

Pricing in Electricity Markets: A Mean Reverting Jump Diffusion Model with Seasonality

ÁLVARO CARTEA & MARCELO G. FIGUEROA

Birkbeck College, University of London, London WC1E 7HX, UK

ABSTRACT *This paper presents a mean-reverting jump diffusion model for the electricity spot price and derives the corresponding forward price in closed-form. Based on historical spot data and forward data from England and Wales the model is calibrated and months, quarters, and seasons-ahead forward surfaces are presented.*

KEY WORDS: Energy derivatives, electricity, forward curve, forward surfaces

1. Introduction

One of the key aspects of a competitive market is deregulation. In most electricity markets this has, however, only occurred recently. Prior to this, price variations were often minimal and heavily controlled by regulators. In England and Wales in particular, prices were set by the Electricity Pool, where, due to centralization and inflexible arrangements, prices failed to reflect falling costs and competition. Deregulation came in the recent introduction on 27 March 2001 of NETA (New Electricity Trade Arrangement), removing price controls and openly encouraging competition.

Price variations have increased significantly as a consequence of the introduction of competition, encouraging the pricing of a new breed of energy-based financial products to hedge the inherent risk, both physical and financial, in this market. Most of the current transactions of instruments in the electricity markets are carried out through bilateral contracts ahead of time although electricity is also traded on forward and futures markets and through power exchanges.

One of the most striking differences that singles out electricity markets is that electricity is very difficult or too expensive to store, hence markets must be kept in balance on a second-by-second basis. In England and Wales, this is done by the

Correspondence Address: Álvaro Cartea, Birkbeck College, University of London, Malet Street, Bloomsbury, London WC1E 7HX, UK. Email: a.carte@bbk.ac.uk

National Grid Company which operates a balancing mechanism to ensure system security.¹ Moreover, although power markets may have certain similarities with other markets, they present intrinsic characteristics which distinguish them. Two distinctive features are present in energy markets in general, and are very evident in electricity markets in particular: the mean reverting nature of spot prices and the existence of jumps or spikes in the prices.

In stock markets, prices are allowed to evolve 'freely', but this is not true for electricity prices; these will generally gravitate around the cost of production. Under abnormal market conditions, price spreads are observed in the short run, but in the long run supply will be adjusted and prices will move towards the level dictated by the cost of production. This adjustment can be captured by mean-reverting processes, which in turn may be combined with jumps to account for the observed spikes.

Therefore, to price energy derivatives it is essential that the most important characteristics of the evolution of the spot, and consequently the forward, are captured. Several approaches may be taken, generally falling into two classes of models: spot-based models and forward-based models. Spot models are appealing since they tend to be quite tractable and also allow for a good mathematical description of the problem in question. Significant contributions have been made by Schwartz (1997), for instance, where the author introduces an Ornstein-Uhlenbeck type of model which accounts for the mean reversion of prices. Lucia and Schwartz (2002) extend the range of these models to two-factor models which incorporate a deterministic seasonal component. On the other hand forward-based models have been used largely in the Nord Pool Market of the Scandinavian countries. These rely heavily, however, on a large data set, which is a limiting constraint in the case of England and Wales. Finally, it must also be pointed out that the choice of model may sometimes be driven by what kind of information is required. For example, pricing interruptible contracts would require a spot-based model while pricing Asian options on a basket of electricity monthly and seasonal forwards calls for forward-based models.

The spot models described by Schwartz (1997) and Lucia and Schwartz (2002) capture the mean reverting nature of electricity prices, but they fail to account for the huge and non-negligible observed spikes in the market. A natural extension is then to incorporate a jump component in the model. This class of jump-diffusion models was first introduced by Merton (2001) to model equity dynamics. Applying these jump-diffusion-type models in electricity is attractive since solutions for the pricing of European options are available in closed-form. Nevertheless, it fails to incorporate both mean reversion and jump diffusion at the same time. Clewlow *et al.* (2001) describe an extension to Merton's model which accounts for both the mean reversion and the jumps but they do not provide a closed-form solution for the forward. A similar model to the one we present, although not specific to the analysis of electricity spot prices, has been analysed in Benth *et al.* (2003).

The main contribution of this paper is twofold. First, we present a model that captures the most important characteristics of electricity spot prices such as mean reversion, jumps and seasonality and calibrate the parameters to the England and Wales market. Second, since we are able to calculate an expression for the forward curve in closed-form and recognizing the lack of sufficient data for robust parameter

estimation, we estimate the model parameters exploiting the fact that we can use both historical spot data and current forward prices (using the closed-form expression for the forward).²

The remaining of this paper is structured as follows. In Section 2 we present data analysis to support the use of a model which incorporates both mean reversion and jumps. In Section 3 we present details of the spot model and derive in closed-form the expression for the forward curve. In Section 4 we discuss the calibration of the model to data from England and Wales. In Section 5 we present forward surfaces reflecting the months, quarters and seasons-ahead prices. Section 6 concludes.

2. Data Analysis

For over three decades most equity models have tried to ‘fix’ the main drawback from assuming Gaussian returns. A clear example is the wealth of literature that deals with stochastic volatility models, jump-diffusion and more recently, the use of Lévy processes. One of the main reasons to adopt these alternative models is that Gaussian shocks attach very little probability to large movements in the underlying that are, on the contrary, frequently observed in financial markets. In this section we will see that in electricity spot markets assuming Gaussian shocks to explain the evolution of the spot dynamics is even a poorer assumption than in equity markets.

Electricity markets exhibit their own intrinsic complexities. There is a strong evidence of mean reversion and of spikes in spot prices, which in general are much more pronounced than in stock markets. The former can be observed by simple inspection of the data in both markets. Figure 1 shows daily closes of the FTSE100 index from 2/01/90 to 18/06/04. The nature of the price path can be seen as a combination of a deterministic trend together with random shocks. In contrast, Figure 2 shows that for electricity spot prices in England and Wales there is a strong mean reversion.³ This is, prices tend to oscillate or revert around a mean level, with extraordinary periods of volatility. These extraordinary periods of high volatility are reflected in the characteristic spikes observed in these markets.

Normality Tests

In the Black–Scholes model prices are assumed to be log-normally distributed, which is equivalent to saying that the returns of the prices have a Gaussian or Normal distribution.⁴ Although fat tails are observed in data from stock markets, indicating the probability of rare events being more frequent than predicted by a Normal distribution, models based on this assumption have been largely used as a benchmark, albeit modified in order to account for fat tails.

For electricity though, the departure from Normality is more extreme. Figure 3 shows a Normality test for the electricity spot price from 2/04/01 to 3/03/04. If the returns were indeed Normally distributed the graph would be a straight line. We can clearly observe this is not the case, as evidenced from the fat tails. For instance, corresponding to a probability of 0.003 we have returns which are higher than -0.5 ; instead if the data were perfectly Normally distributed, the dotted lines suggests the probability of such returns should be virtually zero.

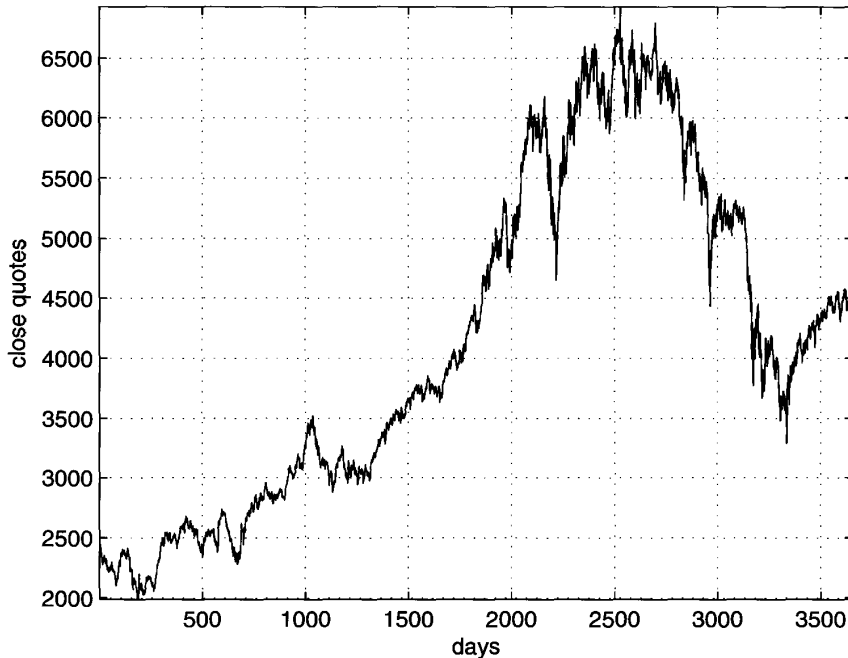


Figure 1. FTSE100 daily closes from 2/01/90 to 18/06/04

Deseasonalization

One important assumption of the Black–Scholes model is that returns are assumed to be independently distributed. This can be easily evaluated with an autocorrelation test. If the data were in fact independently distributed, the correlation coefficient would be close to zero. A strong level of autocorrelation is evident in electricity markets, as can be seen from Figure 4. As explained for instance in Pindyck and Rubinfeld (1998) the evidence of autocorrelation manifests an underlying seasonality. Furthermore, the lag of days between highly correlated points in the series reveals the nature of the seasonality. In this case, we may observe that the returns show significant correlation every 7 days (there is data for weekends also); which suggests some intra-week seasonality.

In order to estimate the parameters of the model, we strip the returns from this seasonality. Although there are several ways of deseasonalizing the data, we follow a common approach which is to subtract the mean of every day across the series according to

$$R_t = r_t - \bar{r}_d \quad (1)$$

where R_t is the defined deseasonalized return at time t , r_t the return at time t and \bar{r}_d is the corresponding mean (throughout the series) of the particular




Figure 2. Averaged daily prices in England and Wales from 2/04/01 to 3/03/04

day r_t represents. Figure 5 shows the autocorrelation test performed on the deseasonalized returns. As expected, the strong autocorrelation is no longer evidenced.

Jumps

As seen from the Normality test, the existence of fat tails suggest the probability of rare events occurring is actually much higher than predicted by a Gaussian distribution. By simple inspection of Figure 2 we can easily be convinced that the spikes in electricity data cannot be captured by simple Gaussian shocks.

We extract the jumps from the original series of returns by writing a numerical algorithm that filters returns with absolute values greater than three times the standard deviation of the returns of the series at that specific iteration.⁵ On the second iteration, the standard deviation of the remaining series (stripped from the first filtered returns) is again calculated; those returns which are now greater than three times this last standard deviation are filtered again. The process is repeated until no further returns can be filtered. This algorithm allows us to estimate the cumulative frequency of jumps and other statistical information of relevance for calibrating the model.⁶

The relevance of the jumps in the electricity market is further demonstrated by comparing Figure 6 to Figure 3; where we can clearly observe that after stripping the returns from the jumps, the Normality test improves notoriously.




Figure 3. Normal probability test for returns of electricity prices from 2/04/01 to 3/03/04




Figure 4. Autocorrelation test for returns of electricity prices from 2/04/01 to 3/03/04




Figure 5. Autocorrelation test for deseasonalized returns of electricity prices




Figure 6. Normal probability test for filtered returns of electricity prices

3. The Model: Mean-reversion and Jump Diffusion in the Electricity Spot

When modelling the electricity market two distinct approaches may be taken: modelling the spot market or modelling the entire forward curve. As mentioned earlier, one of the appeals for using spot models relies on the fact that it is simple to incorporate the observed characteristics of the electricity market. On the other hand, forward based models rely more heavily on the amount of historical data available. Since data for electricity prices in England and Wales are only regarded to be liquid and 'well priced' since the incorporation of NETA on 27 March 2001, the amount of data available is limited. This lack of sufficient data motivates the use of spot-based models rather than modelling the entire forward curve in the particular case of this market. It is worth emphasising that different power markets, although similar in some aspects, exhibit their own properties and characteristics. Hence, based on the manifest existence of mean-reversion and jumps on the data for England and Wales presented in the previous section, we propose a one-factor mean-reversion jump diffusion model; adjusted to incorporate seasonality effects.

Electricity can be bought in the spot market, but once purchased it must be used almost immediately, since in most cases electricity cannot be stored, at least not cheaply. Hedging strategies which typically involve holding certain amounts of the underlying (in this case electricity) are not possible, therefore in electricity markets forwards on the spot are typically used instead. As a consequence, it turns out it is extremely useful to be able to extract a closed-form formula for the forward curve from the spot-based model, which we are able to do for the model proposed here.

From the data analysis of the previous section we have concluded that two distinctive characteristics of electricity markets should be accounted for in the model; the mean reversion of the price and the sudden fluctuations in supply and low elasticity in demand which are reflected in price spikes. Moreover, it would also be important to incorporate some seasonality component which would be reflected in a varying long-term level of mean reversion.

Schwartz (1997) accounts for the mean reversion, and Lucia and Schwartz (2002) extend the mean reverting model to account for a deterministic seasonality. However, these models do not incorporate jumps. We propose in this paper a similar model extended to account for the observed jumps.

As in Lucia and Schwarz (2002) let us assume that the log-price process, $\ln S_t$, can be written as

$$\ln S_t = g(t) + Y_t \quad (2)$$

such that the spot price can be represented as

$$S_t = G(t)e^{Y_t} \quad (3)$$

where $G(t) \equiv e^{g(t)}$ is a deterministic seasonality function and Y_t is a stochastic process whose dynamics are given by

$$dY_t = -\alpha Y_t dt + \sigma(t)dZ_t + \ln J dq_t \quad (4)$$

In (4) Y_t is a zero level mean-reverting jump diffusion process for the underlying electricity spot price S_t , α is the speed of mean reversion, $\sigma(t)$ the time dependent

volatility, J the proportional random jump size, dZ_t is the increment of the standard Brownian motion and dq_t a Poisson process such that

$$dq_t = \begin{cases} 1 & \text{with probability } l dt \\ 0 & \text{with probability } 1 - l dt \end{cases} \quad (5)$$

where l is the intensity or frequency of the process.⁷ Moreover, J , dq_t and dZ_t are independent.

Regarding the jump size, J , the following assumptions are made:

- J is log-Normal, i.e. $\ln J \sim N(\mu_J, \sigma_J^2)$.
- The risk introduced by the jumps is non-systematic and so diversifiable; furthermore, by assuming $\mathbb{E}[J] \equiv 1$ we guarantee there is no excess reward for it.

With the assumptions made above, the properties of J can be summarized as follows:

$$J = e^\phi, \phi \sim N\left(-\frac{\sigma_J^2}{2}, \sigma_J^2\right) \quad (6)$$

$$\mathbb{E}[J] = 1 \quad (7)$$

$$\mathbb{E}[\ln J] = -\frac{\sigma_J^2}{2} \quad (8)$$

$$\text{Var}[\ln J] = \sigma_J^2 \quad (9)$$

Now, from (3) and (4) we can write the SDE for S_t , namely

$$dS_t = \alpha(\rho(t) - \ln S_t)S_t dt + \sigma(t)S_t dZ_t + S_t(J - 1)dq_t \quad (10)$$

where the time-dependent mean reverting level is given by

$$\rho(t) = \frac{1}{\alpha} \left(\frac{dg(t)}{dt} + \frac{1}{2} \sigma^2(t) \right) + g(t) \quad (11)$$

The interpretation of (10) is as follows. Most of the time $dq_t=0$, so we simply have the mean reverting diffusion process. At random times however, S_t will jump from the previous value S_{t-} to the new value JS_{t-} . Therefore the term $S_{t-}(J - 1)$ gives us the change after and before the jump, $\Delta S_t = JS_{t-} - S_{t-}$.

Forward Price

The price at time t of the forward expiring at time T is obtained as the expected value of the spot price at expiry under an equivalent \mathcal{Q} -martingale measure, conditional on the information set available up to time t ; namely

$$F(t, T) = \mathbb{E}_t^{\mathcal{Q}}[S_T | \mathcal{F}_t] \quad (12)$$

Thus, we need to integrate first the SDE in (10) in order to extract S_T and later calculate the expectation.

For the first task we define the log-return as $x_t \equiv \ln S_t$ and apply Itô's Lemma to (10) to arrive at

$$dx_t = \alpha(\mu(t) - x_t)dt + \sigma(t)dZ_t + \ln J dq_t \quad (13)$$

where

$$\mu(t) = \frac{1}{\alpha} \frac{dg}{dt} + g(t) \quad (14)$$

is the time-dependent mean reverting level which depends on the seasonality function.

Regarding the expectation, we must calculate it under an equivalent \mathcal{Q} -martingale measure. In a complete market this measure is unique, ensuring only one arbitrage-free price of the forward. However, in incomplete markets (such as the electricity market) this measure is not unique, thus we are left with the difficult task of selecting an appropriate measure for the particular market in question. Yet another approach, common in the literature, is simply to assume that we are already under an equivalent measure, and thus proceed to perform the pricing directly. This latter approach would rely however on calibrating the model through implied parameters from a liquid market. This is certainly difficult to do in young markets, as in the market of electricity in England and Wales, where there is no liquidity of instruments which would enable us to do this.

We follow instead Lucia and Schwartz' (2002) approach, which consists of incorporating a market price of risk in the drift, such that $\hat{\mu}(t) \equiv \mu(t) - \lambda^*$ and $\lambda^* \equiv \lambda \frac{\sigma(t)}{\alpha}$, where λ denotes the market price of risk per unit risk linked to the state variable x_t . This market price of risk, to be calibrated from market information, pins down the choice of one particular martingale measure. Under this measure we may then rewrite the stochastic process in (13) for x_t as

$$dx_t = \alpha(\hat{\mu}(t) - x_t)dt + \sigma(t)d\hat{Z}_t + \ln J dq_t \quad (15)$$

where

$$\hat{\mu}(t) = \frac{1}{\alpha} \frac{dg}{dt} + g(t) - \lambda \frac{\sigma(t)}{\alpha} \quad (16)$$

and $d\hat{Z}_t$ is the increment of a Brownian motion in the \mathcal{Q} -measure specified by the choice of λ .⁸

In order to integrate the process we multiply (15) by a suitable integrating factor and integrate between times t and T to arrive at

$$\begin{aligned}
x_T &= \overline{g}(T) + (x_t - g(t))e^{-\alpha(T-t)} - \lambda \int_t^T \sigma(s)e^{-\alpha(T-s)} ds \\
&+ \int_t^T \sigma(s)e^{-\alpha(T-s)} d\hat{Z}_s + \int_t^T e^{-\alpha(T-s)} \ln J dq_s
\end{aligned} \tag{17}$$

Now, since $S_T = e^{x_T}$, we can replace (17) into (12) to obtain

$$\begin{aligned}
F(t, T) &= \mathbb{E}_t[S_T | \mathcal{F}_t] \\
&= \hat{\lambda}_t^T G(T) \left(\frac{S(t)}{G(t)} \right)^{e^{-\alpha(T-t)}} \mathbb{E}_t \left[e^{\int_t^T \sigma(s)e^{-\alpha(T-s)} d\hat{Z}_s} | \mathcal{F}_t \right] \mathbb{E}_t \left[e^{\int_t^T e^{-\alpha(T-s)} \ln J dq_s} | \mathcal{F}_t \right] \\
&= \hat{\lambda}_t^T G(T) \left(\frac{S(t)}{G(t)} \right)^{e^{-\alpha(T-t)}} e^{\frac{1}{2} \int_t^T \sigma(s)^2 e^{-2\alpha(T-s)} ds} \mathbb{E}_t \left[e^{\int_t^T e^{-\alpha(T-s)} \ln J dq_s} | \mathcal{F}_t \right]
\end{aligned} \tag{18}$$

where $\hat{\lambda}_t^T \equiv e^{-\lambda \int_t^T \sigma(s)e^{-\alpha(T-s)} ds}$ and expectations are taken under the risk-neutral measure. In the Appendix we prove that the expectation in (18) is

$$\mathbb{E}_t \left[e^{\int_t^T e^{-\alpha(T-s)} \ln J dq_s} | \mathcal{F}_t \right] = \exp \left[\int_t^T e^{-\frac{\sigma^2}{2} e^{-\alpha(T-s)} + \frac{\sigma^2}{2} e^{-2\alpha(T-s)}} l ds - (T-t)l \right] \tag{19}$$

Finally, replacing $\hat{\lambda}_t^T$ and (19) into (18) we obtain the price of the forward as

$$F(t, T) = G(T) \left(\frac{S(t)}{G(t)} \right)^{e^{-\alpha(T-t)}} e^{\int_t^T \left[\frac{1}{2} \sigma^2(s) e^{-2\alpha(T-s)} - \lambda \sigma(s) e^{-\alpha(T-s)} \right] ds + \int_t^T \zeta(\sigma_J, \alpha, T, s) l ds - l(T-t)} \tag{20}$$

where $\zeta(\sigma_J, \alpha, T, s) \equiv e^{-\frac{\sigma_J^2}{2} e^{-\alpha(T-s)} + \frac{\sigma_J^2}{2} e^{-2\alpha(T-s)}}$

4. Calibration

One of the arguments in favour of spot-based models is that they can provide a reliable description of the evolution of electricity prices. Moreover, these models are versatile in the sense that it is relatively simple to aggregate ‘characteristics’ to an existing family or class of models like for example adding a seasonality function. On the other hand, one of the drawbacks of these models is that it is quite difficult to estimate parameters given the relatively large number of parameters combined with a very small sample data (Clewlow and Strickland, 2000; Eydeland and Wolyniec, 2003; Knittel and Roberts, 2001).

One approach is to estimate all the parameters involved from historical data using maximum likelihood estimators (MLE) through the approximations presented by Ball and Torous (1983, 1985).⁹ However, for the England and Wales data this method yielded incorrect estimates, i.e. negative values for certain parameters that should otherwise be positive and estimates which depended heavily on the starting value of the parameters. We believe this is mainly due to the scarcity of data in this market.

As an alternative we propose a ‘hybrid’ approach that uses both historical spot data and forward market data. The former is used to calculate the seasonality component, the rolling historical volatility, the mean reversion rate and the frequency and standard deviation of the jumps.¹⁰ The latter is used to estimate the market price of risk.

Spot-based Estimates

Seasonality function. In (3), $G(t)$ is a deterministic function which accounts for the observed seasonality in power markets. The form of this seasonality function inevitably depends on the market in question. For instance, some electricity markets will exhibit a discernible pattern between summer and winter months. In such cases a sinusoidal function could be suitable (as suggested by Pilipović (1998)). Other alternatives include a constant piece-wise function, as for instance in Knittel and Roberts (2001). Furthermore, Lucía and Schwartz (2002) introduce a deterministic function which discerns between weekdays and a monthly seasonal component.

However sophisticated these functions may be, they all rely on the inclusion of dummy variables and on being able to calibrate them correctly from the sample of historical data. As discussed earlier, this might be a serious constraint when dealing with markets with not enough historical data. Moreover, although it is reasonable to assume that there might be a distinguishable pattern between summer and winters in England and Wales, this is yet not evident from the available data.

Hence, including a seasonality function dependent on parameters to estimate from historical data would only add difficulty and unreliability to the already difficult calibration of the model. Instead, we have chosen to introduce a deterministic seasonality function which is a fit of the monthly averages of the available historical data with a Fourier series of order 5. In this way, we introduce a seasonality component into the model, but do not accentuate even further the problems involved in the calibration.¹¹ The seasonality function is shown in Figure 7.

Rolling historical volatility. It can easily be shown that volatility is not constant across time in electricity markets. One common approach then, is to use as an estimate a rolling (or moving) historical volatility, as described in Eydeland and Wolyniec (2003) for instance. In this case, we use a yearly averaged rolling historical volatility with a window of 30 days.

Mean reversion rate. The mean reversion is usually estimated using a linear regression. In this case we regressed the returns Δx_t versus the series of returns x_t of the log-spot price.

Jump parameters. In order to estimate the parameters of the jump component of the spot dynamics, we filtered the data of returns using the code that was explained earlier. As an output of the code, we estimated the standard deviation of the jumps, σ_J , and the frequency of the jumps, l , which is defined as the total number of jumps divided by the annualised number of observations.




Figure 7. Seasonality function based on historical averaged months

Forward-based Estimate

We estimate the remaining parameter, the market price of risk λ , by minimising the square distances of the theoretical forward curve for different maturities (obtained through (20)) to given market prices of equal maturities.¹²

Results

The results obtained are summarized in Table 1.

Based on the result obtained for the standard deviation of the jumps through the filtering process discussed previously we could not conclude that the relationship imposed in the model between the mean and the variance of the logarithm of the jumps (through (8) and (9)) holds in each iteration. However, as mentioned earlier, this condition can easily be relaxed. This would lead, nonetheless, to the inclusion of an extra parameter to be estimated (the mean of the logarithm of the jumps). At this point, one must compromise between the imposed assumptions and the feasibility of calibrating a model dependent on too many parameters. The estimated frequency of

Table 1. Annualized estimates for the standard deviation of the jumps σ_J , frequency of the jumps l , mean reversion rate α and average (denoted by [.]) market price of risk per unit risk λ^* . When available, the 95% confidence bounds are presented in parenthesis

σ_J	l	α	$[\lambda^*]$ (%)
0.67	8.58	0.2853 (0.2431, 0.3274)	-0.2481 (-0.2550, -0.2413)

jumps suggests that there are between eight and nine jumps per year, which is in agreement with observed historical data.

The estimated mean reversion rate represents a daily estimate. To interpret what the estimated value of mean reversion implies, let us re-write (15) in an Euler discretized form in a period Δt where no random shocks or jumps have occurred, namely

$$x_{t+1} = \alpha(\hat{\mu}(t) - x_t)\Delta t + x_t \quad (21)$$

We may easily see that when we multiply the daily estimate by the appropriate annualization factor (in this case 365), and since $\Delta t = 1/365$; when $\alpha = 1$ we have $x_{t+1} = \hat{\mu}(t)$.

This is, when $\alpha = 1$ the process mean reverts to its equilibrium level over the next time step. In our case, the estimated parameter suggests it mean reverts very rapidly, in 0.84 days, this is, almost in a day. This is not surprising in electricity markets, and may already be inferred by the nature of the spot price series, as seen in Figure 2.

Escribano *et al.* (2002) and Knittel *et al.* (2001) have extensively calibrated mean reverting jump diffusion models to electricity data for different markets. In both cases, they calibrate discrete-time parameters. The connection between the time continuous parameters and the discrete version can be seen by writing the MRJD process defined in (15) as

$$x_t = \theta_t + \beta x_{t-1} + \eta_t \quad (22)$$

where

$$\theta_t \equiv \hat{\mu}(t)(1 - e^{-\alpha}) \quad \text{and} \quad \beta \equiv e^{-\alpha} \quad (23)$$

and η_t represents the integral of the Brownian motion and the jump component between times $t-1$ and t .

From (23) we may recover the discrete-time parameter corresponding to the mean reversion rate, which gives $\beta = 0.7518$; which is such that $|\beta| < 1$, guaranteeing that the process mean reverts back to its non-constant mean. Moreover, our estimate of β is entirely compatible with the estimates presented in (1) and (12) for different electricity markets.

Finally, let us interpret the results obtained for the market price of risk per unit risk. In Table 1 we have shown an average value of the market price of risk; the average value results from taking an average historical volatility of the rolling volatility we have estimated.¹³ Through (15) and (16) we note that the drift of x_t is given mainly by $\alpha(g(t) - \lambda^*)$, since the term $\frac{1}{\alpha} \frac{dg}{dt}$ is practically zero. Hence, the drift is being pushed upwards by our market price of risk. The fact that this market price of risk is negative, does not seem uncommon in some energy markets, and in electricity markets in particular. In fact, Botterud *et al.* [6] make an empirical study of the risk premium in the Scandinavian electricity market and find negative values for their estimates. They explain the risk premium in terms of the difference in the number of participants on the supply and demand sides. In this context, a negative risk premium would be consequence of an excess demand for futures contracts.

5. Applications

Pricing a European call option on a forward was first addressed by Black in 1976. Based solely on arbitrage arguments one can obtain the price of a forward contract under a GBM very easily, simple arguments then lead to a closed-form solution for a call option written on a forward, which is widely known as Black's formula.¹⁴

However, when departing from the very idealized GBM, and incorporating both mean reversion and jump-diffusion to the process; closed-form solutions are very hard, if possible at all, to obtain. Duffie *et al.* (2000) are able to extract semi-closed-form solutions provided that the underlying follows an affine jump-diffusion (AJD); which they define basically as a jump-diffusion process in which the drift vector, instantaneous covariance matrix and jump intensities all have affine dependence on the state vector.

On the other hand, without imposing these dependencies, a closed-form analytical solution might prove significantly harder to obtain. Hence, the pricing of these models is generally done numerically. Regardless of the numerical method employed, ultimately the performance of the model relies on the capability of successfully capturing the discussed characteristics of this market.

For instance, the model (once calibrated) must yield price paths for the price of electricity which resemble those observed in the market. In Figure 8 we show a simulated random walk which results from discretizing (13) and later recovering the spot price as $S_t = e^{S_t}$; subject to the calibration discussed in the preceding section. Here we observe that the price path succeeds in capturing the mean reversion and incorporating the jumps, which are mostly (as desired) upwards. Moreover, the monthly averages of the simulated price path closely resembles the seasonality function, which is evidence that the process is mean reverting towards a time-dependent equilibrium level dictated mainly by the seasonality function, as expected from (14).

In order to further test the validity of the model, we show in Figure 9 the calibrated forward curve with its 95% confidence interval, the averaged months from the calibrated forward curve and the monthly market forwards. We can observe that the forward curve sticks on average to levels close to the market curve; albeit showing a great degree of flexibility. By this we mean that the curve exhibits all the variety of shapes observed commonly in the market; which are commonly known as *backwardation* (decrease in prices with maturity), *contango* (increase in prices with maturity) and *seasonal* (a combination of both).

In Figure 10 we show a forward surface for 5 months ahead.¹⁵ To understand this graph better, let us concentrate on the first month of July. For each day in June 2004 we calculate the forwards with starting date t_i , $i \in (1, 30)$ with maturities T_k , $k \in (1, 31)$; where i sweeps across the days in June and k across maturities in July. The forward for each day in June then is calculated as the average of the forwards of maturity T_k , thus reflecting the price of a forward contract of electricity for the entire month of July, as quoted on the i th day of June. Similarly, in Figures 11 and 12 we show forward surface for quarters and seasons ahead.

As can be seen from Figure 10 for instance, the surface evolves in accordance to the monthly seasonalities, sticking to higher prices towards the end of the winter of 2004. This is again observed in Figure 11, where the prices for quarter 4, 2005 are




Figure 8. Simulated price path




Figure 9. Optimized forward curve: the circles represent the forward; the lower triangles the upper bound of the estimated forward; the upper triangles the lower bound of the estimated forward; the solid line the monthly average of the estimated forward and the dotted line the market forward

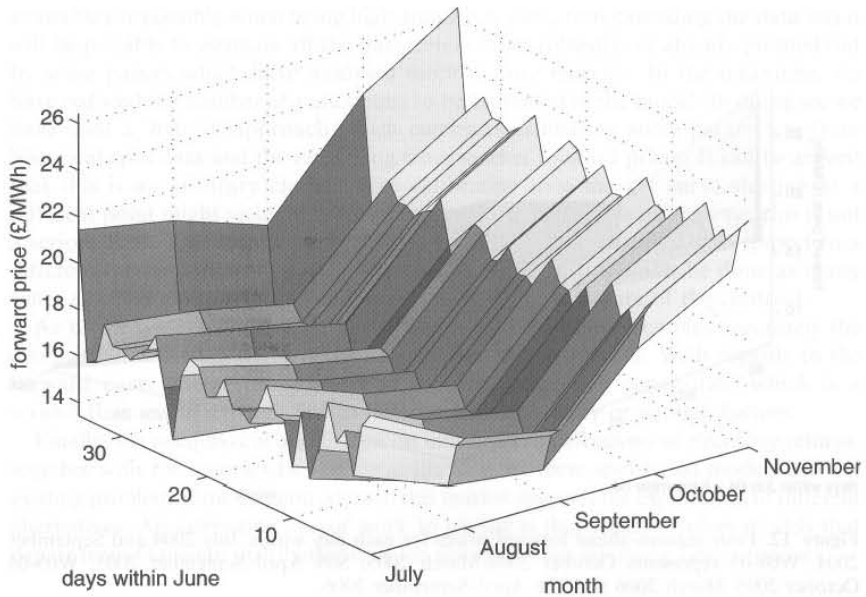


Figure 10. 5-months-ahead forward prices for each day in June 2004

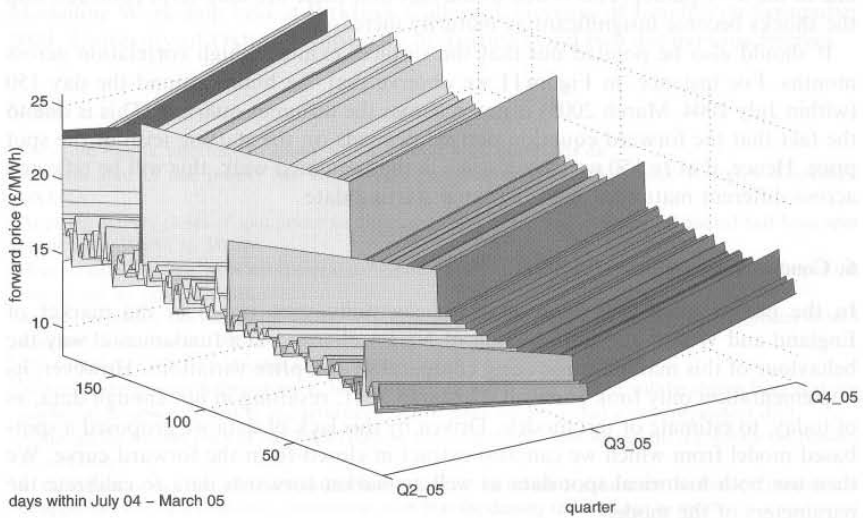


Figure 11. Three-quarters-ahead forward prices for each day within July 2004 and March 2005. Q2_05 represents April–June 2005, Q3_05 July–September 2005 and Q4_05 October–December 2005

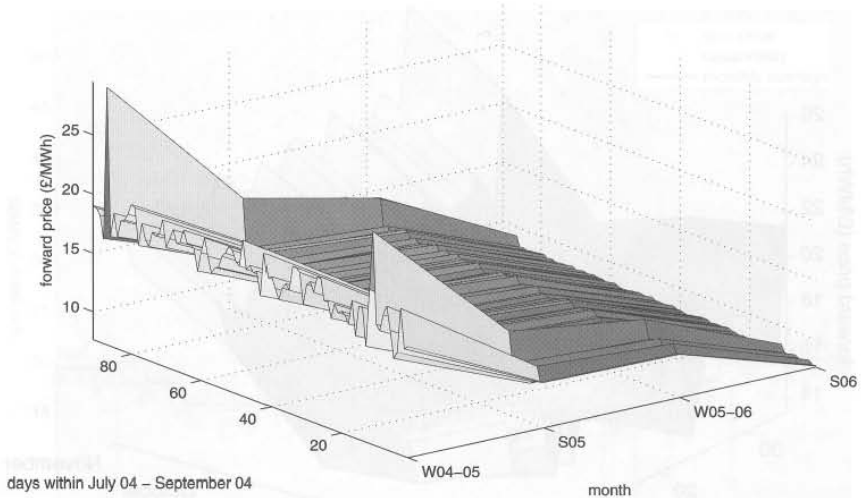


Figure 12. Four-seasons-ahead forward prices for each day within July 2004 and September 2004. W04-05 represents October 2004–March 2005; S05 April–September 2005, W05-06 October 2005–March 2006 and S06 April–September 2006

higher, as expected. In Figure 12, we observe that for the second and third season ahead the calculated forward price exhibits little variation (seen as an almost straight line in the x - y plane). This is due to the fact that these are long-term contracts and the shocks become insignificant as maturity increases.

It should also be pointed out that the surfaces exhibit a high correlation across months. For instance, in Figure 11 we observe that the hump around the day 150 (within July 2004–March 2005) is noted across the different quarters. This is due to the fact that the forward equation derived depends on the starting level of the spot price. Hence, if at $t=150$ we have a spike in the simulated walk, this will be reflected across different maturities with the same starting date.

6. Conclusions

In the present paper we have analysed electricity spot prices in the market of England and Wales. The introduction of NETA changed in a fundamental way the behaviour of this market introducing competition and price variations. However, its implementation only took place on 27 March 2001, resulting in not enough data, as of today, to estimate or test models. Driven by this lack of data we proposed a spot-based model from which we can also extract in closed-form the forward curve. We then use both historical spot data as well as market forwards data to calibrate the parameters of the model.

Regarding the calibration of the model, we have circumvented a known drawback in electricity spot-based models, which is the overwhelming dependence on a great number of parameters to estimate. As the market evolves and more data becomes

available (or possibly when using high-frequency data, thus extending the data set) it will be possible to estimate all the parameters more robustly; as already pointed out by some papers which have analysed more mature markets. In the meantime, we have reduced the number of parameters to be estimated in the model. In doing so, we have used a 'hybrid' approach which combines estimating some parameters from historical spot data and the remaining from market forward prices. It can be argued that this is an arbitrary choice, since calibrating to a market curve starting at a different point might yield different parameters. Even if this were the case, this is not a serious flaw. This would imply re-calibrating the forward curve with respect to a different market curve. In a dynamic hedging-strategy this could be done as many times as necessary, depending on the exposure and the nature of the contract.

As to the output of the model, the simulated price path resembles accurately the evolution of electricity spot prices as observed in this market. With regards to the forward curve shown, it succeeds in capturing changing convexities, which is a serious flaw in models that fail to incorporate seasonality or enough factors.

Finally, the unequivocal evidence of fat tails in the distributions of electricity returns, together with the complexities on the calibration of these spot-based models and the existing problem of the exiguous data in this market suggests the exploration of different alternatives. An interesting line of work to pursue is that which involves models that depart from Gaussian distributions, as for instance those involving Lévy processes.

Acknowledgements

We are grateful for comments from Andrew Shaw, Gareth Davies, Murray Hartley, Raymond Brummelhuis, Ron Smith, seminar participants at the Commodities Modelling Workshop held at Birkbeck College, University of London on September 2004, University of Oxford, University of Toronto, and two anonymous referees.

Notes

¹ For more specific information about NETA consult www.ofgem.gov.uk.

² All data used in this project has been kindly provided by Oxford Economic Research Associates, OXERA.

³ As proxy to daily closes of spot prices we have used the daily average of historical quoted half-hour spot prices from 2/04/01 to 3/03/04.

⁴ Here we define 'return' as in the classical definition; $r_t = \ln(S_{t+1}/S_t)$. Note that this is also referred to as the 'log-return' by other authors.

⁵ As can be readily calculated, the probability in a Normal distribution of having returns greater than 3 standard deviations is 0.0027.

⁶ The calibration will be addressed in Section 4.

⁷ Although the process followed by Y_t mean reverts around a zero level, it will be shown later that the stochastic process followed by the returns of S_t will mean revert around a time-dependent drift.

⁸ Although the market price of risk itself could be time-dependant, here we assume it constant for reasons of simplicity.

⁹ In these papers they demonstrate that for low values of the intensity parameter the Poisson process can be approximated by a Bernoulli distribution, such that the density function can be written as a mixture of Normals.

¹⁰ By the restriction imposed in (7) we have reduced the need to calibrate the mean of the jumps in the spot.

¹¹ Although in electricity there is also evidence of intra-day seasonality, it is not necessary to account for it in this model since we take as spot prices the average of intra-day half-hour prices. The weekly pattern of

seasonality however, could be accounted for, albeit at the cost of introducing yet another parameter into the model.

- ¹² The market quotes in reference were obtained from 'Argus' and represent forward prices at 7 May 2004 for the next six months.
- ¹³ When reconstructing the spot and forward prices however, we multiply at each time-step the market price of risk by the appropriate volatility at that time, as indicated by (16).
- ¹⁴ This derivation can be found in many textbooks, for a simple and intuitive explanation see Bjork (2004).
- ¹⁵ The forward surfaces have been calculated with a considerably lower mean reversion rate in order to capture the dynamics of larger maturities more realistically.
- ¹⁶ See pages 176–177 in this reference for more details.

References

- Ball, C. A. and Torous, W. N. (1983) A simplified jump process for common stock returns, *The Journal of Finance and Quantitative Analysis*, 18(1), pp. 53–65.
- Ball, C. A. and Torous, W. N. (1985) On jumps in common stock prices and their impact on call option pricing, *The Journal of Finance*, 40(1), pp. 155–173.
- Benth, F. E., Ekeland, L., Hauge, R. and Nielsen, B. F. (2003) A note on arbitrage-free pricing of forward contracts in energy markets, *Applied Mathematical Finance*, 10(4), pp. 325–336.
- Björk, T. (2004) *Arbitrage Theory in Continuous Time* (London: Oxford University Press).
- Botterud, A., Bhattacharyya, A. K. and Ilic, M. (2002) Futures and spot prices - an analysis of the Scandinavian electricity market, *mimeo*, Massachusetts Institute of Technology and Norwegian University of Science and Technology.
- Clelow, L. and Strickland, C. (2000) *Energy Derivatives, Pricing and Risk Management* (Lacima Publications).
- Clelow, L., Strickland, C. and Kaminski, V. (2001) Extending mean-reversion jump diffusion, *Energy Power Risk Management, Risk Waters Group*, February.
- Duffie, D., Pan, J. and Singleton, K. (2000) Transform analysis and asset pricing for affine jump-diffusions, *Econometrica*, 68(6), pp. 1343–1376.
- Escribano, Á., Peña, J. I. and Villaplana, P. (2002) Modelling electricity prices: international evidence, Universidad Carlos III de Madrid – Working paper 02–27, June.
- Etheridge, A. (2002) *A Course in Financial Calculus*, 1st edn (Cambridge University Press).
- Eydeland, A. and Wolyniec, K. (2003) *Energy and Power Risk Management*, 1st edn (John Wiley & Sons).
- Knittel, C. R. and Roberts, M. R. (2001) An empirical examination of deregulated electricity prices, POWER working paper PWP-087.
- Lucia, J. J. and Schwartz, E. S. (2002) Electricity prices and power derivatives: evidence from the nordic power exchange, *Review of Derivatives Research*, 5, pp. 5–50.
- Merton, R. C. (2001) *Continuous-Time Finance*, 1st revised edn (Blackwell).
- Pilipović, D. (1998) *Energy Risk, Valuing and Managing Energy Derivatives* (Mc Graw-Hill).
- Pindyck, R. S. and Rubinfeld, D. L. (1998) *Econometric Models and Economic Forecasts*, 4th edn (McGraw-Hill).
- Schwartz, E. S. (1997) The stochastic behavior of commodity prices: implications for valuation and hedging, *The Journal of Finance*, 52(3), pp. 923–973.
- Schwartz, E. S. and Smith, J. E. (2000) Short-term variations and long-term dynamics in commodity prices, *Management Science*, 46(7), pp. 893–911.

Appendix A Proof of Expected Value in Forward Equation

We want to evaluate

$$\mathbb{E}_T \left[e^{\int_t^T e^{-\alpha(T-s)} \ln J_s dq_s} \right] = \mathbb{E}_T \left[e^{\int_t^T \alpha_s dq_s} \right] \quad (\text{A } 1)$$

where

$$\alpha_s \equiv e^{-\alpha(T-s)} \ln J_s \quad (\text{A } 2)$$

We will first calculate (A1) in the interval $[0, t]$ to later extend the calculation to the interval $[t, T]$.

Let us start by defining L_t such that

$$\begin{aligned} L_t &\equiv e^{\int_0^t \alpha_s dq_s} \\ &\equiv e^{m_t} \end{aligned} \quad (\text{A } 3)$$

where m_t is then

$$m_t = \int_0^t \alpha_s dq_s \quad (\text{A } 4)$$

and equivalently

$$dm_t = \alpha_t dq_t \quad (\text{A } 5)$$

In order to write the SDE followed by L_t for the process defined in (A 5) we need to generalize Itô's Lemma in order to incorporate the jumps. We will use the generalization followed by Etheridge (2002) to write the SDE followed by L_t as¹⁶

$$dL_t = \frac{\partial L_t(m_{t^-})}{\partial m_t} dm_t - \frac{\partial L_t(m_{t^-})}{\partial m_t} (m_t - m_{t^-}) dq + (L_t - L_{t^-}) dq \quad (\text{A } 6)$$

where we have not included any second derivative since the process defined by (A 5) is only a pure jump process.

In order to evaluate (A 6) let us first clarify the notation. If there is a jump in $\{m_t\}_{t>0}$ it is of size α_t and such that

$$m_t = m_{t^-} + \alpha_t \quad (\text{A } 7)$$

where if a jump takes place at time t , the time t^- indicates the time interval just before the jump has occurred.

Hence by (A 4) we can also write (A7) as

$$m_t = \int_0^t \alpha_s dq_s = \int_0^{t^-} \alpha_s dq_s + \alpha_t \quad (\text{A } 8)$$

Using (A 7) we can rewrite (A 3) as

$$\begin{aligned} L_t &= e^{m_t} \\ &= e^{m_{t^-} + \alpha_t} \\ &= L_{t^-} e^{\alpha_t} \end{aligned} \quad (\text{A } 9)$$

Noting that $\frac{\partial L_t(m_t^-)}{\partial m_t} = L_t^-$ and replacing back (A5), (A7) and (A9) into (A6) we get

$$dL_t = L_t^- (e^{z_t} - 1) dq_t \quad (\text{A } 10)$$

which we can integrate between 0 and t to obtain

$$L_t = 1 + \int_0^t L_s (e^{z_s} - 1) dq_s \quad (\text{A } 11)$$

where we have used that $L_0 = 1$.

By taking expectations of the above equation we arrive at

$$\mathbb{E}_0[L_t] = 1 + \int_0^t \mathbb{E}_0[L_s] (\mathbb{E}_0[e^{z_s}] - 1) l ds \quad (\text{A } 12)$$

where we are using the fact that $\mathbb{E}_0[dq] = l dt$ and l is the intensity of the Poisson process as had been defined in (5).

Defining now $\mathbb{E}_0[L_t] = n_t$ we can rewrite (A12) as

$$n_t = 1 + \int_0^t n_s (\mathbb{E}_0[e^{z_s}] - 1) l ds \quad (\text{A } 13)$$

which we can differentiate with respect to t to obtain

$$\frac{dn_t}{dt} = n_t (\mathbb{E}_0[e^{z_t}] - 1) l \quad (\text{A } 14)$$

Integrating now over the interval $[0, t]$ we get

$$\int_0^t \frac{dn_t}{n_t} = \int_0^t (\mathbb{E}_0[e^{z_s}] - 1) l ds \quad (\text{A } 15)$$

Finally, upon integrating and noting that $n_0 = L_0 = 1$ and replacing the definitions of n_t and L_t we obtain

$$\mathbb{E}_0 \left[e^{\int_0^t z_s dq_s} \right] = e^{\int_0^t (\mathbb{E}_0[e^{z_s}] - 1) l ds} \quad (\text{A } 16)$$

It is then straightforward to show that

$$\mathbb{E}_t \left[e^{\int_t^T z_s dq_s} \right] = e^{\int_t^T (\mathbb{E}_0[e^{z_s}] - 1) l ds} \quad (\text{A } 17)$$

which proves (A1).

Alternatively we can show the result in the following way. Note that the process $\int_0^t \ln J_s dq_s$ is a compound Poisson process, hence it is a Lévy process. Let $\int_0^t \ln J_s dq_s = \int_0^t dL_s$ with moment generating function, based on the Lévy–Khintchine representation,

$$\mathbb{E} [e^{\theta L_t}] = e^{t(\Psi_{\ln J}(\theta) - 1)} \quad (\text{A } 18)$$

where $\Psi_{\ln J}(\theta)$ is the moment generating function of the jumps $\ln J$. It is a well-known fact that for a deterministic function $f(t)$ and a Lévy process \tilde{L}_t the moment generating function of the process $\int_0^t f(s) d\tilde{L}_s$, when it exists, is given by

$$\mathbb{E} \left[e^{\theta \int_0^t f(s) d\tilde{L}_s} \right] = e^{\int_0^t \Psi(f(s)\theta) ds} \quad (\text{A } 19)$$

where $\Psi(\theta)$ is the log-moment generating function of the Lévy process \tilde{L}_t . Therefore

$$\mathbb{E} \left[e^{\theta \int_0^t e^{-\gamma(t-s)} \ln J dq} \right] = e^{\int_0^t (\Psi_{\ln J}(e^{-\gamma(t-s)}\theta) - 1) ds} \quad (\text{A } 20)$$

and by evaluating at $\theta=1$ delivers the desired result.

Evaluating the Integral

In order to evaluate (A17) we must calculate first the expected value of e^{z_s} . Thus, from (A2) we wish to calculate

$$\mathbb{E}_0[e^{z_s}] = \mathbb{E}_0 \left[e^{e^{-\alpha(T-s)} \ln J_s} \right] \quad (\text{A } 21)$$

and calling $h(s) = e^{-\alpha(T-s)}$ then

$$\begin{aligned} \mathbb{E}_0[e^{z_s}] &= \mathbb{E}_0 \left[e^{h(s) \ln J_s} \right] \\ &= \mathbb{E}_0 \left[e^{h(s)\phi} \right] \end{aligned} \quad (\text{A } 22)$$

since we had defined that the jumps J were drawn from a Normal distribution, and by requiring that $\mathbb{E}[J]=1$ we had that $\phi \sim N\left(-\frac{\sigma_J^2}{2}, \sigma_J^2\right)$, where σ_J is the standard deviation of the jumps.

Thus, (A22) yields

$$\mathbb{E}_0[e^{z_s}] = e^{-\frac{\sigma_J^2}{2} h(s) + \frac{\sigma_J^2}{2} h^2(s)} \quad (\text{A } 23)$$

and therefore (A17) becomes

$$\begin{aligned} \mathbb{E}_t \left[e^{\int_t^T z_s ds} \right] &= e^{\int_t^T (\mathbb{E}_0[e^{z_s}] - 1) ds} \\ &= \exp \left[\int_t^T e^{-\frac{\sigma_J^2}{2} h(s) + \frac{\sigma_J^2}{2} h^2(s)} ds - \int_t^T ds \right] \\ &= \exp \left[\int_t^T e^{-\frac{\sigma_J^2}{2} e^{-\gamma(T-s)} + \frac{\sigma_J^2}{2} e^{-2\gamma(T-s)}} ds - l(T-t) \right] \end{aligned} \quad (\text{A } 24)$$