Financial Modelling Agency

Contents

Financial Modelling Agency

Tel: +27-11-447-2901
Fax: +27-11-447-2901

Website:
www.finmod.co.za

General E-mail:
graeme@finmod.co.za

Physical Address:
19 First Avenue East
Parktown North
Gauteng
2193
South Africa

24 September 2009

Authors:
Graeme West
graeme@finmod.co.za

A SUMMARY OF OUR APPROACHES TO THE SABR MODEL

1 The need for a stochastic volatility model

Vanilla OTC European options or European futures options are priced and often hedged using respectively the Black-Scholes or Black model. In these models there is a one-to-one relation between the price of the option and the volatility parameter σ, and option prices are often quoted by stating the implied volatility σ_{imp}, the unique value of the volatility which yields the option price when used in the formula. In the classical Black-Scholes-Merton world, volatility is a constant. But in reality, options with different strikes require different volatilities to match their market prices. This is the market skew or smile.

Typically, although not always, the word skew is reserved for the slope of the volatility/strike function, and smile for its curvature.

Handling these market skews and smiles correctly is critical for hedging. One would like to have a coherent estimate of volatility risk, across all the different strikes and maturities of the positions in the book.

To resolve this problem, in Hagan et al. [2002] the SABR model is derived. The model allows the market price and the market risks, including vanna and volga risks, to be obtained immediately from Black's formula. It also provides good, and sometimes spectacular, fits to the implied volatility curves observed in the marketplace. More importantly, the SABR model captures the correct dynamics of the smile, and thus yields stable hedges.

2 Building the model

The pricing of an option is via the ordinary SAFEX Black formula, with a skew volatility input. The skew volatility can be written as a function

$$\sigma_X = f(F, X, \tau, \sigma_{\text{atm}}, \beta, \rho, v)$$

where F is the futures level and σ_{atm} is the at the money volatility level, X is the strike, τ is the term in years of the option, and β, ρ and v are the specific model parameters.

Our first task then is to fit the parameters β, ρ and v.

Market smiles can generally be fit more or less equally well with any specific choice of β. In particular, β cannot be determined by fitting a market smile since this would clearly amount to "fitting the noise".

To use a value of $\beta \approx 1$ is most natural in equity markets, but it implies that the at the money volatility moves horizontally as the market increases or decreases. A value of $\beta < 1$ would indicate that the volatility decreases/increases as the market increases/decreases. Therefore such a value would be preferred.

Using the calibration methods prescribed in Hagan et al. [2002] one finds that it is appropriate to use a value of $\beta = 0.7$ for all expiries in the South African market. See West [2005].

ρ is the correlation between the underlying and the volatility. As such, it is negative. This parameter principally causes the skew in the curve.

ν is the volatility of volatility. This parameter principally causes the smile in the curve.

3 Calibrating the model

The values of ρ and v need to be fitted. For this, there are several possibilities, three of which we now elaborate.

The appropriate mechanism to compare a modelled skew to other skews is via the ability to reproduce accurately the pricing performance. Comparing the actual volatilities in each skew is meaningless if a significant number of structures have traded, as opposed to outright options. Given a record of trades, we can compare

- The prices that are or were recorded in the market.
- The prices that the model estimates.
- The prices that any other competing model of the skew estimates.

See Figure 1.

Subsequent to that, models compete on the basis of

- Hedge performance, which may be determined by ...
- ... parameter stability.

3.1 Fitting to a given market skew

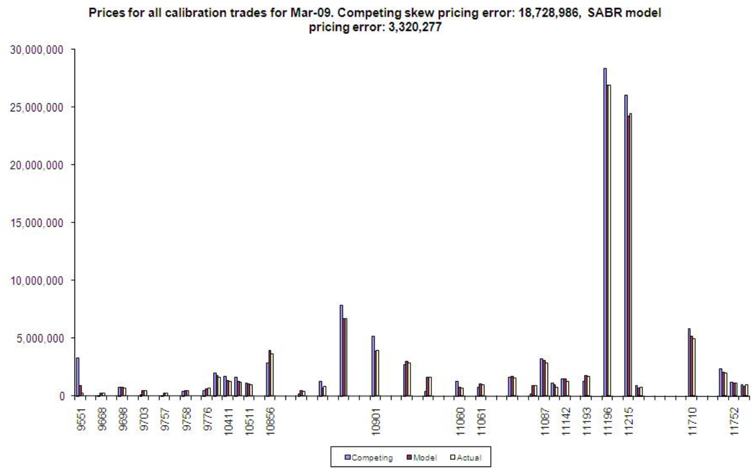

FIGURE 1: *Price replication*

It is possible to simply specify a discrete skew (input by the dealer, as observed in the market) and find the SABR model which best fits it. This means find the values (ρ, v) which minimise the distance from that SABR model to the dealer input.

The question arises as to what is meant by minimising the distance between the skews. We can aim to minimise the error in pricing on the SABR skew versus pricing on the input skew. An alternative would simply be to minimise the distances from the actual volatilities on the SABR skew to the volatilities on the trader skew.

Indeed, the input might actually be a bid skew and an offer skew. This time the error expression per input might be the distance to the closer of the bid or offer price if the price is outside that double, and zero if inside it. See Figure 2.

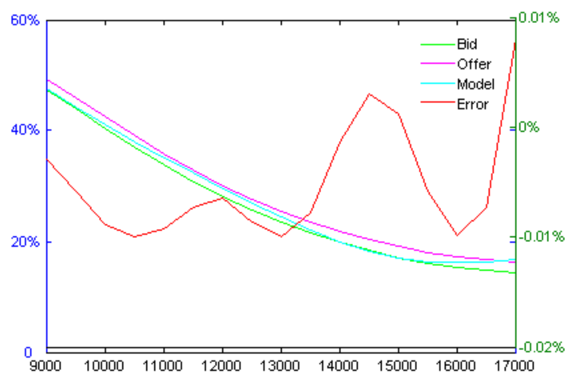

FIGURE 2: *Given an input dealer skew, we can find the SABR model which best fits the skew.*

3.2 Fitting to market data

We can set ourselves the task of finding the SABR model which best fits given traded data, independently of any dealer input as to the skew.

For any input pair (ρ, v), we determine an error expression $\text{err}_{\rho,v}$, which per trade is the distance between

- the currency cost that the trade was done at;
- the currency cost that the trade would have been done at if all the legs of the trade had been done on the SABR skew with that ρ and υ.

The total error across all trades is some sum of these errors. Thus

$$\text{err}_{\rho,\nu} = \sum_i |V_i(\rho,\nu) - P_i| \tag{1}$$

where P_i is the actual price of the i^{th} structure and $V_i(\rho,\nu)$ is the hypothetical price if the structure had been priced off the skew with parameters ρ and ν.

The trades that have been observed in the market may be weighted for age, for example, by using an exponential decay factor: the further in the past the trade is, the less contribution it makes to the optimisation. Trades which are far in the past might simply be ignored.

The pair of parameters (ρ,υ) are found which are most reasonable i.e. minimise the residual error $\text{err}_{\rho,\upsilon}$. Mechanisms for achieving this are developed in West [2005] and West [2006]. As one can see in Figure 3 - this result is typical - the choice of parameters is fairly robust, with the minimum found at the bottom of a shallow valley.

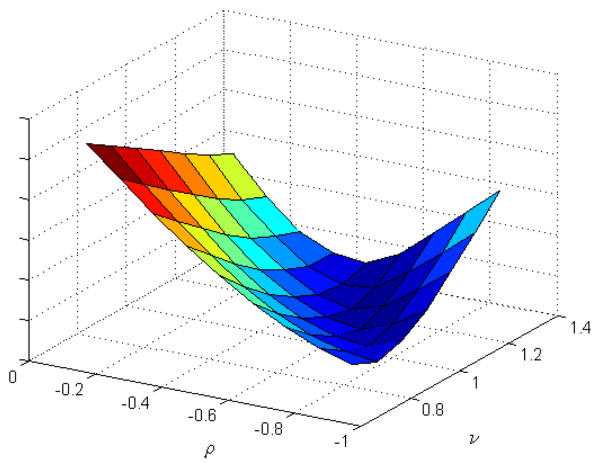

FIGURE 3: *The error quantities for ρ and ν: March 2007 calibration.*

As pointed out in Hagan et al. [2002], the idea is that the parameter selection change infrequently (perhaps only once or twice a month) whereas the input values of F and σ_{atm} change as frequently as they are observed. This is in order to ensure hedge efficiency.

Having found the skew, we can then recalibrate the history of trades that have been used to build the model to that skew. This is a fairly tricky task, dealt with in detail in West [2005]. See Figure 4.

3.3 Fitting to bid-offer market quotes of structures

Typically brokers offer a bid-offer spread on a variety of deals: both outright deals, but also various structures. The bid-offer double will be quoted in terms of volatility, but of course this is easily converted to a double on prices.

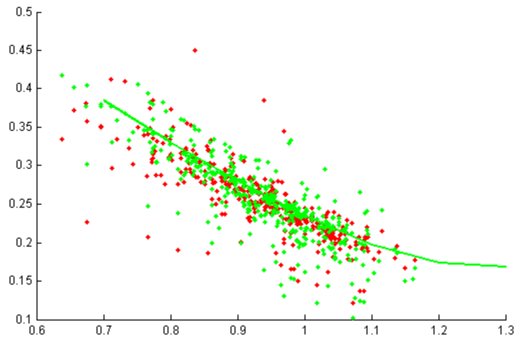

FIGURE 4: *The SABR model for March 2005 expiry, with traded (quoted) volatilities (green), and with strategies recalibrated to a fitted skew (blue), and the fitted skew itself (solid line).*

We can take this set of quotes as being the universe of trades for inclusion in a model which, from an implementation approach, is very similar to the previous one. The error expression will again be the distance to the closer of the bid or offer price if the price is outside that double, and zero if inside it.

4 SABR in the risk process

4.1 Mark to Market process

We apply the process seen expiry by expiry.

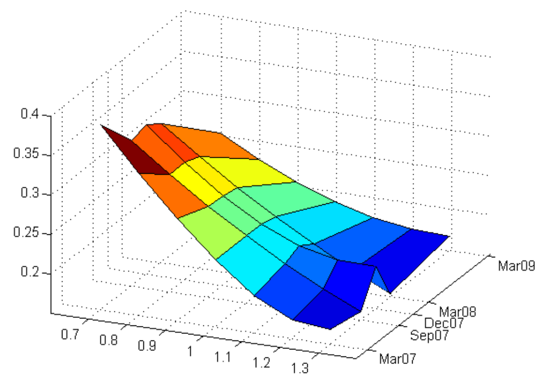

FIGURE 5: *The skew surface corresponding to several instances of the process already seen.*

The SABR model can be used as a very effective surface interpolator. In order to find volatilities at any strike or tenor, one finds interpolated F, β, ρ and ν parameters and then applies the formula as given. For this linear interpolation is completely satisfactory.

In particular, deals for non-standard expiries (or illiquid expiries, and so not yet part of the model) can be priced with a reasonable degree of confidence. For options with any sort of path dependence, the dynamic SABR model needs to be implemented Hagan et al. [2002].

Skew volatilities for non-index products - for example, single equities, or baskets of equities - can be modelled. If the underlying is sufficiently 'index-like' then the β, ρ and v parameters can be taken from the index, the futures level can be taken to be the forward level, only the at the money volatility will need to be modelled. With sufficient liquidity one can take this to be the traded volatility. Alternatively, the proposal is to define

$$\sigma^P_{\text{atm}} := \frac{\sigma^P_h}{\sigma^I_h}\sigma^I_{\text{atm}} \tag{2}$$

where σ_h is the historical volatility of the underlying product (single equity or basket) P or the index I. To calculate this historical volatility relies on having a significant history of both the stock and the index (at least two years of data).

4.2 Stress process

In stress experiments only the spot or futures level and the ATM volatility need to be stressed, other variables are left alone, the parametric form of the skew takes care of itself.

Thus, one calculates F_{stress} directly or as $S_{\text{stress}}e^{r_{\text{stress}}\tau}$ and $\sigma_{\text{atm, stress}}$. One assumes that the values of β, ρ and ν remain unchanged. This is a completely reasonable assumption in the short time period that is typical in stress and VaR calculations; typically these parameters only change every few weeks.

4.3 VaR process

VaR can be calculated using the simulated values of the futures level (or spot level) and the at the money level, and the assumption can be made that the parameters β, ρ and v will not change. For the short time period under consideration, this assumption is perfectly reasonable. Thus, one has a rich skew model - determining skew volatilities for all the different relative strikes under consideration in the VaR experiments.

Let us consider the use of the SABR skew in value at risk calculations in more detail.[1] At time $t+1$, the volatility to use will be a function of time $t+1$ values of F, σ_{atm}, τ, β, ρ and ν. For this,

- τ is updated by decrementing by one business day, i.e. $\tau = T - \text{NBD}(t)$ where T is the expiry date of the option and $\text{NBD}(t)$ is the next business day of the valuation date t. Here τ is measured in years.
- F^i will be found using the chosen VaR method. This might be using the futures market, or it might involve spot and risk free rate updates, an assumption about dividends, and the assumption that forwards and futures coincide.
- At the money volatility will be given by the historical experiment

$$\sigma^i_{\text{atm}} = \sigma_{\text{atm}}(t)\frac{\sigma_{\text{atm}}(i)}{\sigma_{\text{atm}}(i-1)}$$

[1] Our notation will be as follows: Today will be assumed as a business day and will be denoted as time t, the previous business day time $t-1$. the business day before that $t-2$, etc. The next business day will be denoted time $t+1$, the business day after that time $t+2$, etc. These are business days, not calender days.

Suppose our number of historical experiments required is 250, so we have 251 many data points available.

i will denote an arbitrary date ie. i is a variable. The time that any time series starts will be date $i = 0$. This might be different for different variables, this is not a problem, as the calculations that involve these indices are separate for each variable. But it will be assumed that we have observations for at least dates $i = t - 250, \ldots, t$ ie. at least today's and the 250 previous day's observations, for a total of 251 observations. This means that we have at least 250 returns (one less because from n prices we have $n-1$ returns) which are the essential ingredients for 250 historical based scenarios.

Given a market variable $x(\cdot)$, let x^i be the value of the variable at time $t+1$ given the history of $x(\cdot)$ up to an including time t, the value generated by the occurrence at time i, and the choice of historical updating method: classical historical, Hull and White [1998], etc.

- We assume that the values of β, ρ and ν remain unchanged.
- We find the skew volatility as in Section 2.
- Apply the appropriate Black-type function using the skew volatility as input.

References

Patrick S. Hagan, Deep Kumar, Andrew S. Lesniewski, and Diana E. Woodward. Managing smile risk. *WILMOTT Magazine*, September:84–108, 2002. URL http://www.wilmott.com/pdfs/021118_smile.pdf. 1, 2, 4, 5

John Hull and Alan White. Incorporating volatility up-dating into the historical simulation method for V@R. *Journal of Risk*, 1(Fall), 1998. 6

Graeme West. Calibration of the SABR model in illiquid markets. *Applied Mathematical Finance*, 12(4):371–385, 2005. 2, 4

Graeme West. Another two algorithms for calibration of the SABR model in illiquid markets, 2006. URL http://www.finmod.co.za/SABRanother.pdf. 4