

# To store or not to store

Here we describe the optimal operation and valuation of gas storage based on a real option methodology. Using Zeebrugge gas prices as a practical example, *Cyriel de Jong* and *Kasper Walet* clarify the optionality in gas storage, analyse its valuation and discuss hedging strategies to secure part of the storage value

**S**torage plays an increasingly important role in matching supply with demand in natural gas markets. However, companies waste significant storage value, if optimal investment and operational decisions are solely based on the seasonality pattern in gas prices: daily volatility creates significant additional value. Storage facility owners can reap this value using a recently developed real option technique.

## Optimisation

Optimal operation of a natural gas storage facility comes down to finding the right time to withdraw and inject gas, depending on current and expected gas prices<sup>1</sup>. In short, two characteristics of gas prices allow a storage operator to maximise its value: predictable price movements (seasonality) and unpredictable price movements (volatility).

Seasonality is generally considered the main source of gas price profits – for example, buying in summer when gas is cheap and selling in winter when prices rise. However, completely relying on expected price movements ignores the real option – or flexibility – value of storage. For example, even though we observe high prices today, prices may climb even higher in a few days. Similarly, we may store some extra gas to be prepared for some (unexpected) further profitable days of high prices.

Why is storage optimisation a complex problem? The complexity arises from a few factors:

- *Seasonality in gas prices*: this means that what would be an optimal decision in December (under otherwise similar circumstances) is different from what would be an optimal decision in May.
- *Maximum and minimum volume of storage (and other operational constraints)*: this means that what constitutes an optimal action varies depending on whether storage is nearly full or nearly empty.
- *Non-normal price characteristics, such as mean reversion and time-varying volatility and jumps*: this factor makes expected price movements and their magnitude vary from day to day.
- *Long time horizons*: due to yearly seasonality, at least a full year – and preferably a few years – should be analysed.

A combination of these factors makes the problem a complex compound real option problem, for which we need to use advanced option techniques.

For example, global energy consultancy Maycroft Consultancy Services uses Monte Carlo simulations in its MayStore storage model, since simulations do not impose a rigid structure on the spot prices and allow for accurate decision-making for both short and long time horizons. More specifically, MayStore is based on the recently developed least-squares Monte Carlo technique that combines flexible Monte Carlo simulations with fast and accurate least-squares regressions (Longstaff and Schwartz (2001)).

<sup>1</sup> In this article we describe the valuation of natural gas storage. Storage of other commodities, such as oil and liquefied natural gas, may be valued with the same approach, depending on the availability of reasonably liquid spot markets and seasonality and volatility in prices

The approach requires a relatively long computation time, but is flexible and precise. Carriere (1996) set out the idea for it some years ago, and it was later popularised by Longstaff and Schwartz (2001). Since then, analysts have increasingly used it to value complex financial and physical options. We go on to describe the model and discuss a case study.

## Case description

The case centres on a realistic, but stylised example of a gas storage facility that is connected to the Zeebrugge gas hub in Belgium. Characteristics of storage may vary in terms of minimum and maximum working volumes, injection and withdrawal (production) rates and operational costs. We look at a storage facility with a working volume of 8 million gigajoules (GJ), an injection capacity of 60,000 GJ a day and a production capacity of 250,000 GJ a day. This means it takes around 133 days to fill the facility and 32 days to eject the total working volume.

We take a horizon of two years, and assume that by the end of the two years, the volume in storage should be at least as big as the starting volume. For the starting volume, we assume half the working volume. For simplicity, we do not include operational costs or trading costs (the bid-ask spread), but the model allows for their inclusion.

Our aim is to arrive at a strategy that consists of daily trading decisions, which depend on actual price and volume levels and expectations about future levels. On average, the strategy should maximise the total discounted revenue across all possible price scenarios over a given time horizon, with an annual discount rate ( $r$ ) of 5%<sup>2</sup>. If we denote individual days by  $t$ , the terminal trading day by  $T$ , the storage level at time  $t$  by  $L_t$  and the spot price by  $P_t$ , then our goal is to maximise the following expected storage value:

$$V = E \left[ \sum_{t=1}^T (L_t - L_{t+1}) \cdot P_t \cdot e^{-rt} \right] \quad (1)$$

## Gas spot price characteristics

MayStore uses a pricing benchmark that determines how the storage facility is operated. This could be a day-ahead or within-day spot price at a physical gas hub or non-physical trading point. It is important that the prices of the benchmark closely resemble the prices for which the gas in the facility can be traded on a daily basis.

We use Zeebrugge day-ahead prices, since Zeebrugge is the most liquid hub on the European continent. We use Platt's day-ahead prices, measured in €/GJ from March 2000 till August 22, 2003 (figure 1). Typical gas characteristics become immediately apparent, including high volatility, seasonality and a tendency – after extreme prices – to revert to 'normal' price levels (mean reversion).

Prices are available on the usual trading days – 252 days per full year, on average – and we assume injections and withdrawals only take place on these days. The basis simulation model in MayStore is the

<sup>2</sup> At the end of the article we discuss hedging opportunities to secure (part of) the storage value. Depending on the effectiveness of hedging, we may employ a discount rate close to the risk-free rate

mean-reverting model that has proved successful for oil and gas (Schwartz (1997)). However, more advanced models are often desirable – for example, to allow for jumps, time-varying seasonal volatility and time-varying autoregressive volatility (Garch). Such price modelling flexibility is one of the strengths of the least-squares Monte Carlo approach. Here we follow the basic mean-reverting model, for ease of explanation. It assumes the following form:

$$\ln S_t = \ln S_{t-1} + \alpha (\mu - \ln S_{t-1}) + \sigma \varepsilon_t \quad \text{where } \varepsilon_t \sim N(0,1) \quad (2)$$

In this model, there are two crucial parameters: the mean-reversion rate  $\alpha$ , which determines how quickly prices revert from unusual to normal levels; and the volatility  $\sigma$ , which determines the magnitude of the unexpected price fluctuations. We also estimate a seasonal function to capture readily observable seasonality. Although this function will later be replaced by the (forward-looking) seasonality derived from the forward market curve, we need to improve the estimates of mean-reversion and volatility. The highest prices are observed in winter months, the lowest in summer, with a maximum difference of 1.12 €/GJ.

The mean-reversion rate of 0.0678 means the half-life is 3.88 – meaning the spot price takes around four trading days to revert from actual levels halfway back to its expected level as determined by the seasonal pattern. The residuals of the spot price process ( $\varepsilon_t$  in equation 2) have a daily standard deviation of 8.4%, which equals a yearly volatility of 133%. This confirms that volatility levels of natural gas spot prices are quite high in comparison to those of most other commodities – electricity being one exception – meaning gas storage has significant value in terms of options.

Since the model is calibrated with historical data, it lacks the most recent trading information. As a result, we re-align the spot price model to market forward or futures prices. If markets are efficient and risk premiums negligible, the forward curve represents the average expectation of market participants as to future spot levels. Platt's reported forward prices on August 22, 2003 for the next three months, as well as eight quarters ahead (see table).

We replace our historically calibrated seasonal function with a seasonal pattern that matches the forward mid prices. To obtain greater precision, we decompose the quarterly forwards in three monthly forwards based on historical monthly averages. We can make a similar adjustment to the volatility pattern, which may be aligned to implied volatilities derived from natural gas options.

### Optimal strategy

We start with an optimal strategy that only uses the forward market, thereby profiting from the predictable seasonality pattern in gas prices. The reader may verify that an optimal strategy consists of selling as much gas as possible in the two most expensive quarters (Q1 2004 and Q1 2005) and buying enough gas in the cheaper periods. This ensures the storage is full – that is, contains 8 million GJ – at the start of the most expensive quarters and attains the required minimum level of 4,000 GJ two years from now. The strategy yields a total value of €15.9 million.

Although this approach is intuitive, it ignores an important characteristic of gas storage: the flexibility to react to changing market conditions. To translate the unexpected price movements into profits, we need a real option model. To understand how the real option model works, we start with a simple example.

Suppose our storage is contains 5 million GJ and we observe a spot price of 3 €/GJ today and have three choices: do nothing, inject 0.06 million GJ or withdraw 0.25 million GJ, as we assume in our example. Each choice will lead to immediate cashflows and a value of the storage next period as follows:

### Natural gas forward prices in € per gigajoule

	Bid	Ask	Mid
September 2003	2.81	2.84	2.83
October 2003	2.79	2.82	2.81
November 2003	3.38	3.41	3.40
Q4, 2003	3.27	3.30	3.29
Q1, 2004	3.73	3.76	3.75
Q2, 2004	2.73	2.79	2.76
Q3, 2004	2.79	2.86	2.83
Q4, 2004	3.36	3.39	3.38
Q1, 2005	3.73	3.75	3.74
Q2, 2005	2.66	2.73	2.70
Q3, 2005	2.73	2.79	2.76

*Source: Platt's, Moneyline*

- Do nothing: Value = Value of having 5 million GJ next period
- Inject value = Value of having 5.06 million GJ next period – €180,000
- Withdraw value = Value of having 4.75 million GJ next period + €750,000

For example, we will choose to inject if the extra volume (0.06 million GJ) is worth more than €180,000, and choose to withdraw when the storage decline (0.25 million GJ) is worth less than €750,000. More specifically, in each time step we calculate three values, ignoring discounting, for each of the choices above and choose the maximum:

$$VN_t = E[V_{t+1}(L_t)] \quad (3)$$

$$VI_t = E[V_{t+1}(L_t + IR)] - P_t \cdot IR \quad (4)$$

$$VW_t = E[V_{t+1}(L_t - WR)] + P_t \cdot WR \quad (5)$$

$$V_t = \max(VN_t, VI_t, VW_t) \quad (6)$$

### The role of storage

Storage plays a vital role in competitive natural gas markets, because the average variability in the consumption of natural gas is much greater than the average variability in production. Historically, natural gas storage was used for two main functions.

First, it provides local distribution companies with adequate supply during periods of heavy demand by supplementing pipeline capacity and serving as back-up supply in case of an interruption in wellhead production.

Second, storage enables greater system efficiency: instead of satisfying winter demand by adding new production facilities, the industry can maintain production at a much more constant level throughout the year.

In the liberalisation process, the natural gas storage service is unbundled from the sales and transportation services, meaning storage is offered as a distinct, separately charged service. In combination with the development of active spot and futures markets, it becomes possible to adjust trading decisions to price conditions.

In other words, buyers and sellers of natural gas have the possibility to use storage capacity to take advantages of the volatility in prices.

In equations 3 to 6,  $IR$  and  $WR$  are the injection and withdrawal rates.  $VN$ ,  $VI$  and  $VW$  are the storage values if we choose to do nothing, inject and withdraw respectively.

The difficulty in making this decision is that we do not know the next period's value, since it depends on future price levels. The only exception is the last trading day of the current evaluation horizon, when our decision should ensure we reach the level of 4 million GJ. For the other days, we can only calculate an estimate of future storage value by simulating many possible price paths. This is what the least-squares Monte Carlo approach entails: simulating possible price paths, starting with a decision on the last day, and then working backwards to derive an optimal strategy that maximises total value. The estimates of future storage value are based on least-squares regressions from a cross-section of the prices across all simulations. The fact that the least-squares regressions are precise and executed very efficiently is the strength of the least-squares Monte Carlo method.

Next we generate a large number of simulations with the spot price process. We then work backwards and perform the aforementioned regressions (for  $VN$ ,  $VI$  and  $VW$ ) on each trading day and for each possible volume level. The regressions provide us with an estimate of the storage value if we do nothing, inject or withdraw. We can use this information to derive the optimal strategy for any price path.

Most importantly, for general day-to-day trading, Maycroft's MayStore determines an optimal strategy for today's market prices. With the current storage level known, MayStore advises withdrawing or injecting above and below certain price levels. Figure 2 depicts this option exercise boundary.

## Strategy analysis

To assess the effect of the decision rules on profitability, cashflows and storage levels over time, we generate a new set of price simulations<sup>3</sup> and determine the optimal strategy along each price path. We thus obtain a distribution of storage value, storage volume and cash flows. Using 1,000 simulations, the average storage value for the next two years equals just over €29.4 million: nearly twice the 'lock-in' value. Although there is quite some dispersion of values across the simulations, the value is below the lock-in value in only three instances. (See figure 3.)

It can be inconvenient for schedulers not to know the volume levels in advance, as this makes it difficult to contract pipeline capacity forward. However, the model generates a distribution of volume levels over time. They are not completely random – they follow the seasonality pattern on average. Schedulers can use these volumetric statistics to project the need for pipeline capacity and so can easily identify constraints in pipeline capacity. For example, traders will most

probably buy natural gas in the first 50 days – between the end of August and the end of October – so schedulers may contract transport (inflow) capacity in advance for this period. Schedulers can feed this information back into the model by adapting injection and withdrawal capacities in constrained periods. (See figure 4.)

## Sensitivity analysis

Our case discussed one particular storage facility based on one set of historical prices and forward levels. It is interesting to analyse how sensitive the outcomes are with regard to variations in the spot model parameters and storage characteristics.

We start with adaptations to the spot model: the mean-reversion rate, volatility and seasonality. Without any mean reversion, price changes would be completely unpredictable, apart from their seasonality. In such a situation, locking in prices is the best strategy. Hence, volatility alone is not enough to create option value for storage: we also need mean reversion<sup>4</sup>. This becomes clear when we vary the mean-reversion rate between 0 and 0.10 (see figure 5). Without mean reversion, the asset contains no option value: its value equals the value that can be readily locked in using the forward curve. Slightly increasing the mean-reversion level quickly generates extra profits, but increasing mean reversion (above approximately 0.05) introduces a counter-effect: a rate that is too high dampens out price fluctuations very quickly, and thus reduces the width of 'unexpected cycles'. Above this threshold, therefore, the real option value starts to diminish progressively.

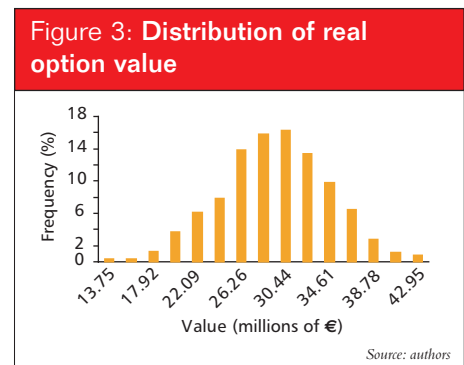

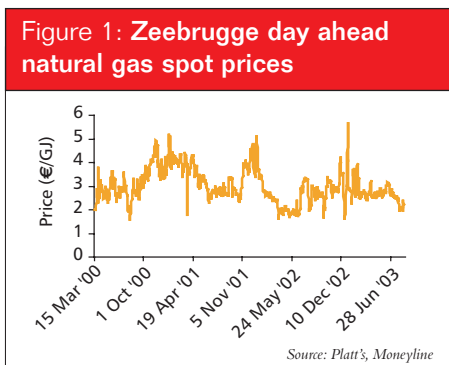
The relationship between volatility and real option value is close to linear. In our example, each percentage point in yearly volatility (daily volatility multiplied by the square root of 252) yields around 1% in additional option value. Although volatility has a positive effect on asset value, we should bear in mind that actual payouts become more uncertain as a result, as do expected volume levels.

Storage facilities have varying levels of operational flexibility, which are best reflected by the speed of injection and withdrawal. The higher the level of flexibility, the more quickly the facility can respond to both expected price cycles (due to seasonality) and unexpected price cycles (due to mean reversion and volatility). If we cut flexibility by a factor of two, and keep the injection/withdrawal ratio constant – that is, 30,000 GJ injection and 125,000 GJ withdrawal – asset value drops by 38%. We can attribute this drop to a reduction due to lower seasonality profits (22%) and a reduction due to lower real option value (16%). A flexibility twice as high as the base case – that is, 120,000 GJ injection and 500,000 GJ withdrawal – only has a minor impact on the seasonality value (+2%), but a considerable impact on the real option value (+57%).


So, above a certain level, we do not need additional flexibility to

<sup>3</sup> We generate a new set of simulations to perform an out-of-sample assessment, thereby avoiding any in-sample bias. This is one of the advantages compared to a tree-building approach, where we can make no distinction between in- and out-of-sample


<sup>4</sup> This is different from standard call and put options, because a storage facility exploits 'cycles' in gas prices to buy low and sell high and not price fluctuations per se. Cycles can be produced either by predictable movements (seasonality) or unpredictable movements (volatility plus mean reversion)




**Figure 4: Distribution of storage volume over time**



**Figure 5: Sensitivity of real option value to mean reversion and volatility**



**Figure 6: Delta-sensitivity of storage value to changes in forward prices**



exploit the forward curve much further, since one month – the smallest forward period in our example – suffices to fill or empty the facility completely. Yet additional flexibility allows for a quicker reaction on shorter-lived unexpected price fluctuations at all levels of flexibility.

### Hedging

When we follow a strategy on the spot market, we do not know the value of the storage facility in advance, which may fluctuate over time. From a risk-return point of view, it is often desirable to minimise these fluctuations. This is possible by trading gas derivative contracts such as forwards, futures and options. We can find the optimal contracts by calculating the Greeks of the asset, which are its sensitivities to changes in the underlying market prices. For example, the model calculates the delta of the asset towards changes in each of the forward prices. The delta tells a trader which positions neutralise the portfolio's value (storage plus contracts) to small changes in one of the forward contract prices.

The deltas in figure 6 are directly related to the expected storage development in figure 4. In periods where we would expect gas to be bought (September to December 2003), the deltas are most negative: higher forward prices make expected purchase prices more expensive and thus reduce storage value. Similarly, in periods where we would expect gas to be sold – notably the first quarter of 2004 and 2005, which have the highest expected price levels – higher forward prices enhance the storage value.

Mathematically, the deltas represent the first-order derivative of the storage value with respect to the underlying forward. This means the storage value is expected to decrease by around €21,000 (2.1 million x 0.01) in response to an increase of the October 2003 forward by €0.01. To make the portfolio value independent to small changes in this forward contract, which has a current value of €2.81,

a trader should go long in 0.75 million (2.1 million/2.81) GJ October 2003 contracts.

### Conclusion

Storage facilities will play an increasingly important role in liberalised natural gas markets in matching time-varying demand to supply. For optimal investment decisions, we must accurately determine the value of storage, taking into account both the intrinsic value (due to seasonality) and the option or flexibility value (due to volatility). Accurate valuation based on the flexible least-squares Monte Carlo approach may reduce investment uncertainty. Such increased certainty, combined with the substantial real option value – as we might find in a Zeebrugge storage facility – may trigger even more investments in this area. **EPRM**

**Cyriel de Jong** is director of research and **Kasper Walet** managing director at global energy consultancy Maycroft Consultancy Services in Amsterdam. **email:** dejong@maycroft.com; walet@maycroft.com

### References

- Carriere, J**, Valuation of early exercise price of options using simulations and nonparametric regression, *Insurance: Mathematics and Economics*, 19, pages 19–30, 1996.
- Dixit, AK and RS Pindyck**, *Investment under uncertainty*, Princeton University Press, Princeton NJ, 1994.
- Longstaff, FA and ES Schwartz**, Valuing American options by simulation: a simple least-square approach, *Review of Financial Studies*, 14:1, pages 113–147, 2001.
- Schwartz, ES**, The stochastic behavior of commodity prices: implications for valuation and hedging, *Journal of Finance*, 52:3, pages 923–973, 1997.

AD