Smile dynamics

Traditionally, smile models have been assessed according to how well they fit market option prices across strikes and maturities. However, the pricing of most recent exotic structures, such as reverse cliquets or Napoleons, is more dependent on the assumptions made for the future dynamics of implied volatilities than on today's vanilla option prices. Here, Lorenzo Bergomi studies examples of some popular classes of models, such as stochastic volatility and jump/Lévy models, to highlight some structural features of their dynamic properties

This article focuses on the dynamic properties of smile models. In the Black-Scholes model, by construction, implied volatilities for different strikes are equal and frozen. Over the years, several alternative models, starting with local volatility, have emerged that aim to fit market implied volatilities across strikes and maturities.

This capability is a desirable feature of any smile model: the model price then incorporates by construction the cost of trading vanilla options to hedge the exotic option's vega risk – at least for the initial trade. Otherwise, the price has to be manually adjusted to reflect hedging costs, that is, the difference between market and model prices of vanilla options used for the hedge. This may be sufficient if the vega hedge is stable, which is usually the case for barrier options.

However, most of the recent exotic structures, such as Napoleons and reverse cliquets[1], require rebalancing of the vega hedge when the underlier or its implied volatilities move substantially. To ensure that future hedging costs are priced-in correctly, the model has to be designed so that it incorporates from the start a dynamic for implied volatilities that is consistent with the historically experienced one.

Stated differently, for this type of options, $\partial^2 P/\partial\hat{\sigma}^2$ and $\partial^2 P/\partial S\partial\hat{\sigma}$ are sizeable and a suitable model needs to price in a theta to match these gammas. In our view, this issue is far more important than the model's ability to exactly reproduce today's smile surface.

An example. As an illustration, let us consider the following example of a Napoleon option with a maturity of six years. The client initially invests 100, then gets a 6% coupon for the first two years and, at the end of years three, four, five and six, an annual coupon of 8% augmented by the worst of the 12 monthly performances of the Eurostoxx 50 index observed each year, with the coupon floored at zero. At maturity, he also gets his 100 back. The payout for the last four coupons is designed so that their value at inception is very small, thereby financing the 'large' fixed initial coupons[2], which we remove from the option in what follows.

Figure 1 shows on the left the Black-Scholes value of the option at time $t = 0$, as a function of volatility. As we can see, the Napoleon is in effect a put option on long (one-year) forward volatility, for which no time value has been appropriated for in the Black-Scholes price (no theta matching $\partial^2 P/\partial\hat{\sigma}^2$).

Now let us move to the end of the first month of year three. The right-hand side of figure 1 shows the vega of the coupon of year three at 20% volatility, as a function of the spot price, assuming the spot value at the beginning of the year was 100. It is a decreasing function of the spot and goes to zero for low spot values, as the coupon becomes worthless. Now, as the spot decreases, the option seller will need to buy back vega. However, moves in spot prices are historically negatively correlated with moves in implied volatilities, resulting in a negative profit and loss to the seller, not accounted for in the Black-Scholes price (no theta matching $\partial^2 P/\partial S\partial\hat{\sigma}$).

The Black-Scholes price should thus be adjusted for the effect of the two cross-gammas mentioned, as well as for the one-month forward skew contribution.

Local volatility models (Dupire, 1994), whose *raison d'être* is their ability to exactly fit observed market smiles, have historically been used to price skew-sensitive options. Even though implied volatilities do move in these models, their motion is driven purely by the spot and is dictated by the shape of the market smile used for calibration. This also materialises in the fact that forward smiles depend substantially on the forward date and the spot value at the forward date.

It would be desirable to be able to independently calibrate today's market smile and specify its future dynamics. One can attempt to directly specify an *ab initio* joint process for implied volatilities and the spot. This approach has been explored (Schönbucher, 1999) but is hampered by the difficulty of ensuring no arbitrage in future smiles.

In this article, we focus on models based on a specification of the spot process. We consider some of the most popular models and characterise the dynamics of implied volatilities that they generate. Our purpose is not to be exhaustive; rather, we select examples of models and products to point out specific properties of the models at hand and, more importantly, structural features that are shared by classes of models. We comment on the pricing of specific products.

This article is organised as follows. In the next section, we set up a simple pricing and hedging framework for models in incomplete markets, specialising to the case of stochastic volatility and jump and Lévy processes. We discuss pricing equations and deltas. We then deal with the Heston model, typical of one-factor stochastic volatility models. Jump and Lévy processes and one of their stochastic volatility extensions are then covered. The concluding section summarises and presents our views on future work.

[1] *See review article by Jeffery (2004)*
[2] *As well as the distributor's fee, typically 1% a year*

1. Price and vega of Napoleon

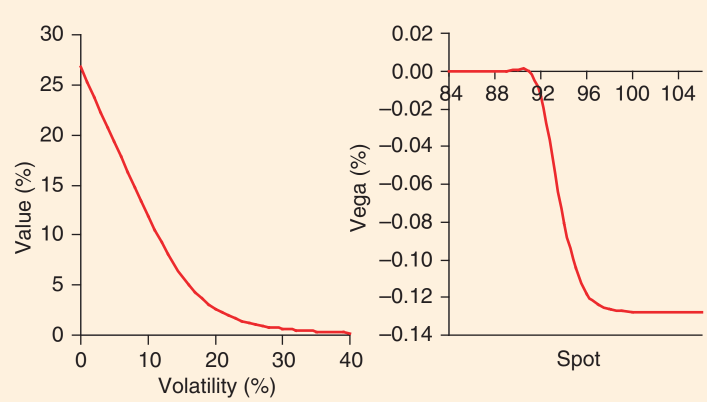

Note: (left) initial value of coupons of years three, four, five and six as a function of volatility; (right) vega of a coupon at the end of the first month, as a function of the spot price

Pricing and hedging

Pricing and hedging is in essence a stochastic control problem: once a measure of the replication risk has been specified, what is the optimal hedging strategy, and what price should be quoted?

In the usual Black-Scholes and local volatility framework, the only source of randomness is the spot process, which is diffusive. It turns out that the delta strategy not only minimises the replication risk, it eliminates it. This peculiar feature is typical of this framework and is not generic. Actually, the variance of the hedger's final profit and loss will be finite only because trading in the underlier does not occur continuously in time.

In more general settings, the variance of the final profit and loss (P&L) will be finite even though trading occurs continuously, either because the spot process is not continuous (this is typical of jump and Lévy processes) or because additional sources of randomness are present (as in stochastic volatility models) or both.

In this article, we derive pricing equations assuming that we only trade in the underlier. Our criterion is to minimise the variance of the hedger's discounted final profit and loss[3], which, for a European-style option reads:

$$P\,\&\,L = -e^{-r(T-t)} f\left(S_T\right) + \int_t^T e^{-r(\tau-t)} \Delta\left(\tau, S, \ldots\right)\left(dS_\tau - \left(r-q\right) S d\tau\right) \quad (1)$$

Here, f denotes the payout function, r is the interest rate, T is the maturity, and q incorporates both repo cost and dividend yield. Δ is a function of S and t, and may depend on other variables. Δ is determined by requiring that it minimises the variance of the profit and loss. We then define the price of the option as $-E[P\&L]$.

In contrast with approaches based on utility functions, we do not adjust the price for the residual risk. One reason is that, in practice, the option will be added to an existing book: the marginal variation in the risk upon adding an extra option depends on the existing book. The other reason is that, for the sake of simplicity, we want pricing to remain a linear operation: the price of a book is the sum of the prices of each option in the book. In the two following sections, we carry out this analysis, first for the Heston model, typical of one-factor stochastic volatility models, then for a jump model.

□ **Stochastic volatility – the Heston model.** In the Heston model (Heston, 1993), the historical dynamics for the spot process is:

$$\begin{aligned} dS &= \mu S dt + \sqrt{V} S dZ_t \\ dV &= -k\left(V - V_0\right) dt + \sigma \sqrt{V} dW_t \end{aligned} \quad (2)$$

where W, Z are Brownian motions with correlation ρ.

Let $m_\Delta(t, S, V)$ be the expectation and $v_\Delta(t, S, V)$ the variance of the hedger's discounted final profit and loss assuming zero initial wealth at time t. The subscript Δ indicates that m and v depend on the – as yet unknown – function $\Delta(t, S, V)$. In the Hamilton-Jacobi-Bellman (HJB) stochastic control formalism, one derives a partial differential equation for the 'value function' J. Here, the part of J is played by v_Δ, the control being Δ. In contrast to the usual HJB setting, the equation for v_Δ is not autonomous; it has to be supplemented with an equation for m_Δ. In what follows, we will drop the Δ subscripts for notational economy. From the dynamics (2), we derive the following coupled equations for m and v:

$$\frac{\partial m}{\partial t} + \mathcal{L} m - rm = -\left(\mu - r + q\right) S \Delta$$

$$\frac{\partial v}{\partial t} + \mathcal{L} v - 2rv = -VS^2 \left(\Delta + \frac{\partial m}{\partial S} + \frac{\rho\sigma}{S}\frac{\partial m}{\partial V}\right)^2 - \left(1-\rho^2\right)\sigma^2 V \left(\frac{\partial m}{\partial V}\right)^2$$

where the differential operator $\mathcal{L}$ reads:

$$\mathcal{L} = \mu S \frac{\partial}{\partial S} - k\left(V - V_0\right)\frac{\partial}{\partial V} + \frac{1}{2} V S^2 \frac{\partial^2}{\partial S^2} + \frac{1}{2}\sigma^2 V \frac{\partial^2}{\partial V^2} + \rho\sigma S V \frac{\partial^2}{\partial S \partial V}$$

At maturity, the profit and loss has no uncertainty anymore, hence the boundary conditions for m and v:

$$\begin{aligned} m\left(T, S, V\right) &= -f\left(S\right) \\ v\left(T, S, V\right) &= 0 \end{aligned}$$

As expected, the source term in the equation for m only involves the difference between the historical drift μ and $r-q$, which is the cost of trading in the underlying. The source term for v is the sum of two positive contributions: one generated by the spot, the other generated by the portion of volatility degrees of freedom that cannot be hedged by the spot. By variationally differentiating v with respect to Δ and requiring that v be minimal, we get the following expression for Δ:

$$\Delta = -\frac{\partial m}{\partial S} - \frac{\rho\sigma}{S}\frac{\partial m}{\partial V}$$

This expression of Δ makes the first source term in the equation for v cancel out. The second term remains: the variance of the final profit and loss does not vanish and there is no risk-neutral price for the option. We define the price P as $P = -m$.

By plugging the expression of Δ in the equation for m, we get the following equation for P:

$$\begin{aligned} &\frac{\partial P}{\partial t} + \left(r-q\right) S \frac{\partial P}{\partial S} - k\left(V - \overline{V_0}\right)\frac{\partial P}{\partial V} \\ &+ \frac{1}{2} V S^2 \frac{\partial^2 P}{\partial S^2} + \frac{1}{2}\sigma^2 V \frac{\partial^2 P}{\partial V^2} + \rho\sigma S V \frac{\partial^2 P}{\partial S \partial V} = rP \end{aligned} \quad (3)$$

where:

$$\overline{V_0} = V_0 - \frac{\left(\mu - r + q\right)\rho\sigma}{k}$$

Δ is given by:

$$\Delta = \frac{\partial P}{\partial S} + \frac{\rho\sigma}{S}\frac{\partial P}{\partial V} \quad (4)$$

A few observations are in order:

- As expected, the pricing drift for the spot is its financing cost $r-q$.
- The Black-Scholes delta and price are recovered when σ tends to zero.
- The second portion of the delta is the ratio of the covariance of V and S increments to the variance of S increments.
- V_0 is renormalised. This is due to the fact that the volatility degree of freedom is partially hedged by trading in the underlying. Note that $\overline{V_0}$ keeps the same functional form as V_0 (here a constant) so that the pricing equation keeps its usual form. In other stochastic volatility models, the functional form for the pricing drift of V as a function of S and V will be different, unless $\mu = r - q$.

We will use the above pricing equation in the sequel and replace $\overline{V_0}$ with V_0 for notational economy. As in the Black-Scholes framework, the pricing equation generalises to path-dependent options.

□ **Jump models.** Now we apply the same ideas to jump models (Merton, 1976). Let the process for the spot be a jump-diffusion process where σ is the volatility, λ the intensity of the Poisson process and J the size of the jumps, itself a random variable:

$$dS = \mu S dt + \sigma S dZ_t + JS dq_t$$

Now Δ will depend solely on S and t. The equations for m and v are:

$$\frac{\partial m}{\partial t} + \mathcal{L} m - rm = -\left(\mu - r + q + \lambda\overline{J}\right) S \Delta$$

$$\frac{\partial v}{\partial t} + \mathcal{L} v - 2rv = -\sigma^2 S^2 \left(\frac{\partial m}{\partial S} + \Delta\right)^2 - \lambda\left(\overline{\delta m^2} + \overline{J^2}\left(\Delta S\right)^2 + 2\overline{\delta m J}\Delta S\right)$$

with the same boundary conditions as in the previous section. We have used the following notation: $\delta m = m(S(1+J), t) - m(S, t)$ and $\overline{f} = E[f]$, where the expectation is taken over J, the amplitude of the jump. The integro-differential operator $\mathcal{L}$ is defined as:

$$\mathcal{L} f = \mu S \frac{\partial f}{\partial S} + \lambda\overline{\delta f} + \frac{\sigma^2 S^2}{2}\frac{\partial^2 f}{\partial S^2}$$

[3] *See Bouchaud & Potters (2000) for a treatment in discrete time*

2. Fitted values of V, V_0, σ and ρ

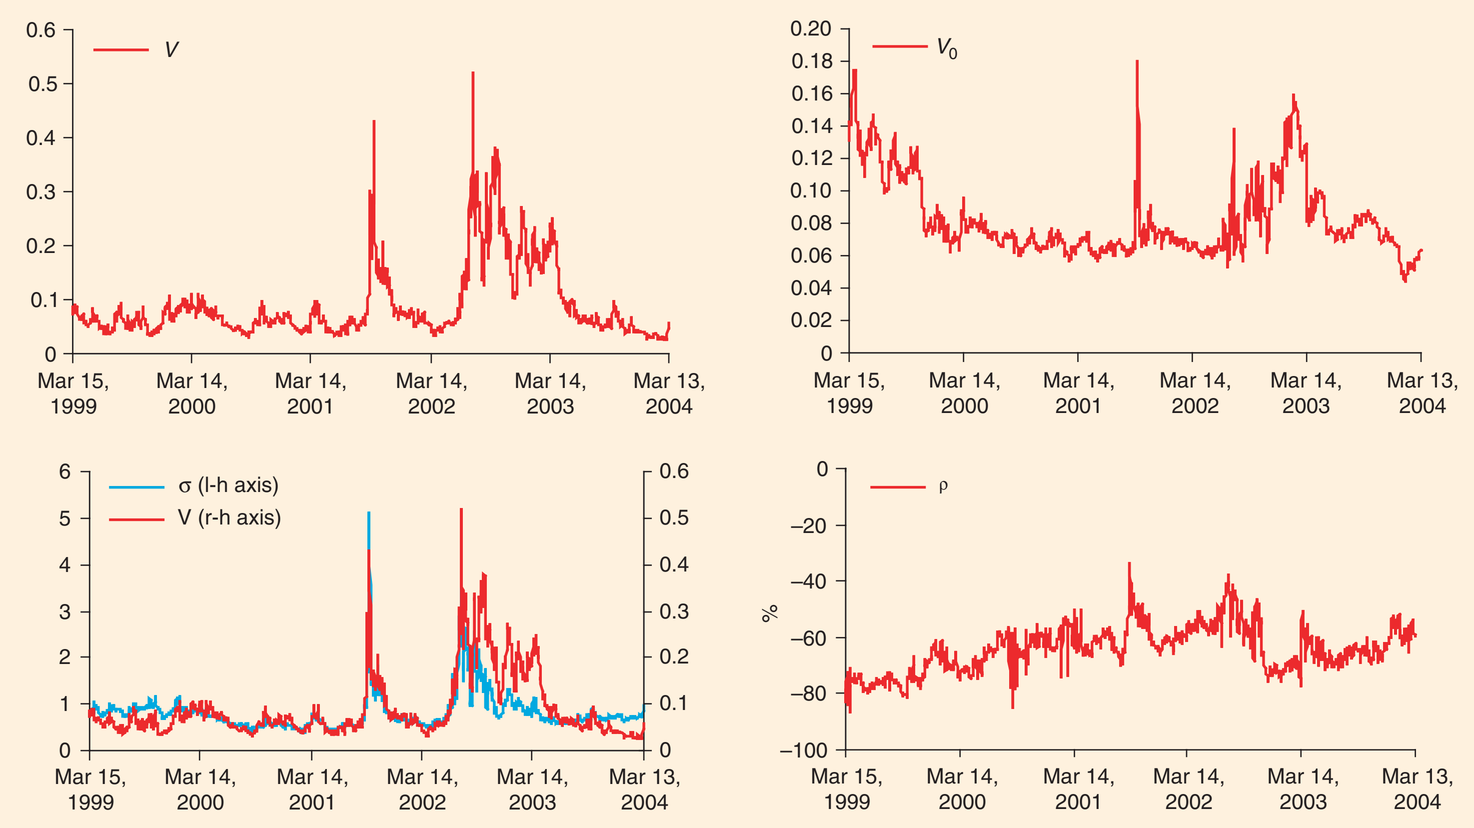

Differentiating with respect to Δ yields the following expression:

$$\Delta = -\frac{\sigma^2 \frac{\partial m}{\partial S} + \lambda \overline{J^2}\,\frac{\overline{\delta m J}}{S\overline{J^2}}}{\sigma^2 + \lambda\overline{J^2}} \tag{5}$$

which is readily interpreted as the ratio of the covariance of the price and spot increments – either generated by diffusion or by jumps – to the variance of the spot increments. Let us now set $P = -m$ and use the above expression for Δ.

A power expansion in the size of jumps yields at the lowest non-trivial order:

$$\Delta \approx \frac{\partial P}{\partial S} + \frac{1}{2}\frac{\lambda\overline{J^3}}{\sigma^2 + \lambda\overline{J^2}} S \frac{\partial^2 P}{\partial S^2} \tag{6}$$

Because Δ is different from $\partial P/\partial S$, one can see in the equation for m that μ remains in the pricing equation. Thus by using the delta in equation (5) we are making a bet on the realised historical drift. This materialises in the fact that our profit and loss between two re-hedges comprises a linear term of the form $(\Delta - \partial P/\partial S)\delta S$, where δS is the variation of the spot. This is the price we pay for having an 'optimal' delta that takes jumps into account: in exchange for reducing the contribution of jumps to the variance of the profit and loss, we increase the contribution of the 'normal' diffusive behaviour.

Thus it will be sensible to use the delta in equation (5) only if the specification of the jump model is in agreement with the historical dynamics of the underlying. As this is not guaranteed, it may be wiser to choose $\Delta = -\partial m/\partial S$ so as to remove the linear contribution in the profit and loss. This is the choice we make here.[4] The pricing equation then reads:

$$\frac{\partial P}{\partial t} + (r-q)S\frac{\partial P}{\partial S} + \lambda\left(\overline{\delta P} - \overline{J}S\frac{\partial P}{\partial S}\right) + \frac{\sigma^2 S^2}{2}\frac{\partial^2 P}{\partial S^2} = rP \tag{7}$$

and the associated delta is:

$$\Delta = \frac{\partial P}{\partial S} \tag{8}$$

This can be generalised to Lévy processes and path-dependent options.

Dynamic properties – stochastic volatility (the Heston model)

Here, we examine the Heston model, a typical example within the class of one-factor stochastic volatility models. First, we characterise its static properties. Next we compare the model-generated dynamics of implied volatilities with their historical dynamics. Then we comment on the pricing of forward-start options and end with a discussion of the delta and a comparison with local volatility models.

□ **The Heston model.** The Heston model has five parameters, V, V_0, ρ, σ and k, among which k plays a special role: $\tau = 1/k$ is a cutoff that separates short and long maturities. The Heston model is homogeneous: implied volatilities are a function of V and moneyness: $\hat{\sigma} = f(K/F, V)$, where F is the forward price. Perturbation of the pricing equation at first order in σ yields the following expressions for the skew and at-the-money-forward volatility:

- $T \ll \tau$, at order zero in T:

$$\hat{\sigma}_F = \sqrt{V}, \qquad \left.\frac{d\hat{\sigma}}{d\ln K}\right|_F = \frac{\rho\sigma}{4\sqrt{V}} \tag{9}$$

- $T \gg \tau$, at order one in $1/T$:

$$\begin{aligned} \hat{\sigma}_F &= \sqrt{V_0}\left(1 + \frac{\rho\sigma}{4k}\right) + \frac{\sqrt{V_0}}{2kT}\left(\frac{V - V_0}{V_0} + \frac{\rho\sigma}{4k}\frac{V - 3V_0}{V_0}\right), \\ \left.\frac{d\hat{\sigma}}{d\ln K}\right|_F &= \frac{\rho\sigma}{2kT\sqrt{V_0}} \end{aligned} \tag{10}$$

[4] *In a context where jumps are used to model rare and extreme events with the purpose of reducing the size of the profit and loss upon a jump, the delta in equation (5) would be used*

3. $\sqrt{V}$ and one-month ATM volatility, and $\hat{\sigma}_{VS}$ and one-year ATM volatility

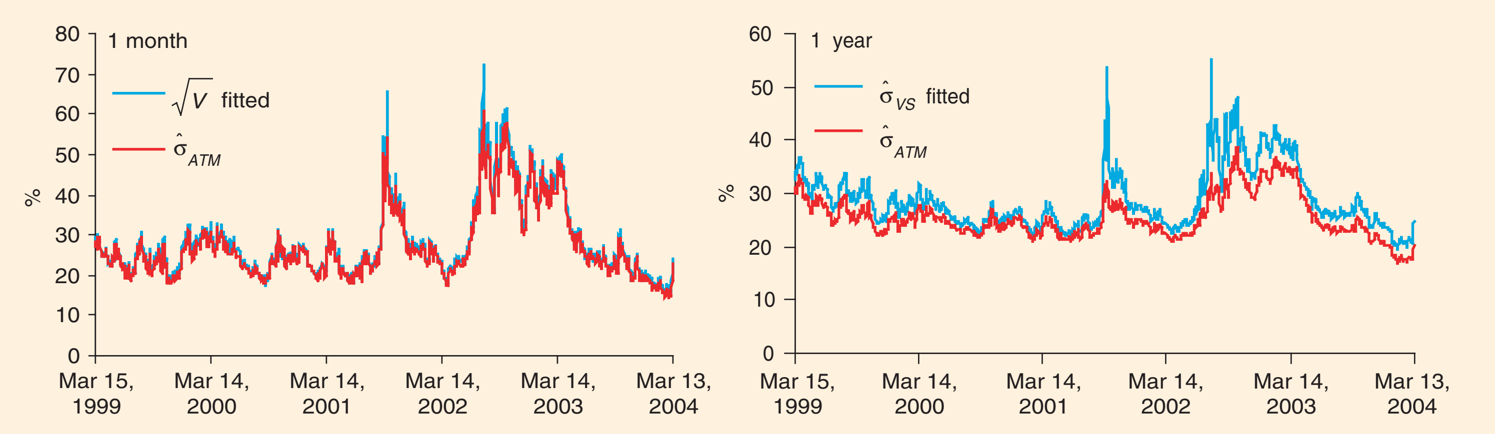

The long-term behaviour of the skew is what we expect: in a stochastic volatility model with mean reversion, increments of $\ln(S)$ become stationary and independent over long periods. Thus the skewness of $\ln(S)$ scales like $1/\sqrt{T}$; consequently[5], the skew decreases like $1/T$.

Let us write the expression of the variance swap volatility $\hat{\sigma}_{VS}(T)$, defined such that $T\hat{\sigma}^2_{VS}(T)$ is the expectation of the realised variance for maturity T:

$$\hat{\sigma}^2_{VS}(T) = V_0 + (V - V_0)\frac{1-e^{-kT}}{kT} \tag{11}$$

□ **Dynamics of implied volatilities.** We have calibrated the market implied volatilities of the Eurostoxx 50 index from March 12, 1999, to March 12, 2004, for options with maturities of one month, three months, six months and one year.

Although the dynamics of both short and long implied volatilities in the model is driven by V, equation (11) shows that the dynamics of V is mostly reflected in that of short volatilities. We thus choose $k = 2$ and fit all other parameters. The daily historical values for V, V_0, σ and ρ are shown in figure 2.

We can see surges in volatility on September 11, 2001, in summer 2002 following the WorldCom collapse and in spring 2003 at the beginning of the second Gulf war.

Figure 3 shows how well levels of short and long implied volatilities are tracked. The graph on the left shows the at-the-money one-month implied volatility and $\sqrt{V}$.

The right-hand graph in figure 3 shows the one-year at-the-money volatility as well as the one-year variance swap volatility, calculated from V and V_0 using equation (11). We see that, as we would expect for equity smiles, the variance swap volatility lies higher than the at-the-money volatility. Here, too, the calibration is satisfactory.

■ Discussion. In the Heston model, while S and V are dynamic, V_0, ρ and σ are supposed to be constant. Their dynamics is not priced-in by the model. Figure 2 shows that: V_0 moves, but this is expected as the model fits both short and long implied volatilities; ρ is fairly stable, and does not seem correlated with other parameters; and σ is the most interesting parameter. We have superimposed the graph of V with a scale 10 times larger. We see that σ varies substantially and seems closely correlated with V.

The last observation can be accounted for by looking at the approximate expression for the short-term skew. Equation (9) shows that in the Heston model it is inversely proportional to $\sqrt{V}$, which is approximately equal to the at-the-money volatility. The fact that fitted values for σ are roughly proportional to V suggests that market skews are proportional to at-the-money volatilities, rather than inversely proportional.

In this respect the model is mis-specified, since it is not pricing in the observed correlation between V and σ. This correlation is very visible in graphs for V and σ, mostly for extreme events. However, it is high even in more normal regimes. For example, daily variations of V and σ measured from March 15, 1999, to September 10, 2001, have a correlation of 59%.

The recent past shows different behaviour: starting in summer 2003, while at-the-money volatilities decreased, skews steepened sharply, an effect that the Heston model naturally generates. Figure 2 indeed shows that during that period σ remains stable while V decreases.

Let us now turn to the dynamics of implied volatilities generated by the model, as compared with the historical one. In the Heston model, the implied volatility dynamics is determined, by construction, by that of S and V.

We can use daily values for the couple (S, V) to check whether their dynamics is consistent with the model specification (2). Let us calculate the following averages, which in theory should all be equal to one:

$$R_S = \left\langle \frac{\delta S^2}{S^2 V\delta t} \right\rangle = 0.75 \quad R_V = \left\langle \frac{\delta V^2}{\sigma^2 V} \right\rangle = 0.4 \quad R_{SV} = \left\langle \frac{\delta S \delta V}{\rho\sigma S V\delta t} \right\rangle = 0.6$$

where brackets denote historical averages using daily variations.

From these numbers, we estimate that:

$$\frac{\sigma_{realised}}{\sigma_{implied}} = \sqrt{R_V} = 0.63 \qquad \frac{\rho_{realised}}{\rho_{implied}} = \frac{R_{SV}}{\sqrt{R_V R_S}} = 1.1$$

suggesting that calibration on market smiles overestimates the volatility of volatility σ by 40%, while the value of the spot/volatility correlation ρ is captured with acceptable accuracy.

Surprisingly, R_S is notably different from one, showing that short implied volatilities overestimated historical volatility by 13% on our historical sample, possibly accounting for the enduring popularity of dispersion trades.

It is possible that these global averages are excessively affected by extreme events. Let us then look at running monthly averages. Figure 4 shows the results for the six following quantities:

$$V^{real}_{\frac{\delta S}{S}} = \left\langle \frac{\delta S^2}{S^2} \right\rangle \quad \text{and} \quad V^{impl}_{\frac{\delta S}{S}} = \langle V\delta t \rangle$$

$$V^{real}_{\delta V} = \langle \delta V^2 \rangle \quad \text{and} \quad V^{impl}_{\delta V} = \langle \sigma^2 V\delta t \rangle$$

$$C^{real}_{\frac{\delta S}{S}\delta V} = \left\langle \frac{\delta S}{S}\delta V \right\rangle \quad \text{and} \quad C^{impl}_{\frac{\delta S}{S}\delta V} = \langle \rho\sigma V\delta t \rangle$$

where brackets now denote running monthly averages. The sign of:

$$C^{real}_{\frac{\delta S}{S}\delta V} \quad \text{and} \quad C^{impl}_{\frac{\delta S}{S}\delta V}$$

has been changed.

We see that even during normal market conditions, the difference between realised and implied quantities is substantial. For example, using monthly running averages estimated on data from March 15, 1999, to Sep-

[5] *See Backus* et al *(1997)*

4. $V^{real}_{\delta S/S}$ and $V^{impl}_{\delta S/S}$, $V^{real}_{\delta V}$ and $V^{impl}_{\delta V}$, $C^{real}_{(\delta S/S)\delta V}$ and $C^{impl}_{(\delta S/S)\delta V}$

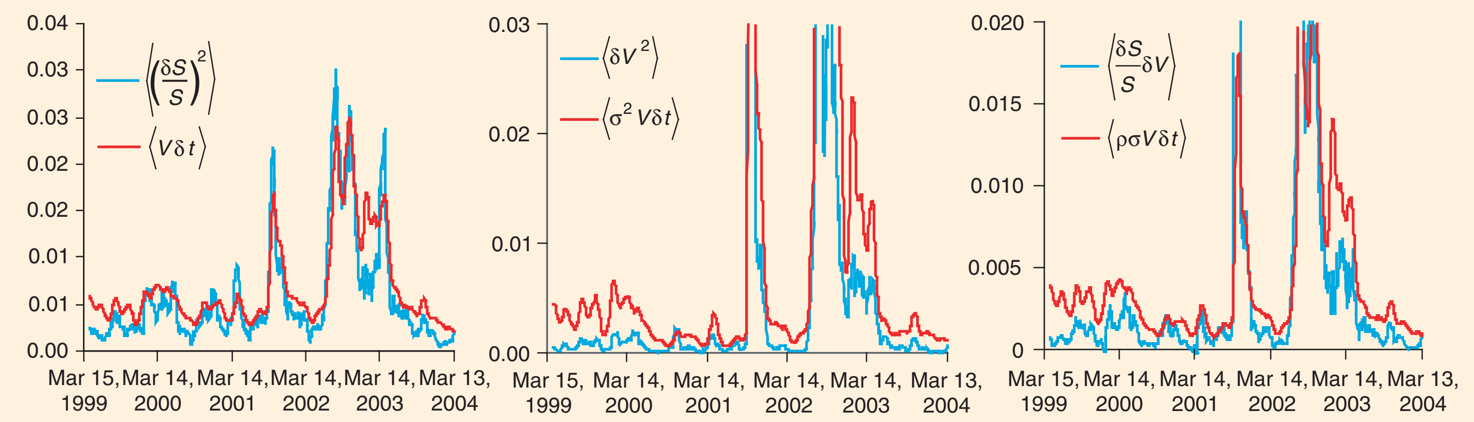

tember 10, 2001, gives the following numbers:

$$R_S = 0.73, \quad R_V = 0.30, \quad R_{SV} = 0.44$$

corresponding to the following ratios:

$$\frac{\sigma_{realised}}{\sigma_{implied}} = 0.54, \qquad \frac{\rho_{realised}}{\rho_{implied}} = 0.95$$

again showing that, while the spot/volatility correlation ρ is well captured by market smiles, the volatility of volatility σ is overestimated by roughly a factor of two.

This means that the model is pricing in a volatility of volatility for one-month at-the-money volatilities that is twice as large as its historical value: future vega re-hedging costs are not properly priced-in. It also implies that the delta:

$$\Delta = \frac{\partial P}{\partial S} + \frac{\rho\sigma}{S}\frac{\partial P}{\partial V}$$

is not efficient, as it over-hedges the systematic impact of spot on volatility.

The main results of our historical analysis are: σ and V are closely correlated; and the value of σ determined from calibration on market smiles is larger by a factor of two than its historical value.

While the first of these could be solved by altering the model's specification, the second is structural. Indeed, we have only one device in the model – namely the volatility of volatility σ – to achieve two different objectives, one static, the other dynamic: create skewness in the distribution of $\ln(S)$ so as to match market smiles, and drive the dynamics of implied volatilities in a way that is consistent with their historical behaviour. It is natural that we are unable to fulfil both objectives. We view this as a structural limitation of any one-factor stochastic volatility model.

We have concentrated here on the dynamics of short-term volatilities. Space prevents us from examining the crucial issue of the term structure of the volatility of volatilities, which is controlled by the correlation function of V.

□ **Forward-start options.** Here we consider a one-period forward call option that pays:

$$\left(\frac{S_{T_1+\theta}}{S_{T_1}} - \xi\right)^+$$

at date $T_1 + \theta$, for different values of moneyness ξ. From the model-generated price of the forward-start option we imply Black-Scholes volatilities to get what is called the forward smile $\hat{\sigma}(\xi)$.

Figure 5 shows the forward smile calculated using the following typical values: $V = V_0 = 0.1$, $\sigma = 1$, $\rho = -0.7$ and $k = 2$, for two values of θ: 0.25 (three months) and one (one year). Today's smile ($T_1 = 0$) is also plotted for reference.

Note that forward smiles are more convex than today's smile: since the price of a call option is an increasing and convex function of its implied volatility, uncertainty in the value of future implied volatility increases the option price.

As T_1 is more distant, the distribution for V becomes stationary in the Heston model. Thus forward smiles collapse on to a single curve for $T_1 >>$ 6 months, in our example. This is manifest in figure 5.

The graphs also show that the increased convexity with respect to today's smile is larger for strikes $\xi > 100\%$ than for strikes $\xi < 100\%$. This can be traced to the dependence of the skew on the level of at-the-money volatility. Since the short-term skew is inversely proportional to the at-the-money volatility, implied volatilities for strikes above 100% will move more than those for symmetrical strikes below 100%. This is specific to the Heston model.

While the forward smile is a global measure of the distribution of implied volatilities at a forward date, it is instructive to look at the distribution itself. Let $T_1 >> 1/k$. The density of V has the following stationary form:

$$\rho(V) \propto V^{\left(\frac{2kV_0}{\sigma^2}-1\right)} e^{-\frac{2k}{\sigma^2}V}$$

Using the parameter values listed above, we find that $(2kV_0/\sigma^2) - 1 = -0.6$, that is, the density for V diverges for small values of V.

Thus even simple cliquets are substantially affected by the model specification. The practical conclusion for pricing is that, for short-term forward-start options, the Heston model is likely to overemphasise low at-the-money volatility/high skew scenarios.

□ **Local dynamics and delta.** We here study the local dynamics of the Heston model: how do implied volatilities move when the spot moves? This sheds light on the model's delta since its deviation from the Black-Scholes value is related to the model's expected shift in implied volatilities when the spot moves.

In local volatility models, the motion of implied volatilities is driven by the spot. From the expression of the local volatility (Dupire, 1994), with the assumption of short maturity and weak skew, one can derive the following well-known relationship linking the skew to the dynamics of the at-the-money volatility as a function of the spot:

$$\frac{d\hat{\sigma}_{K=S}}{d\ln S} = 2\left.\frac{d\hat{\sigma}}{d\ln K}\right|_{K=S}$$

showing that $\hat{\sigma}_{K=S}$ moves 'twice as fast' as the skew.

In stochastic volatility models, while implied volatilities are not a function of S, they are correlated with S. This is what the second piece of the delta in equation (4) hedges against. Conditional on a small move of the spot δS, V moves on average by $\delta V = (\rho\sigma/S)\delta S$.

Let us calculate the expected variation in $\hat{\sigma}_F$, for short and long maturities:

■ For $T << \tau$ we use expressions (9), correct at order zero in T. At this

5. $\hat{\sigma}(\xi)$ for one-year and three-month maturity

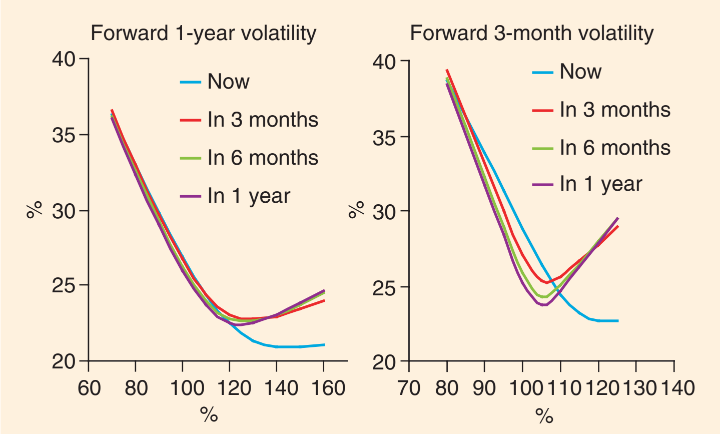

order, F and S can be identified. The expression for σ_F gives:

$$E\left[\delta\hat{\sigma}_{K=S}\right]=\frac{\rho\sigma}{2\sqrt{V}}\frac{\delta S}{S}$$

Looking at the expression for the skew, we notice that:

$$\frac{E\left[\delta\hat{\sigma}_{K=S}\right]}{\delta\ln S}=2\left.\frac{d\hat{\sigma}}{d\ln K}\right|_{K=S} \tag{12}$$

This shows that, locally, the shift in implied volatilities expected by the Heston model when the spot moves is identical to that of a local volatility model. Thus the deltas of vanilla options for strikes near-the-money will be the same in both models – at order one in σ. This result is generic and holds for all stochastic volatility models.

■ For $T >> \tau$ we use expressions (10), correct at order one in $1/T$. We get, keeping only terms linear in σ:

$$E\left[\delta\hat{\sigma}_{K=F}\right]=\frac{\rho\sigma}{2kT\sqrt{V_0}}\frac{\delta S}{S}$$

Comparing with the expression of the skew in equation (10), we see that:

$$\frac{E\left[\delta\hat{\sigma}_{K=F}\right]}{\delta\ln S}=\left.\frac{d\hat{\sigma}}{d\ln K}\right|_{K=F}$$

The at-the-money-forward volatility slides along the smile and the Heston model behaves like a sticky-strike model: implied volatilities for fixed strikes do not move as the spot moves. Thus the deltas of vanilla options for strikes near the forward will be equal to their Black-Scholes deltas – again at order one in σ. The possible extension to other stochastic volatility models is left for future work.

These results are obtained for the Heston model at first order in σ and are relevant for equity smiles. If ρ is small, as is the case for currency smiles, the contribution from terms of order σ^2 dominates, altering the conclusions: for example, the similarity to local volatility models for short maturities will be lost.

Dynamic properties – jump/Lévy models

□ **Jump/Lévy models.** Here, we consider jump models for which the size of the relative jump experienced by the spot does not depend on the spot level. Such models are homogeneous: implied volatilities are a function of moneyness $\hat{\sigma}(K, S) = \hat{\sigma}(K/S)$.

The spot is the only degree of freedom in the model. As it moves, the smile experiences a translation along with it: for a fixed moneyness, implied volatilities are frozen. This has two consequences:

■ Forward smiles do not depend on the forward date and are the same as today's smile: a graph similar to figure 5 would show all smiles collapsing on to a single curve. When pricing a cliquet, this is equivalent to giving all forward-start options the same smile cost.

■ The deltas for vanilla options are model-independent and can be read off the smile directly. The delta for strike K is given by:

$$\Delta_K=\Delta_K^{BS}-\frac{1}{S}Vega_K^{BS}\frac{d\hat{\sigma}_K}{d\ln K}$$

where Δ_K^{BS} and $Vega_K^{BS}$ are the Black-Scholes delta and vega of the vanilla option of strike K calculated with its implied volatility $\hat{\sigma}_K$.

In jump/Lévy models, increments of $\ln(S)$ are independent, so the skewness of $\ln(S_T)$ scales like $1/\sqrt{T}$, and, at first order in the skewness, the skew decreases as $1/T$, too fast in comparison with market smiles.

Stochastic volatility models generate a smile by starting with a process for $\ln(S)$, which is Gaussian at short time scales, and making volatility stochastic and correlated with the spot process. In contrast, jump/Lévy models generate a smile without additional degrees of freedom by starting with a process for $\ln(S)$ at short time scales with sufficient embedded skewness and kurtosis so that both are still large enough at longer time scales to generate a smile, even though they scale like $1/\sqrt{T}$ and $1/T$, respectively.

In the next section, we use the example of variance swaps to illustrate how the behaviour of jump/Lévy models at short time scales affects the price of very path-dependent options.

□ **Variance swaps.** A variance swap (VS) is a now standard option that pays at maturity the realised variance of the spot, measured as the sum of squared returns observed at discrete dates – usually daily.

If the observations are frequent enough, the price P_{VS} is just the discounted expected variance by construction:

$$P_{VS}=e^{-rT}\hat{\sigma}_{VS}^2$$

We now introduce the log swap volatility $\hat{\sigma}_{LS}(T)$. This is the implied volatility of the log swap, which is the European-style payout $-2\ln(S)$. This profile, when delta-hedged, generates a gamma profit and loss that is equal[6] to the squared return of the spot between two re-hedging dates. Because this statically replicates the payout of a VS, VSs are usually priced using $\hat{\sigma}_{LS}(T)$. In the Black-Scholes model, with the limit of very frequent observations, $\hat{\sigma}_{LS} = \hat{\sigma}_{VS} = \sigma$.

The value of $\hat{\sigma}_{LS}(T)$ is the implied volatility of a European-style payout; it is thus model-independent and is derived from the market smile. For equity smiles, $\hat{\sigma}_{LS}(T)$ usually lies higher than the at-the-money volatility. For example, in early March of this year, because of the high skew/low volatility context, the one-year $\hat{\sigma}_{VS}$ for the Eurostoxx 50 index was about four points higher than the at-the-money volatility, which was around 18%.

In the Heston model, direct calculation yields $\hat{\sigma}_{VS}(T) = \hat{\sigma}_{LS}(T)$. This self-consistency can be shown to hold for all diffusive models.

In jump/Lévy models, however, $\hat{\sigma}_{VS}$ is usually lower than $\hat{\sigma}_{LS}$ and even lower than $\hat{\sigma}_{ATM}$. For example, with the limit of frequent jumps of small amplitude, the following relationship can be derived, at first order in the skewness:

$$\hat{\sigma}_{K=F}-\hat{\sigma}_{VS}=3\left(\hat{\sigma}_{LS}-\hat{\sigma}_{K=F}\right)$$

where $\hat{\sigma}_{K=F}$ is the volatility for a strike equal to the forward.

The question then is: to price VSs, should we use $\hat{\sigma}_{VS}$ or $\hat{\sigma}_{LS}$ or yet another volatility? To understand the difference, imagine hedging the profile $-2\ln(S)$ with the Black-Scholes delta calculated with an implied volatility $\hat{\sigma}$. If there are no dividends, the delta is independent of the volatility, equal to $-2/S$. The gamma portion of the gamma/theta profit and loss realised during Δt, stopping at third-order terms in ΔS reads:

$$\left(\frac{\Delta S}{S}\right)^2-\frac{2}{3}\left(\frac{\Delta S}{S}\right)^3$$

Introducing the volatility σ, given by $\sigma^2\Delta t = E[(\Delta S/S)^2]$, and the skewness $\mathcal{S}_{\Delta t}$ of $\Delta S/S$, we can write the expectation of this profit and loss as:

[6] *Except if dividends are modelled as discrete cash amounts*

$$\sigma^2\Delta t\left(1-\frac{2\mathcal{S}_{\Delta t}}{3}\sigma\sqrt{\Delta t}\right)$$

Let us take the limit $\Delta t \to 0$.

■ In stochastic volatility models, as $\Delta t \to 0$, returns become Gaussian and $\mathcal{S}_{\Delta t} \to 0$. Thus the profit and loss generated by delta-hedging the log swap profile is exactly the realised variance. This explains why $\hat{\sigma}_{LS}$ and $\hat{\sigma}_{VS}$ are the same.

■ In jump/Lévy models, because $\mathcal{S}_{\Delta t} \propto 1/\sqrt{\Delta t}$, the third-order term contribution tends to a finite constant as $\Delta t \to 0$. Delta-hedging the log swap profile generates an additional contribution from third-order terms.[7]

For equity smiles $\mathcal{S}$ is negative. Delta-hedging the log swap profile then generates, in addition to the realised variance, a spurious positive profit and loss. Thus, the *VS* should be priced using a volatility lower than $\hat{\sigma}_{LS}$: $\hat{\sigma}_{VS} < \hat{\sigma}_{LS}$.

If real underliers behaved according to the jump/Lévy model specification, we should price *VS*s using $\hat{\sigma}_{VS}$. The daily returns of the Eurostoxx 50 index show, however, that $\mathcal{S}_{\Delta t}$ is a number of order one. Using a daily volatility of 2% gives an estimation of the contribution of the third-order term about 50 times smaller than that of the second-order term, in sharp contrast with the model's estimation.

The conclusion for the pricing of *VS*s is that it will be more appropriate to use $\hat{\sigma}_{LS}$. More generally, we have to be aware of the fact that, once their parameters are calibrated to market smiles, jump/Lévy models will predict excessive skews at short time scales; this behaviour is structural.

□ **Stochastic volatility extensions to jump/Lévy models.** A simple way of adding dynamics to implied volatilities in a jump/Lévy model is to make the flow of time stochastic: replace t with a non-decreasing process τ_t and evaluate the Lévy process L at τ. This is a particular case of a subordinated process. If the characteristic functions of both L_t and τ_t are known, then the characteristic function of L_τ is also known and an inverse Laplace transform yields European-style option prices. Carr *et al* (2003) choose τ_t as the integral of a Cox-Ingersoll-Ross process:

$$\tau_t = \int_0^t \lambda_u du \qquad d\lambda = -k\left(\lambda-\lambda_0\right)dt + \sigma\sqrt{\lambda}dZ_t$$

What is the dynamics of implied volatilities in such a model? Here, we look at short-term options. The shape of the smile for maturity T is determined by the distribution of $\ln(S_T)$. Given the variance $\mathcal{V}$ and the skewness $\mathcal{S}$ of a distribution for $\ln(S_T)$, perturbation at first order in $\mathcal{S}$ gives (Backus *et al*, 1997):

$$\hat{\sigma}_{K=F} = \sqrt{\frac{\mathcal{V}}{T}} \tag{13}$$

$$K\left.\frac{d\hat{\sigma}}{dK}\right|_{K=F} = \frac{\mathcal{S}}{6\sqrt{T}} \tag{14}$$

where F is the forward of maturity T.

Because λ_t is a continuous process, for short maturities $L_{\tau_T} \approx L_{\lambda T}$. In other words, λ acts as a pure scale factor on time. Since the cumulants of L all scale linearly with time, we have:

$$\mathcal{V} \propto \lambda T \qquad \mathcal{S} \propto \frac{1}{\sqrt{\lambda T}}$$

Plugging these expressions in equations (13) and (14), we get the following form for the at-the-money-forward volatility and skew, for short maturities:

$$\hat{\sigma}_{K=F} \propto \sqrt{\lambda} \qquad K\left.\frac{d\hat{\sigma}}{dK}\right|_{K=F} \propto \frac{1}{\sqrt{\lambda}T}$$

Let us examine the scaling behaviour of these expressions. The dependence of volatility and skew on T is what we would expect. More interesting is the dependence on λ. Combining both equations yields the following:

$$K\left.\frac{d\hat{\sigma}}{dK}\right|_{K=F} \propto \frac{1}{\hat{\sigma}_{K=F}}$$

Thus, for short maturities, the skew is approximately inversely proportional to the at-the-money-forward volatility.

This result is interesting in that it is general for the class of models considered: it depends neither on the choice of Lévy process nor on the process for λ. Thus, affecting time with a stochastic scale factor allows implied volatilities to move but with a fixed dependence of the short-term skew on the level of at-the-money forward volatility. As noted earlier, this feature is also shared by the Heston model, for very different reasons. To get different behaviour, we would need to make the parameters of the Lévy process λ-dependent, probably losing the analytical tractability of the model.

Conclusion

As mentioned in the introduction, we believe that analysing and controlling the dynamics of implied volatilities is a central issue in the construction of models for pricing the recent breed of exotic structures as well as general path-dependent options that cannot be hedged statically.

We have studied some aspects of the dynamic properties of implied volatilities for two of the most popular classes of models: stochastic volatility and jump/Lévy models. We have also pointed out some of the structural implications of choosing a particular type of driving process for the spot.

It is our assessment that, in addition to the spot process, at least another driving process is needed to model the implied volatility dynamics, and presumably more than one if our objective is to correctly match the term-structure of the volatility of volatilities. To avoid the inconsistencies noted in the analysis of the dynamic properties of the Heston model, such a model would have 'state variables', calibrated to market smiles, whose dynamics is priced-in, and 'structural parameters', either calibrated to the historical dynamics of implied volatilities, or chosen by the trader. Within the set of model parameters, it would then be very useful to be able to precisely identify those governing static features and those governing dynamic features of the model-generated smiles. In this respect, associating stochastic volatility with Lévy processes seems a promising avenue of research. ■

Lorenzo Bergomi is head of quantitative research in the equity derivatives department at Société Générale. He is based in Paris. Email: lorenzo.bergomi@sgcib.com

REFERENCES

Backus D, S Foresi, K Li and L Wu, 1997
Accounting for biases in Black-Scholes
Unpublished

Bouchaud J-P and M Potters, 2000
Theory of financial risks
Cambridge University Press

Carr P, H Geman, D Madan and M Yor, 2003
Stochastic volatility for Levy processes
Mathematical Finance 13(3), July, pages 345–382

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

Heston S, 1993
A closed-form solution for options with stochastic volatility with applications to bond and currency options
Review of Financial Studies 6, pages 327–343

Jeffery C, 2004
Reverse cliquets: end of the road?
Risk February, pages 20–22

Merton R, 1976
Option pricing with discontinuous returns
Bell Journal of Financial Economics 3, pages 145–166

Schönbucher P, 1999
A market model for stochastic implied volatility
Philosophical Transactions of the Royal Society, Series A, 357(1,758), August, pages 2,071–2,092

[7] *Higher-order terms would also yield at each order a non-vanishing contribution. For example, the contribution from the fourth-order term would have included a factor* $\Delta t(3 + \kappa_{\Delta t})$, *where* $\kappa_{\Delta t}$ *is the kurtosis of returns. Since, as mentioned in the previous section,* $\kappa_{\Delta t} \propto 1/\Delta t$, *the contribution from the fourth order is finite as* $\Delta t \to 0$