Modeling Spreads in Natural Gas Markets

By: David Soronow, Financial Engineering Associates, Inc.

Options based on the spread between two commodity prices are some of the most common types of derivatives found in energy commodity markets. Almost all energy companies have direct exposure to spread risk in one form or another including location spreads, refining margins, quality spreads and time spreads. Risk managers face the challenge of modeling spreads for earnings, profits, cash flow or value at risk calculations. As any trader can attest, the behaviour of a particular spread depends on the type of spread in question and each can exhibit remarkable complexity in terms of its stochastic and deterministic evolution through time. To tackle this problem, the risk management practitioner needs a host of modeling assumptions that can be called upon for each individual situation.

This article is the second part in a two part series investigating the main assumptions, approaches, and techniques applied to modeling spreads for derivatives pricing and risk management in energy markets. The first article reviewed the key implications behind different modeling assumptions, with a focus on locational spreads in natural gas markets. In this article we focus on a type of spread that is notoriously difficult to hedge: the quantity spread.

Variable Quantity Derivatives

A difficult case for derivatives pricing and risk management in energy markets are derivatives with a payoff based not only on the price of commodity, but also on the amount taken, which itself can be considered a stochastic variable. In reality this "spread" isn't truly a spread at all; at least not in the traditional sense of the word. Nevertheless, the risk associated with a variable-quantity type of spread is similar, and in many ways a magnification, of other spread risk found in energy markets.

To illustrate this, let's consider the case of a natural gas retailer who contracts with end-user consumers to supply gas at a pre-specified fixed price, with an amount that is determined at the discretion of the consumer. The "amount" or "volume" is variable in the sense that the end user can take as much or a little gas as they choose. The retailer simply checks the consumer's gas meter at the end of the month and bills them accordingly.

$$\text{Payoff}(t) = (S(t)_{Gas} - K) * N(t)$$

Payoff(t) = payoff at time t
$S(t)_{Gas}$ is the spot price of natural gas at time t
K is the contract fixed price per unit of gas
N is the total number of units of gas purchase at time t.

The idea of a variable quantity payoff presents somewhat of a dilemma for traditional derivatives pricing techniques, which typically rely on simplistic assumptions of ruthless exercise and unlimited liquidity. Here, the retailer knows that the consumer cannot simply turn around and resell the gas on the open market so the quantity taken by a consumer is determined by non-ruthless factors such as weather, season of the year, availability of substitutes, etc.

So how does this relate to spread risk and modeling spreads?

The astute reader will note that this retailer is faced not with outright risk associated to price changes but with something even more opaque, spread risk associated with "quantity". Let's look at this situation in more detail from a hedging perspective.

Variable quantity types of derivatives can be some of the most difficult to deal with from a hedging standpoint. In some cases the retailer may be able to mitigate a portion of this quantity risk through the use of a weather derivative, but more often than not, the only available hedging instruments are those that are designed to offset fixed quantity price risk. The NYMEX Henry Hub futures contract is an example of a fixed quantity hedging instrument. The naïve hedging approach would be to hedge the expected volume on a one-to-one basis, thus leaving an enormous amount of residual price and quantity risk.

To tackle this problem, the first step is to propose a stochastic model for volume. But what assumptions should we make?

Modeling Volume as a Stochastic Process

One approach that has proven popular with some practitioners is to model volume as a normally distributed, mean reverting random variable with a mean and variance parameterized as periodic functions of time in order to capture the seasonality in consumer demand.
This approach has several advantages and we can start by taking a closer look at the assumption of normality. Unlike price, volume has no theoretical justification for log normality other than the fact that volume cannot become negative (at least not in this case since end users can't sell gas back to the retailer). There is no compelling evidence, empirical or otherwise, to suggest that log changes in volume have equal probability of occurring for values on either side of the mean value. Rather, volume tends to evolve with reference to the absolute number of units.

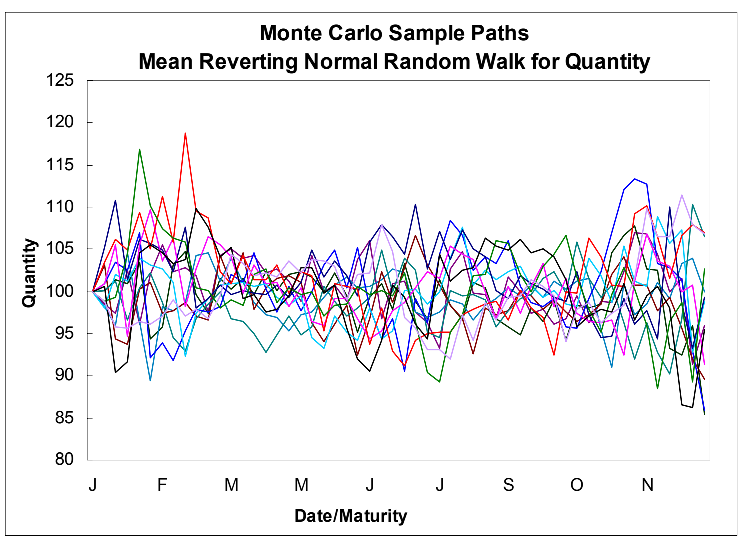

In addition to the assumption of normality, we also observe that volume doesn't follow a pure random walk where each change is independent of previous changes.

We intuitively know that variance in volume doesn't grow unbounded with time, but rather is driven by short term, and often fleeting influences such as weather. If an unpredictable cold spell in December causes a spike in natural gas consumption, we know that the spike is temporary and volume will quickly revert back to "normal" levels once the unusual weather subsides. This provides a fairly compelling argument for incorporating mean reversion into our stochastic model for volume.

To review, most Brownian motion models used by practitioners assume that the stochastic factor's variance grows linearly with time. In contrast, the essence of mean reversion is state dependent drift, which intuitively means that if volume is higher (lower) than the mean value it will eventually gravitate down (up) towards the mean or "equilibrium" level. This is exactly what we find in natural gas markets.

The last part of our model is a parameterization issue. To incorporate seasonality in both the expected value and variance, we can specify these parameters as periodic functions of time. Just as we have a forward price curve and volatility curve for describing the distribution of prices through time, we could also have a forward volume curve and volume volatility curve to describe the distributional characteristics of quantity. Naturally the expected value and variance would be highest in winter months when consumption is most prone to sudden drops or spikes.

Simulating the Mean Reverting Normal Quantity

$$V_T = e^{-a(T-t)}[V_t - \overline{V}] + \overline{V} + \sigma\sqrt{T-t}\varepsilon$$

Where:
V_T is the volume at time T
V_t is the volume at time t
$\overline{V}$ is the equilibrium volume value
a is the mean reversion rate
σ is the dollar volatility of the spread
ε is a drawing from a standard normal distribution

By now you should be aware that we are modeling quantity and price as two separate, yet correlated stochastic factors. For volume, we can use the model proposed above and for price we could use a model such as the mean reverting, lognormal model described in the previous article.

Correlation

In order to place structure on the relationship between price and quantity we need a metric summarizing the relationship between these values. The natural choice for such metric is the correlation coefficient, which measures the strength of the relationship between two random variables.
In most cases, we expect to find positive correlation between price and volume since prices tend to be high when demand is high, and low when demand is low.

Hedging and Risk management

In derivatives pricing we often make use of the so-called “risk neutral” valuation method, which in effect assumes that the market price of risk is zero, thus simplifying the derivatives valuation problem considerably. That said, risk managers require real world models of risk that utilize real world probabilities and not the risk neutral probabilities we conveniently use in derivatives pricing. The first place to start is with the recognition that perfect hedges are few and far between. This is nowhere more true than with variable quantity types of derivatives such as the situation faced by the natural gas retailer.

As stated earlier, the retailer only has fixed quantity futures contracts available as hedging instruments. A naïve hedging approach would suggest a one-to-one hedge with the expected value, thus exposing the company to residual risk of changes in price multiplied by changes in quantity. The possibilities fall along two dimensions and can be categorized into four main outcomes.

The first case is when the market price is lower than the fixed contract price and the actual volume is lower than the expected volume. In this case, the hedger will earn a loss on the long futures positions which is only partially offset by a smaller than expected gain on the short variable quantity forward.

The second case is when the market price is higher than the fixed contract price and the actual volume is lower than the expected volume. In this case, the hedger will experience a gain on the long futures position, and a lower than expected loss on the short variable quantity forward.

The third case is when the market price is lower than the fixed contract price and the actual volume is higher than the expected volume. In this case, the hedger will earn a loss on the long futures position which may be offset with a larger than expected gain on the short variable quantity forward.

The last case is when the market price is higher than the fixed contract price and the actual volume is higher than the expected volume. In this case, the hedger will earn a gain on the long futures position and a larger than expected loss on the short variable quantity forward. This case presents the most danger due to the potential for unlimited losses.

Clearly the naïve hedge approach has serious drawbacks especially in the case that presents the most danger. This provides motivation for finding an optimal hedge and we see that with our model in place for both volume and price, we can derive the optimal hedge ratio directly from our derivative pricing model. As one would expect, if volume and price are positively correlated, the optimal hedge calls for a hedge ratio in excess of one and vice versa.

Conclusion

Building a valuation and risk model for different types of spreads is a complicated task and requires analysis of the spread in question. This analysis coupled with our toolkit of modeling assumptions and techniques provides us with a framework for addressing valuation and risk related issues. In this article we looked at "quantity" spreads arising from variable quantity derivatives and used the case of a natural gas retailer as a reference.

David Soronow is a Senior Derivatives Specialists in the Global Support and Professional Services Group at FEA. *Soronow thanks Dr. Carlos Blanco, Dr. King Wang, and Dr. Angelo Barbieri for their advice and assistance.*