
Probability distributions of future asset prices implied by option prices

By Bhupinder Bahra of the Bank's Monetary Instruments and Markets Division.

The most widely used measure of the market's views about the future value of an asset is the mean or average price expectation—a point estimate. This article shows how this information set can be extended by using option prices to estimate the market's entire probability distribution of a future asset price. It also illustrates the potential value of this type of information to the policy-maker in assessing monetary conditions, monetary credibility, the timing and effectiveness of monetary operations, and in identifying anomalous market prices. Finally, the article looks at the limitations in data availability and details some areas for future research.

Introduction

Many monetary authorities routinely use the forward-looking information that is embedded in financial asset prices to help in formulating and implementing monetary policy. For example, they typically look at changes in the forward rate curve implied by government bond prices to assess changes in market perceptions of future short-term interest rates.⁽¹⁾ But, although implied forward rates are informative about the market's mean expectation for future interest rates, they tell us nothing about the range of expected outcomes around such estimates. For this, we can turn to options markets.

An option on a given underlying asset is a contract that gives the holder the right, but not the obligation, to buy or sell that asset at a certain date in the future at a predetermined price. Options that give the holder the right to buy the underlying asset are known as *call* options, while those that give the holder the right to sell the underlying asset are known as *put* options. The predetermined price at which the underlying asset is bought or sold, which is stipulated in an option contract, is known as the *exercise price* or *strike price*. The date at which an option expires is known as the *maturity date*, *exercise date* or *terminal date*. Options that can be exercised only on the maturity date are known as *European* options, while those that can be exercised at any time up to and including the maturity date are known as *American* options.⁽²⁾

If the option holder decides to take up his/her right to buy or sell the underlying asset then he/she would exercise the option against the person with which the contract was agreed (known as the writer of the option). So, for example, if the holder of a call option were to exercise that option against its writer, the writer would be obliged to supply the underlying asset to the holder at the pre-agreed exercise price. Of course, the holder of a call option would consider

exercising it only if the price of the underlying asset lay above the strike price at that time.

Consider a set of European options on the same underlying asset, with the same time-to-maturity, but with different exercise prices. The prices of such options are related to the probabilities attached by the market to the possible values of the underlying security on the maturity date of the options. Intuitively, this can be seen by noting that the difference in the price of two options with adjacent exercise prices will reflect the value attached to the ability to exercise the options when the price of the underlying asset lies between the two exercise prices. This price difference in turn depends on the probability of the underlying asset price lying in this interval.

Such probabilities can be estimated, using the full range of exercise prices, from observed options prices in the form of a risk-neutral probability density (RND) function. A probability density is a measure of the frequency with which a particular event occurs. The area under a probability density function for a given range of possible outcomes gives the probability of the eventual outcome being in that range. Since probabilities must sum to one, the total area under a probability density function must be one. Risk neutral, as used here, means that the probability density function depicts the weights attached by a representative risk-neutral market participant to the possible future values of the underlying asset.

This article describes a technique for estimating implied risk-neutral probability density functions from options prices, and illustrates how the information they provide is additional to mean estimates of future asset prices. Further details on the theory, and a comparison of different techniques for estimating implied RND functions will be given in a forthcoming *Bank of England Working Paper* on the topic.⁽³⁾

(1) See, for example, Breedon (1995) and Deacon and Derry (1994).

(2) For further details about options and other derivative securities, see Hull (1993).

(3) Bahra, B (1996), 'Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Application', *Bank of England Working Paper series*, forthcoming.

How are option prices determined?

The current price of a European option on a non dividend paying asset depends on five underlying parameters:

- (i) the current price of the underlying asset on which the option is written;
- (ii) the time remaining until the maturity date of the option;
- (iii) the (annualised) risk-free rate of interest over the remaining life of the option;
- (iv) the exercise price of the option; and
- (v) the annualised volatility of the underlying asset price over the remaining life of the option.

In order to calculate an option's price, one has to make an assumption about how the price of the underlying asset evolves over the life of the option. Many option pricing models specify a *stochastic* process for the price.⁽¹⁾ Each stochastic process is consistent with a particular RND function for the price of the underlying asset on the expiry date of the option. For example, the classic Black-Scholes (1973) option pricing model assumes that the price of the underlying asset evolves according to the stochastic process known as *geometric Brownian motion*. This implies that the underlying asset is expected to earn a constant rate of return, although the price is subject to independent, normally distributed shocks over the life of the option. Under the geometric Brownian motion assumption, the risk-neutral probabilities attached to various possible outcomes for the price of the underlying asset on the maturity date of the option take the form of a lognormal distribution: a lognormal distribution is a tilted, or 'skewed', bell-shaped curve.

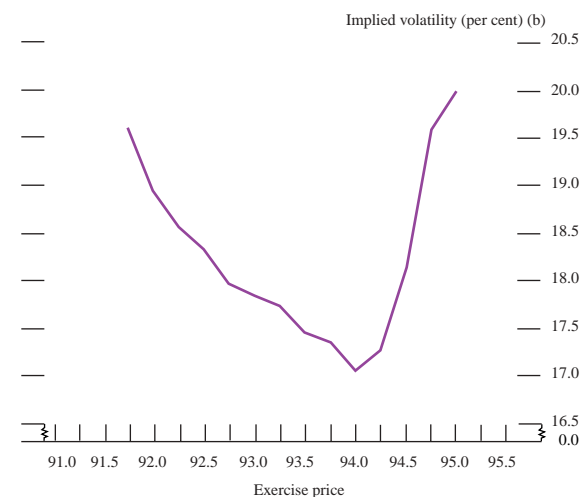
Implied volatilities

Out of the five parameters that determine the price of an option, the only one that is currently unobservable is the future volatility of the underlying asset price. But an estimate of this can be inferred from the prices of options traded in the market: given an option price, one can solve an appropriate option pricing model to obtain a market estimate of the future volatility of the underlying asset price. This type of volatility estimate is known as *implied volatility*.

Under the Black-Scholes assumption that the price of the underlying asset evolves according to geometric Brownian motion, the implied volatility ought to be the same across all exercise prices of options on the same underlying asset and with the same maturity date. But the implied volatilities observed in the market typically vary with the exercise price. In particular, the implied volatilities associated with

exercise prices a long way from the current price of the underlying asset tend to be higher than those associated with exercise prices which are closer to the current price of the underlying asset. The relationship between implied volatility and exercise price is described by what is known as the implied volatility *smile* curve, as illustrated in Chart 1.

Chart 1
Implied volatility smile curve for LIFFE December 1996 options on the short sterling future^(a)



(a) As at 16 April 1996.

(b) Implied volatility is an annualised estimate of the instantaneous standard deviation of the return on the underlying asset over the remaining life of the option.

The existence of the volatility smile curve indicates that market participants make more complex assumptions than geometric Brownian motion about the path of the underlying asset price. And as a result, they attach different probabilities to terminal values of the underlying asset price than those that are consistent with a lognormal distribution. The extent of the curvature of the smile curve indicates the degree to which the market RND function differs from the Black-Scholes (lognormal) RND function. In particular, the greater the curvature, the greater the probability the market attaches to extreme outcomes. This causes the market RND function to have 'fatter tails' than a lognormal density function. In addition, the direction in which the smile curve slopes reflects the direction in which the market RND function is skewed.⁽²⁾

Any variations in the shape of the smile curve are mirrored by corresponding changes in the curvature of the *call pricing function*—the plot of call prices across exercise prices for options on the same underlying asset and with the same time-to-maturity. The slope and curvature of the smile curve, or of the call pricing function, can be translated into probability space to reveal the market's (non-lognormal) implied terminal RND function. There are a number of techniques for doing this, all of which can be related to an approach first taken by Breeden and Litzenberger (1978).

(1) A variable whose value changes over time in an uncertain way is said to follow a stochastic process.

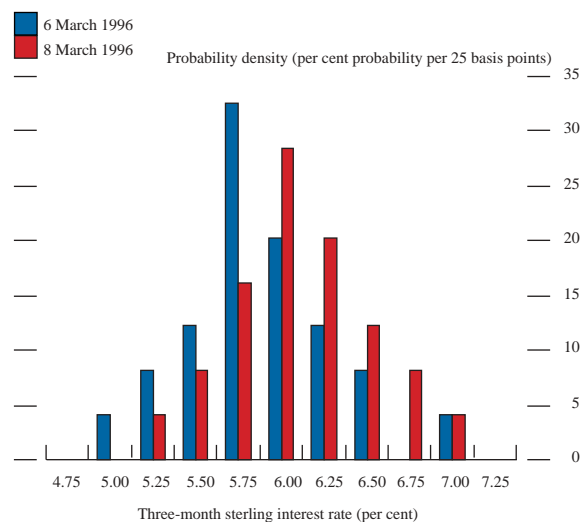
(2) The skewness of a probability density function characterises the distribution of probability either side of the mean.

The Breeden and Litzenberger approach

Breeden and Litzenberger (1978) derived a relationship linking the curvature of the call pricing function to the terminal RND function of the price of the underlying asset. In particular, they showed that the second partial derivative of the call pricing function with respect to the exercise price is directly proportional to the terminal RND function. Details about the derivation of the Breeden and Litzenberger result are given in Bahra (1996). The rest of this article focuses on how this result can be applied in order to estimate market RND functions for short-term interest rates in the future and how such RND functions can be used for policy analysis.

The simplest approach to estimating RND functions is to approximate the second derivative of the call pricing function by calculating the second difference of actual call prices observed across a range of exercise prices.⁽¹⁾ This approach produces the implied risk-neutral histogram of the price of the underlying asset at the maturity date of the options.⁽²⁾ Chart 2 shows how the implied histogram for the three-month sterling interest rate on 19 June 1996 (as implied by the June short sterling futures price) changed between 6 March and 8 March 1996, a period which included a cut of 25 basis points in official UK interest rates and the publication of stronger-than-expected US non-farm payrolls data.⁽³⁾

Chart 2
Implied risk-neutral histograms for the three-month sterling interest rate in June 1996^(a)



(a) Derived using LIFFE June 1996 options on the short sterling future, as at 6 March and 8 March 1996.

The main drawback of this approach is that it does not smooth out irregularities in observed call pricing functions. These may be due, in cases where bid-ask spreads are

observed instead of actual traded prices, to measurement errors arising from using middle prices. Irregular call pricing functions may also arise if readings are taken at slightly different times. Such irregularities can result in negative implied probabilities. Also, the procedure provides no systematic way of modelling the tails of the probability distributions, which are not observable due to the limited range of exercise prices traded in the market.

But sensible continuous RND functions can be obtained by smoothing the call pricing function in a way that places less weight on data irregularities while preserving its overall form under the assumption of no arbitrage. Since option prices are only observed at discrete intervals across a limited range of exercise prices, the procedures for doing this essentially amount to interpolating between observed exercise prices, and extrapolating outside their range to model the tail probabilities.

Three related approaches have been used in the literature:

- (i) the RND function is derived directly from a particular specification of the call pricing function (or of the implied volatility smile curve);⁽⁴⁾
- (ii) assumptions are made about the stochastic process that governs the price of the underlying asset and the RND function is inferred from it;⁽⁵⁾ and
- (iii) an assumption is made about the form of the RND function itself and its parameters are recovered by minimising the distance between the observed option prices and those that are generated by the assumed functional form.⁽⁶⁾

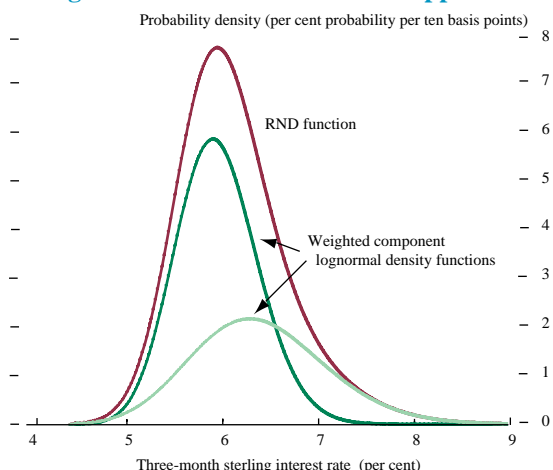
The lognormal mixture distribution approach

In our research we have adopted the third approach, which focuses directly on the RND function. This means we impose a minimum of structure on the stochastic process of the price of the underlying asset. For the purposes of policy analysis, the functional form assumed for the RND function should be relatively flexible. In particular, it should be able to capture the main contributions to the smile curve, namely the skewness and the kurtosis (ie fatness of the tails) of the underlying distribution. In light of these criteria, we assume that the RND function is a weighted sum of two independent lognormal density functions and we then estimate their parameters from observed option prices.⁽⁷⁾ Each lognormal density function is completely defined by two parameters. The values of these parameters, and the relative weighting applied to the two density functions, together determine the overall shape of the implied RND function.

(1) Such second difference estimates are directly proportional to the probabilities attached by the market to the underlying asset price lying in a fixed interval around each of the strike prices when the options expire. The constant of proportionality is the present value of a zero-coupon bond that pays £1 at maturity, with the discount rate being the risk-free rate of interest.
(2) For further examples of this approach, see Neuhaus (1995).
(3) The histograms were calculated using data for the LIFFE June 1996 option on the short sterling future. The LIFFE settlement prices were used to avoid the problems associated with asynchronous data.
(4) See Bates (1991), Jarrow and Rudd (1982), Longstaff (1992, 1995), Malz (1995a) and Shimko (1993).
(5) See Bates (1991, 1995), and Malz (1995b).
(6) See Bahra (1996), Jackwerth and Rubinstein (1995), Melick and Thomas (1994), and Rubinstein (1994).
(7) Details of the minimisation problem are given in the Technical annex.

Chart 3 shows an example of an implied RND function derived from LIFFE options on the short sterling future using the method described above, which we call the two-lognormal mixture distribution approach. It also shows the (weighted) component lognormal density functions of the RND function.

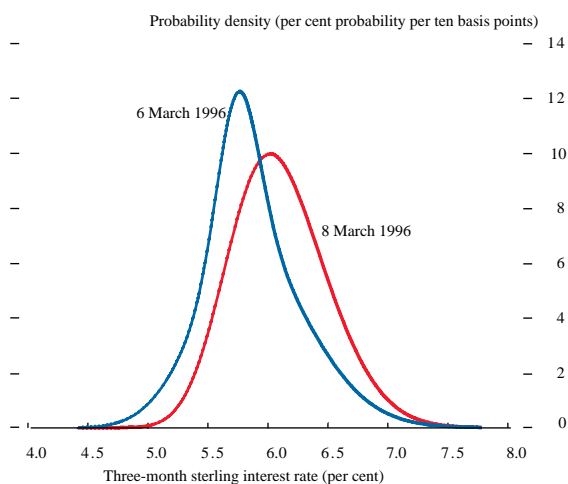
Chart 3
An implied RND function derived using the two-lognormal mixture distribution approach^(a)



(a) Shown with its (weighted) component lognormal density functions. This RND function was derived using LIFFE December 1996 options on the short sterling future as at 10 June 1996.

Chart 4 further illustrates the method, showing how the implied RND function for the three-month sterling interest rate on 19 June 1996 changed between 6 March and

Chart 4
Implied RND functions for the three-month sterling interest rate in June 1996^(a)



(a) Derived using LIFFE June 1996 options on the short sterling future, as at 6 March and 8 March 1996. These graphs illustrate the result when using the two-lognormal mixture distribution approach.

8 March 1996. The left axis of the chart depicts probability density. This is a measure of the frequency with which events occur. On the chart, the probability density associated with a given future interest rate is approximately equal to the probability of the outcome lying in a corridor of five basis points either side of that rate.

The shape of the distribution would be expected to vary over time as news arrives and option prices adjust to incorporate changing beliefs about future events. The two-lognormal mixture distribution can incorporate a wide variety of possible functional forms which, in turn, are able to accommodate a wide range of possible scenarios. This includes, for example, a situation in which the market believes that the terminal price of the underlying asset is most likely to take one of two possible values, in which case it attaches high probabilities around those levels, giving rise to an implied RND function with two modes, ie the market has a bi-modal view.

It is important to remember that the implied density functions derived are risk neutral, that is, they are equivalent to the true market density functions only when investors are risk neutral. In reality investors are likely to be risk averse, and option prices will incorporate these preferences towards risk as well as beliefs about future outcomes. To distinguish between these two factors would require specification of the aggregate market utility function (which is unobservable) and estimation of the corresponding coefficient of risk aversion. But, even if the market does demand a premium for taking on risk, the true market implied density function may not differ very much from the RND function, at least for some markets.⁽¹⁾ Moreover, on the assumption that the market's aversion to risk is relatively stable over time, changes in the RND function from one day to the next should mainly reflect changes in investors' beliefs about future outcomes for the price of the underlying asset.

Using the information contained in implied RND functions

We now illustrate how the information contained in implied RND functions may be used in formulating and implementing monetary policy. We begin by describing various summary measures for density functions and then suggest a way to validate the two-lognormal mixture distribution approach. Next, we outline different ways in which implied RND functions may be used by the policy-maker. Finally, we discuss some caveats and limitations in data availability, and detail some areas for future research.

Summary statistics

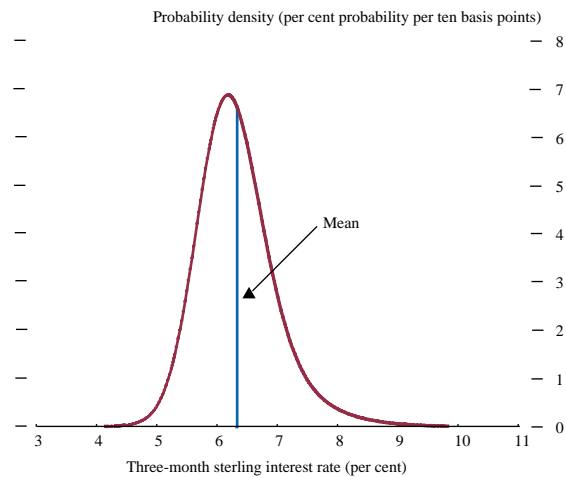
Much of the information contained in RND functions can be captured through a range of summary statistics. For example, the mean is the expected future value of the underlying asset, or the average value of all possible future outcomes. Forward-looking information, whether derived directly from futures prices, or indirectly via bond yields is typically based on the mean. The median, which has 50% of the distribution on either side of it, is an alternative measure of the centre of a distribution. The mode, on the other hand, is the most likely future outcome. The standard deviation of an implied RND function is a measure of the

(1) For example, Rubinstein (1994) converts an RND function for an equity index to a 'subjective' density function under the assumption that the representative investor maximises his/her expected utility of wealth with constant relative risk aversion (CRRA). He finds that for assumed market risk premia of between 3.3% and 5%, the subjective distribution is only slightly shifted to the right relative to the risk-neutral distribution, and that the qualitative shapes of the two distributions are quite similar.

uncertainty around the mean and is analogous to the implied volatility measure derived from options prices. An alternative dispersion statistic is the interquartile range (IQR). This gives the distance between the 25% quartile and the 75% quartile, that is, the central 50% of the distribution lies within it. Skewness characterises the distribution of probability either side of the mean. A positively skewed distribution is one for which there is less probability attached to outcomes higher than the mean than to outcomes below the mean. Kurtosis is a measure of the 'peakedness' of a distribution and/or the likelihood of extreme outcomes: the greater this likelihood, the fatter the tails of the distribution. These summary statistics provide a useful way of tracking the behaviour of RND functions over the life of a single contract and of making comparisons across contracts.


Charts 5 and 6 show the RND functions, as at 4 June 1996, for the three-month sterling interest rate in December 1996

Chart 5
Implied RND function for the three-month sterling interest rate in December 1996^(a)



(a) Derived using LIFFE December 1996 options on the short sterling future, as at 4 June 1996.


Chart 6
Implied RND function for the three-month sterling interest rate in March 1997^(a)



(a) Derived using LIFFE March 1997 options on the short sterling future, as at 4 June 1996.


and in March 1997. Charts 7 and 8 depict the RND functions, also as at 4 June 1996, for the three-month Deutsche Mark interest rate in the same months. Table A shows the summary statistics for these four distributions.

Chart 7
Implied RND function for the three-month Deutsche Mark interest rate in December 1996^(a)



(a) Derived using LIFFE December 1996 options on the Euromark future, as at 4 June 1996.

Chart 8
Implied RND function for the three-month Deutsche Mark interest rate in March 1997^(a)



(a) Derived using LIFFE March 1997 options on the Euromark future, as at 4 June 1996.

The means of the distributions are equivalent to the interest rates implied by the current prices of the relevant futures contracts, and are lower in Germany than in the United Kingdom.⁽¹⁾ For both countries, the dispersion statistics (standard deviation and IQR) are higher for the March 1997 contract than for the December 1996 contract. One would expect this since, over longer time horizons, there is more uncertainty about the expected outcome. Chart 9 confirms this, showing the upper and lower quartiles with the mean and the mode for the three-month sterling interest rate on four different option maturity dates as at 15 May 1996. It can be seen that the IQR is higher for contracts with longer maturities. Also, the standard deviations of the two distributions for the sterling rate are higher than the

(1) The mean of an implied RND function should equal the forward value of the underlying asset. In this case the underlying assets are short-term interest rate futures contracts. The expected growth rate of a futures price in a risk-neutral world is zero. Hence, the means of the implied RND functions are equal to the interest rates implied by the respective current futures prices.

Table A
Summary statistics for the three-month sterling and Deutsche Mark interest rates in December 1996 and March 1997^(a)

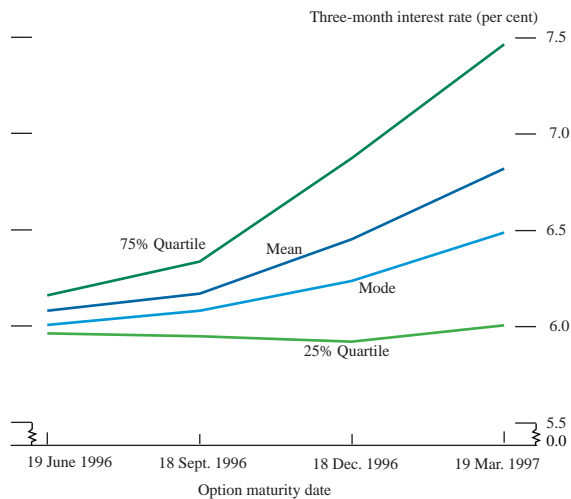
Sterling	December 1996	March 1997
Mean	6.33	6.66
Mode	6.18	6.43
Median	6.27	6.56
Standard deviation	0.66	1.01
Interquartile range	0.80	1.19
Skewness	0.83	0.76
Kurtosis (b)	4.96	4.67

Deutsche Mark	December 1996	March 1997
Mean	3.45	3.73
Mode	3.29	3.47
Median	3.39	3.62
Standard deviation	0.55	0.84
Interquartile range	0.69	0.95
Skewness	0.75	1.16
Kurtosis	4.27	6.06

(a) Derived using LIFFE December 1996 and March 1997 options on the short sterling and Euromark futures, as at 4 June 1996.
 (b) A normal distribution has a fixed kurtosis of three.

corresponding standard deviations of those for the Deutsche Mark rate, suggesting greater uncertainty about the level of future short-term rates in the United Kingdom than in Germany. Another feature of all four distributions is that they are positively skewed, indicating that there is less probability to the right of each of the means than to their left. The fact that the mode is to the left of the mean is usually also indicative of a positive skew. This feature is discussed in greater detail below.

Chart 9
Implied RND summary statistics for the three-month sterling interest rate on four different option maturity dates^(a)



(a) Derived using LIFFE options on the short sterling future, as at 15 May 1996.

Validation

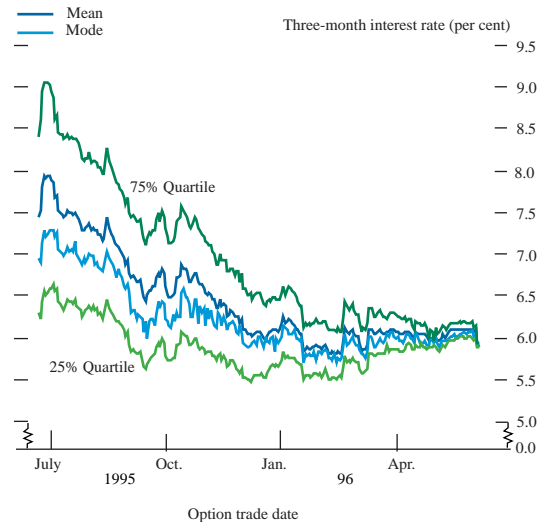
In deciding whether to place reliance on the information extracted using a new technique, one not only needs to be confident of the theory, but must also test whether in practice changes in the expectations depicted are believable in light of the news reaching the market. In the case of short-term interest rate expectations, we sought to do this by examining the way RND functions for short-term sterling interest rates change over time, and by comparing the RND functions for short-term sterling interest rates with those

from Germany, a country with different macroeconomic conditions and monetary history.

Analysing changes in implied RND functions over time

Charts 10 and 11 show a convenient way of representing the evolution of implied RND functions over the life of a single option contract. Chart 10 shows the market's views of the three-month sterling interest rate on 19 June 1996 (as implied by the prices of LIFFE June short sterling futures

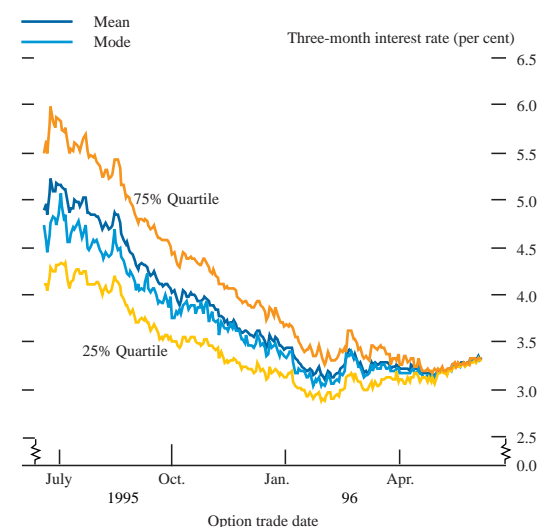
Chart 10
Implied RND summary statistics for the three-month sterling interest rate in June 1996^(a)



(a) Derived using LIFFE June 1996 options on the short sterling future.

options) between 22 June 1995 and 7 June 1996. Chart 11 shows the same type of information for the three-month Deutsche Mark interest rate on 17 June 1996 (as implied by

Chart 11
Implied RND summary statistics for the three-month Deutsche Mark interest rate in June 1996^(a)



(a) Derived using LIFFE June 1996 options on the Euromark future.

the prices of LIFFE June Euromark futures options) between 20 June 1995 and 7 June 1996. Both charts depict the mean, mode, and the lower (25%) and upper (75%) quartiles of the distributions.

These time-series representations of implied RND functions convey how market uncertainty about the expected outcome changed over time; an increase in the distance between the lower and upper quartiles indicates that the market became more uncertain about the expected outcome. Charts 10 and 11 also convey information about changes in the skewness of the implied distributions. For example, the location of the mean relative to the lower and upper quartiles is informative of the direction and extent of the skew. Movements in the mean relative to the mode are also indicative of changes in skewness.

Generally, both sets of implied RND functions depict falling forward rates over the period analysed, as evidenced by the downward trend in the mean and mode statistics. At the same time, the gaps between these measures narrowed, suggesting that the distribution of market participants' expectations was becoming more symmetrical as the time horizon shortened. Charts 10 and 11 also show that as the maturity date of a contract is approached, the distributions typically become less dispersed causing the quartiles to converge upon the mean. This is because as the time horizon becomes shorter, the market, all other things being equal, becomes more certain about the terminal outcome due to the smaller likelihood of extreme events occurring. Another feature of the distributions is that the mode is persistently below the mean expectation in both countries, indicating a positive skew to expectations of future interest rates. In the United Kingdom, this might be interpreted as reflecting political uncertainty, with the market attaching some probability to much higher short-term rates in the future. However, in Germany the macroeconomic and political conditions are different and yet the RND functions are also positively skewed.

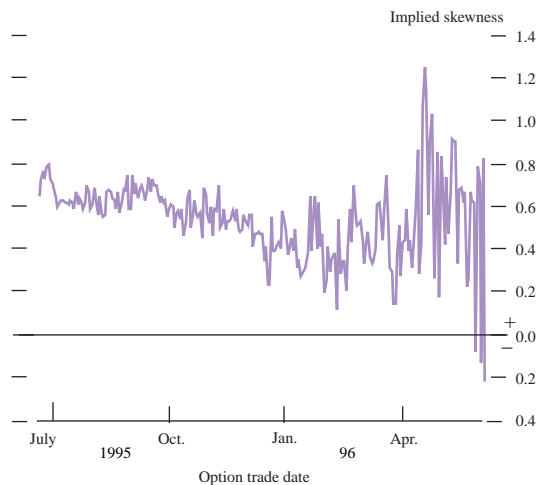
One possible explanation is that the market perceives there to be a lower bound on nominal interest rates at zero. In this case, the range of possible outcomes below the current rate is restricted, whereas the range of possible outcomes above the current rate is, in principle, unlimited. If market participants are generally uncertain, that is, they attach positive probabilities to a wide range of possible outcomes, the lower bound may naturally result in the RND function having a positive skew. Moreover, the lower the current level of rates, the more positive this skew may be for a given degree of uncertainty.

Charts 12 and 13 show how the skewness and kurtosis for the three-month sterling interest rate on 19 June 1996 changed between 22 June 1995 and 7 June 1996. It is notable that, unlike the measures of dispersion, these statistics exhibit no clear trend over their life cycles. Also, they appear to become more volatile towards the end of the contract's life.

Analysing changes in implied RND functions around specific events

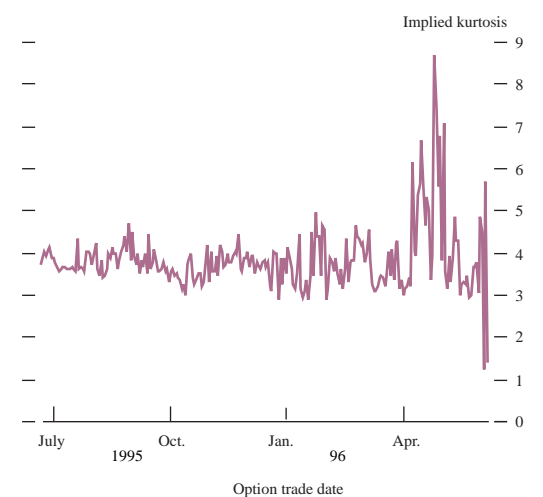
A detailed example of a change in perceptions following a particular news event is given in Chart 14 which shows the

Chart 12
Implied skewness for the three-month sterling interest rate in June 1996^(a)



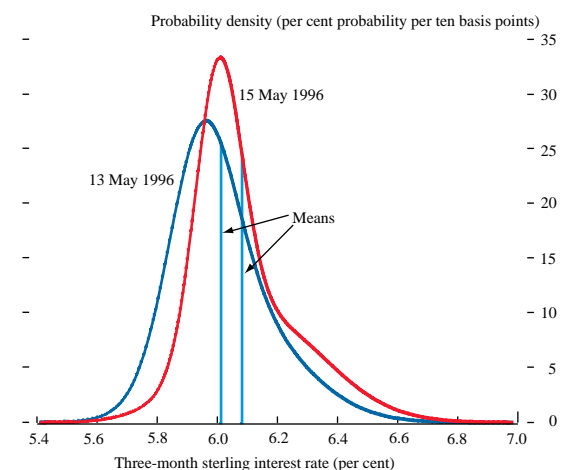
(a) Derived using LIFFE June 1996 options on the short sterling future.

Chart 13
Implied kurtosis for the three-month sterling interest rate in June 1996^(a)



(a) Derived using LIFFE June 1996 options on the short sterling future.

Chart 14
Change in the implied RND function for the three-month sterling interest rate in June 1996 around the publication of the May 1996 Inflation Report^(a)

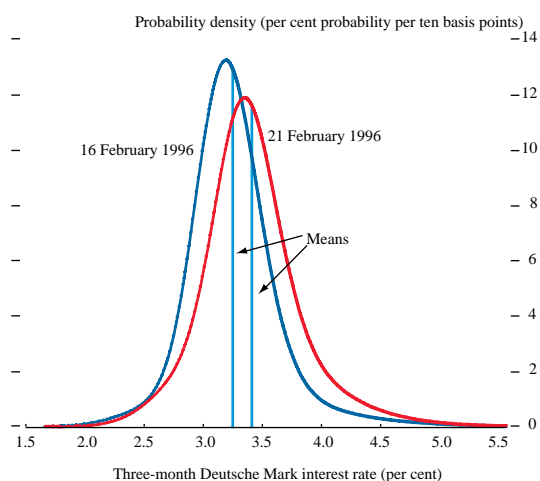


(a) Derived using LIFFE June 1996 options on the short sterling future, as at 13 May and 15 May 1996.

change in the shape of the implied RND function for the three-month sterling interest rate in June 1996 around the publication of the May 1996 *Inflation Report* on 14 May. The *Inflation Report* concluded that it was marginally more likely than not that inflation would be above 2.5% in two years' time were official rates to remain unchanged throughout that period. This was followed by an upward revision of the market's mean expectation for short-term interest rates between 13 May and 15 May. However, it seems that this upward move was not driven so much by a rightward *shift* in the distribution as by a change in the entire *shape* of the distribution; a reallocation of probability from outcomes between 5.6% and 5.9% to outcomes between 5.9% and 6.6% resulted in a fatter right tail which was in part responsible for the upward movement in the mean. This type of change in the shape of implied RND functions is illustrative of how they can add value to existing measures of market expectations such as the mean.

A similar change in market sentiment can be observed in Germany between 16 and 21 February 1996, ahead of the publication of the German M3 figure on 23 February. Chart 15 shows how the implied RND function for the three-month Deutsche Mark interest rate in June 1996 changed between these two dates. There was a significant shift in probability from outcomes between 2.5% and 3.3% to outcomes between 3.3% and 4.5%, apparently driven by market speculation ahead of the publication of the data. In particular, on 21 February the market attached a much higher probability to short-term rates being around 4% in June than it did on 16 February.

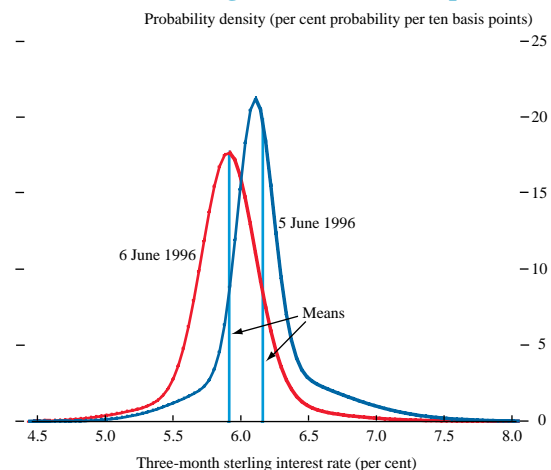
Chart 15
Change in the implied RND function for the three-month Deutsche Mark interest rate in June 1996^(a)



(a) Derived using LIFFE June 1996 options on the Euromark future, as at 16 February and 21 February 1996.

The cut in UK official interest rates on 6 June 1996 provides an illustration of how market perceptions may change around a monetary policy decision. Chart 16 shows the change in the shape of the implied RND function for the three-month sterling interest rate in September 1996 between 5 and 6 June 1996. Table B shows the summary statistics for the RND functions on each of these dates.

Chart 16
Change in the implied RND function for the three-month sterling interest rate in September 1996^(a)



(a) Derived using LIFFE September 1996 options on the short sterling future, as at 5 June and 6 June 1996.

Table B
Summary statistics for the RND functions in Chart 16

	5 June 1996	6 June 1996
Mean	6.16	5.91
Mode	6.11	5.91
Median	6.12	5.91
Mean minus mode	0.05	0.01
Standard deviation	0.35	0.30
Interquartile range	0.27	0.31
Lower quartile	6.00	5.76
Upper quartile	6.27	6.07
Skewness	0.85	0.22
Kurtosis (a)	6.61	7.02

(a) A normal distribution has a fixed kurtosis of three.

The first point to note is that the mean moved down by 25 basis points, which was the size of the interest rate cut. Second, the distribution on 6 June was more symmetrical (in fact the mean was almost equal to the mode) and had a higher standard deviation compared to the previous day; ie the market was more uncertain on 6 June than on 5 June about the short-term interest rate in September, and attached the same weight to it being above the mean as to it being below the mean. The change in the degree of skewness can also be seen by the shift in probability from outcomes between 6% and 7% to outcomes between 5.5% and 6%, resulting in a much thinner right tail and a left tail which was only slightly fatter. By comparison with other day-to-day movements, this particular change in the shape of the implied distribution was quite large indicating the extent to which the market was surprised by the rate cut.

The above examples suggest that the two-lognormal mixture distribution approach is validated by recent market developments in the United Kingdom and in Germany. Although the mean expectation remains a key summary statistic, on the basis of these and other examples there is no reason to doubt that implied RND functions can add to our understanding of short-term interest rate expectations.

Use of implied RND functions by monetary authorities

We now discuss four ways in which the policy-maker may use implied RND functions.

Assessing monetary conditions

Assuming that financial market expectations are indicative of those in the economy as a whole, RND functions have the potential to improve the authorities' ability to assess monetary conditions on a day-to-day basis.

In principle, the whole probability distribution of future short-term interest rates is relevant to the determination of economic agents' behaviour. A lot of this information is captured in the mean of the distribution, which can already be observed directly from the yield curve or forward rates, but other summary statistics may add explanatory power. For example, suppose that agents tend to place less weight on extreme interest rate outcomes when taking investment or consumption decisions than is assumed in the mean of the interest rate probability distribution. In this case, a trimmed mean—in which the probabilities attached to extreme outcomes are ignored or given reduced weight—may reflect the information used by agents better than the standard mean, and so may provide a better indication of monetary conditions for the monetary authorities. Much of the time the standard mean and the trimmed mean may move together, but one could envisage circumstances in which the standard mean is influenced by an increase in the probabilities attached to very unlikely outcomes, while the trimmed mean is less affected. Similar issues would arise if investors or consumers placed *more* weight on extreme interest rate outcomes than allowed for in the standard mean.

At present, this kind of scenario is entirely speculative. Further empirical research is required to assess whether summary statistics such as an adjusted mean, the mode, median, interquartile range, skewness and kurtosis can add explanatory power to the standard mean interest rate in conventional economic models.

RND functions may also provide evidence of special situations influencing the formation of asset price expectations. For example, if two distinct economic or political scenarios meant that asset prices would take very different values according to which scenario occurred, then this might be revealed in bi-modal probability distributions for various asset prices.

Assessing monetary credibility

A monetary strategy to achieve a particular inflation target can be described as credible if the public believes that the government will carry out its plans. So, a relative measure of credibility is the difference between the market's perceived distribution of the future rate of inflation and that of the authorities.⁽¹⁾ Some information on this is already available in the United Kingdom in the form of implied forward inflation rates, calculated from the yields of index-linked and conventional gilts. But this only gives us the mean of the market's probability distribution for future inflation. Even if this mean were the same as the

authorities' target, this could mask a lack of credibility if the market placed higher weights on much lower and much higher inflation outcomes than the authorities.

Unfortunately, there are at present no instruments which enable the extraction of an RND function for inflation. Future research on implied probability distributions for long-term interest rates revealed by options on the long gilt future may, however, help in this respect, to the extent that most of the uncertainty over long-term interest rates—and hence news in the shape of a long gilt RND function—may plausibly be attributed to uncertainty over future inflation.

Assessing the timing and effectiveness of monetary operations

Implied RND functions from options on short-term interest rates indicate the probabilities the market attaches to various near-term monetary policy actions. These probabilities are in turn determined by market participants' expectations about news and their view of the authorities' reaction function.

In this context, implied RND summary statistics may help the authorities to assess the market's likely reaction to particular policy actions. For example, a decision to raise short-term interest rates may have a different impact on market perceptions of policy when the market appears to be very certain that rates will remain unchanged, (as evidenced by a narrow and symmetric RND function for future interest rates) from when the mean of the probability distribution for future rates is the same, but the market already attaches non-trivial probabilities to sharply higher rates, albeit counterbalanced by higher probabilities attached to certain lower rates.

Equally, implied RND functions may help in the *ex post* analysis of policy actions. For example, if the shape and location of the implied RND function for short-term interest rates three months ahead remains the same following a change in base rates, this suggests, all other things being equal, that the market fully expected the change in monetary stance. By contrast a constant mean is less informative because it could disguise significant changes in skewness and kurtosis.

Implied probability distributions may also be useful for analysing market reactions to money-market operations which do not involve a change in official rates, or events such as government bond auctions. These can be assessed either directly by looking at probability distributions from the markets concerned, or indirectly by looking at related markets.

Identifying market anomalies

All of the above uses of RND data assume that markets are perfectly competitive and that market participants are rational. But provided one has overall confidence in the

(1) For further explanation, see King (1995).

technique used, RND functions may help to identify occasional situations where one or other of these assumptions does not hold, essentially because the story being told is not believable.

For example, in the face of an ‘abnormal’ asset price movement—such as a stock market crash or a sharp jump in the nominal exchange rate, which is not easily explained by news hitting the market—the information embedded in options prices for this and related assets may help the authorities to understand whether the movement in question is likely to be sustained with consequent macroeconomic effects, or whether it reflects a temporary phenomenon, possibly due to market failure. For example, if RND functions suggest that the market factored in the possibility of the very large asset price movement because it purchased insurance against the move in advance, then the amount of news required to trigger the change might reasonably be expected to be less than in the situation where there was no advance knowledge. This in turn might make it more believable that the move reflected fundamentals and hence would be sustained.

Limitations in data availability

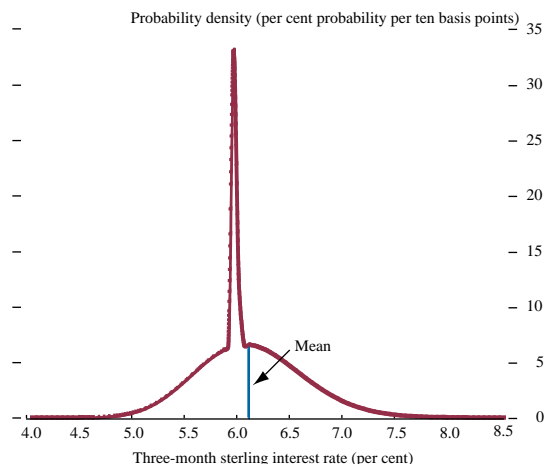
The most important limitation, from the point of view of a monetary authority, is that there are no markets that allow us directly to assess uncertainty about future inflation. To learn about the market’s future inflation distribution would require a market in options on inflation, for example, options on annual changes in the retail prices index (RPI), or a market in options on real rates, as in index-linked bond futures. This would reveal what price agents were willing to pay to insure themselves against the risks to the inflation outcome, and hence the probabilities they attached to various future inflationary outcomes. Neither inflation options, nor options on index-linked bond futures are traded on exchanges anywhere in the world. However, such instruments could conceivably be available in the future.

Another limitation is that the technique is restricted to European options, whilst many of the more liquid exchange-traded options are often American.⁽¹⁾ This restriction is a feature of most of the existing techniques for deriving RND functions. Fairly complex extensions of these techniques are required to estimate terminal RND functions from the prices of American options.⁽²⁾ Even then the RND function can only be derived within a bound that allows for the possibility that the options may be exercised at any time before the maturity date.

There are also limitations to the quality of the data that is available. Some option contracts are fairly illiquid, particularly at those strike prices which are a long way above or below the prevailing market price of the underlying asset. The prices of such contracts may be less informative

about market expectations, or may not be available. This data limitation sometimes results in sudden changes in the degree of curvature of the option pricing function. The two-lognormal mixture distribution approach may in turn be sensitive to this. Chart 17 shows an example of the sort of (implausibly) spiked RND function that has on occasion resulted when there are relatively few data observations across strike prices.

Chart 17
Implied RND function for the three-month sterling interest rate in September 1996^(a)



(a) Derived using LIFFE September 1996 options on the short sterling future, as at 8 May 1996.

To derive implied RND functions we need options prices across the widest possible range of strike prices. To ensure that they are representative of the market’s views, and that they can be estimated regularly, we use exchange-traded options contracts. But these have a limited number of fixed maturity dates, which is problematic when deriving time series of distributions and when assessing changes in market perceptions of short-term rates in the very near future. For example, if there are three months remaining until the nearest option maturity date, it is not possible to determine the market’s perceptions of the short-term rate in one month’s time. Also, because it is not possible with exchange-traded options to ensure that intra-day call and put prices are observable across exercise prices at the same time, only (end-of-day) settlement prices are usable in practice.

Conclusions

This article has shown how the information contained in implied risk-neutral probability density functions estimated from options prices can add to the type of forward-looking information available to policy-makers. To the extent that the distribution around the mean is observed to change in shape over time, measures such as the standard deviation, mode, interquartile range, skewness and kurtosis are useful in quantifying these changes in market perceptions. But, a

(1) LIFFE options on interest rate futures, although American, can be treated as European. This is because they are margined daily, which means that the buyer is not required to pay the option premium up front. So, the buyer can keep the position open at zero cost for as long as favourable movements in the underlying price generate positive cash flows into his/her margin account, whilst losses can be mitigated by closing out the position. This means it is never optimal for the buyer to exercise such options early.

(2) See, for example, Melick and Thomas (1994).

good deal of further research, including event studies and the use of RND summary statistics in addition to the mean in classic economic models, is required to extract the maximum benefit from such information.

As a first step, it is important to be able to identify when a particular change in an implied probability distribution is significant by historical standards. One way of doing this is to establish suitable benchmarks. This would enable a large change in the shape of an RND function to be compared with changes in market perceptions at the time of a significant economic event in the past. In addition, RND

functions could be estimated over the life cycles of many historical contracts for the same underlying asset in order to calculate average values for their summary statistics at particular points in the life cycle. These average values would identify the characteristics of a typical implied RND function during its life cycle. The Bank plans to calculate this information for the implied RND functions of short-term sterling and Deutsche Mark interest rates. It is also in the process of implementing the technique discussed in this article for options on long-term interest rate futures and for currency and equity options.

This annex describes the objective function that we minimise in the two-lognormal mixture distribution approach to estimating RND functions.

The price of a European call option can be written as the discounted sum of all expected future payoffs in a risk-neutral world, that is,

$$c(X) = e^{-r(T-t)} \int_X^{\infty} q(S_T)(S_T - X) dS_T \quad (1)$$

where $c(X)$ is the price of a European call option with exercise price X , r is the (annualised) risk-free rate of interest over the remaining life of the option, $T - t$ is the time remaining until the maturity date, T , of the option, S_T is the price of the underlying asset on the maturity date, and $q(S_T)$ is the RND function of S_T .

In theory any functional form for the RND function, $q(S_T)$, can be specified in equation (1), and its parameters estimated by numerical optimisation. But, given that options are traded across a finite range of exercise prices, there are limits on the number of distributional parameters that can be estimated from the data. As noted in the article, a flexible and numerically tractable parametric specification for $q(S_T)$, which is consistent with observed financial asset price distributions, is a weighted sum of two independent lognormal density functions. Under this assumption the call pricing equation becomes:

$$c(X) = e^{-r(T-t)} \int_X^{\infty} [\theta L(\alpha_1, \beta_1; S_T) + (1-\theta)L(\alpha_2, \beta_2; S_T)](S_T - X) dS_T \quad (2)$$

where the weight parameter, θ , lies between zero and one, and $L(\alpha_i, \beta_i; S_T)$ denotes a lognormal density function for variable S_T with parameters α_i and β_i .

For fixed values of X and $T - t$, and for a set of values for the five distributional parameters and r , equation (2) can be used to provide a fitted value of $c(X)$. This calculation can be applied across all exercise prices to minimise the sum of squared errors, with respect to the five distributional parameters and r , between the option prices generated by the mixture distribution model and those actually observed in the market. In practice, since we can observe interest rates which closely approximate r , we use this information to fix r , and thereby reduce the complexity of the minimisation problem. So, the minimisation is carried out with respect to the five distributional parameters only.

Since both call and put options are priced off the same underlying distribution, we include both sets of prices in the minimisation. Also, in the absence of arbitrage opportunities, the mean of the implied RND function should equal the forward price of the underlying asset. In this sense we can use the incremental information provided by the forward price of the underlying asset by including it as an additional observation in the minimisation procedure. The minimisation problem is:

$$\text{Min}_{\alpha_1, \alpha_2, \beta_1, \beta_2, \theta} \sum_{i=1}^n [c(X_i) - \hat{c}_i]^2 + \sum_{i=1}^n [p(X_i) - \hat{p}_i]^2 + \left[\theta e^{\alpha_1 + \frac{1}{2}\beta_1^2} + (1-\theta)e^{\alpha_2 + \frac{1}{2}\beta_2^2} - e^{r(T-t)} S_t \right]^2 \quad (3)$$

subject to $\beta_1, \beta_2 > 0$ and $0 < \theta < 1$, over the observed strike range $X_1, X_2, X_3, \dots, X_n$, where $c(X_i)$ and $p(X_i)$ are the observed prices of call and put options, respectively, with exercise prices X_i , and S_t is the time- t (current) price of the underlying asset. The (weighted) sum of the first two exponential terms in the last bracket in equation (3) represents the mean of the mixture RND function.

References

- Bahra, B** (1996), 'Implied Risk-Neutral Probability Density Functions From Option Prices: Theory and Application', *Bank of England Working Paper series*, forthcoming.
- Bates, D S** (1991), 'The Crash of '87: Was it Expected? The Evidence from Options Markets', *Journal of Finance*, 46(3), pages 1,009–44.
- Bates, D S** (1995), 'Post-87 Crash Fears in S&P 500 Futures Options', *Wharton Business School*, Working Paper.
- Black, F and Scholes, M** (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy*, 81, pages 637–59.
- Breedon, D T and Litzenberger, R H** (1978), 'Prices of State-Contingent Claims Implicit in Option Prices', *Journal of Business*, Vol 51, No 4, pages 621–51.
- Breedon, F** (1995), 'Bond prices and market expectations of inflation', *Bank of England Quarterly Bulletin*, Vol 35, May, pages 160–65.
- Deacon, M and Derry, A** (1994), 'Estimating market interest rate and inflation expectations from the prices of UK government bonds', *Bank of England Quarterly Bulletin*, Vol 34, August, pages 232–40.
- Hull, J C** (1993), 'Options, Futures, and Other Derivative Securities', 2nd ed, *Prentice Hall International*.
- Jackwerth, J C and Rubinstein, M** (1995), 'Implied Probability Distributions: Empirical Analysis', *Haas School of Business, University of California*, Working Paper, No 250.
- Jarrow, R and Rudd, A** (1982), 'Approximate Option Valuation for Arbitrary Stochastic Processes', *Journal of Financial Economics*, 10, pages 347–69.
- King, M** (1995), 'Credibility and monetary policy: theory and evidence', *Bank of England Quarterly Bulletin*, Vol 35, February, pages 84–91.
- Longstaff, F** (1992), 'An Empirical Examination of the Risk-Neutral Valuation Model', *College of Business, Ohio State University, and the Anderson Graduate School of Management, UCLA*, Working Paper.
- Longstaff, F** (1995), 'Option Pricing and the Martingale Restriction', *Review of Financial Studies*, Vol 8, No 4, pages 1,091–124.
- Malz, A** (1995a), 'Recovering the Probability Distribution of Future Exchange Rates From Option Prices', *Federal Reserve Bank of New York*, mimeo.
- Malz, A** (1995b), 'Using Option Prices to Estimate Realignment Probabilities in the European Monetary System', *Federal Reserve Bank of New York Staff Reports*, No 5.
- Melick, W R and Thomas, C P** (1994), 'Recovering an Asset's Implied PDF From Option Prices: An Application to Crude Oil During the Gulf Crisis', *Federal Reserve Board, Washington*, Working Paper.
- Neuhaus, H** (1995), 'The Information Content of Derivatives for Monetary Policy: Implied Volatilities and Probabilities', *Deutsche Bundesbank Economic Research Group*, Discussion Paper No 3/95.
- Rubinstein, M** (1994), 'Implied Binomial Trees', *Journal of Finance*, Vol LXIX, No 3, pages 771–818.
- Shimko, D** (1993), 'Bounds of Probability', *RISK*, Vol 6, No 4.