Quantifying Cross Commodity Risk in Portfolios of Futures Contracts

By Ted Kury, Senior Structuring and Pricing Analyst
The Energy Authority, Inc.

The risk management process varies greatly from organization to organization, but regardless of whether it is simple or complex, the crucial first step is always the same. Before risk can be managed, risk must be quantified. The volatility of forward price curves can be a key contributor to the risk of any entity that uses a value at risk metric or stop loss limits in the risk management process, or is subject to margin requirements due to exchange traded financial instruments.

It is important to differentiate between the modeling of spot prices and the modeling of forward prices. Each has its uses, but the models used cannot be substituted for one another, as they have different temporal characteristics. The behavior of spot commodity prices has been studied at length, and a number of different models, such as the single and multi- factor mean reverting models of Pindyck (1999) and Schwartz (1997), the mean reverting with jump diffusion models of Clewlow and Strickland (2000) and Clewlow, et al. (2001), and regime switching models derived from Hamilton (1994) have emerged. These models can help to determine the distribution of prices when the spot day or spot hour occurs, but provide little insight before then. For example, if you want an idea of the distribution of natural gas prices on March 14, 2006, these spot price models may all help you. If, however, you are interested in the distribution of prices for the March 2006 natural gas futures contract 5 days from now, they are useless.

Spot price uncertainty grows proportionally with the square root of time in all of these models, while futures prices behave a bit differently. The front months of the forward curve tend to be much more volatile than the back months. Generally, forward price volatility increases as time to expiry decreases. This behavior is a consequence of the mean reverting behavior of the spot price, (Clewlow and Strickland, 2004). A model of forward prices should incorporate this market behavior.

Once the term structure of the forward price volatility has been properly modeled, it is also important to capture the relationship between futures contracts of the same commodity and across commodities. In the market, the correlation between the prompt month futures contract and the remainder of the forward curve tends to change, as the contract gets closer to expiry. This behavior should be reflected in the model.

Volatility models with a constant term structure are often used to model the expected volatility of futures contracts. These volatility estimates are assumed to be equal to the standard deviation of the log price returns over some trailing time period. Due to the tendency for the volatility of futures prices to increase as time to expiry decreases, this approach may prove inadequate for holding periods longer than one day, or as contracts get closer to expiry.

The derivation of a model for forward prices requires some assumptions about the behavior of the spot price of the commodity. For our purposes, we are going to use the single factor mean reverting framework of Pindyck. The spot price of a commodity is assumed to follow the following form:

Where:

$$\ln S_t = \ln S_{t-1} + \alpha(\mu - \ln S_{t-1}) + \varepsilon_t$$

$$\varepsilon_t \sim N(0, \sigma^2)$$

α is the mean reversion rate, and
μ is the log of the long run equilibrium price.

From this model, Clewlow and Strickland (1999) and Lucia and Schwartz (2001) have shown that the volatility of the forward price can be expressed as:

$$\sigma_t = \frac{\sigma}{2\alpha t} * \left(1 - e^{-2\alpha t}\right)$$

That is, that forward volatility at any time t is inversely proportional to both the mean reversion rate of the commodity and the time to expiry of the futures contract, and directly proportional to the theoretical volatility at expiration. We can use comparative statics to illustrate this relationship. A change in sigma, holding alpha constant, results in a level shift of the volatility curve, while a change in alpha holding sigma constant, results in a change in the slope.

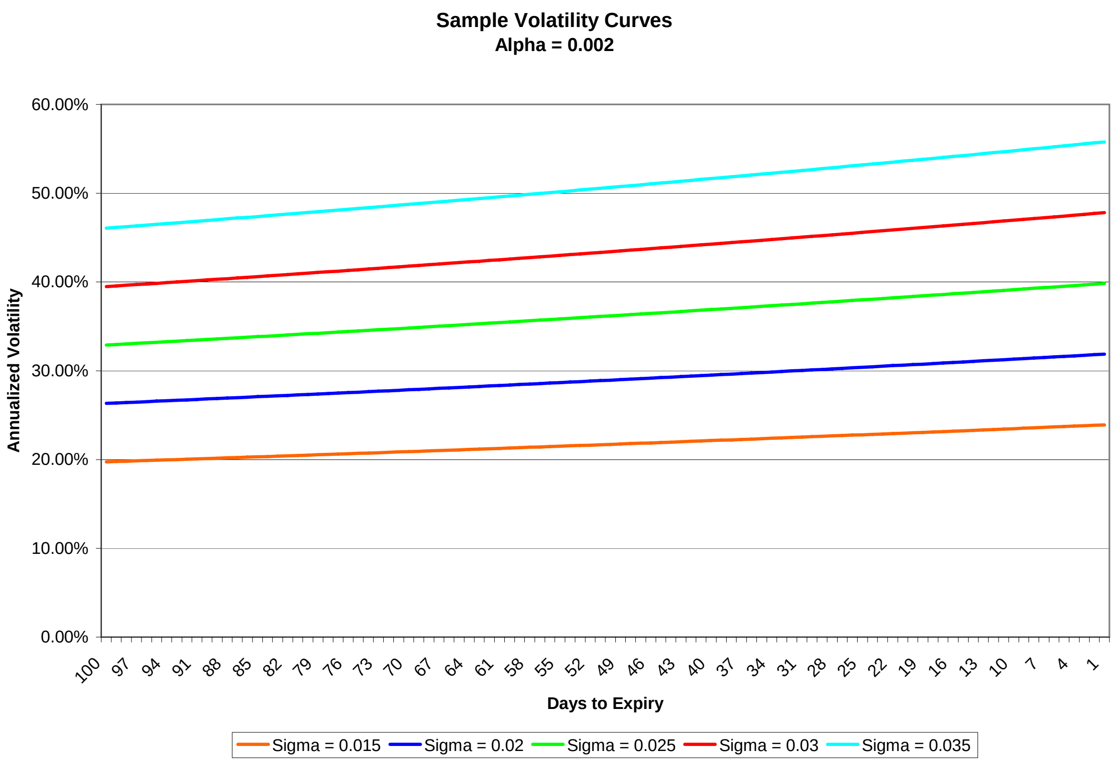

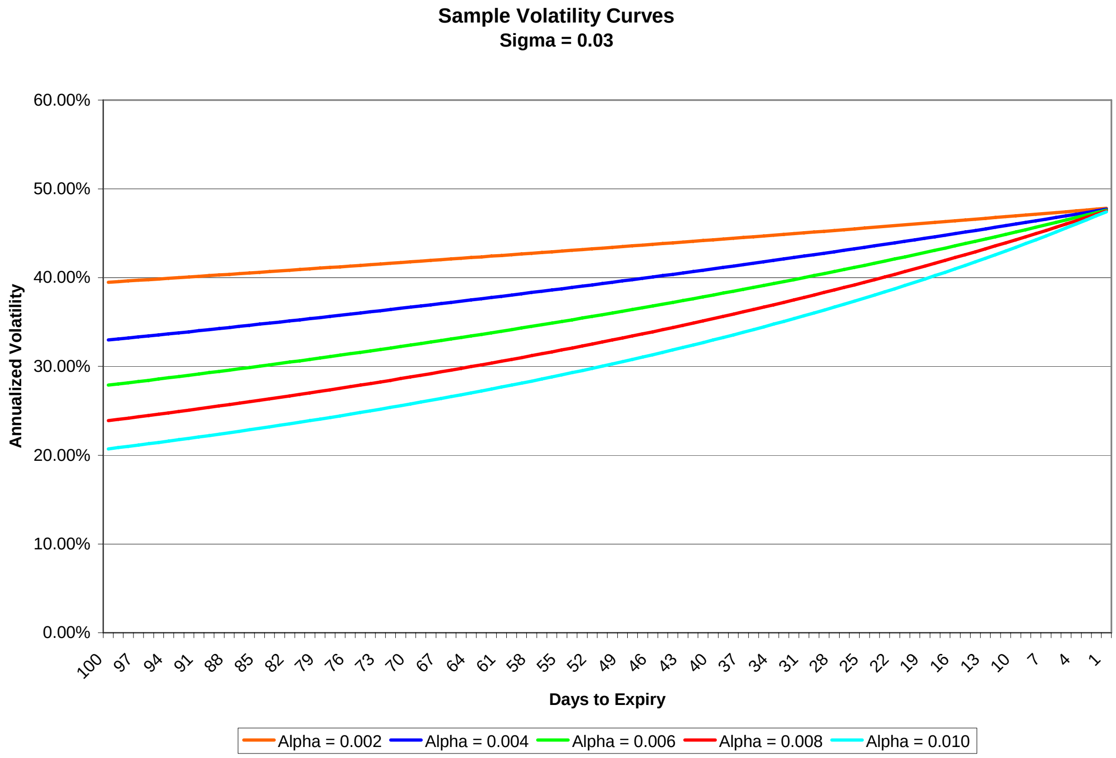

It is important to stress that these volatility curves exist distinctly for each forward contract, as each contract may exhibit its own volatility structure. For example, the price of electricity tends to revert to the mean more quickly than the price of natural gas, and thus will exhibit a steeper volatility curve. Natural gas prices may revert more quickly than oil prices in some months, but not in others. Further, summer electricity prices tend to be more volatile - and thus have a higher sigma - than other months. Winter natural gas prices exhibit similar behavior. Thus, each contract should be allowed its own volatility curve.

This chart shows the historical volatility of NYMEX natural gas contracts with the three most volatile contracts in each year highlighted in blue. Note that there is little pattern as to which contracts will be the most volatile. Frequently it is the winter contracts, but some spring, summer, and fall contracts make the top three in certain years.

	Annualized Volatilities				
	2000	2001	2002	2003	2004
Jan	20.03%	28.27%	32.73%	31.90%	31.68%
Feb	19.21%	32.69%	33.41%	33.33%	32.98%
Mar	18.76%	33.67%	33.50%	38.68%	31.87%
Apr	17.41%	28.76%	33.90%	38.29%	27.28%
May	16.18%	25.55%	34.54%	35.44%	26.44%
Jun	16.82%	26.41%	34.61%	34.42%	26.45%
Jul	20.35%	27.80%	34.83%	33.60%	25.59%
Aug	20.93%	29.06%	34.79%	32.25%	24.85%
Sep	21.30%	31.59%	35.90%	32.43%	24.93%
Oct	20.97%	32.59%	37.13%	32.27%	25.61%
Nov	21.54%	33.09%	34.39%	30.97%	25.52%
Dec	22.77%	32.50%	32.42%	29.92%	25.44%

The absolute value of the daily log price returns of each futures contract can be fit to this volatility function using non-linear optimization techniques. The volatility function is continuous and twice differentiable, and because there are only two free parameters, the matrix of second derivatives will not be unwieldy. Therefore, we can use one of the hill-climbing methodologies, such as Newton-Raphson, Gauss-Newton, Goldfeld et al. (1966), Broyden-Fletcher-Goldfarb-Shanno (BFGS), or Berndt-Hall-Hall-Hausman (BHHH 1974), to perform the non-linear optimization. The benefit of these hill-climbing algorithms is that they are straightforward to implement. However, they are also local search algorithms, so it is necessary to perform the optimization with a range of initial points to ensure that your solution converges to a global, rather than local, maximum. Once we have determined the correct area of the volatility surface in which to search, adding daily data points should not change the surface significantly.

Non-linear optimization can be a tricky process, and it is best to be as parsimonious as possible. We have tested a number of variables that might improve the fit of the volatility function such as seasonal parameters for all commodities, dummy variables for Mondays (as the first trading day after two non-trading days) and Thursdays (the day the EIA storage report is published) for natural gas, and none improve the fit sufficiently for the added computational complexity.

The volatility curves and absolute value of the daily log price returns of the April 2005 and November 2006 NYMEX natural gas contracts are shown below. Note that the volatility curve for the product nearer to expiry is much steeper.

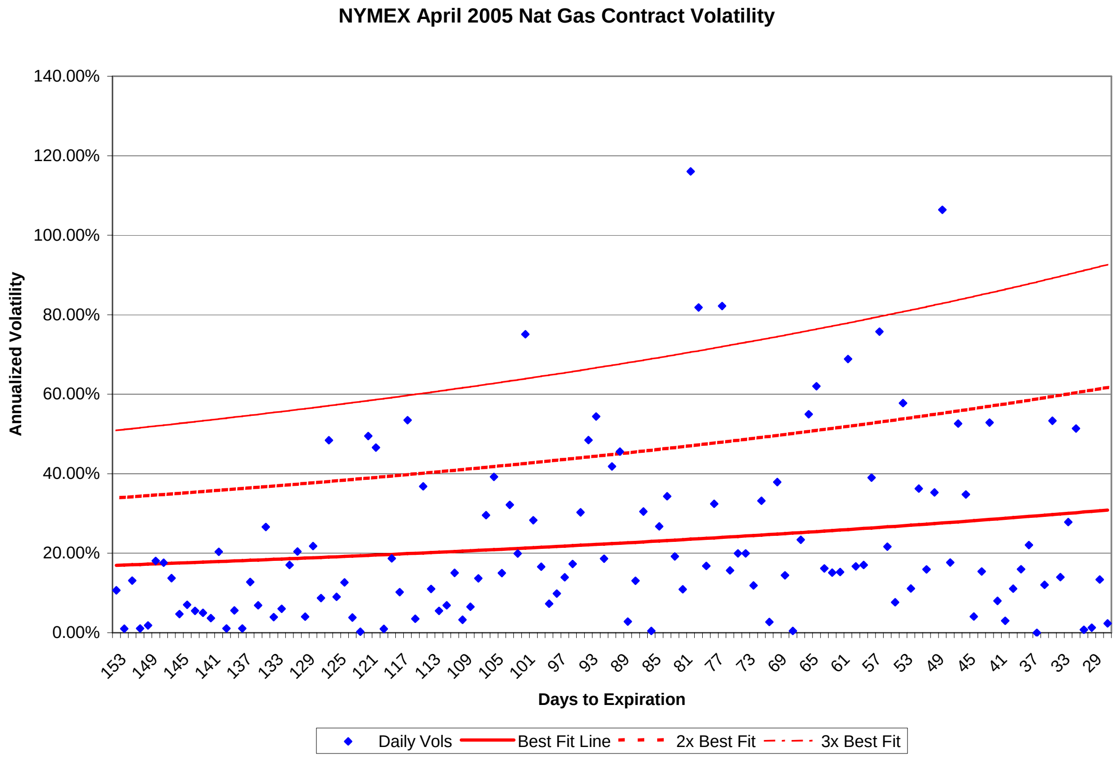

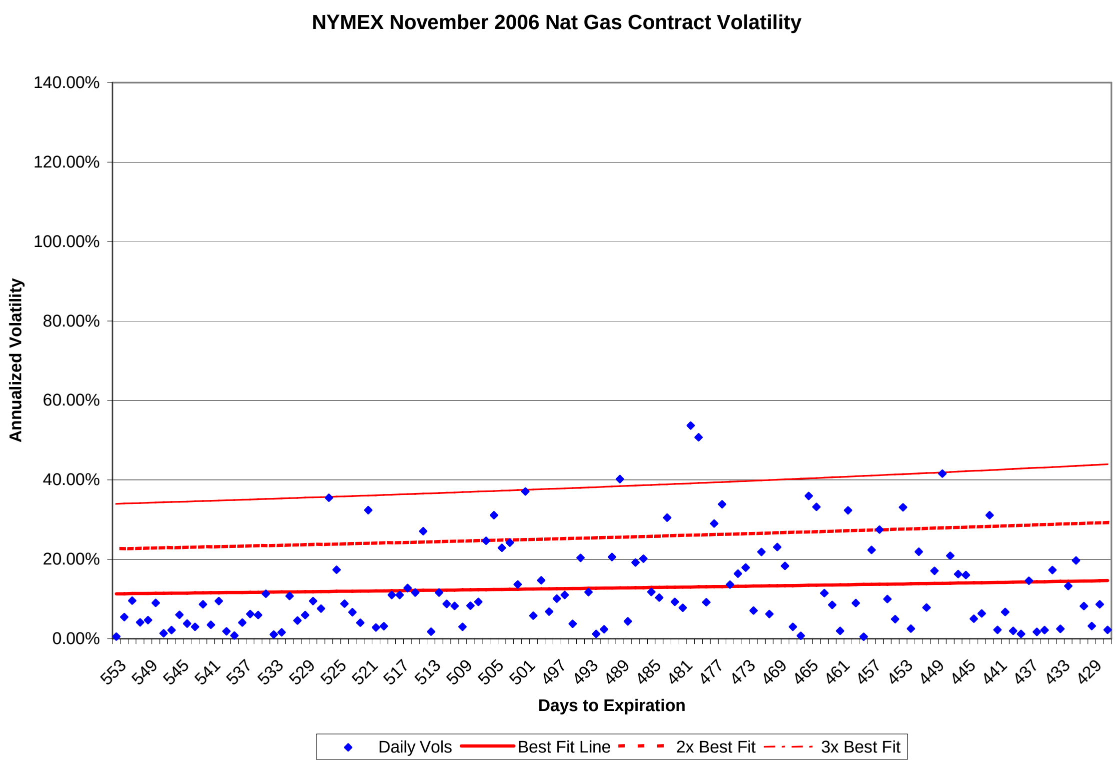

In each case, approximately 35% of the volatilities fall above the best fit line, 13% fall above twice the best fit line, and 3% fall above three times the best fit line. This is slightly more than would be expected from a normal distribution, suggesting that there may be excess kurtosis in the price returns.

A magnification of the April 2005 contract may be used to illustrate the difference between this methodology and forward contract volatility based on standard deviation of the price returns. The dashed red line is the expected forward volatility used in simulations of our model.

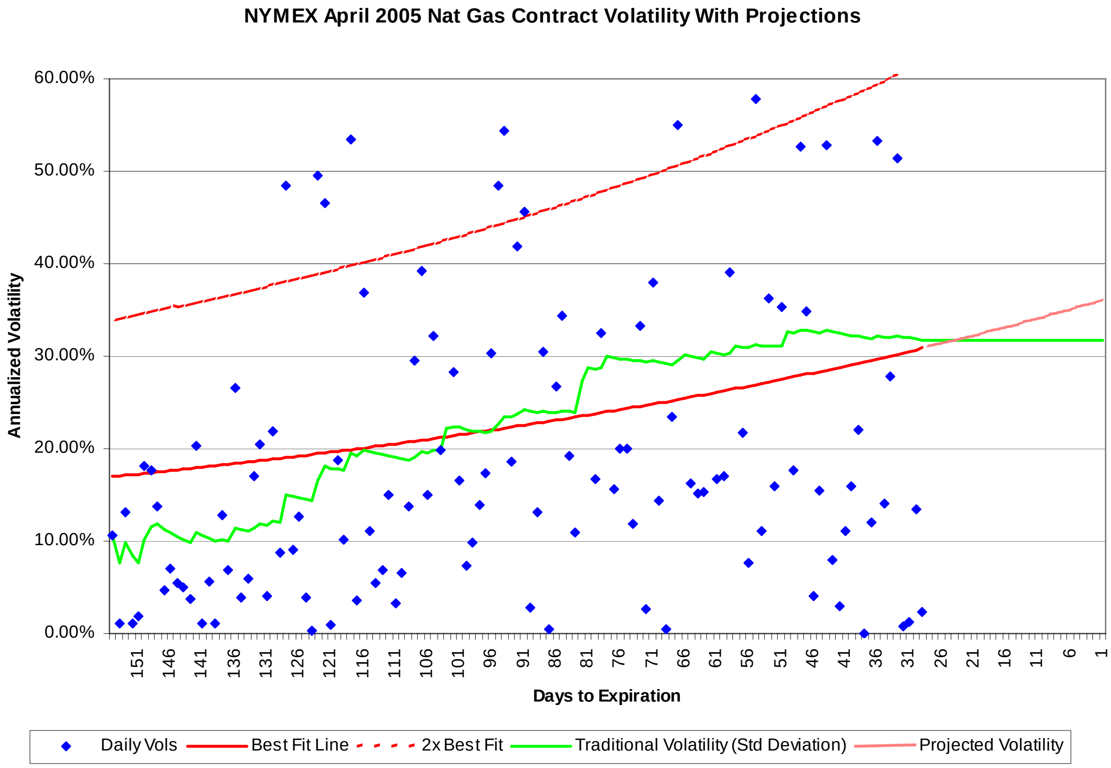

It may not appear that the difference between the two methodologies is significant, but consider that just a 3% difference in assumed volatility will swing the value at risk of a $20 million portfolio by approximately $1 million at the 95th percentile.

Once we have parameterized the volatility term structure for each contract, we turn our attention to modeling the interrelationships of the forward contracts. We are not only trying to capture the relationships between contracts of the same commodity, but between other commodities as well. We are also trying to accomplish this as efficiently as possible. One general methodology involves the decomposition of the correlation matrix of normalized shocks. This is an important concept. We will construct the correlation matrix of the daily price returns divided by the expected volatility for that particular contract on that particular day. The normalized shocks will be analogous to the standard normal random variables that will generate in our simulations of the forward prices. In this way, we will not overestimate the correlation between the prompt month contract and the rest of the forward price curve.

Data in the energy industry can be difficult to acquire. If the data set includes data series with different time frames, the correlation matrix may not be positive semi-definite. If the correlation matrix is not positive semi-definite, then Cholesky decomposition will not work. Cholesky decomposition will also capture all of the contract interrelationships, some of which may not be meaningful, that is, in the noise. For these reasons, principal components analysis with eigen decomposition may be the preferable technique. First, eigen decomposition is absolutely foolproof for any symmetric - like a correlation - matrix. This makes it a handy component in any automated system. Second, using the principal components of the correlation matrix allows us to explain the relationships in a large matrix with a relatively small number of equations and to model as much of the interrelationship as necessary. The result of the eigen decomposition will be a series of N eigenvectors and their associated eigenvalues. The eigenvectors describe linear relationships between the variables and their respective eigenvalues correspond to the relative importance of each relationship. As such, we can choose eigenvectors to explain 95%, 99% or even 100% of the variation in the correlation matrix.

For example, the first three principal components of the crude oil forward curve explain over 99% of the relationships in the correlation matrix.

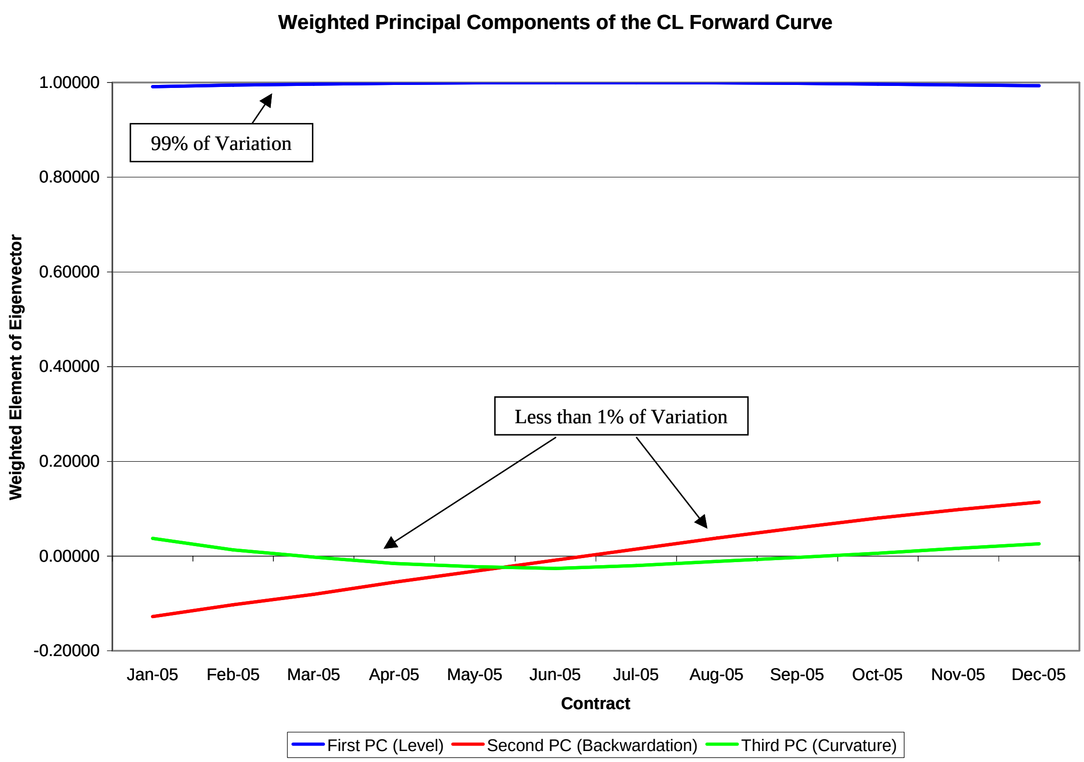

The first principal component is the level component, and reflects the tendency of the crude oil forward contracts to move up or down together. The second principal component reflects the tendency of the curve to change its tilt, while the third reflects the tendency to change curvature.

Similarly, the first four principal components of the natural gas forward curve explain 99% of the relationship in that correlation matrix.

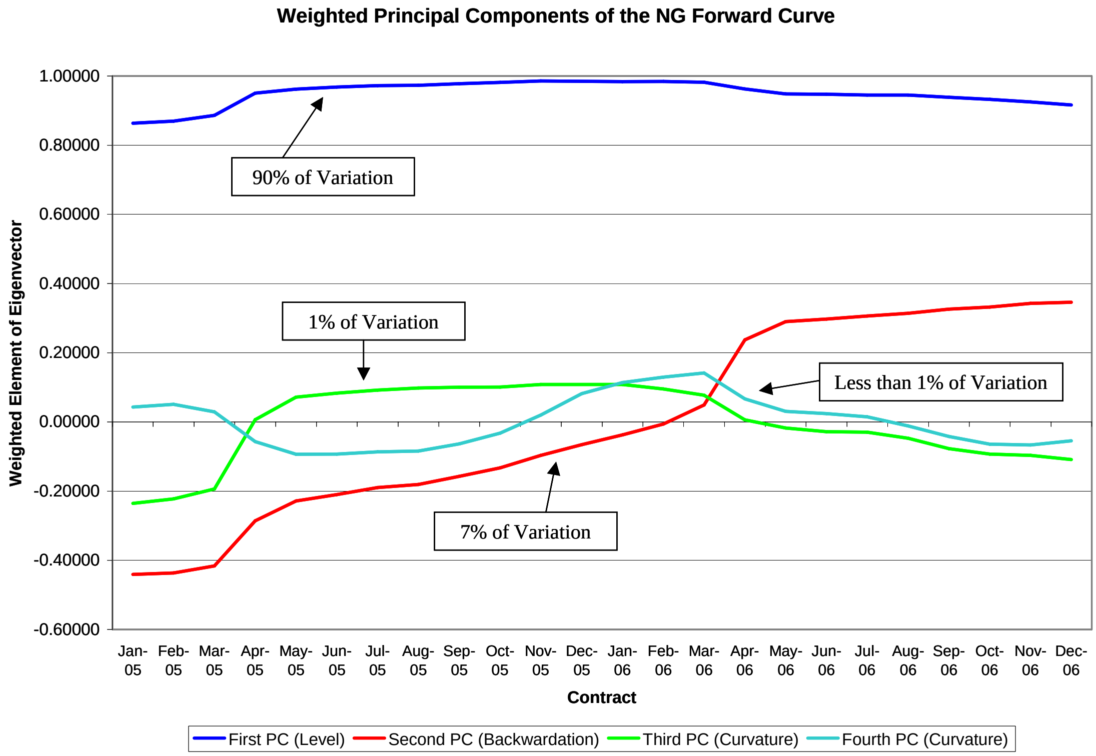

As with crude oil, the first principal component is the level component, but we can see that there is less of a tendency for natural gas contracts to move up or down together. The second principal component again reflects the tendency of the curve to change its tilt, while the third and fourth reflects the tendency to change curvature. The third principal component shows the tendency of the April 2005 through May 2006 strip to move counter to the rest of the curve. The fourth component measures the changes in the relationship between winter and summer contracts.

Finally, we can explain 99% of the relationship between the natural gas and crude oil curves with five principal components.

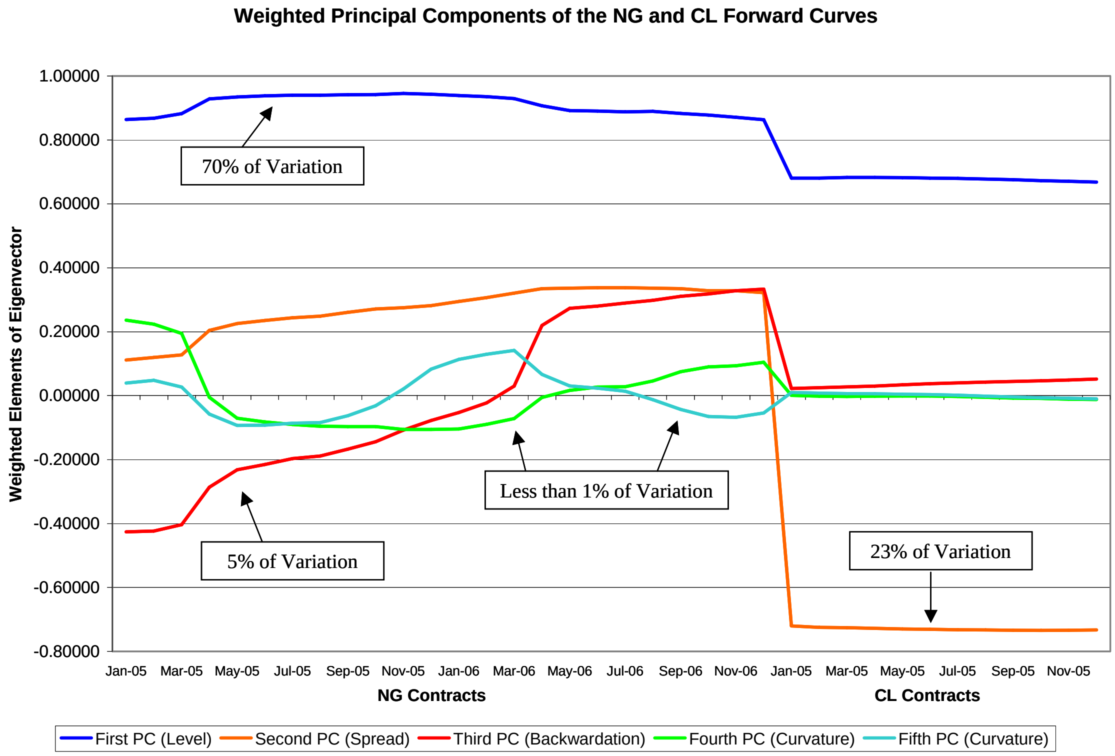

Four of these components should look familiar; they are the level, tilt and curvature components that we saw in the first two curves. However, we have added a new component, and it explains 23% of the variation of the two curves. The first principal component, the level component, explains much less variation when two commodities are considered. While different contracts of the same commodity may exhibit a greater tendency to move together; among different commodities it is less so. This second principal component reflects the tendency of these commodities to move in opposite directions.

But how does all of this work together? Let's consider a simple example of an electricity producer that owns natural gas and oil-fired generation. This entity hedges price exposure in the physical fuels markets with the purchase of crude oil and natural gas futures on the NYMEX. This hedging strategy itself is subject to corporate value at risk limits, while the futures contracts are subject to exchange margin requirements. While considering a hedging strategy, then, this entity has to not only find one that mitigates price risk, but one that can still be maintained while conforming to value at risk and margin requirements. These last two considerations are often overlooked. Even if you, as hedging specialist, devise the greatest hedging strategy in financial history, constraints outside of your immediate control can force liquidation of the portfolio, and no one would ever know.

So, it's early December, and this entity has decided to purchase 50 natural gas contracts for each month of 2005, 25 natural gas contracts for each month of 2006, and 5 crude oil contracts for each month of 2005. This portfolio has a notional value of approximately $62 million. The salient questions are: where are the risk limits, where are the stop limits, and how much capital is in reserve for possible margin calls?

We can use our volatility term structure model and the principal components of the correlation matrix to simulate natural gas and crude oil curves. Five such natural gas price curves may look like this:

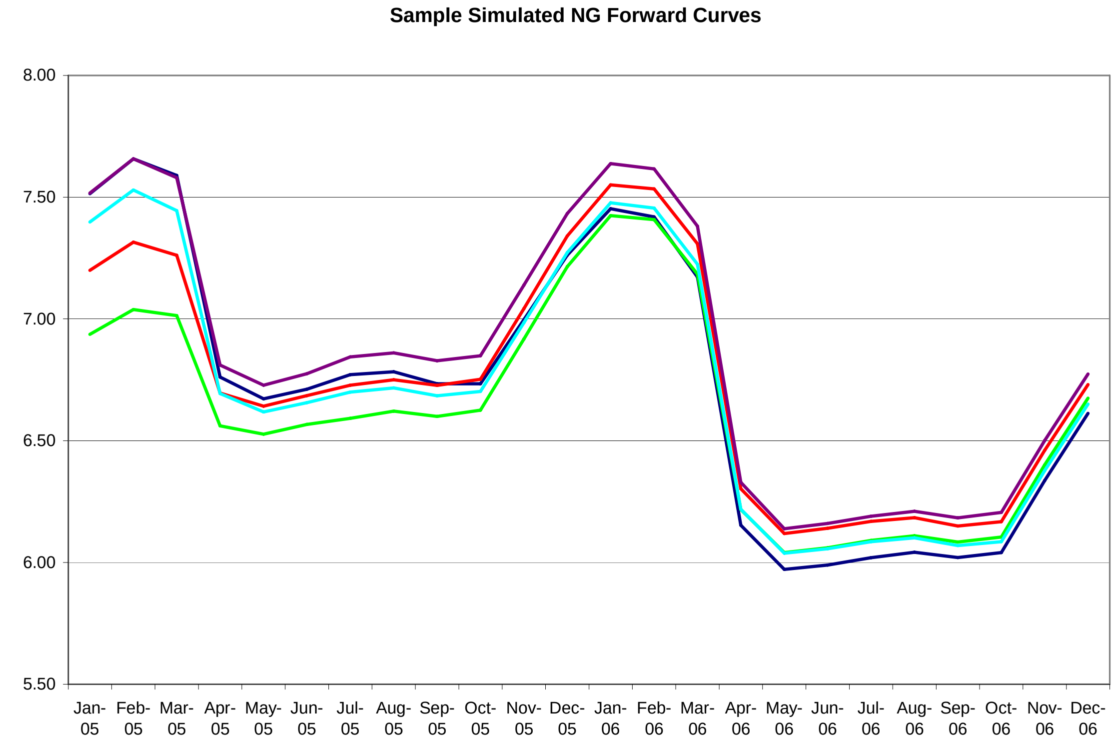

And five simulated crude oil curves may look something like this:

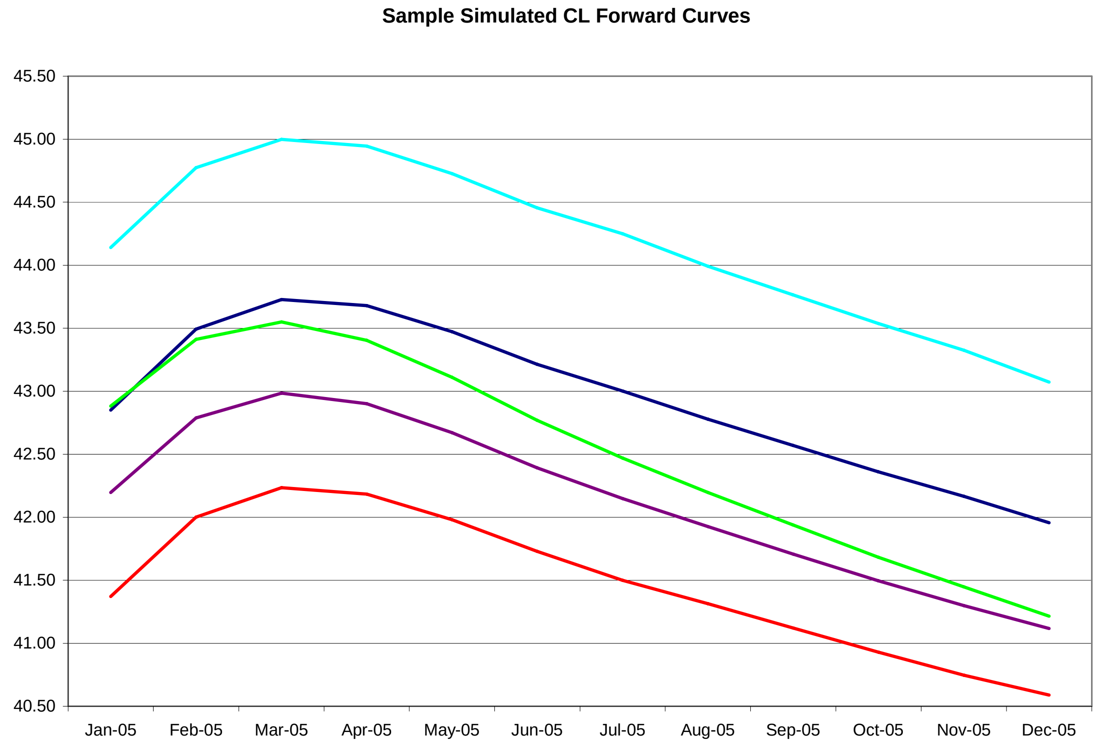

After simulating curves for five days, the confidence bands for the natural gas and crude oil forward curves as of mid-December 2004 might look something like this:

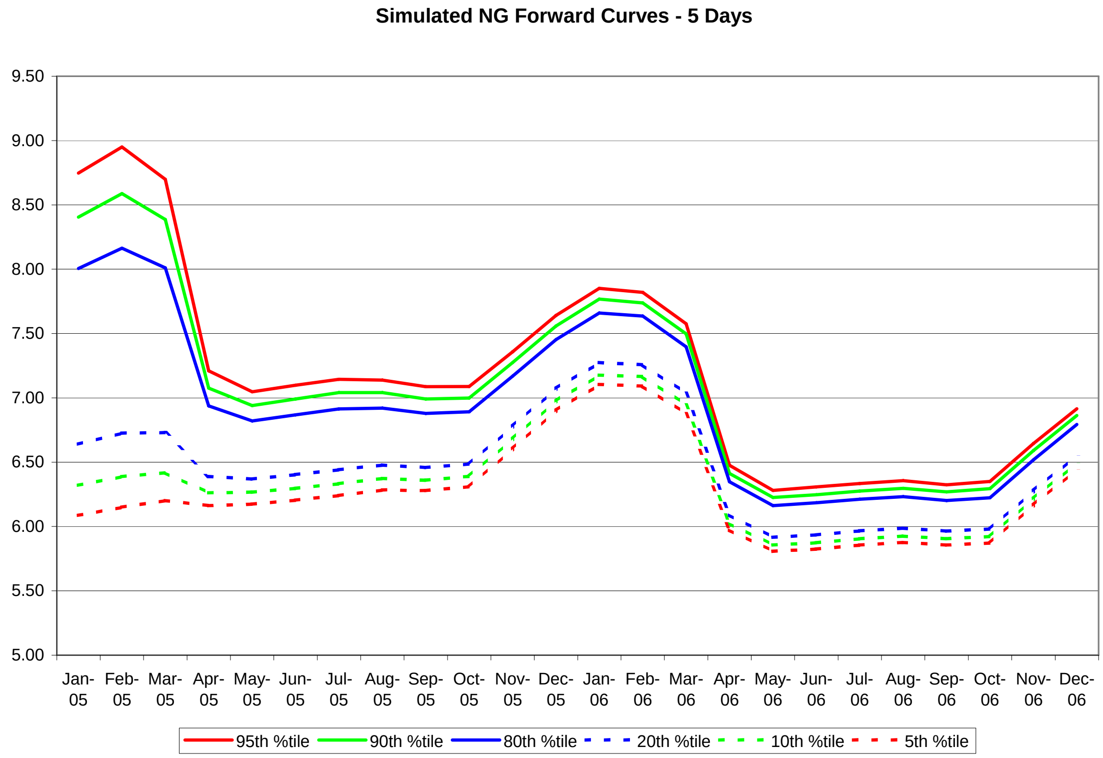

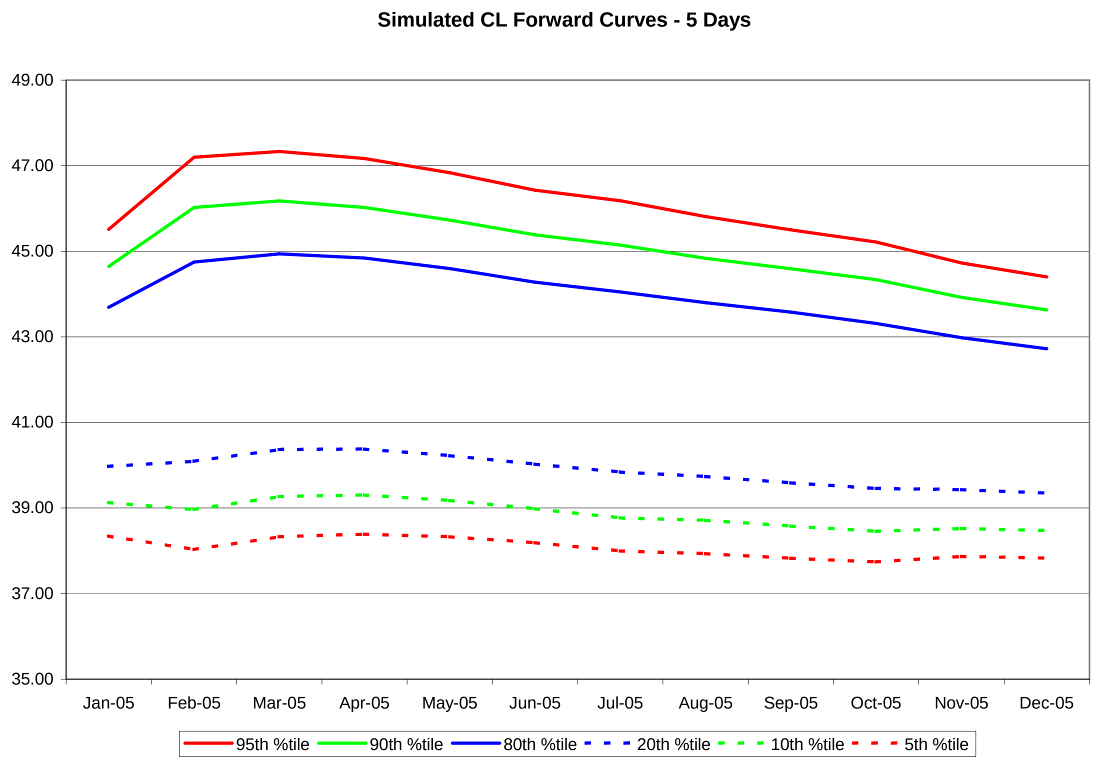

Note the considerably wider confidence bands for the front months of the simulated natural gas curves. The crude oil curves, with their greater tendency to move together, don't exhibit this behavior to the same degree. The confidence bands for the mark to market value of this portfolio might look something like this:

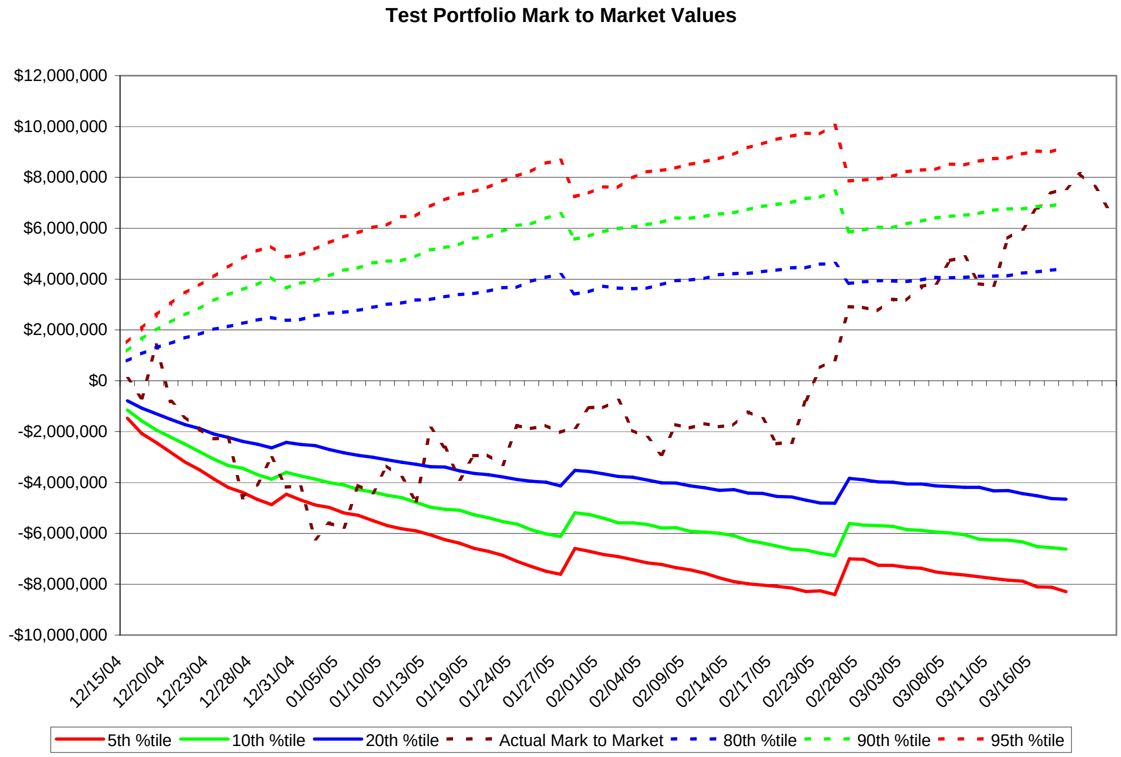

The jumps in the bands are the days that the natural gas positions expire. There are similar jumps in the middle of the months when the crude oil contracts expire, but they are a relatively small component of the overall portfolio and therefore do not stand out. The dotted brown line is the actual mark to market of the portfolio through the first quarter of 2005. During this time span, 41% of the actual values fell outside the 60% confidence band, 18% fell outside the 80% confidence band, and 6% fell outside the 90% confidence band. This is consistent with our expectations, but a successful backtest is not the point of this exercise. When purchased, the notional value of this portfolio was about $62M, and if you were not prepared to suffer mark to market losses of $6.3M at some point (as happened in the middle of January), the positions would not be in the money as of the end of March, they would have been liquidated.

So where do we go from here? Pindyck (2004) and Alexander (2004), for example, have worked to fit GARCH models to futures prices, but daily simulation of entire forward curves under such a framework could be an arduous task. Our own opinion is that it is more important to model the time-varying volatility than any autoregressive behavior that might exist. We are concerned, however, about possible excess kurtosis in the daily return distribution. The next refinement we would make is to fit the volatility functions with either a generalized error distribution or Student's t-distribution. The strength of this model remains the ability to incorporate time-varying volatility in a tractable framework.

References

Alexander, C., 2004, "Correlation in Crude Oil and Natural Gas Markets", in V. Kaminski (ed.) *Managing Energy Price Risk*, Third Edition, Risk Books, pp. 573-601

Berndt, E.K., B.H. Hall, R.E. Hall, and J.A. Hausman, 1974, "Estimation and Inference in Nonlinear Structural Models", *Annals of Economic and Social Measurement*, Vol 3/4, pp. 653-665

Clewlow, L. and C. Strickland, 1999, "Valuing Energy Options in a One Factor Model Fitted to Forward Prices", *working paper*, University of Technology, Sydney

Clewlow, L. and C. Strickland, 2000, *Energy Derivatives: Pricing and risk management,* Lacima Publications.

Clewlow, L., C. Strickland, and V. Kaminski, 2001, "Extending Mean-Reversion Jump Diffusion", *Energy Power Risk Management*, February 2001.

Clewlow, L. and C. Strickland, 2004, "Simulating Spots", *Energy Risk*, May 2004, pp. 48-51

de Jong, C. and R. Huismann, 2002, "Option Formulas for Mean-Reverting Power Prices with Spikes", *working paper*, Energy Global/Erasmus University Rotterdam

Goldfeld, S., R. Quandt, and H. Trotter, 1966, "Maximization by Quadratic Hill Climbing", *Econometrica* 34, pp. 541-551

Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press.

Lucia, J. and E. Schwartz, 2001, "Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange", *working paper*, UCLA

Pindyck, R., 1999, "The Long Run Evolution of Energy Prices", *The Energy Journal* 20:2, pp. 1-27

Pindyck, R., 2004, "Volatility in Natural Gas and Oil Markets", *working paper*, MIT.

Schwartz, E., 1997, "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging", *Journal of Finance* 52:3, pp. 923-973

