Central Bankers and Uncertainty

By

Charles Goodhart

SPECIAL PAPER 106

October 1998

FINANCIAL MARKETS GROUP
AN ESRC RESEARCH CENTRE

LONDON SCHOOL OF ECONOMICS

ISSN 1359-9151-106

Central Bankers and Uncertainty

by Professor C.A.E. Goodhart
Financial Markets Group
London School of Economics

For the majority of my professional life I have had the good fortune to be simultaneously involved both as a participant in, and as an academic observer of, Central Banks. Today, and as is suitable for this occasion and audience, I shall be primarily emphasizing my academic observations. Nevertheless, my study of Central Bank behaviour is inevitably informed and coloured by my previous years as a Bank official, and current position as an external member of the Monetary Policy Committee, but my comments today are unauthorised, not necessarily representative of any of my colleagues or of other Central Bankers, independent and, I trust, reasonably objective - and where they are mistaken I have no one to blame but myself, except of course for the econometrics where I have had help from the Bank staff.

Let me plunge into the central policy issue. The key decision that the monetary authorities take each month is whether, and by how much, to change the short-term interest rate. There was a time when a vocal segment of the academic community advocated a notably different operating mechanism, of monetary base control, but that debate has faded.

The question has, instead, become how Central Banks actually do, and how they should, vary

interest rates in response to economic developments. The suggestion has now been made by a number of academics, notably by John Taylor, that most Central Bank reaction functions, (except for those pegging their exchange rates and hence their interest rates to some other country), can in practice be reasonably well described by a relatively simple function, often now termed the Taylor rule, and that this rule approximates quite closely to the social welfare optimum, when examined in the context of a variety of models established for a variety of countries, (Taylor, 1998 a, b and c, and papers at June 1998 Stockholm Conference). Under such a Taylor reaction function the nominal level of the interest rate is determined by the current level of two variables, the rate of inflation and an, inherently somewhat uncertain, measure of the output gap, the deviation of actual output from potential, so

$$i_t = a + b_1\eth_t + b_2\,(Y_t - Y*)$$

where á is the equilibrium real interest rate, (usually about 2 or 3%).[1]

My first point is that virtually all attempts to estimate the Taylor rule empirically require the addition of a lagged dependent variable, i.e. the interest rate in the previous period, in order to fit well. Moreover, with monthly, or quarterly data, the coefficient on the lagged dependent variable is usually close to, and in some estimated cases greater than, unity. This means that Central Banks have historically changed rates by only a small fraction of their ultimate cumulative reaction in response to an inflationary shock or to a deviation of output from potential. Thus, the equation actually fitted becomes:-

[1] Indeed in some cases, notably Germany, evidence has been presented that such a reaction function fits the observed data better than the explanations given by the Central Bank of its own behaviour. Thus Clarida and Gertler (1997) show that the addition of monetary variables to a Taylor rule reaction function for Germany adds nothing to the explanatory power of that equation.

$$i_t = a + (1-\rho)b_1\pi_t + (1-\rho)b_2(Y_t - Y^*) + \rho i_{t-1}$$

My main theme today is to enquire further into this phenomenon whereby virtually all Central Banks change interest rates, in response to shocks, by a series of small steps in the same direction, rather than attempting more aggressively to offset that shock quickly in order to return the economy to equilibrium.

Some academics studying this subject deal with this issue by positing that changes in interest rates enter the authorities' loss function. But why should that be so? One can easily understand the social loss arising from inflation and deviations of output from potential, but what exactly is the social loss arising from changes in interest rates themselves? We will attempt to pursue this question further soon, but in the interim I want to raise a few points about the use of such a reaction function and its application to the UK.

First, the generally quite good fit of an estimated Taylor rule is not to say that in some countries over some time periods one cannot improve the fit by adding other variables. In small, open economies, especially those pegging their exchange rate, the interest rate in the home country will also respond significantly to interest rates in its larger neighbour (Peersman and Smets, 1998). Nor, of course, are the coefficients closely similar for all countries, (and over all time periods), in such estimated reaction functions.

One of the curious lacuna of this literature has been the failure so far to integrate the Taylor

reaction function literature with the literature on Central Bank Independence.[2] I would expect the measure of independence to be positively associated with the size, and perhaps the speed, of the authorities' reaction to inflation shocks. Moreover I would expect the coefficients to be regime varying, with the scale, and perhaps the speed[3], of reaction to inflation shocks rising with greater independence. There is some partial and preliminary evidence that this conjecture is correct. For example, Stephen Wright at Cambridge (1997) tested such reaction functions for Germany, US and UK over the time period 1961 Q1 till 1994 Q4 and found that over this time period the estimated cumulative response of the monetary authorities in the UK to an inflationary shock, i.e. the size of the coefficient b_1, at 0.8, was both considerably less than that of the Federal Reserve and of the Bundesbank, and also below the value of unity required to guarantee price stability. But when I asked Stephen to rerun his equation over the last decade, he obtained the much higher value for b_1 of 1.6 for the b_1 coefficient in the UK, as large as that in Germany, and slightly larger than the standard value of 1.5 incorporated in the normative versions of the Taylor rule.

Similarly a preliminary study of a number of separate, and quite short, monetary regimes in the UK, undertaken in the Bank by Ed Nelson (1998), has found the coefficients in the Taylor reaction function, especially that of the b1, inflation response, coefficient to be strongly time varying, as shown below:-

---

[2] This void is being rapidly filled now, see for example Murchison and Siklos (1998).

[3] Though there is evidence that the Bundesbank, and perhaps other more independent central banks react as, or more, slowly than those that have been more subservient, see Goodhart (1997) and A.M. Fischer (1996).

Table 1

UK	b_1	b_2	ñ	
1972/76	0	0.69*	0.7*	Qtrly
1976/79	0.44[x]	0.58	0.7*	Mnthly
1979/87	0.46[x]*	0.08	0.75*	Mnthly
1987/90	-ve	0.25*	0.66*	Mnthly
1992/97	1.3[x]*	0.24	0.4*	Qtrly

[x] Forward looking; Using Instrumental Variables

* Significant, t > 2

One of the most visible and widely remarked aspects of current Central Banking mores is that they, especially when independent, are supposed to give absolute primacy to the achievement of price stability. The level of output is not supposed to enter, for example, the objective function of the ECB or of the Bank of England. Yet, as described, the revealed preference of all monetary authorities appears to be to respond both to current inflation and to the current output gap. Actually this seeming conundrum is very simply resolved. There are two ways to answer this question. The first is that these two variables, i.e. current inflation and the current output gap, are the critical variables needed to forecast future inflation. A regression of current inflation for the U.K. on the levels of inflation and a measure of the output gap one year previously, which measure is as always somewhat arbitrary and uncertain, gives the following result:-

$$\eth_t = 0.010 + 0.840\ \eth_{t-1} + 0.527\ (y - y^*)_{t-1}$$
$$(0.011)\ (0.113) \qquad (0.199)$$

$R^2 = 0.739$, SEE = 0.029, 1974-1997 annual data.

This is not to say that the vast efforts put in by the Bank staff and others to construct the inflation forecast do not add value to our estimates of future inflation, but it does suggest that knowledge of current inflation and where the country stands on the output gap, or equivalently using Okun's Law with respect to the Natural Rate of Unemployment, can take one most of the way there. Given that lags in the transmission mechanism mean that the authorities can only reasonably target an inflation forecast (Svensson, 1997, a and b, and Svensson and Rudebusch, 1998), appearing to respond to current inflation and to the current output gap may well appear superficially much the same as targeting a pure inflation forecast.

The second leg of the answer, which was discussed at greater depth by Mervyn King in his 1997 Financial Markets Group lecture, is that even if we knew exactly how our economies worked, subject only to additive, stochastic shocks with mean zero, such shocks would still, from time to time, drive us away from our longer-term objectives of holding output close to productive potential with low, or zero, inflation. As is well known, the problem is particularly acute with supply shocks. That gives rise to the well-understood complication that, if one tries to restore inflation back very rapidly to its equilibrium, the lagged effects of monetary policy can lead both to large-scale, `excessive', variations in output (around productive potential), and in many cases also to instrument instability, (when the changes in interest rates needed to offset last times' disequilibrium become explosively greater over time). On the other hand, enormous concern for

preventing any large deviation of output from its equilibrium can lead to continuing and excessive deviations of inflation from target. This leads to a trade-off between output-variability and inflation variability of the following general form [Chart 1].

Fortunately for the Monetary Policy Committee, the empirical evidence for the UK currently indicates that this is not a serious problem. The work of Haldane, Batini and Whitley (1997) suggests that, if one chooses an appropriate horizon for returning inflation to its target, one will achieve about as good an outcome for both inflation and output variability together as is practicably possible. Thus in Chart [2], the choice of lag length (j in the Chart) for returning inflation to its target simultaneously more or less minimises both inflation deviations and output variability following a shock. In another independent exercise, my discussant, Charlie Bean (1998), estimated such a policy frontier between the standard deviations of inflation and output (Chart 3). He then wrote, and I quote,

> "The most striking thing about these frontiers are how sharply curved they are - indeed they are almost rectangular - and how close together are the optimal points for relative weights in the range 1:3 to 3:1. This rectangular quality is also found in the work of Haldane and Batini (1998),...., suggesting that it is not simply an artefact of the rather simple model structure employed here. This rectangularity has an important implication: a wide range of possible weights on output *vis-a-vis* inflation lead to the selection of rather similar points on the policy frontier. Hence little is lost by the government being able to write only an incomplete contract with the central bank, which does not explicitly prescribe the relative weight the central bank is supposed to place on output volatility versus inflation volatility; the central bank only needs to know that preferences are not extreme. Furthermore such an incomplete contract is likely to lead to a better outcome than a more completely specified contract that encourages the central bank to select a policy that is at the upper end of the policy frontier. One interpretation of the UK inflation remit is that it is precisely such an incomplete contract."

So the evidence suggests that the short-term trade-off between the variance of inflation and output, over which so much blood has been spilt, is, in the UK at least, in practice not such a

difficult and troublesome issue. The key point is that the MPC should choose an appropriate future horizon at which to aim to return to the inflation target set by the Chancellor. By doing so, they should be able to minimize the variance of both output and inflation. Given that horizon, how then should the monetary authorities operate, according to the principles that flow from our models of the economy, always remembering, and I really want to emphasize this, that in most of these models the only uncertainty in the system is of an additive, stochastic nature?

The answer to that conditional question is fairly clear. We should each month alter interest rates so that the expected value of our target, the forecast rate of inflation at the appropriate horizon about 18 months to two years hence, should exactly equal the desired rate of 2½%. Lars Svensson has written several papers on the optimality of such a procedure. If we start from an initial position in which the predicted forecast value of inflation is already close to the objective, then as a first approximation we should expect interest rates to respond to the unanticipated element in the incoming news. Since this is, by definition a martingale series, often somewhat loosely termed, a random walk, then, on these assumptions, an optimally conducted interest rate path also ought to be nearly random walk, as should also, of course, be the voting pattern of individual members of the MPC. This is, broadly, what the generality of our economic models imply.

I shall shortly demonstrate how, and why, no Central Bank actually does behave in such a random walk fashion. But before I do so, I want to contrast the normative theory inherent in our basic models with the public perception that such random walk behaviour is not optimal in practice. Thus under the headline on Thursday, June 11th, in The Times on "Anger grows at Bank's U-turn", (p. 29), Janet Bush and Anne Ashworth state that,

> "Critics of the increase described the Bank's apparent shift in policy as `almost laughable'. One said: 'It is like a drunk staggering from side to side down the street.'"

You will appreciate that this latter is an almost perfect description of a random walk path. Similarly the Sunday Business main leader of June 7th was entitled `The fickleness of hawks today and doves tomorrow', the unnamed writer commented,

> "Where the committee lost credibility last week is in its inconsistency.... What is the outside world meant to make of members who can change their view so readily? It suggests a fickle committee, influenced by the latest anecdotal or statistical evidence, swaying its opinions one way or the other and back again."

One of the arguments used by Wim Duisenberg, the President of the ECB, in rejecting the publication not only of individual voting records but also of minutes for some long duration is apparently, and this passage is in direct quotes in Robert Chote's FT article on June 1st, (p. 10), that,

> "Publication of the minutes soon after decisions have been taken or meetings have taken place will - and this is only human - make it more difficult for individual participants in the discussion to change their minds and be convinced of the arguments of others."

Now this struck a particular chord with me; for example yet another commentator, Jonathan Loynes, writing in Greenwell Gilt Weekly on May 18th, wrote,

> "Of course, this does not mean that Professor Goodhart cannot switch *back* to the Hawks. If his change of heart was driven by recent softer earnings numbers then the latest pick-up could cause him to think again. But an immediate about-turn is most unlikely, if only for reasons of credibility."

Wim Duisenberg presumably now doubts my humanity, Jonathan Loynes my credibility. Yet let me reprise once again. If policy is roughly on course to deliver the desired objective, then policy should be finely balanced and should react to incoming unanticipated news in an approximately

random walk fashion. A committee, or an individual within that, who consistently votes the same way for month after month either has got the balance of policy seriously wrong, or individually must think that that balance is seriously wrong.

I previously qualified the term random walk behaviour with the adverb "approximately". The first point to make is that the dynamic structure of the economy involves strong serial correlation and long lags in monetary policy effects. If we seek to optimise monetary policy in a model with such inherent lags, even if we still use a certainty equivalent model just involving additive stochastic uncertainty, then we could expect to find some degree of serial correlation in the path of interest rates. The dynamic structure of the economy itself can account for part of the observed persistence in the directional movement of interest rates. To repeat, interest rates should not be random walk even under certainty equivalence. But the degree, the extent, of gradualism exhibited in interest rate policy is far higher than the dynamic structure of serial correlation in the economy alone can justify.

An excellent paper by Brian Sack (1998a, also see 1998b) of the staff in the Fed's Board examined, by using a VAR model, initially with additive uncertainty, what the expected policy in adjusting the Fed Funds rate would have been if policy was to be optimised. He found (p. 4) that,

> "The optimal policy displays a tendency to move in a particular direction over sustained periods of time, as found in the data. Still, the optimal policy responds more aggressively to changes in the state of the economy than the observed policy. As a result, the funds rate path under the expected policy is more volatile than the actual funds rate. Moreover, the observed policy tends to lag behind the expected policy, limiting any changes in the funds rate and gradually moving towards the optimal policy over a period of 6 months. The actual policy is therefore described by an excessive amount of interest rate smoothing that cannot be explained strictly by the dynamic behavior of the variables to which the Fed is responding.
>
> The interest rate smoothing that is observed indicates that the analysis under additive

uncertainty ignores an important element of policymaking."

One way of visually expressing this difference is to compare the path of the calculated `optimal', and actual Fed Funds rate as Sack does in his Figure 2, here Chart 4. You can see that the darker and thicker optimal expected line is more jagged, with more reversals of direction than the actual Fed Funds path. As you can see from the time path of the actual planned target rate, most of the changes amount to small steps in the same direction, Chart 5. The cumulative distribution for the expected optimal policy with additive uncertainty is very different from that of the actual policy followed.

There are, however, some technical problems relating to the estimation and assessment of the calculated optimal interest rate change at any time. For example, should this be done on a one-step-ahead basis, starting from the actual level of interest rates in the preceding period, or on a dynamic basis starting from what would have been the optimal level of interest rates in the preceding period? In practice, when the actual level of interest rates is not too far from the estimated optimal level, the results are qualitatively pretty similar. I shall present both sets of results in my tables, but only describe in detail the one-step-ahead numbers, with the dynamic results in brackets.

Anyhow the main results are shown in Table 1. This compares the actual changes in interest rates in each month in the US with those that would have been made under the optimal policy rule, assuming stochastic additive uncertainty. The interest changes, which in the model can take any size, are here grouped into `bins', whereby any optimal change between plus and minus 12½ basis points is counted as a no change decision, any optimal change between 12½ and 37½ basis points

is grouped into the 25 basis point (¼%) change `bin', and so on. You should also note that, for reasons that will become increasingly obvious, I have grouped all changes that were continuations of an existing direction of change on the left of the table and all changes that reversed the direction of movement on the right hand side. Let me draw three features to your attention. First, under the optimal policy there would have been 57 (57) changes over this time period of ½% or more; in reality there were 23. So policy is less aggressive than the model would suggest was optimal. Second, no change was made in practice more than twice as often as this model indicated would be optimal. Third, whereas the number of continuations in the model, 58 (62), almost exactly equalled the number actually made, the number of reversals in the model, 55 (58) was over five times those made historically (10). Compared with the model predictions, the Fed has a bias to make no change, appeared extraordinarily reluctant to reverse the direction of change and tended to eschew large, aggressive movements.

Because of the importance I attach to this kind of analysis, I have been encouraging the Bank staff to complete a companion study for the UK to that done by Sack for the USA, not that they needed much encouragement from me; it was already on their agenda. Unfortunately the estimation of satisfactory VAR models for the UK is a much more complex exercise. The UK is a more open economy, which requires a model with a larger dimension; policy regime changes have been more frequent and more drastic; and the price puzzle has been even more stubbornly pervasive in UK, than in US models. I had hoped that I could unveil an accompanying table for the UK to supplement that for the US, but alas that was not to be. In an earlier, and in some respects unsatisfactory, VAR model for the UK, the results were, however, qualitatively extremely similar to those obtained by Sack for the USA. And I am prepared to conjecture that when the Bank econometric work in this field is in a fit state to be published, that qualitatively

similar results will again occur.

So the gist of my assessment is that, both in the UK and the USA, there are more small steps in the same continuing direction, more no change decisions, somewhat fewer large changes in interest rates and many fewer reversals of direction than might appear optimal under a certainty equivalent model. Moreover this is not just an Anglo-Saxon phenomenon. The general dislike of making large aggressive changes in interest rates, and the bias towards no change decisions, is well documented for all developed countries. What I would like to emphasise here is that a concern to avoid reversals of direction is also well-nigh universal, as documented in the latest 1998 BIS Annual Report. This Report comments (p. 68), and I quote,

> "There is some evidence that a dislike of reversals of this sort is not uncommon in the industrial countries. Central banks generally move interest rates several times in the same direction before reversing policy. Moreover, the interval between policy adjustments is typically considerably longer when the direction is changed (table IV.1). As the size of the steps at turning-points is not systematically larger than at other times, this pattern of adjustments risks being interpreted as a tendency to move “too little, too late”. One possible rationalisation for such behaviour is uncertainty about the policy impulses. Such uncertainty is likely to be greatest at the turning-points of the interest rate cycle. A further reason for wishing to avoid frequent interest rate reversals is the desire to provide clear guidance to markets, both to strengthen the pass-through along the yield curve and to avoid destabilising markets."

Table 2

Policy rate adjustments

	Sequence of adjustment											
	Number of changes				Average Duration[1]				Average change[2]			
	+ +	+ -	- +	- -	+ +	+ -	- +	- -	+ +	+ -	- +	- -
US	6	1	2	22	41	108	321	39	0.46	0.25	0.25	0.28
Germany	65	31	31	107	22	24	34	14	0.25	0.19	0.12	0.15
France	8	5	6	86	47	72	77	31	0.51	0.40	0.83	0.21
Italy	9	6	6	24	122	182	121	83	1.31	0.88	0.96	0.73
UK	28	17	18	84	36	69	49	23	0.94	0.50	0.77	0.37
Canada	10	1	2	21	22	57	103	21	0.43	0.25	0.25	0.25
Spain	4	5	4	33	56	72	67	35	0.42	0.24	0.35	0.38
Australia	2	1	1	17	43	413	264	67	1.00	0.50	0.75	0.79
Netherlands	55	27	28	108	16	15	32	15	0.42	0.53	0.40	0.21
Belgium	9	7	8	82	17	10	82	10	0.45	0.24	0.34	0.14
Sweden	14	1	2	24	16	132	146	10	0.12	0.25	0.27	0.18
Austria	15	1	1	48	70	42	150	34	0.38	0.50	0.25	0.16

the Taylor rule in the guise of the near unitary value of the lagged dependent variable.

John Taylor, of the eponymous rule, has studied the comparative virtues of rules of this kind both with, and without, smoothing of the form empirically observed, in simulations carried out in some 10 models of various economies. His conclusions (1998d, p. 11) are that,

> "Comparing such rules [with smoothing] with the two rules that do not respond to the lagged interest rate shows that neither type of rule dominates across all models. However, for a number of models the rules with lagged interest rates have very poor performance with extraordinarily large variances. These could be Great Depression or Great Inflation scenarios in some models. It turns out, however, that the models that predict very poor performance for the lagged interest rate rules are models without rational expectations, or in which forward looking effects are not strong in the model. Why? Interest rate rules which respond with a lag exploit people's forward-looking behavior; these rules assume that people will expect later increases in interest rates if such increases are needed to reduce inflation."

Put another way, it is alright for the authorities to act slowly in a series of cautious small steps just as long as a forward-looking public can effectively undo such cautious lags by immediate anticipation. In a similar vein Marvin Goodfriend has argued that an anticipated series of small steps in short rates will trigger off a large change in longer-term bond yields when the sequence starts, and that it may be the latter that has more effect in some economies in influencing demand. This may be particularly the case in countries where the objectives, and forecasts of the likelihood of reaching those objectives, are not regularly and publicly quantified.

It surely must be the case that the eventual determination to vary interest rates enough to defeat inflation is more important than the speed, or path, by which this is done: the Bundesbank, for example, is even more prone to smoothing than has been the case in the UK. When the reputation for determination is in place, then the ultimate measures will probably be broadly

anticipated by the public. But even if it can thus be claimed that smoothing is, in general, a fairly harmless exercise, it still leaves the question of why the monetary authorities in virtually all major countries have adhered to this behaviour pattern so determinedly. What have we failed to understand?

The failings, of course, lie far more in the standard economic models than in the practical behaviour of Central Bankers. One of the central problems is that uncertainty is far more complex, insidious and pervasive than represented by the additive error terms in standard models. The true uncertainty is multiplicative, that is attached to the coefficients in the models, or, in simpler terms, we do not know the true workings of the economic system. In some cases we do not even know which coefficients are non-zero, that is which variables are relevant. But even when we do know which variables to include in our equations, we certainly do not know what the true value of their coefficients may be.

Let me give you just two topical examples of such uncertainties. First, in an open economy, one of the main ways in which interest rate changes have an impact on the economy is via their effect on exchange rates. But can anyone, you, me, the MPC, predict the market's response at all accurately in advance? Second, to revert to the Taylor rule, discussed earlier, life would be so much easier if we knew exactly, when we come to take decisions, what was the sign of the output gap, or of its kissing cousin, the natural rate of unemployment, let alone their true arithmetic values. The regressions on the Taylor rule, that I showed you earlier, were predicated on the assumption that the way we estimate the underlying rate of productive potential is absolutely correct, and known with certainty. Whereas in practice most governments' supply-side measures are intended to give a beneficial shift to the growth of productive potential and to the

natural rate of unemployment. Moreover, it is patently obvious that such supply side factors have varied over time, though, as in Continental Europe, not always for the better.

As the Governor recently said in his speech to the TUC,

> "The truth is that neither we, nor they, nor anyone else, can know with any great certainty precisely where demand is in relation to capacity in the economy as a whole. Still less do we know where it is likely to be over the next couple of years -- and that is the more relevant consideration, given the time it takes before changes in interest rates have their full effects."

What even is the current sign of the output gap? As is evidenced by our differing votes, we in the MPC can and do individually see the same underlying data having different implications for that gap. Even in the USA where the natural rate has been historically most stable, there are always arguments that new developments, a new paradigm, may have caused significant shifts in underlying productivity and the natural rate.

Such uncertainty would matter less if it was not for the associated stylized fact that policy actions, notably monetary policy, only take effect with long lags. In the presence of multiplicative uncertainty, it would seem optimal to proceed cautiously, as Bill Brainard first demonstrated. Indeed, but if there were not such long lags, then the sensibly cautious tendency to underdo the dosage would become rapidly apparent, and just as rapidly rectified. But the problem is that it can take so long for cautious moves to become recognized as such, that the inherent dynamic of the economy can lead to inflationary, or deflationary, momentum building up in the meantime. Or in simpler terms, excessive caution, even though entirely understandable in an uncertain world Tw31 Tw (sÔut t) Tj 39

It is, perhaps, in this latter context that the publication of a central bank's inflation forecast becomes so crucial. Despite being properly hedged around with probability distributions, where our uncertainties decently peep out from under our fan charts, and with, of course, the repeated mantra that we never take the forecast either literally or slavishly, nevertheless the publication of the forecast acts as a discipline on us. Against the natural tendency to defer action in an uncertain context, the publication of the forecast holds the MPC's feet to the fire. If the projected outcome for prospective inflation is significantly different from the target, (and please allow me just for today to duck the question of how one might assess exactly what is a "significant" difference), then the MPC comes under strong pressure to rectify the situation. We all know that forecasts are fallible, but without a published forecast, in a world of long lags, the tendency towards `too little, too late', would become much worse.

`Too little, too late' could, in principle, be perfectly symmetric, in the sense that the response to deflationary pressures could be just as delayed and hesitant as the response to inflationary pressures. And we can all think of episodes, though mostly in other countries, where we might have preferred a more aggressively expansionary response to deflationary pressures. Yet it is my personal opinion that this syndrome is likely to be somewhat asymmetric. Interest rate increases are rarely popular, while expansionary measures are so. In a world of uncertainty, where what you surely know is that you do not know either the future, or even really the present state of the economy, there is, in my view, an absolutely natural, and perfectly human, tendency towards delaying restrictive action for longer than expansionary measures. I must, however, add an equally common public perception is that Central Bankers so hunger for `credibility' that they have an asymmetric bias towards tightening. Perhaps the two biasses roughly balance out?

Again my discussant, Charlie Bean, got the analysis absolutely right. Having, correctly in my view, largely dismissed the idea that politicians underhandedly try to aim for output levels intentionally in excess of the equilibrium, he goes on to say,

> "A far more plausible explanation as to why governments might be inclined to push output above the natural rate is that they are expected to deliver a high level of output through the *whole range*

Let me revert to my central concern about the nature of uncertainty. Unless there is a good reason, and there usually is not, to believe that there is negative inverse correlation between the additive and multiplicative sources of uncertainty, then the existence of multiplicative uncertainty will generally cause the authorities to move in smaller steps. On average they should underdo the dosage, since a larger change in the instrument, given multiplicative uncertainty, will add to the variance of outcomes. Given the loss function, there is a trade-off between getting as near as possible to the deserved value of the target variable on the one hand and increasing the prospective variance of the target variable(s).

From my personal viewpoint, the essential features of the economy that both set the agenda for, and complicate the life of, the monetary authorities are the interaction between the effects, and implications, of multiplicative uncertainty on the one hand and long lags in the effects of monetary policy on the other. I need hardly remind you that virtually all analysis of monetary policy games, going well beyond textbooks to what are presumed to be state-of-the-art articles, has been based on models in which neither feature appears at all.

We all know that, in principle, such multiplicative, Brainard, uncertainty should lead to greater caution in varying policy instruments, here interest rates, because a large change in rates will have an uncertain effect on outcomes, and hence raise the possibility of potentially large social losses. But a problem for practitioners is that no one, until recently has made much empirical study of how quantitatively important such Brainard uncertainty should be regarded in practice. Let me put it another way; the manner in which monetary authorities around the world appear to vary interest rates in a series of consecutive small steps of the same sign might be optimal if, and very likely only if, multiplicative uncertainty was indeed a problem of the first order of importance.

Is it such? Even if practical Central Bankers may not have known that they were talking prose all their lives, have they in practice been acting almost optimally? Until recently there was no serious attempt to measure this empirically. But now Brian Sack of the staff of the Feds' Board of Governors has made an excellent first stab of doing just that in the article that I have already quoted. He uses a five variable VAR model with production, unemployment, inflation and commodity prices as the non-policy variable and the federal funds rate as the policy variable. This exercise can both incorporate the long lags involved, and allow one to estimate the variance/co-variance matrix for the coefficients, and hence the extent of multiplicative uncertainty.

Not surprisingly, he found that such an exercise brought the actual historical conduct of US monetary policy much closer into line with what the model indicated would be optimal, see for example his Figure 5, Chart 6 here. Thus he concluded (p. 28),

> "Gradual movements in the federal funds rate do not necessarily indicate that the Federal Reserve has an interest rate smoothing incentive. Dynamic structure and parameter uncertainty can account for a considerable portion of the gradual funds rate movements that are observed. The intertemporal behavior of the targeted variables causes the funds rate to move in a particular direction over substantial periods of time. However, under additive uncertainty, expected path of the funds rate is much more volatile and reacts to changes in the economy more aggressively than the observed funds rate. This smoothing of the interest rate can be explained by the fact that the Fed does not perfectly know the structure of the economy. Uncertainty arising from imprecise estimation of the VAR coefficients is minimized at the level of the funds rate predicted by the policy rule that has been historically implemented. An aggressive policy would result in high expected variance for the targeted variables because the Fed has traditionally smoothed the funds rate. The policy rule that accounts for parameter uncertainty therefore reacts to changes in the state of the economy with gradual movements in the funds rate, which reduces the excess volatility of the expected policy and limits the deviation of this policy from the observed level of the funds rate.
>
> Although the uncertain dynamic structure results in gradual funds rate movements, there remains an element of interest rate smoothing that cannot be explained in this exercise."

Nonetheless there are still several remaining differences between such central bank behaviour in practice, and those actions that would appear optimal even after taking account of multiplicative Brainard uncertainty. Let me revert to the Table showing the implied distribution of interest rate changes that I showed earlier, but this time including also the result with multiplicative Brainard uncertainty.

What this table, Table 3, shows is that, once one takes Brainard uncertainty into account, the paucity of large aggressive jumps in interest rates becomes largely explained. With Brainard uncertainty there would only have been 24 (34) changes of 50 basis points, or more, in the US case compared to the 23 found historically.

What, however, the empirical application of Brainard uncertainty still largely fails to explain is the small number of reversals. Under our VAR models, with or without Brainard uncertainty, the number of reversals of direction of policy should have been some three to five times as common as found in practice, depending on whether one uses as the basis for judgement the one-step-ahead or the dynamic prediction from the model. Let me present some bar charts that illustrate this.

Once again, I have been encouraging the Bank staff to replicate this same study for the UK, but, as already noted, this exercise has not yet been satisfactorily completed. As before, in an earlier, and in some respects unsatisfactory, VAR model of the UK economy, the same qualitative results appeared. These are, essentially, that Brainard, multiplicative uncertainty can satisfactorily explain the small size of the standard, usual interest rate changes, but not the apparent reluctance to reverse the direction of change. And let me emphasize and repeat that I do not think that this

is just an Anglo-Saxon propensity. It is, I believe, common to all major Central Banks.

Moreover, the distributions from such a VAR model probably provide an upper-bound on the degree of caution, and interest rate smoothing, that should theoretically be undertaken, because the construction of this model completely leaves out the advantage that can be obtained from more aggressive action whereby one then learns more about the working of the economy, which should, in principle, reduce future uncertainty, see for example Sack (1998b). Thus Volker Wieland (1998, p. 2) wrote,

> "There are a number of reasons to believe that such a Brainard-type analysis overstates the case for gradualism. For example, Caplin and Leahy (1996) show that in a game between a policymaker who attempts to stimulate the economy and potential investors, a cautious policy move may be ineffectual, because investors anticipate lower interest rates in the future. Another reason, investigated in this paper, is that a more aggressive policy move may generate more information, which would improve the precision of future estimates and thereby future policy performance."

Indeed two eminent American economists, Tom Sargent (1998) and James Stock (1998), have recently argued that a Central Bank seeking to insure against the worst risks coming about (a minimax strategy) in the context of multiplicative uncertainty should actually be more aggressive, not less. The implied corollary, of course, is that if such aggression should prove to have been unnecessary, the measures can be reversed in a subsequent period. But such a reversal of policy is just what Central Banks appear, on this evidence, loathe to do.

Not only the evidence that I have presented here, but also other anecdotal reports, suggest that Central Bankers are, as a class, extremely reluctant to make a move on interest rates that might subsequently need to be reversed, and much more so than our currently best models suggest would be optimal.

There are two reasons, not mutually exclusive, why this might be so. The first I owe mainly to Michael Woodford (1998). Assume that for some reason the Central Bank wants to reduce the variance of the level of short term interest rates. Nevertheless the Central Bank wants to maintain the ability to have a quick and strong effect on the economy at a time of a major shock hitting the economy. If the Central Bank can commit to behave in such a way that any small reversal in direction of change will be followed by several similar steps in the same direction, then forward-looking rational agents will make large changes to their behaviour whenever reversals occur. But the downside for the Central Bank, the corollary, is that it must be cautious about reversing direction in the face of minor shocks, since too many short-lived reversals would limit its power to combat major shocks, given of course the initial reluctance to increase the variance of short-term rates.

The second reason is tied up with the credibility issue. As I explained earlier, when policy is already just about on course, so that the decision is finely balanced, it might, indeed, be technically optimal to change one's views and one's decisions, and the direction of movement of interest rates, as news comes in, even from month to month, certainly from quarter to quarter. It seems difficult to explain this to outside commentators, who often perceive such reversals as evidence of inconsistency, patent error, and irresolution. We all react to criticism. So long as commentators castigate the monetary authorities for moves that turn out after the event to have been inappropriate and unnecessary, then that will tend to reinforce the tendency towards `too little, too late'. The lessons from such outside criticism on changing one's mind, is that no change in interest rates should be made unless and until the probability is quite strong that a subsequent change in the same direction will also subsequently be needed. That is, I would argue, not the optimal way of conduct policy, but it is, I believe, what happens around the world.

To conclude, there is an absolute yawning gap between the general perception of non-economist outsiders that reversals of policy, changes of mind, are to be deplored and castigated as evidence of error, irresolution and general incompetence, and the apparent findings from our economic models that such reversals should optimally occur some four, or so, times more frequently than they do in practice. Maybe our models are missing something important. If not, we have then singularly failed to explain to the world at large how policy should be carried out. Either way there is till an enormous amount of work to be done.