

Ronald Huisman (The Netherlands), Mehtap Kilic (The Netherlands)

Extreme changes in prices of electricity futures

Abstract

The purpose of this paper is to analyze the occurrence of extreme price change in power delivery forward and futures contracts. The results indicate that the distribution of price changes are significantly fatter tailed than a normal distribution function and the authors discuss that risk managers in the power industry can obtain better insight in the amount of risk their companies face by applying extreme value theory.

Keywords: extreme value theory, electricity futures prices, energy derivatives.

Introduction

In this paper we focus on the occurrence of extreme price changes in power delivery forward or futures contracts. These contracts are traded on exchanges worldwide and energy companies use these contracts to hedge themselves against market risk. For instance, an energy company that needs to deliver power to clients in the year 2011, can buy a power futures contract somewhere in 2010 and fixate the price against which it will purchase power for its clients. We refer to these contracts as both power delivery forward and futures contracts in the remainder of this paper. The pricing of these contracts is not as straightforward as pricing futures contracts on stocks, for instance. As discussed by Fama and French (1987) and many others traders use the availability of storage capacity to value futures contracts. A trader that sells a futures contract can make his position risk free by purchasing the commodity on the spot market. As a result, the futures price should reflect the spot price of the commodity plus interest forgone, storage costs, and a convenience yield that reflects the value that can be derived out of having the commodity physically. Power is not yet economically storable and, as a consequence, the power futures prices reflect expectations and risk premiums (see Fama and French, 1987; Lucia and Schwartz, 2002; Eydeland and Wolyniec, 2003; and Huisman, 2009 among others). Power futures prices do not necessarily depend on the spot price of power and therefore their price dynamics should be modeled as a stand-alone process.

Risk managers in the power sector use these contracts to actively manage market risk. For instance, consider a power company that has agreed to deliver power against a fixed price to clients in 2011. When the company will buy the power during the delivery period 2011 in the spot market, it faces the risk that the average price, paid in the spot market, is higher than what is agreed with the clients. By purchasing a power forward delivery contract, the risk manager can fixate the price against which the company will purchase power in the market in 2011 and by doing

so price risk is reduced. However, the timing of when to purchase these forward contracts is a difficult decision. One can buy such a contract today or perhaps tomorrow when prices might be lower. It depends on the risk of a potential price increase that might occur between today and tomorrow, whether a company wants to purchase today or wait.

In this paper, we focus on this price risk. We examine to what extent changes in power delivery futures prices can be modeled using a normal distribution function or whether another method should be applied. We apply extreme value theory to assess the level of tail-fatness, i.e., the frequency with which large price movements occur, such that we can observe whether these price changes can be modeled using a normal distribution or not. Bernhardt et al. (2008) apply extreme value theory to estimate high quantiles dynamically for day-ahead electricity prices in Singapore. Byström (2005) applies extreme value theory to model electricity prices on the NordPool market, making quantile (*VaR*) forecasts allowing both for fat tails and time-varying volatility. Both papers find strong support for the existence of fat tails in day-ahead prices and for the superior quantile estimates that extreme value theory produces. Ren and Giles (2007) present an extreme value analysis of daily Canadian crude oil prices and find strong support for fat tails. Although the amount of tail-fatness is examined in oil markets and for day-ahead power prices, it has never been examined for changes in the price of power futures delivery contracts. This is the goal set in this paper.

We analyze the occurrence of extreme price change in power delivery forward and futures contracts. Our results indicate that the distribution of price changes are significantly fatter tailed than a normal distribution function and we discuss that risk managers in the power industry can obtain better insight in the amount of risk their companies face by applying extreme value theory.

1. Extreme value theory

Extreme value theory is a field within statistics that deals with the frequency with which extreme observations occur. We follow Hull (2007) and Huisman

(2009) in discussing extreme value theory and we start with the key result in extreme value theory found by Gnedenko (1943). Suppose $F(v)$ is the cdf of a variable v : $F(v) = Pr\{V \leq v\}$. As extreme value theory focuses on the structure of the tail, consider a value u that is a value of v somewhere in the right tail of the distribution function of v . The probability that v lies between u and $u + y$ equals $F(u + y) - F(u)$ for $y > 0$. Define $F_u(y)$ as the probability that v lies between u and $u + y$ conditional on $v > u$. Thus, $F_u(y) = Pr\{u \leq v \leq u + y | v > u\}$. Gnedenko (1943) shows that for large values for u , $F_u(y)$ converges to the generalized Pareto distribution for many probability distribution functions $F(\cdot)$. The generalized Pareto distribution $G_{\alpha,\beta}(y)$ is:

$$G_{\alpha,\beta}(y) = 1 - \left(1 + \frac{y}{\alpha\beta}\right)^{-\alpha} \quad (1)$$

Hull (2007) continues reasoning that the probability that $v > u + y$ given that $v > u$, $1 - F_u(y)$, then equals $1 - G_{\alpha,\beta}(y)$. Furthermore, the probability that $v > u$ is $1 - F(u)$. The unconditional probability that v exceeds a value x , $Pr\{v > x\}$ equals:

$$Pr\{v > x\} = Pr\{v > u\}Pr\{v > u + (x - u) | v > u\} = (1 - F(u))(1 - G_{\alpha,\beta}(x - u)) \quad (2)$$

This result of Gnedenko (1943) implies that many distribution functions follow a generalized Pareto distribution in the tails. When we approximate $1 - F(u)$ by its empirical counterpart $\frac{n_u}{n}$, where n is the number of observations in the sample and n_u is the number of observations that exceed the value u , equation (2) can be written as:

$$Pr\{v > x\} = \frac{n_u}{n} (1 - G_{\alpha,\beta}(x - u)) = \frac{n_u}{n} \left(1 + \frac{x - u}{\alpha\beta}\right)^{-\alpha} \quad (3)$$

If we now set $u = \beta\alpha$ and $K = \frac{n_u}{n} \left(\frac{1}{\alpha\beta}\right)^{-\alpha}$, then we obtain what is called the power law:

$$Pr\{X > x\} = Kx^{-\alpha} \quad (4)$$

We have now formalized the main ideas within extreme value theory. Beyond a certain threshold, fat-tailed distribution functions exhibit power decay. The speed of decay is measured by α in equation (4). This parameter is called the tail-index. The bigger α is, i.e., the steeper the decay, the thinner the tails become and vice versa. The normal distribution, being

a thin tailed distribution function, exhibits exponential decay, which is obtained when $\alpha \rightarrow \infty$.

The goal of this paper is to examine the tail structure of log-price changes of electricity forward prices. To do so, we estimate the tail-index α using the procedure outlined in the following paragraph.

1.1. Estimating the tail-index α . Let's focus on estimating the tail-index of the right tail of the distribution function. Let k be the number of tail observations that we include in the estimation, such as the k highest or lowest returns. Let x_i be the i^{th} order statistic, such that $x_i \geq x_{i-1}$. Hill (1975) shows that the estimate of the inverse of the tail-index, $\gamma = \frac{1}{\alpha}$, for k tail-observations equals:

$$\gamma_k = \frac{1}{k} \sum_{j=1}^k \ln(x_{n-j+1}) - \ln(x_{n-k}), \quad (5)$$

where n is the total number of observations in the entire sample. How to select k , the number of tail observations to include in the estimate? Initially, researchers calculated estimates for α for different values for k and then state their conclusions in terms of the average result. Others tried to approximate the optimal k by assuming that the data came from some distribution function and then select k that would lead to the best results in a simulation study. Examples of these approaches are (among others) Jansen and de Vries (1991), Koedijk and Kool (1994), and Kearns and Pagan (1997) who estimated the tail-index for the returns distributions of exchange rates and stocks. The results indicated fat-tails, but one is left with the uncomfortable feeling that the tail-index estimates suffer from a bias in choosing k . One way to limit the influence of this bias is proposed by Huisman et al. (2001), a method that we apply in this paper. It is an extension of the Hill (1975) estimator. Huisman et al. (2001) observe that the expected value for the estimate γ_k equals $\frac{1}{\alpha}$ plus some function $f()$ that depends on k :

$$E(\gamma_k) = \frac{1}{\alpha} + f(k) \quad (6)$$

Huisman et al. (2001) show for several distribution functions, among them the Student-t, that the function $f()$ is almost linear. They formulate the following regression equation:

$$\gamma_k = \beta_0 + \beta_1 k + \varepsilon_k, \quad (7)$$

and the estimate for β_0 is then an accurate estimate for $\gamma = \frac{1}{\alpha}$. Basically, the Huisman et al. (2001)

estimator combines information from different choices of k to reduce the bias in the Hill (1975) estimator. Still, a number k of observations needs to be chosen, however Huisman et al. (2001) show that the estimates of γ are not that sensitive to wrong

choice for k . In this paper, we set $k = \frac{n}{4}$ as sug-

gested by Huisman et al. (2001)¹. They applied their method to changes in the values of exchange rates, finding values for α between 3 and 5.

2. Data and descriptive analysis

The primary data for this study consists of daily forward closing prices for two markets, the European Energy Exchange (EEX) in Germany and the Nordic Power Exchange (NordPool), which is the single power market for Norway, Denmark, Sweden and Finland. The forward contracts for the EEX market include the base- and peakload delivery contracts for the years of 2009, 2010, and 2011². These contracts are traded for several years before delivery on the exchanges. We limit ourselves to study the forward prices obtained in the period between one year before maturity until the last trading day before delivery starts as commonly the next-year's delivery contract is the most liquid. Therefore, we study the prices as quoted in 2008 for the 2009 delivery contracts, the prices quoted in 2009 for the 2010 delivery contracts and the prices quoted in 2010 for the 2011 delivery period. Our dataset spans the trading days between January 1, 2008 through December 17, 2010, having approximately 250 daily forward price observations per year. Table 1 contains descriptive statistics for the daily changes in the natural logarithms of the forward prices for the EEX and NPX base- and peakload forward contracts.

Table 1. Statistics for the daily log forward price changes observed

	2009		2010		2011	
	Base	Peak	Base	Peak	Base	Peak
EEX						
Mean	0.000	0.000	-0.001	-0.001	0.000	-0.001
Median	0.000	0.000	-0.001	-0.002	-0.002	-0.002
Max	0.065	0.049	0.052	0.045	0.037	0.038
Min	-0.059	-0.057	-0.046	-0.035	-0.033	-0.025
St.dev.	0.015	0.014	0.295	0.012	0.010	0.010
Skew	-0.229	-0.618	0.295	0.395	0.558	0.696
Kurt	2.386	2.574	1.926	2.029	0.926	1.253
n	251	251	251	251	247	247

¹ We refer to Huisman et al. (2001) for the weighted least squares method to estimate the tail-index and for the procedure to obtain standard errors.

² For instance, the baseload 2009 contract involves the delivery of 1MW of power in any hour of the calendar year 2009 and the peakload 2009 contract involves the delivery of 1MW of power in any hour on weekdays between 8 a.m. and 8 p.m. in 2009.

	2009		2010		2011	
	Base	Peak	Base	Peak	Base	Peak
NPX						
Mean	-0.001	-0.001	0.000	-0.001	0.001	-0.001
Median	0.001	0.000	0.000	-0.002	0.002	-0.002
Max	0.067	0.059	0.092	0.045	0.064	0.038
Min	-0.090	-0.090	-0.068	-0.086	-0.056	-0.059
St.dev.	0.023	0.023	0.023	0.024	0.017	0.020
Skew	-0.72	-0.736	0.248	0.336	0.099	0.327
Kurt	2.094	1.587	1.339	1.957	1.040	1.292
n	249	248	248	244	243	243

Table 1 shows that the daily mean log-price change was about -0.001 in 2010, or -0.1%, for the 2011 peakload delivery contract. The maximum price change was 3.8% on one day and the minimum was 2.5%. The daily log-price changes for the peakload 2011 contract were positively skewed, 0.696, and exhibit excess kurtosis of 1.253 (in excess of the normal distribution function) indicating fatter tails than a normal distribution function. All excess kurtosis values are positive, which is a sign of fat tails in all years. On average, it seems that the tails of the distribution of log-price changes for EEX contracts are fatter than for the distribution of log-prices changes in the NPX.

3. Results

This section shows the tail-index estimates for the power delivery forward contracts. Table 2 shows the tail-index estimates for the baseload contracts and Table 3 shows those estimates for the peakload contracts. Let's focus on the baseload results in Table 2 first. The γ estimate for the 2009 delivery contract as traded on the EEX is significantly different from zero, being 0.267 with a standard error of 0.118. This γ estimate yields a value of 3.748 for α . The left tail of the empirical distribution of log-price changes of the EEX 2009 delivery contract has a γ estimate of 0.286 and the right tail has a γ of 0.386, implying that the right tail is fatter than the left tail, i.e., more extreme positive than negative price changes occurred for the EEX 2009 delivery contract.

The first result we learned from Tables 2 and 3 is that the tail-index estimates (in terms of α) vary between 1.837 and 6.609 for the EEX baseload contracts and between 2.461 and 12.040 for the NPX baseload contracts³. For the peakload contracts these values vary between 2.313 and 6.399 for the EEX and 3.015 and 31.707 for the NPX contracts. These levels are in line with the tail-index estimates as observed for returns on stocks and exchange rates⁴. The empirical distribution of log-price changes on power delivery forward contract are clearly fatter tailed than a normal distribution.

³ We ignore the α estimate of -42.433 here, which is perhaps due to some estimation error.

⁴ See for instance Jansen and de Vries (1991), Koedijk and Kool (1994), Kearns and Pagan (1997), and Huisman et al. (2001).

The second result we learned from these tables is that there is no clear relation between the level of the tail-index and the specific market in which the forward delivers. Neither it seems that there is an apparent difference in tail-index values between the baseload and peakload contracts or between the left and the right tail of the distribution.

Table 2. Tail fatness estimates for EEX and NPX forward base returns

2009	γ	α	γ_l	α	γ_r	α
EEX	0.267 (0.118)	3.748	0.286 (0.051)	3.491	0.386 (0.016)	2.593
NPX	0.281 (3.027)	3.559	0.406 (0.014)	2.461	0.284 (0.115)	3.521
2010	γ	α	γ_l	α	γ_r	α
EEX	0.151 (0.021)	6.609	-0.024 (0.003)	42.433	0.275 (0.443)	3.635
NPX	0.089 (0.012)	11.185	0.083 (0.023)	12.040	0.122 (0.119)	8.182
2011	γ	α	γ_l	α	γ_r	α
EEX	0.402 (0.008)	2.489	0.544 (0.007)	1.837	0.284 (0.058)	3.527
NPX	0.258 (0.029)	3.882	0.166 (0.017)	6.025	0.359 (0.022)	2.781

Notes: Standard errors are in parenthesis. γ reflects the tail-index for both tails; γ_l for the left tail, γ_r for the right tail; α is calculated as $1/\gamma$.

Table 3: Tail fatness estimates for EEX and NPX forward peak returns

2009	γ	α	γ_l	α	γ_r	α
EEX	0.343 (0.015)	2.918	0.300 (0.197)	3.335	0.357 (0.022)	2.801
NPX	0.173 (0.022)	5.787	0.161 (2.433)	6.218	0.161 (0.058)	6.207
2010	γ	α	γ_l	α	γ_r	α
EEX	0.242 (0.099)	4.139	0.164 (0.070)	6.102	0.267 (0.023)	3.747
NPX	0.300 (0.073)	3.333	0.239 (0.209)	4.179	0.332 (0.033)	3.015
2011	γ	α	γ_l	α	γ_r	α
EEX	0.370 (0.011)	2.702	0.432 (0.010)	2.313	0.156 (0.032)	6.399
NPX	0.119 (0.014)	8.436	0.181 (0.018)	5.514	0.032 (0.021)	31.707

Notes: Standard errors are in parenthesis. γ reflects the tail-index for both tails; γ_l for the left tail, γ_r for the right tail; α is calculated as $1/\gamma$.

Discussion and concluding remarks

In this paper, we have shown that the empirical distributions of log-price changes (or returns) of power forward prices exhibit significant fat tails. This implies that extreme price movements (both up and down) occur more frequently than what a normal distribution function would express. This is a result too important to ignore for risk managers, for instance, as they can-

not use normal distributions to calculate their risk measures or values for options and other derivative contracts. If they would do so, they would underestimate the level of risk. With this in mind they can improve their estimates of value at risk, the average loss beyond the value at risk measure, or expected maximum losses using extreme value theory. To see this, we briefly discuss in the extreme value theory way of calculating value at risk, based on Huisman (2009). Suppose a risk manager likes to measure the 99% one-day Value at Risk (VaR) faced on an open position in the 2011 power baseload delivery contract on the EEX. Per definition, the one-day 99% VaR is that price increase that is being exceeded in only 1% of all days. Let r be the one-day percentage return. Then

$$Pr\{r > 99\%VaR\} = 0.01. \tag{8}$$

From equation (3), we can derive

$$Pr\{r > VaR\} = \frac{n_u}{n} \left(1 + \frac{VaR - u}{\alpha\beta} \right)^{-\alpha} = 0.01. \tag{9}$$

Rewriting this yields

$$VaR = u + \beta\alpha \left(\left(\frac{0.01n}{n_u} \right)^{\frac{-1}{\alpha}} - 1 \right). \tag{10}$$

Using the estimates for α we have obtained before and if we choose a proper value for u , we can easily measure VaR using extreme value theory. Hull (2007) suggests to set u equal to the 95% quantile of the distribution function obtained from historical observations, such that $\frac{n_u}{n} = 0.05$. Still, the risk manager has to estimate β . Huisman et al. (1998) show an alternative way of calculating VaR using extreme value theory. They argue that the degrees of freedom ν in a Student-t distribution, which is fat-tailed, equals the tail-index α . They apply extreme value theory to estimate the degrees of freedom ν and then they read off the value at risk from the Student-t distribution. The advantage of this method is that one does not need to estimate β in equation (10) or to choose a proper value u . They show that their method provides better VaR estimates for stocks and bonds compared to VaR based on the normal distribution.

This paper shows that the price changes of power delivery forward contracts are fatter tailed than a normal distribution function. A risk manager in the power industry can, therefore, obtain more accurate risk calculations by applying extreme value theory than by applying measures based on normal distributions.

References

1. C. Bernhardt, C. Klüppelberg, and T. Meyer-Brandis (2008). Estimating high quantiles for electricity prices by stable linear models, *Journal of Energy Markets*, 1 (1), pp. 3-19.
2. H.N.E. Byström (2005). Extreme value theory and extremely large electricity price changes, *International Review of Economics and Finance*, 14, pp. 41-55.
3. A. Eydeland and K. Wolyniec (2003). Energy and power risk management: new developments in modeling, pricing, and hedging, John Wiley & Sons, New Jersey, U.S.A.
4. E.F. Fama and K.R. French (1987). Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage, *Journal of Business*, 60 (1), pp. 55-73.
5. D.V. Gnedenko (1943). Sur la distribution limitée du terme d'une série aléatoire. *Annals of Mathematics*, 44, pp. 423-453.
6. B. Hill (1975). A simple general approach to inference about the tail of a distribution, *The Annals of Mathematical Statistics*, 3, pp. 1163-1174.
7. R. Huisman (2009). An introduction to models for the energy markets, Risk books.
8. R. Huisman, C.G. Koedijk, and R. Pownall (1998). Var-x: fat tails in financial risk management, *The Journal of Risk*, 1(1), pp. 43-61.
9. R. Huisman, C.G. Koedijk, C. Kool, and F. Palm (2001). Tail-index estimates in small samples, *Journal of Business and Economic Statistics*, 19, pp. 208-216.
10. J.C. Hull (2007). The power law. In Jon Danielsson, (ed.), The value at risk reference. Key issues in the implementation of market risk. Risk books.
11. D. Jansen and C. de Vries (1991). On the frequency of large stock returns: putting booms and busts in perspective. *Review of Economics and Statistics*, 73, pp. 18-24.
12. P. Kearns and A. Pagan (1997). Estimating the density tail index for financial time series, *Review of Economics and Statistics*, 79, pp. 171-175.
13. C.G. Koedijk and C. Kool (1994). Tail estimates and the ems target zone, *Review of international economics*, 2, pp. 153-165.
14. J. Lucia and E.S. Schwartz (2002). Electricity prices and power derivatives: evidence from the Nordic power exchange, *Review of Derivatives Research*, 5, pp. 5-50.
15. F. Ren and D.E. Giles (2007). Extreme value analysis of daily Canadian crude oil prices, *Technical Report EWP0708*, University of Victoria.