Ernst Eberlein and *Gerhard Stahl* analyse price series of 25 energy spot rates simultaneously using Lévy models. This model class allows the capture of stochastic behaviour of these financial instruments. The implications of this analysis will form the backbone of the German regulatory requirements for electricity risk

Both sides of the fence

★ The energy industry for a long time considered monopolistic structures the natural framework for guaranteeing electricity supply. Many countries began to question this basic assumption in the early 1990s. In Germany, the Energy Industry Act (*Energiewirtschaftsgesetz*) of April 29, 1998 completely liberalised the electricity market. Hence, this market has undergone a radical change. Although a base portion of the power demand is predictable, and supply can be secured by long-term contracts, a substantial part varies considerably. For this part, a day-ahead and spot market is more appropriate than long-term contracts. In summer 2000, two independent exchanges started up: the Leipzig Power Exchange (LPX) and the Frankfurt-based European Energy Exchange (EEX). The exchanges contributed considerably to the transparency of the electricity market in Europe and to fair prices for consumers. They merged in 2002.

Being exposed to the risk of adverse price movements is an inherent component of trading activities at exchanges. To manage this risk requires us to quantify it. Given the short history of exchange-traded electricity in Germany, our analysis will not be based on German data. The data used consists of 25 time series of daily prices recorded by the Nordic region's power exchange, Nord Pool, between January 1, 1996 and February 8, 2000. The prices are day-ahead prices, where the first 24 time series of length 1,500 each correspond to the 24 one-hour intervals of the day, recorded for 365 days a year. The last time series represents the arithmetic mean of the 24 one-hour contracts. Prices are given in Norwegian kroner per megawatt hour (MWh). Figures 1 and 2 show time series 8 and 17, which represent historical prices for the corresponding hours of the day. For example, series 8 corresponds to the eighth hour of the day – 7am to 8am. Compared with financial price histories, the striking difference is that there are spikes, which occur from time to time.

There are various reasons for spikes: for example, changing weather conditions such as a heat wave, or the outage of a major power plant. To understand why such an event can drive the price from 100 to 1,000 – note the biggest spike in time series 9, with a maximum price of 1,808 Nkr/MWh, as shown in table 1 – we must be aware that the marginal cost of generation of power increases sharply beyond a certain production level. Beyond that level, additional electricity can only be produced by adding generators with a high fixed cost.

Comparing the 24 time series, we can see that the spikes are less pronounced for night hours than for the time between 8am and 8pm. Jumps – although not of this size – also occur in financial time series and led researchers to introduce jump-diffusion models (see, for example, Bakshi, Cao and Chen (1997) for an advanced version) and more recently the more powerful pure jump models driven by Lévy processes (see Eberlein (2001)). The difference is that electricity prices usually jump back to the former level within a short time interval. This happens as soon as the heat wave

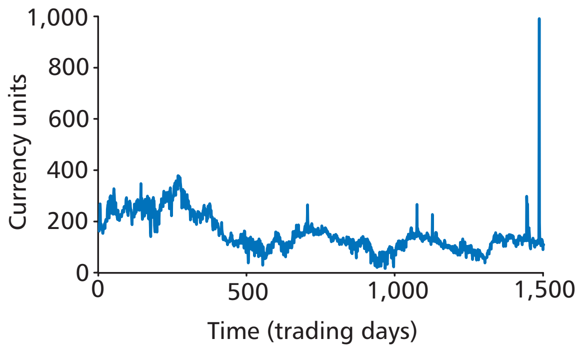

F1. Time series 8

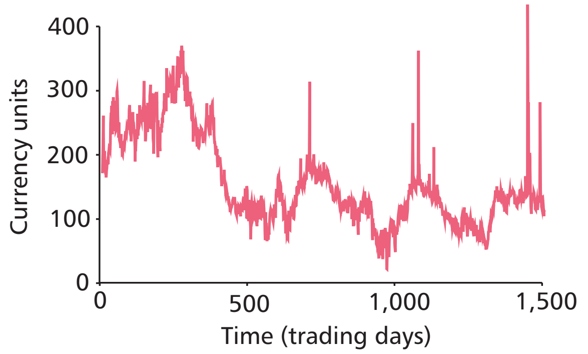

F2. Time series 17

is over, the outage comes to an end or additional electricity is available from elsewhere.

How volatile are electricity spot rates?

It is not only the spike phenomenon that contributes to electricity contract risk. Substantial price changes from one day to the next or within a few days happen with higher frequency and are greater with regard to energy than time series from finance. Given a price series $(S_t)_{t\geq 0}$, we compare – as in finance – log returns

$$X_t = ln\, S_t - ln\, S_{t-1}.$$

X_t reflects the geometric – or multiplicative – character of price changes. Note that electricity time series record prices for seven days a week – not for the usual five trading days. Hence, to derive annualised volatilities from the standard deviation of daily log returns, we have to multiply by the square root of 365. As table 1 shows, the annualised historical volatilities from the 24 series are between 120% and 376%. This is a full order of magnitude higher than the level of volatility typical for financial time series or even commodity prices.

Estimating the long-term volatility from the full-length time series averages out the volatility of the volatilities. Since the risk of a contract is determined by the conditional volatility, we derive historical volatilities from shorter time intervals. This will reveal how volatile volatility is. Figure 3 shows 30 volatilities that have been estimated from non-overlapping consecutive windows of length 50 each. The underlying data is series number 9. The average value we get is 270%, with a minimum of 123% and a maximum of more than 800%.

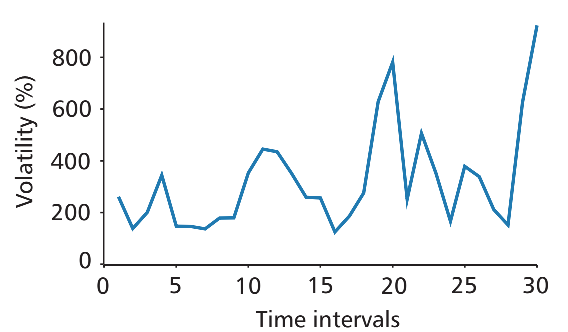

F3. Annualised volatility

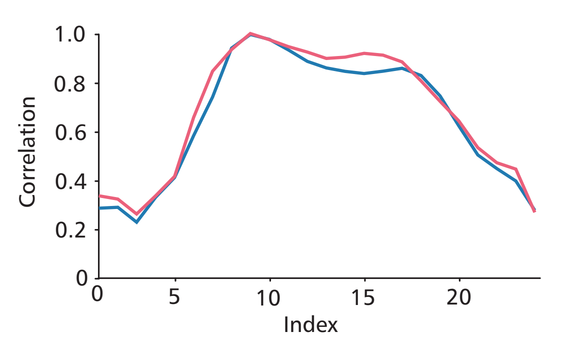

F4. Correlations of full data set and 1998 only

Since the volatility of electricity prices is rather volatile, it is of interest to examine how stable correlations are. From the application of internal models for measuring market risk, we know that changing volatility scenarios may destroy the correlation structure. This turns out not to be the case for the data analysed here. The correlation structure remains the same, regardless of changes of volatility. As an example, we show in figure 4 the correlation of contract number 9 with all the other contracts. One of the two lines refers to the full time interval from 1996 to 2000, the other one to 1998 only, when there were much higher volatility levels. The correlations are almost perfectly correlated. The corresponding coefficient is 0.98.

Figure 5 shows the autocorrelations derived from daily returns – there is a clear weekly pattern. As figure 6 shows, the autocorrelations of squared returns are much less pronounced. This is another difference from the statistical behaviour of returns from stock prices. The autocorrelations of the latter ones do not have cyclical patterns, and autocorrelations of squared returns persist up to a certain lag order, whereas those of the daily returns themselves are typically negligible.

Lucia and Schwartz (2002) also observed the high positive autocorrelations at time lags that are multiples of seven. They draw the conclusion that the increments of the log-prices from one day to the next can be used to predict to a significant degree the increment for the same couple of consecutive days up to several weeks into the future. As figures 7 and 8 show, the positive autocorrelation is – in our data – a consequence of the fact that contracts for the weekend days are traded on Friday together with those for Monday. In other words, prices for Saturday, Sunday and Monday are built on the basis of the same information, namely that available on Friday. Weekend prices do not contain significant new information. If we remove weekend prices from the time series and consider only a five-day week – that is, the five trading days – the autocorrelation, and thus the predictable part, disappears completely. This holds for returns as well as for squared returns.

We just highlight one aspect of energy markets that is related to model risk, to show that regulatory standardised methods cannot be the first choice for determining regulatory capital. This is that over-the-counter contracts of limited liquidity are the core instruments of the market. As a result, capital reserves are common, to account for the model risk caused by the uncertainty over the implied volatility.

Under the supplementary law of the German Banking Act (GBA), financial services firms, as well as banks, must fulfil the regulatory requirements. Even smaller broker firms fall under the GBA, if they actively trade. It is a challenge to create prudential rules that keep the regulatory burden within an acceptable range, especially for smaller market participants. We show in the final section of this paper that the generalised hyperbolic model for measuring electricity risk fulfils this practical constraint and at the same time provides very accurate forecasts for electricity time series.

What is an adequate regulatory capital charge?

For an open position in commodities, the standard regulatory capital charge is 15% (section 16, paragraph 4, sentence 1, principle 1 of the German implementation of the capital adequacy directive (Cad)). According to our analysis above, here we cannot apply such a standard approach characterised by a fixed percentage, since daily volatilities are already close to 15%.

Before we introduce our approach, we take a closer look at how power exchange clearing houses treat the problem of margin calculation. Nord Pool, for example, uses a standard portfolio analysis of risk (Span), as developed by the Chicago Mercantile Exchange. Span calculates the maximum one-day loss potential for a portfolio consisting of futures, forwards and options on the basis of, say, a finite number of scenarios. The calculation involves scenarios for anticipated price movements, theoretical values of instruments and a fixed correlation matrix that takes into account how prices for various delivery periods move together. From the perspective of margin calculations, Span-like structures seem to be appealing blueprints for the formulation of a regulatory framework. In addition to the obvious closeness of these procedures to widespread market standards, Span is an archetype for coherent risk measures (see Artzner et al (1999)).

However, a Span-based regulatory regime would face three drawbacks. The first is the lengthy description of the rules for calculating margins – this is a problem, given that the energy market is just one of many financial markets covered by Cad. Second, energy contracts mirror regional specifics related to the energy production process. Hence, setting a universal standard for volatility and time spreads is difficult. Third, non-banks – the bulk of firms that fall under the GBA regulation – should be motivated to build up state-of-the art risk management systems in a volatile and deregulated market environment.

The analysis above shows that the changes in volatility are the dominant factor in quantifying electricity risk. Hence we propose a portfolio approach that relies on volatility as a risk measure and is based on historical simulation as a standardised method. For calculating regulatory capital, the resulting volatility must be multiplied by a safety factor. Given the complexity of dealing with energy contracts, we must – in addition to carrying out the simulations – run appropriate stress tests, install an adequate limit system and set an appropriate level of qualitative standards.

Series vol	Min price	Max price	Daily vol	Annual
1	20.24	360.89	0.0735	1.4038
2	19.23	352.40	0.0959	1.8314
3	18.49	488.84	0.1238	2.3659
4	18.35	349.86	0.1111	2.1227
5	18.05	356.83	0.1137	2.1722
6	17.97	369.11	0.1223	2.3356
7	18.11	380.34	0.1460	2.7887
8	18.31	993.03	0.1776	3.3924
9	19.33	1,808.66	0.1969	3.7613
10	20.98	1,359.39	0.1646	3.1449
11	22.25	792.84	0.1338	2.5555
12	23.02	380.22	0.1110	2.1202
13	22.83	375.43	0.1058	2.0218
14	22.09	371.10	0.1077	2.0572
15	21.85	373.83	0.1090	2.0834
16	21.64	369.43	0.1061	2.0278
17	22.12	434.20	0.1063	2.0305
18	22.29	983.81	0.1243	2.3752
19	22.64	374.86	0.0901	1.7216
20	22.64	380.44	0.0756	1.4436
21	21.70	380.70	0.0729	1.3918
22	22.45	377.42	0.0710	1.3570
23	22.10	373.30	0.0653	1.2480
24	20.52	361.06	0.0648	1.2386
25	21.27	387.78	0.0952	1.8184

Comparison of daily Nord Pool price change data

How do we calculate the capital cushion?

So how do we calculate the capital cushion? Given a portfolio consisting of linear and non-linear instruments, we consider the current positions. We calculate historical portfolio values $S_{t-N+1}, \ldots,$

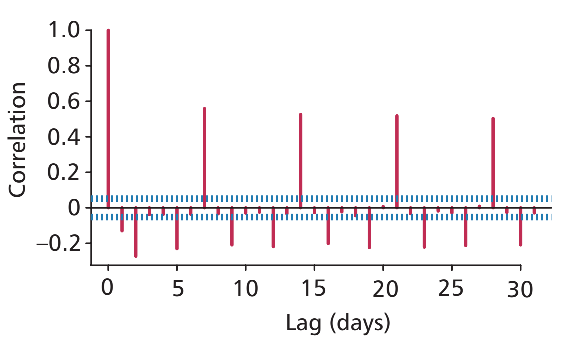

F5. Autocorrelation function of daily returns of series 8

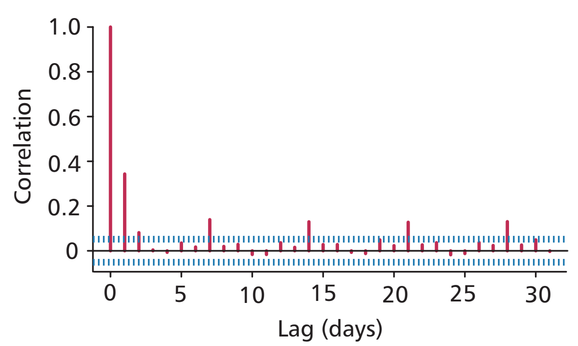

F6. Autocorrelation function of squared returns of series number 8

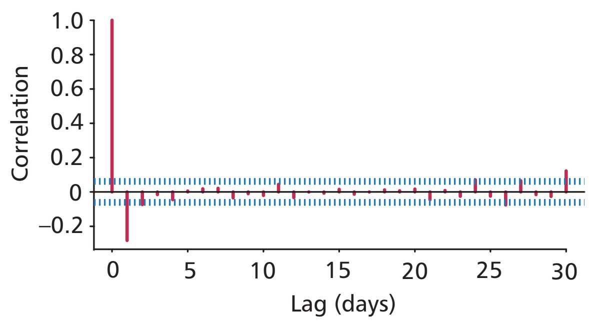

F7. Autocorrelation function of daily returns of series no. 8, weekend prices removed

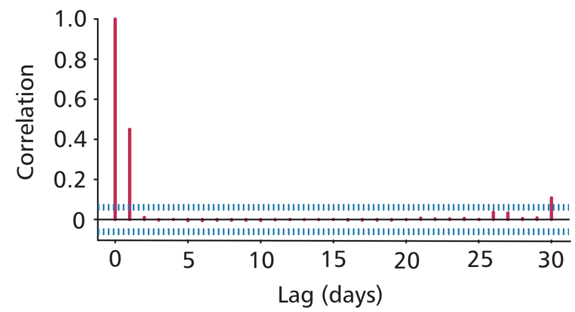

F8. Autocorrelation function of squared returns of series no. 8, weekend prices removed

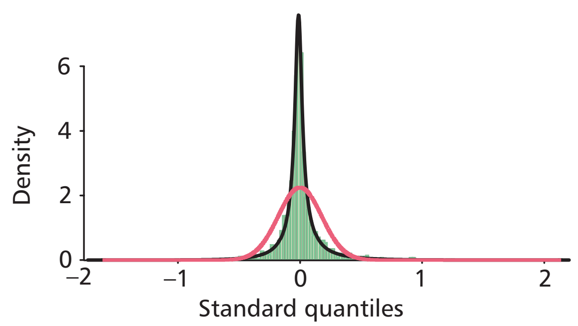

F9. Empirical density from series 8, fitted generalised hyperbolic and fitted normal density

S_{t-1}, S_t ($N=51$) on the basis of (theoretical) prices of the instruments during the past 51 trading days. We derive the current risk from this sequence of portfolio values. The relatively short price history of 51 trading days – which results in 50 log returns – guarantees that the resulting standard deviation reflects the current volatility. For the (absolute) risk – that is, the monetary amount one can lose within a certain time interval Δ – we have to consider the change in the portfolio's value from its present value S_t to the future value $S_{t+\Delta}$ at time $t+\Delta$. Δ will typically have the value 1. If we use daily price data, this means we look one day ahead.

Here we consider two ways to derive the portfolio standard deviation. The first is very straightforward. For a given sequence of historic portfolio values, we apply the standard estimator for the empirical variance,

$$\sigma_0^2 = \frac{1}{N-1}\sum_{i=0}^{N-1}(S_{t-i}-\bar{S}_t)^2$$

where

$$\bar{S}_t = \frac{1}{N}\sum_{i=0}^{N-1} S_{t-i}$$

denotes the mean. The risk is given by the empirical standard deviation σ_0, the square root of the variance σ_0^2.

The second approach is more sophisticated and will result in more accurate risk values. Given the historical portfolio values, we derive the log returns as described above. We then fit a parametric distribution to the empirical return distribution. Generalised hyperbolic distributions are flexible enough to fit the empirical distributions that occur here. Generalised hyperbolic distributions contain many of the classical statistical distributions as sub-classes or as limiting cases. Examples are Student t, hyperbolic, normal inverse Gaussian, normal, Cauchy and variance gamma distributions. Densities of generalised hyperbolic distributions are given by

$$d_{GH}(x;\lambda,\alpha,\beta,\delta,\mu) = a(\lambda,\alpha,\beta,\delta)\left(\delta^2+(x-\mu)^2\right)^{(\lambda-\frac{1}{2})/2} \times K_{\lambda-\frac{1}{2}}\left(\alpha\sqrt{\delta^2+(x-\mu)^2}\right)\exp\left(\beta(x-\mu)\right),$$

where

$$a(\lambda,\alpha,\beta,\delta) = \frac{\left(\alpha^2-\beta^2\right)^{\lambda/2}}{\sqrt{2\pi}\alpha^{\lambda-\frac{1}{2}}\delta^{\lambda}K_{\lambda}\left(\delta\sqrt{\alpha^2-\beta^2}\right)}$$

is a normalising constant and K_ν denotes the modified Bessel function of the third kind with index ν. Each distribution is determined by five parameters: $\alpha>0$ and the class parameter $\lambda\in\mathbb{R}$ determine the shape, β with $0\leq|\beta|<\alpha$ the skewness, $\mu\in\mathbb{R}$ the location, and $\delta>0$ is a scaling parameter comparable to σ in the normal distribution. There are several alternative parameterisations with intuitive interpretations.

Figure 9 shows in green the empirical density resulting from series number 8 and the fitted generalised hyperbolic density. The fitted normal density is given in red to allow comparison. To show that the generalised hyperbolic fit is also quite good in the tails, we plot the density on a logarithmic scale (figure 10) and show the corresponding qq-plot (figure 11). The normal qq-plot (figure 12) underlines how far the empirical distribution is from the normal distribution.

Since our data analysis is based on log returns, the natural model for the price process – or portfolio value – is

$$S_t = S_0\exp\left(X_t\right)$$

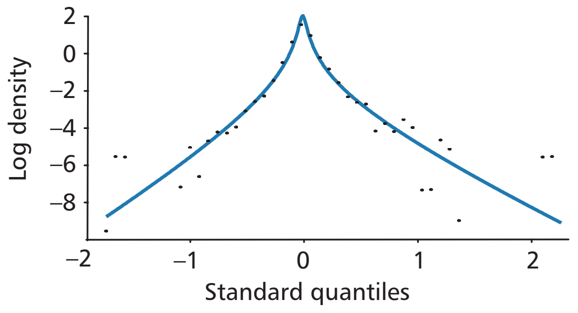

F10. Generalised hyperbolic fitted to empirical density no. 8 on log scale

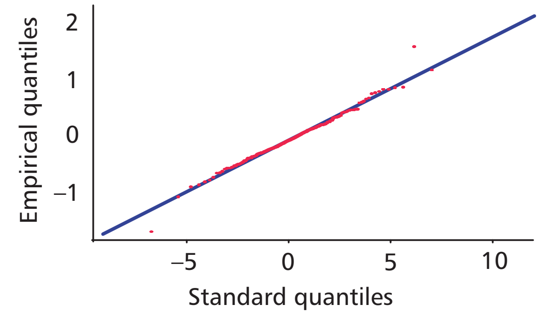

F11. Qq-plot of empirical (no. 8) versus generalised hyperbolic distribution

where $(X_t)_{t\geq 0}$ is the Lévy process generated by the generalised hyperbolic distribution that was fitted to the data. In other words, $(X_t)_{t\geq 0}$ is the unique process with stationary and independent increments such that increments along time intervals of length 1, $X_t - X_{t-1}$, are distributed according to the fitted generalised hyperbolic distribution. We should note that the distribution of the log returns derived from this model is exactly the distribution that was fitted to the log returns from the price data. If the analysis is based on daily data, one day in real time corresponds to time length 1 in the model above.

Based on this dynamic model, we could derive risk numbers for any time horizon ahead and not only for one-day periods. This is because we know the distributions for any time horizon – intra-day as well as many days ahead. Hence the model is highly consistent with respect to changes in the time scale. If the fitted distribution is the Gaussian law – that is, the generated Lévy process is usual Brownian motion – we recover the well known 'square root of time' rule. For simplicity, let us assume $\Delta = 1$, then given the present portfolio value S_t, the risk one day ahead will be given by the standard deviation of $S_{t+1} - S_t$, as derived from the model above. For the variance, we get

$$\begin{aligned}\mathrm{Var}\left(S_{t+1} - S_t\right) &= S_t^2 \,\mathrm{Var}\left(\exp\left(X_{t+1} - X_t\right)\right)\\ &= S_t^2 \,\mathrm{Var}\left(\exp\left(X_1\right)\right)\\ &= S_t^2\left(E\left[\exp\left(2X_1\right)\right] - \left(E\left[\exp\left(X_1\right)\right]\right)^2\right)\end{aligned}$$

The second equality follows from the stationarity of the increments. By construction, the distribution of X_1 is the fitted generalised hyperbolic distribution. We now introduce the moment-generating function

$$M_{\mathrm{GH}}(u) = \int_{-\infty}^{+\infty} \exp(ux)\, d_{\mathrm{GH}}(x)\,dx$$

defined by a generalised hyperbolic distribution. We can easily derive it as

$$M_{\mathrm{GH}}(u) = e^{\mu u}\left(\frac{\alpha^2 - \beta^2}{\alpha^2 - (\beta + u)^2}\right)^{\lambda/2} \frac{K_\lambda\left(\delta\sqrt{\alpha^2 - (\beta+u)^2}\right)}{K_\lambda\left(\delta\sqrt{\alpha^2 - \beta^2}\right)}$$

$M_{\mathrm{GH}}(u)$ exists for $|\beta + u| < \alpha$. Now we can write

$$\mathrm{Var}\left(S_{t+1} - S_t\right) = S_t^2\left(M_{\mathrm{GH}}(2) - M_{\mathrm{GH}}^2(1)\right).$$

Therefore, the risk is given by

$$\sigma\left(S_{t+1} - S_t\right) = S_t\left(M_{\mathrm{GH}}(2) - M_{\mathrm{GH}}^2(1)\right)^{1/2}$$

We should note that the moment-generating function $M_{X_\Delta}(u)$ of the distribution of an increment X_Δ along a time interval of length Δ is given by $M_{X_\Delta}(u) = (M_{X_1}(u))^\Delta$. As a result, we get for any time horizon Δ

$$\sigma\left(S_{t+\Delta} - S_t\right) = S_t\left(\left(M_{\mathrm{GH}}(2)\right)^\Delta - \left(M_{\mathrm{GH}}(1)\right)^{2\Delta}\right)^{1/2}$$

Figure 13 shows the result of the two estimation procedures for time series number 7. The blue line results from the classical empirical variance estimator. The red line shows the corresponding numbers we get from fitting the generalised hyperbolic model. We can see that the hyperbolic model can monitor risk more sensitively than the classical empirical variance estimator. Since the model has been fully calibrated, it is easy to derive the value-at-risk (VaR) as an alternative risk measure and to perform a backtest according to the Basel rules for portfolios in finance.

Figure 14 shows the result of the backtest for series number 9, starting at time point 51. We used the first 50 data points to estimate the initial parameters. The red line is the estimated 99% VaR number one day ahead. The blue values show the actual loss observed in case a loss occurred. There are 21 values – marked along the top – where the actual loss exceeds the forecast. Compared with the theoretical number of 14.5 exceedances in a period of 1,450 days, this is a good result. We can further improve the VaR estimation by using standardised data that takes stochastic volatility into account (see Eberlein, Kallsen, Kristen (2003)). To illustrate, we give in figure 15 the backtest results for series number 17, which shows a perfect result, with 11 exceedances over 1,000 days.

The core question that industry practitioners as well as regulators pose to statistical forecast models is: does the model provide forecasts of adequate quality? The answer to this question gives a final judgement on both the validity and the usefulness of the

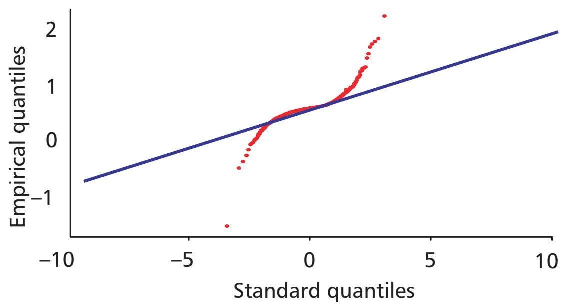

F12. Qq-plot of empirical (no. 8) versus normal distribution

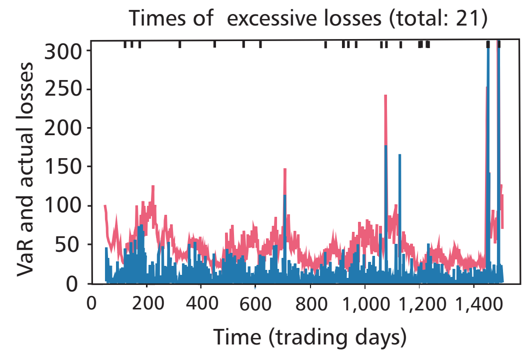

F14. Backtest of value-at-risk for series number 9

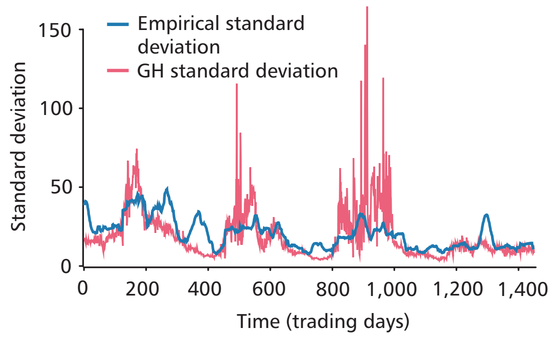

F13. Classical and model-based estimates of the volatility

model. Figures 13–15 provide the empirical evidence for a positive answer. They depict different aspects of the predictive quality of the hyperbolic model.

More specifically, figure 13 highlights the different dynamics of volatility estimates either based on the hyperbolic model or on a simple, non-parametric one. Figures 14 and 15 show that the forecast quality of the hyperbolic model – measured in the VaR metric – is very high, even in the non-standard market of electricity. Hence, the hyperbolic model contributes significantly to the forecast quality of a statistical model and, as a consequence, to the usefulness of a risk management system.

Conclusion

Based on the empirical findings from energy spot prices, we have discussed several aspects of measuring risk in the energy market. Our analysis shows that volatility is the relevant risk parameter. To take portfolio effects into account, we propose to calculate volatility by means of a historical simulation-based approach. The volatility may be determined either by a classical estimator or a more sophisticated method.

Compared with standardised methods based on crude statistical assumptions, the proposed approach has several advantages. It is easy to understand and implement, sets minimum standards for measuring risk, is applicable for internal risk management purposes, leaves space for new instruments, allows for backtesting, is closely related to Span and opens the door for advanced methods. And this balanced approach meets the needs of both the market participants and the regulators. This regulatory framework puts the focus on risk and not on regulation. ER

Ernst Eberlein is a professor in the Department of Mathematical Stochastics and the Freiburg Centre for Data Analysis and Modelling at the University of Freiburg, Germany.
Gerhard Stahl is head of group risk modelling at the Federal Financial Supervisory Authority in Bonn.
email: eberlein@stochastik.uni-freiburg.de; gerhard.stahl@bafin.de

The authors would like to thank EA von Hammerstein, J Kallsen, and F Özkan for their support in preparing the graphs. Gerhard Stahl's views are not necessarily those of the Federal Financial Supervisory Authority.

References

Artzner, P, F Delbaen, J-M Eber, and D Heath: Coherent measures of risk, *Mathematical Finance*, 9, pages 203–228, 1999

Bakshi, G, C Cao, and Z Chen: Empirical performance of alternative option pricing models, *Journal of Finance* 52, pages 2003–2049, 1997

Eberlein, E: Application of generalised hyperbolic Lévy motions to finance, in Barndorff-Nielsen O E, T Mikosch, and S Resnick (eds), *Lévy Processes: Theory and Applications*, Birkhäuser, pages 319–337, 2001

Eberlein, E, J Kallsen, and J Kristen: Risk management based on stochastic volatility, *The Journal of Risk* 5, pages 19–44, 2003

Geman, H and O Vasicék: Plugging into electricity, *Risk*, August 2001, pages 93–97

Lucia, J and E Schwartz: Electricity prices and power derivatives: evidence from the Nordic Power Exchange, *Review of Derivatives Research* 5, pages 5–50, 2002

Nordic electricity clearing technical document: *Security calculation and settlement of financial power contracts*, Nord Pool, Nordic electricity clearing, 1999