A SUMMARY OF THE APPROACHES TO THE SABR MODEL FOR EQUITY DERIVATIVE SMILES

GRAEME WEST, RISKWORX

1. The need for a stochastic volatility model

European options or fully margined (SAFEX) American options are priced and often hedged using the Black-Scholes or SAFEX 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.

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.

Date: December 1, 2005.

To resolve this problem, in (Hagan, Kumar, Lesniewski & Woodward 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

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

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).

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

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

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

2.1. **Fitting to a given market skew.** 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 1.

2.2. **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.

As already discussed, we fix in advance the value of β. Then, for any input pair (ρ, v), we determine an error expression $\mathrm{err}_{\rho,v}$, 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 v.

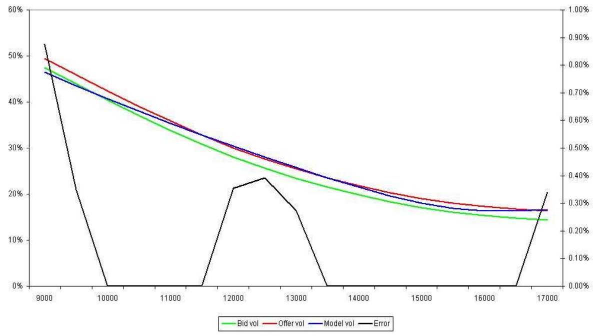

FIGURE 1. Given an input dealer skew, we can find the SABR model which best fits the skew.

The total error across all trades is some sum of these errors. 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, v) are found which are most reasonable i.e. minimise the residual error $\mathrm{err}_{\rho,v}$. This is achieved in (West 2005).

In Figure 2 we see how the pair of parameters are found uniquely at the bottom of a reasonably smooth and shallow valley.

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 3.

2.3. **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

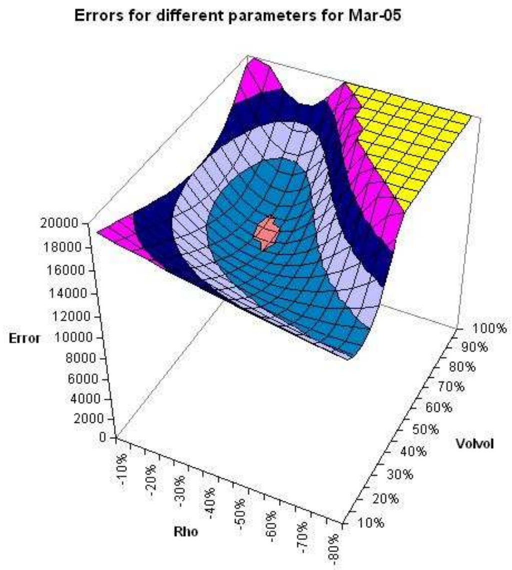

Figure 2. The error quantities for ρ and v.

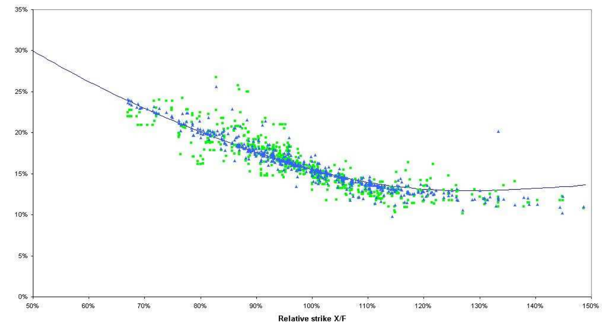

Figure 3. The SABR model for March 2005 expiry, with traded (quoted) volatilities (green), and with strategies recalibrated to a fitted skew (blue), and the fitted skew itself (solid line).

in terms of volatility, but of course this is easily converted to a double on prices.

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.

3. Hedging

As already suggested, a significant problem that arises with static (local) volatility models is that of hedging. Since different models are being used for different strikes, it is not clear that the delta and vega risks calculated at one strike are consistent with the same risks calculated at other strikes. For example, suppose that our option book is long high strike options with a total rand delta risk of 1,000,000, and is long low strike options with a total rand delta risk -1,000,000. Is our option book really delta-neutral, or do we have residual delta risk that needs to be hedged? Since different models are used at each strike, it is not clear that the risks offset each other. Consolidating vega risk raises similar concerns. Should we assume parallel or proportional shifts in volatility to calculate the total vega risk of our book? These questions are critical to effective book management, since this requires consolidating the delta and vega risks of all options on a given asset before hedging, so that only the net exposure of the book is hedged. Clearly one cannot answer these questions without a model that works for all strikes.

(Hagan et al. 2002) suggests that the parameters ρ and v are very stable (and β is assumed to be a given constant), and need to be re-fit only every few weeks. This stability may be because the SABR model reproduces the usual dynamics of smiles and skews. In contrast, the at-the-money volatility σ_{atm} will need to be updated at least daily, possibly every few hours if the market becomes especially fast-paced.

Since the SABR model is a single self-consistent model for all strikes X, the risks calculated at one strike are consistent with the risks calculated at other strikes.

Therefore the risks of all the options on the same asset can be added together, and only the residual risk needs to be hedged.

4. Risk management

The benefits of this model to the entire bank risk system are manifest

- The dealer has a sound and robust tool for deal pricing and hedging.
- Interpolation of the parameters, rather than the volatilities, is possible, and shown to be sound in (Hagan et al. 2002). In particular, deals for non-SAFEX expiries (or for SAFEX expiries which are illiquid, 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 can 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

$$\sigma^P_{\mathrm{atm}} := \frac{\sigma^P_h}{\sigma^I_h}\sigma^I_{\mathrm{atm}}$$

  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).
- Mark to market can be achieved by front, middle, back office and risk, all arriving at the same answer - no mean feat - because the parameters (which form a parsimonious set) can be calculated using fixed algorithms and stored in a common data warehouse.
- The skew can be engineered into common dealer and risk management systems.
- VaR can be calculated using the historical or Monte Carlo 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.
- Stress experiments can be performed on the other parameters.

References

Hagan, P. S., Kumar, D., Lesniewski, A. S. & Woodward, D. E. (2002), 'Managing smile risk', *WILMOTT Magazine* **September**, 84–108.
*http://www.wilmott.com/pdfs/021118_smile.pdf

West, G. (2005), 'Calibration of the SABR model in illiquid markets', *Applied Mathematical Finance* **12**(4).
*http://www.finmod.co.za/sabrilliquid.pdf

RiskWorX, P.O.Box 1075, Pinegowrie, 2123, South Africa.
www.riskworx.com