

TESTING TECHNIQUES FOR ESTIMATING IMPLIED RNDs FROM THE PRICES OF EUROPEAN-STYLE OPTIONS

Abstract:

This paper examines two approaches to estimating implied risk-neutral probability density functions from the prices of European-style options. It sets up a monte carlo test to evaluate alternative techniques' ability to recover simulated distributions based on Heston's (1993) stochastic volatility model. The paper tests both for the accuracy and stability of the estimated summary statistics from RNDs. We find that a method based on interpolating the volatility smile outperforms the commonly used parametric approach that uses a mixture of lognormals.

July 2nd 1999

Neil Cooper
Bank of England

The author would like to thank Robert Bliss, Nikolaos Panigirtzoglou, James Proudman, Paul Soderlind, Gary Xu and participants (especially the discussants Jan Marc Berk and Holger Neuhaus) at the BIS workshop on implied PDFs for helpful comments. The views expressed in this paper are those of the author and do not necessarily reflect those of the Bank of England.

1. Introduction

In the last five years, there has been great interest amongst policy-makers in extracting information from the prices of financial assets. Options prices, in particular, have proved to be a particular rich source of information since they enable the extraction of a complete implied risk-neutral probability density function (RNDs) for the assets, interest rates and commodity prices upon which they trade. These RNDs have proven particularly useful in interpreting the market's assessment of the balance of risks associated with future movements in asset prices.

Reflecting this interest, a relatively large number of papers have been published that set out alternative techniques for the estimation of implied RNDs with examples of their application to particular markets. Despite this wide range of papers, nearly all are based on one of three basic approaches:

- estimating the parameters of a particular stochastic process for the underlying asset price from options prices and constructing the implied RND from the estimated process - see Malz(1995) and Bates(1996) for examples that incorporate jump processes;
- fitting a particular parametric functional form for the terminal asset price, for example a mixture of lognormals directly to options prices - see Bahra (1996,1997) and Melick and Thomas (1997);
- interpolating across the the call pricing function or the volatility smile, following Shimko (1993), and employing the Breeden and Litzenberger (1978) result that the implied distribution may be extracted by calculating the second partial derivative of that function with respect to the strike price.

The first approach has the disadvantage that it is based on a particular stochastic process: we cannot observe whether the assumed process can capture the density functions that are implicit within options' prices. In this paper we focus on the second and third approaches which are more flexible since by trying to estimate the density function directly they are consistent with many different stochastic processes.

Given these alternative techniques, a natural question is: “Which technique performs the best?” A key concern is the accuracy and stability of the estimated RNDs. Suppose we observe an estimated RND that displays bi-modality or “spikes.”¹ Should we interpret this as reflecting actual expectations or estimation errors? If we believe it to be the latter then the value of using implied RNDs is seriously diminished.

This paper attempts to address these concerns. It examines the empirical performance of two approaches to RND estimation by testing the ability of alternative techniques’ ability to recover the implied density function from a set of simulated prices. The simulated prices are generated from a quite general stochastic volatility model set out in Heston (1993). By using simulated prices, rather than actual prices, we can compare estimated RNDs against the “true” RND implied by the underlying price process. We test not just the stability of estimated RNDs and their robustness to small errors as in Bliss and Panigirtzoglou (1999), but also their ability to closely recover the summary statistics from the true density function given sufficient data.

The paper is organised as follows. Section two sets out the two estimation techniques that we compare. Section three sets out the approach we will use for assessing the performance of the alternative methodologies. Section four presents results for European-style options and section five concludes.

2. Alternative Techniques for Estimating Implied RNDs from Options’ Prices

2.1 Underlying Economics

In this section we examine the two estimation approaches that are tested within this paper. Both may be derived from the Cox and Ross (1976) pricing model. This model gives current time t European-style call option prices as the risk-neutral expected payoff of the option at expiry T , discounted back at the risk-free rate:

$$C(S, X, t) = e^{-rt} \int_X^{\infty} (S_T - X) g(S_T) dS_T \quad (1)$$

¹ See Bahra(1996) for examples of such spiked distributions when using the mixture of lognormals approach.

where S_T is the terminal underlying asset price at T , $g(S_T)$ is its RND, X is the strike price and r and $\tau=T-t$ are the risk-free rate and the maturity of the option respectively. The put price can be recovered either through put-call-parity or by replacing the payoff of the call (S_T-X) with the payoff of the put ($X-S_T$) in the above formula and by integrating from zero to the strike price.

The first estimation approach tested in this paper involves specifying a particular parametric functional form for the RND $g(S_T)$ and fitting this distribution to the observed range of strike prices via non-linear least squares. Although a range of functional forms have been suggested, the most commonly used is a mixture of two lognormals². The form chosen should be sufficiently flexible to capture the features of distributions that we might expect to find implicit within the data - excess kurtosis, either positive or negative skewness, and perhaps bi-modality. The mixture of lognormals is parsimonious because it matches these criteria with just five parameters to be estimated.

The mixture lognormal is given by:

$$g(S_T) = qL(a_1, b_1) + (1-q)L(a_2, b_2) \quad (2)$$

where q, a_1, a_2, b_1, b_2 are the parameters to be estimated. The fitted call and put prices are given by³:

$$\begin{aligned} \hat{C}(S, X_i, t) &= e^{-rt} \int_{X_i}^{\infty} (S_T - X_i) (qL(a_1, b_1) + (1-q)L(a_2, b_2)) dS_T \\ \hat{P}(S, X_i, t) &= e^{-rt} \int_0^{X_i} (X_i - S_T) (qL(a_1, b_1) + (1-q)L(a_2, b_2)) dS_T . \end{aligned} \quad (3)$$

To fit the parameters of the RND we minimise the following:

$$\min_{a_1, b_1, a_2, b_2, q} \sum_{i=1}^m (\hat{C}_{i,t} - C_{i,t})^2 + \sum_{j=1}^n (\hat{P}_{j,t} - P_{j,t})^2 \quad (4)$$

² See Bahra (1996,1997) and Melick and Thomas (1997)

³ As explained in Bahra (1997) for the futures options traded at LIFFE that have futures-style margining, the discount factor disappears.

The second approach to estimating implied RNDs that we test here which we term the “smile” approach, exploits the result derived by Breeden and Litzenberger (1978) that the RND can be recovered by calculating the second partial derivative of the call-pricing function with respect to the strike price. This result can be derived simply by taking the second partial derivative of equation (1) with respect to the strike price to get:

$$\frac{\partial^2 C}{\partial K^2} = e^{-rt} g(S_T) \quad (5)$$

So we just have to adjust up the second partial derivative by $\exp(r\tau)$ to get the RND $g(S_T)$. In practice we only have a discrete set of strike prices. So to obtain an estimate of the continuous call-pricing function we need to interpolate across the discrete set of prices. Following Shimko(1993) this interpolation can be done by interpolating across the volatility smile and using Black-Scholes to transform this back to prices. The reason for doing this rather than interpolating the call-pricing function directly is that it is difficult to fit accurately the shape of the latter. And since we are interested in the convexity of that function any small errors will tend to be magnified into large errors in the final estimated RND.

Shimko (1993) used a quadratic functional form to interpolate across the volatility smile. Instead, we follow Bliss and Panigirtzoglou (1999) and use a cubic smoothing spline to interpolate in a similar way to Campa and Chang (1998). This is a more flexible non-parametric curve that gives us control on the amount of smoothing of the volatility smile, and hence the smoothness of the estimated RND. But following Malz (1997), Bliss and Panigirtzoglou (1999) also first calculate the Black-Scholes deltas of the options and use delta rather than strike to measure the money-ness of options. In practice this makes interpolation of the volatility smile even easier, since it becomes a simpler shape to approximate in “delta-space”. Finally, to generate the implied RND we calculate the second partial derivative with respect to strike price numerically as for (5) and adjust for the effect of the discount factor.

So summarising, estimation via the smile-based approach proceeds by:

- calculating implied volatilities of the call and put options;
- calculating the Black-Scholes deltas of the options using those implied volatilities;

- constructing the volatility smile by joining the implied volatilities for out-of-the-money calls with those of the out-of-the-money puts⁴;
- interpolating across the volatility smile in “delta-space” via a cubic smoothing spline;
- transforming back to a price function using the Black-Scholes model;
- taking the second partial derivation of that function with respect to strike and adjusting for the discount factor within equation (5) to generate the final estimated RND.

3. A Monte Carlo Approach to Testing PDF Estimation Techniques

This section of the paper explains the testing procedures we will use to assess the performance of the two estimation approaches set out above. One approach to testing these techniques is to examine how closely they fit actual options data (for example see the approaches taken by Campa and Chang (1998), Jondeau and Rockinger (1998) and Bliss and Panigirtzoglou (1999). But in doing so it is difficult to assess which of the estimated RNDs most closely match the true risk-neutral density since this is unobservable. In the absence of knowledge of what the true density function is, it is difficult to judge this.

Instead we use simulated artificial options price data. We can simulate options prices that correspond to a given risk-neutral density function and see whether the estimation techniques can recover the RND. In addition following Bliss and Panigirtzoglou (1999), we also test whether the estimation technique is robust to small errors in prices that might result in the real world from the existence of discrete tick size intervals.

Any good RND estimation technique should be able to recover the true RND under a wide range of market conditions: that is conditions of high and low volatility; where the true density function has either positive or negative skews; and where we use options across the full range of maturities that are encountered in practice - anything from one week out to a year. So we need a way of generating options data that match this range of conditions.

⁴ We use out-of-the-money options because traded volumes concentrate on at-the-money and out-of-the-money options. Also the out-of-the-money option value is composed entirely of the time value of the option rather than its intrinsic value as for in-the-money options. It is the time value of the options only that reflects the shape of the RND.

To generate sufficiently interesting “true” risk-neutral densities that incorporate the features discussed above, we use Heston’s (1993) stochastic volatility model to generate prices. For European options, this model has a closed form solution. Under Heston’s model, the underlying asset price dynamics are described by the following stochastic differential equations:

$$\begin{aligned} dS &= \mu S dt + \sqrt{v_t} S dz_1 \\ dv_t &= k(q - v_t) dt + s_v \sqrt{v_t} dz_2 \end{aligned} \quad (12)$$

Here the volatility of the underlying asset $\sqrt{v_t}$ is also stochastic. The conditional variance v_t follows a mean reverting process such that the volatility mean-reverts to a long run of \sqrt{q} at a rate dictated by k . The term s_v sets the volatility of the volatility. Finally, the two Wiener processes dz_1 and dz_2 have a correlation given by r . By changing the correlation parameter we can generate skewness in asset returns. Suppose we have a negative correlation between shocks to the asset price and volatility. This means that as we get negative shocks to the price, volatility will tend to increase. This increase in volatility then increases the chance that we can get further large downwards movements. Thus a negative correlation can generate negative skewness in the unconditional distribution of returns. This will be reflected in a downwards volatility smile in the options generated under these parameters. A positive correlation between volatility and the asset price has the opposite effect⁵.

Heston shows that for European call options⁶ on assets that behave according to (12) it is possible to calculate prices with the following formula:

$$\begin{aligned} C(S, v_t, X) &= SP_1 - Xe^{-rt} P_2 \\ P_j(y, v_t, T; \ln(X)) &= \frac{1}{2} + \frac{1}{P_0} \int_0^\infty \text{Re} \left[\frac{e^{-if \ln(X)} f_i(y, v, T; \mathbf{f})}{i\mathbf{f}} \right] d\mathbf{f} \end{aligned} \quad (13)$$

where X is the strike price, $y = \ln(S)$, $i = \sqrt{-1}$,

⁵ See Das and Sundaresan (1998) for more details on the relation between conditional skewness and kurtosis and the parameters of this stochastic volatility model.

⁶ Put prices can be generated simply via put-call-parity.

$$f_j(y, v_t, t; \mathbf{f}) = e^{C(t, \mathbf{f}) + D(t; \mathbf{V})v_t + i\mathbf{f}y},$$

$$C(t; \mathbf{f}) = r\mathbf{f}t + \frac{a}{\mathbf{s}_v^2} \left\{ (b_j - r\mathbf{s}_v\mathbf{f}t + d)t - 2 \ln \left[\frac{1 - ge^{dt}}{1 - g} \right] \right\},$$

$$D(t; \mathbf{f}) = \frac{b_j - r\mathbf{s}_v\mathbf{f}t + d}{\mathbf{s}_v^2} \left[\frac{1 - e^{dt}}{1 - ge^{dt}} \right],$$

$$g = \frac{b_j - r\mathbf{s}_v\mathbf{f}t + d}{b_j - r\mathbf{s}_v\mathbf{f}t - d},$$

$$d = \sqrt{(\mathbf{s}_v\mathbf{f}t - b_j)^2 - \mathbf{s}_v^2(2u_j\mathbf{f}t - \mathbf{f}^2)}.$$

and $a = \kappa\theta$, $b_1 = \kappa + \lambda - \rho\sigma_v$, $b_2 = \kappa + \lambda$.

To generate the true density function and its associated summary statistics we simply apply equation (5) to (13). Figures 1 and 2 show the effect of changing ρ on the terminal asset price distribution and on the volatility smile for options generated under this model with current and long run volatility of 30%, mean reversion $\kappa=2$ and volatility of volatility σ_v of 40%. We can see that the Heston model can generate the sorts of shapes of both the volatility smile and the underlying asset distributions that can be observed in the real world.

Figure 1: Implied RNDs Under Alternative Correlation Parameters





Figure 2: Volatility Smile Under Alternative Correlation Parameters



An additional feature of the real world that we want to incorporate is the existence of errors that are the result of discrete tick size intervals (and possibly any small violations of arbitrage within the settlement prices used for estimation). We want our estimation methodology to be robust to these small errors in the prices. So we perform the following test of the two RND estimation techniques.

We first establish a set of six scenarios corresponding to low and high volatility and three levels of skewness. For each scenario we generate a set of options prices with strikes ranging from 30% out-of-the-money to 40% in-the-money. Then for each combination of scenario and maturity we use the approach developed by Bliss and Panigirtzoglou (1999) to first shock each price by a random number uniformly distributed from $-1/2$ to $+1/2$ a “tick size”. This tick size was chosen as 0.05 to reflect the sorts of tick sizes that are typically found for exchange-traded options. Given these shocked prices we fit RNDs using the two techniques described in section two and calculate the summary statistics. We repeat this procedure of shocking the prices and then fitting the RNDs 100 times for each scenario and maturity combination. Finally we calculate in each case the mean and standard deviation of the calculated summary statistics and the squared pricing errors. In essence, this technique simply amounts to a monte carlo test of the finite sample properties of the two estimators of the sort that is commonly used within standard econometrics - see Greene (1997) Ch.5 or Davidson and McKinnon (1993) Ch. 18.

We then assess the two techniques by comparing the mean estimated summary statistics with the true summary statistic. We are looking for a technique that has both mean estimates of the statistics that are close to the true ones and one that has small standard deviations for the calculated statistics in the presence of the small errors within the options prices used i.e. it is stable. We also want an estimation procedure that performs well across the range of scenarios and maturities. The next section performs these tests for European-style options.

4. Results

This section includes the results of the tests that we described above for the two estimation approaches. As described above, we test performance across a range of six scenarios. The Heston model parameters used for each scenario are set out in table 1 below. These were chosen to generate true RNDs that corresponded to situations of negative skewness, and weak and strong positive skewness in the terminal asset price and also conditions of low and high volatility. To generate these

differing levels of skewness in the terminal asset price distributions, we use three different levels of the correlation parameter -0.9, 0 and 0.9. The long run volatilities of 30% for the high volatility scenarios were chosen on the basis of the levels of implied volatility typically seen within equity markets. The low volatility (10%) scenarios can be thought of as consistent with levels often seen within FX and interest rate markets⁷.

Table 1: Model parameters used under each scenario

	Strong Negative Skew		Strong Positive Skew
Low Volatility	Scenario 1 $\kappa=2$, $\theta^2=0.1$, $\sigma_v=0.1$, $\rho=-0.9$	Scenario 2 $\kappa=2$, $\theta^2=0.1$, $\sigma_v=0.1$, $\rho=0$	Scenario 3 $\kappa=2$, $\theta^2=0.1$, $\sigma_v=0.1$, $\rho=0.9$
High Volatility	Scenario 4 $\kappa=2$, $\theta^2=0.3$, $\sigma_v=0.4$, $\rho=-0.9$	Scenario 5 $\kappa=2$, $\theta^2=0.3$, $\sigma_v=0.4$, $\rho=0$	Scenario 6 $\kappa=2$, $\theta^2=0.3$, $\sigma_v=0.4$, $\rho=0.9$

We test the performance of the two estimation techniques under each of these scenarios across four different maturities - 2 weeks, 1 month, 3 months and 6 months. For each scenario and maturity pairing we first generate the true RND and calculate their summary statistics - their mean, standard deviation, skewness (the third central moment) and kurtosis (the fourth central moment). Table 2 sets out the true summary statistics for all the maturity and scenario combinations.

⁷ We also assume that the market price of volatility risk is zero and that the time t conditional volatility is equal to the long run volatility.

Table 2: True Summary Statistics					
Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	100.000	100.000	100.000	100.000
	2	100.000	100.000	100.000	100.000
	3	100.000	100.000	100.000	100.000
	4	100.000	99.999	99.999	99.994
	5	99.999	100.000	99.999	99.998
	6	99.988	99.997	99.995	99.961
Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	1.958	2.878	4.956	6.966
	2	1.962	2.888	5.004	7.081
	3	1.965	2.898	5.052	7.201
	4	5.849	8.555	14.524	20.099
	5	5.888	8.677	15.093	21.485
	6	5.921	8.799	15.687	22.900
Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	-0.198	-0.280	-0.418	-0.474
	2	0.060	0.089	0.159	0.231
	3	0.318	0.459	0.743	0.956
	4	-0.166	-0.228	-0.301	-0.265
	5	0.180	0.272	0.504	0.756
	6	0.523	0.776	1.346	1.840
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	3.037	3.081	3.178	3.221
	2	3.038	3.086	3.221	3.355
	3	3.160	3.344	3.930	4.599
	4	2.984	2.962	2.878	2.748
	5	3.119	3.269	3.809	4.616
	6	3.426	4.036	6.272	9.026

For all combinations the futures price has been set at 100, so the true mean of the distributions are equal to 100⁸. As we would expect, the standard deviation of the true RNDs increases with maturity and as volatility is increased. Scenarios 1 and 4 which have a negative correlation between the underlying asset price and volatility display negative skewness in the terminal asset price distribution. Except for scenario four⁹, the kurtosis of the terminal asset price distribution is greater than three and increases with maturity.

⁸ The slight differences from 100 are caused by error in the numerical integration used to calculate the summary statistics.

⁹ For this scenario the combination of a high volatility of volatility and the negative correlation between volatility and the asset price appears to reduce the probabilities attached to extreme outcomes.

Given these true summary statistics, we proceed to test the two estimation approaches to recover them from simulated options prices. We use equation (13) to generate European-style call and put futures options for all the scenario and maturity pairings for strikes ranging from 70 to 140 with strikes spaced at intervals of one apart from each other. For each pairing we then test the two techniques by shocking the prices for the tick size errors as described above and estimating RNDs and their summary statistics. This is repeated a hundred times for each pairing. Then for each pairing, and each summary statistic we calculate a measure of the mean estimate and measures of the estimates' stability (the standard deviation, and the distance between the five and ninety five percentiles) from the sets of a hundred estimated summary statistics under each technique. A good technique should have a mean estimate of each of the summary statistics that is close to the true ones and so may be said to be unbiased. A low standard deviation of the estimated summary statistics across the full range of scenarios and maturities indicates that the estimation technique is stable in the presence of small errors within the prices.

Table 3 below gives the mean estimated summary statistics for both approaches across all the different scenario and maturity pairings. To assess the unbiasedness of the two estimators, however, we are interested in the difference between the true summary statistics and the mean estimates from each approach. So table 4 calculates the difference between the true and the mean estimated statistics as a percentage of the true value of the summary statistic: $(\text{true} - \text{mean}) / \text{true}$.

Examining the top two panels of table 4 we can see that the mean of the smile approach is always almost always exactly equal to the true mean. The reason for this accuracy is that when we transform from the volatility smile to the pricing function using Black-Scholes this constrains the mean of the RND to be equal to the forward rate. The mixture lognormal estimation as described by equation (4) does not explicitly constrain the mean of the RND to be equal to the forward rate, so we get small errors between the true and actual means. This could be eliminated if we added an extra constraint to (4) to ensure that the mean equalled the forward rate as described in Bahra (1997), but this may come at the cost of extra instability in the fitted RNDs.

Table 3: Mean Estimated Summary Statistics											
Smile Technique					Mixture Lognormal Technique						
Mean	Scenario	2 weeks	1 month	3 month	6 month	Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	100.0000	100.0000	100.0000	100.0000		1	100.1680	100.0450	100.0312	100.0525
	2	100.0000	100.0000	100.0000	100.0000		2	99.9950	99.8531	99.8428	100.0006
	3	100.0000	100.0000	100.0000	100.0000		3	101.9543	100.4185	100.0332	99.9818
	4	100.0000	100.0000	100.0000	100.0000		4	99.8328	99.6520	100.7073	99.8851
	5	100.0000	100.0000	100.0000	100.0000		5	100.0027	100.0113	100.0024	100.0000
6	100.0000	100.0000	100.0000	100.0000	6	100.2045	98.8045	99.9303	99.8592		
Std Dev	Scenario	2 weeks	1 month	3 month	6 month	Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	1.9595	2.8787	4.9588	6.9676		1	2.6334	2.8864	4.9503	6.9544
	2	2.0615	2.9891	5.0970	7.1651		2	1.9732	2.9732	5.0552	7.1094
	3	1.9685	2.9035	5.0581	7.2073		3	7.4278	3.7705	5.0581	7.1893
	4	5.8485	8.5588	14.5087	20.0582		4	5.8645	8.5775	15.0537	20.1169
	5	5.8905	8.6778	15.0693	21.3463		5	5.8861	8.6659	15.0867	21.4762
6	5.9316	8.8035	15.5988	22.6372	6	5.9116	9.4515	15.5956	22.5237		
Skewness	Scenario	2 weeks	1 month	3 month	6 month	Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	-0.1848	-0.2654	-0.3915	-0.4466		1	-0.0585	-0.2709	-0.3695	-0.4148
	2	0.1968	0.1898	0.2081	0.2548		2	0.0594	0.1232	0.2165	0.2408
	3	0.3132	0.4421	0.7130	0.9128		3	0.1855	0.4236	0.6167	0.8436
	4	-0.1521	-0.2076	-0.2578	-0.1960		4	0.1764	0.2543	-0.2479	-0.1381
	5	0.1784	0.2692	0.4798	0.6825		5	0.1776	0.2603	0.4995	0.7504
6	0.5124	0.7505	1.2121	1.5861	6	0.1821	0.7642	1.1701	1.4949		
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month	Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	3.0163	3.0363	3.0701	3.0744		1	3.0942	3.0672	3.0629	3.0668
	2	3.8713	3.6523	3.4221	3.3638		2	3.0720	3.0711	3.1870	3.4901
	3	3.1615	3.2812	3.6934	4.1546		3	2.3902	3.2004	3.5323	4.0006
	4	2.9720	2.9415	2.8054	2.6707		4	3.0554	3.1175	2.7957	2.7313
	5	3.0781	3.1813	3.5137	3.9345		5	3.0848	3.1811	3.7677	4.5488
6	3.3829	3.8291	5.1522	6.8006	6	3.2202	3.6502	4.9617	6.0907		

The second set of panels gives the results for the estimated standard deviations. For scenarios 4 to 6 and for all scenarios with maturities of three months and above the mean estimates are close to the true standard deviations for both techniques. In most of these cases the mean errors are less than 1%. For scenarios 1 and 3 and for the two week and one month maturities however, the mixture lognormal appears to perform significantly worse. For low times to maturity and low volatility, the mixture lognormal over-estimates the true standard deviation on average.

The results for the higher moments are much more variable. The absolute size of the mean errors as a proportion of the true statistic are much higher than for the first two moments. For skewness, these figures partly over-state the problems, however, because the true skewness is close to zero for at least scenarios 2 and 5. Compared to the mixture lognormal, the smile-based technique has less biased results for skewness under scenarios 1, 3, 4 and 6 - those that display more extreme levels of skewness. But for scenarios 2 and 5 the smile-based approach does better for some maturities but worse than the mixture lognormal for others. Like skewness, the mean errors of the kurtosis estimates are larger and more variable than for the mean or standard deviation. Broadly, the mixture lognormal

mean estimates are poorest for scenarios 3 and 6 in which skewness is strongest. The smile-based approach has the poorest results for scenarios 5 and 6 when maturity is three months or above.

Table 4: Difference Between True and Mean Estimated Summary Statistics (as % of the True)											
Smile Technique						Mixture Lognormal Technique					
Mean	Scenario	2 weeks	1 month	3 month	6 month	Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	0.00%	0.00%	0.00%	0.00%		1	-0.17%	-0.05%	-0.03%	-0.05%
	2	0.00%	0.00%	0.00%	0.00%		2	0.00%	0.15%	0.16%	0.00%
	3	0.00%	0.00%	0.00%	0.00%		3	-1.95%	-0.42%	-0.03%	0.02%
	4	0.00%	0.00%	0.00%	-0.01%		4	0.17%	0.35%	-0.71%	0.11%
	5	0.00%	0.00%	0.00%	0.00%		5	0.00%	-0.01%	0.00%	0.00%
6	-0.01%	0.00%	0.00%	-0.04%	6	-0.22%	1.19%	0.07%	0.10%		
Std Dev	Scenario	2 weeks	1 month	3 month	6 month	Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	-0.05%	-0.04%	-0.05%	-0.03%		1	-34.47%	-0.31%	0.12%	0.16%
	2	-5.09%	-3.51%	-1.87%	-1.19%		2	-0.58%	-2.96%	-1.03%	-0.40%
	3	-0.18%	-0.19%	-0.12%	-0.09%		3	-277.99%	-30.11%	-0.12%	0.16%
	4	0.01%	-0.05%	0.10%	0.20%		4	-0.26%	-0.27%	-3.65%	-0.09%
	5	-0.04%	0.00%	0.16%	0.65%		5	0.03%	0.13%	0.04%	0.04%
6	-0.18%	-0.05%	0.56%	1.15%	6	0.15%	-7.42%	0.59%	1.64%		
Skewness	Scenario	2 weeks	1 month	3 month	6 month	Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	6.83%	5.39%	6.30%	5.81%		1	70.51%	3.44%	11.55%	12.51%
	2	-229.59%	-113.88%	-30.89%	-10.39%		2	0.51%	-38.86%	-36.15%	-4.36%
	3	1.57%	3.72%	3.97%	4.51%		3	41.69%	7.73%	16.95%	11.75%
	4	8.51%	8.82%	14.22%	26.03%		4	206.12%	211.70%	17.50%	47.91%
	5	0.90%	1.20%	4.84%	9.77%		5	1.35%	4.45%	0.93%	0.80%
6	1.93%	3.24%	9.93%	13.82%	6	65.16%	1.46%	13.05%	18.78%		
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month	Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	0.68%	1.44%	3.41%	4.54%		1	-1.89%	0.43%	3.64%	4.78%
	2	-27.42%	-18.34%	-6.23%	-0.27%		2	-1.11%	0.49%	1.07%	-4.03%
	3	-0.04%	1.89%	6.02%	9.65%		3	24.36%	4.30%	10.11%	13.00%
	4	0.40%	0.70%	2.53%	2.81%		4	-2.39%	-5.24%	2.86%	0.60%
	5	1.31%	2.68%	7.75%	14.76%		5	1.09%	2.69%	1.08%	1.45%
6	1.27%	5.14%	17.86%	24.66%	6	6.02%	9.57%	20.90%	32.52%		

On the basis of these tests for the ability on average to estimate the true summary statistics, it is not immediately obvious that one of the techniques is better than the other. The smile-based approach does appear to do marginally better in estimating the first two moments particularly at short maturities with low volatility. For the third and fourth moments, however, neither technique obviously outperforms the other.

But when we look at the stability of the estimates, the story is far more clear cut. Table 5 sets out the standard deviations of the estimated summary statistics. High standard deviations of the summary statistics are indicative of instability in the estimated RNDs. For nearly all the scenarios the mixture lognormal has much higher standard deviations of the estimates for all statistics than for the smile-based approach. This mirrors Bliss and Panigirtzoglou's (1999) findings that the mixture lognormal is unstable using actual options data for FTSE and 3 month sterling interest rates. In particular, the

mixture lognormal appears to perform badly under scenarios 1 and 3 and to some extent scenario 2, when maturities of the options are one month or less. This suggests that the mixture lognormal approach is unstable when volatility is low, the maturity of the options is low and when there is a strong negative or positive skew.

Table 5: Standard Deviation of Summary Statistics											
Smile Technique						Mixture Lognormal Technique					
Mean	Scenario	2 weeks	1 month	3 month	6 month	Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	0.0000	0.0000	0.0000	0.0000		1	4.2206	0.2585	0.0198	0.0176
	2	0.0000	0.0000	0.0000	0.0000		2	0.1927	0.9256	0.7257	0.0149
	3	0.0000	0.0000	0.0000	0.0000		3	11.9560	6.2271	0.3363	0.0198
	4	0.0000	0.0000	0.0000	0.0000		4	0.0133	0.0390	5.7981	2.4340
	5	0.0000	0.0000	0.0000	0.0000		5	0.0121	0.0137	0.0185	0.0082
	6	0.0000	0.0000	0.0000	0.0000		6	0.0859	4.2933	0.0153	0.0093
Std Dev	Scenario	2 weeks	1 month	3 month	6 month	Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	0.0123	0.0110	0.0088	0.0091		1	4.3470	0.0730	0.0102	0.0107
	2	0.0144	0.0137	0.0112	0.0094		2	0.0693	0.5684	0.2252	0.0126
	3	0.0139	0.0123	0.0112	0.0100		3	12.2235	7.5644	0.0994	0.0116
	4	0.0093	0.0095	0.0062	0.0063		4	0.0099	0.0093	5.4027	0.5017
	5	0.0104	0.0080	0.0075	0.0065		5	0.0117	0.0092	0.0135	0.0276
	6	0.0097	0.0079	0.0080	0.0068		6	0.0111	2.4101	0.0088	0.0088
Skewness	Scenario	2 weeks	1 month	3 month	6 month	Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	0.0204	0.0192	0.0130	0.0085		1	0.2348	0.1663	0.0248	0.0142
	2	0.0201	0.0234	0.0104	0.0068		2	0.1721	0.2341	0.2684	0.0165
	3	0.0191	0.0166	0.0106	0.0080		3	0.7527	0.1899	0.1975	0.0130
	4	0.0096	0.0064	0.0030	0.0021		4	0.0003	0.0458	0.0644	0.2577
	5	0.0091	0.0061	0.0035	0.0027		5	0.0015	0.0055	0.0115	0.0167
	6	0.0102	0.0066	0.0038	0.0028		6	0.0423	0.1839	0.0049	0.0037
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month	Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	0.0175	0.0156	0.0141	0.0100		1	1.5428	0.1002	0.0532	0.0335
	2	0.0645	0.0333	0.0163	0.0101		2	0.0832	0.1835	0.0827	0.1190
	3	0.0517	0.0296	0.0215	0.0189		3	0.9664	0.3374	0.2073	0.0565
	4	0.0092	0.0065	0.0035	0.0022		4	0.0002	0.0185	0.1992	0.2575
	5	0.0100	0.0078	0.0076	0.0069		5	0.0239	0.0274	0.0677	0.1572
	6	0.0157	0.0150	0.0143	0.0139		6	0.0271	0.3777	0.0249	0.0227

Table 5 indicates that the “smile”-based estimation is far more stable than the mixture lognormal approach. But how does the instability of the mixture lognormal estimates manifest itself? To see this we look in detail at the estimated RNDs under scenario 3, one of the scenarios in which mixture lognormal is most unstable, and compare them to the true RND. Figures 3,4 and 5 compare the true RND with estimated RNDs using the two techniques for maturities of two weeks, 1 month and 3 months respectively. In each figure, the top panel displays the true RND, while the bottom two panels each display thirty of the RNDs estimated from the previous tests for the smile-based approach and the mixture lognormal technique.

It is immediately clear that the smile-based RND estimates are far more stable than the mixture lognormal RNDs. The former match the shapes of the true RND closely, particularly for the longer maturities. At the two week maturity there is greater variation in the fitted RNDs but this is to be expected given that the tick size errors that are added to the prices will have a greater proportionate impact on the time value of these shorter maturity options.

The mixture lognormal distributions are highly unstable at the two week and one month maturities. The most common - but not only - cause of this instability is the existence of “spikes” in the distribution. The spikes occur when the variance of one of the distributions collapses. The mixture of two distributions then looks like a single lognormal distribution with a spike, usually towards the centre. Clearly such a spike is not contained within the true distribution and reflects estimation errors. As the maturity of the options increases, the proportion of these spiked distributions falls. In addition to these spiked cases, there are a few mixture lognormal RNDs which are not spiked but which display skewness which is quite different to the true RND.




Examining further the one month mixture lognormal distributions, we can see that for a larger proportion of cases, the mixture lognormal technique manages to get a close fit to the true distribution than at two weeks. As we move to three months we get a higher proportion still of fitted mixture lognormal RNDs that are close to the true distribution. At this maturity, the optimisation used to fit the RND appears to be flipping between between two minima - one which closely matches the true RND and one which contains a spike and hence severely mis-estimates the RND. It is the existence of these spiked distributions that causes the increased standard deviation of the estimated summary statistics compared with the smile-based approach.

What appears to be the key difference between the two estimation approaches is that the small errors in the prices cause only small local errors in the estimated RNDs under the smile approach, while for the mixture lognormal non-linear least squares estimation, the errors can be sufficient for the minimisation to reach very different parameter estimates with large changes in the shape of the estimated RND as a result. The end result of this is that while the bias of the mixture lognormal estimator does not appear to be much larger than the smile-based estimator (at least for the third and fourth moments) it is far more unstable. Since in practice we are often concerned with changes in the

PDF from one day to another this instability is a concern and reduces the value of the mixture lognormal technique as a practical tool.

Figure 3

True RND, Scenario 3, 2 Weeks Maturity

**"Smile" Estimated RNDs - Scenario 3, 2 Weeks Maturity****Mixture Lognormal Estimated RNDs - Scenario 3, 2 Weeks Maturity**

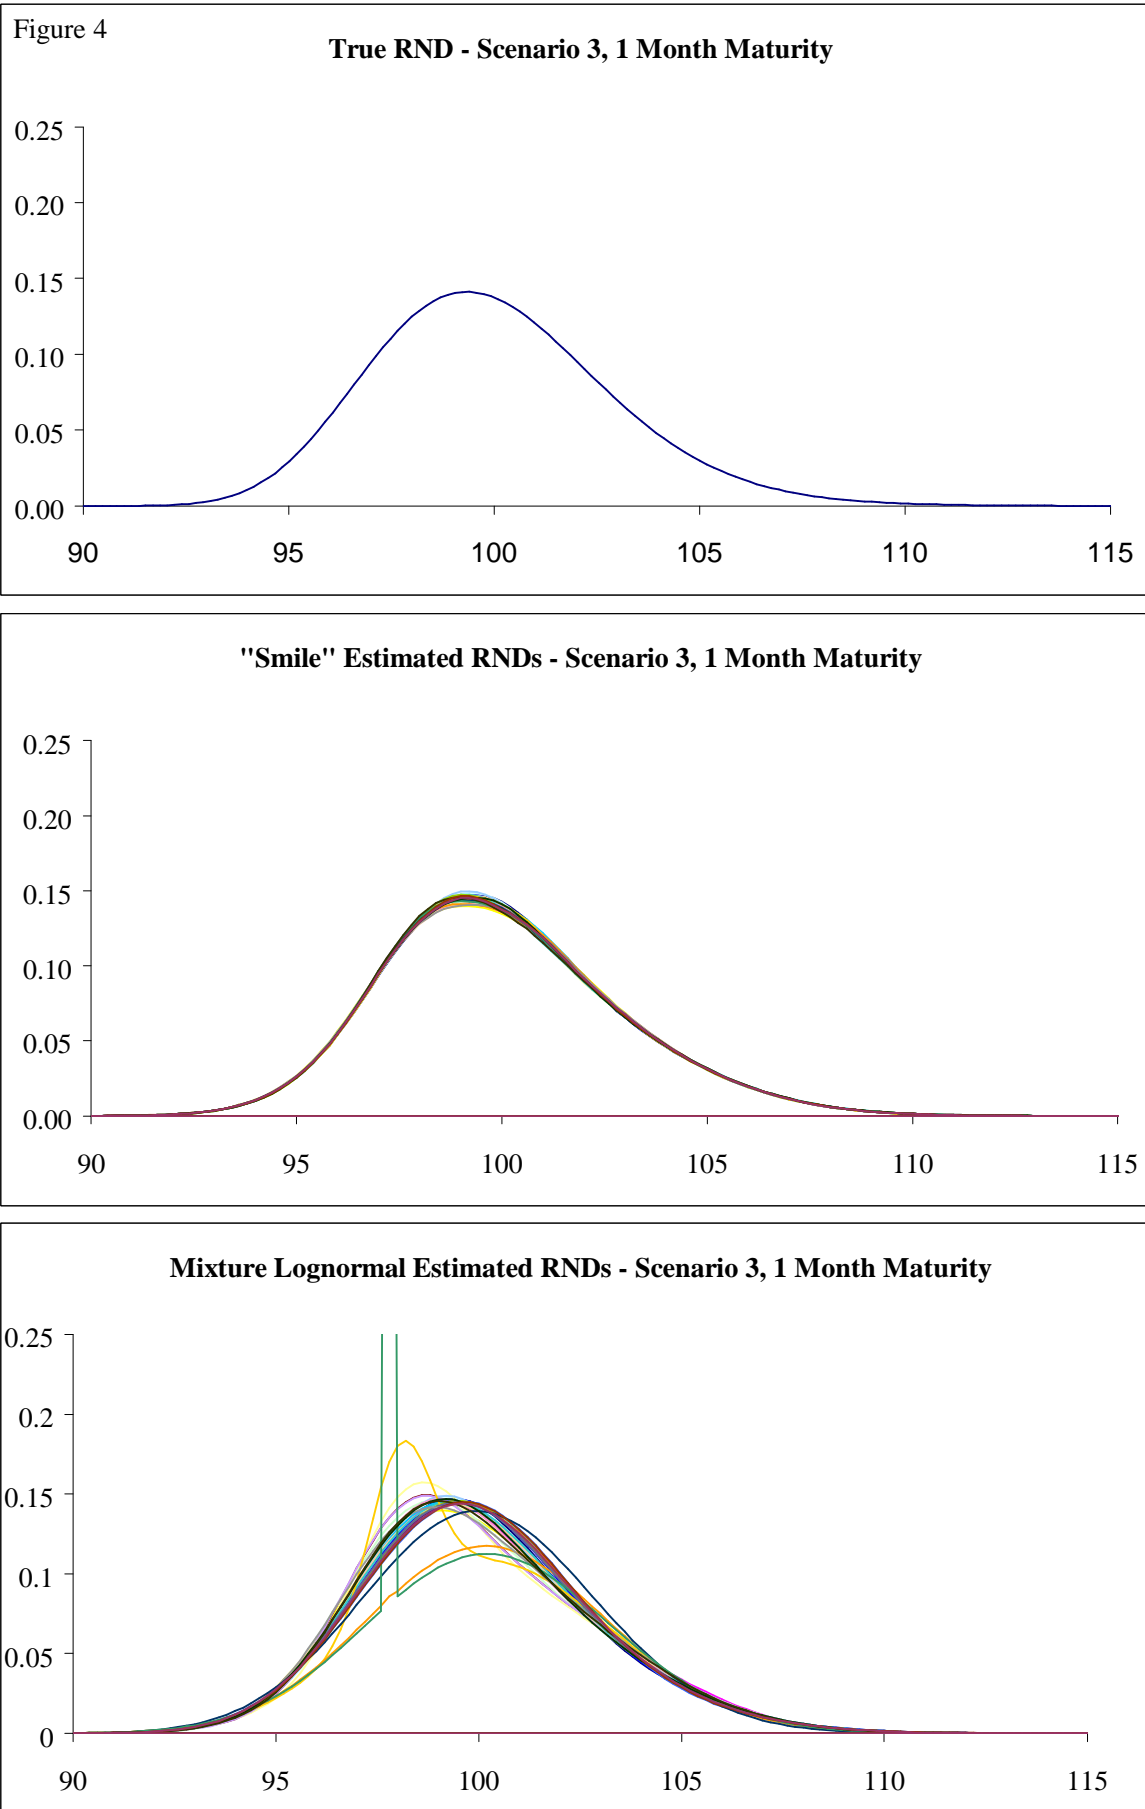





Figure 5

True RND - Scenario 3, 3 Months Maturity**"Smile" Estimated PDFs - Scenario 3, 3 Months Maturity****Mixture Lognormal Estimated RNDs - Scenario 3, 3 Months Maturity**

5. Assessing Alternative Approaches

This paper has examined two alternative approaches to estimating implied RNDs from European-style options. The first was the commonly used mixture lognormal approach which uses non-linear least squares estimation to fit a parametric form to observed options prices. The second approach interpolated across the volatility smile using a cubic smoothing spline and then employed the Breeden and Litzenberger result to recover the RND by calculating the second partial derivative of the call pricing function with respect to the strike price.

The monte carlo tests of the two estimators in section four demonstrated that the second “smile”-based approach performed a little better in terms of its ability to match the first two moments of the true RND. We also saw that the higher moments appear to be much more difficult to estimate accurately with both techniques often resulting in estimates that are on average quite a long way from the true ones.

But we also observed that the smile-based technique was far more stable than the mixture lognormal approach. The latter technique has severe mis-estimation problems when using options on low volatility assets or when using low maturity (less than three months) options. This mis-estimation most often shows up as a “spiked” distribution when one of the lognormal distribution’s estimated variance falls to a very low level. In contrast the “smile”-based estimation appears to perform well across all scenarios and maturities (although the existence of discrete tick size errors does create increased instability at maturities below 1 month). These results suggest the use of the smile-based approach over the mixture lognormal by practitioners and researchers alike. They also suggest that where the mixture lognormal is still used that the results have to be interpreted with great caution given the tendency of the estimation approach to severely mis-estimate the true RND.

Future work at the Bank will use the monte carlo tests set out here to examine the empirical performance of RND estimators that use American-style options. As well as examining the accuracy and stability of the techniques, this forthcoming work will examine how important it is to take account of the early exercise premium when using these options.

Bibliography

Bates, D (1996), “Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options”, *Review of Financial Studies*, Vol 9, No. 1 pp 69-107

Bahra, B. (1996), “Probability Distributions of Future Asset Prices Implied By Option Prices”, *Bank of England Quarterly Bulletin*, August 1996, pages 299-311.

Bahra, B. (1997), “Implied Risk-Neutral Probability Density Functions from Options Prices: Theory and Application”, *Bank of England Working Paper* No.66

Bliss, R and Panigirtzoglou, N (1999). “Testing the Stability of Implied PDFs”, *Bank of England mimeo*

Breeden, D.T., and R.H. Litzenberger (1978), “Prices of State-Contingent Claims Implicit in Option Prices”, *Journal of Business*, 51, 4, 621-651.

Campa, J., Chang, K. and Reider, R. (1998) “Implied Exchange Rate Distributions: Evidence From OTC Option Markets”, *Journal of International Money and Finance*

Cox, J.C. and Ross, S.A (1976) “The Valuation of Options for Alternative Stochastic Processes”, *Journal of Financial Economics* 3, 145-66

Das, S and Sundaresan, R. (1998), “Of Smiles and Smirks: A Term Structure Perspective”, *Working Paper*

Davidson and McKinnon (1993) “Estimation and Inference in Econometrics” Oxford University Press

Heston, S (1993), “A Closed-Form Solution for Options with Stochastic Volatility With Applications to Bond and Currency Options”, *The Review of Financial Studies*, Vol 6, No 2, pp 327-343.

Greene, W. (1997) “Econometric Analysis”, Prentice Hall

Jondeau, E. and Rockinger, M (1998), “Reading the Smile: The Message Conveyed By Methods Which Infer Risk Neutral Densities”, *CEPR Discussion Paper* No 2009.

Kurpiel, A. and Roncalli, T. (1999) “Hopscotch Methods For Two-State Financial Models”, *FERC Working Paper, City University*

Malz, A (1995), “Using Options Prices to Estimate ReAlignment Probabilities in the European Monetary System”, *Federal Reserve Bank of New York Staff Reports*, No. 5

Malz, A (1997), “Estimating the Probability Distribution of Future Exchange Rates from Options Prices”, *Journal of Derivatives*

Melick and Thomas (1997), “Recovering an Asset’s Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis”, *Journal of Financial and Quantitative Analysis*, Vol 32 No1 March 1997

Shimko (1993), “Bounds of Probability”, *RISK*, Vol 6, No.4.


Discussion of Neil Cooper's paper:
TESTING TECHNIQUES FOR ESTIMATING IMPLIED RNDs
FROM THE PRICES OF EUROPEAN-STYLE OPTIONS

by Holger Neuhaus¹

What I have to say requires fifteen minutes. I have ten. So, fasten your seat belts, we are about to take off.

The motivation for Neil Cooper's paper is comparably old: in 1995 when I was just finishing my research paper I was talking a lot to Bhupi Bahra at the Bank of England and he was uncertain whether to implement Shimko's smile approach or Melick's and Thomas' mixture of lognormals to derive implied probability density functions. Well - you all know what the Bank eventually decided, but now Neil Cooper tries to investigate in an objective way which of the two methods is the better one and he does that by the following methodology:

Fig. 1 True risk-neutral density function




Neil imposes probability density distributions like the one in figure 1 – although his are fancier, of course. He then uses this function to derive corresponding option prices which, in turn, serve as input for the models for backing out the *implied* probability distribution. Eventually, the results can be compared with the imposed distributions he started with. Moreover Neil shocks the option prices by “half a tick size” to find out how sensitive the models are to inaccurate option prices, the inaccuracy being caused by discrete tick-sizes.

¹ The views expressed represent exclusively the opinion of the author and do not necessarily correspond to those of the European Central Bank.

All in all, this is a promising approach as it is the only method that allows a comparison of the estimates with a known probability density function. However, a relevant question is whether the implementation of the method is close enough to reality to allow for a fair comparison of the different approaches for backing out implied probabilities.

In this context, one observation frequently made in the real world is that one may not always have a sufficient number of strikes and option prices to cover the whole distribution – as illustrated in figure 2.

Fig. 2 True risk-neutral density function with missing strikes




This presentation is of course an exaggeration, but it drives home my point that if only a part of the density function is backed out, it can be difficult to correctly allocate the missing probability mass. When looking at the smile technique, should one, for example, extrapolate the smile or keep the implied volatility constant at the tails?

Some models have more problems with these errors than others do and Neil's dataset is "too" complete and thus leaves out an important aspect in the evaluation of different approaches chosen. As a matter of fact, this is one issue where, as we have already discussed earlier today, the same model can yield (slightly) different results when implemented in a (slightly) different manner.

In this context I also keep mentioning that one means for deciding on how to allocate the missing probability mass is not to estimate the probability density function but to estimate the first derivative of the option price with respect to the strike price, i.e. the implied (*cumulative*) probability distribution. If the lowest and highest cumulative probability derived from the data are, say, 1% and 98% respectively, it can be inferred that 1% of the probability mass is missing at the lower end and 2% at the right tail. This could, in particular, be valid for the smile technique.

Fig. 3 True risk-neutral density and cumulative distribution function with missing strikes



When comparing different approaches to recover implied probabilities, it is also important to mimic realistically another feature of the relevant option market: the number of and distance between strike prices. On the foreign exchange market, for example, few strike prices exist, while, say, for short-term interest rates, some derivative exchanges provide a large number of options. The number of, and distance between, strike prices are determined by certain rules aiming at striking a balance between having a choice between a large number of strike prices and liquid option contracts. Looking at Neil's paper, I have the impression that the intervals chosen between the strike prices are comparably close, which could, in my view, favour the smile approach.

Anyway, looking at Neil's paper more closely, in particular at the summary statistics he generated after he had backed out the (shocked) implied probability density 100 times for each of his scenarios/maturities, one can conclude the following. For facilitating the comparison, I categorised the

results as g or b. G is green and good, b is brown and bad², depending on the relative performance of the different methods' estimates.

Table 1: Mean of estimates

Table 3: Mean Estimated Summary Statistics					
Smile Technique					
Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	100.000	100.000	100.000	100.000
	2	100.000	100.000	100.000	100.000
	3	100.000	100.000	100.000	100.000
	4	100.000	100.000	100.000	100.000
	5	100.000	100.000	100.000	100.000
	6	100.000	100.000	100.000	100.000
Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	1.9595	2.8787	4.9588	6.9676
	2	2.0615	2.9891	5.0970	7.1651
	3	1.9685	2.9035	5.0581	7.2073
	4	5.8485	8.5588	14.5087	20.0582
	5	5.8905	8.6778	15.0693	21.3463
	6	5.9316	8.8035	15.5988	22.6372
Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	-0.1848	-0.2654	-0.3915	-0.4466
	2	0.1968	0.1898	0.2081	0.2548
	3	0.3132	0.4421	0.7130	0.9128
	4	-0.1521	-0.2076	-0.2578	-0.1960
	5	0.1784	0.2692	0.4798	0.6825
	6	0.5124	0.7505	1.2121	1.5861
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	3.0163	3.0363	3.0701	3.0744
	2	3.8713	3.6523	3.4221	3.3638
	3	3.1615	3.2812	3.6934	4.1546
	4	2.9720	2.9415	2.8054	2.6707
	5	3.0781	3.1813	3.5137	3.9345
	6	3.3829	3.8291	5.1522	6.8006
Mixture Lognormal Technique					
Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	100.168	100.045	100.031	100.053
	2	99.995	99.853	99.843	100.001
	3	101.954	100.419	100.033	99.982
	4	99.833	99.652	100.707	99.885
	5	100.003	100.011	100.002	100.000
	6	100.204	98.804	99.930	99.859
Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	2.6334	2.8864	4.9503	6.9544
	2	1.9732	2.9732	5.0552	7.1094
	3	7.4278	3.7705	5.0581	7.1893
	4	5.8645	8.5775	15.0537	20.1169
	5	5.8861	8.6659	15.0867	21.4762
	6	5.9116	9.4515	15.5956	22.5237
Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	-0.0585	-0.2709	-0.3695	-0.4148
	2	0.0594	0.1232	0.2165	0.2408
	3	0.1855	0.4236	0.6167	0.8436
	4	0.1764	0.2543	-0.2479	-0.1381
	5	0.1776	0.2603	0.4995	0.7504
	6	0.1821	0.7642	1.1701	1.4949
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	3.0942	3.0672	3.0629	3.0668
	2	3.0720	3.0711	3.1870	3.4901
	3	2.3902	3.2004	3.5323	4.0006
	4	3.0554	3.1175	2.7957	2.7313
	5	3.0848	3.1811	3.7677	4.5488
	6	3.2202	3.6502	4.9617	6.0907

Indeed, the smile outperforms the mixture of lognormals at most occasions. However, the difference in quality between the models is not that clear-cut. Sometimes the smile yields better results but only marginally and the question has to be addressed: how big is big, i.e. when are results significantly different. In this context, I would like to mention that some of the percentage mistakes presented in the paper are misleading, in particular those where the benchmark value is close to zero. A point of greater substance is that, as we had already discussed today, the measurement of the skewness and kurtosis by calculating the 3rd and 4th moment is not recommendable but could be replaced by other measures.

Looking at the expected value of the smile technique, I would like to note that, even when using the shocked option prices, the smile leads to more reliable estimates of the true distribution than the figures provided in what Neil calls the “true” summary statistics of the imposed probability distributions. Because of the numerical integration method used, the “true” mean is slightly different from 100, while the smile always generates an expected value of 100. That looks very - or even too - robust to me.

And indeed, looking at the standard deviation of the estimated summary statistics (mean, standard deviation, skewness, and kurtosis), the smile looks very robust and fares better than the mixture of lognormals approach.

² In black and white copies the brown fields are the darker ones.

Table 2: Standard deviation of estimates

Smile Technique						Mixture Lognormal Technique					
Mean	Scenario	2 weeks	1 month	3 month	6 month	Mean	Scenario	2 weeks	1 month	3 month	6 month
	1	0.0000	0.0000	0.0000	0.0000		1	4.220623	0.258529	0.019781	0.01758
	2	0.0000	0.0000	0.0000	0.0000		2	0.19266	0.925638	0.725677	0.014939
	3	0.0000	0.0000	0.0000	0.0000		3	11.95599	6.227077	0.336293	0.019803
	4	0.0000	0.0000	0.0000	0.0000		4	0.013332	0.03904	5.798117	2.433972
	5	0.0000	0.0000	0.0000	0.0000		5	0.012113	0.01368	0.01845	0.008227
	6	0.0000	0.0000	0.0000	0.0000		6	0.085873	4.29325	0.015288	0.009329
Std Dev	Scenario	2 weeks	1 month	3 month	6 month	Std Dev	Scenario	2 weeks	1 month	3 month	6 month
	1	0.012303	0.010959	0.008757	0.009063		1	4.347047	0.072974	0.010168	0.010725
	2	0.014392	0.013679	0.011181	0.009371		2	0.069271	0.568389	0.225205	0.012629
	3	0.013898	0.012333	0.011182	0.010014		3	12.22352	7.564441	0.09944	0.011641
	4	0.009302	0.009502	0.006178	0.006323		4	0.009903	0.009298	5.402745	0.501706
	5	0.01037	0.008047	0.007502	0.006537		5	0.011721	0.009154	0.013471	0.027603
	6	0.009686	0.007899	0.008014	0.006796		6	0.011146	2.410094	0.008808	0.008811
Skewness	Scenario	2 weeks	1 month	3 month	6 month	Skewness	Scenario	2 weeks	1 month	3 month	6 month
	1	0.020438	0.019151	0.012986	0.008499		1	0.234815	0.166297	0.024775	0.014154
	2	0.020113	0.023413	0.010409	0.006759		2	0.172144	0.234108	0.268442	0.016485
	3	0.019066	0.016583	0.010591	0.008017		3	0.752712	0.189888	0.1975	0.012989
	4	0.009557	0.006367	0.00304	0.002144		4	0.000301	0.045771	0.064426	0.25774
	5	0.00911	0.006055	0.003462	0.002719		5	0.00151	0.005481	0.011528	0.016678
	6	0.0102	0.006642	0.003844	0.002834		6	0.042258	0.183944	0.004942	0.003734
Kurtosis	Scenario	2 weeks	1 month	3 month	6 month	Kurtosis	Scenario	2 weeks	1 month	3 month	6 month
	1	0.017536	0.015553	0.014076	0.010039		1	1.542771	0.10017	0.053217	0.03349
	2	0.064452	0.033344	0.016313	0.010056		2	0.083219	0.183497	0.082698	0.119001
	3	0.051657	0.029554	0.021507	0.018863		3	0.966391	0.337422	0.207318	0.056505
	4	0.009222	0.006492	0.003471	0.00217		4	0.00019	0.018513	0.199159	0.257488
	5	0.010033	0.007826	0.00757	0.006854		5	0.023942	0.027396	0.067714	0.157168
	6	0.015728	0.015041	0.014292	0.013867		6	0.027087	0.377703	0.024928	0.022667

Again, the issue is: how big is big? And: why is the smile approach looking so good? One reason is, of course, its design: option prices are translated into implied volatilities, a smooth function is calculated that fits the smile (in delta space) and is used to back out the density function. When the prices are shocked by small amounts, in this case by half a tick, the shocks have a small impact on the implied volatility, which is then smoothed away by the curve fitted to the smile. The mixture of lognormals is more sensitive with respect to these shocks as the option prices are directly used to estimate the density function. That is why I had a look at the option prices used, by processing some of Neil's data with my own model. A further reason – to be honest – was that I had marvellous results as regards the mean and standard deviation for my own model (but less promising estimates of skew and kurtosis).

The equation I used is an approximation of the first derivative of the option price with respect to the strike, which yields the implied cumulative probability distribution at that strike (K_i), which I approximate with a simple difference quotient, which in this audience does not require a lot of explanation.³

³ Cf. Holger Neuhaus (1995), The information content of derivatives for monetary policy – implied volatilities and probabilities, Deutsche Bundesbank Economic Research Group, Discussion paper 3/95 (July 1995).

C_i , K_i and F_T are the price of option i , its strike price and the value of the futures at the expiry of the option. To be precise, the option in this equation should be either margined or C should already be adjusted for the discount factor (as is the case here).

$$(1) \quad C = \int_{-\infty}^{+\infty} w(F_T) \max(0, F_T - K) dF_T$$

$$(2) \quad C_K = - \int_K^{+\infty} w(F_T) dF_T$$

$$(3) \quad p(F_T > K_i) = \frac{C_{i-1} - C_{i+1}}{K_{i+1} - K_{i-1}}$$

Even if one does not recognise the quotient (consisting of the difference between two option prices divided by the difference in their strike prices) as a cumulative probability distribution, it is clear that deep in-the-money options going deeper into the money by one unit will increase in value by the same amount, i.e. the ratio should be exactly unity. Is this always the case in the paper? No: as one can see in table 3, there are small problems with the true prices (having taken into account the discount factor)

Table 3: True prices and implied cumulative probability distribution values

Strike	C	C compounded	Probability distribution
71	28.94429	29.000015	1.0000
72	27.94621	28.000015	1.0000
73	26.94813	27.000015	1.0000
74	25.95006	26.000015	1.0000
75	24.95198	25.000015	1.0000
76	23.9539	24.000015	1.0000
77	22.95582	23.000015	1.0000
78	21.95774	22.000015	1.0000
79	20.95966	21.000015	1.0000
80	19.96159	20.000015	1.0000
81	18.96351	19.000015	1.0000
82	17.96543	18.000015	1.0000
83	16.96735	17.000015	1.0000005009626
84	15.96927	16.000014	1.0000
85	14.97119	15.000015	1.0000
86	13.97311	14.000014	1.0000
87	12.97504	13.000015	1.0000
88	11.97696	12.000014	1.0000005009626
89	10.97888	11.000014	1.0000001502888
90	9.980799	10.000014	0.9999999499037
91	8.98272	9.000014	1.0000001502888

in particular, have an influence on the results of the mixture of lognormals approach. However, does it imply that this approach is indeed too sensitive? Not necessarily, as the prices used allow for arbitrage and are thus “wrong”, albeit to a limited extent only. A way forward in the research would be to just shock prices of options that are not *deep* in or out of the money.

Table 4: Some shocked prices and implied cumulative probability distribution values

Shock	Strikes				
	78	79	80	81	82
1	0.9963	0.993189	1.002543	1.011356	0.99394
2	0.999621	1.024057	1.008464	0.977107	0.993295
3	1.007501	1.004483	0.994608	1.008015	1.012959
4	0.999495	1.017243	1.005898	1.003024	0.994951
5	0.995052	0.990734	0.990985	1.001388	0.99782
6	1.010791	1.004175	0.990104	0.983198	0.992493
7	0.997608	0.984138	1.00091	1.019346	1.001663
8	1.000668	0.998463	0.984069	1.000081	1.001444
9	1.019999	0.991685	0.998495	1.014912	1.00225
10	1.011092	1.018889	0.981598	0.996102	1.011387
11	1.002029	1.01328	0.999109	0.993804	0.985691
12	1.016841	0.989879	0.990558	1.013711	1.003615
13	1.001403	0.98645	1.00682	0.994389	0.986779
14	1.004532	1.015066	0.98888	0.993103	1.015303
15	1.017755	0.984908	0.997246	0.996538	0.999667
16	1.003631	1.003485	1.002598	1.004908	0.995692
17	0.992376	0.998081	0.998523	1.004808	1.015937
18	1.00167	0.987514	1.00775	1.011708	0.989414
19	0.98171	1.0117	1.009924	0.978308	1.00168
20	0.999093	1.002282	1.01776	1.002007	1.00061
21	0.981344	0.997791	1.006633	1.002047	1.009586
22	1.013006	1.000001	0.983824	0.998456	1.02053
23	1.006047	0.985566	0.996316	1.008509	0.984099
24	0.993576	1.001908	1.003176	0.985739	0.988618
25	0.988323	1.005294	0.988197	1.00032	1.013966
26	0.993232	1.003898	1.010038	1.010592	0.998971
27	0.99807	0.990528	1.01003	1.006907	0.987676
28	1.021555	0.991778	0.981613	1.002087	1.009585
29	0.977072	1.000689	1.013373	0.996852	0.98623
30	0.996151	0.999982	0.997678	1.019486	1.003256

Conclusions

To summarise what I found out in the short time I had to look at the paper, I think that it is interesting and a step in the right direction. Nevertheless, I have some food for thought:

- Is there “too much” data?
 - Some models are designed to generate results with a very limited data set (strike prices) only, some require more. The choice of model is likely to depend on the market to be monitored.
 - The number of and interval between strike prices used in the comparison should reflect the features of the market that should be monitored.
- Should the probability *distribution* be estimated rather than the *density*?
- Is the smile technique “too” stable?
- Is the mixture lognormals too sensitive as regards errors in prices?
- Data must not allow for arbitrage.
 - This also holds true for “shocked” data.
 - Only the prices of at-the-money options but not of far in-the-money or out-of-the-money options should be shocked.
- How big is big?
 - The size of the errors has to be put into perspective. Criteria used should involve also, for example, computational costs in a broad sense (computer and software requirements, robustness of the estimates).
 - In particular, some percentage errors shown may be misleading.
- The use of the third and fourth moment is debatable.

Discussion of Neil Cooper's paper:

**Testing techniques for estimating implied RNDs
from the prices of European-style options**

Discussant: Jan Marc Berk

BIS, 14 June 1999

1. My contribution is structured as follows. I will start by giving a brief summary of the paper. This is then followed by some comments, and I conclude by sketching some paths for future work on the subject at hand.

Summary

2. The paper aims to compare two methods for calculating PDFs. Both methods are applied on both European and American-style options. The performance of both methods is tested by means of Monte Carlo analysis, although the current version of the paper only deals with comparing the methods applied to European options. Innovative aspects of the paper include the variant used for calculating volatility-smile-based PDFs, and the Monte Carlo experiment.
3. The methods used to construct PDFs for *European* options are the well-known mixture of lognormals approach (MLN), as documented by, for example, Bahra (1997), and a method based on interpolation of the volatility smile (IVS), as introduced by Shimko (1993). The paper slightly amends the Shimko approach, as it uses cubic splines in stead of quadratic forms, and interpolates in delta space in stead of volatility space. These amendments are in line with, for example, Malz (1997).
4. PDFs for *American*-style options are the MLN variant introduced by Melick and Thomas (1997), and the early exercise premium is taken into account within the IVS method by using the approximation of Barone-Adesi and Whaley (1987).
5. The performance of both methods is compared in a Monte Carlo experiment, using simulated artificial data in stead of observed prices. However, by using the stochastic volatility model of Heston (1993), the author generates quite realistic data, whilst retaining the advantage of knowing the 'true' PDF.
6. The results from the Monte Carlo analysis, in the current version of the paper applied only to European options, are that, on average, there is no clear winner between MLN and IVS. However, the latter method provides far more stable estimates. The instability of MLN estimates are due to spikes, and reflect estimation errors. Instability increases with volatility and skewness, and decreases with

time to maturity. Based on these Monte Carlo analysis, the author expresses a preference of IVS over MLN.

Comments

7. I find the paper of Neil Cooper very interesting, and as it is work in progress, I suspect it will become even more interesting. The paper reflects my own experiences, or should I say frustations, with the (in)stability of the MLN method. Without meaning to detract from the quality of the paper, there are some points that, in my view, deserve some further consideration. Given the time constraint, I will only briefly touch upon them here:
8. Whilst the application of the MLN method in the paper is fairly standard, the version of the IVS method employed is more innovative. The paper could benefit from a more extensive discussion on the effects of the amendments *vis-à-vis* the Shimko approach.
9. In a similar vein, no mention is made in the paper of possible drawbacks of the IVS method, such as the problem of fitting the tails of the PDF (ie outside the observed range of strikes), and negative probabilities.
10. The comparison of both methods is based on Monte Carlo analysis. Yet, given the results of Melick and Thomas (1998), who find widely different results for simulations based on Monte Carlo and bootstrap methods, and given the the assumptions underlying the MC-method (independent errors, regularity conditions) *vis-à-vis* actual options prices, some attention to the validity of Monte Carlo as a tool for comparison seems in order.
11. The comparison of both methods uses artificial data, so there is no distinction between exchange-traded and otc data. Campa, Chang and Reider (1997) compare MLN and IVS methods using otc data, and find that they yield similar results. Could or should the choice of method (MLN versus IVS) be dependent on the type of data used?
12. The focus of the paper is primarily technical and not economic, which is understandable given its objective. However, more attention to the economic aspects would seem in order, as it could provide an answer to the question as to how important the instability of MLN based PDFs is. Clearly, this answer depends on the purpose of the analysis using PDFs.

Way to proceed

13. Based on my, admittedly limited, knowledge of the estimation and use of PDFs, there are two basic questions which in my view remain to be answered in a convincing way. First, regarding the method used to calculate PDFs, do we really need to impose so much structure? Second, regarding the estimation of PDFs, do the data allow us to impose so much structure?

14. My personal opinion on these questions is that we should use different methods for different purposes, also taking the amount and types of data into account. As an economist, I would tend to say that economic considerations should govern the purpose of the analysis, as well as that the results of the analysis should be useful to economists. As an economist working in a monetary policy department, I will go even one step further and state that the results of the analysis should be useful to policy makers. Given the fairly technical nature of work involving PDFs, it is my own experience that translating the results of PDF based analyses to policy messages is by no means an easy task.
15. Data considerations are also of importance in the choice of method. I already touched upon the difference between exchange-traded and otc data and possible implications for the choice of calculation method. Moreover, we all encounter situations when only a few data points are available, or that only a limited subset of a larger set of prices reflect sufficient liquidity. In these situations, I found entropy-based (Bayesian) methods for calculating PDFs useful. Moreover, the field of maximum entropy econometrics has a firm statistical foundation, and provides a natural metric for evaluating different methods. Finally, when not even a limited set of data on options prices is available, it may be still possible to extract a PDF, using alternative methods (Hördahl, 1999).

References

- Bahra, B (1997): "Implied risk-neutral probability density functions from option prices: theory and application". *Working Paper no 66*, Bank of England, London.
- Barone-Adesi, G and R E Whaley (1987): "Efficient analytic approximation of American option values". *Journal of Finance*, 42, pp. 301-20.
- Campa, J.M., P H K Chang and R L Reider (1997): "Implied exchange rate distributions: evidence from otc option markets". *Working Paper no 6179*, NBER, Cambridge (MA).
- Heston, S (1993): "A closed-form solution for options with stochastic volatility with applications to bond and currency options". *The Review of Financial Studies*, 6, pp. 327-43.
- Hördahl, P (1999): "Estimating the implied distribution of the future short term interest rate using the Longstaff-Schwartz Model". Paper presented at the ECB Workshop on yield curve modeling, Frankfurt, March 1999.
- Malz, A. (1997): "Estimating the probability distribution of the future exchange rate from option prices". *The Journal of Derivatives*, Winter 1997, pp. 20-36.
- Melick, W R and C P Thomas (1997): "Recovering an asset's implied pdf from option prices: an application to crude oil during the Gulf Crisis". *Journal of Financial and Quantitative Analysis*, 32, pp. 91-115.
- Melick, W R and C P Thomas (1998): "Confidence intervals and constant maturity series for probability measures extracted from option prices". Paper presented at the conference 'Information Contained in Prices of Financial Assets', Bank of Canada.
- Shimko, D (1993): "Bounds of probability". *Risk*, 6, pp. 33-37.