Gas Storage Valuation Using a Monte Carlo Method

Alexander Boogert and Cyriel de Jong

Alexander Boogert is a senior quantitative analyst at Essent Energy Trading in 's-Hertogenbosch, The Netherlands. He is pursuing a Ph.D. degree in the Commodities Finance Centre at Birkbeck College, The University of London.
alexander.boogert@essent.nl

Cyriel de Jong is partner at Maycroft Consulting and affiliated with the Rotterdam School of Management, Erasmus University, The Netherlands.
dejong@maycroft.com

Developed countries increasingly rely on gas storage for security of supply. Widespread deregulation has created markets that help assign an objective value to existing and planned storages. Storage valuation is nevertheless a challenging task if we consider both the financial and physical aspects of storage. In this article, we develop a Monte Carlo valuation method, which can incorporate realistic gas price dynamics and complex physical constraints. Specifically, we extend the Least Squares Monte Carlo method for American options to storage valuation. We include numerical results and show ways to improve computational speed.

The natural gas industry has undergone market reforms in many countries around the world. As part of this process, the role of storage in the gas market has changed. Traditionally, storages were owned by utilities for balancing the variability in demand of their customers. As a result of the deregulation in the U.S. and Europe, the natural gas storage service is now unbundled from the sales and transportation services, meaning that storage is offered as a distinct, separately charged service. In combination with the development of active spot and futures markets, it has become possible to adjust storage trading decisions to price conditions.

Storage is ultimately needed to ensure security of supply. The seasonal demand for gas is traditionally linked to gas heating of houses, resulting in higher gas demand in winter than in summer. Other seasonal patterns are mainly related to the operation of gas-fired power plants. These plants are often designed for peak delivery (creating a day-night pattern in gas demand) and perform real-time balancing (creating an intraday pattern in gas demand). The flexibility in supply needed to match the demand can be delivered in several ways. Besides the usage of gas storages, one option is to regulate the output from gas fields to match the current demand. This can only be done as long as the fields are reasonably full and close to the gas grid. Finally, on a short-term basis line packing can be used. This means that the volume in the pipeline system is temporarily increased.

While gas demand in the U.S., Europe and Asia is growing year on year, the indigenous production flexibility is steadily falling. This explains the growing interest in investing in new gas storage facilities. The International Energy Agency [2004] estimates that the global underground storage capacity will double in the next 30 years (2000–2030), requiring an estimated annual investment in storages of between \$10 and \$20 billion. Additional flexibility in supply will come from the fast growing sector of liquified natural gas (LNG). The LNG-chain consists of liquefaction plants, LNG-ships, LNG-storage tanks and regasification plants. The combined investment in this flexible LNG supply chain is expected to be $25 to $40 billion.

These market developments call for accurate investment analysis methodologies, incorporating the various operating characteristics of storages and the random nature of natural gas prices. In this article we explore such a methodology, applicable both to real physical storage facilities as well as financial storage flexibility contracts. We will adapt existing option pricing methodologies to the more difficult problem of gas storage. The method is a generalization of the Least Squares Monte Carlo (LSM) approach for the valuation of American options described in Carrière [1996] and Longstaff and Schwartz [2001]. LSM is particularly suited to energy market applications, because one can separate the asset optimization (exercise strategy) from the price evolution model. The latter is often fairly complex and may be commodity specific due to the very volatile nature of energy markets. Simultaneously, LSM allows for the incorporation of various types of complex physical constraints into the exercise strategy.

In addition to facilitating the trading of gas storage, the development of new storage facilities forms an important application of the proposed methodology. Before the actual development, it should be decided how large the facility should be and how much to invest in compressors, which will, in turn, determine the injection and withdrawal rates. These decisions cannot be made without a clear understanding of the impact of operational characteristics on the valuation of storage.

Our approach clarifies how the storage value depends on operational constraints such as working volume (effective capacity) and injection and withdrawal rates (flexibility). Although the exact operational constraints depend on the storage under consideration, general settings may be distinguished, dependent on the four different types of storage (see Maragos [2002]): depleted gas and oil fields are often large but not so flexible; aquifers are typically small and flexible; salt caverns fall somewhat in between; and LNG-storages are very small and very flexible. Our approach also shows the dependence of the storage value on the dynamics of market prices, in particular, spot price volatility, mean reversion, and seasonality.

The issue of storage valuation and optimal operation is not limited to gas markets. Storages also play a significant balancing role in, for example, oil markets, soft commodity markets, and even electricity markets (through pumped-storage hydropower stations). The principles of our approach are applicable to those markets as well, as long as there is fairly liquid spot trading in the underlying commodity, and as long as the spot prices exhibit a fair amount of mean reversion, a condition we will clarify further on.

The article proceeds as follows. In the following section we review the literature, formally define a storage contract, and present our valuation approach. The next section is dedicated to numerical results. Then, we discuss computational issues and ways to improve performance. In the last section we conclude.

VALUATION

In this section we discuss the valuation of a storage contract. In the first two subsections, we review the available literature and make the choice for a spot-based Monte Carlo approach. In the next two subsections, we define the storage valuation contract and show its relationship to other options. Then, we introduce the valuation problem and explain our methodology. We end the section with theoretical and empirical convergence results which are available in the literature.

Forward-based Valuation

A storage operator owns the flexibility to inject and withdraw gas at any moment in the future. The task is to find the optimal operation (injection and withdrawal) of the storage, depending on current and expected gas prices. In short, two characteristics of gas prices allow a storage operator to maximize its value: predictable price movements (seasonality) and unpredictable price fluctuations (volatility). The basic approach to storage valuation is to calculate the optimal position given the available forward curve and take this position on the forward market. This *intrinsic value* approach captures the predictable seasonal pattern in gas prices and secures a sure profit. Additional value (extrinsic value) can be created by reacting to the price fluctuations on the spot market. A storage operator must thus choose between operating on the forward market (and perform a forward-based valuation) and operating on the spot market (and perform a spot-based valuation), or a combination of these.

In this article we follow the typical gas and power market convention that defines spot trades as day-ahead. In practice, a storage operator can switch operation from injection to withdrawal within just a few hours. This may create additional value based on decisions for each individual hour during the day. Although our method is also valid for hourly valuation, actual implementation is difficult due to (very) limited liquidity in the hourly market.

To value a storage contract, we follow a spot trading strategy, because it captures most of the contract's flexibility. This can be related to an observation from Maragos [2002] on the relationship between the storage value and the risk of the operating strategy employed. He notes that the value of storage increases with the amount of risk taken. Or alternatively, risk decreases with the number of hedges put in place. The spot-based valuation yields high return under high risk, while the rolling intrinsic valuation yields a medium return under very low risk. Because volatility and mean-reversion in gas spot markets are much larger than in gas forward markets, larger profits can potentially be made in the spot market.

How much additional extrinsic value above the intrinsic can be created using the spot-based approach depends on the market and operational parameters. In practice, most operators use a mixed strategy—after having reserved part of the storage for internal portfolio management, part of the remaining trading capacity is secured in the forward market and the remaining capacity is traded on the spot market. In the next section, we demonstrate through an example that the additional value of spot trading is often substantial. A trader should be aware, however, that an active spot trading strategy may involve more trading costs and that liquidity in the spot market may be more limited than in the forward market.

Before we continue, let us compare our approach to methods employing the forward or futures market. Gray and Khandelwal [2004] propose a rolling intrinsic approach. They propose that the holder captures the intrinsic value at the start of the contract. When new forward prices arrive, the holder calculates whether the profit of (partially) unwinding his position and taking the optimal position based on these new prices outweighs the transaction costs.

The rolling intrinsic strategy yields extrinsic value if a spread between different tradable months or quarters changes sign and if it makes sense to swap trading decisions. This strategy, though very safe, is not profit maximizing. First, prices from adjacent months and quarters, where a swap may make sense, are often strongly correlated. Second, because volatility in the forward market is limited, the magnitude of the swap change will typically remain small. Maragos [2002] describes a variant of this rolling intrinsic approach by adjusting only the spot trades to new information in the forward curve.

Spot-based valuation

For the valuation of the spot-based strategy, two approaches are popular—stochastic control and Monte Carlo. Examples of the stochastic control approach include Thompson, Davison, and Rasmussen [2003] and Weston [2002]. Their models are based on Bellman equations, which are solved by means of a finite difference method. In the stochastic control approach there is a direct link between the stochastic price process and the optimal strategy. In the Monte Carlo approach the two are separated. This means we can quickly experiment with different spot price processes, which is an advantage. Another advantage of the Monte Carlo approach is the ease to incorporate additional operational constraints. Such operational constraints often stem from the physical nature of the storage. Examples include protection against rock deformation, start-up and turning restrictions between injection and withdrawal, and volume-dependent injection and withdrawal rates.[1] A disadvantage is that the Monte Carlo approach is relatively slow. Computational issues will therefore deserve particular attention. We discuss computational issues later in the article.

Simulation techniques to price European-style options have been in use for quite some time and were introduced even before the binomial option pricing model (see Boyle [1977]). Simulations are especially attractive in situations with multiple stochastic factors. For quite some time, however, the pricing has been restricted to European, not American options. The first proposed solution for finding the optimal early exercise strategy was Tilley [1993]. Broadie and Detemple [2004] and Glasserman [2004] provide a survey of valuation methods for both European and American options. Most methods determine a lower bound on the option price based on the supremum over the payoff from all potential exercise strategies. Recently, however, a few articles have considered the dual problem of finding a positive-biased estimate for the option price (see Haugh and Kogan [2001]; Rogers [2002]; Andersen and Broadie [2004], and Meinshausen and Hambly [2004]). The downside is that this results in a higher computational load, which remains an issue in simulation-based methods.

In this article we have chosen to work with the Least Squares Monte Carlo (LSM) method. The LSM method has become a popular method to solve American option pricing by simulation since the article by Longstaff and Schwartz [2001]. Earlier variations of regression-based simulation were given by Carrière [1996] and Tsitsiklis and van Roy [2001]. To our knowledge, LSM has been applied to three different problems in the energy sector. De Jong and Walet [2003] value storage, but do not mention details on their implementation. Peterson and Gray [2004] and Tseng and Barz [2002] show an application in the valuation and optimal dispatching of power plants, whereas Ghiuvea [2001]; Keppo [2004], and Thanawalla [2005] value swing options.

Definition of a Storage Contract

We consider a financial storage contract from the perspective of the holder of the contract. We assume the contract is signed at time $t = 0$ and settled at time $t = T + 1$. The contract allows the holder to take a desired action at any discrete date $t = 1, \ldots, T$ after the spot price $S(t)$ is revealed. For convenience we will call the basic time unit a day and assume the spot price is known at the beginning of the day. Every day the holder of a storage contract can choose to inject gas, do nothing, or withdraw gas, within certain volumetric limitations. In the remainder of this subsection the contract is specified.

The holder of the contract faces different payoffs during the lifetime of the contract depending on his strategy. We take an injection at day t as a positive volume change, $\Delta v(t)$, and a withdrawal as a negative volume change, $\Delta v(t)$. Note that at time $t = 0$ the holder cannot take an action, and $\Delta v(0) := 0$, by definition. Positive volume changes have to be bought at the market and represent costs. Negative volume changes can be sold at the market and represent profits. We will assume the market has incorporated all knowledge about the optimal usage of storage into the prevailing forward curve and that the market will not be influenced by our trades. Note these points are of particular importance for a development decision where the impact of an additional storage facility on prices should be estimated.

We denote the (accumulated) volume in storage at the *start* of day t by $v(t)$. Thus, the holder can take his first action at day 1, which results in a volume $v(2)$ at the end of day 1, which equals the volume at the start of day 2. After noting $\Delta v(t) = v(t + 1) - v(t)$, we see the volume at any day t can be expressed in terms of $v(0)$ and the strategy followed of injecting and withdrawing:

$$v(t) := v(0) + \sum_{i=1}^{t} \Delta v(i-1) \tag{1}$$

We denote the payoff at day t by $h(S(t), \Delta v(t))$ and define for $t = 0, \ldots, T$:

$$h(S(t), \Delta v(t)) := \begin{cases} -c(S(t))\Delta v(t) & \text{inject at day } t \\ 0 & \text{do nothing at day } t \\ -p(S(t))\Delta v(t) & \text{withdraw at day } t \end{cases} \tag{2}$$

where $c(S(t))$ and $p(S(t))$ represent the cost of injection and profit of withdrawal. We allow these costs and profits to include both transaction costs a and bid–ask spreads b:

$$c(S(t)) := (1 + a_1)S(t) + b_1 \tag{3}$$

$$p(S(t)) := (1 - a_2)S(t) - b_2 \tag{4}$$

for certain pre-defined constants $a_1, a_2, b_1, b_2 \in \mathbf{R}_0^+$ such that $c(S(t)) \geq 0$ and $p(S(t)) \geq 0$. These conditions ensure that the payoff $h(S(t), \Delta v(t))$ has the opposite sign of $\Delta v(t)$. Note that in a world without transaction costs or bid–ask spreads, we will have $c(S(t)) = p(S(t)) = S(t)$. We denote the interest rate by δ.

We assume the contract is settled the day after the last trading date, that is at day $T + 1$. At settlement the holder receives a potential penalty, denoted by $q(S(T + 1), v(T + 1))$, which may depend on (the lack of) remaining gas in the storage, current price level, and so on.

We assume two volumetric limitations on the strategy are followed. First, the volume in storage must stay between a minimum level $v^{min}(t) \geq 0$ and a maximum level $v^{max}(t)$ for $t = 1, \ldots, T + 1$:

$$v^{min}(t) \leq v(t) \leq v^{max}(t) \tag{5}$$

Allowing $v^{min}(t)$ and $v^{max}(t)$ to be time dependent makes it possible to incorporate physical requirements. For example, to protect against rock deformations we could require a higher minimum level during a certain period in the year. In a traded storage contract, $v^{min}(t)$ and $v^{max}(t)$ are usually constant. For convenience we call the minimum allowed level the zero level, that is, $\min_t [v^{min}(t)] := 0$.

Second, the injection or withdrawal is limited per day. Let us write

$$i^{min}(t, v(t)) \leq \Delta v(t) \leq i^{max}(t, v(t)) \tag{6}$$

where $i^{min}(\cdot)$ and $i^{max}(\cdot)$ are pre-defined functions of the volume in storage $v(t)$ and time t. In practice, these functions are often constant with deviations at low and high volumes. This is due to the fact that for low storage volumes it is harder to withdraw gas, while for high storage volumes it is harder to inject gas. Note $i^{\min}(\cdot)$ will normally be negative and associated with withdrawal, whereas $i^{\max}(\cdot)$ will normally be positive and associated with injection.

Relation to American and Swing Options

The holder of a storage contract is faced with a timing problem of when to inject and when to withdraw. A similar timing problem can be found in an American option where the holder has to decide when to exercise his option.[2] Compared to an American option, the storage contract has the following special features:

- Storage holder must decide among multiple actions at all moments in time
- Storage holder has to comply with volumetric restrictions
- Storage contract can have payoffs at several moments in time
- Storage contract can have positive and negative payoffs

An American option is a special case of a storage contract. To see this, consider a storage with the following characteristics: $v(0) := 1$, $v^{\min}(t) := 0$, $v^{\max}(t) := 1$. Suppose there is no possibility to inject, that is, $i^{\max} := 0$, and we have to withdraw one unit if we decide to withdraw, that is, $i^{\min} := -1$. If we then set $a_2 := 0$, $b_2 := K$ in Equation (4) and refrain from a penalty at settlement, we are valuing an American option with strike K.

Another related option is the swing option. This option allows the holder multiple exercises (rights) during the contract. If the holder uses a right, he can decide whether he would like to get a positive or negative adjustment on the contracted volume. As in storage, there is often a penalty if the final accumulated volume adjustment deviates from zero. We can see a storage contract as a swing option with as many rights as exercise dates, but with a restriction on the accumulated adjustments.

Valuation of a Storage Contract

The value of a storage contract is the expected value of the accumulated future payoffs $h(S(t), \Delta v(t))$ under the most optimal strategy π. Thus, we need to consider the following pricing problem:

$$\sup_{\pi} \mathbf{E}\left[\sum_{t=0}^{T} e^{-\delta t} h(S(t), \Delta v(t)) + e^{-\delta(T+1)} q(S(T+1), v(T+1))\right] \tag{7}$$

This above expectation is assumed to be under a risk-neutral pricing measure. In complete markets this measure is unique, ensuring only one arbitrage-free price of the storage contract. To hedge a storage contract with basic securities other than a physical storage is not possible. From one perspective, we may conclude that the gas market is incomplete. From another perspective, one should remember that the purpose of physical storages and financial storage contracts is to cover demand variation, which helps to complete the market. Whatever the level of market completeness, in our valuation we will use risk-neutral pricing and assume there exists a risk premium to compensate the holders of the residual risk. Specifically, we will assume the risk premium is incorporated in the drift of the spot price process, as is a common approach in energy markets (see, e.g., Lucia and Schwartz [2002]).

As storage has volumetric restrictions, we explicitly incorporate dependence on volume into the optimal strategy $\pi = \{\pi(1, S(1), v(1)), \ldots, \pi(T, S(T), v(T))\}$ where $\pi(t, S(t), v(t))$ is the decision rule at day t with spot price $S(t)$ being at volume $v(t)$. The problem will be formulated as a dynamic program. Because the dynamic program will be solved backward in time, the holder does not know which volume levels will be visited. This means he has to find a decision rule at every time t for all possible volume levels.

For convenience, let us introduce two sets. The set of *allowed volume levels* at day t is denoted by $\mathcal{V}(t)$,

$$\mathcal{V}(t) := \{v \mid v^{\min}(t) \le v \le v^{\max}(t)\} \tag{8}$$

We denote the set of all *allowed actions* on day t being at volume $v(t)$ by $\mathcal{D}(t, v(t))$,

$$\mathcal{D}(t, v(t)) := \{\Delta v \mid v^{\min}(t+1) \le v(t) - \Delta v \le v^{\max}(t+1),\ i^{\min}(t, v(t)) \le \Delta v \le i^{\max}(t, v(t))\} \tag{9}$$

Let us denote the value of a storage contract starting on day t at volume level $v(t)$ by $U(t, S(t), v(t))$ and define the continuation value, $C(t, S(t), v(t), \Delta v)$, as the value we attach to the contract after taking an allowed action, that is, $\Delta v \in \mathcal{D}(t, v(t))$ as defined in Equation (9):

$$C(t, S(t), v(t), \Delta v) := \mathbf{E}\, [e^{-\delta} U(t + 1, S(t + 1), v(t) + \Delta v)] \tag{10}$$

Now the value $U(t, S(t), v(t))$ satisfies the following dynamic program:

$$U(T + 1, S(T + 1), v(T + 1)) = q(S(T + 1), v(T + 1)) \tag{11}$$

$$U(t,S(t),v(t)) = \max_{\Delta v \in \mathcal{D}(t,v(t))} \{h(S(t),\Delta v) + C(t,S(t),v(t),\Delta v)\} \tag{12}$$

for all t. Remember, we assume our contract is settled at $T + 1$.

The dynamic program states that the holder of the contract weighs different actions against one another on the basis of their direct payoff plus expected future payoff. In general, an action with a high direct payoff (selling gas) has a lower expected future payoff (lower inventory level). In order to make a decision, the holder needs to consider the continuation values $C(t, S(t), v(t), \Delta v)$ for all allowed actions $\Delta v \in \mathcal{D}(t, v(t))$.

This dynamic program is actually closely related to the dynamic program of an American option. If we interpret continuation value as the value of not exercising the option at day t, assuming it has not been exercised before, we find the usual dynamic program for the American option for $t = 1, \ldots, T - 1$ (normally the American option is settled at the last trading date, $t = T$):

$$U(T, S(T)) = h(S(T)) \tag{13}$$

$$U(t, S(t)) = \max\{h(S(t)), C(t, S(t))\} \tag{14}$$

One way to find an approximate solution for this dynamic program is by regression-based simulation. From the available regression-based methods, we use the Least Squares Monte Carlo (LSM) method. Underlying reasons for this choice and some related literature were discussed previously.

In the next subsection we discuss our extension of the LSM method for American options to a storage contract. The resulting pricing algorithm is presented at the end. We refer to Longstaff and Schwartz [2001] for a simple numerical example of an American put. Another example is brought forward by Moreno and Navas [2003], which emphasizes the possible impact of having too few simulation paths. In their example, the price of the American option becomes lower than the European option. This pitfall can be avoided with enough simulation paths, which requires an efficient pricing algorithm as we will discuss.

Least Squares Monte Carlo

According to the dynamic program defined by Equations (11) and (12), the holder of the storage contract has to calculate continuation values for all t for all possible actions $\Delta v \in \mathcal{D}(t, v(t))$ for allowed volume levels $v(t) \in \mathcal{V}(t)$. The assumption behind regression-based techniques is that one can approximate continuation values by a (finite) linear combination of known basis functions $\phi_q(t, S(t), \mathcal{V}(t), \Delta v)$ of the current state. This means we approximate

$$C(t, S(t), v(t), \Delta v) \approx \sum_{q=1}^{Q} \phi_q(t, S(t), v(t), \Delta v)\beta_{q,t} \tag{15}$$

for certain constants $\beta_{q,t} \in \mathbf{R}$ and $Q \in \mathbf{N}$. It is a priori not clear which value to assume for Q and from which family of basis functions to choose $\phi_q(t, S(t), v(t), \Delta v)$. The original method works as follows: first, the coefficients in this expansion are estimated by a least-squares regression of known continuation values on the current state variables. Then, a continuation value is approximated by substituting the regression coefficients and a decision rule can be determined.

According to our definition, the continuation value is dependent on time t, spot price $S(t)$, volume level $v(t)$, and volume change Δv. Before we discuss the LSM method in more detail, let us show our approach to reduce the dimensionality of the problem towards a point that the basis functions need only depend on $S(t)$.

A first observation is that we can reduce the dimensionality from 4 to 3, because the dependency of continuation value is actually on the *sum* of volume level and volume change, instead of volume level *and* volume change. To see the improvement, suppose we are at volume level k and consider decision $\Delta v = 0$ so that we need to approximate $C(t, S(t), k, 0)$. Now if we are at volume level $k - \alpha$ and consider decision $\Delta v = \alpha$, we have to calculate $C(t, S(t), k - \alpha, \alpha)$. Because these two actions both reach volume level k at time $t + 1$, they will lead to the same cash flows later. That means that

$$C(t, S(t), k, 0) = C(t, S(t), k - \alpha, \alpha) \quad (16)$$

For this reason, we can write continuation values as a three-dimensional vector: $C(t, S(t), v(t) + \Delta v)$ or, equivalently, $C(t, S(t), v(t + 1))$.

In the LSM method it is possible to include the time unit in the regressions, but it is more customary and in line with the original method of Longstaff and Schwartz to run separate regressions for each time unit. Similar to the time variable, one has the option to include the volume level $v(t + 1)$ in the regression together with the price, $S(t)$. We tested various ways to incorporate both volume and price in bivariate basis functions, but found that even with high-order polynomials, convergence results were unsatisfactory. In many different types of empirically tested storage problems, the relationship between volume, price, and continuation value appeared to be not very smooth. Hence, it is better to run a separate regression per volume level. For this reason we discretized the volume into $n = 1, \ldots, N - 1$ units of fixed width α. We estimate separately a different set of regression parameters for each volume point $v(t + 1; n) := (n - 1)\alpha$. This leads to an estimate for the continuation value for all discrete volume points, denoted by $C(t, S(t), v(t + 1; n))$.

Thus we have reduced the dimension of the basis to one. This means the only stochastic part involved is the spot price and we can write $\phi(S(t))$, similar to Longstaff and Schwartz. Of course, the precise functional form of $\phi(S(t))$ remains a discussion and we will come back to this issue in the next section.

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

$$C^b(t, S^b(t), v(t+1; n)) \approx e^{-\delta} Y^b(t+1, S^b(t+1), v(t+1; n)) \quad (17)$$

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 $S^b(t + 1)$.[3] With approximation (17), 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 (15), we get an approximation $\hat{C}^b(t, S^b(t), v(t; n) + \Delta v)$ of the continuation value for all volume points $v(t; n)$.

Later we will discuss a more efficient approach for finding an estimate of the continuation value for all allowed volume levels for all allowed actions. Right now, we will assume that we discretize both volume and volume actions into a fine grid, which will enable us to approximate the continuation value.

With an approximation of the continuation value for all allowed actions, Δv, we can determine a decision $\hat{\pi}^b(t, S^b(t), v(t))$, for all allowed volume levels:

$$\hat{\pi}^b(t, S^b(t), v(t)) = \arg \max_{\Delta v \in \mathcal{D}(t, v(t))} \left\{ h(S^b(t), \Delta v) + \hat{C}^b(t, S^b(t), v(t) + \Delta v) \right\} \quad (18)$$

This means we make the decision that we expect will have the highest accumulated payoff. Now we are able to approximate the value of future accumulated cash flows, $\hat{Y}^b(t + 1, S^b(t + 1), v(t + 1))$:

$$\begin{aligned} &\hat{Y}^b(t+1, S^b(t+1), v(t+1)) \\ &\quad = h(S^b(t+1), \hat{\pi}^b(t+1, S^b(t+1), v(t+1))) \\ &\qquad + e^{-\delta} Y^b(t+2, S^b(t+2), v(t+2) \\ &\qquad + \hat{\pi}^b(t+1, S^b(t+1), v(t+1))) \end{aligned} \quad (19)$$

This procedure is started at day $T + 1$ where the continuation value is zero or equal to a penalty function, $q(S^b(T + 1), v(T + 1))$, which may depend on (the lack of) remaining gas in storage, current price level, and so on.

$$Y^b(T+1, S^b(T+1), v(T+1)) = q(S^b(T+1), v(T+1)) \quad (20)$$

Then, stepping backwards, we determine a decision rule for all points in the time-volume grid. Finally, at the valuation date, that is $t = 0$, we set the value of the storage contract equal to the average value of the future accumulated cash flows. That is,

$$\hat{U}(0,S(0),v(0)) = \frac{1}{M}\sum_{b=1}^{M} Y^b(1,S^b(1),v(0)) \qquad (21)$$

Before we continue, let us summarize our pricing algorithm:

1. Simulate M independent price paths $S^b(1), \ldots, S^b(T+1)$ for $b = 1, \ldots, M$ starting at given $S(0)$
2. Assign a value to the contract at maturity according to Equation (20)
3. Apply backward induction for $t = T, \ldots, 1$. For each t, step over N allowed volume levels $v(t; n) \in \mathcal{V}(t)$:
   - Run an OLS regression to find an approximation of the continuation value $\hat{C}^b(t, S^b(t), v(t+1; n))$ according to Approximation (15)
   - Combine the different continuation values, $\hat{C}$, into a decision rule, $\hat{\pi}^b(t, S^b(t), v(t))$, according to Equation (18)
   - Implement the decision rule to calculate the accumulated future cash flows, $Y^b(t, S^b(t), v(t))$, according to Equation (19)
4. Storage value is the average accumulated future cash flow over all price paths according to Equation (21)

Convergence Results in the Literature

A natural question in a simulation setting is whether the algorithm converges and whether it converges to the correct value. In this subsection, we provide the convergence results available in the literature for the LSM method. Because the underlying in our storage contract is the spot price, the available results on the LSM method with one underlying are of specific interest. Our conclusion is that the literature indicates that the LSM method converges, although it is worth checking the results of using different basis functions and different numbers of paths in the LSM method. Our own numerical tests will be illustrated in the next section.

The theoretical convergence of the LSM method for the American option has been discussed in an asymptotic sense in three different articles. Longstaff and Schwartz [2001, Proposition 1] proposed to increase the number of basis functions until the option price does not increase any more for a fixed number of paths. This was extended by Clément, Lamberton, and Protter [2002], who note that the algorithm converges when both the number of paths and the number of basis functions go to infinity. The resulting price distribution will then be Gaussian. Recently, Stentoft [2004b, Theorum 1] noted that the smoothness of the payoff also plays a role. More precisely, he claims that, assuming some smoothness in the payoff function, one must increase both the number of paths and the number of basis functions; let there be Q basis functions and M paths, so that if $Q \to \infty$ and $Q^3/M \to 0$ then the price estimate converges to the true price in a mean-square sense.

From a numerical perspective, testing has been performed for many different options. Besides the American put, Longstaff and Schwartz showed the application of the LSM method for more complicated payoff structures like American-Bermuda-Asian options, cancelable index amortizing swaps (an interest rate product), American options on jump diffusions, swaptions, and a maximum option. They reached similar results as for the American put: for an increasing amount of basis functions the option price converges. They also mention that the results are robust to the choice of basis functions.

These numerical claims were checked by Moreno and Navas [2003] and Stentoft [2004a]. Moreno and Navas use ten different polynomials (power, Legendre, Laguerre, two types of Hermite, and five of types Chebyshev) to price an American put. Their results show that Proposition 1 of Longstaff and Schwartz can be violated in numerical examples, but it typically holds. Stentoft [2004a] finds a similar numerical counterexample.

Concerning the normality of the resulting price distribution, Kircher [2004] has a remarkable finding for the American put option. He found asymmetric fat tails when, for example, a fifth order polynomial is used. Moreover, "asymmetric fat tails of the distribution appear only if more than four terms are included in the regression function. This effect remains when more paths are included in the regression" (Kircher [2004], p. 31).

EMPIRICAL RESULTS

In this section we discuss numerical results of the pricing algorithm, implemented in Matlab. After introducing our spot price process, we clarify the impact of market dynamics and operational characteristics on storage value. At the end, we demonstrate empirically the convergence of the pricing algorithm.

Spot Price Simulation

Because building gas storage facilities requires considerable investment, gas storage is in limited supply and costly to rent. Consequently, sometimes supply has more difficulty in matching demand than at other times, exemplified in seasonal patterns and mean reversion in spot prices. We show the standard process that captures these dynamics. It will be used for the simulation of daily gas prices in the rest of this section, and it can easily be extended to more complex processes.

We simulate M independent price paths according to the one-factor Schwartz [1997] model with time-dependent drift where we use antithetic sampling (see, e.g., Glasserman [2004], Chapter 4.2). A natural extension would be to consider a two-factor model like Schwartz and Smith [2000]. Our price process is given by

$$\frac{dS(t)}{S(t)} = \kappa[\mu(t) - \ln S(t)]dt + \sigma dW(t) \tag{22}$$

Note the long term level $\mu(t)$ is a time-varying function, whereas the mean reversion rate, κ, and volatility, σ, are assumed constants. Using $X(t) := \ln S(t)$, one rewrites Equation (22) into

$$dX(t) = \kappa\left[\mu(t) - X(t) - \frac{\sigma^2}{2\kappa}\right]dt + \sigma dW(t) \tag{23}$$

Standard Storage Contract

As an example, we will use a storage contract on a small salt cavern connected to the Dutch gas system with characteristics given in Exhibit 1.

For simplicity, we have decided not to take into account a bid–ask spread or interest rates; that is, we take $a_1 = a_2 = b_1 = b_2 = \delta = 0$ in Equations (3) and (4). For this example we set the penalty function high enough to ensure that all exercise strategies will terminate at the desired end-volume. The valuation has been processed with the following discretization: $\alpha = 2500$, $N = 101$, $T = 365$. As a standard we use the first three power basis functions: $\{1, x, x^2, x^3\}$.

We simulate $M = 500$ price paths based on a forward curve of the Dutch TTF market at June 10, 2005. On that day, the highest market forward price of 25.44 €/MWh

EXHIBIT 1
Characteristics of the Standard Storage Contract

The volume can also be measured in m³ = 0.01076 MWh.

Nomination	daily
Time to maturity	$T = 1$ year
Start date	07/01/2005
End date	06/30/2006
Min. volume	$v^{\min}(t) = 0$ MWh for all t
Max. volume	$v^{\max}(t) = 250{,}000$ MWh for all t
Start volume	$v(0) = 100{,}000$ MWh
End volume	$v(T) = 100{,}000$ MWh
Max. withdrawal	$i^{\min}(t, v(t)) = \max(v^{\min}(t) - v(t); -7{,}500)$ MWh/day for all t, $v(t)$
Max. injection	$i^{\max}(t, v(t)) = \min(v^{\max}(t) - v(t); 2{,}500)$ MWh/day for all t, $v(t)$

was for February 2006, whereas the lowest price of 14.88 €/MWh was for July 2005. For the simulation, we transformed the available monthly and quarterly contracts to a daily forward curve based on a fixed weekly pattern creating a week–weekend spread.

The estimation of the process parameters, σ and κ, are not the focus of this article. We initially study two representative parameter cases, both with a daily mean-reversion rate of $\kappa = 0.05$. The high volatility case has $\sigma = 9.45\%$, whereas the low volatility case has $\sigma = 3.15\%$. Later, we vary these parameters further to better clarify their impact.

We have previously introduced the concepts of intrinsic and extrinsic value. The intrinsic value is the value that can be captured with the current forward curve. The full storage value also incorporates the extrinsic value, which monetizes the volatility in the gas spot prices. The intrinsic value of our storage is €2.7 million, primarily based on a single full cycle over the year (see the middle line in Exhibit 2) at an average spread of slightly less than €11. This value is achieved by first injecting in July and August 2005 at an average price of €15.01 and a total cost of €2.2 million. Then, until January 12, 2006, the storage level is kept almost constant at its maximum of 250,000 MWh. It takes until March 8 to fully empty the storage at an average sales price of €25.83 because on lower priced weekends extra gas is injected. Total revenue in this period is €6.5 million. Finally, with some temporary variations, gas is re-injected to the required level of 100,000 MWh by the end of June at an average price of €15.73.

The intrinsic value assumes no random variation in spot prices. With a fairly low daily volatility of 3.15%, storage value increases by 15% to a total of €3.1 million. Nevertheless, the storage's flexibility can be really exploited in the high volatility case of 9.45%, which doubles the storage value, compared to the intrinsic value, to

Exhibit 2

Minimum and Maximum Attained Volume

Minimum and maximum attained volume are calculated using both the extrinsic approach (solid) and the intrinsic approach (dashed) for the low and high volatility cases from a single run. Note that the intrinsic path and the maximum volume of the extrinsic approach coincide in the left side of the graphs.

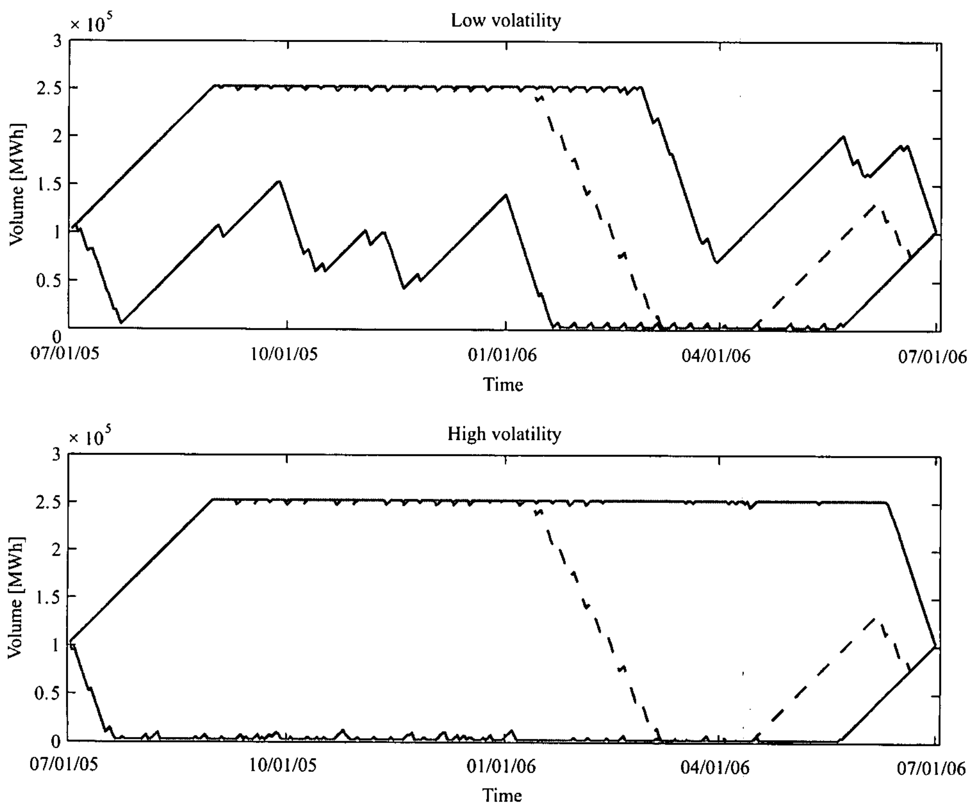

EXHIBIT 3
Value Distribution from Extrinsic Approach with Intrinsic Value

Values are for the low and high volatility case from a single run. Vertical lines indicate the intrinsic value (€2.7 million), the mean in the low volatility case (€3.1 million) and the mean in the high volatility case (€5.4 million).

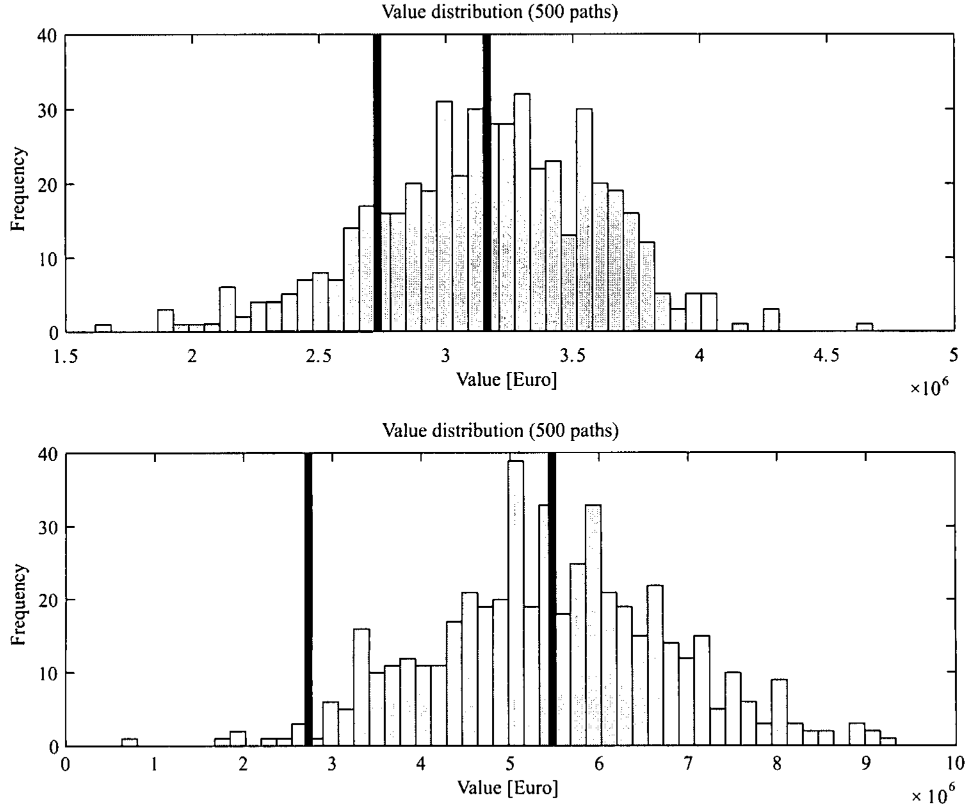

€5.4 million. Apparently, the extrinsic value varies greatly with the spot market volatility. In a volatile market, the seasonal component is no longer the major price driver and a flexible storage can benefit from this volatility.

The volume and price distributions from a single run of 500 simulations are shown in Exhibit 2 and 3. We see the high volatility case has a wider price distribution, which translates into a higher storage value. The volume distribution in the high volatility case is a large band occupying almost all possible volumes. In the low volatility case, the seasonal spread is the major driving factor for the gas price. We observe therefore that the realized volumes in the low volatility case follow more closely the intrinsic volume path. Another issue worth noting is the weekday–weekend pattern in the volume paths: the exploitation of lower weekend prices leads to saw-tooth volume paths. As remarked by De Jong and Walet [2003], the expected volume distribution can serve as an input to the transport booking.

Impact of Market Dynamics

The difference between the high and low volatility cases already hinted at the impact of market dynamics on storage value. It is interesting to analyze this further. In Exhibit 4 we show how the value changes with mean reversion and volatility, the two main drivers of gas spot prices. As expected, the values increase with volatility; the relationship is almost linear.

With respect to mean reversion, two effects play a role in opposite directions. On the one hand, a higher mean reversion makes price movements more predictable and the successful timing of purchases and sales easier. This effect raises storage value. On the other hand, a higher mean reversion ensures that prices are pulled back faster to an average level. This limits the possibility of large price swings and decreases storage value. Exhibit 4 shows that the second effect dominates for small mean-reversion rates, but the first effect takes over at an inflexion

EXHIBIT 4

Impact of Changing Mean Reversion (top) and Volatility (bottom) Parameters

The vertical lines indicate the parameter values in the high volatility case.

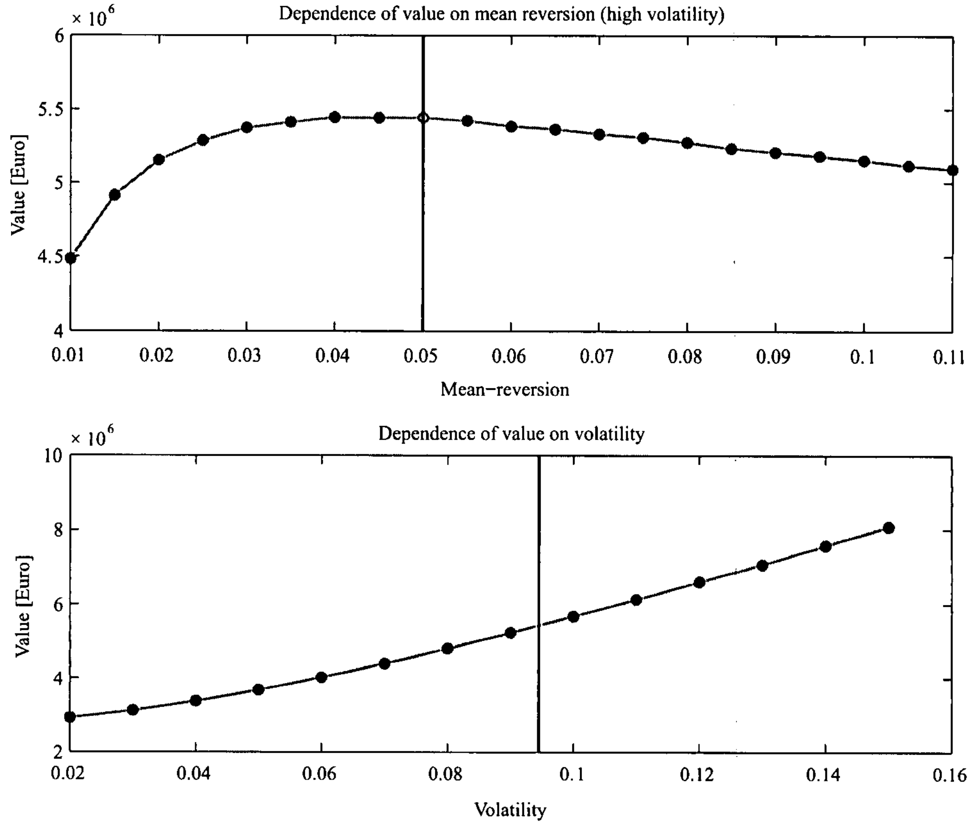

point of around $\kappa = 0.04$. We verified that this inflexion point lies further to the right for more flexible storages (with high injection and withdrawal rates), because they especially exploit small short-lived price swings. In general, within a reasonable range for the mean-reversion rate of around 0.02 to 0.11 our storage value is fairly stable. We therefore conclude that volatility often has a larger impact than mean reversion.

Impact of Operational Characteristics

Besides the market dynamics, the operational characteristics can dramatically change the value of storage. In Exhibit 5, we show a sensitivity analysis of the working volume, injection, and withdrawal rates, changing one parameter at a time. The top panel demonstrates maximum withdrawal rates ranging from –1,000 to –20,000 MWh/day. The middle panel shows the maximum injection rates from 500 to 15,000 MWh/day. Finally, for the bottom panel we changed the working volume from 100,000 to 700,000 MWh. The general pattern is that storage value increases most strongly if flexibility (injection or withdrawal) is raised from very low levels, but less so if it is raised from higher levels. When the working volume is raised we observe a similar flattening off of the value increase. At some point, somewhere above 600,000 MWh, the full working volume cannot even be cycled in a year with the available injection and withdrawal rates, and higher working volumes have no effect on storage value.

We conclude that the magnitude of the marginal change in storage value depends on the current characteristics of the storage. Changing the factor that is most binding has the best payoff. In our case, for example, raising the injection rate with an extra MWh per day has a far larger marginal effect than raising the withdrawal rate or working volume with an extra MWh. These marginal changes should be compared with the additional investment costs to find the optimal expansion decision.

EXHIBIT 5
Impact of Changing Operational Parameters in the High-Volatility Case

The top graph shows the impact of changing the withdrawal rate, the middle graph shows the impact of changing the injection rate, and the bottom graph shows the impact of changing the maximum volume. The vertical lines indicate the parameter value in the standard storage contract.

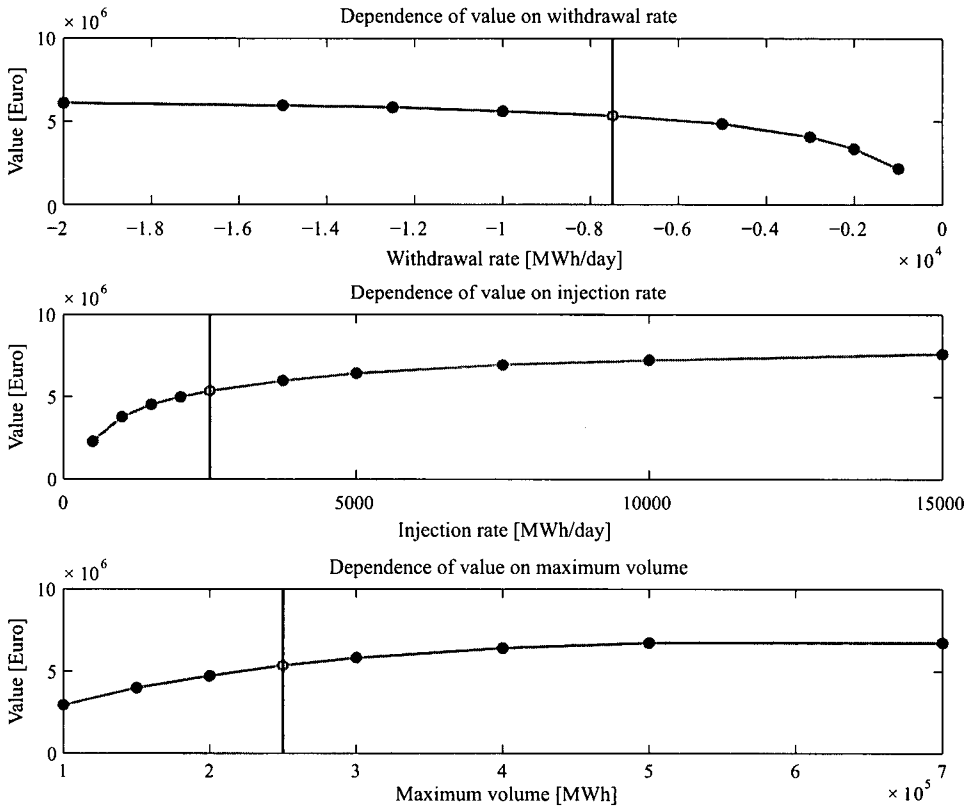

Convergence, Normality, and the Empirical Error

In this and the following subsection, we empirically test the convergence of the proposed pricing algorithm. In this subsection we consider the normality of the price distribution. According to the asymptotic results by Clément et al. [2001], it should be a normal distribution. According to Kircher [2004], asymmetric fat tails can occur if we use more than four basis functions. The normality of the price distribution is important for the determination of a confidence interval. We calculate the standard deviation as the standard error of the mean values from ten different runs.

Upon initial inspection, the price distributions in Exhibit 3 (both from a single run) appear normally distributed. The Lilliefors test is a statistical test of whether data is normally distributed with unknown mean and variance and was used in this instance. We find that at a 5% level we cannot reject the hypothesis of normality for the high volatility case. Increasing the simulation to 1,000 instead of 500 paths, we find we cannot reject the hypothesis for either volatility case.

Impact of Regression Settings

In the regression method, it is not clear a priori which kind and how many basis functions to choose, and how many paths to use. In this subsection we present the results based on choosing from three different families: power, weighted Laguerre, and b-splines.[4] The first two are polynomial families, the last one joins piecewise cubic polynomial segments. An alternative is to regularize the fit. This can be done by adding a penalizing term to control the smoothness of the fit, often the second derivative. Such an approach was used by Thanawalla [2005] for

EXHIBIT 6

Test Results for Different Basis Functions

The impact of different families of basis functions (power P and b-splines b), different number of basis functions, and different number of paths (50, 500 or 5000). For example, P_4 indicates the power basis $\{1, x, x^2, x^3, x^4\}$. Regressions for a specific number of paths use the same set of simulated paths. We present the high volatility case and calculate the mean and standard deviation on ten runs. Running time indicates the average running time of the program over the ten runs.

Basis	Number of paths	Mean	Std	Running time (sec)
P_2	50	5469006	67536	1.0
	500	5496365	22884	3.8
	5000	5502115	5341	40.9
P_3	50	5449743	68298	1.0
	500	5433701	24246	3.8
	5000	5436252	3368	41.2
P_4	50	5463046	70319	1.1
	500	5414859	24213	3.8
	5000	5412658	4294	41.4
P_5	50	5473720	69226	1.1
	500	5401761	22241	3.9
	5000	5397023	4171	41.7
b_5	50	5466878	71008	3.4
	500	5414515	24477	5.7
	5000	5413977	4295	66.1
b_6	50	5474144	69125	7.0
	500	5403416	21749	5.9
	5000	5398939	4462	68.7

swing options. Ramsay (see Endnote 4) offers the possibility of smoothing the fit to become monotonic.

During the implementation we encountered numerical instabilities for the weighted Laguerre. We found poor fits and even values well below the intrinsic value. For this reason we do not report the values for the weighted Laguerre. The implementation also produced numerical instabilities when increasing the number of basis functions to seven in both the power and the b-splines regression.

The results of tests for the power and b-splines can be found in Exhibit 6. The exhibit presents the mean and standard deviation based on ten runs and the total running time for different basis functions, different numbers of paths and different numbers of basis functions. To enable a correct comparison, the simulations are fixed per number of paths (e.g., one set of 10 * 500 simulations). We verified that our b-splines, with no interior knots and the power basis $\{1, x, x^2, x^3\}$, yield the same fit. The presented numbers are calculated using in-sample valuation as we will present in the next section.

In Exhibit 6 we see a fast increment in time with the number of paths while the empirical error is decreasing. The mean values are surprisingly stable between the power and b-splines for the different numbers of paths and different number of basis functions (minimum of 5,397,023 and maximum of 5,502,115). From the test we conclude the storage valuation is quite robust to the regression settings.

In comparison to the literature, our results are strikingly good given the number of paths we use. For example, Longstaff and Schwartz used 100,000 paths to price an American put. By changing the parameters, we found this phenomenon was not due to the specific parameter setting in this example. We have no full explanation for this behavior, but believe it is partly attributable to the reduction in dimension described earlier, which makes the pricing algorithm relatively precise. Another part can be attributed to the high mean-reversion rate in the gas market. As a result, the price simulations do not move far off the initial forward curve. Especially with a one-factor process, which essentially assumes that only the short end of the curve is variable, the price distribution is not so wide. Internal testing has shown that a two-factor price process requires a higher number of paths, though 5,000 is certainly still enough. Increasing the number of paths to 50,000 did not really impact the value.

IMPROVING COMPUTATIONAL EFFICIENCY

The valuation of a storage contract using a Monte Carlo method is a computationally demanding task. We have seen calculation time exploding with the number of paths and the chosen discretization in time or volume, limiting the number of cases to be studied. Also, in order to incorporate more physical constraints, an efficient program is essential. In this section, we discuss two ways to improve the performance of the program: use in-sample valuation and reformulate the pricing algorithm.

In the literature, more ways are available to improve the speed of convergence, especially in regard to price simulation. Glasserman [2004] offers an overview of variance reduction techniques, and Stentoft [2004a] finds that antithetic sampling does not always outperform the standard simulation.

In-Sample Valuation

An interesting diagnostic test to study the convergence of the simulation algorithm is to compare in-sample and out-of-sample valuation. The simulation algorithm is said to perform well if the two valuations are close to each other. The pricing algorithm described in the previous section is an in-sample valuation. The out-of-sample valuation follows from implementing the decision rules on a new set of simulated price paths. Thus, the out-of-sample valuation contains both a backward and forward induction, while the in-sample valuation contains only a backward induction. In mathematical terms we have (for comparison reasons we replicate Equation (21))

$$U_{in} := \hat{U}(0,S(0),v(0)) = \frac{1}{M}\sum_{b=1}^{M} Y^b(1,S^b(1),v(0)) \quad (24)$$

$$U_{out} := \frac{1}{M}\sum_{b=1}^{M}\sum_{t=1}^{T+1} h(S^b(t),\hat{\pi}(t,S^b(t),v(t))) \quad (25)$$

Longstaff and Schwartz ([2001], Table 2) found, in the case of the American put, positive and negative differences between the in-sample and out-of-sample values, and that only 5% of the values are larger than two standard errors. We found no references in the literature for swing or storage showing both in-sample and out-of-sample valuations.

In the test we compare the in-sample and out-of-sample values for five runs in the low volatility and five runs in the high volatility case. The results are shown in Exhibit 7. In the exhibit we see both positive and negative differences with a maximum relative difference of 1.51%. In this test we used the usual first three power basis functions.

From this test we see that the resulting price distributions from an in-sample and out-of-sample valuation are similar and have approximately the same mean. Empirical tests showed that an increasing number of simulations

EXHIBIT 7
In-Sample and Out-of-Sample Values

Values are for five runs in the high volatility case and five runs in the low volatility case and their relative difference.

	High			Low		
Run	In	Out	Rel. diff.	In	Out	Rel. diff.
1	5428806	5367698	1.14%	3162855	3170827	-0.25%
2	5459147	5377910	1.51%	3183351	3164578	0.59%
3	5445448	5382609	1.17%	3166992	3158039	0.28%
4	5445076	5454934	-0.18%	3171856	3157683	0.45%
5	5387498	5341233	0.87%	3168123	3160245	0.25%

decreases the difference between an in-sample and out-of-sample valuation.

The out-of-sample valuation has an extra computational demand on top of the in-sample valuation in terms of memory usage and calculation time. For the out-of-sample valuation, during backward induction we have to store a decision rule for all volume grid points and all time periods. This means we store a three-dimensional matrix of regression parameters and use it during forward induction to obtain the out-of-sample value with a new set of price simulations. For the in-sample valuation, it suffices to store the regression parameters for a specific grid point until the valuation for that grid point has been done, after which we can forget it.[5]

It is worth noticing that we can avoid three-dimensional matrices entirely if we run an in-sample valuation. It is enough to note that the pricing algorithm requires the accumulated future cash flows and not all individual cash flows. Therefore, it is sufficient to discount during the backward induction, leaving out the time dimension.

Thus we can observe that the use of three-dimensional matrices is directly connected to a discussion about out-of-sample versus in-sample valuation. We conclude that in-sample valuation saves calculation time and the regression parameters and the intermediate cash-flows do not need to be stored. The only negative aspect of in-sample valuation is that we are no longer able to get insight into the possible volume paths. For valuation purposes the possible volume paths are often not relevant, though for managing physical storages they might be important. In those cases we advise running an in-sample valuation (based upon a large number of paths) for valuation purposes with an out-of-sample valuation (based upon a small number of paths) to create a rough volume distribution. All numbers presented in this article are based upon in-sample valuation.

Reformulate Pricing Algorithm

We have argued that it is more efficient to run separate regressions per time unit and volume point than include time and volume in the basis functions of the regressions. We also showed that continuation value depends on the *sum* of current volume and volume change, rather than on either alone. All this leads to basis functions that depend only on $S(t)$.

Discretization of both the volume level and the possible actions was one way to reduce the dimensionality. But in order to incorporate all volume levels that can possibly be attained, a very fine grid of volume points may be required. We perform an interpolation as a solution to limit the size of the grid and at the same time obtain accurate estimates of continuation values. More precisely, the expected continuation values for intermediate volume levels are calculated as the distance weighted average of the expected continuation values at adjacent volume points, denoted by $v(t+1; n^*)$ and $v(t+1; n^*+1)$. In mathematical terms we can write,

$$n^* := \sup\left\{n \in \{1, \ldots, N\} \middle| (n-1)\alpha \leq v(t+1)\right\} \tag{26}$$

$$w^* := \frac{v(t+1) - v(t+1; n^*)}{\alpha} \tag{27}$$

$$\hat{C}(t, S(t), v(t+1)) := w^*\hat{C}(t, S(t), v(t+1;\, n^*)) + (1-w^*)\hat{C}(t, S(t), v(t+1;\, n^*+1)) \tag{28}$$

This interpolation is especially time saving when the working volume is not a nice multiple of injection and withdrawal rates, which is quite likely in practice. It is even almost inevitable when injection and withdrawal rates are time dependent and/or volume dependent.

The remaining step in the process is to consider how to treat the different possible actions. According to the definition of the decision rule in Equation (18), we should choose the maximum over $\Delta v \in \mathcal{D}(t, v(t))$. One wonders if the optimal decision is to always choose among maximum possible injection, maximum possible withdrawal, or no action. From a computational perspective it is certainly beneficial to be able to consider only three actions and leave out all intermediate possibilities. Such an approach was chosen by De Jong and Walet [2003]. In our experiments we have seen this is most often, but not always, the case. As an example of how this strategy is suboptimal, consider a storage where the working volume is not a multiple of the maximum injection and withdrawal rate. In this case the optimization will miss the opportunity to inject less than the maximum even though the volume is not close to the maximum.

CONCLUSION

We have shown in this article how to value storage contracts using a Monte Carlo method. Our proposed solution method is a generalization of the Least Squares Monte Carlo method proposed by other authors for American options. In gas storage valuation and optimization, the current volume level plays a central role. The level of this continuous variable determines which actions are allowed and makes the problem more complex than standard American puts or standard swing options. We propose a pricing algorithm which incorporates a number of solutions to efficiently deal with this volume variable and the various constraints.

We show empirically that this pricing algorithm converges well compared to in-sample and out-of-sample valuations, but also by using different basis functions and a different number of paths. We observe that as few as 50 simulations often suffice to get a precise storage value. This is in stark contrast to other research where up to 100,000 simulations are mentioned as a requirement for low variance in the option value. We have no full explanation for this behavior. We partially attribute this remarkable result to the strong mean reversion in gas spot markets, which limits the variance of the price distribution. However, we also attribute this to our effective pricing algorithm and the one-factor price model.

Our method assumes a spot trading strategy and a one-factor price process. It is therefore most suitable for valuations with a horizon of only a few years. From a financial perspective, it would be especially valuable to adapt the methodology for long-term valuations. This can be achieved with the same valuation method, but a different price simulation. The new simulations would have to come from a multifactor price model which also describes uncertainty in long-term price level, summer–winter spreads, and interest rates. Such a model would further support the decision-making process in new gas storage investments, in which many billions of dollars will be invested over the next couple of years.

ENDNOTES

The authors want to thank the following people for their input to this article: Alena Musilová, Rutang Thanawalla, Jan Willem Weenink, Kasper Walet, Hans van Dijken, Menke Ubbens, Gido Brouns, Manfred Bartels, Dominique Dupont, the various users of the MayStore storage tool, and participants in the FELab Seminar and the Modelling and Measuring Energy Risk 2005 Conference. The information presented in this article does not necessarily reflect the views of Essent Energy Trading.

[1]The walls of a physical storage are flexible. When pressure becomes low, these walls will move, and rock deformation occurs. See, for example, DeVries [2003] for a descriptive model of this process. Such models provide guidelines for the actual operation of the gas storage. In particular, they may indicate that the inventory level (pressure) has to stay above a minimum average level most of the time, and may never fall below another minimum level. This creates additional optimization and valuation problems.

[2]It is good to realize that the terms *exercise*, *stopping time*, and *strike*, which are common in the literature for American and swing options, are not easily transferred to storage contracts. The holder of the storage contract can take actions and does not actually stop anything. Therefore, a decision rule is a more appropriate concept than a stopping time. As well, the holder compares the expected payoff of an action to the expected payoffs of the other allowed actions, and not to a fixed strike like the holder of an American or swing option. Continuation values are thus an appropriate concept.

[3]To improve the convergence speed, Longstaff and Schwartz [2001] propose using only in-the-money paths in the regression. This recommendation cannot be applied here, because a storage will naturally generate negative payoffs during injection.

[4]The b-spline regression is performed with software provided by Jim Ramsay on his website: ftp://ego.psych.mcgill.ca/pub/ramsay/. We use the functions "create_bspline_basis" and "data2fd" with $order = 4$ and $nbasis = [4,5,6]$. This implies we have 0, 1, and 2 internal knots.

[5]Another way to reduce the size of the decision matrix is to store only the switch points. With b possible actions in a point, while assuming economic explainable results, it would suffce to store only b numbers for each time-volume combination.

REFERENCES

Andersen, L., and M. Broadie. "A Primal-Dual Simulation Algorithm for Pricing Multi-Dimensional American Options." *Management Science*, Vol. 50, No. 9 (2004), pp. 1222–1234.

Boyle, P.P. "Options: A Monte Carlo Approach." *Journal of Financial Economics*, 4 (1977), pp. 323-338.

Broadie, M., and J. Detemple. "Option Pricing: Valuation Models and Applications." *Management Science*, Vol. 50, No. 9 (2004), pp. 1145–1177.

Carrière, J.F. "Valuation of the Early-Exercise Price of Options Using Simulations and Nonparametric Regression." *Insurance: Mathematics and Economics*, 19 (1996), pp. 19–30.

Clément, E., D. Lamberton and P. Protter. "An Analysis of a Least Squares Regression Method for American Option Pricing." *Finance and Stochastics*, Vol. 6, No. 4 (2002), pp. 449–471.

De Jong, C., and K. Walet. "To Store or Not To Store." *EPRM* (October 2003), pp. 8–11.

DeVries, K.L. "Improved Modeling Increases Salt Cavern Storage Working Gas." *GasTIPS* (Winter 2003), pp. 33–36.

Ghiuvea, C. "Pricing of Generalized American Options with Applications to Energy Derivatives." Ph.D. dissertation, Carnegie Mellon University, Pittsburgh, 2001.

Glasserman, P. *Monte Carlo Methods in Financial Engineering*. New York, NY: Springer-Verlag, 2004.

Gray, J., and P. Khandelwal. "Towards a Realistic Gas Storage Model." *Commodities Now* (June 2004), pp. 75–79.

Haugh, M.B., and L. Kogan. "Pricing American Options: A Duality Approach." *Operations Research*, Vol. 52, No. 2 (2001), pp. 258–270.

International Energy Agency. *Natural Gas Information 2004*.

Keppo, J. "Pricing of Electricity Swing Options." *Journal of Derivatives*, 11 (2004), pp. 26–43.

Kircher, S. *Robustness of the Least Square Monte Carlo Method*. Master's thesis, Christ Church College, University of Oxford, 2004.

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

Lucia, J., and E.S. Schwartz. "Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange." *Review of Derivatives Research*, Vol. 5, No. 1 (2002), pp. 5–50.

Maragos, S. "Valuation of the Operational Flexibility of Natural Gas Storage Reservoirs." In E. Ronn, ed., *Real Options and Energy Management*. London: Risk Books, 2002.

Meinshausen, N., and B.M. Hambly. "Monte Carlo Methods for the Valuation of Multiple Exercise Options." *Mathematical Finance*, 14 (2004), pp. 557–583.

Moreno, M., and J.F. Navas. "On the Robustness of Least-Squares Monte Carlo (LSM) for Pricing American Derivatives." *Review of Derivatives Research*, Vol. 6, No. 2 (2003), pp. 107–128.

Peterson, D., and J. Gray. "Managing Physical Asset Risk with a Dynamic Programming Model." *Commodities Now* (March 2004), pp. 74–79.

Rogers, L.C.G. "Monte Carlo Valuation of American Options." *Mathematical Finance*, Vol. 12, No. 3 (2002), pp. 271–286.

Schwartz, E.S. "The Stochastic Behaviour of Commodity Prices: Implications for Valuation and Hedging." *Journal of Finance*, Vol. 52, No. 3 (1997), pp. 923–973.

Schwartz, E.S., and J.E. Smith. "Short-term Variations and Long-term Dynamics in Commodity Prices." *Management Science*, Vol. 46, No. 7 (2000), pp. 893–911.

Stentoft, L. "Assessing the Least Squares Monte-Carlo Approach to American Option Valuation." *Review of Derivatives Research*, Vol. 7, No. 2 (2004a), pp. 129–168.

——. "Convergence of the Least Squares Monte Carlo Approach to American Option Valuation." *Management Science*, Vol. 50, No. 9 (2004b), pp. 1193–1203.

Thanawalla, R.T. "Valuation of Gas Swing Contracts Using Least Squares Monte Carlo." Working Paper, Heriot-Watt University, 2005.

Thompson, M., M. Davison, and H. Rasmussen. "Natural Gas Storage Valuation and Optimization: A Real Options Application." Working Paper, University of Western Ontario, 2003.

Tilley, J.A. "Valuing American Options in a Path Simulation Model." *Transactions of the Society of Actuaries*, 45 (1993), pp. 83–104.

Tseng, C.-L., and G. Barz. "Power Plant Operations and Real Options." In E. Ronn, ed., *Real Options and Energy Management*, London: Risk Books, 2002.

Tsitsiklis, J.N., and B. van Roy. "Regression Methods for Pricing Complex American-Style Options." *IEEE Transactions on Neural Networks*, Vol. 12, No. 4 (2001), pp. 694–703.

Weston, T. "Applying Stochastic Dynamic Programming to the Valuation of Gas Storage and Generation Assets." In E. Ronn, ed., *Real Options and Energy Management*. London: Risk Books, 2002.