



Centre for Risk & Insurance Studies

enhancing the understanding of risk and insurance



Nottingham University
Business School

Backtesting the RPIX Inflation Fan Charts

Kevin Dowd

CRIS Discussion Paper Series – 2007.V

BACKTESTING THE RPIX INFLATION FAN CHARTS

Kevin Dowd¹

Centre for Risk and Insurance Studies, Nottingham University Business School,
Jubilee Campus, Nottingham NG8 1BB, United Kingdom. Email:
Kevin.Dowd@nottingham.ac.uk.

Abstract

This paper provides a systematic evaluation of the forecasting performance of the Bank of England's RPIX inflation fan charts. The tests are carried out on forecast data that have gone through probability integral and then Berkowitz transformations to become standard normal under the null hypothesis that the Bank's forecasts are adequate. These transformations are carried out for each of 9 forecast horizons. Except for the transformed series over the first (i.e., current-period) forecast horizon, these series are not predicted to be independent. The key problem in testing is then how to deal with this dependence structure, and three alternative approaches are suggested to deal with this problem: ignoring the dependence structure, positing an AR(1) dependence structure and positing an MA($k+1$) dependence structure. Results indicate that there are serious problems with the fan chart forecasts whichever approach we take to the dependence structure, and the results for the longer horizon forecasts are especially poor. There can therefore be little doubt that the fan chart forecasts perform badly, especially over longer forecast horizons.

1. Introduction

Ever since 1992, the principal objective of UK monetary policy has been to target inflation, and a key feature of this policy involves the Bank of England taking a forward-looking view of inflationary pressure. At the same time, the Bank has also

¹ The author thanks Ken Wallis for helpful correspondence on the fan charts over the last few years. Any remaining errors are however his own.

regarded as very important its ability to credibly convey its views about prospective future developments in the economy – and, more to the point, its ability to give a credible representation of its own uncertainty about future inflation.

The vehicle chosen for this purpose was the famous inflation ‘fan chart’, a chart showing the central projection of future inflation surrounded by a series of prediction intervals at various levels of probability.² This density forecast is represented graphically in terms of prediction intervals covering 10%, 20%, ..., 90%, of the forecast density function. Each of these intervals is shaded, with the 10%-interval darkest and the shading becoming lighter as we move to broader intervals. Interval forecasts are given for horizons up to 8 quarters ahead, and typically ‘fan out’ and become more dispersed as the horizon increases.²

The first inflation fan charts were published in the February 1996 *Inflation Report*, which gave fan chart gave forecasts of RPIX inflation – the RPIX being the inflation rate which the Bank itself was targeting – and in August 1997, the Bank started to publish the values of the parameters on which these fan chart forecasts were based. This fan chart is reproduced as Figure 1. Together with information on the probability density function used, these forecasts enable independent analysts to reproduce the density forecasts in their entirety, and therefore allow them to carry out independent assessments of the fan charts’ forecasting performance. Since then, the Bank has published inflation fan chart forecasts every succeeding quarter. However, in late 2003 the Bank switched over from an RPIX target to a CPI one and began publishing CPI fan chart forecasts. The Bank then had no further use for RPIX fan charts, and the last set of RPIX forecasts was published in February 2004.

Insert Figure 1 here

² In fact, the Bank published two different types of RPIX inflation fan chart, which have since been replaced with their CPI equivalents. The first is the constant-rate model based on the assumption that short-term market interest rates will remain constant. In February 1998, a second type of fan chart was introduced, based on the assumption that short-term interest rates will follow market expectations over the horizon period. We focus here on the former model, both for expositional ease and because it has a longer history.

A number of recent papers have set out to evaluate the forecasting performance of the inflation fan charts. These include Wallis (2003, 2004), Clements (2004), Dowd (2004, 2006, 2007a) and Elder *et alia* (2005). These studies use a variety of different approaches over (typically) different sample periods, so some degree of difference in their evaluations is perhaps to be expected. None the less, it is still surprising how wide the range of assessments actually is. At one extreme, Elder *et alia* (2005, p. 326) report it is still too soon to draw strong conclusions, but go on to suggest that “inflation ... outcomes have been dispersed broadly in line with the MPC’s fan chart bands” (*loc. cit.*) and so conclude the “fan charts gave a reasonably good guide to the probabilities and risks facing the MPC” (*loc. cit.*). Wallis (2003, 2004) comes to the conclusion that the one-year ahead forecasts significantly over-estimate inflation risk, and the Dowd studies conclude that the fan charts perform very poorly except for short horizons forecasts, and the first of the Dowd studies also reports that forecast performance deteriorates sharply with the length of the forecast horizon. Previous studies also differ considerably in the methodologies and tests used, and one study (Dowd (2007a)) does not report any test results at all.³ So, overall, we can say that most studies report greater or lesser problems with the fan chart forecasts, but even this conclusion is not universally accepted.⁴

The present paper re-examines the performance of the fan charts with a view to applying a more systematic and up-to-date set of tests to all the RPIX inflation fan chart forecasts (i.e., over all forecast horizons). Our intent is to obtain a (one

³ Instead, this study attempts to back out the probability that inflation would have remained in the narrow range it has taken in the RPIX fan chart period, assuming that the fan chart forecasts are correct. It comes up with an answer that that probability is sufficiently low – under 0.006% – as to lead a reasonable person to conclude that the fan chart forecasts lack credibility (Dowd (2007a, p. 101).

⁴ If I can offer my own personal view of the evidence, I believe that the assessments in my own earlier studies are broadly right, and this belief is reinforced by the additional evidence presented in this paper. I have no reason to dispute Wallis’s findings, but he did not look at forecast horizons of over a year, and therefore (I believe) missed the most problematic fan chart forecasts. As for Elder *et alia*, their density-forecast tests were rather limited and their test results do not contradict the evidence of poor performance reported by other studies. I would therefore conclude that the pre-

might hope) fairly definitive and robust set of results. And, to anticipate our conclusions, our results confirm that there are indeed serious problems with the performance of the RPIX inflation fan chart forecasts.

This paper is laid out as follows. Section 2 explains the inflation pdf used by the Bank. Section 3 sets out the forecast-evaluation framework to be used. Section 4 carries out some preliminary data analysis, and Section 5 reports and discusses the tests results. Section 6 concludes.

2. The RPIX Inflation Density Function

The Bank's fan charts are based on an underlying assumption that inflation obeys a two-piece normal (2PN) probability density function.⁵ In the statistical literature the 2PN pdf is usually defined as:

$$(1) \quad f(x; \mu, \sigma_1, \sigma_2) = \begin{cases} C \exp\left[-\frac{1}{2\sigma_1^2}(x - \mu)^2\right], & x \leq \mu \\ C \exp\left[-\frac{1}{2\sigma_2^2}(x - \mu)^2\right], & x \geq \mu \end{cases}$$

where $C = k(\sigma_1 + \sigma_2)^{-1}$, $k = \sqrt{2/\pi}$ and μ is the mode (see, e.g., John (1982) or Johnson, Kotz and Balakrishnan (1994)). This distribution takes the lower half of a normal distribution with parameters μ and σ_1 , and the upper half of a normal with parameters μ and σ_2 , and rescales these to give the same value at the mode. The distribution is negatively skewed if $\sigma_1 > \sigma_2$ and positively skewed if $\sigma_1 < \sigma_2$.

existing literature suggests that there are problems with the RPIX fan chart forecasts even though Elder *et alia* found no major problems with them.

⁵ This explanation of the Bank's 2PN distribution is based on Wallis (2004). The hitherto definitive account, Britton *et alia* (1997), mistakenly reported that σ was the standard deviation, and their pdf formula also had an error in the sign of γ . The latter error was quickly corrected by Wallis (1999), but the former error was only corrected when the Bank (discreetly and without any explanation) revised its guidance notes on the 2PN in 2003.

However, for the RPIX fan charts, the Bank uses a 2PN based on an alternative parameterisation:

$$(2) \quad f(x; \mu, \gamma, \sigma) = \begin{cases} C \exp \left[-\frac{1}{2\sigma^2} (1 + \gamma)(x - \mu)^2 \right], & x \leq \mu \\ C \exp \left[-\frac{1}{2\sigma^2} (1 - \gamma)(x - \mu)^2 \right], & x \geq \mu \end{cases}$$

where $-1 < \gamma < 1$, and γ and σ are related to σ_1 and σ_2 through the relationships:

$$(3) \quad (1 + \gamma)\sigma_1^2 = \sigma^2 \quad \text{and} \quad (1 - \gamma)\sigma_2^2 = \sigma^2 .$$

For each horizon, the Bank publishes forecasts of the mean, median and mode (μ), a skew parameter (which is not γ , but the difference between the mean and μ) and an uncertainty parameter σ , which is not the same as the standard deviation except where the skew is zero.⁶ Once the MPC specifies the values of μ , σ and the skew parameter for each horizon, the model is complete and the density forecasts can be obtained from it.⁷

3. A Forecast Evaluation Framework

Let x_t be the realised value of inflation in quarter t . Each realised value is to be compared against the forecasted density over each forecast horizon k , where $k=0,1,\dots,8$. Both the realised inflation and the forecasted pdf will typically change from one quarter to the next, but we standardize our observations by applying the following transformation:

⁶ The parameter forecast data are downloaded from the Bank of England website at <http://www.bankofengland.co.uk/inflationreport/rpixinternet.xls>.

⁷ Thus, μ and σ are given by the Bank, but the value of γ needs to be derived from these two parameters and the skew parameter. Details of how this can be done are given by Wallis (2004).

$$(2) \quad p_{k,t} = F_{k,t}(x_t)$$

where $F_{k,t}(\cdot)$ is the probability-integral transformation (PIT) that maps the realised RPIX inflation rate in quarter t to its value in terms of its forecasted cumulative density function made k quarters earlier.

Under the null hypothesis that the model is correct, $p_{k,t}$ should be uniformly distributed over the interval $[0,1]$. Furthermore, for $k=0$ (i.e., for current-quarter forecasts) a good risk model also predicts that the $p_{0,t}$ should be independent.⁸ We should therefore subsume this independence condition into the null, so our null hypothesis for $k=0$ is now $H_0^{k=0}: p_{0,t} \sim \text{iid } U[0,1]$. However, a good risk model does *not* predict that $p_{k,t}$ should be iid for positive k (i.e., for forecasts that extend into future quarters) because consecutive observations share common factors.

It is now convenient to put the $p_{k,t}$ through a second (Berkowitz) transformation to make them standard normal under the null of model adequacy, i.e., we apply the following transformation

$$(3) \quad z_{k,t} = \Phi^{-1}(p_{k,t})$$

where $\Phi(\cdot)$ is the standard normal distribution function (see Berkowitz, 2001). This second transformation is helpful because testing for standard normality is more convenient than testing for standard uniformity, and because a normal variable is also more convenient when dealing with temporal dependence. Under the null, each $z_{k,t}$ series will be stationary and distributed as $N(0,1)$, but it will not be independent except where $k=0$.

We now face the problem of testing the $z_{k,t}$ for standard normality in a context where (except for $k=0$) they have a dependence structure.⁹ We now suggest three alternative approaches to this problem:

(a) Ignore the dependence structure

The first and simplest approach is simply to ignore the dependence structure. This allows us to apply ‘off the shelf’ tests of the predictions of standard normality, and the most obvious tests might include:

- t -tests of the prediction that the $z_{k,t}$ have means equal to 0;
- variance-ratio tests of the prediction that the $z_{k,t}$ have variances equal to 1,¹⁰ and
- Jarque-Bera tests of the prediction that the $z_{k,t}$ are normally distributed.

*(b) An AR(1) dependence structure*¹¹

Our second approach is to postulate that $z_{k,t}$ follows a parsimonious ARMA process, and the most obvious choice is a first-order autoregression (AR(1)):

$$(4) \quad z_{k,t} = \rho_k z_{k,t-1} + \varepsilon_{k,t} \quad , \quad |\rho_k| < 1$$

where the errors $\varepsilon_{k,t}$ are iid normal and the autoregression parameters ρ_k are not expected to be 0 except where $k=0$.

⁸ The proof that model adequacy implies that $p_{0,t}$ is iid U(0,1) is found in Diebold *et alia* (1998, pp. 865, 867-869).

⁹ One criticism that I would make of other authors’ studies of the inflation fan charts it is that they either ignore this issue or, at best, do not give it the importance I believe it deserves. One can therefore argue that some of the tests applied are not suitable to the inflation fan charts. I would therefore conclude that more debate is needed over which tests are suitable and which are not.

¹⁰ This test is based on the knowledge that under the null the test statistic $(n-1)s^2 / \sigma^2 = (n-1)s^2$ is distributed as a chi-squared with $n-1$ degrees of freedom

¹¹ AR(1) processes are assumed by Dowd (2004) and Elder *et alia* (2005) in their fan chart studies.

For the $k=0$ case, we can carry out the t -tests, variance-ratio tests and JB tests set out in the previous sub-section. For the cases where $k > 0$, we can now test the mean and variance predictions using the following Monte Carlo procedure:¹²

- We obtain the $p_{k,t}$ and put these through the Berkowitz transformation (3) to obtain a $z_{k,t}$ sample. We denote this by $\hat{z}_{k,t}$ and then fit the AR(1) process (i.e., (4)) to it.
- We use the fitted AR(1) process to simulate a large number m of possible $\hat{z}_{k,t}$ series, each of which has the dependence structure of the fitted AR(1) process, and let us denote the i^{th} such series as $\tilde{z}_{k,t}^i$.
- We estimate the values of the mean and variance of each $\tilde{z}_{k,t}^i$ series. This gives us a ‘sample’ of m mean values and a ‘sample’ of m variance values. The $m \times 0.025^{\text{th}}$ and $m \times 0.975^{\text{th}}$ highest means then give us the 95% confidence interval for the mean test statistic, and the $m \times 0.025^{\text{th}}$ and $m \times 0.975^{\text{th}}$ highest variances give us the 95% confidence interval for the variance test statistic, under the null hypothesis that $\tilde{z}_{k,t}^i$ is standard normal but has the dependence structure of the fitted AR(1) process.
- We estimate the values of the mean and variance of our ‘real’ sample $\hat{z}_{k,t}$; for each of these two statistics, the forecasts pass the test if the sample value lies within the estimated confidence interval, and otherwise fail.

Thus, this second approach allows us to take account of an estimated dependence structure in our testing procedure, but at the cost of making it difficult to take account of possible departures of $z_{k,t}$ from normality.

(c) An MA(k+1) dependence structure

¹² We focus on the mean and variance predictions because the AR(1) framework assumes that the $\varepsilon_{k,t}$ are normal, and this (arguably) undermines any rationale for using it to test for departures from

Our third approach is more powerful, but also somewhat speculative. This approach is based on a conjecture put forward by Dowd (2007b). If $z_{h,t}$ are the Berkowitz-transformed observations for an h -step ahead density forecasting model, then his conjecture is that these $z_{h,t}$ should obey the following MA process:

$$(5) \quad z_{h,t} = h^{-1/2} \sum_{i=1}^h z_{1,t-i+1}$$

where the $z_{1,t}$ are to be understood here as the Berkowitz-transformed observations for a 1-step ahead density forecasting model.¹³ Given that $h=k+1$,¹⁴ we now translate this into our notation as stating that the $z_{k,t}$ should obey:

$$(6) \quad z_{k,t} = \frac{1}{\sqrt{k+1}} \sum_{i=1}^{k+1} z_{0,t-i+1}.$$

This conjecture suggests that $z_{k,t}$ is a moving average of the $z_{0,t}$, and so tells us how the $k+1$ -period ahead forecasts relate to current-period ones.¹⁵

This conjecture allows us to test for standard normality taking account of the (proposed) dependence structure.¹⁶ We start by subtracting from (6) its once-lagged equivalent and then simplify to obtain:

normality. This suggests that it might be unwise to use the approach suggested here to test for skewness or excess kurtosis.

¹³ Dowd (2007, pp. 5-7) is able to prove his conjecture in the special case of a standard normal risk forecasting model, and also gives various reasons in support of it. However, he does not offer a more general proof.

¹⁴ Remember that 1-step ahead refers to forecasts made for the current quarter, 2-steps ahead refers to forecasts made for the next quarter, and so on.

¹⁵ It also says more: since the $z_{0,t}$ are distributed as iid $N(0,1)$, the conjecture also implies that $z_{k+1,t}$ has a unique dependence structure which takes the form of a moving average of $k+1$ iid $N(0,1)$ random variables. This gives us a strong prediction about the form that the dependence structure takes.

$$(7) \quad \Delta z_{k+1,t} = \frac{1}{\sqrt{k+1}} \{z_{0,t} - z_{0,t-k-1}\}$$

where $\Delta z_{k+1,t} \equiv z_{k+1,t} - z_{k+1,t-1}$ is the first-difference of $z_{k+1,t}$. If we treat $-(k+1)^{-1/2} z_{0,t-k-1}$ as if it were an unobserved iid normal noise process $u_{k,t}$, we can rearrange (7) as

$$(8) \quad \Delta z_{k+1,t} = \frac{1}{\sqrt{k+1}} z_{0,t} + u_{k,t}$$

If we then estimate the OLS regression

$$(9) \quad \Delta z_{k+1,t} = \alpha + \beta z_{0,t} + u_{k,t}$$

we can evaluate the model by testing the predictions $\alpha = 0$, $\beta = (k+1)^{-1/2}$ and $u_{k,t} \sim \text{iid } N(0, (k+1)^{-1})$.

These tests are easy to carry out. For example, we can test the α and β predictions using conventional t -tests, we can test the prediction that $u_{k,t}$ has a variance equal to $(k+1)^{-1}$ using a variance-ratio test, we can test the normality of $u_{k,t}$ using a JB test, and we can test the iid-ness of $u_{k,t}$ using any standard test of iid-ness (e.g., portmanteau tests, runs tests, LR tests, etc.).

4. Preliminary Analysis

¹⁶ As before, for the $k=0$ case, we can carry out the same t -, variance-ratio and JB tests set out previously.

Our inflation data consist of quarterly average observations of realized RPIX inflation, as measured over the previous 12 months, and these cover the period 97Q3 to 06Q1.¹⁷ Our density forecast data consist of each quarter's published values of the parameters of the density function for this inflation series, i.e., $[\mu_0, \mu_1, \dots, \mu_8]$, $[\sigma_0, \sigma_1, \dots, \sigma_8]$ and $[skew_0, skew_1, \dots, skew_8]$, where the subscript i refers to the horizon-period (in quarters) to which the relevant parameter applies: thus, μ_0 is the quarter's forecast value of μ over the period until the end of the current quarter, μ_1 is the forecast over the period until the end of the next quarter, etc. Our fan chart parameter data cover the period 97Q3-04Q1 and so span 27 quarters.

The realised RPIX inflation rate turns out to have a mean of 2.30%, which is very close to the old RPIX inflation target of 2.50%. It also has a small standard deviation (0.31%) and a fairly narrow range of [1.87%, 3.0%]. Thus, realised inflation was close to the target and stable.

Figures 1-9 show plots of the $p_{k,t}$ or PIT observations. Each Figure also shows the least-squares slope of the relevant empirical $p_{k,t}$ series, which should be close to 1 under the null. Both the Figures and the reported slopes indicate that the PITs are pretty good for low values of k , but tend to deteriorate markedly as k continues to rise. For example, the slope is 0.9 for $k = 0$, but tends to fall as k gets bigger and is only 0.49 for $k=8$.

Insert Figures 2-10 here

Table 1 shows the sample moments of the $z_{k,t}$ series. This shows that the sample means are negative, the sample variances are less than 1 and the sample skewnesses are positive; and for the most part, these statistics move further away

¹⁷ The RPIX data is National Statistics' series CDKQ, downloaded at www.statistics.gov.uk/downloads/theme_economy/RPIX.pdf.

from their predicted values as k gets larger. These sample values again suggest that the performance of the fan chart forecasts tends to deteriorate as the forecast horizon lengthens.

One is particularly struck by the low variances and by their tendency to fall further as k gets larger. A low sample variance suggests that the fan chart models are over-estimating inflation dispersion, and the fact that the variances fall as the horizon lengthens suggests that the tendency to over-estimate inflation dispersion also rises with the forecast horizon. This latter finding suggests that the fan charts seem to fan out too much.

For their part, the sample kurtoses range from 2.678 to 4.611 and are generally quite close to their predicted value of 3.

5. Test Results

(a) Results based on an assumption that $z_{t,k}$ is iid

We turn now to more formal results, and begin with those obtained ignoring the dependence structure of the $z_{t,k}$ and so implicitly assuming that the $z_{t,k}$ are iid. Accordingly, Table 2 shows the P-values of the mean, variance and JB tests: the forecasts generally pass the mean test, except for $k=7$ and (especially) $k=8$; they generally fail the variance prediction, except for very short horizons; and they generally pass the JB test, except for $k=8$.

Insert Table 2 here

These results also show a marked deterioration as k gets larger: we can say that the fan chart forecast perform well only for $k=0$ and $k=1$; as k gets bigger their performance tends to become increasingly problematic, and the results for $k=5, 6, 7$, and 8 indicate clear ‘fails’, which continue to worsen as k increases further.

The statistical significance of the variance test results and the fact that these results deteriorate as k gets larger would seem to confirm our earlier suggestion that the fan charts do indeed fan out too much.

(b) Results based on an AR(1) dependence structure

Table 3 reports results when an AR(1) process is fitted to the $z_{k,t}$ series. These results show that the AR(1) parameters are low in magnitude and of negligible significance for $k=0$ and $k=1$ and, as an aside, the former result confirms the prediction that model adequacy for the $k=0$ forecasts implies that the $z_{0,t}$ series should be independent. However, for $k>2$, we find that the AR(1) are statistically significant, and both their magnitude and their significance tend to rise as k gets larger. These results suggest that the $k>0$ $z_{k,t}$ seem to have a dependence structure, and also suggest that this dependence structure is (usually) statistically significant. The statistical significance of the dependence structure suggests in turn that we should be careful about carrying out tests that ignore the dependence structure and treat the $z_{k,t}$ as iid.

Insert Table 3 here

Table 3 reports the mean and variance test results obtained making the assumption that $z_{k,t}$ follows an AR(1) process. By and large, these results indicate a slightly better performance; this slight improvement is reflected in the forecasts now passing all the mean tests and in some of the variance test P-values being a little better than they were. However, these improvements are only minor, and the main message is similar to what we obtained before: namely, that the forecasts generally fail the variance prediction, and we only get ‘respectable’ variance results for $k=0$ and $k=1$. Thus, the medium and longer horizon forecasts are again very problematic.

(c) Results based on an MA(k+1) dependence structure

Finally, Table 4 reports results based on the conjecture that $z_{k,t}$ follows the MA process (6). The tests are of the predictions of the model embodied in equation (9), namely, that estimates of (9) would produce parameters that satisfy the predictions $\alpha = 0$ and $\beta = (k + 1)^{-0.5}$, and that the residuals are normally distributed and have a variance equal to $(k + 1)^{-1}$.

Insert Table 4 here

The results of these tests indicate that we can always accept the predictions that $\alpha = 0$ and that the residuals are normally distributed. However, the results of the β tests are always problematic and generally very strongly so: in all but one case, they are significant at well under the 1% level. For their part, the results of the variance test are usually problematic: in the 8 cases considered, we get five results that are significant at well under the 1% level, one result that is significant at the 5% level, and only two results that are not significant at all.

Since our results that are also significant at well under the 1% level for all horizons, these results suggest that the model ‘fails’ at all horizons.¹⁸ However, we should be a little cautious in relying too heavily on these results relative to the others, because they are based on an underlying conjecture that still remains speculative.

6. Conclusions

This paper has evaluated the Bank of England’s RPIX inflation fan chart forecasts using a variety of tests. The tests are carried out on forecast data that have gone through PIT and then Berkowitz transformations to become standard normal under the null hypothesis that the Bank’s forecasts are adequate. These transformations

¹⁸ However, unlike the case with our moments-based tests, the diagnostic interpretation of these results is less clear.

are carried out for each of 9 forecast horizons. Except for the transformed series over the first (i.e., current-period) forecast horizon, these series are not predicted to be independent. The key problem in testing is then how to deal with this dependence structure, and three alternative approaches were suggested to deal with this problem: ignoring the dependence structure, positing an AR(1) dependence structure and positing an MA($k+1$) dependence structure.

Our results indicated that there are serious problems with the fan chart forecasts whichever approach we take to the dependence structure, and the results for the longer horizon forecasts are especially poor. There can therefore be little doubt that the fan chart forecasts perform badly, especially over longer forecast horizons.

Perhaps the most strongly recurring theme is that the forecasts are very poor on the $z_{k,t}$ variance prediction. Given that the sample variances are much lower than predicted and get smaller as the horizon increases, this finding suggests that the fan charts tend to over-estimate the dispersion of inflation, and that this over-estimation tends to rise with the length of the horizon. This finding is interesting in itself, but is also of considerable diagnostic value, as it can be used to help the Bank's modellers modify the model and hopefully lead to improved fan chart forecasts.

One is tempted to suggest that the main problem with the fan charts is, quite simply, the fact that they fan out. And, to end on a provocative note, I would suggest that this conclusion is not only correct but almost obvious from the fan charts themselves: a typical fan chart shows a fairly stable historical inflation rate that fluctuates around its target value – reflecting the fact that the MPC has been successful in its inflation-targetting – followed by a fan chart forecast suggesting that the inflation rate can plausibly go well outside its recent range. Thus, the historical inflation series suggests a mean-reverting process whereas the fan chart forecasts suggest some kind of diffusion process. These are quite different processes and provide convincing evidence (at least to me) that the fan chart forecasts of inflation are incompatible with the actual inflation process.

References

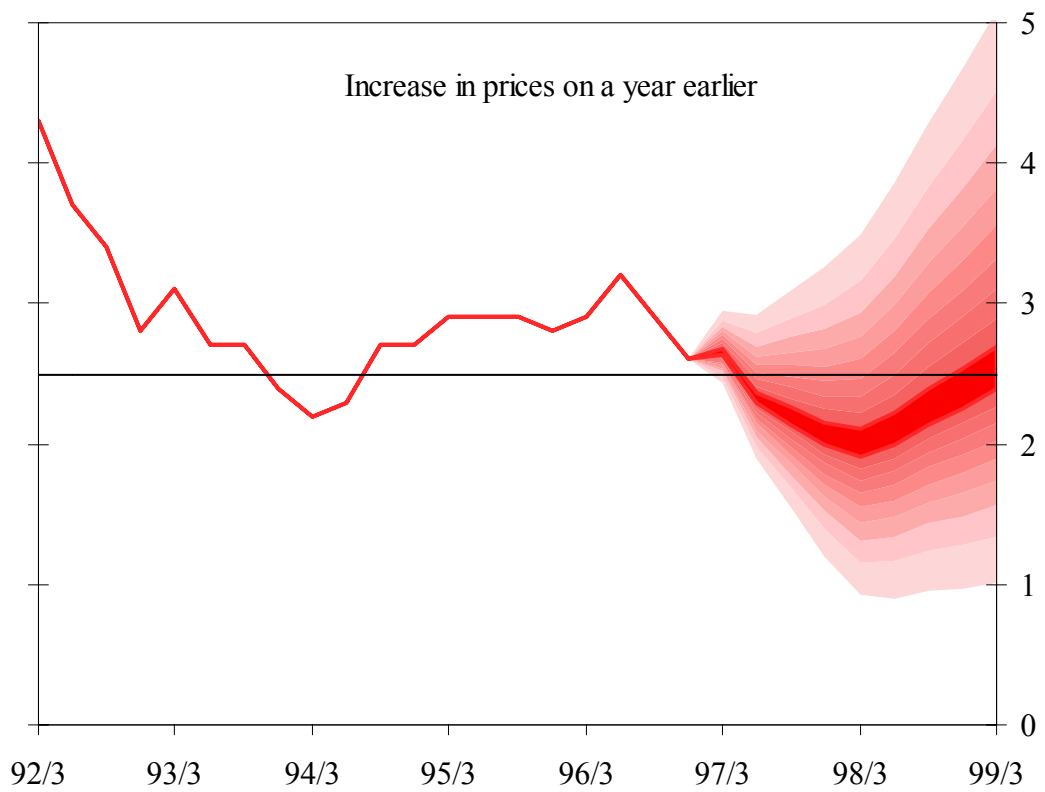
- Berkowitz, J. (2001), "Testing density forecasts, with applications to risk management," *Journal of Business and Economic Statistics*, 19: 465-474.
- Britton, E., Fisher, P., Whitley, J. (1997), "The Inflation Report projections: understanding the fan chart", *Bank of England Quarterly Bulletin*, 38: 30-37.
- Clements, M. P. (2004), "Evaluating the Bank of England density forecasts of inflation," *Economic Journal*, 114, pp. 844-866.
- Elder, R., G. Kapetanios, T. Taylor and T. Yates (2005) "Assessing the MPC's fan charts," *Bank of England Quarterly Bulletin* (Autumn): 326-348.
- Dowd, K. (2004) "The inflation 'fan charts': An evaluation." *Greek Economic Review* 23: 99-111.
- Dowd, K. (2006) "Forecasting inflation: the inflation 'fan charts'," Chapter 7 in K. Matthews and P. Booth (eds) *Issues in Monetary Policy: The Relationship Between Money and Financial Markets*, pp. 80-93. Chichester: John Wiley and Sons, 2006.
- Dowd, K. (2007) "Too good to be true? The (in)credibility of the UK inflation fan charts," *Journal of Macroeconomics*, 29: 91-102. (a)
- Dowd, K. (2007) "Temporal dependence in multi-step density forecasting models," *Journal of Risk Model Validation* 1: 3-22. (b)
- Elder, R., G. Kapetanios, T. Taylor and T. Yates (2005) "Assessing the MPC's fan charts," *Bank of England Quarterly Bulletin* (Autumn): 326-348.
- 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., S. Kotz and N. Balakrishnan (1994), *Continuous Univariate Distributions. Volume I*, New York: Wiley.
- Pagan, A. (2003), *Report on Modelling and Forecasting at the Bank of England*, London: Bank of England.

- 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.
- Wallis, K. F. (2004), “An assessment of Bank of England and National Institute inflation forecast uncertainties”, *National Institute Economic Review*, No. 189: 64-71.

FIGURES


FIGURE 1:

The August 1997 Inflation Fan Chart




Notes: The Figure shows previous realized values and predicted intervals for RPIX inflation, measured as the percentage increase in prices over the previous 12 months (measured against the y -axis), for the constant rate model. The Figure is reproduced with permission from the Bank of England.

FIGURE 2:



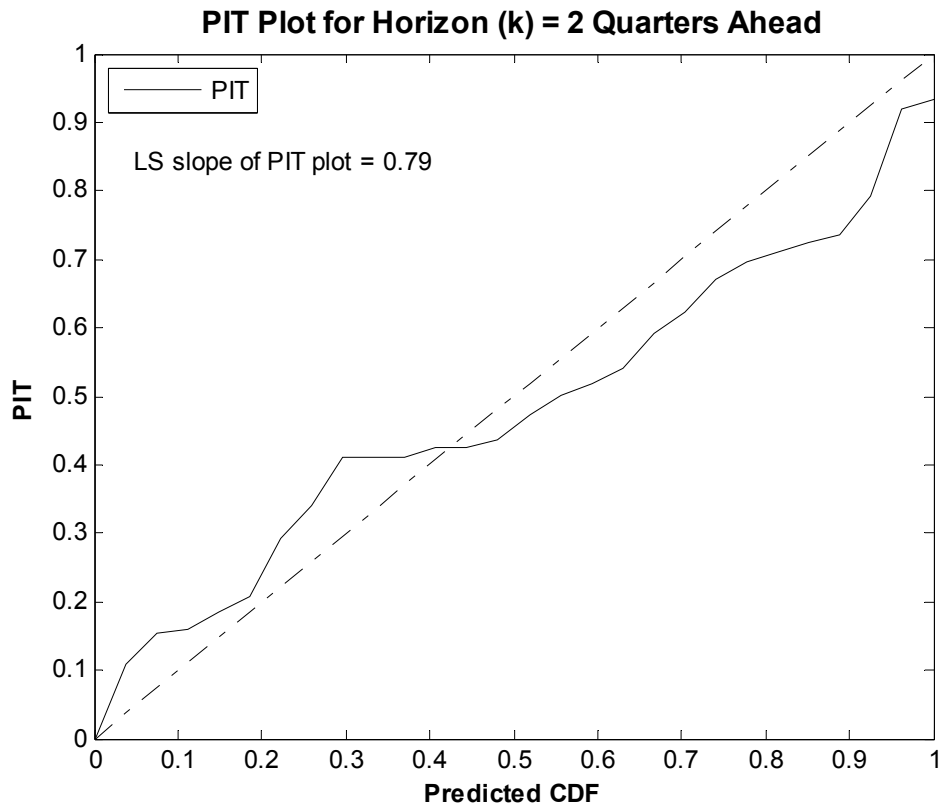
Note: Based on 27 observations of fan chart parameters and RPIX inflation over the period 97Q3 to 04Q1. The PIT values are the $p_{0,t}$, the inflation observations mapped to their cdf values on the forecast 0-period horizon inflation distribution function.

FIGURE 3:




Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 97Q4 to 04Q2. The PIT values are the $p_{1,t}$, the inflation observations mapped to their cdf values on the forecast 1-period horizon inflation distribution function.

FIGURE 4:




Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 98Q1 to 04Q3. The PIT values are the $p_{2,t}$, the inflation observations mapped to their cdf values on the forecast 2-period horizon inflation distribution function.

FIGURE 5:




Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 98Q2 to 04Q4. The PIT values are the $p_{3,t}$, the inflation observations mapped to their cdf values on the forecast 3-period horizon inflation distribution function.

FIGURE 6:



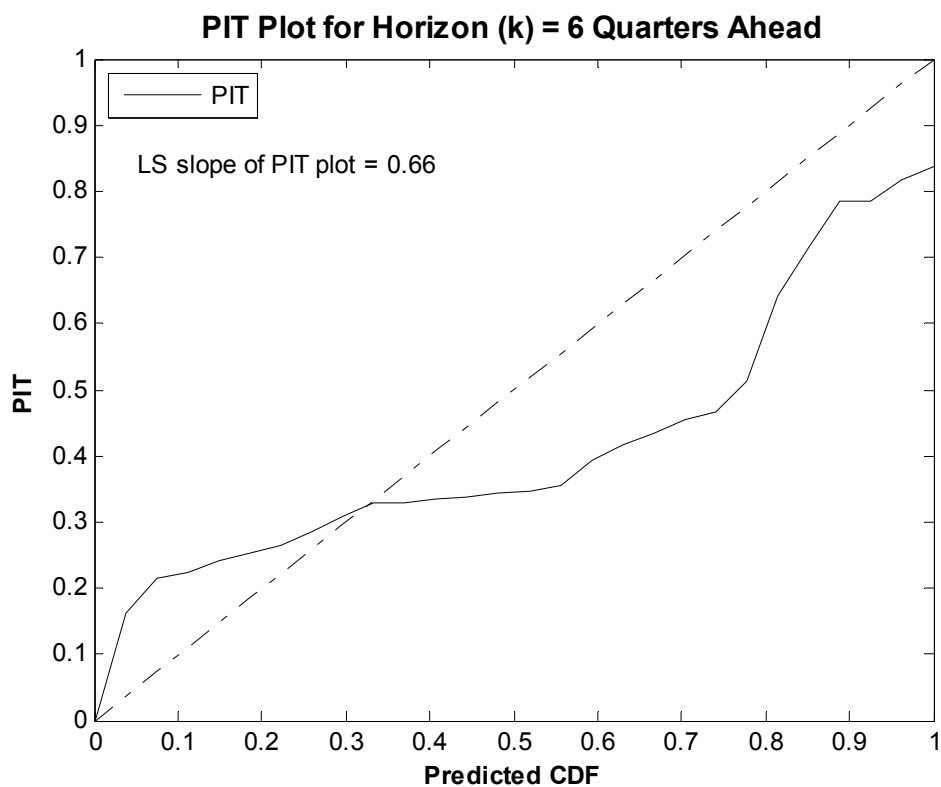
Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 98Q3 to 05Q1. The PIT values are the $p_{4,t}$, the inflation observations mapped to their cdf values on the forecast 4-period horizon inflation distribution function.

FIGURE 7:




Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 98Q4 to 05Q2. The PIT values are the $p_{5,t}$, the inflation observations mapped to their cdf values on the forecast 5-period horizon inflation distribution function.

FIGURE 8:




Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 99Q1 to 05Q3. The PIT values are the $p_{6,t}$, the inflation observations mapped to their cdf values on the forecast 6-period horizon inflation distribution function.

FIGURE 9:



Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 99Q2 to 05Q4. The PIT values are the $p_{7,t}$, the inflation observations mapped to their cdf values on the forecast 7-period horizon inflation distribution function.

FIGURE 10:



Note: Based on 27 observations of fan chart parameters over the period 97Q3 to 04Q1, and 27 observations of RPIX inflation over 99Q3 to 06Q1. The PIT values are the $p_{8,t}$, the inflation observations mapped to their cdf values on the forecast 8-period horizon inflation distribution function.

TABLE 1: SAMPLE MOMENTS OF $z_{k,t}$

Statistic \ Horizon (k)	$k=0$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$	$k=8$
	Summary statistics								
Mean	-0.048	-0.009	-0.001	-0.034	-0.078	-0.112	-0.183	-0.272	-0.368
Variance	0.866	0.679	0.475	0.557	0.513	0.369	0.310	0.260	0.189
Skewness	0.031	0.247	0.283	0.315	0.353	0.655	0.923	0.857	1.503
Kurtosis	3.838	2.977	2.972	2.678	3.028	2.786	2.815	3.108	4.611
n	27	27	27	27	27	27	27	27	27

Notes: Estimated for the constant-rate model using fan chart parameter forecast data over 97Q3 to 04Q1 and RPIX inflation over 97Q3 to 06Q1.

TABLE 2: P-VALUES OF STANDARD NORMALITY PREDICTIONS ASSUMING $z_{k,t}$ IS IID

Test \ Horizon (k)	$k=0$	$k=1$	$k=2$	$K=3$	$k=4$	$k=5$	$k=6$	$k=7$	$k=8$
t -test of mean prediction	0.791	0.956	0.995	0.816	0.578	0.347	0.100	0.010*	0.0002**
Variance ratio test of variance prediction	0.339	0.112	0.011*	0.034*	0.019*	0.001**	0.0003**	0.0001**	0.0000**
JB test of normality prediction	0.878	0.858	0.826	0.656	0.541	0.388	0.168	0.228	0.005**

Notes: Estimated for the constant-rate model using fan chart parameter forecast data over 97Q3 to 04Q1 and RPIX inflation over 97Q3 to 06Q1. Tests carried out using Eviews. * indicates significant at 5% level and ** indicates significance at 1% level.

TABLE 3: ESTIMATED AR(1) PROCESS FOR $z_{k,t}$ AND P-VALUES OF STANDARD NORMALITY PREDICTIONS

Result \ Horizon (k)	$k=0$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$	$k=8$
Estimated AR(1) process									
$\hat{\rho}_k$	0.138	0.104	0.495	0.514	0.561	0.619	0.566	0.572	0.822
P-value	0.486	0.598	0.010*	0.005**	0.003**	0.0007**	0.002**	0.002**	0.0000**
P-values									
Test of mean prediction	0.418	0.484	0.491	0.460	0.405	0.383	0.299	0.220	0.245
Test of variance prediction	0.361	0.121	0.045*	0.000**	0.000**	0.025*	0.006**	0.000**	0.013*

Notes: Estimated for the constant-rate model using fan chart parameter forecast data over 97Q3 to 04Q1 and RPIX inflation over 97Q3 to 06Q1. Estimates of $\hat{\rho}_k$ were obtained using EViews regressions, and estimates of P-values were obtained using Monte Carlo simulation with 20000 simulation trials. * indicates significant at 5% level and ** indicates significance at 1% level.

TABLE 4: P-VALUES OF TESTS ASSOCIATED WITH EQUATION (9) $\Delta z_{k+1,t} = \alpha + \beta z_{0,t} + u_{k,t}$

Prediction \ Horizon (k)	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$	$k=8$
$\alpha = 0$	0.764	0.630	0.518	0.602	0.741	0.903	0.893	0.932
$\beta = (k+1)^{-0.5}$	0.011*	0.0000**	0.0001**	0.0000**	0.0000**	0.008**	0.005**	0.0000**
$\text{var}(u_{k,t}) = (k+1)^{-1}$	0.001**	0.094	0.003**	0.014*	0.0096*	0.0006*	0.0002**	0.455
$u_{k,t}$ is normal	0.069	0.630	0.465	0.294	0.779	0.835	0.592	0.860

Notes: Estimated for the constant-rate model using data from 97Q3 to 04Q1. Estimates were obtained using EViews and all tests are based on chi-squared test statistics. * indicates significant at 5% level and ** indicates significance at 1% level.