Gas Storage:
Spot Optimisation

Authors:

John Breslin
Les Clewlow
Calvin Kwok
Chris Strickland
Daniel van der Zee

Lacima Group

February 2009

This article is the third in our gas storage thought leadership series, which describe the most common methodologies for valuing a gas storage facility. The previous two articles focused on trading strategies using forward or option contracts. In our most recent article we showed how a rolling intrinsic strategy can be used to capture extra value from a storage facility as the market evolves through time. But although this is a dynamic strategy it is suboptimal in that trading decisions are based solely on the prevailing market prices, without regard to potential future opportunities. In this article we describe a spot trading strategy which maximises the total expected cash flow from the storage facility using a trinomial tree based optimisation model. While the trading strategy determines the daily decisions for injection or withdrawal in the spot market it is relatively straightforward to implement an equivalent strategy in the forward market through delta hedging.

The key feature of this valuation methodology is that the underlying spot price model is consistent with and calibrated to the market forward curve dynamics. This ensures that the values obtained will be consistent with the forward based strategies as discussed in our previous articles. In this article our model for the spot price is the single factor mean reverting model (Clewlow and Strickland, 2000), in which the spot price is consistent with the market forward curve and spot volatilities. This model can be described by the equation:

$$\frac{dS(t)}{S(t)} = \left[\frac{\partial \ln F(0,t)}{\partial t} + \alpha\left[\ln F(0,t) - \ln S(t)\right] + \frac{1}{2}\alpha\int_0^t \sigma(u)^2 \exp\left(-2\alpha(t-u)\right)du\right]dt + \sigma(t)dz(t)$$

where $S(t)$ is the spot price at time t, $F(t,T)$ is the forward price at time t with maturity date T, α is the mean reversion rate, $\sigma(t)$ is the spot price volatility at time t, and $dz(t)$ is an increment in a Wiener process. Clewlow and Strickland (2000) describe how a trinomial tree can be constructed to represent this model (see also Clark, et al (2008)).

We have discussed trinomial trees in previous articles of this series – see for example Clark, et al (2008a) and Clark, et al (2008b). Each node in the resulting tree can be identified by a pair of integers (i, j) where $i = 0,\ldots,N$ is the time step and $j = -i,\ldots,i$ is the level of the asset price relative to the initial asset price. At node (i, j) the date is t_i where we assume $t_0 = 0$ and the spot price is $S_{i,j}$ with $S_{0,0} = S(0)$, the initial spot price. Each node has three branches which connect with nodes at the next time step and each branch has an associated transition probability which we denote $p_{u,i,j}$, $p_{m,i,j}$, and $p_{d,i,j}$ for the 'up', 'middle', and 'down' probabilities respectively. The resulting trinomial tree is represented graphically in Figure 1.

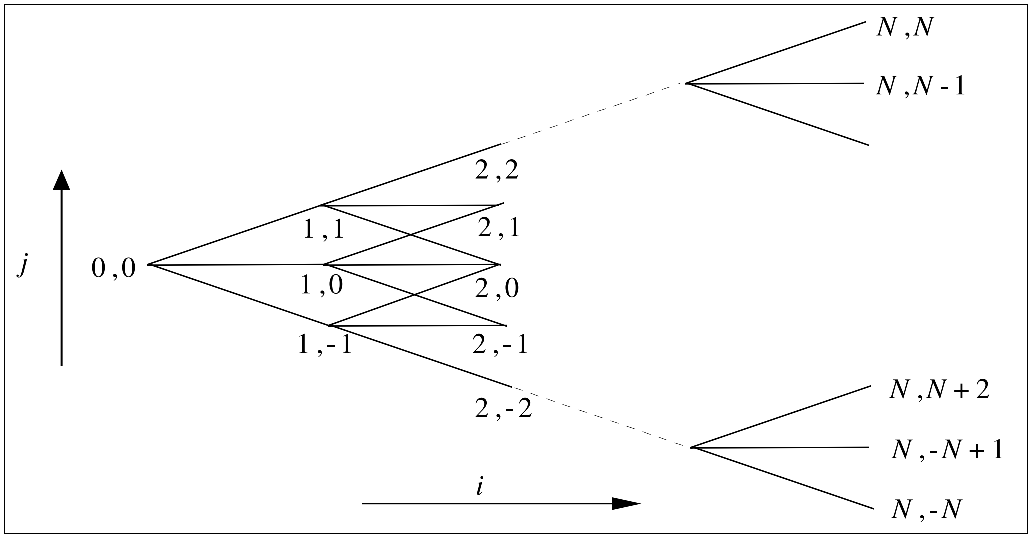

Figure 1 : Nodes in the trinomial tree

In order to evaluate a storage strategy we need to represent the state of the storage asset at each node of the tree. We do this by adding an extra dimension to the tree so that at each node we have a discrete set of possible values for the amount of gas in the storage asset $V_{i,j,k}$ where $k=1,\ldots,N_V(i)$, where $N_V(i)$ is the number of discrete volume levels at time step *i*. We define $V_{MIN}(i)$ and $V_{MAX}(i)$ as the constraints on the minimum and maximum amount of gas which can be in storage on each trading date. The size of the steps in the volume discretisation is given by $\Delta V(i)=\left(V_{MAX}(i)-V_{MIN}(i)\right)/\left(N_V(i)-1\right)$, so we have $V_{i,j,k}=V_{MIN}(i)+(k-1)\Delta V$.

Now for each state of the market (ie, spot price) and storage asset (ie, volume in storage) we need to choose the optimal operating decision, $\Delta V_{i,j,k}$. We assume a discrete set of possible decisions which are typically:

- inject the maximum possible, $V_I(i,k)$, which may be a function of the trading date and the volume of gas in storage,
- do nothing, or
- withdraw the maximum possible, $V_W(i,k)$, which also may be a function of the trading date and the volume of gas in storage.

The evaluation procedure works in the same way as the standard dynamic programming method in a trinomial tree used for path dependent derivatives such as Swing options or American options. The first step is to calculate the final decisions and payoffs on the last trading date. This is straightforward, since there are no future constraints or cashflows to consider, we simply choose the decision which maximises the cashflow on the final date for each spot price $S_{N,j}; j = -i, \ldots, i$ and each volume of gas in storage $V_{N,j,k}; k = 1, \ldots, N_V(N)$. We then step backwards through the tree calculating the decisions for each spot price and volume of gas in storage which maximise the sum of the current cashflow and expected value of the future cashflows. The expected value of the future cashflows is given by:

$$p_{u,i,j} C_u(\Delta V_{i,j,k}) + p_{m,i,j} C_m(\Delta V_{i,j,k}) + p_{d,i,j} C_d(\Delta V_{i,j,k})$$

where $p_{u,i,j}$, $p_{m,i,j}$, $p_{d,i,j}$, are the probabilities for the up, middle and down transitions at node i, j, and $C_u(\Delta V_{i,j,k}), C_m(\Delta V_{i,j,k}), C_d(\Delta V_{i,j,k})$ are the expected values of the storage asset at the nodes at the end of the up, middle and down branches when the volume of gas in storage at the respective up, middle or down node given by the current value $V_{i,j,k}$ plus the current decision $\Delta V_{i,j,k}$. If the volume of gas in storage given by the current decision is not explicitly represented at the up, middle and down nodes, then interpolation can be used to obtain an estimate of the expected value. Figure 2 illustrates the process of searching for the decision which maximises the expected value of future cashflows.

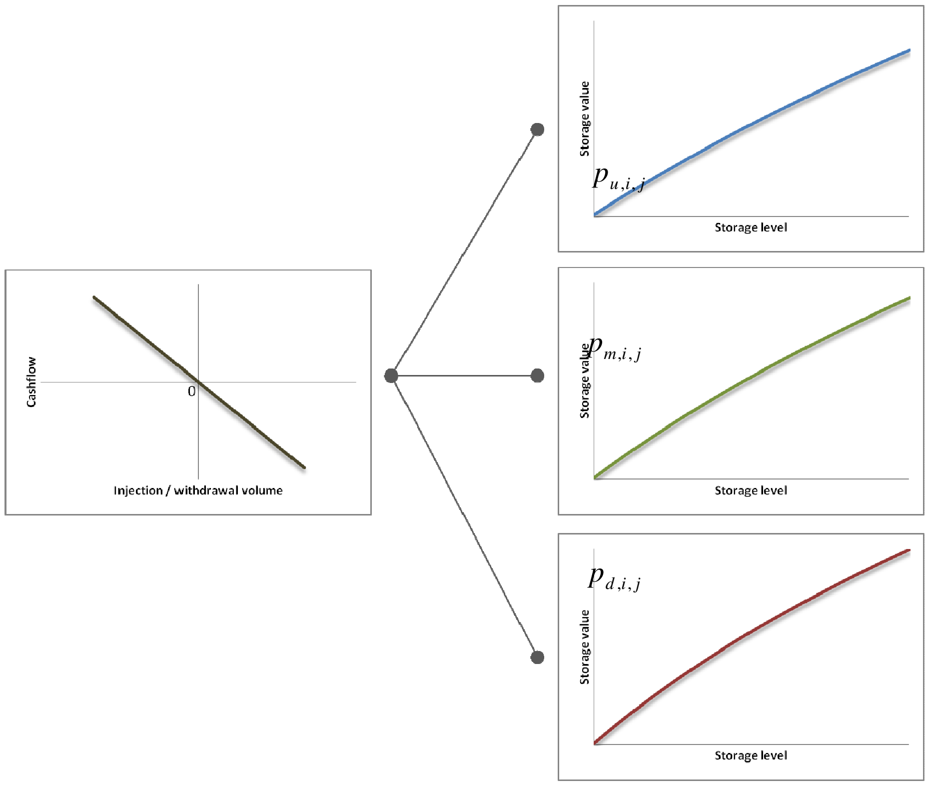

Figure 2 : Illustration of searching for the optimal storage decision

In Figure 2, the left-hand box illustrates the cashflow at the node as a function of the decision. Negative volumes indicating withdrawal of gas giving a positive cashflow from the sale of the withdrawn gas into the market and positive volumes indicating injection with a negative cashflow from the purchase of the injected gas from the market. The boxes on the right-hand side indicate the value of the storage asset as a function of the quantity of gas in the storage asset. It is this method of looking forward and taking account of the probabilities of future outcomes in the tree based optimisation which allows it to outperform the Rolling Intrinsic strategy.

Each decision must also be tested for validity i.e. that it does not result in a volume of gas in storage which violates the minimum or maximum allowed at each time step.

Using the same example storage facility as in our previous article (see Table 1) we have set up an optimisation using a trinomial tree, as described above. We use the NBP gas price from March 31, 2007 (see Figure 3) as our initial forward curve, and have calculated the mean reversion rate and quarterly spot volatilities using the last four years of spot price data.

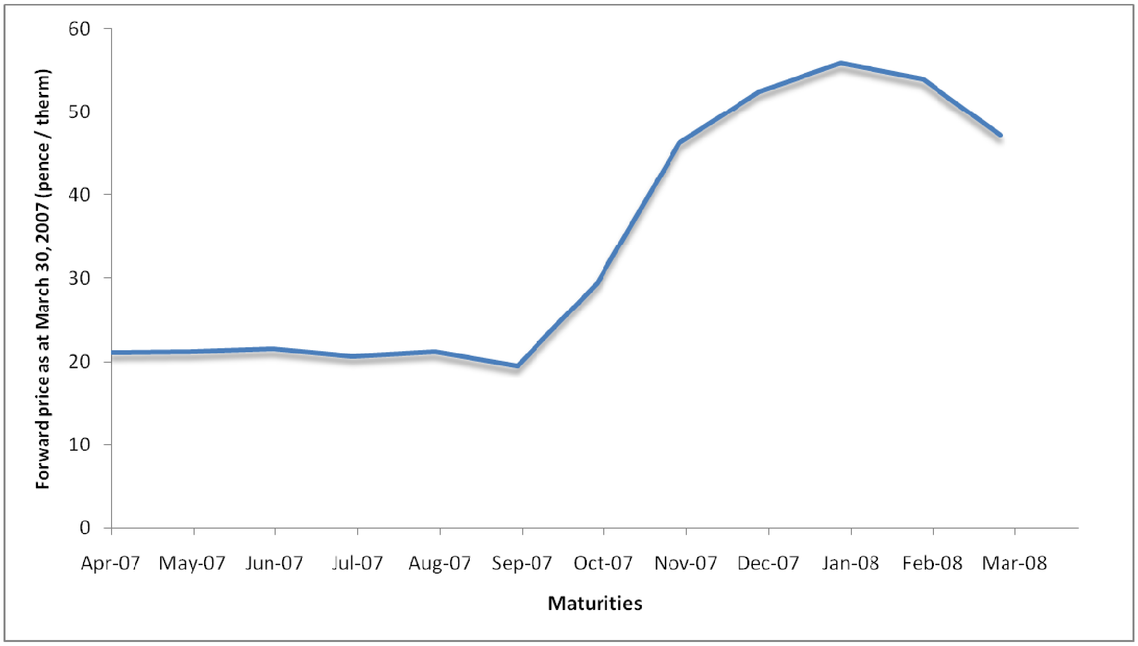

Figure 3 : NBP gas forward prices on March 31, 2007

- Total capacity: 1,000,000 MMBtu (or 1 Bcf)
- Maximum injection rate: 8,197 MMBtu/day
- Maximum withdrawal rate: 16,393 MMBtu/day
- Injection cost: 0.010 pence/therm
- Withdrawal cost: 0.006 pence/therm
- The valuation period is from April 1, 2007 to March 31, 2008, with the valuation being performed as at March 31, 2007.
- The original and terminal constraints are that the facility must be empty on the start and end dates.
- Assume a flat discount rate of 3.5% for the valuation period.

Table 1: Details of the example storage facility

The spot optimisation gives an expected present value for the storage facility of £4,188,597. Figure 4 plots the mean storage level and various percentiles depicting the distribution of those levels for each month in the contract period. In Figure 5 we show a similar representation of the distribution of injection and withdrawal levels for each month. Although the overall patterns of both graphs appear to be quite similar to those with the rolling intrinsic case with injections mainly take place in summer and withdrawals in winter, it is clear that the spot optimisation is able to take advantage of a much wider range of profitable scenarios with significant withdrawals in the Summer period and injections in the Winter period. It is this adaptability of the optimisation approach which allows it to capture additional value over the rolling intrinsic approach.

As part of the valuation routine we also obtain the optimal decision surface. This surface is obtained at each time slice in the tree (i.e. typically daily for gas storage problems) and represents the optimal injection and withdrawal decision for each possible storage level and each possible spot price in the tree at that time slice. The decision surface for January 1, 2008 is depicted in Figure 6. In general the strategy is to inject if the current spot price is relatively low and withdraw if it is relatively high. The key to the optimisation strategy is that the critical spot price at which the decision switches between injection and withdrawal varies with the amount of gas in the storage facility. The critical price decreases as the volume in storage increases reflecting the constraint that the facility must be empty at the end of the evaluation period.

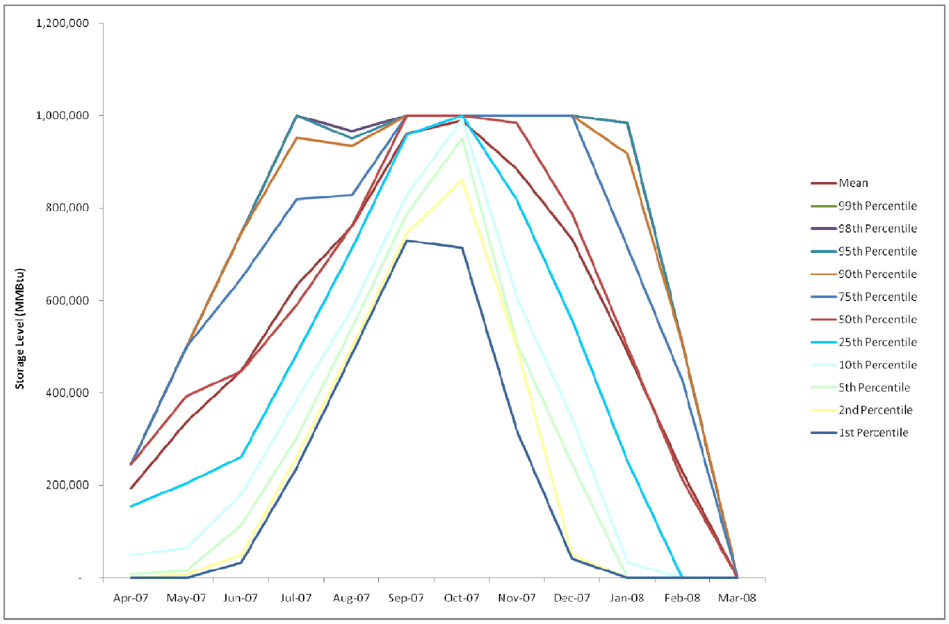

Figure 4 : Distribution of storage levels in each month

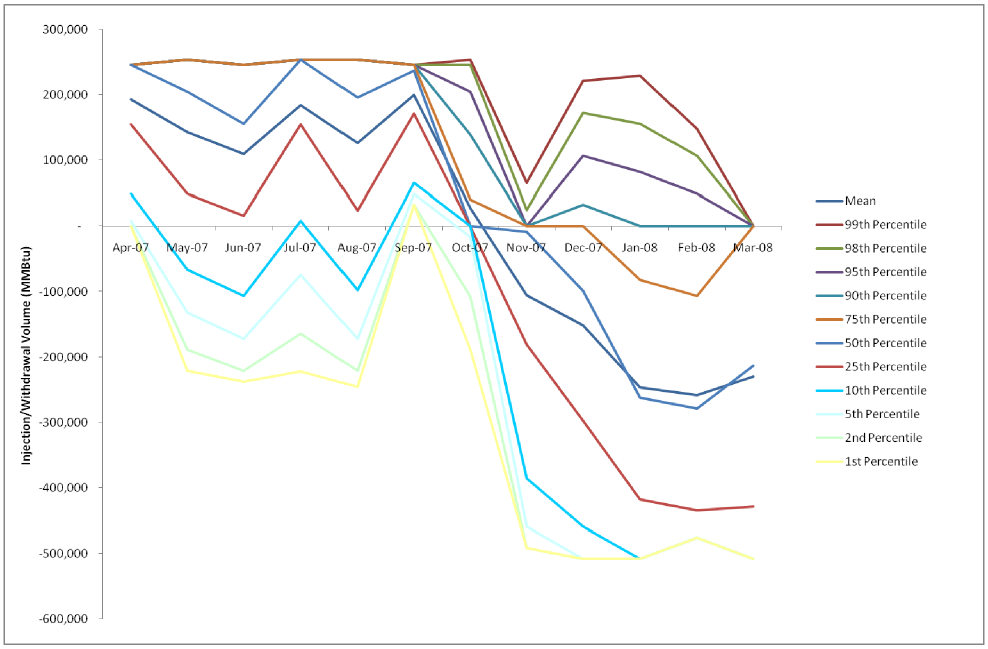

Figure 5 : Distribution of injections/withdrawals in each month

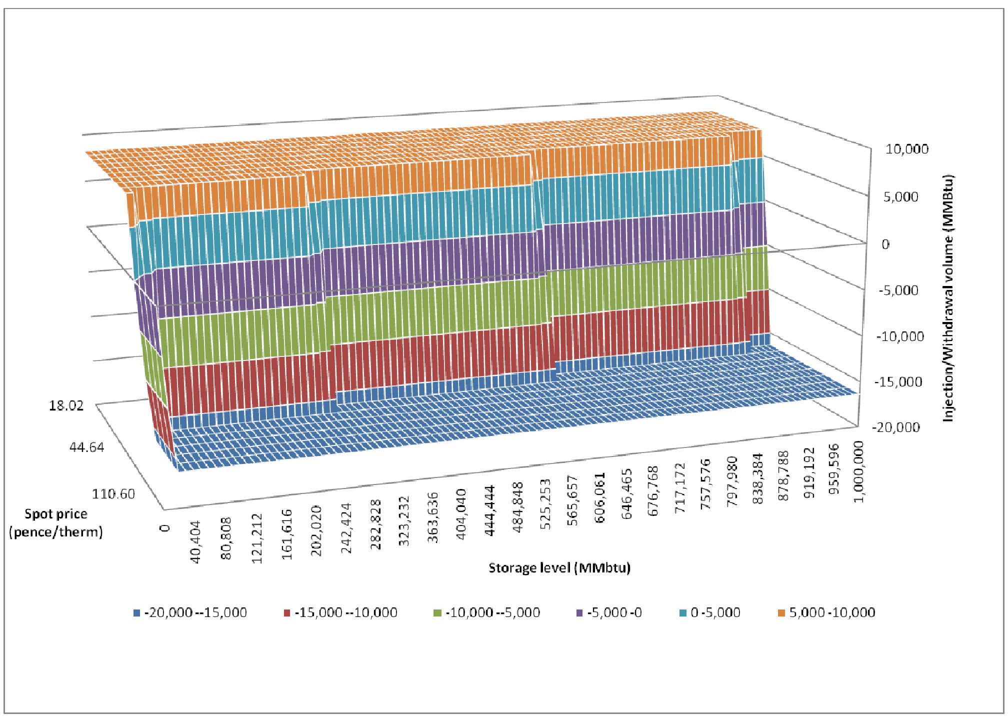

Figure 6 : Decision surface on January 1, 2008 (positive volumes = injection; negative volumes = withdrawal)

Two keys parameters of the single factor model used in building the tree are the time-varying spot volatility and the mean reversion rate. The impact of spot volatility on the expected storage value is shown in Figure 7. In producing this figure, the mean reversion rate is fixed at the estimated value which is 9.51. When the spot volatility is zero, the expected storage value is equal to the intrinsic value simply because the tree used for the valuation is just the forward curve. Other than this, it is clear from the figure that the expected storage value increases as the spot price becomes more volatile because higher volatility increases the probability of larger spreads between the injection and withdrawal prices.

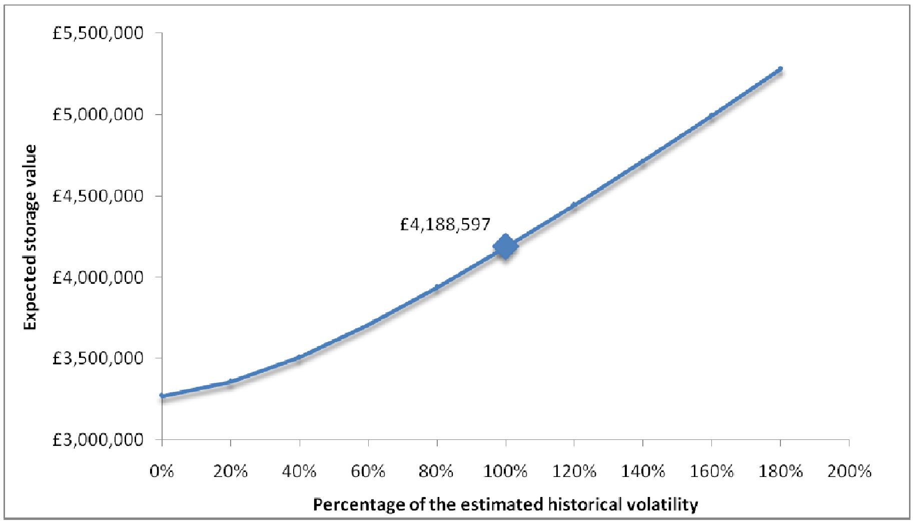

Figure 7: The effect of changing spot volatilities on the expected storage value

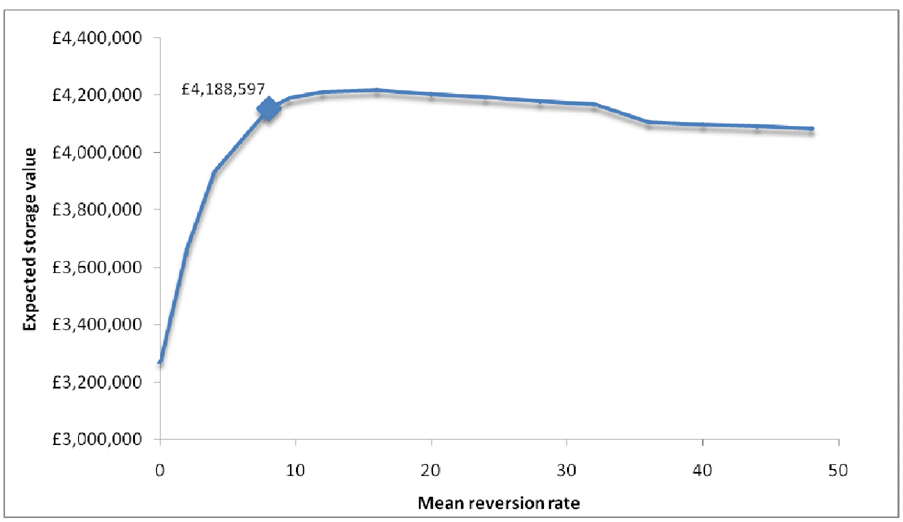

Figure 8: The effect of changing mean reversion rate on the expected storage value

Figure 8 shows the expected storage value as a function of the mean reversion rate while the quarterly volatilities are fixed at the estimated values. It can be seen from the figure that the function is concave and the expected storage value starts at the lowest when the mean reversion rate is zero. This level of mean reversion implies a Black-Scholes type model where the gas spot price at any time step has equal probabilities of moving up and down. In this case there is zero probability of the existing spreads in the forward increasing since the whole forward curve moves in parallel shifts. Therefore, the expected storage value in this case is equal to the intrinsic value.

As the mean reversion rate increases, the probability of the seeing larger spreads than in the initial forward curve increases, as the front end of the forward curve moves more than the long end. The figure shows that the highest expected storage value is attained when the mean reversion rate is about 16. Above this level of mean reversion rate, no additional value can be captured since the increasing mean reversion rate causes the effective volatility to reduce which in turn leads to the decline of the expected storage value. As the mean reversion continues to increase, the spot price tree will eventually collapse to the forward prices and the expected storage value under which will become just the intrinsic value.

The tree based optimisation allows a trading strategy which takes account of the future distribution of spot prices to be efficiently calculated. In order to implement the strategy we need to calculate appropriate positions in the forward contracts. This can be achieved by calculating the forward contract deltas of the storage using the tree. Figure 9 shows the daily for the first month and then monthly forward deltas on 1-Apr-07.

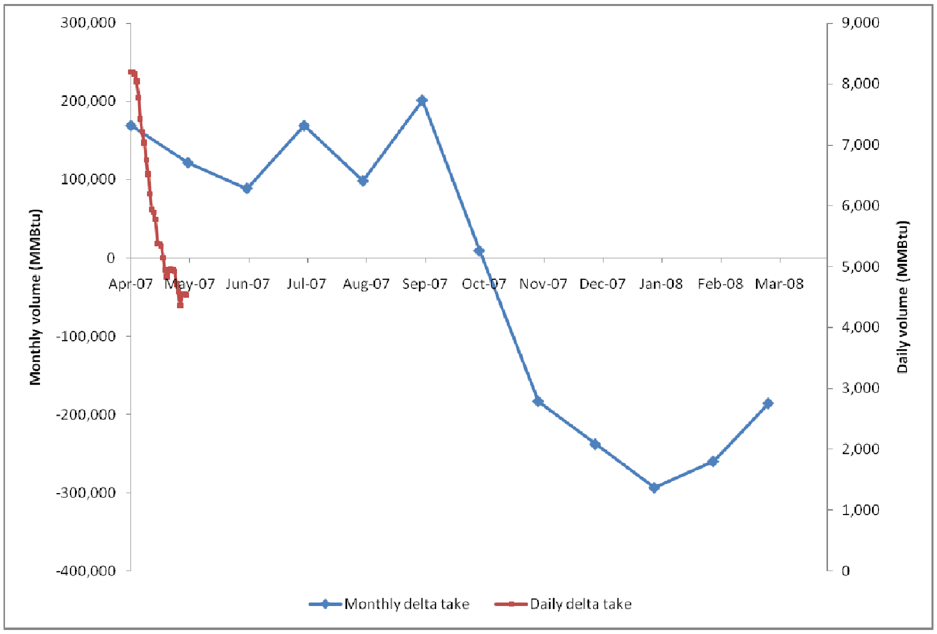

Figure 9 : The daily and monthly forward deltas on the April 1, 2007

The daily forward delta for the next day is very close to the maximum daily injection rate since the volume that will be required to inject is almost certain. As we look further ahead in April the delta positions decline reflecting the fact that the forward curve may move such that injecting the maximum amount may not be optimal. The monthly deltas are qualitatively similar to the intrinsic positions in that they reflect injection over the period Apr-07 to Sep-07 and withdrawal over the period Oct-07 to Mar-08. However, they are importantly different in that they take into account the future forward curve dynamics and that the forward positions will be adjusted regularly to capture additional value.

Figure 10 shows an example of how the daily forward deltas change if the spot price follows the initial forward curve.

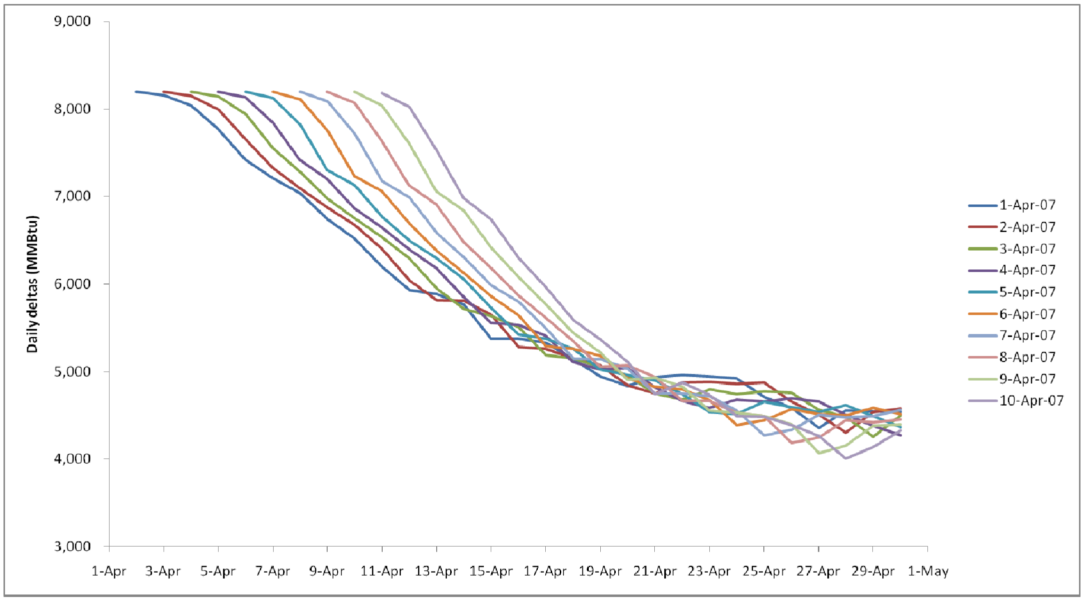

Figure 10 : Example of the evolution of the daily forward deltas from the 1-Apr-07 to 10-Apr-07

Notice how the daily positions for the months are gradually increased to the maximum injection amount as the uncertainty is gradually resolved.

Finally, in order to calculate a value for the storage based on the forward deltas trading strategy we simulate this strategy using the multi factor model described in the previous article. The expected value of the storage facility using this approach is £4,313,125, which is over 11% higher than the greatest value (i.e. £3,858,870 with daily rebalancing of a daily hedge) we could achieve using the rolling intrinsic method.[1]

References

Breslin J, Clewlow L, Elbert T, Kwok C, and Strickland C, 2008a, "*Gas storage: overview and static valuation*", Energy Risk, November 2008

Breslin J, Clewlow L, Elbert T, Kwok C, Strickland C, and van der Zee D, 2008b, "*Gas storage: rolling intrinsic valuation*", Energy Risk, December 2008

Clark J, Clewlow L, and Strickland C, 2008, "*Calibrating trees to the market prices of options*", Energy Risk, August 2008

[1] As discussed in our previous article, rolling intrinsic with daily rebalancing of a daily hedge is generally unrealistic. A more practical case would be rolling intrinsic with 7-day rebalancing period of a monthly hedge which yielded an expected storage value of £3,696,389.

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

About Lacima Group

Lacima Group is a specialist provider of energy and commodity pricing, valuation and risk management software and advisory services. Based on its internationally acclaimed research in energy risk modelling, Lacima's solutions help energy trading organisations to effectively quantify and manage risks associated with structured contracts and physical assets across multiple commodities and regions. For further information, visit www.lacimagroup.com or email info@lacimagroup.com.