Gas Storage Valuation Models

First Draft

Jeffrey Hunt Mantel, David Peterson, and **Martyn Turner**

QuIC Financial Technologies

February 5, 2007

Fix that gauge or you'll run out of gas
A cool operator can make it last
– "Cool 'N' Out", by Joe Strummer & The Mescaleros

1 Introduction

A key factor in the control and valuation of a gas storage facility is the difference between prices of short-dated forward contracts and long-dated contracts. This difference is a reflection of both the contango or backwardation prevailing in the market and seasonal variation with respect to the forward delivery date. One way to create a suitable "wedge" between short- and long-dated gas forward prices is to use a multifactor term structure model for gas forward prices. This model must also capture the pronounced seasonality of gas forward prices. The first part of this document describes two possible approaches to modeling the term structure of natural gas forwards, incorporating both multifactor dynamics and seasonal variation. The second part describes a simulation–based valuation method for a gas storage facility, based on recent developments in the literature. This valuation methodology can be applied to any model of underlying forward prices.

The present incarnation of this work is more a proposal than a white paper. While QuIC Financial Technologies has implemented one of the gas forward models proposed in Section 2 below, this work is based on outdated market data; *e.g.* Henry Hub gas and various North American electricity hubs *circa* 2001–2002. As noted below, the results of this exercise are quite promising, but it is necessary to evaluate the models proposed here using more recent market data in order to assess their general suitability. QuIC has also implemented a gas storage valuation methodology similar to that described in Section 4 below. However, this work is based on a client's proprietary model of gas forward dynamics and cannot be packaged for resale elsewhere. Accordingly, QuIC seeks to develop its own forward modeling methodology and apply it to the pricing and risk analysis of gas storage facilities. To this end, two gas forward modeling approaches and a storage valuation methodology are proposed below.

The main objective of this proposal is to arouse sufficient interest to elicit current gas market data and/or critical input from benevolent readers. With data in hand, QuIC will apply the methodologies presented here to obtain representative pricing results for a gas storage model.

2 Seasonal forward price dynamics

As is customary, let $F(t,T)$ denote the time–t forward price for gas expiring at time T. Clewlow and Strickland (2000, p. 147) seek to account for seasonal variation in gas forward prices by specifying risk–neutral dynamics

$$\frac{dF(t,T)}{F(t,T)} = v(t)\sum_{i=1}^{d}\sigma_i(T-t)dZ_i(t) = v(t)\sigma(T-t)^\top dZ(t), \tag{1}$$

where $\sigma = (\sigma_1, \cdots, \sigma_d)^\top$ is a d–dimensional function and $Z = (Z_1, \cdots, Z_d)^\top$ is a d–dimensional standard Brownian motion.[1] The scalar function $v(t)$ is today's (time t) spot price volatility.

To understand the implications of this approach, suppose that today (t) is February 1, 2007. According to this model, the volatility of the September 2007 forward contract is proportional to spot price volatility today. However, one would not expect to see a substantial link between September forward volatility and February spot volatility. Uncertainty about the demand for September gas is unlikely to be reflected in spot gas volatility in February. This suspicion is borne out by QuIC's study of 2001–2002 natural gas data. In comparison to a textbook nonseasonal model; *i.e.* where $v(t) = 1$, the model (1) does little in terms of reducing the number of principal components required to explain the variation in historical gas forward prices. At least four principal components are required to explain 90 percent of the variation in historical gas forwards, whether an estimate of spot volatility is used for $v(t)$ or whether $v(t)$ is simply held constant. The situation is even more pronounced in the case of electricity prices, where at least seven principal components are required for on–peak Mid-Columbia power prices whether or not forward volatility is scaled by spot volatility. Furthermore, when the PCA is repeated on a rolling basis from one historical observation date to the next, the leading principal components are observed to be quite unstable over time, whether or not forward price volatility is scaled by spot volatility. Accordingly, at least for the 2001–2002 data record available to QuIC, the model (1) is not particularly compelling. We propose two alternatives below.

2.1 Scalar time–to–maturity variation

In both of the approaches proposed here, forward price volatility is represented as the product of a scalar function v and a vector function σ, one of which depends on the forward expiry date T and the other on the time–to–maturity $T - t$. In our first approach, risk–neutral

[1] Z and σ are both column vectors. "$X^\top$" denotes the transpose of a vector X.

forward price dynamics take the form

$$\frac{dF(t,T)}{F(t,T)} = v(T-t)\sigma(T)^{\top} dZ(t), \quad t < T. \tag{2}$$

Here the scalar function $v(T-t)$ depends on the time–to–maturity, while the d–dimensional vector $\sigma(T)$ is a function of the contract expiry date T. Unlike (1), forward volatility is not anchored to today's spot volatility. Instead, the function v "looks ahead" to the forward expiry date. The problem, of course, is that we must extract $v(T)$ from historical data instead of using a simple estimate of the current spot price volatility. A method for doing this is described below.

To simplify things somewhat, we impose the parametrization

$$v(T-t) = e^{-\beta(T-t)} + v_{\infty}, \tag{3}$$

where v_{∞} and β are constants. If β is positive, then as the observation date t approaches the contract expiry date, forward volatility increases. The exponential decay in forward volatilities with respect to $T-t$ is a characteristic of linear mean reversion in spot prices. For forward expiries far in the future, the exponential term contributes a negligible amount to forward volatility. In this case, the forward volatility is largely a function of the calendar month in which the contract expires.

2.2 Calibrating the model

Let $\theta = (v_{\infty}, \beta)$, $v_{\theta}(T-t) = e^{-\beta(T-t)} + v_{\infty}$, and $f(t,T) = \ln F(t,T)$. Then if Δt is small, the scaled increments

$$\left(\frac{\Delta f}{v_{\theta}}\right)(t,T) = \frac{f(t+\Delta t,T) - f(t,T)}{v_{\theta}(T-t)} \tag{4}$$

are stationary, with variance $\|\sigma(T)\|^2 \Delta t$. If we impose the reasonable seasonality assumption that $\|\sigma(T)\| = \|\sigma(T + 1 \text{ year})\|$, it follows that

$$\operatorname{Var}\left[\left(\frac{\Delta f}{v_{\theta}}\right)(t,T)\right] = \operatorname{Var}\left[\left(\frac{\Delta f}{v_{\theta}}\right)(t,T+1 \text{ year})\right]$$

for every expiry T. The Generalized Method of Moments then suggests that we choose $\theta = (v_{\infty}, \beta)$ to minimize the objective

$$\sum_{T_i} \left(\frac{\operatorname{Var}\left[\left(\frac{\Delta f}{v_{\theta}}\right)(t,T_i)\right]}{\operatorname{Var}\left[\left(\frac{\Delta f}{v_{\theta}}\right)(t,T_i + 1 \text{ year})\right]} - 1 \right)^2,$$

where the sum is taken over some subset $\{T_i\}$ of observable expiry dates and "Var[ ]" denotes the (sample) variance over the historical observation window $[t_1, t_2]$. Having estimated θ in this way, we can apply Principal Component Analysis to the scaled price increments (4) to estimate $\sigma(T)$.

Applying this approach to historical Henry Hub gas data over the time period 2001–2002 yields interesting results. The number of principal components required to explain 90 percent of historical forward variation drops from four to two. The first two principal components are also noticeably more stable through time when taking this approach. Results for Mid–Columbia electricity forwards over the same time period are even more compelling. The number of principal components falls from seven to three and the same time stabilization of leading principal components is observed. Accordingly, it seems promising to investigate the performance of this approach on more recent market data.

Another important practical consideration involves fitting the model to option volatilities.[2] A simple step in this direction would be to replace the calibrated $v_\theta(T-t)$ function with a piecewise constant scalar function that replicates at–the–money Black volatilities on the storage facility valuation date $t = 0$. Option data plays a larger role in the calibration of our next model, described below.

2.3 Scalar seasonal variation

One drawback of model (3) is the dependence of the vector function σ on expiry dates. If we are interested in valuing long–lived assets like gas storage facilities, the model must be calibrated many years into the future, perhaps beyond the liquidity horizon of the market. Accordingly, it may be necessary to extrapolate the function σ beyond expiry dates available in the historical data. This presents a challenge, not only due to the multidimensional nature of σ, but also because of its potentially complex seasonal variation. It would be easier if we only had to extrapolate a scalar function. Accordingly, in our second approach, we specify the forward price dynamics

$$\frac{dF(t,T)}{F(t,T)} = v(T)\sigma(T-t)^\top dZ(t), \quad t < T. \tag{5}$$

Here the scalar function $v(T)$ depends on the forward maturity date T while the vector $\sigma(T-t)$ depends on the time to maturity $T-t$. The dependence of σ on the time to maturity is intended to capture two things. First, it captures mean reversion of spot prices via the rate of decrease of $\|\sigma(T-t)\|$ with time to maturity $T-t$. Second, it captures imperfect correlation between forward prices for different maturity dates via the multifactor nature of the model. Moreover, since the components of σ do not exhibit seasonal variation,

[2] Unfortunately, option volatilities were not available to QuIC at the time of the study mentioned above.

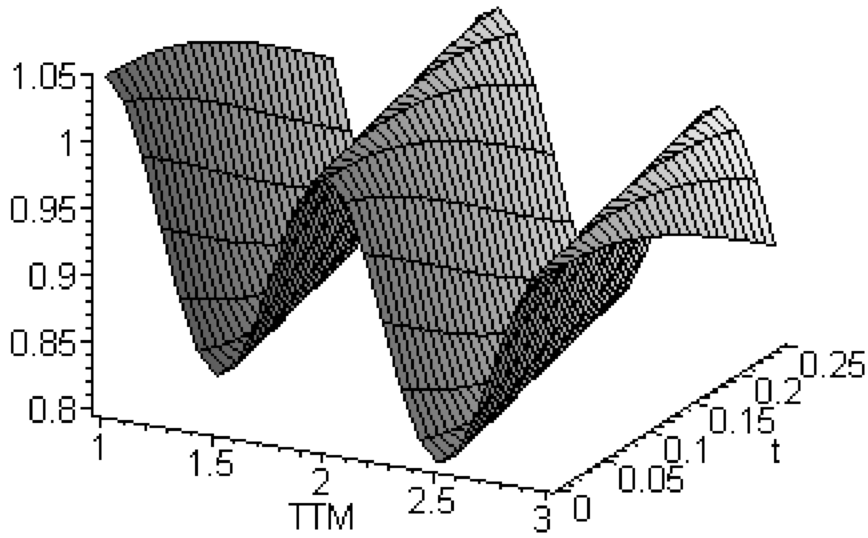

Figure 1: Sample Black volatility function

it is likely more feasible to extrapolate them to expiries T far in the future.[3] Seasonal variations in forward volatilities are captured entirely by the scalar multiplier $v(T)$. We need only worry about seasonal issues when extrapolating the scalar function $v(T)$.

The Black volatility $\sigma_B(t,T)$ corresponding to the observation date t and the maturity date T satisfies

$$\sigma_B(t,T)^2 = \frac{1}{T-t}\int_t^T v(T)^2 \|\sigma(T-u)\|^2 du, \quad 0 < t < T. \tag{6}$$

Figure 1 depicts the evolution of Black volatilities through calendar time t for a simple special case of (5). For every date t on the observation date axis, the Black volatility curve $\sigma_B(t, t+TTM)$ extends for three years along the TTM (time to maturity) axis. Observation dates advance over a three month period along the t axis. Seasonal peaks and troughs in this plot trace regular angles with respect to the horizontal axes, reflecting the fact that when the observation date t advances by a day, the next seasonal peak moves closer to t by a day.[4]

2.4 Calibrating the model

The model is designed to capture market volatilities, seasonal variation, mean reverting spot prices, and imperfect correlation across the forward term structure. It is somewhat difficult

[3]Note that under these dynamics, instantaneous forward correlations have no seasonality. However, *terminal* forward correlations can indeed exhibit seasonality; *cf.* Brigo & Mercurio (2006) §6.6.

[4]This plot corresponds to the extremely simple two–parameter case $v(T) = 1 + A\cos(2\pi T), \sigma(T-t) = e^{-a(T-t)}$, $A = a = 0.1$. Black volatilities having more interesting seasonal and time–to–maturity effects can be obtained by using more complex parametrizations and/or a multidimensional σ function.

to account for all of these phenomena simultaneously as part of a global parameter fitting exercise, so we take an iterative approach. We first deseasonalize historical forward price increments by using historical Black volatilities. The function $\sigma(T-t)$ is chosen to replicate the covariance structure of the deseasonalized forward price data. Then the scalar function $v(T)$ is defined so that the model replicates the Black volatilities observed on the storage facility valuation date $t = 0$.

2.4.1 Step 1 - Deseasonalization

Market data is assumed to consist of the following. On days $t_0 < t_1 < \cdots < t_I < 0$, both forward prices $F(t_i, t_i + \tau_j)$ and Black volatilities $\sigma_B(t_i, t_i + \tau_j)$ are available, where $0 = \tau_0 < \tau_1 < \cdots < \tau_J$ is a grid of times to maturity.[5] We assume that seasonality in Black volatilities mirrors the seasonality in historical forward volatilities. Accordingly, we will use the Black volatilities to deseasonalize forward price increments. However, for this purpose, Black volatilities cannot be applied directly. We must instead extract *instantaneous* volatilities from the Black volatilities. Specifically, we define piecewise constant functions $v(t_i, t_i + \tau)$ such that

$$\sigma_B(t_i, t_i + \tau_j)^2 = \frac{1}{\tau_j} \int_0^{\tau_j} v(t_i, t_i + u)^2 du, \quad 1 \le i \le I, \; 1 \le j \le J. \tag{7}$$

This is fairly easy. Assume that $v(t_i, t_i + \tau)$ is a constant v_{ij} over the time interval $\tau_{j-1} < \tau \le \tau_j$; *i.e.*

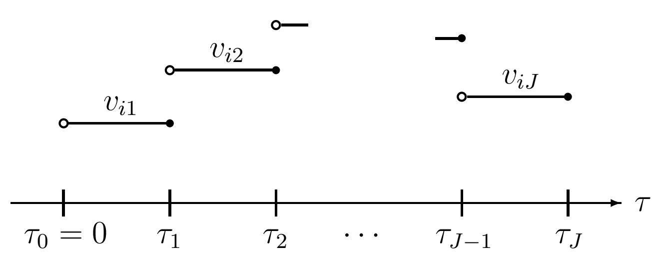

Figure 2: Instantaneous volatility function $v(t_i; t_i + \tau)$

[5] The observation dates t_i would typically consist of consecutive trading days spanning a few months while the times to maturity τ_j would extend as far as the longest-dated liquid forward contract, typically twelve to twenty–four months. The time to maturity $\tau_0 = 0$ is introduced simply for notational convenience. We assume that $\sigma_B(t_i, t_i) = 0$ for all i. Gas market data will typically not correspond exactly to any particular time to maturity grid. Consequently, some form of interpolation must be applied to the data. For more on this point, see Section 2.5 below.

To satisfy (7), we define $v_{ij} > 0$, $1 \le j \le J$, such that

$$\tau_j \sigma_B(t_i, t_i + \tau_j)^2 = \tau_{j-1} \sigma_B(t_i, t_i + \tau_{j-1})^2 + v_{ij}^2 (\tau_j - \tau_{j-1}). \tag{8}$$

We then define deseasonalized log forward price increments

$$x_{ij} = \frac{1}{v_{ij}} \left[\ln F(t_i, t_i + \tau_j) - \ln F(t_{i-1}, t_{i-1} + \tau_j) \right], \tag{9}$$

where $1 \le i \le I, 1 \le j \le J$.

2.4.2 Step 2 – Calibrating to deaseasonalized data

We now turn to the calibration of the $\sigma(T-t)$ function. Given x_{ij} as defined in (9), let

$$\overline{x}_j = \frac{1}{I} \sum_{i=1}^{I} x_{ij}, \quad c_{jk} = \frac{1}{I} \sum_{i=1}^{I} \left(x_{ij} - \overline{x}_j \right) \left(x_{ik} - \overline{x}_k \right),$$

where $1 \le j, k \le J$. The $J \times J$ matrix $C = (c_{jk})$ is the covariance matrix of deseasonalized log forward price increments. Now define vectors $\sigma(\tau_j)$, $1 \le j \le J$, such that $\sigma(\tau_j)^\top \sigma(\tau_k) \Delta t \approx c_{jk}$, Δt being the average time step $t_i - t_{i-1}$ between adjacent historical observation dates. The dimension d of $\sigma(\,\cdot\,)$ can be specified directly (a typical value being $d = 3$) or it can be determined implicitly by a percentage threshold.[6]

Some form of interpolation is required to define $\sigma(\,\cdot\,)$ for all possible times to maturity. In practice, a piecewise constant specification is desirable because it simplifies the calculation of the integral (6) defining the Black volatility. Accordingly, we define

$$\sigma(\tau) = \begin{cases} \sigma(\tau_j) & \text{if } \tau_{j-1} < \tau \le \tau_j \\ \sigma(\tau_J) & \text{if } \tau > \tau_J. \end{cases} \tag{10}$$

[6]This is a standard problem. An eigendecomposition of C yields the representation $C = U^\top D U$, where U is a $J \times J$ orthogonal matrix and $D = \text{diag}[\lambda_1, \cdots, \lambda_J]$ is the $J \times J$ diagonal matrix having the eigenvalues of C along its diagonal. Assume that $\lambda_1 \ge \cdots \ge \lambda_J$. By discarding all but the d largest eigenvalues, we arrive at the approximate representation $C \approx V^\top D^* V$, where $D^* = \frac{\lambda_1 + \cdots + \lambda_J}{\lambda_1 + \cdots + \lambda_d} \text{diag}[\lambda_1, \cdots, \lambda_d]$ and the $d \times J$ matrix V is obtained by discarding the last $J - d$ rows from U. We can then define

$$\sigma(\tau_j)_\ell = V_{\ell j} \sqrt{D^*_{\ell\ell} / \Delta t}, \quad 1 \le j \le J,\ 1 \le \ell \le d,$$

and set $\sigma(\tau_j) = (\sigma(\tau_j)_1, \cdots, \sigma(\tau_j)_d)^\top$. If the integer d is not specified directly, it can be determined implicitly by requiring the fraction $\frac{\lambda_1 + \cdots + \lambda_d}{\lambda_1 + \cdots + \lambda_J}$ to exceed a user–specified percentage.

2.4.3 Step 3 – Fitting today's Black volatilities

Having defined $\sigma(T-t)$ in (10), it is a simple matter to define the scalar function $v(T)$ to fit market observed Black volatilities $\sigma_B(0,\tau_j)$. We assume that $v(T)=\sum_{j=1}^{J} v_j \mathbf{1}_{(\tau_{j-1},\tau_j]}(T)$ has the piecewise constant structure shown in Figure 2, but instead of defining the levels v_j by analogy to (8), we define them such that (6) holds; *i.e.*

$$\tau_j\sigma_B(0,\tau_j)^2=\tau_{j-1}\sigma_B(0,\tau_{j-1})^2+v_j^2\|\sigma(\tau_j)\|^2(\tau_j-\tau_{j-1}), \tag{11}$$

where $v_j>0$ and $1\leq j\leq J$. This completes the calibration process.

In principle, we could now make a second pass, where, given the vectors $\sigma(\tau_j)$, we recalibrate the $v(t_i,\,\cdot\,)$ functions to historical Black volatilities $\sigma_B(t_i,\,\cdot\,)$ in a manner analogous to (11). Specifically, we update the constants v_{ij} such that

$$\tau_j\sigma_B(t_i,t_i+\tau_j)^2=\tau_{j-1}\sigma_B(t_i,t_i+\tau_{j-1})^2+v_{ij}^2\|\sigma(\tau_j)\|^2(\tau_j-\tau_{j-1})$$

and substitute them into (9) to obtain new x_{ij}s. We can now recalculate the $\sigma(\tau_j)$ vectors from the covariance matrix of the updated x_{ij}s and recalculate $v(T)$ with a second application of (11). In fact, we could iterate this process arbitrarily many times. The practical impact of the number of iterations will be assessed in our work with market data in Section 3 below.[7]

2.5 Implementation considerations

Energy price data typically requires some amount of preprocessing. Ideally, the forward price $F(t,T)$ described above will correspond to the delivery of one unit of gas on a single future day T. However, forward prices corresponding to one specific delivery day are not usually available, particularly for long–dated contracts. Instead, observable forward prices usually correspond to a non–negligible delivery period, such as a month, a quarter, or an entire year. Accordingly, some form of "daily scaling" must be applied to the data to extract a reasonable proxy for future daily delivery.[8] Likewise, Black volatilities corresponding to fictitious daily delivery contracts will also be unobservable. Instead, we typically observe the prices of options to take delivery over monthly, quarterly, or annual delivery periods. These types of contracts resemble interest rate caps or swaptions more than standard options, and suitable adjustments to the Black pricing formula must be made.[9] Accordingly, some amount of preprocessing will typically also be required to extract the Black volatilities referred to above.

[7]Please note that unlike (3), QuIC has not yet implemented this model because of the unavailability of historical option market data.

[8]To this end, QuIC has implemented scaling algorithms similar to those described in Forsgren (1998).

[9]See, for example, Brigo & Mercurio (2006) §6.15.

3 Application to market data

Work in progress.

4 Valuing a gas storage facility

Commodity storage facilities provide the capability to take advantage of market price shocks. When commodity prices rise, inventory can be released from storage and sold on the spot market. The inventory can be replenished later when prices fall again. There are several different types of storage technologies available. For example, salt caverns can be used to store natural gas, man–made caverns blasted out of rock can used to store oil, and dammed reservoirs can be used to hold water for discretionary power generation. While the physical characteristics of these examples differ substantially, the business logic and business objectives remain the same – to operate the storage facility in a manner that maximizes its value. Accordingly, while the discussion below is phrased in terms specific to natural gas, the analysis applies in a more general context.

The basic idea behind the valuation model is as follows. Given the gas forward curve at the valuation date $t = 0$, simulate the evolution of the forward curve using the model (2) or (5), generating many thousands of forward curve scenarios through time, out to a valuation horizon $t = T_H$. Now given the gas stored in inventory at $t = 0$, calculate the optimal storage management policy. This involves a backward recursion over the time steps of the forward curve simulation. Least squares regressions are applied to one–step–ahead NPVs at every step of the recursion. From these results, an optimal storage management policy can be determined. The NPV of the storage facility at $t = 0$ is the average discounted value of all simulated cash flows generated under the optimal policy.

4.1 Operational characteristics

There are several key parameters that characterize a gas storage facility, arising either from physical or regulatory constraints. These quantities are, in effect, user–specified static data.

4.1.1 Capacity parameters

In a standard gas storage model, capacity parameters depend on the following system characteristics:

- **Base Gas** – the volume that must be maintained in storage at all times;
- **Maximum Capacity** – the maximum volume of gas that the facility can hold;
- **Working Capacity** – Maximum Capacity minus Base Gas.

The amount of gas in storage may be subject to time–dependent constraints. A minimum amount greater than the Base Gas level may need to be stored at various times for regulatory reasons. Seasonal capacity constraints may arise for purely physical reasons. If the storage facility has been leased from another party for the time period $0 \leq t \leq T_H$, the lease may require that the amount of gas in storage at time T_H be at least as large as the amount initially stored at time $t = 0$. However these limits may arise, there is a time–dependent **minimum inventory** $I_{\min}(t)$ and a time–dependent **maximum inventory** $I_{\max}(t)$.[10]

4.1.2 Withdrawal and injection tables

Gas can be injected or withdrawn from storage. However, the rate with which injection or withdrawal takes place depends on the volume of gas in storage. The more gas in the facility, the higher the pressure, making it easier to withdraw and harder to inject gas. Daily injection and withdrawal tables will be used to specify the amount of gas that can be injected into or withdrawn from the facility over the course of a day for a given level of inventory. These tables specify

- The **withdrawal schedule** $V_-(I)$, the volume of gas that can be withdrawn from the facility over a day given the initial inventory I;
- The **injection schedule** $V_+(I)$, the volume of gas that can be injected over the course of a day given initial inventory I.

It proves convenient to assign V_- a negative sign while V_+ is positive. We also define $V_0(I) = 0$. Gas storage facilities are also subject to a certain degree of seepage. Typically the seepage rate will higher the greater the pressure; *i.e.* the larger the inventory. Therefore we assume the daily **seepage rate** to be a function $V_s(I)$ of the current inventory level.

4.1.3 Cost tables

There are costs associated with storing the gas and pumping it into or out of the storage facility. Like the control parameters described in Section 4.1.2 above, these costs will be represented by user–specified tables. They will specify the following functions:

- The **withdrawal cost** $P_-(I)$, the cost of withdrawing gas from storage;

[10]The default unit of storage in the case of natural gas is one billion cubic feet (Bcf).

- The **injection cost** $P_+(I)$, the cost of pumping gas into storage;
- The **storage cost** $P_0(I)$, the cost of storing gas over the course of a day.

These are unit costs denominated in USD per million BTU per day.

4.1.4 Switching costs

Finally, there may be one–time **setup costs** associated with the activation or deactivation of the pumping facilities required to move gas in or out of storage, regardless of the volume actually pumped. These costs are represented by the off–diagonal entries of the 3×3 matrix

$$S = \begin{pmatrix} 0 & S_{0-} & S_{0+} \\ S_{-0} & 0 & S_{-+} \\ S_{+0} & S_{+-} & 0 \end{pmatrix}.$$

S_{0-}, S_{0+} are the one–time costs involved in activating pumps for withdrawal or injection, respectively. S_{-+}, S_{+-} are the costs involved in reversing the pumps. S_{+0}, S_{-0} are pump shutdown costs. These costs are denominated in USD.

4.2 Valuation approach

The value of a storage facility will depend on the following factors:

- **Time** t.
- **Market risk factor(s)** $X(t)$; *e.g.* a finite set of forward prices. $X(t)$ can be one–dimensional, for example, it might consist of the nearby forward price alone, or it can be multidimensional; *e.g.* it could consist of the nearby forward price together with several longer maturity forward prices. We will assume that $X(t)$ always has the **nearby forward price** as a component, which we write as $\nu(X(t))$.
- The **inventory level** $I(t)$.
- A **control variable** $C(t) \in \{-, 0, +\}$ describing whether we are currently withdrawing gas from storage, holding the inventory level constant, or injecting gas, respectively.
- The **discount factor** $DF(t, t')$ in effect between any two dates $0 \leq t < t' \leq T_H$.

Solving for the storage facility's NPV under general conditions can only be done numerically. Accordingly, we discretize calendar time into steps $0 = t_0 < t_1 < \cdots < t_{\text{nSteps}} = T_H$. Write the storage facility NPV at time t_m and in state $(x, i, c) = (X(t_m), I(t_m), C(t_m))$ as $PV_m(x, i, c)$.

4.2.1 Inventory dynamics

Controls are assumed to be constant between adjacent time steps t_m, t_{m+1}, while risk factors and inventory levels may continuously change. Consequently, given the inventory level $I(t_m) = i$ and the control c put in place at time t_m, the future inventory level $I(t_{m+1})$ is a deterministic function $I_{m+1}(i, c)$ of the arguments i and c. In the absence of storage capacity constraints, we would define $I_{m+1}(i, c) = \mathcal{I}_{m+1}(i, c)$, where

$$\mathcal{I}_{m+1}(i, c) = i + (V_c(I) - V_s(I))\, \Delta_m$$

and $\Delta_m = t_{m+1} - t_m$ is the number of days between adjacent time steps. Some additional care is required in the presence of capacity constraints. For example, suppose that $\mathcal{I}_{m+1}(i, +) \ll I_{\min}(t_{m+1})$. In this case, given the inventory level i at time t_m, it may simply be physically impossible to pump enough gas into the storage facility to meet the minimum storage constraint at t_{m+1}. For example, if there is a lease stating that the storage facility must be returned with at least as much inventory as the original amount $I(0)$, we would have $I_{\min}(t_{\text{nSteps}}) = I(0)$. If $I(0)$ is relatively large but the inventory at $t_{\text{nSteps}-1}$ is close to the minimum tank capacity, it may simply not be possible to inject enough gas over the time period $[t_{\text{nSteps}-1},\, t_{\text{nSteps}}]$ to meet the terminal capacity constraint. On the other hand, if the inventory level i at t_m lies in $[I_{\min}(t_{m+1}),\, I_{\max}(t_{m+1})]$ but $I_{m+1}(i, c)$ does not, then we must have reached the minimum or maximum storage capacity at some time *between* t_m and t_{m+1}. In this case, we can simply cap $I_{m+1}(i, c)$ at the relevant capacity limit.

To account for these subtleties in the definition of $I_m(i, c)$, it is convenient to introduce a non–numeric "death state" † and define $\mathcal{I}_{m+1}(\dagger, c) = \dagger$ for all controls c. Given the considerations above, we then define

$$I_{m+1}(i, c) = \begin{cases} \text{cap}_{m+1}\left(\mathcal{I}_{m+1}(i, c)\right) & \text{if either } i \text{ or } \mathcal{I}_{m+1}(i, c) \in [I_{\min}(t_{m+1}),\, I_{\max}(t_{m+1})] \\ \dagger & \text{otherwise.} \end{cases} \tag{12}$$

Here

$$\text{cap}_{m+1}(i) = \max\{I_{\min}(t_{m+1}),\, \min\{I_{\max}(t_{m+1}),\, i\}\}.$$

Thus, $I_{m+1}(i, c)$ is set to the death state precisely when it is physically impossible to meet the capacity constraint at time t_{m+1}. Under any other circumstance, it is capped at the relevant capacity limit.

4.2.2 Cash flows

Suppose that $X(t_m) = x$ and $I(t_m) = i$. Changing from control c to control c' at time t_m triggers the immediate cash flow $CF_m(c, c', x, i)$. In reality, the actual cash flow may be a stream of several payments occurring between t_m and t_{m+1}. However, assuming that Δ_m is small, it is reasonable to approximate this stream by a single aggregate payment. We first consider some examples where $I_{m+1}(i, c') \neq \dagger$; situations involving an infeasible inventory level at time t_{m+1} will be treated as exceptions; see (13) below. For example, assuming that $I_{m+1}(i, -) \neq \dagger$, we set

$$\begin{aligned} CF_m(0, -, x, i) &= -S_{0-} + (\nu(x) - P_-(i)\Delta_m)\,\eta\,(i - I_{m+1}(i, -)) \\ &\quad -P_0(i)\eta\Delta_m i. \end{aligned}$$

The first term on the right–hand side is the pump activation cost. The second term is the net revenue obtained by selling the incremental volume $i - I_{m+1}(i, -)$ on the spot market, net of unit pumping costs.[11] The final term is the incremental storage cost over the time period between t_m and t_{m+1}. Likewise, assuming that $I_{m+1}(i, +) \neq \dagger$,

$$\begin{aligned} CF_m(0, +, x, i) &= -S_{0+} - (\nu(x) + P_+(i)\Delta_m)\,\eta\,(I_{m+1}(i, +) - i) \\ &\quad -P_0(i)\eta\Delta_m i. \end{aligned}$$

Cash flows in all other cases involving feasible future inventory levels can be determined in the same way. To deal with situations involving infeasible future inventory levels, we define

$$CF_m(c, c', x, i) = -\infty \quad \text{whenever } I_{m+1}(i, c') = \dagger. \qquad (13)$$

4.2.3 Value maximization

The storage facility operator seeks a control schedule that maximizes the facility's NPV subject to the operational constraints described above. As one might expect, at time step t_m, the control c' is chosen to maximize the sum of the immediate cash flow and the discounted expected future NPV of the storage facility. Accordingly, the NPV satisfies the backward recursion

$$\begin{aligned} PV_m(x, i, c) = \max_{c'=-,0,+} \big(&CF_m(c, c', x, i) + \\ &DF(t_m, t_{m+1})\,\mathbb{E}\left[PV_{m+1}(X(t_{m+1}), I(t_{m+1}), c') \,|\, X(t_m) = x, I(t_m) = i\right]\big). \end{aligned} \qquad (14)$$

[11]Here the conversion factor η is the number of MMBTUs per Bcf of gas.

The term on the right–hand side of this equation is the highest value available from the three possible controls $-, 0$, and $+$.[12]

The conditional expectations appearing on the right–hand side of (14) can rarely be expressed in closed form. However, Carmona & Ludkovski (2005), building on work by Longstaff & Schwartz (2001), provide an efficient method for approximating them via Monte Carlo simulation and linear regression. Here we apply the same approach, accounting for inventory constraints described previously. The basic idea is to apply a cross–sectional regression over simulations of the underlying state variables $X(t)$ and $I(t)$. Specifically, given simulations

$$\left(x_0^n, \cdots, x_{\text{nSteps}}^n\right), \quad n = 1, \cdots, \text{nSims},$$

of the sample path $(X(t_0), \cdots, X(t_{\text{nSteps}}))$ and uniformly distributed random draws

$$i^n \in [I_{\min}(t_m),\, I_{\max}(t_m)], \quad n = 1, \cdots, \text{nSims},$$

make the approximation

$$\mathbb{E}\left[PV_{m+1}(X(t_{m+1}), I(t_{m+1}), c) \middle| X(t_m) = x, I(t_m) = i\right] \approx E_{m+1}(x, i, c),$$

whenever $i \neq \dagger$ and

$$E_{m+1}(x, i, c) = \sum_{j=1}^{\text{nBasis}} a_j^* \psi_j(x, i), \quad -\infty < a_j^* < \infty, \tag{15}$$

is a linear combination of user–specified basis functions ψ_j. The a_j^* coefficients solve the linear least squares problem

$$\min_{a_j} \sum_{n=1}^{\text{nSims}} w\left(I_{m+1}(i^n, c)\right) \left[\sum_{j=1}^{\text{nBasis}} a_j \psi_j(x_m^n, i^n) - PV_{m+1}(x_{m+1}^n, I_{m+1}(i^n, c), c)\right]^2, \tag{16}$$

where

$$w(i) = \begin{cases} 1 & \text{if } i \neq \dagger \\ 0 & \text{otherwise.} \end{cases}$$

The ψ_js are typically polynomial products of the arguments x, i. Solving for the coefficients a_j^* is a standard problem in linear algebra; *e.g.* Demmel (1997) §3.2.1. The purpose of the weighting function w is to eliminate situations where the PV_{m+1} term inside the square brackets is $-\infty$. Otherwise the regression would clearly fail to be meaningful.

[12] Here the expectation is taken with respect to the risk–neutral distribution under which the dynamics (2) or (5) were originally formulated. The backward recursion follows from Bellman's principle of optimality, which is the basis of dynamic programming. See Bellman (1957). Note that the convention (13) effectively removes any control leading to an infeasible future inventory level from consideration in the maximization problem.

5 Valuation Algorithm

Here we summarize the previous concepts in a two part valuation algorithm.

5.1 Step 1 – Backward pass

1. **Monte Carlo Simulation.** Use the dynamics specified in (2) or (5) to generate independent sample paths $\left(x_0^n, \cdots, x_{\text{nSteps}}^n\right)$, $n = 1, 2, \cdots, \text{nSims}$.

2. **NPV Initialization.** $m = \text{nSteps}$, $PV_{\text{nSteps}+1}(x, i, c) \equiv 0$.

3. **Inventory simulation.** Draw random variables i^n, $n = 1, 2, \cdots, \text{nSims}$, independently and uniformly from the interval $[I_{\min}(t_m), I_{\max}(t_m)]$.

4. **Regression.** For each $c = -, 0, +$, use the simulations x^n, i^n, $n = 1, \cdots, \text{nSims}$, to calculate the least squares coefficients a_j^* defining the function $E_m(x, i, c)$ in (15).

5. **NPV Update.** Define the function

$$PV_m(x, i, c) = \max_{c'=-,0,+} \left(CF_m(c, c', x, i) + DF(t_m, t_{m+1}) E_m(x, i, c')\right).$$

6. $m \leftarrow m - 1$. If $m > 0$, return to Step 3. Otherwise stop.

At this point, we've calculated coefficients a_j^* for each of the three least squares problems arising at every simulation time step. While this is enough information to fully represent the functions defined in Step 5, it only amounts a small amount of data, since in practice we will never need more than a few tens of basis functions. We will reuse these coefficients in a forward pass below, where we construct the optimal inventory process and the cash flows corresponding to each of the simulated sample paths $\left(x_0^n, \cdots, x_{\text{nSteps}}^n\right)$. Recall that $I(0)$ denotes the initial storage inventory.

5.2 Step 2 – Forward pass

1. **Initialization.** $m = 0$, $i_0^n = I(0)$, $c_{-1}^n = 0$, $DCF^n = 0$, $n = 1, \cdots, \text{nSims}$.

2. **Control selection.** For every $n = 1, \cdots, \text{nSims}$, choose

$$c_m^n \in \text{argmax}_{c'=-,0,+} \left(CF_m(c_{m-1}^n, c', x_m^n, i_m^n) + DF(t_m, t_{m+1}) E_m(x_m^n, i_m^n, c')\right).$$

3. **Inventory and DCF update.** Set $i^n_{m+1} = I_{m+1}(i^n_m, c^n_m)$ and $DCF^n = DCF^n + DF(0, t_m)CF_m(c^n_{m-1}, c^n_m, x^n_m, i^n_m)$, $n = 1, \cdots$, nSims.

4. $m \leftarrow m + 1$. If $m <$ nSteps, return to Step 2.

5. **Valuation.** $NPV = \frac{1}{\text{nSims}} \sum_{n=1}^{\text{nSims}} DCF^n$.

The forward pass iteratively constructs the optimal control policy and corresponding inventory level forward through time for each of the simulation paths. The discounted cash flows generated along the n^{th} simulation path are aggregated over time in the quantity DCF^n. Finally, the storage facility NPV on the valuation date $t = 0$ is defined as the average of the pathwise DCFs.

5.3 Implementation considerations

5.3.1 Choice of basis functions

An important design consideration is the selection of a "reasonable" set of basis functions. The two–pass approach is somewhat robust to this choice in that it does not make a direct attempt to solve Equation (14) in the backward pass. Instead, it only focuses on the control policy in the backward pass, leaving the calculation of the corresponding NPV to the final forward pass. This will introduce a downward bias to the calculated NPV, in that the calculated control policy will be suboptimal. However, the size of this bias is likely to be smaller than it would be should we have attempted to solve (14) directly in the backward pass. The bias can be further reduced by using a different set of Monte Carlo simulations in the forward pass. For related discussion, see Jaeckel (2002), Chapter 9 and Tavella (2002), Chapter 8. A second important consideration is the dimensionality of x; *e.g.* the number of simulated forward contracts. The number of basis functions naturally increases rapidly with this dimension, so it is crucial to keep it to an absolute minimum. Studies in the interest rate exotics market show that in the absence of path dependency, very good results can be obtained using only one or two dimensions. See, for example, Piterbarg (2003). Since the gas storage problem involves path dependency via the inventory level, it may be necessary to regress on three factors in this case; *e.g.* the nearby forward, a long–dated forward, and the inventory level.

5.3.2 Sensitivity calculations

Another important aspect of the implementation is to report price sensitivity to various buckets in the forward term structure. In principle, this is a matter of bumping the initial forward curve and repricing the storage facility. However, well publicized failures[13] mandate that in the calculation of bump sensitivities via Monte Carlo, some care must be taken. For example, to minimize the impact of purely numerical noise, we must ensure that the same underlying set of driving Brownian increments is used for all bumped simulations.

6 Summary

This document describes two forward price models for natural gas, both of which incorporate multifactor dynamics and seasonal variation. These features are key to capturing realistic gas forward market behaviour. A simulation–based valuation method for gas storage facilities is also detailed. This method can be applied to any model of underlying forward prices. While a preliminary study using outdated market data yields promising results, it remains to apply these methods to contemporary market data. There are several future considerations. Another possible modeling approach could be based on a fully parametric model for gas volatities and correlations, in the spirit of the parametric interest rate term structure models described in Brigo & Mercurio (2006) §6.9.1. Should such an approach prove feasible, extrapolation of volatilities to distant time horizons would no longer be an issue. A more ambitious goal is to account for volatility skew. While the value of a gas storage facility is mainly derived by exploiting intertemporal differences in gas forward prices, it may also be important to account for changing volatility across the dimension of underlying prices as well. This may be another important determinant of value if the price thresholds for injecting and withdrawing gas happen to be significantly different.

[13] See, for example, van Deventer & Imai (1997) p. 226, footnote 6.

References

R.E. Bellman (1957), *Dynamic Programming.* Princeton.

D. Brigo & F. Mercurio (2006), *Interest Rate Models – Theory and Practice. Second Ed.* Springer.

R. Carmona & M. Ludkovski (2005), "Gas storage and supply guarantees: an optimal switching approach," submitted to *Management Science*.

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

J.W. Demmel (1997), *Applied Numerical Linear Algebra.* SIAM.

A. Forsgren (1998), "A note on maximum smoothness approximation of forward interest rates," working paper, Department of Mathematics, Royal Institute of Technology, Sweden.

P. Jaeckel (2002), *Monte Carlo Methods in Finance.* Wiley.

F.A. Longstaff & E.S. Schwartz (2001), "Valuing American options by simulation: a simple least squares approach," *Review of Financial Studies* **14**: 113–148.

V. Piterbarg (2003), "A practitioner's guide to pricing and hedging callable LIBOR exotics in forward LIBOR models." SSRN working paper.

D. Tavella (2002). *Quantitative Methods in Derivative Pricing: An Introduction to Computational Finance.* Wiley.

D.R. van Deventer & K. Imai (1997), *Financial Risk Analytics: A Term Structure Model Approach for Banking, Insurance & Investment Management.* Irwin.