

What Is the Monetary Policy Committee Attempting to Achieve?

*Charles Goodhart**

Introduction

In broad terms, the answer to this question is obvious. The Chancellor of the Exchequer initially required the Monetary Policy Committee (MPC) to maintain the level of a version of the Retail Price Index (RPIX) at 2.5 per cent at all times over an indefinite horizon, and thereafter continuously reaffirmed this objective (Brown 2003) until June 2003. Then, as a facet of the study into U.K. membership of the single [euro] currency (HM Treasury 2003), the Chancellor announced that he would switch the MPC's target to the Harmonised Index of Consumer Prices (HICP), with a symmetric objective of 2 per cent, starting later in 2003. Although the MPC has steadily sought to achieve its broad target, there do remain some narrower questions of detail that this paper aims to answer empirically.

First, unforeseen shocks will drive current inflation away from target inflation. How fast does the MPC want to return inflation to target? This will be a function of the inertia in the system, and of the current extent of deviation from target. For example, if current inflation was 15 per cent and the target was 2.5 per cent, it would not, in most people's views, be sensible to aim to return inflation to target within, say two years, even if the lags were such that current interest rate changes had their maximum effect on inflation at a two-year, eight-quarter horizon. Instead, the authorities should, in principle,

* My thanks are due to A. Orphanides, D. Thornton, and the participants at this Festschrift, and to M. Emond for research assistance. All errors and idiosyncracies, however, remain my own.

set out a multi-year declining path for inflation, as has indeed been tried in several transition countries. In more technical terms (Bean 2003), the path should be set so that the (desired) marginal rate of substitution (MRS) between output and inflation equalled their marginal rate of transformation.

In practice, however, the MPC was established in relatively benign economic conditions. The rate of inflation of RPIX was already close to target at the outset in 1997, and has remained so since then (Figure 1). Moreover, the rate of growth of output has also remained stable, on any historical comparison (Figure 2); and unemployment has trended gently downwards (Figure 3). In such propitious conditions, the MPC has felt able to achieve its broad objectives by comparatively small (on a historical basis) changes in interest rates (Figure 4).

The MPC publishes a forecast of its projections for inflation (RPIX) and output over the next two years (actually nine quarters in all, since the current quarter also needs to be forecast). It can adjust its interest rates before publishing the forecast so as to drive inflation into line with its objective at some future horizon (n quarters hence, where n has to be determined). Svensson has argued (1997; also see Svensson and Woodford 1999) that this general procedure, whereby the authorities drive their own inflation *forecast* into line with the target, is optimal (although he has some objections to the particular way that the MPC does this, on which more later).

The hypothesis examined here is that, at a horizon n quarters from the forecast date, two conditions will hold; first, the level of the inflation forecast n quarters ahead *averaged over all forecasts* (from May 1997 to November 2002) will be insignificantly different from 2.5 per cent, and second, that at that horizon the standard deviation of the successive *forecasts* (not of the outcomes) will be at a minimum. Since inflation has, in practice, been close to 2.5 per cent throughout, the average forecast of inflation at $n + x$, $n - x$, quarters ahead may also hardly differ from 2.5 per cent, but, if the MPC is seeking to drive the inflation forecast to 2.5 per cent n quarters hence, the variance of the inflation forecast at $n + x$, $n - x$ should be greater. Thus, by examining the mean and standard deviations of the forecasts (not the outcome) from one to eight quarters ahead, we might be able to pin down the MPC's effective horizon.

In the sticky-price neo-Keynesian forecasting models that the MPC has used (Bank of England 1999; 2000), interest rates affect output before they have their main effect on inflation (there can be a partial immediate effect on inflation via the exchange rate). Hence, if the inflation horizon is n quarters, the horizon for output should be shorter at q quarters, $n - q > 0$. With both inflation and output having behaved benignly, the proposition is that the average forecast level of the rate of growth of output, q quarters hence, will

Figure 1
Annual RPIX

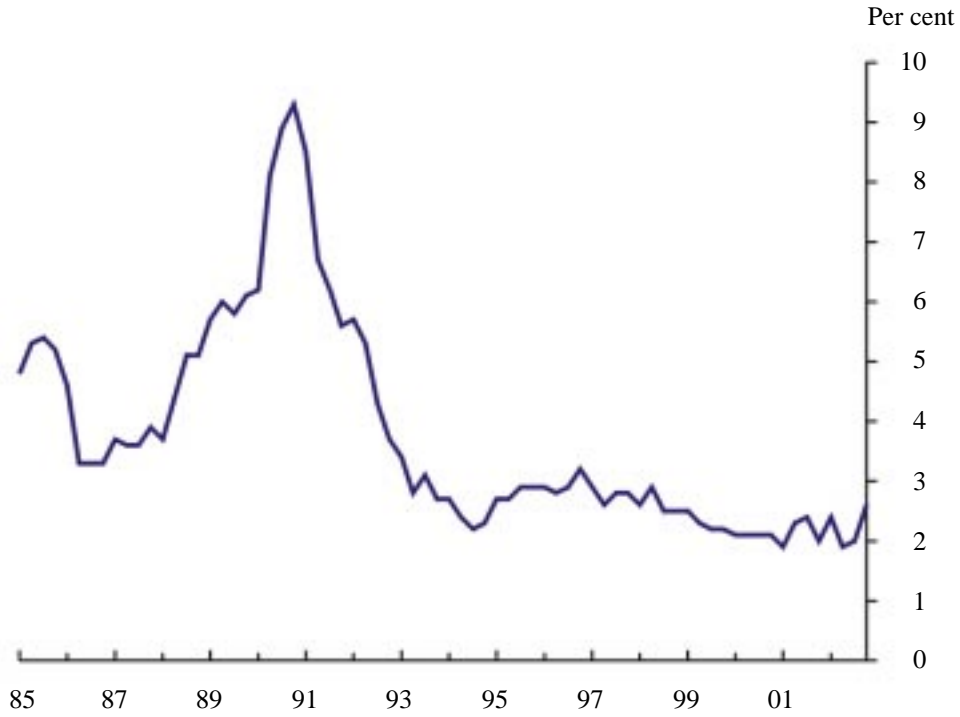


Figure 2
Real GDP growth

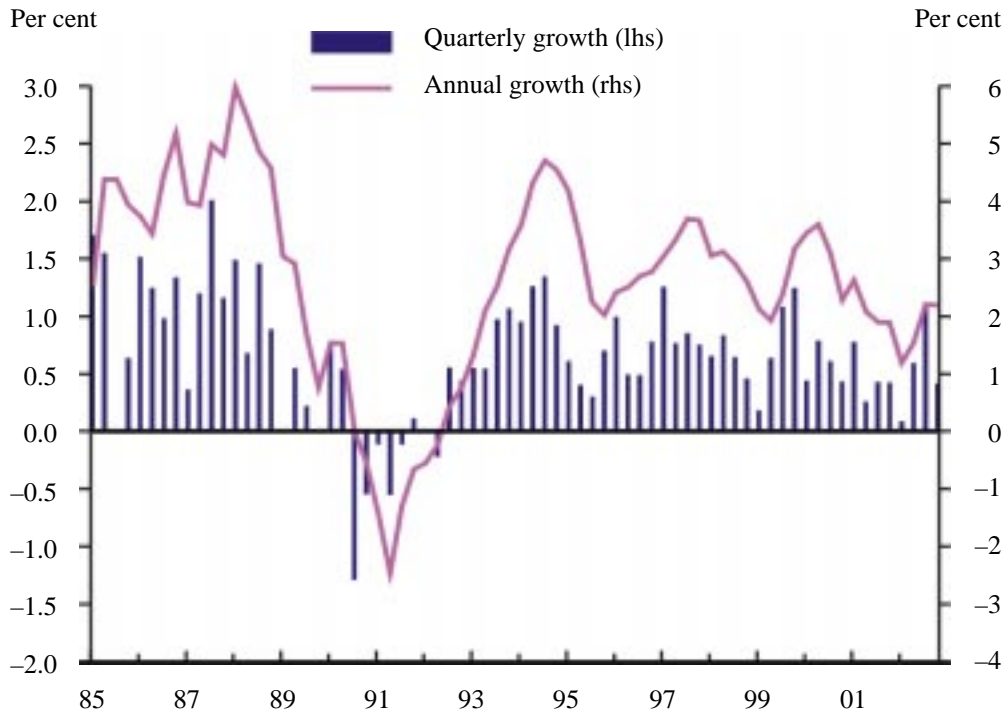


Figure 3
Claimant count unemployment rate

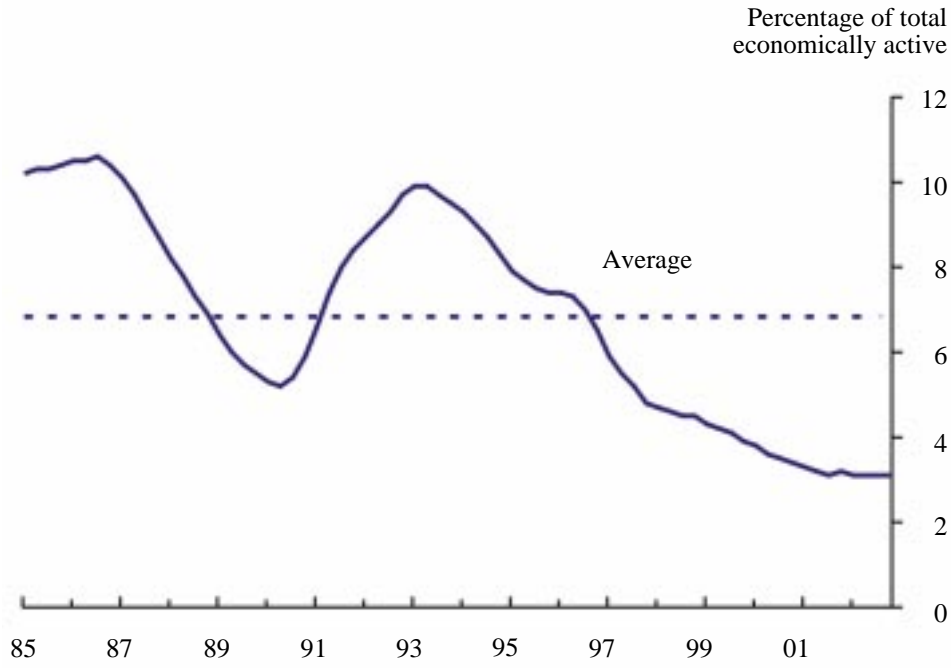
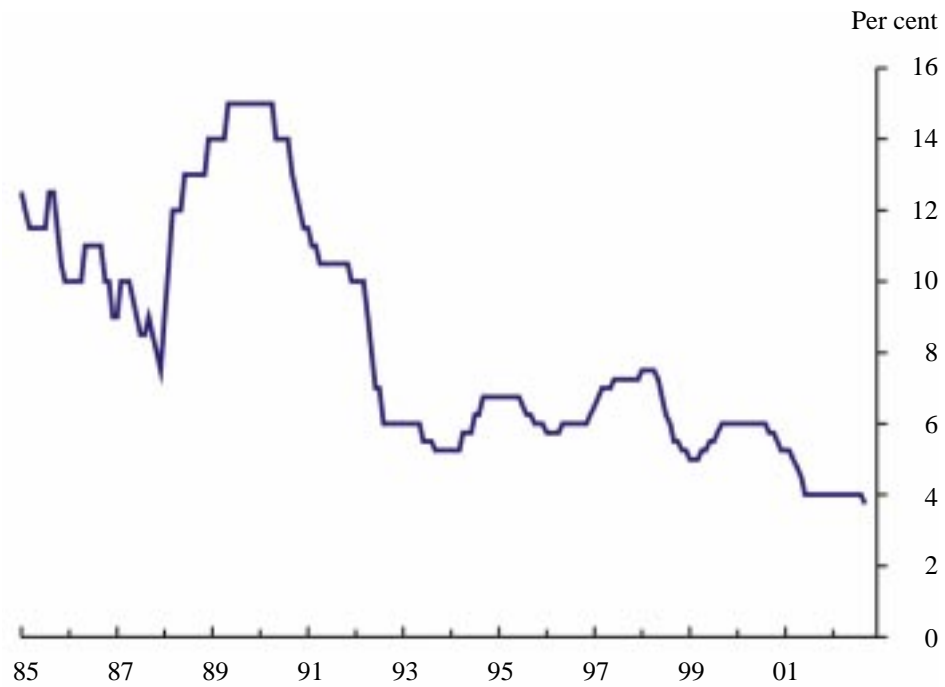


Figure 4
Official short-term interest rate



approximately represent the MPC's estimate of the underlying, sustainable trend rate of growth. The hypothesis is that the standard deviation (SD) of the forecast rate of growth will reach a minimum q quarters hence, where q is less than n . As will be seen later, however, this hypothesis did not fare well. Note that the MPC publishes its forecast for the rate of growth of output, but not for the output gap; although the latter measure is even more judgmental, many commentators regard it as more critical for informing the interest rate decisions.

In the MPC's earlier booklet, "The Transmission Mechanism of Monetary Policy" (1999), it was stated that simulation of the medium-term model suggested that the greatest effect of interest rates on output was to be found at a horizon of about one year, and on inflation at a two-year horizon. To be consistent with that, we would expect q , the local minimum in the SD of output forecasts, to be four quarters hence, and n to be eight quarters hence, with the SD of the inflation forecast falling continuously as the horizon lengthened to its maximum two-year length. As will be seen in section 2, these predictions were correct in the case of inflation, but not for output growth.

The relativity between the SDs of the forecasts for output and inflation is also important. If the MPC's overriding objective is to drive the n quarter forecast for inflation close to 2.5 per cent, then the SD of that forecast will be low. By contrast, trying to adjust output growth q quarters hence, so as to enable RPIX to be at 2.5 per cent $n - q$ quarters afterwards, will depend on the initial conditions for current inflation and output. Hence, the SD for output growth would be expected to be much greater (note that the mean forecast for the output trend, at 2 per cent +, is of much the same numerical size as the inflation target of 2.5 per cent).¹

This hypothesis, that the variance of the forecast for the rate of growth of output q quarter hence is far greater than the variance of the forecast for the rate of inflation n quarters hence, i.e., $SDdy_t + q \gg SDdPt + n$, is strictly different from the zone-quadratic hypothesis of Orphanides and Wieland (2000). In this, they allow (p. 1352):

for a zone-quadratic objective on the part of the policymaker, that is a loss function which assigns quadratic loss to inflation deviations outside an explicit target zone and implies a near zero loss as long as inflation is contained within the zone. As a consequence, if the policymaker assigns at least some weight

1. Again, it is important to remember that these are SDs of forecasts, not of outcomes. Hence, the possibility of short-run shocks to the outcomes for output being different in scale to those affecting inflation is irrelevant.

to output stabilization, the output objective will dominate at times when inflation is within the zone but will recede in importance when inflation is outside the zone. Such preferences would also be consistent with an apparent tendency of central banks to “put out fires”, that is to react to inflation primarily when it becomes a problem but concentrate on other objectives when inflation is under control.

This suggestion is revisited in Orphanides and Wilcox (2002), where the authors propose an “opportunistic² loss function.”

The policy-maker’s loss function is given by

$$L_A = (\pi - \pi^*)^2 + \gamma y^2 + \Psi |y| \quad (1)$$

where $\gamma \geq 0$, $\Psi \geq 0$, and π^* is an intermediate target for inflation (zero being the long-run target). Importantly, the central bank’s loss function depends on the deviation of inflation from the intermediate target, not the long-run target (p. 52), and

The essential point to be noted is that the policy-maker suffers first-order losses from output gaps; this is reflected in the fact that . . . the marginal loss exhibits a discrete jump at the bliss level of 0 and is everywhere bounded away from 0 (pp. 53–4).

Also see Terlizzese (2002).

But if this was so in the case of the MPC, they should be less concerned about the forecast level of inflation, both at n and at any other horizon, so long as it remained between the bounds of 1.5 and 3.5 per cent (these being the limits that, if broken, require the MPC to write a letter of explanation to the Chancellor). By contrast, the MPC should seek to minimize the output gap q quarters ahead. Given that inflation has remained within the target zone since 1997, the implied hypothesis of these latter papers should be that the variance of the forecast for the rate of growth of output q quarters hence should be considerably lower than the variance of the forecast for inflation n quarters from the forecasting data, i.e., $SDdy_t + q \ll SDdPt + n$. As will be seen in section 2, the Orphanides-Wieland model does not hold for the MPC in the United Kingdom (though it has not been similarly tested in other countries).

2. “‘Opportunistic’ is the word used to describe a policy stance that some have argued was maintained by the Fed at times since 1979” Orphanides and Wilcox (2002, 48).

So far, we have used the phrase, “the forecast” as if there was a single value for that forecast. When the projected distribution of outcomes is symmetric, that phraseology is, perhaps, pardonable, since all measures of central tendency are identical: mode, median, and mean. Even though the projected *variance* of the distribution of outcomes, a measure of uncertainty, will vary from forecast to forecast, under most assessments of behaviour, certainty equivalence will hold in cases of symmetric distributions (see, for example, al-Nowaihi and Stracca 2003; Bray and Goodhart 2002; Chadha and Schellekens 1999).

In practice, asymmetric risk is often present, and that will drive wedges between the forecast average mode and median, and median and mean. One aim of this paper was to use such wedges, at the estimated horizon of n quarters, to estimate directly the utility function of the MPC. Assume that the utility function of the MPC is a function of $|(\pi - 2.5)|^i$, where the function is symmetric and i is an exponent to be estimated. If the MPC has a quadratic loss function, $i = 2$. As will be shown later, the average risk over this period was an upward skew on inflation. So, if $i = 2$, the mean forecast for inflation should equal 2.5 per cent, with the median and, even more so, the mode below 2.5 per cent. With an exponent of one, so that perceived losses are linear in the extent of deviations, we should find the median forecast equal to 2.5 per cent, with the mean above and the mode below. If the MPC felt that any deviation was unacceptable (i.e., an exponent near zero), with the loss function being strictly concave (so that any added deviation had a minimal effect on utility), then the mode would equal 2.5 per cent, with the median and, even more so, the mean lying above it. Kahneman and Tversky (2000) have studied human behaviour in this respect and have argued that their tests indicate an exponent slightly below 1, which would suggest that in the MPC’s forecasts, the mode should be somewhat below 2.5 per cent and the median slightly above it, n quarters hence.

In practice, the average extent of skew in the inflation forecast has been slight, so the average values of the three measures of central tendency (n quarters hence) are quite close together, and so the power (significance) of the test is not high. Even so, the result shown in section 2 was surprising (to me), which is that the mode was marginally above 2.5 per cent with the median and mean even higher.

We discuss a number of possible reasons for this finding, though we cannot easily distinguish between them. One possible line of argument, which I choose to preview here, is that the MPC’s forecasts are conditioned on certain prior assumptions. One key assumption is that, after the current change in interest rates, if any, interest rates are then assumed to be held constant for the next n quarters (Goodhart 2001). The MPC could then take

the view that, should the perceived upside risk to inflation actually become apparent, they would still have time subsequently to raise interest rates to head that off. In other words, examining the wedges between central tendencies as a measure of the utility function of the MPC is, in practice, a joint test both of the key conditioning assumptions of the forecast and of the utility function. While I do not see how to unscramble these two elements, the exercise is not, I hope, without interest.

Section 1 briefly sets out the source of the data; section 2 describes the results; and the final section, again briefly, concludes.

1 The Data

The Inflation Forecast is purposefully published without any reference to point estimates for any measure of central tendency. Instead, the focus is on the fan charts for inflation (graphically known as “the rivers of blood” for its combination of delta shape and red colouring) and output. The aim is to emphasize probabilities and uncertainty, and to reduce the concentration on point estimates, which are bound in reality to be missed.

Nevertheless, measures of such central tendencies exist behind the charts. After a decorous interval, they are published on the Bank’s Web site, <www.bankofengland.co.uk>. This is what has been used here. An example, for the forecast for RPIX made in May 2000 (my own last appearances on the MPC), for the current and next eight quarters is given in Table 1.

Table 1
May 2000 forecast

Inflation									
	Current	<i>t</i> + 1	<i>t</i> + 2	<i>t</i> + 3	<i>t</i> + 4	<i>t</i> + 5	<i>t</i> + 6	<i>t</i> + 7	<i>t</i> + 8
Mode	1.88	1.93	2.10	2.20	2.47	2.53	2.53	2.53	2.56
Median	1.89	1.95	2.13	2.24	2.51	2.55	2.53	2.51	2.52
Mean	1.89	1.96	2.14	2.25	2.52	2.56	2.53	2.50	2.51
Output growth									
	Current	<i>t</i> + 1	<i>t</i> + 2	<i>t</i> + 3	<i>t</i> + 4	<i>t</i> + 5	<i>t</i> + 6	<i>t</i> + 7	<i>t</i> + 8
Mode	2.94	2.58	2.45	2.57	2.61	2.65	2.70	2.74	2.70
Median	2.92	2.54	2.39	2.50	2.53	2.53	2.51	2.50	2.46
Mean	2.91	2.52	2.37	2.48	2.51	2.49	2.46	2.44	2.40

Similar forecast estimates from May 1997 to November 2002, for RPIX and for annual growth rates for GDP, giving 23 forecasts in all, can be taken directly from the Bank’s Web site. Consequently, this study can be immediately replicated by anyone else.

2 Empirical Results

We start with the forecast for inflation. The average forecast for the predicted mode of inflation has shown a gentle decline over the first four quarters ($t = 1$ until $t = 4$, where $t = 0$ is the current quarter) from 2.318 at $t = 0$ to 2.262 at $t = 4$. There is then a more rapid rise in forecast inflation over the subsequent year to its high point of 2.510 at $t = 8$. This is shown in Figure 5. The most interesting feature is that forecast inflation only comes very close into line with target right at the very end of the period $t = 7$, $t = 8$. Forecasts over shorter horizons have on average been below, though not much below, barely 25 basis points, the target.

We now go on to examine the SDs of these forecasts at each horizon (that is, we look at each forecast for $t + i$ quarters ahead, and calculate the SD of the $t + i^{\text{th}}$ forecast). Here we see that the SD of the forecasts themselves is broadly flat from $t = 0$ to $t = 2$ (over the current and next two quarters), declines gently from there to $t = 5$, and then declines very sharply indeed, by almost two-thirds in the final four quarters (from $t = 5$, 0.251, to $t = 8$, 0.096). This is consistent with a story in which the authorities reckon to have little influence on short-run inflationary shocks, but get an increasing grip on inflation, as the horizon lengthens to that chosen to drive the inflation forecast into line with target. As the horizon lengthens to $t = 7$ and $t = 8$, so the modal forecasts at each forecasting round bunch much closer together. This is shown in Figure 6.

Next, we do the identical exercises for the forecast of annual GDP growth. First, we look at the mean of the modal forecasts at each forecasting horizon. This is shown in Figure 7. The pattern is rather akin to that of inflation. Initially there is a gentle decline in forecast output (on average) from 2.257, currently at $t = 0$, to a low point two quarters ahead, 2.130 at $t = 2$. Then the forecast for output increases (at a roughly constant 0.1 per cent per quarter) over the next six quarters, to a rate of 2.608 at $t = 8$.

Unlike inflation, where there was a quantitative target set, we do not know what the sustainable equilibrium rate of growth of output is, or whether it is constant, or what the MPC may have thought it to be at any time. There is some qualitative evidence that the authorities may have seen this sustainable rate as lying between 2.25 and 2.50 per cent. In the latest (2003) Budget, HM Treasury makes a (conservative) assessment that the underlying rate of improvement in productivity is 2 per cent per annum, and there is a small expected increase in the working population. I had hoped that I could throw some light on this, if there was an interior minimum for the SD of the forecast of dy at a horizon less than $t = 8$. This would have suggested that the authorities were trying to vary interest rates to stabilize the growth rate

Figure 5
Average of modal forecast of inflation

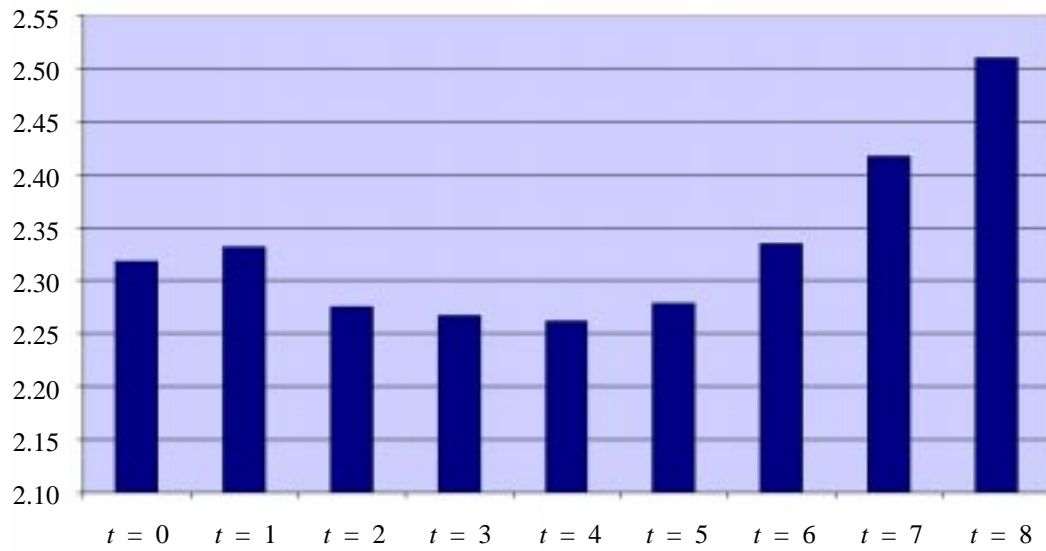


Figure 6
SD of modal forecast of inflation

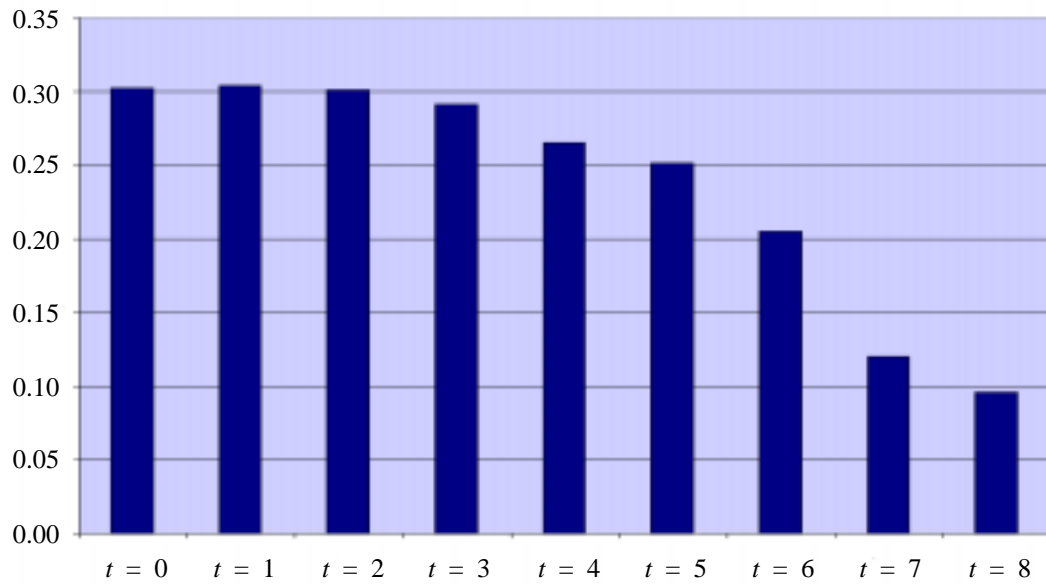
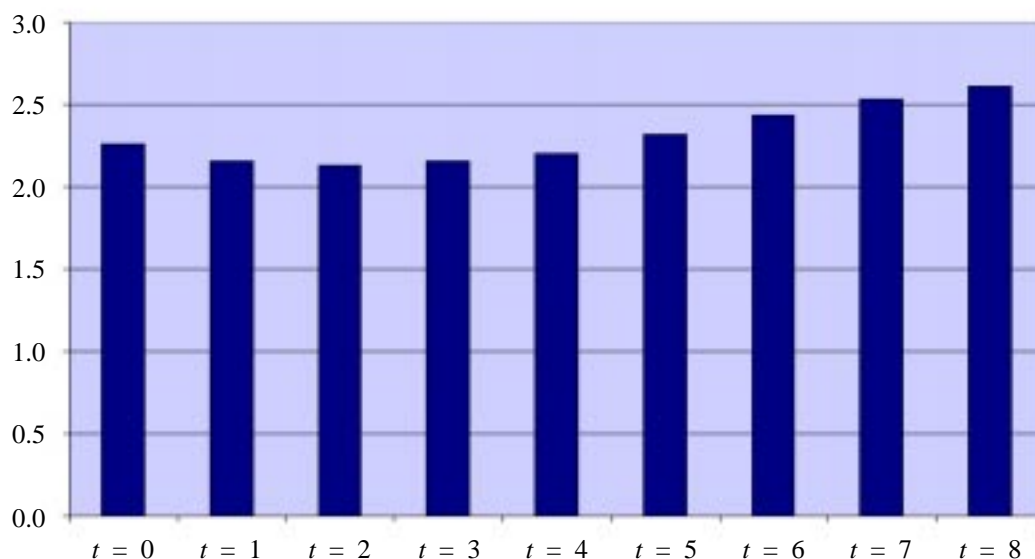


Figure 7
Mean of modal forecast of output growth



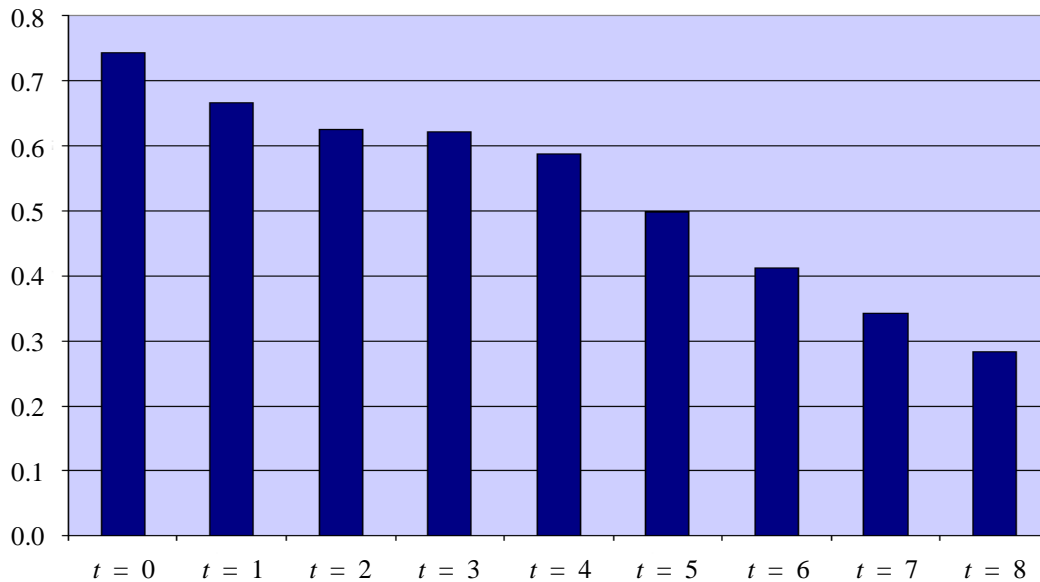
of output at its equilibrium rate at this horizon. But this was not to be. As shown in Figure 8, the SDs of the forecasts themselves decline monotonically as the forecasting horizon lengthens. Indeed, as in the case of forecasting inflation, the degree of bunching of the modal forecasts increases as the horizon lengthens. Thus the SD falls from 0.74 at $t = 0$, to 0.59 at $t = 5$, and then down to 0.28 at $t = 8$.

This result is not consistent with my earlier hypothesis that the authorities were trying to drive the rate of growth of output towards some equilibrium level at some horizon shorter than $t = 8$. What is going on? The implication of the mean forecasts (for inflation and output) is that, on average, the authorities saw themselves as dealing with some mildly deflationary shocks (1997–98; 2000–02) and were setting the level of interest rates so that, once these latter (interest rates) began to have their full effect on output, output would on average be rising steadily throughout, with a tendency to overshoot the equilibrium rate at the (published) horizon.

If you assume a variety of initial shocks, a constant interest rate assumption translating into a stable forecast rate of change in the growth rate of output (i.e., the second derivative), and the objective to hit the inflation target at $t = 8$, then you would get this kind of pattern in the SDs of the forecasts for output growth as the horizon lengthens from $t = 1$ to $t = 8$.

Next we test for the relationship between the SDs of the inflation forecasts and those of the output forecasts, at all horizons. The Orphanides-Wieland-Wilcox hypothesis has the authorities caring much more about output

Figure 8
SD of modal forecast of output growth



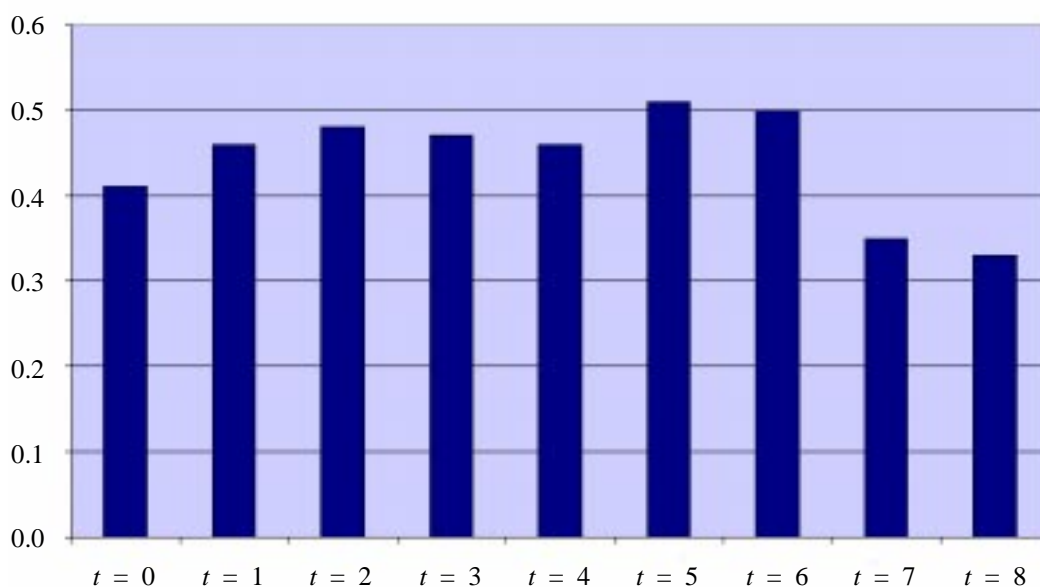
deviations than inflation deviations, so long as inflation is within acceptable bounds. Assuming such bounds relate to the 1 per cent letter-writing points in the U.K. case (1.5 per cent to 3.5 per cent), this latter acceptable range holds throughout. Under that hypothesis I would have expected the forecast SD for output growth to be below that of inflation, especially over the horizon where interest rates are presumed to have their main effect on output, say horizons of $t = 3$ to $t = 6$.

Per contra, if the MPC were indeed sticking to the Chancellor's remit, one would expect the SD of the inflation forecast to be decisively below that of the output-growth forecast, and perhaps increasingly so as the horizon lengthened to the apparent target horizon at $t = 7$ or 8.

The results are shown in Figure 9. They show that the SD of the inflation forecast is always considerably less than that of the output forecast. For the shorter forecast horizons, $t = 1$ to $t = 6$, the SD of forecast inflation remains at a roughly constant ratio to (half) that of the output forecast. Then at $t = 7$ and $t = 8$, the ratio drops, showing that at these horizons the extent of bunching in the modal forecast of inflation much exceeds the (similar, but less marked) bunching in the modal output-growth forecast.

This result, I claim, is entirely inconsistent with the Orphanides-Wieland-Wilcox hypothesis, and entirely consistent with the hypothesis that the MPC's overriding objective has been to drive forecast inflation back into line with target at $t = 7$ and $t = 8$, particularly the latter.

Figure 9
Ratio SD inflation divided by SD output



With these results being consistent with, supportive of, the hypothesis that the MPC has aimed to vary interest rates so as to drive the inflation forecast back to target at a horizon, $t = 8$, broadly two years ahead, we can, in principle, use evidence of asymmetries in the average distribution of the forecasts to test for the nature of the MPC's loss function. On average in these forecasts, the skews (the asymmetric risks) were on the upside, i.e., the chance of a really large upward deviation of inflation from the modal forecast at $t = 8$ was greater than that of a really large downward deviation. If so, and other things being equal, with a quadratic loss function, $L = (\pi - 2.5)^2$, then the MPC should vary interest rates so that the *mean* forecast of inflation, at $t = 8$, should be 2.5 per cent, with the median, and even more so the modal, forecast lying below 2.5 per cent. If the loss function was linear, $L = (\pi - 2.5)$, then the MPC should vary interest rates so that the *median* forecast of inflation, at $t = 8$, should be 2.5 per cent, with the mean forecast above 2.5 per cent and the modal forecast below it. If other things are equal, then in the context of skewed expected outcomes, one can use the relative levels of the MPC's mean, median, and modal forecasts, at $t = 8$, to estimate the exponent in their loss function.

In practice, the average forecast outcomes at $t = 8$ were as follows:

Table 2
Forecast outcomes at $t = 8$ for inflation

	Mean forecast	Median forecast	Modal forecast
Average	2.609	2.583	2.510
SD	0.185	0.146	0.096

Taken literally, this would suggest that the MPC tried to get the modal (most likely) forecast bang on target, despite asymmetric risk. This would suggest (superficially) an exponent in the loss function of zero, below even the Kahneman-Tversky behavioural lab experiments, and a long way from the usual quadratic assumption. There are, however, a lot of caveats.

The first is that the *average* upward skew of inflation was very mild, so that the three measures of central tendency are not significantly different from each other. Thus, this empirical test of the MPC's loss function has very little power. Second, one cannot distinguish between the hypothesis that the MPC did have a quadratic loss function, but also preferred a slightly higher than required inflation rate, from the hypothesis that they aimed to put the modal forecast on the button. Note that this exercise gives exactly the opposite impression to a simple comparison of actual, ex post inflation with the target (where actual inflation, on average, marginally undershot the target). Looking at the forecasts instead (which I would argue was a better way of assessing the MPC's intentions), gives an impression of a committee more willing to accept an upside, than a downside, risk to inflation.

I would contend that the key, persistent feature of these forecasts over this time period was the remarkably high level of the United Kingdom's real effective exchange rate, and that the consequential continuing risk was of a sharp drop in the exchange rate to what would seem a more sustainable level. This concern about the exchange rate can, I believe, be documented, but it would take too long to do so here. This apparent disequilibrium in the exchange rate, I claim, accounts for the upside asymmetric risk to the inflation forecast.

If there is such a (probably small) risk of a collapse in the exchange rate in future, does it make any kind of sense to raise interest rates *now*, thereby probably worsening the current exchange rate disequilibrium? Although such a course has on occasion been advocated, what would seem (to me) a better solution would be to ignore the (slight) risk until clearer evidence emerged that its probability of occurrence had risen (or even until it

happened).³ There should still be time to raise interest rates (above their projected path), to deal with the perceived risk should that transpire.

As already noted, the MPC forecasts are conditioned on an assumption of unchanged interest rates over the forecasting horizon. Another way of interpreting these findings is that this conditional forecasting assumption is itself conditioned on the assumption that the perceived asymmetric risks lying ahead do *not* come about during the forecast horizon. If they did, an interest rate response would be forthcoming. Since Lars Svensson does not like the conditional assumption of constant future interest rates (2003), he may be slightly mollified by this minor qualification.

Besides the fact that the difference between the mean, median, and modal forecasts at an horizon of $t = 8$ is insignificantly small, what is interesting is that there are several differing possible explanations of the finding that the mean was above 2.5 per cent and the mode equal to 2.5 per cent. Let me repeat these, and give my view of their prior probability.

- (i) Exponent in the loss function equals zero; assessed probability, almost zero.
- (ii) MPC prefers upside inflation to downside inflation risk; assessed probability, one-quarter.
- (iii) Method of handling the prospect of a low probability, but major future asymmetric risk; assessed probability, three-quarters.

Conclusion

This paper has sought to demonstrate that the way to identify and indeed to quantify the objectives of the MPC is to examine their forecasts, not the ex post actual outcomes that are distorted by unforeseen shocks. What I claim to have established is that the MPC has indeed aimed to drive the inflation forecast into line with target at a two-year horizon, with this latter horizon being well determined empirically. The Orphanides-Wieland-Wilcox “opportunistic loss function” does not hold in the United Kingdom. A comparison of the average mean, median, and modal forecasts at this horizon, $t = 8$, gives some surprising results, though the differences are small. My best judgment is that this latter finding arises from an intelligent response of the MPC to asymmetric risk.

It would be interesting to do a companion exercise for other countries, but I do not know of any countries where comparable data would make that feasible.

3. This policy appears completely at odds with the mini-max, “robust,” approach to policy making, whereby the policy-maker tries to minimize the likelihood of the worst possible outcome. I address this issue in accompanying work with Margaret Bray (see Bray and Goodhart 2002).

References

- al-Nowaihi, A. and L. Stracca. 2003. "Behavioural Central Bank Loss Functions, Skewed Risks and Certainty Equivalence." *Manchester School* 71 (Supplement): 21–38. Also chapter 2 in Stracca's Ph.D. thesis, Leicester University; and European Central Bank Working Paper No. 129 (2002).
- Bank of England. 1999. "Economic Models at the Bank of England." Pamphlet. April.
- . 2000. "Economic Models at the Bank of England, September 2000 Update." Pamphlet.
- Bean, C. 2003. "Asset Prices, Financial Imbalances and Monetary Policy: Are Inflation Targets Enough?" Presented at BIS Conference on *Monetary Stability, Financial Stability and the Business Cycle*. Basel, Switzerland.
- Bray, M. and C. Goodhart. 2002. "You Might as Well Be Hung for a Sheep as a Lamb: The Loss Function of an Agent." LSE Financial Markets Group, Discussion Paper No. 418.
- Brown, G., Chancellor of the Exchequer. 2003. "Remit for the Monetary Policy Committee." Letter to Sir Edward George, Governor of the Bank of England (April).
- Chadha, J.S. and P. Schellekens. 1999. "Utility Functions for Central Bankers: The Not So Drastic Quadratic." LSE Financial Markets Group Discussion Paper No. 308. Also available as "Monetary Policy Loss Functions: Two Cheers for the Quadratic." Bank of England Working Paper No. 101.
- Goodhart, C. 2001. "Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates." *Federal Reserve Bank of St. Louis Review* 83 (4): 165–81 (July/August).
- Kahneman, D. and A. Tversky (ed.). 2000. *Choices, Values, and Frames*. Cambridge: Cambridge University Press.
- Monetary Policy Committee, Bank of England. 1999. "The Transmission Mechanism of Monetary Policy." Pamphlet. Bank of England.
- Orphanides, A. and V. Wieland. 2000. "Inflation Zone Targeting." *European Economic Review* 44 (7): 1351–87.
- Orphanides, A. and D.W. Wilcox. 2002. "The Opportunistic Approach to Disinflation." *International Finance* 5 (1): 47–71.
- Svensson, L.E.O. 1997. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." *European Economic Review* 41 (6): 1111–46.

- Svensson, L.E.O. 2003. "The Inflation Forecast and the Loss Function." Chapter 5, in *Monetary Policy, Central Banks and Financial Regulation*, Cheltenham, United Kingdom: Edward Elgar.
- Svensson, L. and M. Woodford. 1999. "Implementing Optimal Policy through Inflation-Forecast Targeting." Princeton Economics Department. Photocopy. Also NBER Working Paper No. 9747.
- Terlizzese, D. 2002. "A Note on Lexicographic Ordering and Monetary Policy." Research Department, Banca d'Italia. Photocopy.
- Treasury, HM. 2003. *UK Membership of the Single Currency: An Assessment of the Five Economic Tests*. Monograph (Cm 5776). London: HM Treasury, June 2003.