Chapter 8: Modelling of the Electricity Futures Market

Tables and Figures

Table 8.1 Delivery period structure for market models with time measured in years

	τ_c^b	τ_c^e	C
W_c - Week	$(c-1)/52$	$c/52$	6
M_c - Month	$(c-1)/12$	$c/12$	6
Q_c - Quarter	$(c-1)/4$	$c/4$	8
Y_c - Year	$c-1$	c	3

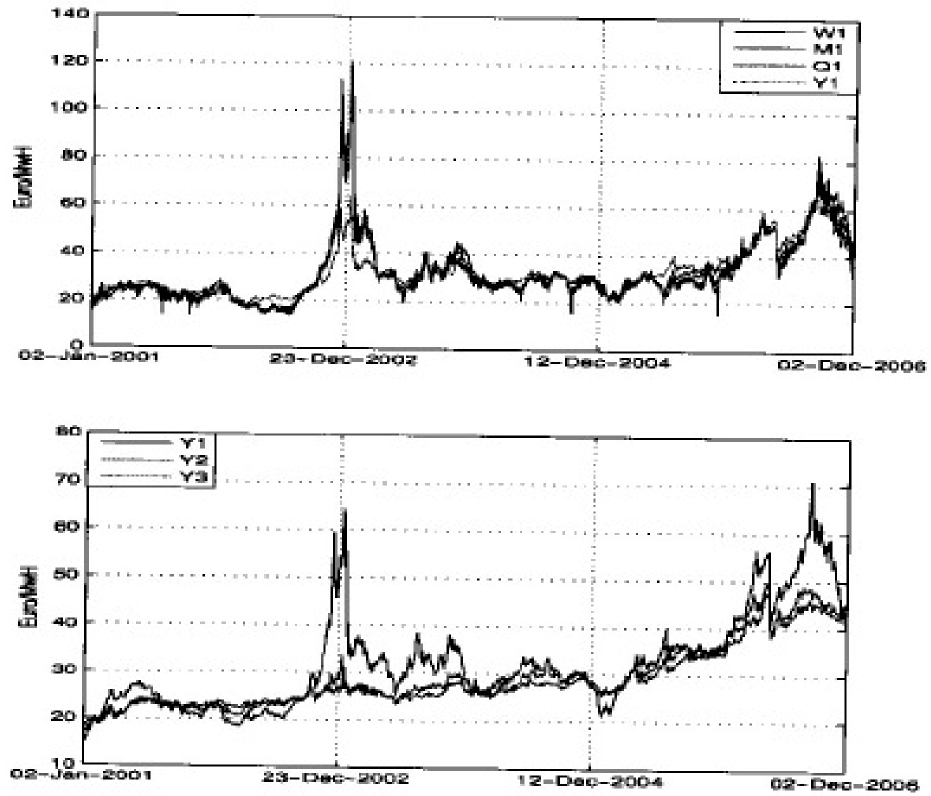

Fig. 8.1 Time series of electricity futures price data. Contracts with immediate delivery and varying delivery period - week (W_1), month (M_1), quarter (Q_1) and year (Y_1) on the upper panel. Contracts with varying time to delivery and yearly delivery period - year (Y_1), year (Y_2) and (Y_3) on the bottom panel.

Table 8.2 Descriptive statistics of price levels from 2 January 2001 untill 1 December 2006, a total of 1,479 trading days

	Mean	Variance	Skewness	Kurtosis
Week				
W_1	31.71	179.62	2.02	8.94
W_2	32.83	205.07	2.25	10.61
W_3	33.24	212.72	2.11	9.45
W_4	33.50	214.47	2.00	8.72
W_5	33.62	213.94	1.93	8.29
W_6	33.66	209.98	1.82	7.60
Month				
M_1	32.95	202.27	2.07	9.17
M_2	33.71	205.14	1.70	6.89
M_3	33.92	195.36	1.45	5.50
M_4	33.82	188.37	1.38	5.09
M_5	33.49	172.88	1.32	5.02
M_6	32.89	150.64	1.32	5.25
Quarter				
Q_1	33.52	197.18	1.72	6.98
Q_2	33.40	165.26	1.30	4.92
Q_3	31.69	106.68	1.10	3.96
Q_4	30.36	69.95	0.96	3.52
Q_5	29.79	69.65	0.83	2.86
Q_6	29.52	79.26	0.90	3.31
Q_7	29.24	68.01	1.07	3.54
Q_8	29.11	51.66	1.04	3.35
Year				
Y_1	32.23	110.04	1.10	3.83
Y_2	29.41	58.88	1.01	3.09
Y_3	28.75	46.22	1.14	3.29

Table 8.3 Descriptive statistics of electricity futures price returns. Volatility is annualized using 250 trading days a year.

	Mean	Volatility	Skewness	Kurtosis
Week				
W_1	0.00	90.2 %	0.21	25.01
W_2	0.00	59.4 %	0.60	19.76
W_3	0.00	56.2 %	0.01	25.39
W_4	0.00	55.8 %	−0.28	26.95
W_5	0.00	56.5 %	−0.12	29.48
W_6	0.00	54.7 %	−0.35	26.02
Month				
M_1	0.00	52.6 %	0.40	22.02
M_2	0.00	52.1 %	−0.92	22.54
M_3	0.00	50.3 %	−1.48	19.47
M_4	0.00	48.3 %	−1.34	18.42
M_5	0.00	48.1 %	−1.00	14.54
M_6	0.00	47.5 %	−0.86	13.18
Quarter				
Q_1	0.00	48.9 %	−0.76	22.32
Q_2	0.00	45.1 %	−1.06	15.12
Q_3	0.00	40.7 %	−0.41	39.43
Q_4	0.00	36.4 %	−0.22	20.22
Q_5	0.00	30.6 %	−0.58	12.93
Q_6	0.00	29.3 %	−0.84	14.21
Q_7	0.00	29.4 %	−1.19	17.96
Q_8	0.00	25.0 %	−0.32	9.24
Year				
Y_1	0.00	38.2 %	−0.52	14.32
Y_2	0.00	22.3 %	−0.94	12.26
Y_3	0.00	21.3 %	−0.38	8.28

Table 5.4 Volatility decomposed by season. Volatility is annualized assuming 250 trading days a year. "Const. vol" is the average volatility when all observations in the sample period are used. The rest of the table shows volatility calculated for individual months.

	Const. vol	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
Week													
W_1	90 %	102 %	75 %	62 %	60 %	84 %	79 %	103 %	74 %	98 %	86 %	135 %	97 %
W_2	59 %	105 %	70 %	40 %	38 %	44 %	45 %	40 %	50 %	35 %	52 %	42 %	106 %
W_3	56 %	100 %	46 %	44 %	39 %	40 %	39 %	42 %	53 %	36 %	46 %	42 %	105 %
W_4	56 %	93 %	46 %	44 %	44 %	38 %	40 %	42 %	52 %	36 %	44 %	46 %	106 %
W_5	56 %	91 %	46 %	45 %	46 %	42 %	41 %	39 %	51 %	36 %	44 %	49 %	111 %
W_6	55 %	90 %	44 %	44 %	47 %	41 %	40 %	36 %	51 %	36 %	44 %	49 %	101 %
Month													
M_1	53 %	90 %	47 %	39 %	38 %	37 %	42 %	40 %	48 %	35 %	43 %	46 %	94 %
M_2	52 %	90 %	44 %	44 %	46 %	38 %	37 %	35 %	51 %	36 %	42 %	45 %	88 %
M_3	50 %	94 %	47 %	42 %	42 %	36 %	33 %	36 %	50 %	37 %	40 %	46 %	69 %
M_4	48 %	84 %	43 %	41 %	45 %	35 %	29 %	36 %	47 %	35 %	39 %	49 %	70 %
M_5	48 %	82 %	39 %	44 %	45 %	34 %	32 %	35 %	43 %	35 %	39 %	48 %	75 %
M_6	47 %	80 %	38 %	37 %	42 %	34 %	28 %	32 %	42 %	34 %	44 %	47 %	81 %
Quarter													
Q_1	49 %	89 %	47 %	39 %	40 %	32 %	33 %	33 %	48 %	33 %	40 %	40 %	81 %
Q_2	45 %	78 %	37 %	37 %	41 %	32 %	28 %	32 %	42 %	32 %	41 %	47 %	69 %
Q_3	50 %	69 %	41 %	33 %	39 %	30 %	24 %	39 %	40 %	35 %	38 %	87 %	81 %
Q_4	36 %	43 %	28 %	35 %	41 %	27 %	25 %	36 %	26 %	26 %	29 %	35 %	72 %
Q_5	31 %	48 %	24 %	27 %	38 %	29 %	24 %	28 %	26 %	19 %	30 %	24 %	40 %
Q_6	29 %	48 %	29 %	30 %	32 %	31 %	22 %	23 %	19 %	30 %	21 %	18 %	33 %
Q_7	29 %	56 %	32 %	25 %	24 %	24 %	20 %	26 %	21 %	23 %	24 %	23 %	36 %
Q_8	25 %	35 %	20 %	20 %	27 %	22 %	20 %	26 %	25 %	31 %	22 %	21 %	24 %
Year													
Y_1	38 %	65 %	30 %	32 %	36 %	27 %	24 %	29 %	37 %	26 %	33 %	34 %	65 %
Y_2	22 %	37 %	22 %	22 %	26 %	22 %	16 %	20 %	18 %	17 %	19 %	15 %	26 %
Y_3	21 %	33 %	22 %	20 %	30 %	23 %	20 %	21 %	15 %	16 %	15 %	15 %	17 %

A Market Model for Electricity Futures

8.4 A market model for electricity futures

We consider a simple market model as discussed in Sect. 6.4, and recall it together with some notations. Assume that market participants trade C different electricity futures contracts with non-overlapping delivery periods. The price at time t for an electricity futures with delivery period $[\tau_c^b, \tau_c^e]$ is denoted by $F_c(t) = F_c(t, \tau_c^b, \tau_c^e)$, $c = 1, \ldots, C$. Assume that under the real world measure the price dynamics of $F_c(t)$ is lognormal, that is,

$$F_c(t) = F_c(0) \exp\left(\int_0^t A_c(u) du + \sum_{k=1}^{p} \int_0^t \Sigma_{c,k}(u)\, dB_k(u) \right), \qquad (8.2)$$

for $t < \tau_c^b$, with B_k, $k = 1, \ldots, p$ being independent Brownian motions and $\Sigma_{c,k}$ and A_c continuous functions on $[0, \tau_c^e]$. We assume that $p \leq C$, that is, the number of traded contracts is at least as many as we have Brownian motions driving the swap price dynamics. This implies in particular that the market is complete. We further note that the proposed model does not deal with any idiosyncratic risk. If $p < C$, (8.2) may allow for arbitrage opportunities. Including more Brownian motions will model the idiosyncratic risk, and also remove the possibility of arbitrage in our model. Our focus in the coming empirical study of the model (8.2) is on common risk, maturity

[5] [Benth and Koekebakker (2005)] include low liquidity daily contracts to their sample. Their argument is that including the daily contracts yields more correct short-term volatility estimates.

A Market Model for Electricity Futures

and seasonality effects in the volatility, thus the assumption $p \leq C$.

The logreturn over the period $[t_{n-1}, t_n]$ of the contract F_c is defined as

$$x_{n,c} \triangleq \ln\left(\frac{F_c(t_n)}{F_c(t_{n-1})}\right).$$

Using (8.2), we have

$$x_{n,c} = \int_{t_{n-1}}^{t_n} A_c(u)\,du + \sum_{k=1}^{p} \int_{t_{n-1}}^{t_n} \Sigma_{c,k}(u)\,dB_k(u). \tag{8.3}$$

Hence, the logreturns are normally distributed under the real world probability. By the Girsanov transform,[6] the drift will be altered under an equivalent martingale measure, while the volatility remains unchanged.

With $N+1$ trading days in our sample, the $N \times C$ data matrix $\mathbf{X}_{N\times C}$ is specified as

$$\mathbf{X}_{N\times C} = \begin{bmatrix} x_{1,1} & x_{1,2} & \cdots & x_{1,C} \\ x_{2,1} & x_{2,2} & \cdots & x_{2,C} \\ \vdots & \vdots & \ddots & \vdots \\ x_{N,1} & x_{N,2} & \cdots & x_{N,C} \end{bmatrix}.$$

In the next subsections we will analyse the factor dynamics, first by PCA, and next by estimating a multi-factor parametric model for electricity futures price returns.

8.5. Principal Component Analysis

In this section we will investigate the changes in the term structure of electricity futures price returns. PCA is utilised for the identification of structure within a set of interrelated variables. It establishes dimensions within the data, and serves as a data reduction technique. The aim is to determine factors (that is, principal components) in order to explain as much of the total variation in the data as possible.

We have a total of N observations of C return series, and collect time series of each contract in N-dimensional vectors $\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_C$. The data matrix $\mathbf{X}_{N\times C}$ is then

$$\mathbf{X}_{N\times C} = \left[\mathbf{x}_1 \; \mathbf{x}_2 \cdots \mathbf{x}_C\right]. \tag{8.4}$$

8.5. Principal Component Analysis

The corresponding sample covariance matrix of dimension $C \times C$ is denoted $\mathbf{\Omega}$. The orthogonal decomposition of the covariance matrix is

$$\mathbf{\Omega} = \mathbf{P}\mathbf{\Lambda}\mathbf{P}', \tag{8.5}$$

where

$$\mathbf{P} = [\mathbf{p}_1 \; \mathbf{p}_2 \cdots \mathbf{p}_C] = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1C} \\ p_{21} & p_{22} & \cdots & p_{2C} \\ \vdots & \vdots & \ddots & \vdots \\ p_{C1} & p_{C2} & \cdots & p_{CC} \end{bmatrix},$$

and $\mathbf{\Lambda}$ is a diagonal matrix with the eigenvalues $\lambda_1, \lambda_2, ..., \lambda_C$ on the diagonal. The matrix $\mathbf{P}$ is orthogonal, with ith column, $\mathbf{p}_i$, being the eigenvector corresponding to λ_i. $\mathbf{P}'$ denotes the transpose of $\mathbf{P}$. The matrix $\mathbf{Z} = \mathbf{XP}$ is called the matrix of principal components, while $\mathbf{P}$ the matrix of factor loadings. The eigenvectors on the diagonal of $\mathbf{\Lambda}$ are by convention ordered so that $\lambda_1 \geq \lambda_2 \geq ... \geq \lambda_C$. To explain all the variation in $\mathbf{X}$, we need C principal components. Since the objective of our analysis is to explain as much as possible of the covariance structure with just a few factors, we approximate the theoretical covariance matrix in (8.5) using only the first $M < C$ eigenvalues in $\mathbf{\Lambda}$ while putting the remaining equal to zero. The proportion of total variance accounted for by the first M factors is

8.5. Principal Component Analysis

$$\text{Cumulative contribution of first } M \text{ factors} = \frac{\sum_{i=1}^{M} \lambda_i}{\sum_{i=1}^{C} \lambda_i}.$$

The M factors should explain a "big" part of the total covariance of the underlying variables. In empirical studies, one is typically choosing M so that around 95% of the variation is explained.

Now we present the results from the PCA. First, we consider the complete data set. Recall that it consists of the 16 contracts representing the total market, $\{W_1, ..., W_4, M_2, ..., M_6, Q_3, ..., Q_8, Y_3\}$. Our results are comparable to the investigations in [Koekebakker and Ollmar (2005)] and [Frestad (2007a)]. Next, we analyse dynamics within each particular market segment (week, month, quarter and year). The descriptive statistics showed evidence of seasonality. Therefore all the return series have been normalised prior to the PCA analysis. Each return series is sorted according to the observation month, and then normalised by subtracting the mean of the series and dividing by the standard deviation.

8.5.1 *Principal component analysis of the total data set*

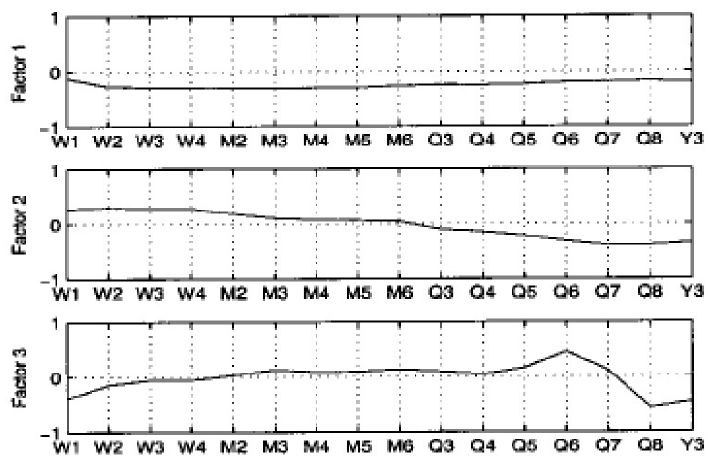

Fig. 8.2 First three factor loadings for the total data set.

In Fig. 8.2 we present factor loadings across contracts for the full data set. We see that the first factor can clearly be identified as a shifting factor. A shock to this factor shifts all contracts in the same direction. Factor two is the tilting factor. A shock to this factor moves weekly and monthly contracts in one direction, and the quarterly and yearly contracts in the opposite direction. The third factor is less clear, but it might perhaps be interpreted as a bending factor. The factor loadings change sign twice along the term structure, shifting the short and the long ends in one direction and the middle contracts in the opposite one. However, factor loadings are

Table 8.5 Individual and cumluative variance explained from PCA for the total data set

#	%-explained	%-cumulative
1	54 %	54 %
2	10 %	64 %
3	6 %	70 %
4	5 %	76 %
5	5 %	80 %
6	4 %	85 %
7	4 %	88 %
8	3 %	91 %
9	2 %	93 %
10	1 %	95 %
11	1 %	96 %
12	1 %	98 %
13	1 %	99 %
14	1 %	99 %
15	0 %	100 %
16	0 %	100 %

Table 8.6 Correlation matrix for total data set

	W_1	W_2	W_3	W_4	M_2	M_3	M_4	M_5	M_6	Q_3	Q_4	Q_5	Q_6	Q_7	Q_8	Y_5
W_1	1															
W_2	0.44	1														
W_3	0.31	0.87	1													
W_4	0.34	0.80	0.83	1												
M_2	0.35	0.69	0.76	0.78	1											
M_3	0.25	0.66	0.73	0.74	0.83	1										
M_4	0.32	0.61	0.68	0.68	0.76	0.86	1									
M_5	0.28	0.59	0.65	0.65	0.72	0.73	0.84	1								
M_6	0.19	0.53	0.59	0.60	0.61	0.66	0.65	0.83	1							
Q_3	0.17	0.48	0.53	0.53	0.55	0.56	0.57	0.56	0.55	1						
Q_4	0.22	0.48	0.52	0.53	0.59	0.58	0.57	0.54	0.50	0.65	1					
Q_5	0.20	0.42	0.45	0.46	0.54	0.57	0.55	0.53	0.49	0.46	0.67	1				
Q_6	0.09	0.34	0.38	0.37	0.42	0.49	0.48	0.48	0.40	0.45	0.41	0.60	1			
Q_7	0.12	0.32	0.35	0.35	0.36	0.42	0.44	0.41	0.36	0.41	0.44	0.39	0.67	1		
Q_8	0.14	0.32	0.33	0.33	0.33	0.36	0.38	0.37	0.32	0.35	0.39	0.37	0.23	0.53	1	
Y_5	0.15	0.33	0.37	0.37	0.41	0.43	0.45	0.44	0.42	0.44	0.46	0.48	0.35	0.38	0.70	1

8.5.2 *Principal component analysis for individual market segments*

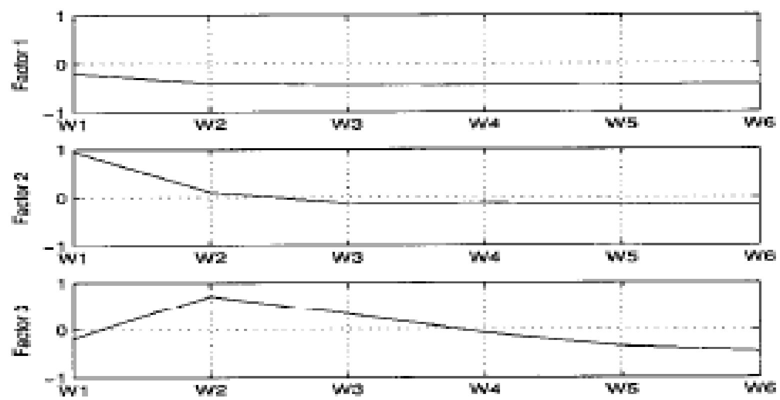

Fig. 8.3 First three factor loadings for weekly contracts.

In Fig. 8.3 we plotted factor loadings for the six weekly contracts. At first glance, it seems that the shapes of the first three factors correspond to shifting, bending and tilting. But knowing that there are low correlations between W_1 and the other weekly contracts, we go into more detail. Note that the first factor affects W_1 less than the other contracts. The second factor has an effect on W_1, and nearly zero effect on all the other contracts, while the third factor is a bending factor for all contracts except W_1, being basically not affected at all by the third factor. In the top panel of Ta-

Table 8.7 Factor analysis and correlation matrices for weekly and monthly contracts

#	%-explained	%-cumulative	Correlation matrix – weekly						
				W_1	W_2	W_3	W_4	W_5	W_6
1	76 %	76 %	W_1	1					
2	14 %	91 %	W_2	0.44	1				
3	6 %	97 %	W_3	0.31	0.87	1			
4	2 %	99 %	W_4	0.34	0.80	0.93	1		
5	1 %	100 %	W_5	0.33	0.75	0.87	0.95	1	
6	0 %	100 %	W_6	0.31	0.72	0.83	0.88	0.95	1

#	%-explained	%-cumulative	Correlation matrix – monthly						
				M_1	M_2	M_3	M_4	M_5	M_6
1	79.1 %	79.1 %	M_1	1					
2	8.7 %	87.8 %	M_2	0.80	1				
3	4.7 %	92.5 %	M_3	0.74	0.84	1			
4	3.7 %	96.2 %	M_4	0.69	0.75	0.83	1		
5	2.3 %	98.5 %	M_5	0.67	0.70	0.76	0.86	1	
6	1.5 %	100.0 %	M_6	0.63	0.66	0.72	0.73	0.84	1

Table 8.8 Factor analysis and correlation matrices for quarterly and yearly contracts

#	%-explained	%-cumulative	Correlation matrix - quarterly								
				Q_1	Q_2	Q_3	Q_4	Q_5	Q_6	Q_7	Q_8
1	58.7 %	58.7 %	Q_1	1							
2	12.5 %	71.2 %	Q_2	0.80	1						
3	8.1 %	79.3 %	Q_3	0.59	0.67	1					
4	6.9 %	86.1 %	Q_4	0.56	0.60	0.64	1				
5	5.6 %	91.7 %	Q_5	0.58	0.60	0.48	0.65	1			
6	4.0 %	95.7 %	Q_6	0.52	0.57	0.46	0.46	0.67	1		
7	2.3 %	98.0 %	Q_7	0.40	0.50	0.38	0.45	0.41	0.60	1	
8	2.0 %	100.0 %	Q_8	0.36	0.44	0.36	0.41	0.44	0.39	0.67	1

#	%-explained	%-cumulative	Correlation matrix - yearly			
				Y_1	Y_2	Y_3
1	72 %	72 %	Y_1	1		
2	18 %	90 %	Y_2	0.70	1	
3	10 %	100 %	Y_3	0.50	0.52	1

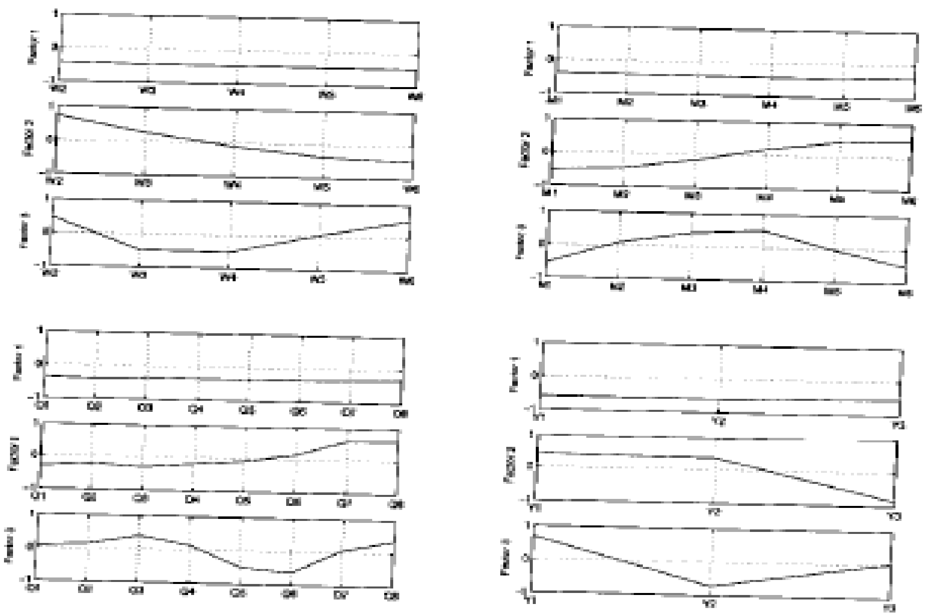

Fig. 8.4 First three factor loadings for the four different market segments: weekly contracts (upper left), monthly contracts (upper right), quarterly contracts (lower left), yearly contracts (lower right). The weekly contract with immediate delivery W_1 is excluded from the sample (see text for explanation).

In Fig. 8.4 we also plot the first three factor loadings for the other market segments; monthly contracts (upper left), quarterly contracts (lower left) and yearly contracts (lower right). For all market segments the first factor can be identified as a shifting factor, the second as a bending factor, and the third as a tilting factor. In the bottom panel of Table 8.7 and in Table 8.8 we report the

8.6. Estimating a Parametric Multi-Factor Market Model

In this section we estimate a parametric market model for each market segment. For the weekly contracts we exclude W_1 from the analysis. Further, we assume that the dynamics of the electricity futures price for each market segment can be described by (8.2) with $p = 3$, that is, a three-factor model. In addition, we assume that $A_c(u)$ is constant for each c. Such a specification corresponds to deterministic market prices of risk which might be different across the contracts. This is of course a simplification, but our main interest lies in the volatility dynamics, and not in the nature of the market price of risk.

We also assume that the factor volatilities $\Sigma_{k,c}$ can be factorised into a common seasonal function $\sigma_S(t)$ (with t representing the time of year) and a maturity dependent function $\widetilde{\sigma}_k(\tau_c^b - t)$. The latter function depends on the time to the start of the delivery period, $\tau_c^b - t$. Thus, the factor volatilities can be represented as

$$\Sigma_{k,c}(t) = \sigma_S(t)\widetilde{\sigma}_k(\tau_c^b - t)\,, \tag{8.6}$$

for $k = 1, 2, 3$. Note that $\Sigma_{k,c}$ implicitly depends on the delivery period as well, since we perform an empirical analysis for each market segment, where all contracts within a segment have the same length of delivery (week,

8.6. Estimating a Parametric Multi-Factor Market Model

month, quarter or year). The seasonal volatility is assumed to have the form

$$\sigma_S^2(t) = c_1 + \sum_{l=1}^{L} \left[c_{2l} \cos\left(\frac{2l\pi t}{365}\right) + c_{2l+1} \sin\left(\frac{2l\pi t}{365}\right)\right] . \qquad (8.7)$$

If $L = 1$, the seasonal variation is symmetric, meaning that a peak in seasonal volatility produces an equally low variance exactly six months later. By increasing L, we allow for asymmetric seasonal variance. For reasons that will become clearer later, we chose $L = 4$ in the estimation procedure. The maturity function is specified as

$$\widetilde{\sigma}_k(\tau_c^b - t) = \sigma_0 + \left(\sigma_1 + \sigma_2(\tau_c^b - t)\right) e^{-\kappa(\tau_c^b - t)} . \qquad (8.8)$$

This form is the same for all factors, and it is chosen for its simplicity and flexibility. The functional form in (8.8) is chosen to allow for bends and humps in the term structure of volatility. Other functions could be used instead. One alternative specification is

$$\widetilde{\sigma}_k(\tau_c^b - t) = \sigma_0 + \sigma_1 e^{-\kappa_1(\tau_c^b - t)} + \sigma_2 e^{-\kappa_2(\tau_c^b - t)} .$$

This specification is used by [De Jong, Driessen and Pelsser (2004)] to model the volatility term structure in fixed income securities. A humped term structure of volatility can be accomplished by allowing both positive and negative values for the parameters σ_0, σ_1 and σ_2. Our choice (8.8) is adopted from the popular Nelson-Siegel model for the yield curve in interest rate theory (see, for example, [James and Webber (2000)]).

8.6. Estimating a Parametric Multi-Factor Market Model

We can now compute daily normalised logreturns (with time measured in days) as

$$\overline{x}_{n,c} = (x_{n,c} - A_c) / \sigma_S(t_{n-1}) . \qquad (8.9)$$

The market price of risk A_c corresponds to the estimated average logreturn. From (8.2) and the assumed structure of the factor volatilities, the normalised logreturns become independent and centered normally distributed, with approximative factor volatilities given by $\widetilde{\sigma}_k(\tau_c^b - t_{n-1})$. An alternative version of this model can be expressed in terms of principal components in continuous time (see, for instance, Ch. 16 in [James and Webber (2000)]), yielding that the empirical factor volatilities can be written as $\sqrt{\lambda_k} p_k$, where λ_k are the eigenvalues and $p_k(\tau_c^b - t_n)$ are the eigenvectors of the covariance matrix of normalised logreturns.

We estimate the model in four steps.

8.6. Estimating a Parametric Multi-Factor Market Model

(1) Estimate the deterministic seasonal volatility $\sigma_S(t)$.
(2) Normalise electricity futures logreturns using (8.9).
(3) Compute principal components from the normalised returns.
(4) For each factor, estimate the parameters of the maturity function $\widetilde{\sigma}_k$ from the empirical factor volatilities obtained via the values achieved in the previous step.

The results from the empirical analysis are discussed in the following subsections.

8.6.1. Seasonal Volatility

To estimate the seasonal volatility σ_S, we first find the empirical mean and volatility for each contract c, denoted by $\widehat{m}_c$ and $\widehat{\sigma}_c$, respectively. Next, each price logreturn series are normalised (ignoring seasonality) using

$$\varepsilon_c(t) = (x_c(t) - \widehat{m}_c) / \widehat{\sigma}_c \,. \tag{8.10}$$

Obviously, $\varepsilon_c(t)$ will have an unconditional variance $\mathbb{E}\left[\varepsilon_c^2(t)\right] = 1$.

The deterministic variance is assumed to be constant across maturities. We therefore compute a series of average normalised squared returns

$$\overline{\varepsilon}_c^2(t) = \frac{1}{C} \sum_{c=1}^{C} \varepsilon_c^2(t) \,. \tag{8.11}$$

8.6.1. Seasonal Volatility

The parameters of σ_S are estimated by minimising, in a least squares sense, the difference between theoretical and empirical variance, that is, by finding the $\widehat{\mathbf{c}}_S$ which solves

$$\min_{\mathbf{c}_S} \frac{1}{T} \sum_{t=1}^{T} \left(\sigma_S^2(t) - \overline{\varepsilon}_c^2(t) \right)^2 ,$$

where T is the sample size and $\mathbf{c}_S = (c_1, c_2, ..., c_{2M+1})$ is the vector of parameters.

To account for asymmetric variance (high variance in December and January), we experimented with different values of L, and settled for $L = 4$ for all market segments as a reasonable choice. A lower value of L did not capture the asymmetry well, whereas a higher one did not give significantly better fit. This resulted in nine parameters to estimate for each data set. The parameter estimates are given in Table 8.9 and the variance functions are plotted in Fig. 8.5. We see that the asymmetric seasonality is clearly

8.6.1. Seasonal Volatility

Table 8.9 Fitted parameters for the seasonal volatility given in (8.7)

	Week	Month	Quarter	Year
c_1	1.0434	1.0420	1.0400	1.0380
c_2	0.8511	0.8007	0.6260	0.4686
c_3	−0.0786	−0.0430	−0.0098	0.2255
c_4	0.6640	0.5596	0.4474	0.3525
c_5	−0.1023	−0.0129	−0.0813	−0.0190
c_6	0.4270	0.4989	0.4541	0.4760
c_7	−0.1578	−0.0905	−0.1451	−0.0645
c_8	0.0863	0.1945	0.2396	0.2926
c_9	−0.0713	0.0563	0.0292	0.2109

8.6.2. Maturity Volatility

Next we estimate the parameters of the maturity volatilities $\widetilde{\sigma}_k(\tau_c^b - t)$.

First, we normalise logreturns using

$$\overline{x}_{n,c} = (x_{n,c} - \widehat{m}_c) / \widehat{\sigma}_S(t_{n-1}), \quad (8.12)$$

where $\widehat{\sigma}_S$ is the seasonal volatility obtained from the estimated parameter values $\widehat{\mathbf{c}}_S$ in Subsect. 8.6.1. Next, we estimate parameters by minimising,

8.6.2. Maturity Volatility

in a least squares sense, the difference between theoretical and empirical maturity volatilities obtained from a PCA of $\overline{x}_{n,c}$, that is,

$$\min_{\mathbf{m}_k} \frac{1}{C} \sum_{c=1}^{C} \left(\widetilde{\sigma}_k(\tau_c^b - t_n) - \sqrt{\widehat{\lambda}_k}\, \widehat{p}_k(\tau_c^b - t_n) \right)^2 .$$

Here, $\widehat{\lambda}_k$ are the empirical eigenvalues and $\widehat{p}_k(\tau_c^b - t_n)$ are the empirical eigenvectors. Furthermore, $\mathbf{m}_k = (\sigma_0, \sigma_1, \sigma_2, \kappa)$ is the vector of parameters for the maturity function.

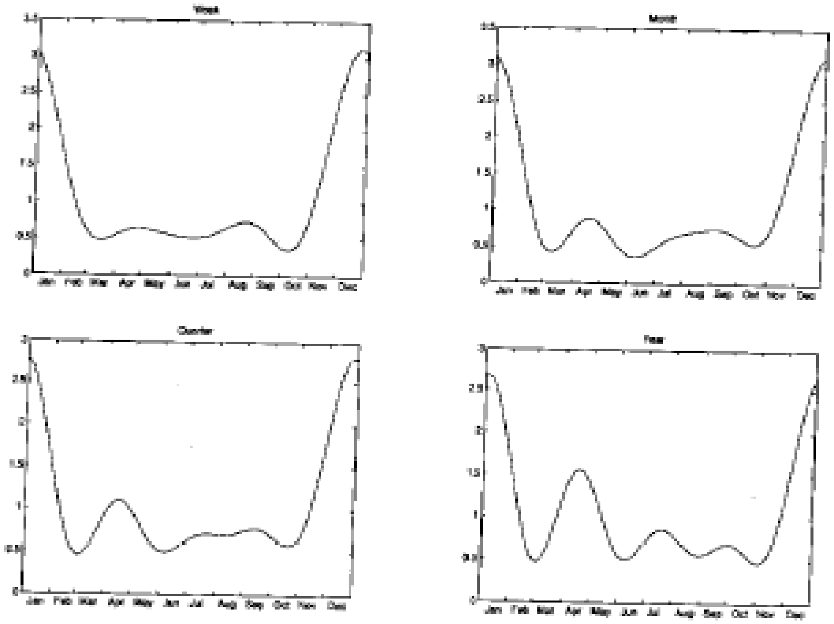

Fig. 8.5 Estimated seasonal functions for four market segments: weekly contracts (upper left), monthly contracts (upper right), quarterly contracts (lower left) and yearly contracts (lower right). W_1 is excluded from the weekly data set.

Table 8.10 Parameter estimates for a three-factor model. The function is the same for all three factors, and it is given in (8.8)

	σ_0	σ_1	σ_2	κ
Week				
Factor 1	−2.83	3.32	13.28	3.25
Factor 2	6.06	−5.56	−37.41	3.99
Factor 3	0.30	0.73	−57.85	33.23
Month				
Factor 1	−3.83	4.28	2.65	0.58
Factor 2	17.67	−17.88	4.63	−0.21
Factor 3	−26.85	26.67	21.27	0.68
Quarter				
Factor 1	0.02	0.35	0.00	0.50
Factor 2	13.47	−13.65	0.72	−0.04
Factor 3	−28.62	28.55	3.50	0.11
Year				
Factor 1	0.17	0.17	−0.07	1.00
Factor 2	0.21	−0.34	−1.34	1.87
Factor 3	−0.06	−0.09	1.13	1.94

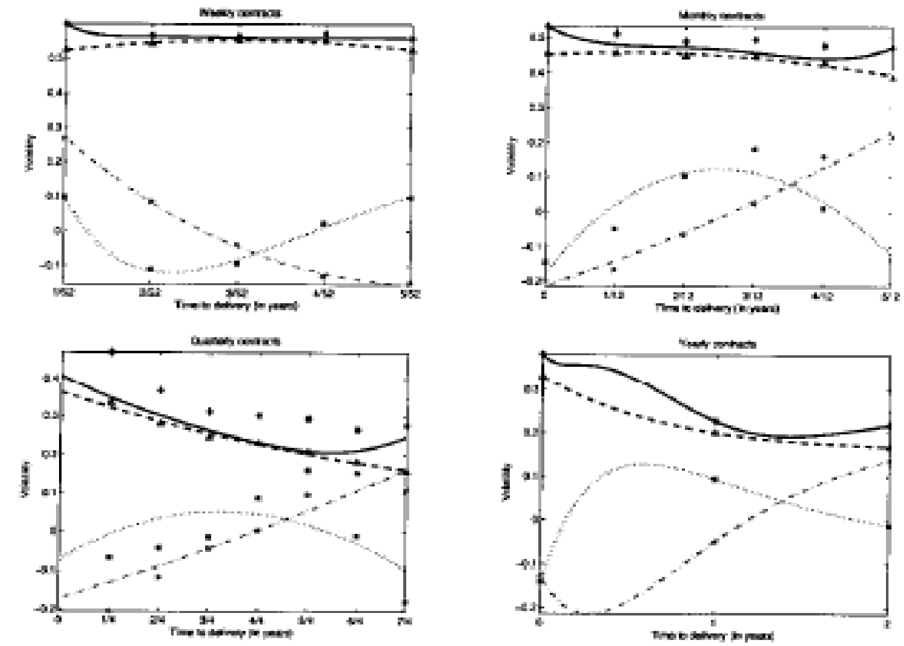

Fig. 8.6 Estimated maturity functions in a three-factor model: weekly contracts (upper left), monthly contracts (upper right), quarterly contracts (lower left) and yearly contracts (lower right). The first factor is the dashed line (triangles mark empirical volatility), the second factor is the dashed-dotted line (circles mark empirical volatility), the third factor is the dotted line (empirical volatility is marked by squares) and overall volatility is the solid line (empirical volatility is marked by diamonds). W_1 is excluded from the weekly data set.

8.7. Normalized Logreturns and Heavy Tails

In this Chapter we started out quite ambitiously with building a model that instantaneously incorporates all market segments (contracts with different delivery period length). Then we pursued a less ambitious task, by isolating each market segment and estimating multi-factor models for each segment (week, month, quarter or year). In this final section, we will not attempt to model joint dynamics at all, but instead focus on *single* contracts. That is, we pick a single contract with a specified time to delivery and a specified delivery period. We analyse the distributional properties of normalised logreturns, and show that they are far from being Gaussian. Clear signs of heavy tails are detected in all market segments, and we show that the NIG distribution models the stylised facts of the normalised logreturns in an excellent way. Our results are in line with the studies of [Frestad, Benth and Koekebakker (2007)].

8.7. Normalized Logreturns and Heavy Tails

Since we do not assume any particular parametric form for seasonality or maturity, we apply a different normalising routine than the one in the previous Section. For each contract in the market, we transform logreturns according to

$$x_{n,c} = (x_{n,c} - m_{S,c}) / \sigma_{S,c} ,$$

where $m_{S,c}$ and $\sigma_{S,c}$ are (seasonal) mean and standard deviation for $c = 1, ..., C$, respectively. We assume that $m_{S,c}$ and $\sigma_{S,c}$ are constant within each month. Therefore each data series is sorted according to month, normalised by subtracting the mean and dividing by standard deviation. Since we do this for each contract, both the seasonality and the maturity effects are removed from the data. In Table 8.11 the estimated parameters for the NIG distribution are presented for the six weekly contracts, starting with immediate delivery, and then delivery starting next week, two weeks later and so on. In addition, we include the estimates of the shape triangle parameters (defined in (2.32)), where in particular we observe that ξ is

Table 8.11 NIG parameters estimated for weekly contracts

Week	μ	α	β	δ	ξ	χ
W1	0.013	0.62	−0.014	0.62	0.85	−0.019
W2	−0.005	0.85	0.005	0.84	0.76	0.004
W3	0.008	0.83	−0.008	0.82	0.77	−0.007
W4	0.016	0.84	−0.016	0.82	0.77	−0.015
W5	0.022	0.79	−0.022	0.78	0.79	−0.022
W6	0.025	0.80	−0.026	0.78	0.78	−0.025

The estimates for the monthly contracts are found in Table 8.12. We find estimates being similar to the weekly ones, with maybe lower ξ parameter in the shape triangle. The parameter χ is close to zero for all six months, a reflection of hardly any skewness in the data. As a consequence of the data normalisation, the estimates of μ are close to zero. It seems to be a tendency of increasing α with the month (that is, time to maturity), whereas ξ seems to be slightly decreasing with the month. It is not easy to tell the maturity effect on δ, where estimates vary a bit up and down in an unclear pattern.

Table 8.12 NIG parameters estimated for monthly contracts

Month	μ	α	β	δ	ξ	χ
M1	−0.075	0.82	0.076	0.81	0.78	0.072
M2	0.037	0.78	0.037	0.78	0.79	−0.036
M3	0.019	0.89	−0.019	0.90	0.74	−0.016
M4	0.087	0.91	−0.089	0.89	0.74	−0.072
M5	−0.044	0.91	0.045	0.89	0.74	0.037
M6	0.066	0.95	−0.067	0.93	0.73	−0.061

The results for the quarterly contracts are presented in Table 8.13.

Table 8.13 NIG parameters estimated for quarterly contracts

Quarter	μ	α	β	δ	ξ	χ
Q1	−0.018	0.84	0.019	0.84	0.77	0.017
Q2	0.031	0.90	−0.031	0.88	0.75	−0.026
Q3	0.051	0.81	−0.053	0.79	0.78	−0.051
Q4	0.094	0.94	−0.096	0.91	0.73	−0.075
Q5	0.054	0.96	−0.055	0.95	0.72	0.041
Q6	0.035	0.89	−0.036	0.86	0.75	−0.030
Q7	0.088	0.88	−0.089	0.87	0.75	−0.076
Q8	0.022	0.97	−0.022	0.96	0.72	−0.016

for the yearly contracts there is a tendency towards a smaller ξ, in comparison to the contracts with quartely delivery period. It is hard to draw

Table 8.14 NIG parameters estimated for yearly contracts

Year	μ	α	β	δ	ξ	χ
Y1	0.035	0.95	−0.035	0.94	0.73	−0.027
Y2	0.105	1.10	−0.108	1.06	0.68	−0.067
Y3	0.100	0.99	−0.101	0.97	0.71	−0.073

In the panel plot depicted in Fig. 8.7, we show the fitted NIG distribution together with the empirical and standard normal. The chosen contracts are W2 (top left), M2 (top right), Q2 (bottom left) and Y2 (bottom right). We see that the tails are heavy, and that the NIG distribution is superior to the normal in fitting the data along length of the delivery period and time to delivery. The center of the empirical distribution is more peaky

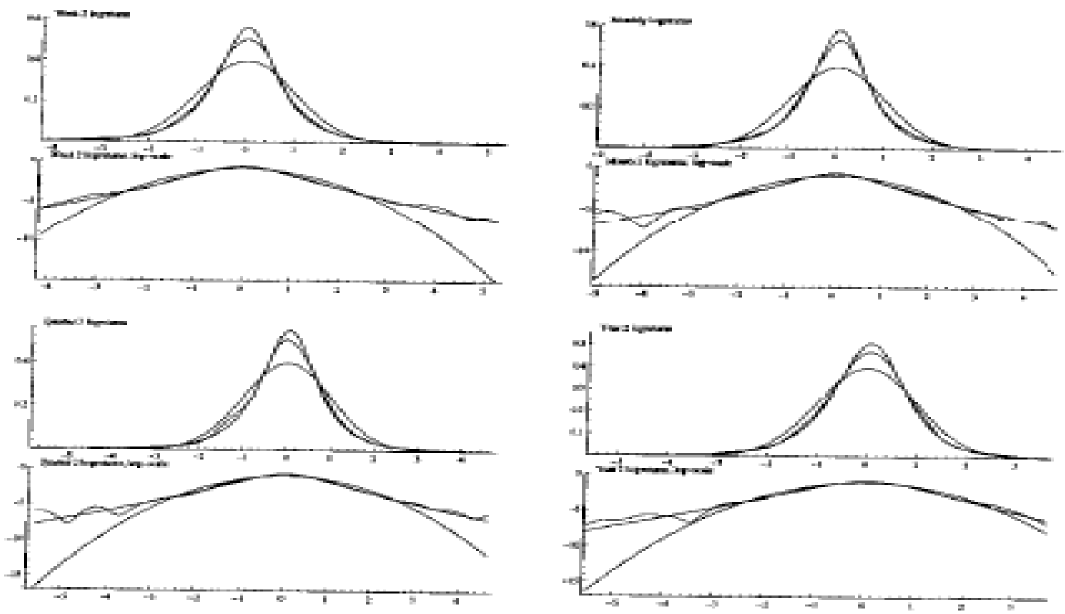

Fig. 8.7 Density plots of the empirical, NIG and standard normal distributions for normalised logreturns of electricity futures prices. The top row contains plots for the second week (left) and the second month (right). The bottom row contains plots of the second quarter (left) and the second year (right). For each electricity futures, we included the density plot on logarithmic scale to highlight the tails.

The Empirical Analysis Suggests a Market Model Including Jump Processes

$$F_c(t) = F_c(0)\exp(\Upsilon_c J(t)),$$

where J is a NIG Lévy process. The scaling factor Υ_c can be interpreted as a volatility structure. It may be hard to determine *one* set of parameters for J

The Empirical Analysis Suggests a Market Model Including Jump Processes

and scaling functions Υ_c which matches the estimated NIG distributions for the logreturns in question. Also, it is unlikely that we have a close to perfect dependency between the contracts, as discussed by [Frestad (2007b)]. The alternative is to model the electricity futures price dynamics by

$$F_c(t) = F_c(0)\exp(J_c(t)),$$

with one NIG Lévy process J_c per contract c. This makes it simple to estimate the characteristics of J_c directly from data. The next step then is of course to introduce a possible dependency structure on the contracts. If we choose a multivariate NIG, we would need to estimate the distribution parameters on all data, a difficult numerical task taking into account the dimension of the market and amount of data available. A copula structure is another possibility, as discussed in Subsect. 6.4.1.

8.8. Final Remarks

8.8 Final remarks

In this chapter we conducted an empirical investigation of market models at Nord Pool. This research is, at the time of writing this book, still in its infancy, and we will probably see a lot of new developments.

There seems to be low correlation between the very short end of the term structure (spot price), and the financial contracts trading at the exchange. If spot price models are to be used, they should be estimated on the traded contracts, and volatility must be estimated on implied volatility or empirical term structure volatility. But then a market model approach seems to be far superior to spot price models. However, a market model approach does not solve all our problems. Modelling all contracts simultaneously is a daunting task. Financial electricity contracts seem to behave more idiosyncratic than what we usually see in other commodity markets. Modelling different parts of the term structure individually may be a better idea than modelling all contracts simultaneously in the market. Of course this does not help us if we need a model for the whole market (which is the case, for example, when analysing portfolio Value at Risk for a trading department). From an empirical perspective, the models investigated in this Chapter are still a long way from being satisfactory term structure models.

Go, Canada, Go!

TM/MC

vancouver 2010