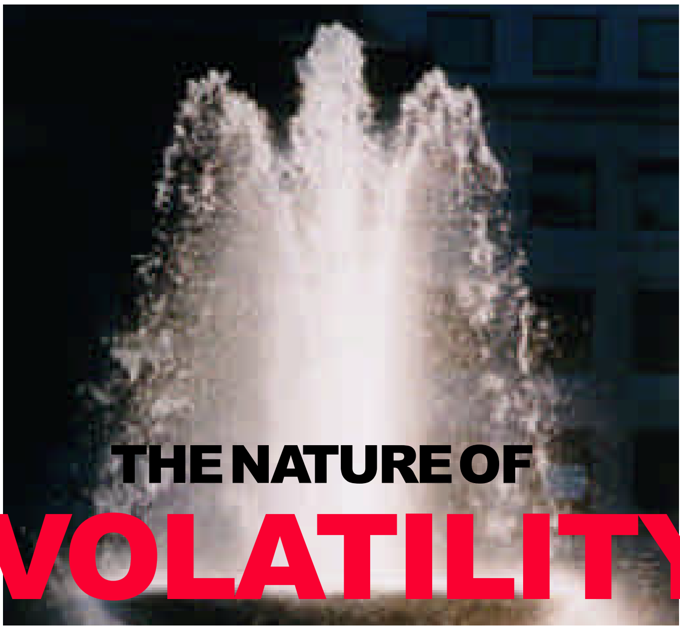

THE NATURE OF VOLATILITY

A little volatility can be a dangerous thing. When we build a new volatility model do we have to gauge its performance against realized volatility or at-the-money implied volatility? **Kirill Ilinski** looks at what affects our choice and why

Among all factors affecting an option price, volatility is notoriously difficult to account for accurately. The reason is that it is not readily measurable in the market. Apart from over the counter products (OTC), which are traded at high margin and some illiquid exchange-traded volatility futures, the volatility itself is not tradable or observable.

Due to this 'intriguing' property of being 'virtual', the volatility can be used as a fitting parameter in derivative-pricing formulae to match market prices when they are ready available. The values that give the correct market prices for vanilla calls and puts are the market-implied volatilities widely used to characterize the option market.

Now here come the questions. What future short-term volatility shall we model to price OTC derivatives, the implied one or the actual one? When we back test our volatility model shall we do it against realized or implied at-the-money vol? In other words, if we calibrate a pricing model with stochastic volatility, shall we hope to obtain volatility model parameters to be close to the statistical estimates calculated from implied or realized volatility? These are not irrelevant questions. Both mean-reversion speed and volatility of volatility are considerably higher for the implied volatility than for the realized one, while the correlation between short-term price changes and volatility changes is strongly negative for the realized vol and unstable (with boundaries +40 per cent and -60 per cent) for the implied one.

CHOOSING THE RIGHT VOL

From a pure no-arbitrage point of view, the answer is straightforward - one has to model the future-realized vol. If one believes that the implied vol is a good proxy for future realized vol then one takes the implied one and put it in the model. This is a standard argument for the use of the implied volatilities in the constant vol picture.

The same logic works for the class of volatility models dependent on the price level and time. Dupire/Derman-Kani local volatility model gives the best market estimate of future-realized vols in this class because it is the mathematical expectation of future volatility conditional on a certain price level and calculated from the risk-neutral distribution implied by the option market.[z]

However, as soon as we decide to build a new class of volatility models, possibly with vol diffusion and jumps, the only logical conclusion consistent with the no-arbitrage assumption is to gauge the model performance against the historic-realized volatility. This is what most people usually do (when they calibrate Heston model with the spot/vol correlation -80 per cent and smile happily). But does it really make sense?

Say, we want to price a cliquet, the one-year tenor call starting in one year with the strike at-the-money. The price of this option is proportional to the one-year ATM implied volatility as will be found in one year's time. If we build a model to capture the realized volatility we model a wrong quantity. As another example, let's consider any OTC derivative that we want to hedge dynamically with vanilla options (as is the case for most OTC derivatives). The 'fair' price of the derivative will be equal to the cost of the hedging strategy, which, on its own, will depend on the dynamics of the implied volatility rather than the actual one. Therefore, the choice of the implied volatility versus the realized one is essentially dictated by the nature of the product and by our choice of the hedging strategy.

NOTING THE DIFFERENCES

There are two further issues related to the use of realized volatility versus implied one in the context of derivative pricing:

• calculations of historic-realized volatilities ignore price seasonalities (therefore ignoring any kind of technical information available) explicitly implying random walk

Figure 1:VDAX Implied volatility index

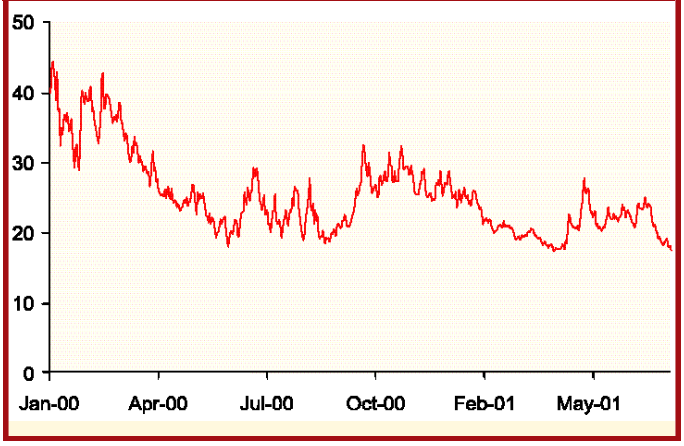

Figure 2: VIX Implied volatility index

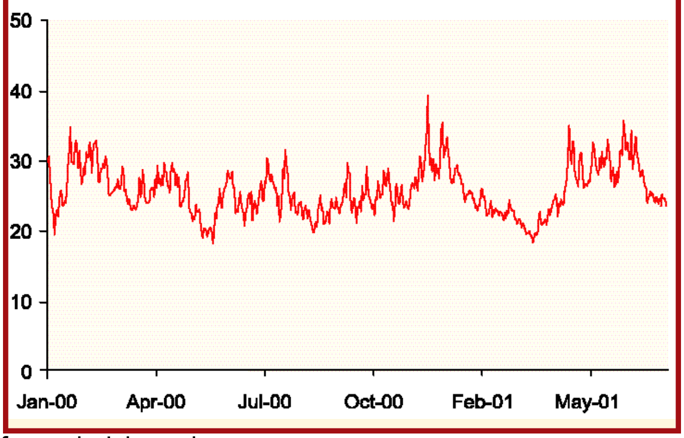

for underlying prices

• use of the actual volatility in the context of no-arbitrage is effectively equivalent to neglecting supply and demand as price forming factors for the derivatives

The implied volatility is essentially a different way to quote an option price (using an accepted market convention such as Black 76 formula). Therefore, it is not restricted by assumptions about the underlying price process or the no-arbitrage assumption. The essence of the no-arbitrage assumption is the statement that derivative prices are controlled by arbitrageurs who are strong enough to absorb speculative flows. In reality it is never the case. The reason is that by participating in the arbitrage operations, the arbitrageurs take the model risk betting on their pricing model. If they consider the premium for taking the risk insufficient they are reluctant to take the risk and, as a result, the mispricing remains in the market. The implied volatility is the only way in the Black-Scholes framework to account for the situation. Therefore, generally, to get realistic pricing models accounting for speculative pressure, this is the implied volatility that has to be used in the no-arbitrage analysis.

One can think about the volatility modeling from two different perspectives. One is to start with a (multi-) stochastic factor model and apply the no-arbitrage constraint: all exchange-traded derivatives have to be priced correctly. This gives a certain self-consistency to pricing deals in a book and is generally preferred by book-runners. When the underlying asset price is identified as the factor in the single-factor model one arrives at the Dupire expression for the local volatility. As we already mentioned above, this volatility model does not perform well in back testing and considerably underestimates future volatility skews (in contrast to stochastic volatility models), but formally gives a complete market model and allows fast and simple implementations. In the case of the multi-factor model, the no-arbitrage constraints are so restrictive and technically untreatable that they almost exclude direct practical applications of the model. The alternative is to come up with an empirically realistic model that, however, will produce prices diverging from the ones of exchange-traded options (hence allowing a formal arbitrage). The latter class of models are known as the arbitrage models. As we have already pointed out, this sort of (formal) arbitrage should not surprise anyone if there are not enough market participants who wish to take up the model risk: a small number of arbitrageurs will cause the mispricing to stay and, therefore, undermining the equilibrium model, increasing the model risk and supporting further their reluctance to step in.

As soon as we are ready to accept that the implied volatility, at least partially, reflects deviations of real markets from the equilibrium market assumptions, the next step would be to account in the model for the efficient and inefficient factors separately. Shiller's volatility test can give us a lead to follow.

The rational behind volatility tests is that in the efficient market the volatility σ of a stock price has to be related to the volatility of some fundamental information such as expectations of a firm's future dividends predictions. Starting with the discount present value of future dividends $D_{t'}$ the theoretical price S_t^{theor} can be written as

$$S_t^{\text{theor}} = \sum_{i=1}^{n-1} \delta^i E_t(D_{t+i}) + \delta^n E_t(S_{t_n}) \ .$$

Having in hand the actual dividends from the past and the current price (to use as S_{t_n}) Shiller found that the actual price in the past was more volatile than its theoretical counterpart. Although there are a number of econometric issues related to the variance tests (such as stationarity assumption for the time series and small sample biases), the variance tests show that the stock prices are generally more volatile than their fundamentals and the excessive volatility is related to the pressure generated by speculators.

USING THE FIGURES

Let's now have a look at Figures 1 and 2, which show the VIX and VDAX volatility indices calculated from ATM short-term volatilities on S&P 100 and DAX exchange-traded options. The reader can see that the volatility indices have a clear oscillatory pattern with a typical period of a few weeks. Occasional sharp peaks are followed by fast decay and a next wave. This sort of behavior clearly contradicts our intuitive picture of the underlying slow economic dynamics and would be rather attributed to the demand and supply factors, the speculative flows in the option market.

Figures 1 and 2 allow us to make another interesting observation: for both indices there exist volatility floors that are touched occasionally (usually in summer months or low market activity). One can associate the floor in the volatility with the fundamental low boundary in uncertainty that is dictated by the economic activity (in contrast to the speculative trading impact). The floor changes slowly with time and can be taken as a constant on the time-scale of several years. In this picture the residual volatility fluctuations are a pure result of the speculative pressure and can be modeled statistically.

Formalizing the above argument, one can propose, for example, the following short-term stochastic volatility model:

• there is a low boundary (which we call the floor and denote as θ) for the short-term implied volatility Σ which is related to the underlying economic activity, varies slowly and can be taken as constant on the time-scale of several years.

• the varying part of the implied volatility - $\sigma = \Sigma - \theta$ - is positive and mean-reverting. The natural choice for the stochastic model for VIX or VDAX would be

$$d\sigma = -\lambda(\sigma_0 - \sigma)dt + \sigma l dZ \ . \qquad (1)$$

where λ is the mean-reverting parameter, σ_0 is the mean-reverting level, which reflects an average impact of the speculative fluctuations in volatility (Shiller's difference between actual and theoretical volatilities) and l is the volatility of volatility (often referred as Vvol).

The volatility enters the stochastic differential equation for the underlying price S

$$\frac{dS}{S} = r_0 dt + \Sigma dW \ , \ \ \Sigma = \sigma + \theta, \qquad (2)$$

with which it is assumed to be correlated (or rather anti-correlated in the case of equity markets)

$$\langle dW, dZ \rangle = \rho dt. \qquad (3)$$

Formulas 1,2 and 3 constitute the complete set of equations governing the short-term price/volatility dynamics. It is interesting to note that in the case of $\sigma_0 = 0$ it is possible to find exact solutions for European option prices and European volatility derivatives - Kalinin/Ilinski, Equity Derivatives Research working paper, Chase Manhattan International, (2000).

One can try other short-term models for the speculative volatility σ such as a generalization of the Heston's model or technical/behavioral models. The important point is to model the underlying economic floor and the speculative volatility separately.

FOOTNOTES

1. The fact that the local vol model strongly underestimates the future skew (i.e one-year implied vol curve for options starting in one-year's time will be steeper than predicted by the local model) says that the price level is not the most important factor.

In fact, even for a short time horizon (from a few days to a few weeks) local volatility is not particularly accurate and performs particularly badly during market shake-ups.