Recent developments in extracting information from options markets

By Roger Clews, Nikolaos Panigirtzoglou and James Proudman of the Bank's Monetary Instruments and Markets Division.

The Monetary Policy Committee is provided with information from options markets to quantify market uncertainty about the future course of financial asset prices. For short-term interest rates, this is shown in the Inflation Report's *blue fan chart. Similar information can be obtained from a wide range of other assets. This article compares the performance of alternative techniques for extracting information from options prices. Using a technique for estimating uncertainty about interest rates at a constant horizon a short way into the future, we consider how this uncertainty has evolved since the Bank was granted operational independence in May 1997.*

Introduction

Virtually all financial assets pay out in the future. So the prices at which different assets trade can tell us something about the market's view of future states of the world. For example, the prices of bonds of different maturities contain information about the expected course of interest rates between maturity dates (see, for example, Anderson and Sleath (1999)).

Options are contracts giving the right (but not the obligation) to buy or sell an asset at a point in the future at a price set now (the strike price).(1) Options to buy (call options)(2) are only valuable if there is a chance that when the option comes to be exercised the underlying asset will be worth more than the strike price. So if we look at options to buy a particular asset at a particular point in the future but at different strike prices, the prices at which such contracts trade now tell us something about the market's view of the chances that the price of the underlying asset will be above the various strike prices. So options tell us something about the probability the market attaches to an asset being within a range of possible prices at some future date.

Over the last few years, there has been considerable interest among academics, market participants and policy-makers in extracting information of this kind from options prices. The techniques used are described more fully below, but a common way of displaying the information extracted is as an implied risk-neutral probability density function (pdf) for the asset upon which the contract trades.

Chart 1 shows a pdf derived from contracts based on a short-term interest rate (three-month Libor). Possible levels of the interest rate are measured horizontally; probability is measured vertically. The area under the curve sums to 100%. The shaded area for example depicts the probability

Chart 1
March 2000 short sterling implied pdf; 5 January 2000

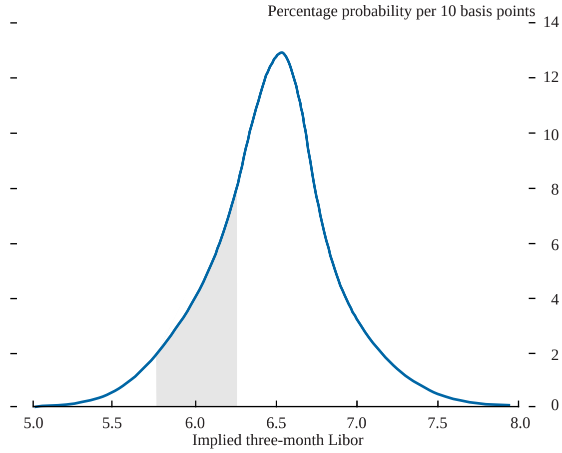

that the interest rate will lie between 5.75% and 6.25%. The area is 24% of the area under the curve; thus there is estimated to be a 24% probability that the interest rate will lie in that range when the contract settles.

The Monetary Policy Committee is provided with information from options prices to assess the degree of market uncertainty. For example, pdfs have proved useful in estimating the market's assessment of the balance of risks associated with future movements in asset prices. Market uncertainty about UK short-term interest rates is shown in the *Inflation Report's* blue fan chart. Chart 2 shows a fan chart using data as at 5 January 2000. This is built up from the risk-neutral pdfs of three-month sterling interest rates, derived from the prices of options on each of the short sterling futures contracts settling at three-month intervals up to December 2000.

The fan chart is rather like a contour map, looking down on the pdf 'hills'. At any given point in time, the depth of the

(1) The glossary on page 59 explains the key technical terms used in this article.
(2) Options to sell are put options.

Chart 2
Implied distribution for sterling three-month interest rates

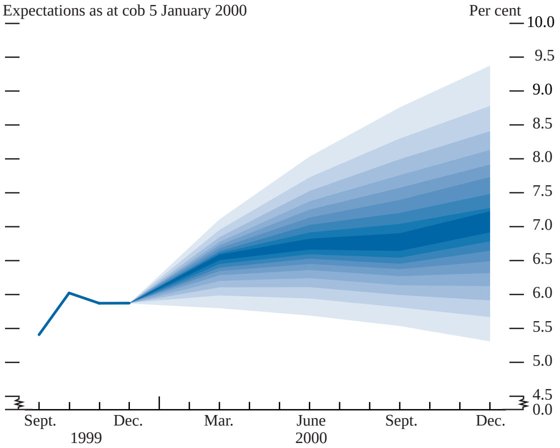

shading represents the height of the pdf implied by the markets over the range of potential outcomes for short-term interest rates. Assuming risk-neutrality, the markets judge that there is a 10% chance of interest rates being within the darkest, central band at any date. Each successive pair of bands covers a further 10% of the probability distribution until 90% of the distribution is covered.

The Bank also estimates pdfs from options for a range of other financial assets. Pdfs for FTSE 100 index options and euribor futures options are estimated from contracts traded on the London International Financial Futures and Options Exchange (LIFFE). A range of pdfs—such as for the S&P 500 index, the Nikkei 225 index, eurodollar and euroyen futures options—is derived from options traded on the Chicago Metal Exchange.(1) Pdfs can also be estimated for physical commodities. For example, pdfs for crude oil and gold prices can be extracted from futures options traded on the New York Mercantile Exchange.

In recent years, the pdfs used at the Bank have been estimated using a parametric technique, the mixture of two lognormals, described in Bahra (1996 and 1997). In the following sections of this article, we review recent research carried out in the Bank to evaluate the performance of this technique.(2) First, we discuss the quality of the data used to estimate pdfs. Next, we evaluate the parametric technique against a new non-parametric method, the 'smile interpolation', discussed in Bliss and Panigirtzoglou (2000) and Cooper (2000).

Exchange-traded options (for which data are most readily available) settle on particular days in the year. This means that the maturity of a pdf from such a contract gradually gets shorter as time passes. The pdf that we estimate today is not quite comparable with the one we estimated from the same contract yesterday and is much less comparable with the one we estimated a month or a year ago. So in the latter sections of the article we show how the new technique can be used to construct a pdf with a constant-maturity horizon. This can help us to answer a range of questions of interest to policy-makers, such as whether the degree of market uncertainty about short-term interest rates has altered since the Bank was granted operational independence in May 1997.

Extracting information from options prices

As noted above, options prices can provide us with a guide to the likelihood the market attaches to future values of asset prices. By comparing options with different strike prices, it is possible to infer the probabilities that the market attaches to different levels of the underlying asset price.

Breedon and Litzenberger (1978) derived the result that the probabilities attached to different levels of the underlying asset price may be derived from options prices, if one assumes that investors are risk-neutral. To see intuitively why we would expect the prices of options to reflect these probabilities, suppose that we observe a set of call option prices with the same maturity but with different strike prices. A call option with a lower strike price will always be worth more than a higher strike option. This is because the option with the lower strike price will have a higher pay-off if exercised and has a higher probability of delivering a positive pay-off. This additional probability reflects the chances that the underlying asset price will lie between these strikes. If we have option prices for a range of strikes, it is possible to infer what the probabilities are of the underlying asset price at maturity lying between each of them, by examining the relative prices of options with adjacent strikes.

Under the assumption of risk-neutrality, the distribution is the set of probabilities that investors would attach to future asset prices in a world in which they were risk-neutral. But if investors are risk-averse, risk premia will drive a wedge between the probabilities inferred from options and the true probabilities they attach to alternative values of the underlying asset price. The mean of the risk-neutral pdf, for example, will not equal the expected price of the underlying asset at maturity.

This potential bias may affect the way in which we interpret estimated pdfs, especially those based on equity index futures options.

Sources of options data

The accuracy of the pdfs that we estimate depends crucially on the quality of the options prices used as the inputs into the estimation process. One source of estimate instability may be that the end-of-day settlement prices we obtain from

(1) These options are American options, which can be exercised at any time before maturity. We adjust for the early exercise premium using the Barone-Adesi and Whaley (1987) approximation.
(2) This article summarises work done in the Bank over the last year, and draws on the contributions of Robert Bliss, Neil Cooper and Gary Xu.

the exchanges, such as LIFFE, may not reflect true market prices across all strike prices.

One reason for this may be low trading activity. To examine this, we analysed the trading patterns of two options contracts—the FTSE 100 index and the short sterling futures options contracts for the period 1990–97. Our main findings were as follows.

- The markets for these contracts were not very liquid. The average daily number of trades for FTSE 100 index options was only 155, for all calls and puts across all strike prices and maturities. The comparable number for short sterling was 80. But there were periods when liquidity was consistently higher than average. For example, the trading volume of short sterling futures options was much higher during the second half of 1998 when, arguably, uncertainty about future short-term interest rates was relatively high.

- Trading was heavily concentrated in options whose strike price was close to the current futures price (near-the-money) or in call (or put) options whose strike price was above (or below) the futures price (out-of-the-money). We illustrate this in Chart 3. The chart shows the number of options contracts on the March 2000 short sterling futures traded at different strike prices on 5 January. It is typical of trading patterns in these contracts. The short sterling futures contract settles at a price equal to 100 minus three-month Libor on the last trading day for the contract. So a higher interest rate means a lower price for the short sterling futures contract. The red central line in the chart denotes the interest rate implied by the current futures price; options contracts with a strike price at this level would be at-the-money (atm). The chart shows that some call option contracts traded with strike prices close to this (near-the-money), but most traded at higher strike prices, ie at an implied interest rate below the current level (out-of-the-money). Some put options also traded near-the-money but again most traded out-of-the-money, ie at an implied interest rate above the current level. Note that no contracts were traded very far away from the current price of the underlying asset; ie no contracts traded on strike prices in the tails of the estimated distribution.

Chart 3
March 2000 short sterling; 5 January 2000

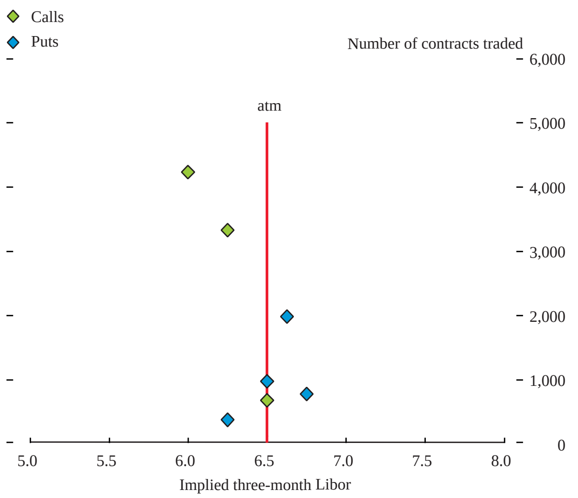

- Trading was concentrated in the options contract closest to maturity. For options contracts in which there had been no trading during the day, settlement prices were assigned by LIFFE using a pricing model. The absence of traded prices was a problem, particularly for in-the-money and deep out-of-the-money options and options with a long time to maturity.

- The range of strike prices for which trading was observed was greater for FTSE 100 than short sterling contracts. For example, the median number of strike prices in which trading was observed was 16 for FTSE 100 index options with a time to maturity of around one month, compared with 5 for short sterling options.

These results suggest that some of the options prices we obtain from the financial futures exchanges may be distorted by factors associated with low liquidity. To reduce these distortions, the pdfs we discuss below were estimated using only the near-the-money and out-of-the-money call and put options prices, which are generally traded in more liquid markets.

Alternative techniques for estimating pdfs

As discussed above, the value of a call (or put) option depends on the probability of the asset price lying above (or below) the strike price and the value of the pay-off at the expiry of the option. More formally, the call price function relates the price of a call option to its underlying parameters, such as maturity, the strike price of the option and the pdf of the underlying asset price at the expiry of the option.

Many different econometric techniques have been developed to derive pdfs from a range of call options prices of the same maturity. Some of these techniques involve specifying the parameters of a statistical process for the underlying asset price. The parameters of the process can be used to generate a pdf for the underlying asset at the maturity of the option. This in turn can be used to generate a call price function. The parameters of the stochastic process can then be chosen to make the implied call price function as similar as possible to that observed in the data (see, for example, Malz (1995) and Bates (1996)).

Another technique for estimating pdfs is to assume a specific parametric form for the pdf. The parameters of the pdf are estimated in such a way that the implied call price function is again as close as possible to the one actually observed in the data. The mixture of two lognormals is an example (see Melick and Thomas (1997)).

This form is sufficiently flexible to capture features that we might expect to find in the data, such as fatness in the tails of the distribution (excess kurtosis), positive or negative skewness, or bimodality. And the mixture of two lognormals method is parsimonious, in the sense that it can be derived by estimating only five parameters. This parametric approach has been used in recent years at the Bank of England to estimate pdfs, as described in Bahra (1996 and 1997). We describe the technique in more detail in the technical appendix on pages 58–59.

But the parametric approach has, in practice, proved to have some undesirable properties. One is that it can generate pdfs characterised by sharp spikes. This occurs when one of the two lognormals is estimated to have a very small standard deviation, as illustrated in Chart 4. A pdf estimated using an alternative technique—the smile interpolation (see below)—is included for comparison. Another problem is that the parametric method can generate implausibly large changes in the shape of the pdf between consecutive days. This is true particularly for measures of the skewness and kurtosis of the distribution. These problems led us to conclude that the parametric method does not always fit the data well and to consider whether more robust estimation methods exist.

Chart 4
March 2000 short sterling implied pdf; 25 January 1999

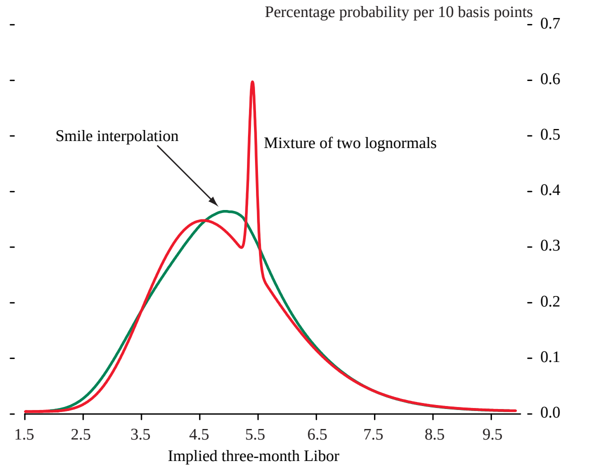

An alternative technique—such as the smile interpolation—would be to estimate a smooth and continuous call price function directly, by interpolating across the call prices we observe for different strike prices. We could then exploit the Breedon and Litzenberger (1978) result that we can infer underlying probabilities directly from the call price function. However, the call price function has a large curvature for options near-the-money and very little curvature for options far away-from-the-money. This can make direct interpolation across the call price function difficult. To avoid this practical problem, we transform the call price function into a particular form of 'volatility smile', estimate a smooth smile, convert it back into a call price function and use that to derive the pdf.

To convert a call price function into the relevant volatility smile (and *vice versa*) involves transforming both axes in a non-linear way. We convert option prices into implied volatilities. The implied volatility is the volatility of the underlying asset price implied by the Black-Scholes (1973) model and is a non-linear transformation of the option price. A conventional volatility smile plots implied volatility against the strike price, but such smiles can vary in smoothness from day to day, making consistent interpolation problematic. We choose to interpolate implied volatilities across deltas rather than strikes, as illustrated in Chart 5. The delta of an option is the rate of change of the option price with respect to the underlying asset price and is a non-linear transformation of the strike price. These 'delta smiles' have a more stable degree of smoothness from day to day.

Chart 5
Implied volatility smile of the September 1999 short sterling contract; 27 January 1999

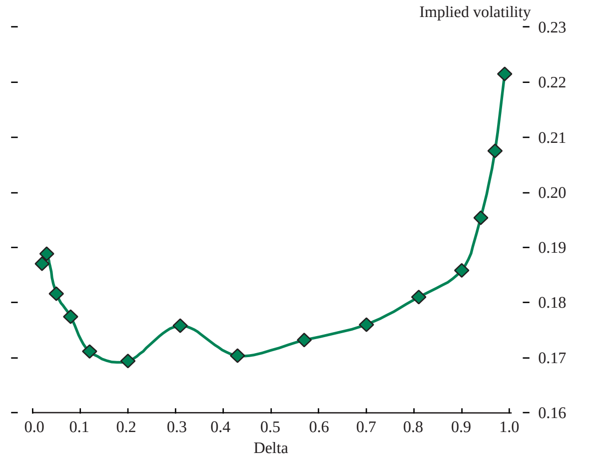

The interpolation across the delta smile as in Chart 5 is done using a smoothing spline, which is a flexible non-parametric technique. A smoothing spline is a piecewise cubic polynomial, the smoothness of which is controlled by a single parameter, the smoothness parameter. Because we interpolate across delta space we can hold the smoothing parameter constant from day to day. This means that changes in pdfs from day to day reflect changes in the underlying data, and not in the estimation technique.(1)

To compare the parametric and non-parametric techniques, we examined their performance with respect to two criteria, which we discuss in the next section.

Comparing the stability of the techniques

Our comparisons of the stability of the implied pdfs derived from the parametric and non-parametric methods are

(1) The complete estimation process is described in detail in Bliss and Panigirtzoglou (2000) and Cooper (2000), and is also summarised in the technical appendix.

discussed in detail in Bliss and Panigirtzoglou (2000) and Cooper (2000).

Given the problems, noted above, with the prices used in estimation, the first criterion was that our technique for estimating pdfs should be robust to small and random errors in the underlying options prices. We therefore examined the extent to which small perturbations in actual options prices generated large changes in the estimated pdfs.(1)

The exercise was carried out on more than 700 short sterling futures options contracts and on 1,400 FTSE 100 index options contracts. The test involved repeatedly perturbing the set of options prices for each contract in our sample by a small and random amount. The perturbations were drawn from a uniform distribution between plus and minus one half of one tick-size. The tick-size was chosen as the range of the distribution because this is the smallest observable difference in quoted prices. So prices within one tick-size are observationally equivalent to each other.(2)

For each set of simulated prices, we estimated pdfs using both the parametric and the non-parametric methods. This process was repeated 100 times for each futures contract. The parametric and non-parametric methods were then evaluated by comparing the sample distributions of a number of summary statistics—such as the standard deviation, skewness and kurtosis—estimated for each pdf.

The results are summarised in Table A. The dispersion of the summary statistics for the non-parametric method was smaller than that of the summary statistics estimated using the parametric method for all three measures presented. These results suggest that the non-parametric technique is more robust to small and random errors.

Table A
Standard deviations of estimated summary statistics

	Short sterling		FTSE 100 index	
	Parametric	Non-parametric	Parametric	Non-parametric
Standard deviation	0.020	0.004	2.140	0.150
Skewness	0.192	0.068	0.050	0.005
Kurtosis	1.199	0.231	0.165	0.018

Note: The results shown are after filtering potential outliers, defined as any value outside the 0.5 to 99.5 percentiles of the respective distribution.

The second criterion we considered for evaluating the performance of the techniques was their ability to recover accurately a pdf from a set of simulated prices.(3) By using simulated prices, rather than actual prices, we can compare the estimated pdfs against the 'true' pdf implied by the underlying stochastic process.

The simulated prices were generated from a general stochastic volatility model, set out in Heston (1993). This is an attractive model because it allows us to simulate option prices drawn from a wide range of underlying pdfs; with high or low volatility and kurtosis, and positive or negative skewness. In this model, the dynamics of the asset price are given by the following stochastic differential equation:

$$dS = \mu S dt + \sqrt{v_t} S dZ_1 \qquad \textbf{(1)}$$

and

$$dv_t = k(\theta - v_t)dt + \sigma_v \sqrt{v_t} dZ_2 \qquad \textbf{(2)}$$

where S is the asset price, μ is its mean rate of drift and v_t its conditional variance at time t. This follows a mean-reverting process such that the variance reverts to a long-run mean of θ at a rate k. The parameter σ_v is its standard deviation. Finally, Z_1 and Z_2 are Wiener processes whose correlation is given by a value ρ. By changing ρ, we can generate skewness in the distribution of the asset price. For example, 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 downward movements. So a negative correlation will generate negative skewness in the asset price distribution. A positive correlation has the opposite effect.

Using this model, we established a set of six scenarios corresponding to low and high volatility, and positive, negative and no skew. For each scenario, we generated a set of options prices over a range of strike prices and maturities. For each combination of scenario and maturity, we shocked each price by a small and random amount, in the same way as described above. We fitted pdfs to each set of perturbed prices using the parametric and non-parametric techniques and calculated summary statistics associated with each pdf. We repeated this process 100 times for each scenario and maturity.

We assessed the two techniques by comparing the standard deviations of the estimated summary statistics. Table B presents the standard deviations of the estimated summary statistics for one-month pdfs, estimated from both the parametric and non-parametric approaches.

Larger standard deviations of the summary statistics indicate greater instability in the estimated pdfs. For nearly all the scenarios, the parametric technique has larger standard deviations for the three statistics than does the non-parametric method.

The research suggested that the parametric technique is less stable than the non-parametric one evaluated against both the criteria discussed above. This instability is likely to reduce the value of the parametric technique as a practical tool, compared with the non-parametric.

(1) See Bliss and Panigirtzoglou (2000).
(2) The tick-size was defined as 0.05.
(3) See Cooper (2000).

Table B
Standard deviations of estimated summary statistics for one-month pdfs

Summary statistic	Scenario	Parametric	Non-parametric
Standard deviation	1	0.0730	0.0110
	2	0.5684	0.0137
	3	7.5644	0.0123
	4	0.0093	0.0095
	5	0.0092	0.0080
	6	2.4101	0.0079
Skewness	1	0.1663	0.0192
	2	0.2341	0.0234
	3	0.1899	0.0166
	4	0.0458	0.0064
	5	0.0055	0.0061
	6	0.1839	0.0066
Kurtosis	1	0.1002	0.0156
	2	0.1835	0.0333
	3	0.3374	0.0296
	4	0.0185	0.0065
	5	0.0274	0.0078
	6	0.3777	0.0150

Note: Scenarios are: (1) negative skew, low volatility; (2) no skew, low volatility; (3) positive skew, low volatility; (4) negative skew, high volatility; (5) no skew, high volatility; and (6) positive skew, high volatility.

Intuitively, the greater instability of the parametric approach arises because the pdf is estimated using only five parameters. Small changes in the price of one option can affect the value of these parameters, and hence the shape of the whole pdf. In contrast, when we use the non-parametric method, the effect of changes in the price of one option on the shape of the pdf is more localised. A similar result was found when comparing the stability of parametric and non-parametric techniques for fitting yield curve data (see Anderson and Sleath (1999)).

Estimating pdfs over a constant horizon

Options contracts traded on financial futures exchanges, such as LIFFE, have fixed expiry dates corresponding to the maturity of the underlying futures contracts: March, June, September and December. This feature can make comparing pdfs over time difficult. This is because the degree of uncertainty about the price of the underlying futures contract at the expiry date of the option naturally decreases as the expiry date approaches. So the implied volatilities and variances of the pdfs that we estimate diminish over time, without any real change in the degree of uncertainty about the asset. Normally the implied volatility of each contract drifts downwards through the operation of this ‘time-to-maturity’ effect. The pattern of decaying implied volatilities for successive short sterling contracts is shown in Chart 6. But volatilities can also be shocked by some external event.

To discern more clearly such underlying changes we need a method for stripping out the ‘time-to-maturity’ effect. Our method for doing this is based on—and is consistent with—the non-parametric technique discussed above. There we interpolated across the implied volatilities of options with different deltas but with the same maturity. Here we interpolate across the implied volatilities of contracts with the same delta but with different maturities. Here too there is an advantage in using deltas rather than strike prices. The range of possible values of delta runs from 0 to 1 and for every maturity there are contracts with deltas covering most of the range. The range of available strike prices by contrast

Chart 6
Short sterling atm implied volatility

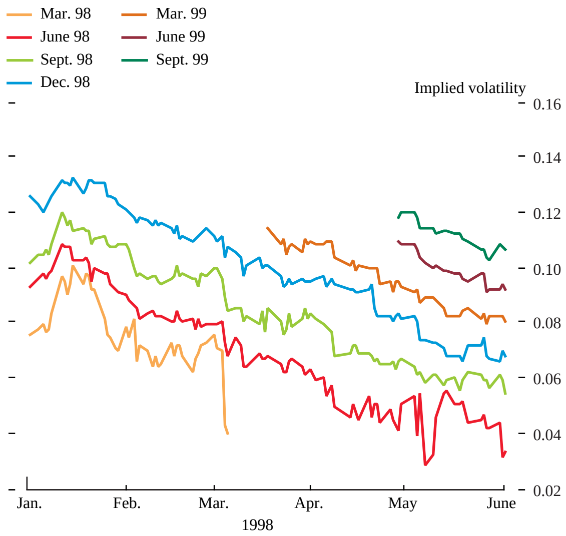

Chart 7
Short sterling implied volatilities at different maturities; 27 January 1999 (delta = 0.6)

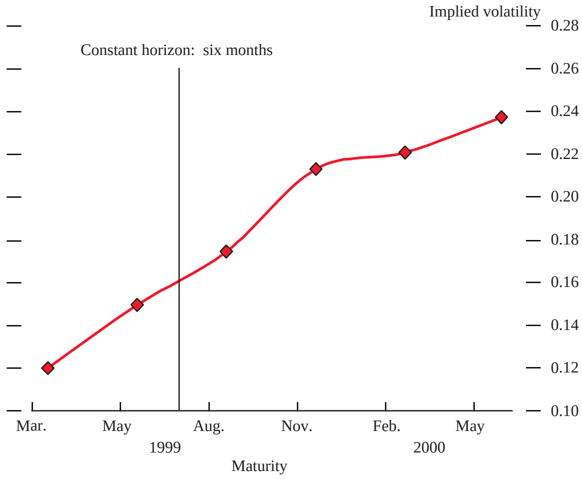

is often quite different at different maturities. An example of the relationship between implied volatilities and maturities of options with the same delta is given in Chart 7.

The volatility of a hypothetical contract with a delta of 0.6 and a six-month maturity can easily be read off the chart. A complete ‘delta smile’ for hypothetical six-month contracts can be built up from similar charts for contracts with different deltas. And then a pdf with six-month maturity can be constructed from the smile in the usual way. The whole process can be repeated on the next day. Even though actual contracts will have a maturity that is one day shorter, the maturity of the constructed pdf will remain at six months.

In fact, we can construct a surface of implied volatility, as shown in Chart 8. The surface is estimated from the implied volatilities from contracts on all available deltas and maturities. The implied volatility smile of a constant-horizon pdf can be thought of as a cross-section of the surface at a particular date.

Chart 8
Short sterling implied volatility surface; 27 January 1999

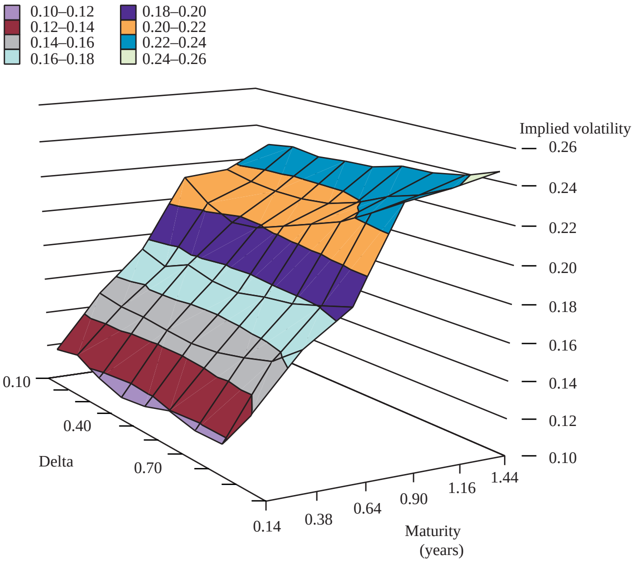

Interpolating across maturities introduces a source of potential measurement error into the estimated pdfs. One problem is that when the contract closest to maturity expires, it is replaced by another—longer-dated—contract. This may induce instability in the constant-horizon pdfs, particularly at shorter horizons. To test for the size of this effect we estimated the absolute daily changes in a number of the summary statistics of the constant-horizon pdfs, both for when there was contract switching and for when there was none. We then constructed two samples and tested the null hypothesis of equal means.

For all the summary statistics for short sterling, the differences in the means of the two samples were significantly different from zero. Except for the variance, they were also significantly different for all the FTSE 100 summary statistics. However, the differences in the means were very small.

The evolution of short-term interest rate uncertainty in the United Kingdom

Investors' uncertainty about the future path of short-term interest rates may partly be related to uncertainty about the monetary authorities' reaction function. But it will be influenced by uncertainty about the shocks to which the monetary authorities react. So changes in market uncertainty may reflect a perceived change in the monetary policy reaction function, and/or a perceived change in the nature of the exogenous uncertainty facing the economy.

Constant-horizon pdfs are a useful tool for evaluating changes in market uncertainty over long periods of time. For example, we can examine the time series properties of the summary statistics. In this section, we consider what constant-horizon pdfs can tell us about the evolution of short-term interest rate uncertainty since January 1997. Did the markets become less uncertain about the outlook for short-term interest rates following the introduction of operational independence for the Bank of England in May 1997? How uncertain about the future course of monetary policy were the markets in the wake of the Long Term Capital Management (LTCM) crisis in autumn 1998?

Chart 9 plots the level and the standard deviation of three-month sterling Libor implied by the constant six-month horizon pdf, from January 1997 to December 1999. The standard deviation is a measure of how dispersed the implied level of the interest rate was seen to be, and hence of uncertainty.

Chart 9
Summary statistics of the six-month constant-horizon short sterling pdf

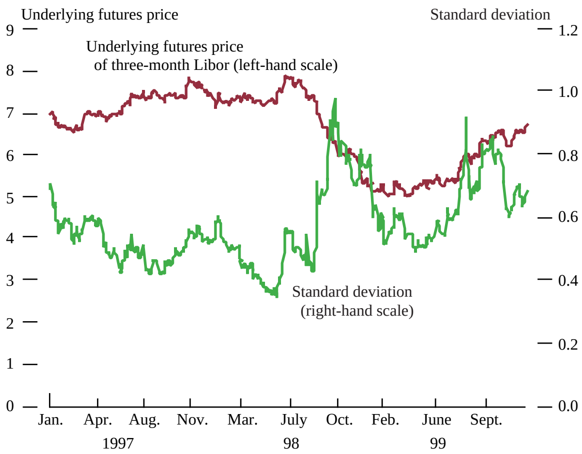

The series of the standard deviation is fairly volatile, but was little changed overall during the period. It rose a little in anticipation of the General Election in May 1997. Uncertainty then fell modestly in the months following the granting of operational independence to the Bank. It rose slightly following the market's response to the 25 basis point rise in the Bank's repo rate in June 1998. But these moves were dwarfed by the large rise in market uncertainty during the period of financial turbulence in the late summer and autumn of 1998.

We can also use the higher moments of the six-month constant-horizon pdf for short sterling to consider the evolution of the balance of risks to monetary policy in the United Kingdom. In Chart 10, we plot the kurtosis and skewness of the same pdf over the same sample period.

Skewness is a measure of the balance of risks attached by the market to different outcomes. Positive skewness occurs when the market attaches a higher probability to a sharp upward movement in short-term interest rates than to a comparable downward movement. Because short-term interest rates are bounded from below by zero, measures of skewness tend to be positive. But since May 1997, the degree of skewness has fallen towards zero, the level at which the pdf is symmetrical. One interpretation is that operational independence for the Bank reduced the

Chart 10
Summary statistics of the six-month constant-horizon short sterling pdf

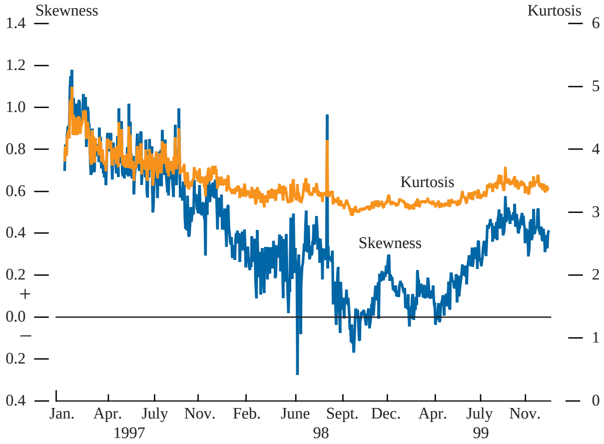

probability the market attached to a sharp upward movement in interest rates. The LTCM crisis may have further reduced the perceived likelihood of a sharp rise in rates. In contrast, the unexpected rise in repo rates in June 1998 was followed by a period of increased variability in the skew. This could indicate market difficulty in assessing the likely consequences of the decision to raise rates. Skewness increased again during the first half of 1999, as the probability of a further sharp easing in rates diminished with the recovery in UK economic activity.

Kurtosis measures the probability the market attaches to extreme levels of interest rates, either up or down. Levels of kurtosis above three indicate that the market attaches higher probability to extreme outcomes than would be implied by a normal distribution. Since May 1997, the level of kurtosis has fallen slightly, indicating that the market has come to attach a lower probability to extreme values of short-term interest rates. This may again be associated with the change in monetary policy regime. Kurtosis increased sharply—but very briefly—following the LTCM crisis.

Conclusion

Research provides evidence that a non-parametric technique for estimating pdfs is an improvement upon the parametric one that has been used at the Bank over recent years. This conclusion mirrors a result found in tests on the yield curve (see Anderson and Sleath (1999)).

We can also use a non-parametric technique to estimate pdfs over a constant-maturity horizon. As we illustrate with a simple example, this tool can be helpful for addressing questions such as the evolution over time of market uncertainty about the outlook for short-term interest rates. Using this technique we show that there has been little change overall since 1997 in our measure of market uncertainty, despite the sharp rise following the financial turbulence in autumn 1998. There is also evidence of a fall in the probability the market attaches to sharp upward movements in rates.

In due course, we intend to make our data on pdfs available on the Bank's Internet site, at www.bankofengland.co.uk

Technical appendix

Constructing fixed expiry-date and constant-horizon probability density functions from exchange-traded options prices

Constructing fixed expiry-date pdfs

The two estimation techniques discussed in this article—the parametric and non-parametric approaches—may be derived from the Cox and Ross (1976) pricing model. This model yields the call option price C_t at time t as the risk-neutral expected pay-off of the option at expiry T, discounted back at the risk-free rate:

$$C(S,X,\tau) = e^{-r\tau}\int_X^{\infty}(S_T - X)g(S_T)dS_T \qquad \textbf{(A1)}$$

where S_T is the terminal underlying asset price at T, $g(S_T)$ is its risk-neutral pdf, X is the option's 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 pay-off of the call $S_T - X$ with the pay-off of the put $X - S_T$ in equation **(A1)** and by integrating from zero to the strike price.

The parametric method

The parametric estimation approach involves specifying a particular functional form for the pdf—$g(S_T)$—and fitting this distribution to the observed range of strike prices via non-linear least squares. Although a range of functional forms has been suggested, the most commonly used is the mixture of two lognormal distributions, as discussed in Bahra (1996 and 1997).

This form is sufficiently flexible to capture the features of distributions we might expect to find implicit in the data. And the mixture of two lognormals is parsimonious because it matches these criteria with just five parameters.

The mixture of two lognormals is given by:

$$g(S_T) = \theta L(\alpha_1, \beta_1) + (1 - \theta)L(\alpha_2, \beta_2) \qquad \textbf{(A2)}$$

where θ, α_1, α_2, β_1 and β_2 are the parameters to be estimated. The fitted call and put prices are given by:

$$\hat{C}_t(S,X_i,\tau) = e^{-r\tau}\int_{X_i}^{\infty}(S_T - X_i)(\theta L(\alpha_1,\beta_1) + (1-\theta)L(\alpha_2,\beta_2))dS_T$$

$$\hat{P}_t(S,X_i,\tau) = e^{-r\tau}\int_{0}^{X_i}(X_i - S_T)(\theta L(\alpha_1,\beta_1) + (1-\theta)L(\alpha_2,\beta_2))dS_T$$

$$\textbf{(A3)}$$

To fit the parameters of the pdf, we minimise the following expression:

$$\min_{\alpha_1,\beta_1,\alpha_2,\beta_2,\theta}\sum_{i=1}^{m}(\hat{C}_{i,t} - C_{i,t})^2 + \sum_{j=1}^{n}(\hat{P}_{j,t} - P_{j,t})^2 \qquad \textbf{(A4)}$$

The non-parametric method

The non-parametric technique for estimating fixed expiry-date pdfs—described in Bliss and Panigirtzoglou (2000) and Cooper (2000)—exploits the result derived by Breedon and Litzenberger (1978) that the pdf can be recovered by calculating the second partial derivative of the call price function with respect to the strike price. This result can be derived simply by taking the second partial derivative of the call price function **(A1)** with respect to the strike price X to get:

$$\frac{\partial^2 C}{\partial X^2} = e^{-r\tau}g(S_T) \qquad \textbf{(A5)}$$

So we just have to adjust the probabilities by $e^{-r\tau}$ to get $g(S_T)$. 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 across the volatility smile, using the Black-Scholes formula to transform this back to prices. The reason for doing this rather than interpolating the call price 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 pdf.

Shimko (1993) used a quadratic functional form to interpolate across the implied volatility smile. Instead, we use a cubic smoothing spline. This is a more flexible non-parametric curve that gives us control of the amount of smoothing of the volatility smile and hence the smoothness of the estimated pdf. But following Malz (1997), we first calculate the Black-Scholes deltas of the options. This is because in practice it is usually easier to interpolate across the volatility smile in 'delta space' than in 'strike price space'. Finally, to generate the implied pdf, we calculate the second partial derivative with respect to the strike price numerically and adjust for the effect of the discount factor.

The method for estimating pdfs with a constant horizon

Our technique for estimating constant-horizon pdfs is based on the non-parametric technique for estimating fixed expiry-date pdfs. The technique involves interpolating across the implied volatilities of options with the same delta, but on different maturities. We interpolate across the implied volatilities for each particular delta rather than each strike price because the range of possible values of delta—between 0 and 1—is not maturity-dependent.

Consider two sets of options contracts on FTSE 100 index futures: the first with two months to maturity and strike prices in the range 5000–6500; and the second with six months to maturity and strike prices in the range 4000–7000. If we wanted to interpolate across implied volatilities for each strike price, the implied volatilities on the second contract corresponding to strike prices outside the range 5000–6500 could not be used. In other words, information on the second contract would be lost if we were to interpolate across implied volatilities for each strike price. This problem does not occur if we interpolate across implied volatilities for each delta, since both contracts have a range of deltas between 0 and 1.

Our technique for estimating constant-horizon pdfs involves the following steps:

- For each delta, we interpolate across the implied volatilities of the options on the different LIFFE contracts using a cubic smoothing spline. We then select the point on the interpolated curve corresponding to the desired maturity of the constant-horizon pdf. For example, to generate the implied volatility corresponding to a value for delta of 0.6 for a six-month constant horizon pdf, we interpolate across the implied volatilities corresponding to a value for delta of 0.6 for the LIFFE contracts with different maturities. We then select the six-month point on the interpolated curve.

- We repeat the process for different values of delta and hence construct a curve of implied volatility against delta—an 'implied volatility smile'—for hypothetical options with six months to maturity.

- We then use the implied volatility smile to generate the constant-horizon pdf using the same non-parametric interpolation method we use for generating fixed expiry-date pdfs.

Glossary of technical terms

The *call price function* relates the prices of call options of the same maturity to their strike price.

The *delta* of an option is the rate of change of the option price with respect to the underlying asset price and is a non-linear transformation of the strike price.

The *implied volatility* is the volatility of the underlying asset price implied by the Black-Scholes (1973) model. The implied volatility is a non-linear transformation of the option price.

Kurtosis is defined as the fourth central moment of a probability distribution, normalised by the fourth power of its standard deviation.

Moneyness is *at-the-money*, *near-the-money*, *in-the-money*, *out-of-the-money*. Options which give the right to buy (ie calls) or sell (ie puts) at a level equal to (or close to) the current futures price of the underlying asset are said to be 'at-the-money' (or 'near-the-money'). Call options which give the right to buy at a level higher (or lower) than the current futures price of the underlying asset are said to be 'out-of-the-money' (or 'in-the-money'). Put options which give the right to sell at a level higher (or lower) than the current price of the underlying asset are said to be 'in-the-money' (or 'out-of-the-money').

Skewness is defined as the third central moment of a probability distribution, normalised by the third power of its standard deviation.

The *strike price* of an option is the price at which the investor can exercise the option.

References

Anderson, N and Sleath, J (1999), 'New estimates of the UK real and nominal yield curves', *Bank of England Quarterly Bulletin*, Vol 39(4), pages 384–92.

Bahra, B (1996), 'Probability distributions of future asset prices implied by option prices', *Bank of England Quarterly Bulletin*, Vol 36(3), 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.

Barone-Adesi, G and Whaley, R (1987), 'Efficient analytic approximation of American option values', *Journal of Finance*, Vol XLII, No 2, pages 301–20.

Bates, D (1996), 'Jumps and stochastic volatility: exchange rate processes in Deutsche Mark options', *Review of Financial Studies*, Vol 9, No 1, pages 69–107.

Black, F and Scholes, M (1973), 'The pricing of options and corporate liabilities', *Journal of Political Economy*, 81, pages 637–54.

Bliss, R and Panigirtzoglou, N (2000), 'Testing the stability of implied pdfs', *Bank of England Working Paper*, forthcoming.

Breedon, D and Litzenberger, R (1978), 'Prices of state-contingent claims implicit in options prices', *Journal of Business*, Vol 51, No 4, pages 621–51.

Cooper, N (2000), 'Testing techniques for estimating implied risk-neutral densities from the prices of European and American-style options', *mimeo*.

Cox, J and Ross, S (1976), 'The valuation of options for alternative stochastic processes', *Journal of Financial Economics*, 3, pages 145–66.

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, 2, pages 327–43.

Malz, A (1995), 'Using options prices to estimate re-alignment 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*, Winter, pages 18–36.

Melick, W R and Thomas, C P (1997), 'Recovering an asset's implied pdf from options prices: an application to crude oil during the Gulf Crisis', *Journal of Financial and Quantitative Analysis*, Vol 32, 1, pages 91–115.

Shimko, D C (1993), 'Bounds of probability', *RISK*, Vol 6, 4, pages 33–37.