

The *Inflation Report* projections: understanding the fan chart

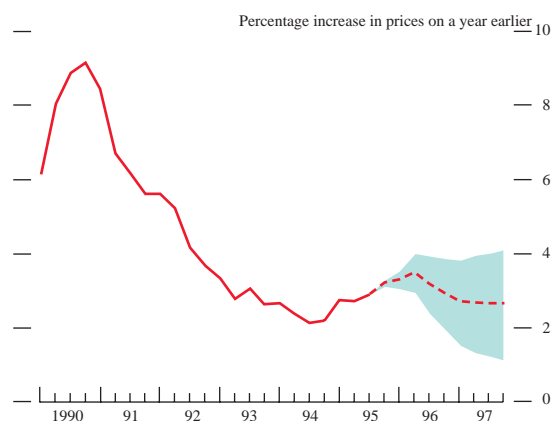
By Erik Britton, Paul Fisher and John Whitley of the Bank’s Conjunctural Assessment and Projections Division.

Since February 1996, the Bank’s inflation forecast has been published in the form of a probability distribution—presented in what is now known as ‘the fan chart’. This article discusses the motivation for the change, describes how the chart is produced and explains how it reflects the forecast process.

Introduction

The introduction of an inflation-targeting regime for UK monetary policy in 1992 has placed more emphasis on taking a forward view of inflationary pressure.⁽¹⁾ That forward view is inevitably best described by producing and presenting an explicit forecast for inflation. Between February 1993 (when the *Inflation Report* was first published) and February 1996, the Bank of England published a two years ahead forecast for the inflation rate in the form of a chart (see Chart 1) showing a path for the central projection of inflation. That chart also gave a measure of the range of uncertainty, as indicated by a blue shaded area around the central projection. The range of uncertainty was based on forecast errors from the previous ten years. The edges of this shaded area were derived by adding to (and subtracting from) the central projection the average absolute value of past forecast errors. Normally, one would have expected the outcome for inflation to lie within the blue area just over half of the time.

Chart 1
November 1995 RPIX inflation projection, showing symmetric error bands^(a)



(a) The range of the error bands is defined as the central projection plus or minus the average absolute error on RPIX inflation projections since 1985.

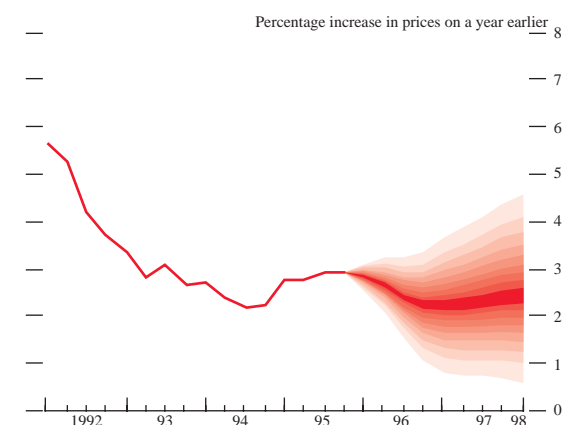
That chart was not completely satisfactory. It gave no weight to the discussion of risks to the forecast (or alternative scenarios) and encouraged the reader to concentrate on an apparently precise central projection,

ignoring the very wide degree of uncertainty surrounding it. Hence, small changes in the projection were given too much prominence relative to the risk assessment. Internal discussions at the Bank have typically spent at least as much time discussing the ways in which the central projection might be wrong as in fine-tuning that central case. In addition, the shaded area itself was often misread as indicating upper and lower bounds for the forecast, rather than the representation of probabilities that it actually showed.

Since February 1996, the Bank’s inflation forecast has been published explicitly in the form of a probability distribution—presented in what is now known as ‘the fan chart’. The aim of the fan chart has been to convey to the reader a more accurate representation of the Bank’s subjective assessment of medium-term inflationary pressures, without suggesting a degree of precision that would be spurious.

Chart 2 shows the first fan chart, published in February 1996. The rest of this article explains the derivation of the chart, how it should be read, and the forecast process that underlies it. The article also discusses the new chart for GDP growth introduced in the November 1997 *Report*.

Chart 2
RPIX inflation projection in February 1996



Sources: CSO and Bank of England.

(1) See the Governor’s 1996 speech at Loughborough University, reprinted in the *Quarterly Bulletin*, February 1997, pages 98–103.

Judgment has always been key to the forecast process in the Bank. But whose judgment and whose forecast? A distinctive feature of the *Report* process prior to May 1997 was the involvement of the Governors and Directors of the Bank in agreeing key assumptions and risks, on the basis of advice from Bank staff.⁽¹⁾ With the advent of the Monetary Policy Committee (MPC), the *Report* and the forecast represent the views of the MPC members, again aided by advice from Bank staff. Since the MPC has adopted and adapted the forecast process, this article describes that process as the MPC's, without prejudice to further changes that the MPC might wish to make in future.

The MPC builds up its assessment of risks by discussing in detail the major economic issues affecting the forecast. The multiple models approach to forecasting⁽²⁾ allows the Bank to develop and use the appropriate models for each issue. This eclectic approach means that a wide variety of views can be explored and no school of thought is automatically excluded.

The forecast process

The February 1998 *Report* is the third complete forecast round undertaken by the MPC. The process will inevitably develop further over time. This brief snapshot describes how the forecasts have been prepared to date.

There is a series of meetings between the MPC and the Bank staff. At the first meeting, roughly a month before *Report* publication, the key assumptions, the main issues and the starting-point for the risk assessment are discussed. At this stage, no forecast is presented. Following this meeting, the forecast team map the decisions of the MPC onto a central projection and risk distribution. A second meeting with the MPC considers this draft forecast. The quantification of the mapping from each assumption and risk assessment is reviewed, new data are incorporated and changes are requested. A third meeting gives the MPC an opportunity to fine-tune the revised forecast distribution and bring it up to date. The final forecast, published in the *Report*, includes adjustment in response to the advent of market-related data in the period up to the relevant monthly MPC meeting, and reflects any change in interest rates made by the Committee in that meeting.

It is important to note that the Bank's published forecasts have assumed unchanged UK short-term interest rates during the forecast period, and that the fan chart does not reflect any uncertainty about UK interest rates. The MPC's forecast distribution is not easily comparable with forecasts that allow interest rates to vary, such as those surveyed in the *Report*.

The forecast distribution

The fan chart portrays a probability distribution that approximates to the MPC's subjective assessment of

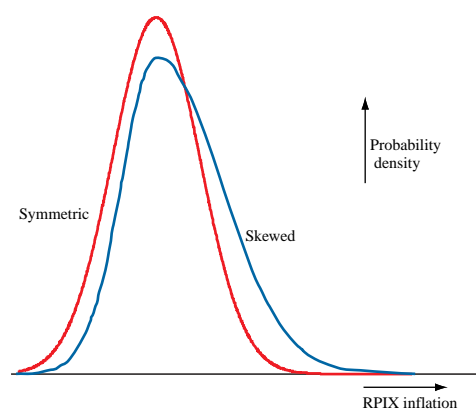
inflationary pressures evolving through time, based on a central view and the risks surrounding it. Whatever the mix of judgment and statistics used in this assessment, the process needs to be as rigorous as possible: the MPC needs to be able to explain exactly why the chart looks as it does and why it changes between *Reports*. This is vitally important both for the consistency of policy-making and for the presentation of the analysis.

For any particular forecast, one can think of the projection as being based on a model that maps choices about economic assumptions onto an inflation forecast. To generate the probability distribution, one would ideally like to evaluate all the possible alternative assumptions numerically using this model. In practice, this process is approximated by assuming a known functional form for the distribution and evaluating a limited number of alternative assumptions. These alternatives are sufficient to calibrate the key parameters of the distribution.

The choice of distribution

One might perceive the possible outcomes for inflation as being roughly symmetrically dispersed around a central, most probable value, with the values closer to the centre being more likely than those further away. That would suggest that the forecast distribution should be based on the normal 'bell-shaped' distribution widely used in statistical analysis. But the assessment of likely alternative outcomes sometimes suggests that forecast error is more likely to be in one direction than the other. This has led to the choice of a particular form of statistical distribution (a 'two-piece' normal) which has a degree of asymmetry in the form of a variable skew. The chosen form of distribution is given in full in the Appendix. A graphical representation of the distribution for a single point in time is shown in Chart 3 for both a symmetric (when it is equivalent to a normal distribution) and a skewed case. The distribution allows for the possibility of negative inflation rates.

Chart 3
RPIX inflation symmetric and skewed probability distributions



(1) See the Governor's 1996 speech to Loughborough University, reprinted in the February 1997 *Quarterly Bulletin*, especially page 101, and the article 'Economic models and policy-making' by John Whitley in the May 1997 *Quarterly Bulletin*, pages 163–73.

(2) Outlined in the article by John Whitley in the May 1997 *Quarterly Bulletin*, *op cit*.

Choosing a particular form for the distribution does not rule out the possibility of changing that choice between *Reports*. For example, suppose that two quite distinct scenarios, which are nearly equally likely, are considered. A bi-modal distribution could then be chosen so as to show two ‘humps’ instead of the usual one.

To derive the MPC’s forecast distribution, three parameters need to be evaluated. First, a measure of the central tendency for inflation—usually expressed as a particular projected path. Second, a view on the degree of uncertainty (the variance). Third, a view on the balance of the risks, to get a measure of the skew. We discuss these three aspects of the forecast in turn.

The central view

No single projection of inflation at a future date has much chance of matching the subsequent outcome. Policy discussions need to take account of the full range of possibilities. The Bank’s preference has been to start with a set of key assumptions consistent with the most likely view of developments in the economy. The central projection of inflation is then interpreted as being the ‘mode’ of the statistical distribution—it is the single most likely outcome based on current knowledge and judgment, even if the actual chance of it matching the eventual outcome is small. This central projection is based on a consistent set of assumptions about economic behaviour that provides the foundations for subsequent assessment of how the economy is evolving relative to the forecast.

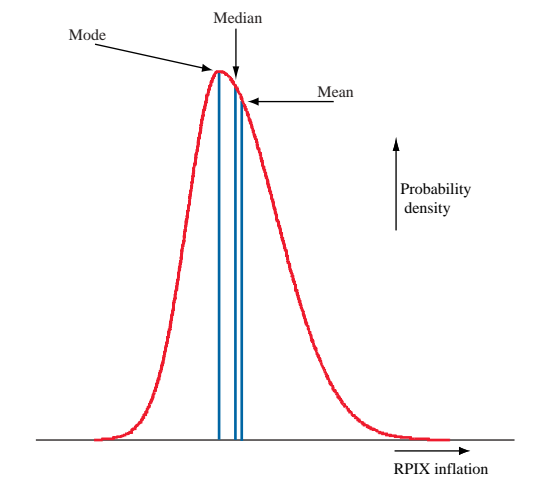
If the risks are symmetrically distributed around the central view, this will also provide a view of the average outcome (or mean forecast). But when the risks are unbalanced, the average of all the alternatives is unlikely to be the same as the single most likely case, and the mean forecast will differ from the mode. We return to this with an example when discussing the balance of risks (the skew of the distribution).

Neither the most likely nor the average view will necessarily split the forecast distribution in half. The point in the distribution that has 50% probability on either side is known as the median. The relationship between the mode, median and mean is important in interpreting the fan chart. Chart 4 shows how the three measures relate to each other when the distribution is skewed. When the balance of risks is on the upside, the mean inflation forecast will be higher than the median, which will be higher than the mode.

The degree of uncertainty

The uncertainty in the subjective assessment of inflation relates to how likely it is that the future events will differ from the central view. It is therefore a forward-looking view of the risks to the forecast, not a mechanical extrapolation of past uncertainty. Nevertheless, the initial calibration of uncertainty is based on the experience of

Chart 4
Central tendencies of a skewed distribution



forecast errors from the previous ten years (as in the shaded area shown in Chart 1). So the fan chart approach then requires the MPC to form a view as to whether or not uncertainty looking forward is greater or less than in the past.

The degree of uncertainty (the degree of dispersion in the distribution) can be measured by a variety of statistics such as variance, mean absolute error or inter-quartile range. The Bank uses a variance measure.

It is always tempting when forecasting to assume that the current degree of uncertainty is greater than usual. Knowledge of the current state of the economy is strictly limited, unusual shocks are always occurring, and statistical models based on the past are always likely to break down in the face of new developments. In practice, it has been shown that, though forecasting is indeed notoriously uncertain in an absolute sense, the track record of forecasts is rather better than one would suppose from simply evaluating the uncertainty inherent in statistical models.⁽¹⁾ And since 1994, inflation has been much less volatile than previously, helping to improve forecast accuracy. When evaluating the risks, the MPC may decide to vary the degree of forecast uncertainty to reflect the prospective or recent variability of economic developments.

The balance of risks

In deciding upon central assumptions and risks across key components of the forecast, it may become clear that the risks are unbalanced. A good example of this is the effect of ‘windfall’ gains to consumers from the conversion of several building societies to banks in 1997.⁽²⁾ Uncertainty about the pace at which the windfalls would be spent represented a risk to the forecast of consumer spending. The Bank’s theoretical analysis suggested that only a small proportion of these gains would be spent in the first year, and correspondingly took this as a central view. In the Bank’s judgment, the risks were much greater than actual

(1) Not just the Bank’s own—see Wallis, K F and Whitley, J D, 1991, ‘Sources of Error in Forecasts and Expectations: UK Economic Models 1984–88’, *Journal of Forecasting* 10, pages 231–53.

(2) See the *Inflation Report*, February 1997, page 22.

expenditure would be in excess of the central forecast assumption than that it would be less. This was an upside risk to the forecast during most of 1997.

In order to produce the fan chart, only one number is needed to summarise the degree of skewness (the balance of risks). Just as with the central view and the degree of uncertainty, there is more than one possible choice of parameter. The Bank's analysis focuses on the difference between the mean and the mode of the forecast distribution to be presented in the *Report*. This difference is of interest as a summary statistic of the balance of risks, and it provides a practical way of calibrating the distribution.

The Bank has concentrated on systematically building up a forecast distribution of inflation in a manner that reflects its subjective judgment, and now that of the MPC. What matters is that the MPC should be content with both the fan chart and the recorded mapping from its discussions of the issues. Because the current procedure is an approximation, evaluating a limited number of alternative scenarios, a certain amount of iteration in the discussions between the MPC and the forecasting team is needed until the fan chart is agreed.

The mapping

To evaluate the complete forecast distribution, one would ideally want to assess a potentially unlimited number of shocks that might affect the inflation forecast. But to keep the process tractable, one has to focus on the major issues of the day, while ensuring a comprehensive review of the economic situation as a whole. The first MPC meeting confirms the selection of major issues and provides the economic analysis to form the basis of the forecast. For each observed shock, the MPC forms a central view of its size and consequences, and considers how that view might be wrong. The degree of uncertainty and the balance of risks for each shock are then calibrated. In most cases this is done by examining various alternative models to assess what the consequences might be if the central view is mistaken. Eventually, a judgment has to be made about whether the risks are skewed and by how much, and whether the uncertainty about the relevant relationships is more or less than in the past.

The central projection represents a mapping of the central assumptions onto an inflation projection, using an economic model. In order to understand the issues of particular relevance in any one forecast round, the MPC considers several different ways of looking at the economy before selecting the set of relationships—or model—that represent its view for that forecast round. Hence, the econometric model used to ensure consistency of variables in the forecast is not set in stone, but changes from one forecast to the next.

The variance of inflation can be derived from the underlying variances of the basic shocks, using the mapping provided by the economic model. To make this tractable, simulations are used to identify the contribution of the relevant basic variances to the variance of the inflation forecast. For independent shocks, the inflation forecast variance could be obtained by a weighted sum of the individual variances.⁽¹⁾ But rather than add up all the variances, the past inflation forecast error variance is taken as a starting point and then adjusted upwards or downwards, based on changes to a limited number of variance assumptions.⁽²⁾ By adjusting the basic variances, the forecast variance of inflation is thus changed to match the degree of uncertainty as viewed by the MPC.

Wherever discussions suggest that there might be an unbalanced risk, a plausible alternative assumption is made for the relevant parameter or shock in the direction of the identified skew, and a model (or models) is simulated to show how the forecast for inflation at different horizons would change under that assumption, allowing for all relevant feedbacks. The MPC attaches a probability weight to that alternative, which scales up or down the effect of the alternative assumption on the inflation forecast distribution. The MPC must make a judgment about that probability weight, which they can do by reference to the underlying analysis, or by reviewing the impact on the inflation forecast. Attaching probability weights in this way approximates the task of simulating all possible forecast variants with a limited number.

Once the individual risks are evaluated, they must be aggregated to give an overall balance of risks. In practice, this means adding together the individual simulations of the alternative assumptions, each weighted by their probability. This gives an overall balance of risks, expressed in terms of an effect on the inflation forecast.

The balance of risks is interpreted as giving the difference between the mean and the mode of the forecast. Why is that? If we take a probability-weighted average across a range of alternative forecasts, then we have implicitly estimated the mean of the forecast distribution.⁽³⁾ Hence, the balance of risks estimated in this way is used to calibrate the forecast distribution.

The MPC judges not only the individual components of the forecast, but whether the final result for the inflation distribution is felt to be consistent with its analysis of the issues. If not, then the forecast team will be asked to change the nature of the assumptions or the probability of alternative assumptions, or to reassess the simulation responses of the models used. In this way, the fan chart is made consistent with the MPC's judgments—both 'bottom up' and 'top down'.

(1) The risks are identified as independent shocks that have zero covariance, though this assumption can be relaxed if necessary.

(2) Technically, the assumptions about the degree of uncertainty and its skew, as expressed by the difference between the mode and mean, may not be independent. In practice, it is possible to alter the distribution so as to preserve one while adjusting the other, if that is appropriate.

(3) This can be thought of as a partial numerical integration of the distribution.

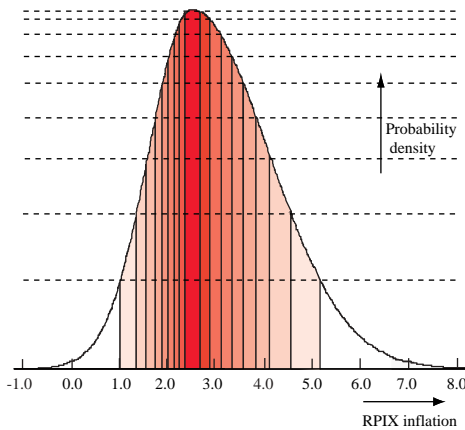
Drawing the fan chart

The distribution to be plotted is generated by an iterative procedure, given the central projection as a mode, the variance and balance of risks. The chosen distribution is adjusted until the required variance and the required difference between mean and mode are obtained, and appropriately adjusted to ensure that the probability area is equal to one, as required by a probability distribution. The skew and variances are evaluated for one year ahead and two years ahead projections and then interpolated for the quarterly fan chart. This gives the distribution to plot.

Having obtained the forecast distribution for inflation at each point in a nine quarter ahead forecast, its graphical representation remains a matter of choice. The fan chart was chosen to meet the criterion that it should give information on the whole of the forecast distribution, without claiming a spuriously high degree of precision.

The fan chart itself is best understood by looking at Charts 5 and 6. A vertical section of the chart at any time period (for example as shown by the thick black line in Chart 6) corresponds to a bird's-eye view of the underlying probability density function (pdf) for that period. This pdf is shown conventionally in Chart 5. The height of the pdf is proportional to the probability of inflation being a particular value in that time period. Hence, the central projection corresponds to the peak of the distribution, as it is associated with the mode.

Chart 5 Cross-sectional probability distribution of RPIX inflation with 10% confidence bands^(a)



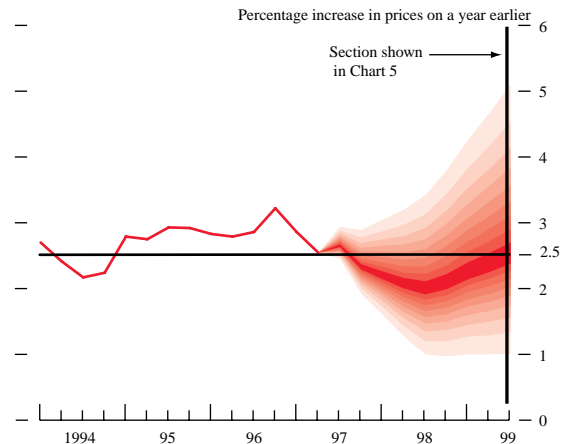
The dashed lines show how the edges of the bands are drawn—see text.

(a) Taken from the August 1997 *Inflation Report* fan chart.

The style of the chart is to make the shade of red reflect the relative probability of inflation lying in a particular band. To draw the bands, the following rule is used.⁽¹⁾ Two points of equal probability density are shown, one on either side of the mode. The two points are then moved away from the centre simultaneously, keeping the values of the probability density the same, until there is 10% of the distribution in a single central band, with these two points marking the

outside edges. That band is coloured the deepest shade of red. The two points are moved outwards again on either side of the first band (still keeping equal probability density) until another 10% of the distribution has been added, this time marking a pair of bands, one on either side of the centre. These two bands are shaded the same colour as each other, but are lighter than the central band. Pairs of bands continue to be added until 90% of the distribution is covered.

Chart 6 RPIX inflation projection in August 1997



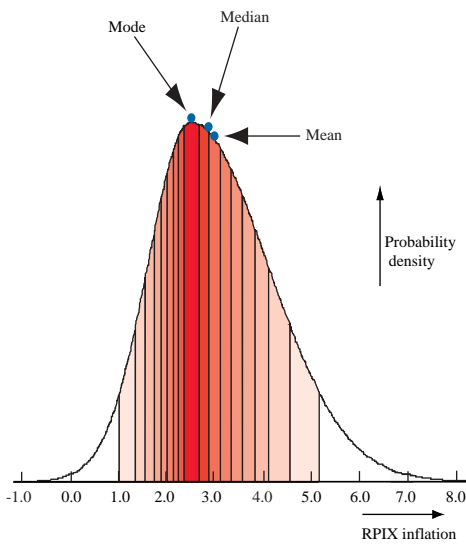
The chart shows the relative likelihood of possible outcomes. The central band, coloured deep red, includes the central projection: there is judged to be a 10% chance that inflation will be within that central band at any date. The next deepest shade, on both sides of the central band, takes the distribution out to 20%; and so on, in steps of 10 percentage points. Of course, it is impossible to assess the probabilities with any precision, but this represents the MPC's best estimate. The more uncertainty there is about the inflation outcome at any particular time horizon, the wider the bands, and the more gradually the colour fades. And if the risks are more on one side than the other, then the remaining bands will be wider on that side of the central band.

The fan chart always has the following features. There is an equal number of red bands on either side of the central band (eight). Each pair of bands covers 10% of the distribution but, if the risks are unbalanced, the same colour bands are not of equal width (representing unequal probability intervals). The distribution is truncated, so that there is an implicit ninth and final pair of bands, occupying the white space outside the 90% covered. The central projection is, by construction, always in the deepest red band since it is associated with the mode. For heavily unbalanced risks, the mean and median may not be in the deepest red band, as shown in Chart 7.

At any particular date in the forecast period, the shading gets lighter as the probability of inflation lying in bands further away from the central projection diminishes. But as uncertainty increases with the forecast horizon (compare the example distributions for years one and two in Chart 8), we could also vary the shading over time. This is done explicitly in Chart 9 with a three-dimensional chart, which includes the probability on a vertical axis. The shading of this chart emphasises that there is much less certainty about the outcome two years ahead than about the outcome one quarter ahead. In the two-dimensional chart, this increased uncertainty is shown by the widening of the bands.

(1) The rule might need to change if the distribution changed shape (eg if it were bi-modal).

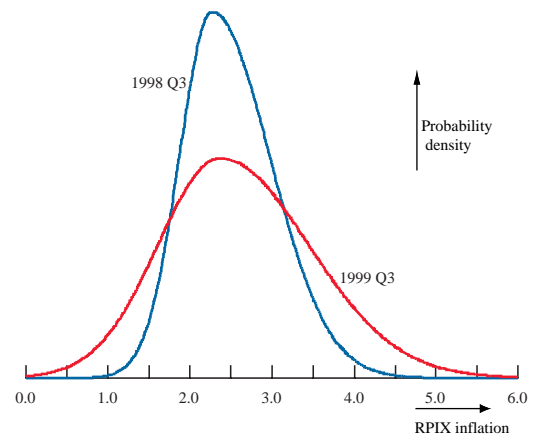
Chart 7
Central tendencies of the forecast probability distributions^(a)



(a) Taken from the August 1997 *Inflation Report* fan chart.

There are some common misconceptions about the chart. First, as explained above, the ‘fan’ does not cover 100% of the probability. Second, the central projection (the mode) is not necessarily the centre of the deepest red band—although

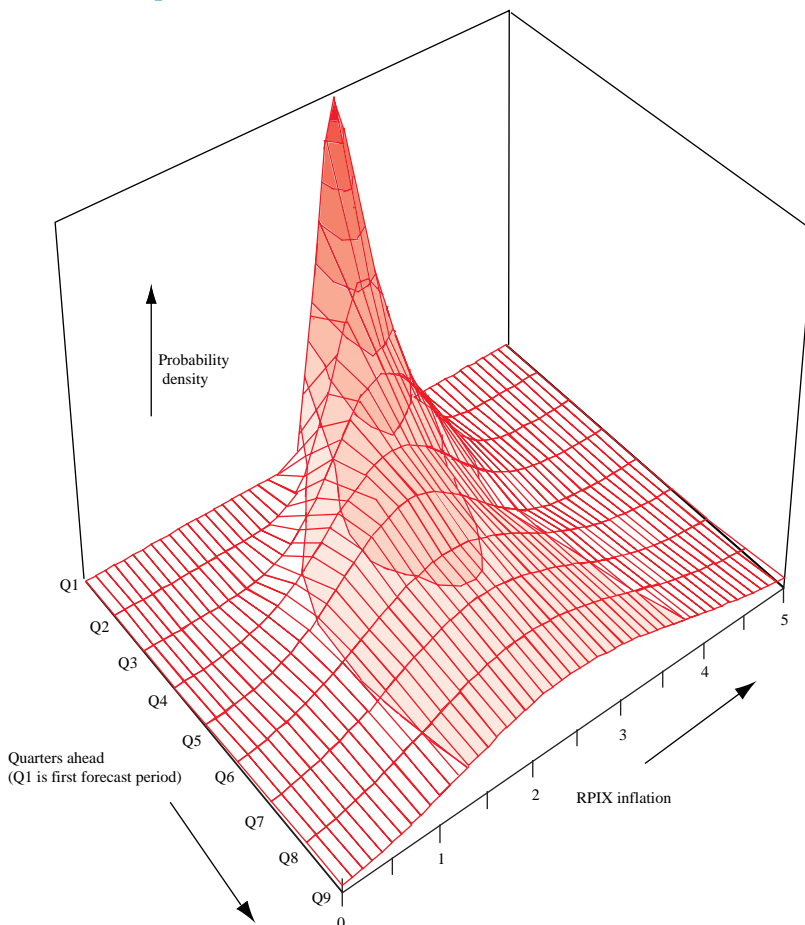
Chart 8
Cross section of the fan chart (August 1997)—one and two year ahead forecast distributions^(a)



(a) Illustrating increasing uncertainty as projection horizon becomes more distant.

it is always within it and is usually close to the centre. Third, though the fan chart could be used to represent a forecast distribution generated by purely statistical methods such as stochastic simulation of a model,⁽¹⁾ the Bank’s approach is to represent a subjective distribution for its inflation projection based on economic analysis and the judgment of the MPC.

Chart 9
Three-dimensional plot of RPIX inflation forecast distributions^(a)



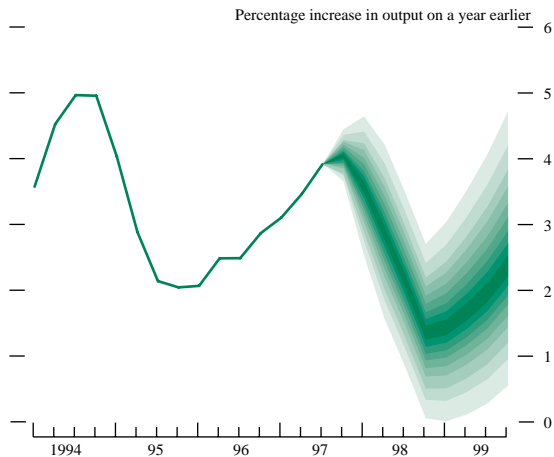
(a) Based on the August 1997 *Inflation Report* fan chart.

(1) See Blake, ‘Forecast Error Bands by Stochastic Simulation’, *National Institute Economic Review*, May 1986, pages 72–9.

The GDP fan chart

In principle, the process used above could be used to derive a fan chart for any forecast variable, as long as all the relevant risks are considered. The November 1997 *Report* presented a fan chart for GDP growth for the first time, shown in Chart 10.

Chart 10
November 1997 GDP projection



The chart shows the relative likelihood of possible outcomes. The central band, coloured deep green, includes the central projection: there is judged to be a 10% chance that output growth will be within that central band at any date. The next deepest shade, on both sides of the central band, takes the distribution out to 20%; and so on, in steps of 10 percentage points. Of course, it is impossible to assess the probabilities with any precision, but this represents the MPC's best estimate. The more uncertainty there is about the output growth at any particular time horizon, the wider the bands, and the more gradually the colour fades. And if the risks are more on one side than the other, then the remaining bands will be wider on that side of the central band.

The process of producing a GDP fan chart is exactly the same as the inflation fan chart: they are part of the same forecast. Consistency of the distributions is ensured by using the same models and judgments when mapping from alternative assumptions onto the projections for both output and inflation. In general, one cannot alter the variance or skewness of one chart without altering the other. But except for the mean and the mode, positions on one chart cannot necessarily be associated with particular points on the other, without knowing the specific risks leading to these outcomes.

A shock to demand will usually result in a positive co-movement in inflation and output. Hence, if the

assessment is that there is an unbalanced risk arising from a demand shock, both distributions will be skewed upwards. But a supply shock (eg a sudden rise in world oil prices) will usually generate a negative co-movement in output and inflation. If the assessment is that the risks are unbalanced because of supply-side factors, then the charts will be skewed in opposite directions.

What has been gained from the fan chart?

The original objectives set for the fan chart were principally to improve presentation: to focus attention on the whole of the forecast distribution, rather than on small changes to the central projection. It was hoped that this would promote discussion of the risks to the economic outlook, and thus contribute to a wider debate about economic policy. The fan chart helps to make it clear that monetary policy is about making decisions in an uncertain world, and that the MPC does not pretend to know with certainty the exact rate of inflation in two years' time.

The process used to produce the fan chart has also had a major impact on the Bank's approach to forecasting. The process forces the MPC to consider not just a single possible outcome for the economy, but a range of possibilities in areas where the central view is most likely to be wrong. In turn, this should promote better economic analysis of the underlying issues, and a necessary focus on the shocks hitting the economy.

The process also gives a rigorous accounting framework for essentially judgmental decisions. For each forecast, there is an explicit account of every discussion and decision that generated a component of the forecast for the central view, the degree of uncertainty and the balance of risks. These discussions and decisions are reflected in the text of the *Report*, and one can move from *Report* to *Report* to trace developments in the MPC's assessment of risks. In due course, one should be able to draw on all this information to judge how successful this assessment has been in identifying and calibrating the relevant shocks and their effects on UK inflation at a two-year forecasting horizon.

Appendix

The functional form for a normal distribution is as follows:

$$pdf = \frac{1}{\sqrt{2\pi\sigma^2}} e^{\left[\frac{-(x-\mu)^2}{2\sigma^2} \right]}$$

where μ is the mean of the distribution, σ^2 its variance and x is the normally distributed random variable. The fan chart distribution incorporates an extra parameter γ , to measure its skewness (where γ lies between 1 and -1). The parameter γ is incorporated into the normal distribution as follows:

$$S = \frac{1}{\sqrt{2\pi\sigma^2}} e^{\left[\frac{-1}{2\sigma^2} \left\{ (x-\mu)^2 + \gamma \left(\frac{x-\mu}{|x-\mu|} \right) (x-\mu)^2 \right\} \right]}$$

But with non-zero skewness, the integral of this function is not equal to one. So the distribution also requires a multiplicative area correction to ensure that, whatever the value of γ , the integral equals one:

$$A = \frac{2}{\left(\left(1/\sqrt{1-\gamma} \right) + \left(1/\sqrt{1+\gamma} \right) \right)}$$

So the *pdf* for the fan chart is equal to:

$$pdf = AS = \frac{2}{\left(\left(1/\sqrt{1-\gamma} \right) + \left(1/\sqrt{1+\gamma} \right) \right)} \frac{1}{\sqrt{2\pi\sigma^2}} e^{\left[\frac{-1}{2\sigma^2} \left\{ (x-\mu)^2 + \gamma \left(\frac{x-\mu}{|x-\mu|} \right) (x-\mu)^2 \right\} \right]}$$

This distribution is known as a ‘two-piece’ normal and is briefly discussed in Johnson, Kotz and Balakrishnan (1994), *Continuous Univariate Distributions*, Vol 1, page 173.