Calibrating and pricing with embedded local volatility models

Consistently fitting vanilla option surfaces when pricing volatility derivatives such as Vix options or interest rate/equity hybrids is an important issue. Here, ***Yong Ren***, ***Dilip Madan*** *and* ***Michael Qian Qian*** *show how this can be accomplished, using a stochastic local volatility model as the main example. They also give, for the first time, quanto corrections in local volatility models*

Local volatility models introduced by Dupire (1994) and Derman & Kani (1994) are now widely used to price and manage the risks of structured products. The dimensionality of risks to be simultaneously managed continues to expand with the demand for hybrid products and the growth of markets directly trading volatility. The formulation and implementation of local volatility models in these higher-dimensional Markov contexts is now becoming an important issue. Of particular interest to the financial industry are the accommodation of stochastic volatility, stochastic interest rates, and the pricing of options on foreign stocks, quantos and baskets, in the presence of volatility skew. The general recipe for models based on a Brownian filtration was provided by Gyöngy (1986), who showed how to associate with a general Itô process a Markov process with the same marginal distributions.

We illustrate the required computations with particular emphasis on the presence of stochastic volatility as the additional dimension. Stochastic volatility in a local volatility context permits the exact calibration of vanilla options while at the same time addressing the exposure of financial contracts to the rate of mean reversion in volatility and the volatility of volatility. We implement the algorithm developed for options on realised variance and options on the Vix index.

Gyöngy and the matching of one-dimensional marginals

First, we briefly summarise the important result in Gyöngy (1986). Consider a general n dimensional Itô process of the form:

$$d\xi(t) = \beta(t,\omega)dt + \delta(t,\omega)dW(t)$$

This process gives rise to marginal distributions of the random variables $\xi(t)$ for each t. Gyöngy then shows that there is a Markov process $x(t)$ with the same marginal distributions. The explicit construction is given by:

$$dx(t) = b(t,x(t))dt + \sigma(t,x(t))dW(t)$$

where:

$$\sigma^2(t,x) = E\left[\delta(t,\omega)\delta^T(t,\omega)\middle|\xi(t)=x\right]$$

$$b(t,x) = E\left[\beta(t,\omega)\middle|\xi(t)=x\right]$$

In the rest of this article, we will repeatedly find this result useful in identifying the local volatility function σ and the local drift function b of the one-dimensional process with the same marginal distributions as the true high-dimensional dynamics.

The stochastic local volatility model

Here, we consider an extension of the Dupire (1994) local volatility model that incorporates an independent stochastic component to volatility. We develop dynamics under a risk-neutral measure where the stock price process ($S(t)$, $t \geq 0$) has a drift equal to the risk-free interest rate r less the dividend yield q and in our case an average volatility given by a deterministic function of the stock price and calendar time, $\sigma(S, t)$. The independent stochastic component of volatility is modelled by a stochastic process ($Z(t)$, $t \geq 0$) that starts by assumption at $Z(0) = 1$. Hence the evolution of the stock price is given by:

$$dS = (r-q)Sdt + \sigma(S,t)Z(t)SdW_S(t) \qquad (1)$$

where ($W_S(t)$, $t \geq 0$) is a standard Brownian motion.

We illustrate using a mean-reverting lognormal model for $Z(t)$, the stochastic component of volatility. We suppose that:

$$d\ln Z = \kappa(\theta(t) - \ln Z)dt + \lambda dW_Z(t) \qquad (2)$$

where κ is the rate of mean reversion and λ is the volatility of volatility. There is a long-term deterministic drift given by ($\theta(t)$, $t \geq 0$). With a view to interpreting $\sigma^2(S, t)$ as the average local variance, we force the unconditional expectation of $Z(t)^2$ to be unity by requiring that:

$$\theta(t) = -\frac{\lambda^2}{2\kappa}\left(1 + e^{-2\kappa t}\right)$$

The dependence of volatility on the stock price is already captured in the leverage function $\sigma(S, t)$, so we assume that the Brownian motion driving the stochastic component of volatility W_Z is uncorrelated with the Brownian motion driving the stock price W_S. We could extend the model to a non-zero correlation, which ideally we can observe from the market, then some of the skew will come from the correlation and some from the leverage function (note that if σ has no dependence on S, the model is very Heston-like). Here we will only describe the simplest model to implement and calibrate. Other approaches at making local volatility stochastic include Derman & Kani (1998) and Dupire (2004).

The fact that volatilities are stochastic and capable of rising without a movement in spot prices is now widely recognised and practical considerations of risk management require an assessment of the magnitude of this exposure for a variety of structured products. This necessitates the introduction of an independent stochastic volatility unrelated to the stochasticity introduced by leverage functions and permits the assessment of exposure to the volatility of volatility that is now independently parameterised here by the coefficient λ. Other stochastic volatility models based on the Heston model and its generalisations also permit such an assessment but do not capture the surface of implied volatilities as precisely as a local volatility model. This is why there is interest in introducing stochastic volatility into a local volatility context.

One-dimensional Markov process for the stock price marginals

From Gyöngy (1986), we see that the one-dimensional Markov process with the same marginal distributions as those of our model is given by:

$$dS = (r-q)Sdt + \sigma_{LV}(S,t)Sd\tilde{W}_S(t) \tag{3}$$

where we must have that:

$$\sigma_{LV}^2(K,T) = \sigma^2(K,T)E\left[Z(T)^2 \middle| S(T)=K\right]$$

For the one-dimensional Markov process (3), the Dupire (1994) and Derman & Kani (1994) relationship holds between call option prices of strike K and maturity T, $C(K, T)$, and the one-dimensional leverage function $\sigma_{LV}(S, T)$. In particular, we have:

$$\sigma_{LV}^2(K,T) = 2\frac{C_T + (r-q)C_K + qC}{K^2 C_{KK}}$$

Hence, we may now recover our leverage function from option prices provided we have synthesised the function $\psi(K, T) = E[Z(T)^2|S(T) = K]$ by:

$$\sigma^2(K,T) = \frac{\sigma_{LV}^2(K,T)}{\psi(K,T)} \tag{4}$$

We now address the issue of the simultaneous solution of $\psi(K, T)$ and $\sigma^2(K, T)$ satisfying equation (4). For this, we introduce the forward joint transition density $p(x, y, T)$ for the logarithm of the stock price X, to reach the level x and the logarithm of the stochastic component of volatility $Y = \ln(Z)$ to reach the level y at time t, starting from $S = S(0)$ and $Z = 1$ respectively at time zero. This density element satisfies the Kolmogorov forward equation:

$$\begin{aligned} &-\frac{\partial p}{\partial t} - \frac{\partial}{\partial x}\left[\left(r-q-\frac{e^{2y}\sigma^2(e^x,t)}{2}\right)p\right] \\ &+\frac{\partial^2}{\partial x^2}\left[\frac{e^{2y}\sigma^2(e^x,t)}{2}p\right] - \frac{\partial}{\partial y}\left[\kappa(\theta(t)-y)p\right] \\ &+\frac{\partial^2}{\partial y^2}\left(\frac{\lambda^2}{2}p\right) = 0 \end{aligned} \tag{5}$$

subject to the boundary condition that at time zero we start at $X = \ln(S(0))$ and $Y = \ln Z(0) = 0$ with probability one. We then recover the function ψ as:

$$\psi(K,T) = \frac{\int_0^\infty e^{2y} p(e^K, y, T)dy}{\int_0^\infty p(e^K, y, T)dy} = \sum_{j=1}^{n_Y} p_j e^{2y_j} / \sum_{j=1}^{n_Y} p_j \tag{6}$$

The strategy is to solve the forward equation in the density element (5) forward one step at a time, recognising that at the start we have $\sigma(S(0), 0)$ from the local volatility equation (4) since $\psi(S(0), 0) = 1$. We then use (6) at the next time step to determine the function ψ at this point and that allows us to infer from (4) the function $\sigma(K, T)$ at this time step. We then proceed in this fashion through time to recover simultaneously both the functions $\psi(K, T)$ and $\sigma^2(K, T)$. Along the way we have also solved for all the density elements $p(x, y, t)$ for all time points t:

$$\begin{aligned} &\psi(K,0)=1, \quad \sigma(S,0)=\sigma(S(0),0) \\ &\downarrow \\ &p(x,y,\Delta t) \\ &\downarrow \\ &\sigma(S,\Delta t) \\ &\downarrow \\ &p(x,y,2\Delta t) \\ &\downarrow \\ &\sigma(S,2\Delta t) \\ &\downarrow \\ &p(x,y,3\Delta t) \\ &\downarrow \\ &\sigma(S,3\Delta t) \\ &\ldots \end{aligned}$$

Our particular example uses alternating direction implicit methods to solve for the probability element. With p_{ij}^t denoting $p(x_i, y_j, t)$ for points (x_i, y_j, t) on the space time grid, and using operator splitting with a fully implicit method in both directions, we get:

$$\begin{aligned} &p_{ij}^{1/2} + \frac{\Delta t}{2\Delta X}\left\{\left[r-q-\frac{e^{2Y_j}\sigma^2(e^{X_{i+1}},t)}{2}\right]p_{i+1,j}^{1/2}\right. \\ &\qquad \left.-\left[r-q-\frac{e^{2Y_j}\sigma^2(e^{X_{i-1}},t)}{2}\right]p_{i-1,j}^{1/2}\right\} \\ &-\frac{\Delta t}{2\Delta X^2}\left[e^{2Y_j}\sigma(e^{X_{i+1}},t)p_{i+1,j}^{1/2} - 2e^{2Y_j}\right. \\ &\qquad \left.\sigma(e^{X_i},t)p_{ij}^{1/2} + e^{2Y_j}\sigma(e^{X_{i-1}},t)p_{i-1,j}^{1/2}\right] = p_{ij}^{t-\Delta t} \\ &p_{ij}^t + \frac{\Delta t}{2\Delta Y}\left\{\kappa\left[\theta(t)-Y_{i,j+1}\right]p_{i,j+1}^t - \kappa\left[\theta(t)-Y_{i,j-1}\right]p_{i,j-1}^t\right\} \\ &-\frac{\lambda^2\Delta t}{2\Delta Y^2}\left(p_{i,j+1}^t - 2p_{ij}^t + p_{i,j-1}^t\right) = p_{ij}^{1/2} \end{aligned}$$

One first solves for $p_{ij}^{1/2}$ from $p_{ij}^{t-\Delta t}$ using a tridiagonal solver, then p_{ij}^t can be solved from $p_{ij}^{1/2}$.

The advantage of such an approach is that the probabilities solved for in this way are very smooth and stable (just like the vanilla option values one calculates by partial differential equation (PDE) methods). For a similar but less rigorous numerical approach in the context of foreign exchange markets and a square-root volatility process solved on trinomial trees, we refer the reader to Jex, Henderson & Wang (1999).

Once $\sigma(K, T)$ has been calculated, they can be used in either Monte Carlo or PDEs to price securities that are sensitive to stochastic volatility. It is natural to price options on realised variance using the Monte Carlo approach where the realised variance can be summed up along simulated paths. Because the option on the

1 $\zeta(K, T)$ calibrated to vanilla option surface

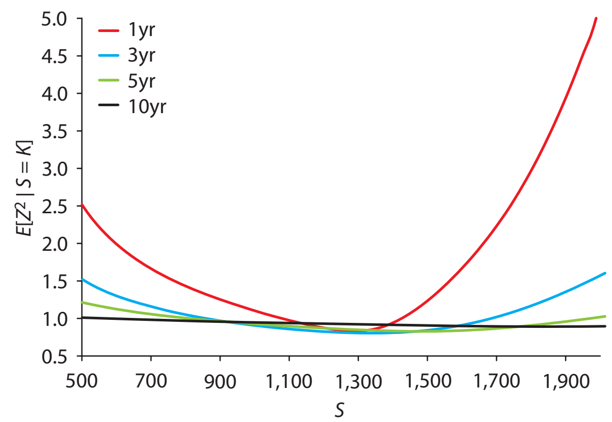

2 Calibration of SLV to SPX market implied volatilities

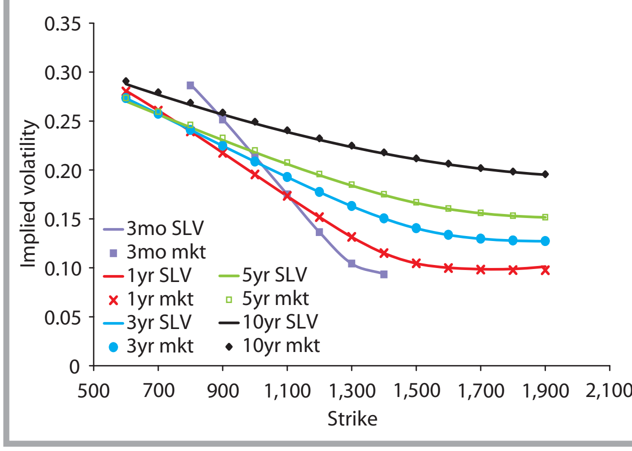

Vix index is written on the sum of expected future variances, the pricing can be easily implemented on the PDE grid. One may evaluate the value of the forward Vix level at time t on the grid using transition probabilities. This is given by:

$$Vix_t = \sqrt{E_t\left[\frac{1}{h}\int_t^{t+h} Z_u^2\sigma^2\left(S(u),u\right)du\right]}$$

To do so, we notice that we have already calculated $\sigma(K, T)$, and Z is simply e^Y, one of the dimensions of our PDE grid. We start at time $t + h$, and put the value of $Z^2\sigma^2(S, t)$ on the grid. We then use a PDE solver to propagate backwards on a space time grid, a conditional expectation at time $t + h$ to the grid points at time $t + h - \Delta t$. As we propagate backwards on the grid, the value on the grid becomes the expected future value of the variance at $t + h$ (care must be taken that the value is not automatically discounted). We add to this quantity the value of $Z^2\sigma^2(S, t)$ at the current grid point, and propagate this total value back to $t + h - 2\Delta t$, and so on. When we reach time t, we have on the grid the value for the expected forward variance, $E_t[\int_t^{t+h} Z_u^2\sigma^2(S(u), u)du]$. Or expressed as a formula, the 30-day expected variance is calculated as:

$$Var_{30}(t) = E_t\left[\frac{1}{h}\int_t^{t+h} Z_u^2\sigma^2(u)du\right]$$

$$= E_t\left[\sum_{i=1}^{30}\frac{1}{30}Z_{t+i}^2\sigma_{t+i}^2\right]$$

$$= \frac{1}{30}E_t\left[Z_{t+1}^2\sigma_{t+1}^2 + E_{t+1}\left[\sum_{i=2}^{30}Z_{t+i}^2\sigma_{t+i}^2\right]\right]$$

$$= \frac{1}{30}E_t\left[Z_{t+1}^2\sigma_{t+1}^2 + E_{t+1}\left[Z_{t+2}^2\sigma_{t+2}^2 + \cdots + E_{t+29}\left[Z_{t+30}^2\sigma_{t+30}^2\right]\right]\right]$$

All that remains is to define the payout for a call option on Vix_t of strike k as $(Vix_t - k)^+$ and propagate with discounting to the valuation time. Since we have calibrated our model to all plain vanilla options, Vix contracts that can be regarded as the underlying for options on the Vix index are priced correctly as well, according to the argument in Carr & Madan (1998). Thus our model is able to price volatility derivatives such as options on the Vix index consistently with vanilla options and volatility contracts (variance swaps, Vix futures). This is done in a context that incorporates leverage effects on volatility along with independent shocks to volatility.

Stochastic local volatility results

We consider a lognormal stochastic volatility with mean reversion $\kappa = 0.5$ and a volatility of volatility of $\lambda = 0.5$. We recognise that calibration of the vanilla option surface alone typically poorly identifies the rate of mean reversion and the volatility of volatility in previous stochastic volatility models. In our model, the local volatility function may be used to calibrate the entire surface of implied volatilities, at any choice for the mean-reversion rate and the volatility of volatility. However, the mean-reversion rate and the volatility of volatility do jointly affect the term structure of volatility options, which we discuss at the end of this section. Here, we focus attention on the sensitivity of volatility products, such as options on realised variance and options on the Vix index, to these parameters.

We take 20 vanilla options at each of 13 maturities from two weeks to 10 years, and solve simultaneously for the leverage function $\sigma^2(K, T)$ and:

$$\zeta(K,T) = E\left[Z(T)^2 \middle| S(T) = K\right]$$

The vanilla options used are on the S&P 500 index as at October 3, 2005.

Since the leverage function is just the ratio of the usual local volatility function divided by ζ, we present a graph in figure 1 of the function ζ as a function of K at the one-, three-, five- and 10-year point. We observe that these expectations are convex in K near the at-the-money point, where they dip below unity and they rise above unity in both tails. Additionally, in figure 2 we present the market-implied volatilities and the calibrated implied volatilities from the stochastic local volatility model.

The calibration speed on a present-day Linux box was about 20 seconds for maturities up to three years. For maturities up to 10 years, the corresponding speed was 60 seconds.

The use of stochastic local volatility models arises in assessing the nature of product exposure to the presence of an independent stochastic volatility. The model formulated here introduces two additional parameters connected with stochastic volatility – the rate of mean reversion in volatility and the volatility of volatility. As an example, we present the exposure of at-the-money one-year straddles on realised variance to these components of stochastic volatility. For the calibration date of December 12, 2006, we show the resulting valuations in table A.

We first observe from table A that the stochastic volatility model reduces to local volatility at a zero volatility of volatility.

We further observe that the at-the-money one-year realised variance straddle is convex in the volatility of volatility and one may therefore be exposed to an undervaluation if in fact the volatility of volatility is itself stochastic. With respect to the rate of mean reversion, we have concavity and an overvaluation if there is stochasticity in the rate of mean reversion.

In table B, we also present a sample of six-month at-the-money straddles on Vix_t for a range of κ and λ values and the calibration date of December 12, 2006.

We observe that options on the Vix rise in value with an increase in volatility and are convex in the volatility. They fall with an increase in the rate of mean reversion and are concave in this direction. Stochastic volatility with a zero volatility of volatility again reduces to the local volatility model.

As we previously noted, as realised variance options are written on an average of past squared returns, they are calculated using Monte Carlo from the calibrated model. The options on the Vix index are based on averages of future squared returns and these may be calculated as explained easily by backward propagation on the PDE grid.

While we have given here the qualitative effect of our model parameters on a selected set of volatility derivatives, we have not attempted a thorough study of simultaneously calibrating for the examples of the Vix surface and the equity option surface. At this stage, we have a fairly advanced equity option model embedded into an elementary, almost Black-Scholes-type volatility model. We anticipate that the joint calibration of both surfaces will eventually involve multi-factor models for the volatility and may go as far as incorporating a local volatility-of-volatility surface. We leave these matters for future research.

However, it is instructive to observe that in our model, as:

$$Z(T) = \exp\left[-\frac{\lambda^2}{2\kappa}\left(1 - e^{-2\kappa T}\right) + \int_0^T e^{-\kappa(T-t)} \lambda dW \right]$$

and the realised variance till maturity T is given by:

$$\frac{1}{T}\int_0^T Z^2(t)\sigma^2(S,t)dt$$

the variance of realised variance quickly decays in the presence of a large mean-reversion rate. In such cases, the innovations become independent and the variance of the realised variance will be proportional to $1/T$ as the maturity of the option gets longer. On the other hand, for a negligible mean-reversion rate, one is averaging sums and this does not diminish as fast. These are simple observations on the term structure of the variance of realised variance. Models with more factors may exhibit richer term dynamics. The term structure of the variance of realised variance is embedded in the prices of options on realised variance and these may be used to calibrate the dynamics of volatility.

Other extensions of local volatility

Apart from stochastic volatility extensions of local volatility models, there are other extensions well understood in the Black-Merton-Scholes context of a constant volatility that are somewhat more involved when we come to a local volatility formulation. These include the pricing of foreign stock options in the domestic currency that is inclusive of both the stock and exchange rate risk. The pricing of quantos that shaves out the exchange rate risk is also well known for constant volatility, and we show here the precise adjustments needed in the presence of local volatility models

A. Valuations of one-year at-the-money straddles on realised variance

κ	Local volatility	Volatility of volatility λ		
		0%	50%	100%
0.5	2.52	2.52	2.99	3.94
1.0	2.52	2.52	3.16	4.30
2.0	2.52	2.52	3.23	4.45
4.0	2.52	2.52	3.13	4.30

B. Valuation of half-year at-the-money straddles on the Vix

κ	Volatility of volatility λ		
	0%	50%	100%
0.5	3.67	4.81	6.70
1.0	3.67	4.58	6.24
2.0	3.67	4.26	5.53
4.0	3.67	3.96	4.71

for both the stock and the exchange rate. Next we consider options on baskets and, finally, the case of stochastic interest rates. These topics are taken up in separate subsections.

■ **Foreign stock.** Consider a two-dimensional Markov process for the foreign price of stock and the exchange rate, where both models are of the local volatility form:

$$dS = \left(r_f - q\right)Sdt + \sigma_s\left(S(t),t\right)S(t)dW_S(t)$$
$$dX = \left(r - r_f\right)Xdt + \sigma_x\left(X(t),t\right)X(t)dW_X(t)$$
$$dW_S dW_X = \rho\left(S(t),X(t),t\right)dt$$

where r_f is the foreign interest rate, r is the domestic interest rate and q is dividend yield, while σ_s, σ_x are the two local volatility functions, and $(W_S(t), W_X(t), t \geq 0)$ are standard Brownian motions with instantaneous correlation ρ.

Let $Y(t)$ be the domestic price of foreign stock or:

$$Y(t) = S(t)X(t)$$

For the domestic currency of the dollar, this is a dollar-denominated asset, so it has a risk-neutral evolution on this filtration given by the martingale representation theorem, and we may also write $Y(t)$ as an Itô process with the representation:

$$dY = (r-q)Ydt + \sigma_S YdW_S(t) + \sigma_X YdW_X(t)$$

By Gyöngy's result, there is a one-dimensional Markov process with the same one-dimensional marginals and the evolution:

$$dY = (r-q)Ydt + \sigma\left(Y(t),t\right)YdW_Y(t)$$

where:

$$\begin{aligned}\sigma_y^2(y,t) = E\Big[&\left(\sigma_s^2(S,t) + \sigma_x^2(X,t)\right.\\ &\left.+2\rho\left(S(t),X(t),t\right)\sigma_s(S,t)\sigma_x(X,t)\right)\Big|Y(t)=y\Big]\end{aligned} \qquad (7)$$

When the volatilities are strike-independent, we recover the well-known Black-Scholes result. Similar relation also holds for the cross exchange rate and the two corresponding currency exchange rates. The expression (7) is particularly useful in a joint stock and exchange rate local volatility context because it relates the skew of the stock price process and the exchange rate process to that of the foreign stock process. In the absence of such an expression, we are left with the need to price a call option on $Y(t)$ and use the Dupire equation to infer the required local volatility function. Equation (7) provides an alternative way to directly infer the local volatility from the terminal joint density, which can be available as copula functions or calculated by other means.

Quantoed stock. We now note that the quantity $S(t)$ on which we may write quantoed options is given by:

$$S(t)=\frac{Y(t)}{X(t)}$$

We have the risk-neutral law with respect to the dollar numeraire of $X(t)$ and $Y(t)$. Specifically, we write:

$$dY=(r-q)Ydt+\sigma_s(S(t),t)YdW_S(t)+\sigma_x(X(t),t)YdW_X(t)$$
$$dX=(r-r_f)Xdt+\sigma_x(X(t),t)X(t)dW_X(t)$$

We then write the law for S from an Itô analysis of the ratio as:

$$dS=\sigma_s(S(t),t)SdW_S(t)$$
$$+\Big[r-\Big(q+(r-r_f)+\sigma_s(S(t),t)\sigma_x(X(t),t)\rho(S(t),X(t),t)\Big)\Big]dt$$

Applying Gyöngy, we obtain a one-dimensional Markov process with the dynamics:

$$\begin{aligned}dS=\Big[r-\Big(q+(r-r_f)+\sigma_s(S(t),t)\\ E\big[\sigma_x(X(t),t)\rho(S(t),X(t),t)\big|S(t)\big]\Big)\Big]dt\\ +\sigma_s(S(t),t)SdW_S(t)\end{aligned} \tag{8}$$

Here, we have an example of how the local volatility in the exchange rate transfers into a localised drift in the one-dimensional law for the quantoed stock. Again, we see if σ_x and ρ have no dependence on X or S, this reverts to the well-known quanto adjustment. This result is non-trivial as there is no more direct way of calculating the adjustment.

Options on baskets. We now consider the case of a weighted basket of stocks S_i, each of which has a local volatility model. The Brownian motions are correlated with instantaneous correlation $\rho(S_i, S_j, t)$. The stock dynamics are:

$$dS_i=(r-q_i)S_idt+\sigma_i(S_i,t)S_idW_i$$

Now define the basket by the sum with $S(t)=\Sigma_i w_iS_i(t)$ and we may develop the expression:

$$dS=\left(rS-\sum_i w_iq_iS_i\right)dt+\sum_i w_i\sigma_i(S_i,t)S_idW_i$$

An application of Gyöngy's result yields that:

$$dS=(r-q(S,t))Sdt+\sigma(S(t),t)SdW$$

where:

$$Kq(K,t)=E\left[\sum_i w_iq_iS_i(t)\Big|S(t)=K\right] \tag{9}$$

$$K^2\sigma^2(K,t)=E\left[\begin{matrix}\sum_i w_i^2\sigma_i^2(S_i,t)S_i^2(t)+2\sum_{i,j}w_iw_j\\ \rho(S_i,S_j,t)S_i(t)S_j(t)\sigma_i(S_i,t)\sigma_j(S_j,t)\end{matrix}\Bigg|S(t)=K\right] \tag{10}$$

Again, the expression (10) relates the skews in the individual assets to the basket skew. This is an important consideration in pricing options on baskets. We also notice that by setting the weights to be negative, the result also applies to outperformance or Margrabe options. For the implementation of an approximation to equation (10), the reader is referred to Avellaneda *et al* (2002).

The results contained in this section on the pricing of foreign stock, quanto and basket options are easily derived under the Black-Scholes assumption, but are not so obvious in a local volatility setting. They not only elucidate the role of skew in the dynamics, but are also of practical interest. First, approximations can be made in these formulas for practical use; second, expectations in these formulas can be evaluated using either known densities, or by Monte Carlo and PDE methods. Once the drift and local volatilities are known, the dimension of the dynamics of the asset is effectively reduced to one. While the dimensionality reduction is less necessary in the pricing of simple options on these assets, it is crucial in more complex structures where, for example, PDE methods become impractical with the higher dimensionality of the problem. Dimensional reduction will also help in many risk management applications where speed is crucial. Not surprisingly, these are the situations where the results derived in Black-Scholes settings are often used (and abused).

Stochastic interest rates. Our last example involves a stochastic interest rate economy with local volatility stock dynamics. We notice attempts along a similar line have been made by Atlan (2006). Our approach follows from the same calibration method we used for the stochastic local volatility case:

$$\frac{dS}{S}=(r(t)-q)dt+\sigma(S(t),t)dW_S(t) \tag{11}$$

Applying Gyöngy directly to equation (11), we observe that the Markov dynamics embedded in the marginal stock return distributions are given by:

$$\frac{dS}{S}=\big(E\big[r(t)\big|S(t)=S\big]-q\big)dt+\sigma(S(t),t)dW_S(t)$$

and hence the effects of stochastic interest rates on equity options are fully captured on determining the conditional expectation of the spot rate given that the stock prices reach the level S at time t. For calibration of the local volatility function, we analyse the call prices.

In this case, we may write the option prices as:

$$C(K,T)=E\left[\exp\left(-\int_0^T r(u)du\right)(S(T)-K)^+\right]$$

Elementary manipulations show, on applying the Meyer-Tanaka formula to the discounted payout and simplifying, that:

$$\sigma^2(K,T)=$$
$$2\frac{C_T+qC-qKC_K+KE\left[\exp\left(-\int_0^T r(v)dv\right)r(T)\mathbf{1}_{S(T)>K}\right]}{K^2C_{KK}}$$

We therefore need, in addition to option prices, the price of the contract that pays $r(T)\mathbf{1}_{S(T)>K}$. Using the T forward measure, we may write this price as:

$$w(T)=P(0,T)\tilde{E}_T\left[r(T)\mathbf{1}_{S(T)>K}\right]$$

Hence, we need to develop the joint law of $r(T)$, $S(T)$ under the T forward measure.

By way of an example, we work using a Heath-Jarrow-Morton model for the fixed-income dynamics with the pure discount bond prices $P(t, T)$ with dynamics given by:

$$\frac{dP(t,T)}{P(t,T)} = r(t)dt + \sum_\alpha \left(\int_t^T \sigma_\alpha(t,u)du\right)dW_\alpha(t)$$

for a set of correlated Brownian motions with correlations $\rho_{\alpha\beta}$.

We now shift to the forward measure denoted $\widetilde{E}_T(\cdot)$, with forward date T where T is the option maturity. The dynamics of the stock price under this measure with correlation $dW_S dz_\alpha = \rho_a$ are given by:

$$\begin{aligned}\frac{dS}{S} &= \left(r(t) - q + \sum_\alpha \rho_\alpha \left(\int_t^T \sigma_\alpha(t,u)du\right)\sigma(S(t),t)\right)dt \\ &\quad + \sigma(S(t),t)dZ_S\end{aligned} \tag{12}$$

where Z_S is a standard Brownian motion under the forward measure.

The dynamics of the spot rate under the T forward measure may be calculated to be:

$$\begin{aligned} r(t) = f(0,t) - \int_0^t \sum_{\alpha\beta} \sigma_\alpha(s,t)\left(\int_t^T \sigma_\beta(s,w)dw\right) \\ \rho_{\alpha\beta}ds - \sum_\alpha \int_0^t \sigma_\alpha(s,t)dZ_a(s)\end{aligned} \tag{13}$$

where Z_α are the new forward motion Brownian motions.

We may now apply Gyöngy to the joint equations (13) and the solution to (12) to construct the two-dimensional Markov process in $(r(t), S(t))$ with the same marginals as these Itô processes:

$$dr = \eta_r(t,r,S)dt + \sum_\alpha \beta_{r\alpha}(t,r,S)dZ_\alpha$$

$$dS = \eta_S(t,r,S)dt + \sigma(S(t),t)dZ_S$$

$$\eta_r = \frac{\partial}{\partial t} f(0,t)$$

$$-E\left[\begin{array}{l} \sum_{\alpha\beta}\sigma_\alpha(t,t)\left(\int_t^T \sigma_\beta(t,w)dw\right)\rho_{\alpha\beta} \\ -\int_0^t \sum_{\alpha\beta}\sigma_\alpha(s,t)\sigma_\beta(s,t)dw\rho_{\alpha\beta}ds \qquad \Big| r(t)=r, S(t)=S \\ +\int_0^t \sum_{\alpha\beta}\partial_t\sigma_\alpha(s,t)\left(\int_t^T \sigma_\beta(s,w)dw\right)\rho_{\alpha\beta}ds \end{array}\right]$$

$$\beta_{r\alpha}^2 = E\left[\sigma_\alpha^2(t,t) \middle| r(t)=r, S(t)=S\right]$$

$$\eta_S = S\left(r - q + E\left[\sum_\alpha \rho_\alpha\left(\int_t^T \sigma_\alpha(t,u)du\right)\middle| r(t)=r, S(t)=S\right]\sigma(S,t)\right)$$

To calibrate the local volatility in the presence of stochastic interest rates, the Heath-Jarrow-Morton interest volatilities are first calibrated to interest rate markets. We are then left with the stock price and spot rate dimensions with the dynamics of the spot rate fully determined. This is very similar to our stochastic local volatility case. We then calculate, using the Kolmogorov forward equation, the density distribution $\rho(S, r, t)$, which is then used in calculating the conditional drift of the stock price $r(T)\mathbf{1}_{S(T) > K}$ that goes into the local volatility calculation. We repeat this procedure going forward in time until all values of local volatility $\sigma(S, T)$ are known. Then the original hybrid model can be implemented using either Monte Carlo or a PDE.

Conclusion

We exploit Gyöngy's (1986) result to represent the marginal laws of Itô processes by Markov processes in a local volatility context. This representation leads to generalisations of the Dupire (1994) and Derman & Kani (1994) equations for determination of local volatility or leverage functions jointly with certain conditional expectations of the Markov process. These may be solved by PDE methods that extract the joint density using a solution of the Green's function.

The method is applied in particular to a stochastic volatility model in a local volatility context that permits an exact calibration of vanilla options while at the same time addressing questions on contract exposure to the volatility of volatility. Calibration and repricing speeds are observed to be around 20 seconds for maturities extending up to three years. For maturities up to 10 years, the speed was 60 seconds. These are reasonable speeds, making the method both tractable and useful for the investigation of high-dimensional extensions to the base local volatility model. We illustrate the required calculations for options on realised variance and options on the Vix index. We further give the results for pricing options on foreign stocks, quantos and baskets in a local volatility setting. Finally, we address the problem of calibrating hybrid models with stochastic interest rates. ■

Yong Ren was executive director in the equity analytical modelling group of Morgan Stanley in New York and is currently a director at Saba Principal Strategies, Deutsche Bank in New York. Dilip Madan is professor of finance at the Smith School of Business and serves as a consultant to Morgan Stanley. Michael Qian Qian is executive director in the equity analytical modelling group of Morgan Stanley in New York. Email: yren@alumni.princeton.edu, dbm@rhsmith.umd.edu, michael.qian@morganstanley.com

References

Atlan M, 2006
Localizing volatilities
Working paper, Laboratoire de Probabilités, Université Pierre et Marie Curie

Avellaneda M, D Boyer-Olson, J Busca and P Friz, 2002
Reconstructing volatility
Risk October, pages 87–91

Carr P and D Madan, 1998
Toward a theory of volatility trading
In Volatility, edited by Robert Jarrow, Risk Publications

Derman E and I Kani, 1994
Riding on a smile
Risk February, pages 32–39

Derman E and I Kani, 1998
Stochastic implied trees: arbitrage pricing with stochastic term and strike structure
International Journal of Theoretical and Applied Finance 1, pages 61–110

Dupire B, 1994
Pricing with a smile
Risk January, pages 18–20

Dupire B, 2004
A unified theory of volatility
In Derivatives Pricing: The Classic Collection, edited by Peter Carr, Risk Publications

Gyöngy I, 1986
Mimicking the one-dimensional marginal distributions of processes having an Ito differential
Probability Theory and Related Fields 71, pages 501–516

Jex M, R Henderson and D Wang, 1999
Pricing exotics under the smile
Risk November, pages 72–75