The beta stochastic volatility model

Local stochastic volatility models combine perfect calibration at time zero with realistic price dynamics. But traditional methods tend to underestimate the forward skew, and mis-price exotics such as cliquets as a result. ***Piotr Karasinski*** *and* ***Artur Sepp*** *introduce a new model that uses the sensitivity of the at-the-money implied volatility to spot as a key input to ensure forward skew is better fit and exotics are correctly priced*

The traditional approach to stochastic volatility (SV) modelling begins with the specification of an SV process, typically on the grounds of its analytical tractability (see, for example, Heston, 1993). Then, after a closed-form solution for vanilla options has been derived and implemented, the parameters of the SV process are calibrated to the implied volatility surface of vanilla options using complicated non-linear optimisation methods. The drawback of this approach for business applications is that the calibrated parameters are unstable from day to day irrespective of changes in the implied volatility surface such as in the at-the-money (ATM) implied volatility, skew or convexity. Also, key variables, such as forward skews and volatility of volatility, are outputs in these models and cannot be tweaked to reproduce market observables or proprietary views.

The problem of modelling forward skews was relaxed by the introduction of local stochastic volatility (LSV) models (Lipton, 2002), which combine local volatility to fit the vanilla surface and the SV dynamics to model forward skews. The advantage is that as the vanilla surface is matched by construction for any set of SV dynamics, the modeller has the freedom to specify parameters to match either historical data or exotic options. The typical driver for the SV part is chosen to be either the square root process, as in the Heston model, or an exponential Ornstein-Uhlenbeck process (Bergomi, 2004).

While LSV models became popular in foreign exchange markets, in equity markets they are less popular. One disadvantage of LSV models is their tendency to underprice volatility products such as cliquets because they imply high convexity for the forward smile and as a result overprice caps and underprice floors of forward-starting products. This can be relaxed by introducing jumps in the model dynamics (Sepp, 2011) at the expense of adding extra parameters, one for the jump intensity and two for jump sizes in the price and volatility. Another disadvantage is the difficulty in interpreting the SV parameters such as equity-volatility correlation, volatility of volatility and mean-reversion rate, even before including jumps, and their impact on forward skews.

In this article, we introduce a stochastic volatility model. The key parameter of the model, β, can be naturally interpreted as the rate of change in the short-term at-the-money (ATM) volatility given change in the spot price and easily implied from historical and current market data. The remaining two parameters – the idiosyncratic volatility of volatility and the mean-reversion rate – have less impact on forward skews and also can be estimated from historical data without the need to apply time-expensive and obscure non-linear fits. The beta stochastic volatility for the logarithm of the volatility was first implemented by Bardhan & Karasinski (1993). A similar model, but in the discrete time setting, was developed by Langnau (2004). Here we describe the model dynamics incorporating the local volatility with calibration that can be implemented using conventional partial differential equation (PDE) methods. In the limiting case of the deterministic local volatility, we derive an accurate approximation for prices of vanilla options in this model. Finally, we provide illustrations of model calibration and its implications.

The dynamics

We consider two specifications of model dynamics for the underlying price $S(t)$ and its instantaneous stochastic volatility. The first is an econometric version applied for model analysis using time series, and the second is a pricing version applied for risk-neutral valuation purposes in the context of local stochastic volatility models.

The econometric model is expressed in terms of the stochastic volatility process $V(t)$:

$$\begin{aligned} \frac{dS(t)}{S(t)} &= \hat{\mu}(t)dt + V(t)dW^{(0)}, \quad S(0) = S \\ dV(t) &= \hat{\kappa}\left(\hat{\theta} - V(t)\right)dt + \hat{\beta}V(t)dW^{(0)} + \hat{\varepsilon}dW^{(1)}, \quad V(0) = V_0 \end{aligned} \tag{1}$$

where $\hat{\beta} < 0$ is the rate of change in the volatility corresponding to change in the spot price, $\hat{\varepsilon}$ is the idiosyncratic volatility of volatility, $\hat{\kappa}$ is the mean-reversion rate, $\hat{\theta}$ is the mean level of volatility, $W^{(0)}(t)$ and $W^{(1)}(t)$ are two Brownian motions with $d\langle W^{(0)}, W^{(1)}\rangle_t = 0$ and $\hat{\mu}(t)$ is the drift.

The pricing model is expressed in terms of the normalised volatility factor $Y(t)$:

$$\begin{aligned} \frac{dS(t)}{S(t)} &= \mu(t)dt + \sigma\left(1 + Y(t)\right)dW^{(0)}, \quad S(0) = S \\ dY(t) &= -\kappa Y(t)dt + \beta\sigma\left(1 + Y(t)\right)dW^{(0)} + \varepsilon dW^{(1)}, \quad Y(0) = 0 \end{aligned} \tag{2}$$

where σ is the overall level of volatility (we assume that σ is set to either constant volatility σ_{CV}, deterministic volatility $\sigma_{DV}(t)$ or local stochastic volatility $\sigma_{LSV}(t, S)$), $W^{(0)}(t)$ and $W^{(1)}(t)$ are two Brownian motions with $d\langle W^{(0)}, W^{(1)}\rangle_t = 0$ and $\mu(t)$ is the risk-neutral drift.

For constant volatility σ_{CV}, the following relationship holds between parameters of the econometric and pricing models:

$$Y(t) = \frac{V(t)}{\sigma_{CV}} - 1, \quad \beta = \frac{\hat{\beta}}{\sigma_{CV}}, \quad \varepsilon = \frac{\hat{\varepsilon}}{\sigma_{CV}}, \quad \kappa = \hat{\kappa} \tag{3}$$

We note that in this case the diffusion terms in the dynamics (2) are linear in $Y(t)$ so a unique global solution exists for this pair of stochastic differential equations, unlike some conventional exponential Ornstein-Uhlenbeck models, which is important for

1 Scatter plots

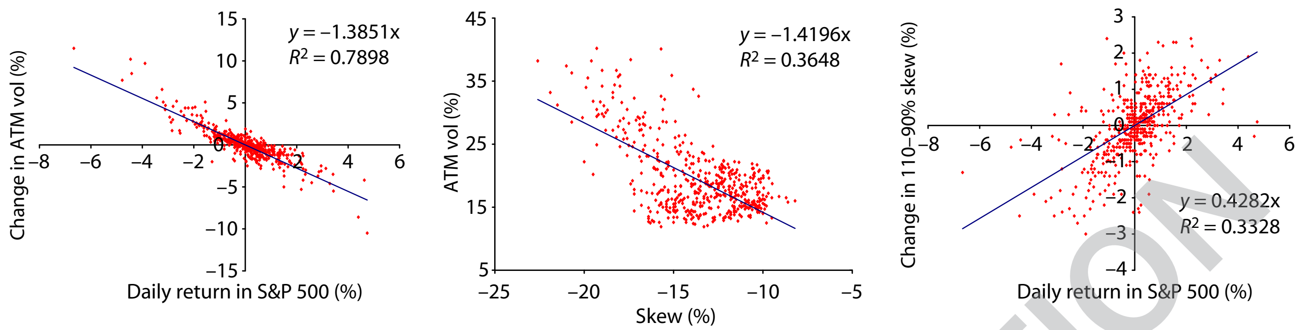

Left: the scatter plot of daily changes in one-month ATM volatility against daily returns in the S&P 500 index; centre: the scatter plot of one-month 110–90% skew against the one-month ATM volatility; right: the scatter plot of daily changes in one-month 110–90% skew against daily returns in the S&P 500 index using market data from October 2009 to October 2011

the stability of Monte Carlo and PDE methods.

Intuition. To get an intuition of the dynamics (1) of the econometric model, suppose there is no mean-reversion or idiosyncratic volatility of volatility, taking $\hat{\kappa} = \hat{\varepsilon} = 0$ and $\hat{\mu}(t) = 0$. Assume that the dynamics of the ATM implied volatility for an infinitesimal maturity is specified by $\sigma_{ATM}(t) = V(t)$, so that using (1) we obtain:

$$d\sigma_{ATM}(t) = \hat{\beta}\frac{dS(t)}{S(t)} \tag{4}$$

As a result, the dynamics (4) can be interpreted as a regression model. As $\hat{\beta} < 0$, if the spot price decreases and vice versa, recreating the leverage effect observed in the market.

The coefficient $\hat{\beta}$ can be estimated using the time series for $\sigma_{ATM}(t)$ and $S(t)$ by means of standard regression methods as we describe below. Empirically, there is strong correlation between changes in $\sigma_{ATM}(t)$ and $S(t)$. In figure 1 (left), we show the scatter plot daily changes in one-month ATM volatility corresponding to daily returns in the S&P 500 index using market data from October 2009–October 2011, and the corresponding regression model. We see that the regression model explains about 80% of daily variations in $\sigma_{ATM}(t)$. The rest of the variation is modelled by the idiosyncratic volatility of volatility $\hat{\varepsilon}$. Empirically, the residual is independent of $S(t)$ so we can safely assume the correlation between the two Brownian motions in the model dynamics (2) is zero to reduce the number of model parameters.

Comparison with the SABR and other stochastic volatility models. Now we consider the dynamics of the pricing model (2) again with no mean-reversion and idiosyncratic volatility of volatility, $\kappa = \varepsilon = 0$. Using notation as in the stochastic alpha-beta-rho (SABR) model of Hagan *et al* (2002) with $\beta = 1$, where β is the constant elasticity of volatility exponent, we obtain the following specification of SABR model parameters in terms of the beta stochastic volatility model:

$$\hat{a} = \sigma_{CV}\left(1 + Y(t)\right), \quad \rho = -1, \quad \nu = -\beta\sigma_{CV}, \quad \alpha = \sigma_{CV} \tag{5}$$

In particular, using formula (3.1a) of Hagan *et al* (2002) for a short maturity and small log-moneyness k, $k = \ln(K/S_0)$, we obtain the following relationship for the Black-Scholes-Merton implied volatility $\sigma_{IMP}(k)$:

$$\sigma_{IMP}(k) = \sigma_{CV} + \frac{1}{2}\sigma_{CV}\beta k - \frac{1}{12}\sigma_{CV}(\beta k)^2 \tag{6}$$

The main similarity between the beta SV model and the SABR model is that both assume the lognormal process for the instantaneous volatility (in the restricted case (5)). In general, the volatility process in the dynamics (2) is a mixture of lognormal and normal components.

The difference between the two models is that, in the lognormal case of the SABR model, the covariance between spot and volatility moves and the variance of idiosyncratic modes in the volatility are modelled through the same set of parameters, ν, $\hat{a}$ and ρ. In the beta SV model, these are modelled through independent parameters so we have more freedom to model the skew and curvature of the implied volatility.

Importantly, the key difference between the beta SV model and the SABR model is that the latter is mostly used for the parameterisation of the implied volatility surface, while to use a stochastic volatility model for pricing path-dependent options we require a mean-reversion. The intuition behind the mean-reversion is that when the ATM volatility is small it will tend to rise, while when it is too large it will tend to decline. Thus, over long time horizons (say two years or three years), the realised volatility of volatility is expected to stay in a range and not increase with maturity time, as implied by models with no mean-reversion. The concept of mean-reversion is related to the existence of the steady-state distribution and the ergodicity of the process (the long-term distribution does not depend on maturity time). The steady-state distribution of the volatility does not exist in the SABR model. However, it does exist in the beta SV model (see equation (20)).

Interpretation of parameter β. Following the same assumptions as in the previous section, we use formula (6) as an alternative way of calibrating the parameter β from the implied market volatility $\sigma_{IMP}(k)$ and skew:

$$\beta = \frac{Skew}{\sigma_{IMP}(0)}, \quad Skew = \frac{\sigma_{IMP}(k_+) - \sigma_{IMP}(k_-)}{\frac{1}{2}(k_+ - k_-)} \tag{7}$$

where k_+ and k_- are the log-strikes for out-of-the-money (OTM) calls and puts typically 10% either side, respectively. We find the model implies the following approximate relationship:

$$Skew(t) = \beta\sigma_{ATM}(t) \tag{8}$$

This relationship is observed empirically. In figure 1 (centre), we show the scatter plot of one-month 110–90% skew (unscaled) against one-month ATM volatility. The relationship is relatively strong. In particular, high ATM volatilities are associated with steep skews.

A. Market and model historical data – econometric model

Year	SP500	Vix	R SP500	R Vix	St SP500	St Vix	Beta	Kappa	St Resid	R2	Correl	Skew	State
1990	334.60	23.08%	–8.46%	9.42%	15.95%	26.97%	–1.00	14.66	21.17%	39%	–60%	–10.05%	Risk Off
1991	376.55	18.35%	25.57%	–7.67%	14.27%	19.39%	–0.74	10.12	15.82%	33%	–57%	–7.38%	Risk On
1992	415.82	15.41%	4.72%	–5.59%	9.72%	11.71%	–0.60	15.99	9.96%	28%	–50%	–5.98%	Risk On
1993	451.73	12.68%	7.05%	–0.79%	8.62%	11.45%	–0.64	31.04	9.57%	30%	–50%	–6.38%	Risk On
1994	460.39	13.94%	–0.89%	1.68%	9.84%	15.04%	–1.07	23.54	10.15%	55%	–71%	–10.71%	Range
1995	542.36	12.41%	30.48%	–2.06%	7.83%	9.94%	–0.50	54.27	8.46%	28%	–40%	–4.98%	Risk On
1996	670.95	16.50%	17.87%	8.95%	11.82%	15.71%	–0.87	30.78	11.06%	51%	–67%	–8.71%	Risk On
1997	874.37	22.39%	29.65%	2.28%	18.16%	19.69%	–0.75	6.86	13.79%	51%	–71%	–7.46%	Risk On
1998	1,086.51	25.61%	25.14%	2.75%	20.29%	29.63%	–1.24	5.57	15.06%	74%	–86%	–12.37%	Risk On
1999	1,328.23	24.35%	18.60%	–1.96%	18.10%	21.85%	–0.97	17.68	12.08%	69%	–82%	–9.65%	Risk On
2000	1,426.54	23.32%	–10.08%	5.78%	22.38%	21.73%	–0.76	11.70	12.98%	64%	–79%	–7.62%	Risk Off
2001	1,193.66	25.74%	–8.32%	–7.28%	21.21%	23.98%	–0.93	4.68	13.76%	67%	–82%	–9.27%	Risk Off
2002	992.96	27.29%	–20.49%	2.68%	26.24%	27.47%	–0.87	2.81	14.79%	71%	–84%	–8.74%	Risk Off
2003	966.02	21.95%	21.24%	–7.17%	16.76%	13.96%	–0.54	3.54	10.49%	44%	–65%	–5.43%	Risk On
2004	1,131.02	15.46%	8.72%	–4.14%	11.12%	12.76%	–0.86	16.44	8.02%	61%	–76%	–8.64%	Risk On
2005	1,207.49	12.80%	5.94%	–2.94%	10.38%	11.34%	–0.90	13.91	6.16%	71%	–83%	–8.96%	Risk On
2006	1,311.05	12.81%	11.51%	0.90%	9.88%	14.18%	–1.13	14.91	8.26%	66%	–80%	–11.28%	Risk On
2007	1,477.31	17.58%	3.42%	11.13%	16.02%	25.77%	–1.38	6.66	12.69%	76%	–87%	–13.81%	Range
2008	1,218.01	32.76%	–35.54%	16.02%	41.17%	56.77%	–1.20	2.28	27.54%	77%	–87%	–11.98%	Risk Off
2009	948.84	31.40%	23.23%	–19.15%	27.15%	32.27%	–0.93	5.60	19.67%	63%	–79%	–9.35%	Risk On
2010	1,140.52	22.54%	13.18%	–2.43%	18.02%	31.93%	–1.46	15.12	17.15%	71%	–83%	–14.56%	Risk On
2011	1,267.66	24.22%	3.12%	5.36%	23.30%	39.92%	–1.51	4.01	18.48%	79%	–89%	–15.06%	Range

In the Heston SV model, the short-term skew is approximately given by $\rho\varepsilon/(2\sqrt{V(0)})$, where ρ is the correlation between the variance V and log-returns, and ε is the volatility of variance (Bergomi, 2004). So the Heston model implies the inverse relationship between short-term ATM volatility and skew. As a result, through calibration to steep skews and high ATM volatilities, the Heston model implies very high volatility of variance ε (see figure 2 in Bergomi, 2004), as $\rho \approx -1$ and the only way for the model to fit the skew is by increasing ε. By contrast, in the beta SV model, idiosyncratic volatility of volatility ε is stable over different market conditions.

Finally, (4) and (8) imply the following approximate dynamics for the short-term skew in the econometric model:

$$dSkew(t) = \hat{\beta}^2 \frac{dS(t)}{S(t)} \qquad (9)$$

So the model produces explicit dependence between changes in the skew given changes in the spot price, unlike conventional SV models. The dependence between changes in the skew and spot returns is observed empirically. In figure 1 (right), we show the scatter plot daily changes in one-month 110–90% skew corresponding to daily changes in the S&P 500 index, and the corresponding regression model. Although it is less strong than that of the ATM volatility, the corresponding regression model explains 33% of daily variations in the skew. The slope is positive as negative changes in the S&P 500 increase hedging demand for OTM puts, thus increasing the skew.

Time series estimation

First we discuss how to estimate model parameters $\hat{\beta}$, $\hat{\kappa}$ and $\hat{\varepsilon}$ and examine the time series of these parameters. We use the time series of daily returns $R(t_i) = (S(t_i)/S(t_{i-1})) - 1$, on the S&P 500 index and the Vix, which we use as a proxy for one-month ATM implied volatilities, $\sigma_{ATM}(t_i)$. We assume that $V(t_i)$ is proxied by the ATM volatility $V(t_i) = \sigma_{ATM}(t_i) \equiv \mathrm{Vix}(t_i)$.

The continuous time dynamics (2) can be approximated in discrete time using the following regression:

$$V(t_i) - V(t_{i-1}) = \hat{\beta} R(t_i) + \eta\left(\bar{\sigma}_{ATM} - V(t_{i-1})\right) + \vartheta(t_i) \qquad (10)$$

where $\bar{\sigma}_{ATM}$ is the average of the ATM volatilities, $\eta \equiv \hat{\kappa}dt$ with $dt = 1/252$ and $\vartheta(t_i)$ are independent normally distributed residuals. The standard deviation of residuals, $\vartheta(t_i)$, is used to estimate $\hat{\varepsilon}$, after applying the annualisation factor 252.

For illustration, we use the time series of the S&P 500 index and Vix from January 1990–January 2012. For each year in the sample, we apply the regression model (10) for daily returns to estimate the model parameters using data corresponding to any given year. Our results are reported in table A. Here, we use the following notation: 'SP500' and 'Vix' denote averages of the S&P 500 and Vix closing levels, respectively, during the given year. A prefix 'R' or 'St' stands for the average daily return and standard deviation, both annualised, on the S&P 500 index (arithmetic return) and the Vix (absolute return), respectively. 'Beta' and 'Kappa' are estimates for the parameters $\hat{\beta}$ and $\hat{\kappa}$. 'St Resid' stands for standard deviation of residuals, annualised. 'Correl' denotes correlation between returns on the S&P 500 index and Vix, and 'Skew' is the implied skew calculated as $10\%\hat{\beta}\bar{\sigma}_{ATM}$. Finally, 'State' stands for the market regime with three states: 'Risk Off' and 'Risk On' when average annual returns on the S&P 500 index are negative and positive respectively, and 'Range' when returns are within 3%.

We see that the Vix increases during Risk Off years and declines during Risk On years, while the skew typically declines right before Risk Off years or during Range years, when the demand for protective OTM puts increases relative to ATM options. In contrast, the Vix increases during Risk Off years, when demand for both ATM puts and calls increases relative for OTM puts, as the latter become prohibitively expensive given high levels of ATM volatility.

We see that $\hat{\kappa}$ increases during Risk On years as any shock in volatility dissipates quickly driven by 'buy the dip' mentality. But it is very low right after Risk Off or Range years, implying that the volatility declines very slowly from elevated levels given the prevalent risk-aversion. Similarly, $\hat{\beta}$ increases during Risk On years for the same reasons. The standard deviations increase during Risk Off years. Remarkably, the idiosyncratic volatility $\hat{\varepsilon}$ is relatively stable, ranging from 10–20%, in contrast with the volatility of volatility in the Heston model, which is known to be highly unstable.

Finally, we note that the model R^2 tends to increase during Risk On and Range years but decreases right after Risk Off years, as declines in the volatility tend to lag increases in stock prices. Remarkably, we observe the increasing growth in R^2 and absolute value of the correlation during the past decade, in line with intensified correlations in cross-asset returns.

In conclusion, we find that the model R^2 and idiosyncratic volatility $\hat{\varepsilon}$ are relatively stable, while changes in $\hat{\beta}$ and $\hat{\kappa}$ can be explained by prevalent market conditions. For pricing applications, we can use a term structure of β to fit forward skews or reflect proprietary views on the market dynamics.

Pricing equation and calibration

The model dynamics (2) do not have a closed-form solution for vanilla options even with constant volatility. We obtain an approximate solution for call option values with a constant σ_{CV} or deterministic volatility $\sigma_{DV}(t)$. We illustrate that the proposed formula is in good agreement with numerical PDE solution across different maturities and strikes, especially for near-ATM strikes (available upon request). Nevertheless, the calibration of local volatility can only be implemented by means of PDE methods. Here we derive the necessary equations.

Using dynamics (2) for the log-spot, $X(t) = \ln(S(t)/S(0))$, we obtain:

$$\begin{aligned} dX(t) &= \mu(t)dt - \frac{1}{2}\sigma^2(1+Y(t))^2 dt + \sigma(1+Y(t))dW^{(0)}, \quad X(0)=0 \\ dY(t) &= \beta\sigma(1+Y(t))dW^{(0)} - \kappa Y(t)dt + \varepsilon dW^{(1)}, \quad Y(0)=0 \end{aligned} \tag{11}$$

with:

$$dY(t)dY(t) = \left(\varepsilon^2 + \beta^2\sigma^2(1+Y(t))^2\right)dt$$

$$dX(t)dY(t) = \beta\sigma^2(1+Y(t))^2 dt$$

Finally, the pricing equation for value function $U(t, T, X, Y)$ has the form:

$$\begin{aligned} &U_t + \frac{1}{2}\sigma^2\left(1+2Y+Y^2\right)\left[U_{XX} - U_X\right] + \mu(t)U_X \\ &+\frac{1}{2}\left(\varepsilon^2 + \beta^2\sigma^2\left(1+2Y+Y^2\right)\right)U_{YY} - \kappa Y U_Y \\ &+\beta\sigma^2\left(1+2Y+Y^2\right)U_{XY} - r(t)U = 0 \end{aligned} \tag{12}$$

where $r(t)$ is the discount rate and subscripts denote partial derivatives.

The parameters of the stochastic volatility, β, ε and κ are specified before the calibration. Given that these parameters are specified, we calibrate the local volatility $\sigma \equiv \sigma_{LSV}(t, S)$ so that the vanilla surface is matched by construction. We use the standard relationship for local SV models:

$$\sigma^2_{LSV}(T,K)\mathbb{E}\left[(1+Y(T))^2 \middle| S(T)=K\right] = \sigma^2_{LV}(T,K) \tag{13}$$

where $\sigma^2_{LV}(T, K)$ is the local Dupire volatility.

The above expectation is calculated by solving the forward PDE corresponding to pricing PDE (12) using finite-difference methods and calculating $\sigma^2_{LSV}(T, K)$ stepping forward in time (see Sepp, 2011, for details). Once $\sigma_{LSV}(t, S)$ is calibrated, we use either backward PDEs or Monte Carlo simulation for valuation of exotic options.

Properties of the volatility process

Here, we consider the pricing model (2) with constant volatility σ_{CV}.

The instantaneous variance of $Y(t)$ is given by:

$$dY(t)dY(t) = \left(\beta^2\sigma^2_{CV}(1+Y(t))^2 + \varepsilon^2\right)dt \tag{14}$$

which has the systemic part and idiosyncratic part ε. In a stress regime, for large values of $Y(t)$, the variance is dominated by $\beta^2\sigma^2_{CV}Y^2(t)$. While in a normal regime, the variance is nearly a constant $(\beta^2\sigma^2_{CV} + \varepsilon^2)$. The variance of variance can be increased by increasing β (so the model approaches a pure SV model) or ε (so the model approaches a local SV model). Regimes with high volatility of volatility, say with volatility of the Vix of 100%, are reproduced by a combination of high values of σ_{CV} (local volatility σ^2_{LSV} in the local beta SV model), β and ε.

Now, we show that the volatility process has steady-state volatility, so that the volatility approaches stationary distribution in the long run. We consider $\Upsilon(t) = Y^2(t)$ using equation (2):

$$d\Upsilon(t) = \left(-2\kappa\Upsilon(t) + \varepsilon^2 + \beta^2\sigma^2_{CV}\left(1+2Y(t)+\Upsilon(t)\right)\right)dt + \ldots \tag{15}$$

We define $\bar{Y}(t) = \mathbb{E}[Y(t)|\, Y(0) = 0]$ and $\bar{\Upsilon}(t) = \mathbb{E}[\Upsilon(t)|\, Y(0) = 0]$. Using (15), we obtain $\bar{Y}(t) = 0$ and:

$$\bar{\Upsilon}(t) = \frac{\varepsilon^2 + \beta^2\sigma^2_{CV}}{2\kappa - \beta^2\sigma^2_{CV}}\left(1 - e^{-\left(2\kappa - \beta^2\sigma^2_{CV}\right)t}\right) \tag{16}$$

The effective mean-reversion for the volatility of variance is:

$$2\kappa - \beta^2\sigma^2_{CV}$$

so that we need to enforce $\kappa > ½\beta^2\sigma^2_{CV}$. Also the steady-state variance of volatility is:

$$\frac{\varepsilon^2 + \beta^2\sigma^2_{CV}}{2\kappa - \beta^2\sigma^2_{CV}}$$

For high volatility σ_{CV}, the mean-reversion rate decreases, while the volatility of variance increases, which is in line with observed behaviour following Risk Off regimes, as explained above. For the SV model based on the Ornstein-Uhlenbeck process, the steady-state volatility of variance is given by $\varepsilon^2_{OU}/2\kappa_{OU}$, so that we obtain the following relationship:

$$\kappa_{OU} = \kappa - \frac{1}{2}\beta^2\sigma^2_{CV}, \quad \varepsilon^2_{OU} = \varepsilon^2 + \beta^2\sigma^2_{CV}$$

■ **Instantaneous correlation.** We consider the instantaneous correlation between $dY(t)$ and $dX(t)$ using (11):

$$\rho\left(dX(t)dY(t)\right) = \frac{\beta\sigma^2_{CV}(1+Y(t))^2}{\sqrt{\left(\varepsilon^2 + \beta^2\sigma^2_{CV}(1+Y(t))^2\right)}\sqrt{\sigma^2_{CV}(1+Y(t))^2}} \tag{17}$$

2 Impact of β and ε on the term structure of the current skew and the forward three-year three-month skew

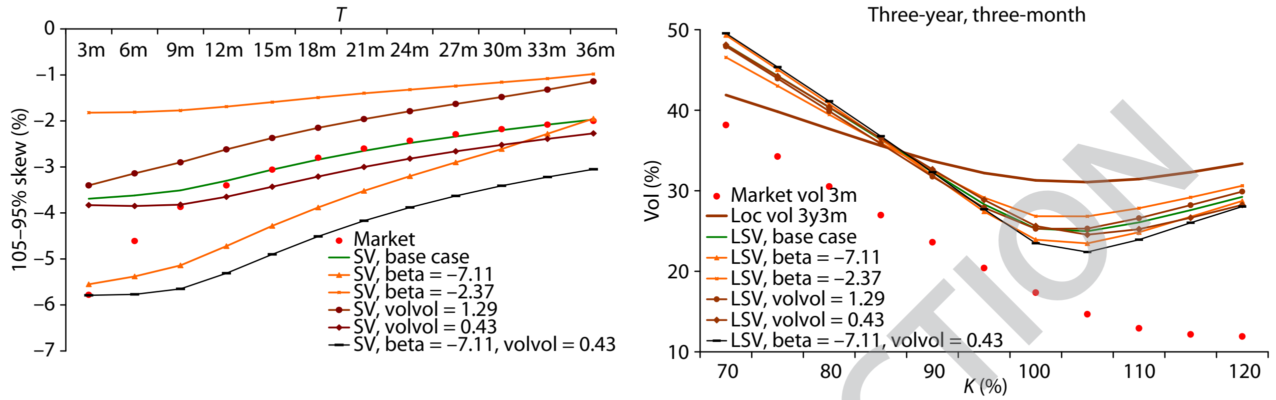

Left: impact on skew from β, ε and their combination on the term structure of the model implied 105–95% skew using model parameters in (25) as the base case; right: impact on the forward three-year, three-month skew

The first thing to note is that the correlations are state-dependent – in contrast to the Heston and Ornstein-Uhlenbeck-based SV models, in which it is constant – through the role of $Y(t)$ in the functional form. But asymptotes for high and low volatility regimes can be derived.

A high-volatility regime can be considered by letting $Y(t) \to \infty$ in (17), to find $\rho(dX(t)dY(t)) = -1$, as $\beta < 0$. On the other hand, in a normal regime, $Y(t) \approx 0$, so:

$$\begin{aligned} \rho\big(dX(t)dY(t)\big)\Big|_{Y(t)\approx 0} &= -\frac{1}{\sqrt{\left(\frac{\varepsilon^2}{\beta^2\sigma_{CV}^2}+1\right)}}, \\ \rho\big(dX(t)dY(t)\big)\Big|_{Y(t)\approx \infty} &= -1 \end{aligned} \tag{18}$$

Now consider the low-volatility regime case when $Y(t)$ is close to -1, $Y(t) = -1 + \delta$, with $|\delta|$ small. Making this substitution yields:

$$\rho\big(dX(t)dY(t)\big)\Big|_{Y(t)=-1+\delta} = -\frac{1}{\sqrt{\left(\frac{\varepsilon^2}{\delta^2\beta^2\sigma_{CV}^2}+1\right)}} \to -1 \text{ as } \delta \to 0$$

So the correlation approaches -1 for small instantaneous volatility as well. The asymptote is obtained in the limiting case of lognormal volatility, $\varepsilon = 0$. Since ε is small, the probability of $Y(t)$ going below -1 is negligible.

Finally, we note that we can use the first equation in (18) to estimate the idiosyncratic volatility of volatility ε given a specified spot-volatility correlation ρ^* by equating it to a given correlation ρ^* to obtain:

$$\varepsilon^2 = \sigma_{CV}^2\beta^2\frac{1-\left(\rho^*\right)^2}{\left(\rho^*\right)^2} \tag{19}$$

■ **The steady-state density.** The steady-state density function $G(Y)$ of the volatility factor $Y(t)$ in dynamics (2) solves the following equation:

$$\frac{1}{2}\Big[\Big(\varepsilon^2+\beta^2\sigma_{CV}^2\big(1+2Y+Y^2\big)\Big)G\Big]_{YY}+\big[\kappa YG\big]_Y=0 \tag{20}$$

We can show that $G(Y)$ exhibits the power-like behaviour for large values of Y:

$$\lim_{Y\to+\infty} G(Y) = Y^{-\alpha}, \quad \alpha = 2\left(1+\frac{\kappa}{\left(\beta\sigma_{CV}\right)^2}\right) \tag{21}$$

This power-like behaviour contrasts with Heston and exponential volatility models that imply exponential tails for the steady-state density of the volatility. The beta SV model predicts much higher probabilities of large values of instantaneous volatility. This requires some care for numerical PDE implementation. For numerical PDE methods, we specify the upper bound Y_∞ so that:

$$\mathbb{P}\big[Y > Y_\infty\big] \equiv \frac{1}{\alpha-1}Y_\infty^{1-\alpha} \le \varepsilon, \quad \text{so } Y_\infty \ge \big[(\alpha-1)\varepsilon\big]^{\frac{1}{\alpha-1}} \tag{22}$$

To find the lower bound, we can show that, in the vicinity of $Y = -1$, $G(Y)$ is Gaussian:

$$\lim_{Y\to-1} G(Y) = \mathbf{n}\left(\frac{Y+1}{\eta}\right), \quad \eta^2 = \frac{\varepsilon^2}{2\left(\kappa+\left(\beta\sigma_{CV}\right)^2\right)} \tag{23}$$

where $\mathbf{n}(x)$ is the normal probability density function. So the lower bound Y_0 can be specified by $Y_0 \le \eta\mathbf{N}^{-1}(\varepsilon) - 1$, where $\mathbf{N}^{-1}(x)$ is the inverse of the normal cumulative distribution function. For typical model parameters and $\varepsilon = 10^{-4}$, $Y_\infty \approx 15$ and $Y_0 \approx -3$. For our PDE solver, we use a refined grid with 800 points in the Y direction and 250 in the time and spot directions.

Illustrations

■ **Model calibration.** The quality of fit of the beta SV model is similar to one-factor SV models. The model fits longer-term skews well but is unable to fit short-term skews of up to one year unless the beta parameter β is large. To model the short-term skews, we could resort to jumps, but choose to introduce local volatility.

To calibrate the model, we first use the pricing model with deterministic volatility and set the parameter β so that the model fits the market-implied long-term 105–95% skew. We found that the following empirical rule based on equation (7) works well:

$$\beta = \frac{\sigma_{IMP}(6m,5\%) - \sigma_{IMP}(6m,-5\%)}{0.05\sigma_{IMP}(6m,0\%)} \tag{24}$$

where $\sigma_{IMP}(6\text{m}, k\%)$ is six-month implied volatility for a forward-based log-strike k.

Then we use equation (19) to specify ε given beta and σ_{CV} =

$\sigma_{DV}(1y)$. A couple of iterations are needed as $\sigma_{DV}(1y)$ mildly depends on ε. We set the correlation to an approximate historical average, $\rho^* = -0.80$. The mean-reversion κ is fixed so that the shape of the model-implied 105–95% skew fits the market above one year. Applying this to market options data on the S&P 500 index as of June 20, 2012, we obtain:

$$\beta = -4.74, \quad \varepsilon = 0.86, \quad \kappa = 1.10 \tag{25}$$

Impact of model parameters on term structure of current skew. In figure 2 (left), we show the impact of parameters β, ε and both of them, respectively, on the term structure of the model-implied 105–95% skew using model parameters in (25) as the base case. We use the beta SV model with deterministic $\sigma_{DV}(t)$ fitted to match the term structure of ATM volatilities for any combination of model parameters. Here, we shift β down and up by 50% to β = –7.11 and β = –2.37, respectively, and ε up and down by 50% to ε = 1.29 and ε = 1.43, respectively.

We see that, in the base case, the beta SV with deterministic volatility fits the term structure of the skew above one year very well. Decreasing β in absolute value leads to a smaller overall level of skew. Increasing β leads to steeper skew in the short and medium term but the effect is not parallel as the convexity effect reduces the skew for longer maturities. Decreasing (increasing) ε leads to a roughly parallel downward (upward) shift in the term structure of the skew as the idiosyncratic component of the volatility of volatility declines (grows), increasing (decreasing) the skew. For shorter maturities, ε has little impact. Finally, increasing β in absolute value while decreasing ε leads to a roughly parallel downward shift in the model-implied skew as a smaller value of the idiosyncratic volatility ε mitigates the convexity coming from a high value of β. As a result, by changing values of β and ε, we can generate a variety of shapes for the term structure of the model-implied skew.

Impact of model parameters on forward skew. In figure 2 (right), we show the impact of parameters β, ε and both of them, respectively, on the forward three-year three-month skews for forward-start options. Here, we use the beta SV model with local volatility so that the model produces the same vanilla skews for different combinations of SV model parameters. In comparison, we plot the current three-month implied skew, and the forward skews calculated by the local volatility. In general, we see that the local volatility predicts fattening skews; in contrast, the beta SV model produces steep forward skews.

We see that for larger (in absolute value) β, the forward skews are steeper. Increasing ε, the convexity increases and the skew becomes U-shaped. However, the impact of ε is less pronounced for forward skews of OTM puts. Finally, increasing β and reducing ε leads to steep forward skews with forward three-year three-month 105–95% skew being nearly as steep as the current three-month skew. As a result, by changing the values of the parameters β and ε, we can produce different shapes of forward skews.

Pricing cliquets. Now we apply the beta SV model for pricing cliquets using parameters as in (25) with local volatility calibrated to the vanilla surface. For comparison, we provide prices calculated by the local volatility model and the Heston model calibrated to skews at maturities of one year, two years and three years. For the beta SV model, we use three sets of parameters: LSV beta (I) corresponds to the base case with β = –4.74, ε = 0.86; LSV beta (II) corresponds to the model with increased β, β = –7.11, ε = 0.86; and LSV beta (III) corresponds to the model with increased β and decreased ε, β = –7.11, ε = 0.43.

Our example is the cliquet option with the following payout:

$$\sum_{i=1}^{N} \max\left[\min\left\{\frac{S(t_i)}{S(t_{i-1})} - 1, C\right\}, F\right]$$

In table B, we show prices of cliquet options with quarterly fixings and local floor (F) and cap (C) of –5% and 5%, respectively. This product depends heavily on the forward skew. In the local volatility model, the cap is overpriced while the floor is underpriced so that the cliquet price is too small. The Heston model produces markedly higher prices than the local volatility model, but it is known that observed market prices of cliquets are higher than those that can be produced by the Heston model. The base case model (I) underprices cliques compared with Heston, while increasing β produces higher prices in model (II), while the strongest effect is produced with decreased ε in model (III).

B. Pricing cliquets (%)

	Local vol	Heston	LSV beta (I)	LSV beta (II)	LSV beta (III)
1y	1.03	2.10	1.38	1.93	1.82
2y	0.67	3.70	2.25	2.82	3.69
3y	0.01	5.09	3.16	4.10	5.99

Note: LSV beta (I) corresponds to β = –4.74, ε = 0.86; LSV beta (II) corresponds to β = –7.11, ε = 0.86; LSV beta (III) corresponds to β = –7.11, ε = 0.43

Conclusion

We have presented the beta stochastic volatility model. One of the important features of this model is that its key parameter β has a natural interpretation as the rate of change in the short-term ATM volatility given change in the stock price. As a result, empirical estimation of model parameters is easy to implement and interpret. In general, the beta SV model is augmented with local volatility, so that users can concentrate on modelling of forward skews with the vanilla surface being matched by construction. In the case of deterministic volatility, we have derived an accurate approximation for call prices in this model. We have shown than the beta SV model produces steeper forward skews and can be better suited for modelling of the forward volatility and related products. ■

Piotr Karasinski is a senior adviser at the European Bank for Reconstruction and Development in London. Artur Sepp is a vice-president of the global quantitative group at Bank of America Merrill Lynch in London. They are
grateful to their colleagues and two anonymous referees for helpful and instructive comments. Artur is particularly grateful to Hassan El Hady for bringing to his attention some early work on this model by Piotr.
The views expressed in this article are those of the authors alone. Email: piotr.g.karasinski@gmail.com, artur.sepp@baml.com

References

Bardhan I and P Karasinski, 1993
Random volatility model implied by S&P option prices
Internal paper, Goldman Sachs

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

Hagan P, D Kumar, A Lesniewski and D Woodward, 2002
Managing smile risk
Wilmott Magazine, September, pages 84–108

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

Langnau A, 2004
Introduction to the stochastic implied beta model
Internal paper, Merrill Lynch

Lipton A, 2002
The vol smile problem
Risk February, pages 81–85

Sepp A, 2011
Efficient numerical PDE methods to solve calibration and pricing problems in local stochastic volatility models
Global Derivatives Conference, Paris