Unconstrained Fitting of Implied Volatility Surfaces Using a Mixture of Normals

Riccardo Rebonato[+]* and Maria Teresa Cardoso*
*RBS Quantitative Research Centre (QUARC)
[+]Oxford University - OCIAM

July 6, 2004

Abstract

The paper presents a method to estimate a smooth volatility surface from the market prices of plain-vanilla options. Obtaining relibale surfaces is important for risk management purposes and for calibrating models. The method we suggest uses a mixture of (log)-normals that are combined in an unconstrained fashion to produce risk-neutral densities that display both kurtosis and skew. The method is shown to provide excellent …ts to a variety of model or market volatility surgaces.

1 Statement of the Problem

Obtaining reliable smile surfaces from the market prices of plain-vanilla options is of great importance both for pricing and for risk management. In the risk management arena, obtaining such a reliable surface minimizes the risk of mis-marking options for which prices are not directly visible. As for pricing, one of the most important problems in calibrating option models (eg, stochastic-volatility, local-volatility, jump-di¤usion, etc) is obtaining a reliable smile surface from the (often noisy and non-contemporaneous) market prices of plain-vanilla options. For some models it can be debated whether using 'undoctored' noisy prices might be better than regularizing the input, but, for some approaches, such as local-volatility models, having a smooth (and di¤erentiable) input smile surface is a must (see the brief discussion below).

Many techniques have been proposed to obtain from the observed market prices of plain-vanilla options a smooth volatility surface[1]. Most of these approaches fall in either of three categories: smooth interpolation of (possibly

[1] A systematic survey would require a full review article in itself. Without attempting to cover the literature systematically, we simply mention Shimko (1993) as a representative example of a …tting procedure directly to prices, Jacquier and Jarrow (1995), Madan, Carr and Chang (1998) and Elliot, Lahaie and Madan (1995) as examples of …tting to transformed prices, Ait-Sahalia (2002) for a direct …tting of the implied volatility curve, Avellaneda (1998) for an application of the minimum-entropy methodology, Mirfendereski and Rebonato (2001) for direct modelling of the risk-neutral density.

transformed) prices, smooth interpolation of implied volatilities, or smooth fitting of risk-neutral densities. We argue that modelling directly the density is the most desirable approach, because a smooth and 'plausible' density will certainly produce (by integration) smooth and well-behaved price and implied volatility surfaces, but the converse is not true. See, for instance, the discussion in Mirfendereski and Rebonato (2001). The weakness of price-based smoothing methods is that, if one directly interpolates across prices, one can easily obtain wildly oscillating densities. This is because the latter are obtained, modulo discounting, from the call prices by double differentiation with respect to strike of the price function[2]:

$$\phi(S) \propto \frac{\partial^2 C}{\partial K^2} \tag{1}$$

Obtaining such a wildly oscillating density is undesirable for several reasons. For instance, Willmot (1998) highlights that, if the implied risk-neutral density is not well behaved, local-volatility solutions can become extremely noisy. This is easy to see, if one recalls Dupire's (1994) well-known result:

$$\sigma^2_{K,T} = 2\frac{\frac{\partial C_{K,T}(0,S)}{\partial T} + (r - d)K\frac{\partial C_{K,T}(0,S)}{\partial K} + dC_{K,T}(0,S)}{\frac{\partial^2 C_{K,T}(0,S)}{\partial K^2}K^2} \tag{2}$$

Note, in fact, that the denominator in the expression above is simply linked to the risk-neutral density. Such a 'noisy' local-volatility solution is unlikely to have a financial justification, and is more likely to reflect the ill-posed nature of the inverse problem of recovering densities from prices. See, eg, Engl (1993).

Moving to methods based on the direct interpolation of implied volatility surfaces, one might argue that, as long as the implied volatility has been smoothly interpolated, the resulting risk-neutral densities will be smooth as well. Figs 1 and 2 show that this is not the case: Fig. 1 displays two optically indistinguishable implied volatility curves, and Fig 2 the associated risk-neutral densities. It is clear that relying on the smoothness of the implied volatility curve would not guarantee, for instance, that the local volatility extracted using Equation 2 would have desirable properties. More generally, Rebonato (2004) argues that obtaining smooth and well-behaved density functions is always important when the pricing of non-European options is required.

The only situation when direct modelling of the prices might be desirable is when the trader or the risk manager considers the market prices to be known with 'infinite precision', and therefore requires their perfect recovery by construction. Such an exact recovery cannot be guaranteed (and, indeed, will not in general be achieved) with the method we propose here. We show however that in most situations we can get extremely close to the prices to be recovered, and we submit that, in the presence of bid-offer spreads and of non-synchronous quotes, perfect recovery of an 'observed' price may be neither desirable nor, indeed, meaningful.

[2] It is important to observe that no model-dependent assumptions are needed to obtain Equation 1: one simply requires that the price of a European option can be written as a discounted expectation of the terminal payoff over some risk-neutral density.

This paper therefore sets as its goal to obtain an excellent fit to the market prices of plain-vanilla options by directly modelling the risk-neutral density. We choose to do so by extending the mixture-of-normals approach presented in Alexander (2001). The advantages of our extension is that the resulting risk-neutral density can show non-zero skew and still satisfy the risk-neutral forward condition, while retaining an unconstrained numerical search.

The paper is organized as follows. In Section 2 we set the problem in context by recalling known results. In Section 3 we present a numerical technique that will allow us to find the solutions in an efficient and unconstrained manner, while automatically satisfying all the financial and mathematical requirements. Section 4 then presents the numerical results, both for 'synthetic' model prices and for market prices. Section 5 presents a word of caution against 'over-interpreting' the meaning of the derivative of the price with respect to the underlying, and we argue that it should not be regarded as a 'delta'. Section 6 concludes the paper.

2 Summary of Known Results

If one wants to find the risk-neutral density, $\psi(\ln S_T)$, implied by a set of market prices for a given time horizon T, one can express it as a linear combination of normal densities, $\varphi(\mu_i; \sigma_i^2)$:

$$\psi(\ln S_T) = \sum_i w_i \varphi(\mu_i; \sigma_i^2) \tag{3}$$

If we want the resulting density to integrate to one, we must impose the normalization condition

$$\sum_i w_i = 1 \tag{4}$$

This is the approach followed, for instance, by Alexander (2001). The mean and the variance of the mixture, μ_ψ and σ_ψ^2; are related to the means and variances of the original normal distribution by the relationships

$$\mu_\psi = \sum_i w_i \mu_i \tag{5}$$

$$\sigma_\psi^2 = \sum_i w_i \sigma_i^2 + \left[\sum_i w_i \mu_i^2 - \left(\sum_i w_i \mu_i \right)^2 \right] \tag{6}$$

Note carefully that, since one is fitting risk-neutral densities, the first moment is not a free-fitting parameter, but must recover the forward condition, ie, the expectation in the risk-neutral measure of the 'stock' price must equal its forward value. Very often this condition has been enforced in the literature (see, eg, Alexander (2001)), by requiring that all the μ_i should be equal to the risk-neutral drift. This is, however, unnecessarily restrictive, because it forces the distribution of log-prices to display no skew. We show below how this feature can be naturally incorporated in the method we propose.

Since one of the most common features of empirical distributions is their leptokurtosis, it is also useful to give an expression (see Alexander (2001)) for the excess kurtosis for a mixture of normals. In the case when all the terms μ_i are equal to the risk-neutral drift one obtains:

$$\kappa_{\phi} = 3\left[\frac{\sum_i w_i \sigma_i^4}{\left(\sum_i w_i \sigma_i^2\right)^2} - 1\right] \tag{7}$$

From Equation 7 it is clear that the density of a mixture of normals (with same means μ_i!) will always have a positive excess kurtosis (ie, will be more leptokurtic than a normal density). This is because, for any non-degenerate case,

$$\sum_i w_i \sigma_i^4 > \left(\sum_i w_i \sigma_i^2\right)^2 \tag{8}$$

If one wanted to use a mixture of normals to fit an empirical price densities there are therefore two main routes:

1. one can directly estimate the first four moments of the empirical distribution of the logarithms of the price density, select two normals as the 'basis set' and fit the four moments exactly. With the fitted mixture-of-normals distribution the prices for the calls, C_K^T, can be determined and compared with the market values. The procedure is very straightforward, but, since it is based on the historical (backward-looking) distribution, the fit to the (forward-looking) market prices is unlikely to be very good.

2. one can determine the optimal weights $\{w\}$ by means of a least-square fit to the option prices after converting the density into call prices. This procedure is made easy by the fact that the option price C_K^T is simply given by a linear combination with the same weights $\{w\}$ of Black-and-Scholes formulae.

As mentioned above, however, if all the normal densities in the mixture are 'centered' (in log space) around the forward value, one is automatically guaranteed to recover the no-arbitrage forward pricing condition, but the resulting pricing density will display no skewness. This is at odds with empirical findings (see, eg, Madan et al (1998)). Skewness can be easily obtained by allowing the different constituent Gaussian densities to be centered around different location coefficients, μ_i. By so doing, however, some care must be given to recovering the first moment of the density exactly, since this is linked to the no-arbitrage cash-and-carry forward condition. Furthermore, if the weights are left unconstrained (apart from Equation 4), there is no guarantee that the resulting density will be everywhere positive. The following section shows how both these problems can be overcome.

3 Fitting the Risk-Neutral Density Function: Mixture of Normals

3.1 Ensuring the Normalization and Forward Constraints

Denote by S_i the price of the stock at time T_i: $S(T_i) = S_i$. If we denote by $\Phi(S_i)$ its risk-neutral probability density, we want to write

$$\Phi(S_i) = \sum_k w_k^i \varphi(S_k^i) \tag{9}$$

where

$$\varphi(S_k^i) = LN(\mu_{ik}; \sigma_{ik}^2; S_0) \tag{10}$$

and $LN(\mu_{ik}; \sigma_{ik}^2; S_0)$ denotes a log-normal density with

$$E(S_k^i) = S_0 \exp(\mu_{ik} T_i) \tag{11}$$

$$Var(S_k^i) = [S_0 \exp(\mu_{ik} T_i)]^2 \left[\exp\left(\sigma_{ik}^2 T_i\right) - 1\right] \tag{12}$$

By this expression, the risk-neutral[3] density for the stock price is expressed as a sum of log-normal densities, and therefore the stock price density is not log-normal.

In order to ensure that the density is everywhere positive, we require that all the weights should be positive. This can be achieved by imposing:

$$w_k^i = \left(\alpha_k^i\right)^2 \tag{13}$$

The normalization condition, which requires that

$$\sum_k w_k^i = 1 \tag{14}$$

therefore becomes

$$\sum_k \left(\alpha_k^i\right)^2 = 1 \tag{15}$$

This condition can always be satisfied by requiring that the coefficients α_k^i should be the polar co-ordinates of a unit-radius hyper-sphere. Therefore we can write

$$\alpha_k^i = f(\theta_1^i; \theta_2^i; \ldots; \theta_{n-1}^i) \tag{16}$$

For instance, for $n = 2$ one simply has

$$\alpha_1^i = \sin(\theta_1^i) \tag{17}$$

$$\alpha_2^i = \cos(\theta_1^i) \tag{18}$$

[3] To lighten the prose, the qualifier 'risk-neutral' is often omitted in the following where there is no risk of ambiguity.

This is certainly acceptable, because, for any angle θ_1^i,

$$\sin(\theta_1^i)^2 + \cos(\theta_1^i)^2 = \left(\alpha_1^i\right)^2 + \left(\alpha_2^i\right)^2 = 1 \tag{19}$$

and Equation 15 is satisfied. In the more general case the coefficients α_k^i are given by:

$$\alpha_k^i = \cos\theta_k^i \prod_{j=1}^{k-1} \sin\theta_j^i \qquad k = 1, 2, \ldots, m-1 \tag{20}$$

$$\alpha_k^i = \prod_{j=1}^{k-1} \sin\theta_j^i \qquad k = m \tag{21}$$

The reason for expressing the coefficients α_k^i in terms of 'angles' is that we will want to optimize the model density over the weights w_k^i in an unconstrained manner, while automatically resting assured that the resulting linear combination is a possible density. This will be the case only if 15 is always satisfied. In general, ie, if one tried to optimize directly over the weights w_k^i, one would have to carry out a heavily constrained numerical search: not only every weight w_k^i would have to be greater than zero but smaller than one, but also every partial sum over the weights would have to be strictly positive[4] and less than one. The procedure suggested above automatically ensures that this will always be the case, and therefore allows one to carry out an unconstrained optimization over the angle(s) θ[5].

Apart from the requirements in Equation 15, there is at least one more constraints. The no-arbitrage forward condition

$$E[S_i] = \int \phi(S_i = \hat{S}_i | S_0 = \hat{S}_0)\hat{S}_i d\hat{S}_i = S_0 \exp(rT_i) \tag{22}$$

must in fact always be satisfied exactly under penalty of arbitrage[6]. This could be trivially achieved by imposing

$$\mu_{ik} = r \qquad \text{for any } k \tag{23}$$

This, however, would give rise to densities with kurtosis but no skew. (Indeed, this is the approach suggested by Alexander (2001)). To obtain skew, we want to allow the various basis functions to be centred around different locations in $\log S$-space, but we want to do so while retaining the forward-pricing condition. This can be achieved as follows. From the relationships above we can write

$$E[S_i] = E\left[\sum_k w_k^i S_i^k\right] = \sum_k w_k^i E\left[S_i^k\right] = S_0 \exp(rT_i) \tag{24}$$

[4] If, in the real-world measure, we assume that all positive values of the underlying are possible, the density cannot go to zero under any equivalent measure.

[5] The range of the angles $\{\theta\}$ is unconstrained, becasue of the periodic nature of the sine and cosine functions.

[6] To lighten notation, we denote by r either the short rate or the difference between the short rate and the dividend yield or the difference between the domestic and the foreign short rates, according to whether one is dealing with the case of a non-dividend paying asset, of a dividend-paying stock or of an FX rate, respectively.

Recalling that

$$E\left[S_k^i\right] = S_0 \exp(\mu_{ik} T_i) \tag{25}$$

one finds that

$$\exp(rT_i) = \sum_k w_k^i \exp(\mu_{ik} T_i) \tag{26}$$

The summation over the number of basis functions, k, can be split into the first term, and the sum, $\sum'$, over the remaining terms:

$$\exp(rT_i) = w_1^i \exp(\mu_{i1} T_i) + \sum_k{}' w_k^i \exp(\mu_{ik} T_i)$$

which can be solved for μ_{i1}:

$$\exp(\mu_{i1} T_i) = \frac{\exp(rT_i) - \sum_k' w_k^i \exp(\mu_{ik} T_i)}{w_1^i} \;\rightarrow$$

$$\mu_{i1} = \frac{\ln\left[\frac{\exp(rT_i) - \sum_k' w_k^i \exp(\mu_{ik} T_i)}{w_1^i}\right]}{T_i} \tag{27}$$

In other words, if, for any maturity T_i, we choose the first location coefficient according to the expression above, we can always rest assured that the forward condition will be automatically satisfied. Note, however, that, a priori, there is no guarantee that the argument of the logarithm will always be positive. In practice, we have never found this to be a problem.

So, for any set of angles $\{\theta\}$, for any set of $\sigma_{ik}; k = 1, 2, \ldots, n$ and for any set of $\mu_{ik}; k = 2, 3, \ldots, n$ the forward and the normalization conditions will always be satisfied if μ_1^i is chosen to be given by Equation 27. In the following we will always assume that this choice has been made.

3.2 The Fitting Procedure

Let $Call_{K_j}^{T_i}(mod)$ be the model value of the call expiring at time T_i for strike K_j, and $Call_{K_j}^{T_i}(mkt)$ the corresponding market prices. The quantity $Call_{K_j}^{T_i}(mod)$ is given by

$$Call_{K_j}^{T_i}(mod) = \sum_{k=1,n} \alpha_k^i(\theta)^2 \int \varphi(S_i^k) G(S_i^k; K_j) dS_i^k \tag{28}$$

where $G(S_i; K_j)$ is the pay-off function. In the case of a call this is given by:

$$G(S_i^k; K_j) = \left[S_i^k - K_j\right]_+ \tag{29}$$

and $\varphi(S_i^k)$ is the log-normal density:

$$\varphi(S_i^k) = \frac{1}{S_i^k \sigma_{ik} \sqrt{2\pi T_i}} \exp\left[-\frac{1}{2}\left(\frac{\ln\left(\frac{S_i^k}{S_0 \exp(\mu_{ik} T_i)}\right) + \frac{1}{2}\sigma_{ik}^2 T_i}{\sigma_{ik}\sqrt{T_i}}\right)^2\right] \tag{30}$$

Note that each term under the integral sign is simply equal to the value of a Black-and-Scholes call when the 'risk-less rate' is equal to μ_{ik} and the volatility is equal to σ_{ik}. Therefore one can write:

$$\text{Call}_{K_j}^{T_i}(\text{mod}) = \sum_{k=1,n} \alpha_k^i(\theta)^2 \text{Call}_{BS}(\mu_{ik}, K_j, T_i, \sigma_{ik}) \tag{31}$$

where μ_{i1} is fixed from the previous forward relationship. Equation 31 lends itself to a simple interpretation: for any given strike, K_j, the model price is expressed as a linear combination of Black-and-Scholes prices, with the same strike, time to maturity and volatility, but with risk-less rate equal to μ_{ik} and volatility equal to σ_{ik}.

Now define χ^2 as

$$\chi^2 = \sum \left[\text{Call}_{K_j}^{T_i}(\text{mod}) - \text{Call}_{K_j}^{T_i}(\text{mkt}) \right]^2 \tag{32}$$

Then, in order to obtain the optimal fit to the observed set of market prices we simply have to carry out an unconstrained minimization of χ^2 over the $(n-1)$ angles $\{\theta\}$; the n volatilities $\sigma_{ik}, k = 1, 2, ..., n$, the $(n-1)$ location coefficients $\mu_{ik}, k = 2, 3, ..., n$ and with μ_{1i} given by (27). Therefore, for each expiry we have at our disposal $3n - 2$ coefficients. (For $n = 1$, we simply have one coefficient, ie one volatility. For $n = 2$ we have four coefficients, ie, two volatilities, one weight (ie one angle θ), and one location coefficient; etc).

4 Numerical Results

4.1 Description of the Numerical Tests

In this section we explore how well the mixture-of-normals method works in practice. We do so by looking both at theoretical densities and at market prices. The theoretical densities are obtained from three important models that will be discussed in the following, ie the jump-diffusion, the stochastic-volatility and the variance-gamma model. We have sometimes used rather 'extreme' choices of parameters in order to test the robustness and flexibility of the approach.

Finally, for simplicity we assumed a non-dividend-paying stock (interest rate at 5%) with spot price of \$100, and we looked at maturities of 0.5, 1, 2 and 4 years. Longer maturities, because of the Central Limit Theorem, would actually produce an easier test. All the optimized coefficients are reported in Tabs I to III. As for the market prices, they were obtained from the GBP caplet market in March 2003.

Tab I: The means, standard deviations and weights obtained for the fits to the jump-diffusion, stochastic-volatility and variance gamma models discussed in the text using a mixture of three log-normals. The chi squared statistics is also displayed.

Tab II: Â^2 statistics and maximum error for the mixtures described in Tab I

Tab III: Same as Tab I for the market data used (GBP caplet implied volatilities, March 2003). The Â^2 statistics is also displayed.

4.2 Fitting to Theoretical Prices: Stochastic-Volatility Density

The simplest test is probably that of a stochastic-volatility process for the underlying, because we know (Hull and White (1987)) that, in this case, the process for the logarithm of the price directly generates a terminal risk-neutral density which is made up of a mixture of normals[7]. A mean-reverting process was chosen for the volatility, with an initial value of the volatility equal to the reversion level (12.13%). The volatility of the volatility was given a very high value (100%) to 'stress' the test. We assumed no correlation between the Brownian shocks a¤ecting the underlying and the volatility.

The smiles produced by these parameters are shown in Fig 3. The theoretical density, the …tted density and the log-normal density matched to the …rst two moments are shown in Fig. 5. The match is excellent everywhere even with just three basis functions. The resulting …t to the smile is shown in Fig. 4.

Fig 3: The smiles produced by the parameters discussed in the text in the stochastic-volatility case.

Fig 4: The …t to the 0.5-year stochastic-volatility smile obtained using three log-normals.

Fig 5: The theoretical density, the …tted density and the moment-matched log-normal density in the stochastic-volatility case.

It is interesting to point out that we found that, after optimization, the three basis distributions in the mixture turned out to have the same (risk-neutral) mean even if they were not required to be so centered. This is consistent with the assumption of the underlying following a stochastic-volatility process, which automatically produces a mixture of identically-centered distributions. This results in a non-skewed density (see Figure 5). On the other hand, the resulting distribution has positive kurtosis, correctly reproduced by the …tting procedure.

4.3 Fitting to Theoretical Prices: Variance-Gamma Density

For this model, we consider the whole set of parameters estimated by Madan, Carr and Chang (1998) for the risk-neutral density of the S&P: ¾ = 12.13%,

[7] Each normal density component would have as variance the square of the root-mean-sqaured volatility encountered along each volatility path.

º = 16.86%, µ =-0.1436. Madan, Carr and Chang (1998) show that, in the risk-neutral world, the hypothesis of zero skewness can be rejected. Their risk-neutral density will therefore provide the …rst test for our method when the underlying distribution is skewed. Fig 8 displays the …t to the 0.5-year density obtained with just three basis functions, together with a moment-matched log-normal …t. Fig 7 displays the …t to the two-year smile. It is clear from the …gure that the model prices are everywhere recovered well within bid-o¤er spread[8].

Fig 6: The smiles produced by the parameters discussed in the text in the variance-gamma case

Fig 7: The …t to the 0.5-year variance-gamma smile obtained using three log-normals.

Fig 8: The theoretical density, the …tted density and the moment-matched log-normal density in the variance-gamma case.

4.4 Fitting to Theoretical Prices: Jump-Di¤usion Density

The last theoretical smile we consider is the 'stress case' of a log-normal jump-di¤usion process with parameters chosen so as to produce a multi-modal risk-neutral density for some maturities. Under what circumstances a jump-di¤usion process can give rise to multi-modal densities, and what this implies for the associated smiles is discussed in Rebonato (2004). The model parameters used were 1 jump/year for the jump frequency, and an expectation and volatility of the jump amplitude ratios of 0.7 and 1%, respectively. The volatility of the di¤usive part was taken to be constant at 12.13%.

The theoretical density (expiry 0.5 years) is shown in Fig 11 with a thin continuous line. The same …gure also shows for comparison a moment-matched log-normal density. The line with markers then shows the risk-neutral density obtained with a mixture of three log-normals. It is clear that, even for such a di¢cult-to-match theoretical risk neutral density, a very good agreement is obtained everywhere (with the exception of the very-low-strike region) with as few as three log-normals. Fig 12 shows that the …t to the risk-neutral density becomes virtually perfect everywhere with 5 basis functions. The resulting theoretical and …tted smiles for expiry 0.5 years (the most challenging one) are shown in Fig 10. One can observe that everywhere the target and …tted implied volatilities coincide to well within bid-o¤er spread. The largest discrepancy was found to be 0.7 basis points in volatility (in these units one vega would be 100 basis points).

Fig 9: The smiles produced by the parameters discussed in the text in the jump-di¤usion case.

[8] The bid-o¤er spread was assumed to be half a vega (50 basis points in volatility).

Fig 10: The …t to the 0.5-year jump-di¤usion smile obtained using three log-normals.

Fig 11: The theoretical density, the …tted density and the moment-matched log-normal density in the jump-di¤usion case (three lognormals). Note the relatively poor recovery of the theoretical density in the far left tail.

Fig 12: The theoretical density, the …tted density and the moment-matched log-normal density in the jump-di¤usion case (…ve lognormals). The density is now well recovered everywhere.

4.5 Fitting to Market Prices

In order to test the method with real market prices we looked at the smile curves for caplet prices (GBP, March 2003) for di¤erent expiries. See Fig. 13. Again, the shortest maturity provided the most challenging test. In Fig 14 we show the …ts obtained for all the maturities (0.5, 1, 2, 4 and 8 years). Also in this case, the …t is virtually perfect everywhere with four log-normal basis functions.

On the basis of these results, one can conclude that the mixture-of-normals approach provides a simple and robust method to …t even very complex price patterns. The fact that the model price is expressed as a linear combination of Black-and-Scholes prices makes it very easy to calculate the derivative $\frac{@C}{@S}$. It also makes it very tempting to interpret this derivative as the 'delta' statistic, ie, as the amount of stock that will allow to hedge (instantaneously) against movements in the underlying, and to replicate a payo¤ by expiry. This interpretation is however unwarranted, as we discuss below.

Fig 13: The market caplet smiles for several maturities (GBP, March 2003)

Fig 14: The …t to the market data using four log-normals for expiries of 0.5, 1, 2, 4, and 8 years.

5 Is the Term $\frac{@C}{@S}$ Really a Delta?

When one uses a …tting procedure like the mixture-of-normals approach one expresses the marginal (unconditional) price densities in terms of one or more basis functions. The terminal densities therefore display a parametric dependence on the value of the 'stock' price today, S_0. Since the plain-vanilla option prices (denoted by C for brevity in this section) are obtainable as integrals over these probability densities, also these prices will display a parametric dependence on the initial value of the stock. It is therefore possible to evaluate the quantity $\frac{@C}{@S}$. Furthermore, in the mixture-of-normal case one can do so analytically. If one wanted to, one could also evaluate $\frac{@C^2}{@S^2}$, $\frac{@C}{@t}$, etc. This has led to the statement, often found in the literature, that these closed-form expressions for the price, the 'delta', the 'gamma' etc of the plain-vanilla option constitute an alternative pricing model, di¤erent from, but on a conceptual par with, say, the Black-and-Scholes model. Some authors speak of 'pricing' and 'hedging' with the mixture-of-normals approach, and refer to the pricing of options as the 'mixture models'.

This statement should be quali…ed very carefully. This is because the quantity $\frac{\partial C}{\partial S}$ plays in the Black-and-Scholes world not just the role of the derivative of the call price with respect to the stock price today, but also of the amount of stock we have to hold in order to be …rst-order neutral to the stochastic stock price movements. So, in the Black-and-Scholes world we posit a process for the underlying, and we obtain that, if this process is correct, we can create a riskless portfolio by holding a $\frac{\partial C}{\partial S}$ amount of stock. By following this delta-neutral strategy to expiry we can replicate the terminal payo¤. It is the replication property that allows us to identify the fair price of the option with the price of the replicating portfolio. The delta quantity is however only correct only because we have assumed that the process was known (a geometric di¤usion in the Black-and-Scholes case). This is in general not true for the same quantity $\frac{\partial C}{\partial S}$ when seen in the context of the …tting procedures presented above. This is because the market prices, which our method could recover almost perfectly, might have been produced by a more complex process (perhaps by a process that does not allow perfect replication by dealing in the underlying). If we want, we can still call the terms $\frac{\partial C}{\partial S}$, $\frac{\partial C^2}{\partial S^2}$, $\frac{\partial C}{\partial t}$'delta', 'gamma', 'theta' etc, but their …nancial interpretation in terms of a self-…nancing dynamic trading strategy is not warranted. The problem (see, eg, Piterbarg (2003)) is that the process for the underlying is not uniquely speci…ed by the marginal densities. So, even if these were perfectly recovered by the …tting procedure, the trader would still not know what process has generated them.

We would like to note in passing that Brigo and Mercurio (2000) circumvent this problem by carrying out the mixture-of-normals …tting in conjunction with a local-volatility approach, that is designed to recover exactly any exogenous set of option prices.

6 Conclusions

We have shown the e¤ectiveness of mixture of normals as basis functions in terms of which a great variety of risk-neutral densities can be expanded. The contribution of this paper is to show how the positivity and forward-pricing conditions can be automatically and simply satis…ed without requiring that all the basis Gaussian functions should have the same mean. This would imply (in log space) absence of skew, which, at least in the risk-neutral measure, is a well-established and essential feature of the market smile.

References

Ait-Sahalia Y, (2002), 'Telling from Discrete Data Whether the Underlying Continuous-Time Data is a Di¤usion', Journal of Finance, Vol 57, No 5, October

2002, 2075-2112

Avellaneda M (1998) 'Minimum-Realtive-Entropy Calibration of Asset-Pricing Models', International Journal of Theoretical and Applied Finance, Vol 1, No. 4, 447-472

Alexander C (2001) 'Market Models - A Guide to Financial Data Analysis', John Wiley & Sons, Chicester, New York, Weinheim, Brisbane, Singapore, Toronto

Bahra, B, (1997), 'Implied Risk-Neutral Probability Density Functions from Option Prices: Theory and Applications', Bank of England Working Paper Series, No. 66

Brigo D, Mercurio F, (2000), 'A Mixed Up Smile', Risk, September 2000, 87-92

Jacquier E, Jarrow R A (1995) 'Dynamic Evaluation of Contінegent Claim Models: An Analysis of Model Error', Working Paper, Johnson Graduate School of Management, Cornell Univerity, Ithaca, NY

Elliot R J, Lahaie C H, Madan D B (1995) 'Filtering Derivative Security Valuations from Market Prices', Proceedings of the Isaac Newton Workshop in Financial Mathematics, Cambdridge University Press, Cambridge

Engl H E, (1993), Regularization Methods for the Stable Soultion of Inverse Problems', Surveys of Mathematics for Industry, 3, 71-143

Hull J, White A, (1987), 'The Pricing of Options on Assets with Stochastic Volatilities', Journal of Finance, XLII, (2), 281-300

Madan D B, Carr P P, Chang E C, (1998) 'The Variance Gamma Process and Option Pricing', European Finance Review, 2, 79-105

Madan D B, Seneta E, (1990), 'The Variance Gamma (VG) Model for Share Market Returns', Journal of Business, Vol. 63, no. 4, 511-524

Piterbarg V V, (2003), 'Mixture of Models: A Simple Recipe for a ... Hangover?', Bank of America Working Paper

Rebonato R, (2004), 'Volatility and Correlation: The Perfect Hedger and the Fox', Second Edition, Wiley, Chichester