A Primer on Natural Gas Storage Valuation

by

Christopher Athaide
Cygnet Risk Group, Ltd.

Disclaimer

This document is provided for information purposes only. The views expressed in this document represent those of the author and Cygnet Risk Group, Ltd. The author and Cygnet Risk Group, Ltd. do not warrant that any of the material presented herein is accurate or timely. This document is not intended to provide advice on any aspect of natural gas storage valuation. Readers relying on this document to make any investment or trading decisions are doing so at their own risk.

About Cygnet Risk Group, Ltd.

Cygnet Risk Group, Ltd. is a boutique firm offering software products and consulting services to energy companies worldwide. The company was founded in January 2000 and is based in Houston, Texas.

Since its inception, Cygnet has done work for a number of companies in the areas of natural gas storage valuation, forward curve modeling, value-at-risk and credit risk modeling, power plant valuation, power scheduling and energy derivative modeling.

Chapter 1: Introduction to Natural Gas Storage

Types of Storage

The three main types of underground storage currently being used are: (1) Depleted reservoirs in oil/gas fields, (2) aquifers, and (3) salt caverns. There is a fourth type of storage facility called LNG facility (sometimes also called peak shaving, but we shall not be discussing them in this book.) Table 1 provides a summary of the three main types of storage used.

The ultimate goal of this document is to help the reader understand the various issues that arise when one attempts to determine the value of natural gas storage. For the purpose of this discussion, we are not using the term value from the point of view of a purchaser or a seller of the facility. Rather, we are looking at it from the viewpoint of a lessee and a lessor, and the price that a lessee would be willing to pay a lessor for the use of the facility for a fixed time period. When we refer to the value of a Natural Gas Storage Facility (NGSF), we are making an implicit reference to a finite time interval over which this value is being realized. In other words, each time we say “the value of an NGSF is ...”, what we are actually saying is that “the value of an NGSF over the time interval specified by the start date and end date is ...”.

The easiest way to think about natural gas storage is to regard it as an option instrument, albeit a complex one. The demand charge reflects the payment for the use of the facility — this payment can be made upfront, or at periodic intervals. The demand charge represents the option premium. The value over the term of the lease is the expected amount of money that can be made by trading gas in and out of the facility. Financial theorists would argue that the premium should equal the value, otherwise there is an arbitrage. In reality, the premium is always less than the value (in North America, at least.) The participants in the natural gas market who own storage facilities typically do not trade storage for regulatory or other reasons. There are administrative and trading costs (and risk) associated with trading storage and so the premiums are lower than the theoretical value to reflect these factors.

For the storage trader, when deciding whether or not to lease a NGSF, the question to ask is: What is the value of the NGSF, and what is the associated premium? The value must be viewed relative to the demand charge before any decision is made regarding the lease.

The value of a NGSF is derived from a number of factors. The two key factors in determining this value are 1) the physical characteristics of the facility, and 2) the forward curves at the facility location.

Storage Facility	Depleted Oil/Gas Reservoir	Aquifer	Salt Caverns
Description	Depleted Reservoirs that used to contain oil/gas are modified for natural gas storage.	Underground natural water migration traps can be used for natural gas storage.	Cavity created in the ground by pumping fresh water into a salt formation and pumping out brine.
Characteristics	Partial infrastructure already exists. Geology already known. Gas may already be present to use as Base Gas. Base Gas is approximately 50% of Total Capacity. Shorter development time. Low Cycling Rates.	Expensive to develop. Geology unknown. Infrastructure must be installed. Base Gas is approximately 80% of Total Capacity. Longer development time. Low Cycling Rates. EPA requirements on groundwater contamination.	Expensive to develop. Smaller acreage compared to the other two types. Base Gas is approximately 20-25% of Total Capacity. Longer development time. High Cycling Rates. Environmental constraints related to brine disposal.
Cycling	Between 1-4 cycles/year	Between 1-4 cycles/year	Between 6-16 cycles/year

Table 1—**Different Types of Storage Facilities**

Storage Constraints

There are several restrictions that affect the ability of the trader to inject and withdraw gas from storage. Most of these are related to the geological characteristics of the facility. These restrictions manifest themselves in the form of different classes of constraints.

- **Daily Injection and Withdrawal Ratchets**

  The daily injection and withdrawal tables are rules that determine how much gas can be injected or withdrawn from the facility per day for a given level of inventory. The daily injection ratchet table is a nonincreasing function of inventory because as the level of inventory increases, the ability to inject gas into the facility decreases, or stays constant. The daily withdrawal ratchet table is a nondecreasing function of inventory because as the inventory level decreases, the ability to withdraw gas from the facilility does not increase.

- **Monthly Injection/Withdrawal Constraints.**

  In addition to daily injection and withdrawal constraints, there are some storage facilities that have monthly injection and withdrawal constraints. These facilities are typically reservoir or aquifer facilities, and in order for the constraints to be binding, the monthly injections and withdrawal limit must be less than the cumulative daily sums for injections and withdrawals, respectively. These constraints are usually necessary to maintain the geological integrity of the facility. In the case of aquifers, the gas has to be constantly cycled otherwise it moves to the periphery of the facility and becomes difficult to recover. Sometimes the facility operator decides that there will be no injections or withdrawals during a given month because the facility has to be overhauled, geological studies have to be conducted etc.

- **Time Based Inventory Constraints.**

  In this case, the constraints are imposed on the level of inventory at a particular point in time. For example, there are facilities that have the constraint that amount of gas in the facility cannot be more than 60% of the maximum capacity on June 30 of each year. Again, these constraints are imposed to maintain the geological integrity of the facility.

- **Fuel and Commodity Charges**

  There are two types of fuel charges: 1) injection fuel charges, and 2) withdrawal fuel charges.The same categories hold for commodity charges. A number of storage operators impose a charge on just one leg of the transaction, typically during injections. The idea is that if the lessee starts and ends the lease with the same level of inventory, the total injections must equal the total withdrawals. Administratively, it is simpler to account for this charge either during injection or withdrawal. (This is similar to toll collection systems at some bridges and tunnels. The toll is collected in one direction of the traffic.) Storage operators levy fuel and commodity charges to offset the energy and administrative costs associated with these trades.

  These charges are usually broken down into these two categories for historical reasons. In the United States, where most storage rates are regulated, the storage operator has to get permission from the FERC before a rate change can go into effect. The fuel charge is levied because of the energy needed to operate the compressors when the gas is pumped into the storage facility at high pressure. The fuel charge allows the operator to quote the cost as a percentage of the volume injected or withdrawn. By accepting the fuel payment in-kind, the operator is assuming less price risk and passing the risk on to the lessee. The volume of fuel required to operate the compressor is the same whether gas is priced at $2.00/mmBtu or $7.00/mmBtu. The commodity charge allows the operator to assign some of the costs that are less sensitive to the price of natural gas such as personnel salaries.

- **Ad Valorem Taxes**

  Certain jurisdictions impose a tax on the value of gas in the NGSF on a fixed date. In this case, the trading strategy would have to include this effect in the decision of the appropriate level of inventory on the Ad Valorem date.

The Value of Storage

In an economic system, where the supply and demand are subject to shocks or uncertainties, the presence of a storage (or inventory) facility serves as a buffer that helps to stabilize the system. It acts as a decoupling agent that insulates the supply side from the demand side and mitigates shocks in the system. We will provide a brief overview of three types of systems: 1) manufacturing systems, 2) service systems, and 3) energy systems.

Manufacturing Systems

In manufacturing supply chains, the presence of a storage supply helps meet any limited imbalances in supply due to a number of factors (e.g. strikes, shortages of raw materials, power outages etc.) or sudden demand surges (people purchasing extra quantities of food and accessories before a hurricane etc.) Here the value of a storage facility is viewed less as a source of profits and more as something that is necessary to prevent the occurrence of some undesirable event. For example, automakers carry inventories of parts to be used in the automobile plant assembly lines because the nonavailability of these parts may lead to a shutdown of the entire auto plant (not desirable.) The absence of a certain brand of consumer product on the store shelves may lead to customers trying out another brand, and this may result in a permanent loss of market share. If one wanted to assign a value to storage in such systems, a cost would have to be assigned to each undesirable outcome. One has to multiply this cost by the probability of such an undesirable event to obtain a weighted cost, take the sum of these weighted costs over all such events, and one has a more tangible measure of the value of the inventory buffer. In practice, it is a little more complicated but this is the general idea.

In these systems, the cost of storage is usually low because the products are typically stored in warehouses which are located on some inexpensive land, and do not cost a lot to build. The products themselves could be expensive depending on what they are (e.g. automobiles vs.

pins), and manufacturers accordingly try to reduce their working capital requirements by managing the amount of inventory in storage. The sales/inventory ratio (also called the number of inventory turns ratio) is commonly used as a measure of the effectiveness of a given manufacturing operation in turning over its inventory in a given year.

Service Systems

In service systems, the situation is different. Examples of service systems include hotels, telephone services, public transportation systems like airlines and subways, and so on. The first question to be answered in any such system is: What is the product? In the case of airlines, the product is a seat on a particular airline flight on a particular day. In the case of the hotel industry, the product is a given hotel room on a given day, and for a telephone system, it is an unused line path that is used to route a call between two locations. The key feature in all these systems is the following: the product cannot be inventoried or stored because the product is perishable. If a seat on Flight AAA between Houston and New York is empty, that particular seat cannot be “stored” for use on some future flight. If a particular hotel room is not booked on a given day, the revenue for that day can never be recovered.

Another characteristic of service systems is that they have to be built to handle peak capacity. A subway system has to be able to handle the number of people during rush hour. It cannot be built to handle average capacity because in this case, the waiting times during peak periods will be very long, and this is unacceptable to the public. As you might guess, adding peak capacity is very costly. As units get added on to meet peak capacity requirements, the average utilization rate of the system decreases and the operating cost per unit of capacity would increase. In such systems, price is often used as a mechanism to reduce peak demand. In some telephone systems, one sees lower rates in off peak periods, and in the hotel industry, the off season has lower rates. The idea in all these cases is to shift the demand from peak periods to off peak periods, and reduce the peak capacity requirements. In manufacturing systems, one can reduce the peak capacity requirements by running the production plant 24 hour a day in three shifts.

An important goal in service systems is not to maximize the utilization of the system per se, but to maximize the total revenue from different groups of customers. This is a field in itself, and it goes under different names such as revenue management, yield management etc.

Energy Systems

In the case of natural gas, the product is the physical gas. In natural gas systems, the supply of storage is limited and in general it is expensive relative to the cost of gas. The supply is limited by geology, geography, financial and engineering constraints, environmental regulation, zoning requirements and governmental regulation. The amount of gas that can be stored is limited by the storage capacity available. Unlike manufacturing or service systems, storage in the energy market takes on a whole new meaning. When one evaluates a physical commodity like natural gas, there are three main attributes that one has to consider: (1) Price, (2) Physical Location and (3) Time. Assuming that the gas is pipeline quality, the molecules are indistinguishable and the product is truly a commodity (unlike different retail brands of gasoline, for example.) A consumer in New York is indifferent as to whether the gas is sourced from a well in Louisiana or an LNG terminal in Georgia. The magic of storage, the thing that gives storage its value, is the property that it allows you to transform a molecule of natural gas across time. In other words, if you own a certain quantity of gas that you are looking to sell at a particular time, and are not willing to sell the gas at the prices that are being offered, one could inject the gas into a storage facility and choose to withdraw it at some later point in time when the price is more favorable.The gas would have to be traded at the same physical location. Storage does not allow you to transform a molecule across location — this type of transformation is enabled by a transport option e.g. a pipeline connecting two points, A and B, allows you to move gas from A to B and thus you have transformed the asset across location. If you owned a quantity of natural gas at a particular location and time, there are four things that you could do with it at that time in the absence of a storage facility at that location: (a) Consume it, (b) Sell it to another party, (c) transport the gas to another location by pipeline or LNG tanker if such facilities are available, or (d) Burn it into the atmosphere (the flame that you see from many oil production facilities — there is nowhere to store and no ability to ship the gas someplace else.)

The neat thing about storage is that at the time the gas is injected into the facility, you do not have to make a decision about when you are going to withdraw the gas. You can sell this gas in the forward market using a futures contract if you choose to do so. Before the contract has expired, you can buy back the contract and sell this gas forward in some other months. This can happen multiple times before the gas is finally withdrawn from the facility. The ability to trade the forward curve and profit before any actual injections and withdrawals take place contributes to a significant portion of the value of storage, and it is this feature that complicates the valuation process considerably.

The importance of storage can be placed into perspective when one looks at markets where storage facilities are limited or absent. This document is not about power markets, but we will make brief mention here. Electricity markets exhibit a number of the characteristics of service systems. Due to its physical nature, electricity cannot be stored, and to handle peak capacity, a given market or region has a number of “peaking” units that can be turned on and off at short notice and manage peak demand. In these markets, on days with extreme weather conditions that lead to surges in demand, the inability to store electricity leads to temporary price spikes that are well outside the range of a few standard deviations. The primary reason for this phenomenon is the inability to store electricity.

Chapter 2: Valuing Storage

Attributes of Natural Gas Storage Valuation

When one tries to model a system, there are typically several variables that affect its behavior. The modeler tries to extract a minimal subset of variables that can satisfactorily explain most of the effects that are being studied.

In this case, the system that we are studying is a natural gas storage facility and we would like to determine its value over a fixed time period. We would like (a) to identify the attributes, or inputs, that are the key determinants of this value, and (b) to determine how the value changes as a function of the inputs.

Before we go about trying to value natural gas storage, we should try and come up with a list of the things that would be important to consider in any valuation process. We have listed these items below.

- The valuation method should use
  - Forward prices.
  - Implied volatilities (instead of historical volatilities.)
- The method should be able to capture the intricacies of most NGSF's including:
  - Injection/Withdrawal Ratchets: In this case, injections and withdrawals depend on inventory levels.
  - Time-based injection/withdrawal constraints (e.g. in the month of August, injections are restricted to lie between 10% and 30% of capacity)
  - Time based inventory constraints (e.g. on August 31, the facility cannot be more than 80% full.)
  - Injection/Withdrawal Fuel Factors and Commodity Rates.
  - Ad Valorem taxes.

There are a few other requirements that we would like to impose on a valuation methodology: the data requirements should be modest and the calculation time should be reasonable.

Different Methods for Valuing Storage

At this point, a few general comments are in order:

- In most storage contracts, the lessee begins and ends the lease with the same level of inventory. Consequently, for each injection, there is a corresponding withdrawal. For any given injection/withdrawal schedule where the starting and ending inventories are equal, one can pair the injections and withdrawals, and create a set of spreads. Using this approach helps one understand storage valuation better at an intuitive level, although modeling storage in this manner can complicate the valuation process.

There are five main methods for valuing storage. They are listed below:

- Alternate Cost Method.
- Intrinsic Value Method.
- Spread Option Method.
- Monte Carlo Simulation Method.
- Probability Tree Method.

We will now discuss each of the five methods for valuing storage.

Alternate Cost Method

This method is not precisely defined, and whether or not it merits mention in a separate category is debatable. It provides a range of values to indicate what storage could or should be worth.

The logic behind this method is the following: How can one replicate the services that one currently has without storage, and how much would it cost to provide such services.

For example instead of storage, one could use long haul tranportation, coupled with a LNG peaking contract to achieve the deliverability requirements. Then one could compare the cost of these services versus the cost of purchasing storage.

The drawbacks of this method are

- It is hard to completely replicate the services offered by storage because storage is quite a unique animal.
- The valuation problem now belongs to someone else, because the entity providing the services would almost certainly have to use storage as a backstop.

Nevertheless, this method is useful because it does provide a window as to what the storage asset class should be worth relative to the alternative services available in the marketplace. When storage developers and investors are looking to construct new storage facilities, the cost of these units should be compared to the alternatives available.

Intrinsic Value Method

The intrinsic value method is the most common method that is used to value storage and it is the benchmark upon which all other valuation is based. The idea behind the intrinsic method is simple: Suppose we assume that the forward curve is fixed or static and we wish to determine the value of storage at the current time instant. The optimal value obtained using the static curve is the intrinsic value of storage. Another way of looking at the intrinsic value of storage is as follows: If one had to choose the storage hedges today, and cannot alter them in the future, the maximum value that can be locked in with certainty is the intrinsic value.

We will illustrate the application of the intrinsic value method with a simple example. The example is small in order to keep it understandable, but the ideas can be extended to larger cases as well.

Before we proceed to the example, it would be useful to look at how a decision is modeled. We are going to model each decision as a spread. We use two time periods when defining a spread. e.g. Aug03-Jan04. We will use the following rules when we define a spread.

- When we buy 1 unit of the spread (or the position is +1), the injection of 1 unit will occur in the first time interval and the withdrawal will occur in the second time interval. When we sell 1 unit of the spread (or the position is -1), the withdrawal will occur before the injection (assuming of course that the gas exists in the facility to make the earlier withdrawal possible)
- If the day in the time period is not specified, we are referring to the first day of the time period.

The intrinsic method has the following advantages:

- It is easy to implement.

- Most parties can agree on what this value is.
- It is a useful benchmark to gauge the valuation using any other method.

The biggest drawback of the intrinsic value method is that it fails to assign any value to storage due to the price volatility. As a result, the values generated by this method are not realistic and therefore of limited significance. Despite its shortcomings, the intrinsic value is calculated as a part of any valuation process.

In the example that we have used to illustrate the intrinsic method (Example 1), there are a few things to be noted.

- We are assuming that each of the spreads can be traded at the prices that are listed in the table denoting the static spread matrix. In practice, one would have to consider the bid/ask differential when valuing these spreads, but this does not change any of the arguments.
- In most cases, the spreads are not the financial instruments that are traded in the marketplace. For each spread, one has to calculate the equivalent forward contracts and execute the trades for the required amounts of the forward contracts. The spread is an artifact that we have constructed for this problem. Instead one could have looked directly at the forward contracts to value this problem to exactly the same answer. By presenting the argument in this manner, it becomes easier to illustrate the spread option method which we shall describe in the next section.

Spread Option Value Method

The Spread Option method is a natural extension to the intrinsic value method described earlier. One of the drawbacks of the intrinsic method is that it fails to assign any value to the optionality of storage. The Spread Option method attempts to overcome this limitation by looking at the option value of each spread rather than the static value.

We will look at the same example as we did when we studied the intrinsic value method. The only difference here is that we are going to replace the intrinsic spread with the spread option value. To calculate the spread option value, we will use the Margrabe formula.

In this example, the optimal solution appears to be +100 units of the Sep03-Nov03 spread. If one passes this problem to a linear programming solver, the optimal solution is the following:

Buy 100 units of the Sep03-Oct03 spread.

Buy 100 units of the Oct03-Nov03 spread.

Start Date: 01-Aug-03.

End Date: 30-Nov-03

Injection Fuel Factor (IFF): 1.00%. Withdrawal Fuel Factor (WFF): 0%

Injection/Withdrawal Quantity: 100 units/month.

Minimum Storage Capacity: 0 units. Maximum Storage Capacity: 400 units.

Injection/Withdrawal Commodity Rate: 0.

	Forward Price
Aug-03	$ 5.000
Sep-03	$ 4.970
Oct-03	$ 5.100
Nov-03	$ 5.080

Table 1—Prices

The value of each spread is calculated according to the following formula:

$$\text{With. Price}*(1\text{-WFF}) - \text{Inj. Price}*(1\text{-IFF})^{-1} - \text{Inj Commodity Price} - \text{With Commodity Price}$$

	Aug-03	Sep-03	Oct-03	Nov-03
Aug-03		($ 0.081)	$ 0.049	$ 0.029
Sep-03			$ 0.080	$ 0.060
Oct-03				($ 0.072)
Nov-03				

Table 2—Static Spread Matrix

	Aug-03	Sep-03	Oct-03	Nov-03
Aug-03		0	0	0
Sep-03			+100	0
Oct-03				0
Nov-03				

Table 3— Optimal Spreads

Optimal Injection/Withdrawal Schedule:

Inject 100 units in Sep-03.

Withdraw 100 units in Oct-03

Example 1: Intrinsic Value of Storage Calculation

If one takes the net positions from these spreads, the optimal solution is the following:

Buy 100 units of the Sep03-Nov03 spread.

Now the question is: Why did the model not choose the Sep03-Nov03 spread as the optimal solution? Because the value obtained from exercising the Sep03-Nov03 spread directly is only $11.70 whereas the value obtained by exercising this basket spreads is $13.02. In the following discussion, we illustrate some of the issues that we have to deal with when using the spread option method.

Spread Option matrix should be arbitrage free

This example illustrates a very important point: It is difficult to recognize *a priori* whether the spread values are consistent or arbitrage-free. In other words, the spread option matrix should contain values such that the value of any basket of spreads should be equal to the value of the net spread position. One would think that the values of the Sep03-Nov03 spread should be equal to the values obtained by any other combination of spreads. In this example, we have done something that seems perfectly reasonable: we have merely applied the Margrabe formula to each pair of spreads. When one uses a correlation matrix that is obtained using historical data, there is no guarantee that the spread option matrix being generated is arbitrage free. One needs to be very careful when using a correlation matrix for the spreads because even if it works for the volatilities obtained from the implied volatility strip for one day, there is no guarantee that it will work for another set of volatility values. Frequently, one hears arguments to the effect: If such a matrix is generated correctly, this problem will not arise. Maybe so, but one would have to verify that this condition is satisfied for every volatility curve that is being used. Suppose one finds that the arbitrage-free condition is not satisfied, how does one tweak the numbers so that the matrix that satisfies this condition? If one finds such a matrix, are the new correlation values consistent with the data? Or does it indicate some more fundamental problems with the spread option approach?

What is the tenor of each option?

For modeling purposes, the tenor, or the time to maturity of the spread option is the earlier time instant of the two time instant pairs. For example, in the Aug03-Jan04 option, the expiration date of the option is 01-Aug-03.

	Aug-03	Sep-03	Oct-03	Nov-03
Aug-03		$0.0	$ 0.061	$ 0.058
Sep-03	$0.003		$ 0.082	$ 0.117
Oct-03	$0.001	$0.0		$ 0.048
Nov-03	$0.005	$0.026	$0.064	

Table 1—**Spread Option Matrix**

	Aug-03	Sep-03	Oct-03	Nov-03
Aug-03		0	0	0
Sep-03			100	0
Oct-03				100
Nov-03				

Table 2— **Optimal Spreads**

Example 2: Valuation using the Spread Option Method

What is the granularity of the time grid that is being considered?

This is an important question. In a one year deal, does one consider monthly time buckets, or daily time buckets? The argument in favor of daily time buckets is that the decisions are being made on a daily basis. The disadvantage of using daily buckets is that the problem now becomes very large because the data requirements have increased considerably. For a one year deal, instead of a 12x12 correlation and spread matrix, we have 365x365 matrices.

	Forward Price
Aug-03	$ 6.000
Sep-03	$ 5.750
Oct-03	$ 5.500
Nov-03	$ 5.250

Table 1—**Prices**

	Aug-03	Sep-03	Oct-03	Nov-03
Aug-03		$0.000	$ 0.000	$ 0.000
Sep-03	$0.192		$ 0.000	$ 0.001
Oct-03	$0.444	$0.194		$ 0.0007
Nov-03	$0.667	$0.449	$0.216	

Table 2—**Spread Option Matrix**

	Aug-03	Sep-03	Oct-03	Nov-03
Aug-03		0	0	100
Sep-03	0		0	0
Oct-03	0	100		0
Nov-03	0	0	0	

Table 3— **Optimal Spreads**

Example 3: Spread Option Method - Curve in Backwardation

Valuation of options requires a correlation matrix

The spread option method requires a correlation matrix which is difficult to generate. It is important to examine what the correlation number is: Suppose the current date is 01-June-03 and we are looking at the Aug03-Jan04 spread. The correct number is the average price correlation between the Aug03 forward and Jan04 forward for the time period starting 01-Jun-03 and ending 01-Aug-03. For the Sep03-Jan04 spread, the correlation is between the price forwards of Sep03 and Jan04 for the time period starting 01-Jun-03 and ending 01-Sep-03. It is important to take this into consideration when determining this matrix, or ask how this matrix was generated.

Correlation does not have a term structure property

A good question to ask is the following: How are the correlation numbers generated? Most users use a 40 day price history, evaluate the logarithm of the price ratios between consecutive days to create a series of returns, and then use the pairs of time series to calculate the correlation matrix. But as we had described earlier, this is not what the correlation number represents. The correlation number should be the average correlation between the price returns for the two forwards that are being considered. Here is an example that will make the issue more clear: Suppose the current date is 01-Jun-03 and we are looking at the correlation between the prices of Jan05 and May05. In June 2003, both these forwards are towards the back end of the curve and so the correlation will be close to 1.0. As time elapses and the Jan05 forward gets closer to the front month, the correlation will decrease. The correlation will tend to decrease over time as the nearer month spread reaches maturity.

What are the correct volatilities to use for the spreads?

The data that is widely available is the implied volatility strip. For a given option, the implied volatility is the average volatility over the life of the option. In practice, the volatility is low at first and tends to spike up as the option approaches maturity. Merely using the flat volatility may give rise to the wrong value of the spread option.

Probability of option exercise is not considered

The problems with the spread option model are more obvious when the curve is in backwardation. Here is an example that illustrates the point (see Example 3.)

The optimal solution is:

- Buy 100 units of the Aug-Nov spread., i.e. buy 100 units in August, sell 100 units in November.
- Sell 100 units of the Sep-Oct spread., i.e., sell 100 units in September, buy 100 units in October.

The problem with this approach is that the probability of exercise does not enter into the calculation. Even though the model says that we are long 100 units of the Aug-Nov spread (i.e. buy 100 units in August, sell 100 units in November), this does not mean that the spread can be exercised at this price. In fact, for this specific spread, the probability of exercise (as given by the $N(d_2)$ term in the Margrabe formula) is close to zero. The sale of 100 units of the Sep-Oct spread (i.e. sell 100 units in September, buy 100 units in October) requires 100 units to be in the facility so that the gas can be withdrawn in September. (Storage operators are usually unwilling to allow users to withdraw gas in excess of their inventory levels because that amounts to a loan which is accompanied by credit risk.)

If one attempts to incorporate these probabilities into the model, the situation gets hopeless. In order to do this, one has to use a decision tree (this is different from the probability price tree that we will talk about later.)

- Case I: Aug-Nov spread gets exercised. The Sep-Oct spread can be exercised.

  a) If the Sep-Oct spread is exercised, (another tree branch is required) then one can sell the Oct-Nov spread.

  b) If the Sep-Oct spread is not exercised, then one can buy the Oct-Nov spread.

- Case II: Aug-Nov spread does not get exercised.

  a) In this case, the Sep-Oct and Sep-Nov spread can be bought. Depending on which spread is exercised, new tree branches will have to be created.

The point to be taken away from the above discussion is that the example that we are looking at is a simple four period example, and the decision tree is already quite complicated. As the size of the problem gets larger, the problem grows so quickly that it is almost impossible to solve.

If an option is exercised, additional options come into play

As we have seen in the example above, if the Aug-Nov spread is exercised, then the user is can sell the Sep-Oct spread. Sometimes it is possible that the solution generated can be inconsistent. For example, the optimal solution could be to purchase +100 units of the Sep-Nov spread. But if this option were exercised, the Oct-Nov option comes into play.

Option value is highly sensitive to the correlation coefficient

The option value changes substantially as a function of the correlation coefficient. This is especially true for spreads that are at-the-money.

Model assumes one time exercise of options

Since the method uses the Margrabe formula, the implicit assumption is that the option is exercised at maturity. Everything else being equal, this would understate the value of the deal.

Mean reversion is not be considered

The Margrabe model does not assume mean reversion in the price processes for the two assets that are being considered. It is important to remember that not only are the prices mean reverting (this is taken care off by the exponential decay of implied volatility curve), but the spreads are mean reverting as well. That implies that for longer term options, the value of the spread option will be overstated. One has to reduce the correlation coefficient to get the value of the spread option down to a more realistic level.

The covariance matrix must be positive semi-definite.

The definition for positive semi-definitiveness is a bit complicated. Intuitively it means the following: if V is the variance-covariance matrix for a portfolio of asset returns, then the variance for any portfolio must be nonnegative. If this assumption were not satisified, one could imagine a stock portfolio where the variance of the return was negative. This property is something we take for granted. When a user is sometimes trying to price a multi-year deal, the correlation matrix needs to be expanded on the fly, and there is no guarantee that this condition will be met. The point to be made here is that the data requirements are substantial.

By replacing the static value of the spread by the option value, the spread option method is trying to capture the dynamic behavior of storage. But it fails to capture the true dynamic behavior associated with storage — the changing nature of the correlation, the non-uniformity of the implied volatilities over time, early option exercise, and the complex interdependence between the spread options. The data requirements of the model are large, and it is difficult to have a good feel for the correlation numbers. For modeling a ten year deal, the user has to make up the numbers while ensuring that the matrix remains positive semidefinite. In such a scenario, mean reversion becomes paramount otherwise the value of the option becomes very large. The user then has to tweak the correlation coefficient so that the spread is a reasonable value. At this point, the problem has moved out of the control of the user.

Despite these limitations, the spread option method is widely used for a number of reasons. It can be built relatively quickly, the inputs can be seen and the process is transparent. Users develop a certain degree of comfort with these models and continue to use them.

Monte Carlo Simulation Method

In a Monte Carlo simulation, one attempts to generate prices according to a distribution. We have the forward price and volatility curves. This translates into a distribution for a price vector that has a mean (forward price curve), a standard deviation (forward volatility curve) and a correlation matrix. Using a random number generator, one can generate a large number of sample vectors according to the distribution. For each of these price vectors, we have to value the storage deal. Taking the average over all these sample values provides us with the value of storage.

Large number of simulations required to reduce sample error.

Let the true value be denoted by V_{true}. The true value is the value of storage that would be obtained by taking an infinite number of independent samples. As the number of samples increases, we would expect the sample mean to converge to the true value. Let the value using the price vector s_1 be denoted by $V(s_1)$. Then the sample error $\varepsilon = |V(s_1) - V_{true}|$ is the difference between the true mean and the sample mean. (which in this case is $V(s_1)$.)

Suppose we generate the price vector n times to get a series of price vectors $s_1, s_2, \ldots s_n$. Then the values for each price vector are $V(s_1), V(s_2), \ldots, V(s_n)$ and the sample mean

$$\bar{V} = \frac{V_1 + \ldots + V_n}{n}$$

In this case, the size of the error term is

$$|\bar{V} - V_{true}| = \frac{\varepsilon}{\sqrt{n}}$$

So we see that as the number of simulation runs increases, the size of the sample error decreases by the square root of the number of runs. As a result, it takes a long time to reduce the size of the error. There are techniques that are used in Monte Carlo simulation (e.g. Control Variate method) that can be used to reduce the sample error, but it is hard to apply the control variate method to the case of natural gas storage.

Calculating the moments requires special care.

When one uses the Monte Carlo simulation to calculate the price deltas (or first moments), one needs to take special care. The way the delta is usually calculated is the following: The value of the deal is calculated, then the price is “tweaked” by a small amount and the new value is calculated. The difference yields the price delta. The problem that arises is that each of the calculations has a sample error. How much of the price difference is due to the change in prices and how much is due to the sample error? There are ways in which these problems can be reduced, but it is still something that you should be aware off.

Monte Carlo methods are inappropriate for valuing American options.

A fundamental problem arises with the use of Monte Carlo simuation for valuing American options. The problem arises from the way the prices are generated in a Monte Carlo simulation. In a Monte Carlo simulation, the prices are generated *ex post* — in other words, the model assumes that the deal has been completed, it then looks at the settlement prices for each of the time periods, and uses this price vector to determine the optimal injection and withdrawal schedule. In an American option where an option can be exercised early, for each given possible sample price path, the holder cannot "peek" along the path into the future to determine what the settlement price will be in order to determine whether or not to exercise.The user has to look at the information that is available only at the time of exercise. Hence, the value obtained using the Monte Carlo method will always be higher because of the advantage that this method has over what is the case in reality. The modeler now has to adjust the parameters so that the values are consistent with those being seen in practice.

Inaccurate representation of how the value of storage is monetized.

In the Monte Carlo method, the assumption is that the information is available to the model *ex post*, i.e. after the fact. So the algorithm goes through the following steps:

- Generate a random price path.
- Determine the optimal injection withdrawal schedule for each path.
- Perform this step a number of times i.e. 10,000 times.
- Take the average over all such paths to obtain the value.

As you can see, the algorithm looks at a price path, determines the optimal schedule and obtains the value. This is not the case in practice. The problem here is that all the decisions are not made simultaneously. On this sample path, one cannot implement the rolling intrinsic hedge strategy. If you were required to be delta neutral, this would require that any gas injected into the ground has a corresponding sale out in the future, but one could change the timing of the sale anytime before the gas was actually withdrawn from the ground.

In conclusion, the Monte Carlo simulation fails to capture the true dynamics of storage valuation.

Probability Tree Method

The probability tree method is the most rigorous among all the methods, and also the most difficult to implement. The probability tree method uses the following inputs:

1) Forward Prices

2) Forward Volatilities

3) Tree Parameters.

The first step in this process is to construct a probability tree that is consistent with the forward price curve and the forward volatility curve. If one were using this tree to price a forward contract, one should get the same price as the forward price that was used as an input. If one used this tree to value a simple option i.e. call, put or straddle, one should get the same price as one would get pricing such an option using the implied volatility in the standard Black-Scholes formula.

At this point, the tree consists of a set of nodes (which represent prices) and arcs which represent the transition probabilities of moving from one node to another. The next step is to define at each node, a set of inventory states which represent the level of inventory at that node, i.e. price/time period, to create a forest. A transition from one inventory state to another represents an injection or withdrawal depending on the states. One can then use value the problem on the tree using the stochastic dynamic programming approach.

Is Geometric Brownian Motion the right process to use for modeling commodities prices?

The brownian motion is one of the most beautiful mathematical processes ever studied. When we use the term geometric brownian motion, we are modeling the return, or the natural logarithm of the price process which is assumed to follow a brownian motion. If the logarithm of the price process is a brownian motion, then the price process itself is said to follow a geometric brownian motion. A simple way to view brownian motion is to think of a random walk in continuous time. It has the following properties: 1) stationarity, 2) martingale property, and c) Markov property. These properties have made the analysis of brownian motion a lot more tractable — there are a few other processes that possess all of these three properties (e.g. Poisson process.) Brownian motion has been studied extensively for about a hundred years, and some of the best mathematicians of the twentieth century have worked on this problem. If one were to make up a list of properties that a probability distribution for price returns should have, it could be:

- Additivity: The distribution should be stable. Suppose R_1 and R_2 are the returns for time periods 1 and 2 respectively, each of them has a distribution D, and R_1 and R_2 are independent. If D is stable, then $R_1 + R_2$ has the distribution D. In other words, the sum of the returns has the same distribution as the individual returns if the returns are independent. This property makes life a lot easier because the only thing that changes are the parameters of the distribution.
- Heavy Tails: The distribution should have a kurtosis (i.e. fourth moment) greater than that of the normal distribution. It has been observed that extreme events in commodities markets occur more frequently than that predicted by a normal distribution.

- Finite Variance: The variance of the distribution should be finite. If the variance is not finite, then it would become harder to use the information about implied volatilities that is generated by the options market. The option values implied by the distribution could still be finite, but one would need to create a measure of dispersion (other than the standard deviation) in order to track the volatility of prices.

If the returns follow a brownian motion, it can be shown that they are normally distributed. If one uses the brownian motion, one can leverage off the work that has been done in the last century. The most common complaint about the normal distribution is that it does not exhibit heavy tails. Before one tries to come up with an alternative, it would be helpful to keep the following facts in mind.

- The normal distribution is the **only** stable distribution where the mean, variance and all higher moments are finite. For all other stable distributions that exhibit heavy tails, the tails are so heavy that they have an infinite variance.
- There does **not** exist a distribution that has all three properties: additivity, heavy tails and finite variance. The normal distribution does not include the heavy tailed assumption. If one wants to include the heavy tailed property, one will have to sacrifice one of the other two properties. If one wants to model heavy tails using the volatility smile, then the additivity assumption will have to be abandoned. This may be okay when one is using a tree to fit the distribution using the market data. However before one builds the tree, there are a number of analytical formulas that have to be derived, and it is difficult to make much progress without the additivity assumption.

Brownian motion is not perfect for modeling commodity prices, but it is pretty good when you look at the alternatives.

Difficult to model correctly

Simply building the tree so that the prices are consistent with option and forward prices does not uniquely define the tree. For a given set of parameters, one can fit the tree to the forward price and volatility curves, and still get different values for a storage facility. In the financial markets, one can use additional information available from the captions and swaptions markets to obtain more clues about the tree parameters. It is more difficult to obtain this information in the energy markets due to the lack of liquidity. Consequently the tree parameters have to be determined from historical data.

How many factors are appropriate?

The theoretical answer: the more the number of factors, the better the fit. Increasing the number of factors can only increase the goodness of the fit. One needs to be careful about overfitting the data. A given model with a large number of factors may fit a particular data set very well, but when tested with another data set, the fit may be poor. Having a large number of factors also means that there are a greater number of parameters to be estimated and there could be errors in parameter estimation.

The practical answer: as few factors as possible. A good model should be able to capture most of the desired effects with a few factors. What the appropriate number should be is open to debate.

How does the model implement daily volatilities?

If one uses the implied volatility data from a futures exchange, e.g. NYMEX in the US, one has to examine the underlying index to see whether it is consistent with the storage market. In the US, the underlying natural gas futures contract is a monthly contract. Taking physical delivery under this contract implies that the underlying physical gas is delivered prorata over the course of the month. Hence the price is the average of the daily spot prices that trade for that month. If one uses these volatility numbers to value storage, the value will most likely be understated because the implied volatility on the monthly index will be lower than the implied volatility on the daily prices. One would have to determine the daily volatility from the monthly number using an appropriate model.

How does the model take advantage of the cash vs. futures arbitrage

The model should be indifferent as to whether the price came from the cash market or the futures market. The only thing that the model should care about is the following: For the time period in question, can one buy and sell gas in the physical market at the stated price? If not, the input should be a forward price that satisfies this property. One way that this can be done is to add a basis or physical premium to the futures curve.

Valuation tends to be less transparent

In many cases, the development of a tree based model takes several years. The creators of these products are hesitant to publish their work because it represents a substantial investment of intellectual capital. That being said, the fact that a model is proprietary does not imply that it is good. Unfortunately, the user will have to evaluate the products and decide the level of confidence that they have in a given model.

References

[1] Bernstein, Peter L., *Capital Ideas: The Improbable Origins of Modern Wall Street*, Free Press, 1993.

[2] Dacarogna, M.M., et al., *The Distribution of Extremal Foreign Exchange Rate Returns in Extremely Large Data Sets,* 1995.

Conclusion

We hope that you have found this document useful. We would appreciate any feedback that you have.

We can be contacted by regular mail, email or telephone.

Contact Information

Mailing Address:	Cygnet Risk Group, Ltd. 2476 Bolsover St, #495 Houston, TX 77005
Email address:	cathaide@cygnetriskgroup.com
Telephone:	(713) 668-8170

Please contact us if you would like additional copies of this document.

