Continuing our series on applications of Monte Carlo simulation to applied problems in energy risk management, *Les Clewlow*, *Chris Strickland*, *Oleg Zakharov*, and *Scott Browne* look at potential future exposure and the analogous measure of expected credit exposure and how these mitigate OTC counterparty risk

Double exposure

★ Potential Future Exposure (PFE) – sometimes called potential credit exposure (PCE) – and the analogous measure of expected credit exposure (ECE) – allow a risk manager to measure and manage counterparty risks of over-the-counter instruments across an entire portfolio.

Suppose, for example, a financial analysis of counterparty X leads the risk group to be comfortable with a level of $25 million in unsecured exposure. Using ECE or PFE (depending on the risk tolerance of the company) it can be determined whether the exposure will exceed this level with a chosen probability. If the ECE was $30 million and the company still wants to complete the deal, then they can ask for a letter of credit for $5 million, or if the counterparty was unwilling to provide the guarantee, purchase a credit default swap on the counterparty for $5 million. Another common use of the exposure measure is to be able to assess the potential margin that may need to be posted under the contract.

In the following we will look at the practical implementations of PFE for an electrical swap contract. The swap that we will use as an example is an at-the-money contract for physical delivery during the months October 2004 through to December 2008 where the origination date for the swap is May 17, 2004. The swap has a specified delivery schedule with a total of approximately 10 million MWh over this 51-month period. Figure 1 plots both the load profile and the PJM forward curve observed on 17 May 2004.

In order to evaluate the exposure calculation we need a model to describe the future evolution of the PJM forward and spot prices. In this article, we will consider two practical models of a particular case of a general multi-factor lognormal process that is detailed in Clewlow and Strickland [2000]. In terms of forward price evolution this general model can be represented by the following stochastic differential equation:

$$\frac{dF(t,T)}{F(t,T)} = \sum_{i=1}^{n} \sigma_i(t,T)dz_i(t) \qquad (1)$$

In equation (1) $F(t,T)$ describes the forward curve for T-maturity observed at calendar time t, and $\sigma_i(T,t)$ the n general volatility functions, with $dz_i(t)$ the corresponding random Brownian innovations. We will calculate the exposure levels under both a single factor and a two factor restriction of equation (1), both of which can be represented by the following;

$$\frac{dF(t,T)}{F(t,T)} = \sigma_S(t)\cdot(\sigma_1(T-t)dz_1(t) + \sigma_2(T-t)dz_2(t)) \qquad (2)$$

Here $\sigma_S(t)$ is a spot volatility (we will use the prompt month price as a proxy for spot in the following analysis) which we allow to be seasonally dependent (by month). The volatility parameterisation in equation (2) implies a certain pattern for the forward volatility seasonality – it suggests that forward volatilities

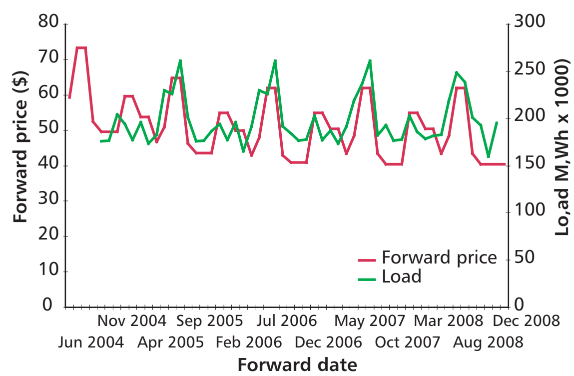

F1. Forward curve and load profile

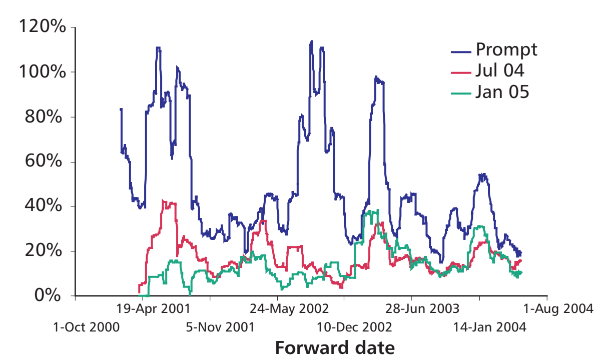

F2. 30-day rolling volatilities

should increase when spot volatility increases. A quick analysis of historical data justifies this assumption. Figure 2 plots both the seasonal spot volatility and the seasonal 'long' volatilities as represented by a rolling 30-day window of the prompt month contract and the July 2004 and January 2005 contracts respectively.

This figure illustrates that the seasonal features of the forward volatilities indeed closely resemble the spot volatility. The first model we consider is a single factor version of equation (2) with the following choices for the two volatility functions:

$$\begin{aligned} \sigma_1(T-t) &= (1-\sigma_l)e^{-\alpha(T-t)} + \sigma_l \\ \sigma_2(T-t) &= 0 \end{aligned} \qquad (3)$$

In equation (3) α plays the role of attenuating forward volatility with increasing maturity; σ_i the role of long-term forward volatility, expressed as a percentage. Both the simple mean-reverting model[1] of ($\sigma_i = 0$) and the Black (1976) model ($\alpha = 0$, $\sigma_i = 0$) are special cases of this model. It can be shown that this process is equivalent to the mean reverting process for the underlying spot price generalised to include long-term forward volatility. Since the evolution in this model is driven by a single stochastic factor, all the forward prices are instantaneously perfectly correlated, but, since the front end of the curve reacts differently to the long end when the random shocks are applied, the slope of the curve may change. Figure 3 plots the forward price curve on 1 October 2003, 2 February 2004, and 12 May 2004. The variability of the short and long end of the curve suggests that such a one factor model with decaying volatility may be a reasonable approximation for the historical data and a good starting point for our analysis.

In order to calibrate the model defined by equations (2) and (3) we need to derive a seasonal spot volatility, $\sigma_s(t)$, a rate of attenuation α and a long term forward volatility σ_i. For the analysis we used a history of PJM forward prices from January 2001 to May 2004 for monthly peak contracts. To derive the spot volatility input we analysed the prompt month forward data, after removing the balance of month prices for all the contracts and disregarding the returns on the contract roll date. This cleansing is performed to remove jumps from the analysis of price returns caused solely by the contract rollovers.

In figure 4 we plot the term structure of the forward price return volatility against months to maturity using varying lengths of historical data; the whole period of our dataset as well as the sub-periods of the last two years, last one year and last six months of data.

In order to calibrate $\sigma_i(T-t)$ we use a least squares methodology to minimise the sum of squared differentials with the historical data. This produces parameter values of roughly 7.4 and 32% respectively for α and σ_i respectively.

The second model we consider is a two-factor version of equation (2) where the shape of the first two volatility functions, and their weights, are obtained by applying principal components analysis to the historical forward price returns. These two factors are plotted in figure 5.

1. *See Clewlow, Strickland, Kaminski, Nov 2000, 'Making the Most of Mean Reversion'*

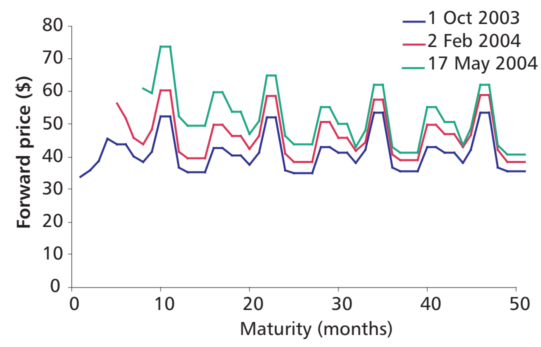

F3. Forward price curve on selected dates

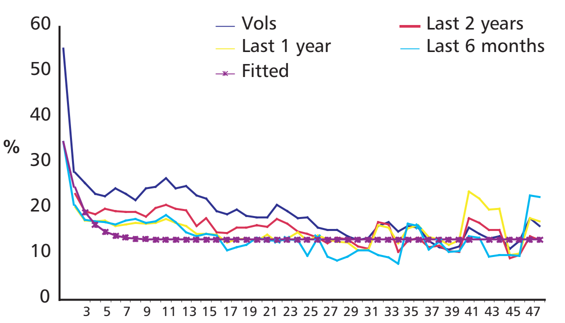

F4. Term structure of forward volatility

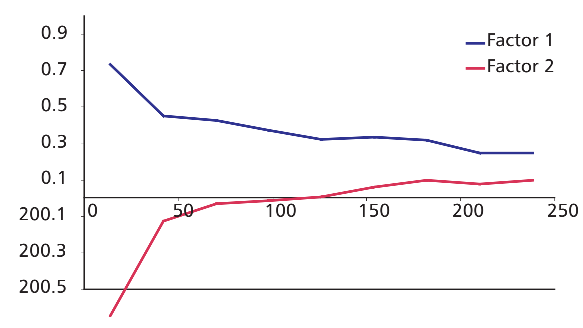

F5. Two factors obtained from PCA

To compare the results using the single factor and the two factor approaches better, the volatility functions in figure 5 are renormalised to fit the total volatility of the single factor model.

The credit exposure calculations are computed by simulating the forward curve according to our two models, on a daily basis from May 17, 2004 until the end of December 2008. Because we will report the exposure on a monthly basis at the beginning of each month we calculate, at these times, the present value of the

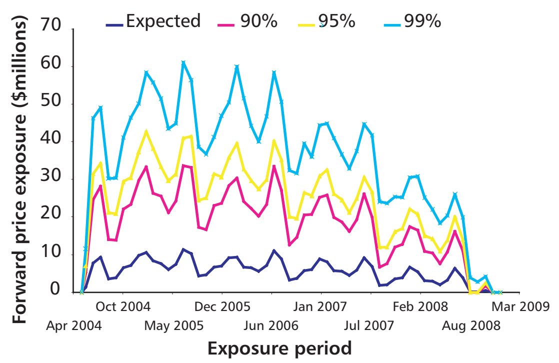

F6. Exposure profiles for the single-factor model

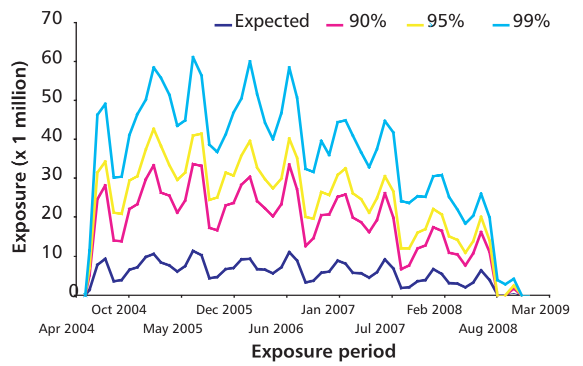

F7. Exposure profiles for the two-factor model

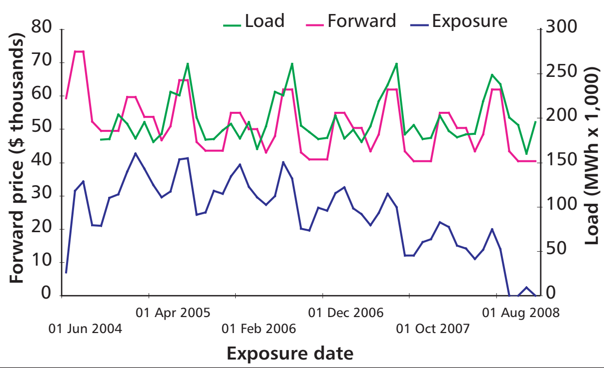

F8. Carbon and energy intensity trends in the US

remaining life of the swap for each simulated forward curve. The present value statistics are then used to compute the expected credit exposure (ECE) and the exposure profiles at the 90th, 95th, and 99th percentiles. In this example, we limit our calculations to the mark-to-market or 'replacement' exposure. Our results for the single-factor model, and using 1,000 simulations of the forward curve, are presented in figure 6.

Figure 6 shows that the exposure is represented by a regular sawtooth type profile with maximum exposure points on the swap payment dates, and where the maximum ECE is given by $11,341,410 and the maximum PFE at the 95th percentile is given by $42,741,921.

Analogous results for the two factor model are presented in figure 7 where the maximum levels of the ECE and PFE are given by $10,603,673 and $38,222,161 respectively (with 1,000 simulations).

To get a sense check on the exposure levels we note that a parallel shift in the original forward curve by $1 would result in the net present value change of $10,000,000 = 10,000,000 MWh x $1 (without discounting). By looking at the long end of the forward curve we can see that it moved by almost $3 between February 2004 and May 2004 – such a move would have resulted in an exposure of around $30,000,000, and even more if we take into account our observation that front-end prices move further.

We now turn our analysis to the shape of the exposure profile. The exposure increases quickly over the first few months of the analysis, leveling off – although with the sawtooth pattern noted earlier – and then gradually decreasing until the exposure is zero when the last rest date of the swap has settled. In order to better understand the shape of the profile through time, figure 8 plots the 95th percentile exposure – under the single-factor model – as well as the initial forward curve and the load shape of the swap.

We consider three possible sources for the shape of the exposure profile – the load shape itself, the initial forward curve profile, and our assumptions about spot volatility. Figures 9, 10, and 11 show the exposure profile recalculated with each of these sources (load shape, forward curve, and spot volatility respectively) set to their average level in turn.

The main conclusions that we draw from this sensitivity analysis are the following:

- Load shape does not significantly influence the exposure shape
- Forward curve shape is the biggest determinant of the exposure profile for this swap
- Spot volatility shape is the second biggest factor
- There is some minor statistical noise from the Monte Carlo simulations

Our analysis, although informative, hides some of the complex interactions that are driving the exposure profile. We find that, because of sharp increases in the volatility for the short-term months, these months contribute more to the total exposure. The summer months, with their higher prices, result in significantly higher absolute price changes (remember that the underlying process is lognormal) that in turn lead to significant increases in the exposure. As the summer months roll off the swap, and winter months become the prompt months, the price moves become smaller, resulting in the decreasing exposures. Without increases in the volatility for the prompt months this effect would average out, and would not have been so pronounced. Similarly, when the seasonal volatility of the prompt months is higher we would expect a higher total exposure number because the volatility of the whole curve increases.

Looking at the differences in exposure calculations for the single-

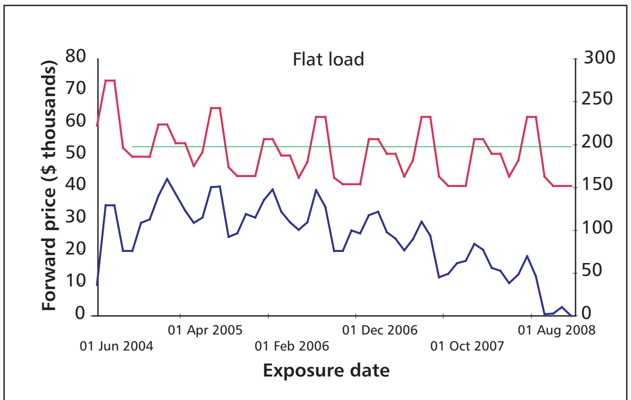

F9. Exposure profile recalculated: load shape

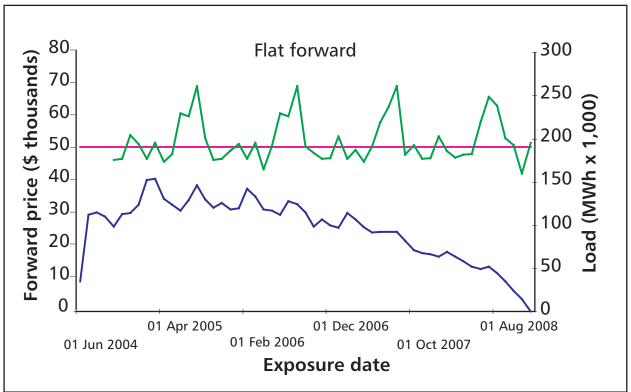

F10. Exposure profile recalculated:, forward curve

factor versus the two-factor case, we note that the two-factor model adds a twist to the curve dynamics allowing for short-term and long-term prices to move in opposite directions, thus decoupling the perfect correlations of the single-factor model. Intuitively, it is clear why adding this factor would reduce the exposure, since the curve now does not necessarily move as a whole. At the same time, the historical data indicates that the contribution of such a twist to

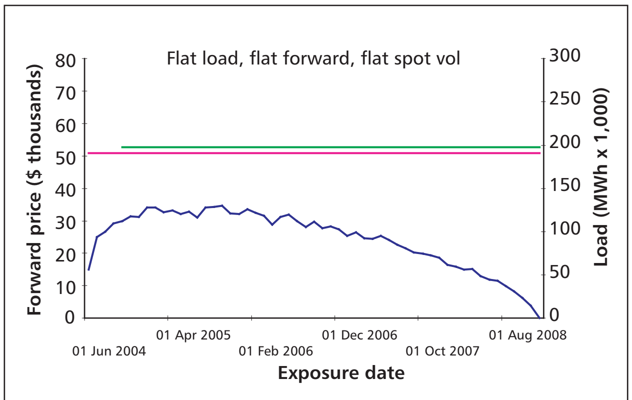

F11. Exposure profile recalculated: spot volatility

the curve dynamics is not significant, resulting in only minor reduction of the exposure numbers.

Finally, we note that the two models we have applied in this article have been chosen as examples only – although our data analysis shows that they are consistent with some of the key properties in the data. It is straightforward to extend either of these models to include jumps – one of the things that we sometimes find when back-testing is that the actual exposure gets outside of the bounds of the PFE due to large shifts in the longer dated part of the curve. ER

Les Clewlow and **Chris Strickland** are principals of Lacima Group, and authors of the forthcoming book *Energy Risk Management: Applications Using Simulation.* **Oleg Zakharov** is general manager, Europe and North America for Lacima Group. **Scott Browne** is director, enterprise risk management at Public Service Enterprise Group in New Jersey **Email:** chris@lacimagroup.com

References

Clewlow, L, and Strickland, C, 2000, *Energy Derivatives: Pricing and Risk Management*, Lacima Publications