Lecture 3: Asymptotics and Dynamics of the Volatility Skew

Jim Gatheral, Merrill Lynch*

Case Studies in Financial Modelling Course Notes,
Courant Institute of Mathematical Sciences,
Fall Term, 2001

*I am grateful to Peter Friz for carefully reading these notes, providing corrections and suggesting useful improvements.

9 Asymptotic Behaviour of the Volatility Skew

9.1 Short Expirations

We start by rewriting our original general stochastic volatility SDEs (1) and (2) in terms of the log-moneyness $x = \log\left(\frac{F}{K}\right)$ and under the risk neutral measure, specializing to the case where α and β do not depend on S or t.

$$\begin{aligned} dx &= -\frac{v}{2}dt + \sqrt{v}\,dZ_1 \\ dv &= \alpha(v)\,dt + \eta\sqrt{v}\beta(v)\,dZ_2 \end{aligned} \tag{42}$$

We may rewrite

$$dZ_2 = \rho dZ_1 + \varphi\,dZ_1^*$$

with $\varphi = \sqrt{1-\rho^2}$ and $\langle dZ_1^*, dZ_1\rangle = 0$. Eliminating $\sqrt{v}dZ_1$, we get

$$dv = \alpha(v,t)\,dt + \rho\eta\,\beta(v,t)\,\left\{dx + \frac{v}{2}dt\right\} + \varphi\eta\,\beta(v)\,\sqrt{v}\,dZ_1^*$$

Then,

$$\mathbf{E}\left[v + dv\,|dx\right] = v + \alpha(v)\,dt + \rho\eta\,\beta(v)\,\left\{dx + \frac{v}{2}dt\right\}$$

so for small times to expiration (relative to the variation of $\alpha(v)$ and $\beta(v)$), we have

$$\begin{aligned} v_{loc}(x,t) &= \mathbf{E}\left[v_t\,|x_t = x\right] \\ &\approx v_0 + \left[\alpha(v_0) + \rho\eta\frac{v_0}{2}\beta(v_0)\right]t + \rho\eta\beta(v_0)\,x \end{aligned} \tag{43}$$

The coefficient of x (the slope of the skew) here agrees with that derived by Lee (2001) using a perturbation expansion approach.

To extend the result to implied volatility, we need the following lemma:

Lemma

The local volatility skew is twice as steep as the implied volatility skew for short times to expiration.

Proof

From Section 5.2, we know that BS implied total variance is the integral of local variance along the most probable path from the stock price on the valuation date to the strike price at expiration. This path is approximately a straight line (see Figure 1). Also, from Equation 43, we see that the slope of the local variance skew is a roughly constant $\beta(v_0)$ for short times. The BS implied variance skew, being the average of the local variance skews, is one half of the local variance skew. Formally,

$$\begin{aligned}\sigma_{BS}(K,T)^2 &\approx \frac{1}{T}\int_0^T v_{loc}(\tilde{x}_t,t)dt \\ &\approx \text{const.} + \frac{1}{T}\int_0^T \rho\eta\beta(v_0)\tilde{x}_t dt \\ &\approx \text{const.} + \frac{1}{T}\int_0^T \rho\eta\beta(v_0)x_T\frac{t}{T}dt \\ &= \text{const.} + \frac{1}{2}\rho\eta\beta(v_0)\,x_T\end{aligned}$$

where $\tilde{x}$ represents the "most probable" path from the stock price at time zero to the strike price at expiration. □

We conclude that for short times to expiration, the BS implied variance skew is given by

$$\frac{\partial}{\partial x}\sigma_{BS}(x,t)^2 = \frac{\rho\eta}{2}\beta(v_0) \tag{44}$$

Recall that in the Heston model, $\beta(v) = 1$; we see that equation 44 is consistent with the short-dated volatility skew behavior that we derived earlier in Section 6.2 for the Heston model.

Note that the short-dated volatility skew is not explicitly time-dependent; it depends only on the form of the SDE for volatility. In contrast, as we shall see, local volatility models imply short-dated skews which decay rapidly as time advances. So even if we find a stochastic volatility model and a local volatility model that price all European options identically today, forward-starting options (that is options whose strikes are to be set some time in the future) cannot possibly be priced identically by these two models. Both models fit the options market today but the volatility surface dynamics implied by the two models are quite different.

Figure 1: Integrating local variance to get implied variance

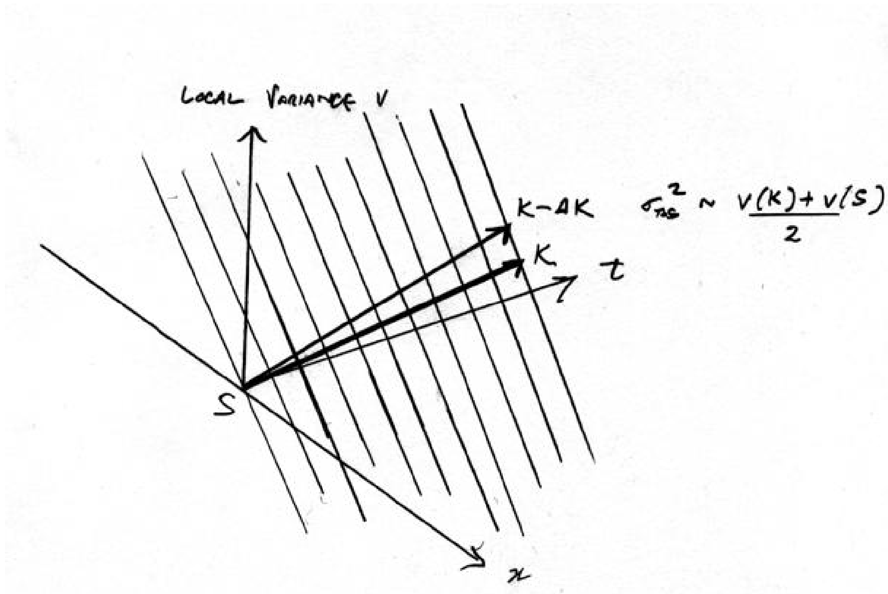

Equation 44 suggests a wild generalization: perhaps all stochastic volatility models, whether analytically tractable or not, have similar implications for the BS implied volatility skew up to a factor of $\beta(v)$. By investigating the behavior of the volatility skew at long expirations, we will present evidence which makes this claim more plausible.

9.2 Long Expirations

Fouque, Papanicolaou, and Sircar (1999) and Fouque, Papanicolaou, and Sircar (2000) show using a perturbation expansion approach that in any stochastic volatility model where volatility is mean-reverting, Black-Scholes implied volatility can be well approximated by a simple function of log-moneyness and time to expiration for long-dated options. In particular, they study a model where the log-volatility is a Orenstein-Uhlenbeck process (log-OU for short). That is:

$$
\begin{aligned}
dx &= -\frac{\sigma^2}{2}dt + \sigma\, dZ_1 \\
d\log(\sigma) &= -\lambda[\log(\sigma) - \overline{\log(\sigma)}\,]dt + \xi dZ_2
\end{aligned}
$$

They find that the slope of the BS implied volatility skew is given (for large λT) by

$$\frac{\partial}{\partial x}\sigma_{BS}(x,T) \approx \frac{\rho\xi}{\lambda T} \tag{45}$$

To recast this in terms of v to be consistent with the form of the generic process we wrote down in Equation 42, we note that (considering random terms only), $dv \sim 2\,\sigma\,d\sigma$ and in the log-OU model,

$$d\sigma \sim \xi\sigma dZ_2$$

So

$$dv \sim 2\xi v dZ_2$$

Then $\beta(v)$ as defined in Equation 42 is given by

$$\eta\beta(v) = 2\xi\sqrt{v}$$

and, from Equation 45, the BS implied variance skew is given by

$$\frac{\partial}{\partial x}\sigma_{BS}(x,T)^2 \approx \frac{2\rho\xi\sqrt{v}}{\lambda T} = \frac{\rho\eta\beta(v)}{\lambda T}$$

Looking back at section 6.2 again, we see that the Heston skew (where $\beta(v) = 1$) has the same behavior for large λT. We now have enough evidence to make our generalization more plausible: it seems that both for long and short expirations, the skew behavior may be identical for all stochastic volatility models up to a factor of $\beta(v)$. Supposing this claim were true, what would be the natural way to interpolate the asymptotic skew behaviors between long and short expirations?

Clearly, the most plausible interpolation function between short expiration and long expiration volatility skews is the one we already derived for the Heston model in Section 6.2 and

$$\frac{\partial}{\partial x}\sigma_{BS}(x,T)^2 \approx \frac{\rho\eta\,\beta(v)}{\lambda' T}\left\{1 - \frac{\left(1 - e^{-\lambda' T}\right)}{\lambda' T}\right\} \tag{46}$$

with $\lambda' = \lambda - \frac{1}{2}\rho\eta\beta(v)$.

9.3 Dynamics of the Volatility Skew under Stochastic Volatility

At first it might seem that a result that says that all stochastic volatility models have essentially the same implications for the shape of the volatility surface would it make it hard to differentiate between models. That would certainly be the case if we were to confine our attention to the shape of the volatility surface today. However, if instead we were to study the dynamics of the volatility skew – in particular, how the observed volatility skew depends on the overall level of volatility, we would be able to differentiate between models.

Empirical studies of the dynamics of the volatility skew show that $\frac{\partial}{\partial x}\sigma(x,t)$ is approximately independent of volatility level over time. Translating this into a statement about the implied variance skew, we get

$$\frac{\partial}{\partial x}\sigma_{BS}(x,t)^2 \;=\; 2\sigma_{BS}(x,t)\frac{\partial}{\partial x}\sigma_{BS}(x,t) \;\sim\; \sqrt{v(x,t)}.$$

This in turn implies that $\beta(v) \;\sim\; \sqrt{v}$ and that v is approximately lognormal in contrast to the square root process assumed by Heston. This makes intuitive sense given that we would expect volatility to be more volatile if the volatility level is high than if the volatility level itself is low.

Does it matter whether we model variance as a square root process or as lognormal? In certain cases it does. After all, we are using our model to hedge and the hedge should approximately generate the correct payoff at the boundary. If the payoff that we are hedging depends (directly or indirectly) on the volatility skew, and our assumption is that the variance skew is independent of the volatility level, we could end up losing a lot of money if that's not how the market actually behaves.

Is any stochastic volatility model better than none at all? The answer here has to be yes because whereas having the wrong stochastic volatility model will cause the hedger to generate a payoff corresponding to a skew that may perhaps be off by a factor of 1.5 if volatility doubles, having only a local volatility model will cause the hedger to generate a payoff that corresponds to almost no forward skew at all. We will now show this.

9.4 Dynamics of the Volatility Skew under Local Volatility

Empirically, the slope of the volatility skew decreases with time to expiration. From the above, in the case of mean-reverting stochastic volatility, the term structure of the BS implied variance skew will look something like Equation 46. In particular, the slope of the volatility skew will decay over time according to the time behaviour of the coefficient $\frac{1}{\lambda' T}\left\{1-\frac{\left(1-e^{-\lambda' T}\right)}{\lambda' T}\right\}$.

Recall from Section 2.3 the formula for local volatility in terms of implied volatility:

$$v_{loc}=\frac{\frac{\partial w}{\partial T}}{1-\frac{x}{w}\frac{\partial w}{\partial x}+\frac{1}{4}\left(-\frac{1}{4}-\frac{1}{w}+\frac{x^2}{w^2}\right)\left(\frac{\partial w}{\partial x}\right)^2+\frac{1}{2}\frac{\partial^2 w}{\partial x^2}}$$

Differentiating with respect to x and considering only the leading term in $\frac{\partial w}{\partial x}$ (which is small for large T) , we find

$$\frac{\partial v_{loc}}{\partial x}\approx\frac{\partial}{\partial T}\frac{\partial w}{\partial x}+\frac{1}{w}\frac{\partial w}{\partial T}\frac{\partial w}{\partial x}$$

That is, the local variance skew $\frac{\partial v_{loc}}{\partial x}$ decays with the BS implied total variance skew $\frac{\partial w}{\partial x}$.

To get the forward volatility surface from the local volatility surface in a local volatility model, we integrate over the local volatilities from the (forward) valuation date to the expiration of the option along the most probable path joining the current stock price to the strike price using the trick presented in Section 5.2. It is obvious that the forward implied volatility surface will be substantially flatter than today's because the forward local volatility skews are all flatter.

Contrast this with a stochastic volatility model where implied volatility skews are approximately time-homogeneous. In other words, local volatility models imply that future BS implied volatility surfaces will be flat (relative to today's) and stochastic volatility models imply that future BS implied volatility surfaces will look like today's.

10 Digital Options and Digital Cliquets

In our first investigation of actual derivative contracts, we choose to study digital options because their valuation involves the volatility skew directly.

10.1 Valuing Digital Options

A digital (call) option $D(K,T)$ pays 1 if the stock price S_T at expiration T is greater than the strike price K and zero otherwise. It may be valued as the limit of a call spread as the spread between the strikes is reduced to zero.

$$D(K,T) = -\frac{\partial C(K,T)}{\partial K} \tag{47}$$

where $C(K,T)$ represents the price of a European call option with strike K expiring at time T.

To see that its price is very sensitive to the volatility skew, we rewrite the European call price in Equation 47 in terms of its Black-Scholes implied volatility $\sigma_{BS}(K,T)$.

$$\begin{aligned} D(K,T) &= -\frac{\partial}{\partial K} C_{BS}\left(K, T, \sigma_{BS}(K,T)\right) \\ &= -\frac{\partial C_{BS}}{\partial K} - \frac{\partial C_{BS}}{\partial \sigma_{BS}} \frac{\partial \sigma_{BS}}{\partial K} \end{aligned}$$

To get an idea of the impact of the skew in practice, consider our usual idealized market with zero interest rate and dividends and a one year digital option struck at-the-money. Suppose further that at-the-money volatility is 25% and the volatility skew (typical of SPX for example) is 3% per 10% change in strike. Its value is given by:

$$\begin{aligned} D(1,1) &= -\frac{\partial C_{BS}}{\partial K} - \frac{\partial C_{BS}}{\partial \sigma_{BS}} \frac{\partial \sigma_{BS}}{\partial K} \\ &= N\left(-\frac{\sigma}{2}\right) - \text{vega} \times \text{skew} \\ &= N\left(-\frac{\sigma}{2}\right) + \frac{1}{\sqrt{2\pi}} e^{-\frac{d_1^2}{2}} \times 0.3 \\ &\approx N\left(-\frac{\sigma}{2}\right) + 0.4 \times 0.3 \end{aligned}$$

If we had ignored the skew contribution, we would have got the price of the digital option wrong by 12% of notional!

10.2 Digital Cliquets

For an example of an actual digital cliquet contract, see the addendum. Here is a description of the Cliquet from the IFCI site at http://risk.ifci.ch:

"The French like the sound of 'cliquet' and seem prepared to apply the term to any remotely appropriate option structure. (1) Originally a periodic reset option with multiple payouts or a ratchet option (from vilbrequin à cliquet – ratchet brace). Also called Ratchet Option. See Multi-period Strike Reset Option (MSRO), Stock Market Annual Reset Term (SMART) Note. See also Coupon Indexed Note. (2) See Ladder Option or Note (diagram). Also called Lock-Step Option. See also Stock Upside Note Security (SUNS). (3) Less commonly, a rolling spread with strike price resets, usually at regular intervals. (4) An exploding or knockout option such as CAPS (from cliqueter – to knock)."

Their payoff diagram shown in Figure 2 is also a work of art. For our

Figure 2: Illustration of a Cliquet Payoff Courtesy of IFCI.

purposes, a cliquet is just a series of options whose strikes are set on a sequence of futures dates. In particular, a digital cliquet is a sequence of digital options whose strikes will be set (usually) at the prevailing stock price on the relevant reset date. Denoting the set of reset dates by $\{t_1, t_2, ..., t_n\}$, the digital cliquet pays Coupon $\times\, \theta\big(\mathrm{S}_{\mathrm{t_i}} - \mathrm{S}_{\mathrm{t_{i-1}}}\big)$ at t_i where $\theta(.)$ represents the Heaviside function.

One can see immediately that the package consisting of a zero coupon bond together with a digital cliquet makes a very natural product for a risk-averse retail investor – he typically gets an above market coupon if the underlying stock index is up for the period (usually a year) and a below market coupon (usually zero) if the underlying stock index is down. Not surprisingly, this product was and is very popular and as a result, many equity derivatives dealers have digital cliquets on their books.

From the foregoing, the price of a digital cliquet may vary very substantially depending on the modeling assumptions made by the seller. Those sellers using local volatility models will certainly value a digital cliquet at a lower price than sellers using a stochastic volatility (or more practically, those guessing that the forward skew should look like today's). Perversely then, those sellers using an inadequate model will almost certainly win the deal and end up short a portfolio of misvalued forward-starting digital options. Or even worse, a dealer could have an appropriate valuation approach but be pushed internally by the salespeople to match (mistaken) competitor's lower prices. The homework assignment deals with exactly this set of circumstances.

How wrong could the price of the digital cliquet be? Taking the example of the deal documented in the addendum, neglecting the first coupon (because we suppose that all dealers can price a digital which sets today), the error could be up to 12% of the sum of the remaining coupons (52%) or 6.24% of Notional. In the actual deal, the digital are struck out-of-the-money and interest rates and dividends are not zero. Nevertheless, a pricing error of this magnitude is a big multiple of the typical margin on such a trade and would cause the dealer a substantial loss.

References

Fouque, Jean-Pierre, George Papanicolaou, and K. Ronnie Sircar, 1999, Financial modeling in a fast mean-reverting stochastic volatility environment, *Asia-Pacific Financial Markets* 6, 37–48.

——— , 2000, Mean-reverting stochastic volatility, *International Journal of Theoretical and Applied Finance* 3, 101–142.

Lee, Roger W., 2001, Implied and local volatilities under stochastic volatility, *International Journal of Theoretical and Applied Finance* 4, 45–89.