

Ghostbusting: The UK Experience of Inflation Targeting

by

**Andrew Haldane
Bank of England**

I. INTRODUCTION

The number of inflation-targeting central banks is now certainly in double-digits. Although only just over a decade old, inflation targeting has clearly grown rapidly in importance as a monetary framework. Looking at prospective inflation-targeting recruits, these numbers may well increase further in the near future.

While initially confined largely to Anglo-Saxon countries, the inflation-targeting net is now widening. Increasingly, it is embracing emerging as well as developed countries. As a result, inflation targeting is now operating in virtually every continent in the world: in Australasia (Australia and New Zealand); in Eastern and Western Europe (the UK, Sweden, Poland and the Czech Republic); in Africa (South Africa); in Latin America (Chile and Brazil); in North America (Canada); and (at least prospectively) in Asia (Korea, Thailand and maybe in time Japan).

The two largest currency areas—the euro-area and the U.S.—have so far bucked this trend. But there has been no shortage of academic and policy debate on the merits of these countries shifting from their combination strategies—whether inflation-cum-monetary targeting, or inflation-cum-output targeting—toward formalized inflation targeting (see, for example, Bernanke et al (1999)). Time will tell.

This paper considers some of the issues that are central to the design of an inflation-targeting regime. The particular issues the paper discusses are: the institutional framework for monetary policy; specification of the inflation target; dealing with monetary transmission lags; dealing with (in particular) inflation uncertainties; dealing with output objectives, in particular following supply shocks; transparency about monetary policymaking; and dealing with the exchange rate.

The paper discusses these issues through the particular lens of the United Kingdom's experience over the last eight or so years. No one monetary framework should be a rigid template for other countries. But UK experience is perhaps illustrative along several dimensions. And one theme runs through the various technical issues discussed below: "ghostbusting." Arguably, all successful monetary regimes ought to embody some ghostbusting features. By that I mean they take seriously the need to be pre-emptive in setting monetary policy, offsetting incipient inflationary pressures. Like ghosts, these pressures will be invisible to the general public at the time policy measures need to be taken. Claims of sightings will be met with widespread derision and disbelief. But the central bank's job is to spot the ghosts and to exorcise them early. A successful monetary policy framework is ultimately one in which the general public is not haunted by inflationary shocks.

II. THE INSTITUTIONAL FRAMEWORK FOR INFLATION TARGETING

There seems to be an increasing consensus across central banks about the necessary institutional conditions for the effective pursuit of price-stability. One such necessary condition is a degree of operational autonomy on the part of the central bank in the course of setting monetary

policy. In the United Kingdom, such autonomy has only recently been granted and enshrined in statute in the Bank of England Act 1998. This Act confers instrument-independence on the Bank of England, though the goals of policy continue to be set by the government. There is, in the jargon, goal-dependence but instrument independence. Among these goals, price-stability is given primacy.

Monetary policy decisions in the UK are made on a regular monthly cycle by a nine-person Monetary Policy Committee (MPC). The MPC comprises 5 “insiders” to the Bank of England (the Governor, two Deputy Governors and two of the Bank’s Executive Directors) and 4 “outsiders,” with prior technical expertise in the monetary policy area. Seven of the nine members are appointed by government, two by the Bank itself. The minutes of the MPC’s monthly deliberations are published, with a two week lag. The minutes disclose the voting patterns of each of the individual MPC members. The discussion in the minutes is unattributed, though the minutes provide a detailed account of the analysis discussed by the MPC in the course of reaching its monetary policy decisions.

Though operationally independent in their choice of the interest rate, the MPC are held to account for their decisions in a variety of ways: through the published minutes of their monthly meetings; through published quarterly Inflation Reports; through appearances by MPC members before Parliamentary committees (including after the publication of the Inflation Report); through the Bank writing “Open Letters” to the Chancellor in the event of the inflation target being breached by one percentage point in either direction (discussed further below); and through an “override” (of the Bank’s interest rate decision) clause, which allows the Chancellor to exercise interest rate control “in extreme economic circumstances.”

This institutional architecture, and the inflation-targeting regime more generally, seem to have yielded some tangible benefits. Credibility, as measured by inflation expectations relative to target, have improved throughout the period of the regime. Chart 1 shows a set of term structures of inflation expectations on a variety of dates, derived from estimated yield curves for nominal and real bond yields in the United Kingdom. Immediately following the UK’s exit from the Exchange Rate Mechanism (ERM) in September 1992, inflation expectations were between 5 percent and 7 percent at maturities 10 to 20 years ahead—well above the inflation target at the time of 1-4 percent. By April 1997, five years into the regime, inflation expectation had ratcheted down to just over 4 percent. A credibility gap remained but it had narrowed markedly. The announcement of operational independence for the Bank of England in May 1997 caused a further fall in inflation expectation across all maturities, by around 50bp. Independence delivered an immediate credibility bonus. By the end of 1998, inflation expectations at all maturities along the inflation term structure were around the UK’s 2.5 percent inflation target. They have remained at that level since then.

What lessons do we draw? That the improvement in the credibility of the inflation-targeting regime has been gradual, taking seven or eight years to establish. That the institutional framework for monetary policy in the UK has helped buttress this credibility—for example,

through instrument-independence for the Bank of England. But that inflation expectations do now appear to be anchored at the inflation target, all the way along the term structure. The credibility bonus of inflation targeting was no windfall; it was earned rather than endowed. But the inflation target and the accompanying institutional infrastructure is now serving an important nominal anchoring role.

III. SPECIFICATION OF THE INFLATION TARGET

In the UK, the inflation target itself is affirmed on an annual basis by the government. The price-stability objective is, however, enshrined in statute and the target itself is timeless. Currently, the UK has a point target of 2.5 percent. This raises two technical questions about the specification of this target: why 2.5 percent?; and why a point target?

The reasons for choosing 2.5 percent are manifold. First, there are well-known measurement biases in the UK's CPI, as there are in all other countries. In the UK, we estimate these biases to be around 1 percent on an annual basis (Cunningham (1996)), again in line with other developed countries. Second, an inflation rate of 2-3 percent is in line with the current inflation norm in other developed countries, including the other inflation-targeters. Third, and importantly, 2-3 percent seems to be close to the general public's preferred inflation rate in the UK. Chart 2 shows the distribution of survey responses to a Bank of England questionnaire which asked the public to suggest their preferred rate of inflation. As the chart shows, almost all respondents wanted inflation to lie above 0 percent and below 5 percent. And the vast majority suggested an optimal inflation rate of around 2-3 percent. Public acceptability is an important factor in ensuring the ongoing legitimacy of any country's monetary framework. Inflation targeting is no exception.

Many emerging markets are of course coming from a different inflationary background than that in the UK, though in many cases recent inflation rates have fallen to close to UK levels. An interesting question in this context is: what benefits are conferred by pushing down inflation that few extra percentage points? Might not these benefits be outweighed by the disinflationary costs?

On the face of it, there appears to be some academic support for this proposition. It is, for example, difficult to find convincing evidence of inflation having a damaging effect on growth at rates of inflation below 10 percent (Barro (1995)). There is, however, evidence of inflation volatility having an effect on growth at rates of inflation below 10 percent (Judson and Orphanides (1996)). Moreover, inflation need only adversely affect the level of GDP—not necessarily its growth rate—for it to have a significant welfare cost. A particularly interesting example of the size of such costs, even at low inflation rates, has recently been provided by Feldstein (1997). Feldstein considers the welfare costs of inflation's interaction with the unindexed capital income taxation system in the U.S. Feldstein estimates the GDP-equivalent welfare gains from reducing inflation from 2 percent to zero in the U.S. to be around 1 percent of GDP. Similar-sized numbers—sometimes a littler larger, sometimes a little smaller—have recently been found for the UK, Spain, Germany, New Zealand and elsewhere (Table 1).

Moreover, the benefits of low inflation are permanent, whereas the costs of disinflation are transient, provided we believe in a vertical long-run Phillips curve. So the net present welfare value of even a small reduction in inflation can be quite large. For example, assuming a 5 percent discount rate and trend growth of 2.5 percent, Feldstein's estimates would deliver a net present value welfare gain of anywhere between 10-60 percent of initial GDP for developed countries. That is much larger than any plausible estimate of the output cost of a two percentage point disinflation.

Turning to the issue of inflation target points versus ranges, there is clearly no consensus on this issue across the existing inflation-targeters. For example, the UK, Brazil, Chile, Finland, Sweden and Australia all have point targets; whereas Canada, New Zealand, Israel, and the Czech Republic all operate with inflation target ranges.

There are several reasons for the choice of a point target in the UK. First, it serves as a clear ex-ante fix-point for monetary policy decision-making. There is a danger that an inflation range becomes a "band-of-indifference," within which the authorities do not adjust policy. If this is the case, then the probability of breaching the band becomes high and inflation will spend long(er) periods deviating from the target mid-point. An inflation target needs a clear trigger, above or below which policy will respond. A point target makes transparent that trigger.

Second, a point inflation target can also serve as a useful fix-point for private sector agents' inflation expectations. The UK provides a good case study here. It was striking, for example, that inflation expectations in the UK became "stuck" at the top of the UK's 1-4 percent inflation range between 1992-96. There was, in the jargon, "range bias." More recently, as the UK has moved from an inflation band to a 2.5 percent point target, expectations have ratcheted down to that point (Chart 3). A third benefit of a point inflation target is that it makes transparent the symmetry of monetary policy actions once inflation is at its long-run target. An inflation-targeting framework is as much a safeguard against deflation as inflation, which is why some people have recently advocated an inflation target in Japan. A point inflation target insures equally against the output costs of both inflation and deflation.

Using the inflation target as a reference point for expectations is also important during the transition to low inflation. The target then serves as a means of guiding downwards inflation expectations over time. That, for example, was the way the inflation targets in the UK, Canada and New Zealand operated when they were first put in place, with a downwards transition path built in. It is also the way that inflation targets are currently being operated in Israel, the Czech Republic, and Brazil. There is, however, an important difference in how an inflation-targeting regime ought to operate during the period of disinflationary transition. Along this path, an asymmetric—or "opportunistic"—approach to monetary policy is more appropriate. What this means in practice is that adverse inflation outcomes are still vigorously offset through monetary policy; but that favorable inflation shocks are accommodated. The reason for this asymmetry in response is that reflating the economy following a favorable inflation shock would mean inflicting a further disinflation on the economy at some later stage. Rather than do this, it may be better to

pocket the lucky inflation shock—to seize the opportunity—and accept temporary over-achievement of the inflation target. Many countries in recent years have found themselves needing to behave in this asymmetric fashion, most recently Israel and the Czech Republic.

There is an ex-post as well as ex-ante dimension to the specification of the UK's inflation target. In particular, there is a fluctuation margin of +/-1 percentage point around the UK's inflation target. This is meant to accommodate inevitable ex-post inflation variability. Breaches of this fluctuation margin have to be accounted for in an open letter from the MPC to the government. This ex-post/ex-ante dimension to the specification of the inflation target in the UK is a useful way of helping sharpen ex-ante policy actions and public inflation expectations, while at the same time recognizing the inherent ex-post difficulties of inflation-control.

IV. DEALING WITH MONETARY TRANSMISSION LAGS

There are well-known “long and variable lags” between changes in the instruments of monetary policy and their ultimate impact on output and inflation. For example, in some developed countries the maximum marginal impact of a change in short-term interest rates on inflation comes after around two years. In transition and emerging economies the lag is probably shorter, for example because of a greater degree of price flexibility in these countries. These lags in policy do mean, however, that inflation-targeting needs to have a forward-looking dimension. Responding to past inflation would mean that policy was always acting too late to prevent inflationary pressures taking hold.

It is for this reason that, in some inflation-targeting countries, monetary policy is based explicitly around inflation forecasts. In effect, there is “inflation-forecast-targeting”—or ghostbusting. Those countries which explicitly target inflation forecasts include Canada, New Zealand, the United Kingdom and Brazil. In the UK, for example, inflation forecasts up to two years ahead are published in the Bank of England's quarterly Inflation Report. These forecasts can be thought of as the intermediate target for monetary policy (see Svensson (1997), Haldane (1997)). For example, the Bank of England's “policy rule” can be thought to take the generic form:

$$(1) \quad E_t \pi_{t+2} / i_t = \pi^*$$

That is, interest rates today (i_t) are set at such a level that expected inflation two years ahead ($E_t \pi_{t+2}$) is in line with the inflation target (π^*). Policymaking in practice, of course, is never quite that mechanical. But as a pedagogical device, this approach of “inflation-forecast-targeting” has some clear attractions, as a response to the transmission lags in monetary policy. It also means that publication of the inflation forecast becomes of paramount importance, in facilitating outside agents' understanding and monitoring of the authorities' policy rule.

There are at least two questions left open by the operation of an inflation-forecast-targeting policy rule such as equation (1). First, how are the forecasts themselves produced? (What is the ghost-detecting machinery?) And second, what determines the two-year-ahead

forecasting horizon for policy? On the first question, the starting point for the quarterly forecasting round is a core, small structural macro-model (see Bank of England (1999)). This is a conventional open-economy sticky-price IS/LM model. The outputs from this core model are augmented with various pieces of “off-model” information. This off-model information includes simulations and forecasts from other models (ranging from pure time-series models to theoretical general equilibrium models (again see Bank of England (1999)); the MPC’s subjective judgement on various issues; survey data; various financial market information etc. This information and analysis is important to understand and embody when monitoring inflationary dynamics, but is difficult to encompass in a single small model framework. Importantly, the published inflation forecasts reflect and embody the views of the MPC—they are MPC rather than staff forecasts.

A forecasting model—or, more accurately, models—is important to the operation of inflation targeting in the UK, just as it is in New Zealand, Canada, Sweden and some other countries. But the importance of a model for forecasting purposes is easily over-emphasized. There are other than model-based techniques for arriving at measures of inflation expectations. Survey and financial-market based measures of inflation expectations can play a surrogate role. In Israel, for example, market-based measures of inflation expectations play a key role in the setting of monetary policy. So while a well-developed suite of macromodels is a useful tool to have in the kit, it is by no means an essential prerequisite for the effective implementation of inflation targeting. Models are desirable (certainly), but are not necessary and most definitely not sufficient conditions for effective inflation-forecast-targeting.

On the choice of inflation forecast horizon, two years is around the horizon which the Bank has found to be optimal using model-based policy simulations. Chart 4 shows the results of a typical such simulation. The dotted line in both panels shows the “optimal control” response of inflation following a one percentage point inflation disturbance in the first period. Inflation returns to target after roughly 8-10 quarters. The first and second panels show the results from using, respectively, horizons which are “too short” (2-3 quarters) and “too long” (5 years). In the first case, policy is forced to do “too much, too soon” and itself generates secondary inflationary cycles. In the second case, the disinflation path is too gradual—policy is doing “too little, too late”—and inflation as a result spends a protracted period away from target.

The precise optimal forecasting horizon depends on the nature of the shocks to the economy; the authorities' relative inflation/output preferences (discussed further below); and the dynamics of the monetary transmission mechanism. But taking illustrative UK values of these parameters, two years seems to be roughly about the right horizon. Similar forecast horizons are targeted in Canada and New Zealand. For the emerging economies, transmission lags are shorter but the economy is more prone to supply shocks. The net effect of these two offsetting influences upon the optimal inflation forecasting horizon would require careful empirical evaluation.

V. DEALING WITH UNCERTAINTY

Basing policy around inflation forecasts is not without its costs. Inflation forecast errors are substantial. The historical mean absolute error of UK inflation forecasts one-year-ahead is, for

example, over one percentage point (Haldane (1997)). In the UK, there are two approaches to dealing with these forecast uncertainties—again, one ex-ante, the other ex-post.

The ex-ante dimension amounts to making clear up front the full extent of the uncertainties attaching to any inflation forecast. The way this is done at the Bank of England is by publishing a “fan chart” for inflation over the next two years—a full probability density function for future inflation. An example of this is shown in Chart 5, taken from the Bank’s November 1999 Inflation Report. There is a 10 percent probability of inflation lying in the deepest-shaded area. The progressively lighter-shaded areas fanning out on either side of the deepest-shaded area are 10 percent probability contours. The entire shaded area is meant to cover 90 percent of the distribution of inflationary outcomes up to two years ahead.

The benefits of presenting an inflation forecast in this way are essentially four-fold. First, publishing a forecast distribution means that the general public are not focussed on a single point inflation forecast. The central bank is thereby not open to criticism for having got the forecast “wrong” when, given shocks, this is an inevitable part of the forecasting process. Second, the distribution quantifies the extent of likely forecast uncertainties—how likely it is that inflation will miss its target. Third, the distribution also embodies asymmetries. It allows inflationary risks at different horizons to be unbalanced. For example, in Chart 5 near-term risks to the inflation projection were on the downside, while over the medium-term they were on the upside. This feature is important because, often in policymaking, the distribution of the underlying variables may be believed to be asymmetric. For example, asset prices may be thought to be under or overvalued or external risks to world growth may be unbalanced etc. The fan chart allows those unbalanced risks to be embodied explicitly in the inflation projection and hence in policy decision-making. Fourth, and related to all of the above, a quantified distribution allows policy to be exercised in an explicitly probabilistic fashion. In an inflation-forecast-targeting framework, statements such as “inflation will be at target x periods ahead” make no sense, given inflation shocks. The most that can be said is that “there is an x percent chance of inflation lying between y-z percent;” or “there is a z percent chance of the inflation target being breached.” The published inflation distribution allows such an ex-ante probabilistic quantification. The fan chart approach—using a slightly different methodology—is also used in Sweden, Norway, and Brazil.

There is an ex-post—accountability rather than transparency—dimension to dealing with inflation control errors in the UK. An independent central bank (the agent) needs to be able to explain deviations from the target to the government and the public (the principals). In the UK, if inflation deviates from the inflation target by more than one percentage point in either direction, then the MPC have to write an open letter to the Chancellor. So far, inflation has yet to breach this range. But in that event, the open letter system is an accountability device, ensuring target misses are properly explained and appropriate remedial policy action is taken. The open letter has to contain three elements: an explanation of why the deviation from target has occurred; an explanation of what the MPC intends to do about it; and a statement of the time horizon over which inflation is to be returned to target. It is through the third of these aspects—the choice of

time horizon—that output objectives come to be factored into policy choices under an inflation-targeting framework. We now turn to a discussion of that issue.

VI. DEALING WITH OUTPUT

Although the centerpiece of an inflation-targeting regime is the inflation target itself, this does not mean that real objectives—for example, for output and employment—are disregarded in the setting of policy under such a regime. The Bank of England Act 1998, for example, states that the Bank's objectives shall be: "(a) to maintain price stability; and (b) subject to that, to support the economic policy of the government, including its objectives for growth and employment." This is consistent with the statutes of the European Central Bank; and the Humphrey-Hawkins Act in the U.S. has of course explicitly dual objectives of price-stability and maintaining full employment.

But how is this concern for output and employment objectives made operational within an inflation-targeting regime? The first point to make is that, because the "output gap" is usually seen as a useful predictor of future inflation, basing policy on expectations of future inflation means that output and employment are already implicitly entering inflation-targeters' policy rule. In other words, a policy rule such as equation (1) could be rewritten to have the output gap on the right-hand-side instead of expected inflation (see, e.g., Svensson (1997)). The absence of output terms from an inflation-forecast-targeting rule may be more apparent than real.

Second, in a world where there were only demand shocks, hitting an inflation target while simultaneously smoothing output relative to trend would pose few dilemmas for monetary policymakers. Examples here would include shocks to fiscal policy, to external demand and to investors' "animal spirits." In those situations, inflation is likely to end up above target at just the same time as output is above target (or vice-versa). So to smooth both output (relative to trend) and inflation (relative to target) the right response is to tighten policy (or to loosen it in the event of an opposite shock). There is no necessary conflict of output/inflation objectives.

Third, supply shocks pose an altogether different problem, however—for example, shocks to the NAIRU, the terms of trade or indirect taxes. These shocks will tend to shift output and inflation in opposite directions, and so pose a dilemma for policymakers. Should policy be tightened or loosened? Existing inflation-targeting regimes well recognize these supply shock difficulties. They are likely to be especially acute in small open economies (where terms of trade shocks will be more prevalent) and economies undergoing structural transition. So dealing with supply shocks is likely to be an issue of considerable importance for all of the existing inflation-targeters, but particularly the emerging economies.

Approaches to dealing with supply shocks differ across the inflation-targeters. In each case, however, the aim is to accommodate the first-round effects of the shock, thereby cushioning any adverse effects on output of adhering to the inflation target. In New Zealand, for example, the Policy Targets Agreement allows "significant" supply shocks from a pre-specified list to be exempted from the measure of inflation. Another approach is to allow greater inflation variation

around the inflation target before remedial policy adjustment takes place. Inflation target bands, the type of which operate in Canada, allow this cushioning. In both of these cases, the intention is to allow short-run accommodation of the supply shock, while at the same time ensuring inflation expectations remain anchored over the medium term.

A third approach to supply shock accommodation, and the one employed in the UK, is to use the flexibility built into the inflation forecast horizon: that is, by allowing inflation to return to target over a longer horizon when large supply shocks strike. Using policy simulations, it is possible to trace out an inflation/output variability trade-off defined in terms of different inflation forecast horizons. An example is given in Chart 6. The trade-off is L-shaped, so extreme forecast horizons are unlikely to be optimal. The optimal point on the trade-off, based on this UK calibration, seems to be roughly around 6 quarters, which is broadly in line with the inflation forecast horizon used by the Bank of England when setting policy.

The same simulation work has shown that, by judicious choice of forecast horizon, inflation-forecast-targeting regimes can achieve as much by way of output smoothing as policy rules which target output explicitly (Batini and Haldane (1998)). The vehicle through which the Bank of England exercises this leverage over the inflation target horizon is the open letter system. Reflecting its concern about the profile for output as well as inflation, the Bank of England publishes in the Inflation Report a “fan chart” for output growth as well as inflation.

There is again a key distinction to be made, however, between countries in a stationary inflation state and those still in the process of disinflationary transition. There is less scope for accommodation of inflation shocks during the process of disinflationary transition. In these circumstances, positive inflationary impulses which are not offset risk damaging credibility and stoking-up inflationary expectations. The same is not, however, true of negative inflation shocks which lower inflation. During the disinflationary process, these negative shocks should be fully accommodated and locked-in. In other words, during disinflationary transition, shocks should be dealt with “opportunistically”—the favorable shock presenting an opportunity which ought to be seized through accommodation, but the unfavorable shock not being accommodated in this way. As well as being a ghost-buster, the inflation-targeting central banker needs also to be an opportunist. Such behavior is not designed to endear central banks to politicians or to the general public. It looks inconsistent. But the logic is clear, at least during the period of reputation-building.

VII. TRANSPARENCY

A defining feature of inflation-targeting countries is the extent of their monetary policy transparency. In some ways this is also a necessary feature, because the policy rule under this regime—equation (1)—is undoubtedly a complex one. For example, without publication of the inflation forecast and some notion of its origin, it would be difficult for the public to monitor whether monetary policymakers were adhering to the rule. Credibility may be damaged. That is one reason why inflation forecasts are published in the UK, New Zealand, Brazil and elsewhere.

There are a variety of other ways in which monetary policy transparency is enhanced in the UK. The minutes of the monthly MPC meetings are now published with a two week lag; and the Bank's Inflation Report is produced on a quarterly cycle. These publications give an insight into the analysis underlying the Bank's inflation outlook; and the way in which this analysis in turn affects the distribution of future inflation and, ultimately, the MPC's decision-making. Although the UK is perhaps at one end of the transparency spectrum, it is significant that all the inflation-targeting central banks—indeed, some non-inflation targeters too, such as Norway—are now publishing inflation or monetary policy reports on a regular cycle.

It is worth considering some of the benefits that enhanced transparency confers. There are both internal and external benefits. Internally, exposing policy analysis to outside scrutiny acts as a powerful incentive to ensure this analysis is of the highest possible standard. It exposes internal thinking to a powerful external discipline. Externally, transparency serves as a useful accountability device—a political-economy benefit. It simplifies the principal/agent monitoring process. Transparency also, however, potentially delivers macroeconomic benefits. For example, it in principle ought to increase the predictability of the authorities' reaction function, and to help stabilize the yield curve.

To understand this latter benefit, take a limiting case. Imagine a world in which the policymakers' rule was perfectly predictable and the authorities' inflation target was fully-credible. We would then see the yield curve moving around in response to macroeconomic “news,” to the extent that this news affected the central banks' inflation forecast and hence the settings of its policy instrument. But we would not expect to see any response from the yield curve in response to monetary policy changes. These policy changes would contain no extra information; they would be perfectly anticipated because the policy rule was fully transparent.

There is evidence of just this having happened—if on a less dramatic scale—in the period since the introduction of the UK's inflation target. Some evidence is given in Table 2 (taken from Haldane and Read (2000)). The numbers show the percentage point response of different maturity forward interest rates—from three to twelve months—to a one percentage point change in official interest rates in the UK. The first column shows the average response of forward rates over the period October 1990 to September 1992, ahead of the introduction of inflation targeting. The second column shows this average response over the whole of the period since the UK inflation target was introduced (October 1992 to March 1999). And the third column gives the average yield curve response over the period since operational independence and the introduction of the MPC in the UK in May 1997.

Two points are significant. First, policy “surprises” along the yield curve have clearly been smaller under the inflation-targeting regime than under earlier regimes. There is evidence of much greater conditional stability in the yield curve after 1992, with surprises on average around one third as large. This might go some way towards explaining the move towards greater openness among all central banks over recent years. Second, the committee-based decision-making procedures of the MPC do not appear to have increased the incidence of monetary policy

surprises. The MPC have locked-in the gains of the inflation-targeting regime, and have if anything further dampened conditional yield curve volatility.

VIII. DEALING WITH THE EXCHANGE RATE

The inflation-targeting countries are all archetypical small, open economies. This means that they inevitably face difficult questions about how best to respond to movements in the exchange rate, given its important influence on inflation and output dynamics. There are broadly two approaches which inflation-targeting central banks have used to analyze the effects of exchange rate movements on monetary conditions.

The first is the construction of Monetary Conditions Indices (MCI), weighting together interest and exchange rate movements, with weights depending upon their relative importance for output/inflation dynamics. Such indices have been widely used by the Reserve Bank of New Zealand and the Bank of Canada, though both countries have recently reduced their reliance on these indices when formally setting policy.

The second approach is what we might call the “spot-the-shock” approach. This is broadly the approach used in the UK. It begins by asking the question: what shock caused the exchange rate to move? Was it temporary or permanent? Was it domestic or foreign? Was the shock real or monetary in nature? The reason for asking these questions is that, depending on the shock, exchange rate changes may have very different implications for monetary conditions and hence for the setting of monetary policy.

Perhaps the best way of illustrating this approach is with an example. Between August 1996 and the middle of 1998, sterling rose by over 20 percent. Why? Did it simply reflect the effect of the (relative) tightening of UK monetary policy over the period—a temporary effect? Or was the shift due to something structural and hence persistent in nature? One means of identifying the “monetary” component of exchange rate movements is to look at relative yield curve movements. Using the uncovered interest parity theorem (and some subsidiary assumptions), it is possible to calculate how much of a given exchange rate change is induced by relative yield curve news (see Bridgen, Martin and Salmon (1997)). An example of this type of exercise, from August 1996 onwards, is shown in Chart 7. The light line shows the cumulative actual change in the exchange rate; the darker lines, the range of appreciation that can be potentially explained by relative yield curve movements—“monetary news”—on the assumption that uncovered interest parity holds. As is clear, monetary shocks did not seem capable of explaining more than 2-3 percent of the 20 percent appreciation over the period 1996–98.

An alternative explanation of Chart 7 is that the assumption of uncovered interest parity may be invalid. But the same technique has shown that monetary news can account for a substantial proportion of large exchange rate movements. For example, following sterling's exit from the ERM in September 1992, movements in relative interest rates—monetary shocks—plausibly accounted for most of the subsequent 10 percent depreciation.

But if most of sterling's appreciation between 1996 and mid-1998 was not sourced in monetary policy, what was its source? Another possible explanation was a fall in the sterling risk premium—for example, resulting from EMU uncertainties which affected the EU countries against whom sterling appreciated most strongly. Again, some calculations can be conducted using the uncovered interest parity condition and survey-based measures of exchange rate expectations. Chart 8 shows the estimated risk premium on sterling assets versus a synthetic euro asset (comprising France, Germany and Italy). This went from being positive in mid-1996 to strongly negative at the end of 1996 and through 1997. This is consistent with part of sterling's appreciation being driven by a lowering of its perceived riskiness relative to other euro currencies in the run-up to EMU. Toward the end of the sample, as EMU uncertainty abated, the estimated risk premium rose again—and sterling fell.

A third potential factor behind sterling's appreciation was a real exchange rate shift - for example, some exogenous improvement in the demand for UK exports. Historical variance decompositions from structural VARs suggested this may have been an important part of the explanation (Astley and Garratt (1996)). And if so, it would have important implications for the evaluation of monetary conditions. The shift is then more likely to be permanent, or at least persistent. And the appreciation need not then be as damaging to competitiveness and hence to output, since it represents a rise in the equilibrium real exchange rate. Over the last few years, the persisting strength of sterling lends some credence to these structural stories (Wadhvani (1999)).

This example serves to illustrate some of the merits of the “spotting the shocks” approach when generating quantified inflation forecasts. Indeed, without an appreciation and understanding of these shocks, a mechanical feeding through of the exchange rate is likely to prove misleading.

IX. SUMMARY

This paper has discussed some of the technical issues which arise when operating monetary policy according to an explicit inflation target. It has also discussed the approaches which some inflation-targeting central banks—and the Bank of England in particular—have taken to deal with these issues. Some of these approaches, if not the issues themselves, are new and there is scope for developing them further. Ghostbusting is an evolutionary science.

To date, inflation targeting has largely been taken up by developed countries. But all that may be about to change. Certainly, none of the technical issues discussed in this paper would preclude the effective use of an inflation-targeting regime by a transitional or emerging economy. Witness the success of inflation targeting in countries as diverse as Chile, Israel, and the Czech Republic. Indeed, the flexibility, transparency and clarity of such a regime mean that it may be well suited to the emerging economies. For these reasons, it is attracting some attention across much of Asia, including among the crisis countries. Looking forward—the essence of inflation targeting—we might expect further recruits to the inflation-targeting club.

Chart 1:




Chart 2

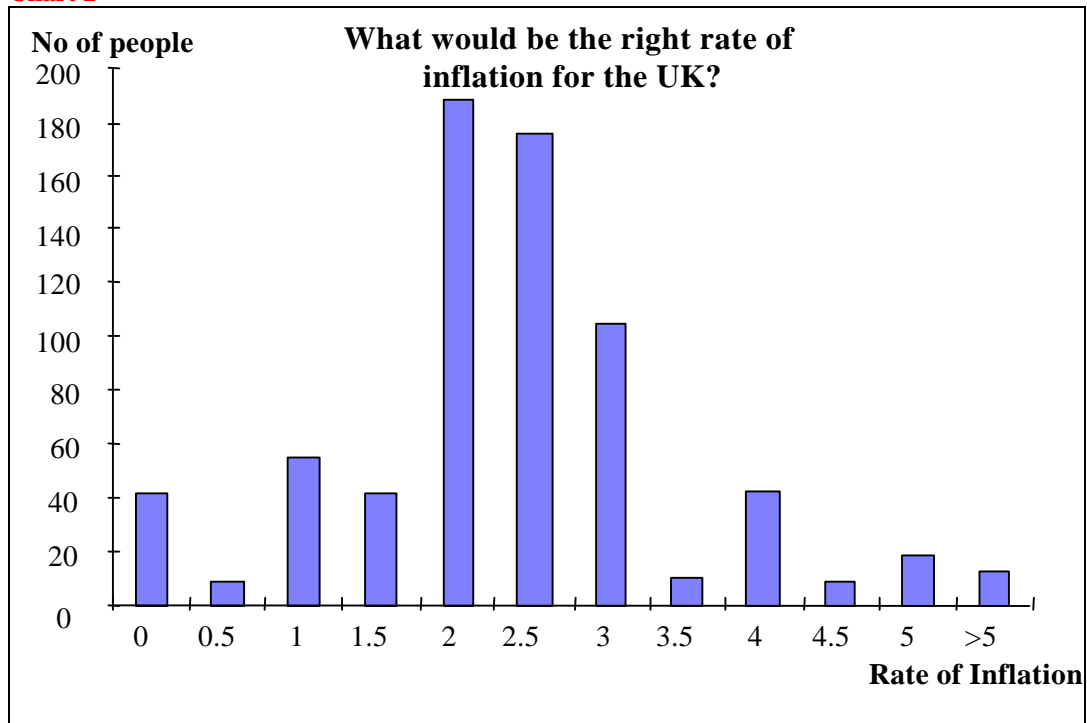


Chart 3:

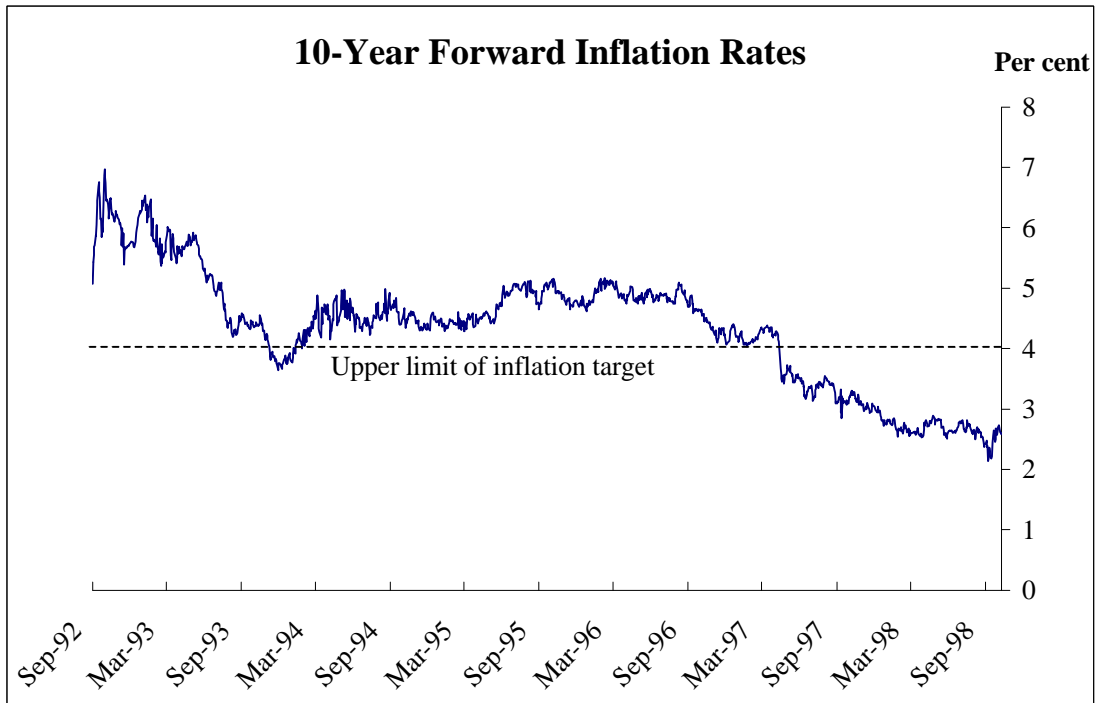


Chart 4 (a):

Chart 4 (b):

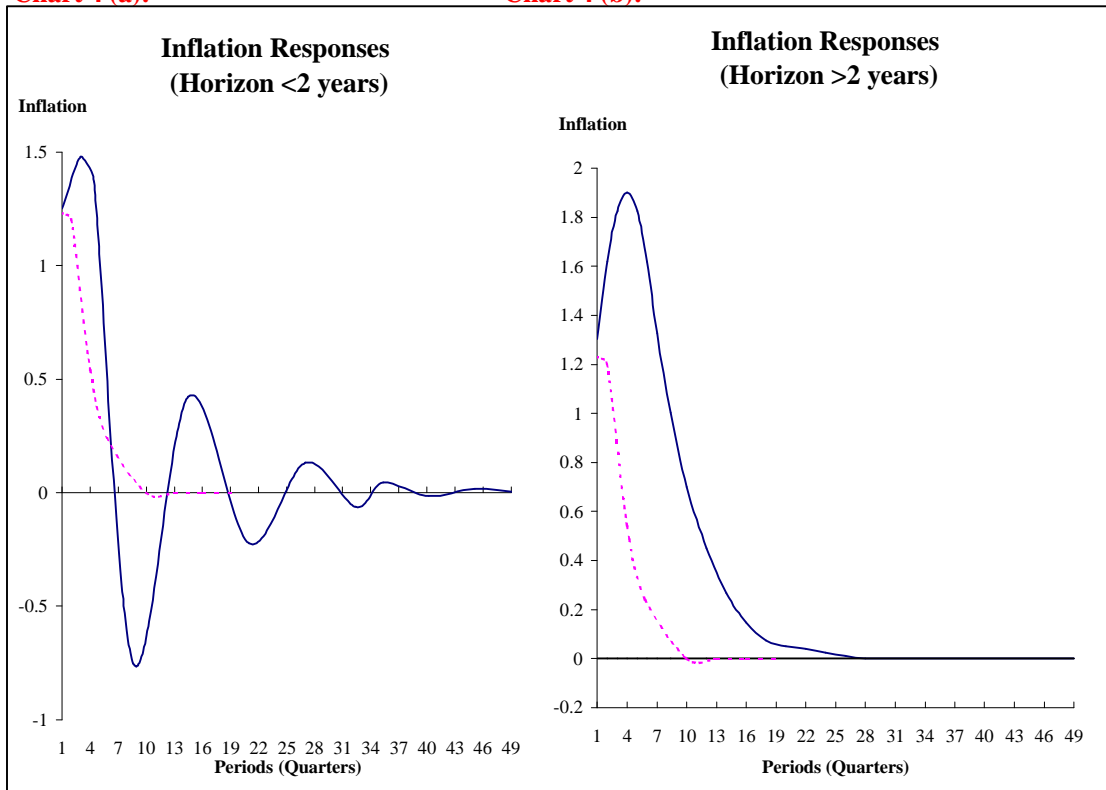


Chart 5: Inflation Forecasts

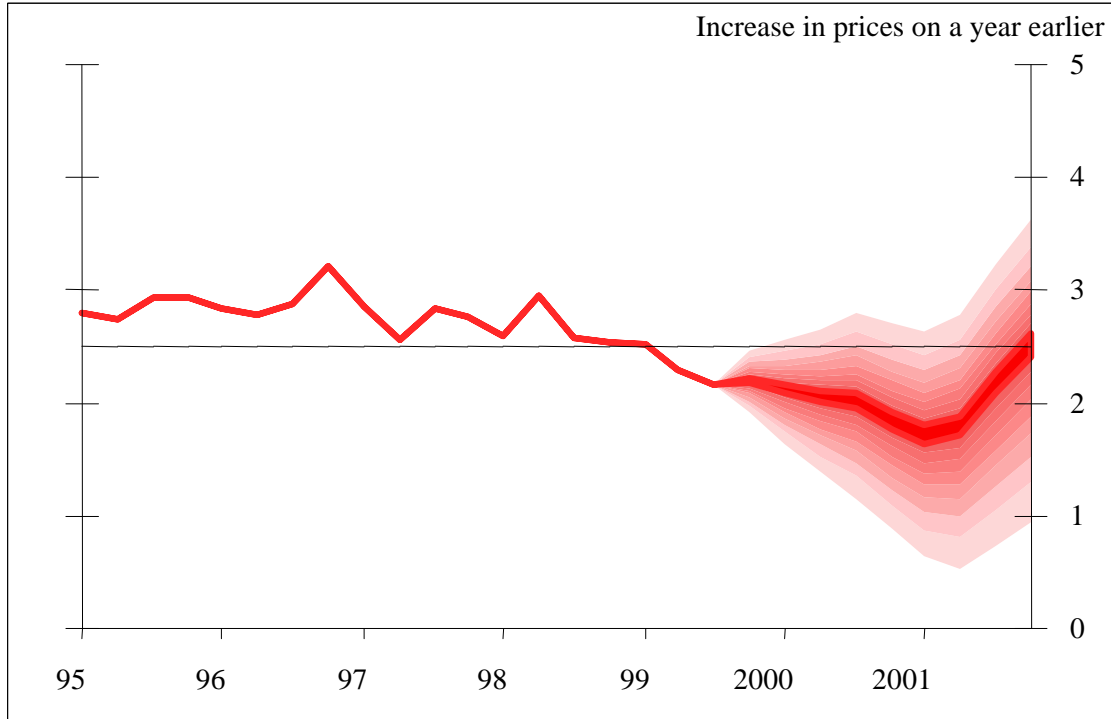


Chart 6: Inflation forecast horizons

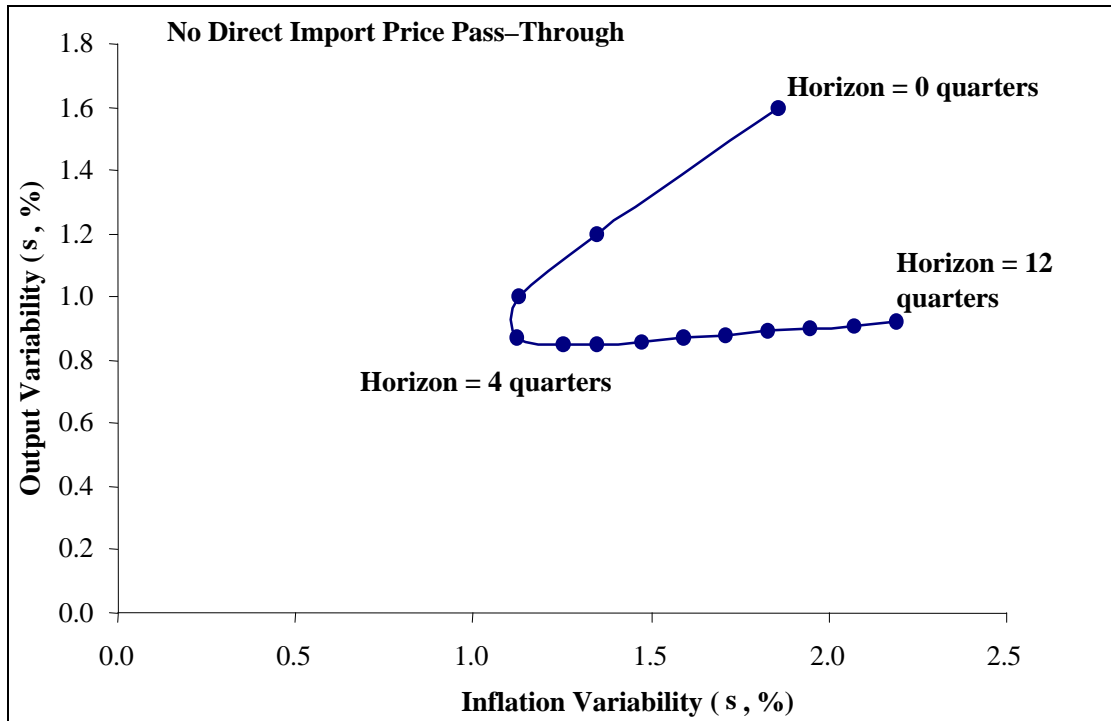


Chart 7:

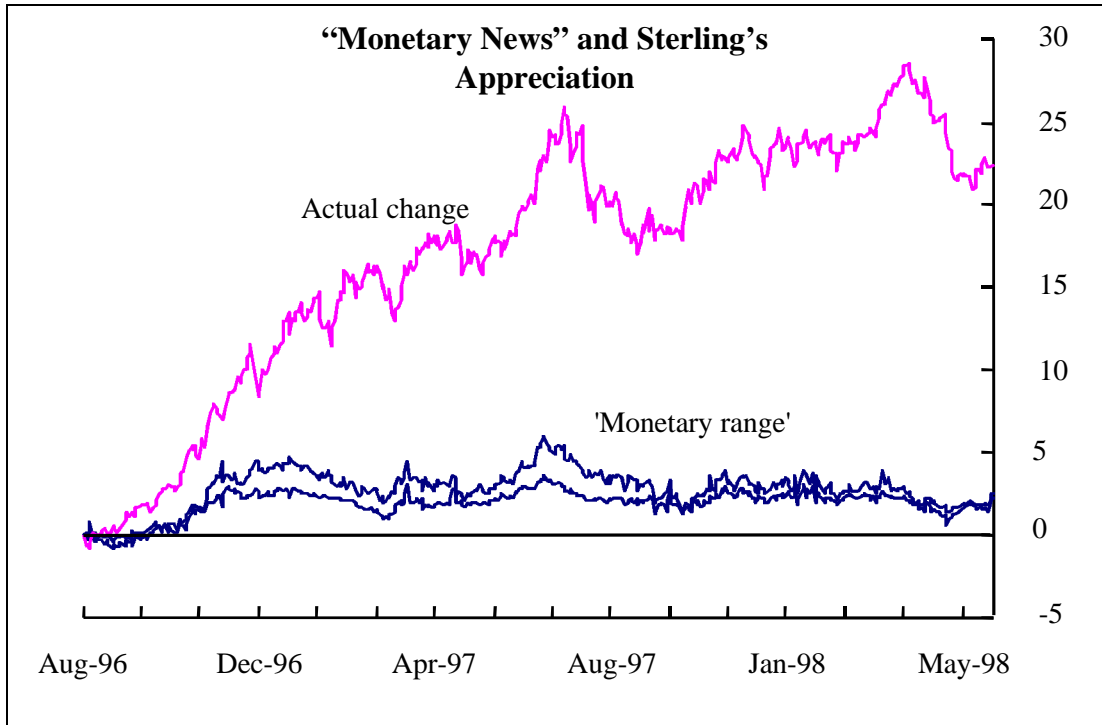


Chart 8: Risk Premium on Sterling -v- Euro




Table 1: Welfare Effects of a 2 percentage point fall in inflation

(In percent of GDP)

US	UK	Germany	Spain
0.68	0.21	0.85	1.47

Table 2: Yield Curve "Surprises"

	ERM (Oct 90 - Sept 92)	Inflation-Targeting (Oct 92 - Mar 99)	MPC (May 97 - Mar 99)
Maturity			
3 months	0.52	0.12	0.10
6 months	0.50	0.13	0.11
1 year	0.44	0.22	0.08

References

- Astley, M., and A. Garrett, 1996, "Interpreting Sterling Exchange Rate Movements", *Bank of England Quarterly Bulletin*, pp. 394–404.
- Bank of England (1999), *Economic Models at the Bank of England*, Bank of England.
- Barro, R., 1995, "Inflation and Economic Growth", *Bank of England Quarterly Bulletin*, Vol. 35, pp. 166–176.
- Batini, N., and A. G. Haldane, 1999, "Forward-looking Rules for Monetary Policy," *Bank of England Working Paper No. 92*.
- Bernanke, B., T. Laubach, F. Mishkin, and A. Posen, 1999, "Inflation Targeting."
- Brigden, A., B. Martin, and C. Salmon, 1998, "Decomposing Exchange Rate Movements According to the Uncovered Interest Rate parity Condition," *Bank of England Quarterly Bulletin*, Vol. 37, No. 4, pp. 377–389.
- Cunningham, A. W. F., 1996, "Measurement Biases in Price Indices: An Application to the UK's RPI", *Bank of England Working Paper No. 47*.
- Feldstein, M., 1996, "The Costs and Benefits of Going from Low Inflation to Price Stability," *NBER Working Paper No. 5469*.
- Haldane, A. G., 1995, *Targeting Inflation*, Bank of England.
- , 1997, "On Inflation Targeting in the United Kingdom", *Scottish Journal of Political Economy*, Vol. 45, pp. 1–32.
- Haldane, A. G. and V. Read, 2000, "Monetary Policy Surprises and the Yield Curve," *Bank of England Working Paper, No. 106*.
- Judson, R., and A. Orphanides, 1996, "Inflation, Volatility and Growth", Board of Governors of the Federal Reserve System Finance and Economics Discussion Paper No. 96-119.
- Leiderman, L., and L. Svensson, 1995, *Inflation Targets*, Centre for Economic Policy Research, London.
- Lowe, P., 1997, *Monetary Policy and Inflation Targeting*, Reserve Bank of Australia.
- Mahadeva, L., 1998, "Forecasting in the Central Bank", *photocopy*, Bank of England.

Svensson, L. E. O., 1997, "Inflation-Forecast-Targeting: Implementing and Monitoring Inflation Targets", *European Economic Review*, Vol. 41, pp. 1,111–46.

Wadhvani, S., 1999, "Sterling's Puzzling Behavior", *Bank of England Quarterly Bulletin*, Vol. 39, pp. 416–427.

Whitley, J. D., 1998, "The Inflation Report Projections: Understanding the Fan Chart," *Bank of England Quarterly Bulletin*, Vol. 38, No.1, pp. 30–37.