Smile dynamics II

In an article published in *Risk* in September 2004, Lorenzo Bergomi highlighted how traditional stochastic volatility and jump/Lévy models impose structural constraints on the relationship between the forward skew, the spot/volatility correlation and the term structure of the volatility of volatility. Here, he proposes a model that enables them to be controlled separately and also prices options on realised variance consistently. He presents pricing examples for a reverse cliquet, a Napoleon, an accumulator and an option on variance

A common feature of the recent breed of exotic options such as Napoleons and reverse cliquets is that their price depends on assumptions made for the joint dynamics of the underlying and its implied volatilities. These fall into three categories:

- The dynamics of implied volatilities, and more specifically the term structure of the volatility of volatility.
- The forward skew.
- The spot/volatility correlation.

In a previous article (Bergomi, 2004), we analysed popular stochastic volatility and jump/Lévy models and pointed out that, although these models produce prices that include an estimation of the three effects listed above, they impose structural constraints on how these features of the joint dynamics of the spot and implied volatilities are related. Another of their drawbacks is that they are based on a specification of the spot process and they fail to take into account the fact that variance swaps (VSs) should be considered as hedge instruments too, and be endowed with their own dynamics.

This article is a natural continuation of our first one: here, we propose a model that aims at pricing both standard exotic options and general options on variance in a consistent manner, and lets us independently set requirements for:

- The dynamics of VS volatilities.
- The level of short-term forward skew.
- The correlation between the underlying and short and long VS volatilities.

The article is organised as follows. First, we set up a general framework for the dynamics of forward variance swap variances (FVs). Then we specify a dynamics for the underlying that is consistent with that of variances. In the next section, we specify a particular choice for the dynamics of forward variances and the underlying. We then focus on practical features of the model, such as the term structure of the volatility of volatility and the term structure of the skew. Then a section focuses on pricing examples: we consider a reverse cliquet, a Napoleon, an accumulator and a call on variance. The concluding section summarises our work.

Modelling variances

A VS pays at maturity $V^h_{tT} - V^T_t$, where V^h_{tT} is the annualised variance of the spot, realised over the interval $[t, T]$ and V^T_t is the implied VS variance, observed at time t for maturity T. Because VSs are statically replicable by vanilla options, V^T_t depends only on the implied volatilities seen at time t for maturity T.[1] Because of the definition of V^T_t, the VS contract has zero value at inception.

Now consider the FV V^{T_1,T_2}_t defined as:

$$V^{T_1,T_2}_t = \frac{(T_2 - t)V^{T_2}_t - (T_1 - t)V^{T_1}_t}{T_2 - T_1}$$

where $T_1, T_2 > t$.

To write a pricing equation for an option on V^{T_1,T_2}_t we first need to know the cost of entering a trade whose payout at time $t + dt$ is linear in $V^{T_1,T_2}_{t+dt} - V^{T_1,T_2}_t$. Let us buy $(T_2 - t)/(T_2 - T_1)e^{r(T_2 - t)}$ VS of maturity T_2 and sell $(T_1 - t)/(T_2 - T_1)e^{r(T_1 - t)}$ VS of maturity T_1. This is done at no cost; our profit and loss at time $t' = t + dt$ is:

$$P\,\&\,L = \frac{T_2 - t}{T_2 - T_1}\left(\frac{V^h_{tt'}(t'-t) + V^{T_2}_{t'}(T_2 - t')}{T_2 - t} - V^{T_2}_t\right)e^{r(T_2-t)}e^{-r(T_2-t')}$$

$$-\frac{T_1 - t}{T_2 - T_1}\left(\frac{V^h_{tt'}(t'-t) + V^{T_1}_{t'}(T_1 - t)}{T_1 - t} - V^{T_1}_t\right)e^{r(T_1-t)}e^{-r(T_1-t')}$$

$$= \left(V^{T_1,T_2}_{t'} - V^{T_1,T_2}_t\right)e^{-r(t'-t)} = \left(V^{T_1,T_2}_{t+dt} - V^{T_1,T_2}_t\right)(1 - r\,dt)$$

This position generates a profit and loss linear in $V^{T_1,T_2}_{t+dt} - V^{T_1,T_2}_t$ at lowest order in dt, at zero initial cost. Thus the pricing drift of any forward FV V^{T_1,T_2}_t is zero.[2]

We now specify a dynamics for the VS curve. Let us introduce $\xi^T_t = V^{T,T}_t$, the value of the variance for date T, observed at time t.

- **A one-factor model.** We are free to specify any dynamics on the $\xi^T(t)$ that complies with the requirement that $\xi^T(t)$ be drift-less. However, for practical pricing purposes, we would like to drive the dynamics of all of the $\xi^T(t)$ with a small number of factors. Here, we show how this can be done by carefully choosing the volatility function of $\xi^T(t)$.

Let us assume $\xi^T(t)$ is lognormally distributed and that its volatility is a function of $T - t$ so that the model is translationally invariant through time:

$$d\xi^T = \omega(T - t)\xi^T dU_t$$

where U_t is a Brownian motion. Let us choose the form $\omega(\tau) = \omega e^{-k_1\tau}$.

$\xi^T(t)$ can be written as:

$$\xi^T(t) = \xi^T(0)e^{\left(\omega e^{-k_1(T-t)}X_t - \frac{\omega^2}{2}e^{-2k_1(T-t)}E\left[X_t^2\right]\right)} \tag{1}$$

where X_t is an Ornstein-Ühlenbeck process:

$$X_t = \int_0^t e^{-k_1(t-u)}dU_u$$

whose dynamics reads:

$$dX_t = -k_1 X_t dt + dU_t$$
$$X_0 = 0$$

$\xi^T(t)$ is drift-less by construction. Knowing X_t, we can generate $X_{t+\delta}$ through:

$$X_{t+\delta} = e^{-k_1\delta}X_t + x_\delta$$

where x_δ is a centred Gaussian random variable such that $E[x_\delta^2] = (1 - e^{-2k_1\delta})/(2k_1)$.

Starting from known values for X_t and $E[X_t^2]$ at time t we can generate the FV curve $\xi^T(t + \delta)$ at time $t + \delta$ by using the following relationship:

[1] *As well as on how dividends are modelled and assumptions on interest rate volatility*

[2] *The drift-less nature of forward VS variances had been noticed before (see Dupire, 1996)*

$$X_{t+\delta} = e^{-k_1\delta}X_t + x_\delta$$

$$E\left[X_{t+\delta}^2\right] = e^{-2k_1\delta}E\left[X_t^2\right] + \frac{1-e^{-2k_1\delta}}{2k_1}$$

and expression (1).

Thus, by choosing an exponentially decaying form for $\omega(\tau)$ the model becomes Markovian: all $\xi^T(t)$ are functions of just one Gaussian factor X_t.

☐ **A two-factor model.** To achieve greater flexibility in the range of term structures of volatilities of volatilities that can be generated, we prefer to work with two factors. We then write:

$$d\xi^T = \omega\xi^T\left(e^{-k_1(T-t)}dU_t + \theta e^{-k_2(T-t)}dW_t\right)$$

where W_t is a Brownian motion. Its correlation with U_t is ρ. We can run through the same derivation as above. $\xi^T(t)$ now reads:

$$\xi^T(t) = \xi^T(0)\exp\Big(\omega\left[e^{-k_1(T-t)}X_t + \theta e^{-k_2(T-t)}Y_t\right] - \frac{\omega^2}{2}\Big[e^{-2k_1(T-t)}E\left[X_t^2\right] + \theta^2 e^{-2k_2(T-t)}E\left[Y_t^2\right] + 2\theta e^{-(k_1+k_2)(T-t)}E\left[X_tY_t\right]\Big]\Big) \quad (2)$$

As in the one-factor case, if X_t, Y_t, $E[X_t^2]$, $E[Y_t^2]$ and $E[X_tY_t]$ are known at time t, they can be generated at time $t + \delta$ through the following relationships:

$$X_{t+\delta} = e^{-k_1\delta}X_t + x_\delta$$
$$Y_{t+\delta} = e^{-k_2\delta}Y_t + y_\delta$$

and:

$$E\left[X_{t+\delta}^2\right] = e^{-2k_1\delta}E\left[X_t^2\right] + \frac{1-e^{-2k_1\delta}}{2k_1}$$
$$E\left[Y_{t+\delta}^2\right] = e^{-2k_2\delta}E\left[Y_t^2\right] + \frac{1-e^{-2k_2\delta}}{2k_2}$$
$$E\left[X_{t+\delta}Y_{t+\delta}\right] = e^{-(k_1+k_2)\delta}E\left[X_tY_t\right] + \rho\frac{1-e^{-(k_1+k_2)\delta}}{k_1+k_2}$$

where, in the right-hand terms, the second component is, respectively, the variance of x_δ, the variance of y_δ and the covariance of x_δ and y_δ. Starting from time $t = 0$ we can easily generate a FV curve at any future time t by simulating two Gaussian factors. We choose $k_1 > k_2$ and call X_t the short factor and Y_t the long factor.

☐ **A discrete structure.** Instead of modelling the set of all instantaneous forward variances, it may be useful to set up a tenor structure and model the dynamics of forward variances for discrete time intervals, in a way that is analogous to Libor market models.

In the fixed-income world, this is motivated by the fact that forward Libor rates are the actual underliers over which options are written. In our case, it is motivated by the fact that we want to control the skew for a given time scale.

Let us define a set of equally spaced dates $T_i = t_0 + i\Delta$, starting from t_0, today's date. We will model the dynamics of FVs defined over intervals of width Δ: define $\xi^i(t) = V_t^{t_0+i\Delta,\, t_0+(i+1)\Delta}$ for $t \le t_0 + i\Delta$. $\xi^i(t)$ is the value at time t of the FV for the interval $[t_0 + i\Delta, t_0 + (i+1)\Delta]$.

$\xi^i(t)$ is a random process until $t = t_0 + i\Delta$.When t reaches $t_0 + i\Delta$, the VS variance for time interval $[t, t + \Delta]$ is known and is equal to $\xi^i(t = t_0 + i\Delta)$.

We model the ξ^i in the same way as their continuous counterparts:

$$\xi^i(t) = \xi^i(0)\exp\Big(\omega\left[e^{-k_1(T_i-t)}X_t + \theta e^{-k_2(T_i-t)}Y_t\right] - \frac{\omega^2}{2}\Big[e^{-2k_1(T_i-t)}E\left[X_t^2\right] + \theta^2 e^{-2k_2(T_i-t)}E\left[Y_t^2\right] + 2\theta e^{-(k_1+k_2)(T_i-t)}E\left[X_tY_t\right]\Big]\Big) \quad (3)$$

where we use the same recursions as above for X_t, Y_t, $E[X_t^2]$, $E[Y_t^2]$ and $E[X_tY_t]$.

While this set-up for the dynamics of the ξ^i is reminiscent of the Libor market models used in fixed income, there are as yet no market quotes for prices of caps/floors and swaptions on forward variances, on which to calibrate volatilities and correlations for the ξ^i.

☐ **An *N*-factor model.** We may generally write:

$$\xi^i(t) = \xi^i(0)e^{\omega_i Z_t^i - \frac{\omega_i^2 t}{2}}$$

where ω_i and $\rho(Z_i, Z_j)$ are chosen at will. Later in this article, we will compare pricing results obtained in the two-factor model with those obtained in an N-factor model for which $\omega_i = \omega$, a constant, and the correlation structure of the Z^i is:

$$\rho\left(Z_i, Z_j\right) = \theta\rho_0 + (1-\theta)\beta^{|j-i|} \quad (4)$$

where $\theta, \rho_0, \beta \in [0, 1]$.

It should be noted that, when pricing an option of maturity T, in contrast with the two-factor model, the number of factors driving the dynamics of variances in the N-factor model is proportional to T, thus the pricing time will grow like T^2.

Specifying a joint dynamics for the spot

☐ **A continuous setting.** We could use the dynamics of instantaneous forward variances specified in equation (2) and write the following lognormal dynamics on the underlying:

$$dS = (r-q)Sdt + \sqrt{\xi^t(t)}SdZ_t$$

with correlations ρ_{SX} and ρ_{SY} between Z and, respectively U and W. This yields a stochastic volatility model whose differences with standard models are:

☐ It has two factors.

☐ It is calibrated by construction to the term structure of VS volatilities.

In such a model, the level of forward skew is determined by ρ_{SX}, ρ_{SY}, ρ, ω, k_1, k_2 and θ with no way of controlling it separately, just like in standard stochastic volatility models.

☐ **A discrete setting.** Here we achieve our objective of controlling the forward skew – or, in other words, the skewness of the spot process for time scale Δ – by using the discrete tenor structure defined above and the dynamics of forward variances given by expression (3).

At time $t = T_i$ the VS volatility $\hat{\sigma}_{VS}$ for maturity $T_i + \Delta$ is known. It is given by:

$$\hat{\sigma}_{VS} = \sqrt{\xi^i(t = T_i)}$$

To be able to specify the spot process over the interval $[T_i, T_i + \Delta]$, we make two more assumptions:

☐ that the spot process over the time interval $[T_i, T_i + \Delta]$ is homogeneous: the distribution of $S_{T_i+\Delta}/S_{T_i}$ does not depend on S_{T_i}. The reason for this requirement is that we want to decouple the short forward skew and the spot/volatility correlation. Imposing this condition makes the skew of maturity Δ independent on the spot level. Thus the prices of cliquets with period Δ will not depend on the level of spot/volatility correlation.

☐ that the at-the-money-forward (ATMF) skew $(d\hat{\sigma}_K/d\ln K)|_F$ for maturity $T_i + \Delta$ be a deterministic function of $\hat{\sigma}_{VS}$ or $\hat{\sigma}_{ATMF}$. In this article, we impose that it is constant or proportional to $\hat{\sigma}_{ATMF}$. Other specifications for the dependence of the ATMF skew on $\hat{\sigma}_{VS}$ or $\hat{\sigma}_{ATMF}$ are easily accommodated in our framework.

There are many processes available for fulfilling these objectives – note that we also need to correlate the spot process with that of forward variances ξ^j for $j > i$. We could use a Lévy process, especially one that has an expression in terms of a subordinated Brownian motion.[3] Here we decide to use a constant elasticity of variance (CEV) form of local volatility. Over

[3] *For example, the variance gamma and normal inverse Gaussian processes*

the interval $[T_i, T_i + \Delta]$ the dynamics of S_t reads:

$$dS = (r_t - q_t)Sdt + \sigma_0 \left(\frac{S}{S_{T_i}} \right)^{1-\beta} SdZ_t \qquad (5)$$

where $\sigma_0(\hat{\sigma}_{VS})$, $\beta(\hat{\sigma}_{VS})$ are functions of $\hat{\sigma}_{VS} = \sqrt{\xi^i(T_i)}$ calibrated so that the VS volatility of maturity $T_i + \Delta$ is $\hat{\sigma}_{VS}$ and the condition on the ATMF skew is fulfilled. r_t and q_t are, respectively, the interest rate and the repo, inclusive of dividend yield. Note that instead of – or in addition to – controlling the skew we could have controlled the convexity of the smile near the money. This would be relevant in the forex or fixed-income world. In this article, we restrict our attention to the skew.

This completely specifies our model and the pricing algorithm. We can write the corresponding pricing equation as:

$$\frac{dP}{dt} + (r_t - q_t)S\frac{dP}{dS} + \frac{\sigma\left(S_{T_{i_0}}, \xi^{i_0}, S\right)^2}{2} S^2 \frac{d^2P}{dS^2}$$
$$+ \frac{1}{2}\sum_{i,j>i_0} \rho_{ij}\omega_i\omega_j\xi^i\xi^j \frac{d^2P}{d\xi^i d\xi^j} + \sum_{i>i_0} \rho_{Si}\omega_i\sigma\left(S_{T_{i_0}}, \xi^{i_0}, S\right)S\xi^i \frac{d^2P}{dSd\xi^i} = rP$$

where $i_0(t)$ is such that $t \in [T_{i_0}, T_{i_0} + \Delta[$, ω_i is the volatility of the ξ^i and ρ_{ij} is their correlations.

In the N-factor model, $\omega_i = \omega$ and $\rho_{ij} = \rho(Z_i, Z_j)$. In the two-factor model, the dynamics of the ξ^i is driven by the processes X and Y. The pricing equation can then be written more economically as:

$$\frac{dP}{dt} + (r_t - q_t)S\frac{dP}{dS} - k_1X\frac{dP}{dX} - k_2Y\frac{dP}{dY} + \frac{\sigma(\cdots,S)^2}{2}S^2\frac{d^2P}{dS^2}$$
$$+ \frac{1}{2}\left(\frac{d^2P}{dX^2} + \frac{d^2P}{dY^2} + 2\rho\frac{d^2P}{dXdY}\right) + \sigma(\cdots,S)S\left(\rho_{SX}\frac{d^2P}{dSdX} + \rho_{SY}\frac{d^2P}{dSdY}\right) = rP$$

where ρ_{SX} and ρ_{SY} are, respectively, the correlation between Brownian motions U_t and Z_t and the correlation between W_t and Z_t. $\sigma(\cdots, S)$ is a shorthand notation for:

$$\sigma(\cdots,S) \equiv \sigma\left(S_{T_{i_0}}, \xi^{i_0}\left(X_{T_{i_0}}, Y_{T_{i_0}}\right), S\right)$$

Pricing

We now turn to using the model for pricing, focusing on the two-factor model. In what follows, we choose as time scale Δ = 1 month. By construction, the model is calibrated at time t_0 to the FV curve for all maturities $t_0 + i\Delta$. We specify, in this order:

- values for k_1, k_2, ω, ρ, θ
- a value for the ATMF skew
- values for ρ_{SX} and ρ_{SY}.

These steps are discussed in the next three sections.

Setting a dynamics for implied VS volatilities. Our aim is to price options whose price is a very non-linear function of volatility; as we roll towards the option's maturity, the maturity of the volatilities we are sensitive to shrinks as well. We thus need to be able to control the term structure of the volatilities of volatilities, be they ATMF or VS volatilities. In our framework, it is more natural to work with VS volatilities. In our model, the dynamics of VS volatilities is controlled by k_1, k_2, ω, ρ and θ. As there is currently no active market for options on forward ATM or VS volatility, these parameters cannot be calibrated on market prices. Thus their values have to be chosen so that the level and term structure of volatility of volatility are reasonably conservative when compared with historically observed volatilities of implied volatilities.[4]

Here we choose the following values:

$$\omega = 2.827, \rho = 0, \theta = 30\%, k_1 = 6(2\ months), k_2 = 0.25(4\ years) \qquad (6)$$

so that the volatility of volatility for a one-month horizon is about 120% for the one-month VS volatility, 45% for the one-year volatility and 25% for the five-year volatility. Figure 1 displays the term structure of the volatilities of VS volatilities generated by the two-factor model with a flat initial VS term structure at 20% volatility using these parameter values. We graph:

$$\frac{1}{\sqrt{\Delta t}} StDev\left[\ln\left(\frac{\sqrt{V_{\Delta t}^{\Delta t,\Delta t+\tau}}}{\sqrt{V_0^{\Delta t,\Delta t+\tau}}}\right)\right]$$

for a range of values of τ from one month to five years. We have picked Δt = 1 month. The value of ω is chosen so that, over the interval of Δt = 1 month, the volatility of the one-month VS volatility is 120%.

We also display the term structure generated by the N-factor model using the following parameters:

$$\sigma = 240\%, \theta = 40\%, \rho_0 = 5\%, \beta = 10\%$$

These values are chosen so that, for Δt = 1 month, the term structure of the two-factor model is matched. Now let us measure volatilities over a time interval of one year, instead of one month (see figure 2).

They are very different. Although both models would yield similar prices for options on VS variances observed one month from now, they would

1. Term structure of the volatilities of VS volatilities for a one-month interval

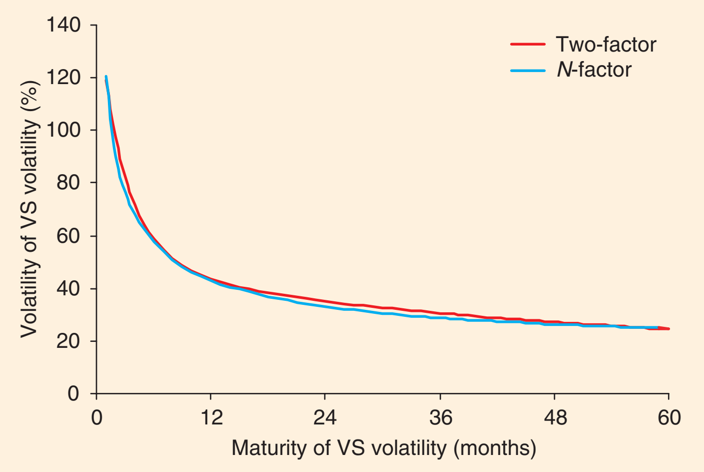

2. Term structure of the volatility of VS volatilities for a one-year interval

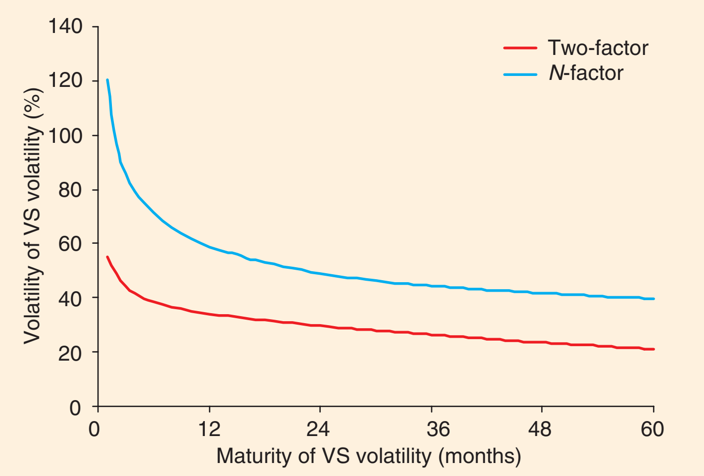

[4] *Dealers trading Napoleons and reverse cliquets typically accumulate a negative gamma position on volatility. Ideally, bid and offer term structures of volatility of volatility will be used for pricing*

3. β and σ_0 as a function of $\hat{\sigma}_{VS}$ in the case of a constant 5% skew

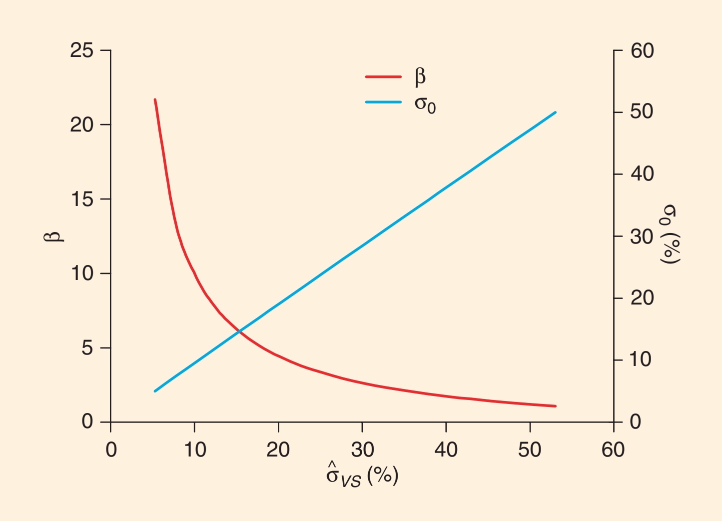

price differently options on VS variances observed in one year. In the two-factor model, volatilities of volatility will tend to decrease as the time scale over which they are measured increases, due to the mean-reverting nature of the driving processes. In the N-factor model, by contrast, they increase: this is due to the fact that forward variances are lognormal. The term structure would be unchanged if forward variances were normal.

☐ **Setting the short forward skew.** We calibrate the dependence of σ_0 and β to $\hat{\sigma}_{VS}$, so that the one-month ATMF skew has a constant value – say, 5%. We use the 95–105% skew:

$$\hat{\sigma}_{95\%}-\hat{\sigma}_{105\%}\cong-\frac{1}{10}\frac{d\hat{\sigma}_K}{d\ln K}\bigg|_F$$

instead of the local derivative $(d\hat{\sigma}_K)/(d\ln K)$. This defines the functions $\sigma_0(\hat{\sigma}_{VS})$ and $\beta(\hat{\sigma}_{VS})$. This calibration is easily done numerically; we can also use analytical approximations.[5]

If needed, individual calibration of $\sigma_0(\hat{\sigma}_{VS})$ and $\beta(\hat{\sigma}_{VS})$ can be performed for each interval $[T_i, T_i+\Delta]$. Typically the same calibration will be used for all intervals except the first one, for which a specific calibration is performed so as to match the short vanilla skew. Here we use the same calibration for all intervals. Figure 3 shows functions $\sigma_0(\hat{\sigma}_{VS})$ and $\beta(\hat{\sigma}_{VS})$ for the case of a constant 95–105% skew equal to 5%.

The level of 95–105% skew can either be selected by the trader or chosen so that the market prices of call spread cliquets of period Δ (here one month) are matched.

☐ **Setting correlations between the spot and short/long factors – the term skew.** ρ_{SX} and ρ_{SY} cannot be chosen independently, since X and Y have correlation ρ. We use the following parametrisation:

$$\rho_{SY}=\rho_{SX}\rho+\chi\sqrt{1-\rho_{SX}^2}\sqrt{1-\rho^2}$$

with $\chi\in[-1,1]$. ρ_{SX} and ρ_{SY} control both the correlation between spot and short and long VS volatilities, and the term structure of the skew of vanilla options. They can be chosen, calibrated to the market prices of call spread cliquets of a period larger than Δ, or calibrated to the vanilla skew for the maturity of the option considered. The dependence of the term skew on ρ_{SX} and ρ_{SY} is made explicit in the following section.

In the N-factor model, we need to specify correlations between the spot process and all forward variances, in a manner that is consistent with correlations of variances, a non-trivial task that we leave outside the scope of this article.

☐ **The term skew.** To shed light on how our model generates skew, we derive an approximate expression for the ATMF skew as a function of maturity, for the case of a flat term structure of VS volatilities, at order one in both ω and the skew $(d\hat{\sigma}_K/d\ln K)|_F$ at time scale Δ, which we denote $Skew_\Delta$. Given the skewness $\mathcal{S}_T$ of the distribution of $\ln(S_T/F_T)$, the ATMF skew is given at first order in $\mathcal{S}_T$ by (Backus et al, 1997):

$$Skew_T=\frac{\mathcal{S}_T}{6\sqrt{T}} \tag{7}$$

where F_T is the forward for maturity T.

Consider a maturity $T=N\Delta$ and let us calculate the second and third moments of $\ln(S_T/F_T)=\Sigma_{i=1}^N r_i$ where returns r_i are defined as:

$$r_i=\ln\left(\frac{S_{i\Delta}}{F_\Delta}\right)-\ln\left(\frac{S_{(i-1)\Delta}}{F_{(i-1)\Delta}}\right)$$

While returns are not independent, they are uncorrelated. Thus, assuming that Δ is small, so that the drift term in $E[r_i]$ is negligible with respect to the random term:

$$M_3^T=\left\langle\left(\sum_{i=1}^N r_i\right)^3\right\rangle=\sum_i\left\langle r_i^3\right\rangle+3\sum_{j>i}\left\langle r_i r_j^2\right\rangle$$

Let us work at the lowest order in Δ. To derive an expression of the third moment at order one in ω and S_Δ, we can use the following approximations:

$$r_j^2=\Delta\xi^j\left(T_j\right)$$
$$r_i=\sqrt{\xi^i\left(T_i\right)}\int_{T_i}^{T_i+\Delta}dZ_t$$

Let us denote as ξ the constant value of the VS volatilities at time zero. We get, at order one in ω:

$$M_3^T=\sum_i\left\langle r_i^3\right\rangle+3\sum_{j>i}\Delta\left\langle\sqrt{\xi^i\left(T_i\right)}\int_{T_i}^{T_i+\Delta}dZ_t\,\xi^j\left(T_j\right)\right\rangle$$

$$=\sum_i\left\langle r_i^3\right\rangle+3\sum_{j>i}\Delta\left\langle\sqrt{\xi^i\left(T_i\right)}\int_{T_i}^{T_i+\Delta}dZ_t\,\xi^j(0)\begin{pmatrix}1+\omega\int_0^{T_j}e^{-k_1\left(T_j-u\right)}dU_u\\+\theta\omega\int_0^{T_j}e^{-k_2\left(T_j-u\right)}dW_u\end{pmatrix}\right\rangle$$

$$=\sum_i\left\langle r_i^3\right\rangle$$
$$+3\sum_{j>i}\Delta\omega\xi^j(0)\sqrt{\xi^i(0)}\left\langle\int_{T_i}^{T_i+\Delta}dZ_t\int_0^{T_j}\left(e^{-k_1\left(T_j-u\right)}dU_u+\theta e^{-k_2\left(T_j-u\right)}dW_u\right)\right\rangle$$
$$=\sum_i\left\langle r_i^3\right\rangle+\rho\omega\xi^{\frac{3}{2}}\Delta^2N^2\left[\rho_{SX}\zeta\left(k_1\Delta,N\right)+\theta\rho_{SY}\zeta\left(k_2\Delta,N\right)\right]$$

where $\zeta(x,N)$ is defined by:

$$\zeta(x,N)=\left(\frac{1-e^{-x}}{x}\right)\frac{\sum_{\tau=1}^{N-1}(N-\tau)e^{-(\tau-1)x}}{N^2}$$

Since we have set the short skew to a value that is independent on the level of variance, expression (7) shows that the skewness of r_i is constant. Thus:

$$\sum_i\left\langle r_i^3\right\rangle=\mathcal{S}_\Delta\sum_i\left\langle\left(\Delta\xi^i\left(T_i\right)\right)^{\frac{3}{2}}\right\rangle=N\mathcal{S}_\Delta(\Delta\xi)^{\frac{3}{2}}$$

where $\mathcal{S}_\Delta$ is the skewness at time scale Δ. We then get:

$$M_3^T=N\mathcal{S}_\Delta(\xi\Delta)^{\frac{3}{2}}+\rho\omega\sqrt{\Delta}(\xi\Delta)^{\frac{3}{2}}N^2\left(\rho_{SX}\zeta\left(k_1\Delta,N\right)+\theta\rho_{SY}\zeta\left(k_2\Delta,N\right)\right)$$

[5] *See, for example, Zhou (2003)*

At order zero in $\mathcal{S}_\Delta$ and ω:

$$M_2^T = \left\langle \left(\sum_{i=1}^{n} r_i \right)^2 \right\rangle = N\xi\Delta$$

hence the following expression for $\mathcal{S}_T = M_3^T/(M_2^T)^{\frac{3}{2}}$:

$$\mathcal{S}_T = \frac{\mathcal{S}_\Delta}{\sqrt{N}} + \sqrt{N}\omega\sqrt{\Delta}\left[\rho_{SX}\zeta(k_1\Delta, N) + \theta\rho_{SY}\zeta(k_2\Delta, N)\right]$$

Using equation (7) again, we finally get the expression of $Skew_{N\Delta}$ at order one in $Skew_\Delta$ and ω:

$$Skew_{N\Delta} = \frac{Skew_\Delta}{N} + \frac{\omega}{2}\left[\rho_{SX}\zeta(k_1\Delta, N) + \theta\rho_{SY}\zeta(k_2\Delta, N)\right] \qquad (8)$$

This expression is instructive as it makes apparent how much of the skew of maturity T is contributed on the one hand by the intrinsic skewness of the spot process at time scale Δ and on the other hand by the spot/volatility correlation.

When $\omega = 0$, the skew decays as $1/T$, as expected for a process of independent increments. The fact that volatility is stochastic and correlated with the spot alters this behaviour. Inspection of the definition of function ζ in equation (8) shows that for $N\Delta >> 1/k_1, 1/k_2$, $\zeta(x, N) \propto 1/N$, so that $Skew_{N\Delta} \propto 1/N$, again what we would expect.

Equation (8) also shows how ρ_{SX} and ρ_{SY} can naturally be used to control the term structure of the skew.

Figure 4 shows how the approximate skew in equation (8) compares with the actual skew. We have chosen the following values: $\Delta = 1$ month, the one-month 95%/105% skew is 5%, $\omega = 2.827$, $\rho = 0$, $\theta = 30\%$, $k_1 = 6$, and $k_2 = 0.25$. The spot/volatility correlations are: $\rho_{SX} = -70\%$, $\rho_{SY} = -35.7\%$ ($\chi = -50\%$). Even though ω and $Skew_\Delta$ are both large, the agreement is very satisfactory.

The two contributions to $Skew_{N\Delta}$ in equation (8) are illustrated in figure 5. 'Intrinsic' denotes the first piece and 'spot/vol correlation' denotes the second piece in equation (8). While the contribution of $Skew_\Delta$ to $Skew_{N\Delta}$ is monotonically decreasing, the contribution of the spot/volatility correlation is not, as it starts from zero at short time scales. Depending on the relative magnitude of both terms, the term structure of the skew can be non-monotonic.

We have derived expression (8) for the case of a flat VS term structure but the general case poses no particular difficulty.

Pricing examples

Here, we use our model to price a reverse cliquet, a Napoleon, an accumulator and a call on realised variance, and analyse the relative contribution of forward skew, volatility-of-volatility and spot/volatility correlation effects to prices. We use zero interest rates and dividend yield.

For the sake of comparing prices, we need to specify how we calibrate model parameters. While it is natural to calibrate to the vanilla smile when pricing options that can be reasonably hedged with a static position in vanilla options, it is more natural to calibrate to call spread cliquets and ATM cliquets when pricing Napoleons and reverse cliquets, which have a high sensitivity to forward volatility and skew.

These products are also very sensitive to volatility of volatility. They are usually designed so that their price at inception is small but increases significantly if implied volatility decreases.[6] As there is as yet no active market for options on variance, we use the volatility-of-volatility parameter values listed in (6).

Unless forward skew is turned off, the constant 95–105% one-month skew is calibrated so that the price of a three-year 95–105% one-month call spread cliquet has a constant value, equal to its price when volatility of volatility is turned off and the one-month 95–105% skew is 5%, which is equal to 191.6%.

In all cases the level of the flat VS volatility has been calibrated so that the implied volatility of the three-year one-month ATM cliquet is 20%.

The values for ρ_{SX} and ρ_{SY} are $\rho_{SX} = -70\%$, $\rho_{SY} = -35.7\%$ ($\chi = -50\%$).

4. The 95–105% skew as a function of maturity

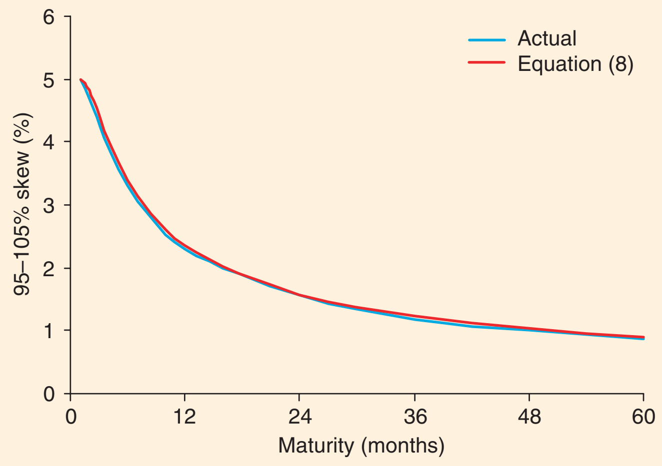

5. The two contributions to the 95–105% skew in equation (8)

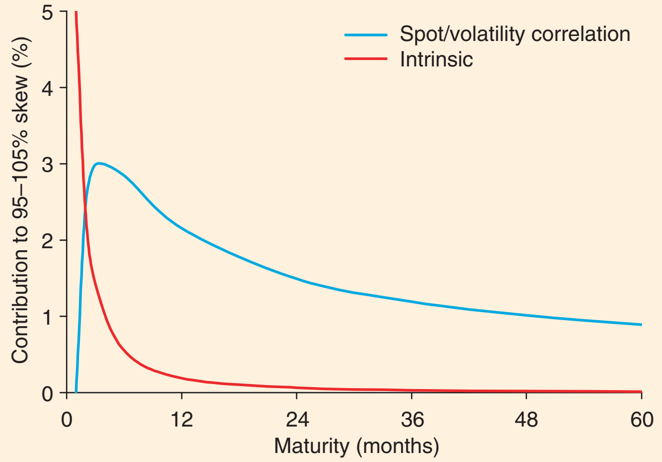

A. Model prices

Model	Reverse cliquet	Napoleon	Accumulator
Black-Scholes	0.25%	2.10%	1.90%
With forward skew	0.56%	2.13%	4.32%
With volatility of volatility	2.92%	4.71%	1.90%
Full	3.81%	4.45%	5.06%

B. Model prices

Model	Reverse cliquet	Napoleon	Accumulator
Full – correlations halved	3.10%	4.01%	5.04%
Full – proportional skew	3.05%	4.30%	4.15%

[6] *See figure 1 in Bergomi (2004)*

6. Implied volatility of a call option on realised variance as a function of maturity in the two-factor model

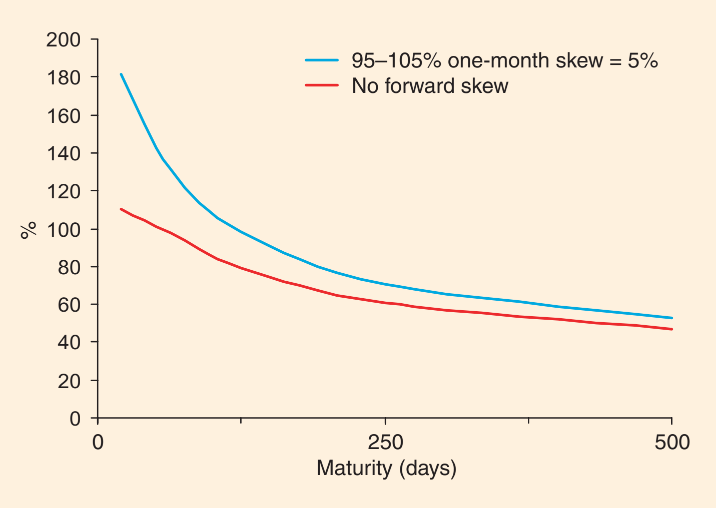

7. Implied volatility of a call option on realised variance as a function of maturity in the two-factor and N-factor models

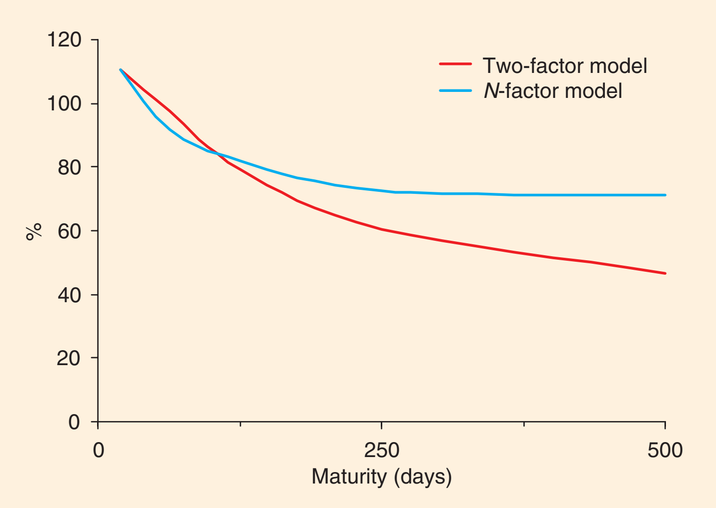

The corresponding term skew is that of figure 4.

In addition to the Black-Scholes price, we calculate three other prices by switching on either the one-month forward skew ($\hat{\sigma}_{95\%} - \hat{\sigma}_{105\%} \neq 0$, $\omega = 0$) or the volatility of volatility ($\hat{\sigma}_{95\%} - \hat{\sigma}_{105\%} = 0$, $\omega \neq 0$) or both (full). These prices are listed in table A. We give the definition of each product and comment on pricing results in the following paragraphs.

□ **Reverse cliquet.** Here we consider a globally floored locally capped cliquet, which pays once at maturity:

$$\max\left(0, C + \sum_{i=1}^{N} r_i^-\right)$$

The maturity is three years, returns r_i are observed on a monthly basis ($N = 36$), $r_i^- = \min(r_i, 0)$ and the value of the coupon is $C = 50\%$.

Notice that corrections to the Black-Scholes price are by no means small, the contribution of volatility of volatility being the largest. The fact that volatility of volatility makes the reverse cliquet more expensive is expected: this option, as well as the Napoleon, is in essence a put on volatility (Bergomi, 2004).

To understand why forward skew increases the price, consider first that, in the four cases listed above, $E[\Sigma_{i=1}^{N} r_i^-]$ is constant, by calibration on the ATM cliquet. Next consider the last period of the reverse cliquet. The final payout is a function of the final return; it is a call spread whose low and high strikes are, respectively, $-C + \Sigma_{i<N} r_i^-$ – if it is negative – and zero. When forward skew is turned on, the implied volatility of the ATM strike is unchanged, by calibration, while the implied volatility for lower strikes increases, making the call spread more expensive. The same argument holds for returns prior to the last one.

□ **Napoleon.** The maturity is still three years and the option pays at the end of each year a coupon given by:

$$\max\left(0, C + \min_i r_i\right)$$

where r_i are the 12 monthly returns observed each year. Here we use $C = 8\%$.

Again, we notice that volatility of volatility accounts for most of the price. Forward skew seems to have no sizeable impact, though this is not generic; its magnitude and sign depend on the coupon size. While the payout is still a call spread as a function of the final return, both strikes lie below the money. Also, in contrast to the case of the reverse cliquet, $E[\min_i r_i]$ is not constant in the four cases considered.

□ **Accumulator.** The maturity is again three years with one final payout, given as a function of the 36 monthly returns r_i by:

$$\max\left(0, \sum_{i=1}^{N} \max\left(\min\left(r_i, cap\right), floor\right)\right)$$

where floor = –1% and cap = 1% – a standard product.

The largest contribution comes from forward skew. Notice that switching on the volatility of volatility in the case when there is no forward skew has no material impact on the price while it does when forward skew is switched on. To understand this, observe that, in Black-Scholes, when both strikes are priced with the same volatility, a 99–101% one-month call spread has negligible vega. However, when the call spread is priced with a downward sloping skew, it acquires positive convexity with respect to volatility shifts.

□ **Effect of spot/volatility correlation – decoupling of the short forward skew.** In standard stochastic volatility models, changing the spot/volatility correlation changes the forward skew and thus the price of cliquets. In this model, because of the specification chosen for the spot dynamics in equation (5), changing the spot/volatility correlation does not change the value of one-month cliquets. It only alters the term skew.

Prices quoted above have been calculated using $\rho_{SX} = -70\%$ and $\rho_{SY} = -35.7\%$. Figure 4 shows that, with these values the three-year 95–105% skew is 1.25%.

Let us now halve the spot/volatility correlation: $\rho_{SX} = -35\%$ and $\rho_{SY} = -18\%$ ($\chi = -19.2\%$). The three-year 95–105% skew is now 0.75% – almost halved. The implied volatility of the three-year cliquet of one-month ATM calls remains 20% and the price of a 95–105% one-month call spread cliquet is unchanged, at 191.6%. The new prices appear on the first line of table B. The difference with prices on the fourth line of table A measures the impact of the term skew, all else – in particular cliquet prices – being kept constant. The fact that prices decrease when the spot/volatility correlation is less negative is in line with the shape of the Black-Scholes vega as a function of the spot value.[7]

□ **Making other assumptions on the short skew.** Here, we want to highlight how a different model for the short skew alters prices, using the three payout examples studied above. We now calibrate functions $\sigma_0(\hat{\sigma}_{VS})$ and $\beta(\hat{\sigma}_{VS})$ so that, instead of being constant, the 95–105% skew for ma-

[7] *See figure 1 in Bergomi (2004)*

turity Δ is proportional to the ATMF volatility for maturity Δ.

We have calibrated the proportionality coefficient so that the three-year cliquet of one-month 95–105% call spreads has the same value as before. The flat VS volatility is still chosen so that the implied volatility of the three-year cliquet of a one-month ATM call is 20%. Prices are listed on the second line of table B.

The accumulator is now sizeably cheaper. One can check that, in Black-Scholes, the value of a symmetrical call spread as a function of ATM volatility, when the skew is kept proportional to the ATM volatility, is almost a linear function of volatility – in contrast with the case when volatilities are shifted in parallel fashion, where it is a convex function of volatility – thus suggesting why volatility of volatility has much less impact than in the constant skew case.

☐ **Option on realised variance.** Here we consider a call option that pays at maturity:

$$\frac{1}{2\hat{\sigma}_K}\max\left(\sigma_h^2-\hat{\sigma}_K^2,0\right)$$

where volatility $\hat{\sigma}_K$ is the strike and σ_h^2 is the annualised variance measured using daily log returns. We assume there are 250 daily observations in a year, equally spaced. The variance of the distribution of σ_h^2 has two sources:

☐ The dynamics of VS variances.

☐ The fact that observations are discrete. In the case where VS variances are static, it is the only contribution and it is determined by the distribution of spot returns, in particular its kurtosis, which depends on assumptions made for the short-maturity smile – in our context, the value of β. In the general case, it affects short-maturity options most noticeably.

Prices are expressed as implied volatilities calculated with the Black-Scholes formula with zero rate and repo. The underlying is the VS variance for the maturity of the option, whose initial value is given by the VS term structure observed at the pricing date.

In our model, daily returns are generated by the stochastic local volatility function form in equation (5). Their conditional kurtosis is a function of β, a parameter we use to control the short-term skew. The prices of options on variance will thus depend on assumptions we make for the skew at time scale Δ. Figure 6 shows implied volatilities of call options on variance, using a flat term structure of VS volatilities at 20%, the same correlations as in the examples above ($\rho_{SX} = -70\%$, $\rho_{SY} = -35.7\%$), for the two cases $\hat{\sigma}_{95\%} - \hat{\sigma}_{105\%} = 5\%$ and $\hat{\sigma}_{95\%} - \hat{\sigma}_{105\%} = 0\%$.

Figure 6 illustrates how assumptions for the forward skew affect the distribution of returns significantly and thus the price of options on variance, mostly for short options. The shortest maturity in the graph corresponds to options with a maturity of one month (20 days). Since we have taken $\Delta = 1$ month, the distribution of σ_h^2 does not depend on the dynamics of variances ξ^i – it only depends on β.

Note that, in our model, VS volatilities for maturities shorter than Δ are not frozen. Instead of being driven by equation (3), their dynamics is set by the value of β.

☐ **Using the N-factor model.** It is instructive to compare prices of options on realised variance generated by the two-factor and N-factor models. As figures 1 and 2 illustrate, even though the dynamics of VS volatilities in both models are similar for the short term, they become different for longer horizons.

Here we price the same option on variance considered above using the N-factor model of forward variances. Parameter values for the dynamics of forward variances are the same as those used in figures 1 and 2. We have taken no forward skew ($\hat{\sigma}_{95\%} - \hat{\sigma}_{105\%} = 0$). Also, to make prices comparable with those obtained in the two-factor model, we have taken zero correlation between spot and forward variances. Implied volatilities for both models are shown in figure 7.

Because $\Delta = 1$ month, for the shortest maturity considered – 20 days – the implied volatilities for both models coincide. For longer maturities, the fact that implied volatilities are higher in the N-factor model is in agreement with figure 2, which shows that, for longer horizons, the volatilities of forward variances in the N-factor model are larger than in the two-factor model.

Finally, in addition to the effects discussed above, prices of calls on variance have to be adjusted to take into account bid/offer spreads on the VS hedge. These can be approximately included by shifting the level of volatility of volatility (Leland, 1985).

Conclusion

We have proposed a new model that, in contrast to popular stochastic volatility and jump/Lévy models, gives us the flexibility to independently control:

☐ The term structure of the volatility of volatility.

☐ The short-term skew.

☐ The correlation of spot and volatilities.

This model consistently prices general exotic options and options on variance or volatility. We achieve this by choosing a time scale Δ, specifying how the forward skew for maturity Δ depends on the level of volatility, and modelling the dynamics of VS variances. The model can be simultaneously calibrated on the short vanilla skew, the long vanilla skew, an ATM cliquet and a call spread cliquet of period Δ, while letting us specify freely the dynamics of forward variances as well as how the short forward skew depends on volatility.

By directly controlling the short forward skew we are able to adjust the amount of skewness of the distribution of $\ln(S_T)$ generated on the one hand by the intrinsic skewness of the process at short time scales, and the spot/volatility correlation on the other hand. Handling this issue appropriately is, in our view, an essential task in the design of models that accurately capture the three effects mentioned above.

Even though the choice of the time scale Δ is natural for many payouts – for example, Napoleons and reverse cliquets – it is more arbitrary for other options, for example, options on variance. It would be useful to have scaling relationships relating parameter sets for different values of Δ so that some model features remain unchanged, for example, the skew of vanillas or the implied forward skew of cliquets. This is left for future work.

As of today there are no market prices for caps/floors/swaptions on forward variances. Choosing a value for the parameters governing the dynamics of forward variances is thus a trading decision. It is the hope of the author that the liquidity of options on volatility and variance increases so that we will soon be able to trade the smile of the volatility of volatility! ■

Lorenzo Bergomi is head of quantitative research in the equity derivatives department at Société Générale. He wishes to thank members of his team for helpful suggestions. Email: lorenzo.bergomi@sgcib.com

REFERENCES

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

Bergomi L, 2004
Smile dynamics
Risk September, pages 117–123

Dupire B, 1996
A unified theory of volatility
Unpublished

Leland H, 1985
Option replication with transaction costs
Journal of Finance 40(5), pages 1,283–1,301

Zhou F, 2003
Black smirks
Risk January, pages 87–91