Alternative Valuation Methods for Swaptions: The Devil is in the Details

By Carlos Blanco, Josh Gray and Marc Hazzard

Introduction

Options on swaps (or swaptions) are one of the most difficult derivative instruments to value. The main reason is that their prices are heavily influenced by the properties of the forward contracts with maturities over the term of the underlying swap. These properties include:

- A term structure of implied volatility for each of the forward prices that are part of the swap
- A term structure of co-movement between contracts with different maturities

In this paper, we will present the most commonly used methods to value swaptions (e.g. one-factor, multi-factor and alternative approaches), and address some of their key problems and pitfalls. We will also show that a multi-factor forward curve evolution model can more naturally address these problems, while accommodating the necessary set of parameters needed to calibrate the model.

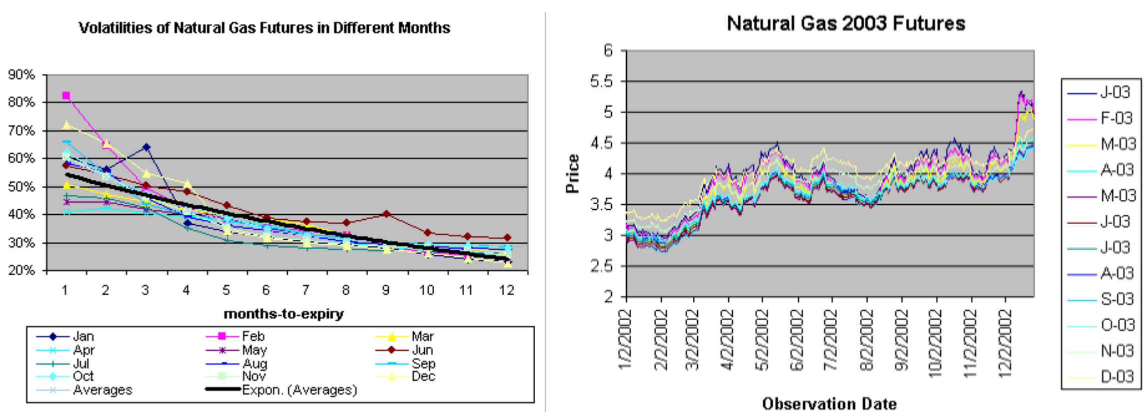

Figure 1: Term structure of Natural Gas futures' volatility and Historic futures prices

One-factor models: Use with extreme caution

The usual analytical valuation of a swaption relies on a one-factor forward curve evolution model. The one-factor model assumes that the dynamics of the full forward curve are driven solely by the spot price. In other words, the spot and all forward contracts are perfectly and positively correlated with the same stochastic component. As such, this underlying assumption ruins the reality of co-evolution of the forward curve for a variety of reasons:

1. While futures contracts do generally move in tandem, they certainly do not exhibit *perfect* correlation. In fact, the observed correlation is *not* itself a stable parameter over long time periods and exhibits persuasive levels of time-dependence.

2. As essentially a spot price model, it is difficult to consistently and robustly incorporate information about the term structure of implied volatility (as derived from options on the forward contracts), and the volatility of the spot price (as possibly derived from Gas Daily-type options)[1].

3. Furthermore, there is an ambiguity in calibrating the necessary model parameters. It is unclear whether to use market information about the spot price (as the one-factor model suggests) or market information about the relevant forward contracts.

In the case of an option on a calendar swap with monthly forward prices, there are 66 correlation coefficients (a 12x12 symmetric correlation matrix minus the unit diagonals) and 12 implied volatilities that we cannot fit directly into the model. Indeed to provide a monthly adjustment of these parameters requires 12 sets of these parameters!

As most energy commodities are tied to some cost of production, shocks to the spot price of the asset are dissipated away over certain characteristic time spans. This time period ranges from hours (power prices) to years (oil prices) depending on the setting in which it is observed and the subsequent manner of its calibration. In this calibration, as is popular with assets that are physically settled or where historical forward prices and option quotes are scarce, regression analysis of the historical spot price can be used to determine the mean-reversion rate. For calibration of the more subtle long-term mean-reversion effect, as is observed in a strongly backwardating implied volatility curve, we return to historical or implied forward contract volatility for option prices for the parameter estimation.

Regardless of the manner in which it is calibrated, the confusing matter is that there are in fact two mean-reversion rates, the mean reversion of the shocks to the spot price and the mean reversion of the forward price volatilities. This poses additional problems that one must face in applying a one-factor model.

For consistency, one would hope that these two parameters would agree, at least approximately. However, there is no guarantee that these two parameters should even be in the same order of magnitude! The common example is from the natural gas markets where both effects can be observed. Spot price analysis can show calibrated rates showing reversion to the mean on the order of one-month to several days, depending on the location. The backwardation fit can also be performed, which yields values showing

[1] It is worth noting for clarification, that with the one-factor model, a term structure of implied volatility can be used to price a *strip of options* on forward contracts of different maturities (say, coinciding with the expiry dates of the options in the strip). However, this should not be confused with the use of a term structure of implied volatility in the valuation of an option on a *strip of forward* contracts.

very little (i.e. the Black ’76 limit) to very slow signs of mean-reversion (on the order of years or several months).

Multi-factor models: A natural solution to the limitations of one-factor models

Multi-factor models allow for establishing different “drivers” that determine the evolution throughout time of the spot and the different forward prices that comprise the swap and determine the value of a swaption.

The starting point for the model is to establish the historical variance and covariance profiles of each of the forward contracts. Recognizing that energy commodities exhibit some degree of seasonality, these profiles are established from historical data on a seasonal basis. For example, isolating observations by month for each historical year is a reasonable starting point. In each of these monthly bins, we can then derive the scalar correlation between forward contracts of different maturities and the average level of volatility of each of the different forward contracts.

At this point, we find it convenient to identify forward contracts of different maturities independent of a fixed maturity date. In this “fixed maturity” basis, we identify the prompt-month contract as the “1st nearby”, and so on. So for instance, in the historical October 2000 observation record, the contract slated for deliver in November 2000 is the 1st nearby, and the contract delivering through December 2000 is the 2nd nearby, etc. As we roll into historical November 2000 observations, December 2000 is promoted to the 1st nearby status, and the November 2000 contract is now in delivery (i.e. some combination of spot per observation date, and balance-of-the-month for remainder of November). Further, the November *2001* contract is 12th nearby. This identification helps us then identify how the volatilities and correlations of contracts in this rolling basis behave from month-to-month, over a period of several years.

Having completed this procedure over all the available years of historical data, we will have established seasonal correlation matrices and volatilities. While the further calibration to implied correlation is typically problematic (option markets on spreads may not be observable in the market), the calibration of the term structure of volatility to current market levels is necessary.

To perform this fit to levels of volatility currently implied by the market, we track the seasonal volatilities on a fixed maturity basis. Let’s say that on August 3, 2003 we have acquired a set of implied volatilities from options on forwards beginning with the prompt-month (September) and going out several months. Beginning with the August historical volatilities, we scale the volatility of the 1st nearby (September) to the level suggested by September’s implied volatility. Going out then to the following contract’s implied volatility, we must match this to both the volatility of the 2nd nearby observed in August and the 1st nearby observed in September. We do this in such a way as to conserve the volatility of the maturity of the option through a time-weighted averaging.

This re-calibration of the seasonal volatility to the levels currently seen in the market is then done for every available point. Given the newly re-scaled seasonal volatility and the seasonal correlation matrices, we could then realize a seasonal multi-factor simulation model through the correspondingly derived seasonal variance-covariance matrices.

From this point of view, the crux of the work has been resolved, but a further computational simplification can be performed. Returning to the example wherein 12 forward contracts are relevant (e.g. a calendar swap), the resulting 12 x 12 VC-matrix gives all the correlated pieces need for a joint Monte Carlo simulation. However, we have found that in most cases, a reduction of the problem to a set of uncorrelated factors through Principal Component Analysis (PCA) yields a computationally simpler problem, while allowing for a majority of the variance (90%+) to be accounted for through 3 to 5 of these factors (the principal components).

From seasonal VC-matrices we may derive a seasonally-based PCA. Interestingly enough, seasonal effects that were implicitly contained within the seasonal VC-matrices but were somewhat obscured, are eloquently captured through a glimpse of the seasonal PCA calibration. For example, one might expect that at certain times of the year the maximum magnitude of spreads between two forward contracts can change to another pair. As an illustration consider that in April the September-December spread is the greatest. This could be explained through the PCA by showing that the second factor demonstrates a large separation between the September and December contracts' effective volatility. Moreover, moving into May the separation between September and November is the greatest, leading to a greater widening of the spread than that of the Sept-Dec pair in the previous month.

What now remains is to establish consistency among the mean-reversion rates. In the one-factor model, a fit of the time-to-maturity effect to a decaying exponential was sufficient to establish the mean-reversion rate. Through the exponential form, a link between the term structure of spot volatility and the volatility of forward contracts was then forged. In the multi-factor world, we do not need to make an attempt to reconcile the calibration and fit of the forward contract's volatilities with the term structure of spot volatility. Indeed, so long as the time-to-maturity effect is truly seen in the historical calibration, it is encoded in the VC-matrix and the subsequent PCA analysis. A term structure of spot price volatility can be independently considered if deemed necessary. Furthermore, even if the time-to-maturity effect is not 100% valid (e.g. power markets), the specific form of the term structure is immaterial.

The mean-reversion rate that describes the effect of the spot price relationship with the prompt month contract can be accounted for in the multi-factor model through other means[2]. FEA has developed a two-factor model of the spot price in conjunction with the multi-factor forward curve evolution model. These two factors are an independent random variable (perhaps, *the* factor in a one-factor model) and a variable correlated to

[2] Although we've described a model of the full forward curve that is independent of the spot price model, a coincident spot simulation is sometimes needed to resolve balance-of-the-month issues in the context of swaptions.

the random shocks to the prompt-month contract. In this way, the spot price model contains both the effect of mean-reversion to the prompt-month contract, and a component that embodies the remaining and typically strong correlation between physical delivery and financial contracts. A full term structure of spot price volatility can be incorporated as well, allowing for all the details of the one-factor model as well as the rich behavior described by the multi-factor approach.

Alternative Valuations

In addition to the one-factor and multi-factor swaption pricing approaches we have described above, there is an alternative valuation that is pursued by traders. In the instance when the market quotes swap prices, as opposed to the strip forward prices, it is tempting to estimate the volatility of the swap itself. Then the standard Black commodity option pricing formula can be applied to the swap directly.

The main benefit of this approach is that we only have one unknown parameter (volatility) until the maturity of the swaption. If the underlying swap is quoted in the market, the swaption can be delta hedged. And if there is a liquid market for that swaption, a single implied volatility can be obtained from option quotes.

However, there are several pitfalls in applying this simplistic pricing method. The foremost being that the market for most swaptions is fairly illiquid and therefore the volatility of the underlying swap contract cannot be reliably implied. Alternatively, estimating swap volatility from historical analysis is difficult if swap prices are not routinely quoted. In both instances, proponents of the Black approach typically use the implied volatilities of the individual forward contracts. Then, combining the sum of the squares of these volatilities with cross-terms proportional to the correlation between the forward contracts, then the analysis yields the approximate lognormal volatility of the swap. This is generally a good estimate of the swap's volatility if the swap correctly reflects the fair-value (i.e. with no fixed-price or risk-adjustments). Being able to reconcile the swap quote to the prices of the ingredient forward contracts should suffice to determine the validity of this approximation.

Of course, estimating a full correlation matrix for the forward curve in an opaque market is not an easy task. Implied correlations from calendar spread options are also very difficult to obtain due to lack of liquidity. If historical forward prices are available, the correlation can be estimated, but given a substantial amount of historical information, the multi-factor model could be used.

Comparisons of Results Obtained from Different Models

In looking at actual swaption pricing, given an arbitrary amount of market data, there are actually several ways one could arrive at a swaption value in both the one- and multi-factor worlds. We will look at the 5 most consistent, and prevalent models:

1. **Full-blown seasonal Multi-Factor Model** ("MF")- Makes use of seasonal volatilities and correlations for the different forward contracts.

2. **Reduced Form Multi-Factor Model** ("MF-R")**-** Uses an average correlation matrix, an average term structure of volatility and possibly a mean-reversion rate that was derived from the seasonal analysis for the multi-factor model.

3. **Forward-driven One-Factor Model with Mean Reversion** ("1F-M fwd")**-** Uses a term structure of volatilities and a single mean reversion rate calibrated from the volatility of different forward contracts

4. **Spot-driven One-Factor Model with Mean Reversion** ("1F-M spot) - The parameters are calibrated from historical spot information.

5. **Black '76** ("Black")**-** Treating the forward curve as a unique underlying that determines each swap price and deriving the lognormal volatility as the sum of the squares of the average volatilities with cross-terms proportional to the average correlation.

A comparison of the results obtained through the multi-factor and a one-factor forward curve evolution model can be made after some attention is brought to the model parameter calibration. Given the term structure of volatilities of forwards (calibrated either from a purely historical analysis, or through an implied-adjusted fit thereof) the one-factor model's spot price volatility curve and mean-reversion rate can be derived. At this point, any comparison between the two models is appropriate.

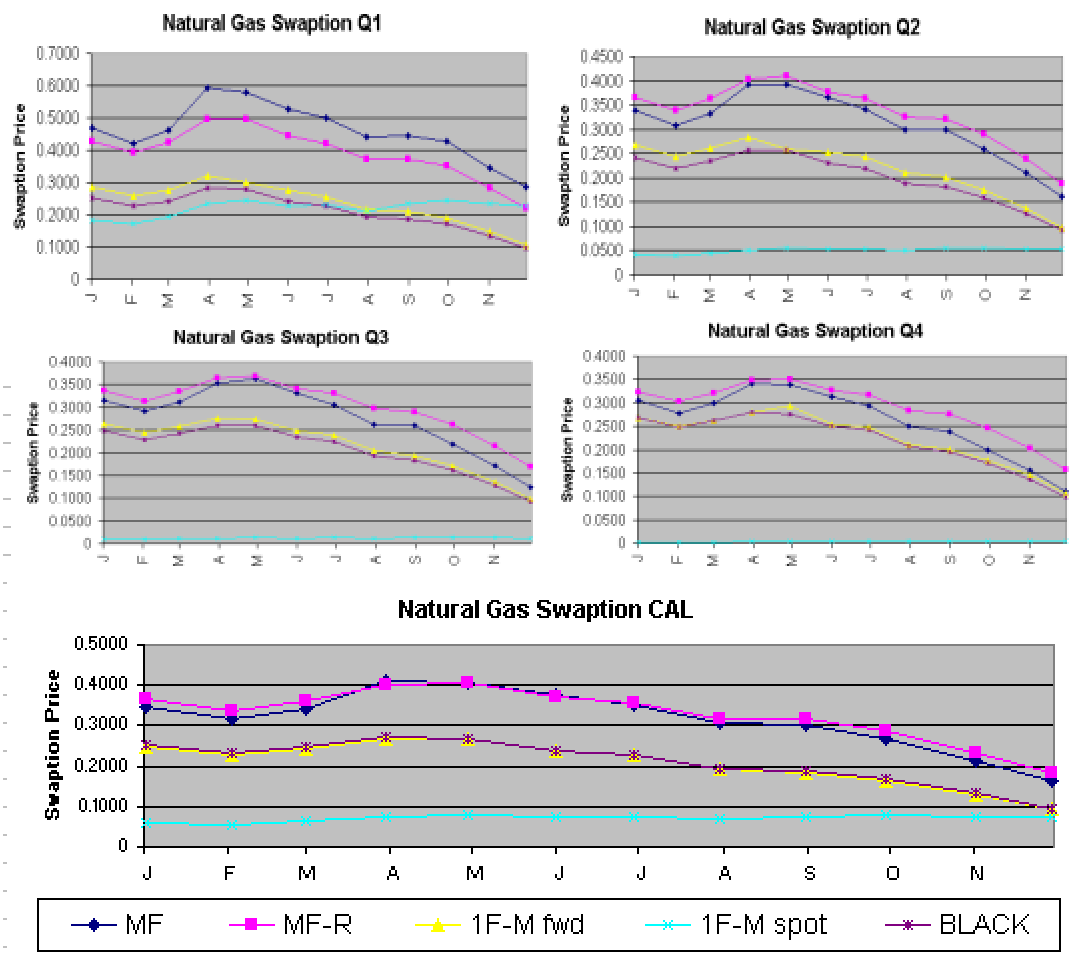

Figure 2: Comparison of pricing models for natural gas swaptions

Our model comparisons look at these five models on options with one-year to one-month maturities on four quarterly and one calendar natural gas swap. The full multi-factor model ("MF") and its reduced form ("MF-R") separate themselves from the others and present relatively larger swaption prices with significant elements of seasonality. One might expect that the *multi*-factors that drive this forward curve model are responsible for the higher relative prices. On the other hand, the forward-tuned, one-factor model ("1F-M fwd") and the "Black" model have nearly identical pricing behavior. Given that both are one-factor models of the swap, it is not surprising that their results are similar and of a lesser relative value than the multi-factor models. Finally, the one-factor model tuned to the spot price dynamics gives much lower swaption values consistently through the term of the contract. Again, this is due to an overestimate of the relevant mean-reversion rate through historic spot prices, and a subsequent substantial reduction in the swap's volatility.

Summary

In commodities where a forward market can be observed with a regular frequency (oil and natural gas and, in a limited context, power markets), a multi-factor model of the forward price curve can be calibrated and used to capture the known correlation and volatility effects.

As the underlying of a swaption is the full forward curve, capturing the richness of the movement of the forward curve highlights the need for a multi-factor simulation model for pricing. In comparison, applications of simpler approaches, while expedient, miss the volatility and correlation effects that can easily be specified in a multi-factor setting. While we then lose analytical tractability and forced into a Monte Carlo valuation setting, use of Principal Components Analysis (PCA) can dramatically reduce the dimensionality of the problem and result in a faster analysis.

Dr. Carlos Blanco directs the Global Support and Professional Services Group at FEA, a Barra Company. Dr. Josh Gray and Marc Hazzard are Senior Derivatives Specialists in the same group. The authors thank Craig Jimenez, David Soronow, Michael Pierce, Rajnish Kamat, and Dr. Angelo Barbieri for advice and assistance.