Les Cahiers du GERAD ISSN: 0711–2440

A simple and tractable time-series model of electricity prices

J.-G. Simonato
M. Denault

G–2015–08

February 2015

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

A simple and tractable time-series model of electricity prices

Jean-Guy Simonato[a]

Michel Denault[b]

[a] *Department of Finance, HEC Montréal, Montréal (Québec) Canada, H3T 2A7*

[b] *GERAD & Department of Decision Sciences, HEC Montréal, Montréal (Québec) Canada, H3T 2A7*

jean-guy.simonato@hec.ca
michel.denault@hec.ca

February 2015

Les Cahiers du GERAD

G–2015–08

Abstract: In deregulated markets, electricity prices are typically characterized by four key features: seasonality, mean-reversion, the possibility of large downward or upward unexpected spikes, and volatility clustering. We propose a time-series price model with skewed and leptokurtic shocks, which displays all four features above. Importantly, the model fundamentally relies on a continuous and monotone transformation of a one-dimensional normal random variable, which is of considerable interest when the electricity price model is only a part of a larger problem whose solution requires the use of numerical integration and/or simulation techniques. Using a maximum likelihood approach, we compare the proposed model with other specifications of the electricity price dynamics. The estimation is done with data from Nord Pool Spot, NYISO and the United States EIA. The results reveal that the proposed model provides an adequate fit and summarizes well peak period electricity price data.

Résumé : Les prix de l'électricité dans les marchés déréglementés possèdent quatre caractéristiques principales : un caractère saisonnier, un retour à la moyenne, la possibilité de hausses et baisses soudaines et inattendues, et des poches de volatilité élevée. Nous proposons un modèle de série chronologique dont les innovations sont asymétriques et leptocurtiques, et qui présente les quatre caractéristiques recherchées. Ce modèle repose à la base sur une transformation continue et monotone d'une variable gaussienne univariée, ce qui présente un intérêt considérable lorsque le modèle de prix n'est qu'une partie d'un problème plus large résolu par intégration numérique ou simulation. À l'aide d'un maximum de vraisemblance, nous comparons notre modèle avec d'autres spécifications de prix de l'électricité tirées de la littérature. L'estimation repose sur des données de Nord Pool Spot, du NYISO et de l'EIA américain. Les résultats indiquent que notre modèle obtient une bonne adéquation aux données et qu'il les résume bien.

1 Introduction

In the last two decades, many countries have opted to deregulate their electricity markets. For example, the United States and Nordic European countries did so in the early nineties and, more recently, Great Britain, Australia and parts of Canada. On these deregulated markets, electricity prices typically show important variations and often spike up or down unexpectedly. Such characteristics are caused, in part, by the inelastic short-run demand and by the limited ability to store power, precluding the use of inventory to smooth out production and prices.

There is considerable interest in modelling electricity prices **with stochastic processes**: for purely trading the various types of contracts (including derivatives), for risk management purposes, but also for power plant managers, who adapt productions levels to the market prices. In most cases, the "best" price model should be easy to estimate, adequate in fit, and simple to incorporate in a larger problem. The main objective of this study is thus to propose and examine a time-series model which is rich enough to reproduce the main characteristics of peak period electricity prices, but that can also be conveniently used in problems requiring numerical integration and/or simulation.

In the literature, four distinctive features have been found important to characterize electricity prices. See, for example, Knittel and Roberts (2005) and Liu and Shi (2013). A first feature is seasonality, which is the portion of the power prices rising because of behavioral and natural (weather) regularities. For electricity markets, seasonality is multiscale: daily variations associated to the day and night activities (peak vs off-peak periods); week vs weekend variations associated to weekly working activities; and monthly variations induced by increases or decreases in electricity demand due to seasonal temperatures. A second important feature is a mean-reverting behavior, as demand shocks are absorbed on the longer run by variations in production capacity. For example, high prices due to unanticipated cold weather eventually go back to lower levels because of increases in supply by different producers with flexible production capacity. A third key feature of electricity prices is the occurrence of large unpredictable upward or downward movements. Such changes in power prices are associated with unexpected events like sudden unexpected variations in weather conditions or production decreases due to equipment failure. Finally, a fourth characteristic is volatility clustering. Periods of high (low) volatility tends to be followed by periods of high (low) volatility. Associated with this clustering, Knittel and Roberts (2005) also find an inverse leverage effect, i.e. positive shocks tend to increase price volatility more than negative shocks.

The power price literature comprises both continuous-time and discrete-time models. The main advantages of continuous-time models is the possibility to obtain closed-form solutions for derivatives prices. See, for example, Lucia and Schwartz (2002), Cartea et al. (2005), Barlow (2002) and Geman and Roncoroni (2006). On the other hand, discrete-time processes are widely used both to model prices directly or as approximations of continuous-time models. A full array of numerical techniques apply readily to discrete time models, and many problems are naturally discrete in nature, power generation decisions for example. Note also that data is inevitably available at discrete intervals, and parameter estimation is usually easier for a discrete time model. GARCH models have often been proposed in power price modelling (see Liu and Shi (2013), or Fowowe (2013) for example) as well as in other energy prices modelling (see e.g. Hung et al. (2011), Ji and Fan (2011), Nguyen and Nabney (2010)). Regime-switching models are also popular, see Janczura and Weron (2010) for example.

We propose a discrete-time, continuous-state price model that captures the key characteristics mentioned above, and that remains highly tractable. The model includes seasonality, mean-reversion and volatility clustering specified as a non-linear asymmetric GARCH process (NGARCH) with skewed and leptokurtic shocks. In the GARCH literature, Choi and Nam (2008) and Simonato (2012) propose the use of Johnson-S_u shocks to model exchange rates and stock returns. We adapt this specification to electricity prices. Using a maximum likelihood approach, we show that this specification of shocks can accommodate the large price swings observed in electricity markets and perform better than alternative models with jumps or Student distributed shocks. Our specification is also theoretically appealing since models of electricity prices predict asymmetric distributions in periods of high volatility (see for example Bessembinder and Lemmon (2002)). Finally, the skewed and leptokurtic Johnson S_u random variables used in this model are continuous and monotone transformation of standard normal random variables. This is an important advantage in contexts where the price model is an ingredient of a larger problem: for example, energy portfolio optimization, contract design, or power storage management. For such problems, reliable and efficient numerical tools exist that are based on normal random variables; using these tools then requires little or no adjustment at all. **Our contribution can also be seen as an extension of Knittel and Roberts (2005) to the non-gaussian shock framework. In their study, models with the key characteristics mentioned earlier are examined for various U.S. electricity markets. However, their specifications use Gaussian shocks, which have been found by many to be inaccurate descriptions of financial and commodity price data.**

The paper is organized as follows. Section 2 describes the electricity processes examined in this study. Three specifications allowing for clustering, jumps and non-normal shocks are examined. Section 3 examines the fit of the proposed model with peak electricity prices from several markets. Section 4 provides a brief discussion on how the time series model can be used in the numerical work typically required in larger problems. Section 5 concludes.

2 The electricity price model

We first assume that the power price over the next time interval is formed of two components: a known and deterministic seasonal component and a stochastic mean-reverting component i.e.

$$S_t = f_t + s_t$$

where f_t is the seasonal component and s_t is the mean-reverting component. For the sake of simplicity, **we use daily averages of peak daytime period data**, with weekends and holidays removed. Hence, seasonality is essentially brought by monthly weather changes. To capture these regularities, various techniques have been proposed, see for example Nowotarski et al (2013). For the sake of simplicity, we proceed as in Lucia and Schwartz (2002) and use dummy variables to write the seasonal component as:

$$f_t = \pi_1 \mathbf{1}_{\{t \in \text{Jan}\}} + \cdots + \pi_{12} \mathbf{1}_{\{t \in \text{Dec}\}} \tag{1}$$

where $\mathbf{1}_{\{\cdot\}}$ is an indicator function and $\pi = [\pi_1, ..., \pi_{12}]'$ is a vector of fixed parameters.

For the stochastic mean-reverting component, three different specifications are examined. All include mean-reversion and NGARCH variances, therefore capturing the volatility clustering phenomenon.[1] With respect to the yet unaccounted feature, the unexpected large changes in price, the first specification uses jumps while the second and third ones use leptokurtic and/or non-normal shocks. It should be noticed that all time series specifications examined have known transition densities and distributions. The transition density allows us to define the likelihood function with which parameter estimates can be computed.

2.1 Model 1: NGARCH volatility with jumps

A first approach that can be used to capture the large unexpected changes in electricity prices is the addition of jumps to the typical process used to model mean-reverting behavior. As in Das (2002), the mean-reverting price with jumps is given by the following discrete-time process:

$$s_{t+1} = (1 - \kappa)\, s_t + z_{t+1}\sqrt{h_{t+1}} + J \times D_{t+1} \tag{2}$$

with the variance model

$$h_{t+1} = \beta_0 + \beta_1 h_t + \beta_2 h_t (e_t/h_t - \theta)^2$$

where $\kappa > 0$ is a fixed parameter interpreted as the speed of mean-reversion, z_{t+1} is a $N(0{,}1)$ noise, h_{t+1} is the conditional variance of the seasonally-adjusted price and $e_{t+1} = s_{t+1} - E\left(s_{t+1} \mid s_t\right)$. We specify a mean of zero, since the seasonal component captures the average price for a given time point. The conditional

[1] We also tested a model with jumps but no GARCH shocks. We do not report results for this model, as its statistics were uniformly worse than the three other.

variance follows a non-linear asymmetric GARCH (NGARCH) process of Engle and Ng (1993) with the typical GARCH parameter restrictions: $\beta_0 > 0, \beta_1 \geq 0, \beta_2 \geq 0$. Parameter θ determines the "leverage effect". A positive θ combined with a negative shock increases the variance next period. In the stock price literature, this behavior is referred to as the leverage effect.[2] With a negative value of θ, a positive shock will increase the variance. A negative θ is labeled "inverse leverage effect" by Knittel and Roberts (2005). Other GARCH specifications allowing for the leverage effect could be used. For example, the GJR-GARCH of Glosten et al (1993) or the EGARCH of Nelson (1991). However, the NGARCH is more parsimonious, with four parameters compared to five with these other specifications, and provides equivalent performances.

It should be noticed that the process defined above can yield negative electricity prices. Unlike many commodities, electricity prices can become negative, as explained in Sewalt and De Jong (2003). Negative prices are possible because there must always be a balance between supply and demand on a power network. During off peak periods, power supply can be higher than demand. Negative prices are acceptable to power suppliers because the opportunity costs of a shutdown period can be very high.

The large changes in price are captured by the last term in (2). There, J, is a jump shock with a $N\left(\alpha_J,\ \sigma_J^2\right)$ distribution that is independent of z_{t+1}, and D_{t+1} is a discrete-time Poisson increment. This Poisson process is approximated by a Bernoulli distribution with parameter $\gamma = \lambda dt + O\left(dt\right)$ where λ is the annual jump intensity parameter and dt the length of a discrete time interval. As a result, the expected value of the future price, conditional on the current price, is

$$E\left(s_{t+1} \mid s_t\right) = \gamma\left[\left(1-\kappa\right)s_t + \alpha_J\right] + \left(1-\gamma\right)\left[\left(1-\kappa\right)s_t\right].$$

Because the limit of the Bernoulli process is governed by a Poisson distribution, the density of the Poisson model is well approximated with a Bernoulli mixture of normal. This approximation, examined in Ball and Torous (1983), Das (2002), and Knittel and Roberts (2005), amounts to the assumption that in each time interval, no more than one jump can occur. This is not constraining for problems with short frequencies such as the one examined here. As found in Ball and Torous (1983) and Das (2002), this approximation has the advantage of providing a tractable, stable and convergent estimation procedure with a maximum likelihood approach. For this purpose, the transition density of the seasonally-adjusted price is a simple combination of normal densities (with and without jump) and is given by:

$$\begin{aligned}\phi\left(s_{t+1} \mid s_t, \Psi\right) = {} & \gamma \frac{1}{\sqrt{2\pi\left(h_{t+1}+\sigma_J^2\right)}} \exp\left(-\frac{1}{2}\frac{\left(s_{t+1}-\left(1-\kappa\right)s_t-\alpha_J\right)^2}{h_{t+1}+\sigma_J^2}\right) \\ & + \left(1-\gamma\right)\frac{1}{\sqrt{2\pi h_{t+1}}}\exp\left(-\frac{1}{2}\frac{\left(s_{t+1}-\left(1-\kappa\right)s_t\right)^2}{h_{t+1}}\right)\end{aligned}$$

where $\Psi = \left[\kappa, \alpha_J, \sigma_J, \gamma, \beta_0, \beta_1, \beta_2, \theta\right]$. With a time series s_t for $t = 1$ to T, the parameters can be estimated by finding the parameter values maximizing the log likelihood function of this model, which is obtained from the above transition density.

2.2 Model 2: NGARCH volatility with standardized Students shocks

A second approach for capturing large unexpected jumps is to use shocks with leptokurtic distribution. In the GARCH framework, Bollerslev (1987) has suggested the use of standardized Student shocks with GARCH volatility processes. Such a distribution has fatter tails than the normal distribution and can thus capture the large sudden shocks associated with the inelastic demand of the electricity market.

Here, s_t, the seasonally adjusted mean-reverting price is

$$s_{t+1} = \left(1-\kappa\right)s_t + e_{t+1}\sqrt{h_{t+1}} \tag{3}$$

[2] For stock prices, a negative shock decreases the overall equity value of the firm. This decrease in equity value increases the proportion of debt in the firm i.e. increases the leverage. Higher leverage values produces more volatile cash flows. Hence, a negative shock should have a greater impact on volatility than a positive shock.

with

$$h_{t+1} = \beta_0 + \beta_1 h_t + \beta_2 h_t (e_t - \theta)^2$$

where e_t is a standardized $Student\,(\delta)$ noise with $\delta > 2$ degrees of liberty, and h_t is the conditional variance of the seasonally-adjusted price. As discussed in Bollerslev (1987), the transition density for this process is given by :

$$\phi\left(s_{t+1} \mid s_t, \Psi\right) = \frac{\Gamma\left(\left(\delta+1\right)/2\right)}{\Gamma\left(\delta/2\right)\sqrt{\pi\left(\delta-2\right)h_{t+1}}} \times \left(1 + \frac{\left[\left(s_{t+1} - \left(1-\kappa\right)s_t\right)/\sqrt{h_{t+1}}\right]^2}{\delta - 2}\right)^{\frac{-(1+\delta)}{2}}$$

where $\Gamma\left(\cdot\right)$ is the gamma distribution and Ψ is the vector of parameters

$$\Psi = \left[\kappa, \beta_0, \beta_1, \beta_2, \theta, \delta\right].$$

For large values of δ, the distribution has an excess kurtosis of zero and converges to the standard normal distribution as δ goes to infinity. The Student distribution is symmetric. This might be a disadvantage since theoretical models of electricity prices predict asymmetric distributions in periods of high volatility (see Bessembinder and Lemmon (2002)). Another disadvantage of the Student distribution is the possibility of infinite higher moments. For example, δ must be higher than four for the kurtosis to be well defined. As with the first model, the parameters can be estimated by maximizing the log likelihood function.

2.3 Model 3: NGARCH volatility with Johnson S_u shocks

Finally, we consider a model with shocks based on a leptokurtic Johnson S_u- normal distribution, as proposed in Choi and Nam (2008) and Simonato (2012). Unlike the Student distribution examined above, this distribution can accommodate a wide variety of skewness and kurtosis combinations for the error terms, all with finite moments. This (potential) asymmetry is important because models of electricity prices predict asymmetric distributions in periods of high volatility (see for example Bessembinder and Lemmon (2002)).

In addition, the distribution function associated to the Johnson S_u-normal is computed with the standard normal distribution function, and is therefore as easy to evaluate. We adapt the NGARCH model with standardized Johnson S_u-normal shocks proposed in Simonato (2012). In our context, the stochastic mean-reverting price is:

$$s_{t+1} = \left(1-\kappa\right)s_t + \varepsilon_{t+1}\sqrt{h_{t+1}} \tag{4}$$

with

$$h_{t+1} = \beta_0 + \beta_1 h_t + \beta_2 h_t (\varepsilon_t - \theta)^2. \tag{5}$$

Here, the ε_t term,

$$\varepsilon_t = \frac{y_t - M_y}{\sqrt{V_y}} \quad \text{with} \quad y_t = \sinh\left(\frac{z_t - a}{b}\right), \tag{6}$$

is a standardized Johnson S_u-normal shock, which is a transformation of a standard normal error term z_t with parameters $-\infty < a < +\infty$ and $b > 0$ which control the skewness and kurtosis of the distribution. Quantities M_y, V_y and sinh are defined in Appendix A which describes in more length the standard Johnson S_u-normal random variable. Using such random variables, the transition density function for s_{t+1} can be written as

$$\phi\left(s_{t+1} \mid s_t, \Psi\right) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\left(a + b\cdot\sinh^{-1}\left(M_y + \varepsilon_{t+1}\sqrt{V_y}\right)\right)^2} \times \frac{\sqrt{V_y}b}{\sqrt{h_{t+1}}\sqrt{\left(M_y + \varepsilon_{t+1}\sqrt{V_y}\right)^2 + 1}},$$

with

$$\varepsilon_{t+1} = \frac{s_{t+1} - (1-\kappa)\, s_t}{\sqrt{h_{t+1}}}$$

and a parameter vector $\Psi = [\kappa, \beta_0, \beta_1, \beta_2, \theta, a, b]$ which can be estimated by maximum likelihood.

3 Case study

In this section, we present comparative results for the three price models, on three different datasets. After a few words on the estimation procedure, we present the datasets in Section 3.1 and comment on the results in Section 3.2.

In theory, the use of the maximum likelihood technique requires a maximization over all parameters concurrently. In practice, when the number of parameters is large, the estimation is often done in two steps to simplify the numerical optimization problems.

We thus use a two-step procedure and first estimate the twelve parameters of the seasonal dummy variables with an ordinary least-square procedure on the electricity prices. As specified in our model, the residual vector from this linear regression is the stochastic mean-reverting component. The second step thus uses these residuals in conjunction with a maximum likelihood approach to estimate the parameters of the mean-reverting component by nonlinear optimization. This explains why the likelihood functions of Section 2 were provided for the seasonally-adjusted prices and not the original prices. The log-likelihood functions of the models described earlier are given by

$$\sum_{t=2}^{T} \ln \phi\left(s_t \mid s_{t-1}, \Psi\right)$$

where $\phi\left(s_t \mid s_{t-1}, \Psi\right)$ is the transition density for mean-reverting portion and Ψ is a vector containing the parameters described earlier.

3.1 Datasets

To provide broader support for our conclusions, the models were compared on three different datasets, described below. The first two, Nord Pool Spot and NYISO, are market operators who are mandated to organize power exchanges on their respective territories. The third, the EIA, is an agency that collects energy data in the United States. In all cases, to simplify the seasonality treatment, weekend and holiday data were removed from the time series.

3.1.1 Nord Pool Spot Dataset

Nord Pool Spot operates the day-ahead and intraday power exchanges in Nordic Europe. It is one of the best established power markets in the world, in operation for close to twenty years. Details for our data are as follows.

Price nature: Day-ahead system prices, determined as the intersection of the aggregate supply and demand curves.

Time period: Peak time periods only (from 8 am to 8 pm). Our daily price is the simple average of the peak hours prices for the day.

Time horizon: From 1 January 2009 to 14 June 2011 (two and a half years)

Data origin: Web site of Nord Pool Spot, www.nordpoolspot.com.

3.1.2 NYISO Dataset

NYISO is the Independent System Operator of the state of New York. As such, it operates a daily market of wholesale electricity. Details on the data are as follows.

Price nature: Day-ahead prices.

Time period: Peak time periods only (from 8 am to 8 pm). Our daily price is the simple average of the peak hours prices for the day.

Time horizon: From 3 January 2006 to 31 December 2012 (seven years).

Locational prices: Four sub-markets (zonal prices) where examined, out of fourteen: Centrl, North, Nyc and West.

Data origin: Web site of NYISO, www.nyiso.com.

3.1.3 EIA Dataset

The United States Energy Information Administration (EIA) is the statistical and analytical agency within the U.S. Department of Energy. Details on the data are as follows.

Price nature: Day-ahead prices. Prices are power indices taking into account the volume of transactions. The power indices definition is available from the EIA website.

Time period: Peak time periods only (from 8 am to 8 pm). Our daily price is the simple average of the peak hours prices for the day.

Time horizon: From 5 January 2006 to 2 January 2013 (seven years).

Locational prices: Five hub prices where examined, out of eight: Indiana Hub (Midwest region); Mass Hub (New England region); Mid-C (Mid-Columbia, Northwest region); Palo Verde (Southwest region); SP-15 (Southern California). Only the hubs with times series of more than 230 data points each year where selected.

Data origin: Web site of the EIA, www.eia.gov.

3.2 Estimation results

Because the estimates show great similarities within a dataset, we present detailed results and analyses for three time series only (one from each data set): Nord Pool Spot, the "Centrl" price of NYISO and the "Indiana Hub" price of EIA. Detailed results for the other locations in NYISO and EIA are available from the authors.

Table 1 presents the seasonal parameter estimates for the three datasets mentioned above. All parameters are significant at the five percent level. Standard errors are provided in parentheses. Substantial differences are observed between the monthly dummy parameters. In Scandinavia, as expected, colder months are

Table 1: Seasonal parameters for three datasets

	Nord Pool		NYISO Centrl		EIA Indiana Hub	
	π_i	std	π_i	std	π_i	std
January	55.1	(1.5)	63.3	(1.6)	48.0	(1.4)
February	61.8	(1.6)	55.9	(1.7)	47.1	(1.5)
March	48.6	(1.6)	50.8	(1.6)	43.7	(1.3)
April	45.1	(1.7)	52.1	(1.6)	46.8	(1.4)
May	44.8	(1.7)	52.8	(1.6)	47.2	(1.4)
June	43.2	(1.6)	58.7	(1.6)	55.3	(1.4)
July	36.3	(1.6)	67.8	(1.6)	59.1	(1.4)
August	38.3	(1.6)	59.5	(1.6)	55.7	(1.3)
September	37.3	(1.6)	50.2	(1.6)	39.8	(1.4)
October	39.7	(1.6)	51.2	(1.6)	40.9	(1.3)
November	45.0	(1.6)	54.2	(1.6)	41.7	(1.4)
December	58.2	(1.6)	55.9	(1.6)	43.4	(1.4)

The table reports the maximum loglikelihood parameters for the seasonal component of prices (see equation (1)) for three datasets. Weekend and holiday prices are not considered, so that only twelve parameters are required. Standard deviation provided in parentheses.

associated with higher average prices, while in the U.S., air conditioning drives a peak in prices during the summer months. Note that the two U.S. markets differ with respect to heating demand in January: the NYISO Centrl has a clear peak, much subdued at the Indiana Hub.

Tables 2, 3 and 4 present detailed results on the maximum log-likelihood parameter estimates and several statistics for the three datasets Nord Pool Spot, NYISO (Centrl) and EIA (Indiana Hub). For all three tables, the columns present the results associated to the three models of Section 2: model 1 is the NGARCH with normal shocks and jumps, model 2 is the NGARCH with Student shocks, and model 3 is the NGARCH with Johnson-S_u shocks. In all three tables, the parameter estimates are presented first, with standard errors in parentheses, and are followed by several statistics:

Observations is the number of observation in the time series;

Log-likelihood is the log-likelihood value; Akaike is the Akaike information criterion;

JB is the Jarque-Bera normality test for the standard normal residuals with [p-values].

***Q*(20)** is the Ljung-Box portmanteau test for up to 20th-order serial correlation in the standard normal residuals with [p-values].

Q^2(20) is the same test for the squared standardized residuals with [p-values];

Stavarc is the stationarity condition for the NGARCH variance process which is computed as $\beta_1+\beta_2(1+\theta^2)$. A value higher than one is an indication of non-stationary variance.

While the numbers themselves vary slightly, the analysis for the three datasets is similar, so that we do not distinguish between datasets in what follows, with a few explicit exceptions. The plots are for the NYISO "Centrl" dataset.

All reported parameter estimates are significant at the 5 % level. For all price models, there is a strong evidence of mean-reversion with a positive speed of mean-reversion parameter estimate for κ. For the model with jumps, the estimates of the average jump values α_J are positive as expected, with large standard deviations σ_J. The estimated probabilities of a jump over the next day γ are, in general, between 5% and 10%.

For all three models, the leverage parameter θ has a negative sign coherent with the "reverse leverage" effect discussed in Section 2.1. The estimated number of degrees of freedom δ in the NGARCH-Student model is smaller than four, implying that the fourth moment of the error term is not defined. This is worse for the Nord Pool Spot dataset ($\delta = 2.2$) and better of the IEA case ($\delta = 3.55$). For Nord Pool, the third moment is also undefined since the estimate of δ is smaller than three.

With respect to the likelihoods, the NGARCH-Student improves markedly on the NGARCH-Jump, and the NGARCH-Johnson even more so, showing the importance of the added flexibility given by allowing for skewness and kurtosis. Because these models are not nested, it is not possible to test for the best specification using the likelihood values. A better comparison can be given by the Akaike information criterion, which takes into account the number of parameters. With this criterion, the better model is the one with the minimum value. Again the NGARCH-Johnson specification stands out with the lowest Akaike values in all cases.

With respect to the distributional assumptions, Figure 1 presents the quantile to quantile plots (qqplots) of the standardized residuals associated to the three models in the NYISO "Centrl" case. Given that the model with jumps is a combination of independent normals, the expected residual is used in the plots, where the expectation is taken with the estimated jump probability. As shown in this graph, the normality assumption is clearly violated in the tails of the distributions for this model. For the model with Student errors, the qqplots shows a much better fit than the preceding model. Some deficiencies in both tails of the distribution can however be noticed. Finally, for the Johnson case, the qqplot of the estimated z_t's, which have a standard normal distribution, are shown to correspond adequately to the normality assumption of the implied standard normal error term at the source of the Johnson S_u random variable (Equation (9) in the annex). For models 1 and 3, the Jarque-Bera test of normality confirms the qqplots and strongly rejects normality for model 1, while the assumption cannot be rejected at any reasonable level for the Johnson.

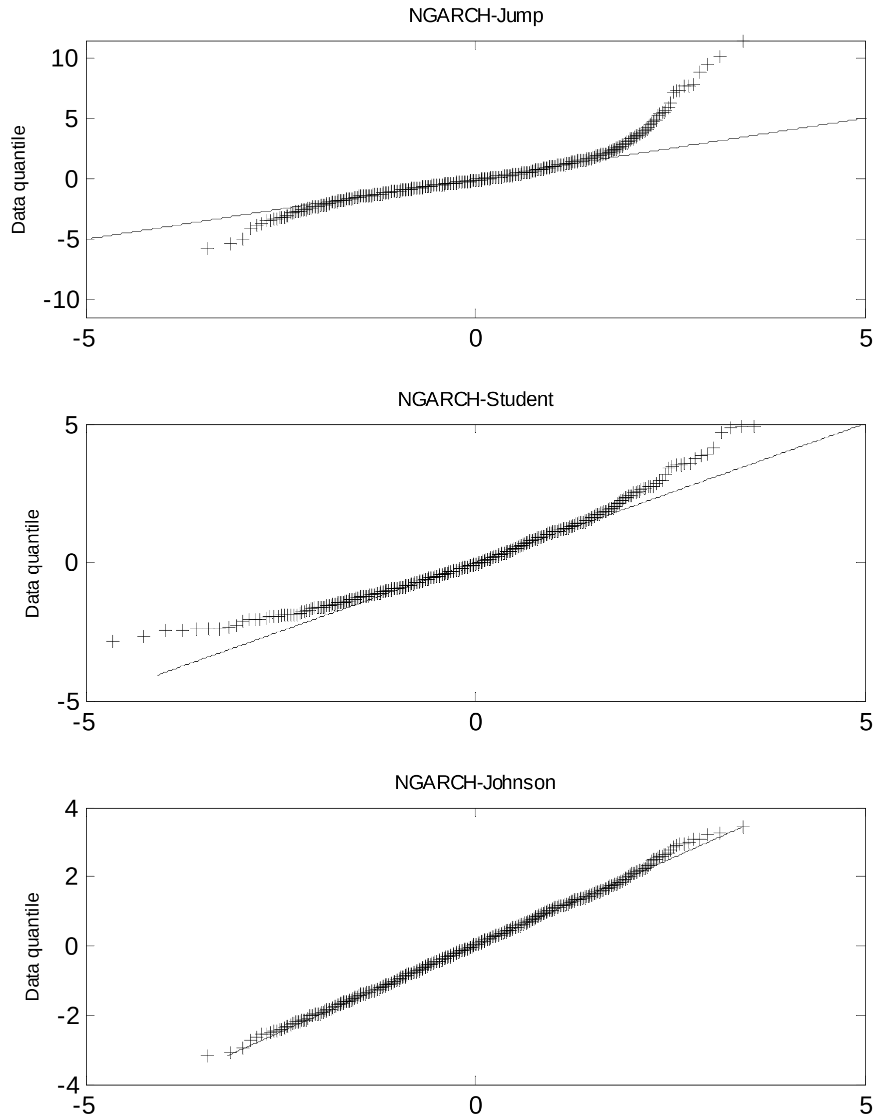

Figure 1: Quantile to quantile plot of standardized residuals for NYISO Centrl

The Ljung-Box statistics indicate that, for some of the time series, significant autocorrelations are still present in the residuals. A detailed look at the residuals indicates that these are caused by small, but significant, autocorrelation coefficients in the first few lags. Figure 2, which shows the sample autocorrelation function with confidence bands (two standard deviations), illustrate this for NYISO "Centrl" and the three models. For the squared residuals, a similar phenomena occurs (very small but significant auto-correlations) for three out of ten time series, showing that the NGARCH specification captures most of the predictability in the variance.

With respect to the stationarity of the variance process, notice first that the sum of β_1 and β_2 is close to or larger than one for the NGARCH-Student model, slightly less so for NGARCH-Johnson, and much less so for NGARCH-Jump. This is an indication of an integrated GARCH process (or close to integrated process). In this case, the variance process is non-stationary and could explode for very long maturities. The stationary variance condition ("Stavarc" in the tables) supports the same conclusions of a non-stationary variance, , with the most severe case for the NGARCH-Student model, the least severe for the NGARCH-jump model, and the NGARCH-Johnson S_u in between. For the NGARCH-Johnson model, the largest "Stavarc" value was observed for the EIA (Mass Hub) dataset at 1.20, while the largest "Stavarc" value for the NGARCH-Student model was 2.16, for the Nord Pool dataset. For the NGARCH-Johnson model, eight series out of

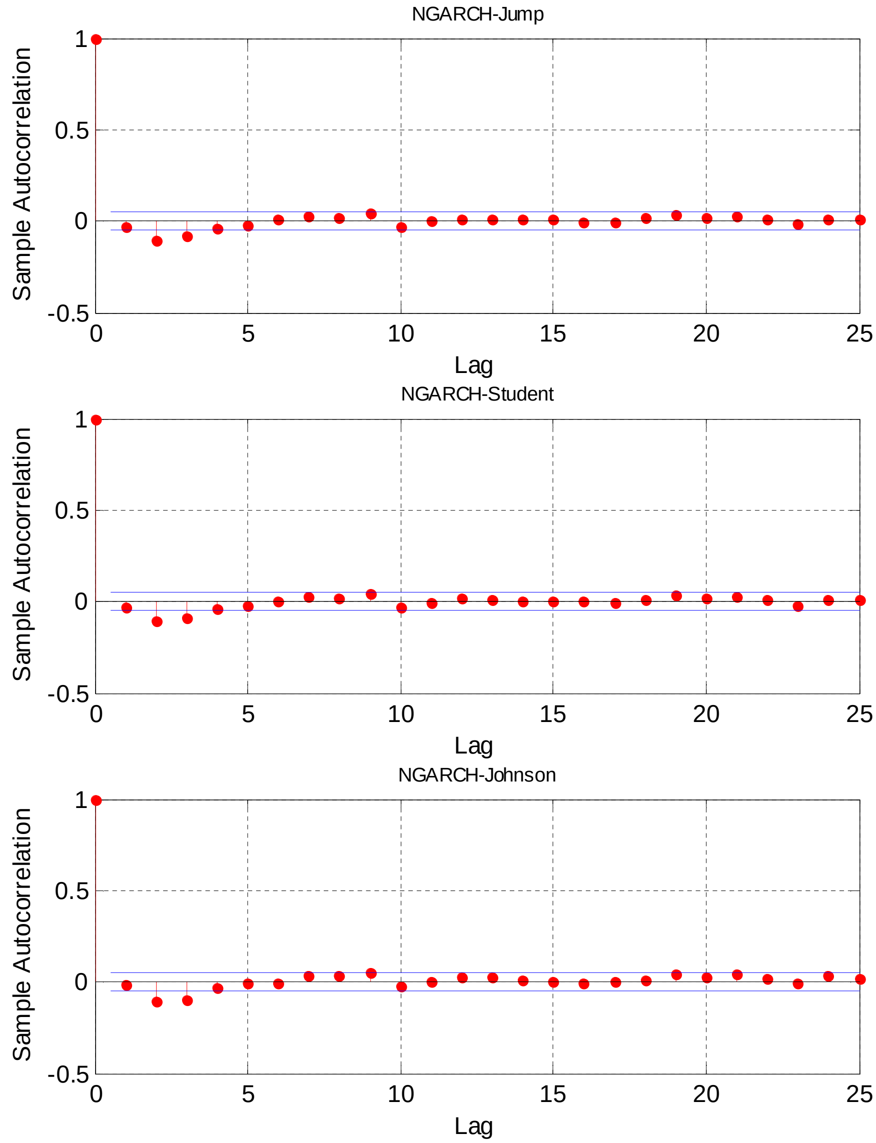

Figure 2: Sample autocorrelations with confidence bounds for NYISO Centrl

the ten that have been examined show a non-stationary variance. Note that a logarithmic transformation of the price series is capable of tempering this problem with fewer cases of non-stationary variance. Tables 5 and 6 report the parameter estimates and statistics for this log-transformed model for all ten time series. As shown in these tables, the model performs similarly on log prices about parameter estimates and statistics, but with four out of ten cases with non-stationary variance. Such a specification is, however, not able to generate negative prices. A visual examination of the qqplots and estimated autocorrelation coefficients of the residuals show characteristics similar to those obtained without the logarithmic transformation (plots and coefficients not shown)

Based on the above analysis, the preferred specification is the NGARCH with Johnson S_u shocks. It offers the best likelihood values and its residuals are coherent with the distributional hypothesis. It violates the variance stationarity assumption, but never by a large margin, and a logarithmic transformation of the price can temper the problem. Finally, on the theoretical side, it is appealing for its capacity to represent asymmetric price distributions, which the other two models cannot do.

Table 2: Parameters and statistics for the dataset "Nord Pool"

	NGARCH-Jump	NGARCH-Student	NGARCH-Johnson
κ	0.0119	0.0091	0.0088
	(0.0048)	(0.0046)	(0.0042)
α_J	10.7990	—	—
	(3.2153)		
σ_J	17.6237	—	—
	(1.0553)		
γ	0.0678	—	—
	(0.0109)		
β_0	0.7677	4.2326	1.7265
	(0.1003)	(3.0714)	(0.3820)
β_1	0.4496	0.4975	0.5199
	(0.0310)	(0.0436)	(0.0396)
β_2	0.3014	1.6563	0.6482
	(0.0345)	(1.1967)	(0.1397)
θ	-0.5004	-0.0809	-0.1681
	(0.1262)	(0.0593)	(0.0834)
δ	—	2.2047	—
		(0.1654)	
a	—	—	-0.2145
			(0.0411)
b	—	—	0.9196
			(0.0467)
Observations	1031	1031	1031
Log-likelihood	-2584	-2536	-2518
Akaike	5184	5085	5050
JB	13005	—	1.57
	[0.0000]		[0.4436]
$Q(20)$	21	19	18
	[0.4257]	[0.5031]	[0.5821]
$Q^2(20)$	8	5	33
	[0.9900]	[0.9999]	[0.0375]
Stavarc	0.83	2.16	1.19

This table reports the maximum likelihood parameters estimates for the deseasonalized time series. Standard errors are reported in parenthesis. Also reported are several statistics. Observations is the number of observation; Log-likelihood is the log-likelihood value; Akaike is the Akaike information criterion; Stavarc is the stationarity condition for the NGARCH variance process; JB is the Jarque-Bera statistic for the standard normal residuals; $Q(20)$ is the Ljung-Box portmanteau statistic for up to 20th-order serial correlation in the standard normal residuals while $Q^2(20)$ is the same statistic for the squared standardized residuals. The p-values for these statistics are reported in square brackets.

The next section discusses how this model can be used in numerical work associated to the analysis of problems broader than simple price modelling.

4 Simulation and integration of Johnson S_u random variables

As discussed in the introduction, an important feature of the NGARCH Johnson S_u model that we propose is the convenience associated with simulation and numerical integration, which are tools often required when the price model is part of a larger problem. For example, computing risk measures or prices of derivative securities on electricity often involves simulation and/or numerical integration. Power plant management or optimal hydroelectric production are other examples of larger problems involving dynamic programming, which requires the capacity to repeatedly solve for expected value with numerical integration. Because a Johnson S_u random variable is a continuous and monotone transformation of a standard normal variables, most of the numerical tools and techniques that are widely available for the normal case can be used for the Johnson S_u case.

Table 3: Parameters and statistics for the dataset "NYISO Centrl"

	NGARCH-Jump	NGARCH-Student	NGARCH-Johhson
κ	0.0373 (0.0044)	0.0369 (0.0051)	0.0426 (0.0048)
α_J	7.4608 (1.7096)	—	—
σ_J	14.4643 (0.8995)	—	—
γ	0.0832 (0.0123)	—	—
β_0	1.7628 (0.2203)	2.8316 (0.5724)	3.1224 (0.5183)
β_1	0.4853 (0.0315)	0.5892 (0.0371)	0.5266 (0.0374)
β_2	0.1443 (0.0147)	0.2825 (0.0493)	0.2669 (0.0381)
θ	-1.5493 (0.1293)	-0.8606 (0.1100)	-0.9750 (0.1085)
δ	—	3.0377 (0.2397)	—
a	—	—	-0.3317 (0.0547)
b	—	—	1.1843 (0.0595)
Observations	1763	1763	1763
Log-likelihood	-5589	-5583	-5559
Akaike	11194	11177	11132
JB	9908 [0.0000]	—	0.95 [0.6250]
$Q(20)$	48 [0.0004]	49 [0.0003]	57 [0.0000]
$Q^2(20)$	8 [0.9903]	8 [0.9903]	29 [0.0872]
Stavarc	0.98	1.08	1.05

This table reports the maximum likelihood parameters estimates for the deseasonalized time series. Standard errors are reported in parenthesis. Also reported are several statistics. Observations is the number of observation; Log-likelihood is the log-likelihood value; Akaike is the Akaike information criterion; Stavarc is the stationarity condition for the NGARCH variance process; JB is the Jarque-Bera statistic for the standard normal residuals; $Q(20)$ is the Ljung-Box portmanteau statistic for up to 20th-order serial correlation in the standard normal residuals while $Q^2(20)$ is the same statistic for the squared standardized residuals. The p-values for these statistics are reported in square brackets.

In problems involving dynamic programming, gaussian quadratures are widely used to perform numerical integration of random variables. For example, Gauss-Hermite quadrature is specifically adapted to the Gaussian density function, and performs numerical integrations very accurately and efficiently. Because a Johnson S_u is a monotone, continuous and invertible transformation of a normal random variable, one can use the change of variable theorem and Gauss-Hermite quadrature to repeatedly perform the numerical integrations required to solve a dynamic program. As shown in Appendix B , the expected value of a function $\varphi(\varepsilon)$ of a Johnson variable can can be written as

$$E[\varphi(\varepsilon)] = \int_{-\infty}^{+\infty} \varphi(g(z))\ f(z)\ dz$$

where z is a standard normal random variable, $f(z)$ is the density of the standard normal, and $g(z)$ is a generic expression for equation (6) (the Johnson random variable written in terms of the standard normal). Using the change of variable described in Judd (1998), the integral can then be easily rewritten in a way that makes it compatible with the use of the Gauss-Hermite quadrature approach.

Table 4: Parameters and statistics for the dataset "EIA Indiana Hub"

	NGARCH-Jump	NGARCH-Student	NGARCH-Johnson
κ	0.0398	0.0457	0.0574
	(0.0054)	(0.0063)	(0.0059)
α_J	10.2645	—	—
	(3.3649)		
σ_J	9.1805	—	—
	(1.8593)		
γ	0.0736	—	—
	(0.0222)		
β_0	1.2247	1.2914	1.7715
	(0.1663)	(0.2978)	(0.3204)
β_1	0.5879	0.7266	0.6683
	(0.0283)	(0.0285)	(0.0285)
β_2	0.1390	0.2210	0.2239
	(0.0152)	(0.0379)	(0.0320)
θ	-1.3963	-0.6602	-0.7815
	(0.1267)	(0.1041)	(0.1020)
δ	—	3.5550	—
		(0.3521)	
a	—	—	-0.5490
			(0.0858)
b	—	—	1.3508
			(0.0883)
Observations	1729	1729	1729
Log-likelihood	-5389	-5408	-5368
Akaike	10795	10829	10749
JB	6163	—	1.04
	[0.0000]		[0.5898]
$Q(20)$	112	109	98
	[0.0000]	[0.0000]	[0.0000]
$Q^2(20)$	10	11	10
	[0.9718]	[0.9416]	[0.9688]
Stavarc	1.00	1.04	1.03

This table reports the maximum likelihood parameters estimates for the deseasonalized time series. Standard errors are reported in parenthesis. Also reported are several statistics. Observations is the number of observation; Log-likelihood is the log-likelihood value; Akaike is the Akaike information criterion; Stavarc is the stationarity condition for the NGARCH variance process; JB is the Jarque-Bera statistic for the standard normal residuals; $Q(20)$ is the Ljung-Box portmanteau statistic for up to 20th-order serial correlation in the standard normal residuals while $Q^2(20)$ is the same statistic for the squared standardized residuals. The p-values for these statistics are reported in square brackets.

Another approach which is used for solving dynamic programming problems in Gaussian context is the Markov chain approximation (see Judd (1998)). For example, as shown in Duan and Simonato (1998) and Denault et al. (2013), dynamic programming in a normal context (with constant or GARCH volatilities) can be performed using a Markov chain approximation. One of the main requirement of such an approach is the transition probability matrix of the Markov chain, whose elements can be computed with expressions involving the distribution function of the random shocks associated with the process. For processes with normal shocks, the distribution function (cumulative density function) of a standard normal random variable is used. For such a function, fast and precise algorithms are available in most software environment. For Johnson S_u random variables, the distribution function can also be computed quickly since it only involves the distribution function of a standard normal random variable. More specifically, in the present context, the distribution function of the standard normal Johnson random variable is given by

$$\Phi\left(a + b \cdot \sinh^{-1}\left(M_y + \varepsilon\sqrt{V_y}\right)\right)$$

Table 5: Parameters estimates – NGARCH-Johnson on log prices

	κ_1	β_0	β_1	β_2	θ	a	b
NYISO West	0.050	0.001	0.759	0.167	-0.420	-0.261	1.270
	(0.007)	(0.000)	(0.035)	(0.032)	(0.134)	(0.051)	(0.066)
NYISO Centrl	0.039	0.001	0.667	0.220	-0.509	-0.183	1.306
	(0.006)	(0.000)	(0.041)	(0.038)	(0.114)	(0.055)	(0.072)
NYISO North	0.038	0.001	0.689	0.280	-0.303	-0.062	1.146
	(0.006)	(0.000)	(0.035)	(0.048)	(0.095)	(0.042)	(0.059)
NYISO NYC	0.023	0.002	0.484	0.281	-0.919	-0.182	1.278
	(0.005)	(0.000)	(0.040)	(0.047)	(0.112)	(0.062)	(0.078)
Nord Pool	0.009	0.001	0.573	0.624	0.029	-0.173	0.939
	(0.005)	(0.000)	(0.046)	(0.142)	(0.077)	(0.042)	(0.056)
IEA Mass Hub	0.021	0.001	0.539	0.349	-0.645	-0.430	1.377
	(0.003)	(0.000)	(0.025)	(0.037)	(0.064)	(0.060)	(0.063)
IEA Indiana Hub	0.063	0.001	0.813	0.129	-0.415	-0.419	1.550
	(0.008)	(0.000)	(0.028)	(0.023)	(0.131)	(0.092)	(0.122)
IEA Mid-C	0.019	0.001	0.564	0.473	0.167	0.053	1.274
	(0.004)	(0.000)	(0.027)	(0.051)	(0.048)	(0.042)	(0.058)
IEA SP-15	0.075	0.001	0.822	0.105	-0.275	-0.295	1.358
	(0.016)	(0.000)	(0.052)	(0.038)	(0.297)	(0.096)	(0.122)
IEA Palo Verde	0.017	0.001	0.733	0.180	-0.176	-0.281	1.662
	(0.004)	(0.000)	(0.028)	(0.025)	(0.088)	(0.076)	(0.096)

This table reports the maximum likelihood parameters estimates of the NGARCH-Johnson model for the log transformed prices. Standard errors are reported in parenthesis.

Table 6: Statistics – NGARCH-Johnson on log prices

	Observations	Log-likelihood	Akaike	JB	$Q(20)$	$Q^2(20)$	Stavarc
NYISO West	1762	1279	-2544	0.5219	58.9186	17.4998	0.9557
				[0.7586]	[0.0000]	[0.6203]	
NYISO Centrl	1762	1462	-2909	0.0265	69.5061	25.2470	0.9441
				[0.9920]	[0.0000]	[0.1921]	
NYISO North	1762	1350	-2686	0.5810	65.4958	30.1171	0.9939
				[0.7484]	[0.0000]	[0.0680]	
NYISO NYC	1762	1213	-2412	4.0397	74.5425	35.9678	1.0016
				[0.1176]	[0.0000]	[0.0155]	
Nord Pool	1030	1369	-2724	0.8982	16.3310	24.0083	1.1972
				[0.6199]	[0.6959]	[0.2420]	
IEA Mass Hub	2783	2584	-5153	2.1973	121.6754	29.2369	1.0323
				[0.3240]	[0.0000]	[0.0832]	
IEA Indiana Hub	1728	1248	-2482	0.3602	108.2609	9.0630	0.9640
				[0.8243]	[0.0000]	[0.9822]	
IEA Mid-C	2625	2217	-4421	0.3573	73.6925	24.6392	1.0493
				[0.8237]	[0.0000]	[0.2156]	
IEA SP-15	711	782	-1549	0.0598	61.9563	64.1338	0.9341
				[0.9800]	[0.0000]	[0.0000]	
IEA Palo Verde	2624	3074	-6133	1.0873	117.4490	28.7076	0.9188
				[0.5962]	[0.0000]	[0.0937]	

This table reports the statistics obtained for the NGARCH-Johnson model for the log transformed prices. Observations is the number of observation; Log-likelihood is the log-likelihood value; Akaike is the Akaike information criterion; Stavarc is the stationarity condition for the NGARCH variance process; JB is the Jarque-Bera statistic for the standard normal residuals; $Q(20)$ is the Ljung-Box portmanteau statistic for up to 20th-order serial correlation in the standard normal residuals while $Q^2(20)$ is the same statistic for the squared standardized residuals. The p-values for these statistics are reported in square brackets.

where $\Phi(\cdot)$ is the standard normal distribution function, and the expressions for $\sinh^{-1}(\cdot)$, M_y and V_y are availlable in Appendix A.

Finally, in problems requiring simulated paths of electricity prices, standardized Johnson S_u random variables are simply obtained by drawing a standard normal random variable and applying the transformations outlined in equation (6) i.e.

$$\varepsilon_t = \frac{y_t - M_y}{\sqrt{V_y}} \quad \text{with} \quad y_t = \sinh\left(\frac{z_t - a}{b}\right),$$

where z_t is a standard normal, and where the other quantities are defined in Section 2.3 and Appendix A.

5 Conclusion

We introduce in this paper an electricity price model based on an NGARCH with Johnson-S_u shocks (but no jumps). The model has two main advantages. First, it offers a good fit to data, as a case study indicates. The model is suited for seasonality behaviour, mean-reversion, large, unpredictable swings, and volatility clustering. It also allows for asymmetric price distributions. The second main advantage is that the model fundamentally relies on the standard normal distribution. This means that when the price model is integrated in a broader problem, such as risk-management, dynamic optimization or dam operation, classical techniques such as Markov chains, Gaussian quadratures and simulation, all apply easily.

A The standard Johnson S_u random variable

The standard Johnson S_u shock ε_t is defined as:

$$\varepsilon_t = \frac{y_t - M_y}{\sqrt{V_y}} \quad \text{with} \quad y_t = \sinh\left(\frac{z_t - a}{b}\right),$$

where sinh is the hyperbolic sine function ($\sinh(x) \triangleq (\exp(x) - \exp(-x))/2$) and z_t is N(0,1). The parameters M_y and V_y are the mean and variance of the Johnson random variable y_t and can be computed with

$$M_y = -w^{\frac{1}{2}} \sinh(\Omega), \tag{7}$$

$$V_y = \frac{1}{2}(w-1)(w\cosh(2\Omega) + 1), \tag{8}$$

where $w = e^{\frac{1}{b^2}}$, $\Omega = \frac{a}{b}$ and $\cosh(x) \triangleq (e^x + e^{-x})/2$. The standard normal error term can be recovered with

$$z_t = a + b \cdot \sinh^{-1}\left(M_y + \varepsilon_t\sqrt{V_y}\right) \tag{9}$$

and the inverse hyperbolic sine function given by $\sinh^{-1}(x) \triangleq \ln\left(x + \sqrt{x^2+1}\right)$. As discussed in Simonato (2012), the density function for ε_t can be written as

$$f(\varepsilon_t) = f(z_t) \times \left|\frac{\partial z_t}{\partial \varepsilon_t}\right|$$

where f_{z_t} is the density function of a standard normal random variable, which yields, after some manipulations, the transition density shown in Section 2.3.

B Expected value of a function of a Johnson S_u

Denote a Johnson S_u random variable as a continuous, monotonically increasing and invertible function of a standard normal random variable z :

$$\varepsilon = g(z) \quad \leftrightarrow \quad z = g^{-1}(\varepsilon).$$

The expected value of a function $\varphi(\cdot)$ of a standard Johnson S_u random variable can be written as

$$E[\varphi(\varepsilon)] = \int_{-\infty}^{+\infty} \varphi(\varepsilon)\ f(\varepsilon)\ d\varepsilon$$

where $f(\varepsilon)$ is the density function of the standard Johnson S_u-normal random variable. Using a change of variable, the Johnson density can be written as

$$f(\varepsilon) = f(z) \times \left| \frac{\partial g^{-1}(\varepsilon)}{\partial \varepsilon} \right|$$

where $f(z)$ is the density function of a standard normal random variable, which can be substituted in the above expression to yield:

$$E[\varphi(\varepsilon)] = \int_{-\infty}^{+\infty} \varphi(\varepsilon) \times f(z) \times \left| \frac{\partial g^{-1}(\varepsilon)}{\partial \varepsilon} \right| d\varepsilon.$$

Using the change of variable theorem with $\varepsilon = g(z)$ obtains the following expression :

$$E[\varphi(\varepsilon)] = \int_{-\infty}^{+\infty} \varphi(g(z)) \times f(z) \times \left| \frac{\partial g^{-1}(\varepsilon)}{\partial \varepsilon} \right| \times \frac{\partial g(z)}{\partial z} dz.$$

Given that $\frac{\partial g(z)}{\partial z} > 0$ for all z because $g(\cdot)$ is a monotonically increasing function, we have that

$$\left| \frac{\partial g^{-1}(\varepsilon)}{\partial \varepsilon} \right| \times \frac{\partial g(z)}{\partial z} = \left| \frac{\partial z}{\partial \varepsilon} \right| \times \frac{\partial \varepsilon}{\partial z} = 1$$

and

$$E[\varphi(\varepsilon)] = \int_{-\infty}^{+\infty} \varphi(g(z)) \times f(z)\, dz.$$

References

[1] Ball, C., Torous, W.N., 1983, A simplified jump process for common stock returns, Journal of Financial and Quantitative Analysis 18, 53–65.

[2] Barlow, M.T., 2002, A diffusion model for electricity prices, Mathematical Finance 12, 287–298.

[3] Bessembinder, H., Lemmon, M.J., 2002, Equilibrium pricing and optimal hedging in electricity forward markets, Journal of Finance 57, 1347–1382.

[4] Bollerslev, T., 1987, A conditionally heteroskedastic time series model for speculative prices and rates of return, Review of Economics and Statistics 69, 542–547.

[5] Cartea, A., Figueroa, M.G., 2005, Pricing in electricity markets: a mean-reverting jump diffusion model with seasonality, Applied Mathematical Finance 12, 313–335.

[6] Choi, P., Nam, K., 2008, Asymmetric and leptokurtic distribution for heteroscedastic asset returns: the s_u-normal distribution, Journal of Empirical finance 15, 41–63.

[7] Das, S.R., 2002, The surprise element: Jumps in interest rates, Journal of Econometrics 106, 27–65.

[8] Denault, M., Simonato, J.G., Stentoft, L., 2013, A simulation-and-regression approach for stochastic dynamic programs with endogenous state variables, Computers & Operations Research 40, 2760–2769.

[9] Duan, J.C., Simonato, J.G., 2001, American option pricing under GARCH by a Markov chain approximation, Journal of Economic Dynamics and Control 25, 1689–1718.

[10] Engle, R., Ng, V., 1993, Measuring and testing of the impact of news on volatility, Journal of Finance 48, 1749–1778.

[11] Geman, H., Roncoroni, A., 2006, Understanding the fine structure of electricity prices, Journal of Business 79, 1225–1261.

[12] Glosten, L., R. Jagannathan, Runkle, D., 1993, Relationship between the expected value and the volatility of the nominal excess return on stocks, Journal of Finance 48, 1779–1801.

[13] Facci, A., Andreassi, L., Ubertini, S., 2014, Optimization of CHCP (combined heat power and cooling) systems operation strategy using dynamic programming, Energy 66, 387–400.

[14] Fowowe, B., 2013, Jump dynamics in the relationship between oil prices and the stock market: Evidence from Nigeria, Energy 56, 31–38.

[15] Hung, J., Wang, Y., Chang, M., Shih, K., Kao, H., 2011, Minimum variance hedging with bivariate regime-switching model for WTI crude oil, Energy 36, 3050–3057.

[16] Janczura, J., Weron, R., 2010, An empirical comparison of alternate regime-switching models for electricity spot prices, Energy Economics 32, 1059–1073.

[17] Ji, Q., Fan, Y., 2011, A dynamic hedging approach for refineries in multiproduct oil markets, Energy 36, 881–887.

[18] Judd, K., 1998, Numerical methods in economics, MIT Press, Cambridge.

[19] Knittel, C., Roberts, M., 2005, An empirical examination of restructured electricity prices, Energy Economics 27, 791–817.

[20] Liu, H., Shi, J., 2013, Applying ARMA-GARCH approaches to forecasting short-term electricity prices. Energy Economics 37, 152–166.

[21] Lucia, J., Schwartz, E., 2002, Electricity prices and power derivatives: evidence from the Nordic power exchange, Review of Derivatives Research 5, 5–50.

[22] Nelson, D., 1991, Conditional heteroskedasticity in asset returns: a new approach, Econometrica 59, 347–370.

[23] Nguyen, H., Nabney, I., 2010, Short-term electricity demand and gas price forecasts using wavelet transforms and adaptive model, Energy 35, 3674–3685.

[24] Nowotarski, J., Tomczyk, J., Weron, R., 2013, Robust estimation and forecasting of the long-term seasonal component of electricity spot prices, Energy Economics 39, 13–27.

[25] Sewalt, M., De Jong, C., 2003, Negative prices in electricity markets, Commodities Now, June, 74–77.