In this Masterclass article, *John Breslin, Les Clewlow, Calvin Kwok, Chris Strickland* and *Matthias Pfau* continue their discussion of the multi-factor multi-commodity modelling framework, by considering its practical application to the valuation of complex spread options

Application of MFMC to spread option valuation

★ Spread options are used in many markets for trading and risk management. One of the main activities in energy markets is trading the differences between individual commodities. For example, these differences can be based on the spread between prices in different calendar months for the same commodity (calendar spreads), between prices of the same commodity at different locations (locational spreads), between prices of crude oil and refined products such as heating oil (crack spreads), or between the price of power and the fuel used to generate it (spark spreads).

For some simple European options under very restrictive distributional assumptions, there exist closed-form solutions for the valuation of options on these spreads. For examples, see Ravindran (1995), Kirk (1995), Hikspoors and Jaimungal (2007). However, despite their importance and widespread use, there does not exist a consistent framework for pricing and hedging general spread options in a modelling framework that is relevant for energy prices. For this reason, and also for cases where the spread option payoff is relatively complex, a Monte Carlo simulation approach to valuation is often the only practical method available.

As noted in our first multi-factor multi-commodity (MFMC) article, if enough data is available, then using a general MFMC model for the price dynamics allows us to account for a very rich set of modelling assumptions. In this article, we describe how the MFMC model can be used as the basis of a Monte Carlo simulation for the valuation of complicated spread options that depend on a basket of energy prices.

Firstly, we describe the basic principle of spread options. The simplest European spread option is an option on the difference between two underlying assets, which have values at time T denoted by $S_1(T)$ and $S_2(T)$. Similar to the payoff for a simple European call option the payoff of the simple spread option at maturity T is given by:

$$\text{Payoff} = \max\left[S_2(T) - S_1(T) - K, 0\right] \quad (1)$$

where K is the strike price of the option. If the difference in the asset prices is greater than K, then the holder would exercise the option, otherwise it will expire worthless. Of course, the payoff definition need not be as simple as the difference between the two asset prices and a strike price.

Many market participants price power tolling agreements and financial heat rate options, as well as view generation assets as European spread options, or as a strip of such options. For example, a standard approach to modelling a single fuel (such as gas) generation asset is to use a strip of hourly spread call options with the following payout:

$$\text{Payoff} = \max\left[Power(T) - Gas(T) \times HR(T) - K, 0\right] \quad (2)$$

where:
- $Power(T)$ defines the power price at time T;
- $Gas(T)$ defines the contemporaneous gas price;
- $HR(T)$ is the heat rate that defines the efficiency factor of the plant;
- K is used to account for the costs of generation.

In general, a spread option refers to an option where any number of commodity prices or indices can be combined to define the payoff calculation. Many energy producers or consumers are exposed to a basket of commodities. Long-term interruptible gas supply contracts – which are common in Europe for pricing natural gas, for example – often depend on a basket of commodities such as gas oil, fuel oil and coal, and so hedging contracts that play on the spread between the gas price and the basket of fuels are common. One reason for this type of arrangement is that some gas markets can be relatively illiquid, especially where they are dominated by a handful of gas producers, and so using reference prices from more liquid markets (such as oil) provides greater price transparency and minimises the risk of price manipulation.

As an example, we consider the spread between the gas forward price (based on the UK's National Balancing Point (NBP)) and the forward value of an index of a basket of fuels. The fuels we consider are gas oil (GO) and fuel oil (FO). An additional complication is that the spread payoff involves the sum of the spread between the gas price and the fuel index over a specified number of delivery months. The payoff is defined as:

$$\text{Payoff} = \max\left[\sum_{i=1}^{n} M_i \cdot \left(F_{NBP}(T,T_i) - A(T,T_i)\right), K\right] \quad (3)$$

where:
- T is the expiry date of the option;
- M_i is the gas volume for the i^{th} delivery month (i=1, …, n) expressed in megawatt hours (MWh);
- $F_{NBP}(T,T_i)$ is the forward NBP gas price for the delivery month T_i on date T expressed in €/MWh;
- $A(T,T_i)$ is the forward value of the fuel index for month T_i on date T expressed in €/MWh. The sum is taken over all delivery months.

The fuel index used in this example is given by

$$\begin{aligned} A(T) &= P_0 + w_1 c_1\left(F_{GO}^{*/*/*/*}(T) - GO_0\right) \\ &+ w_2 c_2\left(F_{FO}^{*/*/*/*}(T) - FO_0\right) \end{aligned} \quad (4)$$

where P_0, w_i, c_i, GO_0, and FO_0 are constants. The first constant, P_0, can be thought of as the strike price of the NBP gas component of the option; if the two terms relating to the other fuels were set to zero, then equation (3) would reduce to a strip of simple call options on NBP gas.

The weightings, w_i, determine the relative contribution of the price of the corresponding fuel to the index, while the constants GO_0 and FO_0 represent a reference price for each of these commodities. The constants c_i are conversion factors to convert all prices into the same units. The terms $F_X^{*/*/*/*}(T)$ denote an averaged price for the fuel X (either GO or FO). The superscript */*/*/* defines a convention used to calculate average values from the historical prices of the commodity. From the perspective of the purchaser of these contracts, the averaging provides a level of certainty about the price they will be faced with when they take delivery of the commodity. This is discussed in more detail below. Each commodity can have its own averaging method.

Essentially, the averaging method is used to define seasonal averages, where a season can be defined differently for each commodity depending on its particular characteristics. The notations for the averaging methods have the following format:

$$a/b/c/d,$$

where:
- a is a reference month (between 1 and 12) that defines the start of an averaging cycle (and is not relevant when b=1);
- b is the number of months with the same value, which defines a cycle length.

The average value calculated will be used for the reference month and the following $b-1$ months. The parameter b must be one of 1, 2, 3, 4, 6, or 12.

Note that if b=1 then the parameter a is not relevant, as each cycle is only one month in length, *i.e.* the average is calculated for each month;
- c represents the number of monthly prices used to calculate the average;
- d is the lag between the first month in the current cycle and the last month in the average.

Clearly, the calculation of the spread option payoff defined in equations (3) and (4) is not straightforward, and deriving an analytical expression for valuing such an option is not possible. The only option left in such a case is to use a Monte Carlo-type simulation to calculate the expected payoff and hence derive a value for the option.

This also provides an additional benefit, as a distribution of the option payoffs can be obtained, which can be used for risk management purposes. The MFMC framework is perfectly suited for tackling this example – since it can capture the

Defining seasonal averages

For example, if the current month is April 2008 and the averaging is defined by 3/2/4/2, we then have the following:
- The cycles start in March;
- Two consecutive months have the same average value. Together with $a=3$, it implies that the following month couples have the same values and define the following cycle: March/April, May/June, July/August, September/October, November/December, January/February for the next year. For this example, our April 2008 average will be the same as for March 2008;
- There are four months in the average calculation;
- The first month in the current cycle is March and the lag is two months. This means that the last month in the average is January 2008. As there are four months in the average, it means that October 2007, November 2007, December 2007 and January 2008 are used to calculate the average value for the March/April 2008 component of the cycle.

F1. Forward curves used for the simulation

dynamics of each commodity from the observed market data, it can capture the correlations between the commodities, and it allows for the simulation of the prices necessary for the complex payoff calculation and averaging methods.

To illustrate the use of the MFMC model for this problem, we consider an option on 12 monthly deliveries as defined by the following dates:

T1. Option dates

Valuation date	May 1, 2008
Expiration date	October 1, 2008
Delivery start month	October 2008
Delivery end month	September 2009

The parameters used in this example to define equations (3) and (4) are as in table 2.

Finally, the averaging methods used for each of the fuels are in table 3.

The constants GO_0, and FO_0 are set equal to the average forward price over the delivery months. The constants c_1 and c_2 are set such that when applied to the average forward price over the delivery months makes the result equal to the average NBP forward price over those months. The weights are typical values for fuel indices where gasoil makes a relatively larger contribution than fuel oil.

As described in our previous article, to simulate the commodity prices using the MFMC model we require a forward curve for each commodity, a set of volatility functions for each commodity and a correlation matrix for the Brownian shocks. In figure 1, we plot the forward curves used in this example for each commodity. In figure 2, we plot the first three volatility factors for each commodity.

While the definition of the option used in this article is quite complex, the key result of the MFMC valuation is simply the expected value of the option. For risk management, another key output is the distribution of values. For trading these instruments, the expected value is obviously of prime importance, as it allows the buyer or seller to determine the true value for the option, taking into account the volatility and correlations of the underlying commodities.

A natural question to ask then is how sensitive is the value to different modelling assumptions? To answer this, we have set up four valuation models. These are described in table 4.

The single factor models are derived by calculating the single factor for each commodity that is equivalent to the three factors used in the multi-factor simulations. In table 5 we show the expected value for each case, using 10,000 simulations.

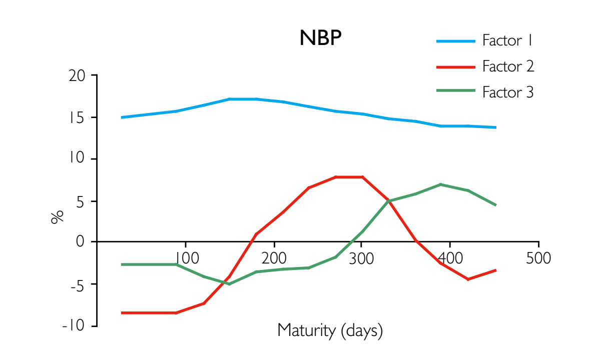

F2.1. Volatility factors for gas forward price

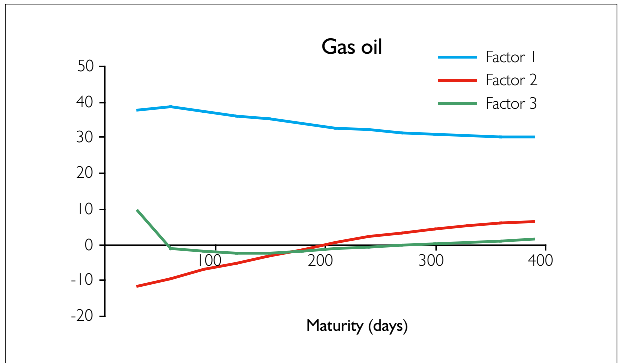

F2.2. Volatility factors for gas oil

T2. Option parameters

M_i	1	GO_0	748
K	0	FO_0	2.29
P_0	0		
w_1	60%	c_1	1.55
w_2	40%	c_2	504

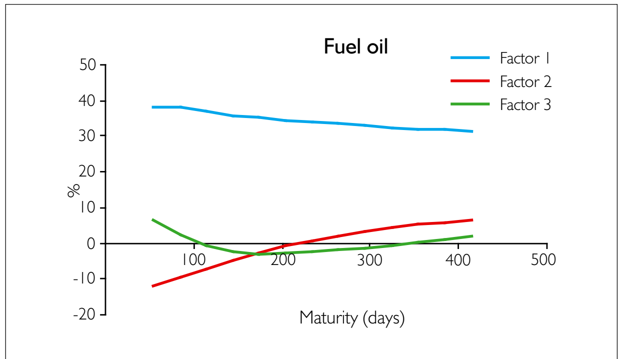

F2.3. Volatility factors for fuel oil

T3. Averaging methods

Fuel	Method
GO	1/3/6/1
FO	1/3/6/1

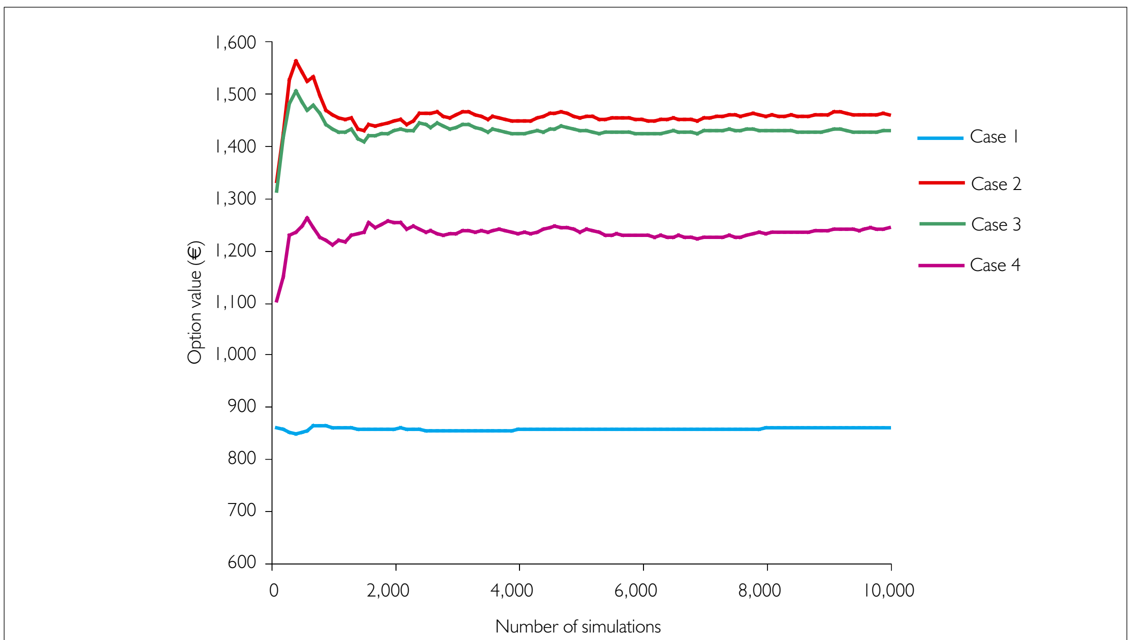

F3. Option value as a function of the number of simulations for each test case

T4. Cases used to analyse model performance	
Case	Description
1	Use a single factor model for NBP and deterministic (forward) prices for the other fuels
2	Use a single factor model for all commodities
3	Use the full MFMC model for all commodities
4	Same as Case 3, but with correlations between all commodities set to zero

There are several points to note here. First, there is a significant difference between the full MFMC model (*Case 3*) compared to *Case 1*. This is expected since *Case 1* does not simulate any variation in the price of the basket of fuels and so the true value in the spread option is not captured. The difference between *Case 2* and *Case 3* is also reasonably significant (2%), which indicates that modelling the dynamics of these commodities in a full multi-factor framework is important to capture the value and risk accurately. Finally, we note the large difference between the *Case 4* (with correlations between the commodities set to zero) and *Case 3* results (13%). This highlights the importance of including correlations in the model, especially for an instrument like the one being considered here, where the payoff depends on the underlying prices of multiple correlated commodities.

For Monte Carlo simulations, it is also of interest to look at the rate of convergence of the results. In figure 3 we show the values for each test case as a function of the number of simulations. It is clear that for examples like this where several commodities are being considered then a large number of simulations is required to sufficiently sample the problem space. This is especially noticeable for the *Case 4* results, which appear to require significantly more simulations for the result to converge. Due to the zero correlations between commodities in this case, the simulated values can explore a larger volume in the problem space and so it requires more simulations to sample it sufficiently. Of course methods such as antithetic resampling can improve the efficiency of Monte Carlo simulations, or distributed computation can be used if the hardware is available. ER

T5. Results for the four test cases	
Case	Value (€)
1	857.43
2	1458.46
3	1425.66
4	1240.52

Les Clewlow and Chris Strickland are the founders and directors of Lacima Group, where John Breslin is a principal and Calvin Kwok a quantitative analyst. Mattias Pfau is director of pricing at E.ON Energy Trading.
Email: info@lacimagroup.com

References

Breslin J, L Clewlow, C Kwok and C Strickland, 2008
Gaining from complexity: MFMC models
Energy Risk, April 2008

Clewlow L and C Strickland, 1998
Implementing derivative models
John Wiley, London

Ravindran K, 1995
Low-fat spreads
Over the rainbow, Risk Publications, pp. 141–142

Kirk E, 1995
Correlation in the Energy Markets
Managing Energy Price Risk, Risk Publications and Enron, pp. 71-78

Hikspoors S and S Jaimungal, 2007
Energy spot price models and spread options pricing
IJTAF 10 (7), pp. 1,111–1.135