

Macro uncertainty and monetary policy

Bridget Rosewell

Volterra Consulting

July 2005

browsewell@volterra.co.uk

Paper prepared for the Brisbane Club conference on complexity, University of Queensland, July 2005

1 Background

Monetary policy has been credited with many powers in its long life. Governments have always had some form of monetary policy – even if it consisted of how most effectively to debase the currency. And governments have always borrowed and been active in money markets.

But it is only in the twentieth century that monetary policy has been upgraded to a role in managing the state of the macro-economy and governments have been assigned the task of running such policies. Even now fundamental debates persist about what monetary policy actually is and how it is transmitted through the economy. For some, it is about control of the money supply in direct quantitative terms. For others, it is about managing the price of credit to affect the willingness of all groups (consumers, businesses and indeed governments) to borrow.

What however is largely a consensus is that monetary conditions have an effect on the economy and such conditions have a more immediate impact on the cyclical state of the economy than most other measures. Changing the monetary lever of choice will have some kind of effect, but for this to be a successful outcome, in the sense of being able to bring about a more desirable outcome than otherwise would have occurred, the authorities need to be able to both:

- Make forecasts of the underlying conditions which are reasonably accurate in a systematic way over time
- Understand with reasonable accuracy the effect of changes in policy instruments on the economy

Unless the authorities know with reasonable confidence what the state of the economy will be in, say, one year's time, it is not possible to say what action is required now in order to bring about a more desirable outcome. And unless the authorities understand the impact of their actions, it is not possible to know what should be done in order to bring about any desired outcome.

This paper considers:

- The ability of the authorities to make reasonably accurate forecasts
- The ability of the authorities to understand reasonably accurately the impact of policy changes
- How to cope with the implications of the above two points

2 Predictability and the business cycle

By scientific standards, the macroeconomic forecasting record is poor. As examples of the one-year ahead forecasting record for GDP growth, for the US economy recessions have not generally been forecast prior to their occurrence, and the recessions following the 1974 and 1981 peaks in the level of output were not recognised even as they took place [1]. Further, growth has generally been overestimated during slowdowns and recessions whilst underestimates occurred during recoveries and booms [2]. For the UK, the predictions of the Treasury over the 1971-1996 period have been at least as good as those of other forecasters, but the mean absolute annual forecast error for these one-year ahead predictions was 1.45% of GDP, compared to an actual mean absolute change of 2.10% [3]. In 13 European countries over the 1971-1995 period, the average absolute error was 1.43% of GDP, compared to the average annual change of 2.91% [4].

In general, the forecasting record exhibits a certain degree of accuracy in that the average error over time is smaller than the size of the variable being predicted. But the error is still large compared to the actual data, and most of the accurate forecasts were made when economic conditions were relatively stable [2].

The poor forecasting record of GDP growth by economists appears to be due to inherent characteristics of the data, and cannot be improved substantially no matter what economic theory or statistical technique is used to generate them. Over what is thought of as the time period of the business cycle in economics, in other words the period over which any regularity of behaviour of the growth of GDP might be postulated to exist, the genuine information content of correlations over time in the data is low.

This result is demonstrated formally in [5]. A delay matrix of time-series data on the overall rate of growth of the economy is formed, with lags spanning the period over which any regularity of behaviour is postulated by economists to exist. Methods of random matrix theory are used to analyse the correlation matrix of the delay matrix. This is done for annual data from 1871 to 1994 for 17 economies, and for post-war quarterly data for the US and the UK. The properties of the eigenstates of these correlation matrices are similar, though not identical, to those implied by random matrix theory. This suggests that the genuine information content in economic growth data is low, and that the time-series data on GDP growth is very similar to genuinely random data. Ref [6] shows similar results for the change in the rate of inflation in the UK.

Increasingly, this issue is being recognised. The May 2002 Bank of England *Inflation Report* stated that 'Nobody can predict the future evolution of the economy with absolute certainty. It is more realistic for forecasters to recognise that uncertainty when describing their projections' (p.48). This is a very welcome advance in accepting the reality of the economy, and the Bank has reacted to this by presenting their forecasts as a range rather than a point estimate. However, it raises a further question how measures of that uncertainty should be prepared and how such uncertainty interacts with policy levers.

3 Understanding the impacts of policy changes

While prediction might be difficult, this should not be confused with the ability to explain. Prediction of an outcome which is then the subject of a policy reaction will in any case produce the desire precisely to falsify the original prediction. Nonetheless, we may still be able to understand how the economy is evolving and how the policy makers decision operate upon it.

The standard instruments of macroeconomic policy - policy designed to influence the behaviour of the economy at the aggregate level - have come to be seen since the Second World War as variables such as public expenditure, taxation and interest rates. The view

that governments, or monetary authorities, can set these in order to control the course which the economy follows is still widespread.

However, despite an intensive programme of research in applied macroeconomics, spanning more than three decades, very little progress has been made in understanding the behaviour of the economy with macroeconomic models. Refs [7, 8] show that over a three year horizon and beyond, macro models of the UK do not even agree on the *sign* of the basic public expenditure multiplier.

Further by way of example, the aggregate consumption function is one of, if not the, most intensively researched in the whole of applied macro economics. But even now, there is no agreement on its specification.

Anyone who has been involved in macroeconomic forecasting on a regular basis soon learns that individual equations break down much more frequently and much more dramatically when confronted with genuine out-of-sample data than standard statistical theory suggests.

An important reason for this is that at the micro-economic, individual agent level, tastes and preferences are not fixed. Agents alter their behaviour in part according to the behaviour of other agents, in varying degrees across different markets. This is of course a hallmark of a complex system. So in general the behaviour of agents is not time-invariant, which is reflected at the system-wide level.

At any point in time, those setting policy will disagree about the potential effect of any proposed change, since the outcome will be the result of a range of differential reactions which are not well or even at all captured in standard linear multiple regression macro economic models. For example, at a recent hearing of the UK Parliament Treasury Select Committee considering the conduct of monetary policy, the experts were questioned about the role of recent movements in commodity prices in inflation. They failed to agree on whether the commodity price movement would last, whether exchange rates could adjust to equilibrate these movements or whether there would be a major impact on UK inflation.

A possible conclusion from this is that monetary policy is inoperable. However, some kind of policy decision is inevitable, since the authorities must set some kind of base borrowing rate for the banking system. So if there are limitations both on the ability to predict and on the ability to understand it becomes imperative to understand the scope of these limitations and the ranges of economic outcomes which are plausible.

4 Stress testing under uncertainty

The Bank of England has made a start in this area in its introduction of fan charts around the published forecasts for real GDP growth and inflation. These represent the Monetary Policy Committee's view of the probability of observing outcomes which differ from those of the central forecasts. Specifically, they represent the range within which the Committee expect the outcome to take place 90 per cent of the time. They are based on judgmental scenarios but have been shown to represent the experience of the relevant variables over the most recent ten year moving window of data¹.

This is a welcome step in recognising the inherent uncertainty which surrounds forecasts of the macro economy. But it raises the question of how to measure this uncertainty and what shape the fan charts should take.

The fan charts reflect an implied probability density function of the outcomes which is based on the judgement of Monetary Policy Committee members. In this section of the paper we compare the empirical evidence with the judgmental views of the Monetary Policy Committee.

We do this by analysing post-war data on growth and inflation in the UK economy, and use a variety of standard techniques to calibrate empirical probability density functions. We conduct a similar exercise for Australian data. Specifically, we do this by a) linear interpolation b) kernel density estimation [9] c) calibrating around a hypothesised

¹ K Wallis, An Assessment of Bank of England and National Institute Inflation Uncertainties, National Institute Economic Review, July 2004

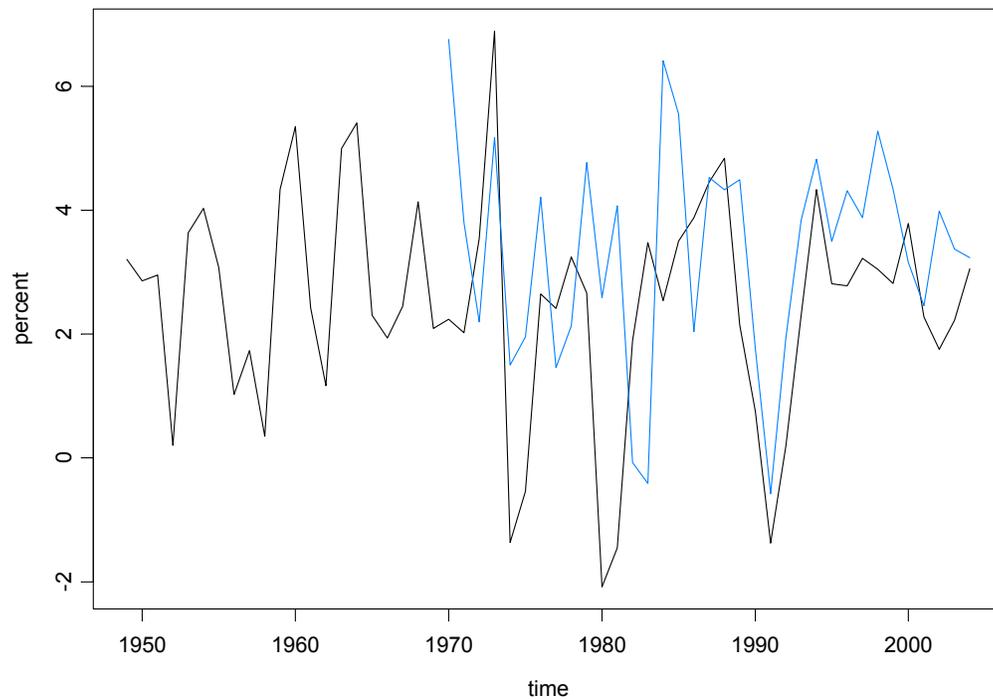
distribution, using the general Kolmogorov-Smirnov test to test the null hypothesis that the relevant data follows a hypothesised distribution.

4.1 The data

GDP

Data on real GDP in the UK is available online from the Office for National Statistics on an annual basis from 1948 to 2004². The rate of growth of this variable is plotted over the 1949-2004 period. Data for Australia comes from the Australian Bureau of Statistics, for the period 1970-2004.

Figure 1: *Annual GDP growth UK (black) Australia (blue)*



² <http://www.statistics.gov.uk/statbase/tsdintro.asp>

An obvious first step with the data in Figure 1 might seem to be to obtain estimates of the 5 and 95 percentile points in its probability distribution. These would give us the 90 percentage point spread. There are several options available as to how to compute these, and they are discussed in the Appendix.

However, there is, as always with economic time series data, a question of what is the relevant part of the sample of data 1949-2004 to use. The year 1973, for example, is often used to mark a watershed in post-war economic history, with the whole post-war period being divided into the period up to and including 1973, and the period afterwards³. For Australia we only have the period after this.

For the UK, we therefore split the data period into two, from 1949-1973 and 1974-2004, and carried out two statistical tests. First, a test of the null hypothesis that the mean of the growth rate was the same in the two periods and, second, a test of the null hypothesis that the variance of the growth rate was the same in the two periods⁴. The null hypothesis of equality of the means was rejected at a p value of 0.081 and of the variances at $p = 0.54$. In other words, the null hypotheses were not rejected at the conventional level of significance $p = 0.05$. The null hypothesis that the two distributions are the same more generally is rejected on a Kolmogorov-Smirnov test only at $p = 0.68$.

We then advanced the first period year-by-year, comparing 1949-1974 with 1975-2004, and so on up to 1949-1998 and 1999-2004. The null hypothesis of equality of means was never rejected at $p = 0.05$. The null hypothesis of equality of the variances was never rejected until the periods 1949-1991 and 1992-2004. During the latter short period of not much more than a decade, there have of course been no economic recessions in the UK, the last being in 1991.

³ see, for example, A.Maddison, *Monitoring the World Economy 1820-1992*, 1995, OECD, Paris

⁴ these were carried out using the commands 't.test' and 'var.test' in the statistical package S-Plus. In general, with t.test, the option that the variance in the two periods was not necessarily equal was used. Snedecor, G. W. and Cochran, W. G., *Statistical Methods*, 7th ed., 1980, Ames, Iowa: Iowa State University Press.

An additional point to consider is the autocorrelation function of real GDP growth. If there is strong autocorrelation, this would need to be taken into account in assessing the probability of any particular outcome in the future, given the recent past history of growth. The empirical autocorrelation function at lag 1 has a value of 0.281, and the null hypothesis that this is zero is rejected at $p = 0.036$. However, at no other lag of the function is the null hypothesis that the coefficient is zero rejected at $p = 0.05$, computing the function from 1 to 15 lags. There is therefore some weak, positive autocorrelation at lag 1, but the effect is decidedly weak. Further, the initial estimate for GDP growth in 2004 is just over 3.0 per cent, not too dissimilar from the mean over the entire period of 2.5 per cent.

For the Australia data the autocorrelation function at lag 1 has a value of 0.103 and the null hypothesis that this is zero is not rejected ($p = 0.539$). Only at lag 7 do we find that the null hypothesis that the coefficient is zero is rejected, however since no other lag shows significant autocorrelation this is ignored.

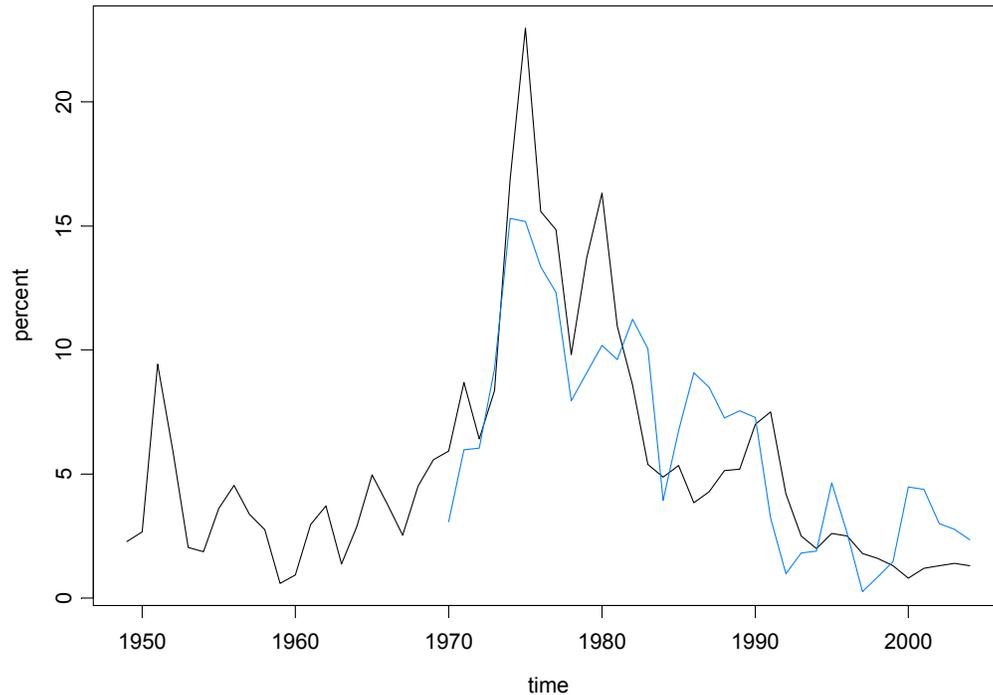
Inflation

The inflation variable which is targeted by the MPC has changed over the course of the MPC's existence. From being RPIX – prices excluding mortgage costs – it had become the Harmonised Index, now known as CPI. This excludes all forms of housing costs since the EU cannot agree how to include them. The Office for National Statistics has data on this particular variable only from 1989. For the period 1949-1988, we used the rate of change of the overall consumer expenditure deflator.⁵ For Australia, we use the CPI Index published by the ABS.

Inflation presents us with more issues about the choice of sample period over which to calibrate a probability density function than does GDP growth. Figure 2 plots the rate of inflation.

⁵ data for 1963-1988 is available on the Office for National Statistics website, and the deflator for 1948-1962 was taken from C.H. Feinstein *National Income, Output and Expenditure in the United Kingdom 1855-1965*, CUP, 1972

Figure 2: *Consumer Price Inflation UK(black) Australia (blue)*



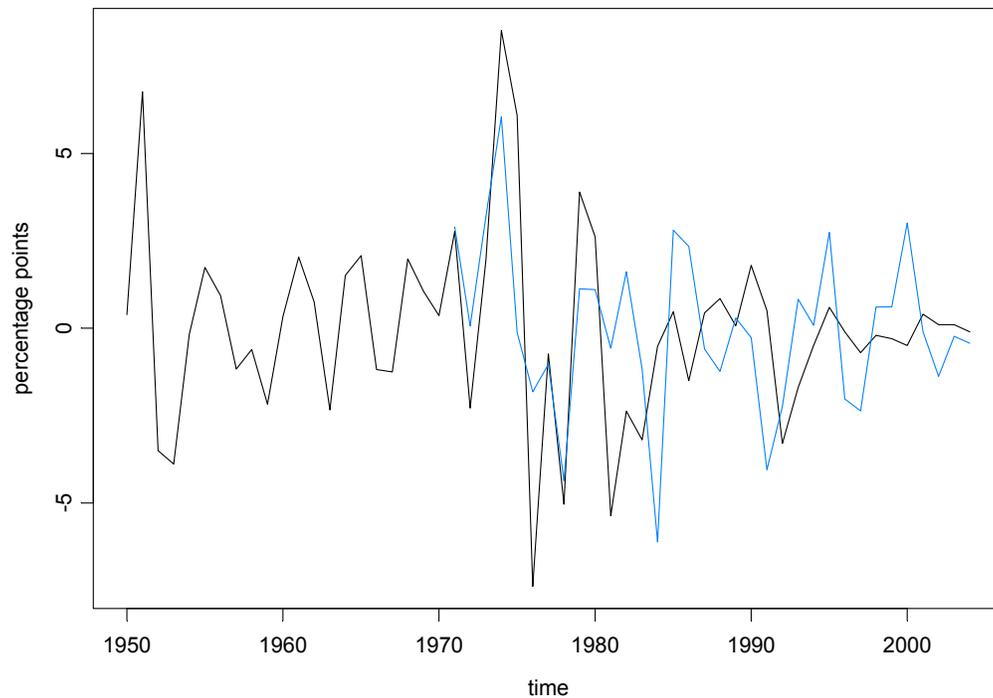
An immediate point is that there is very strong positive autocorrelation of the inflation rate in both series. This is entirely to be expected, because changes in costs today feed through into price and wage changes in the near future, introducing substantial inertia into the process of inflation. At lag 1, for example, the value of the autocorrelation function is 0.83 for the UK and 0.82 in Australia.

However, this question can be dealt with in a straightforward way. The simple difference of the inflation rate gives an autocorrelation function which is effectively zero from lags 1 through 15, with the null hypothesis that the coefficient is zero only being rejected for lag 2, where the estimated coefficient is -0.29 , and the null is rejected at $p = 0.032$ (UK).

For the Australian data the simple difference of the inflation rate gives an autocorrelation function which, as in the UK, is effectively zero from lags 1 through 15, with the null hypothesis that the coefficient is zero never being rejected.

Much more problematic is the choice of sample period. Figure 3 plots the difference in the annual inflation rate 1950-2004.

Figure 3: *Difference in Annual Consumer Price Inflation in Percentage points UK(black) and Australia (blue)*



We can identify three years in which there was an external shock which had the effect of increasing inflation. This is not to say that the external shock was the only factor operating on inflation in these years, but on each occasion it was important. First, the Korean War in 1951 in the UK data; second, the first oil shock in 1974 and third the second oil shock in 1979/80 in both countries. In each case, the increase in inflation triggered responses in the economy, in part by deliberate policy and in part by the endogenous movement of variables, which led to the sharp reductions shown in Figure 3 which followed the initial upward shocks.

The variance of the change in inflation is dramatically different depending upon whether the whole sample period is chosen, or whether the periods of inflationary shock and their aftermath (1951-53, 1974-83) are excluded.

For the UK we err on the conservative side and restrict our focus to the two more stable periods of historic inflation, namely 1954-1973 and 1984-2004. For the purpose of deriving confidence levels around inflation forecasts this restriction narrows the range of possible bands. The Australian calculations use the whole period.

The null hypothesis that the means of the change in inflation are the same in the periods 1954-73 and 1984-2004 is only rejected at $p = 0.24$, and the null hypothesis that the variances are the same is rejected at $p = 0.06$, so it is legitimate to combine these two periods into a single sample for the UK.

4.3 Results

The fan charts in Charts 1 and 2 of the *Inflation Report* show ranges of GDP growth and inflation over the forecast period. The outturns are expected to lie within these bands with a probability of 90 per cent. We do not have the exact numbers which correspond to the bands, but reading off the published charts, we can see that for 2005 the approximate ranges for the variables within the 90 per cent probability bands are:

GDP growth	1.5 to 3.5 per cent
Inflation	1.0 to 2.0 per cent

GDP

Consider first of all the growth rate of real GDP. In the Appendix, we note three potential ways of estimating a probability density function: linear interpolation, kernel density estimation, and using a hypothesised distribution. The latter can only be used if the relevant data can be shown to be compatible with a known statistical distribution.

Figure 4 plots a histogram of real UK GDP growth 1949-2004.

Figure 4: *Histogram of annual UK GDP growth, 1949-2004*

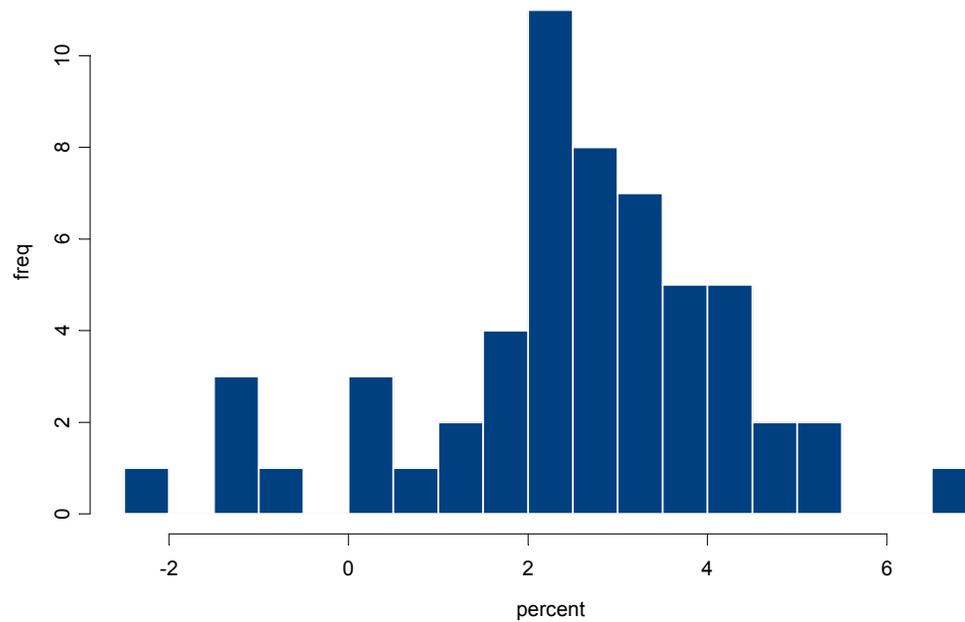
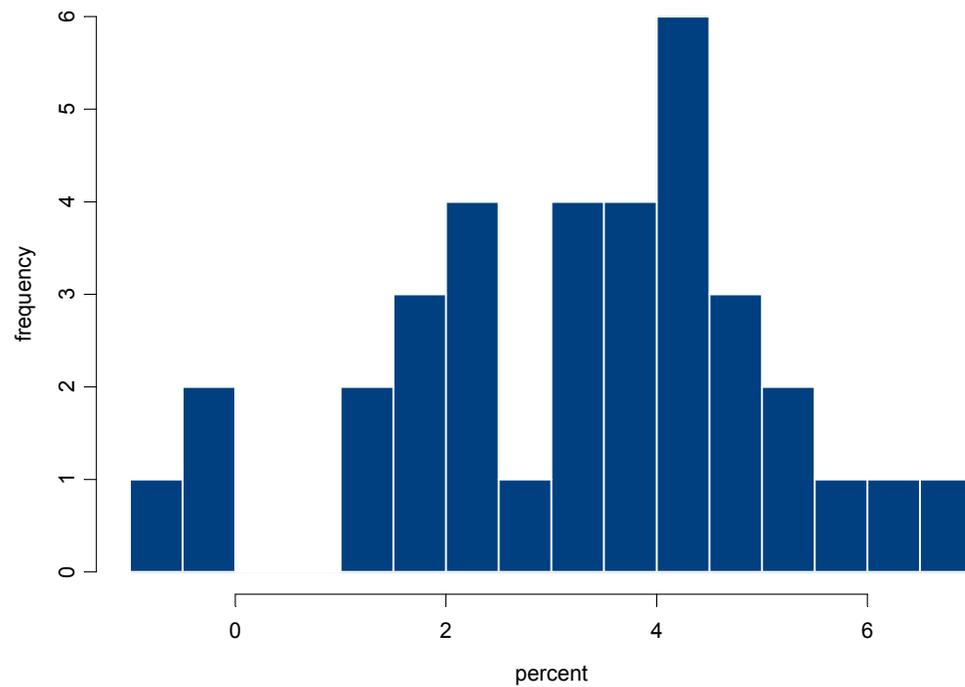


Figure 5: *Histogram of Australia GDP growth*



The data appears to follow as a broad approximation a normal distribution, but somewhat skewed to the left. In other words, there are more years with weaker growth than would be the case if the data were genuinely distributed normally. This is confirmed

by a Kolmogorov-Smirnov test. The null hypothesis of normality is rejected at $p = 0.013$.

We therefore use linear interpolation and kernel density estimation. Of the two, the latter should be regarded as being in general the more accurate.

Table 1: Range for real annual GDP growth rate (per cent) in the UK in 2005: 90 per cent range

	Low	High	Probability of growth < 0
Monetary Policy Committee judgement	1.5	3.5	$\cong 0$
Linear interpolation of probability density function of data 1949-2004	-1.4	5.1	0.1
Kernel density estimation of pdf 1949-2004	-1.3	5.4	0.11
Australian data	-0.2	6.2	0.06

The statistically based estimates of the 90 per cent range of the data are much higher than the judgmentally based ones of the MPC. It may be thought implausible, for example, that a recession could take place in 2005. But anticipating turning points, and recessions in particular, is well known to be the weakest aspect of macroeconomic forecasting⁶. In recession years, forecasts made early in the year have often projected positive growth.

In Australia, the likelihood of recession is smaller than in the UK, and the upper band is also higher. The probability of growth being less than zero is only 6 % in Australia, compared to 11% in the UK.

⁶ see, for example, H.Stekler and R.Fildes, (1999), 'The state of macroeconomic forecasting', George Washington University, Center for Economic Research Discussion Paper 99-04; V.Zarnowitz and P.Braun, (1992), 'Twenty-two years of the NBER-ASA Quarterly Outlook Surveys: aspects and comparisons of forecasting performance', NBER Working Paper 3965

The fan chart drawn up by the MPC extends to the year 2008. The bands widen as the time horizon is extended, indicating increasing uncertainty about outcomes. It must be said that whilst this has a certain initial plausibility, it is not clear that the actual macroeconomic forecasting record bears this out. We are not aware of any systematic, published comparison of forecasts more than one year out but our impression of two of experienced forecasters, with over 50 years practical experience of macro forecasting between them, is that the forecast errors do not in general increase much when the horizon is extended beyond one year. This is because the one year ahead forecasts themselves are often so inaccurate.

The MPC estimates that the 90 per cent band in 2008 is between 1.5 and 5.0 per cent. The upper end is consistent with the evidence of the post-war period. But the lower end remains very much higher. Indeed, a figure of +1.5 per cent growth corresponds to the 24th percentile, using the kernel density estimation technique. So the range for 2008 put forward by the MPC is of the order of 66 percentage points and not 90.

Inflation

In terms of inflation, the current rate is around 1.5 per cent, which the MPC sees as persisting during 2005, within a 90 per cent probability band of 1 to 2 per cent.

Table 2 sets out the various estimates for the 90 per cent probability band for 2005 using the data on the change in inflation. With the inflation data we cannot reject the hypothesis of the data being normally distributed, with a Kolmogorov –Smirnov normality test returning a p-value of 0.5.

Table 2: Range for inflation (per cent) in the UK in 2005 : 90 per cent range

	Low	High	Probability of inflation < 0
Monetary Policy Committee judgement	1.0	3.0	$\cong 0$
Linear interpolation of probability density function of data	-0.8	3.5	0.14
Kernel density estimation of data	-1.0	3.8	0.13
Hypothesised distribution of data (normal distribution)	-0.7	3.8	0.12
Australian data	-2.2	6.4	0.16

(note that since Inflation is highly autocorrelated all calculations are carried out on the difference of inflation, i.e this is actually 2005 predictions calculated using 2004 value of inflation)

It may at first sight seem surprising that negative inflation is possible given a history in which this has yet to occur. The current level of inflation is, however, very low, and thus a further downward shock could take it below zero. The chances are not high, being between 12 and 14 per cent depending upon which method of estimation is used. But the MPC is effectively ignoring this possibility completely.

Yet in fairly substantial sectors of the economy, negative inflation has been a feature for several years. Manufacturing producer prices, for example, fell in every year between 1996 and 2002. The annual rate of increase in average earnings in the construction sector, after having been around 5 per cent for several years, has fallen to zero because of an increased labour supply from the new EU countries.

In Australia, while the probability of inflation falling below zero is only slightly larger than in the UK, the range is larger at the 90% level. This is largely because the whole data set has been used, while some of the peaks have been excluded from the UK data set.

Explaining the differences

The analysis of economic history, both for GDP growth and inflation provides much wider bands for both variables than recent experience suggests for either country. It is also wider than the judgmental bands drawn up by members of the MPC. Moreover, some commentators have suggested that the MPC itself has overestimated the range of potential outcomes⁷.

One possibility is that history has ceased to be relevant. There may have been a decisive shift of behaviour in the economy which means that most of the post-war data is irrelevant to the future. In the UK the creation of the MPC may or may not be one of the factors behind such a hypothesised shift of behaviour. It is simply not possible to provide decisive evidence on this, not least since this argument means that there are very few data points in the 'new' regime.

If we take only the last ten years data the picture is indeed different.

Table 3 Range for real annual GDP growth rate (per cent) in 2005: 90 per cent range

	Low	High	Probability of growth < 0
Australian data	2.5	5.3	≈0
UK data	1.7	4.3	≈0

Table 4 Range for inflation (per cent) in 2005 : 90 per cent range

	Low	High	Probability of inflation < 0
Australian data	-0.2	5.6	0.08
UK data	0.6	1.9	≈0

⁷ K Wallis, Chi-squared Tests of Interval and Density Forecasts, and the Bank of England's Fan Charts., warwick university

But the onus here is on those who believe there has been a decisive shift to attempt to demonstrate it, particularly given the fact that over much of the post-war period the statistical distribution of GDP growth has been stable with respect to the choice of most possible sample sub-periods, and that the same is true of the change in inflation except for the periods following the major external shocks of the Korean War and the two oil price shocks.

For what it is worth, the opinion of the present author is that the supply side reforms of the 1980s did raise the sustainable long-term growth rate of the economy. But the differences between the probability bands for the last ten years and the bands estimated from economic history are so substantial that it is implausible that the performance of the economy has improved to such a dramatic extent.

Finally, we must return to the difference between forecasting and understanding. The 90 per cent bands drawn up by the UK's MPC are prepared on the basis of market interest rate expectations being fulfilled. These are of course contingent on whatever the market's current view of policy is. Forecasts contingent on a particular policy stance do not tell us what the policy response might be. Relaxing this assumption would undoubtedly widen the band, though it might make them less useful as a guide to policy.

5 Conclusion

We know that there is uncertainty both about the future and the policy levers by which any future may be managed. This paper has looked at the range of these uncertainties both over the longer term and a shorter one. We compare this with the range published by one of the monetary authorities – the UK.

The Monetary Policy Committee in the UK takes into account the uncertainty around any given macroeconomic forecast for GDP growth and inflation. The Committee does this by forming a judgement about the range within which they expect the outcome to occur 90 per cent of the time.

We compare the judgmental bands with the evidence from post-war data on growth and inflation in the UK. We estimate empirically the 90 per cent probability ranges for both

growth and the change in inflation. There are several technical ways of doing this and the results are very robust with respect to the techniques used.

There are issues concerning the choice of sample period of the estimates, particularly for inflation. The estimates are very robust with respect to choice of sample period for GDP growth. In our estimates for inflation, we exclude the inflationary shocks of the Korean War and the two oil shocks, and their immediate aftermaths. The inclusion of these periods in the estimates would widen even further the 90 per cent probability range.

We find that the empirically estimated ranges are considerably wider than the judgmental ones of the MPC. For growth, the MPC's 90 per cent probability range for 2005 is between 1.5 and 3.5 per cent, whereas the empirically based estimates are between -1 and +5 per cent. For inflation, the MPC's range is +1 to +3 per cent, compared to around -1 and +3.5 per cent.

This evidence is consistent with the actual macroeconomic forecasting record. The level of accuracy even of one-year ahead forecasts is low, but it is particularly bad at turning points. Forecasters seem to find it especially hard to anticipate changes in recent trends.

The results do not of course mean that either growth or inflation will be below zero, but that the 90 per cent probability range encompasses values below zero for both these variables. The MPC appear to be assigning far too low a probability to the chance of either an economic recession or deflation, when inflation becomes negative. We estimate that the probability of growth being below zero is approximately 11 per cent (6 per cent in Australia), and of inflation becoming negative of 13 per cent (16 per cent in Australia).

The risk here is of hubris. History has meaning and cannot be ignored. Policy makers have to make a judgement about the economic regime which is relevant to the times they live in. They also form that regime and are part of it. It is impossible for them to grasp the dimensions of the problem since they (and we) are inside it.

1. H.Stekler and R.Fildes, (1999), 'The state of macroeconomic forecasting', George Washington University, Center for Economic Research Discussion Paper99-04
2. V.Zarnowitz and P.Braun, (1992), 'Twenty-two years of the NBER-ASA Quarterly Outlook Surveys: aspects and comparisons of forecasting performance', NBER Working Paper 3965
3. C.Mellis and R.Whittaker, (1998), 'The Treasury forecasting record: some new results', *National Institute Economic Review*, 164, pp.65-79
4. L-E.Öller and B.Barot, (1999), 'Comparing the accuracy of European GDP forecasts', National Institute of Economic Research, Stockholm, Sweden
5. P.Ormerod and C.Mounfield, (2000),'Random Matrix Theory and the Failure of Macro-economic Forecasting', *Physica A*, 280, 497-504
6. P.Ormerod, (2001), *The Illusion of Control: Why Membership of the Monetary Policy Committee is a Sinecure*, Manchester Statistical Society
7. K.B.Church, P.R.Mitchell, P.N.Smith and K.Wallis (1993), 'Comparative Properties of Models of the UK Economy', *National Institute Economic Review*,145
8. J.Bray, S.Hall, A.Kuleshov, J.Nixon and P.Westaway (1995), 'The Interfaces between Policy Makers, Markets and Modellers', *Economic Journal*, vol.105, pp.989-1000
9. B Silverman, (1996), *Density Estimation for Statistics and Data Analysis*, Chapman and Hall