Gas storage valuation using a multi-factor price process

Alexander Boogert*   Cyriel de Jong†

September 1, 2011

Abstract

In this paper we discuss an extension to a popular gas storage valuation method called the spot approach. Least-Squares Monte Carlo, which is the basis for the spot approach, allows for multi-factor price processes. Such price processes can capture more realistically the actual price behavior present in energy markets. In this paper we demonstrate the application of multi-factor Least-Squares Monte Carlo to gas storage valuation. We study the impact of using multi-factor price processes on different aspects of the valuation such as convergence, average storage value and distribution of storage values in a numerical example. We find a counter example to the idea that an increase in market volatility leads to an increase in storage value. As well, we find a counter example to the idea that the natural hedging strategy of the spot approach is no hedge: a simple static financial hedge can reduce the inherent risk of the spot approach. Finally, we study the impact of model error related to the price process.

1 Introduction

Gas storages have traditionally been used to match supply and demand throughout the year. In the current environment of liberalized gas markets including third-party access to gas storage, valuation and hedging of gas storages deserves our attention. Gas storages are managed by utilities and merchants all around the world. Their motive is to either use gas storages for portfolio purposes, pure market trading or a combination of these. In this paper we take the market-based valuation perspective, as it can create an independent benchmark.

In practice three different valuation techniques are used to operate a gas storage: (rolling) intrinsic, spread options, and the spot approach. In the intrinsic approach the current forward curve is taken, and an optimal volume path is determined. An important decision is which forward to use for the intrinsic valuation. One extreme assumption is that each individual day in the future can be traded as a separate forward contract. A more realistic assumption is

*University of London, Birkbeck College, Commodities Finance Centre & EnergyQuants, Amsterdam, The Netherlands. Corresponding author

†Kyos Energy Consulting, Haarlem, The Netherlands

that only truly tradable products can be traded. We should realize that these different assumptions might lead to quite different results, especially when additional constraints are to be satisfied or when the intrinsic trading strategy is applied dynamically in the rolling intrinsic approach.

The rolling intrinsic approach adjusts the intrinsic hedge over time whenever additional money can be made. For example, when the December contract was initially sold forward, but the February contract becomes more expensive, a swap between the two months creates additional money. How much additional money can be made depends on the movements of the forward curve and is especially dependent on switches in the curve. Single-factor gas price models lead to very few switches, and therefore do not demonstrate a lot of value. Note that the assumption of a daily forward curve in combination with a rolling intrinsic strategy (e.g. made in Bjerksund et al. [2]) leads to unrealistically many switches; the correspondingly high storage value cannot actually be captured in the market and should be considered rather theoretically.

A second valuation approach is based on spread options. The idea is to capture the payoff of potential switches of the forward curve by a spread option. In a world where virtually no spread options are being traded (as in the North-American and European gas markets), the spread options approach becomes a delta-hedging approach. An extended description can be found in Lai et al. [7]. The third valuation approach is the spot trading approach; it seeks to capture short-term volatility by trading only in the spot, which is most volatile and exhibits mean-reversion. This spot approach is the focus of the current paper.

If we consider price developments in commodity markets, we find they contain a wide variety of changes in the spot and forward curve. It is impossible to incorporate such changes into a one-factor stochastic price model. Researchers and practitioners have therefore introduced multi-factor stochastic price models (see e.g. Schwartz [13], Hirsch [6], Lai et al. [7]). The introduction of multi-factor price models has raised the question how to deal with them in the optimization of different commodity derivatives. For example, a trader of gas storage tries to benefit from the changing shape in the complete forward curve, and wants to maximize his benefit given the different stochastic factors. Gas storage valuation is thus an interesting candidate to study the impact of multiple stochastic forward curve drivers.

Least-Squares Monte Carlo (or: LSMC) allows for the incorporation of multiple stochastic factors in the trading strategy (Longstaff & Schwartz [10]), but this has not yet been demonstrated in commodity markets. Therefore we decided to study the LSMC method with multiple stochastic factors using gas storage valuation as an example. Our starting point for such analysis will be the earlier work Boogert & De Jong [3], where we adjusted the LSMC method for gas storage valuation using the spot approach. The price process used was a 1-factor mean-reverting price process. In this paper we create the extension to multi-factor price processes, and it leads to various new and unexpected insights. We first extend the original LSMC storage valuation method to a multi-factor price process. Second, we analyze the impact on storage value and separate the price assumptions in the backward valuation step from the price assumptions in the forward valuation step.

In particular, three new insights are derived in this paper concerning volatility and gas storage optimization. It is generally valid that an increase in market volatility enhances the value of options. In this paper we find a counter example: when the level of the forward curve, as influenced by a long-term stochastic factor, becomes more uncertain, storage value actually decreases. This can be explained by the fact that a storage trader does not benefit from movements in the level of the forward curve, but rather from movements in the shape of the forward curve. More uncertainty about the level blurs the trader's assessment of whether he should inject or withdraw. Apart from creating additional risk, measured by a wider value distribution, the long-term stochastic factor reduces storage value on average.

Second, we quantify the impact of model error related to the price process underlying the simulations. The separation of the LSMC approach in a backward (which yields the trading / exercise strategy) and a forward valuation step (which yields the valuation) allows us to study the impact of model error on the valuation. We find that as long as the backward and forward valuation step are based on the same set of stochastic factors, the storage value is about the same with one, two or three factor price processes. When we introduce model error, we find instead a clear value reduction. More specifically, if we use in the backward valuation step a reduced set of stochastic factors, and apply this strategy in the forward valuation step with the full set of stochastic factors, we find a clear loss of value. We thus demonstrate the impact of model error can be significant.

Third, we propose a simple and effective hedging strategy for the spot approach. In earlier literature, the spot trading approach was always presented as a rather risky strategy. In this paper we find that a simple static financial hedge taken at the inception of the contract decreases the expected standard deviation. In our numerical example the static financial hedge resulted in 50 percent less standard deviation. Thus, we conclude the spot approach can be traded with less risk than previously assumed.

Several authors have recently discussed and applied the LSMC method into different promising directions. To our knowledge Li [9] is the only reference in which a multi-factor price process is used in the optimization of a spot-based model. Independently, we have introduced a multi-factor price process in the storage valuation model in a similar fashion. It is surprising that the combination of the spot approach with a multi-factor price process has not been described in detail so far. One potential reason could be that the spot approach in principle only performs spot actions, which can lead to the (incorrect) idea that for explaining this action, only a spot factor is relevant. In this paper we will show that this interpretation yields incorrect valuations, and that one should not apply it in practice.

Felix & Weber [5] compare the spot approach based on recombining trees and the spot approach based on LSMC. In the case considered, they find that the resulting valuations are very similar. They conclude the recombining trees are faster, but less flexible with respect to different price simulations. To us, especially flexibility in price simulations are an important reason to choose for the LSMC method.

Hirsch [6] points to the importance of price dynamics and speed of optimization while applying the LSMC method to hourly electricity swing options. In optimization, this problem is closely related to the gas storage valuation problem, because up and down swings are allowed within a certain given volume range. When only up-swings are allowed (as in the example of the paper), the problems even coincide. It is important to realize a multi-factor price process is used, while only a spot factor is included in the regression. In our numerical example we will show this can be improved significantly by adding more factors to the regression.

Neumann & Zachmann [12] applied the spot approach to the German Dötlingen storage, for which (an exception in Europe) utilization rates are published. They find the model-implied and observed volume development do not match, and conclude German storage sites are not operated on a purely profit maximizing basis.

2 Valuation method

In this section we discuss the new valuation method for gas storage. We start in section 2.1 with a description of a storage contract and in section 2.2 with the spot optimization for the one-factor price process. Both were introduced before in Boogert & De Jong [3]. In section 2.3 we present the focus of this paper: the extension for the multi-factor price process.

2.1 Description of a storage contract

We denote the (accumulated) volume in storage at the start of day t by $v(t)$. We take an injection at day t as a positive volume change $\Delta v(t)$ and a withdrawal as a negative volume change $\Delta v(t)$. We denote the payoff at day t by $h(S(t), \Delta v(t))$ and define for $t = 0, \ldots, T$:

$$h(S(t), \Delta v(t)) := \begin{cases} -c(S(t))\Delta v(t & \text{inject at day } t \\ 0 & \text{do nothing at day } t \\ -p(S(t))\Delta v(t) & \text{withdraw at day } t \end{cases} \quad (1)$$

where $c(S(t))$ and $p(S(t))$ represent the cost of injection and profit of withdrawal, which can include both transaction costs and bid-ask spreads.

There can be different volumetric limitations on the followed strategy stemming from the physical nature of a gas storage. Standard limitations include a minimum and maximum volume level, and maximum injection and withdrawal rates. The case of injection and withdrawal dependent on the accumulated volume (also called ratchets) can be included easily in terms of the discretization of the accumulated volume. We denote the set of allowed volume levels at day t by $\mathcal{V}(t)$, and the set of all allowed actions on day t being at volume $v(t)$ by $\mathcal{D}(t, v(t))$.

2.2 Spot optimization for the one-factor price model

Similar to American option valuation, a central role in the storage valuation is played by the continuation value. We define the continuation value as the value we attach to the contract after taking an allowed action, that is $\Delta v \in \mathcal{D}(t, v(t))$.

The valuation is performed by discretizing both volume and volume actions into a fine grid. In Boogert & De Jong [3] it was motivated that it is more efficient to consider continuation values on the time-volume grid instead of per time unit. The continuation value is thus denoted by $C(t, S(t), v(t+1; n))$. In order to handle situations where actions fall outside the grid, we perform an interpolation where continuation value for intermediate volume levels are calculated as the distance weighted average of the continuation values at adjacent volume points.

If we denote the value of a storage contract starting on day t at volume level $v(t)$ by $U(t, S(t), v(t))$, then $U(t, S(t), v(t))$ satisfies the following dynamic program:

$$U(t, S(t), v(t)) = \max_{\Delta v \in \mathcal{D}(t, v(t))} \{h(S(t), \Delta v) + C(t, S(t), v(t), \Delta v)\} \tag{2}$$

for all t. The dynamic program can be initialized at the end of the contract using either zero costs (the desired end volume is reached) or a penalty cost (the desired end volume is not reached).

The assumption behind regression-based Monte Carlo techniques such as LSMC is that one can approximate continuation values by a (finite) linear combination of known basis functions $\phi_q(S(t))$ of the current state. Because the continuation values are taken on the time-volume grid, the state in the basis functions can be summarized by a single variable: the spot price $S(t)$. This means we approximate

$$C(t, S(t), v(t+1; n)) \approx \sum_{q=1}^{Q} \phi_q(S(t)) \beta_{q,t} \tag{3}$$

for certain constants $\beta_{q,t} \in \mathbf{R}$ and $Q \in \mathbf{N}$.

In practice the continuation values are unknown, and we follow Longstaff & Schwartz [10] in assuming the following approximation for all M independent paths from the simulation, $b = 1, \ldots, M$:

$$C^b(t, S^b(t), v(t+1; n)) \approx e^{-\delta} Y^b(t+1, S^b(t+1), v(t+1; n)) \tag{4}$$

where we denote by $Y^b(t+1, S^b(t+1), v(t+1; n))$ the accumulated value of future realized cash flows in path b following optimal decisions starting at time $t+1$ being at volume level $v(t+1)$ and price $S^b(t+1)$. With approximation (4), we can estimate the best regression coefficients $\hat{\beta}$ by an ordinary least-squares (OLS) regression. If we substitute these $\hat{\beta}$ back into Equation (3), we get an approximation $\hat{C}^b(t, S^b(t), v(t; n) + \Delta v)$ of the continuation value for all volume points $v(t; n)$. This allows us to determine the optimal action for all volume points $v(t; n)$.

2.3 Spot optimization for the multi-factor price model

The key ingredient for the spot optimization is to find a good approximation for the expected continuation value based upon the simulated spot prices. In a one-factor model the single spot factor explained the behavior of the complete forward curve. In a multi-factor price model there are several factors driving the forward curve.

A natural extension in case of multiple factors is to include them all as explanatory variables in the basis. This idea was already proposed in the original paper by Longstaff & Schwartz [10]. For example, Longstaff & Schwartz value a cancelable index amortizing swap where the swap term structure is driven by two independent processes: X, Y. In the regression they use 9 basis functions: a constant, the first three powers of the value of the underlying non-cancelable swap, and X, X^2, Y, Y^2, XY.

In case more than one factor is introduced, not only the spot factor, but also the forward factors and combinations between them can be used inside the basis. The number of basis functions naturally grows by the inclusion of more factors. One alternative, e.g. used in Stentoft [15], uses a complete set of polynomials of order K. This means that all products and cross products of order less than or equal K are included. In our numerical example with a three-dimensional price process, $K = 2$ implies a basis of size 10, and $K = 3$ implies a basis of size 25 (using Stentoft [15, Eq. 20] with $L = 3$). This means in practice K is limited to either 2 or 3. Another alternative, e.g. used in Longstaff & Schwartz [10], uses a limited but arbitrary combination of them.

In a multi-factor setting it becomes more important how to perform the regressions in a stable manner (a similar point was made in Hirsch [6]). In our numerical example we found that using a complete set of polynomials of order 3 led to problems with inversion of matrices, whereas a complete set of polynomials of order 2 underperformed. Therefor we have chosen to work with a limited but arbitrary combination of polynomials.

The usual practice in LSMC applications is to keep the number of basis functions constant over time. We compare this method to a greedy heuristic which at every time step takes a large basis and reduces the basis until the inverse becomes stable. This heuristic is inspired by Longstaff & Schwartz [10, Prop. 1] which led to the criterion to increase the number of basis functions for a given number of simulations until the value no longer increases. Stentoft [16] motivates that instead both the number of basis functions and the number of simulations should be increased. In our greedy heuristic we keep the number of simulations fixed. The only difference between the usual method and the greedy heuristic concerns how the regression step is performed. It is worth noting that the greedy heuristic adds only limited calculation time. One can determine whether the inverse is stable without actually performing a computationally expensive inverse.

First, suppose there is only a single spot factor. We let P the vector of simulated spot prices of length M, and define the (M, R_{spot}) regressor matrix X using $X := [1, P^1, P^2, \ldots, P^{R_{spot}-1}]$ where P^i denotes the i-th member of some polynomial family, e.g. power. We assume a linear relation between the M-vector continuation value V and the regressor matrix X such that $V = XB$, where B is a vector of constants of length R_{spot}. Then, B can be estimated by the Moore-Penrose pseudo-inverse (also known as generalized inverse):

$$B = (X'X)^{-1}X'V \tag{5}$$

where $^{-1}$ denotes the inverse of square matrix $X'X$. The greedy heuristic sets R_{spot} to a large number and decreases R_{spot} until the conditional number (measured by Matlab function *rcond*) crosses a pre-set barrier.

Next, suppose there are multiple price factors, and we need to approximate the continuation value using both spot and forward prices. We propose to use the above greedy heuristic in a sequential way. First we perform the above approximation using only the spot factor, and then sequentially consider whether we can improve the approximation by including the other price factors and/or cross-factors. For example, a basis could become $[1, P^1_{spot}, P^2_{spot}, P^3_{spot}, P^1_{LT}, P^2_{LT}, P_{WS}, P^1_{spot*LT}]$ if we start with $R_{spot} = 4$, $R_{LT} = R_{WS} = 3$, $R_{spot*LT} = R_{spot*WS} = R_{LT*WS} = 1$. Here $R_{\cdot}$ is the maximum number of members included of type $\cdot$. This is the setting we use in part of our numerical example.

The greedy heuristic leads to basis $\hat{X}$, which can then be inverted again. This means we have established the expected continuation value for all grid points and all simulated prices, and we are back into our original setting. The next step is to find the action which yields the highest total value (equal to current payoff + expected continuation value).

An alternative extension, presented by Li [9], extends the one-factor spot approach using the following two step method.[1] First, estimate loading factors to describe the forward curve movements present inside the simulation using principal components analysis (or: PCA). Second, use the loading factors as state variables in the regression.

The main difference between the two methods lies in the first step. We propose to use the original factors of the price simulation instead of estimated loading factors. One advantage is that that the price process can combine time-to-maturity and time-of-maturity effects, whereas PCA focusses on time-to-maturity effects. An example how time-of-maturity effects can impact PCA can be found in Koekebakker & Ollmar [8]. They find that in the Nordic electricity market more than 10 PCA factors were needed to explain 95% of the term structure variation. In our numerical example we use a price process containing both time-to-maturity and time-of-maturity effects. Another difference is that Li [9] estimates one pair of loading factors, whereas we use the time-varying price factors. An alternative to the method by Li would be to re-estimate the loading factors at each time step in order to capture time-varying effects. This would be especially relevant if the simulations contain structural breaks. In

[1] Although Li [9] also mentions the extension of the spot approach with bid-ask spreads, we note these were already present in Boogert & De Jong [3, Eq. 3+4].

addition to the different factor method, we study the convergence behavior and the impact of using different basis functions, which are not present in Li [9].

3 Empirical results

In the LSMC method it is not a priori clear which kind of basis functions, how many basis functions and how many paths to use. We investigate these questions in this section for a specific numerical example. In section 3.1 we describe the price process and in 3.2 the specific storage we will study. In section 3.3, respectively 3.4 and 3.5, we study the impact of different trading strategies, respectively the impact of different basis functions and market parameters. In 3.6 we discuss the results for a more flexible storage.

3.1 Price process

In this numerical example we use a three-factor price model. It is an extension of the mean-reverting one-factor Schwartz [13] model used in Boogert & De Jong [3]. A common extension to that model is to introduce a second factor which moves all prices on the forward curve. An example is the long-term/short-term model by Schwartz & Smith [14]. Our third factor moves the winter-summer spread in the forward and spot market. To our knowledge, this explicit modeling of the winter-summer spread has not been published in the literature so far. The advantage is that the winter-summer spread is naturally monitored in the gas forward market and thereby provides a direct link to the market behavior.

The starting point for our three-factor model is the original mean-reverting one-factor Schwartz [13] model, also referred to as a discrete-time Ornstein-Uhlenbeck process. In log-terms it is given by

$$d\ln P^{ST}(t) = \kappa\left[\mu(t) - \ln P^{ST} - \frac{(\sigma^{ST})^2}{2\kappa}\right]dt + \sigma^{ST}dW^{ST}(t) \tag{6}$$

where $P^{ST}(t)$ is the spot price and the mean level $\mu(t)$ is a deterministically time varying function. The daily mean-reversion rate κ and volatility σ^{ST} are assumed to be constant.

For the three-factor model we make $\mu(t)$ stochastic using two additional factor returns: a long-term (LT) and a winter-summer (WS) factor. We define $\mu(t,T)$ as the mean level T periods from now, measured at time t. The variable $\mu(t,t)$ then naturally becomes the spot mean level. The two additional factor returns affect both the spot price and the forward mean level. The long-term factor affects all prices on the curve equally. This move is independent from where we are on the forward curve (T). The impact of the third factor, however, depends on where we are in the season: a positive return moves winter prices up, summer prices down and leaves prices in between less affected. This dependency is captured by the function $P^{seas}(t)$. We assume there is no correlation between the three stochastic Brownian motions $W^{ST}(t)$, $W^{LT}(t)$ and $W^{WS}(t)$.

$$
\begin{aligned}
d\ln P^{ST}(t) &= \kappa\left[\mu(t,t) - \ln P^{ST} - \frac{(\sigma^{ST})^2}{2\kappa}\right]dt \\
&+ \sigma^{ST} dW^{ST}(t) + \sigma^{LT} dW^{LT}(t) + \sigma^{WS} P^{seas}(t) dW^{WS}(t) \quad (7) \\
\mu(t,t) &= \mu(t,0) + \sigma^{LT} \sum_{s=1}^{t} dW^{LT}(s) + \sigma^{WS} P^{seas}(t) \sum_{s=1}^{t} dW^{WS}(s) \quad (8)
\end{aligned}
$$

The seasonal dependency function $P^{seas}(t)$ is constructed as follows. First, we identify the highest point on the forward curve with the winter (e.g. first of February), and 6 months later with the summer. Next, we desire that if the winter-summer spread changes due to a stochastic move, these movements have an appropriate effect on the contracts around the summer and the winter date. We model this effect using a sine-function $P^{seas}(t)$ with a periodicity of one year. The highest point of 0.5 is on the winter date, and the lowest point of -0.5 is on the summer date. In this way the complete forward curve changes with the winter-summer spread. The stochastic winter-summer movement thus changes the winter-summer spread.

The three-factor price process combines a time-of-maturity effect with a time-to-maturity effect. There are four parameters to be estimated from external data: κ, σ^{ST}, σ^{LT} and σ^{WS}. The long-term volatility almost equals the volatility of an annual contract further on the curve. Such a contract, e.g. 1 year ahead, is not dependent on the winter-summer variations and, assuming sufficient mean-reversion, it is hardly dependent on the short-term (spot) returns.

To derive the winter-summer volatility we use a Q1-forward and a Q3-forward contract. The difference in return between the Q1 and Q3 leads to a series of winter-summer returns. Using the integral over the appropriate 3-month period of the seasonal sine-function, the Q1 and Q3 contract have a weight of approximately 0.9 on the 'peak' winter-summer volatility. Consequently, σ^{WS} equals the volatility of the Q1-Q3 return series divided by 0.9.

To derive the spot mean-reversion rate and short-term volatility we use a history of spot prices. First, we clean the returns from the long-term and seasonal dependence (derived from the forward price returns). The problem then is that the mean-level is stochastic and unobserved. One could use a filtering technique to jointly derive the development of $\mu(t,t)$ and the spot parameters. However, for practical reasons, in this calibration, we assume that $\mu(t,t)$ can be approximated by the log of the month-ahead forward price. For a more detailed explanation we refer to De Jong & Schneider [4].

As a base case in the remainder of this paper we will be using $\kappa = 12\%$, $\sigma^{ST} = 100\%$, $\sigma^{LT} = 20\%$ and $\sigma^{WS} = 20\%$. These parameter estimates correspond with rounded numbers of estimates obtained from Dutch TTF data for the period 2008-2010.

3.2 Example setting

In our numerical example we study the following test storage:

- minimum volume = 0 units, maximum volume = 100 units
- start volume = end volume = 0 units
- maximum injection = maximum withdrawal = 1 unit/day
- trading costs: 0
- trading period: 1 year
- spot volatility $\sigma^{ST} = 100\%$, long-term volatility $\sigma^{LT} = 20\%$, winter-summer-volatility $\sigma^{WS} = 20\%$, spot mean-reversion rate $\kappa = 12\%$

To study the impact of using a multi-factor price process and a multi-factor optimization, we consider three different cases as in Table 1. Case 1 represents the approach in Boogert & De Jong [3] where a one-factor model (spot factor) was used to value a gas storage. The only basis function in case 1 is the spot factor. Case 2 represents an extension to case 1 by incorporating a multi-factor price process (spot, long-term and winter-summer factor) for the simulation of prices. Meanwhile the approach from Boogert & De Jong is maintained to optimize the storage. The resulting spot prices are simply taken and used in the basis for the optimization. Thus, while there are additional price factors used for the simulation of prices these are neglected in the basis. In this way the only basis function in case 2 is the spot factor. Case 3 represents an extension to case 2 as it changes the approach used to optimize the gas storage. In case 3 the additional price factors (spot, long-term and winter-summer factor) are included in the basis in the optimization, such that we have three basis functions. Case 3 represents our suggested approach to use the three factors both in the simulation of prices and in the optimization.

Case	Simulation			Optimization		
	spot	lt	ws	spot	lt	ws
1	yes	no	no	yes	no	no
2	yes	yes	yes	yes	no	no
3	yes	yes	yes	yes	yes	yes

Table 1: Characteristics of the three study cases. We use either 1 or 3 factors for the simulated price process, and either 1 or 3 factors for the basis functions in the optimization

In Figure 1 we show the employed forward curve for the numerical example. In the figure also the development of the spot prices is indicated for case 2 and 3 (including long-term and winter-summer volatility) by means of the 10% and 90% percentiles of the simulated spot prices.

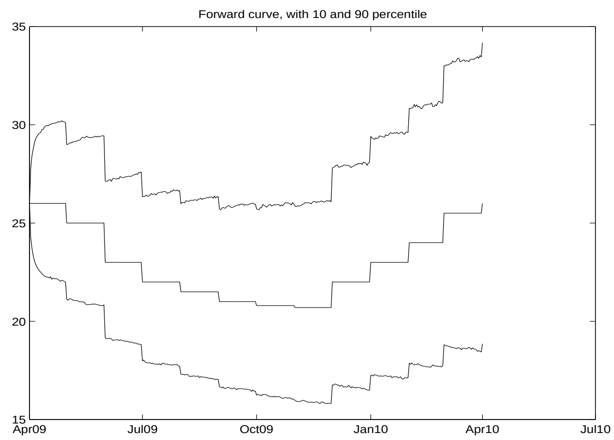

Figure 1: Forward curve (middle) together with the 10% and 90% percentiles of the simulated spot prices for storage case 2 and 3 with 10000 simulations

Whenever sensitivity analysis is presented in this paper, the results are based on the same 60 different seeds.

3.3 Impact of different trading strategies

We implemented the discussed algorithm in Matlab. As a start we valuate the storage using the different strategies mentioned in the introduction. The results are given in Table 2. We see the storage trader can make a profit of 306 assuming he can trade every day in the future individually. Assuming only the truly tradable products can be traded (in this case assumed to be 6 months, 4 quarters, 3 seasons and 3 years), the value decreases to 275. If we perform a tradable rolling intrinsic strategy (here we assume once a month rehedging) we capture 275 at the start date and an expected additional 102 extrinsic value over the life time of the contract. The spot strategy delivers an expected 665, albeit with a larger standard deviation than the rolling intrinsic (270 versus 44).

We found out it is possible to strongly reduce the standard deviation of the spot strategy in our example by inclusion of a static financial hedge on the forward market. Because we assume no costs in this example and the price process does not contain any risk premium, the expected value of the static financial hedge is zero. The expected value including the static financial hedge thus remains 665. The standard deviation however reduces from 270 to 114. In case transaction costs would be present, these would influence the value of both the spot strategy and the static financial hedge. If we assume the bid-ask spread at the spot and forward market to be 0.20, the expected value of the static financial hedge reduces from 0 to -26 (130 volume units are traded at the forward market and later unwound at the spot market). The application of a static financial hedge in combination with the spot approach appears to be a new result in the spot approach literature, and we see hedging as an interesting direction for future research. For example, one could investigate the possibility to maintain a dynamic hedge.

In our example, we constructed the static financial hedge as follows. First we performed the spot approach as proposed above. This results in an expected action for each day in the contract period which we aggregate into monthly actions. The hedge already sells gas at day 0 on the forward market when the expected monthly action is negative, and already buys gas at day 0 on the forward market when the expected monthly action is positive. When the forward hedge goes into delivery, we take the opposite position in the spot market. There are two things worth pointing out. In the first place, note that the static financial hedge does not influence our spot approach. The spot approach during the contract period functions independent of the existing forward hedge position. In the second place, note that we used in our example a static financial hedge. The position is taken at time zero, and is not changed regardless of the subsequent price development during the contract period.

Strategy	Mean	St. dev.
Daily intrinsic value	306	
Tradable intrinsic value	275	
Spot value	665	270
Spot value incl. hedge	665	114
Rolling intrinsic value	377	44

Table 2: Valuation results for storage case 3 with 10000 simulations for different trading strategies

Next, let us compare the results from storage case 3 (use multiple factors in simulation and optimization) with the results from storage case 1 (use only spot factor in simulation and optimization) and 2 (use multiple factors in simulation, and only spot factor in optimization). The resulting valuation is given in Table 3 and the resulting histograms are shown in Figure 2. It is obvious that not including multiple price factors into the basis has a negative effect on the valuation: the expected value decreases from 665 (storage case 3) to 497 (storage case 2)! From the histograms we see a large negative tail in case 2, which is not present in case 3. This can be explained by the fact that spreads are not appropriately monitored in storage case 2. If for example, the forward curve moves up, an initially high spot price might not be high anymore. This can lead to wrong decisions, and negative outcomes as indicated by the large negative tail. The lowest diagram indicates the static financial hedge works on both sides of the distribution: the static financial hedge cancels out both large negative and large positive results from the spot strategy.

If we compare storage case 1 with storage case 3, we compare the appropriate valuations of a one-factor with a multi-factor price process (use as many factors in the optimization as in the simulation). We see that the expected value increases from 654 in case 1 to 665 in case 3, while the standard deviation increases from 96 to 270 in case 3 without a hedge (respectively 114 for case 3 with a hedge). It is not surprising the standard deviation increases because there is more uncertainty. Taking this uncertainty into account, the spot strategy creates an additional profit of 11. We conclude that taking into account uncertainty might yield additional profits, as long as this uncertainty is taken into account in the optimization.

The above results were produced using a single initial seed. Let us consider whether these results are consistent when we use a different number of simulations or a different initial seed. For this reason we compare results for 8 different numbers of simulations and 60 different initial seeds. The results are summarized in Figure 3, where we show the mean of the 60 storage valuations for the 8 different numbers of simulations. From Figure 3 we conclude the above results are consistent. It also shows that starting at 5000 simulations the results appear stable, especially in the relevant storage case 3. The standard convergence of a Monte Carlo method is $\frac{1}{\sqrt{N}}$ if N is the number of simulations. In order words, the logarithm of the standard error versus the logarithm of the number of simulations has a slope of -0.5. In the lower panel of Figure 3, we present the log-log plot. The slope was estimated to be -0.51 (case 1), -0.46 (case 2) and -0.50 (case 3), which is close to the expected theoretical value -0.50.

Case	Mean	St. dev.
1	654	96
2	497	259
3	665	270
3 + hedge	665	114

Table 3: Valuation results for storage case 1, 2 and 3 with 10000 simulations for the spot strategy

3.4 Impact of different basis functions

In the previous subsection we discussed how many simulations to use. In this subsection we discuss which kind of basis functions and how many of them to choose. For this reason we compare the storage valuation using five different polynomial families. Moreno & Navas [11] used ten different families and found small differences between them on a simple derivative like an American put but large differences on a difficult derivative like the option on the maximum of five assets. In our test we use five of the ten families employed by Moreno & Navas. These five families are presented in Table 4: besides powers, we employ Chebyshev of the first kind A, Legendre, Laguerre and Hermite A.

Moreno & Navas point to the fact the families are orthogonal on certain intervals, but state “In most of the cases the range of underlying prices is differ-

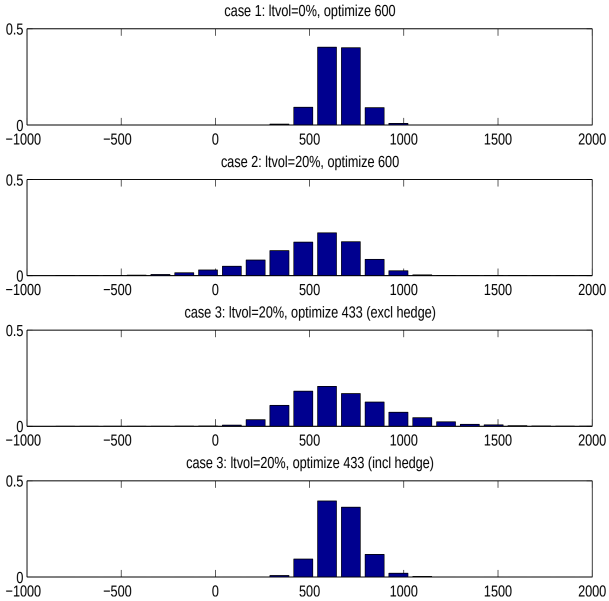

Figure 2: Resulting value histograms from the three different storage cases. The bottom diagram represents case 3 with the static financial hedge included

ent from the interval $[a, b]$, so that the basis functions will not be orthonormal. Consequently, we should increase the number of terms used in the regression." In this paper we instead standardize the explanatory variables to align with the orthogonality intervals. Our standardization is given in column 4 of the lower table in Table 4. The remainder of the table is taken from Abramowitz & Stegun [1].

In particular, we perform standardization different from most option applications where one normally divides by the strike and does not consider out-of-the-money prices. This leads for example for an American put to the interval $[0, 1]$. This approach is not valid for gas storage valuation because there is no strike. In Figure 4 we consider the impact of standardization in our numerical example. In that figure we show the distribution of the three factors after standardization by substraction of the mean and division by three times the standard deviation. From these histograms it is clear that division by three times the standard deviation yields a transformation to $[-1, 1]$ for most of the data points. This is relevant for Legendre and Chebyshev of the first kind A, which have orthogonality interval $[-1, 1]$.

As indicated by Moreno & Navas [11], to use no weighting implies all the families can be rewritten into ordinary powers, creating the same regression. This means one expects no differences in the valuation from a theoretical perspective although numerical errors can occur. We find this to be correct for

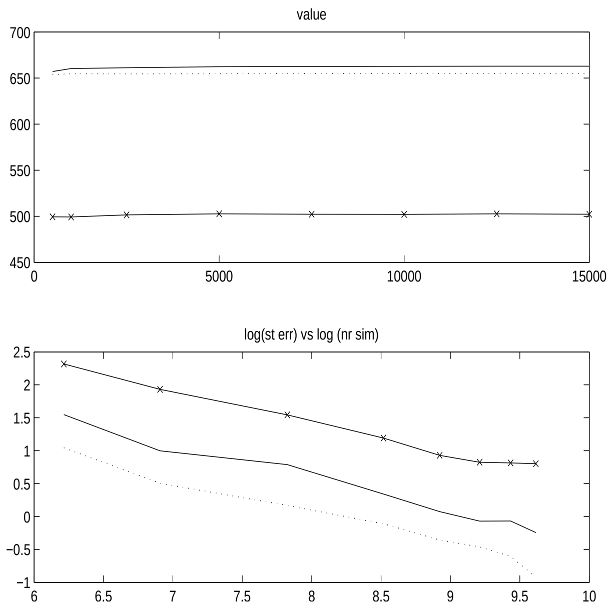

Figure 3: Top diagram: average value for case 1 (dots), case 2 (x) and case 3 (line) for different number of simulations: $500, 1000, 2500, 5000, 7500, 10000, 12500, 15000$. Bottom diagram: logarithm of the standard error versus the logarithm of the number of simulations. Basis: non-greedy powers

all cases except when case 3 is solved with the greedy heuristic. In that case Hermite polynomials clearly under perform, which is in line with the results obtained below.

Therefore we apply weighting as presented in Table 4. To be precise, we first standardize our price factors, then apply the recurrence relation, and finally multiply with the weighting factor. The results are given in Figure 5 for the non-greedy optimization and Figure 6 for the greedy optimization. We can draw a number of conclusions from here:

- Greedy optimization creates lower values than the non-greedy optimization. The differences are especially large when Hermite polynomials are used.
- Chebyshev and Hermite polynomials under perform in all cases in the non-greedy and greedy optimization basis. Surprisingly in the greedy optimization of case 1, Hermite polynomials still perform well.
- Legendre and powers provide exactly the same estimates for the mean value in the non-greedy optimization, which are almost the same as La-

Family	P_0	P_1	P_{n+1}	a	b
Powers	1	x	xP_n	–	–
Chebyshev 1A	1	x	$2xP_n - P_{n-1}$	-1	1
Legendre	1	x	$[(2n+1)xP_n - nP_{n-1}]/(n+1)$	-1	1
Laguerre	1	$1-x$	$[(2n+1-x)P_n - nP_{n-1}]/(n+1)$	0	∞
Hermite A	1	$2x$	$2xP_n - 2nP_{n-1}$	$-\infty$	∞

Family	weight	result	$\hat{x}$
Powers	–	–	$(x-\mu)/\sigma$
Chebyshev 1A	$(1-x^2)^{-1/2}$	$\pi/2$	$(x-\mu)/(3\sigma)$
Legendre	1	$2/(2n+1)$	$(x-\mu)/(3\sigma)$
Laguerre	e^{-x}	1	$1+(x-\mu)/(3\sigma)$
Hermite A	e^{-x^2}	$\sqrt{\pi}2^n n!$	$(x-\mu)/\sigma$

Table 4: Description of the different basis families. In the upper table we provide for five different basis functions, the first two members (column 2 and 3) together with the recurrence relation (column 4) and the orthogonality interval (column 5 and 6). In the lower table we provide the weight function (column 2), the result of integration of two non-identical members (column 3) and our standardization (column 4). μ and σ are the mean respectively standard deviation of x

guerre. In the greedy optimization Legendre creates slightly higher mean value and shows better standard deviations than the powers and Laguerre.

- Convergence of Legendre and Laguerre is similar to powers, and the gradient of the log-log plot remains around -0.50 for all three in greedy and non-greedy optimization, similar as we found earlier for powers in the previous subsection. However, in the greedy optimization, the standard deviations of Laguerre increases again when we increase the number of simulations above 10000 whereas the others continue to go down.
- Because powers take less computational time than Legendre we conclude powers are the preferred choice of basis functions.

3.5 Impact of market parameters

We compare the results if the different market parameters change: long-term volatility, short-term volatility, winter-summer volatility and mean-reversion. Results are obtained by changing a single market parameter, and by averaging the result of 60 initial seeds for the case of 10000 simulations. The results are presented in Figure 7.

We note the special behavior which the storage value shows with respect to volatility: storage value is increasing in both spot and summer-winter volatility but is decreasing in the long-term volatility. This goes against the common

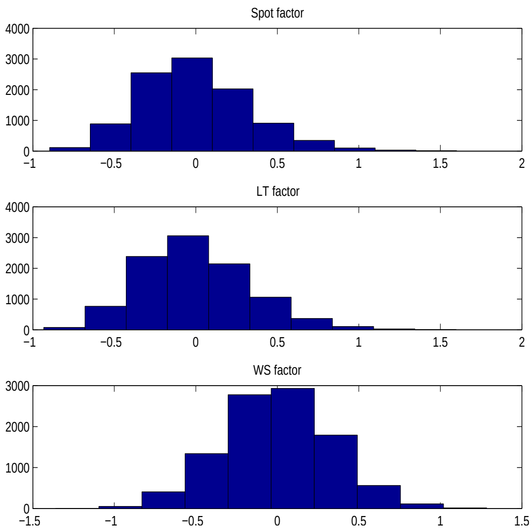

Figure 4: Impact of standardization by subtracting the mean and dividing by three times the standard deviation of each of the three different factors: spot, long-term and winter-summer. Prices taken from case 3 at day 200

intuition that all volatility is good for the storage trader. This exception can be understood if we consider what long-term volatility means for the storage trader. With increasing uncertainty about the long-term level it becomes harder to decide whether prices are high or low today. At the same time, the storage trader does not benefit if the level of the forward curve changes, only if the shape of the forward curve changes.

3.6 A fast storage

The setting for the above example was a relatively slow storage, which could represent a depleted field gas storage. In this subsection we perform a robustness check: we consider here a fast storage, which could represent a salt cavern gas storage. For this reason we change the injection to 2 units (previously 1 unit) and withdrawal rate to 5 units (previously 1 unit), while keeping the same settings for the remainder of the example. In particular, the maximum volume remains at 100 units.

This fast storage can create a lot more value than the slow storage as it can perform several cycles going from minimum to maximum volume and back to minimum volume again. It can thus benefit better from short-lived price spikes, which is reflected in the storage valuation: the expected value increased from 665 to 1650.

For the fast storage we have repeated the valuations with different number of simulations and considered sensitivities to the different market parameters.

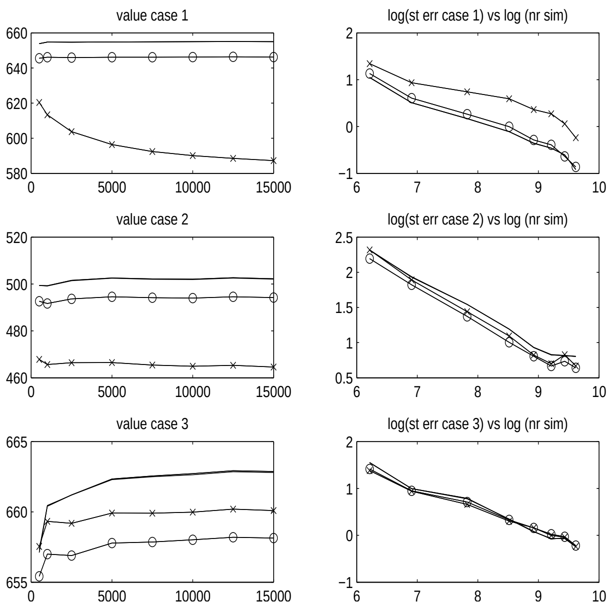

Figure 5: Left diagrams: average value for the three different cases for different numbers of simulations: $500, 1000, 2500, 5000, 7500, 10000, 12500, 15000$. Right diagrams: logarithm of the standard error versus the logarithm of the number of simulations. Chebyshev (x), Hermite (o), polynomial, Legendre and Laguerre (line) overlap. Non-greedy basis

The results can be found in Figure 8 and Figure 9 (for the slow storage it were Figure 5 and Figure 7). We can conclude from these Figures that there is no qualitative difference between the two cases in these aspects: the Hermite and Chebyshev polynomials perform less good than Legendre, powers and weighted Laguerre. The direction of all the sensitivities is the same, where the long-term volatility has a negative slope.

4 Conclusion

In this paper we considered the multi-factor aspect of the Least-Squares Monte Carlo method applied to gas storage valuation. Using multi-factor models we are better able to describe the actual price behavior present in energy markets. As a result, our assessment of the value and risks involved in gas storage trading come closer to reality. We discuss in this paper how to optimally price and hedge a storage asset in such a setting.

We find the application of Least-Squares Monte Carlo works well, and conclude we can include multi-factor price models into gas storage valuation by

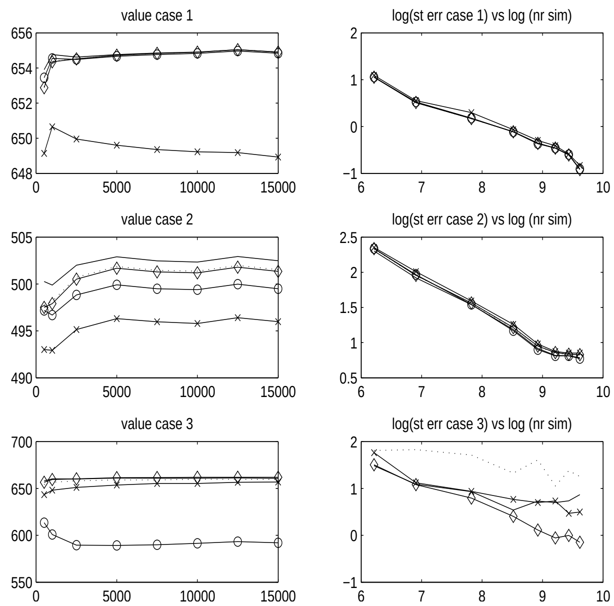

Figure 6: Left diagrams: average value for the three different cases for different numbers of simulations: $500, 1000, 2500, 5000, 7500, 10000, 12500, 15000$. Right diagrams: logarithm of the standard error versus the logarithm of the number of simulations. Chebyshev (x), Hermite (o), polynomial (dots), Legendre (diamond) and Laguerre (line). Greedy basis

a spot approach. When we compare different alternative basis functions, simple powers show a good performance. Going from a one-factor to a multi-factor price model, several observations can be made. First, the number of simulations should be increased, although a limited number like 5000 simulations already provides stable results. Second, we stress it is important to use the same factors for the simulation and optimization. More precisely, with a multi-factor process driving the prices, the model will undervalue the gas storage if the optimization only contains a single spot factor. Third, in a multi-factor setting the common wisdom that volatility is desirable for gas storage traders can break down. Finally, a multi-factor specification is clearly preferred above a single-factor specification: it allows for a better representation of actual price behavior, and shows a more realistic variation in potential storage trading results with and without hedging.

An interesting direction for future research concerns the hedging of the spot approach. In our numerical example we find a static financial hedge can yield a significant reduction of the estimated inherent risk. This new idea partially closes the gap between the inherent risk of the rolling intrinsic and spot approach, which makes the spot approach also interesting for traders with a limited

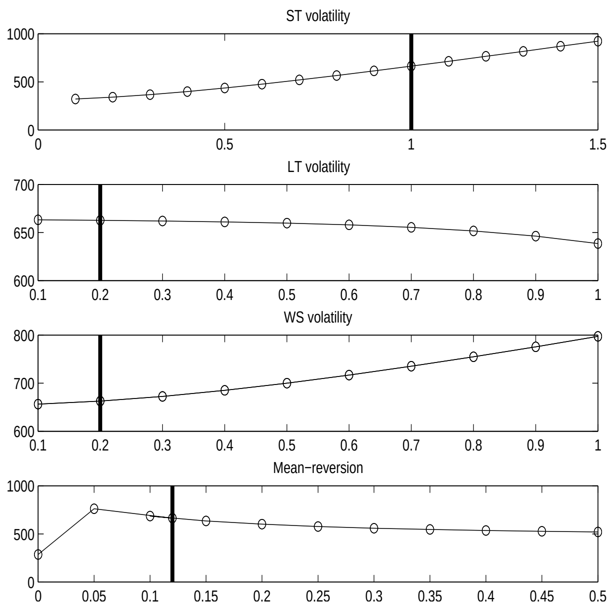

Figure 7: Sensitivity analysis of the four market parameters. Value is calculated by taking the average of 60 seeds for 10000 simulations. A vertical bar indicates the value in the original setting. Basis: non-greedy powers

risk appetite.

Acknowledgements

The authors want to thank María Cortés for her testing on different elements in the multi-factor storage optimization. In addition, we want to thank Björn Naundorf, Kai Tappe, Kay Pilz, Raymond Brummelhuis and the participants in the EURO XXIV Lisbon 2010 conference and the Energy Finance / INREC 2010 conference for their input to the current paper.

References

[1] Abramowitz, W. & I. Stegun, "Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables." Dover Publications, 1972.

[2] Bjerksund, P., G. Stensland & F. Vagstad, "Gas Storage Valuation: Price Modelling v. Optimization Methods." Discussion paper 20/2008, Norwegian School of Economics and Business Administration, 2008.

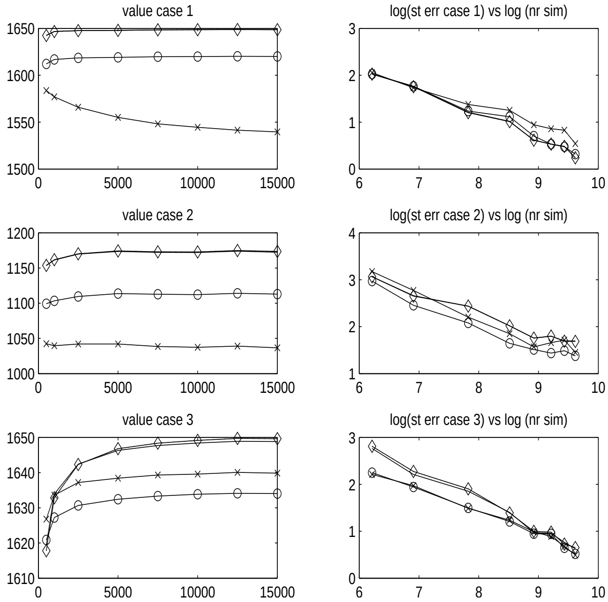

Figure 8: Left diagrams: average value for the three different cases for different numbers of simulations: 500, 1000, 2500, 5000, 7500, 10000, 12500, 15000. Right diagrams: logarithm of the standard error versus the logarithm of the number of simulations. Chebyshev (x), Hermite (o), polynomial, Legendre and Laguerre (line) overlap. Non-greedy basis. Fast storage

[3] Boogert, A. & C. De Jong, "Gas Storage Valuation using a Monte Carlo Method." *Journal of Derivatives*, vol. 15 (3), 2008, pp. 81-91.

[4] De Jong, C. & S. Schneider, "Cointegration between gas and power spot prices." *Journal of Energy Markets*, vol. 2 (3), 2009, pp. 27-46.

[5] Felix, B. & C. Weber, "Gas Storage Valuation: Comparison of Recombining Trees and Least Squares Monte-Carlo Simulation." Engineering Management Conference, IEMC Europe, 2008.

[6] Hirsch, G., "Pricing of hourly exercisable electricity swing options using different price processes." *Journal of Energy Markets*, vol. 2 (2), 2009, pp. 3–46.

[7] Lai, G., F. Margot & N. Secomandi, "An Approximate Dynamic Programming Approach to Benchmark Practice-based Heuristics for Natural Gas Storage Valuation." *Operations Research*, vol. 58, 2010, pp. 564–582.

[8] Koekebakker, S. & F. Ollmar, "Forward curve dynamics in the Nordic electricity market." *Managerial Finance*, vol. 31 (6), 2005, pp. 73–94.

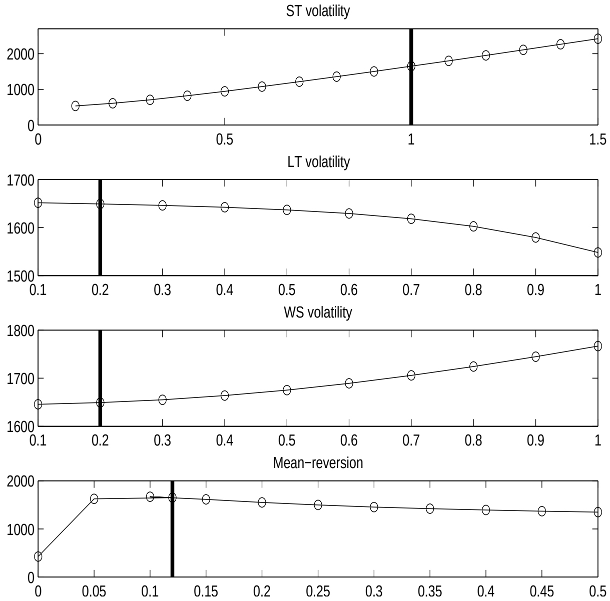

Figure 9: Sensitivity analysis of the four market parameters for the fast storage. Value is calculated by taking the average of 60 seeds for 10000 simulations. A vertical bar indicates the value in the original setting. Basis: non-greedy powers

[9] Li, Y., "Natural Gas Storage Valuation." Master thesis, Georgia Institute of Technology, 2007.

[10] Longstaff, F.A. & E.S. Schwartz, "Valuing American Options by Simulation: A Simple Least-Squares Approach." *The Review of Financial Studies*, vol. 14 (1), 2001, pp. 113–147.

[11] Moreno, M. & J.F. Navas, "On the Robustness of Least-Squares Monte Carlo (LSM) for Pricing American Derivatives." *Review of Derivatives Research*, vol. 6 (2), 2003, pp. 107–128.

[12] Neumann, A. & G. Zachmann, "Expected Vs. Observed Storage Usage: Limits to Intertemporal Arbitrage." In: A. Creti ed., *The Economics of Natural Gas Storage. A European Perspective*, Springer, 2009.

[13] Schwartz, E.S., "The stochastic behaviour of commodity prices: Implications for valuation and hedging." *Journal of Finance*, vol. 52 (3), 1997, pp. 923–973.

[14] Schwartz, E.S. & J.E. Smith, "Short-term Variations and Long-term Dynamics in Commodity Prices." *Management Science*, vol. 46 (7), 2000, pp. 893–911.

[15] Stentoft, L., "Assessing the Least Squares Monte-Carlo Approach to American Option Valuation." *Review of Derivatives Research*, vol. 7 (2), 2004, pp. 129–168.

[16] Stentoft, L., "Convergence of the Least Squares Monte Carlo Approach to American Option Valuation." *Management Science*, vol. 50 (9), 2004, pp. 1193–1203.