A BRIEF HISTORY OF U.S. NATURAL GAS STORAGE RATES

By Mark Houldsworth, PhD and Rich Pastore, CFA

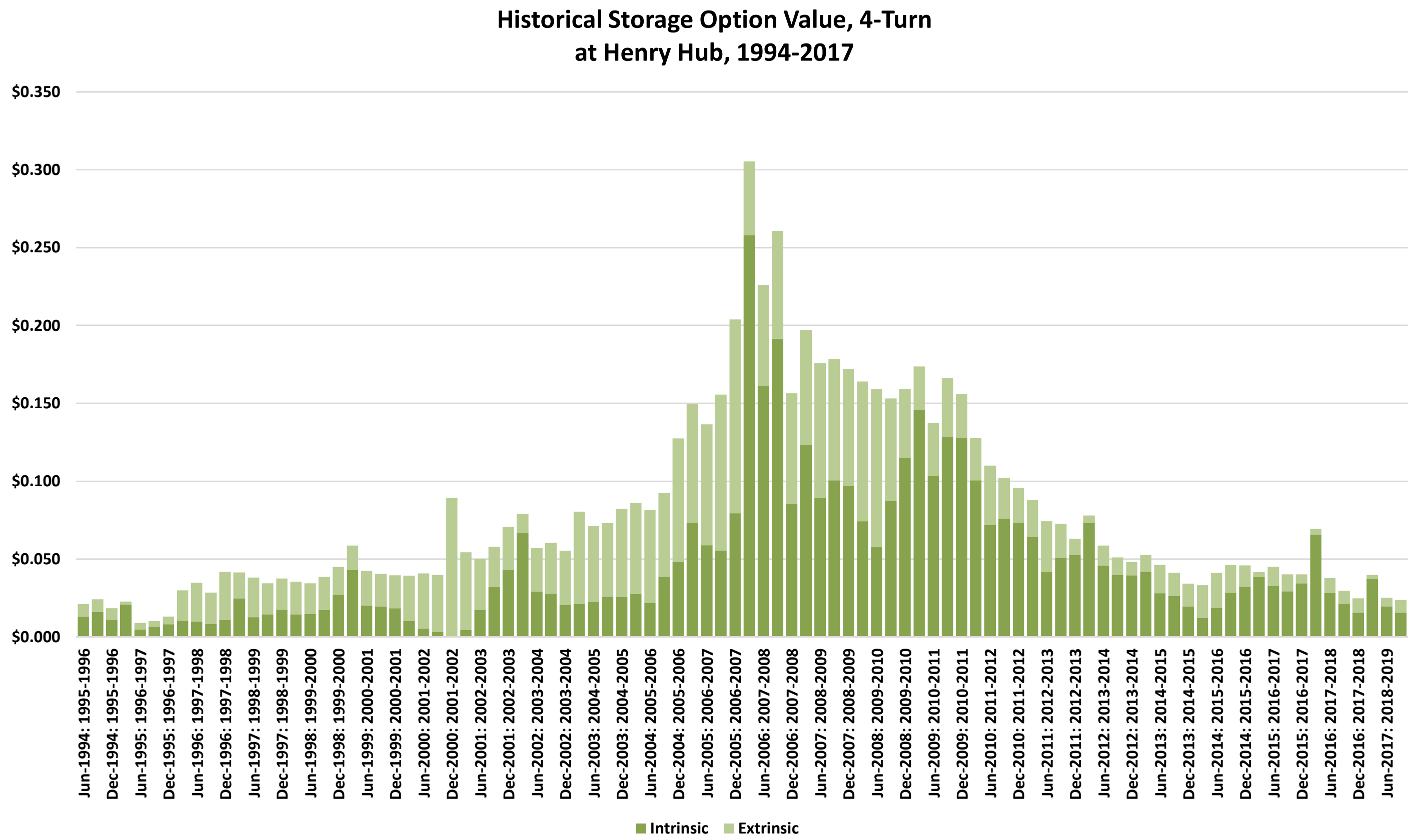

As the chart above clearly indicates, natural gas storage values have been in decline for most of the past 10 years; and storage values on the U.S. Gulf Coast continue to languish at the low end of the historical range. Major changes led to this distressed state, most particularly the vast growth in gas production following the development of shale fracking techniques, as well as large increases in both pipeline and storage infrastructure. Storage values, after falling precipitously, continue to stagger sideways, leaving market participants wondering how storage values will ever climb out of their current depression.

From a model-driven perspective, four key factors have the greatest influence on storage prices. First, the seasonal spread constitutes the self-evident driver of the *intrinsic* value of storage, while the *extrinsic* option value of storage relies on the three remaining factors: volatility, mean reversion, and correlation. A solid understanding of model-driven storage valuation requires a good familiarity with the pricing effects of these three factors.

The chart below illustrates the pricing trend of a hypothetical Gulf Coast storage facility if one were to calculate the option value of an April-March capacity lease as of the 1st of March preceding the respective storage year. This chart is a subset of the bar chart shown on the previous page.

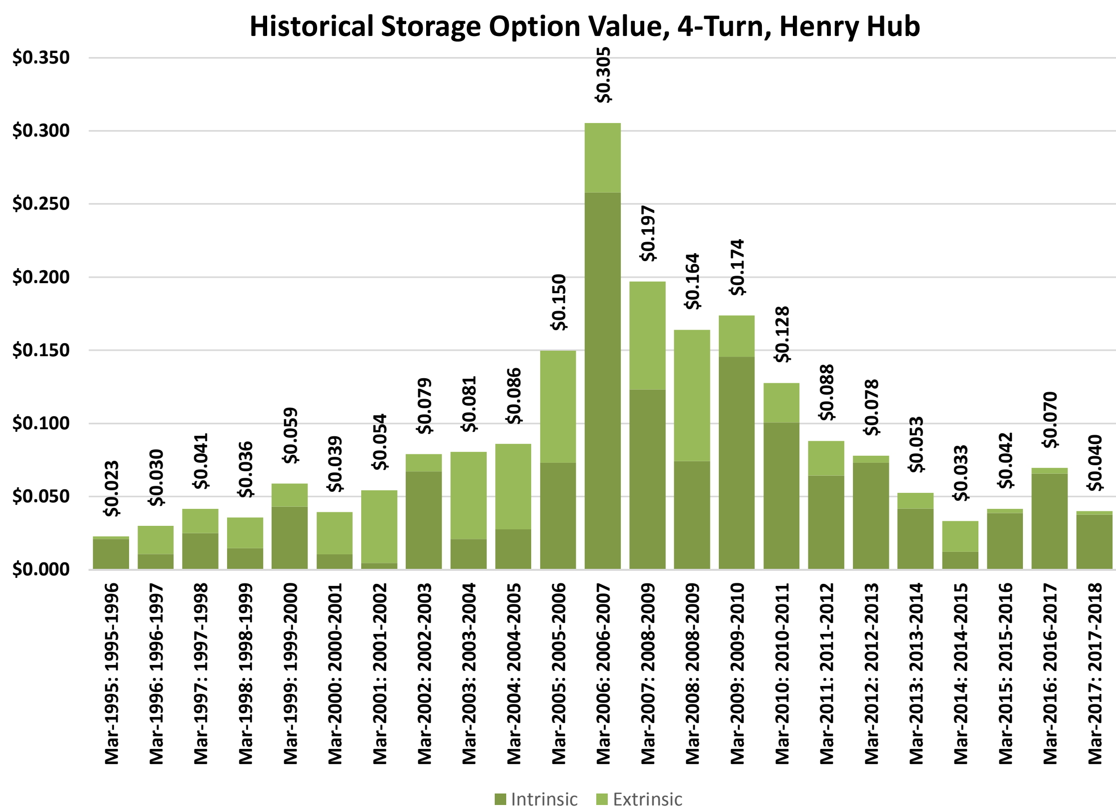

The bar farthest to the right on the chart indicates the value of $0.04 per Dth per month for storage as of March 1, 2017, using the market inputs as they existed at that date. At that point in time, the *intrinsic* value was approximately $0.0375 and the *extrinsic* value was $0.0025. Using this valuation as a starting point, what are the value impacts to *extrinsic* value if one were to turn the dials, so to speak, on volatility, mean reversion, and correlation?

The following bullet points demonstrate the respective storage pricing impacts if one were to turn the dial for each factor in isolation:

- **Volatility** – returning volatility to the historical median setting results in an *extrinsic* value increase from $0.0025 to $0.0089
- **Mean Reversion** – returning mean reversion to the historical median setting results in an *extrinsic* value increase from $0.0025 to $0.0046
- **Correlation** – decreasing all correlations by a nominal 2% results in an *extrinsic* value increase from $0.0025 to $0.0146

If one were to turn the three dials together, the *extrinsic* value increases from $0.0025 to $0.0203. This change is slightly less than the sum of the three effects in isolation because the factors of volatility, mean reversion, and correlation are inter-related.

APPENDIX

Extrinsic Value

The extrinsic value of storage is entirely driven by the spread volatilities over the array of spreads available to the asset.

Suppose the spread volatility for a particular spread is σ_{spread}, the volatility for the first underlying is σ_1, the volatility for the second underlying is σ_2, and the terminal correlation between the two underlyings is $\rho_{1,2}$. Then for linear variance expansion, as in Black, we have.

$$\sigma_{spread} = \sqrt{\sigma_1^2 + \sigma_2^2 - 2\rho_{1,2}\sigma_1\sigma_2}$$

A few illustrations.

Vol 1	Vol 2	Correlation	Spread Vol
0.3	0.3	1	-
0.2	0.5	0.9	0.33
0.2	0.5	0.8	0.36
0.2	0.3	0.9	0.15

If the volatilities are the same and the correlation is 1, there is no spread volatility; and there would be no extrinsic value for this spread option. As the volatilities widen and/or the correlation decreases, the spread volatility rises, thereby increasing the extrinsic value.

What follows is a summary discussion of the different components of the spread volatility.

Mean Reversion

In storage valuation, one is concerned with calendar spread options whose values are significantly determined by the two volatilities of the two relevant underlying forwards. One of those two volatilities, for example σ_2, represents the forward that is more distant (i.e. further in the future) than that of the other volatility. But it is important to remember that the calendar spread option expires at the beginning of the front contract. Mean reversion is the term that takes into account this aspect. In a world with no mean reversion, one could simply use the two Black implied volatilities and value the option. The back volatility would exhibit linear variance expansion and one would only need to stop its expansion at expiry.

On the other hand, if mean reversion in the forwards is present under calibration, as exhibited by calendar swaptions quotes, then an adjustment must be made to the back volatility to guarantee calibration to the market; therefore, the effective volatility for the more distant contract at the expiration of the option will be a little different, impacting the spread volatility and the value of the calendar spread option.

We can specify a simple one factor term structure for the back contract as follows.

$$\sigma_t = \sigma_T e^{-\alpha(T-t)}$$

Where,
$\sigma_t$is local volatility at t<T,
σ_T is the terminal or expiry volatility, and
α is the mean reversion parameter.

As an extension, and setting $\tau = T - t$, we can write the local variance as

$$\sigma_\tau^2 = \sigma_T^2 e^{-2\alpha\tau}$$

Calibration to the market requires the following.

$$\sigma_{Black}^2 T = \int_0^T \sigma_\tau^2 d\tau = \int_0^T \sigma_T^2 e^{-2\alpha\tau} d\tau = \frac{\sigma_T^2}{2\alpha}(1 - e^{-2\alpha T})$$

The Black volatility is known and the α parameter is estimated using swaptions. What remains is the expiry volatility, which can be solved with algebra.

Finally, again with a little algebra one can solve for the appropriate back volatility to use at the front of the front contract.

$$\sigma_{Back|expiration} = \sqrt{\frac{\sigma_{TBack}^2}{2\alpha}(e^{-2\alpha(TBack-TFront)} - e^{-2\alpha TBack})/TFront}$$

The calibration steps herein are important and necessary to match market risk and valuation. Regardless of whether one uses a lattice, least squares Monte Carlo or some other method, calibration is required so that any method can match market spread options and market straddles.

Correlation

As this is not an overtly scholarly discussion, we won't provide theoretical and mathematical foundations of correlation. It is intuitive that some kind of correlation parameter is necessary to properly price a spread option. Moreover, of one is a storage lease owner or a storage asset owner he is short correlation. As we saw in the table above, lower correlations increase the spread volatilities and this increases the value of the spread options.

We will note that if calendar spread options markets were complete one could simply use a Newton Rhapson method and compute the implied correlation and this would be fine for pricing storage calendar spread options. That not being the case one will have to estimate this parameter.

Since we are using European style options we don't really care about either the local volatilities or the local correlations in advance of the expiration. We care about the terminal variance of the two volatilities, and we similarly care about the terminal correlation.

There are nearly as many estimation methods for getting correlations as there are practitioners. We basically impose terminal correlation as an average correlation using historical data for each contract.

So, suppose the model is considering a September – December spread option. Our data set will include the last series of forwards where a September contract expired and we'll companion that with live December data. We'll take this data set back a year and calculate rolling local correlations all the way up to the expiration of the September contract.

We then impose

$$\rho_{terminal} = \frac{\int_t^{expiry} \rho_{local} dt}{(expiry - t)}$$

In this way the correlation used is conditional on both the maturity and the calendar month of expiry.

Other Important Drivers

At this point we've really only dealt with monthly spread options that expire at the beginning of the month. A model limited to these would basically pretend that all trading stops at the expiration and no one trades cash against the prompt or other forward contracts. This is, of course, silly.

To capture this activity, one has to include the appropriate cash volatility for the front month as a parameter and the cash to futures correlation. With these additional parameters one can add cash trading variance to the spread volatility above. Then add a suitable amount of additional maturity to the option. These additions to the spread variance can be material especially for near dated spread options.

Summary

In this appendix we've provided some detail on what drives model-based storage value and how the parameters come together.

Anyone that is long a storage position is basically long the spread volatility. Another way of saying this that storage asset owners love forward uncertainty, but this is only the case because that uncertainty can be hedged to generate far more certain PNL. If you have a good model.

Some quants will look at the material here and say that it's too detailed and more than you really need. And some quants will look at it and say that it isn't detailed enough, or it's the wrong model type. Quants, and you know who I'm talking about, are funny people.

We hope this hasn't been a waste of time for any of you. Call us. We can chat.

Mark Houldsworth
832-443-4010

Rich Pastore
832-545-0243