

Speech by

CHARLES BEAN
CHIEF ECONOMIST & EXECUTIVE DIRECTOR
BANK OF ENGLAND

At Oxonia Distinguished Speakers Seminar,
The Oxford Institute of Economic Policy, Oxford

On Tuesday 22 February 2005

Monetary Policy in an Uncertain World
Speech to Oxonia
Oxford, 22 February 2005

Good evening! The great statesman and scientist Benjamin Franklin – who received an honorary doctorate of civil law from this University in 1762 – is famous for having observed that there are only two certain things in life: death and taxes. But had he been a member of the Monetary Policy Committee, he would have realised there is certitude about something else – namely, uncertainty. For uncertainty is an ever-present feature of the economic landscape that monetary policy makers cannot escape.

Broadly speaking, there are three types of uncertainty that confront us on the MPC: uncertainty about the data; uncertainty about the nature and persistence of shocks; and uncertainty about the structure of the economy. I shall address each of these in turn, drawing out some of the implications for policy on the way. In the context of my discussion of uncertainty about the structure of the economy, I shall also have quite a bit to say about the possible explanations for the unusual stability in macroeconomic performance in recent years. But let me start with some observations on the problems posed by data uncertainty¹.

Uncertainty about data

A former Chancellor of the Exchequer once remarked that steering the economy was like driving along a winding road looking only in the rear-view mirror. In fact there are times when it seems as if it is a great deal more difficult, for that mirror itself is misted up. We do not know where we are, or even where we have just been with any precision. Virtually all the data we rely on are subject to measurement error, either because of sampling error or because they do not correspond exactly to the economic concepts that we are interested in. They appear usually only with a lag and can be subject to considerable revision until well after the fact.

As an illustration of the need to take this issue seriously, Chart 1 shows the revision pattern in the recent history of GDP growth, from its preliminary estimate that is available a couple of months after the end of the quarter through to more recent estimates. On average, first estimates of growth have been revised up by 0.1-0.2 percentage points over the period since the preliminary estimate was first published in 1998 Q3, but more striking is the relatively large variance of the revisions².

I emphasise that this should not be taken as a criticism of those in the ONS and other agencies that collect and assemble the data. Rather it simply reflects the inherent difficulty in measuring the large and complex phenomenon that is the economy, comprising millions of households and businesses, both accurately and in a timely fashion. The ONS use a variety of data sources in constructing, say, their measure of GDP, with more information continually accruing and being integrated into that measure. Of course, they could wait two years or more until all the underlying information were in before producing an estimate, but that would not be much good to us on the MPC.

Instead, the MPC recognises that early releases of data are prone to revision and combines it with other information that has in the past proved useful in predicting the final vintage of ONS data³. The various business surveys produced by the CBI, BCC, CIPS and others provide a source of such independent information. Though the samples for these surveys are typically much smaller than those that the ONS can draw on and the responses are usually qualitative, ie in the form of “balance of ups over downs”, rather than quantitative in nature, statistical tests do suggest that they have some incremental predictive content over and above the early releases of the official data. In addition, sometimes one might want to “aim off” a piece of official data because it is hard to reconcile with movements in other data series that economic theory and previous empirical work suggests ought to move together.

A particularly pernicious form of data uncertainty relates to measuring the level of aggregate demand relative to supply – a key determinant of inflationary pressure in the economy. Aside from the fact that the early estimates of GDP are subject to revision, supply potential is never directly observed but rather must be somehow estimated. A popular approach in business cycle research is to use an appropriate statistical filter to separate output into potential supply and the output gap, ie the deviation of actual output from potential output. But there is a fundamental identification problem here. To distinguish between a change in supply capacity and a change in the intensity of use of factors requires an assumption about the way trend output changes over time. For instance, it is usually assumed that potential output grows smoothly, but it is possible that the rate of growth of technical progress actually varies from quarter to quarter. Moreover, measures obtained from statistical filters are usually subject to an “end-point” problem whereby the absence of any data for the future makes the most recent estimates of potential output – and therefore also the

associated estimate of the output gap – particularly unreliable. But, of course, it is precisely the recent past that the policy maker is most interested in!

In any case, on the MPC we prefer to build up our picture of the inflationary pressures in the economy by considering the pressure of demand relative to supply separately in the product and labour market rather than relating inflation to a single catch-all measure of the aggregate output gap. But unfortunately neither of these is directly observed either.

As far as the product market goes, one wants to be able to compare the volume of output with that which could be produced with the inputs that are currently employed when operated at a normal, or sustainable, rate. So one can look either at business survey questions on capacity utilisation or use econometric estimates of a production function. But, on the one hand, survey-based measures are problematic. They are often only available for manufacturing – indeed it is conceptually hard even to define capacity in some service industries, let alone measure it – and there are usually doubts about exactly how the respondents interpret the questions that they have been asked. On the other hand, measures of utilisation based on econometric estimates of a production function are only as good as the underlying model of the production technology.

But if it is tricky estimating the margin of spare capacity in the product market, it is even harder to reliably assess its labour market equivalent. Much research effort has been expended over the last 25 years trying to explain the movements in the equilibrium, or natural, rate of unemployment in terms of changes in labour market institutions, etc, both in the UK and in the rest of the OECD. While substantial progress has been made, we are still a long way from having a complete story.

Moreover, the picture is muddied by the fact that people find jobs from inactivity as well as from the unemployment pool. And if employers can recruit workers easily from overseas, the whole concept of a domestic supply of labour may not even be very relevant. Measures of vacancies and survey indicators of recruitment difficulties provide a different perspective, but in the end one cannot do much more than look at all the available indicators together and try to form a balanced judgement. However, one needs to be acutely aware that judgement may be badly wrong.

The perils of real-time output gap mismeasurement have been forcefully pointed out by Athanasios Orphanides et al.⁴ and Ed Nelson and Kalin Nikolov⁵. The former conduct an *ex post* evaluation of monetary policy in the United States through the 1970s and 1980s. They use real-time data to estimate the output gap as it might have appeared to policymakers at the time and then show how seriously an interest rate policy that reacted just to inflation and the output gap, ie a conventional Taylor rule, could go wrong as a result of mismeasurement of the output gap. They also show that a policy that reacted instead to inflation and the change in inflation, rather than the output gap, would have performed better (the reason being that the change in inflation is related to the unobservable true output gap and thus serves as a proxy for it in the policy rule).

Nelson and Nikolov adopt a similar approach for the United Kingdom, though employing the forecasts and statements of the Treasury and the National Institute (as a proxy for official Treasury thinking) to construct measures of the output gap as it appeared at the time. They find that monetary policy errors due to output gap mismeasurement contributed 3-7 percentage points to average UK inflation in the 1970s and 1-5 percentage points to inflation in the 1980s. So their analysis also suggests that real-time estimates of the output gap can be seriously misleading.

The bottom line of all this is that, although the output gap may be an indispensable intellectual construct for organising one's thinking, the inherent uncertainty around any given measure means that it should always be employed with a considerable degree of caution and an appropriate degree of scepticism.

Uncertainty about shocks

The second source of uncertainty relates to the nature and persistence of the shocks hitting the economy. Sometimes the shock can be pretty obvious, but its effect may not be. For instance, the Golden Jubilee reduced the number of working days in 2002, complicating the measurement of seasonally adjusted output in the economy. In other cases, it may be not only the impact of the shock, but also its persistence, which is in question, a good example being the recent rise in the oil price. Moreover, in this case what matters for the economy is what the private sector, not the policy maker, believes about the persistence of the shock. Futures prices provide some guide but even they may not coincide with the perception of businessmen.

On other occasions, even the source of the shock may not be apparent. For instance, long-term real forward interest rates around the world are currently at unusually low levels (see Chart 2). But why this should be is not immediately obvious, especially given the large budget deficits in some countries. It could be because the demand for funds to invest is low because the marginal product of capital is low – but there is no sign of the world running out of profitable investment opportunities, especially given the rapid development taking place in China and elsewhere in the Asian sub-continent. Or it could be the consequence of demographics or increased saving by households worried about their living standards in retirement. Getting to the bottom of such puzzles is the daily task of those who work in central banks.

Moreover, the econometric models that all central banks use in forecasting are estimated on past data that incorporate a mixture of different types of shocks. And the shocks that are impinging on the economy now may be different from the average of those that impacted on it in the past. In a first-best world, our models would be specified at a deep enough level that the differential response to different types of shocks would be properly spelt out. Building such models is at the core of much modern macroeconomics, but we are a long, long way from achieving that objective. Indeed, given that the macroeconomic data that we see results from the aggregation of the decisions of millions of different economic agents, each of whom is subject to a myriad of influences, such an ideal is almost certainly unachievable.

Instead, models will inevitably remain as gross simplifications, requiring a range of more or less *ad hoc* features that help them to explain the past. In that case, when a shock occurs the policy maker needs to ask whether or not the response of agents is likely to be the same as it has been in the past. To give a particular example, consider the housing market. As Chart 3 shows, in the past there has been a high correlation between house price inflation and consumption growth, but that correlation appears to have weakened in the last three years or so. The same phenomenon is apparent in consumption functions that include house prices as well as the usual suspects (income, wealth, etc), which typically suggest that consumption should have been stronger than it was over the last three years or so.

The natural explanation for this is that the previous occasions of rapid house price inflation have generally been associated with periods when income expectations became markedly more optimistic and/or financial liberalisation. In such circumstances one would expect to see both consumption and house prices rise

together, driven by these third factors. But this time round, the household saving ratio has been pretty flat over the last three years, suggesting that these factors have not been especially important. Instead, factors more specific to the housing market seem to have been at work, including demographic developments that have raised the demand for housing set against only a moderate rate of growth in the supply of housing, as well as the impact of lower interest rates on the time profile of real mortgage repayments. Of course, that does not imply the complete absence of a structural link between house prices and consumption – housing wealth is, after all, the key source of collateral for most households – but rather that one cannot necessarily rely on past data correlations to be maintained in such circumstances.

Uncertainty about the structure of the economy

The third source of uncertainty, and that which I want to spend most time on this evening, is uncertainty about the structure of the economy. Structural changes, either real or apparent, can arise in just about any part of the economy. But a particularly pertinent question that faces the MPC – as well as central bankers in other developed economies – is whether the low inflation and unusual stability of both inflation and growth in the last decade or so betokens fundamental changes in the way our economies function or whether it is just a temporary aberration. Charts 4 and 5 illustrate this stability by displaying trailing moving standard deviations, calculated over successive overlapping eight-year periods, of GDP growth and inflation respectively for each of the G7 countries, excluding Japan (for which it is difficult to construct a consistent time series for GDP over a long enough period). The extent to which volatility has fallen across all the countries in the sample is striking.

The source of this “Great Stability”, which contrasts so markedly with the “Great Inflation” of the 1970s, has been the subject of an increasing volume of research, mainly but not exclusively focussed on the United States. Essentially three classes of explanation have been advanced: good luck; changes in the structure of the economy; and better macroeconomic policy.

Is it luck?

One possibility is that policy makers have just gotten lucky. According to this view, the shocks impinging on the developed economies have been smaller and less persistent than before. Moreover, they have also been less synchronised across

countries. There is clearly some merit in this view, but the last couple of decades have not been entirely devoid of significant shocks, including the break-up of the USSR, German re-unification, the Iraq wars, the Mexican debt crisis, the Asia crisis, the LTCM crisis, the dotcom boom-bust, 9/11 and so on. And while there have been no major oil price shocks on the scale of OPEC I and OPEC II, the fluctuations in oil prices have nevertheless been significant from time to time, including over the past year.

Jim Stock and Mark Watson⁶ report evidence suggesting that it is indeed such good fortune that accounts for the reduced volatility of the business cycle over the last couple of decades. But other authors, such as Steve Cecchetti et al., find that the factors discussed below – and better monetary policy in particular – should take the lion's share of the credit. The real difficulty with assessing the relative importance of luck compared to structural changes or improved policy is that the shocks are identified with the residuals in econometric equations and there is no way of knowing whether the *size* of those shocks has been reduced as a result of structural changes to the economy or better policy.

Structural changes

Another possibility is that changes in the structure of economies mean that shocks have a smaller or less persistent effect than in the past. There are at least three reasons why this might be so. First, it may be the result of changes in the composition of output of developed economies, away from manufacturing and towards services. Because capital and durable goods are predominantly in manufactured, manufacturing is typically more cyclical than services. So, other things being equal, a more service-intensive economy ought to display less cyclicity than a more manufacturing-intensive one. However, the shift in the composition of output has been so gradual in most industrialised economies it is difficult to believe this is a major factor.

A second, and more plausible, reason can be found in the consequences of financial deregulation, innovation and integration. Greater access to credit allows households and businesses to smooth their spending when times are bad. The corollary is that the saving ratio should be more volatile than it was in the past; that appears to be the case, at least in the United States⁷. Moreover, financial innovation has led to the development of new derivative assets that allow idiosyncratic risk to be diversified more effectively, again making the economy more resilient to shocks⁸. Finally,

international financial integration enhances the scope for risk sharing across countries, though the fact that portfolios are still heavily home-biased suggests that this last effect may be rather weak.

A third possible reason lies in the impact of information and communication technology on the inventory cycle. One might have expected that inventories would act as a damper on cyclical fluctuations as they constitute a buffer between sales and output and thus allow firms to smooth production. But in the past, inventories instead seem to have acted as an amplifying mechanism, with stock levels behaving in a pro- rather than an anti-cyclical fashion. Better inventory management techniques have allowed firms to keep production more closely in line with sales, so reducing the contribution of the inventory cycle to the business cycle⁹. However, at best this only seems likely to constitute a small part of the story. Changes in inventory management have been pretty gradual. And the reduction in volatility in production is similar to the reduction in volatility in sales, not greater as would be expected if this were the correct explanation.

Better policy making

The final explanation for the Great Stability rests on improved policy making, in part reflecting an improved understanding by policy makers of the way the economy functions and the trade-offs that they face. So, and caricaturing only slightly, during the 1960s and early 1970s policy makers saw themselves as confronted with an exploitable Phillips curve trade-off presenting a stable menu of choices between unemployment and inflation: one could have permanently lower unemployment if one were prepared to accept permanently higher inflation. Even though Bill Phillips himself recognised that the position of the curve might be affected by the state of workers' inflation expectations, that qualification was largely ignored by policymakers.

By the start of the 1980s – partly as a result of the juxtaposition of high inflation with rising unemployment in the 1970s and partly as a result of theorising about the underpinning for the Phillips Curve by Milton Friedman, Ned Phelps and Bob Lucas – policy makers had come to realise that no such trade-off existed in the long run and that ultimately monetary policy needed to focus on controlling nominal, not real, magnitudes. A by-product of this change in view was that official interest rates

became more responsive to inflation and, critically, that *real* interest rates rose when inflation rose, in order to dampen demand and so push inflation back down¹⁰.

Parallel to this improvement in economic understanding there has been an improvement in the institutional arrangements for conducting monetary policy, with moves towards greater independence of central banks from their political masters and greater transparency over the objectives of monetary policy. That is obviously the case in this country, with the adoption of an inflation target in 1992 and the creation of the Monetary Policy Committee in 1997, but similar trends can also be observed in a number of other countries.

That better monetary policy might be the explanation for lower and more stable inflation is perhaps not too surprising. More of a puzzle is how that connects to producing lower volatility in output. Indeed if all nominal wages and prices were flexible so that the Classical Dichotomy held, monetary policy would determine inflation outcomes, but would be irrelevant to the behaviour of output.

But, of course, nominal wages and prices are not totally flexible, thus generating a short-run trade-off between inflation and activity. Moreover, the slope of that trade-off does seem to be related to the average rate of inflation. Thus Larry Ball, Greg Mankiw and David Romer¹¹ present cross-country evidence that suggests the short-run trade-off is flatter in low inflation countries than in high inflation ones. Moreover, time-series evidence also suggests that the Phillips Curve may be flatter when inflation is low. Chart 6 plots CPI inflation against unemployment for the G7¹². The changing character of the Phillips relationship is most marked in the United Kingdom, where it has gone from being approximately vertical in the 1970s, to downward sloping in the 1980s, to being approximately flat in the 1990s. But a flattening also appears to be present to varying degrees in most of the other countries. Luca Benati¹³ has explored the changing nature of this relationship in the United Kingdom, showing how it is related to the monetary regime in force, with a flattening taking place in the 1980s and a particularly high degree of stability about the relationship since the adoption of inflation targeting.

This flattening of the apparent short-run relationship between activity and inflation is predicted by some New Keynesian pricing models. Essentially, when inflation is low, firms are likely to change prices less frequently and this leads to a weaker short-term impact of demand fluctuations on prices¹⁴.

But it may also reflect the impact of greater credibility in monetary policy making, as the behaviour of wage and price setters today will be affected by their expectations of the general level of prices over the duration for which those wages and prices are set. If prices are expected to remain stable then, say, a positive shock to demand is less likely to lead to higher wages and prices, than if it promotes expectations of a higher future price level.

That a flatter Phillips Curve might help to explain why inflation has been more stable is not altogether surprising, as fluctuations in aggregate demand would have less effect on inflation. However, other things being equal, there would be a corresponding *increase* in the variability of output in contrast to what has actually happened. So how can we explain the fact that the volatility of output also fell?

There are two possibilities here. First, as I noted earlier, policy makers probably have a better understanding today of how the economy works than was the case 25 years ago. But that has been coupled with greater realism about what monetary policy can, and cannot, do as well as institutional changes that have tended to reduce the extent to which short-term political considerations could dictate interest rate decisions. As a consequence monetary policy has itself been less of a destabilising force than in the past.

Second, as the counter-inflationary credibility of monetary policy increased, so central banks found themselves better able to offset disturbances without the danger of destabilising inflation expectations. To a degree this can happen automatically through asset prices: a positive shock to domestic demand prompts an expectation of higher official interest rates in the future, leading to a rise in longer-term interest rates and an appreciation of the exchange rate, thus counteracting the original shock. In other words, credibility enhances the effectiveness of policy through its impact on expectations.

Policy credibility may enhance the effectiveness of policy not only in the face of demand shocks, but also in the face of supply shocks. When policy is credible and inflation expectations are well anchored, then the chance of an adverse supply shock triggering a wage-price spiral is much less than when people believe that the central bank will accommodate the shock and allow inflation to rise.

The possible importance of this consideration is illustrated by the response of financial markets to the recent increase in oil prices. The spot price of oil rose by around two-thirds between the beginning of 2004 and the autumn, with particularly sharp increases in June-July and September-October (marked with the grey bars in Chart 7). There were a number of reasons for this increase, including rapid growth in demand associated with the global expansion, low stock levels in the United States, geopolitical concerns in the Middle East and interruptions to supply in a number of countries. The increase was almost certainly unanticipated as it was not remotely signalled in the futures price which had been pointing to a modest decline in the oil price to within OPEC's \$22-28pb target range.

Given the experiences of the 1970s, one might have expected the increase in oil prices to lead market participants to expect an increase in inflation and market interest rates to move higher in the expectation of monetary tightening by central banks. And neither the rise in oil prices nor the increase in interest rates could be expected to be good news for equity prices. But what happened? Inflation expectations implied from nominal and indexed bonds hardly moved and market interest rates moved down rather than up (Charts 8 and 9). Moreover, during the September-October period equity prices actually strengthened (see Chart 10). Now there may be other explanations for this constellation of asset price movements, not least as other shocks might have occurred at the same time. But the natural interpretation is that with inflation under control and inflation expectations well-anchored, market participants believed that central banks would be able to pursue a more relaxed monetary policy in order to offset the adverse demand effects of the oil price increase without needing to worry about setting in train a wage-price spiral of the sort seen in the 1970s.

Some implications for policy

The apparently flattening Phillips Curve raises a number of issues for policy. Taken at face value it appears to suggest that we could continue to push down on unemployment with minimal implications for inflation. And it also appears to suggest that if inflation is materially above (below) target then the policy maker needs to engineer a very large recession (boom) in order to bring it back, ie the so-called sacrifice ratio is large. But this would be seriously to misunderstand the nature of the beast and to run the danger of repeating the mistakes of the 1970s.

First, the Phillips Curve plot in Chart 6 does not trace out a structural relationship between unemployment and inflation. The approximate invariance of inflation to the level of unemployment through the 1990s is consistent with a flat structural relationship between the two. But it is also consistent with a conventional downward-sloping short-run structural relationship between the two that has been simultaneously shifting to the left as the natural rate of unemployment fell with monetary policy simultaneously ensuring that inflation remained stable. Now it is reasonable to believe that the structural relationship is actually flatter than in the past, for the reasons laid out earlier. But since the full effects of the labour market reforms of the 1980s and 1990s are likely to take time to work their through, it is equally plausible that the natural rate of unemployment has been falling, while the operation of the inflation targeting regime since 1992 should have helped keep inflation stable. So this is an alternative explanation, observationally equivalent to a flattening in the short-run Phillips Curve.

Seen this way, it is clear that the first and third sources of uncertainty that I have talked about today here interact, with uncertainty about the natural rate of unemployment making it harder to pin down the slope of the short-run trade-off between unemployment and inflation and *vice versa*. As a consequence, there is uncertainty both about the extent of the underlying pressures on inflation and about the impact of changes in interest rates on inflation and activity.

How should policy be set in such circumstances? Very cautiously, I believe. It is possible that the sustainable level of unemployment could be even lower than its current level, which is already pretty low by historical standards. But equally, if that were not the case, then continued attempts to push activity higher would at some stage be likely to generate noticeable upward pressure on inflation as supply bottlenecks become more prevalent. There are therefore good reasons to believe that the Phillips Curve is indeed a curve rather than linear and there is some empirical evidence to support that belief¹⁵. Moreover, if the increase in inflation were substantial, there would be a potential loss to the credibility of the inflation target and a consequent de-anchoring of inflation expectations. And once credibility is lost, it can be costly to regain.

The second observation is that inflation targeting may be a particularly suitable regime when the structure of the economy is changing. Unlike a policy regime based around an intermediate target, such as the money supply, a regime described in terms

of the ultimate goals of policy does not need to be changed whenever the structure of the economy changes – though the analytical processes within the central bank may need to change substantially, of course. Furthermore, future inflation is something that agents intrinsically care about – it affects wage and price setting behaviour and also determines the real interest rate and thus affects aggregate demand too. Even if structural change makes the policymaker less sure about his own ability to meet the target, the existence of the target should itself help to keep private agents' expectations pinned down. That might not be the case if the policy objective were stabilising, say, nominal GDP as agents would become more uncertain about the implications for inflation as a result of uncertainty about the supply-side. So this is another reason to believe that inflation targeting may prove to be a more durable regime than some of its predecessors.

The third observation relates to the link between monetary policy and asset prices. Claudio Borio and Bill White¹⁶ and others have argued that one of the consequences of the anchoring of inflation expectations and the flattening of the Phillips curve is that overheating tends to be manifested in asset price inflation rather than goods price inflation. According to them, the greater likelihood of asset price boom-busts also increases the likelihood of future financial instability and this should be taken account of in the setting of interest rates. Elsewhere, I¹⁷ have argued that this concern can be satisfactorily accommodated within an inflation targeting framework because future financial instabilities are likely to lead to a reduction in growth and a departure of inflation from target, though it does require central bankers to take the long view. There is no doubt, however, that such asset price boom-busts potentially complicate the conduct of policy.

My final observation relates not to the conduct of policy, but rather to a detail of our communication strategy. Each quarter, the MPC presents its projections for GDP growth and inflation in the form of explicit probability distributions or “fan charts”. Reflecting the unusual stability in the economy outturns have tended to be closer to the centre of the forecast distributions than one would have expected. For instance, 14 out of 18 of the two-year-ahead forecasts for CPI inflation made between February 1998 and May 2003 lay in the part of the distribution covering the central 50%. The equivalent figure for the GDP growth projections is 11 out of 18. Some commentators have inferred from this that the variance on our fan charts is too wide. But these fan charts portray the MPC's *subjective* uncertainty over economic prospects. Given that there must be at least some chance that the good performance of recent years is down

to good luck rather than structural changes or better policy, it makes sense not to reflect all of the recent decline in volatility in our assessment of the degree of uncertainty about future outturns.

Let me conclude my remarks by noting that the presence of uncertainty means that it is inevitable that some monetary policy decisions will seem unwise after the fact. Wisdom with hindsight is a wonderful thing, but unfortunately it is a luxury that policy makers do not have. But we can at least strive to reduce that margin of uncertainty by continually seeking to improve our understanding of the forces driving the economy. For, as Benjamin Franklin also remarked, “An investment in knowledge always pays the best interest”.

Endnotes

-
- ¹ For a more extensive treatment of some of these issues, see Bell, M. (2004) ‘Monetary policy, data uncertainty and the supply-side: living with the statistical fog’, speech to the Society of Business Economists, 15 September.
- ² For some analysis of revisions to National Accounts data, see Castle, J. and Ellis, C. (2002) ‘Building a real-time data base for GDP(E)’, *Bank of England Quarterly Bulletin*, Spring, 42-49.
- ³ One might be tempted to argue that the ONS should use this information in constructing their own estimates. But that might involve the input of a considerable amount of economic interpretation and judgement. It is better that the MPC knows what the official data is based on and then chooses how to interpret it in the light of other indicators, rather than the ONS making those adjustments.
- ⁴ Orphanides, A. Porter, R., Reifschneider, D., Tetlow, B. and Finan, F. (2000), ‘Errors in the measurement of the output gap and the design of monetary policy’, *Journal of Economics and Business*, 52, pp. 117-141.
- ⁵ Nelson, E. and Nikolov, K. (2001), ‘UK inflation in the 1970s and 1980s: the role of output gap mismeasurement’, *Bank of England Working Paper*, 148.
- ⁶ Stock, J. and Watson, M. (2003), ‘Has the business cycle changed? Evidence and explanations’, Federal Reserve Board of Kansas City symposium, *Monetary Policy and Uncertainty*. Cecchetti, S., Flores-Laguna, A. and Krause, S., (2004) ‘Has Monetary Policy Become More Efficient? A Cross-Country Analysis’ mimeo, presents a contrary view. The observation that it is difficult to draw any conclusion because the size of the shocks may reflect structural changes or policy improvements is due to Bernanke, B. (2004) ‘The Great Moderation’, speech to the Eastern Economic Association, Washington, 15 February.
- ⁷ See Blanchard, O. and Simon, J. (2001), ‘The long and large decline in US output volatility’, *Brookings Papers on Economic Activity*, 1, 135-164.
- ⁸ See Greenspan, A. (2004), ‘Economic flexibility’, speech to HM Treasury Enterprise Conference, London, England, 26 January.
- ⁹ See McConnell, M. M. and Perez-Quiros, G. (2000), ‘Output fluctuations in the United States: what has changed since the early 1980’s?’, *American Economic Review*, 90(5), 1464-1476.
- ¹⁰ See Clarida, R., Gali, J. and Gertler, M. (2000) ‘Monetary policy rules and macroeconomic stability: evidence and some theory’, *Quarterly Journal of Economics*, 115(1), 147-180, and Taylor, J. (1999) ‘An historical analysis of monetary policy rules’, in Taylor, J. (ed.), *Monetary Policy Rules for evidence on the United States and Nelson, E. (2003) ‘UK monetary policy, 1972-1997: a guide using Taylor rules’, in Mizen, P. (ed.) Central Banking, Monetary Theory and Practice, Essays in Honour of Charles Goodhart, for evidence regarding the United Kingdom.*
- ¹¹ Ball, L., Mankiw, N. G., and Romer, D. (1988): ‘The New Keynesian Economics and the Output-Inflation Trade-Off’, *Brookings Papers on Economic Activity*, 1, 1-65.
- ¹² The original Phillips relationship related wage growth to unemployment. The pictures are similar if nominal wage growth rather CPI inflation is placed on the vertical axis.
- ¹³ Benati, L. (2004), ‘Evolving post-World War II U.K. economic performance’, *Journal of Money, Credit and Banking*, 36(4), 691-717.
- ¹⁴ Though this need not always be so; see Bakshi, H., Khan, H., and Rudolf, B. ‘The Phillips curve under state dependent pricing’ *Bank of England Working Paper*, 227.
- ¹⁵ For some recent evidence, see Dolado, J., Maria-Dolores, R., and Naveira, M. (2005), ‘Are monetary policy reaction functions asymmetric?: The role of non-linearity in the Phillips curve’, *European Economic Review*, 49, 485-503, and references therein.
- ¹⁶ Borio, C. and White, W. (2003), ‘Whither monetary and financial stability? The implications of evolving policy regimes’, Federal Reserve Board of Kansas City symposium, *Monetary Policy and Uncertainty*.
- ¹⁷ Bean, C., (2003) “Asset Prices, Financial Imbalances and Monetary Policy: Are Inflation Targets Enough?” in A. Richards and T. Robinson, eds., *Asset Prices and Monetary Policy*. Sydney: Reserve Bank of Australia, 48-76.