BANK OF ENGLAND

Speech

# What should monetary policy do?

Speech given by

David Miles, External Member of the Monetary Policy Committee, Bank of England

The Institute for Policy Research Lecture 2013, University of Bath

21 February 2013

I would like to thank Jochen Schanz and Arpad Morotz for research assistance, and others for helpful comments. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Monetary Policy Committee.

## 1 Introduction

Monetary policy is about taking calculated risks. So is crossing the road. We live in a world where the consequences of any action cannot be known in advance. I am thankful this is true – to live in a world where everything is entirely predictable would be a nightmare. Knowing the result of every sports game, knowing how every book will end and knowing exactly when you will die do not make for a happy life.

When you cross a road you do not just focus on the most likely outcome – you think about risks. Might that lorry – which appears to be coming at a steady speed of about 30mph – actually be accelerating? Is that dark patch on the road ahead just an area of damp or is it ice? You also – and all this happens in a split second in a series of thought processes that are barely noticeable – ask about the costs and benefits of different outcomes. What is the advantage of crossing now rather than waiting for that lorry to pass? How valuable is it to me to get to the other side in the next 10 seconds rather than in 45 seconds? How bad would it be if the lorry hit me – and might it be a glancing blow or a full on impact?

Monetary policy is similar. How likely is it that a loosening in policy today will create more jobs and output and be consistent with inflation moving towards a target level? Will inflation move towards a desired level far more quickly if policy is made tighter? And at what cost in terms of higher unemployment and lower output?

I want to explore some of the issues about setting policy where there are multiple sources of uncertainty and I want to do so in a way that throws light upon what the right setting for monetary policy is in the UK today.

Focusing on uncertainty is absolutely central to the decisions the Monetary Policy Committee in the UK make. It is one of the reasons why in communicating our views on the economic outlook we produce charts of assessed probabilities of different outcomes – the so called fan charts that appear in the regular *Inflation Reports*. There is, perhaps for understandable reasons, a great deal of focus paid to the single most likely outcomes that these charts show – the mode. (This is the bit in the charts towards the centre of the darkest shaded band; see Figures 1 and 2). But that central path is not of over-riding importance in setting policy. That is for the same reason that focusing just on the most likely outcome is a really bad way to decide how to cross a road. Very, very few people decide when to cross a road by thinking only about what the most likely outcome is. We think about risks – what if the lorry is actually accelerating, even though it probably isn't? What if that is a patch of ice, even though it very probably is not? What would the outcomes be if the lorry is accelerating and I hurry across now taking the shortest route by walking through that patch of stuff which is probably not ice, but might just be? If you don't ask those questions you either lack imagination or don't value your life much.

I want to start by outlining some of the most significant sources of uncertainty today and how they affect my thinking about the best monetary policy to adopt. I will illustrate this with, first, a highly simplified – but I think helpful – way of thinking about the monetary policy choice. I then want to develop something a bit more complete and explore what different views about the possible structure of the economy imply about how monetary policy should be set.

## 2 Uncertainty and its impact on policy decisions

The old joke about the lost tourist who asks for directions has him get the advice that he really shouldn't start from here. I feel much the same about monetary policy in the UK. Output (and productivity) remain far below trend; growth over the past year or so has been close to zero. Inflation remains above the target level of 2% and in the very near term may go a bit higher. The Bank's February *Inflation Report* shows that at essentially unchanged monetary policy[1] the outcomes for this year judged most likely by the MPC are that growth is low (though at least positive) with inflation stuck at around the current level (Figures 1 and 2). Beyond that the most likely outcome is that growth picks up gradually, though does not get back to the average level of recent decades and so the gap between income levels and the trend they appeared to be on before the financial crisis continues to grow. Inflation gradually – but relatively slowly – returns to the target level.

**Figure 1: Level of GDP, past and forecast**

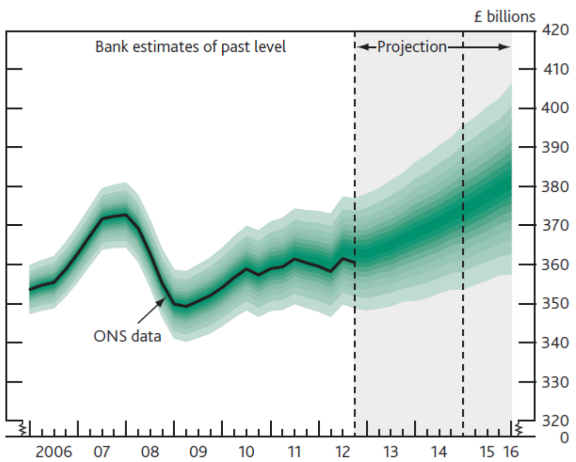


Source: February 2013 *Inflation Report*

Note: The fan charts depict the probability of various outcomes. It has been conditioned on the assumption that the stock of purchased assets financed by the issuance of central bank reserves remains at £375 billion throughout the forecast period.

**Figure 2: Annual CPI inflation**

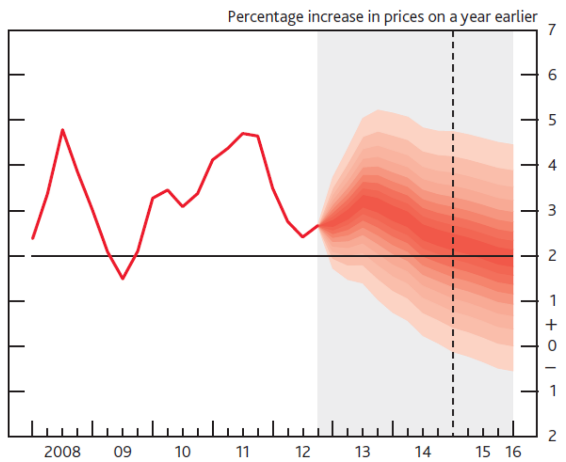


Source: February 2013 *Inflation Report*

If economic circumstances identical to today's were to prevail on 100 occasions, the MPC's best collective judgement is that the mature estimate of GDP growth and the rate of inflation would lie within the darkest central band on only 10 of those occasions.

[1] Figures 1 and 2 show projections based on Bank Rate following a path implied by money market rates which remains close to 0.5% for the next three years. The stock of assets purchased is assumed to remain at £375 billion.

Focusing on the most likely outcomes at unchanged policy does not tell us what the right monetary policy is. The big question remains whether the future, intrinsically uncertain, outcomes are better at unchanged policy or with a more expansionary policy. This depends on a range of factors about which we have limited knowledge. I think two factors are particularly important:

1. To what extent will demand pick up at unchanged monetary policy?
2. How great is the level of spare capacity ("slack") and to what extent will the growth in supply capacity[2] be raised by faster demand or restrained by weak growth?

Imagine for a moment there are just two simple outcomes for each of these two factors. For demand, it will either pick up significantly at unchanged monetary policy or stay close to the sort of levels seen over the past year or so meaning there is no growth. On spare capacity, it is either significant and does not get exhausted by faster growth or is sufficiently limited (and capacity itself unresponsive to growth) so that it would run out at much faster growth. With two possibilities for each factor that gives us 4 possible states of the world. Figure 3 below shows the 4 outcomes and shows what monetary policy might be appropriate in each case. I am assuming here – and this is a big assumption – that monetary policy can increase demand further if it is made more expansionary. I am also implicitly assuming that if growth stays at recent very weak levels inflation will fall back quickly and may move under the target level, while if growth picks up so that spare capacity is soon exhausted inflation does not fall back to target.

**Figure 3: Optimal monetary policy in stylised scenarios**

| | | Spare capacity is significant and supply capacity rises with demand | Supply capacity is limited and does not pick up as demand rises |
|---|---|---|---|
| Does demand pick up significantly at unchanged monetary policy? | Yes | Not clear what policy appropriate | Unchanged policy better |
| | No | More expansionary policy better | More expansionary policy better |

In one of these four cases a more expansionary monetary policy is pretty clearly the wrong strategy – that is where demand would pick up significantly anyway and supply capacity is limited and does not respond positively to more demand. In that case policy should probably be less expansionary – but unchanged policy is certainly better than more expansionary policy. In two situations a more expansionary monetary policy is probably warranted: if demand was not going to pick up then it looks to me that a more expansionary policy is called for, as long as there is at least some slack in the economy, or supply capacity responds positively to

[2] I'll use the term 'supply capacity' to mean the level of output at which cost pressures are sufficiently low for inflation to remain close to target.

demand. If demand was going to pick up and supply (or productivity) also responds positively to demand then it is not so clear whether unchanged or more expansionary policy is better. Let's call that a draw and say that neither policy dominates the other. So in two cases more expansion is better; in one unchanged (or less expansionary) policy is better and in one it is a draw.

This is all very simplistic and may strike you as hopelessly naive. For one thing, I have been vague about what precisely are the goals of monetary policy. And even in the simple example you might say there are really 8 outcomes and not 4 because we should distinguish between the level of spare capacity today and whether supply capacity responds significantly to faster growth (which is an issue to do with the endogeneity of productivity). And of course assuming there are just two outcomes for each factor is highly unrealistic. I have also not attached probabilities to each of the four outcomes – though I still find some value in the table since equally likely outcomes is not such a crazy assumption. If outcomes are equally likely the table implies that a more expansionary policy has a 75% chance of either being better or no worse than an unchanged policy. But I would not want to put much weight on that because of all the simplifications made in the construction of the table. So I now develop a decision making tool that is more realistic and I think does tell us something useful about monetary policy today.

## 3 A more realistic model for assessing optimal policy

I want to take the essential features of the decision making problem outlined above but make it both more realistic and more precise. I aim to describe the economic environment – and crucially the uncertainty about it – in a way that is roughly consistent with the assessment made by the MPC in its latest *Inflation Report* but which also allows an explicit calculation to be made about what optimal monetary policy is. To do that I will need to be explicit about what monetary policy is trying to achieve. But first I want to describe how I think about risks and uncertainty. The model which I am going to use allows for four factors to be random instead of the two in the simple example discussed so far. I am going to allow for:

- Uncertainty about how demand and output will evolve if monetary policy is left unchanged
- Uncertainty about the level of spare capacity today
- Uncertainty about how productive capacity would respond to faster (or slower) growth
- Uncertainty about how a change in monetary policy will affect demand and output

I will let this uncertainty take a plausible form, allowing for more than two outcomes ('yes' or 'no'). I calibrate the model to reflect what I think is a realistic assessment of the degree of uncertainty about these four factors in the UK right now.

Just to be clear at the outset: I use this model as a tool to assess policy in the United Kingdom, right now. I do not claim that it is a reliable description of the economy in other situations.

## 3.1 Structure and calibration of the model

The model that I am going to present is reduced to the bare minimum. It has three parts that describe the evolution of output, inflation and supply capacity. It can be concisely summarised in three equations – but very simple equations and with descriptions of the meaning which I hope makes the economic ideas clear and seem plausible.

I will assume that the relevant horizon for the policy decision taken today is three years – that is I assume that what matters is what happens to inflation and growth over the next three years. That is not because what happens after that does not matter. It is just that I want to focus on the policy setting now and for simplicity I will assume that a policy set now is left in place for some time. So I chose a time horizon long enough for that policy to have effect but not so long that the idea that policy is left there is completely unrealistic. The horizon I chose (3 years) is also that of the fan charts in the *Inflation Report*. This is useful because I want the simple model to be broadly consistent with the assessment of the outlook summarised in those fan charts which are based on many more factors than in my simple, decision-making model.

I'll move quite quickly through the equations and their calibration so that I can tell you sooner what monetary policy should look like according to this model. (More details are in the annex). I'll then show you how optimal policy changes when I choose different calibrations.

**Output**

The first equation describes the impact of changes in monetary policy on the annual percentage growth of demand and output, $\dot{y}_t$ (*t* is the time period and is measured in quarters so, $\dot{y}_t$ is the annual growth in output at time t). Monetary policy is described by an index *M*: the bigger is *M* the more expansionary monetary policy is. Changes in the index of monetary policy *M* (denoted $\Delta M$) could reflect changes in any monetary policy tools: for example, asset purchases, changes in Bank Rate, or in the Funding for Lending Scheme.

$$\dot{y}_t = \beta_0 + \beta_1(\Delta M) \qquad (1)$$

This equation says that if monetary policy is not changed ($\Delta M = 0$) growth will follow a path so that its level over the relevant horizon will be $\beta_0$ – which I take to be random (uncertain). I am going to assume that in the absence of a change in monetary policy, the average rate of annual growth over the relevant horizon will be somewhere between 0% and 3%. I will assume that uncertainty about this – as with uncertainty about all the elements in the model – means that all outcomes in the possible range are equally likely. This assumption is quite common for situations in which one does not have enough information to determine whether one value in this interval is more likely than another. That would mean that at unchanged policy growth will have an expected value of half way between 0% and 3% a year – an average outcome of 1.5% growth a year. 1.5% is roughly the average annual rate of growth over the next three years judged as being the mean outcome in the February Bank of England *Inflation Report*, assuming no material change in monetary policy.

The value of the parameter $\beta_1$ determines how effective monetary policy is in stimulating output growth. At the moment, the Bank's actively used marginal monetary policy tool is asset purchases. So you can think of increases in *M* as originating from additional asset purchases.

I have calibrated the range of the values that $\beta_1$ can take to include the range of estimates that the Bank published about the impact of asset purchases.[3] So $\beta_1$ like $\beta_0$ is a random variable. $\beta_1$ can take a value as low as zero, in which case making monetary policy more expansionary has no impact on demand and output.

**Inflation**

The second equation links annual inflation to the balance between the level of output and the economy's supply capacity. The difference between the economy's supply capacity and output is slack, or the output gap. The more slack there is in the economy – that is, the greater is supply capacity relative to output – the lower are cost pressures for firms, the greater is their reluctance to pass on cost rises and the lower is inflation. The inflation equation is:

$$\pi_t = \alpha_0 + \alpha_1 \pi_{t-1} + \alpha_2 Gap_t \qquad (2)$$

Inflation ($\pi$) depends on its past level and the size of the output gap (or slack). This is a backward looking model of inflation – inflation has inertia and will stay at the value of the last period unless output is above or below the supply capacity of the economy. Expectations play little role – people react to what they see as the current inflation rate and the current amount of slack in the economy. I think this is a realistic way of modelling the behaviour of most people. The key uncertainty that one is facing here is, in my view, about the size of the output gap. The percentage size of this gap is $Gap_t = (y_t^S - y_t)/y_t^S$. I'll be fairly pessimistic and assume that right now it is somewhere between zero and 3%.

To keep the analysis tractable, I assume that the other parameters are known. Inflation is taken to be quite persistent ($\alpha_1 = 0.95$) and the constant ($\alpha_0 = 0.001$) is set so that inflation eventually returns to a target of 2% when the output gap is closed. But the process takes time: any distance between inflation and its target is halved after only 3.5 years. This calibration reflects a degree of what you might call stubbornness in the

[3] Joyce, M, Tong, M, Woods, R (2011), 'The United Kingdom's quantitative easing policy: design, operation and impact', *Bank of England Quarterly Bulletin*, Q3; and Kapetanios, G, Mumtaz, H, Stevens, I, and Theodoridis, K (2011), 'Assessing the economy-wide effects of quantitative easing', *Bank of England Working Paper* No 443.

inflation process which prevents inflation from returning more quickly to target. Finally, I assume that a 1% increase in the output gap reduces the annual inflation rate by 0.1pp each quarter (ie, $\alpha_2 = -0.1$).

**Supply capacity**

The last of the three component parts of the model describes how the economy's supply capacity evolves. We start from a situation in the UK when there is slack in the economy, where productivity has been extraordinarily weak and where demand and output growth has been anaemic. I believe that part of the weakness in productivity is cyclical – in other words it has been a reflection of very weak growth. So the growth in productivity, and with it productive capacity is, to some extent, dependent on growth. To use a bit of jargon, productivity and supply capacity are endogenous. The equation showing the percentage growth of supply capacity ($\dot{y}_t^S$) reflects this:

$$\dot{y}_t^S = \delta_0 + \delta_1(\dot{y}_t - \delta_2) \qquad (3)$$

The key source of uncertainty in this equation is the degree to which higher output growth generates an increase in the economy's supply capacity ($\delta_1$). I will allow for a very wide range over which $\delta_1$ parameter can vary: from zero, where the economy's supply capacity is independent of changes in output, to 1, where a change in the growth of output generates a change in the growth of supply of the same size.

The supply equation implies that if there is no monetary stimulus, and were output then to grow at its expected average rate of $\delta_2 = 1.5\%$ per year, the economy's supply capacity grows only at $\delta_0 = 1\%$ per year. This would mean that the output gap would be expected to close over time when we start from a situation in which output has fallen behind the economy's supply capacity. I have calibrated the model so that on average spare capacity would decline gradually to zero over a 3 year horizon at unchanged monetary policy.

This way of thinking about the growth in capacity means that a monetary policy stimulus, to the extent that it increases output, can increase the growth of the economy's supply capacity, at least to some (possibly very small) extent. Let me be clear about this: I do not believe that monetary policy can affect supply growth in the long run. But starting from the specific situation in which we are today, where productivity appears to have collapsed following the decline in demand, I believe that some of the productivity loss is cyclical and could be regained if demand and output were to pick up again in the near future.

## 3.2 Inflation and output at unchanged policy

Before using the model to assess how monetary policy should be set I want to compare its forecast assuming an unchanged policy with that in the Bank's February *Inflation Report*. Figures 4 and 5 contrast the distributions of inflation outcomes from that *Inflation Report* with that of the model. For the model I assume inflation starts close to 3%. I then use repeated draws from four independent uniform distributions which give particular values for the four uncertain economic features in the model. I use 20,000 different draws of the four factors to generate a whole probability distribution of outcomes. In both figures 4 and 5, the fan chart covers 90% of the distribution of outcomes for inflation. In the *Inflation Report* chart, each pair of identically coloured bands contains 10% of the outcomes. In the simpler fan chart from the stylised model, the dark blue centre covers the central 50% of the distribution. The *Inflation Report* chart is drawn assuming that Bank Rate follows implied market rates at the time the chart was produced. The market did not predict any substantial changes in Bank Rate at the time, so that the *Inflation Report* chart can be interpreted as showing outcomes for CPI inflation at approximately unchanged policy. This makes it comparable to the stylised model's outcomes.

In the very near term the model generates less uncertainty about inflation than the *Inflation Report* fan charts. But further ahead, and particularly at the end of the forecast period (early 2016), the two probability distributions for outcomes are quite similar. The model generates outcomes for inflation that get closer to the target level a little faster than in the *Inflation Report.* Even so, it is not until 2015 that, on average, inflation gets to 2%.

The model is calibrated to generate average growth of output at unchanged policy similar to the average outcome for output growth in the *Inflation Report* (Figure 6). Finally, the model predicts that on average the output gap closes towards the end of the forecast horizon at unchanged policy (Figure 7).

**Figure 4: Annual CPI Inflation, November 2012**

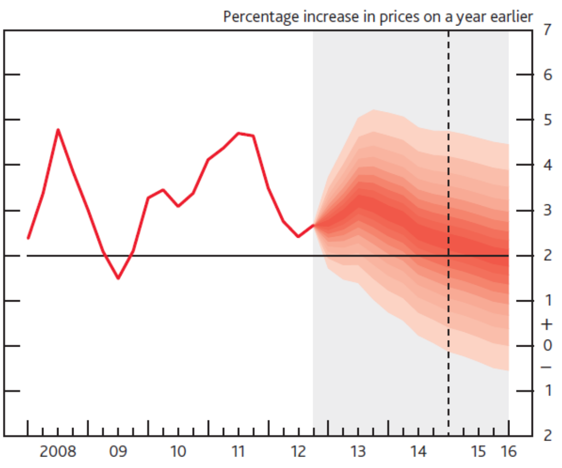


Source: February 2013 *Inflation Report*

**Figure 5: Inflation forecast of stylised model at unchanged policy**

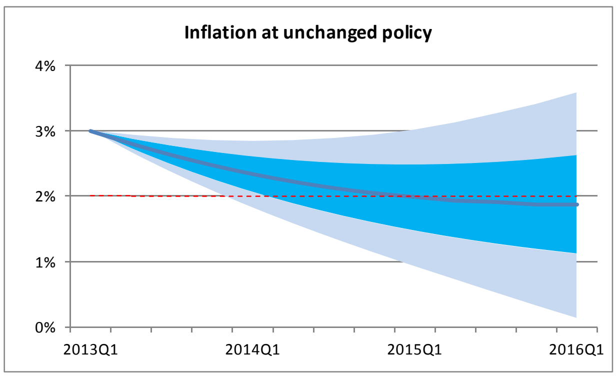


See footnote to Figure 2.

**Figure 6: *Inflation Report* forecast of GDP growth**

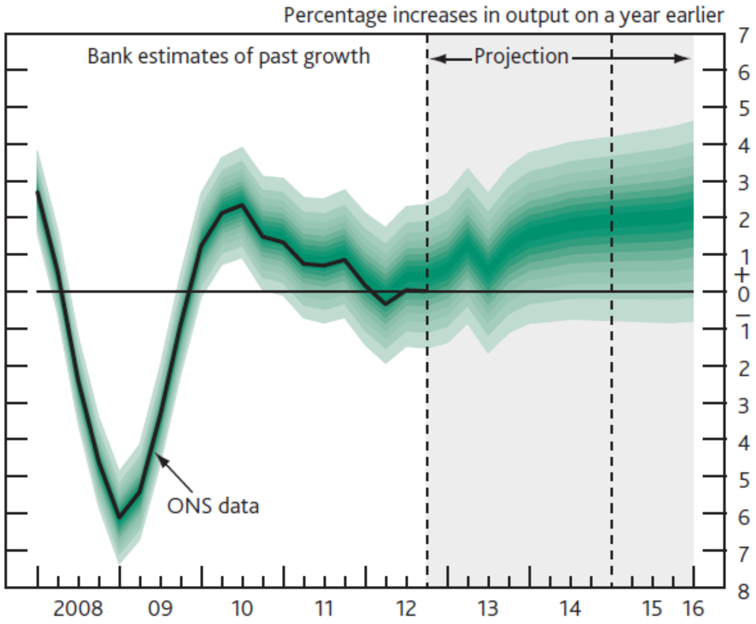


Source: February 2013 *Inflation Report*

Note: The fan chart depicts the probability of various outcomes for GDP growth. It has been conditioned on the assumption that the stock of purchased assets financed by the issuance of central bank reserves remains at £375 billion throughout the forecast period.

**Figure 7: Output gap at unchanged policy**

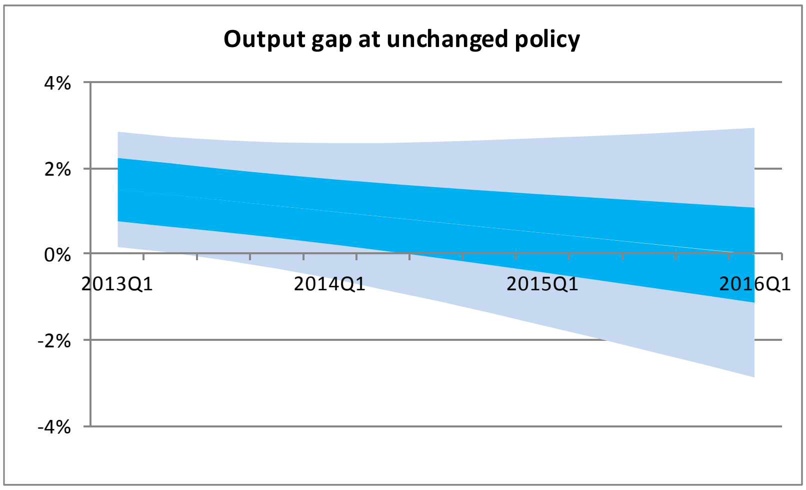

# 4 Optimal policy

## 4.1 Quantifying the Bank of England's remit

The reason for building the model is to see what it has to tell us about optimal monetary policy. To work out what an optimal policy is we need to describe exactly what it is that policy is trying to achieve – in economist's jargon, we need an objective function to maximise (or a loss function to minimise). For the Bank of England, the objective is described by the remit. The key part of this is: [4]

'[...] the objectives of the Bank of England shall be: (a) to maintain price stability; and (b) subject to that, to support the economic policy of Her Majesty's Government, including its objectives for growth and employment. [...]The framework takes into account that any economy at some point can suffer from external events or temporary difficulties, often beyond its control. The framework is based on the recognition that the actual inflation rate will on occasions depart from its target as a result of shocks and disturbances. Attempts to keep inflation at the inflation target in these circumstances may cause undesirable volatility in output.'

The remit defines a flexible inflation target. The MPC can set a policy that does not bring inflation back to target in the shortest possible time if that prevents undue fluctuations in output. This suggests that if we are seeking a precise numerical representation of the central bank's goals it should depend not only on the deviation of inflation from target, but also on the variability of output.

It is common for inflation targeting central banks to have this type of flexibility. This is reflected in the economic literature, where the most common assumption about the central bank's objective is that it minimises the weighted sum of the squared deviation of inflation from its target, $\hat{\pi}$, and of the deviation of output, $y_t$, from the economy's supply capacity, $y_t^s$. So the misery index that should be minimised is something like this:

$$L = \sum_t^T((\pi_t - \hat{\pi})^2 + \lambda(y_t^s/y_t - 1)^2) \tag{4}$$

where we add together over some horizon T the squared deviations of inflation from target and the squared deviations in output from supply capacity. The relative weight attached to output variability is $\lambda$. What the Bank of England's remit does not specify is how quickly inflation should be returned to its target – in other words what sacrifice in terms of extra output variability might be acceptable to bring inflation to target more quickly. It is understandable why the remit is not explicit on this and does not specify what is an acceptable price to pay to bring inflation back to target in all circumstances. It is impossible to anticipate what shocks will hit the economy, and in what sequence.

[4] http://www.bankofengland.co.uk/monetarypolicy/Documents/pdf/chancellorletter120321.pdf

A series of shocks[5] hit the UK economy during 2008/09. Three pushed up inflation: the depreciation of sterling, the increase in indirect taxes, and energy price inflation. At the same time, the economy went through what was probably the worst recession since the Great Depression. It is unlikely that anyone would have anticipated this combination of shocks. It is unrealistic to expect the remit to specify the appropriate trade off between a fast return of inflation to target and of output towards supply capacity in this situation. Monetary policy committees have to use judgement here. But they have to do so in a way that is consistent with the spirit of the remit. The purpose of the letters from the MPC that follow if inflation has moved too far from target is a formal opportunity for the MPC to describe how it judges the path back to target and for the Chancellor to respond. The letter from the MPC is an invitation to assess whether the strategy it is pursuing is consistent with the spirit of the remit.

In the loss function above, the weight given to output stabilisation relative to inflation stabilisation is given by the parameter $\lambda$. I will work here with a value of $\lambda = 1/2$, making deviations from the inflation target more painful for the monetary policy maker compared to deviations of output from supply capacity. Together with the quadratic specification and an inflation target of 2%, this means, for example, that the policymaker judges that if the price of keeping inflation in one period at target is that slack should be at 2.8% that is about as good (or bad) as having inflation at 4% in that period if that means output is 2.8% higher and is in line with supply. It also means that a situation in which inflation is 2% above the target is four times as bad as when inflation is only 1% above target.

One issue with this standard specification of the loss function is that it does not differentiate between situations with the same output gap and inflation but different growth rates of output. If one believed that the growth of the economy's supply capacity is independent of output growth, higher output growth reduces the output gap one for one and increases inflation. However, if growth in the economy's supply capacity depends on output growth, neither slack nor inflation may change significantly in response to higher output growth. In this case, the value of the standard loss function might not significantly change even though different rates of output growth clearly make a difference to wellbeing in the economy.

I'm therefore including output growth, $\dot{y}_t$, as another term in the loss function: the weaker growth, the larger the central bank's loss. This leaves us with the following specification:

$$L = \sum_t^T ((\pi_t - \hat{\pi})^2 + \lambda(y_t^s/y_t - 1)^2 - \gamma\dot{y}_t)$$

I set $\gamma = 1/2$, equal to the weight $\lambda$ on the squared output gap term. This means, for example, that the central bank considers as equally bad: (1) a situation in which inflation is 2pp above target, the output gap is

[5] "shocks" is a rather vague term. What I mean by it is developments that had not been anticipated and were not a result of monetary policy.

zero, and demand grows at 1% and (2) a situation in which inflation is at target, the output gap is 2.6%, and demand is stagnant.

It is hard to know whether the tradeoffs implicit in this are really reasonable. I don't think they are crazy – but that is not a very strong statement. We can at least change the tradeoffs and see what difference that makes to optimal monetary policy. That is what I proceed to do in the next section.

## 4.2 Optimal monetary policy

When evaluating the model I focus on inflation and growth developments over the Monetary Policy Committees forecast horizon of three years, and assume that the MPC can only vary its policy stance at the start of the forecast horizon. Arguably more weight should be attached to the outcomes at the end of the horizon because the terminal state determines what outcomes come after that. But to offset that, and unlike in some of the literature, I have not discounted future costs in calculating the objectives of the central bank.

I calculated the monetary policy that gave the best outcomes – on average – across the 20,000 different possible outcomes for the economic factors that reflect strength in growth, spare capacity, endogeneity of capacity and the impact of monetary policy. So what I have done is calculate the policy that maximised the expected value of our measure of the goals of policy.

At the base calibration, the model suggests that policy should be made more expansionary by 16%. The choice of the value for the impact of changes in monetary policy on output, $\beta_1$, means that this can be interpreted as saying that asset purchases should be increased by 16%, or by about £60bn. (In all tables I express the optimal change in policy in units that can be interpreted as a percent change in asset purchases from a level of £375 billion). Figures 8 and 9 show the distributions of inflation and the output gap under this policy. Towards the end of the forecast horizon, the expected rate of inflation is somewhat higher than at unchanged policy and slightly above the target level, and the output gap is negative, so, the economy is likely to be operating a bit above capacity.

**Figure 8: Annual inflation at optimal policy**

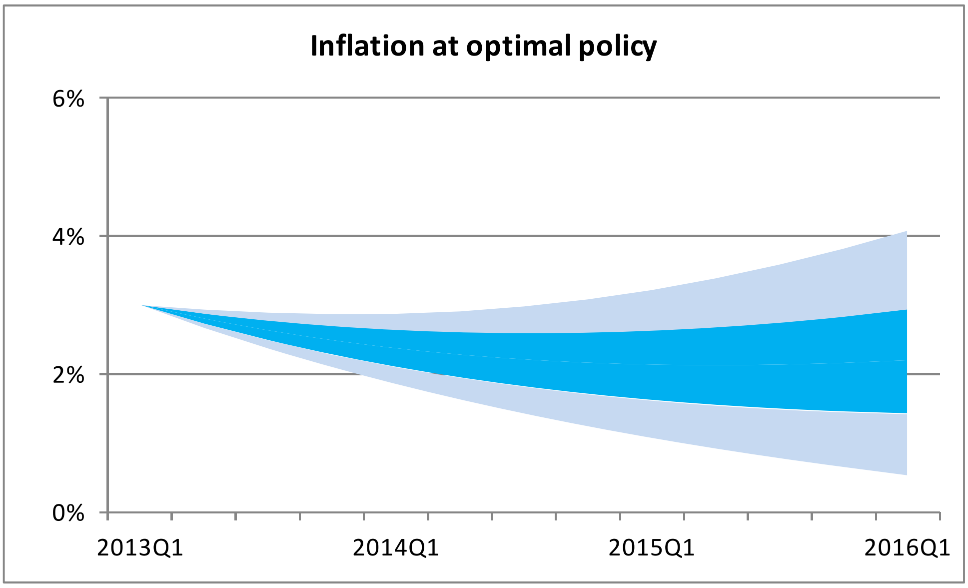


**Figure 9: Output gap at optimal policy**

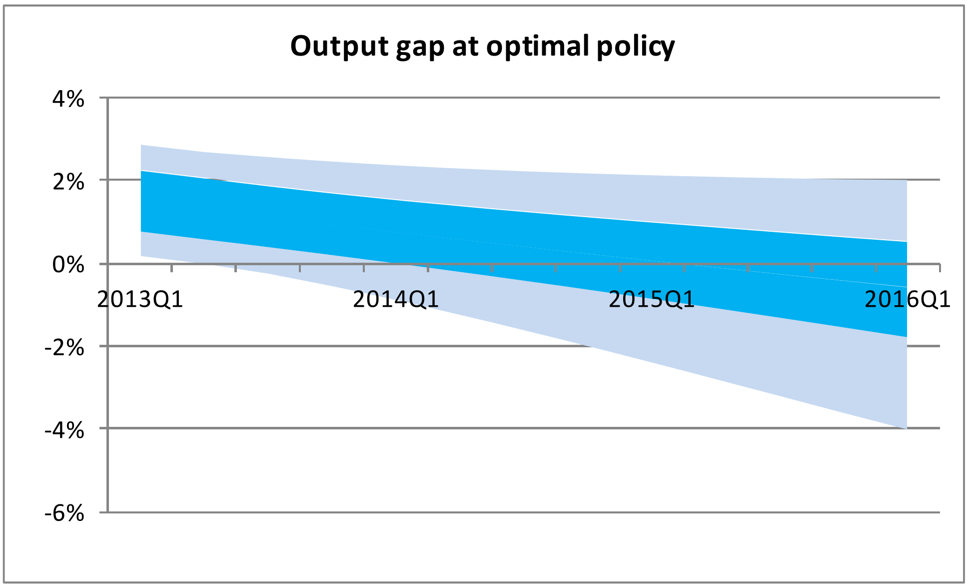


Any model's outcomes depend on how it is calibrated. So I want to look at how sensitive the results on optimal policy and economic outcomes are to different assumptions about the structure of the economy and the aims of policy.

I look first at one of the key parameters, the assumed size of the output gap in the UK right now. Table 1 shows how important the assessment of possible degrees of current spare capacity in the economy is for optimal policy. In the more optimistic case, where the output gap is on average 3% (somewhere in the [0%; 6%] range) instead of averaging 1.5% (somewhere in the [0%; 3%] range) in the base case, the model suggests that the optimal policy is to purchase about £175bn in gilts, 47% more than the Bank already has. In contrast, if the output gap is almost closed now (and is somewhere in the [0%; 1.5%] range), the optimal policy is not to change monetary policy significantly.

**Table 1: Effect of varying the assumed size of the current output gap**

| Output gap $Gap_t = (y_t^S - y_t)/y_t^S$ | Optimal policy |
|---|---|
| **Base case: Uniformly distributed in [0; 3%]** | 16% |
| **More pessimistic case: Uniformly distributed in [0; 1.5%]** | 0.5% |
| **More optimistic case: Uniformly distributed in [0; 6%]** | 47% |

Which scenario is more likely? If we assumed the output gap was close to zero, this would mean that the substantial fall in output, both absolutely but particularly relative to trend, that we have observed during the crisis has been accompanied by an erosion of the economy's supply capacity of about equal magnitude. I find this difficult to imagine.

**Figure 10: Survey indicators of capacity utilisation[(a)]**

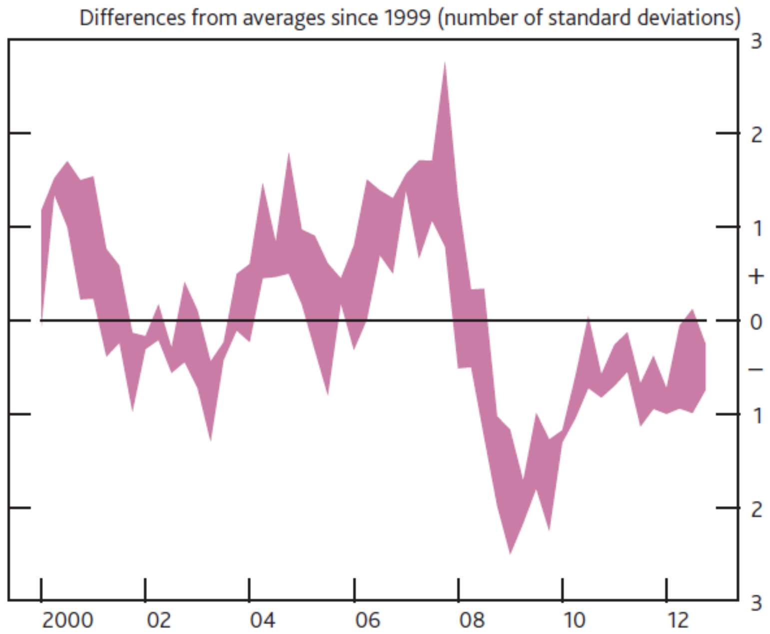


Sources: Bank of England, BCC, CBI, CBI/PwC and ONS.

(a) Three measures are produced by weighting together surveys from the Bank's Agents (manufacturing and services), the BCC (non-services and services) and the CBI (manufacturing, financial services, business/consumer services and distributive trades) using nominal shares in value added. The BCC data are non seasonally adjusted.

Yet some business surveys suggest that there is little spare capacity within firms (Figure 10). But indicators of labour market slack suggest that there is substantial spare capacity in the labour market (Figure 11): if output increased, firms in most sectors could probably hire new employees without substantial increases in wages.

**Figure 11: Selected indicators of slack in the labour market**

| | Averages | | | 2012 | | |
|---|---|---|---|---|---|---|
| | 1998–2007[(a)] | 2010 | 2011 | H1 | Q3 | Q4 |
| LFS unemployment rate[(b)] | 5.3 | 7.9 | 8.1 | 8.1 | 7.8 | 7.7 |
| Claimant count unemployment rate | 3.2 | 4.6 | 4.7 | 4.9 | 4.8 | 4.8 |
| Weighted non-employment rate[(b)(c)] | 7.6 | 9.4 | 9.5 | 9.4 | 9.2 | 9.1 |
| Vacancies/unemployed ratio[(b)(d)] | 0.41 | 0.19 | 0.18 | 0.18 | 0.19 | 0.20 |
| Part-time workers who could not find full-time work[(b)(e)] | 2.2 | 3.8 | 4.3 | 4.8 | 4.8 | 4.7 |

Sources: Bank of England *Inflation Report*, February 2013
(a) Unless otherwise stated.
(b) The figure for 2012 Q4 shows data for the three months to November.
(c) Percentage of the 16–64 population. This measure weights together different types of non-employed by the 1998-2007 averages of quarterly transition rates of each group into employment derived from the LFS.
(d) Number of vacancies (excluding agriculture, forestry and fishing) divided by LFS unemployment. Average is since 2001 Q2.
(e) Number of part-workers reporting to the LFS that they are working part-time because they could not find a full-time job, divided by LFS total employment.

Figure 12 shows a range of estimates for the UK output gap in 2012. Estimates ranged from 0.8% to 5.2%. I think it is conceivable that the output gap is as large as 5%. But I want to avoid making assumptions that could appear to favour an outcome in which monetary policy should be made more expansionary. That is why for the base calibration of the model I only considered values for the output gap between zero and 3%, with all values in that range being equally likely.

**Figure 12: Estimates of the output gap for the UK economy in 2012**[6]

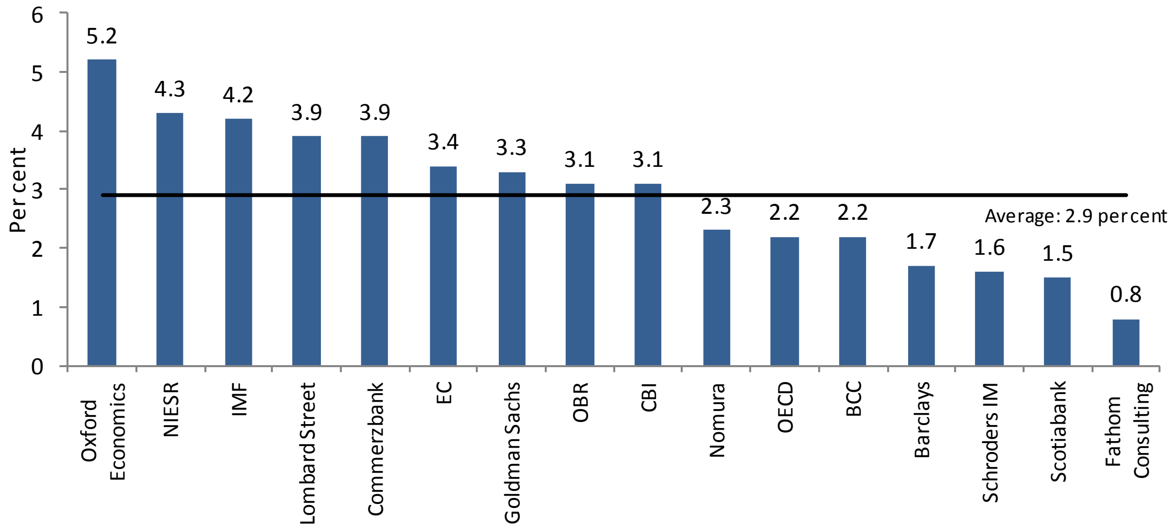


Table 2 shows the importance of judging how output might grow over the forecast horizon at unchanged policy. If output growth is likely to be feeble (and is somewhere in the [0%; 1.5%] range, 0.75% per year on average), then policy should be made much more expansionary. If, instead, growth is expected to be, on average, at about 3% (somewhere in the [0%; 6%] range) then policy should be tightened very significantly immediately.

**Table 2: Effect of varying the assumed output growth at unchanged policy**

| Output growth at unchanged policy ($\beta_0$) | Optimal policy |
|---|---|
| **Base case: Uniformly distributed in [0; 3%]** | 16% |
| **More pessimistic case: Uniformly distributed in [0; 1.5%]** | 34% |
| **More optimistic case: Uniformly distributed in [0; 6%]** | -20% |

The third key factor is the impact of a given change in policy on output growth. I calibrated this parameter for the base case drawing on research by economists in the Bank of England using different techniques for estimating the impact of the first £200bn of asset purchases on GDP.[7] The range of estimates they reported is large. The various models that the authors evaluated suggested that GDP may have increased by between 0.3% and 5.4%. I have calibrated the parameter $\beta_1$ to reflect that range of results[8]. I used a slightly

[6] Source: Office for Budget Responsibility (2012), *Economic and fiscal outlook*, December.
[7] Joyce, M, Tong, M, Woods, R (2011), 'The United Kingdom's quantitative easing policy: design, operation and impact', *Bank of England Quarterly Bulletin*, Q3; and Kapetanios et al (2011), 'Assessing the economy-wide effects of quantitative easing', *Bank of England Working Paper* No 443.
[8] I assume that the monetary conditions index M is measured in a way in a way that lets M increase by 1 unit in response to a 1% increase in asset purchases, by 2 units in response to a 2% increase in asset purchases, etc. 1% of the current level of asset purchases is approximately £4bn. The largest estimate of the impact of asset purchases suggested that the level of output rose by 5.4% in response to £200bn of purchases; the smallest estimate is +0.3%. I take the slightly wider band of 0 to 6%. £4bn are 1/50th of these

larger range of possible values for the effectiveness in policy extending the range down to zero, so that it is possible more expansionary monetary policy has no effect on demand and output. Within the range of possible values for $\beta_1$ all values are assumed to be equally likely.

This way of calibrating $\beta_1$ enables me directly to read off how much the Bank should spend on asset purchases from the change in M that is optimal according to the model. That is useful if I want to change monetary policy using asset purchases. If I want to use another policy instrument, I also need to know how the instrument's effectiveness compares with that of QE in order to interpret the model's result. For example, let's assume that cutting Bank Rate by 10bps has the same impact on output growth as £10bn worth of additional asset purchases, and that the model tells us that it is optimal to increase M by 10 units. Then the optimal policy could be implemented by either a 10% increase in the stock of asset purchases (£40bn) or a reduction of Bank Rate by 40bps.

Unsurprisingly, varying the assumptions about the impact of changes in policy on output growth matters for optimal policy (Table 3). Because there is no cost to varying monetary policy, beyond its impact on inflation and growth, then halving or doubling the average impact of a given change in policy simply doubles or halves the optimal change in policy. Table 3 illustrates this. If one is pessimistic about the impact of asset purchases on growth, one should purchase more assets, not fewer. But if you believed there was a cost (beyond its impact on inflation and growth) in holding policy a long way from normal (as it certainly is today) that result would not prove robust.

**Table 3: Effect of varying the impact of changes in monetary policy on output growth**

| Impact of changes in M on output growth ($\beta_1$) | Optimal policy |
|---|---|
| **Base case: Uniformly distributed in [0; 0.06%]** | 16% |
| **More pessimistic case: Uniformly distributed in [0; 0.03%]** | 32% |
| **More optimistic case: Uniformly distributed in [0; 0.12%]** | 8% |

I now turn to varying assumption about the central bank's objective function. The overarching objective is clear: price stability, specified as 2% inflation in the CPI. But as I explained, the remit allows for flexibility in returning to target after inflation has been blown off course. There are various possible quantifications of the objective function that are consistent with the spirit of the remit. And variations in the central bank's objectives matter. Table 4 shows that the value of the weight on the deviation in output from supply makes a significant difference.

£200bn. This suggests that £4bn worth of purchases would increase the level of output by at most 1/50*6.0% = 0.12%, and so the annual growth of output over the next few years by about half that amount. That would mean that the largest value of $\beta_1$ would be 0.06%. The lowest value is 0. So for the base case we assume $\beta_1$ is somewhere in the 0 to 0.06% range with all outcomes equally likely.

In this model, putting more weight on the output gap reduces the optimal monetary stimulus. This counter-intuitive result arises because we assume that the central bank can only change the monetary policy stance once and for all at the beginning of the forecast horizon. The stronger the monetary stimulus at the start of the forecast horizon, the faster the output gap declines, but the more likely it is that output will be above capacity towards the end of the forecast horizon. The first effect reduces the central bank's losses; the second increases it. On balance, a higher weight on the squared output gap makes optimal policy slightly less expansionary.

**Table 4: Effect of varying the weight on the squared output gap in the central bank's loss function**

| Impact of changes in λ | Optimal policy |
|---|---|
| **Base case: 0.5** | 16% |
| **More weight on output variability: 1** | 13% |
| **Less weight on output variability: 0.25** | 21% |

Table 5 shows that increasing the weight on output growth favours a more expansionary policy. But even eliminating this term from the central bank's loss function (ie, if $\gamma = 0$) still leaves optimal policy as slightly expansionary. And even if we eliminate the weight on the output growth *and* set the weight on the squared output gap to something more akin to state-of-the-art New Keynesian models (say, 0.05), the model suggests that monetary policy should still be made more expansionary by about 4% (equivalent to about £15 billion of asset purchases).

**Table 5: Effect of varying the weight on output growth in the central bank's loss function**

| Impact of changes in γ | Optimal policy |
|---|---|
| **Base case: 0.5** | 16% |
| **More weight on output growth: 1** | 25% |
| **No weight on output growth: 0** | 7% |

## 4.3 The impact of uncertainty on optimal policy

If we took out the uncertainty in the model by assuming that all the four random variables are held at their expected values, the optimal response would be to make policy even more expansionary, by 27% instead of 16%. (So converting this into a equivalent amount of asset purchases it would be £100 billion more rather than £60 billion with uncertainty.) So uncertainty about the structure of the economy and the impact of monetary policy on it makes the policymaker substantially more careful. Why is this?

Let's step back for a moment from this specific model and recall how state-of-the-art models of monetary policy incorporate uncertainty: they add an error term to each of the model equations. This is referred to as additive uncertainty. But what they generally do not do is what I do here, which is to allow for the structure of the model (its parameters) to be uncertain.

And even under the simplifying assumption that the value of the model's parameters are known, most state-of-the-art models are too complex to be solved exactly. So what people do is write down a linear approximation of the model, and then solve this approximation. In combination with a linear-quadratic loss function for the central bank, this means that additive uncertainty does not change the optimal policy. As William Brainard[9] already recognised in 1967, additive uncertainty is 'in the system', and the central bank cannot do anything to reduce it. So the central bank behaves as if it was certain that each error term would take its expected value. This is also referred to as 'certainty equivalence'. This is an unfortunate property of such linearised models.

Intuitively, one would think that uncertainty should have an effect on what decision is optimal. Indeed, Brainard showed that certainty equivalence does not hold when the impact of the central bank's policy action is uncertain. In the model used here, this uncertainty arises because key parameters of the model are uncertain. The impact of variations in policy can then add to the variability of the outcomes that the central bank cares about. This tends to make the policymaker more cautious compared to a situation in which they knew the impact of their actions for certain.[10] (More details on the impact of uncertainty are given in the annex).

## 4.4 Optimal policy – a summary

Let me summarise what I think is the overall lesson we can take from this model. In the majority of cases that I explored, optimal policy is more expansionary than unchanged policy. But not always. And in some cases the degree of extra stimulus is not very big.

An unchanged, or slightly contractionary, policy would be about optimal if the initial output gap is, on average, just 0.75% of supply instead of 1.5%, or when supply capacity only grows at 0.5% a year instead of by 1% per year. Both cases strike me as on the pessimistic side. Withdrawing monetary stimulus would also be optimal if output at unchanged monetary policy, on average, grew by 3% a year over the next 3 years instead of by 1.5%. This strikes me as being on the optimistic side of a balanced judgement; it is certainly

[9] Brainard, W. (1976), 'Uncertainty and the effectiveness of policy', *American Economic Review 57(2)*, 411-425.
[10] Brainard's result, although intuitively appealing, is not totally robust: see, for example, Söderström (2000). In particular, the policymaker may gain information about some parameters if he substantially varies his policy instrument, trading off an increase in short-term uncertainty against the opportunity to learn about his environment. However, Blinder (1998) argued that it is implausible that policymakers will 'experiment' to sharpen their econometric estimates. The concern about worst-case scenarios may also lead to stronger policy responses than in the certainty case (Giannoni, 2002; Onatski and Stock, 2002; and Tetlow and von zur Muehlen, 2002).

more optimistic than the latest prediction taken from the February *Inflation Report*, where the central forecast is that output grows at an annual average rate of about 1.5% over the forecast horizon.

The model would suggest that the additional monetary policy stimulus should be quite small – around 7% (asset purchases of about £25 billion) – if the central bank did not care about growth separately from its influence on inflation and the output gap (ie, if $\gamma = 0$ in its loss function), **or** if the economy's supply capacity did not respond to changes in demand growth.

## 5 Conclusion

Optimal policy depends on making judgements on the relative likelihood of different outcomes; it makes no more sense to just focus on the single most likely outcomes than it would in making decisions on buying insurance or crossing a road.

The current remit is not restrictive. It does not force the MPC to meet (or even try to meet) the inflation target at every point in time – it recognises you can be blown off course and that when that happens returning inflation to target in the shortest possible time is very unlikely to be optimal. The remit allows inflation to deviate from target if this avoids excessive fluctuations in output. There are no explicit guidelines for the period when inflation needs to return to target. But the MPC needs to justify that its strategy is consistent with the spirit of the remit.

The model I have developed here suggests that monetary policy should probably be more expansionary today – but this is not a universal result. You could certainly believe things about the extent of supply capacity and how demand will evolve at unchanged policy that would make you want to NOT change policy. I have also assumed that what matters for inflation is the recent extent of price rises and the amount of slack in the economy – expectations play little role. If you believe that expectations of future events are crucial to current inflation *and* that those expectations are likely to be raised if policy is made more expansionary, one would want to factor that in as a risk.

But based on my views about plausible ranges of outcomes a good case can be made for more expansion. The model does not say that asset purchases are the only way this should be achieved. If there are monetary policy tools that are more reliably effective in boosting demand, they should be used. But it is not clear what these are, which is why I have calibrated the model to reflect my own assessment of the evidence of the impact of asset purchases.

## 6 References

**Blinder, A (1998),** *Central banking in theory and practice*, The MIT Press, Cambridge, Mass.

**Giannoni, M (2002),** 'Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model', *Macroeconomic Dynamics* (6), 111-144.

**Onatski, A, and Stock, J (2002),** 'Robust monetary policy under model uncertainty in a small model of the U.S. economy', *Macroeconomic Dynamics* 6, 85-110.

**Orphanides, A, and Williams, J (2005).** 'Robust monetary policy with imperfect knowledge', *Journal of Monetary Economics* 54, 1406-35.

**Söderström, U (2002).** 'Monetary policy with uncertain parameters', *Scandinavian Journal of Economics* 104, 125-45.

**Tetlow, R, and von zur Muehlen, P (2001),** 'Robust monetary policy with misspecified models: Does model uncertainty always call for attenuated policy?', *Journal of Economic Dynamics and Control* 25, 911-49.

## 7 Annex

This annex describes the model and its calibration, explains how it has been solved, and shows results for different assumptions about the extent of uncertainty.

**Model**

The policymaker chooses a change in monetary policy, $\Delta M$ , in the first period to minimise the loss function

$$L = \sum_{t=1}^{12} ((\pi_t - \hat{\pi})^2 + \lambda(y_t^s/y_t - 1)^2 - \gamma\dot{y}_t)$$

A period is assumed to last for a quarter of a year. $y_t$ is the level of demand (equal to output); $y_t^S$ the economy's level of supply capacity; $\pi_t$ is annual inflation; and $\hat{\pi} = 2\%$ is the inflation target. The time horizon of the policymaker is three years. I abstract from discounting over this short period.

The economy is assumed to evolve according to

$$\pi_t = \alpha_0 + \alpha_1\pi_{t-1} + \alpha_2 Gap_t$$

$$\dot{y}_t = \beta_0 + \beta_1(\Delta M)$$

$$\dot{y}_t^S = \delta_0 + \delta_1(\dot{y}_t - \delta_2)$$

where the dot over a variable denotes its annual (four-quarter) growth rate, $Gap_t = (y_t^s - y_t)/y_t^S$.

**Calibration**

Table A1 contains the calibration choices. $\beta_0$, $\beta_1$, $\delta_1$, and $Gap_{t=0}$ are random variables drawn from four independent uniform distributions.

**Table A1: Calibration choices (base case)**

| **Parameters** | |
|---|---|
| $\alpha_0$ | 0.1% |
| $\alpha_1$ | 0.95 |
| $\alpha_2$ | -0.1 |
| $\beta_0$ | Uniformly distributed in [0, 3%] |
| $\beta_1$ | Uniformly distributed in [0, 0.06%] |
| $\delta_0$ | 1% |
| $\delta_1$ | Uniformly distributed in [0, 1] |
| $\delta_2$ | 1.5% |
| $\lambda$ | 0.5 |
| $\gamma$ | 0.5 |
| | |
| **Starting values** | |
| $Gap_{t=0}$ | Uniformly distributed in [0, 3%] |
| $\pi_0$ | 3% |

**Solution technique**

I assume that the four random variables are independently distributed, and used 20,000 draws for each of the variables. I then computed the value of the loss function, and searched for the monetary stimulus that minimised the loss.

**Impact of uncertainty**

It is easy to see that Brainard's intuition applies. Uncertainty about the initial output gap, and about the growth of demand when policy remains unchanged, are both unaffected by monetary policy. As a consequence, certainty equivalence almost holds: the central bank's policy choice would be the almost the same if it knew that the output gap and output growth at unchanged policy were equal to 1.5% instead of randomly distributed somewhere in the [0%; 3%] range. Certainty equivalence does not exactly hold because the model is not linear.

In contrast, uncertainty about the impact of monetary policy on output growth, and the impact of output growth on supply capacity, are magnified by changes in the monetary policy stance. So the central bank chooses a smaller monetary stimulus in the presence of uncertainty around these parameters because its loss function falls the more volatile are demand and inflation. Table A2 summarises these results.

**Table A2: Effect of eliminating uncertainty around key parameters**

| | **Optimal policy when parameter known to be at mean of uncertainty interval** |
|---|---|
| **Initial output gap: at 1.5% instead of uniformly distributed in [0%; 3%]** | 15.9% |
| **Output growth at unchanged policy ($\beta_0$): at 1.5% instead of uniformly distributed in [0%; 3%]** | 15.5% |
| **Impact of changes in M on output growth ($\beta_1$): at 0.03% instead of uniformly distributed in [0; 0.06%]** | 21.5% |
| **Impact of changes in demand growth on supply growth ($\delta_1$): at 0.5 instead of uniformly distributed in [0; 1]** | 21.7% |

Table A3 shows results on optimal policy as one varies one aspect of the economy (or of the central bank's objectives) holding other values at the base case shown in Table A1.

**Table A3: Optimal setting of policy - % change M (positive number = more expansionary policy)[(a)]**

| Parameter | Description | Values for 'certain' parameters | Ranges in which uncertain parameters are uniformly distributed | Optimal percentage change in monetary policy (positive number corresponds to more expansionary policy) | Comments |
|---|---|---|---|---|---|
| | Base case | | | **16.08** | |
| $\delta_0$ | Exogenous component of growth rate of supply capacity | 1%<br>2%<br>0.5% | | **16.08**<br>**53.93**<br>**-2.56** | If the 'natural' growth rate of supply capacity is faster, monetary policy can be more expansionary. |
| Gap | Output gap at t=0 | | 0% to 3%<br>0% to 6%<br>0% to 1.5% | **16.08**<br>**47.84**<br>**0.45** | More uncertainty and a wider expected output gap induce more expansionary monetary policy. |
| $\beta_0$ | Growth of output at unchanged policy | | 0% to 3%<br>0% to 6%<br>0% to 1.5% | **16.08**<br>**-19.62**<br>**34.48** | The higher output growth at unchanged policy, the less expansionary is monetary policy. |
| $\beta_1$ | Impact of 1% expansion in monetary policy on output | | 0 to 0.0006<br>0 to 0.0012<br>0 to 0.0003 | **16.08**<br>**8.04**<br>**32.15** | If the impact of additional monetary policy on demand is stronger, monetary policy should be less expansionary. |
| $\delta_1$ | Impact of output growth on supply growth | | 0 to 1<br>0 to 0.5<br>0 to 0.1<br>0 to 0.01<br>0 | **16.08**<br>**11.58**<br>**8.91**<br>**8.45**<br>**8.40** | If a change in the growth of demand has a stronger effect on supply, monetary policy should be more expansionary. |

| Parameter | Description | Values for 'certain' parameters | Ranges in which uncertain parameters are uniformly distributed | Optimal percentage change in monetary policy (positive number corresponds to more expansionary policy) | Comments |
|---|---|---|---|---|---|
| $\lambda$ | Weight of squared output gap in CB's loss function | 0.5<br>1<br>0.25 | | **16.08**<br>**12.90**<br>**20.72** | Increasing the weight on the squared output gap in the central bank's loss function makes policy less expansionary (see main text). |
| $\gamma$ | Weight of output growth in CB's loss function | 0.5<br>1<br>0.25<br>0 | | **16.08**<br>**24.91**<br>**11.73**<br>**7.44** | A higher weight on output growth growth in the central bank's utility function suggests that policy should be more expansionary. |
| $\alpha_2$ | Impact of slack on inflation | -0.1<br>-0.5<br>-0.25<br>-0.01<br>-0.005<br>0 | | **16.08**<br>**22.47**<br>**18.54**<br>**19.35**<br>**19.71**<br>**20.08** | The effect of this coefficient is non-monotonic. If slack does not substantially reduce inflation (small absolute value of the coefficient), it is possible to reduce the slack quicker with more stimulus while avoiding a large increase in inflation. In contrast, when the effect of slack on inflation is large, inflation will be lower during the first quarters of the forecast horizon, also suggesting that policy can be made more expansionary. |

(a) The table shows the optimal change in monetary policy when some element of the model is varied while other parameters take on the values shown in Table A1.