Quantifying Natural Gas Storage Optionality: A Two-Factor Tree Model

Cliff Parsons*

Current Draft - August 3rd, 2011

Abstract

We find evidence of strong mean-reversion in U.S. natural gas prices and proceed to test a valuation model of natural gas storage leases based on mean-reversion's effects. The model utilizes a two-factor tree in which both factors mean-revert, and the model calibrates to current market conditions, accounts for volume constraints, and can be applied to historical data. In applying the model to data on U.S. natural gas, we find that the model can consistently capture large amounts of optionality based on the premise of strongly mean-reverting prices: On historical price data spanning 1999 to 2006, simulated trading using our model obtained average values (not including the cost of the lease) of \$1.244 per million British-thermal-units over intrinsic value for fast-cycle storage leases and \$0.397 per million British-thermal-units over intrinsic for slow-cycle leases. Results relating stored inventory to storage lease optionality are also shown.

* *Cliff Parsons, Ph.D., is President of WTM Energy Software, LLC., 6422 Cottonwood Park Lane, Houston, TX 77041. Emails: Inquire@WTMenergy.com; cliff.parsons@sbcglobal.net.* The author gives special thanks to Dr. Stephen Figlewski, Charles Riedhauser, and Dr. Thanassis Tjavaras for many helpful comments and also thanks Dr. Rafael Mendible, Chris Richards, Avill Young, Dr. Wei Zhang, and Dayne Zimmerman for discussions and comments regarding this paper.

1 Introduction

In the U.S. natural gas industry, the modeling and valuation of leases on natural gas storage have been major concerns, especially since the U.S. has several trillion cubic feet of such storage. Firms in the industry recognize these leases possess vast optionality but do not believe they correctly understand and extract it. Heretofore, other authors on this subject have provided some basic understanding of lease optionality, but none has quantified this value using historical data. The goal of this paper is to quantify and better understand the optionality of U.S. natural gas storage leases, and we have two major findings: (a) The mean-reversion in U.S. natural gas spot prices, *which is exploited by storage lease optionality*, appears to be much greater than previously documented, and (b) the actual optionality that could have been extracted from leases for previous years appears to be large.

The U.S. currently has over four-trillion cubic feet of working natural gas storage on annual consumption of over 22-trillion cubic feet. Depleted oil and natural gas reservoirs, accounting for over 80% of U.S. storage, provide the slowest injection/withdrawal rates while the two other main types of storage, salt caverns and aquifers, provide faster rates.[1] In short, depleted reservoirs, relative to the other types of storage, are filled with more sand or debris that inhibits the movement of natural gas stored within, thus slowing the rates of injection into and withdrawal from reservoirs. Thompson, Davison, and Rasmussen (2003) summarizes the main types of storage further.

Owners of storage facilities lease out space within, and a lease-holder has the right to inject into or withdraw from the facility only for a prespecified period of time, usually between each April 1st and the following March 31st, and within prespecified volume constraints, which are described through a "ratchet schedule." A ratchet schedule is a schedule of all possible inventory levels and their associated daily maximum injection and withdrawal rates.[2] As the lease-holder injects or withdraws, two types of transaction costs are typically incurred: a "fuel" charge, which is a percentage of injected or withdrawn gas, and a "commodity" charge, which is a dollar amount per unit of injected or withdrawn gas. These charges mostly exist to cover variable costs of operating the facility, with the chief cost being compressor operation for pushing more natural gas into or out of the facility.

Authors such as Manoliu (2004), Ludkovski and Carmona (2005), Chen and Forsyth (2006), and Boogert and De Jong (2008) posit, as

[1]One can find this information on the Energy Information Administration and the Federal Energy Regulatory Commission websites.

[2]Inventory levels for which these rates change from preceding levels are called "ratchets." Ratchets typically exist since injection rates get slower as facilities fill while withdrawal rates get slower as facilities deplete, just like filling and emptying a balloon.

do we, that storage lease optionality derives from the ability to exploit mean-reverting trends in natural gas spot prices: Only the opportunity to buy spot low, store it as prices mean-revert higher, then sell it high exists.[3] And short-term, current forward prices on natural gas provide information regarding the expected spot price mean-reversion. We believe, as argued in Mastrangelo (2007), that such mean-reversion derives mainly from daily demand characteristics. Natural gas is a major U.S. heating source and colder (warmer) than expected days typically cause demand to increase (decrease). Such weather shocks mean-revert and carry through to demand and spot prices. Seasonality in these prices occurs since U.S. demand is predictably greater in winter than in summer, which typically leads to higher winter forward prices for U.S. natural gas.

If mean-reverting spot trends help create value in storage leases, a natural question concerns the persistence of these trends. Equilibrium arguments concerning convenience yields, e.g., Hull (1997), suggest costless and abundant natural gas storage should eliminate any trends causing excess expected profitability, but natural gas storage is neither costless nor abundant. The fixed costs to develop a storage field can be over $10 million per billion-cubic-feet of capacity, and most storage facilities in the U.S. possess the slowest, most constrained injection and withdrawal capabilities. Furthermore, storage facilities are often sited based on geology and regulation, not demand.

The hypothesis that mean-reversion affects storage lease optionality implies the magnitude of that reversion is important. Studies such as Pilipovic (1998), Clewlow and Strickland (2000), and Benth and Benth (2004) find statistical evidence of weak to moderate mean-reversion in U.S. natural gas prices while Eydeland and Wolyniec (2003) statistically rejects the mean-reversion hypothesis. Parsons (2008) explains why results like these may be biased downward and do not necessarily indicate weak mean-reversion in reality, they may just indicate that the price model used in the estimation is too simplistic regarding the long-run mean. Intuitively, evidence of strong spot mean-reversion in the U.S. market can be seen every day in its forward curve: Forward prices for adjacent months can be over 15% different from each other, which is too much difference for spot prices to overcome in too little time with only slight mean-reversion. In Section 4 we summarize our method of fitting mean-reversion speeds. In short, we find evidence that *very strong* mean-reversion exists in U.S. natural gas prices when the model accounts for a more generalized long-run mean process.

For modeling natural gas storage value, tree models, such as the one in Manoliu (2004), are only one of three prevalent methods. An-

[3] Kjaer and Ronn (2008) claims that trading only forward contracts with storage in some cases can capture values close to such a spot trading strategy.

other method is Monte Carlo simulation, which is outlined in Eydeland and Wolyniec (2003) and Boogert and De Jong (2008). Monte Carlo simulation appears to be the most used of the three at this time in the energy industry. The third, more recent method is the stochastic control approach as seen in Thompson, Davison, and Rasmussen (2003), Ludkovski and Carmona (2005), and Chen and Forsyth (2006). To our knowledge, none of these valuation models has been previously back-tested on historical data for quantifying storage lease optionality.

For building our valuation model, a tree model that we may successfully back-test, we concentrate our efforts in two parts: (a) developing a realistic price model of the (mean-reverting) spot price process, and (b) developing a methodology for capturing all of the optionality in storage leases, given all the constraints and the spot price process. Each part is complex to solve, especially for ensuring that the price model can be calibrated to all market conditions seen in back-testing. We succeed in creating a two-factor tree model of valuation that brings together both parts so that back-testing could commence. The two factors, described in more detail in Section 2, can be thought of as representing a short-term weather effect and a longer-term effect of changing supply and demand forecasts.

In back-testing our valuation model we find it is very successful at consistently capturing vast amounts of optionality based on historical data. We tested the model on two types of storage leases: fast-cycle (the daily injection and withdrawal rates are such that the lease's maximum capacity can be filled and depleted in six cycles per year) and slow-cycle (the cycle-rate is only 1.5 per year). Using historical price data from 1999 to 2006, our simulated daily trading captured average values (not including the cost of the lease) of $1.244 per million British-thermal-units[4] (MMBtu) over intrinsic value for the fast-cycle storage lease and $0.397/MMBtu over intrinsic value for the slow-cycle lease.[5] Corresponding forecasts of those values from the model, before any simulated trading commenced, were $1.149/MMBtu and $0.402/MMBtu, respectively. Further, all values from simulated trading over those years were strictly greater than corresponding initial intrinsic values. We note one caveat about out price data: A time-asynchronicity exists within it, and obtaining data without such asynchronicity is extremely challenging. Our results and the effects of this asynchronicity are explained further in Section 5.

An interesting aspect of our modeled lease optionality concerns its

[4]One million British-thermal-units is approximately equal to 1,000 cubic feet of natural gas.

[5]"Intrinsic value" for storage leases is merely the value that can be risklessly obtained by buying and selling natural gas in the current forward market and using storage to hold any forward-purchased gas through to its subsequent forward sale.

dependence on inventory level. Secomandi (2010) discusses this aspect for simple storage leases. Specifically, correctly positioning the gas inventory to capture spot trends adds to the optionality; this implies storage leases possess follow-on optionality. For example, having zero inventory disallows one from withdrawing now and buying back later should spot prices shock up then trend down; however, being partly filled allows for that trading. Further, inventories associated with higher maximum injection and withdrawal rates, as seen in the ratchet schedule, allow for trading more volume as trends emerge, thus increasing optionality. The preceding suggests that storage lease optionality is typically, but not necessarily, greatest for inventories strictly between full and empty and is optimally extracted by trading spot daily as opportunities arise.

Our valuation model shows that, at any given time, these optimal inventory levels tend to group together into a pocket. Below that pocket, the maximum allowed injection at the current spot price is a positive net present value trade; above, maximum withdrawal is positive net present value. Within the pocket, no injection or withdrawal is recommended. Secomandi (2010) Theorem 1 proves this result for simple storage leases having no ratchets. The pocket is where the change in lease value for a change in inventory, up or down, is close to zero net present value at the given spot price. As the spot price moves, the pocket moves. Thus, for an inventory in the pocket, both directions of spot price movement, up or down, are likely to move the pocket to no longer include that inventory, which then leads to either injection or withdrawal being a positive net present value trade for that inventory. Whereas for an inventory outside the pocket, only one direction of movement increases the corresponding positive net present value for the optimal trade associated with that inventory while the other direction decreases it. Optimal storage trading consists of constantly trading inventory in the direction of the pocket, which is not only a positive net present value trade itself, but also increases the probability that either direction of subsequent spot price movement generates another positive net present value trade.

An intriguing result of our model is that multiple pockets of high-optionality, caused by ratchets, may exist at any one time. This behavior occurs when advantages exist for being on either side of the ratchet. For example, having a higher daily withdrawal rate above a ratchet while having a higher daily injection rate below that ratchet allows for trading higher volumes on either side of the ratchet as particular trends emerge. We illustrate this behavior in Section 5.

The paper proceeds as follows. Section 2 gives the mean-reverting, two-factor spot process used in valuing storage leases, including the reasons for choosing such a process. We also detail the backward recursion methodology used for valuing storage leases, which captures lease optionality for an arbitrary pricing tree while accounting for various trading constraints. Section 3 summarizes the calibration and delta-hedging procedures for the model. Section 4 details our estimation of mean-reversion speeds used in the price model; the estimates give evidence of heavy mean-reversion in U.S. natural gas spot prices. Section 5 gives back-testing results that quantify how much optionality may actually exist in storage leases. The section also details an example from the model relating optionality and inventory. Section 6 concludes with a brief discussion of future research on this topic.

2 A TWO-FACTOR MODEL OF STORAGE LEASE VALUE

U.S. natural gas price processes are complicated as seen in Exhibit 1, which shows daily spot prices[6] from 1999 through 2006 for natural gas delivered at Henry Hub, Louisiana. In this exhibit, spot prices appear to strongly revert toward a long-run mean, which itself appears to wander, but mean-revert. Specifically, these prices commonly fluctuate with over 3% daily volatility, which is quite high; yet, they do not tend to wander in a random walk like stock prices, they tend to cluster together around a price level that slowly wanders within a band of price levels. Further, prices exhibit seasonality, tending to be higher in winter months and lower in summer months.

As previously stated, any good valuation model of natural gas storage leases must (a) account for this complicated mean-reversion and (b) value all the optionality given the volume constraints and this mean-reversion. Our strategy is to develop a two-factor tree model since two factors do well at explaining U.S. natural gas forward price movements (explained just below) and tree models do well at capturing the American-type optionality that storage leases possess. In this section, we first examine a candidate continuous-time price model of spot and forward prices, which is shown to possess very realistic characteristics. Because of this realism, we extend this candidate model into a discrete-time version to be used in a pricing tree. Finally, we detail the backward recursion method for valuing storage optionality using such a tree.

The continuous-time spot process

Principal components analysis routinely shows that two factors of risk explain approximately 95% of movements in U.S. natural gas forward prices. Since forward prices and expected spot prices are related, at

[6] All price data throughout this paper come from Norman's Historical Data.

least theoretically, a two-factor price model is a natural starting point for the spot price process. For our spot process, we settle upon a price model that is close to the one in Pilipovic (1998). In their model they assume a mean-reverting spot price in which the long-run mean is a geometric Brownian motion. We assume a two-component long-run mean with one component being a mean-reverting process, the other being a deterministic process. Such a design appears to better match the data as we show in Section 4. Our price model is as follows:

$$\frac{dS_t}{S_t} = a(ln(L_t) + \mu_t - ln(S_t))dt + \sigma_{S,t}dz_t \tag{1}$$

$$\frac{dL_t}{L_t} = b(ln(\mathcal{L}) - ln(L_t))dt + \sigma_{L,t}dw_t \tag{2}$$

where

S_t	=	gas-daily (spot) price at time t
L_t	=	time t stochastic component of the long-run mean
μ_t	=	time t deterministic component of the long-run mean
$\mathcal{L}$	=	long-run mean of the L_t process
a	=	mean-reversion speed of the S_t process
b	=	mean-reversion speed of the L_t process
$\sigma_{S,t}$	=	time t deterministic volatility of the S_t process
$\sigma_{L,t}$	=	time t deterministic volatility of the L_t process
z_t	=	independent Brownian motion of the S_t process
w_t	=	independent Brownian motion of the L_t process

Before solving this model, intuition about it must be imparted. The choice of a mean-reverting spot model is intuitive since natural gas spot prices appear to do just that. However, the volatility term-structures of forward prices from models in which the long-run mean is solely deterministic (i.e., one-factor models) show volatilities diminishing too rapidly in the forward contract expiration relative to reality. Further, all forward and spot prices from a model in which the long-run mean is deterministic have perfect instantaneous correlation, typically leading to only near-parallel shifts in modeled forward curves from day to day.

The introduction of a stochastic component to the long-run mean as given in (2) overcomes the shortfalls of the one-factor price model and appears more realistic from our analysis on mean-reversion. Obviously, this enhanced setup creates spot and forward prices with less-than-perfect instantaneous correlations, so forward curves can "twist" as well as shift parallel in the model. And (2) also introduces more volatility into long-term forward prices, thus mitigating the steep decline in volatility term-structures from one-factor price models. Further, the enhanced setup's correlation structure has very desirable properties that mimic observed forward price correlations. Schwartz and Smith (2000) and Manoliu (2004) choose similar two-factor price processes for similar reasons. More on our setup is discussed when forward prices from (1) and (2) are introduced.

The introduction of a deterministic component to the long-run mean as given in (1) allows seasonality to be incorporated into the spot process. Deterministic, time-varying volatilities in both (1) and (2) allow for the incorporation of seasonality there, as well. Seasonality in U.S. natural gas forward prices and their volatilities exists and is persistent. These deterministic components help us capture seasonality and facilitate calibration. More on that when calibration is discussed.

Economic intuition regarding the factors is as follows. The shocks and mean-reversion of the S_t process can be surmised as being driven by weather: As current temperatures depart from expectations, the spot price gets shocked, but mean-reversion dissipates those shocks over time. The shocks of the L_t process can be surmised as being driven by supply and demand: As firms adjust their forecasted long-term commodity needs, they tend to exert pressure on spot prices for long periods, which is induced in the model by a stochastic shift in the long-run mean. If these intuitions are correct, then the independence assumption regarding these shocks is probably realistic. Similar arguments are given in Schwartz and Smith (2000).

In solving the spot model, we first apply Ito's lemma to (2) to find the solution for $ln(L_t)$, which is then inserted into (1) to solve (applying Ito's lemma again) for the spot price process.[7] For $T > t$,

$$S_T = S_t^{e^{-a(T-t)}} L_t^{\frac{a}{(a-b)}\left(e^{-b(T-t)}-e^{-a(T-t)}\right)} \cdot A \cdot B \cdot C \cdot D \qquad (3)$$

where

$$\begin{aligned}
S_T &= \text{the gas-daily (spot) price at time } T > t \\
A &= e^{ln(\mathcal{L})\left(1-e^{-a(T-t)}-\frac{a}{(a-b)}\left(e^{-b(T-t)}-e^{-a(T-t)}\right)\right)} \\
B &= e^{\int_t^T ae^{-a(T-t')}\left(\mu_{t'}-\int_t^{t'} e^{-b(t'-t'')}\sigma^2_{L,t''}/2\,dt''\right)dt'} \\
C &= e^{\int_t^T ae^{-a(T-t')}\left(\int_t^{t'} e^{-b(t'-t'')}\sigma_{L,t''}dw_{t''}\right)dt'+\int_t^T e^{-a(T-t')}\sigma_{S,t'}dz_{t'}} \\
D &= e^{-\int_t^T e^{-a(T-t')}\sigma^2_{S,t'}/2\,dt'}
\end{aligned}$$

Obviously, (3) leads to a process affected by both mean-reverting and seasonal trends, trends that can be valuable for holders of natural gas storage leases. We next show that forward prices derived from (1) and (2) possess realistic characteristics.

[7]Concepts from Oksendal (2000) were extremely helpful in solving these equations.

The derived forward price process

Heretofore, we have not discussed the price model in the context of any probability distribution, real or risk-neutral, and we must relate spot and forward prices in this model for calibration and hedging purposes. In a risk-neutral world, forward prices are merely expected spot prices. For natural gas, a hard-to-store commodity in which the no-arbitrage arguments in Hull (1997) do not apply, we believe such behavior exists in the real world as well. Our reasoning is as follows.

Just as producers are long natural gas, consumers (e.g., some manufacturers and anyone with a gas appliance) are short gas. Under risk-aversion and large storage frictions, the former group is willing to accept a discount to expected spot prices for selling forward while the latter is willing to pay a premium for buying forward. Since both groups are fairly well-dispersed and highly competitive, we expect neither has bargaining power over the other in the long-run; thus, we reason that if any risk premium/discount in forward prices exists, it should be small.[8] And if forward prices are martingales in the real world, then the real world and risk-neutral world coincide. Such a situation is convenient when using real-world historical data for any parameter fitting regarding the price model.

In the price model, calibration is performed in a risk-neutral setting; thus, we let $E_t()$ denote the time t conditional expectations operator under the risk-neutral distribution. Also, let $F_{t,T}$ denote a forward price at time t and expiring at time $T > t$. In the risk-neutral world, the following must hold: $F_{t,T} = E_t(S_T)$. Taking the time t expectation of (3) and applying Ito's lemma yields the following result:

$$\frac{dF_{t,T}}{F_{t,T}} = e^{-a(T-t)}\sigma_{S,t}dz_t + \frac{a}{(a-b)}(e^{-b(T-t)} - e^{-a(T-t)})\sigma_{L,t}dw_t \quad (4)$$

The first expression on the right-hand side of (4) is the one-factor representation of forward price changes under mean-reversion. To this term is added another term involving the second factor, and the significance of this factor increases in forward contract expiration for a large and b small. Thus, for a large and b small, (4) mitigates the steep decline in volatility term-structures that one-factor mean-reverting price models possess. Specifically, our volatility term-structure for forward prices starts high, reduces quickly in expiration, then reduces slowly in expiration as the significance of the second term on the right-hand side of (4) increases. This description matches closely the observed volatility term-structures for U.S. natural gas forward contracts.

[8] Eydeland and Wolyniec (2003) show empirical evidence of zero drift in actual power forward prices.

From (4) one may derive correlations of forward prices to one another that are also realistic.[9] Again for a large and b small, instantaneous correlations of log-returns among very long-term contracts go to one, and correlations between near-term and long-term contracts are less than one.

Evidence for a large and b small is presented in Section 4: The spot appears to revert heavily towards its long-run mean in the short-run, and that mean appears to wander over time with slight dampening from mean-reversion.

The price model given in (1) through (4) has many appealing characteristics, so it is chosen to extend into a discrete-time process for modeling storage lease values and accompanying American optionality.

The discrete-time price process

Previously, we developed a continuous-time model of spot and forward prices that possesses realistic characteristics. This section extends that model into the discrete-time process to be used in a pricing tree. The tree is used to value natural gas storage leases and accompanying optionality. The discrete-time price model is as follows.

$$ln(S_{t+\Delta t}) = e^{-a\Delta t} ln(S_t) + (1 - e^{-a\Delta t})(ln(L_t) + \mu_t) - \frac{\eta^2_{t,\Delta t}}{2} + \widetilde{z_{\Delta t}} \quad (5)$$

and

$$ln(L_{t+\Delta t}) = e^{-b\Delta t} ln(L_t) + (1 - e^{-b\Delta t}) ln(\mathcal{L}) - \frac{\xi^2_{t,\Delta t}}{2} + \widetilde{w_{\Delta t}} \quad (6)$$

where

$$\begin{aligned}
\Delta t &= \text{the time-step in years} \\
\eta^2_{t,\Delta t} &= \sigma^2_{S,t} \tfrac{(1-e^{-2a(\Delta t)})}{2a} \\
\xi^2_{t,\Delta t} &= \sigma^2_{L,t} \tfrac{(1-e^{-2b(\Delta t)})}{2b} \\
\widetilde{z_{\Delta t}} &\sim N(0, \eta^2_{\Delta t}) \\
\widetilde{w_{\Delta t}} &\sim N(0, \xi^2_{\Delta t})
\end{aligned}$$

Such a model requires a two-step approach for propagation: Step one holds both L_t and μ_t constant and shocks $ln(S_t)$ as given by (5), step two shocks $ln(L_t)$ as given by (6), and the process repeats. This process is analogous to its counterpart in (1) and (2) where L_t is fixed for the next instantaneous movement in the S_t process, then

[9] Pilipovic (1998) shows historical volatility term-structures and correlations.

L_t moves instantaneously, and the process repeats. Equations (5) and (6) become the basis for our pricing tree, and if Δt is small enough, the discrete-time process will approximate the continuous-time process well.

The tree is three-dimensional across the spot price, the long-run mean, and time with trinomial branches for each of the first two dimensions. The tree is constructed such that the trinomial transition probabilities through time for the S_t process satisfy the following set of equations:

$$E_t(ln(S_{t+\Delta t})) = \Sigma_{i=1}^{3} p_i \big(ln(S_{i,t+\Delta t})\big) \tag{7}$$

$$\eta_{t,\Delta t}^2 = \Sigma_{i=1}^{3} p_i \big(ln(S_{i,t+\Delta t}) - E_t(ln(S_{t+\Delta t}))\big)^2 \tag{8}$$

$$1.0 = \Sigma_{i=1}^{3} p_i \tag{9}$$

where the first two left-hand side expressions come from (5) and (6), while the right-hand side expressions come from the three associated nodes in the tree that branch from the corresponding S_t value. A similar set of equations exist for the L_t process.

Storage contract valuation and the backward recursion methodology

Backward recursion is an excellent method for capturing American optionality. In this section, we discuss the backward recursion used to value the complicated, follow-on optionality in storage leases.

In moving backward through a pricing tree, we calculate storage lease values based on optimal trading and price-taking behavior at each node. Specifically, for each node in the tree, and for each inventory that could possibly be held at each node, we proceed by first calculating two values for each such node-inventory combination: (a) - the value of immediately trading a volume of natural gas at that node's spot price, ensuring that the volume is within the volume constraints pertaining to the inventory; (b) - the continuation value of holding the new gas inventory going into the next period. Needing the continuation value is what drives the use of backward recursion. The optimal spot trade for each such node-inventory combination is the trade that maximizes the sum of these two values (we show this procedure mathematically further below).

The following example will help. Assume we are at some future node in the tree, we are only allowed to trade in increments of 5,000 units if we trade, and the inventories possible at that node are zero units, 5,000 units, and 10,000 units. At that node, we calculate the optimal trades and values for each possible inventory. We start with zero inventory and consider three possible trades: (a) - buy and inject nothing and receive the discounted expected value of carrying zero inventory into next period, (b) - buy and inject 5,000 units and receive the discounted expected value of carrying 5,000 units inventory into next period, and (c) - buy and inject 10,000 units and receive the discounted expected value of carrying 10,000 units inventory into next period. Of the three possible trades, the node's optimal trade for an inventory starting at zero is the trade where the sum of the immediate buy and its corresponding inventory-carryover value is maximized. The same procedure is followed for the other possible inventories at that node, 5,000 and 10,000 units, except that possibly withdrawing and selling is involved with those inventories.

The preceding implies that we must calculate a whole vector, indexed by inventory level, of these optimal values at each node. Storage leases complicate this optimization since both the set of possible inventories held and spot volumes traded are continuous. In the valuation model we approximate continuous inventory levels and trading volumes using a grid of inventory levels ranging from empty to the lease's maximum capacity, and we impose that the possible set of spot trades consists only of those volumes that change inventory from one grid level to another while adhering to all volume constraints in the ratchet schedule. If the grid is fine enough, we will approximate trading continuous values well.

We now state this optimization mathematically. Let each node in our two-factor tree be denoted by a triple of indexes spanning time, a spot price level, and a long-run mean level, respectively. We use the following notation:

$$
\begin{aligned}
X &= \text{the discrete set of inventory levels} \\
Q_x &= \text{the permissible set of trades at the given} \\
&\quad\ \text{inventory level } x \in X \\
q &= \text{a spot trade volume that is positive for} \\
&\quad\ \text{buys (injections) and negative for sales (withdrawals)} \\
S_{t,i,j} &= \text{the spot price associated with being at node } (t,i,j) \\
r_t &= \text{the riskless rate over time } t \\
f &= \text{the fuel transaction charge in percent} \\
c &= \text{the commodity transaction charge in dollars} \\
E_t() &= \text{the time } t \text{ conditional expectations operator} \\
|q| &= \text{the absolute value of } q
\end{aligned}
$$

The optimal value at node (t,i,j), and inventory x is

$$V_{(t,i,j)}(x) = \max_{q \in Q_x} \left(V^{spot}_{(t,i,j)}(q) + V^{continue}_{(t,i,j)}(x+q) \right) \tag{10}$$

where

$$
\begin{aligned}
V_{(t,i,j)}(x) &= \text{the optimal value at node } (t,i,j) \\
&\quad\ \text{with current inventory } x \in X \\
V^{spot}_{(t,i,j)}(q) &= \begin{cases} |q|\left(-S_{t,i,j} - \frac{f}{(100-f)}S_{t,i,j} - c\right) & \text{for buying} \\ |q|\left(S_{t,i,j} - \frac{f}{(100-f)}S_{t,i,j} - c\right) & \text{for selling} \end{cases} \\
V^{continue}_{(t,i,j)}(x+q) &= \frac{E_t(\widetilde{V}_{(t+1,.,.)}(x+q))}{(1+r_{t+1})}
\end{aligned}
$$

Two possible value functions for $V^{spot}_{(t,i,j)}(q)$ exist since one is for buying with paying both the fuel and commodity charges while the other is for selling with paying those charges. The fuel charge is paid on a grossed-up volume such that the desired net volume is delivered once the fuel charge is subtracted. The tilde above the function $V_{(t+1,.,.)}(x+q)$ for the continuation value merely denotes that lease value at time $t+1$ is random at time t. The two dots embedded in the node indexes of that function denote that expectations are to be taken over several nodes at time $t+1$.

The recursion begins at the final-period nodes where all values for continuation in (10) are zero; thus, the optimized value for each inventory in the inventory vector at each of these nodes is merely the value of withdrawing and selling as much inventory as is allowed. To complete the lease valuation, one merely has to recursively perform (10), going backward through each node in the tree, until the beginning period is reached. The value of our lease is the value corresponding to the current inventory level in the inventory vector at the beginning period. This same recursion method for valuing storage leases is shown in Manoliu (2004) and Boogert and De Jong (2008) and is used in the Clewlow-Strickland one-factor storage model. This method is also an extension of the recursion for similar, but less complicated contracts, called swing options, used in Jaillet, Ronn and Tompaidis (2004).

3 CALIBRATION AND HEDGING

In order to calibrate the discrete-time spot process, a bounded, trinomial tree-type structure, described earlier, is imposed upon the price model. The price model itself needs three parameter values prior to calibrations: a, b, and a long-to-short term volatility ratio described just below. Estimates of these values are given in the next section. Once calibration is performed, then derivative contracts may be valued, including forward contracts which are used for hedging. We now examine calibration, then hedging.

Calibrating the discrete-time price process

The major assumptions imposed on the pricing tree are sequentially listed below with discussion of them given after.

Assumption 1 *The deterministic volatilities in (5) are step functions with one step per forward month.*

Assumption 2 *The deterministic volatilities in (6) are merely the corresponding volatilities in (5) multiplied by a constant k, which we call the long-to-short term volatility ratio.*

Assumption 3 *The deterministic component of the long-run mean, μ_t, given in (5) is a step function with one step per forward month.*

Assumption 4 *The following values are assumed for the price model.*

$$\begin{aligned} L_0 &= \textit{The current spot price} \\ \mathcal{L} &= \textit{The current spot price} \\ \Delta t &= \tfrac{1}{365},\ \textit{(one-day)} \end{aligned}$$

Assumptions 1 and 3 exist for calibration: They allow one to calibrate to prices on forward contracts for monthly natural gas delivery, one contract per calendar month, and their volatilities. These forward contracts are prevalent in the U.S. natural gas market and do not exist with finer granularity than monthly. The assumed step functions can result in large changes in the spot process, which is unintuitive. However, these large changes do occur. For example, the forward price for natural gas delivered in December at most any U.S. location can easily be 15% greater than that for delivery in the preceding month.

Assumption 2 aids our effort by halving the number of volatilities needed to calibrate the price model.

Assumption 4 shows three values used in the price model; they enter the model as constants. The third value corresponds to how spot gas is traded in the US: once per day. The first two values are chosen to fix the beginning-state long-run mean, L_0, and start its process with no inherent trend from mean-reversion. Any error in our beginning value for L_0 is ameliorated from calibrating μ_1 to the first forward price, and starting the process with no trend is reasonable since the direction of any inherent trend is unknown to us. The remaining model parameters, a, b, and k were fitted to our historical data as given in Section 4.

The tree is calibrated to match monthly forward prices as seen in the market, starting with the rest-of-month price. This calibration approach is similar to the one in Jaillet, Ronn and Tompaidis (2004) and the Hull-White approach as given in Hull (1997).

The tree is also calibrated to a term-structure of observed volatilities on monthly forward contracts. Since the volatilities needed in the price model are forward volatilities commensurate with a two-factor, mean-reverting price process, the term-structure of observed volatilities is converted accordingly. The conversion is such that total variance to each contract month is the same when calculated from either the observed volatility term-structure or from our model using the converted volatilities associated with it.

We note that our full calibration is not an absolute or squared errors minimization as in Jaillet, Ronn and Tompaidis (2004) and Chen and Forsyth (2006), it is exact.

Hedging storage contract values in discrete-time

In our price model, spot prices are driven by two factors. Subsequently, forward prices and other derivatives are driven by two factors. Since a plethora of traded forward contracts with different expirations exists, any two portfolios of them, except for linearly dependent portfolios, may be chosen to dynamically hedge storage leases. At least this is the way delta-hedging works in continuous-time. However, we have a discrete-time tree in which each node branches into nine other nodes. How may we calculate deltas in this framework?

We use a least-squares approach in which two portfolios of forward contracts are selected for hedging purposes. Specifically, lease value-changes between today and tomorrow from the tree are regressed on value-changes of the two portfolios from the tree. The betas of the two portfolios are the hedging deltas. Since we use only two portfolios of (linear) forward contracts to hedge over nine subsequent nodes of (non-linear) lease values, perfect hedging will not result. However, since ours is a two-factor price model, using two portfolios should be sufficient to hedge away most lease value risk in the model. Back-testing on historical data will confirm whether or not this form of hedging works well in reality.

4 EVIDENCE OF SPOT PRICE MEAN-REVERSION AND PARAMETER ESTIMATES

The effects of spot price mean-reversion cannot be underestimated in valuing storage leases. Mean-reversion creates trends that can generate optionality, adding to the value of storage leases, especially for large magnitudes of mean-reversion. Thus, estimating mean-reversion is important.

Evidence of mean-reversion in U.S. natural gas spot prices is well documented: Prices typically trade within upper and lower bounds; volatility term-structures of forward prices are such that volatilities decrease in forward contract expiration. The former result indicates bounded variances over time, which is consistent with mean-reversion. The latter result indicates that shocks dissipate over time, which occurs with mean-reversion.

Many authors model natural gas prices as mean-reverting, but few have measured its magnitude, and none has done so in the two-factor setting of our price model. The evidence so far is as follows: Pilipovic (1998), Clewlow and Strickland (2000), and Benth and Benth (2004) find that natural gas prices propagate with low to moderate mean-reversion speeds (approximately between 2.0 and 39.0) while Eydeland and Wolyniec (2003) finds no statistically significant mean-reversion. Intuitively, evidence of strong spot mean-reversion in the U.S. market exists almost every day in its forward curve: Forward prices for some adjacent months can be over 15% different from each other, which is too much difference for spot prices to overcome in too little time with only slight mean-reversion. We find evidence of large spot price mean-reversion, detailed below, and we contend that low and insignificant mean-reversion speed estimates in the other studies could be caused by an estimation bias documented in Parsons (2008). This bias occurs in commodities like natural gas in which the long-run mean of prices varies either stochastically or deterministically. Simple estimation procedures, such as the ones used in previous studies, confuse a varying long-run mean with weak mean-reversion, thus leading to low mean-reversion speed estimates even when true mean-reversion is high.

U.S. natural gas prices indicate the presence of a varying long-run mean, e.g., the winter/summer price seasonality, which must be treated carefully when estimating as discussed above. We choose the following approach: Use our two-factor pricing tree, which incorporates a stochastic long-run mean, and previous years of prices to simulate daily trading and hedging on a simple option contract, and we vary the model's parameters, which include the mean-reversion speeds, to best fit initial forecasted trading values with values obtained from simulated daily trading.

The price data are for Henry Hub, consist of both gas-daily (spot) and forward prices, and are for every trading day spanning April 1999 to March 2006. That time span captures the full range of price and volatility regimes. For each "gas-year," defined as April 1st of a chosen year to the following March 31st, we obtain a forecast value and a simulated trading value (including hedges) using our pricing tree applied to a particular natural gas option contract. This simple contract is similar to a storage lease in many respects, but was chosen since, for parameter searches, it is much faster to value than storage leases.[10] The option's forecast value is merely the model's option value calculated on the first day of the gas-year; the simulated trading value is obtained by running the model on each day's historical prices during the gas-year, following the model's daily recommendations for exercising and hedging, and discounting all generated cash flows to the first day of the gas-year and summing.

The two mean-reversion speeds and the long-to-short term volatility ratio, k, are varied to find the least-squares best fit of simulated (dollar) trading value to forecast value. One of those speeds is the mean-reversion speed that the previously cited papers attempt to estimate. Our estimate of this speed, a, our short-term mean-reversion speed, is 87 and is from two to over 35 times greater than what the previous studies document. The estimate of b, our long-term mean-reversion speed, is 0.6. Our estimate of k is 0.9. The two-factor price model in Pilipovic (1998) is essentially a special case of our price model when b is zero, i.e., the long-run mean propagation is geometric Brownian motion. Our estimates indicate that refining the Pilipovic (1998) price model to include a mean-reverting long-run mean appears warranted.

5 RESULTS

In this section we give two sets of results: one for back-testing, and the other for relating optionality on storage leases to storage inventory.

Back-testing

Our valuation model of storage leases was back-tested against historical gas-daily and forward prices for delivery at Henry Hub, Louisiana, spanning 1999 to 2006. Prices were obtained for each business day over that period and are in dollars per MMBtu. The time span was split into seven sub-periods, each one year in length, with sub-periods starting every April 1st and ending on the following March 31st. These sub-periods coincide with the typical terms in the U.S. on one-year natural gas storage leases. Data from before 1999 are not considered since the market structure before then was in transition.

For each sub-period the model was run every business day. Each day's gas-daily price and forward curve were input to the model along with fuel charges, commodity charges, and volume constraints. Our fuel and commodity charges were held constant over the whole sub-period, were typical in magnitude for North America, and were as follows: Injection fuel charge was 1.50%, withdrawal fuel charge was zero,

[10] The option contract is a "swing" put option having 15 exercises and fixed, at-the-money strike prices. See Jaillet, Ronn and Tompaidis (2004) for valuing such contracts.

and both injection and withdrawal commodity charges were \$0.01/MMBtu. No other transactions costs were assumed. Also, only one ratchet was in effect and occurred when the inventory reached half-full.

Other pertinent facts are as follows. First, two hypothetical storage leases were back-tested: One was slow-cycle (cycle-rate of 1.50), and one was fast-cycle (cycle-rate of 6.00). Second, rest-of-month contract prices were not available, but were estimated each day to be a simple average of that day's gas-daily and next-month-forward (prompt-month) prices. Third, volatilities for forward contracts with tenors spanning the sub-period were figured each day using the most recent 30 days of historical log-returns on those contracts and were converted to the volatilities, $\sigma_{S,t}$ and $\sigma_{L,t}, \forall t$, used in the pricing tree.[11] Fourth, our estimates of both mean-reversion speeds, the estimate for the long-to-short term volatility ratio, and the values in Assumption 4 were used. And fifth, for hedging six forward contracts were chosen at the beginning of the sub-period, three with the nearest-term tenors and three with tenors expiring in the last three months of the sub-period, the two middle contracts were successively dropped through the back-test as the diminishing time to the sub-period's end-date would not allow for more, and the volume on the contracts was chosen such that inventory plus long positions equaled short positions. All long position tenors had the same volume, all short positions did, as well, and the total of existing inventory plus long positions did not necessarily equal the total storage capacity of the lease: The total volume was chosen each day to minimize the variance of the daily dollar-change, as given in the pricing tree, for storage value with hedges.[12]

All input was fed to the valuation model each business day during a given sub-period, and the model returned recommended injections, withdrawals and hedges for the day. All recommendations were taken, and all subsequent cash flows over all days were discounted to the sub-period's beginning and summed. This value represents the simulated trading value over that sub-period by using the model. For comparison, the storage value returned from the model at the sub-period's beginning represents the forecast of such trading value over the sub-period. Our hope is that these two values are "close." Lastly, the intrinsic value for the sub-period was calculated at the sub-period's beginning as the maximum value, net of fuel and commodity charges,

[11] Parameter $\sigma_{S,t}$ was also constrained to be greater than or equal to 25% annualized. Doing so aids calibration, hastens run-times, yet is not terribly constraining since these values tend to be much higher.

[12] Since price movements on U.S. natural gas forward contracts are extremely similar for tenors of two months out and beyond, a wide variety of hedging tenor choices will lead to similar back-test results. And we again note for emphasis that in the model (and apparently in reality) forward contracts do not add to value in expectation, they only hedge it.

that could be locked in using forward contracts having tenors spanning the sub-period. Our hope is the simulated trading value consistently beats the intrinsic value. A summary of these results, which are given as dollars per MMBtu, is given in Exhibit 2.

Several conclusions about the results are apparent from the exhibit. First, the average extrinsic values captured (i.e., the average simulated trading value minus the average intrinsic value) are \$1.244/MMBtu and \$0.397/MMBtu for fast-cycle and slow-cycle leases, respectively. Extrinsic value is the optionality captured. Further, every simulated trading value beat corresponding intrinsic value, which indicates optionality is consistently captured by the model. Also, the standard deviation of the difference of simulated and forecast trading values, assuming the expectation of that difference is zero, is \$0.619/MMBtu for fast-cycle and \$0.188/MMBtu for slow-cycle; considering the average forecast values for fast-cycle and slow-cycle are \$1.473/MMBtu and \$0.587/MMBtu, respectively, the standard deviations indicate the model has reasonable forecasting capabilities. Lastly, the number of years the forecast trading value beat the simulated trading value is three out of seven for fast-cycle and five out of seven for slow-cycle. If our null-hypothesis is that forecasts should beat simulated trading 50% of the time, then a simple signs test does not reject this hypothesis at the 95% confidence level, even for a one-tail test. Thus, our forecast values do not appear biased. Admittedly, the sample size is small, seven outcomes for each of the two types of storage leases, but the results are encouraging so far.

Since the price model was developed to incorporate and capture mean-reverting spot trends, these results, along with our results in Section 4, give good evidence of strong mean-reversion in U.S. natural gas prices. Alternatively, if price shocks do not create subsequent trends, then trading based on trends would not lead to more profit over intrinsic in expectation.

The dynamic hedging suggested by the model is moderately successful; however, the simulated trading values can differ from their forecasts by quite a lot in some cases. Stochastic volatility, which is not assumed in the price model, seems to play a role in this.

We present these results and parameter estimates under the caveat that a time-asynchronicity exists in our price data: The gas-daily (spot) prices are realized several hours earlier in the day than are the forward prices. This asynchronicity is an artifact of how prices are collected and reported each day, and obtaining datasets without this problem is extremely challenging.

The effects of this asynchronicity, we suspect, are to cause higher mean-reversion speed estimates, forecasted trading values, and simulated trading values. One can see this by first noting that gas-daily prices are precluded from catching up to forward prices as the latter trade longer through the day. The next day, however, gas-daily prices would appear to "lunge" toward the forward prices realized from the day before, thus giving the appearance of strong mean-reversion. These "lunges" create larger, opportunistic trends for simulated trading to capture, so fitting forecast trading values to such simulated values requires higher mean-reversion speed estimates.

We leave the solution to this problem to future research.

Lease optionality versus storage inventory

As previously explained, one may exploit storage lease optionality by exploiting trends in spot prices, which implies that lease optionality is a function of inventory level: Being partly filled allows for exploiting both up and down trends; inventory levels in the ratchet schedule for which more daily volume can be traded allow for better exploitation of those trends.

Output from the valuation model shows that certain pockets of inventory levels possess higher optionality relative to other levels at any one time. As previously mentioned, these pockets consist of inventory levels for which no injection or withdrawal is recommended, but are poised to yield positive net-present value injections or withdrawals should the current spot price move in either direction. Below (above) those pockets, maximum daily injection (withdrawal) is almost always recommended. Non-zero injections or withdrawals that are less than their maximums infrequently occur, but appear to occur only on the boundary-inventories of those pockets and add little to overall value. The Secomandi (2010) theorem mentioned previously proves such behavior for simple storage leases. This behavior in recommendations implies that the model pushes one towards holding these high-optionality inventories. Exhibit 3 shows this behavior graphically for a lease slightly more complicated than the Secomandi (2010) theorem examines.

This exhibit graphs typical output from the model and shows how multiple pockets of high-optionality inventory levels may exist. The horizontal axis is inventory level in thousands of MMBtus; the vertical axis is the recommended daily injection/withdrawal in thousands of MMBtus with injections being positive and withdrawals being negative. The maximum inventory is one-million MMBtus, and one ratchet exists at 400,000 MMBtus: Below that ratchet, maximum daily injections and withdrawals are 12,000 MMBtus and 9,000 MMBtus, respectively; at or above that ratchet, the numbers are reversed. The exhibit shows that two pockets of high-optionality inventories exist: one just above 200,000 MMBtus, and one just above 400,000 MMBtus. The ratchet, which is the complication violating the assumptions of the Secomandi (2010) theorem, causes this since being below the ratchet level allows for greater daily injection relative to being above the level, while being above the level allows for greater daily withdrawal relative to being below the level. Thus, being on each side of the ratchet level has advantages.

Lastly, the placement of these pockets, which can change a lot from day to day, is sensitive to, among other effects, the shape of the forward curve, storage cycle-rate, and ratchets. And the sensitivity is not obvious since these effects compete with each other for placement of the pockets. We leave discussions on this sensitivity to further research.

Overall, lease optionality is very complicated relative to simple American options, but one observation is clear from repeatedly running our model: With standard American options, exercising decreases optionality; with lease optionality, injection and withdrawal typically occur to increase optionality.

6 CONCLUSION

We present a two-factor tree approach for valuing natural gas storage leases. Although other and similar valuation models have been proposed in the literature, none has been tested against historical data to quantify the optionality in storage leases. Simulated trading on our historical price data shows our model was successful at capturing the vast amounts of optionality forecasted by the model on both fast and slow-cycle leases.

The price model was developed under the assumption of a strongly mean-reverting spot price, an assumption only weakly backed heretofore. We give new evidence to support the assumption, and the valuation model was designed to optimally exploit it. Thus, the mean-reversion evidence we give along with the valuation model's success in consistently capturing vast optionality in our price data is more evidence of strong mean-reversion in U.S. natural gas prices.

However, our price data contain a time-asynchronicity, as previously mentioned, that may skew results. Thus, a direction for future research is to back-test using price data where all prices are essentially collected at the same time each day. Unfortunately, time-synchronized public price data is very hard to come by.

Another research direction is to assess the sensitivities and drivers of storage lease value and optionality. From our initial research, the optionality is not well-behaved: It is dependent on several drivers. Assessing drivers of optionality will lead to more intuition about optimal storage trading.

REFERENCES

Benth, F., and Benth, J. "The Normal Inverse Gaussian Distribution and Spot Price Modelling in Energy Markets." *International Journal of Theoretical and Applied Finance*, Vol. 7, No. 2, 2004, pp 177-192.

Boogert, A., and De Jong, C. "Gas Storage Valuation Using a Monte Carlo Method." *Journal of Derivatives*, Spring (2008), pp 81-98.

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

Chen, Z., and P. Forsyth. "A Semi-Lagrangian Approach for Natural Gas Storage Valuation and Optimal Operation." Working paper, David R. Cheriton School of Computer Science, University of Waterloo, 2006.

Eydeland, A., and K. Wolyniec. *Energy and Power Risk Management*. John Wiley & Sons, Inc., 2003.

Hull, J. *Options, Futures, and Other Derivatives*. 3rd ed., Prentice-Hall, 1997.

Jaillet, P., E.I. Ronn, and S. Tompaidis. "Valuation of Commodity-Based Swing Options." *Management Science*, July (2004), pp 909-921.

Kjaer, M., and E.I. Ronn. "Valuation of a natural gas storage facility." *Journal of Energy Markets*, Winter (2008), pp 3-22.

Longstaff, F., and E. Schwartz. "Valuing American Options by Simulation: A Simple Least-Squares Approach." *Review of Financial Studies*, Spring (2001), pp 113-147.

Ludkovski, M., and R. Carmona. "Gas Storage and Supply Guarantees: An Optimal Switching Approach." Working paper, Department of Mathematics, University of Michigan, 2005.

Mastrangelo, E. "An Analysis of Price Volatility in Natural Gas Markets." Energy Information Administration, August 2007.

Oksendal, B. *Stochastic Differential Equations: An Introduction with Applications*. 5th ed., Springer, 2000.

Manoliu, M. "Storage Options Valuation Using Multilevel Trees and Calendar Spreads." *International Journal of Theoretical and Applied Finance*, June (2004), pp 425-464.

Parsons, C. "Explaining Bias in Mean-Reversion Speed Estimates for Energy Prices." *Energy Risk*, July 2008, pp 74-79.

Pilipovic, D. *Energy Risk: Valuing and Managing Energy Derivatives*. McGraw-Hill, 1998.

Schwartz, E., and J. Smith. "Short-Term Variations and Long-Term Dynamics in Commodity Prices." *Management Science*, July (2000), pp 893-911.

Secomandi, N. "Optimal Commodity Trading with a Capacitated Storage Asset." *Management Science*, March (2010), pp 449-467.

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

EXHIBIT 1: Gas-Daily Spot Prices for Henry Hub Showing Mean-Reversion and Seasonality

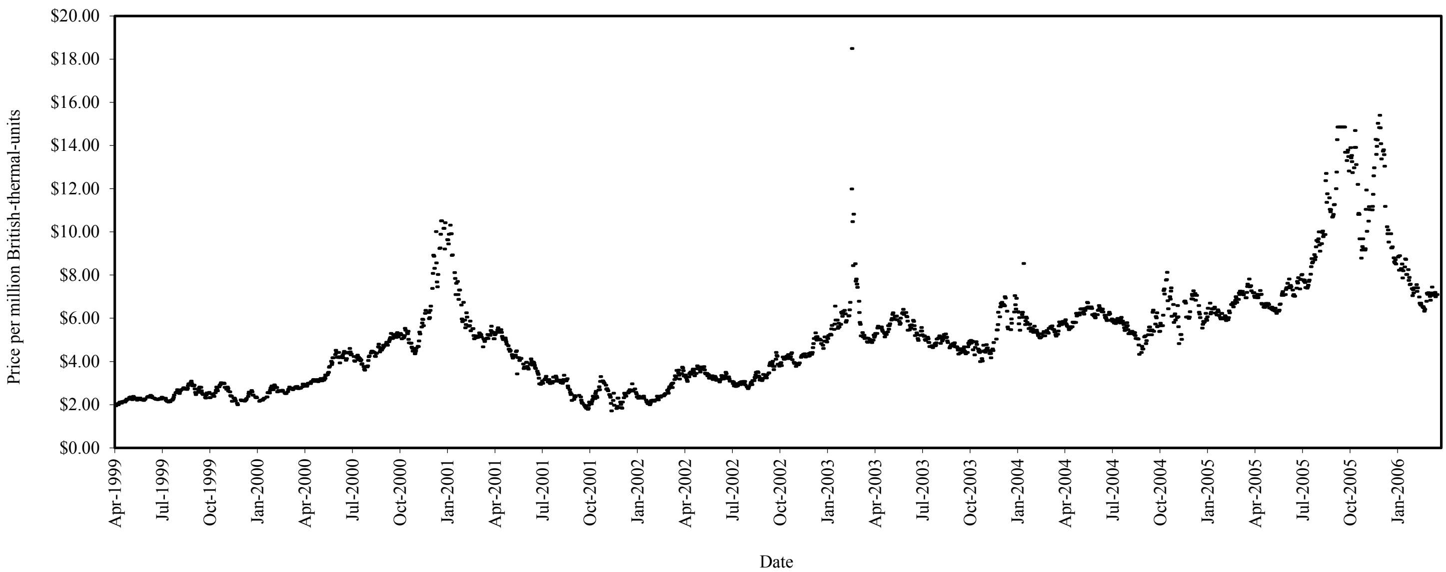

EXHIBIT 2: Model Back-Testing by Historical Year on Henry Hub

(Fast-Cycle Storage: 6-cycle)			
Lease Term	Simulated Trading Value	Forecast Trading Value	Intrinsic Value
1999 / 2000	0.513	0.667	0.393
2000 / 2001	1.157	0.504	0.130
2001 / 2002	1.084	1.982	0.111
2002 / 2003	1.600	2.107	0.418
2003 / 2004	1.681	1.265	0.157
2004 / 2005	1.644	1.466	0.274
2005 / 2006	3.302	2.320	0.788

(Slow-Cycle Storage: 1.5-cycle)			
Lease Term	Simulated Trading Value	Forecast Trading Value	Intrinsic Value
1999 / 2000	0.354	0.363	0.254
2000 / 2001	0.375	0.197	0.061
2001 / 2002	0.453	0.721	0.046
2002 / 2003	0.714	0.774	0.263
2003 / 2004	0.405	0.458	0.056
2004 / 2005	0.428	0.585	0.142
2005 / 2006	1.348	1.014	0.478

Values are in dollars/MMBtu. "Simulated Trading Value" results from the model's daily trade recommendations. "Forecast Trading Value" is the model's expected trading value from the beginning of the trading period. "Intrinsic Value" is the riskless value that could be locked in using forward contracts at the beginning of the trading period.

EXHIBIT 3: Example of Lease Optionality Versus Inventory Level

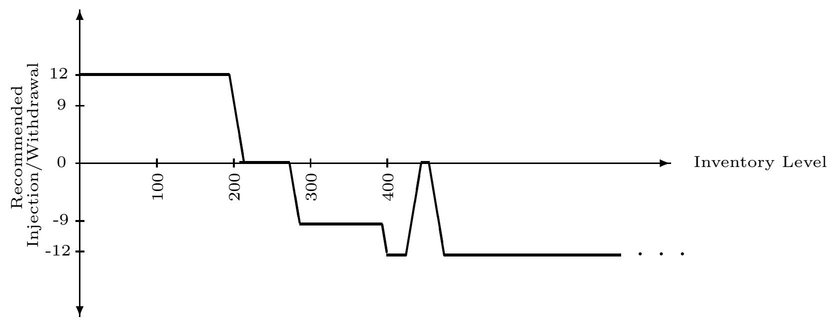

Both inventory and recommended injections/withdrawals are given in '000s-MMBtus. Injections are positive; withdrawals are negative. Inventories for which no injection or withdrawal is recommended represent high-optionality inventories.