# An Assessment of Bank of England and National Institute Inflation Forecast Uncertainties

**Kenneth F. Wallis***

**January 2004**

This paper evaluates the density forecasts of inflation published by the Bank of England and the National Institute of Economic and Social Research. It extends the analysis of the Bank of England's fan charts in an earlier article by considering data up to 2003, quarter 4, and by correcting some technical details in the light of information published on the Bank's website in Summer 2003. National Institute forecasts are also considered, although there are fewer comparable observations. Both groups' central point forecasts are found to be unbiased, but their density forecasts substantially overstated forecast uncertainty.

## 1. Introduction

In February 1996 the Bank of England and the National Institute of Economic and Social Research significantly increased the amount of information they published about the uncertainty surrounding their central projections of inflation. In effect, and in different ways, they each began to publish a density forecast of inflation, that is, an estimate of the probability distribution of possible outcomes for future inflation. The Bank represented this graphically, as a set of forecast intervals covering 10, 20, 30,…, 90 per cent of the probability distribution, coloured red, of lighter shades for the outer bands. This was done for inflation forecasts up to eight quarters ahead, and since the distribution becomes increasingly dispersed and the intervals "fan out" as the forecast horizon increases, the chart became known as the "fan chart" (or, rather more informally, and noting its red colour, the "rivers of blood"). The National Institute represented the distribution as a histogram, in the form of a table reporting the probabilities of inflation falling in various ranges. These intervals, or "bins" of the histogram, have changed from time to time; those used currently are: less than

* Department of Economics, University of Warwick, Coventry CV4 7AL, UK [K.F.Wallis@warwick.ac.uk].
The comments and assistance of Michael Clements and Tim Taylor are gratefully acknowledged.

1.5 per cent, 1.5 to 2.0 per cent, 2.0 to 2.5 per cent, and so on. The forecasts refer to the fourth quarters of the current and following years, and from the beginning have included not only inflation but also real GDP growth. Fan charts for real GDP growth first appeared in the Bank's *Inflation Report* in November 1997.

These advances in the quantification and communication of forecast uncertainty were welcome, and they have contributed to a better-informed discussion of future economic prospects. A more formal justification for the publication of density forecasts as well as point forecasts is provided by the decision theory framework. The decision theory formulation begins with a loss function $L(d,y)$ that describes the consequences of taking decision $d$ today if the future state variable has the value $y$. If the future were known, then the optimal decision would be the one that makes $L$ as small as possible. But if the future outcome is uncertain, then the loss is a random variable, and a common criterion is to choose the decision that minimises the expected loss. To calculate the expected value of $L(d,y)$ for a range of values of $d$, in order to find the minimum, the complete probability distribution of $y$ is needed in general. The special case that justifies restricting attention to a point forecast is the case in which $L$ is a quadratic function of $y$. In this case the certainty equivalence theorem states that the value of $d$ that minimises expected loss $\mathrm{E}[\, L(d,y)\, ]$ is the same as the value that minimises $L\{d, \mathrm{E}(y)\}$, whatever the distribution of $y$ might be. So in this case only a point forecast, specifically the conditional expectation of the unknown future state variable, is required. In practice, however, macroeconomic forecasters have little knowledge of the identity of the users of forecasts, not to mention their loss functions, and the assumption that these are all quadratic is unrealistic. In this framework decision-makers in general require the complete distribution of $y$.

To be useful, forecasts should be reliable. This is no less true for forecasters' statements about the underlying uncertainty than it is for their point forecasts. There is a large, well-established literature on the *ex post* evaluation of *ex ante* point forecasts, and a much smaller, but growing literature on the evaluation of interval and density forecasts. As always, studies of the quality of forecasts may be of interest for their own sake, or they may be explicitly focussed on the improvement of performance in the future. Decision theory considerations suggest that forecasts of all kinds should be evaluated in a specific decision context, in terms of the gains and losses that resulted from using the forecasts to solve a sequence of decision problems. As noted above, however, macroeconomic forecasts are typically published for general use, with little knowledge of users' specific decision contexts, and evaluations are in practice based on the statistical performance of the forecasts. An exception is the evaluation of the Bank's inflation forecasts by Clements (2004), based in part on a loss function that represents the Bank's inflation targeting and reporting objectives.

This article first presents an assessment of the Bank of England's fan chart forecasts of inflation. It extends the analysis in an earlier article (Wallis, 2003) by considering more recent data and by correcting some technical details in the light of background information published on the Bank's website in Summer 2003. The National Institute's forecasts are then considered. This assessment is somewhat briefer since, as explained below, the number of comparable forecasts is relatively small.

## 2. Bank of England Monetary Policy Committee inflation forecasts

The density forecast of inflation first published in the Bank of England's quarterly *Inflation Report* in February 1996 became the responsibility of the Monetary Policy Committee (MPC) on its establishment in 1997, when the Bank was given operational independence. Our evaluation follows the practice of the analyses of the MPC's forecasting record published in the August issue of the *Inflation Report* each year since 1999, by starting from the MPC's first inflation projection published in August 1997, and by focussing on the one-year-ahead forecasts. Strictly speaking, the forecasts are conditional projections, based on the assumption that interest rates remain at the level just agreed by the MPC. They begin with a current-quarter forecast, and extend up to eight quarters ahead. Nevertheless it is argued that the one-year-ahead projections can be evaluated as unconditional forecasts, using standard forecast evaluation procedures, since inflation does not react quickly to changes in the interest rate. On the other hand the inflation outcome two years ahead is likely to be influenced by intervening policy shifts, whose impact is difficult to estimate when comparing the outcome to a forecast with a strong judgemental component, as here. The two-year projection has played an important part in establishing policy credibility, with the central projection seldom deviating far from the inflation target.

The density forecast, like that of the Sveriges Riksbank, assumes the functional form of the two-piece normal distribution (Blix and Sellin, 1998; Britton, Fisher and Whitley, 1998; Wallis, 1999). This has three parameters, which determine its location, scale and skewness; the two parameterisations in use are presented in Box A. The density forecast describes the subjective assessment of inflationary pressures by the MPC, and although the prevailing level of uncertainty is initially assessed with reference to forecast errors over the preceding ten years, the final calibration of the distribution represents the Committee's judgement. In particular, the degree of skewness shows their collective assessment of the balance of risks on the upside and downside of the forecast.

The performance of the density forecasts is assessed by considering the mean, standard deviation and general shape of the forecast distributions in relation to realised inflation. The overall goodness-of-fit of the distributions is studied by calculating, for each inflation outcome, the percentile of the forecast density in which the outcome is located, that is, the forecast probability of observing an inflation outcome no greater than that which actually occurred. For a well-fitting series of density forecasts these "probability integral transforms" of the data, denoted $z$, should be uniformly distributed between 0 and 1. That is, we should expect to see all the percentiles of the forecast densities occupied equally in a long run of correct probability forecasts.

The forecast parameters, inflation outcomes and associated $z$-values for the 22 available observations are shown in Table 1. The definition of inflation is the annual percentage change in the quarterly Retail Prices Index excluding mortgage interest payments (RPIX, Office for National Statistics code CHMK), and the most recent observation available at the time of writing is for the fourth quarter of 2003. This updates the dataset used by Wallis (2003), which ends in the second quarter of 2001. The *Inflation Report* is published in mid-quarter, and the inflation outcome refers to the corresponding quarter one year later. With respect to the asymmetry of the forecast densities, it is seen that 13 of them exhibit positive skewness, with the mean exceeding the mode, whereas five are symmetric and four are negatively skewed. The balance of risks was thought to be on the upside of the forecast more often than not, although the average of the Bank's preferred skew measure (mean minus mode), at 0.075, is small.

Evaluations of point forecasts typically focus on the conditional expectation, the mean of the forecast density, and the *Inflation Report* forecast analyses follow suit, despite the focus on the mode, the most likely outcome, in the MPC's forecast commentary and press releases. With the usual definition of forecast error as outcome minus forecast, the mean forecasts in Table 1 have an average error of zero (0.01, to be precise), thus these forecasts are unbiased. The tendency to overestimate inflation in the early part of the sample is offset by the more recent underestimation. Important contributions to this experience were the unanticipated persistence of the strength of sterling in the early years, followed more recently by surprisingly high house price inflation, which contributes to the housing depreciation component of RPIX inflation.

The standard deviation of the forecast errors is 0.42, indicating that the standard deviation of the fan chart distributions is an overestimate. This has had a downward trend over the period, perhaps in recognition of a decline in the volatility of inflation, although there has been less uncertainty over future inflation than assumed by the MPC at any time. This finding can be expected to dominate assessments of the goodness-of-fit of the complete distributions. A simple approach is to consider the interquartile range, the central 50 per cent probability band of the forecasts, and ask whether these did indeed contain 50 per cent of the outcomes. The answer is that they contained some two-thirds of the outcomes, since 15 of the 22 $z$-values lie between 0.25 and 0.75, thus the forecast interquartile ranges were too wide. More generally the class frequencies in the four classes defined by the quartiles, which are equiprobable under the hypothesis of correct distributions, are 4, 6, 9, 3. Overall there is now little evidence of asymmetry, although it is only the first three outcomes in 2003 that have delivered this finding by falling in the uppermost quartile of the fan charts.

A more complete picture of the correspondence or otherwise of the fan chart forecasts to the correct distribution is given in Figure 1. This compares the sample distribution function of the observed $z$-values with the uniform distribution function, the 45° line representing the hypothesis that the densities are correct. It is again seen that there are fewer observations than there “should” be in the outer ranges of the forecasts, with the sample distribution function being correspondingly steeper than the 45° line in the central region. The fan charts fanned out too much.

The errors in a quarterly series of one-year-ahead point forecasts are expected to exhibit autocorrelation, since the eventual error in the forecast made in the current quarter will reflect shocks to inflation that occur in the following four quarters, three of which will also influence the error in the forecast made in the next quarter. The present series is no exception, with a first-order autocorrelation coefficient of 0.51. Less is known about how the corresponding lack of independence in density forecasts might manifest itself, nevertheless the first-order autocorrelation coefficient of the *z*-values is very similar, namely 0.53. In both cases a hypothesis of forecast efficiency might include the absence of fifth-order autocorrelation, and the available evidence does appear to point in this direction, although for this purpose the series are rather short. In these circumstances the current-quarter forecasts are then of greater interest.

Practical macroeconomic forecasting usually begins with the question, where are we now, as a result of delays in the availability of data. The MPC is no exception. It normally meets on the Wednesday and Thursday following the first Monday of each month, and the quarterly forecasts are produced at the February, May, August and November meetings. The Office for National Statistics has normally released monthly consumer price indexes on the third Tuesday of the following month, hence no current-quarter information on the variable in question has been available to the forecast. Although other economic intelligence on the first month of each quarter is clearly available to the MPC at its mid-quarter meeting, the current-quarter forecasts are nevertheless treated as one-step-ahead forecasts. The year-ahead forecasts are correspondingly treated as five-step-ahead forecasts, as is implicit in the closing discussion of the previous paragraph.

Equivalent data for the current-quarter forecasts are presented in Table 2. The average error of the mean forecasts is again virtually zero (–0.01), so these forecasts also are unbiased. The standard deviation of the forecast errors is 0.16, which is close to the values used to calibrate the fan charts in the earlier part of the sample period. Until the sudden increase in the forecast standard deviation in November 2001 one would have concluded that the current-quarter forecasts were well-calibrated in this respect, and that the overestimation of uncertainty noted above occurred only as the forecasts looked further into the future. Since that change, however, the conclusion is that the fan charts have been too wide throughout their range. The increase might have been a response to the underestimation of second-quarter 2001 inflation, the forecast error in this period being an outlier in the complete series. At the time, however, this was mainly attributed to higher-than-anticipated increases in food prices, reflecting falls in domestic supply – of agricultural crops, following flooding, and of meat, following the outbreak of foot-and-mouth disease. These were not thought to reflect underlying inflationary pressures, and were unlikely to persist. Instead, the November 2001 *Inflation Report* attributed the sudden increase in dispersion to “volatile short-term movements in inflation from month to month”, which overlooks the fact that the fan charts refer to average inflation over the three months of each quarter. Many forecasters might have been tempted to anticipate a sustained increase in economic and political uncertainty linked to possible future repercussions of the September 11 terrorist attacks in the United States, but the MPC explicitly assumed no such increase, correctly so, it appears.

The complete range of the current-quarter density forecasts is better represented in the inflation data than that of the year-ahead forecasts. The frequencies with which the inflation outcome fell in the four classes defined by the quartiles of the forecast densities are 5, 9, 7, 5. There is no evidence of problems of skewness. It is arguable that the increase in dispersion since November 2001 and corresponding widening of the forecast interquartile range has brought two or three outcomes into the central classes which would otherwise have fallen outside. In the absence of the recent increase in forecast dispersion the class frequencies would have been close to equality. The improvement in goodness-of-fit over the year-ahead forecasts is shown in Figure 2, where the sample distribution function of the observed $z$-values is closer to the 45° line than that in Figure 1.

Finally, recalling that consideration of the current-quarter forecasts is partly motivated by the autocorrelation in the errors of the year-ahead forecasts, we report first-order autocorrelation coefficients of $-0.04$ and 0.07 in the errors of the mean forecasts and the $z$-values respectively. With no bias and no significant autocorrelation these point forecasts pass standard tests of weak efficiency, and the density forecasts appear to perform better at shorter horizons.

## 3. National Institute inflation forecasts

The *National Institute Economic Review* has included a density forecast of inflation since February 1996, like the Bank of England’s *Inflation Report*. Unlike the Bank, however, in each of its quarterly publications the National Institute presents forecasts for the fourth quarter of the current year and the fourth quarter of the following year ("this year" and "next year"), but not for the intervening quarters. Thus only the fourth-quarter publication presents a current-quarter and a one-year-ahead forecast, and the number of forecasts comparable to those analysed in the previous section is relatively small. (This has been the October issue of the *Review* throughout our sample period, the quarterly publication schedule having changed in mid-1996.) The intervening point forecasts, at different horizons, could be evaluated using techniques for the analysis of "fixed-event" forecasts developed and applied to earlier National Institute forecasts by Clements (1997). However the extension of these procedures to density forecasts is an open research question.

The forecast density is assumed to be a normal distribution centred on the point forecast, since the hypothesis of unbiased forecasts with normally distributed errors could not be rejected in previous testing. The point forecast is based on the National Institute's domestic macroeconometric model, subject to judgmental adjustment, as is common practice. The standard deviation of the density forecast is set equal to the standard deviation of realised forecast errors at the same horizon over a previous period. The values used to calibrate the forecast histogram have been explicitly published since October 2000. For earlier forecasts the implied values of the standard deviation have been backed out from the published histogram by the present author. That is, the question, what standard deviation would deliver these probabilities, has been answered with reference to standard normal probability tables. Rounding in the published data implies that the results are sometimes imprecise, and where the result is close to a neighbouring year's value they have been set equal, on the presumption that recalibration has not taken place very often.

Data on forecasts and outcomes are presented in Table 3. The standard deviation of both forecasts shows greater fluctuation over time than that of the Bank forecasts, and is generally set at a higher level, implying greater uncertainty. The National Institute forecast is usually completed around the middle of the month of publication of the *Review*, and hence predates the corresponding Bank forecast by about three weeks. This small increase in forecast lead time is not enough to explain the increased uncertainty. Rather, this appears to be the result of using forecast errors further back in time, into episodes of higher inflation, to calculate the standard deviation. For most of the forecasts in Table 3 the historical track record starts in 1982, while the substantial reduction in standard deviation for the 2002 forecasts is a consequence of bringing the starting date forward to 1994.

The point forecasts have tended to underestimate inflation, with positive but small average errors of 0.08 this year and 0.09 next year. These are statistically insignificant, and the unbiasedness hypothesis continues not to be rejected. The standard deviation of the forecast errors is 0.26 this year and 0.37 next year, compared to the MPC's figures in the previous section of 0.16 and 0.42 respectively. In these respects the performance of the two groups is broadly similar. In respect of the National Institute's forecast uncertainty, however, it is seen that this has been considerably overestimated. The recent reduction in forecast standard deviation is a step in the right direction, and has brought the current-quarter forecast close to the observed value of 0.26, although the year-ahead forecast standard deviation is still well above the observed value.

The overall goodness of fit of the density forecasts reflects these findings, despite the small sample size. With a symmetric density forecast a positive error in the point forecast gives a $z$-value above 0.5, thus the greater tendency towards underestimation in the current-quarter point forecasts gives five out of eight $z$-values above 0.5. In the year-ahead forecasts the general, substantial overestimation of uncertainty results in a clustering of $z$-values around 0.5. In neither case is the full range of the forecast distributions occupied. Forecast uncertainty was overestimated, and the forecast densities were much too dispersed.

## 4. Conclusion

The overall conclusion from this assessment of the two groups' density forecasts is that both substantially overstated forecast uncertainty. In both cases there have been recent reductions, but in neither case have these gone far enough. In basing their current estimates of uncertainty on the historical track record both groups go back too far into the past, into monetary policy regimes substantially different from that which has existed through the period under consideration. The Bank noted in the August 2000 *Inflation Report* that outturns that tend to lie close to the centre of the forecast distribution suggest that recent forecast errors have been smaller than in the past, but did not reduce its measure of uncertainty. Whether exaggerated views of uncertainty led to undue caution in the setting of interest rates is an open question.

The use of past forecast performance as an indicator of likely future performance is like any other forecasting problem, now addressed to measures of the dispersion of forecasts, but subject to the same difficulties of forecast failure as point forecasts. Projecting forward from past performance assumes a stable underlying environment, and difficulties arise when this structure changes. If the effects of change can be anticipated, subjective adjustments can be made, just as is the case with point forecasts. But as Pagan (2003, p.84) notes, ultimately this is an issue of whether the decline in the volatility of inflation is a temporary or a permanent phenomenon. Although he can find no convincing explanation of the decline in the volatility of inflation during the 1990s, we note that Milton Friedman long ago claimed an association between the level and variability of inflation. In economics "temporary" and "permanent" are relative terms, and the decline in uncertainty seems sufficiently permanent to be taken on board in forecasting up to two years ahead.

**Box A: Parameterisation of the two-piece normal distribution**

A random variable $X$ has a two-piece normal distribution with parameters $m,\ s_1$ and $s_2$ if it has probability density function (pdf)

$$f(x) = \begin{cases} A\exp\left[-(x-m)^2 / 2s_1^2\right] & x \le m \\ A\exp\left[-(x-m)^2 / 2s_2^2\right] & x \ge m \end{cases} \qquad \text{(A.1)}$$

where $A = \left(\sqrt{2p}\,(s_1 + s_2)/2\right)^{-1}$ (John, 1982; Johnson, Kotz and Balakrishnan, 1994; Wallis, 1999). The distribution is formed by taking the left half of a normal distribution with parameters $(m,\ s_1)$ and the right half of a normal distribution with parameters $(m,\ s_2)$, and scaling them to give the common value $f(m) = A$ at the mode, as in (A.1). The scaling factor applied to the left half of the $N(m, s_1)$ pdf is $2s_1/(s_1 + s_2)$ while that applied to the right half of $N(m, s_2)$ is $2s_2/(s_1 + s_2)$. If $s_2 > s_1$ this reduces the probability mass to the left of the mode to below one-half and correspondingly increases the probability mass above the mode, hence in this case the two-piece normal distribution is positively skewed with mean>median>mode. Equivalently, when $s_1 > s_2$ the distribution is negatively skewed. The mean and variance of the distribution are

$$E(X) = m + \sqrt{\frac{2}{p}}\,(s_2 - s_1) \qquad \text{(A.2)}$$

$$\text{var}(X) = \left(1 - \frac{2}{p}\right)(s_2 - s_1)^2 + s_1 s_2\ . \qquad \text{(A.3)}$$

The two-piece normal distribution is a convenient representation of departures from the symmetry of the normal distribution, since probabilities can be readily calculated by referring to standard normal tables and scaling by the above factors; however, the asymmetric distribution has no convenient multivariate generalisation.

The alternative parameterisation of the distribution used by the Bank of England (Britton, Fisher and Whitley, 1998) is

$$f(x) = \begin{cases} A\exp\left[-(x-m)^2(1+g)/2s^2\right] & x \le m \\ A\exp\left[-(x-m)^2(1-g)/2s^2\right] & x \ge m \end{cases} \qquad \text{(A.4)}$$

where the sign of $g$ has been corrected so that positive values represent positive skewness, as in Britton *et al*.'s discussion, and $A$ is reexpressed in terms of $g$ and $s$ through the equivalences

$$(1+g)s_1^2 = s^2 \text{ and } (1-g)s_2^2 = s^2, \qquad \text{(A.5)}$$

with $-1 < g < 1$. For each forecast the Bank's spreadsheets report the mean, median and mode of the distribution, together with measures of skew and uncertainty. The skew measure is the difference between the mean and the mode. The uncertainty measure corresponds to the parameter $s$ in equation (A.4). This has been made clear in notes to the spreadsheets that first appeared in Summer 2003, whereas researchers including the present author had previously interpreted Britton *et al*. as implying that the reported uncertainty measure was the standard deviation of the distribution as given by the square root of (A.3); this is equal to $s$ only in the symmetric case. An expression for $g$ can be obtained from (A.2) and (A.5); defining a standardised measure of skewness as $s = \left[E(x) - m\right]/s$ , this is

$$g^2 = 1 - 4\left(\frac{\sqrt{1+ps^2}-1}{ps^2}\right)^2$$

where $g$ takes the sign – positive, negative, or zero – of $s$. From values of $g$ and $s$ those of $s_1$ and $s_2$ are given by (A.5), then probability calculations proceed as in the previous paragraph. This account corrects the description of the numerical procedure in the last few lines of Box A of Wallis (1999).

The new notes to the Bank's spreadsheets not only remove the ambiguity in Britton *et al*.'s discussion of the uncertainty measure, but also call attention to revisions in the parameter values of the November 1997 and February 1998 fan charts that reflect a change in their method of construction. As a result three of the mean values in column (2) of Tables 1 and 2 differ from those given by Wallis (2003) for these two forecasts, which are also subject to changes in the uncertainty measure. More generally, the recalculation of the standard deviation following the reinterpretation of the published uncertainty measure implies that the entries in the early part of column (3) of Tables 1 and 2 differ from those given by Wallis (2003) whenever the distribution is asymmetric, although these changes are very small. There are consequential differences, again very small, between the corresponding entries of column (5) of the present tables and those given in the earlier article.

## References


Blix, M. and Sellin, P. (1998). Uncertainty bands for inflation forecasts. Working Paper No.65, Sveriges Riksbank, Stockholm.

Britton, E., Fisher, P.G. and Whitley, J.D. (1998). The *Inflation Report* projections: understanding the fan chart. *Bank of England Quarterly Bulletin*, 38, 30-37.

Clements, M.P. (1997). Evaluating the rationality of fixed-event forecasts. *Journal of Forecasting*, 16, 225-239.

Clements, M.P. (2004). Evaluating the Bank of England density forecasts of inflation. *Economic Journal*, forthcoming.

John, S. (1982). The three-parameter two-piece normal family of distributions and its fitting. *Communications in Statistics – Theory and Methods*, 11, 879-885.

Johnson, N.L., Kotz, S. and Balakrishnan, N. (1994). *Continuous Univariate Distributions*, 2nd ed., vol 1. New York: Wiley.

Pagan, A.R. (2003). Report on modelling and forecasting at the Bank of England. *Bank of England Quarterly Bulletin*, 43, 60-88.

Wallis, K.F. (1999). Asymmetric density forecasts of inflation and the Bank of England's fan chart. *National Institute Economic Review*, No.167, 106-112.

Wallis, K.F. (2003). Chi-squared tests of interval and density forecasts, and the Bank of England's fan charts. *International Journal of Forecasting*, 19, 165-175.

**Table 1. Bank of England Monetary Policy Committee Inflation Forecasts**
One-year-ahead forecasts and outcomes ($n$=22)

| *Inflation Report* | (1) Mode | (2) Mean | (3) Std. Dev. | (4) Outcome | (5) $z$ |
|---|---|---|---|---|---|
| Aug 97 | 1.99 | 2.20 | 0.79 | 2.55 | 0.68 |
| Nov 97 | 2.19 | 2.72 | 0.75 | 2.53 | 0.45 |
| Feb 98 | 2.44 | 2.53 | 0.50 | 2.53 | 0.51 |
| May 98 | 2.37 | 2.15 | 0.66 | 2.30 | 0.56 |
| Aug 98 | 2.86 | 3.00 | 0.62 | 2.17 | 0.08 |
| Nov 98 | 2.59 | 2.72 | 0.64 | 2.16 | 0.19 |
| Feb 99 | 2.52 | 2.58 | 0.62 | 2.09 | 0.22 |
| May 99 | 2.23 | 2.34 | 0.60 | 2.07 | 0.34 |
| Aug 99 | 1.88 | 2.03 | 0.59 | 2.13 | 0.58 |
| Nov 99 | 1.84 | 1.79 | 0.55 | 2.11 | 0.72 |
| Feb 00 | 2.32 | 2.42 | 0.57 | 1.87 | 0.17 |
| May 00 | 2.47 | 2.52 | 0.55 | 2.26 | 0.32 |
| Aug 00 | 2.48 | 2.48 | 0.54 | 2.38 | 0.43 |
| Nov 00 | 2.19 | 2.24 | 0.56 | 1.95 | 0.31 |
| Feb 01 | 2.09 | 2.04 | 0.55 | 2.37 | 0.72 |
| May 01 | 1.94 | 1.89 | 0.55 | 1.86 | 0.47 |
| Aug 01 | 1.96 | 1.96 | 0.55 | 2.00 | 0.52 |
| Nov 01 | 2.06 | 2.26 | 0.60 | 2.61 | 0.73 |
| Feb 02 | 2.13 | 2.33 | 0.59 | 2.89 | 0.83 |
| May 02 | 2.05 | 2.05 | 0.52 | 2.90 | 0.95 |
| Aug 02 | 2.31 | 2.31 | 0.51 | 2.87 | 0.87 |
| Nov 02 | 2.41 | 2.41 | 0.48 | 2.58 | 0.64 |

Notes: see Table 2

**Table 2. Bank of England Monetary Policy Committee Inflation Forecasts**

Current-quarter (one-step-ahead) forecasts and outcomes ($n$=26)

| *Inflation Report* | (1) Mode | (2) Mean | (3) Std. Dev. | (4) Outcome | (5) $z$ |
|---|---|---|---|---|---|
| Aug 97 | 2.65 | 2.69 | 0.16 | 2.81 | 0.79 |
| Nov 97 | 2.60 | 2.71 | 0.15 | 2.80 | 0.74 |
| Feb 98 | 2.60 | 2.64 | 0.20 | 2.59 | 0.43 |
| May 98 | 2.83 | 2.74 | 0.26 | 2.94 | 0.77 |
| Aug 98 | 2.51 | 2.56 | 0.25 | 2.55 | 0.49 |
| Nov 98 | 2.54 | 2.58 | 0.19 | 2.53 | 0.41 |
| Feb 99 | 2.49 | 2.51 | 0.19 | 2.53 | 0.54 |
| May 99 | 2.48 | 2.51 | 0.18 | 2.30 | 0.12 |
| Aug 99 | 2.31 | 2.35 | 0.17 | 2.17 | 0.15 |
| Nov 99 | 2.20 | 2.19 | 0.17 | 2.16 | 0.43 |
| Feb 00 | 1.93 | 1.96 | 0.17 | 2.09 | 0.79 |
| May 00 | 1.88 | 1.89 | 0.17 | 2.07 | 0.85 |
| Aug 00 | 2.38 | 2.38 | 0.16 | 2.13 | 0.06 |
| Nov 00 | 2.36 | 2.37 | 0.17 | 2.11 | 0.06 |
| Feb 01 | 1.94 | 1.92 | 0.17 | 1.87 | 0.37 |
| May 01 | 1.90 | 1.88 | 0.17 | 2.26 | 0.99 |
| Aug 01 | 2.31 | 2.31 | 0.17 | 2.38 | 0.67 |
| Nov 01 | 2.00 | 2.10 | 0.30 | 1.95 | 0.33 |
| Feb 02 | 2.14 | 2.24 | 0.29 | 2.37 | 0.69 |
| May 02 | 2.02 | 2.02 | 0.26 | 1.86 | 0.27 |
| Aug 02 | 1.84 | 1.84 | 0.25 | 2.00 | 0.73 |
| Nov 02 | 2.64 | 2.64 | 0.24 | 2.61 | 0.46 |
| Feb 03 | 2.77 | 2.77 | 0.28 | 2.89 | 0.67 |
| May 03 | 3.09 | 3.14 | 0.26 | 2.90 | 0.18 |
| Aug 03 | 2.85 | 2.85 | 0.25 | 2.87 | 0.53 |
| Nov 03 | 2.72 | 2.72 | 0.24 | 2.58 | 0.29 |

Notes on sources

(1),(2): Bank of England spreadsheets, see
http://www.bankofengland.co.uk/inflationreport/irprobab.htm

(3),(5): calculated using code written in the Gauss Programming Language by Michael Clements. The standard deviation is the square root of expression (A.3) in Box A; $z$ is the probability integral transform of the inflation outcome in the forecast distribution.

(4): annual percentage growth in quarterly RPIX, ONS code CHMK

**Table 3.** **National Institute Inflation Forecasts**
Current-quarter and one-year-ahead forecasts and outcomes

| *NIER* | Point forecast | Standard deviation | Outcome | *z* |
|---|---|---|---|---|
| | **This year** | | | |
| Oct 1996 | 2.9 | 0.3 | 3.22 | 0.86 |
| Oct 1997 | 2.8 | 0.3 | 2.80 | 0.50 |
| Oct 1998 | 2.2 | 0.71 | 2.53 | 0.68 |
| Oct 1999 | 2.1 | 0.71 | 2.16 | 0.53 |
| Oct 2000 | 1.9 | 0.71 | 2.11 | 0.62 |
| Oct 2001 | 2.3 | 0.71 | 1.95 | 0.31 |
| Oct 2002 | 2.3 | 0.27 | 2.61 | 0.87 |
| Oct 2003 | 2.8 | 0.33 | 2.58 | 0.25 |
| | **Next year** | | | |
| Oct 1996 | 2.7 | 2.0 | 2.80 | 0.52 |
| Oct 1997 | 2.4 | 2.0 | 2.53 | 0.53 |
| Oct 1998 | 2.5 | 1.8 | 2.16 | 0.43 |
| Oct 1999 | 2.3 | 1.67 | 2.11 | 0.45 |
| Oct 2000 | 2.0 | 1.67 | 1.95 | 0.49 |
| Oct 2001 | 1.8 | 1.67 | 2.61 | 0.69 |
| Oct 2002 | 2.4 | 0.71 | 2.58 | 0.60 |

**Figure 1. MPC year-ahead forecasts**
Cumulative distribution functions of sample *z*-values (*n*=22) and uniform distribution

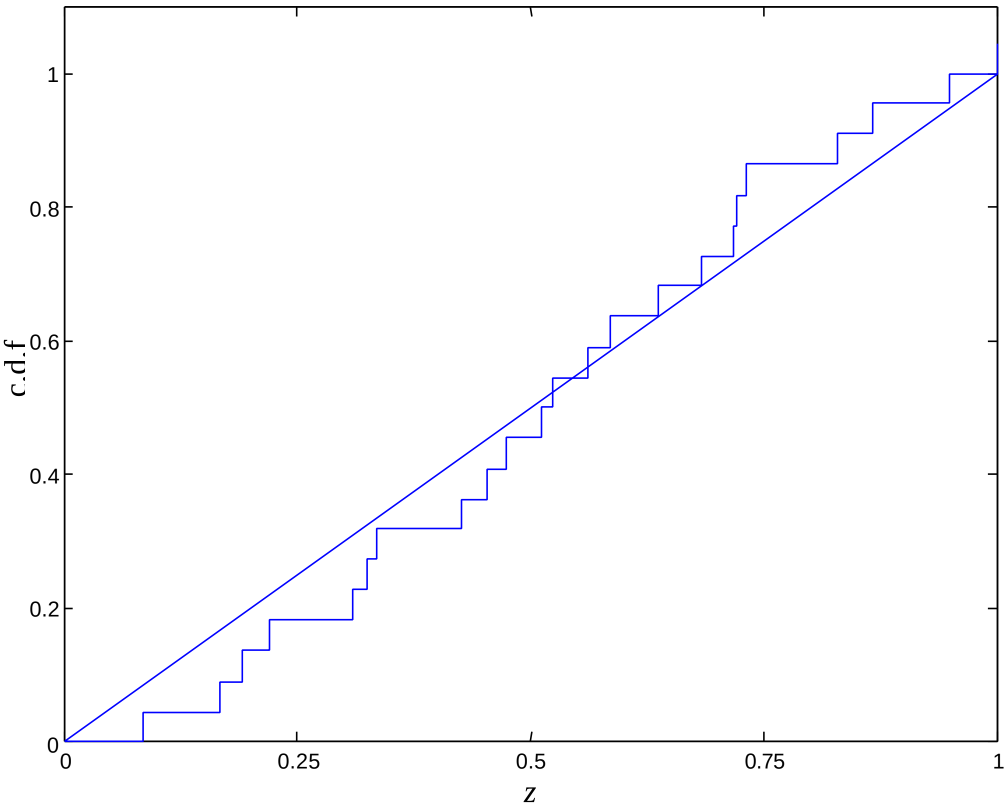

**Figure 2.** **MPC current-quarter forecasts**
Cumulative distribution functions of sample *z*-values (*n*=26) and uniform distribution