

Mervyn King: Monetary policy – practice ahead of theory

Text of the Mais Lecture 2005 by Mr Mervyn King, Governor of the Bank of England, at the Cass Business School, City University, London, 17 May 2005.

The references can be found on the Bank of England's website.

I am indebted to James Proudman, Gertjan Vlieghe, Tony Yates and Richard Harrison who have worked closely with me on this lecture and are effectively co-authors. Alan Mankikar and Tim Taylor provided excellent assistance in preparing the empirical and historical research that support this lecture.

* * *

Introduction

Monetary policy is aimed at maintaining price stability. That may seem self-evident. Thirty years ago it was not. From the end of the second world war until the mid to late 1970s, the majority view of academic economists and policy-makers alike was that monetary policy had rather little to do with inflation, and was largely ineffective as an instrument of demand management.¹ The intellectual basis for that view was never clear. And painful experience taught us that price instability led to costly fluctuations in real output and employment. Far from being ineffective, a monetary policy aimed at price stability has proved to be the key to successful management of aggregate demand. Fortunately, the theory and practice of monetary policy in the UK have changed out of all recognition in the past twenty-five years.² We have moved from the Great Inflation to the Great Stability.

The story of monetary policy in Britain during the intervening period is told by the Mais Lectures. The first Mais Lecture was delivered by my predecessor, Lord Richardson, in 1978, at a point when monetary policy was emerging as the main tool to deal with inflation. Not before time, you might think, since only two years earlier inflation had reached 27%. In 1981 the Chancellor of the Exchequer, Geoffrey Howe, chose as the title of his Mais Lecture: "The Fight Against Inflation". As he said then, with inflation still in double figures, "squeezing inflation out from an economy which has become accustomed to higher rates over a period of years cannot be an easy or painless task. ... the inflationary mentality must be eradicated. ... When we have done that we will find that low inflation or even price stability need not be painful". The conquest of inflation was to prove harder than expected. In the decade that followed Geoffrey Howe's lecture, inflation averaged over 7% a year. Only since 1992 has inflation fallen to levels that could be described as price stability.

In retrospect, two Mais Lectures seem to have been of particular significance: those by Nigel Lawson in 1984 and by Tony Blair, then Leader of the Opposition, ten years ago this month. Despite clear differences of view, what stand out from those two lectures are their similarities. Both emphasised the need for a medium-term framework for monetary and fiscal policy. Over twenty-five years we have moved from monetary targets to an inflation target and from a medium term financial strategy to rules for fiscal policy over the cycle. Yet the essential objective of maintaining monetary and fiscal discipline remains the same. All major political parties in the UK now agree that stability is the key to economic success.

We do not know whether the Great Stability will continue, as it has for more than a decade now. In part, it will depend upon whether our framework of inflation targeting can respond to the economic shocks that will undoubtedly be visited upon us in the years ahead. And that is the subject of my lecture. In only fifteen years inflation targeting has taken the central banking world by storm. Table 1 shows that there are now 22 countries in which monetary policy is based on an inflation targeting regime. So tonight I want to discuss what inflation targeting really means, why it has been successful in Britain and elsewhere, and what challenges it faces in the years ahead.

The lecture tries to answer three questions. First, what can monetary policy do and how has our understanding of that changed over time? Second, what are the challenges for central banks that result from incomplete knowledge of the transmission mechanism of monetary policy? Third, is inflation targeting the answer to those challenges? I believe that it is. Inflation targeting, I shall argue,

¹ This proposition is documented in detail by Batini and Nelson (2005).

² See Capie and Wood (2001).

is the natural way to conduct policy when there is a great deal about its effects that we do not understand. The practice of monetary policy must recognise that monetary theory will continue to evolve. That is why my lecture is subtitled: practice ahead of theory.

What can monetary policy do?

In practice, monetary policy means setting the level of the official interest rate at which the central bank deals with the banking system.³ But ideas about how interest rates should be set, and with what objective, have been subject to radical changes since the 1970s. Let me give three examples. None is new; the subject has moved on. I give them to show that monetary policy operates against an ever-changing backdrop of ideas about the way the economy works, a theme that lies at the heart of my lecture.

First, it is now widely accepted that there is no long-run trade-off between output and inflation. Both theory – following Friedman and Phelps – and practice – particularly in the 1970s – showed that permanently higher inflation does not bring faster growth or higher employment, and may well reduce both. But in the post-war period views were different. In 1959 the Radcliffe Report on the Workings of the Monetary System seemed to support the idea of a permanent trade-off. The objectives of monetary policy included, it argued, “a high and stable level of employment” and “reasonable stability of the internal purchasing power of money”. But it went on, “...there are serious possibilities of conflict between them.”⁴

Second, the rate of inflation in the long run is determined by monetary policy, not by microeconomic factors. Again, that is now taken for granted, but much effort was devoted to the imposition of detailed direct wage and price controls in the 1960s and 1970s. Nicholas Kaldor, adviser to Harold Wilson, wrote in 1971 that “It is also far more generally acknowledged – even by Conservative Prime Ministers – that the process of inflation is 'cost-induced' and not demand-induced', with the evident implication that it can be tackled only by an incomes policy”. Not many Whitehall advisers would give that answer today.

Third, in the short run monetary policy does affect output and employment and so has the potential to be an effective stabilisation tool. Reflecting a post-war consensus that monetary policy was rather ineffective, however, the Radcliffe Report concluded that “... there can be no reliance on this weapon [interest rate policy] as a major short-term stabiliser of demand”⁵ It is now accepted that monetary policy lies at the heart of any attempt to stabilise the economy.

The source of monetary policy's influence over output and employment lies in frictions, which mean that prices and wages do not adjust instantaneously to clear markets whenever demand and supply are out of balance. Firms change prices only irregularly in response to changes in demand; wages adjust only slowly as labour market conditions alter; and expectations are updated only slowly as new information is received. Such frictions generate short-run relationships between money, activity and inflation.⁶ The nature of frictions goes right to the heart of the policy debate over inflation targeting. From time to time shocks will move inflation away from its desired long-run level, and the policy question is how quickly should it be brought back to that level. There is no right or wrong answer to that question. Only an analysis of the nature of the relevant frictions tells us what is the “optimal” monetary policy.⁷

³ For many years there was a debate about whether policy was better seen as setting short-term interest rates or determining the monetary base. That is no longer an issue. For some time, the demand for money has been purely demand-determined. As a result, central banks can set the short-term interest rate either to influence real interest rates or to determine the path of the monetary base or a broader monetary aggregate. Money remains at the heart of the transmission mechanism but since its velocity is unstable most central banks use interest rates as their instrument rather than a monetary aggregate.

⁴ Radcliffe report, cmd. 827, p.18-21.

⁵ Kaldor, N. (1971),

⁶ Radcliffe Report (1959) p.177.

⁷ In a deep sense, only a complete understanding of the nature of the frictions makes it possible to decide on the objectives of monetary policy. Woodford (2003) and others discuss the link between that fundamental analysis and the proposition that monetary policy should aim to stabilise inflation and output.

It is why recent academic analysis portrays monetary policy as a “policy reaction function” which describes the reaction of the official short-term interest rate to any possible configuration of economic shocks that might arise in future. For a given model of frictions it is possible to derive the appropriate policy reaction function which most advances the objectives of the policy-makers. Such a reaction function is a state-contingent monetary policy rule. It describes policy in every situation. There are no exceptions and, by construction, the rule does not change over time.

Monetary policy rules have become a major area of research.⁸ Perhaps the most famous is the so-called Taylor rule, named after John Taylor who has just returned to Stanford after serving as Under Secretary at the US Treasury. The Taylor rule implies that interest rates should rise if inflation is above its target and output is above its trend level, and fall when the converse is true. The path along which inflation should return to its desirable long-run level will therefore vary according to the state of the economy.

A key motivation for the study of monetary policy rules was the insight that if economic agents base their decisions on expectations of the future then the way monetary policy is expected to be conducted in the future affects economic outcomes today. Hence it is very important to think about how policy influences the expectations of the private sector. Consider a simple and stark example. Suppose that a central bank managed to control inflation perfectly by responding to all shocks instantaneously. The outcome would be a constant inflation rate. Households and firms would know that potential movements in inflation would never emerge because all future shocks would be instantly offset by changes in interest rates. Interest rates would change with no apparent link to or effect on inflation. To an observer – whether journalist or econometrician – interest rate changes would appear to have little to do with inflation. The central bank would appear to be behaving almost randomly. But that inference would be false. Indeed, if people did expect the central bank to behave randomly, then the behaviour of households and firms would change and inflation would no longer be stable.

This is what I call the Maradona theory of interest rates. The great Argentine footballer, Diego Maradona, is not usually associated with the theory of monetary policy. But his performance against England in the World Cup in Mexico City in June 1986 when he scored twice is a perfect illustration of my point. Maradona’s first “hand of God” goal was an exercise of the old “mystery and mystique” approach to central banking. His action was unexpected, time-inconsistent and against the rules. He was lucky to get away with it. His second goal, however, was an example of the power of expectations in the modern theory of interest rates. Maradona ran 60 yards from inside his own half beating five players before placing the ball in the English goal. The truly remarkable thing, however, is that, Maradona ran virtually in a straight line. How can you beat five players by running in a straight line? The answer is that the English defenders reacted to what they expected Maradona to do. Because they expected Maradona to move either left or right, he was able to go straight on.

Monetary policy works in a similar way. Market interest rates react to what the central bank is expected to do. In recent years the Bank of England and other central banks have experienced periods in which they have been able to influence the path of the economy without making large moves in official interest rates. They headed in a straight line for their goals. How was that possible? Because financial markets did not expect interest rates to remain constant. They expected that rates would move either up or down. Those expectations were sufficient – at times – to stabilise private spending while official interest rates in fact moved very little. An example of the Maradona theory of interest rates in action is shown in Chart 1. It is a “porcupine” chart which shows the Bank’s official interest rate (the repo rate) as the thick black line together with forward interest rate curves at the time of successive Inflation Reports in 2002. Although by no means a perfect measure, the forward interest rate curve provides an idea of market participants’ expectations of future policy rates. During 2002 the Bank of England was able to achieve its goal by moving on a straight line with unchanged official interest rates. But, although interest rates scarcely moved, expectations of future interest rates – as embodied in the forward curve – did move around as the economic outlook changed from an expectation of a swift recovery to worries about a protracted slowdown. And in turn those changes in expected future rates affected activity and inflation. In other words, monetary policy was able to respond by less than would otherwise have been necessary because it affected expectations.

⁸ An excellent example is the recent book by Michael Woodford (2003) which builds on the ideas of the Swedish economist Knut Wicksell one hundred years ago that the key to price stability lies in thinking about the appropriate path for future nominal interest rates.

That pattern is sometimes described as “the market doing the work for us”. I prefer a different description. It is the framework of monetary policy doing the work for us. Because inflation expectations matter to the behaviour of households and firms, the critical aspect of monetary policy is how the decisions of the central bank influence those expectations. As Michael Woodford has put it, “not only do expectations about policy matter, but, at least under current conditions, very little *else* matters”. Indeed, one can argue that the real influence of monetary policy is less the effect of any individual monthly decision on interest rates and more the ability of the framework of policy to condition inflation expectations. The precise “rule” which central banks follow is less important than their ability to condition expectations. That is a fundamental point on which my later argument will rest.

It should be clear that, just as Maradona could not hope to score in every game by running towards goal in a straight line, so monetary policy cannot hope to meet the inflation target by leaving official interest rates unchanged indefinitely. Rates must always be set in a way that is consistent with the overall strategy of keeping inflation on track to meet the target; sometimes that will imply changes in rates, at other times not.

Learning and its implication for monetary policy

The academic literature on monetary policy rules has performed a great service in emphasising the importance of expectations. But there are two basic problems with the use of rules. The first is that the validity of any given rule depends upon the model of the economy that underlies it being true. The second is that the calculation of the rule – or policy reaction function – is extraordinarily complex. Moreover, these two problems interact, in that the complexity of the decision rule is increased enormously when the possibility of learning about the true model is introduced. So although policy rules offer important insights they do not provide a practical guide to decision-making, and it is useful to examine more deeply why that is the case.

No economist can point to a particular model, and in honesty say “that is how the world works”. A crucial difference between economic and, say, meteorological analysis is that in economics there are no natural constants, not even for the natural rate of unemployment. Our understanding of the economy is incomplete and constantly evolving, sometimes in small steps, sometimes in big leaps. The stock of knowledge is not static. So any monetary policy rule that is judged to be optimal today is likely to be superseded by a new and improved version tomorrow. In other words, there is no time-invariant policy reaction function which could describe the policy intentions of a central bank. Rather, monetary policy in practice is characterised by a continuous process of learning embedded, in the case of the Bank of England, in the rounds of meetings and forecasts that are the daily life of the Monetary Policy Committee.

To convince you of how important learning about key economic relationships is to decisions on monetary policy, let me show you two charts which illustrate some of the challenges facing the Monetary Policy Committee. A basic proposition common to most models of the economy is that if demand exceeds the supply capacity of the economy then there will be upward pressure on wage and price inflation. In the labour market supply capacity is often equated, in the long run, with a particular rate of unemployment. Chart 2 plots the unemployment rate against the inflation rate in the UK over the period 1993-2004. It shows the trade off between unemployment and inflation in the short run, also known as the Phillips curve. Unemployment fell from nearly 10% in 1993 to less than 3% in 2004. But – in stark contrast to the earlier post-war period – inflation remained virtually unchanged. How can we explain this phenomenon? Was it because the natural rate of unemployment also fell – perhaps as a result of labour market reforms enacted in the 1980s and 1990s? Or did the Phillips curve become flatter – perhaps because inflation expectations were anchored on the target so that deviations of unemployment from the natural rate generated less pressure on wages and inflation than before? Or was the outcome the result of a chance sequence of shocks that held inflation down?

Chart 3 shows that the slope of the short-run Phillips curve has moved around during the post-war period, apparently in response to changes in the monetary policy regime. In the 1970s labour market pressure was not offset by tighter monetary policy, leading to a spiral of wage and price inflation. The short-run Phillips curve steepened, with larger inflationary consequences of any deviation from the natural rate of unemployment. As monetary policy became more focused on controlling inflation, the Phillips curve flattened in the latter part of the 1980s and 1990s. Such changes in the monetary policy regime can also be detected in the behaviour of inflation over time. Table 2 shows that the persistence of inflation – measured by the estimated explanatory power of past inflation in predicting current inflation – has fallen quite markedly since the inflation target was introduced in 1992. Was this

because the failure of monetary policy to react quickly to an inflationary shock in the 1970s meant that inflation remained high for some time? And has the prompt response of monetary policy meant that movements in inflation more recently have proved short-lived?

The answers to these questions matter for monetary policy. But the economy is continually evolving, and we can never definitively conclude that one answer is right and the others wrong. So learning about changes in the structure of the economy lies at the heart of the daily work of central banks. To describe monetary policy in terms of a constant rule derived from a known model of the economy is to ignore this process of learning. So how should central banks behave in the light of their ignorance? Two approaches have been suggested.

Interestingly, at one end of the spectrum, both Milton Friedman and Robert Lucas argued that policy should be based on a simple rule precisely because of our ignorance. Central banks, in their view, should have limited ambitions and aim simply at steady growth of the money supply – the so-called *k% rule* under which the money stock rises at a fixed rate, *k%*, each year. As Friedman (1968) put it, “Steady monetary growth would provide a monetary climate favourable to the operation of those basic forces that are the true springs of economic growth. That is the most that we can ask of monetary policy at our present state of knowledge”.⁹ The principle of adopting a strategy that takes into account limits to our knowledge is a sound one. But advocates of a rigid *k% rule* argue that we should ignore all other sources of information (estimates of the output gap, for example) and allow any shocks to the velocity of money to feed through to activity or the price level. In practice, experience in both Europe and the US has shown that velocity shocks can be large and few economists now advocate the use of *k%* rules. So committing to a wholly inflexible rule is likely to be neither desirable nor credible. Our knowledge is neither complete nor constant.

At the other end of the spectrum, rational optimising behaviour can, in principle, generate a policy reaction function which takes into account uncertainty about the economy and the process of learning about economic relationships. Such a reaction function would describe how a central bank would respond to any conceivable shock in the future, and explain how estimates of parameter values and the weights attached to particular models would be updated. But even in very simple examples the cleverest economists find the solution of those decision problems almost impossibly complicated. Fully rational optimising behaviour is unreasonably demanding. In the words of Gerd Gigerenzer (2001), optimisation is for “Laplacean demons” not human beings – a reference to an imaginary being that “...could condense into a single formula the movement of the greatest bodies of the universe and that of the lightest atom...”¹⁰

Both approaches, for very different reasons, end up with a monetary policy rule. The simple rule is not credible because we do know some things and we can learn from the past. The complicated rule is not feasible because it places unrealistic demands on our ability to process information. Given the lack of further guidance from economists as to how to make decisions, central banks have often retreated to the position that setting interest rates requires the exercise of unfettered discretion. But this has problems of its own. As has long been recognised, pure discretion does not keep private sector expectations of inflation in line with the desired rate of inflation. If we are to find our way through the minefield between rules, on the one hand, and pure discretion, on the other, we need to think more carefully about the nature of decision-making in a complex world where the central bank and economic agents alike are learning about their environment.

Human beings, including central bankers, are not “Laplacean demons”. Given the constraints on their scarce time, observation suggests that people follow simple rules of thumb.¹¹ These rules of thumb are sometimes described as “heuristics”. The easiest way to understand a heuristic is to imagine a cricket match. The fielder is standing in the deep when the batsman hits the ball somewhere in his direction – see Chart 4. How should the fielder try to catch the ball? One view – the rational optimisation view – is that the fielder either knows, or behaves as if he knows, the laws of physics. Then he could compute the trajectory of the ball, run to the point at which he could catch it (A in Chart 4), and wait for the ball to arrive. This theory of decision-making has testable implications. The fielder will run in a straight line (the solid line FA), and will normally be stationary when making the catch. But that is not how fielders

⁹ Friedman (1968).

¹⁰ Laplace (1995 translation).

¹¹ Todd (2001).

behave in practice. Various empirical studies of baseball and cricket players suggest that fielders follow simple heuristics. For example, they keep their eye on the ball, adjusting their running speed so that the angle of the gaze – the angle between the eye and the ball – remains roughly constant.¹² The heuristic will guide the fielder to the point at which he can catch the ball, without a need to acquire information about variables such as wind speed and direction, spin or the other relevant factors, nor perform complex calculations on those data. But it means that the fielder will run in a slight arc (the dotted line FA) and be moving when the ball arrives. What is instructive about this example is the ability to distinguish empirically between a simple heuristic and fully optimising behaviour, and that the evidence favours the former.

A useful heuristic has two characteristics. It should be *fast* to compute and *frugal* in its data requirements. New heuristics can be adopted when needed. We might think of a “toolbox” of heuristics from which an appropriate choice can be made according to the task that is to be performed. Experimental evidence in laboratory settings shows that some fast and frugal heuristics can be about as accurate as much more data-intensive, optimisation-based methods such as multiple regression.

What are the implications of heuristics for monetary policy? There are two issues. First, although the central bank will try to be as rational as possible in processing all the relevant information, it may well itself use a range of heuristics. For example, in normal circumstances the heuristic “set interest rates such that expected inflation two years ahead is equal to the target” might serve the Monetary Policy committee well. But in other circumstances, say following a large shock, the heuristic might be “bring inflation back to target over a period of more than two years and explain carefully why the heuristic has changed”. The central bank can adapt its particular policy-setting heuristic to changing circumstances and evolving knowledge, so that the policy regime as a whole is robust to changing views about how the economy works.

Second, we do not know whether – and, if so, to what extent – people use heuristics to make real economic decisions. But a central bank should be alert to the possibility of their doing so. Given the importance of expectations, the more the central bank can do to behave in a way that makes it easy for the private sector to adopt a simple heuristic to guide expectations the better. A good heuristic from that point of view would be “expect inflation to be equal to target”. A bad heuristic would be “if inflation is well away from target expect it to deviate further”. We can encourage people to use the first by announcing targets that are quantitative and useful. We can discourage the second by being open and transparent about the reasons for movements in inflation and decisions on monetary policy. If we have no hidden message, then eventually people will stop looking for it.

Rational optimising behaviour is in many situations too demanding, and actual decisions may reflect the use of heuristics. That must be taken on board in the choice of monetary policy strategy. In turn the strategy may affect the heuristic chosen by economic agents. And a good strategy will not only help agents choose a heuristic but will be robust with respect to that choice. Does inflation targeting meet those criteria?

Inflation targeting as a framework which accommodates learning

So far I have emphasised three key points about monetary policy. First, expectations play a fundamental role in the way monetary policy works. As the Maradona theory of interest rates shows, expectations of future monetary policy actions are at least as important as the level at which the official interest rate is set today. Second, our knowledge of the economy is continuously evolving – as the history of the Mais lecture has itself demonstrated. There simply is no unchanging rule, however complex, that can adequately describe the optimal monetary policy strategy. Third, the complexity of optimising behaviour means that central banks need to allow for the possibility that people use simple rules of thumb.

Taken together, these arguments provide a powerful case for inflation targeting. An inflation targeting framework combines two distinct elements: (a) a precise numerical target for inflation in the medium-term and (b) a response to economic shocks in the short term. The inflation target provides a rule-like framework on which the private sector can anchor its expectations about future inflation. As Gordon Brown put it in his Mais Lecture in 1999, “a credible framework means working within clearly defined

¹² To be precise, the angle of gaze remains within a certain range - reported by Gigerenzer and Selten (2001).

long-term policy objectives, maximum openness and transparency, and clear and accountable divisions of responsibility". It is a natural heuristic around which agents can form their expectations. And the discretion in responding to shocks afforded by inflation targeting allows the central bank to adapt its strategy to new information. That is why inflation targeting is sometimes referred to as a framework of "constrained discretion". Following a shock which moves inflation away from target and output from its normal level, there is discretion about the horizon over which inflation is brought back to target. But the exercise of that discretion must be clearly explained and justified in terms of the need, in the words of the remit of the Monetary Policy Committee, to avoid "undesirable volatility in output". The great attraction of an inflation target is that it is a framework that does not have to be changed each time we learn about aspects of the economy such as the velocity of money or the underlying rate of productivity growth, as was the case in the past with frameworks based on targets for money aggregates or nominal GDP growth. It is a framework designed for a world of learning.

The empirical evidence suggests that inflation targeting has helped to confer tangible benefits. One test of whether inflation expectations are well-anchored is the volatility of long-term interest rates. Chart 5 shows the standard deviation of ten-year forward interest rates in the United Kingdom since 1992 and compares it with the figure for the United States. In both countries, volatility rose in the early 1990s. But whereas volatility has been broadly stable in the United States since the mid-1990s, it has fallen steadily in the United Kingdom. In a comparative study of OECD countries, Levin et al (2004) found that inflation expectations were better anchored in inflation targeting countries in the sense that movements in actual inflation were less likely to cause inflation expectations to change. The clarity and simplicity of an inflation target mean that a natural heuristic for the private sector is "expected inflation equals the inflation target".

Inflation targeting is a framework for making and communicating decisions. It is not a new theory of the transmission mechanism of monetary policy. It does not reflect a new understanding of the laws of economics. But, by anchoring inflation expectations on the target, it can alter the transmission mechanism by reducing the persistence of inflationary shocks. And it does so without pretending to commit to a rule that is incredible because it is not expected to last.

The implications of an inflation target for central bank communications are natural enough. First, the clarity of the inflation target focuses attention on the case for price stability which must be made continuously. Second, each forecast must be accompanied by an explanation of the current thinking behind the MPC's views; in essence the "model" underlying the MPC's thinking is changing all the time. Third, there is no point trying to communicate a time-invariant policy reaction function when that does not exist. The regular commentary on its thinking published in its Minutes and *Inflation Reports* is part of a process by which the MPC communicates with the general public. A reputation for communicating openly and honestly about the range of possible outcomes matters, because it makes it more likely that people will continue to listen.

What are the main challenges for inflation targeting in the future? The most immediate stems from its very success. Although it is now widely accepted that there is no long-run trade off between inflation and output, the ability of monetary policy to affect output in the short run means that there is, in principle, a permanent trade off between the **volatility** of inflation and the **volatility** of output, which might be represented by the line AA in chart 6. The choice of a horizon over which to bring inflation back to target is equivalent to choosing a point on this volatility trade off. The striking change, however, is the remarkable improvement in the trade-off that followed the introduction of inflation targeting, as can be seen in chart 6. The volatility of both inflation and output growth were much lower than in earlier periods.

Part of the improvement may lie in the pattern of shocks over the past decade, although the world economy has hardly contributed to that stability. So the challenge ahead is that if a shock, larger than we have experienced recently but not large relative to historical experience, were to move inflation significantly away from target, then inflation expectations might become dislodged from the target. The behaviour of expectations and so the economy as a whole would change. So far there is little sign that the shocks we have experienced have detached inflation expectations from the target, and that is a source of comfort. But the MPC will continue to pay particular attention to the evidence on inflation expectations. Many of the problems of the past resulted from the failure to take action before expectations had started to drift upwards, and the cost of that inaction proved to be high. When the time comes for me to write an open letter to the Chancellor because inflation has deviated by more than one percentage point from target – and it is very surprising that such a letter has not been required in the eight years since the MPC was set up – I will welcome the opportunity to explain how

we expect to bring inflation back to target and over what horizon. Such letters are an integral part of the policy framework, not an indication of its failure.

Conclusions

In this lecture I have advanced three propositions. First, expectations are of fundamental importance to monetary policy. Second, the strategy of policy is more important than any of the individual monthly decisions on interest rates. Third, in designing a strategy be aware of the likely role of heuristics in forming expectations, and so keep it simple.

From those perspectives inflation targeting appears a natural way to conduct monetary policy. And experience of inflation targeting suggests that a managed monetary standard can lead to stability – of both inflation and the economy as a whole – without the straitjacket of a gold standard, currency board or rigid fixed exchange rate target. Inflation targeting anchors inflation expectations, yet allows a flexible response to economic shocks.

Is inflation targeting the last word in monetary policy? Almost certainly not. Twenty-five years from now, I am confident that one of my successors will be able to look back and explain in his or her Mais Lecture the great improvements that took place between 2005 and 2030. But I like to think that the inflation target framework has the ability to serve us well over that period.

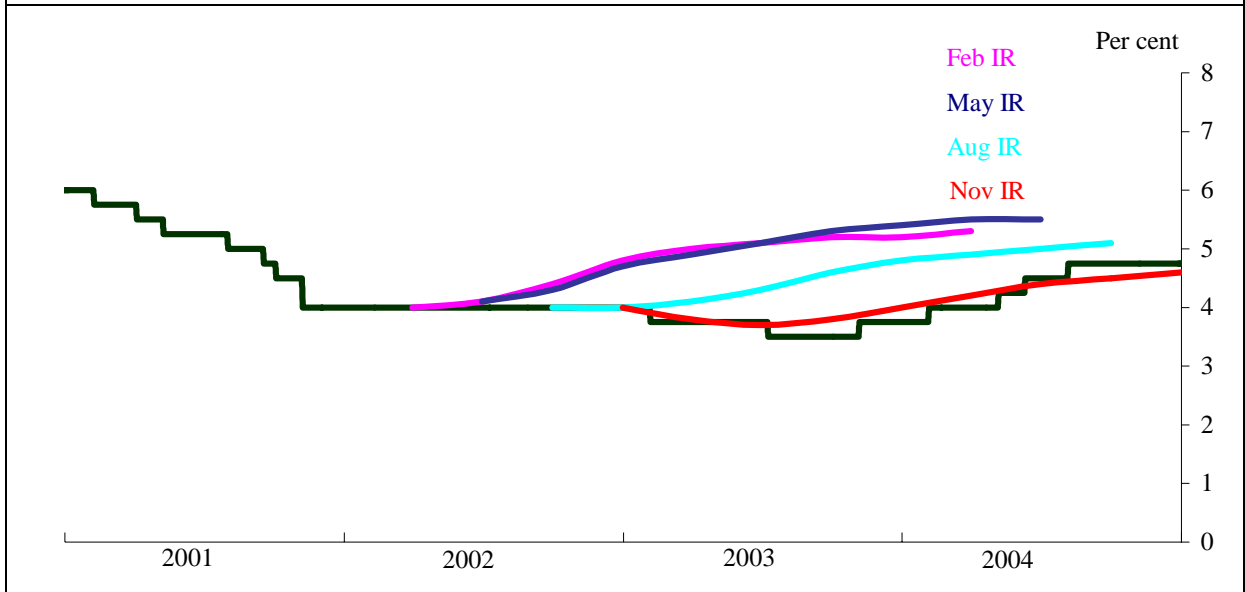
Thirty years ago the theory of monetary policy was ahead of its practice, at least in the United Kingdom. Now I hope that the practice has given the theorists something to think about.

Table 1: Inflation targeting countries

Country	Adoption of Inflation Targeting
New Zealand	Dec. 1989
Chile	Jan. 1991
Canada	Feb. 1991
Israel	Jan. 1992
U.K.	Oct. 1992
Sweden	Jan. 1993
Finland	Feb. 1993
Australia	Mar. 1993
Spain	Jan. 1995
Czech Republic	Apr. 1998
Korea	Apr. 1998
Poland	Oct. 1998
Mexico	Jan. 1999
Brazil	Jun. 1999
Colombia	Sep. 1999
South Africa	Feb. 2000
Thailand	May 2000
Iceland	Mar. 2001
Norway	Mar. 2001
Hungary	July 2001
Peru	Jan. 2002
Philippines	Jan. 2002

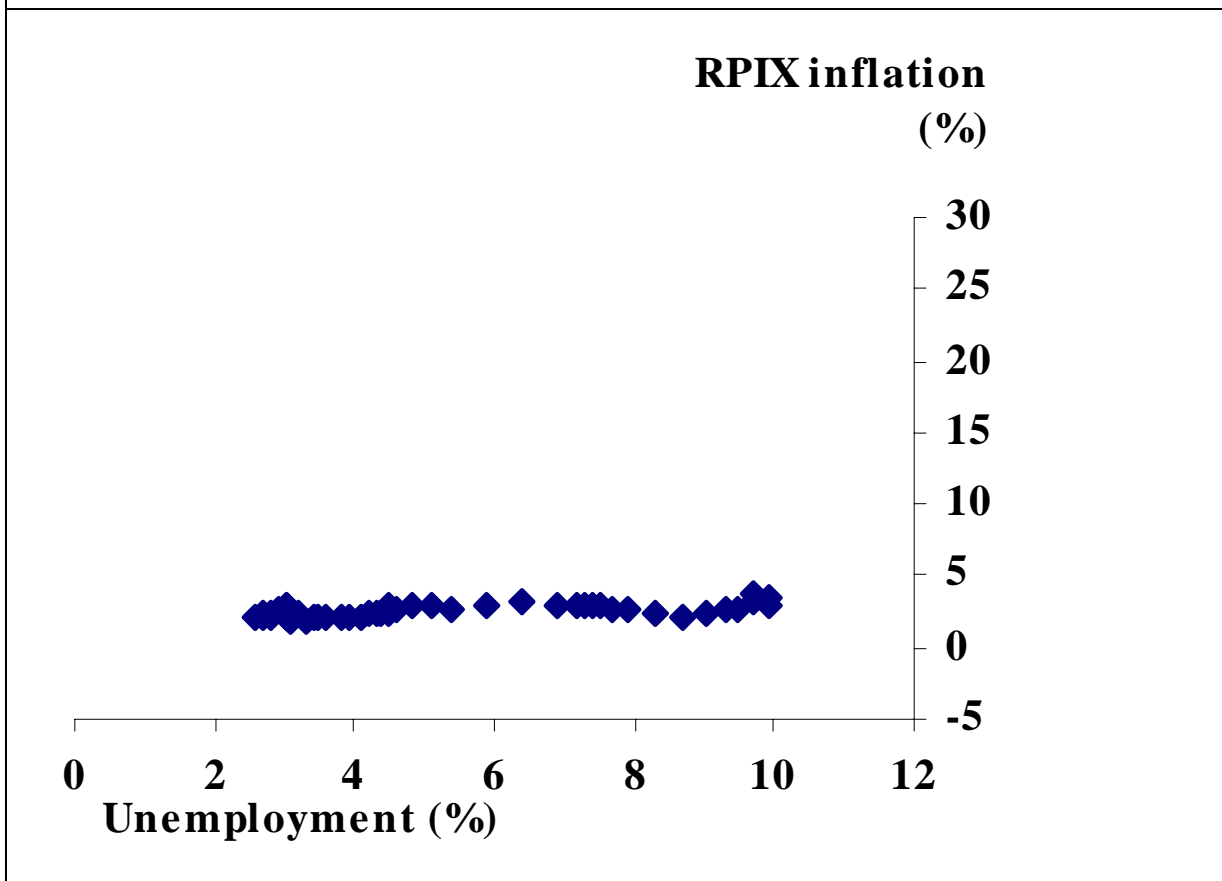
Source: Truman, Edwin (2003), *Inflation Targeting in the World Economy*, Institute for International Economics, Washington DC. The table not only includes current inflation targeting countries, but also Spain and Finland, which have since joined EMU.

Chart 1: The Maradona theory of interest rates in 2002



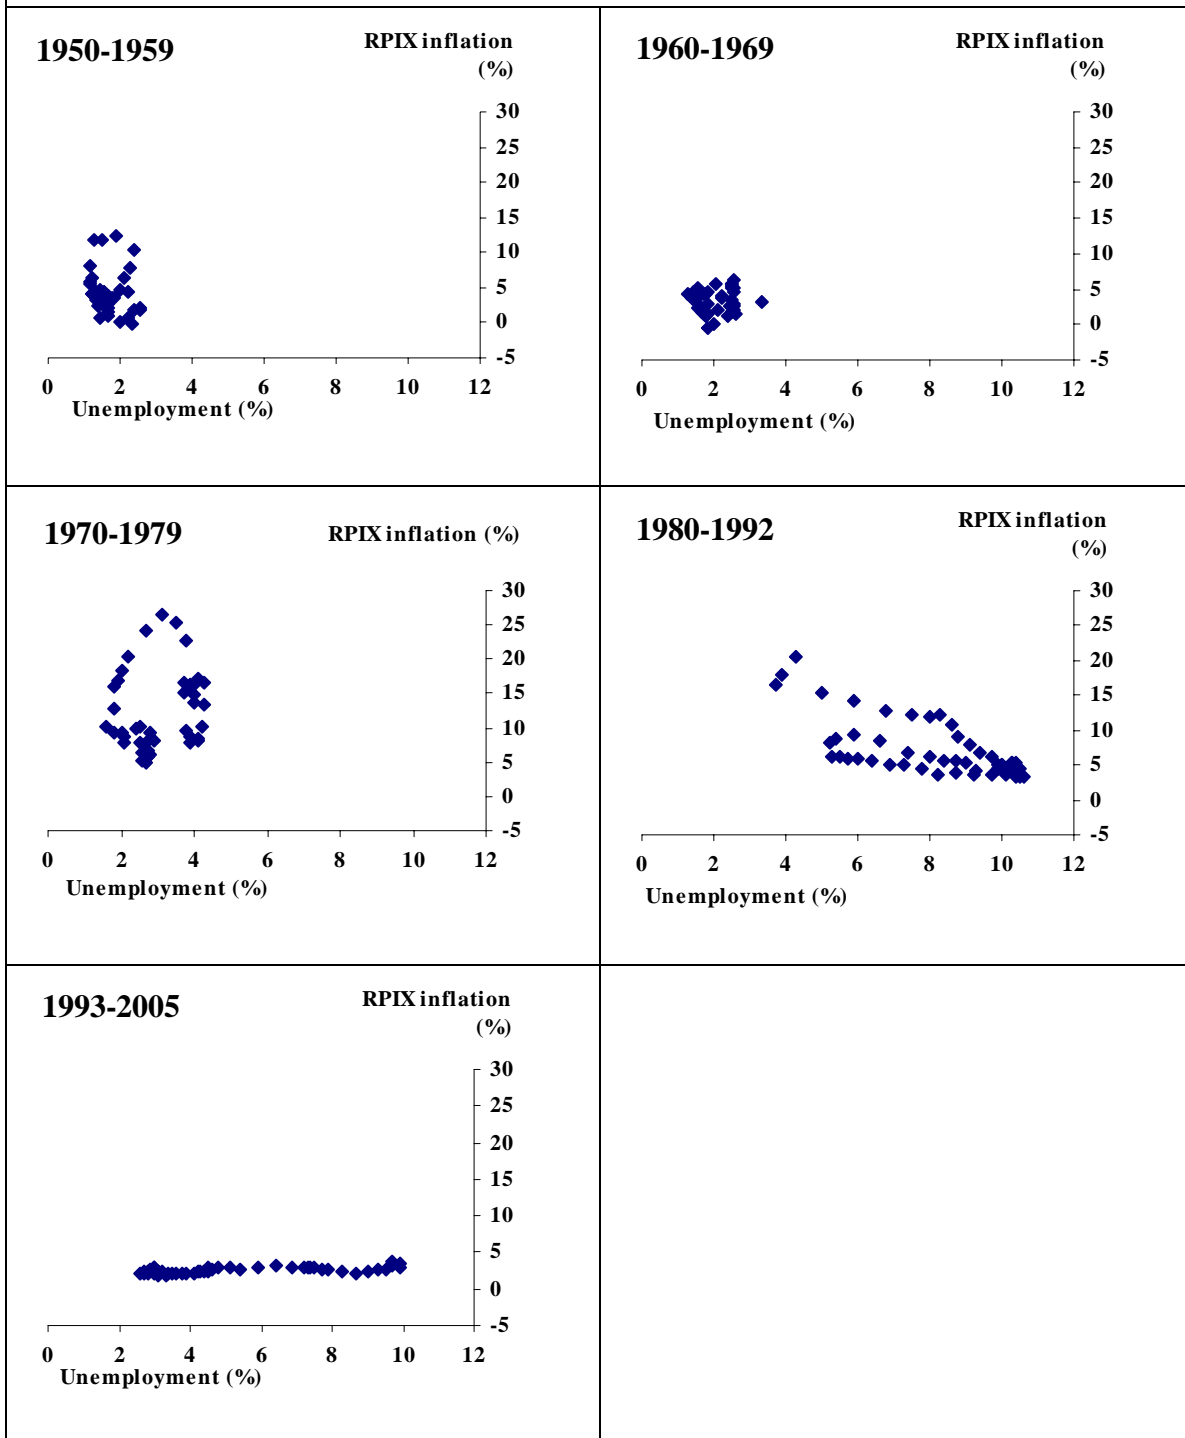
Note: the black line represents the Bank of England official interest rate. The coloured lines represent the market's expectations of future interest rates, as calculated in each of the four *Inflation Reports* published in 2002.

Chart 2: Inflation and Unemployment 1993-2005



Source: ONS. Note: The unemployment rate used here is the Claimant Count measure.

Chart 3: Inflation and Unemployment by Decade



Note: The unemployment rate used here is the Claimant Count measure, published by the ONS from 1971. Unemployment data before 1971 is from Haldane and Quah (1998). The published RPIX series starts in 1976. For observations before 1976, the all-items RPI was used. The RPI series before 1976 did not include mortgage interest payments.

Table 2: The persistence of inflation 1950-2005

	Persistence
1950-59	0.5
1960-69	0.3
1970-79	0.7
1980-92	0.8
1993-2005	0.2

Source: ONS and Bank of England calculations. Note: Persistence in this table is the sum of the coefficients on lagged inflation in a regression of quarterly inflation on a constant and 4 lags. The measure of inflation is RPI before 1976 and RPIX from 1976, seasonally adjusted.

Chart 4: Catching a cricket ball

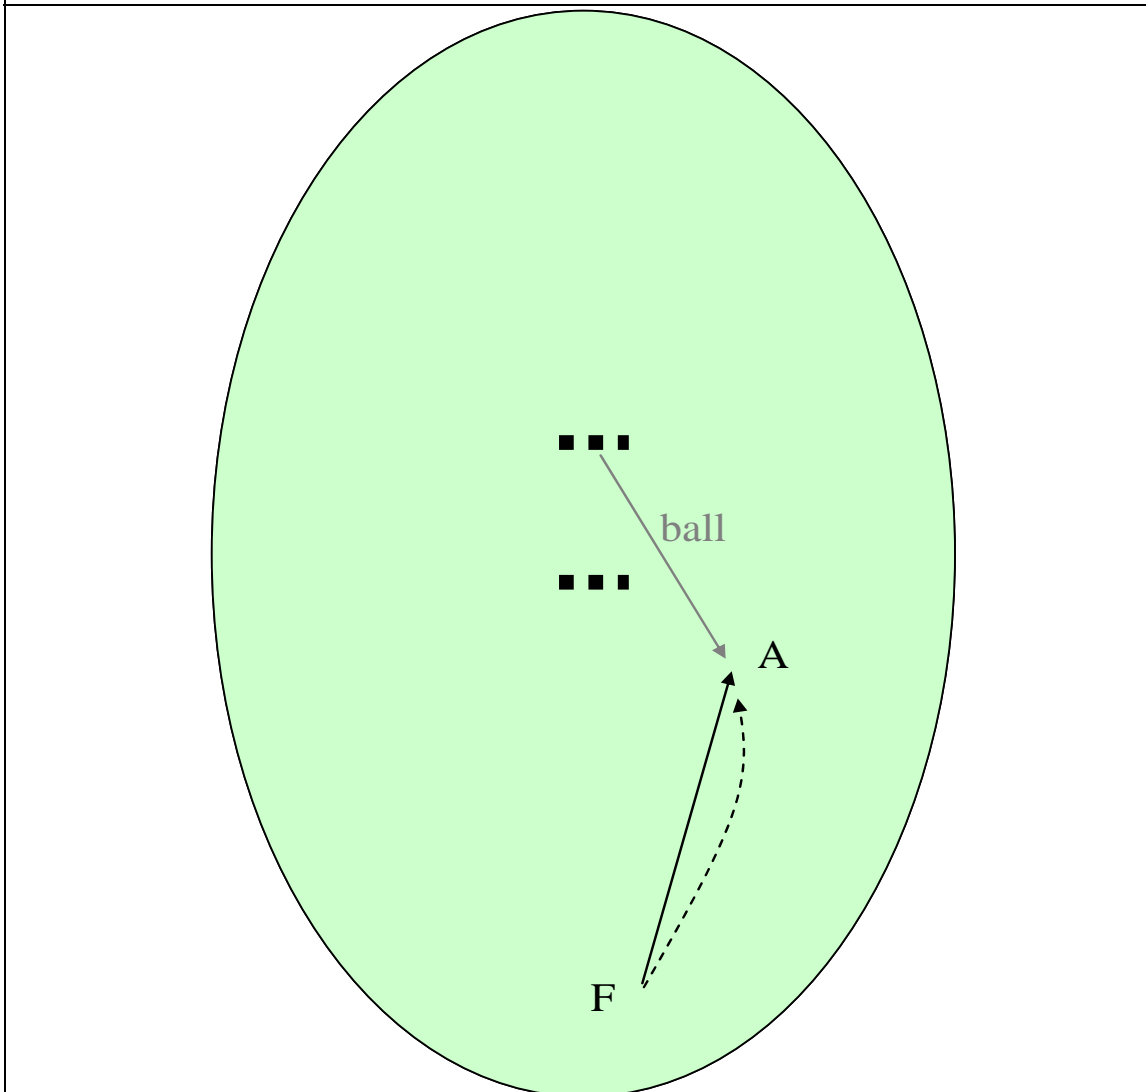
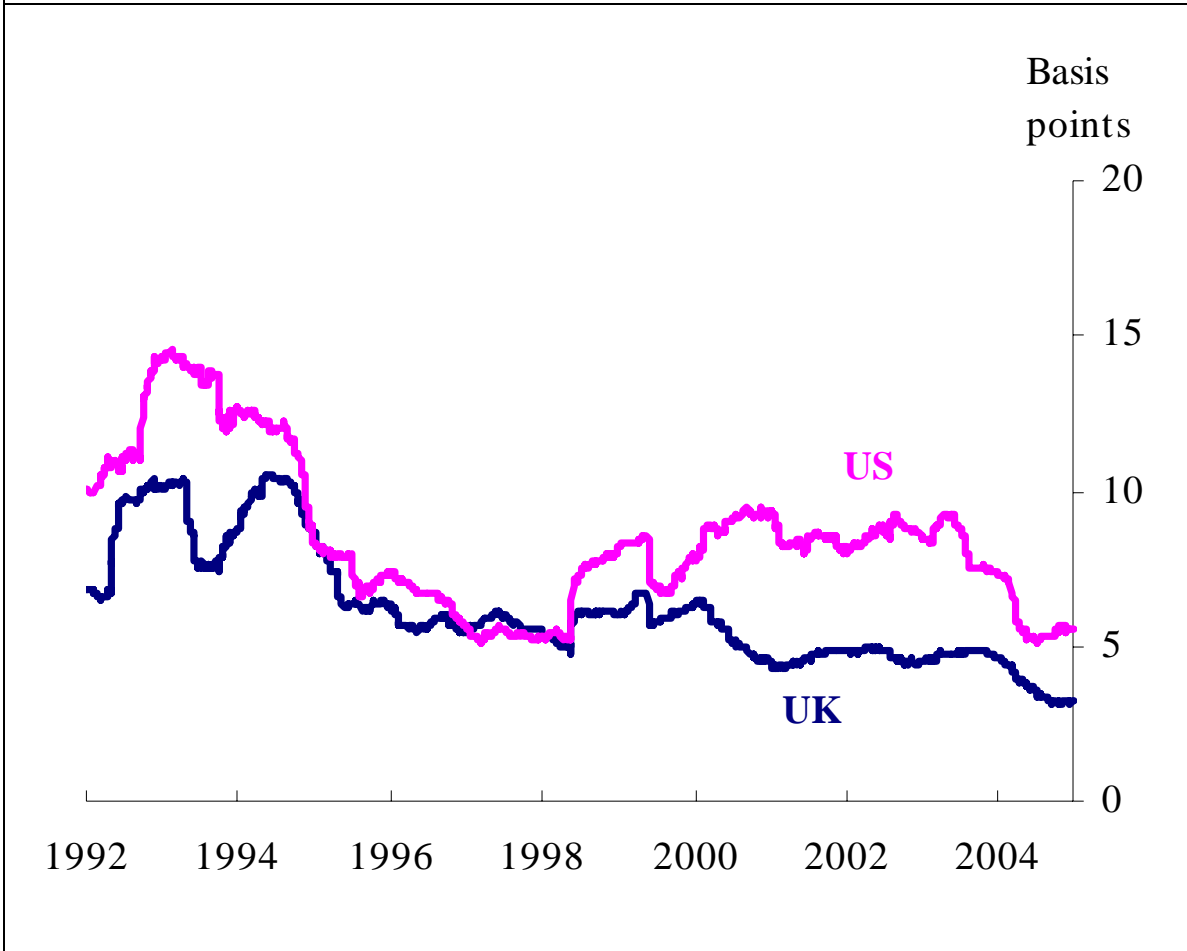
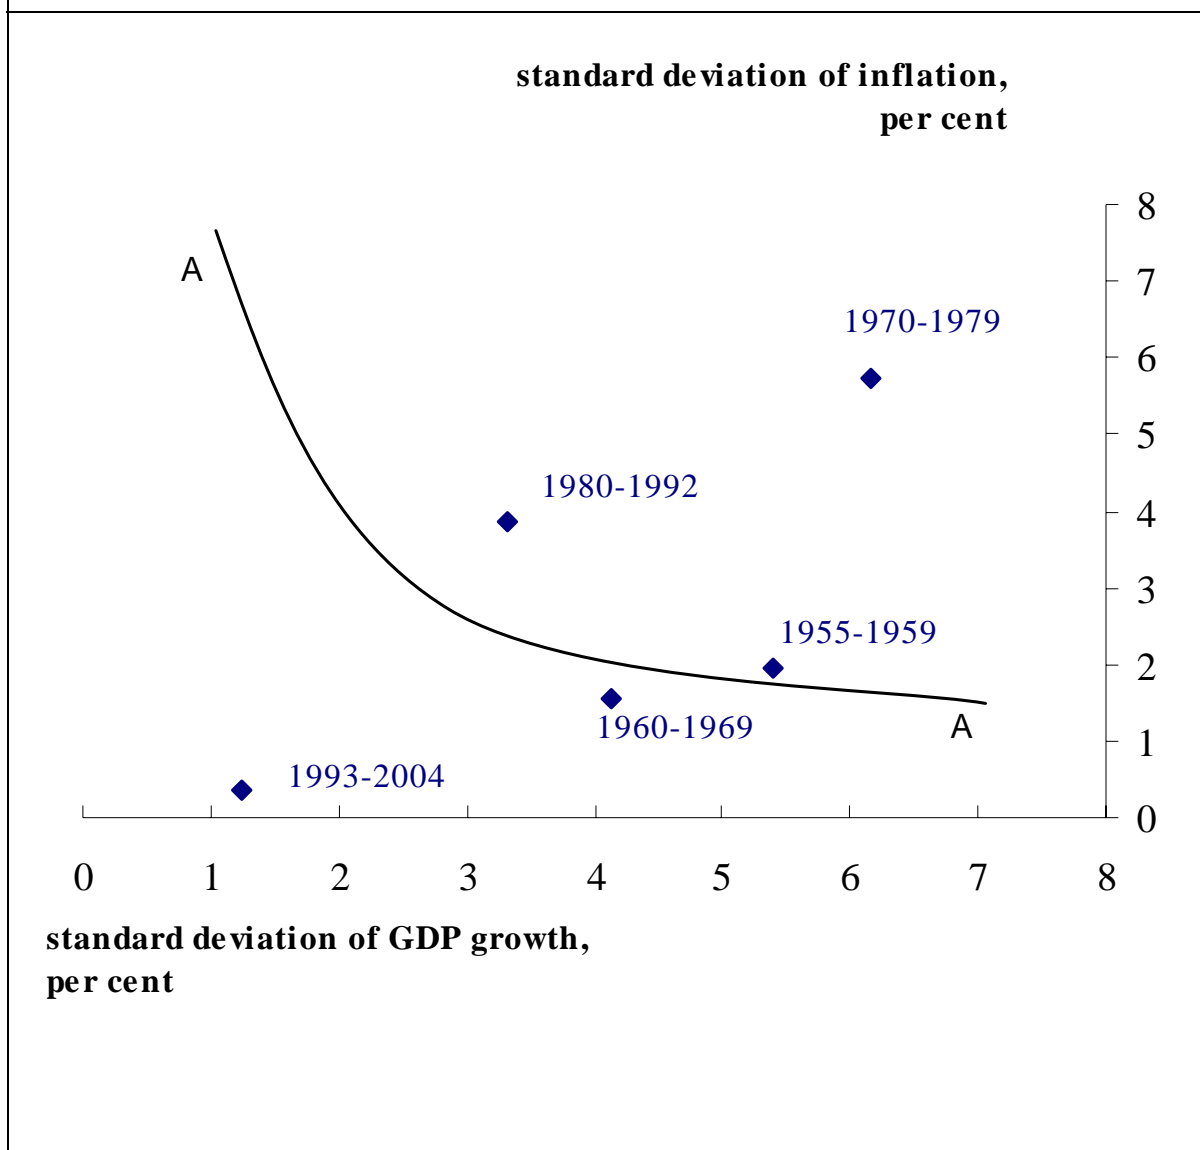


Chart 5: The variability of expected future interest rates, US and UK



Source: Bank of England calculations. Note: variability is calculated as the standard deviation of daily changes in the ten year instantaneous nominal forward rate over a yearly window.

Chart 6: The variability of inflation and output, 1955-2004



Source: ONS and Bank of England calculations. Note: standard deviation of inflation is calculated from quarterly observations of annual inflation; standard deviation of output growth is calculated using annualised quarterly observations of output growth.