conditional Variance, OGARCH Cov. Matrix	0.0247	0.0243	0.0223	0.0213	0.0208	0.0239
% Variance Reduction, BIGARCH Cov. Matrix	–	−7.13	6.16	7.11	5.11	−0.07
% Variance Reduction, OGARCH Cov. Matrix	–	1.11	9.65	13.72	15.82	3.06

Table 10: Percentage reduction of the conditional portfolio variance (average value over the test period) out-of-sample compared to the conditional spot variance. Non-overlapping conditional weekly portfolio variances.

		Naive	OLS	BIGARCH	OGARCH	50-days
% Var. Reduction, BIGARCH Cov. Matrix	series 1	−7.16	6.17	7.11	5.53	0.17
	series 2	−7.32	6.32	7.10	5.34	−0.36
	series 3	−8.42	6.15	7.11	4.27	−0.54
	series 4	−6.93	6.23	7.09	5.22	−0.34
	series 5	−5.75	5.93	7.16	5.13	0.34
% Var. Reduction, OGARCH Cov. Matrix	series 1	−0.04	9.27	13.06	14.71	3.27
	series 2	−1.05	9.09	11.70	13.53	2.50
	series 3	2.11	10.02	14.14	17.11	1.98
	series 4	1.03	9.74	12.86	14.84	2.72
	series 5	3.78	10.24	16.83	18.96	4.60

the best model when it is itself assumed to have generated the data. The OLS hedge is the second best, dominating the other GARCH model, Orthogonal GARCH. On the other hand, when the Orthogonal GARCH model is assumed to have generated the data it is also producing the best hedge. This time the bivariate GARCH hedge performs better than the OLS hedge. The ranking of the very best hedges not surprisingly changes when changing the assumption of the covariance matrix, but those hedges that perform badly in one case, are also shown to perform badly in the other case. The best hedge in the unconditional variance evaluation, the naive hedge, is now instead the worst performer, and while the moving average model earlier

is presented. Instead, looking at the average of the percentage reductions over the test periods gives exactly the same relative performance and relative ranking of the different hedge portfolios.

Table 11: Portfolio return conditional variance and percentage reduction of the conditional variance out-of-sample compared to the non-hedged spot position. The three subperiods.

		Spot	Naive	OLS	BIGARCH	OGARCH	50-days
	subperiod 1	0.0161	0.0175	0.0152	0.0152	0.0154	0.0170
Cond. Variance,	subperiod 2	0.0390	0.0412	0.0366	0.0360	0.0370	0.0386
BIGARCH Cov. Matrix	subperiod 3	0.0195	0.0212	0.0182	0.0181	0.0184	0.0191
	subperiod 1	0.0165	0.0156	0.01469	0.0145	0.0142	0.0160
Cond. Variance,	subperiod 2	0.0367	0.0356	0.0328	0.0303	0.0294	0.0355
OGARCH Cov. Matrix	subperiod 3	0.0204	0.0217	0.0190	0.0187	0.0184	0.0198
	subperiod 1	–	−8.38	5.54	5.82	4.19	−5.66
% Var. Reduction,	subperiod 2	–	−5.70	6.16	7.64	5.30	1.17
BIGARCH Cov. Matrix	subperiod 3	–	−9.07	6.70	7.09	5.50	2.02
	subperiod 1	–	5.74	11.02	11.95	13.73	3.01
% Var. Reduction,	subperiod 2	–	3.06	10.54	17.36	19.95	3.19
OGARCH Cov. Matrix	subperiod 3	–	−6.33	6.86	8.37	9.80	2.87

Table 12: The number of times that the conditional hedge portfolio variance is smaller than the conditional spot variance (out of the 469 weekly variances).

	Naive	OLS	BIGARCH	OGARCH	50-days
number of times, BIGARCH Covariance Matrix	81	468	469	464	347
number of times, OGARCH Covariance Matrix	177	416	428	469	322
% of the full sample, BIGARCH Covariance Matrix	17.27	99.79	100.00	98.93	73.99
% of the full sample, OGARCH Covariance Matrix	37.74	88.70	91.26	100.00	68.66

performed as well as the GARCH models, it now barely reduces the variance at all. In Table 9, I assume, a priori, that return series follow GARCH models and the relative ranking of the different models is therefore of no major interest; the two GARCH models obviously dominate the other models. Instead, the focus should be on the absolute reduction in variance compared to the spot position, as well as the absolute performance of the GARCH models compared to the simpler models. From this, it should appear whether modelling the hedge ratio in this market with GARCH models is worth while. The answer to this question finally depends on how costly a dynamic GARCH hedge is, in terms of transaction costs etc., compared to the simpler static hedges.

As in the unconditional evaluation in the last subsection, we continue by looking at non-overlapping return series as well as subsamples of the whole test period. In Tables 10 and 11, we see how the results from Table 9 remain stable both for the non-overlapping series and the different subsamples[17]; the performance of the different hedges changes very little from series to series in Table 10 and even if there is some change in performance from subsample to subsample in Table 11, this does not affect the relative ranking of the different hedges.

In a similar way as done in Table 8 for the unconditional evaluation, I now look at the number of times the conditional portfolio variance is smaller than the non-hedged spot position variance over the test period. In Table 12, it is observed how the ranking from the earlier Tables 9, 10, and 11 remains unchanged, and that the GARCH model that is assumed to have generated the conditional covariance matrix *always (100%)* reduces the variance compared to the spot position.

When comparing the results from the conditional evaluation with the results from the unconditional evaluation, there are both similarities and differences. In both cases, and throughout the paper, it is quite obvious how hedging in this market can reduce the variance. However, while the naive model performed well in the unconditional evaluation context it is, by far, the worst performer in the conditional context. In some cases, it even increases the variance. The opposite holds for the two GARCH models; while not improving on the simple OLS hedge in the unconditional evaluation, the inclusion of heteroscedasticity and volatility clustering in calculating the hedge ratio clearly contributes towards an optimal hedge when looking at the conditional variance. The disappointing results from the last subsection are therefore reversed and the dynamic modelling of the hedge ratio seems to improve the hedging performance; when spot and futures returns can be modelled as GARCH processes, which is shown to be a plausible assumption in this market, then the hedge ratio should be modified and continuously updated according to these GARCH processes. It is hard to tell from this study however, which GARCH model actually captures most of the variability in the market, since the relative ranking of the two GARCH models strongly depends upon the choice of the true covariance matrix.

5 Conclusions

In this paper, I have looked at a possible scenario of an actor on the Nordic Power Exchange that hedges a position (at a one-week horizon) in the spot market with futures contracts. The traditional unconditional methods of calculating the minimum variance hedge ratio are extended to time-varying moving average and GARCH hedge ratios. A constant correlation bivariate GARCH framework is compared to the new multivariate Orthogonal GARCH approach. Out-

[17] In order to save space, only the variance reductions, not the actual portfolio variances, are presented in Table 10.

of-sample evidence presented in this paper indicates how both the traditional unconditional naive hedge, the unconditional OLS hedge, and the dynamic conditional GARCH hedges reduce the variance of the hedge portfolio compared to the spot position. The relative performance of the different hedges depends on the evaluation measure, however. The OLS hedge and the two GARCH hedges reduce both the unconditional and the conditional variance, while the naive hedge successfully reduces only the unconditional variance. Among the dynamical hedges, the moving average model is dominated by the GARCH models. Overall, there seem to be some gains from including heteroscedasticity and time-varying variances in the calculation of hedge ratios, but the constant OLS hedge ratio is nearly as successful in reducing the portfolio variance.

6 References

Alexander, C.O., and Chibumba, A.M (1998), "Orthogonal Garch: An Empirical Validation on Equities, Foreign Exchange and Interest Rates, unpublished manuscript.

Baillie, R. T., and Myers, R. J. (1991), "Bivariate Garch Estimation of the Optimal Commodity Futures Hedge", *Journal of Applied Econometrics*, 6, 109-124

Bera, A. K., Garcia, P., and Roh, J-S (1997), "Estimation of Time-Varying Hdge Ratios for Corn and Soybeans: BGARCH and Random Coefficient Approaches", OFOR Paper Number 97-06 December 1997

Bollerslev, T. (1986), "Generalized Autoregressive Conditional Heteroskedasticity", *Journal of Econometrics*, 31, 307-327.

Bollerslev, T. (1990), "Modeling the Coherence in Short-run Nominal Exchange Rates: A Multivariate Generalized ARCH Model", *The Review of Economics and Statistics*, 31, 498-505.

Brailsford J. T., and Faff, R.W. (1996), "An evaluation of Volatility Forecasting Techniques", *Journal of Banking and Finance* 20, 419-438

Byström, H., (1999), "Orthogonal GARCH and Covariance Matrix Forecasting in a Stress Scenario: The Nordic Stock Markets During the Asian Financial Crisis 1997-1998", unpublished manuscript, Lund University

Ding, Z., (1994), *Time Series Analysis of Speculative Returns*, PhD Thesis, UCSD.

Duffie, D. (1989), *Futures Markets*, Prentice-Hall

Engle, R.F. (1982), "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom", *Econometrica*, 50, 987-1007

Figlewski, S. (1984), "Hedging Performance and Basis Risk in Stock Index Futures", *Journal of Finance*, 39, 657-669

Kroner, K. F., and Sultan, J. (1993), "Time-varying Distributions and Dynamic Hedging with Foreign Currency Futures", *Journal of Financial and Quantitative Analysis*, 28, 535-551

Markowitz, H. M., (1952), "Portfolio Selection", *Journal of Finance*, 7, 77-91

Myers, R. J. (1991), "Estimating time varying optimal hedge ratios on future markets", *Journal of Futures Markets*, 11, 39-53

Park, T. H., and Switzer, L. N. (1995), "Time-varying distributions and the optimal hedge ratios for stock index futures", *Applied Financial Economics*, 5, 131-137

Pilipovic, D. (1998), *Energy Risk*, McGraw-Hill

Sephton, P. S. (1993), "Optimal hedge ratios at the Winnipeg commodity exchange", *Canadian Journal of Economics*, 26, No. 1, 175-193

Shau, J., and Tu, D. (1995), *The Jackknife and Bootstrap*, Springer

Storset, S., (1995), "Risikostyring i Kraftmarkedet", *Derivatet*, No. 3