# STORING ARB

**HYUNGSOK AHN**, Credit Agricolé Lazard Financial Products & U. of Oxford, UK
**ALBINA DANILOVA**, Imperial College, UK
**GLEN SWINDLE**, Constellation Power Source, USA

In natural gas markets, demand exhibits low price elasticity, while supply follows the market price; a phenomenon typically observed in markets for indispensable goods. A significant portion of gas consumption is due to heating and increasingly power generation, which sets a typical demand pattern: high in winter and summer ('high' seasons) and low in shoulder months ('low' seasons). Thus, the expected spot price exhibits annual or biannual periodicity, where the maxima occur in the high seasons and the minima in the low seasons. Superimposed on this seasonality are the usual fluctuations and occasional dramatic spot price movements due to extreme weather conditions or other unpredictable events. Gas storage provides a vehicle for exploiting both seasonal variations in spot price and short time-scale fluctuations in which storage can be depleted during days of extremely high prices. As will be shown below, much of the value of storage is in the associated ability to monetize short time-scale mean reversion.

This paper presents a theoretical framework for gas storage valuation. In fact, the methods developed in this paper apply to essentially any commodity storage problem with a liquid forward market. Of particular note is pump storage in power markets (this includes air compression facilities and water pumping reservoirs) which can be optimized using our approach, although the injection/withdrawal costs will be substantially higher than in gas storage.

Two types of storages, namely virtual and physical storage, are considered. The latter is a typical storage contract that confines the daily injection/withdrawal rates and charges its holder for each delivery service. Virtual storage, on the other hand, is a purely financial product with no limit on injection/withdrawal rate and zero delivery charge. Injection or withdrawal is merely trading the spot index with a virtual trading account, as in the case of passport options. Virtual storage is interesting not only because such contracts are in fact traded, but also from an analytical perspective in that its value is proportional to its size resulting in the whole price system becomes linear.

The key result is that the value of gas storage depends heavily on the volatility term structure of gas futures, and that evaluation of both virtual and physical storage is possible using basic control theory methods. Both types of storage are valued for the 2001-02 season at Henry Hub using a multi-factor forward curve model calibrated to exchange options. The conclusion is that storage value extracted from the seasonal spreads significantly underestimates the value obtained by optimal injection/withdrawal strategies.

## Valuation Principle

The essence of valuing storage is to identify the benefits of owning it. Some benefits are easy to identify and evaluate. For example, one can extract value from the seasonal spread by purchasing gas forward in a low season and selling it forward in a high season, storing the gas between purchase and sale. However, the profit accounting for funding generated by this simple strategy is typically of the order of 3¢ / MMBtu / month which is three or four times below the market price of the storage. This suggests that the market is sensibly ascribing much more value to storage than merely extracting seasonal spreads.
In order to quantify the true value embedded in gas storage, one must design a self-financing, dynamic trading strategy that extracts the maximal savings obtainable by not having to pay the cost of storage.

# Gas Forward Curve Dynamics

The NYMEX gas futures are highly liquid and provide a natural set of hedging instruments. While our analysis will focus on NYMEX futures, it is important to note that the approach is general and applies to other regions, which have exhibited dramatic price movements (most notably the $15 handles in the Northeast during the winter of 1999 and the recently observed $20 handles in the west). The most important feature of a futures contract is that its settlement price is adjusted daily to reset the contract value to zero. The clearinghouse makes this possible by imposing a variation margin that amounts to a daily profit and loss. As a consequence, futures prices are martingales under risk-neutral measures. In addition to this convenient attribute, several empirical features of the volatility of gas forward curves are apparent: (i) spot prices tend to be significantly more volatile than longer maturity contract prices; (ii) high season spot prices are more volatile than low season spot prices. This latter features is clearly due to the fact that during periods of extremely high demand supply constraints due to transportation and storage withdrawal limits result in supply inelasticity.

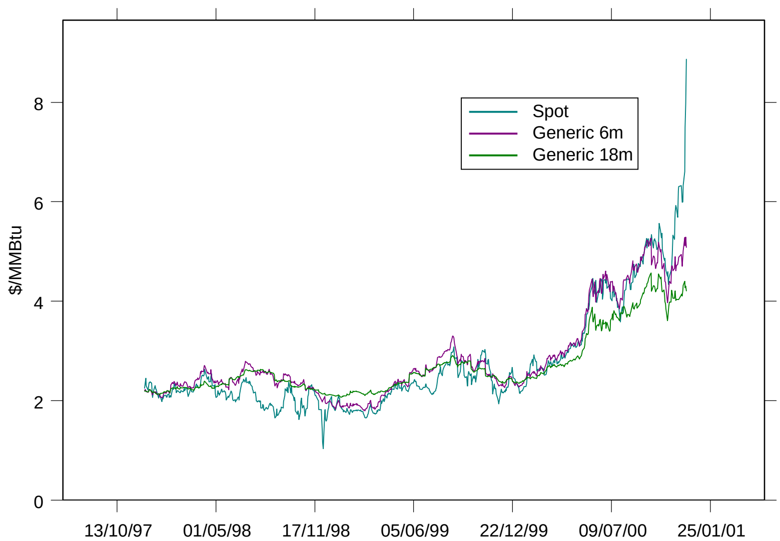


**Figure 1. Henry Hub Spot & NYMEX Futures: Dec 1997 – Dec 2000**

Consequently, the volatility structure should depend not only on time to maturity, but also on maturity date. A commonly used model (see, for example, Clewlow an Strickland (1999)) which has enough flexibility to capture the qualitative features above and which also allows for a finite-dimensional representation of the forward curve is:

$$dF(s,T) = F(s,T) \cdot \sum_{i=1}^{n} \sigma_i(T) e^{-\lambda_i (T-s)} dW_i(s) .$$

Here, $W_i,\ i = 1,\ldots,n,$ form an n-dimensional Brownian motion under the risk-neutral measure, T is maturity, and s is the current time. This particular model for futures prices features mean-reversion and maturity dependent volatility pattern that are manifest in market data. Furthermore, the model is tractable in the sense that we can integrate the above stochastic differential equation (SDE). Using this benefit and ascertaining that futures price converges to spot price, one can describe the spot price as:

$$X(s) = F(s,s) = F(t,s)\cdot\prod_{i=1}^{n}\exp\left\{\sigma_i(s)\int_t^s e^{-\lambda_i(s-u)}\,dW_i(u) - \frac{1}{2}\sigma_i^2(s)\int_t^s e^{-2\lambda_i(s-u)}\,du\right\}.$$

The stochastic integral part of each individual exponent is an Ornstein-Uhlenbeck process that dictates the mean reversion and the deterministic terms (the Itô correction and the forward curve) set the reversion level. As a result, one can describe the whole term structure dynamics with a finite number of state variables, which are Gaussian mean-reverting processes. This property is particularly useful when one tries to build a tree. Finally, the spot price dynamic can be retrieved from the above SDE as well:

$$\frac{dX(s)}{X(s)} = \frac{\partial}{\partial T}\log F(s,T)\Big|_{T=s} ds + \sum_{i=1}^{n}\sigma_i(s)dW_i(s)\,.$$

# Virtual Storage

As mentioned above, virtual storage is a financial contract devoid of constraints on injection/withdrawal rates and without transaction costs. The absence of physical constraints makes valuation simple and intuitive. Assuming that the risk-free rate is a constant r, the value of a self-financing strategy Y satisfies:

$$dY(t) = rY(t)dt + q(t)\cdot\left(dX(t) - rX(t)dt\right).$$

where the variable q denotes the amount of gas in the storage; q varies between zero and the size of the storage, which we will normalize to unity. The last term is the cost of carrying the spot gas without paying any storage cost.

Consequently, the storage value at time t is:

$$\sup_q E_t\left[e^{-r(T-t)}Y(T)\right] = Y(t) + \sup_q E_t\left[\int_t^T e^{-r(s-t)}q(s)\left(dX(s) - rX(s)ds\right)\right].$$

Here, $E_t$ is the expectation under the risk-neutral measure given the information up to time t, and sup is taken over all admissible strategies. The martingale term of the stochastic integral has no effect due to the conditional expectation, and therefore the optimal control is of “bang-bang” type that sets q(s) = 1 (long) whenever the drift of the spot exceeds the risk-free rate and zero (sell off) otherwise. Using the results from earlier section, we obtain:

$$Y(t) + E_t\left[\int_t^T e^{-r(s-t)}\left(\frac{\partial}{\partial T}\log F(s,T)\Big|_{T=s} - r\right)^+ X(s)ds\right].$$

The above expectation can be evaluated analytically:

$$V\left(t,F(t,\cdot),Y\right) = Y + \int_t^T e^{-r(s-t)}F(t,s)\upsilon(t,s)\left[h(t,s)\cdot N\left(h(t,s)\right) + \varphi\left(h(t,s)\right)\right]ds$$

where

$$\upsilon^2(t,s) = \sum_{i=1}^{n}\left(\lambda_i - \frac{\sigma_i'}{\sigma_i}(s)\right)^2\frac{\sigma_i^2(s)}{2\lambda_i}\left(1 - e^{-2\lambda_i(s-t)}\right) \quad\text{and}\quad h(t,s) = \frac{1}{\upsilon(t,s)}\left(\frac{\partial}{\partial T}\log F(t,T)\Big|_{T=s} - r\right).$$

In the above, N and $\varphi$ denote the cumulative distribution function and density for standard normal random variables respectively.

Note that if $\upsilon$ tends to infinity uniformly in s, then h tends to zero and the value of the storage diverges. For that matter, if any one of the model parameters diverges so does $\upsilon$, and the asymptotic behaviour of the value function follows immediately. For example, if $\sigma$ does not depend upon time, we have the following:

- If $\|\lambda\| \to \infty$ and $\sigma$ stays finite (but non-zero), then V is of order $\|\lambda\|^{1/2}$.
- If $\|\lambda\| \to \infty$, $\|\sigma\| \to \infty$, and $\frac{\sigma_i^2}{2\lambda_i}$ stays finite for each $i$, then V is of order $\|\lambda\|$.
- If $\|\sigma\| \to \infty$ and $\lambda$ stays finite, then V is of order $\|\sigma\|$.

An interesting case is when $\lambda$ tends to infinity. Since $\lambda$ governs the speed of mean-reversion, infinite $\lambda$ implies a deterministic spot price. However, the local time of the spot price near the control boundary tends to infinity fast enough so that the value of storage still tends to infinity.

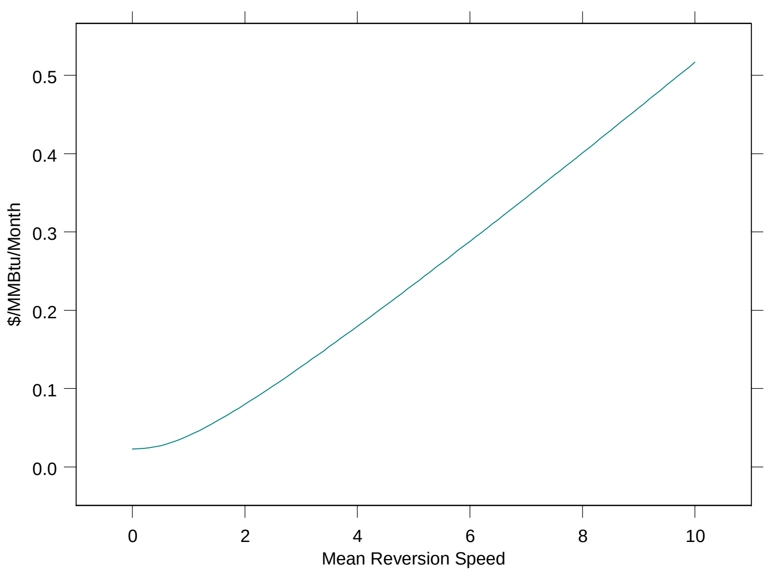


**Figure 2. Virtual Storage Price**

Figure 2 shows the value of virtual storage as a function of mean reversion speed $\lambda$ for the contract period March 2001 through Feb 2002. A two-factor model is considered: (i) factor 1: Flat $\sigma$ (20% per annum) with zero mean reversion; (ii) factor 2: $\sigma$ is calibrated with vanillas, given a mean reversion speed. Increasing the mean reverting speed $\lambda$ elevates the fitted $\sigma$ as well, and the value of the storage grows linearly.

# Physical Storage

A physical storage contract specifies not only the maximum storage quantity but also limits on injection/withdrawal rates a well as volumetric charges. Typically the size of the storage is much larger than the amount of gas that can be transported in a day. In this case, since the whole inventory may not be sold instantly at the current spot price, there is no legitimate way of determining the value of a self-financing strategy using a single state variable. Instead, we must describe the state of the system with a pair: (Y,q), the value of the cash account and the amount of gas in the storage.

Denoting the injection rate by $c(t)$ (c being positive for injection of gas and negative for withdrawal), the retained gas level changes smoothly: $dq(t) = c(t)dt$. Admissible strategies must satisfy: $0 \le q \le Q$ and $-L \cdot 1(q > 0) \le c \le L \cdot 1(q < Q)$ where L is the maximum volumetric gas flow to or from the storage facility. For both injection and withdrawal a delivery charge is imposed. We will assume that the charge is proportional to the delivery size (although the methods below are easily extended to other mechanisms). In this case Y satisfies the following differential equation:

$$dY(t) = \left(rY(t) - c(t)X(t) - \beta|c(t)|\right)dt\ .$$

The first two terms are the risk-free growth of capital and the cost of spot trading. The last term is the delivery charge.

As before, the value of the storage at time t is the risk neutral expectation of the maximal cash amount generated by self-financing strategy:

$$\sup_c E_t\left[e^{-r(T-t)}Y(T)\right]$$

where the sup is taken over the admissible value of c. The value will depend upon $(Y(t), q(t))$ as well as the state variables that govern the forward curve dynamics. One can eliminate the Y-dependency simply by subtracting the cash amount generated so far:

$$U(t,q) = \sup_c E_t\left[e^{-r(T-t)}Y(T) - Y(t)\right].$$

Here, we omitted the state variables that determine the value of U, leaving the choice to implementers. State variables should be chosen so that the whole term structure is recovered from them. A choice is to use a set of futures prices as in the case of virtual storage valuation. But in this case, one has to deal with the correlation structure that is embedded in the model, which slows computation significantly. A more reasonable choice is to take the set of Ornstein-Ulenbeck processes that governs the spot price:

$$Z_i(s) = \int_t^s e^{-\lambda_i(s-u)}\, dW_i(u), \quad \text{for } i = 1,...,n\ .$$

Alternatively, one can use $M_i(s) = e^{\lambda_i s} Z_i(s)$, $i = 1,...,n$, which simplifies tree pricing by eliminating the drift in the state variables. In any case, U satisfies the following dynamic programming equation:

$$U(t,q) = \sup_c E_t\left[(1 - rdt)U(t + dt, q + cdt) - \left(cX(t) - \beta|c|\right)dt\right].$$

Using the Taylor expansion up to first order, we can rewrite the above equation as:

$$U(t,q) = E_t\left[U(t + dt, q)\right] + \sup_c\left[c\left(\frac{\partial U}{\partial q}(t,q) - X(t)\right) - \beta|c|\right]dt - rU(t,q)dt\ .$$

The rationale behind the above equation is that one should inject gas at the highest rate possible whenever the marginal value of storage exceeds the spot price. A partial differential equation of Hamilton-Jacobi-Bellman type can be obtained from above, depending on the choice of state variables. However we will not pursue this here, because the equation above itself is enough for implementing a finite difference scheme or tree.

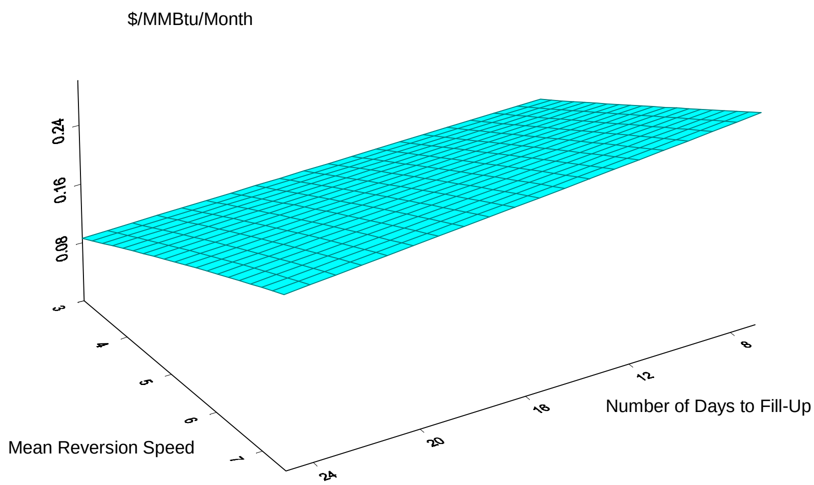


**Figure 3. Physical Storage Price**

In Figure 3, the value of physical storage (Mar 2001 – Feb 2002) is plotted against the number of days to fill the storage. The delivery charge is set to $0.1 per MMBtu. As before, a two-factor model with one flat factor and a fitted mean-reverting factor is considered.

**Remark**. It is worth noting that the valuation above provides an lower bound on the value of storage: if storage is available at a lower cost, such storage can be purchased and the above control strategy yields a risk free profit (an arbitrage). The converse does not work, as physical storage cannot readily be shorted.

# Conclusion

Gas storage has been drawing an unprecedented level of attention in recent times due in part to the remarkable spot price dynamics observed in North America in the past two years. To date, however, the market practice of storage valuation has relied heavily on *ad hoc* methods and is regarded as highly complicated task. The breakthrough presented in this work is formulating the valuation problem in the context of a solvable stochastic control problem originating from a flexible multi-factor description of forward curve dynamics. Using these methods the owner of storage can enhance the value of the facility substantially by utilizing optimal injection and withdrawal strategies.

# Reference


Ahn, H., A. Penaud and P. Wilmott (1999), Various Passport Options and Their Valuation, Applied Mathematical Finance

Danilova, A. (2000), Valuation of Physical and Synthetic Gas Storage Contracts, MSc Dissertation, University of Oxford

Clewlow, L. and C. Strickland (1999), A Multi-Factor Model for Energy Derivatives, Preprint

Juris, A. (1998), Competition in the Natural Gas Industry, Privatesector

Sturm, F. J. (1997), Trading Natural Gas, PennWell Publishing Company