



BANK OF ENGLAND

Speech

Uncertain uncertainty

Speech given by

Martin Weale, External Member of the Monetary Policy Committee, Bank of England

To the Institute and Faculty of Actuaries, London

29 March 2011

I would like to thank Matthew Corder for his assistance and I am also grateful for helpful comments from other colleagues. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Monetary Policy Committee

“The only function of economic forecasting is to make astrology look respectable” J.K. Galbraith

Introduction

It is fair to say that Professor Galbraith’s views on economic forecasting are likely to command strong support even, or perhaps particularly, among economists. Many economists might, if asked, see economic forecasting as the illegitimate offspring of their subject born of a union between bad economic theory and the uncertain application of statistical techniques. And a wide range of commentators loses no opportunity to observe that forecasts are generally wrong, sometimes losing the logic of their arguments in the process. For example, about ten years ago, the *Financial Times* observed that economic forecasts were often least reliable when times were most uncertain. I wonder whether I was the only reader who wondered how one knew when times were most uncertain except by observing that forecasts have large errors.

One might think that similar forecasting problems have been faced by life actuaries. For many years mortality rates of old people decreased slowly. Then, around 1980, the proportionate rate of decline increased sharply. Since this change was without precedent, a natural conclusion was that the pre-1980 pattern would soon re-establish itself. In fact, as we have seen, in more recent years the rate of decline of mortality rates has, if anything accelerated. While I do not have anything to offer on how forecasts of mortality rates might be improved, I can say that if key economic time series behaved like mortality rates, the problems faced by those involved in economic forecasting would be much greater than they actually are.

Of course the reality is that both economic and actuarial forecasts are needed and this is perhaps why, despite the best efforts of their detractors, forecasters seem to be irrepressible. The question of course is not whether they are right or wrong but whether policy is better made, or products such as pensions are better constructed, with forecasts or without them.

One could imagine economic policy being set without forecasts. Indeed some work colleagues and I did over twenty years ago (Weale *at al.*, 1989) looked at policy rules on the assumption that both interest rates and tax rates responded to the data but not to views about what was going to happen in the future. Slightly more recently, the Taylor Rule (Taylor (1993)) suggested a framework where the interest rate was set with reference to the current rate of inflation and the current output gap (although the current output gap cannot sensibly be estimated without a forecast and even with it can hardly be regarded as certain).¹ But, since policy moves such as interest rate changes take some time to have an effect, it makes more sense to set policy not with reference to the current state of the economy but with reference to where it is expected to be or, more explicitly, where inflation would be expected to be at any given interest rate setting (Svensson (1997a, b)).

¹ See Orphanides and van Norden (2005).

On the Monetary Policy Committee we normally aim to set monetary policy so as to achieve the inflation target in two to three years time. This horizon was adopted because the Bank's analysis suggested that interest rate changes had their most powerful influence on inflation about two years after they were made. A degree of flexibility is needed because different types of inflationary shocks work through in different ways. There are, for example, very good arguments that, after large cost shocks to a depressed economy, one should not rush to bring inflation back to target.

More generally, if the MPC aimed to achieve the inflation target too quickly, the process might well match that of someone trying desperately to adjust a shower to the right temperature. By failing to take account of the lag between turning the taps and variations in the temperature of the water that comes out, the bather is condemned to a sequence of water which is too hot followed by water which is too cold and so on. Only a degree of patience can deliver a comfortable shower.

Inevitably, if we aim to achieve our target at some time in the future we have to take account of where we might expect the economy and particularly the inflation rate to be in the absence of any change to economic policy. This can be done only by means of a forecast and the question is not whether the forecast is right or wrong but what is the best means of producing the forecast. Relying only on current information would be best only if the current state of the economy were the best guide we have to the future. Historic evidence suggests that economic forecasts, despite the strength of Professor Galbraith's views, have more predictive power than the assumption that either the current rate of inflation or the current rate of economic growth would continue indefinitely (Poulizac, Weale and Young, 1996).

So forecasts are needed in the making of policy. Given that they are needed, and that they influence policy setting they should be in the public domain. But, given that they are in the public domain, it is important that producers of forecasts should help outsider users of them to understand their limitations. In particular users need to understand that forecasts are neither right nor wrong; they are simply the best that forecasters can do with the information available. Experienced forecasters know better than to claim that they are good at forecasting on the basis of one or two years of good luck. But although it is always important for forecasters to ensure they are explaining their forecast in the clearest possible way, the more urgent need to update and improve the forecast can sometimes crowd out this important work. I therefore wanted to take the opportunity today to talk about ways to explain forecasts – particularly the uncertainty around them.

A focus on the uncertainty surrounding a forecast rather than on a single trajectory is widely regarded as the most appropriate way of communicating the realities of forecasting, whether applied to the economy, to mortality or to any other variable. And there are very good reasons why such information is important both for policy-makers and for other users of forecasts. Businesses which insure people against the risk of longevity are much more likely to be concerned about unexpectedly low rates of mortality than about unexpectedly high rates of mortality. The former may result in the insurance company being unable to meet its promises while the latter lead to high profits for shareholders. While the latter might be welcomed, the

former would put the business model at risk; one could hardly imagine a well-run business feeling that the two risks balanced out. Similarly, for central bankers; after a period like the present, in which inflation has been above target for fifteen² months in a row and for over fifty of the last seventy-two months, it is easy to see why they might be more concerned about inflation being above target at the relevant forecast horizon than about it being below target. The argument is that continuing above-target inflation could lead to inflationary expectations becoming entrenched. On the other hand, if inflation were to be below target after a sustained period above target the public would be unlikely to think that the central bank had lost interest in its target or that an unmanageable deflation threatened. Rather they would think that these sorts of fluctuations were normal in the current choppy environment.

But enough of the reasons why people need to know more about forecasts than simply being given forecast trajectories. Let me now move on to considering how this might best be done. There is a wide variety of possible ways. I therefore want to consider how it has been done in the past before making my own suggestions for possible developments.

Presentation of Uncertainty

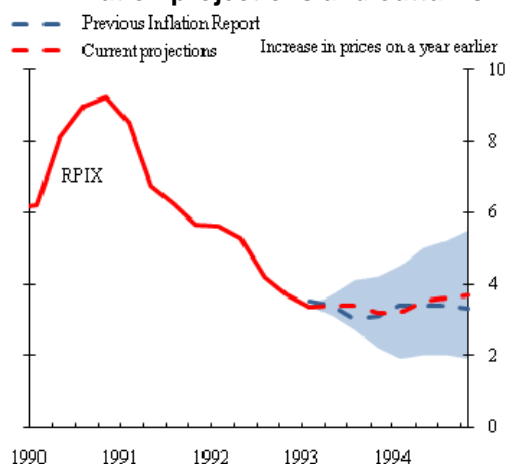
I hesitate the claim that HM Treasury was the first body in the UK to indicate the uncertainty surrounding their forecasts. But the 1976 *Industry Act* obliged them to publish their forecasts and, in 1979, they started to include tables showing the mean absolute errors associated with their forecasts. These tables could have done nothing except convey the (correct) impression that users, even if not themselves aware of the Treasury record, should not expect the forecasts to be accurate. It is questionable how far the presentation of mean absolute error rather than standard error is helpful because, should one want to use the information to parameterise a density function and then derive event probabilities such as the chance that output is stagnant or falling, this is more conveniently done in terms of a standard error. Wallis (1999) explains that if policy-makers regard the costs of the failure to meet their targets as linear in the deviation of the outturn from those targets, then policy should respond to the median outcome and losses should be assessed with reference to the mean absolute deviation from target. But this does not translate into an emphasis on the mean absolute forecast error unless the only costs policy-makers were concerned about arose from being unable to predict the future state of the economy rather than the consequences for the economy itself – a situation which seems most unlikely.

Other UK forecasters did not immediately follow the Treasury lead. However, in 1993, the Bank of England began to produce its *Inflation Report*. From then until November 1995 it published a chart like that shown in Chart 1, which presented the central forecast together with an indication of the mean absolute error.

² The original version of this speech erroneously stated that inflation has been above target for twenty one months in a row.


In February 1996, the Bank began publishing something looking like the current fan chart for the first time. There are two key differences between this and what was done between 1993 and 1995. First, the focus moved away from an analysis based on past forecast errors to one intended to represent the subjective view of the Bank.³ Secondly, it was decided to depart from the assumption that risks are symmetric. In order to help produce an asymmetric distribution, Bank staff use a composite of two normal distributions (Gibbons and Mylroie, 1973) in which the variance of the distribution above the central parameter could differ from that below the central parameter. Of course, like the models used by the Committee to produce the central projection, this distribution is only a tool. It does not necessarily represent all Committee members' view of the true distribution. But, as Britton et al (1998) note in their explanation of the fan chart, the use of a distribution allows the Committee to represent the belief that the distribution of possible outcomes might not be symmetrically distributed about the central reference point without needing to consider the infinite combination of possible scenarios, from which the distribution might be drawn. The most recent fan chart is shown in Chart 2.

Chart 1: May 1993 Inflation Report chart of RPIX inflation projections and outturns^(a)



(a) Chart reproduced using parameters published on Bank of England website

Chart 2: February 2011 *Inflation Report* chart of CPI inflation projections



The National Institute did not start to publish regularly information on its forecast errors until 1993. In April 1996, for the first time it computed the standard deviation of the forecast errors from its past record and used this to compute the probability that inflation and GDP growth would lie in particular ranges, so called event probabilities. This was presented in tabular form as the example from October 2007 shows.

In the late 1990s, work using stochastic simulations of its economic model suggested that a model-based analysis of forecast uncertainty gave results similar to one deduced from past forecast errors. This need not have been the case because, as every forecaster knows, forecasts reflect forecasters' judgements as well as the models that they use.

³ Nevertheless, it is clear from the parameterisation mentioned above that the initial subjective views of the Bank in early 1996 were largely conditioned on the past forecast errors.

Table 1: Event Probabilities for Inflation and GDP from the National Institute forecast for October 2007.

CPI inflation	2007Q4	2008Q4	GDP growth	2007	2008
Central projection	1.7	2.0		3.1	2.2
Root mean squared error	0.29	0.62		0.59	1.12
<i>Probability of 12 month CPI inflation falling in the following ranges</i>			<i>Probability of annual growth rate falling in the following ranges</i>		
less than 1 per cent	1	5	less than 0 per cent	0	3
1 to 1.5 per cent	23	16	0 to 1 per cent	0	12
1.5 to 2 per cent	61	29	1 to 2 per cent	3	29
2 to 2.5 per cent	15	29	2 to 3 per cent	40	33
2.5 to 3 per cent	0	16	3 to 4 per cent	51	18
more than 3 per cent	0	5	more than 4 per cent	6	5
	100	100		100	100

At that time I had certainly hoped that greater attention to forecast uncertainty by the National Institute, in combination with the Bank's approach, would trigger a more general debate by economic forecasters about how best to represent uncertainty. That unfortunately did not happen. However the National Institute's experience did demonstrate the pitfalls of relying on past forecast errors. In the 1970s and 1980s inflation was not only high but also volatile. The National Institute's forecasts of inflation in the 1970s and 1980s inevitably showed fairly large errors and these were translated into high probabilities of inflation being a long way from both the forecast and the Bank's target in the late 1990s. As we know, inflation was in fact extraordinarily stable during this period, with the result that the National Institute substantially over-stated the probabilities of inflation being a long way from its target. Thus a formal study (Mitchell and Hall, (2005)) showed that while the outturns for inflation were coherent with the Bank's fan-chart, their distribution was incompatible with the National Institute's view at that time about the uncertainty surrounding its forecast.

This experience has, of course, a more recent moral as Table 1 might hint. In 2008, GDP was in fact slightly lower than in 2007, something that the National Institute forecast gave a probability of only three per cent. The problem was even worse in 2009 when the UK economy showed the largest contraction in output since 1921. An analysis based on recent forecast errors and the assumption of a normal distribution understated the risk of relatively rare events.

Rare Events

This problem can be understood by asking how often one should expect an economic depression. Over the last hundred years there have been six periods in which the output of the economy has been depressed below its previous peak, ignoring the distorting effects of the two world wars. So what is the risk of a recession starting? If we assume that a recession typically lasts for two years then there have been ninety-four years in which recessions might have started. The probability of a recession starting, given one is not already underway, is 6.4%. But this has a standard error of 2.5%. So a reasonable inference is that there is a one in eight chance of the probability of a recession being almost ten per cent. There must be

some businesses which would pay more attention to the risk of recession if they thought there was a reasonable chance it could be a one in ten year event rather than if they assumed the risk was not much more than half of that.

To use more data in order to obtain more precise estimates of the probabilities of recession is not necessarily a solution. The British economy before the First World War was very different from what it is today and I am not sure that we can learn an enormous amount about the frequency of the risks facing the contemporary economy from a study of Queen Victoria's reign. Indeed, although national accounting data for much of the nineteenth century exist, it is not clear that they are well-suited to identifying a modern economic cycle.

Could the problem be that the nature of the shocks has changed rather than that rare events have turned out to be less rare than thought. On the face of it I am not sure how one might distinguish the two or even that they are distinguishable. These problems have, of course, been appreciated by authors such as Basak and Shapiro (1998), Artzner (1999) and Berkowitz (2001) who have suggested various approaches to the problem of testing models of the probabilities of rare events. Nevertheless, one cannot make a silk purse out of a sow's ear. The probabilities of rare events must be highly uncertain and those who estimate such probabilities should make sure that they are described as such.

As both the Governor and Charlie Bean have noted in the past (King et al (2010) and Bean (2010)), to avoid a spurious degree of precision the MPC does not publish the details of the distribution of the shocks to which it thinks the economy might be subject outside the central 90% of the distribution (although of course it could not publish an expected value for the outcome without making some assumption about the tails of the distribution). But, should such rare events turn out to be less rare than was thought, then the published distribution will seem *ex post* to be wrong.

Uncertain Probabilities More Generally

But it is not only with rare events that one might want to recognize that the probabilities are uncertain. I think it is fair to say that, just as the Monetary Policy Committee feels more than uncomfortable about drawing attention to a single forecast path for the economy, so too it might feel uncomfortable about the idea that it might attach a precise probability to any particular event. A sense of this malaise can be gained from the *Inflation Report*. The Committee, reasonably enough, has felt the need to say what it thought was the probability that inflation would be above or below the two per cent target at some point in the future. A probability by its nature is a precise proportion. And a number like this could be calculated from the assumption that the Committee makes about the dispersion of risks surrounding its central projection.⁴ But,

⁴ The density representing the errors to the forecast is assumed to take the form of a two-piece normal distribution (Johnson, Kotz and Balakrishnan (1994), page 173). The two-piece nature makes it possible to represent a skew should the Committee think that appropriate.

as the Governor discussed in his speech at the Royal Society last year (King et al (2010)), the Committee feels that such precision is inappropriate; instead it shows Chart 3. The reader can infer from this that, in the short term, the MPC thinks that the probability inflation will be above target is close to 100 per cent and that, by late 2012, this probability drops back to around 50 per cent. This 'ribbon', can be drawn directly from the MPC's inflation fan chart, as the boundaries of the ribbon are defined by the edges of the fan chart band into which 2% falls. This has the advantage of allowing the Committee to discuss the balance of risks around the inflation target without requiring agreement on the specific probability distribution used to derive it. The chart's imprecision around the probability of inflation being above target also gives some impression that, as I have indicated above, probabilities are uncertain. But the width of the ribbon may not necessarily be the best way of indicating the MPC's uncertainty around the probability that inflation is above target. I'd like to turn to some other ways of conveying this uncertainty.

Chart 3: The Ribbon Chart – An Indicator of the Probability Inflation will be above the target

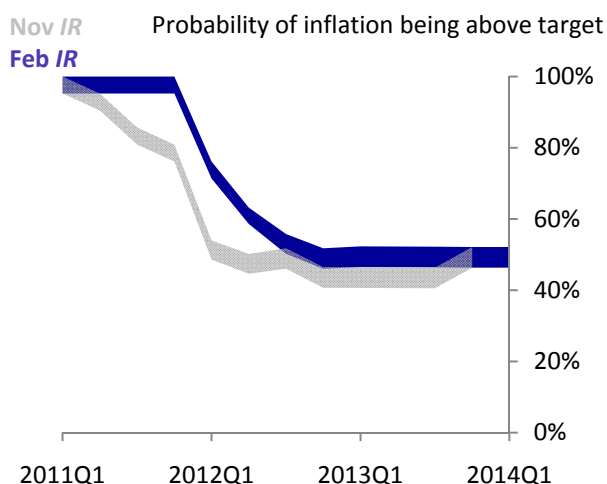
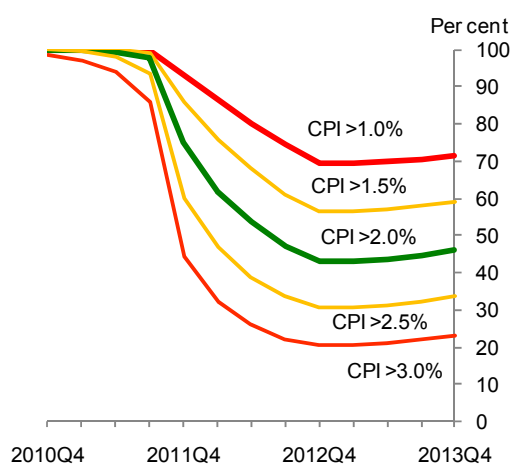


Chart 4: A chart showing probability that inflation will exceed certain thresholds



It is clear that users of the *Inflation Report* feel that the views of the Monetary Policy Committee on the uncertainties surrounding the inflation projection could be set out more clearly. For example Chris Giles of the *Financial Times* has made some helpful suggestions from which Chart 4 is derived. This indicates the probability that the Committee attaches to the inflation being above or below particular threshold values.

Of course it could also be argued that users who want a chart like this can work it out for themselves, as Chris Giles has done, based on the probabilities published with numerical parameters of the *Inflation Report* probability distribution. But this is hardly satisfactory. The *Inflation Report* should convey its message in the form most helpful to its readers. This does not, of course, mean that we should present what readers request if that might distort the message.

The question then arises whether the chart does run the risk of distorting the message the Committee wish to communicate. Not surprisingly Chris is particularly interested in the probability that inflation is going to be

below 1% or above 3%. These are of course the thresholds at which the Governor has to write to the Chancellor. But the Committee has a point target and presenting ranges based on the 1% and 3% threshold may confuse the public into either misinterpreting the Committee's target or into believing that an inflation rate of slightly more than 3% is significantly worse than one of exactly 3%, which is not true. And the Committee does already provide more information on the probabilities of inflation being significantly away from target using illustrations such as Chart 5. For my purposes, today, I will consider these ranges because they are significant departures from the 2% target.

More importantly, as with the earlier discussion of the chance of inflation being above or below target, a diagram like Chart 4 or Chart 5 would convey a precision which the MPC might feel did not accurately represent its views. So the question I want to address is are there ways in which the MPC can helpfully represent the uncertainty it may feel about the uncertainty facing the economy?


Multiple Sources of Uncertainty

One reason why there may be uncertainty about the probabilities of various outcomes is that the parameters of the probability density function on which they are based are themselves uncertain. Perhaps an easy way to think of this is that there is a large number of plausible forecasts, generated from a range of different models and, indeed a range of different people with different views.⁵ Each of these forecasts has a density function associated with it. So each forecast could generate its own fan chart and the uncertainty about the probabilities arises from the reality that we are uncertain which fan chart to use.

Before looking at the implications of such a situation for uncertainty about probabilities, I need to say something about how one might generate an overall fan chart in such a situation. Wallis (2005) suggests that we should draw on the literature associated with mixed distributions. If there are a number of different possible density functions, he suggests giving them equal weights. If this is applied to a situation where there are just two possible density functions, generated by, say two competing models, then the principle is clear enough.

If there is a large number of possible models, verging on a continuum with its own density function, then application of Wallis' suggestion means weighting the density function generated at each point on the

Chart 5: Frequency distribution of CPI inflation based on market interest rate expectations and £200 billion asset purchases



⁵ I am grateful to George Kapetanios for suggesting we should see the problem in this light.

continuum by the probability associated with it. Such a situation might arise if, for example, there were a range of central forecasts and each had a shock around it. If both the forecasts and the shocks were normally distributed, and the latter were independent of the model to which they related and had a known variance, then the resulting overall density function would also match a normal distribution. So in such a situation it would be easy to generate a combined fan chart. But the uncertainty about the probabilities arises from the uncertainty surrounding the parameters of the forecast as a concrete example illustrates. Uncertainty about the appropriate model is a known unknown and the effects of subsequent shocks are unknown unknowns.

Suppose that we were 50% confident that, in the absence of shocks, the inflation rate would lie between 2.4% and 2.7% at some point in the future. And we also assumed that the standard deviation of inflation as a result of unforecastable normally distributed shocks was 1%. The probability that inflation would exceed 3% if the rate in the absence of shocks was 2.4% is 27% while an unperturbed rate of 2.7% leads to a probability of exceeding the 3% threshold of 38%. So in this example, instead of giving a precise probability of inflation being 3% or more, the Committee might wish to say that it was 50% confident that the probability lay between 27% and 38% or perhaps even more generally that it was moderately confident that the risk lay between 25% and 40%. Speaking for myself, I think that this sort of presentation would convey the air of uncertainty surrounding the forecast better than would a single probability – and also that the benefits to be gained from this sort of approach more than outweigh the disadvantage that the model of risks is more complicated than we have at present.

How far should this be taken? If we are to talk about uncertainty surrounding the central value of the distribution one might reasonably well ask whether we should be concerned with uncertainty surrounding the shocks? One might, for example, want to assume that the variance is itself distributed as a χ^2 variable with some assumed number of degrees of freedom. Indeed Tiechroew (1957) sets out the resulting overall distribution if that is the only source of uncertainty. But a moment's thought shows that uncertainty about the variance leads, on its own, to very considerable uncertainty about the extreme probabilities, but none at all about the median. Thus I for one would feel most uncomfortable with relying solely or mainly on uncertainty about the variance as a way of representing the uncertainty I feel about probabilities of particular events. And I do not think that, at present, there is much to be gained from a composite approach in which attention is paid to the effects of uncertainty about the variance as well as uncertainty about the central value – or indeed for worrying about the uncertainty of the skew parameter.

There is a separate question whether the parameters of the density function associated with each of the possible models should be assumed to be related to the parameters of these models. Obviously, when fitted to past data, some models do well and others do badly. A surprising number of forecasters claim that they are better at forecasting than are other forecasters, claims which are most easily supported only on the basis of a small number of lucky out-turns. Indeed, as suggested at the start, there is historic evidence that forecasts produced using macro-economic models do out-perform forecasts generated by the application of

simple statistical techniques. But, for the purposes of conveying uncertainty about uncertainty, I doubt that this is so important as to justify departing from the simple assumption that all forecasts have the same error density function associated with them.

Finally, there is a separate question whether it makes sense to assume that the disturbances associated with each particular forecast in the density function are independent of that forecast. This is an issue which merits further thought; for the time being we have assumed that the parameters of the disturbances for each individual forecast are independent of the value of that forecast. With this structure in mind we can turn to the sort of representations of uncertain uncertainty that a body like the MPC might produce.

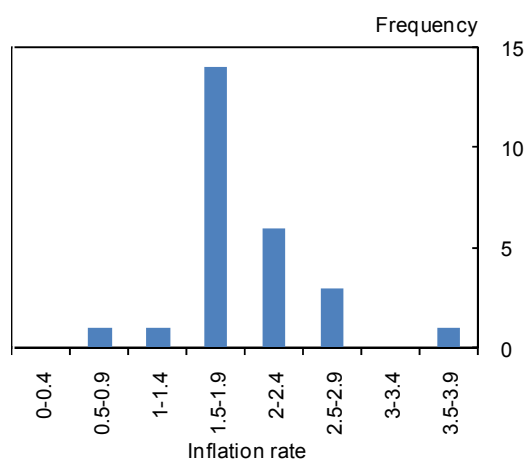
Application to the Fan Chart

In order to illustrate uncertain uncertainty, I have assumed that the dispersion of underlying models is represented by the spread of independent inflation forecasts for the final quarter of 2012 as collected by the Treasury late in February. This is shown in Chart 6.

This chart is in some sense comforting for the MPC. Despite the current elevated inflation rate, it shows that the modal value for inflation in nearly two years time is under 2 per cent. But there are some other features of the distribution which make it of

considerable interest. The mean value of the forecasts is 2.06 per cent. And, as the figure makes clear, the distribution is anything but symmetric. It shows a clear skew to the right. If I attempt to represent this skewed distribution by fitting a two-part normal distribution,⁶ I compute the mode to be just under 1.9 per cent with the standard deviation parameter for points below the mode to be 0.40 and for those above the mode to be 0.6. This analysis where the mode is 0.16 percentage points below the mean can be compared with the most recent forecast produced by the Monetary Policy Committee which shows the mode to be 0.4 percentage points below the mean.

Chart 6: Independent Inflation Forecasts for 2012 Q4



⁶ Although it is not clear that this is a good choice to represent the dispersion of independent forecasts.

I show in Chart 7 a fan chart for the modal forecast on the assumption that the mode has the distribution represented by the two-piece normal distribution⁷ fitted to the forecasts of Chart 6.

To build up an overall distribution from the combination of uncertainty about the model and uncertainty I take this two-piece normal distribution to represent the range of possible modal forecasts that might be envisaged by the MPC. I add to the modal forecasts represented by that distribution a shock, which is assumed to be normally distributed. The variance of this shock is chosen so that the overall variance of the combined density forecast matches that assumed by the MPC.

Chart 7: An Example of an Uncertain Modal Forecast of Inflation




Chart 8: Histogram of inflation forecasts in 2013Q1 derived from February 2011 *Inflation Report* and possible new distribution

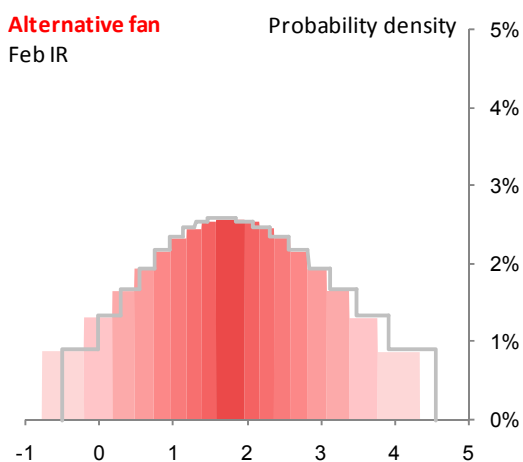



Chart 9: Histogram of inflation forecasts in 2014Q1 derived from February 2011 *Inflation Report* and possible new distribution



A consequence of this is that the skewness of the resulting fan chart is much lower than that adopted by the MPC. Adding on a symmetric disturbance to a skewed distribution like that used in the current fan chart has the effect of reducing the overall skewness of the combined distribution.⁸ This provides an interesting alternative perspective to the scale of the skewness incorporated into the MPC's recent forecast. I show in

⁷ We assume that the standard errors of the quarterly inflation rates are half those of the annual inflation rates in order to build up the chart in the first year.

⁸ The skewness depends on the difference between the standard deviations of the left and right components of the two-piece normal distribution. But, with independence of shocks it is the variances and not the standard deviations of the two distributions which are added together. This has the effect of reducing the gap between the standard deviations of the combined distribution below that of the initial two-piece normal distribution.

Chart 8 and Chart 9 histograms of the existing fan chart and the fan chart generated with the above assumptions nine and thirteen quarters into the forecast.

Representation of Uncertain Uncertainty

I now move on to examine other probabilities that could be presented as uncertain on the assumption that the central forecast has a skewed, two-part normal distribution, while the shocks around that most likely outcome are symmetrically distributed. Chart 10 shows the probability that inflation is going to be above target and is an alternative way of presenting what was shown in Chart 3. My own view is that it suggests a margin of uncertainty around the probability considerably greater than one might infer from Chart 3. I am, of course, conscious that discussing uncertainty around

uncertainty could lead to confusion and some might argue that this chart means that the Committee risks getting into the situation of stating that “It is about as likely as not that inflation is about as likely to be above target as below it in the medium term”. But I would argue that a simpler option would be to state that ‘Although uncertain, inflation is broadly about as likely to be above target as below it in the medium term. In any case, were the MPC to approach the issue in this way, it might well end up with a lower spread than that which arises from using the dispersion of forecasts in Chart 6.

Chart 10: Replacement ‘ribbon’ chart




Chart 11: Probabilities when the mode is drawn from the 25th percentile





Chart 12: Probabilities when the mode is drawn from the 75th percentile



I can also consider the implications of this approach for charts derived from those suggested by Chris Giles. One way of doing this is to show two charts representing the 25th and 75th percentiles. This is done in Chart 11 and Chart 12.

These charts have to be considered together. They allow us to read off, for example, that there is a 50% chance that the probability of inflation being above 3% in 2013Q4 is between 19% and 29%. But processing information can be costly for the receiver, and the fact that the charts have to be read jointly might place too much of the processing burden on the reader and lead to confusion about the message. This highlights the risk of releasing too much information in an unfiltered way – as the Governor put it ‘less can be more’. (King et al (2010)).

An alternative presentation is that shown in Chart 13. Here the two charts blur into one another so that, for example, the 50% probability range for the probability that inflation is above 3% lies in the lower red fan, which shades from slightly red to solid red and back again. Similarly, the probability that inflation will be above 2.5% lies in the lower gold fan, which shades from slightly gold to solid gold. In my view, this is probably a more satisfactory way of conveying the uncertainty surrounding the probabilities

Chart 13: Probability charts with the bands blended together to reflect uncertainty about the mode





Chart 14: Alternative version of Chart 13 where gold bands are omitted



I should, nevertheless, draw attention to one problem associated with Chart 13. There I show only the percentiles comprising the central 50% of the distribution of the probabilities. But even these overlap; if the future became more uncertain this problem would become worse. One solution would be to leave out the gold parts, showing simply the green centre, for the range of probabilities that inflation will be at target, and the red parts. Even greater uncertainty might mean that only the red parts could be shown without any confusion.

Of course these are not the only charts representing probabilities in the *Inflation Report*. If desired it would also be possible to produce versions of the frequency distributions of inflation and growth at particular dates in a way which reflects the uncertain nature of such distributions.

Policy Implications

Where does this leave the practical problem that MPC members face in deciding how to set interest rates? On the one hand the future is uncertain – that is the nature of the future. But on the other hand we need to come to a conclusion and each of us needs to cast a vote each month. We do not have the luxury of being able to spread our votes between different choices.

As you know I have been one of those voting for an increase in the interest rate since our January meeting. I do not want anyone to infer that I believe Chart 6 represents the views of the Monetary Policy Committee. Indeed I certainly have in my mind a number of possible and competing models, as I expect would anyone who seeks to understand what is happening to the economy and where it might be going in the future. But Chart 6 does indicate why there might be a divergence of views in the Committee with different people voting in different ways even though we can agree that something like Chart 7 represents our collective judgement. And obviously I vote in the light of my individual view and not of the collective judgement.

As the minutes of the last meeting made clear, there are three reasons why inflationary pressures might be more marked than the Committee's central view suggests. First of all, expectational effects, although built into our forecast, could easily be more marked. After all, since there is a real risk that inflation will rise close to 5% p.a. can we be confident that pay bargaining and price setting will take place on the assumption that people expect it be appreciably lower than this in the future? Secondly, commodity prices may be more buoyant than has been assumed in our forecasts. And thirdly, there may be more upward pressure on profit margins coming from the aftermath of sterling's depreciation in 2008. Indeed, if prices of traded goods are set in part in international markets, then this is quite likely.

These are all reasons for being more pessimistic about inflation than is our modal forecast which, as I have noted, is very similar to that represented in Chart 7. Obviously there are also reasons for expecting inflation to be lower than the mode and my guess is that some of my colleagues would find themselves in the left-hand part of Chart 6. Such a position is easily understandable; I regard the probability that they may be right as more than trivial. But on balance I find myself more concerned about inflation than the Committee's collective judgement suggests and thus continue to believe that the monetary stance should be tightened.

Nevertheless, an obvious undercurrent of this talk is that people who make decisions on the basis of expectations about the future may find themselves badly wrong. It is perfectly possible that, if economic activity remains weak and the risks which concern me do not materialise then the case I currently see for higher interest rates will fade. One vote does not, on its own, set interest rates. But the preceeding

observation raises an interesting point. It would be bad for the MPC to raise the interest rate and then reduce it again after a few months. But it would also be bad for the MPC to keep the rate unchanged because of this risk and then find that the concerns I have do in fact materialise. Rather than worry about the balance between these two derived risks, perhaps the more important thing is to remember that our task is to bring inflation back towards its target.

Conclusions

Over the years since it was first published, the *Inflation Report* has set standards for the presentation of uncertainty. A key feature of the report has been a reluctance to suggest precision where that precision cannot be justified and for this very good reason the details of the central forecast are not published in the *Inflation Report* but are released one week afterwards. The Monetary Policy Committee has tried to draw attention to risks and uncertainty. And the success with which it has done this is demonstrated by the pressure for clearer information on the probabilities of particular outcomes. But this in itself suggests a degree of precision with which at least this member of the Committee feels uncomfortable.

I have therefore endeavoured to set out a framework which breaks the uncertainty in the forecasts of the *Inflation Report* into two components, one which might be thought associated with uncertainty about the central path and the other which might be linked to the shocks that can hit the economy. This makes it possible both to explain why the probabilities of particular events are uncertain and to replace precise probabilities with ranges.

Were this structure to find favour, there might be much to be said for being less than completely precise when describing the range of probabilities. For example, rather than claim a level of 50% confidence that the probabilities lie between what is represented in Chart 11 and Chart 12, it might be better simply to claim reasonable confidence. That would be in keeping with the way in which I, at least, see forecasting, and it can be done with the structure set out here, but not with the existing fan charts.

Finally, I should remind you that this exercise has no more status than a contribution to a debate. I do not know whether my colleagues might find what is set out attractive, nor whether users might find it useful. As I have discussed above, there are many considerations that need to be weighed when deciding how to communicate about an uncertain future. On the one hand, there is a need to create a coherent framework to help create the illustrations used by the Committee in its explanation of its forecasts. On the other, there is a need for simplicity and a desire to avoid the spurious accuracy implied by using probability density functions as a tool to represent what is ultimately the unmeasurable true distribution of risks around our forecasts. But this is my perspective, and I hope it has, if nothing else, drawn your attention to the importance of communicating uncertainty as clearly as possible. I have yet to hear of astrologers attempting this.

References

- Artzner, P. (1999)**, 'Application of Coherent Risk Measures to Capital Requirements in Insurance', *North American Actuarial Journal*, Vol. 3 No. 2, pages 11–25.
- Barro, R.J. (2006)**, 'Rare Disasters and Asset Markets in the Twentieth Century', *Quarterly Journal of Economics*, pages 823-866.
- Basak, S. and A. Shapiro (1998)**, 'Value-at-Risk Based Risk Management: Optimal Policies and Asset Prices', *Wharton working paper*.
- Bean, C. (2010)**, 'Measuring recession and recovery: an economic perspective', *Speech at the RSS Statistics User Forum Conference on 27 October*, available at <http://www.bankofengland.co.uk/publications/speeches/2010/speech457.pdf>
- Berkowitz, J. (2001)**, 'Testing Density Forecasts with Applications to Risk Management', *Journal of Business and Economic Statistics*, Vol. 19, pages 465-474.
- Britton, E., Fisher, P. and Whiteley, J. (1998)**, 'The *Inflation Report* projections: understanding the fan chart', *Bank of England Quarterly Bulletin*, February, pages 30-37.
- Hall, S.G. and J. Mitchell (2007)**, 'Combining Density Forecasts', *International Journal of Forecasting*, Vol. 23, pages 1-13.
- Johnson, Kotz, and Balakrishnan (1994)**, *Continuous Univariate Distributions, Volumes I and II*, 2nd. Ed., John Wiley and Sons.
- King, M., Aikman, D., Barrett, P., Kapadia, S., Proudman, J., Taylor, T., de Weymarn, I. and Yates, T. (2010)**, 'Uncertainty in macroeconomic policy making: art or science', *Lecture at a Royal Society Conference in London on 22 March*, available at <http://www.bankofengland.co.uk/publications/speeches/2010/speech432.pdf>
- Mitchell, J. and S.G. Hall (2005)**, 'Evaluating, Comparing and Combining Density Forecasts using the KLIC with an Application to the Bank of England and NIESR 'Fan' Charts of Inflation', *Oxford Bulletin of Economics and Statistics*, Vol. 67, pages 995-1033.
- Orphanides, A. and van Norden, S. (2005)**, 'The Reliability of Inflation Forecasts Based on Output Gap Estimates in Real Time', *Journal of Money, Credit and Banking*, Vol. 37 No. 3, pages 583-601.
- Pouzilac, D., M. Weale and G. Young. (1996)**. 'The Performance of National Institute Economic Forecasts'. *National Institute Economic Review*. No 156. Pp55-62.
- Svensson, L. E. O. (1997a)**, 'Inflation-Forecast-Targeting: Implementing and Monitoring Inflation Targets', *European Economic Review*, Vol. 41, pages 1,111-46.
- Svensson, L. E. O. (1997b)**, 'Inflation Targeting: Some Extensions', *NBER Working Paper No 5962*.
- Taylor, J.B. (1993)**, 'Discretion versus policy rules in practice', *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, pages 195-214.
- Teichroew, D. (1957)**, 'The Mixture of Normal Distributions with Different Variances', *Annals of Mathematical Statistics*, Vol. 28, pages 510-512.
- Wallis, K.F. (2003)**, 'Chi-squared tests of Interval and Density Forecasts and the Bank of England's Fan Charts', *International Journal of Forecasting*, Vol. 19, pages 165-175.
- Wallis, K.F. (2005)**, 'Combining Density and Interval Forecasts: a Modest Proposal', *Oxford Bulletin of Economics and Statistics*, Vol. 67, pages 983-994.
- Weale, M., A. Blake, N. Christodoulakis, J. Meade and D. Vines. (1989)**. *Macroeconomic Policy: Inflation, Wealth and the Exchange Rate*. Allen and Unwin. London.