

**FOR 20 2008**

**ISSN: 1500-4066**

OCTOBER 2008

Discussion paper

# Gas Storage Valuation: Price Modelling v. Optimization Methods

BY

**PETTER BJERKSUND, GUNNAR STENSLAND, AND FRANK VAGSTAD**

# Gas Storage Valuation: Price Modelling v. Optimization Methods

By

Petter Bjerksund\*, Gunnar Stensland\*\*, and Frank Vagstad\*\*\*

This version: September 26, 2008

## Abstract

The existence of a financial gas market motivates the analysis of gas storage as a separate asset, using the market value context for utilization and valuation. In the recent literature, gas storage is typically analysed within a framework with a simple one-factor price dynamics that is solved to optimality. We follow a different approach, where the market is represented by a forward curve with daily granularity, the price uncertainty is represented by six factors, and where we impose a simple and intuitive storage strategy that follows from repeated maximization of the intrinsic value.

Based on UK natural gas market price data, we obtain the gas storage value using our approach, and compare with results from one-factor models as well as with perfect foresight. We find that our approach captures much more of the true flexibility value than the one-factor models. Our results indicate that the appropriate framework for analyzing complex assets like gas storage is a rich representation of the price dynamics combined with a simple and intuitive decision rule.

**Key words:** Energy, uncertainty, flexibility, exercise strategy.

**Acknowledgements:** We have benefited from discussions with Statoil representatives, in particular Bente Spissøy and Bård Misund. The access to Statoil's gas market price database is gratefully acknowledged.

---

### Affiliation:

\* Bjerksund is Professor of Finance at the Norwegian School of Economics and Business Administration (NHH).

E-mail: [Petter.Bjerksund@nhh.no](mailto:Petter.Bjerksund@nhh.no)

\*\* Stensland is Professor of Finance at Norwegian School of Economics and Business Administration (NHH).

E-mail: [Gunnar.Stensland@nhh.no](mailto:Gunnar.Stensland@nhh.no)

\*\*\* Vagstad is a Business Analyst with Viz Risk Management.

E-mail: [Frank.Vagstad@viz.no](mailto:Frank.Vagstad@viz.no)

# Gas Storage Valuation: Price Modelling v. Optimization Methods

## 1. Introduction

The analysis of natural gas storage has traditionally been integrated with the valuation of other activities of the company, for instance production, supply, and demand. However, we know from economics and finance that well functioning capital markets pave the way for the separation principle and value additivity. The basic message to the company is to maximize the market value of the production technology, and to use the financial market to manage economic risk.

The existence of a natural gas forward/futures market motivates the use of decision support models from finance. The basic idea is to consider gas storage as a separate asset, and use the market value framework for valuation and utilization of this asset. The company can deal with economic risk by trading in the financial gas market, and cover possible physical imbalances in the spot market.

In the recent literature, several of the methods that are applied from finance are aimed at solving the gas storage problem to optimality. One example is Least Squares Monte Carlo (LSMC) of Longstaff and Schwartz (2001), which is applied to gas storage by for instance Boogert and de Jong (2008).

Longstaff and Schwartz (2001) use the LSMC to evaluate an American stock option to optimality, assuming the usual stock price dynamics. However, the natural gas market consists not only of a spot price but of a whole family of forward prices. Moreover, market data show: (i) clear seasonal price patterns; (ii) considerable price uncertainty; and (iii) more than a few factors are needed to explain the price uncertainty.

In order to solve the gas storage to optimality and still retain computational efficiency, the number of state variables has to be limited, and the problem has to be restrained to accommodate the Markov property. This means that the actual problem has to be simplified substantially. Optimality is at best attained within the simplified framework. So the 1000-dollar question is how much of the “true” gas storage flexibility value is assumed away by reducing the problem to one that can be solved to optimality.

We apply an alternative approach based on a detailed representation of the forward curve and its dynamics, combined with an intuitive and feasible decision rule that follows from repeated maximization of the intrinsic value. At each decision point, the market is represented by an updated forward curve that obeys value additivity, is compatible with the quoted market prices, and reflects historic information (typical time profiles). We apply Principal Component Analysis (PCA) to identify the (six) factors and determine their loads from actual forward curve movements. We use Monte Carlo simulations to model possible realisation of the forward curve dynamics, where the decision that is locked in at each point in time follows from maximizing the intrinsic value of the gas storage.

Our valuation model is tested on market data from the UK National Balancing Point (NBP) covering a two-year period. We compare the results from our approach with the results from

one-factor optimization models as well as with the unfeasible perfect foresight (ex post optimization). Based on our numerical examples, we conclude that our approach, with a rich representation of the price uncertainty combined with a simple and intuitive decision (repeated maximization of the intrinsic value), captures much more of the actual flexibility value than solving gas storage to optimality within a one-factor model framework. Our results indicate that the appropriate framework for analyzing complex assets like gas storage is a multi-factor model combined with a simple decision rule, rather than a model with one (or just a few) factors that is solved to optimality.

## 2. Spot price models

Boogert and de Jong (2008) present a spot price model for valuation of gas storage. They assume that the spot price dynamics is

$$(1) \quad \frac{dS(t)}{S(t)} = \kappa[\mu(t) - \ln S(t)]dt + \sigma dW(t)$$

where the mean reversion rate  $\kappa$  and the instantaneous spot volatility  $\sigma$  are positive constants. The long-term level  $\mu(t)$  is a time-varying function that can be used to fit the spot price process to the forward curve at the initial evaluation date. This model is referred to as log-normal mean-reverting Ornstein-Uhlenbeck price process. Hodges (2004) and Chen and Forsyth (2006) consider a similar price model.

Boogert and de Jong (2008) use (1) above to simulate the spot price paths. These paths, combined with the characteristics of the gas storage facility, are then analyzed by LSMC. Zhuliang and Forsyth (2006) combine this price model with the storage characteristics to set up an optimal control problem. The solution to this problem gives the storage value.

In our opinion, neither of these approaches represents the correct starting point. We will give several reasons.

*Firstly*, the model above is a one-factor spot model. A one-factor model for gas and electricity is unrealistic. Koekebakker and Ollmar (2005) find that up to 10 factors are needed to explain the price movements in the Nord Pool electricity market. Our data indicate that we need six factors to explain 90% of the variation in the NBP gas market in UK.

The reason for choosing a one-factor framework is basically the focus on the solution method. Boogert and de Jong (2008) find stable results for a number of less than 5000 price paths. Increasing the number of factors will dramatically influence this result.

Suppose we extend the model above to a three-factor spot model. Within the LSMC approach, we then need to estimate a decision rule that is made conditional on the value of these three factors for every time step and for every state of the reservoir. Li (2007) suggest the regression

$$(2) \quad E_t \left[ \frac{[V_{t+1}]}{F_t} \right] = \alpha_0 + \alpha_1 \exp(y_t^1) + \alpha_2 \exp(y_t^2) + \alpha_3 \exp(y_t^3) + \alpha_4 \exp(2y_t^1) \\ + \alpha_5 \exp(2y_t^2) + \alpha_6 \exp(2y_t^3) + \alpha_7 \exp(y_t^1 + y_t^2 + y_t^3)$$

where  $y_i^j; i = 1, 2, 3$  are the different factors. The choice might be good but it is still arbitrary. It is an endless number of basis functions of the three factors that are not included above. Another and even more important point is that the best multifactor forward curve models are path-dependent.

The forward curve dynamics we use in this paper requires an even larger state space. Although we use only six factors in the simulation of the forward curve in the next stage, the simulated curve is a function of the previous curve and the random outcome of six factors. Every single factor influences the curve differently for different maturities.

*Secondly*, the initial forward curve used to calculate the trend above plays an unrealistic role. In order to start with a model that is consistent with the current market situation, the trend in the above spot price model is calibrated to the forward market prices that are quoted at the initial date. This creates sensitivity for the initial forward curve that is unrealistic. It will also create problems for a possible hedging strategy.

*And thirdly*, it is unclear how to back test the model over a given period. We may run into problems if we want to compare the value calculated initially with the value we obtain from running the storage following this decision rule. In practice we will never assign that much weight on the initial forward curve. Since market prices do not follow a one-factor model we will observe a new trend function every day. Given that we stick with the old trend function, we take positions on a certain market view. The result will depend on whether this market view is profitable or not. The outcome will be very erratic.

### 3. Valuation model

In the following we describe our method for evaluating the storage. We adopt the standard assumptions from contingent claims analysis of a complete market with no frictions and no riskless arbitrage opportunities, see, e.g., Cox and Ross (1976), Harrison and Kreps (1979) and Harrison and Pliska (1981).

#### 3.1 Forward curve

We assume that the company faces a decision problem with daily granularity. Consequently, the company should use information with the same granularity. The key information in the gas storage problem is the forward curve, hence we need a forward curve with daily granularity.

It follows from economic decision theory that the current forward price for a given day may be interpreted as the certainty equivalent value of that day's future spot price, given the current information. We argue that the forward curve should to be consistent with: (i) the value additivity principle; (ii) updated market information (quoted forward contract prices); and (iii) historic information (typical time profiles).

#### 3.2 Forward curve dynamics

We assume that the risk adjusted dynamics of the forward curve can be represented by a general multifactor model. See, for instance, Heath, Jarrow, and Morton (1992) or Bjerksund, Rasmussen, and Stensland (2000), or Clewlow and Strickland (2000). For an overview of different forward curve methods, see Geman (2005).

The dynamics of the forward curve is represented by the following equation

$$(3) \quad \frac{dF(t,T)}{F(t,T)} = \sum_{i=1}^N \sigma_i(t,T) dZ_i(t)$$

where  $F(t,T)$  represents the forward price at time  $t$  for delivery at time  $T$ . The volatility is represented by  $\sigma_i(t,T)$  where  $i = 1, \dots, N$ . The increments of the  $N$  Brownian motions  $dZ_i(t)$  are assumed independent.

The dynamics just above translates into the future forward price being

$$(4) \quad F(t,T) = F(0,T) \exp \left\{ \sum_{i=1}^N \left( \int_0^t \sigma_i(u,T) dZ_i(u) - \frac{1}{2} \int_0^t \sigma_i(u,T)^2 du \right) \right\}$$

We let the market data decide the number of factors needed. Every day, we construct a forward curve with daily granularity. Next we find a return function for each day as a function of time to delivery. Then we perform principal component analysis (PCA) to find typical curve movements. The volatility functions follow from the loadings in the (PCA). The method is described in Appendix A.

To obtain manageable input estimation of the volatility functions, we assume a time homogeneous model, i.e.,

$$(5) \quad \sigma_i(t,T) = \sigma_i(T-t); i = 1, \dots, N$$

This means that the loadings from the PCA depend only on time to delivery. This might be relaxed if we find ways of estimating PCA-components as a function of calendar time as well as time to delivery.

To perform the simulations we apply the following discrete time representation of (4)

$$(6) \quad F(t + \Delta t, T) = F(t, T) \exp \left\{ \sum_{i=1}^N \left( \sigma_i(T-t) \sqrt{\Delta t} \cdot \tilde{\varepsilon}_i - \frac{1}{2} \sigma_i^2(T-t) \Delta t \right) \right\}$$

Equation (6) is used to simulate the forward curve at time  $t + \Delta t$  conditional on the forward curve at time  $t$ , where we draw  $N$  independent standard normal distributed numbers  $\tilde{\varepsilon}_i$ ,  $i = 1, \dots, N$ .

### 3.3 Intrinsic value

Based on a forward curve with daily granularity we find the optimal deterministic strategy. This is done by solving the following dynamic program

$$(7) \quad V(i, j) = -c(i) + \max_{k=\{\max(-j, -1); 0; \min(J-j, 2)\}} (V(i+1, j+k) e^{-r\Delta t} + p(i) \cdot u \cdot k)$$

where

$V(\cdot, \cdot)$  : value function  
 $i$  : time step  
 $u$  : actual physical size of one unit in the state space  
 $j$  : stock size state, such that  $u \cdot j =$  physical reservoir  
 $k$  : buy/sell strategy,  $k = 2, 0, -1$   
 $p(\cdot)$  : daily price from the forward curve  
 $j_0$  : current reservoir state  
 $J$  : maximum reservoir state  
 $I$  : time horizon  
 $V(I, j)$  : terminal condition  
 $c(\cdot)$  : cost function  
 $V(1, j_0)$  : current value  
 $r$  : interest rate  
 $\Delta t$  : time step size

In the equation above, we need to set the terminal size of the storage. Typically we might choose a long horizon to limit the effect of the terminal condition. One alternative is to choose a horizon where the storage is returned to the initial size. Another alternative is to require the storage to be empty before the filling season.

Some authors (see, e.g., Eydeland and Wolyniec (2003)) refer to the intrinsic value as the optimal initial plan that might be locked in using the tradable products. The implicit assumption seems to be that the granularity of the decision problem corresponds to the product calendar of the market. Now, suppose a problem with a one-year horizon and that only a one-year contract were traded in the market. The intrinsic value of storage would then be zero, according to the above definition. This is clearly meaningless. In most markets, there are traded forward contracts on a monthly delivery (swaps). The definition above would then disregard the possibility of following one strategy during weekdays and another strategy during weekends, which would translate into an unrealistic low reported number for the intrinsic value.

In the following, we define the intrinsic value as the expected value of storage given the risk adjusted dynamics in section (3.2) above, conditional on following the best initial deterministic plan. The reason for our definition of the intrinsic value is as follows. A contract on a monthly delivery can be interpreted as a portfolio of contracts on daily deliveries. However, the fact that there is one quoted marked price for the monthly contract does not necessarily mean that each daily contract commands the same (forward) market price. Although daily contracts are not traded except for in the very short end of the curve, we know from empirical data that there typically will be time-dependent price differences in the future, for instance between weekends and weekdays. Since we start out with a forward curve with a higher granularity (daily) than the product calendar in the market, our intrinsic value will exceed the value of the traditional intrinsic plan.

### 3.4 Repeated intrinsic value method

We start out with the initial forward curve and find the today's intrinsic plan. From this plan we lock in the decision regarding today. If the decision is to fill, the storage is increased and

the cost of filling follows from the current spot price. If the decision is to deplete, the storage reservoir is decreased and the revenue follows from the current spot price. If the decision is to stay put, there is no cash flow and the reservoir is unaltered.

Next we simulate a realisation for tomorrow's forward curve conditional on today's forward curve and the price dynamics given in section (3.2). For this new forward curve, tomorrow's intrinsic plan is determined. This problem is updated for the storage decision that is already locked in as well as the passage of time (one day closer to the terminal date). This plan is used to lock in tomorrow's storage decision. We continue this procedure until we reach the terminal date of our problem.

This procedure gives us one possible realisation of daily cash flows from the storage over the relevant time horizon. We repeat the procedure above to obtain the desired number of possible cash flow realisations. The value of storage is given by the average net present value of the simulated daily cash flows.

### 3.5 Comparing with one-factor optimization models

In this paper we claim that the focus on stochastic optimization methods leads to an oversimplification of the price modelling that severely undervalues the value of storage. In order to compare results, we consider two alternative one-factor methods.

The starting point of the *explicit one-factor model* is the discrete time version of (1), where the spot price dynamics is calibrated to the initial forward curve by (see Appendix B)

$$(8) \quad S(t + \Delta t) = F(0, t + \Delta t) \exp \left\{ e^{-\kappa t} \left[ \ln \left( \frac{S(t)}{F(0, t)} \right) + \frac{1}{2} \frac{\sigma^2}{2\kappa} (1 - e^{-2\kappa t}) \right] - \frac{1}{2} \frac{\sigma^2}{2\kappa} (1 - e^{-2\kappa(t+\Delta t)}) + \sqrt{\frac{\sigma^2}{2\kappa} (1 - e^{-2\kappa\Delta t})} \tilde{\varepsilon} \right\}$$

where  $S(t)$  and  $S(t + \Delta t)$  are the spot prices at time  $t$  and  $t + \Delta t$ ,  $F(0, t)$  and  $F(0, t + \Delta t)$  are the initial forward prices for delivery at time  $t$  and  $t + \Delta t$ , respectively,  $\tilde{\varepsilon}$  is standard normal, and  $\Delta t$  is the time step size. We use (8) to simulate  $M$  possible spot price path realisations. From these paths, a transition matrix is created for each day, and the storage problem is then solved by stochastic dynamic programming.

The starting point of the *implicit one-factor model* is the multi-factor model (6) above, which is used to simulate  $M$  possible spot price path realisations. From these paths, a transition matrix is created for each day. In this way we obtain a one-factor mean-reverting spot price model. The reversion rate is not explicitly estimated but will follow from the PCA-components and the initial forward curve. The storage problem is then solved by stochastic dynamic programming. This might be compared to a model where we use several factors to generate a realistic spot price process but only use the spot price in the optimization problem.

It is claimed that the advantage of the LSMC method compared to optimal control solution techniques is the separation of price paths and the solution method. This is exactly what we are doing here.



## 4. Storage valuation – an example

### 4.1 Data

We use natural gas price data quoted at NBP from October 1<sup>st</sup> 2004 - September 30<sup>th</sup> 2006. The prices are quoted by pence/therm. The traded contracts are spot, day ahead, balance of week, balance of month, as well as calendar months and seasons.

GAS041001.evb (Historic)													
Market: GAS, Area: NBP, Currency: GBP													
GAS	Bid	Ask	Last	Change	Close	Cmp GBP	Cmp EUR	GBP/EUR	User Bid	User Ask	Hours	From	To
W.D.	25.850 p				0.000 p	25.850 p	37.474 c	0.6898			24	1 okt 04	1 okt 04
D.A.	25.900 p				0.000 p	25.900 p	37.538 c	0.6900			24	4 okt 04	4 okt 04
W/END 0	24.700 p				0.000 p	24.700 p	35.803 c	0.6899			48	2 okt 04	3 okt 04
WK/DY 0	25.950 p				0.000 p	25.950 p	37.605 c	0.6901			120	4 okt 04	8 okt 04
Bal Oct-04	27.500 p				0.000 p	27.500 p	39.819 c	0.6906			697	3 okt 04	31 okt 04
Nov-04	39.070 p				0.000 p	39.070 p	56.449 c	0.6921			720	1 nov 04	30 nov 04
Dec-04	52.220 p				0.000 p	52.220 p	75.278 c	0.6937			744	1 des 04	31 des 04
Jan-05	61.080 p				0.000 p	61.080 p	87.847 c	0.6953			744	1 jan 05	31 jan 05
Feb-05	56.050 p				0.000 p	56.050 p	80.436 c	0.6968			672	1 feb 05	28 feb 05
Mar-05	46.000 p				0.000 p	46.000 p	65.869 c	0.6984			743	1 mar 05	31 mar 05
Apr-05	35.750 p				0.000 p	35.750 p	51.077 c	0.6999			720	1 apr 05	30 apr 05
May-05	32.330 p				0.000 p	32.330 p	46.090 c	0.7015			744	1 mai 05	31 mai 05
Jun-05	30.960 p				0.000 p	30.960 p	44.042 c	0.7030			720	1 jun 05	30 jun 05
Jul-05	30.850 p				0.000 p	30.850 p	43.794 c	0.7044			744	1 jul 05	31 jul 05
Aug-05	32.080 p				0.000 p	32.080 p	45.445 c	0.7059			744	1 aug 05	31 aug 05
FWSO-05	32.490 p				0.000 p	32.490 p	46.170 c	0.7037			4392	1 apr 05	30 sep 05
FWV-05	47.010 p				0.000 p	47.010 p	66.040 c	0.7118			4368	1 okt 05	31 mar 06
FWSO-06	30.880 p				0.000 p	30.880 p	42.971 c	0.7186			4392	1 apr 06	30 sep 06
FWV-06	41.170 p				0.000 p	41.170 p	56.818 c	0.7246			4368	1 okt 06	31 mar 07
FWSO-07	27.260 p				0.000 p	27.260 p	37.336 c	0.7301			4392	1 apr 07	30 sep 07
FWV-07	35.500 p				0.000 p	35.500 p	48.289 c	0.7352			4392	1 okt 07	31 mar 08

Figure 1 Snapshot of NBP prices on Oct 1st 2004

We use the Elviz Front Manager software to translate this price information into a forward curve with the following properties: (i) the curve obeys value additivity; (ii) the curve reproduces the market prices of the quoted contracts; and (iii) the curve reflects typical seasonalities over the week and over the year. A detailed outline, which includes a solution method, is given in Benth, Koekebakker and Ollmar (2007).

The forward curve at NBP at October 1<sup>st</sup> 2004 is illustrated in Figure 2 below. Observe that there are price variations within the week (lower prices during weekends) and within the year (higher prices during the winter).


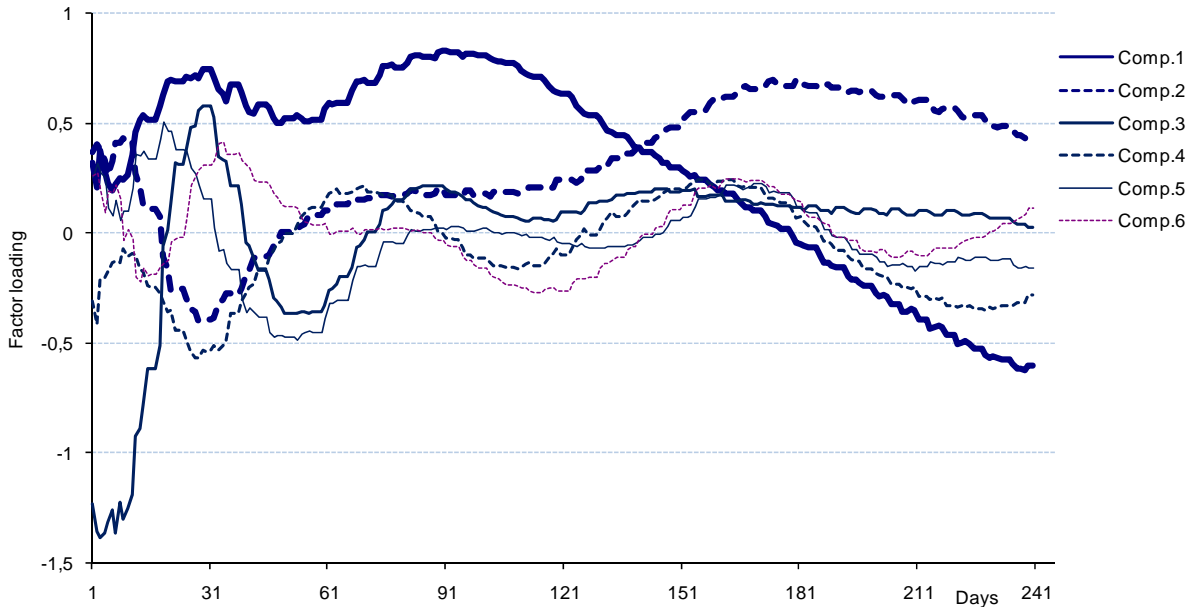



Figure 2 Forward curve on Oct 1st 2004

Based on the forward curve movements the first year, we apply PCA to estimate the loadings of the first six components. The results are given in Figure 3 below. Note that these six factors give rise to a rich class of possible forward curve movements.




**Figure 3 Loadings of the six most important factors**

The explanation power of the components is shown in Figure 4 below. Observe that the first factor explains about 35% of the variation, whereas six factors explain about 90% of the price variation. We conclude that six factors should be sufficient in our analysis.



**Figure 4 Cumulative percentage variance contribution of principal components**

These six factors add up to a volatility curve for the instantaneous movement in the forward curve, see Figure 5. A similar volatility curve, where all factors are included, is provided in Appendix A. In most commodity markets, the volatility curve is a decreasing function of time to delivery. Our data suggests a falling curve with some artifacts. In particular, there are some unexplained peaks such as the volatility in the forward price with delivery in 30 days.



**Figure 5 Volatility term structure**

## 4.2 Storage characteristics

The characteristics of the stylized storage are given in Table 1 below.

**Table 1** Storage characteristics

Initial storage	125 million therms
Terminal storage	125 million therms
Max storage	250 million therms
Injection	2.5 million therms per day
Depletion	2.5 or 5 million therms per day

All costs are set to zero, and we disregard the positive interest rate. We consider a problem with a one-year horizon, and our objective is to obtain the value of the cash flow from operating the storage.

## 4.3 In-sample analysis

We start with an in-sample valuation example, where we consider the value of operating the storage facility the first year (October 1<sup>st</sup> 2004 to September 30<sup>th</sup> 2005), given the initial forward curve (Figure 2) and using the PCA factor load estimated from the same period (Figure 3).

The traditional intrinsic value that can be locked in using the traded monthly products is 58 million £. The intrinsic value using daily resolution of the forward curve as well as the storage strategy is 76 million £.

The explicit one-factor model in Equation (8) (5 000 simulations used to generate transition matrices), optimizing the storage strategy, gives a value of 104 million £. Here we have used the short term volatility of 149% per annum and a mean reverting rate of 0.05 per day, which

translates into a rate of  $\kappa = 0.05 \cdot 365$  per year. The short term volatility is a proxy for the short term volatility in the PCA-components, whereas the mean reverting rate is equal to the rate used by Boogert and de Jong (2008). The implicit one-factor model (10 000 simulations used to generate transition matrices), optimizing the storage strategy, gives a value of 96 million £.

Given the initial forward curve and the estimated factor loadings, the six-factor model with repeated intrinsic value maximization gives a storage value of 187 million £ (1 000 simulations). The standard deviation of this estimate is 1 million £. This means that the value of flexibility is considerably higher than reported by the other models. In order to obtain rather low standard deviation in the estimate using only 1000 simulation we use the value of the financial hedging strategy (described in section 5) as a control variate.

As a comparison, we find that the optimal but unfeasible perfect foresight method gives a storage value 205 million £ (1000 simulations). The standard deviation of this estimate is 2 million £. This result is obtained by simulating all the forward prices paths to the horizon, finding the corresponding spot price path, and the ex post best strategy for each path.

Our results are summarized in Table 2.

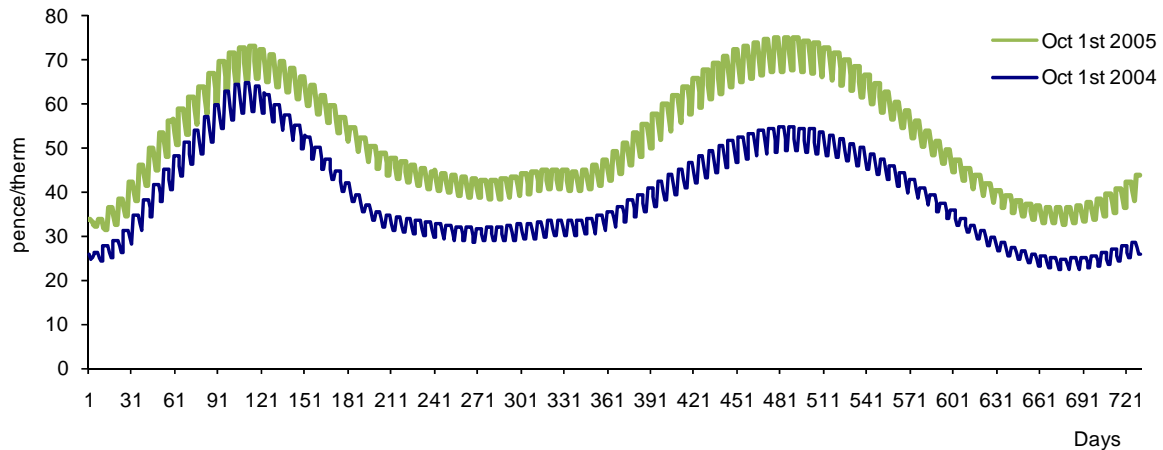
**Table 2** Computed storage value 2004 (million £)

Intrinsic - monthly granularity	58
Intrinsic - daily granularity	76
One-factor model - explicit	104
One-factor model - implicit	96
Dynamic intrinsic model	187
Perfect foresight	205

On this background we claim that the simple repeated intrinsic method is close to the optimal value. We know that the perfect foresight value of 205 is impossible to obtain.

#### 4.4 Out-of-sample analysis

Next we consider an out-of-sample analysis. In this example, we consider the value of operating the storage facility the second year (October 1<sup>st</sup> 2005 to September 30<sup>th</sup> 2006), using the initial forward curve at October 1<sup>st</sup> 2005 and the PCA factor load estimates from the first year (October 1<sup>st</sup> 2004 to September 30<sup>th</sup> 2005).



**Figure 6 Forward curve comparison**

In this case, the traditional intrinsic value that can be locked in using the traded monthly products is 61 million £, whereas the intrinsic value of the storage with daily resolution is 74 million £. These numbers are quite close to the corresponding values of the first year reported in Table 2 above. To explain this, observe that the storage may be interpreted as a complicated calendar spread, and that forward curves at October 1<sup>st</sup> 2004 and October 1<sup>st</sup> 2005 basically have the same shape (relative prices), c.f., Figure 6.

The explicit one-factor model, where we use equation (8) (5000 simulations) to generate transition matrices, and assume the same short-term volatility and mean reverting rate as above, gives a value of 116 million £. The implicit one-factor model gives a storage value of 111 million £.

The repeated intrinsic method gives a value of 228 million £ (1000 simulations). The standard deviation of this estimate is 1.5 million £. As a comparison, we find that the optimal but infeasible perfect foresight model gives 258 million £ (1000 simulations) with a standard deviation of 3 million £.

Our findings are summarized in Table 3.

**Table 3** Calculated storage value 2005 (million £)

Intrinsic - monthly granularity	61
Intrinsic - daily granularity	74
One-factor model - explicit	116
One-factor model - implicit	111
Dynamic intrinsic model	228
Perfect foresight	258

#### 4.5 Conclusions

The two examples show that our six-factor model, combined with repeated intrinsic value maximization, creates a significantly higher value than the one-factor stochastic optimization models. Hence, our results indicate that our approach should be the preferred one.

Moreover, we find that there is a marginal additional value from perfect foresight. Perfect foresight is of course unfeasible, and represents an upper bound to the true storage value. This

indicates that our valuation approach captures most of the storage flexibility value, and suggests that there is a limited potential for improvement of our approach.


### 5. Hedging

A crucial point in option pricing theory is the concept of a replicating portfolio. The rationale behind the Black-Scholes option pricing formula is the existence of a dynamic self-financing trading strategy that gives exactly the same pay-off at the horizon. To rule out arbitrage, the option price must coincide with the cost of creating the replicating portfolio. If we know how to replicate an option, we also know how to hedge it, because the two strategies are opposite ones. However, the beauty of the theory often breaks down in practice. The most important explanation is a mismatch of volatilities between the model and the market.

In the case of risk management of gas storage, it may be instructive to think in terms of a replicating portfolio. It will in general consist of positions in the all of the forward contracts. Since we are not even close to a formula for the gas storage value, the load on every contract has to be simulated.


Consequently, we use an alternative approach to the hedging problem. In particular, we will exploit our dynamic intrinsic plan to define a financial forward trading strategy. The cash flow from this strategy will be highly negatively correlated with the cash flow generated from the gas storage. However, the strategy will not capture the pure flexibility value that accrues the owner of the physical facility. Hence we might call this a sub-replicating hedge portfolio.

Suppose that we can trade daily forward contracts at each point in time. As an illustration, consider October 1<sup>st</sup> 2004. The initial hedge portfolio would then have the following exposure:



**Figure 7 Exposure October 1<sup>st</sup> from trades in the forward market**

Suppose that only monthly contracts are traded. By mapping the daily exposures from the intrinsic storage plan to net monthly exposures, we would have the following exposure at October 1<sup>st</sup> 2004:



**Figure 8 Monthly exposure and initial trades**

We consider a dynamic situation where we utilize the storage according to the storage strategy that follows from repeated intrinsic value maximization. In addition, we follow a dynamic financial hedge strategy. Each day, we liquidate yesterday's hedge portfolio and take new positions that are opposite of the updated intrinsic value maximization storage plan.

Recall that the hedge portfolio consists of forward contract positions. This means that the market value of entering each position (as well as the portfolio) is zero. Next day, new forward prices are quoted. The old hedge portfolio is then liquidated, yielding a positive or negative cash flow. Thereafter, the new hedge portfolio is composed using information from the updated intrinsic value storage plan. And so forth. It follows from above that the market value of launching the dynamic hedge strategy is zero.

Observe that the hedge strategy can be implemented by any market participant. The gas storage ownership per se is no prerequisite for trading in the financial gas market. This asset can of course serve as collateral for the company's financial gas trading and speculation, but that is a different matter.

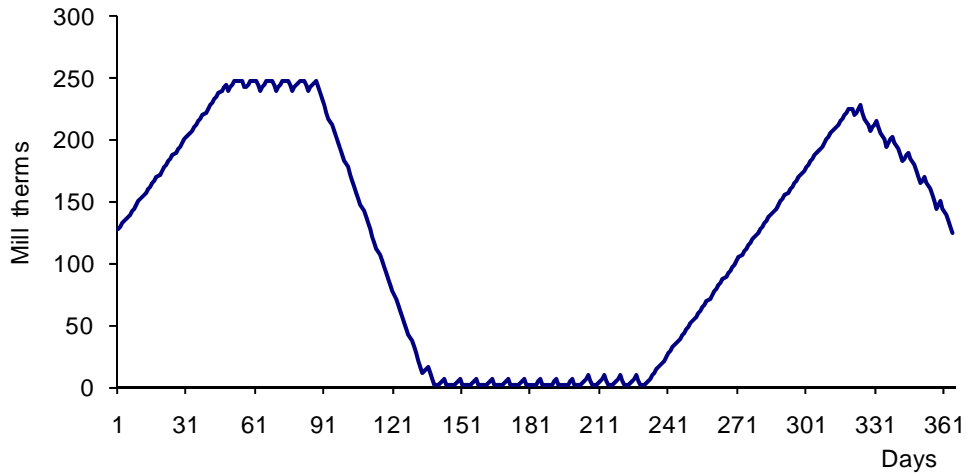
So why bother with such a strategy? The explanation is that the cash flows from the gas storage and the hedge portfolio are highly negatively correlated. One obvious practical application of our hedge strategy is to reduce the risk from owning and utilizing a gas storage facility. Another application is numerical analysis of gas storage. We can use the financial hedge strategy as a control variate to improve the accuracy of our estimates.

## 6. Back testing

In this section, we investigate the performance of our repeated intrinsic value maximization strategy on the spot prices that actually were realized in the market. We assume the same storage characteristics as above and that terminal storage equals the initial storage.


### 6.1 Year 1 - Storage

We start at October 1<sup>st</sup> 2004 and consider the following year. Maximizing the intrinsic value that day gives the strategy that is illustrated in Figure 9.



**Figure 9 Storage from intrinsic plan**

The expected cumulative cash flow from sticking to this strategy is illustrated in the figure below, where we use the property that each day's forward price can be interpreted as the expected future spot price that day (using the equivalent martingale measure).

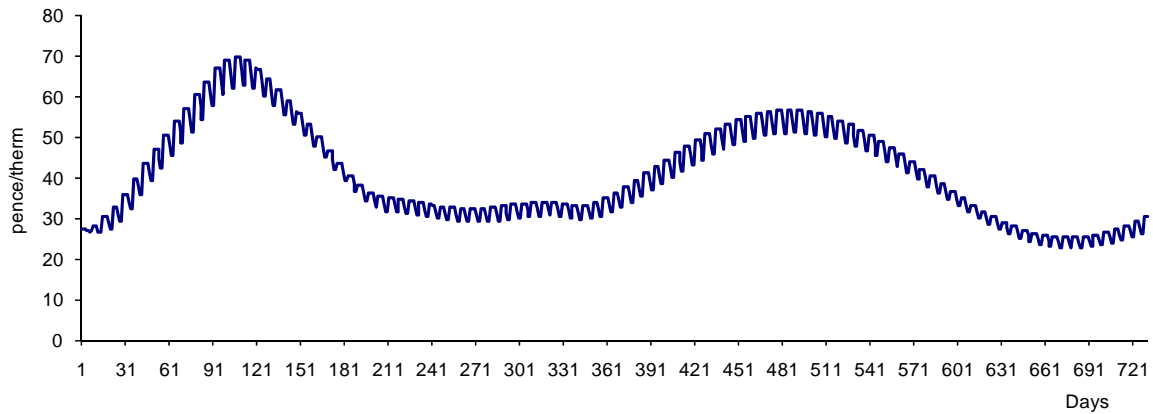


**Figure 10 Accumulated cash flow from intrinsic plan**

The cumulative expected cash flow starts out negative, which reflects that the strategy is to start injecting. The cumulative expected cash flow at the horizon is the initial intrinsic value of 76 million £, c.f. Table 2 above.

We perform an intrinsic value model every day given updated information. In this example, we decide to fill the storage with maximum capacity equal 2.5 million therms this first day. The next day we have the following forward curve






**Figure 11 Forward curve October 2nd 2004**

We now optimize the intrinsic value for the remaining period given updated information. The current reservoir is 127.5 million therms. There are 364 days to the horizon, and the terminal storage must equal 125 million therms.

Note that the strategy is almost unchanged. The best decision day two given that we follow the intrinsic plan is also to inject. We continue in this way for every day. For the weekend we use the curve given on last Friday to take the decision whether to inject or deplete.


We continue in this way until September 30<sup>th</sup> 2005. At this time the reservoir equals 125 million therms. The result for the repeated intrinsic method from October 1<sup>st</sup> 2004 until September 30<sup>th</sup> 2005 is given in the following figures.

The evolution of the storage is given in Figure 12.




**Figure 12 Actual storage size with the dynamic intrinsic method Oct 1st 2004 to Sept 30th 2005**

The spot prices that were realized the first year is shown in Figure 13.



**Figure 13 NBP prices Oct 1st 2004 – Sept 30th 2005**

The development of the accumulated cash flow from following the dynamic intrinsic method is shown in Figure 14.




**Figure 14 Realized accumulated cash flow from the dynamic intrinsic method Oct 1st 2004 – Sept 30th 2005**

The result from following this storage strategy is 48 million £.

## 6.2 Year 1 - Hedged storage

Above, we found that the realised cash flow from storage by implementing the repeated intrinsic value maximization from 1<sup>st</sup> October 2004 to 30<sup>th</sup> September 2005 is 48 million £.

Now, suppose that the owner of the gas storage had launched a dynamic financial hedge strategy as described in the previous section. Assuming a market with daily forward contracts, the realised cash flow generated by the dynamic financial hedge portfolio was 168 million £. This large number is explained by reduced seasonal spread. This adds up to a total cash flow from the hedged storage 216 million £. The development of the cumulated cash is illustrated in Figure 15 below.



**Figure 15 Accumulated physical and financial value**


It may be argued that daily forwards are unrealistic. In order to model a more realistic trading strategy, we assume that all exposure from the optimal intrinsic plan is mapped into net positions in the traded products. For simplicity we assume that the next 10 days are tradable, and thereafter only the relevant trading list from next month and to the horizon. The cash flow generated by the dynamic financial hedge is then reduced to 89 million £. This translates into a total cash flow from the hedged storage of 137 million £. The results of our back test is summarized in Table 4.

**Table 4** Back-testing result 2004 (million £)

	Physical storage	Hedge portfolio	Hedged storage
Daily forwards	48	185	233
Trading calendar	48	89	137

### 6.3 Year 2 - Storage

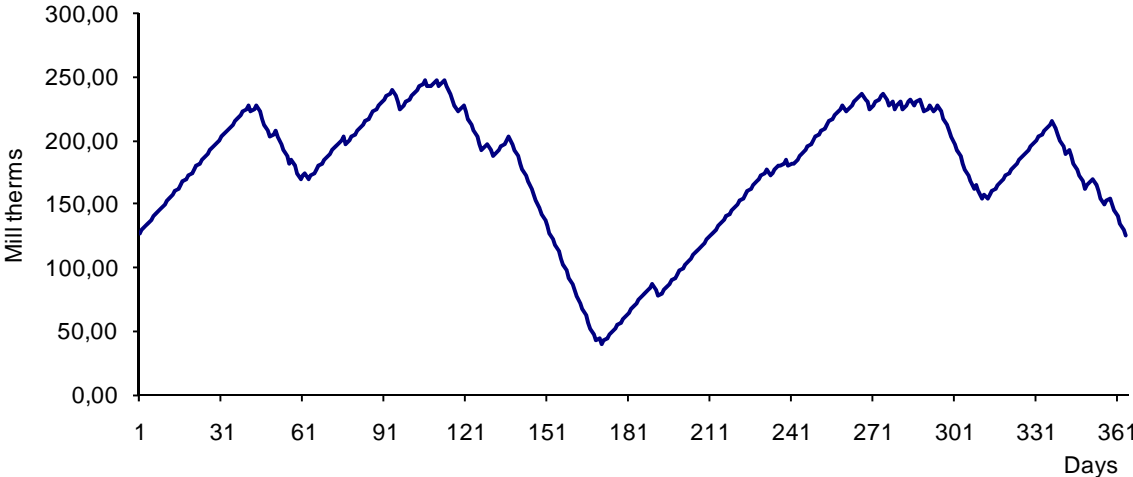
Now we perform the same analysis for the year starting at October 1<sup>st</sup> 2005 with a one-year horizon. The reservoir is the same as in the previous example 125 million therms. The first day we obtain the following intrinsic value results.



**Figure 16 Planned accumulated cash flow from Oct 1st 2005 to the horizon**


We observe that the realized cash flow for the optimal use of the reservoir next year is £74 million. This is close the initial guess the year before. It is the seasonal spread in the forward curve that creates this value. It is therefore fair to say that this pattern is unchanged.

Next we perform a new intrinsic value model every day until September 30<sup>th</sup> 2006. The reservoir is left at the same size as we started with, that is 125 million terms. The result from this strategy is presented below. The evolution of the storage is given in Figure 17.




**Figure 17 Storage size from dynamic intrinsic Oct 1st 2005 to September 30th 2006**

The spot prices that were realized the second year is shown in Figure 18.



**Figure 18 NBP prices from Oct 1st 2005 – Sept 30th 2006**

The development of the accumulated cash flow from following the dynamic intrinsic method is shown in Figure 19.



**Figure 19 Realized accumulated value for the dynamic intrinsic strategy Oct 1st 2005 to September 30th 2006**

The realized cumulated cash flow from storage is 117 million £.

#### 6.4 Year 2 - Hedged storage

Assume a market with daily forward contracts. The dynamic hedging strategy in this case gives a trading profit of £325 mill. So the value of the hedged storage amounts to £ 440 million.

If we only include the 10 first days of forward prices + the trading calendar which starts with the next full month, the result drops to 190 million £. The results of our back test is summarized in Table 5.

**Table 5** Back-testing result 2005 (million £)

	Physical storage	Hedge portfolio	Hedged storage
Daily forwards	115	325	440
Trading calendar	115	75	190

## 7. Conclusion

Storage is traditionally evaluated by models using simple price processes and complicated solution methods. We suggest an alternative approach, with a rich representation of prices (forward curve with daily resolution) and uncertainty (six factors), combined with a simple intuitive decision rule (repeated maximization of intrinsic value).

Based on market price data from the UK gas market, we compare the results from our model with one-factor optimization models and the unfeasible perfect foresight model. We find that our model captures much more of the true flexibility value than the one-factor models.

Our results indicate that the appropriate framework for analyzing complex assets like gas storage is a multi-factor model combined with a simple intuitive decision rule, rather than a model with one (or just a few) factors that is solved to optimality. This suggests that the main

focus for future research is the modelling of realistic prices rather than efficient solution procedure.

## References

Adams, Kenneth J. and Donald R. van Deventer (1994): Fitting Yield Curves and Forward Rates with Maximum Smoothness. *Journal of Fixed Income*, June 1994; pp. 52-62.

Benth, Fred Espen, Steen Koekebakker, and Fridthjof Ollmar: Extracting and Applying Smooth Forward Curves From Average-Based Commodity Contracts with Seasonal Variation. *Journal of Derivatives*, Fall 2007, pp. 52-66.

Bjerksund, Petter, Heine Rasmussen, and Gunnar Stensland (2000): Valuation and Risk Management in the Norwegian Electricity Market. *Working Paper 20/2000, Norwegian School of Economics and Business Administration*.

Bjerksund, Petter and Gunnar Stensland (1996): Utledning av rentens terminstruktur ved maksimum glatthets prinsippet. *Beta* 1996; 1; 2-6. (In Norwegian).

Boogert, Alexander and Cyriel de Jong (2008): Gas Storage Valuation Using a Monte Carlo Method. *Journal of Derivatives*, Spring 2008, pp. 81-98.

Chen, Zhuliang, and Peter Forsyth (2006): Stochastic Models of Natural Gas Prices and Applications to Natural Gas Storage Valuation, Working Paper, David R. Cheriton School of Computer Science, University of Waterloo.

Clewlow, Les and Chris Strickland (2000): *Energy Derivatives Pricing and Risk Management*. Lacima Publications, pp. 142-144.

Cox J. Ross S. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 1976; **3**:145-166.

Eydeland, Alexander and Krzysztof Wolyniec (2003): *Energy and Power Risk Management*. Wiley: Hoboken, NJ.

Geman, Hélyette (2005): *Commodities and Commodity Derivatives*. Wiley.

Harrison JM Kreps D. Martingales and arbitrage in multiperiod security markets. *Journal of Economic Theory* 1979; **20**:381-408.

Harrison JM Pliska S. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and Their Applications* 1981; **11**:313-316.

Heath, D., R.A. Jarrow, and A.J. Morton (1992): "Bond Pricing and The Term structure of Interest Rates", *Econometrica* 60, pp. 77-105.

Hodges, Stewart D. (2004): The value of a Storage Facility. Working Paper, Warwick Business School, University of Warwick.

Li, Yun (2007): Natural Gas Storage Valuation, Master Thesis, Georgia Institutt of Technology.

Koekebakker, Steen and Fridtjof Ollmar (2005): Forward curve dynamics in the Nordic electricity market, *Managerial Finance* 31 (6), 74-95.

Longstaff, Francis A. and Eduardo S. Schwartz (2001): Valuing American Options by Simulation: A Simple Least-Squares Approach. *The Review of Financial Studies*, Vol. 14, No. 1 (Spring, 2001), pp. 113-147.

## **Appendix A: Construction of Forward Curves and the Principle Components Analysis**

In order to understand the way the forward curve changes from day to day, especially how the volatility and correlation of the different maturities behave, some form of simplifications are needed.

It is known from for example stock markets that if you can fit the data to a model (lognormal, bivariate) the stocks future development can be simulated using Monte Carlo techniques, to replicate the historic behavior.

For commodity markets this is more challenging because the traded products are linked to a specific time period. We know from empiric studies that volatility generally changes with time to maturity. Therefore we construct a forward curve with daily granularity. Next we study this term structure of forward prices.

### **A1. PCA Methodology**

Principal components analysis (PCA) is a widely used method for simplifying complex data structures. Instead of attempting to describe everything to a 100% the idea is to filter out the most important factors (principal component) and use them to simulate the market.

It is widely recognized that being able to describe 90-95 % of the variances is acceptable.

Ollmar and Koekebakker (2000) showed that 10 factors were needed to explain 95% of the variations in the electricity market on NordPool.

This study is based on taking out prices from the market forward curve with the same spacing between maturities. Ollmar and Koekebakker (2000) used weekly prices retrieved from the Smoothed Elviz Forward Curves.

Ollmar and Koekebakker used the Elviz curve to first generate a smooth forward curve, keeping all the market quotes inherent in the curve, and pick out weeks in a consistent way. We therefore proposed a solution which takes this idea one step further, as we choose daily resolution.

### **A2. Input / variations**

Transform closing prices to forward curves:

We can not go into the detail of Adam and van Deventer (1994), Bjerksund and Stensland (1996) model of a smooth, arbitrage free forward curve here, but we will quickly outline the principles. A detailed outline, which includes a solution method, is given in Benth, Koekebakker and Ollmar (2007).

First the prices from the forward market (periods can have different length and may also be overlapping) are fed into the smoother. The smoother is a spline function optimiser where the goal function is to maximise smoothness (the same as minimizing the square of the twice differentiated function).


When the curve is smoothed, an hourly profile (historic hourly) is modulated on the curve. The gas market is not hourly as electricity, so we used a daily profile calculating the way prices normally changes within a week. The resulting curve can look something like this:





**Figure A1** The forward curve

We repeat a procedure of inputting daily prices 241 times (the period from Oct 1 2004 to Sep 31 2005) and generated 241 forward curves with daily resolution. These curves are shown in a 3d plot below with 90 days length.



**Figure A2** The first 90 days of the forward curves

The next step is to calculate the daily return using the following formula:


$$\ln[\text{curve2}(t1)/\text{curve1}(t2)],$$

meaning that we divide the first day on a curve with the second day on the previous curve and so on. This is of course also adjusted for the weekends, such that on a Monday the return is  $\ln(\text{Fridaycurve}(4)/\text{Mondaycurve}(1))$ .

The fourth day on the Friday curve is the same day as the first day on the following Monday curve. It can be argued that this return should be weighted differently since it represents the return over three days. But we have found no significant difference in the volatilities caused by this effect, so we have treated them equally.

**A3. Descriptive statistics**

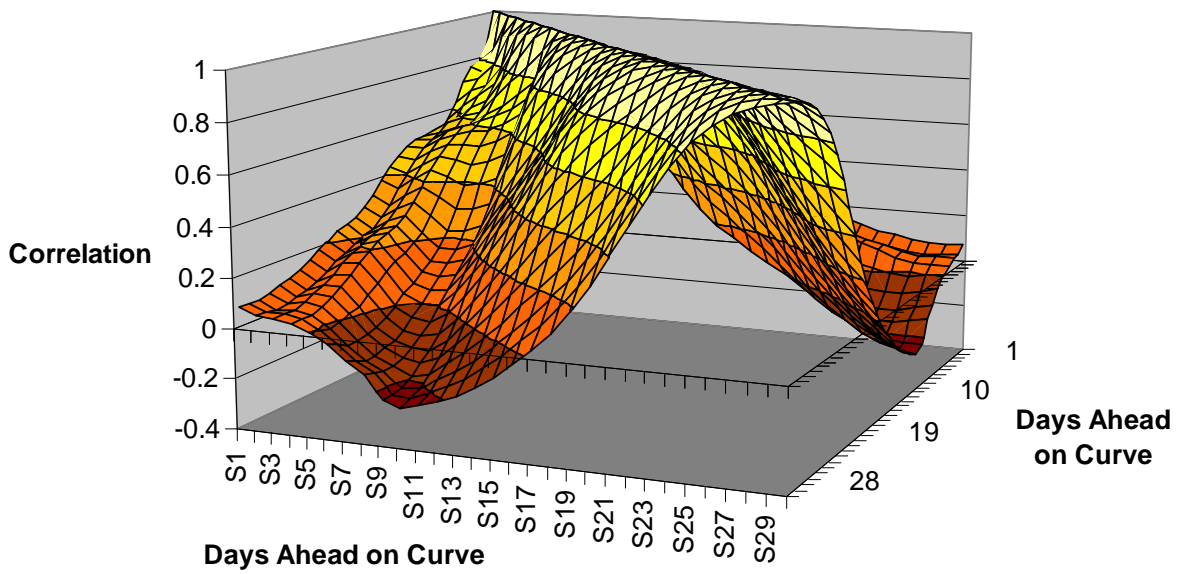
The volatilities of these log returns are calculated for all maturities and presented below. All days refer to the calculation incorporating all the curves. The numbers are annualized using the square root of 250 (number of trading days).



**Figure A3** Annualized volatility term structure

As expected we see a quite sharp fall in volatility from the very short end (1 and 2 days representing the within day WD and day ahead DA products of over 200%, down to 70-50% in the first couple of front months. We also observe a rise 10 days into the curve which have no real explanation. This is attributed to an artifact added by the smoothing algorithm. The falling volatility confirms findings in other markets. Typically the electricity market inherent the same sharp decline in volatility.

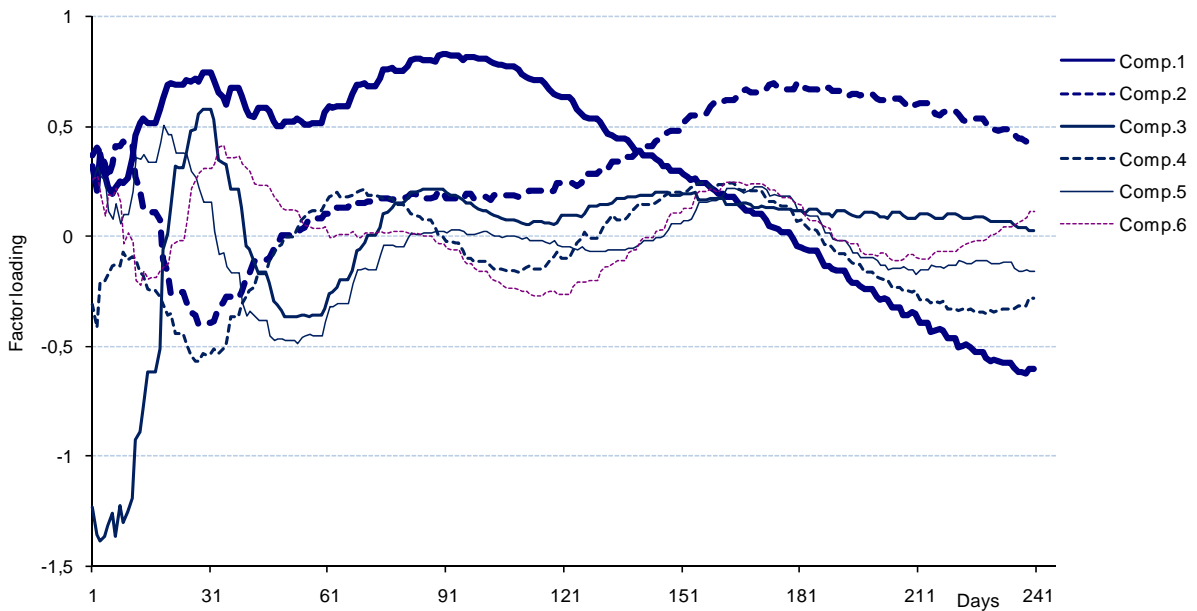
The correlation between different maturities is an important part of the PCA, and is calculated in S-PLUS after import of the forward curves.




**Figure A4** Correlation matrix

The way of calculating the principal components is to find the eigenvalue of the eigenvectors of the correlation matrix, and simply sort them in a falling order. The eigenvalue vector that explains the most of the variations in the forward curves will be the first eigenvalue vector. We call this the principal component.

We can calculate the added contribution by each component, and see how many we need to get a reasonably good description of the variations in the market.



**Figure A5** Volatility attribution of the six first components



**Figure A6.** Relative importance of the components, cumulative attribution

We find that three components explain 75% of the variations, and six components are needed to explain 90%. The first component is recognized as a parallel shift component. The second and third components are more difficult to give intuitive explanations. The second seems to be tilting the curves, but it is far from a linear tilt.

The parallel shifting representing almost half the variations in the curve is the factor that contributes to no value to the gas storage.

## Appendix B: Derivation of the dynamics of the explicit one-factor model

In the following, we want to derive the dynamics of the explicit one-factor model that is stated in Equation (8) above.

### B1. Arithmetic state variable

We start out with the following generalized Ornstein-Uhlenbeck state variable

$$(B1) \quad X(T) = e^{-kT} \left\{ X(0) + \frac{\sigma^2}{2k} \right\} - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=0}^T e^{ks} \mu(s) ds + \sigma e^{-kT} \int_{s=0}^T e^{ks} dZ_s$$

where  $k$  is the mean-reversion parameter,  $\sigma$  is the instantaneous volatility, the function  $\mu(s)$  is the mean-reversion level, and  $dZ_s$  is the increment of a standard Brownian motion.

As seen from date 0, the expectation is

$$(B2) \quad E_0[X(T)] = e^{-kT} \left\{ X(0) + \frac{\sigma^2}{2k} \right\} - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=0}^T e^{ks} \mu(s) ds$$

and the variance is

$$(B3) \quad \begin{aligned} \text{var}_0[X(T)] &= \text{var}_0 \left[ \sigma e^{-kT} \int_{s=0}^T e^{ks} dZ_s \right] \\ &= \sigma^2 e^{-2kT} \int_{s=0}^T e^{2ks} ds \\ &= \sigma^2 e^{-2kT} \left[ \frac{1}{2k} e^{2ks} \right]_{s=0}^T \\ &= \frac{\sigma^2}{2k} (1 - e^{-2kT}) \end{aligned}$$

It follows immediately from (B1) that

$$(B4) \quad \begin{aligned} X(T) &= e^{-k(T-t)} \left\{ e^{-kt} \left( X(0) + \frac{\sigma^2}{2k} \right) + ke^{-kt} \int_{s=0}^t e^{ks} \mu(s) ds + \sigma e^{-kt} \int_{s=0}^t e^{ks} dZ_s \right\} \\ &\quad - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds + \sigma e^{-kT} \int_{s=t}^T e^{ks} dZ_s \\ &= e^{-k(T-t)} \left\{ X(t) + \frac{\sigma^2}{2k} \right\} - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds + \sigma e^{-kT} \int_{s=t}^T e^{ks} dZ_s \end{aligned}$$

Hence, the conditional expectation is

$$(B5) \quad E_t[X(T)] = e^{-k(T-t)} \left\{ X(t) + \frac{\sigma^2}{2k} \right\} - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds$$

and the conditional variance is

$$(B6) \quad \begin{aligned} \text{var}_t[X(T)] &= \text{var}_t \left[ \sigma e^{-kT} \int_{s=t}^T e^{ks} dZ_s \right] \\ &= \sigma^2 e^{-2kT} \int_{s=t}^T e^{2ks} ds \\ &= \sigma^2 e^{-2kT} \left[ \frac{1}{2k} e^{2ks} \right]_{s=t}^T \\ &= \frac{\sigma^2}{2k} (1 - e^{-2k(T-t)}) \end{aligned}$$

The dynamics of this state variable is

$$(B7) \quad \begin{aligned} dX(T) &= -k \left\{ e^{kT} \left( X(0) + \frac{\sigma^2}{2k} \right) + ke^{-kT} \int_{s=0}^T e^{ks} \mu(s) ds + \sigma e^{-kT} \int_{s=0}^T e^{ks} dZ_s \right\} \\ &\quad + k\mu(T)dT + \sigma dZ_T \\ &= -k \left( X(T) + \frac{\sigma^2}{2k} \right) dT + k\mu(T)dT + \sigma dZ_T \\ &= k \left( \left( \mu(T) - \frac{\sigma^2}{2k} \right) - X(T) \right) dT + \sigma dZ_T \end{aligned}$$

## B2. Spot price representation

Now, define the spot price process by

$$(B8) \quad X(T) = \ln S(T) \quad ; \quad S(T) = \exp\{X(T)\}$$

Differensiate

$$dS(T) = \exp\{X(T)\}dX(T) + \frac{1}{2} \exp\{X(T)\}(dX(T))^2$$

and obtain

$$(B9) \quad \begin{aligned} \frac{dS(T)}{S(T)} &= k \left( \left( \mu(T) - \frac{\sigma^2}{2k} \right) - X(T) \right) dT + \sigma dZ_T + \frac{1}{2} \sigma^2 dT \\ &= k(\mu(T) - \ln S(T))dT + \sigma dZ_T \end{aligned}$$

This corresponds to the spot price dynamics used by Cyriel and de Jong (2008).

### B3. Forward prices

The dynamics above are with respect to the equivalent martingale measure. Hence, the current forward price corresponds to the expected future spot price. Moreover, with our spot price representation, the future spot price is lognormal. Consequently, we have that

$$\begin{aligned}
 (B10) \quad F(0,T) &= E_0[S(T)] \\
 &= E_0[\exp\{X(T)\}] \\
 &= \exp\{E_0[X(T)] + \frac{1}{2} \text{var}_0[X(T)]\}
 \end{aligned}$$

use Eqs. (B2) and (B3) to obtain

$$(B11) \quad F(0,T) = \exp\left\{e^{-kT} \left( X(0) + \frac{\sigma^2}{2k} \right) - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=0}^T e^{ks} \mu(s) ds + \frac{1}{2} \frac{\sigma^2}{2k} (1 - e^{-2kT}) \right\}$$

By definition, we have that  $F(0,0) = S(0)$ . Note that once the mean-reversion parameter  $k$  and the instantaneous volatility parameter  $\sigma$  are determined, the drift function  $\mu(s)$  follows from the initial forward curve  $F(0, s)$ .

### B4. Simulating the spot price dynamics

Suppose that we have calibrated the model to market information at date 0, and that we want to build a simulation model given this information.

First, consider the simulation of  $S(t)$  given the information at time 0, and hence corresponds to the first time step. Using log-normality and Eq. (B10), we have

$$\begin{aligned}
 (B12) \quad S(t) &= \exp\{X(t)\} \\
 &= \exp\left\{E_0[X(t)] + \sqrt{\text{var}_0[X(t)]} \tilde{y}\right\} \\
 &= \exp\left\{\ln F(0,t) - \frac{1}{2} \text{var}_0[X(t)] + \sqrt{\text{var}_0[X(t)]} \tilde{y}\right\}
 \end{aligned}$$

where  $\tilde{y}$  is standard normal,  $F(0,t)$  is defined from the initial forward curve, and the variance term is given by (3) above.

Next, we want to consider the simulation of  $S(T)$  conditional on the value that is realised for  $S(t)$ . First, reconsider (B5) above, and take the expectation as of date 0, and use the forward price definitions, to obtain the following identity

$$\begin{aligned}
E_0[E_t[X(T)]] &= e^{-k(T-t)} \left\{ E_0[X(t)] + \frac{\sigma^2}{2k} \right\} - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds \\
ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds &= E_0[X(T)] + \frac{\sigma^2}{2k} - e^{-k(T-t)} \left\{ E_0[X(t)] + \frac{\sigma^2}{2k} \right\} \\
(B13) \quad ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds &= \ln F(0, T) - \frac{1}{2} \text{var}_0[X(T)] + \frac{\sigma^2}{2k} \\
&\quad - e^{-k(T-t)} \left\{ \ln F(0, t) - \frac{1}{2} \text{var}_0[X(t)] + \frac{\sigma^2}{2k} \right\}
\end{aligned}$$

Next, consider the simulation of  $S(T)$  conditional on the simulated realisation of  $S(t)$ .

$$S(T) = \exp\left\{E_t[X(T)] + \sqrt{\text{var}_t[X(T)]} \tilde{z}\right\}$$

where  $\tilde{z}$  is standard normal and independent of  $\tilde{y}$ . Insert Equations (B5) and (B13)

$$\begin{aligned}
(B14) \quad S(T) &= \exp\left\{e^{-k(T-t)} \left( X(t) + \frac{\sigma^2}{2k} \right) - \frac{\sigma^2}{2k} + ke^{-kT} \int_{s=t}^T e^{ks} \mu(s) ds + \sqrt{\text{var}_t[X(T)]} \tilde{z}\right\} \\
&= \exp\left\{e^{-k(T-t)} \left( X(t) + \frac{\sigma^2}{2k} \right) - \frac{\sigma^2}{2k} + \ln F(0, T) - \frac{1}{2} \text{var}_0[X(T)] + \frac{\sigma^2}{2k} \right. \\
&\quad \left. - e^{-k(T-t)} \left( \ln F(0, t) - \frac{1}{2} \text{var}_0[X(t)] + \frac{\sigma^2}{2k} \right) + \sqrt{\text{var}_t[X(T)]} \tilde{z}\right\} \\
&= F(0, T) \exp\left\{e^{-k(T-t)} \left( \ln\left(\frac{S(t)}{F(0, t)}\right) + \frac{1}{2} \text{var}_0[X(t)] \right) \right. \\
&\quad \left. - \frac{1}{2} \text{var}_0[X(T)] + \sqrt{\frac{\sigma^2}{2k} (1 - e^{-2k(T-t)})} \tilde{z}\right\}
\end{aligned}$$

where the variance expressions follows from Eq. (B3).