Preparing to Manage Energy Costs
Basics for Small to Mid-Size Consumers
by Cynthia A. Kase, CMT

Many energy consumers would like to hedge but don't know how to get started. This article veteran energy trader and hedging advisor, Cynthia Kase, addresses the major challenges new hedgers face, how volatility and risk may be understood and the decisions that must be made in order to commence hedging. Though this article will focus on natural gas, Ms. Kase's comments apply equally to hedging propane, #2 and #6 fuel oil and diesel.

Basic Decisions

How Much to Hedge

Before starting a hedging program, a hedger must decide how much exposure to include in the program. Typically 50% to 75% of total exposures are the right amount. This leaves the remainder to either float or to be purchased on a short term, month-by-month basis. So, for example a large end user consuming 24 BCF per year of natural gas would hedge maybe 50% or 12 BCF, and a small end user consuming 0.48 BCF would hedge perhaps a larger amount given the small absolute total, that is, maybe 75% or 0.36 BCF.

A key factor in fine-tuning one's ability to hedge the exposure set aside for a hedging program is whether the company has enough cash-flow or credit to weather any negative mark-to-markets in the hedge. To do this, a value at risk, or VaR calculation must be done, as discussed below. For example, if one's credit line with a particular dealer or market maker is $100,000 and risk calculation shows that there is a 5% chance that the mark-to-market could reach negative $200,000, then the hedge volume anticipated must by cut by half or a new credit line negotiated.

Correlation Issues

The next decision that has to be made is whether to hedge basis NYMEX natural gas or to hedge basis a particular pipeline. The first step is to run a correlation analysis between the pipeline index prices as published by a recognized source, such as Inside FERC, first of the month index. Similar correlations can be done between power (electricity) prices and natural gas. If the correlation is sufficient, which generally means that there is a statistically significant R2 of 0.8 or better, then the hedger may chose to do either. Also, a good correlation means that the hedger can participate in a hedging pool using NYMEX. Otherwise the hedge has to be place over-the-counter, denominated in the index in question.

Scale in or Not

"Scaling in" means to place hedges incrementally over time. Scaling in allows the hedgers to use strategies that avoid making "all at once" decisions, and to choose to place hedges at different times depending on market conditions.

An entity hedging 15 BCF per year has the equivalent of 100 futures contracts per month to hedge. Thus, it's easy for them to break their hedges into 1% increments if they choose, because that's equal to one futures contract. For the small hedger that's not the case. With 0.36 BCF per year, the hedger has only three contracts per month to hedge. So even if it's possible to hedge in one-quarter or one-half contract increments, the hedger's ability to scale in is limited, unless they hedge as part of a consortium.

Hedge Alone or Coop/Consortium

As noted above, large hedgers can hedge on their own and don't need to hedge as part of a consortium in order to aggregate enough volume to scale in, but that's a challenge for smaller hedgers. If groups of smaller hedgers form a hedging coop and aggregate their volumes together, then total volume may be sufficiently large and then suitable for scale in strategies.

Instruments to Use
Hedges can be placed in a variety of ways. Any hedge instrument can be delivered as a "paper" or financial hedge, or as embedded in a physical deal. The most common, simple method is to use fixed price hedging, which again can be separate as a financial hedge or embedded in the physical as a fixed price term purchase. Next in complexity are fixed range hedges, also called collars, which fix a range within prices can float. Option or price protection strategies protect prices from rising but allow them to float below a cap or ceiling price. "Exotics" are combinations of the above. For example, one can hedge with a collar and which puts a floor and a cap on prices. However, if prices fall well below the floor, one can get back into a floating price mode. This is called a three-way collar or just a "three-way".

Most small to mid-size industrial consumers stick to simple fixed price hedges, while those with larger exposures might choose to get more sophisticated in their approach. The key is that there is an inverse relationship between simplicity and upfront cost versus potential all-in costs on the hedge. For example, a fixed price hedge is simple to execute and carries no upfront premium or instrument charge. An option is more complicated to execute and carries an upfront premium. However, after the hedge is executed, if the market drops $3.00 the hedge cost to the fixed price purchaser is $3.00 while the cost to the options buyer is just the up front premium.

Estimating Risk to Reward Ratios
One of the key questions that potential hedgers consider is how much the cost of a fixed price hedge might be versus how much of a benefit could be achieved by placing a hedge. The name for the overall methodology to estimate these amounts is called "Value at Risk" and/or "Profit at Risk". The actual technical methods used to calculate these values vary, but here is a primer on what's behind it.

Simply put, the value at risk is the amount that can be lost based on certain probability, usually either 2.5% or 5% that a certain amount will be lost. For example, no more than a 5% chance that $200,000 will be lost. The mistake potential hedgers make is to misinterpret the data. For example there may be no more than a 5% chance that $200,000 will be lost, but at the same time there could be a 0.1% chance that $360,000 will be lost.

Side Bar Volatility and Risk Basics
The first step in understanding risk is to understand what volatility is. The most straightforward way to understand volatility is to see how it's calculated. To make this process clear, Fig, 1 below will be employed.

Fig. 1

Date	Close	ln	Stddev	Yearly
Aug. 06	P (Price)	ln(P0/P1)	σ, 9	*sqrt(252)
01	10.789			
02	10.856	0.006		
03	10.552	-0.028		
04	10.401	-0.014		
07	10.208	-0.019		
08	10.289	0.008		
09	10.827	0.051		
10	10.657	-0.016		
11	10.480	-0.017		
14	10.318	-0.016	0.024	38%

Fig. 1 lists in the first two columns the first ten trading day in the calendar month of August 2006 during which the NYMEX December natural gas contract traded, and the closing price for that contract on that day.

The third column, marked "ln" shows the logarithmic rate of change, expressed as the logarithm of the ratio price today (day 0 by convention) versus the price yesterday (day 1), for the past 10 days, or nine ratios. The use of a 10-day window is again by convention and the resulting volatility is called "the historical 10-day volatility".

Next, the fourth column, marked "Stddev" is the standard deviation over the 10 day, 9 ratio period. This is also the "raw" or non-annualized volatility. But, volatility is always discussed on an "annualized" basis, so the raw volatility is multiplied by the square root of the number of days in a business year, for purposes of this calculation, 252.

As a refresher, standard deviation is the square root of variance. Variance is the Sum from 0 to N-1 (using the convention above, where the most recent value is value 0) of the value $(X_N - X_M)^2$, where N is the number of observations, X_N is the value of each observation and X_M is the mean or average observation. So clearly volatility is one standard deviations of the logarithmic rate of change on an annualized basis. The fact that standard deviation is the square root of variance explains why volatility is proportional to the square root of time, and in turn the reason that the square root of the days of the year are used as a multiplier.

Understanding how volatility is calculated can clear up a lot of confusion both about risk as well as option premia. For example it is often confusing as to how volatility can be more than 100% in a negative direction, since we don't have negative prices. Let's say volatility is 120% and a particular contract has 28 days to expiration. One-ninth of a 252 day year is 28 days, the square root of nine is three, so the percent change that is represented by one standard deviation after adjusting for the time difference is 120% divided by the square root of nine, that is, three, or 40%.

This also explains why options premia increase with both volatility and time. Let's say today's price is \$7.00, there is a year to go before the contract expires and a buyer wants to purchase a \$10.50 call. In case one the volatility is 25% and in case two 50%. In case one, it would take a two standard deviation move, for prices to change by \$3.50 or 50% and hit the strike price. Statistically, a two standard deviation move has about a 2.5% chance of taking place. In case two it only takes a one standard deviation move for the price to change by 50%, and a one standard deviation move has about a 17% chance of taking place. So the option for case two's scenario is more expensive.

In the same example, lets say the volatility is 25%, the same as in case one, but now in case three the option only has one-quarter to expire. That means that the volatility on a de-annualized basis is 25/sqrt (4) or 12.5%. For prices to change by \$3.50 a four standard deviation move must take place, which has less than a 0.1% probability. Thus the shorter the time frame, the smaller the de-annualized volatility and the less the risk.

Considering Daily Changes

In actuality two things are true. First prices change daily, so in reality using one calculation, such as a 25% change over one year is inaccurate. Also, annualized volatility, on average increases over time. Thus a method called "Ito's Lemma" can help. The formula is named after a Japanese Professor K. Ito who developed it (which in turn is employed in the familiar "Black-Scholes" options pricing model) back in 1951. Using mathematics derived from an understanding of the bell curve and stochastic processes, it can be used to solve for price, if the time to expiration and volatility are known.

In a no bias, or non-trending market, which is the assumption made when doing risk calculations of this sort, risk can be estimated as the difference between the current spot price, P and an estimated future price P_N, where $P_N = e\hat{}(\ln(P) - (m - V^2/2)N \pm KV\sqrt{N})$, and N is the number of days elapsed, V is the non-annualized daily volatility, m is the growth rate of the commodity price and K is a non-negative deviate for a standardized normal distribution corresponding to a specific confidence level. The ± indicates standard deviations below (-) or above (+) the mean.

Let's say a buyer was entering into a swap at \$7.00 for one year, which will expire in one year and ten days, for example a calendar 2007 swap entered into 10 days prior to January becoming the first nearby contract. The first step in calculating the risk over the life of the hedge is break the year up into twelve 21-day increments and to find a representative volatility for each of the twelve months. Kase has done this by calculating the volatility for each month position (13th to expire, 12th to expire, 11th to expire, … first to expire) for the NYMEX natural gas contract over at least the last four years, and then to choose a representative figure to use for our calculations based on a percentile ranking, the 50th percentile, for example, being the median value which might be used to find the VaR under "normal" conditions. A value of the 75th percentile can be used to reflect slightly more volatile conditions, etc.

The value is then incremented on a spreadsheet on a daily basis using the Ito's Lemma formula, changing the volatility figure every 21 days, and solving for price, based on a certain number of standard deviations from the mean. The default we use for VaR in our standard model is 1.65 standard deviations, which reflects the 5th and 95th percentiles of probable price. The risk is the difference between the probable price and the original price, in this case, $7.00. The January contract would be incremented as the second nearby contract for 10 days and then the first for 21 days. The February contract would be incremented as the third nearby contract for 10 days and then second for 21, first for 21, etc. Then all the pluses and minuses for all twelve contracts would be added up to determine risk in either direction.

A calculation was done with the above example, using the 50th percentile of volatility for any declines in price and 75th percentile for increases, with the idea that in today's price environment any prolonged increases could be at a higher than normal volatility. Fig. 2 below shows the risk over time for a period of 260 days. The risk, that is, the VaR, to an unhedged short position, that is unhedged purchases, is rising prices. At a 5% confidence level this value is $3.25, which for one contract equivalent per month, or 120,000 mmBtu per year is $390,120, which takes place about three and a half months into the deal. From that point onward the risk diminishes as the contracts expire and the volume of outstanding contracts declines. For a hedged position, that is a long hedge that offsets the short position, the VaR arises from falling prices that introduce a cost to the hedge. On this flipside the greatest risk averages about $1.73, or a total of $207,350 for the total volume. This value peaks about one month earlier due to the skew of the volatility distributions. In any case after one year and ten days, all volumes expire and both values converge to zero. At the theoretical limit, the maximum that can be lost on a $7.00 long swap is $7.00, while conversely prices can rise to infinity.

Fig. 2

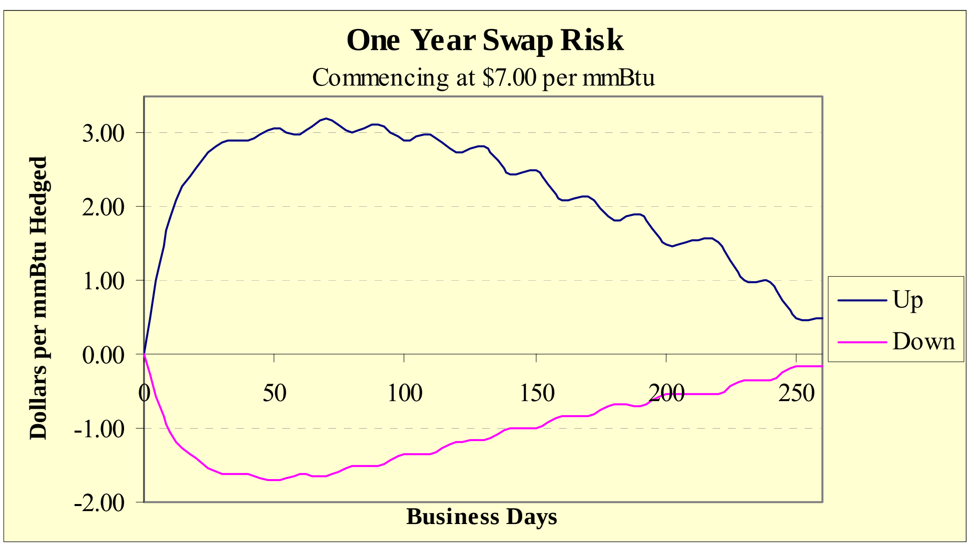

End Sidebar

Developing a Strategy

Once basic decisions have been made, a strategy must be developed for execution. The most common type of strategy is what I call "consensus" hedging. To quote John Kenneth Galbraith – "When it comes to the markets the majority is always wrong". So while it is a great idea to develop a consensus on the goals of a hedging program, its never sensible to use group-think as a strategy, especially in hierarchical organizations in which the boss, who may not be very good at market analysis, rules. A similar approach is to delegate hedging decisions to an individual or group of traders whether inside the organization or outsourced to a marketing company. Here the decision making process is "elevated" to the discretionary (translate speculative) views of people whose incentives, goals and risk appetite may not be in line with the company's risk management program. Even if the traders involved are talented, most traders are both short-term as well as speculatively oriented, and of course if the talent leaves then the program which relied upon them may fail.

Instead of using these methods various types of scale-in strategies, some using statistical methods and probability theory can be used. For small end-users such strategies are of great value because they allow for a uniform method to be used by a coop or consortium where aggregation of volumes can then allow for the scaling in of small volumes over time. The first two methods below also allow for fully automated signals to be generated, that is the signal to hedge is generated by the strategy not by an individual or group of individuals. The value of such programs is that they can be delegated to a gas management firm for execution as part of an aggregated coop where all members follow the same strategy.

Scaling-in Using Volume Cost Averaging

A simple strategy that requires no talent or thought is to do a straight volume averaging over time, hedging a twelve-month strip, for example in advance. Starting at the point when a contract becomes the 13^{th} nearby, one would accumulate $1/250^{th}$ per day of volume and when enough volume has been accumulated to hedge a half- or full-contract, a hedge would be placed. Over the long run such a method would be expected to reflect long-term trends. If over the long run the market has been rising, the scale in method, which is always buying ahead of time, would make money. The reverse would be true of long-term bear markets.

Rule-based Volume Cost Averaging

Instead of blind volume averaging rule-based systems that improve performance of such strategies may be employed. While there are a number of models available from various consultants in North America, Kase's approach can serve as an example. Instead of buying the same amount every day, Kase blocks out certain periods of time in which no hedging takes place.

For example, if the market has been rising for a very long time, the odds increase that a downturn will take place, so hedging ceases. Similarly if the market has been declining for a while, the odds are it will continue to decline, so again, hedging ceases. Also a higher percentage per day is accumulated when prices are low than if prices are high. Our system is called "ezHedge", and the results for this approach are approximately twice as beneficial (one does not use the word "profitable" when hedging) than a straight dollar cost averaging method, not only are the gains larger but the costs (read "losses") one has to undergo in order to gain in the long-run.

Fig. 3 shows how typical volume-averaging systems mighty look. Dots are plotting on the closing price for the twelve-month strip each day. The dots are color-coded and appear in one of five zones created by the four gray lines shown on the screen. The color and the placement in a particular zone tells the hedger how much to execute.

Fig. 3

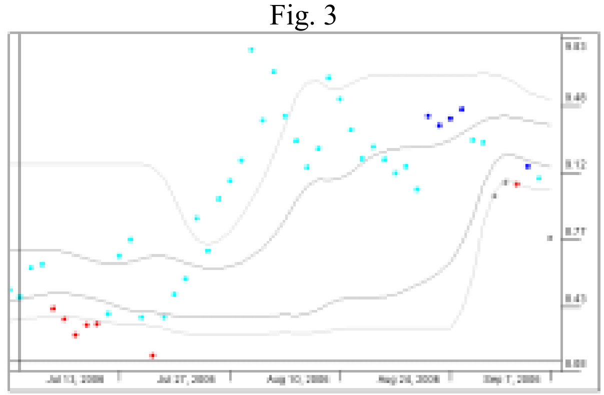

Combining Rule-based Statistical Hedging with Discretion
A third alternative is to combine rule-based hedging with some discretion. Both of the programs can be used in this way, but the primary design is to buy forward on a fixed price basis only and to hold until expiration. If an end user's program allows for discretion relative to instrument use and the removal and replacement of hedges, it then may be advisable to use a rule-based system that allows for some discretion. One such program is Kase's HedgeModel which sets statistically significant points at which to place both fixed price hedges and for purchasing calls. The entry system is fairly mechanical but allows for a range of rule-sets depending on risk appetite instead of a one-size-fits-all approach as well as a range of instruments to choose. For example whenever prices are low and fixed price purchases are allowed, the hedger may use collars, calls, or a range of more complex instruments such as three-way collars. When calls are triggered the buyer can choose how far out-of-the-money to purchase the call as well as fine-tune by using cheaper, more sophisticated strategies such as buying early expiration calls, called "swaptions".

In addition, instead of holding to expiration, such programs allow for rare instances during which the gain in the hedge is sufficient to trigger restructuring or removal altogether. For example, in Fig. 4, the dots, which in this model are not color-coded, as shown in the zone of the lower lines. At that point, according to a non-discretionary rule set chosen for a specific risk appetite, the hedger may have bought forward.

Fig. 4

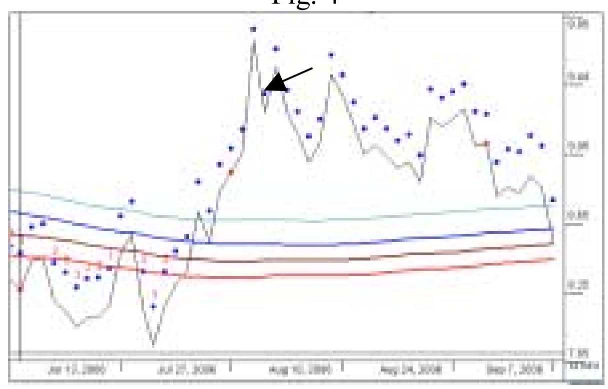

However, later as the market peaked and began to decline as shown by the arrow, according to specific rules, but this time discretionary, the buyer may have chosen to convert the fixed price hedge, now showing a benefit, to a series of out-of-the-money calls and booking the remainder of the funds. For example, if the fixed price hedge was executed at $8.25 and market is now $10.25, the hedger may spend $0.50 on a new series of calls and book the $1.50 remaining. Though more complex to run than a totally non-discretionary system, such models have the potential for avoiding most costs due to the ability to fine-tune strategies and to restructure.

Summary
Hedging can serve to greatly dampen the impact of volatile, rising prices. Many decisions must be made from a business, technical and strategic standpoint. Once done even small hedgers can participate in effective programs though the use of aggregating coops.

Cynthia Kase, CMT holds both a BS and ME in Chemical Engineering. In 1983 after 10 years in that she became an oil trader for a major, then a derivates market maker at a major bank. After a year with the Saudis, she became president of Kase and Company, Inc. (CTA) in 1992, recognized by many as the leading boutique energy hedging advisory firm. She is currently focused on forming hedging cooperatives among groups of small- to mid-size end-users in North America. www.kaseco.com.