Jump Diffusion Processes – Energy Price Processes Used for Derivatives Pricing & Risk Management

This is the third and final article in a three part series exploring the main stochastic price processes used to model energy spot and forward prices for derivatives valuation and risk management. The first article in this series focused on the 'random walk' assumption characterised by the most popular of price processes, Geometric Brownian Motion (GBM), and the Black-Scholes option pricing model that is based on it. In the second article we shifted our attention to the mean reverting process, which incorporates the tendency of energy prices to gravitate towards a 'normal' equilibrium price level that is usually governed by the cost of production and level of demand. In this third article, we introduce a set of processes that attempt to model price jumps, as observed in many energy markets, particularly electricity. By **CARLOS BLANCO** & **DAVID SORONOW**.

WHY MODEL JUMPS? Energy prices, and electricity prices in particular, are characterised by abrupt and unanticipated large changes known as 'jumps' or 'spikes'. Temporary price spikes are the result of supply shocks such as generating or transmission constraints and account for a large part of the total variation of changes in spot prices. In deregulated power markets, firms that are not prepared to manage the risk arising from large price 'spikes' can see their earnings for the whole year evaporate in a few hours.

This third paper in the series will examine the role of jump-enhanced diffusion models for pricing and hedging of energy derivatives. As we pointed out in the previous articles, (see *Commodities Now*, March & June 2001) it is extremely important to take into account price spikes in the modelling framework for risk management and pricing of exotic and deep out-of-the-money options. It is also important for the valuation of certain assets such as peaking plants that are only turned on in extreme price scenarios.

We will argue that it is time to move beyond Black-Scholes types of models based on the assumption of lognormal price changes and model jumps explicitly as an essential component of stochastic diffusion models that describe the evolution of power prices. We will present an intuitive way to represent important features associated with changes in electricity prices through a mean reverting jump-diffusion processes which is a generalisation of the standard Merton (1976) model.

Spot Electricity Price Behaviour – Jumps vs. Spikes

Technically, electricity prices do not jump, but 'spike'. That is, they do not jump to a new level and stay there, but rather quickly revert to their previous levels. The use of mean reversion alongside jumps allows us to simulate this spiking behaviour. Price spikes are especially notorious in on-peak hourly and daily prices. For weekly or monthly averages, the effects of price spikes are usually averaged away in the data.

The main characteristics of electricity spot prices can be summarised as[1] :

1. Mean Reversion – Power prices tend to fluctuate around values determined by the cost of production and the level of demand.
2. Seasonality – Power prices change by time of day, week, month and year in response to cyclical fluctuations in demand.
3. Non-Storability – electricity cannot be stored and once generated it needs to be consumed almost immediately. The lack of storability means that electricity prices do not follow a 'smooth process' as prices of other commodities do.
4. Price Spikes – Power prices exhibit occasional price spikes due to supply shocks such as transmission constraints and unexpected outages.
5. Regional Differences – Due to the fact that electricity is not storable and transmission constraints, spot prices and forward curves may vary drastically from region to region.

Johnson and Barz (1999) evaluated the effectiveness of 4 different stochastic models in describing the evolution of spot prices (Brownian motion, Mean Reversion, Geometric Brownian Motion, and Geometric Mean Reversion). The models were tested with and without jumps. The objective was to try to reproduce electricity price behaviour in several different markets (California, Scandinavia, England and Wales, and Victoria) for the periods of April-May, August-September, and November-December. The authors concluded that the Geometric Mean Reverting model gave the best performance, and that adding jumps to each of the models improved its performance. The authors also concluded that all models without jumps are inappropriate for modelling electricity prices. In fact, according to those models, the least likely event actually observed (the largest jump in the historical data) would be expected to occur less than once every 10 million years.

We can clearly observe this intrinsic characteristic of power markets by looking at historical prices for spot electricity in deregulated markets. For example, in 1998 and 1999, prices spiked dramatically in many areas of the US. The chart

Into Cinergy MWD Daily Index (US$/MWh) On-Peak

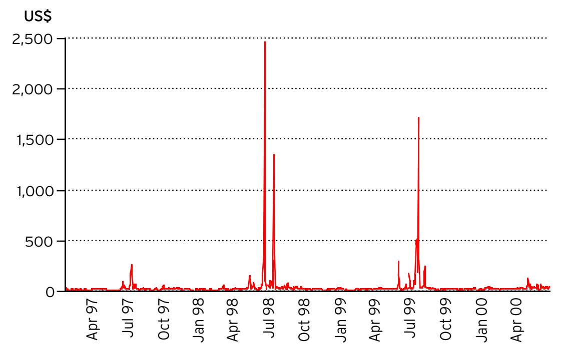

above shows prices at the Cinergy hub in the Mid-continent Area Power Pool market (MAPP). Note that these are daily averages. One MAPP trade in June 1998 was reported at a price of $7,500/MWh, for 50MW volume. Some traders were not well hedged. Losses ran into hundreds of millions of dollars in a single day, and bankruptcies resulted.

Jump-diffusion Processes Applied to Electricity Prices

To price electricity derivatives, it is necessary to characterise the evolution of the price of electricity through time. General diffusion models with time varying volatility (and possibly mean reversion) are the most commonly used by market practitioners, but they fail to capture the higher order moments (which lead to fat-tailed distributions) observed in electricity prices.

In the jump diffusion model, price change dynamics can be divided into two distinct forms:

1. A 'normal', continuous price diffusion process modelled by Geometric Brownian Motion with mean reversion and a volatility term structure. The term structure of forward volatilities coupled with mean reversion allows us to capture electricity price dynamics without spikes.
2. An 'abnormal', discontinuous jump process modelled by a Poisson distribution. These discontinuous price jumps are usually a result of outages, transmission constraints, etc.

The jump diffusion model outlined in Merton's original paper is best suited to the description of equity prices with its assumption of independently and identically distributed returns. This assumption does not incorporate mean reversion, and assumes jumps can be positive or negative; thereby leading to large increases or decreases in prices.

The main limitations of the original Merton 1976 jump diffusion model when applied to power markets is that it assumes the continuous lognormal diffusion process and the Poisson controlled process are independent of one other. This is not the case in electricity. For example, prices are highly unlikely to spike overnight when demand is very low. The nature of price jumps in electricity markets requires a generalisation of Merton's model.

A Random Walk with Mean Reversion & Jumps

In the previous articles, we made an analogy of the mean reverting, random walk process as a stumbling drunk who is being guided home from the bar along an empty street by his trusty German Shepard dog, 'Rusty'. The direction and size of his stumbles are random in nature (a random walk). However, the size of his stumble is bounded by the length of the leash and the direction of his stride tends towards Rusty's position (mean reversion). As the drunk stumbles away from Rusty, he will eventually be pulled back toward Rusty as the leash reaches its limit. Now imagine that on rare occasions a car passes by and Rusty begins to chase after it. The drunk is quickly and violently yanked in the direction of the car. Over time, the drunk is eventually pulled back to a normal position as Rusty's fascination with the vehicle dissipates and the trusty canine reverts to travelling in a homeward direction.

Model Parameters: Jump Frequency, Size & Standard Deviation

One of the great advantages of GBM models is the ease of estimating input parameters. The most difficult parameter to obtain is volatility – the expected future variability of price over time. This is typically estimated using either historical volatility or implied volatilities from current option quotes. Mean reversion models require the estimation of additional unknown parameters such as the mean reversion level (the long-run equilibrium price) and the mean reversion rate (the speed at which prices revert to the mean).

Random Walk with Mean Reversion & Jumps

Mathematically, we capture the phenomena of mean reversion with a modification to the random walk assumption (see formula at top of opposite page).

The discontinuous jump component governs the possibility of a jump, the possible jump size, and the possible jump standard deviation within any short period of time.

The jump-diffusion process introduces three new model parameters that allow us to model outage events. In order to model jumps we need to determine their probability of occurrence or 'jump frequency', their expected size, and their expected variability or standard deviation.

a. Jump Frequency

The jump frequency tells us how often jumps occur on average. The jump frequency is also known as 'arrival frequency', and is expressed as the number of jumps per year.

b. Average Return & Standard Deviation of the Jump

The average jump size, or κ, describes the average return after a jump[2] , while δ describes the dispersion of the actual jump returns around their mean value (technically δ is the standard deviation characterising the normal distribution of the log-changes in prices after jump events).

The jump size standard deviation, also known as jump volatility, is a number greater than or equal to zero specifying the standard deviation of the probability distribution describing jump sizes. For instance, an expected jump size of 200% and a standard deviation of 50% implies that, on average, the price of the underlying doubles after a jump and, within 1 standard deviation, the size of the jump stays between 150% and 250% of the pre-jump price.

The jump diffusion with mean reversion process allows us to make prices return quickly to the expected pre-jump level without the need to impose complex serial correlations onto the jumps. If mean reversion was not included in the model, predicted prices might stay at the higher levels after a jump.

Extracting Jump Parameters

In order to estimate the jump parameters for our model we first need to decide what we mean by normal, non-jumping price behaviour. To do so we specify a threshold and use it to filter out jump events. For instance, depending on the amount of data we are actually using, we might decide that price returns beyond three standard deviations should be considered jump events. Once the jump events have been identified, their frequency can be extracted by simple counting, and the distributional parameters describing the size of the jumps can be obtained by curve fitting. The rare-jump assumption can then be explicitly verified and, if necessary, the volatility of the diffusive part can be re-estimated so that the existence of the jumps does not distort the base volatility value.

Implications for Derivatives Valuation & Risk Management:

Valuing out-of-the-money options with Jump Diffusion vs. Black-Scholes or Mean Reverting type models

Using a jump-diffusion model is particularly important when valuing deep out-of-the-money call options on electricity prices. Out-of-the-money options can be found in many types of electricity contracts. For example forward contracts that contain Interruptible and Pass-Through clauses. Even if those clauses only kick in at high electricity prices, we need to model price spikes directly in order to price and hedge them effectively.

The accuracy of using a pure diffusion process (i.e. one without jumps) to price certain electricity derivatives will be at best questionable and in worst cases (e.g. out-of-the-money, path-dependent options) totally unreliable. While the 'theoretical' premiums of out-of-the-money options might be negligible when we calculate them with traditional valuation models without jumps, the market gives a very different value to those options. In order to appropriately price and hedge certain electricity options, we need to use a model that allows the possibility of occurrence of sizable jumps that could suddenly bring the option into the money.

Modelling price spikes accurately is also important for generation assets, particularly Peaking plants with high running costs whose value might be totally dependent on the existence of price spikes that would allow them to both recover their high marginal costs and recoup their fixed costs over a very short running period. Typically power stations are valued by calculating the discounted value of expected future cashflows, or more sophisticatedly the option value obtained using discounted present value methods or lognormal based option models. However, recent sales in the USA and UK suggest that purchasers have not only taken the option approach to valuation, but have

Mean Reversion with a modification to the random walk assumption

$$\underbrace{S_{t+1} - S_t}_{\text{Change in price between } t+1 \text{ and } t} = \underbrace{\alpha(S^* - S_t)\,\Delta t}_{\text{Mean Reversion Component}} + \underbrace{S_t \sigma\, \varepsilon_{1t} \sqrt{\Delta t}}_{\text{Diffusion Component}} + \underbrace{\eta\left[S_t(\kappa + \delta\varepsilon_{2t})\right]}_{\text{Jump Component}}$$

Where:

S^* is the mean reversion level or long run equilibrium price

S_t is the spot price at time t

α is the mean reversion rate

σ is the volatility of the general diffusion process

ε is the random shock to price from t to $t+1$

η is the variable that takes the value of 1 if there is a jump, and zero otherwise. The value of η depends on the jump frequency λ

κ, δ are the expected jump size and standard deviation of the jump

$\varepsilon_{1t}, \varepsilon_{2t}$ are two normally distributed random variables for the jumps and volatility

FEA's @Energy's Mean Reversion with Jump Diffusion
(extracted from @Energy's User's Guide)

FEA has chosen a jump-diffusion model that represents a generalisation of the standard Merton (1976) model, to the case of a mean reverting price, and provides an intuitive way to represent important features associated with changes in electricity prices. In order to explicitly account for price spikes, the price process for the spot price under the jump diffusion model is governed by the following equation:

$$d\log S_t = a\ \theta_t - \log\ S_t\ \frac{S_t}{F\,(0,_t)}\ dt + \sigma_t dw_t + d\sum_{i=1}^{N(t)} X_i$$

where the last term on the right hand side represents the jump component that can provide a discontinuous price change over an arbitrarily small time interval. $N(t)$ is the Poisson process governing the independent jump events, and X_i is the log-return associated with the i^{th} jump. The X_i are drawn from a normal distribution with mean $\bar{X}$ and standard deviation δ. The X_i corresponding to different jumps are independent and have no correlation to the diffusive process described by dw_t.
As with the standard mean-reversion model, the function θ_t is internally determined by requiring that the expected value of S at time t (determined at time zero), match the input forward price $F\,(0,_t)$.
Source: Financial Engineering Associates

FEA @Energy: Monte Carlo Simulation of Price Process

Market Inputs	Value
Spot Price	
Interpolation	1
Futures Contract	

Simulation Inputs	Value
Simulation sate date	08/01/01
Simulation end date	12/07/01
Sampling Frequency	d

Model Inputs	Value
Model	Jump Diffusion
Mean-Revision rate	365
Cap on Price	
Floor on Price	
Yearly jump frequency	12
Expected jump size	500%
St. Dev. of jump	30%

Other Inputs	
Value Date*	07/01/01
Interest rate	5.00%
Seed*	

Other Inputs	
Simul*	10000
Min-Max	0 - 300

Forward Price Curve

7/1/01	8/1/01	10/1/01	11/1/01	12/1/01	1/1/02	2/1/02	3/1/02
60	50	35	35	35	30	30	40

Volatility Curve

7/1/01	8/1/01	10/1/01	11/1/01	12/1/01	1/1/02	2/1/02	3/1/02
100%	120%	125%	140%	80%	70%	60%	60%

Monte Carlo Sample Price Paths –
Jump Diffusion with Mean Reversion Process

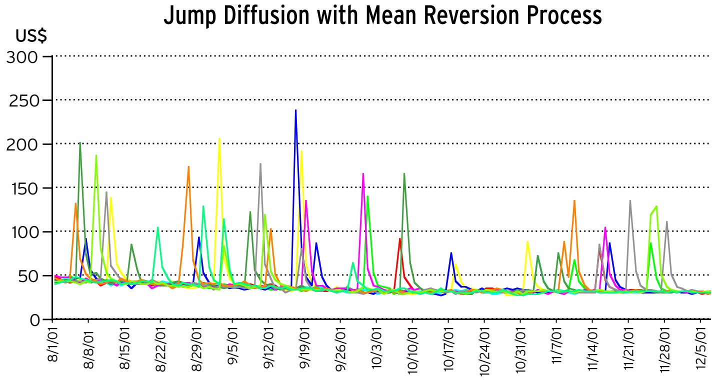

included expectations of price spikes in their valuation of the option.

VaR & Extreme Price Scenarios

If we are dealing in markets where prices may experience sudden spikes, it is necessary to incorporate the possibility of the occurrence of those events in the modelling framework. Standard VaR models do not accurately account for the possibility of extreme events. By using a Jump Diffusion process in the Monte Carlo Simulation, we can incorporate extreme events in our risk simulation. Depending on expected market conditions, we can adjust the jump parameters before conducting the simulation. Historical Simulation can also prove useful in dealing with jumps, but due to the strong mean reverting effects of prices, we need to carefully pre-process the data before conducting the analysis. The existence of jumps and rapid returns can skew the distribution that is used to drive the analysis. We should also be careful when extracting conclusions from a purely historical analysis as future jumps may not necessarily resemble historical jumps. Physical changes to the market, such as upgrading a transmission line, building new power stations, or even different weather patterns may remove the conditions that cause jumps to occur.

Actual vs, Simulated Price with Mean Reverting Jump Diffusion Process

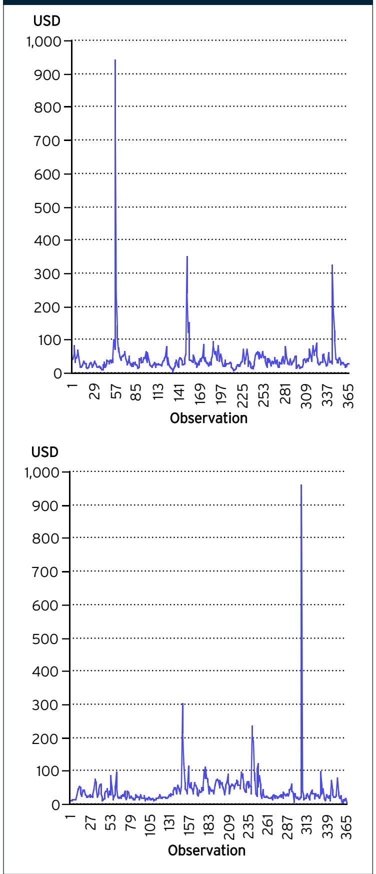

In the charts we can see actual average daily prices from Australia compared to simulated prices obtained using a mean reverting jump diffusion process calibrated with historical data. While the timing of the spikes does not match that of the actual data, the general shape of the price process is very similar in both charts. If the timing of spikes is genuinely random then the simulated prices shown here would work well for contract valuation. However, if the existence of spikes was dependent on other factors, for example they only occur in the very hottest or coldest months, then a more refined model that reflected the seasonality of spikes would be required.

Another way to model price spikes within the VaR framework is to develop and conduct a set of stress tests that include price spikes, assign probabilities of occurrences to those events, and integrate those scenarios as part of the risk simulation in order to obtain a more accurate picture of possible 'states of the world'. For a detailed description of the process see Aragonés and others, 2001.

If the company's earnings also depend on available generation, special attention should be given to the identification and quantification of operational risks such as unit outages, particularly if the firm has to buy electricity in the hourly or daily spot market to meet its contractual obligations. If the combination of market and operational risks could result in losses beyond the tolerance level of the firm, management should consider purchasing some type of insurance (e.g. double-trigger contracts) against those events.

Pitfalls of Using Jump Diffusion with Mean Reversion to Model Energy Price Behaviour

1. Calibration of Jump Diffusion Models with Historical Data

History is often a poor guide in electricity markets. One of the obvious weaknesses of attempting to estimate model parameters based on past data is that the historical record does not incorporate expectations about future jumps. Option markets provide us with the ability to calibrate our models to market prices, but when we have multiple parameters to estimate (see the price process description) we need to make certain assumptions about the value or joint value of those parameters. We cannot calibrate all of them directly, unlike in the case of an implied volatility calculation for a Black-Scholes model.

We can also use forward prices as a guide to the expected level and frequency of jumps in the future. For example, when forward prices for summer delivery in the Midwestern states of the US experience a sudden increase, it is likely that a large part of that increase may be attributable to a change in expectations about the magnitude and frequency of price spikes rather than a general increase in the level of expected prices.

In order to accurately price derivatives with the jump diffusion model, users should combine historical with forward looking estimates of the jump parameters that take into account current market conditions that may influence expected 'spikes' (e.g. excess capacity, weather forecasts, etc.).

2. Jump Parameters are not constant.

The size, frequency and volatilities of the jumps, and the speed at which prices revert to their long run levels may depend on several factors such as the nature, magnitude and length of the event that causes the price shock.

Besides, jump parameters have clear seasonal components. If we calibrate the jump size, frequency and standard deviation for each month of the year using data exclusively from that month, we would find that for most markets each of those parameters would be different as a result of seasonal variations. The most obvious cause of this is that the extreme shortage of available capacity that causes spikes will probably only occur at times of the day/week/year when demand is close to its annual maximum. At other periods there is sufficient excess generation to cover the loss of a power station or transmission line.

3. Mean Reversion Rates may be different for Jumps than for normal price changes.

When electricity prices spike, they tend to return to their mean reversion levels much faster than when they suffer smaller shocks. The Jump Diffusion Process presented here assumes that there is only one speed of mean reversion, but more complex processes assume that the mean reversion is different dependent on the magnitude of the shock and the time of the year. The cause of the spike may also have an effect. For example, if the spike is the result of a failed transmission line, prices may revert to normal immediately the line is brought back into service. If, however, prices are generally high due to abnormal weather conditions and prices are spiking whenever demand peaks, then the speed of return to normality will depend on the weather system.

Summary of Models			
Electricity Price Features	**Lognormal Models**	**Mean Reversion**	**Jump Diffusion with Mean Reversion**
Mean Reversion	Not Captured	Fully Captured	Fully Captured
Seasonality	Captured in the Price levels and Volatilities	Captured in the Price levels and Volatilities	Captured in the Price levels and Volatilities
Price Spikes	Not explicitly captured	Not explicitly captured, although they can be approximated with high volatilities and high mean reversion rates	Explicitly modelled

4. Jump Events are not independent.

Simple Jump Diffusion models assume that jumps are not serially correlated. More sophisticated models such as 'Regime Switching' models provide different 'states of the world' characterised by different price process and/or model parameters, and also a probability of remaining at a particular state once we are in that state in our simulation (i.e. once a jump has occurred). Those models may represent an advantage in terms of increased accuracy through introducing a more realistic depiction of jumps, but their calibration is even more complex, and the probability information between 'states of the world' is usually very difficult to calibrate.

Conclusions

Price process models lie at the heart of derivatives pricing models and risk management systems. If the price process chosen is inappropriate to capture the main characteristics of electricity prices, the results from the model are likely to be unreliable. Models based on Geometric Brownian Motion with mean reversion are unable to capture many of the unique characteristics of power prices.

Jump diffusion models describe what is, for risk management purposes, perhaps the most important energy spot price phenomenon: the discontinuous price jump. The use of these models for electricity is supported by empirical studies of those prices, as well as by basic microeconomic theory.

Energy price jumps are very unstable and difficult to characterise and predict. As more historical data becomes available, and therefore more statistical information can be derived to allow analysts to model jumps, these models will probably gain more widespread acceptance among market practitioners. Jump diffusion models have clear advantages over the Black-Scholes and mean reversion counterparts. However, because calibration of those models to market data is still a very difficult and subjective process, the applicability of those models in trading environments is still relatively limited ■

Footnotes

1. For more information, see e.g. Barz, G. and Johnson, B. in Energy Modelling and the Management of Uncertainty, RISK books (1999).
2. The 'return' is the relative price change from one period to the next.

Carlos Blanco is Manager of Support & Educational Services at Financial Engineering Associates Inc. **David Soronow** is a derivatives specialists at FEA. The authors would like to thank the FEA Financial Engineering & Support teams, and in particular Angelo Barbieri, King Wang and Cheryl Morgan for their helpful comments and software assistance.

References

Aragonés, J., C. Blanco, & K. Dowd. 2001. Incorporating Stress Tests into Market Risk Modelling. Derivatives Quarterly 7 (Spring), 44-49.

Barz, G. and Johnson, B. Energy Modelling and the Management of Uncertainty, RISK books. 1999

Financial Engineering Associates. @Energy User's Guide. Release 2.2. 2001

R. C. Merton, *Option pricing When Underlying stock returns are discontinuous*, Journal of Financial Economics, 3 (1976).