

Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates

Charles A.E. Goodhart

I. INTRODUCTION

Lags and Optimal Control Methods of Setting Interest Rates

Although it is a generally accepted stylized fact that, as Milton Friedman noted, the monetary policy transmission mechanism has “long and variable lags,” most of the early theoretical work on these issues, e.g., the rules versus discretion debate, proceeded on the assumption that monetary policy measures were instantly transmitted to inflation. (See Barro and Gordon, 1983, and Cukierman, 1992.) Thus in his basic model, Cukierman (1992, Chap. 3, p. 28) states that “Abstracting from real shocks, growth and changes in velocity, the rate of inflation is equal to the rate of monetary growth m . Hence inflationary expectations are equal to expected money growth m^e , and the short-run Phillips relation... can be restated as $N - Nn = \alpha(m - m^e)$, where N is employment and Nn the natural rate of employment.”

Since then analysis has become more realistic. Most theorists now accept that the policy instrument that central banks actually adjust is the short-term rate of interest, and not a monetary aggregate (though some still regret this fact). Moreover, the standard work-horse models now in current use, especially Rudebusch and Svensson

(originally 1997, now 1999a), do incorporate monetary transmission lags. The Monetary Policy Committee (MPC) addressed this issue, of the lags in the transmission mechanism, in its report, “The Transmission Mechanism of Monetary Policy,” (Bank of England, 1999)—a publication intended for an audience beyond academic experts in the subject. The existence of such lags raises questions about how interest rates should be set currently at time t to achieve future objectives at time $t + n$. Typically it is supposed that interest rates do not affect current goal variables instantaneously; and, in so far as they can affect certain (intermediate) variables immediately (e.g., exchange rates), they may only do so at the cost of destabilizing other goal variables subsequently (e.g., output).

Academic studies of the way in which interest rates might be set in such a forward-looking context involve a number of dimensions. The three dimensions required for study are the model, the structure of shocks, and the loss function. In such cases, the optimal policy horizon over which inflation returns to target is determined endogenously within the model, as we will discuss subsequently when I report work by Batini and Nelson from the Bank of England. In practice, however, very few central banks, or monetary authorities, set monetary policy by such optimal control methods.

There are no doubt various reasons why this is so. As will be described in Section II, which reports work undertaken within the Bank, such optimal control horizons appear very sensitive to the precise model/shock specification applied. Moreover, all the models on which such exercises have been attempted have been small and simplified. As will be discussed, it remains unclear how successfully, if at all, optimal control techniques could be applied in the context of the larger models used in practice for forecasting purposes.

Last, but not least, such an approach requires specification of a loss function, for which the appropriate discount rate must be included. Because of the complexity and difficulty of such an exercise, governments have not been prepared to do so. Instead they have usually restricted themselves to giving primacy to the objective of price stability, while accepting that some conditions may exist in which the monetary authorities should not aim to return inflation to target excessively quickly, since that might impart undue volatility to output.

In the case of the United Kingdom, the Chancellor set out the proposed target in his letter,

Charles A.E. Goodhart is an economist at the Financial Markets Group, London School of Economics, and thanks Bill Allen, Nicoletta Batini, Spencer Dale, John Flemming, Mervyn King, Ed Nelson, Lars Svensson, Dan Thornton, and John Vickers for helpful comments and suggestions. At the time when the first draft of this paper was written, the author was himself a member of the United Kingdom's Monetary Policy Committee. He has subsequently left the MPC on the expiry of his contract in May 2000. All views and errors, however, remain entirely the personal responsibility of the author and are not necessarily representative of those of any of his other colleagues on the MPC.

June 3, 1998, to the Governor of the Bank of England, writing

The operational target for monetary policy remains an underlying inflation rate (measured by the 12 month increase in the RPI [retail price index] excluding mortgage interest payments) of 2¹/₂ per cent. The inflation target is 2¹/₂ per cent at all times: that is the rate which the MPC is required to achieve and for which it is accountable... 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.

If the government is not prepared to define the loss function, a central bank without goal independence is not constitutionally in a strong position to do so on its own. So, for a combination of reasons central banks have generally given weight to a variety of operational procedures, which do not involve conscious optimization conditional on an assumed loss function. This will involve another set of dimensions, specifically that choice of operational technique (e.g., some kind of feedback rule) and what forecast horizon to set for operating this rule (since the horizon now becomes a choice variable as well). This is discussed further in Section III.

The Rudebusch/Svensson Model and Its Findings on Alternative Policy Horizons

In their excellent paper, Rudebusch and Svensson (R/S) (1999b) consider a variety of such operational techniques in the context of a simple calibrated two-equation model of the U.S. economy, and various alternative versions of their assumed loss function. Then they seek to compare the loss from each procedure with the loss from using the benchmark optimal control approach. The operational techniques that they examine include (i) the Taylor rule, (ii) a forward- (eight quarter) looking Taylor rule, (iii) an instrument rule that responds to a rule-consistent inflation forecast,

(iv) a strict inflation target (where output deviations are given no weight), and (v) a flexible inflation target (where such deviations are given weight). In all cases such techniques are examined both with and without interest rate smoothing.

In the case of the forward-looking Taylor rule, the horizon was, somewhat arbitrarily, set at 8 quarters; while for instrument rule (iii) above and for the strict and flexible inflation targets (iv and v above), the horizons are again somewhat arbitrarily set at 8, 12, and 16 quarters, respectively.¹ R/S look at losses relative to the benchmark optimal control (OC) case in five tables (5.3 through 5.7) relating to various configurations of parameters in the loss function. Out of the 25 permutations (five techniques with differing horizons, five tables with differing loss functions), the shortest horizon gave the *worst* outcome in 21 cases and the best in 2; the 12-quarter horizon gave the worst in 0 cases and the best in 8 cases; and the 4-year, 16-quarter horizon (the longest) was the worst in 4 cases and the best in 15.

Although R/S do not focus strongly on the optimal horizon length, the implication of their findings appears to be that overall economic (adjustment) costs are reduced if the horizon for inflation targetry is considerably longer than normally practiced, at least in the U.K., where the horizon for the inflation fan chart is two years.² There may be several reasons for this difference between their U.S. analytical results and U.K. practice. One might be that the monetary policy transmission lag is somewhat shorter in more open economies, since exchange rates typically adjust faster than real output in response to an interest rate change. Another reason could be that the R/S model is backward-looking; in models with forward-looking elements the transmission process is generally speedier.

¹ The word "horizon" looks simple enough, but, as Svensson has demonstrated in several papers, semantic complexities abound in this field. Horizon can merely mean the end of the forecast, whereas "policy horizon" is usually taken to imply the date at which the objectives of policy are approximately obtained; and "forecast horizon" implies the future date at which the policy makers react to deviations of the objective(s) from their desired level. As described in more detail in Section II, the forecast and policy horizons often differ, and can obviously diverge from the length of the published forecast. In the above case, the R/S paper is examining forecast horizons.

² As described further in Section II, the evidence suggests that this is the MPC's policy horizon, as well as the end-point of the published fan chart.

The Structure of the Remainder of this Paper

Be that as it may, in Section II of this paper I shall review and summarize some current work being done in the Bank on this subject of the optimal *horizon* for monetary policy.

Another feature of U.K. procedures is that the inflation forecast is based on (i.e., conditional on) the assumption that interest rates are held constant thereafter over the remaining two years. As R/S note (p. 13), this was originally done at the Bank because (before operational independence in 1997) it could not presuppose future policy changes by the government. That particular constraint has now gone. In Section III I shall offer a few speculative comments on whether that procedure remains appropriate or whether some alternative *time-path* might instead be adopted. So Section II will concentrate on Bank work on horizons, and then Section III on some less substantial personal musings on time-paths.

A Digression on the Periodicity of Decision-Making Meetings and Forecasts

Before moving on to these issues, however, I want to digress for a moment on the subject of periodicity. The R/S model is a clever combination of annual averages and quarterly data points. In their paper, the periodicity of both the forecasting process and the decision-making process is implicitly assumed to be quarterly. In practice in almost all countries, the periodicity of the decision-making process is higher (more frequent) than that of the forecasting process. This is partly because the decision-makers can subjectively update the forecasts in the light of incoming news at less expense in terms of the use of skilled resources; so it is probably optimal to reconsider decisions with a higher frequency than formal, comprehensive forecasts. Nevertheless formal forecasts *could* be undertaken at almost any frequency, as also could policy decisions.

There seems, on the basis of casual empiricism, to be more international uniformity amongst developed nations on the preferred periodicity of formal forecasting processes for monetary policy purposes (quarterly) than on the preferred frequency for decision-making. The Federal Open Market Committee (FOMC) meets eight times a

year, but can (and does) make inter-meeting policy changes; the Bank of England and the Bank of Japan meet once a month (without inter-meeting changes so far, since the grant of operational independence); and the European Central Bank meets twice a month, but will normally only consider the monetary policy decision in depth at one of these meetings. Given that the frequency of both these processes can be chosen, and is not ineluctably and exogenously determined, there has been notably little analysis on what might be the optimal periodicity for either process. Instead, the frequencies seem to have evolved by practice, convention, and shared experience without much (any?) formal (e.g., cost/benefit) analysis. Perhaps this is desirable?

II. HORIZONS

The Horizon in the U.K.

The Chancellor's letters to the MPC, specifying the target (2.5 percent for RPIX, where RPIX is the standard retail price index excluding X, which are the components of the index directly affected by interest rate changes, specifically the mortgage cost component of housing services prices) that the MPC seeks to achieve, have not set any finite horizon for the achievement of that target, thus requiring it to be met indefinitely. Nevertheless the wording of these letters indicates that some short-run flexibility, to allow for the occurrence of supply shocks and to limit excessive short-run volatility in output, would be both acceptable and desirable. (The relevant wording in the letter of June 3, 1998, was noted earlier in Section I.)

In practice, the MPC has published inflation (and output) fan charts that extend two years into the future. Moreover, observers will have noted that, without exception, all the fan charts for inflation published since the MPC was established have shown RPIX very close to target in the quarters near the terminal date, but deviating somewhat more (though rarely by much) in the prior quarters. The fan charts for inflation for the forecasts in February and May 2000 are shown in Figures 1 and 2. The implication of these findings would seem to be that a horizon of about 18 to 24 months is also the policy horizon of the MPC.

In the meantime, staff at the Bank of England have been exploring in a more formal analytic mode the factors that might determine the choice of an appropriate (perhaps optimal) horizon. The first of the two papers that I shall describe here is

Figure 1

February 2000: Current RPIX Inflation Projection Based on Constant Nominal Interest Rates at 6 Percent





Figure 2

May 2000: Current RPIX Inflation Projection Based on Constant Nominal Interest Rates at 6 Percent



by Batini and Haldane (B/H) (1999).

The Batini/Haldane Study

In this study B/H establish and calibrate a five-equation model for the U.K.; an IS curve with real interest rate and real exchange rate terms; an LM demand for money function; an uncovered interest parity condition for the nominal exchange rate; a supply side equation based on a staggered contracting model, and an equation in which overall inflation relates both to domestic and to imported inflation (import pass through). Against the background of this (numerically calibrated) model, they then explore the (simulated) implications of imposing a class of relatively simple inflation forecast-based rules of the form:

$$r_t = \gamma r_{t-1} + (1 - \gamma)r_t^* + \theta [E_t \pi_{t+j} - \pi^*]$$

where r_t denotes the short-term *ex ante* real rate of interest, $r \equiv i_t - E_t \pi_{t+1}$, where i_t is the nominal interest rate; r_t^* denotes the equilibrium value of the real interest rate;

$$E_t(\cdot) = E(\cdot | \Phi_t),$$

where Φ_t is the information set available at time t and E is the mathematical expectations operator; π_t is inflation ($\pi_t \equiv p_t^c - p_{t-1}^c$) where p_t^c is the log of the consumer price

index); and π^* is the inflation target.

According to the rule, the monetary authorities control deterministically the nominal interest rate (i_t) so as to hit a path for the short-term real interest rate (r_t). The short real rate is in turn set relative to some steady-state value, determined by a weighted combination of lagged and equilibrium real interest rates. The novel feature of the rule, however, is the feedback term. Deviations of *expected* inflation—the feedback variable—from the inflation target—the policy goal—elicit remedial policy actions.

The policy choice variables for the authorities are the parameter triplet $\{j, \theta, \gamma\}$. The parameter γ measures the degree of interest rate smoothing (see Williams, 1997). So, for example, with $\gamma = 0$ there is no instrument smoothing. θ is a policy feedback parameter. Higher values of θ imply a more aggressive policy response for a given deviation of the inflation forecast from its target. Finally, j is the targeting horizon of the central bank when forming its policy. For example, in the United Kingdom, the Bank of England feeds back from an inflation forecast around two years ahead (King, 1997). The horizon of the inflation forecast (j) and

the size of the feedback coefficient (θ), as well as the degree of instrument smoothing (γ), dictate the speed at which inflation is brought back to target following inflationary disturbances. Because they influence the transition path of inflation, these policy parameters clearly also have a bearing on output dynamics. (B/H, pp. 8-9)


Of particular interest, both to the authors B/H and in this context, is what happens to the economy as the coefficient j (the forecast horizon) is varied. One feature of the model (and of the actual experience of open economies) is that the speed of the effect of monetary policy on inflation is very sensitive to the impact of interest rates on exchange rates and thence on imported inflation. So, as B/H state, they experiment with two versions, one “assuming full and immediate import-price pass-through (a shorter transmission lag); the other, no immediate pass-through (a longer transmission lag).” With the other two choice parameters, (γ, θ) set at their assumed 0.5 baseline values, and the calibrated values for the equations in the model, B/H obtain the resulting outcomes for output variability and inflation variability shown in their Chart 2, reproduced here as Figure 3.

As B/H state:

Several points are clear from Chart 2. First, irrespective of the assumed degree of pass-through, the optimal forecast horizon is always positive and lies somewhere between three and six quarters ahead. This forecast horizon secures as good inflation performance as any other, while at the same time delivering lowest output variability. The latter result arises because three to six quarters is around the horizon at which monetary policy has its largest marginal impact. The (integral of the real) interest and exchange rate changes necessary to hit the inflation target is minimised at this horizon. So too, therefore, is the degree of output destabilisation (the integral of output losses). At shorter horizons than this, the adjustment in monetary policy necessary to return inflation to target is that much greater—the upshot of which is a destabilisation of output. Once we allow for the fact that central banks in practice feed back from annual inflation rates,

Figure 3

j-Loci: Full and No Pass-Through Cases



whereas our model-based feedback variable is a quarterly inflation rate, then the optimal forecast horizon implied by our simulations (of three to six quarters) is rather similar to that used by inflation-targeting central banks in practice (of six to eight quarters).³

Second, taking either pass-through assumption, feeding back from a forecast horizon much beyond six quarters leads to worse outcomes for *both* inflation and output variability. This is the flip-side of the arguments used above. Just as short-horizon targeting implies “too much” of a policy response to counteract shocks, long-horizon targeting can equally imply that policy does “too little,” thereby setting in train a destabilising expectational feedback. (p. 28)

They also note further, in their conclusions:

An inflation forecast-based rule, with an appropriately chosen targeting horizon, naturally embodies a degree of output-stabilisation. Moreover, any degree of output-smoothing can be synthetically

³ This comparison is also not exact, because the two definitions of horizon are different: the feedback horizon in the rule and the policy horizon in practice (the point at which expected inflation is in line with the inflation target) are distinct concepts.

recreated by judicious choice of the parameters entering an inflation forecast-based rule. There is no need for any *explicit* output terms to enter this rule. That is evidence of the output-encompassing nature of inflation targeting based around inflation forecasts. (p. 41)

A feature of this exercise that will be explored further in the next section is that the time path of interest rates (both nominal and real) that emanates (and can be predicted to result) from this approach (i.e., a rule-based response of deviations of forecast from target inflation at some future horizon) will neither be constant (i.e., unchanged from the present) nor necessarily, or probably, optimal. B/H did not attempt to specify a loss function, which the authorities might seek to minimize. It is, however, worth noting that a choice of horizon around four to six quarters, based on (annualized) quarterly (rather than average annual) inflation rates, minimized both output and inflation volatility. Using a quarterly, rather than an annual, inflation measure tends to lead to a shorter implied lag measure.

The Batini/Nelson Study

Subsequently in 1999, Batini and Nelson (B/N) have further extended this exercise in their work in progress at the Bank of England (“Optimal Horizons for Inflation Targeting,” Draft July 1999). I hope that a revised version of this will be ready for circulation as a working paper reasonably soon. Unlike B/H, B/N do employ a loss function in which deviations of the output gap and of inflation from its target value enter, in the benchmark case with equal weight; it also incorporates a specific discount factor. To explore the optimal horizon question, they then look at several small models: a one-lag four-equation quarterly vector autoregression (VAR) and several variants of the calibrated model used in B/H.⁴

They approach the question of appropriate horizons using two alternative criteria: First they minimize the loss function to obtain an optimal policy horizon. Here the optimal horizon is not a choice variable, but is determined by the loss minimization exercise. The second approach, rather more akin to B/H, is to examine a simpler feedback rule (i.e., simpler than the OC approach) where interest rates are set in relation to the deviation of forecast future inflation from the target at period k ; thus,

$$i_t = \Psi p \left(E_{t-1} \pi_{t+k} - \pi^* \right)$$

(see B/N, p. 24).

The question then is which is the best value of k , i.e., the value which will minimize the loss function (when a simple rule of this kind is being followed). But the choice of k obviously depends on Ψp , so that they term optimal feedback horizons (OFH) the choice of k when Ψp is also optimal and feedback horizons (FH) when Ψp is not optimal.

The loss that will be suffered in the economy depends on the nature of the shock that disturbs it (e.g., whether a demand, supply, or exchange rate shock) and on the assumed model, in particular whether the economy is supposed to be backward looking, as in most VAR models, or forward looking, as in some variants of the calibrated B/H model. This means, unfortunately, that the optimal policy horizon appears highly sensitive to shock-specific and model-specific factors; see for example their Table 4. By contrast the OFH works off the deviation of forecast from target inflation (whatever may cause that deviation) and so is in a sense immune to the identification of the disturbing shock, though it too is highly sensitive to the specific model of the economy adopted.⁵ Using one of their earlier criteria for this, B/N then examine by what date (policy horizon [PH]) the economy, which follows an OFH, will achieve a convergence of inflation to the target. They find, in their Table 9, that the OFH is, in all cases, shorter than the PH; or, in simpler terms, if you want inflation to converge to target at $t + x$, you need, when using a simple rule adjusting for forecast deviations (from target), to work off deviations at $t + x - q$, where $q > 0$.

It is *not*, however, also necessarily the case that the OFH is shorter than the *optimal* policy horizon (OPH). This OPH depends on the nature of the disturbing shock which has to be identified exactly to estimate an OPH, whereas in the estimation of the PH, B/N assume that the disturbance reflects the

⁴ In assessing (optimal) policy horizons in such a model-based context, any stabilizing rule for interest rate decisions will bring inflation close to, but not exactly equal to, a point target. Consequently B/N define the horizon in various ways, either as the period when inflation converges to within a specified range around a point target, or when a given fraction of the initial shock is permanently eliminated.

⁵ Essentially, the criterion used for deriving the OFH makes B/N select optimal k in the presence of all three shocks simultaneously in stochastic simulations, with the importance of each shock depending both on how strongly it enters the model's structural equations and on the shock variance-covariance matrix.

historical average of all the various shocks. The OPH depends on both the model and the shock. In general, the OFH is again shorter than the OPH in the face of demand or supply shocks, but in three of the four models it is much longer in response to exchange rate shocks.

Preliminary Conclusions

My own reaction to this work included some concern over (i) the multiplicity of variants, some of which (e.g., whether or not interest rate smoothing might be adopted) have been glossed over here, and (ii) the extent to which the numerical results seemed sensitive to the precise specification (e.g., of model, shock, and loss function). These concerns are not that the work is wrong, rather that it seems doubtful how much it can illuminate and ease the difficult decisions facing policymakers. In practice it is rarely clear how far an *actual forecast* deviation of inflation from target is due to one or another shock. How far such an exercise, whether to examine OPHs or OFHs, could be applied to the much larger and more complex models actually used in most real-life forecasting exercises, including those in the Bank, remains an open question. Whereas the B/H exercise gave some comfort (to me at least) that the policy processes in the Bank were pretty much on course, the B/N paper suggests (to me at least) that selection of (optimal) horizons is so model/context specific that little advance can be made unless such studies can be brought to apply to the specific model/context under consideration as used in practice by the MPC. Whether that can be done in practice has yet to be seen.

III. THE PROSPECTIVE FUTURE TIME PATH OF INTEREST RATES

The Prima Facie Case Against a Constant Time Path

As is well known, the MPC's forecast is conditioned on the assumption that, between the starting date of the forecast and the two-year horizon, the short-term nominal policy-determined interest rate will be held constant. Initially this assumption was adopted in some large part because the forecast was done by the Bank, but the future decisions were taken by the Chancellor, and the Bank could not be seen publicly to second-guess what the Chancellor might do in future.

But that quasi-constitutional constraint has gone. In formal constitutional terms the MPC can, if it so wishes, indicate its own future expectations. Meanwhile continuing to adopt a (formally non-binding) assumption might seem to limit artificially the range of options that could be considered, at least in public. It would only be by extraordinary accident that the best, optimal, plan would be to consider a change of interest rates today and then intend to hold rates constant at this (new) level for the next eight quarters.

Moreover, we know that holding nominal interest rates constant tends to lead to Wicksellian instability.⁶ In practice, the rate of change of most variables visible at the two-year horizon in the Bank's forecast generally (though not invariably) tends to persist, and on occasion to accelerate, in the third and subsequent years. Consequently, any simulation, or other model exercise, with a horizon (much) beyond two years requires bolting on to the initial premise of rates being unchanged for two years—an auxiliary assumption that at some point (after the two-year horizon) the unchanged interest rate starting-point is shifted to the adoption of some stabilizing rule for nominal interest rates (e.g., a Taylor rule with appropriately chosen coefficients). Such a joint system might be described as spatchcocked, or at least inelegant (a term of some considerable opprobrium in academic circles!).⁷ There are, however, other reasons to wonder whether the width of the fan charts is quite correct (e.g., it is based largely on historical averages). My own feeling is that, whereas the above criticism is formally valid, it is not a matter of great significance. I can happily live with it.

Finally, it is not surprising that *market expectations* are often for a future time path of interest rates that is far from constant. Not only will an assumed future constant path in the MPC forecast often not be seen as the most likely, or highly

⁶ In the U.K., government expenditure plans are specified in nominal terms. Hence any shift in inflation implies offsetting projections for real government expenditures. This conditioning assumption can strongly dampen any tendency towards Wicksellian instability.

⁷ A minor variant of the same criticism is that the inflation/output fan charts are misleadingly wide, especially in the second year. If inflation did appear to be diverging from target so much, surely the MPC would respond and such response would dampen the divergence (e.g., Flemming, 1999)? That criticism could, in principle, be met by applying a "rule"-based response to divergences (beyond a certain range). But, again, such a combination of a constant-path assumption for the modal (most likely) projection and a rule-based path for divergences could also be described as spatchcocked and inelegant.

credible outcome, but it also causes internal complications for the MPC's own forecasting process. Wherever current behavior depends on future expectations of interest rates, as is the case notably with exchange rates (e.g., under any form of the uncovered interest rate parity [UIP] hypothesis), then the MPC's forecast involves the prediction that the market's expectations will be systematically falsified, based on the assumption that short-term interest rates do remain constant. That means that the MPC forecast must contain a tricky assumption about how the market would react to a (systematic) string of falsified expectations—not an easy matter.

The MPC is fully aware that its conditioning assumption of constant future interest rates does not accord, most of the time, with the implicit forward rate expectations of the market, as for example derivable from a market yield curve. Therefore, it also publishes a fan chart of the projected inflation and output outcomes that it forecasts would result if the markets' (calculated) expected time path for interest rates was to take place (e.g., as shown in the February 2000 Inflation Forecast, Figure 4). Casual observation will, however, show that the deviation of forecast inflation in the MPC's own forecasts from its (2.5 percent) target is generally wider at the 18- to 24-month horizon for the forecast with (calculated) market rates than with the MPC forecast conditioned on constant rates. The implication is that policy making rests on the MPC's conditioned forecast, with the market rate forecast presented also to help inform outside observers of what might happen if the market's implied predictions were to occur.

So the forecast, conditioned on constant interest rates, is the more important for U.K. policy-making purposes. The case against the constant interest rate assumption has several facets: it is unnecessary; it would seem to imply often (if taken as a binding constraint rather than as a conditioning assumption) a probably suboptimal decision; it is often hardly credible as the most likely outcome; and it involves various forecasting and simulation problems.

As the first draft of this paper was being written, this criticism surfaced in an International Monetary Fund working paper by Martijn and Samiei (1999). They argue that

The problem [of the credibility of the Bank's inflation forecast] is compounded by the absence of an explicit assessment

of the likely future path of the interest rate in the Report. While more recently there has been an effort to emphasize alternative views held by MPC members, the primary inflation forecasts are made under the assumption of unchanged interest rates. In principle, there are many interest rate profiles that could deliver an inflation of 2.5 percent two years ahead. Clearly, and as discussed in the previous section, there is no reason to suppose that a policy that holds interest rates unchanged and delivers a two-year ahead inflation of 2.5 percent is necessarily superior to other policies.

The Inflation Report, therefore, appears to lack transparency and credibility in relation to its inflation forecast. It is not obvious to what extent the inflation forecast based on constant interest rates is an expositional or an operational construct. If it is the latter, the MPC indeed does not intend to smooth interest rates, and consistently expects to hit the target at the newly set rate. In that case the framework is transparent, but, given a history of interest rate smoothing, it is not considered credible by market participants, as is evidenced by the deviating market forecasts... On the other hand, the constant interest rate assumption may merely be an expositional tool. In this case, the MPC in fact considers that further interest changes are likely to be necessary, even in the absence of news, implying that the MPC's forecasts lacks transparency as well as credibility. (pp. 15-16)

But What Alternative Procedure Would Be Better?

While the constant rate assumption undoubtedly has certain shortcomings, the question then becomes, What are the alternatives that might be proposed and are they preferable?

One alternative would be to have the MPC decide, and vote, not just on the change in interest rates this month but also on the whole prospective path up to some (arbitrary) horizon in a discretionary mode. But there is a virtually infinite set of possible time paths, delivering convergence to the inflation target at a wide range of policy horizons. The space of choice becomes so great that it is hard to see how a committee could ever reach a

majority for any particular time path. A great advantage of restricting the choice to what to do now, this month, is that it makes the decision relatively simple, even stark. Given the difficulties involved already in achieving majority agreement in the MPC on this simple decision, the idea of trying to choose a complete time path by discretionary choice seems entirely fanciful and counterproductive.

That means that the future time path would have to be determined in some other fashion. A standard academic approach would be by optimal control procedures, i.e., minimizing a loss function applied to the forecast. There are a variety of problems with this. First, formally establishing such a loss function, unless it was agreed by the Chancellor, might be seen as the MPC abrogating the right to select its own (short-term) goals; it could be thought to involve a “democratic deficit.” Against that, it could be argued that the Chancellor’s letter, as outlined previously, does provide some tightly limited room for discretion by the MPC to apply a (short-run) loss function, recognizing 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.” That discretion is limited by the requirement to write a letter to the Chancellor if inflation deviates from target by more than 1 percent. If, in such circumstances, the MPC proposed to return inflation to target more or less rapidly (than the Chancellor wished), given the projected effects of that on the prospective path of output, the Chancellor could tell the MPC to adjust its plans according to his own preferences.

Second, it might be difficult for a committee to agree on any formal functional representation. The coefficients in the function would be somewhat arbitrary (and what would be done about the standard central bank practice of interest rate smoothing?). Moreover, membership of the committee is time-varying, and existing members may find that their views about the (short-run) loss function shift as the context changes. In short, choosing a formal loss function might need to be revisited on each relevant occasion, running into exactly the same problems of complexity for decision-making that were outlined above.

Third, it is not clear that optimal control procedures could be applied in practice to larger, messy forecasting models incorporating a wide variety of

subjective assumptions, residual adjustments, and such other discretionary adjustments as the MPC applies to its own forecast. It is far from clear that such techniques can make the large jump from small (two to six) equation models, with often stylized shocks, to real live forecasting models with a much messier context. That still has to be seen.

Fourth, if such optimal control procedures were applied to the forecast, the resulting outcome of time paths for interest rates, inflation, etc., would become a hideously complex interaction of forecast and OC procedure. It is already difficult enough for MPC members to understand all the nuances of their own forecast, even when conditioned on relatively simple assumptions. Introducing OC procedures as well might lead MPC members to regard the whole exercise as a mysterious “black box” whose entrails were only comprehensible to a tiny number of staff academic specialists.

Fifth, if the MPC should find it more difficult to understand how the resultant outcome for the relevant variables had been determined, how would it be possible to explain it to the public, or to justify the decisions that would hang in part from it? To say that we have done what our model told us was best to do is not very convincing, especially given the track record of fancy economic models.

Those considerations suggest that, for the time being, a no-change assumption for interest rates could *not* be replaced by a full optimal control exercise, but would need to be replaced by some simpler rule, e.g., a forward-looking Taylor rule. That runs into the difficulty that there are a plethora of such potential rules. The R/S paper described in Section I mentions at least six types of rules (inflation deviation, the Taylor rule, a forward-looking Taylor rule, an instrument rule working off a rule-consistent inflation forecast, a strict inflation target, and a flexible inflation target). Each rule would need auxiliary decisions on parameters, including interest-rate smoothing and horizon. The optimality of any rule is almost certainly model dependent. Using a simplified model in the U.S., R/S have done some work on what might be the best buy amongst such rules. What might be the best choice of rule for the U.K. in the context of our actual forecasting model is far from clear. Much more work on how the time paths of the relevant variables (including interest rates) might look if various alternative rules were to be adopted would seem to be required before

one could be confident that one was not jumping out of the frying pan into the fire, if one was to replace the constant assumption by a future “rule.”

Moreover the adoption of any such “rule” would not eliminate some of the criticisms applied to the constant interest rate assumption. For example, the “rule”-related interest rate path would still normally deviate from market expectations, thereby requiring some kind of expectation correction mechanism to be applied to the exchange rate forecast. Second, in some circumstances a “rule”-based forecast for the time path of interest rates might have as little, or less, credibility as a constant forecast, e.g., in the aftermath of an asset price shock. Consider for example whether a “rule”-based forecast would have given a credible projection in autumn 1998 in the aftermath of the Russian default, the Long Term Capital Management (LTCM) hedge fund crisis, and the weakening of business confidence.

Moreover a “rule”-based forecast shares some of the potential faults of an optimal control forecast. It adds to the complexity of the forecast and makes the outcome less transparent to all concerned. It thereby makes it harder to explain and to justify to the general public. Nevertheless the constant conditioning assumption does play some part in guiding the decision on interest rates; at least I know that it did so in my own case. As already noted, there are other paths for interest rates that would not only achieve the target at the 18- to 24-month horizon, but also would at the same time potentially improve some subsidiary objective.

Let me consider two cases. First the time path for output might be projected to be unstable. If interest rates were planned to be, say, lower initially, and then subsequently higher (than the constant) path, it could simultaneously achieve both the inflation target and a smoother path for output. Second, certain asset prices might cause concern because of the potential effect of “boom and bust” on both the economy and on the financial system. Say housing prices were rising much too fast. Then a policy of higher interest rates in the short run, followed by (an expectation of) lower rates thereafter, might be desirable. So the question of whether one should necessarily condition the forecast on an assumption of future constant (short-term) interest rates actually maps into (overlaps with) the issue of whether, and how, one might take asset prices into consideration in determining

interest rate policy (see Cecchetti et al., 2000).

In this context, simplicity is a virtue. One simple approach would be to augment the forecast based on the constant conditioning assumption, with two other modal forecasts for output and inflation.⁸ In these two latter forecasts, the short-term interest rate would first be 0.5 percent lower/higher, respectively, for say the next three quarters,⁹ and then raised over the following five quarters to whatever level was calibrated as necessary to hit the inflation target on average between 18 and 24 months hence.¹⁰

The publication of three alternative interest rate paths, each of them consistent with achieving the inflation target at the 18- to 24-month horizon, might lessen the pressures to act *now*, since even with no further action immediately in prospect the target could still be hit. As we shall discuss further later, this reduces the discipline on the MPC itself. In return, and in exchange, it would provide the MPC with more flexibility to pursue subsidiary objectives (as already noted). Finally, although it is a simple and almost Spartan proposed mechanism (involving Charts A through C, a mock-up of which is shown in Figure 4) that is based on simple (but arbitrary) conditioning assumptions, it could still be subject to the criticism of fine-tuning.

Flexibility and Commitment

If one adopts a “rule,” there is some implication that one will follow that path. The constant, no change, assumption is clearly exactly that—an assumption not a rule. No one infers any commitment from the MPC to abide by that assumption in the future, nor is the credibility of the MPC damaged when, having made this assumption in a forecast one month, it decides to change interest

⁸ One could reasonably assume that the variance and skews would be much the same in all three cases.

⁹ Two quarters would, perhaps, be too short to show much relevant difference. Four quarters would probably require too sharp a subsequent reversal to hit the inflation target given lags in the transmission mechanism.

¹⁰ It should, however, be noted that this proposal, which emphasizes the future prospective *reversal* of a current move, is almost the exact opposite of Woodford’s (2000) proposal for emphasizing inertial prospective continuations. He places weight on the importance of influencing long rates by commitment to a series of moves in short rates. If the above prospective reversals were credible, then the temporary prospective move in short rates would have very little, or no, effect on long rates. Whether that matters depends, in some large part, on whether the transmission mechanism runs largely via changes further out along the yield curve.

rates even in the next month if the “news” should appear to warrant it.

There *is* something special about zero, no change. Any indication that the MPC is formally indicating a future specific change in rates (e.g., as driven by a “rule”-based formula) would be taken to indicate some degree of commitment. It is my view that the few experiences in the past where the MPC has given *any* measure of commitment (e.g., in August 1997) were soon felt to be burdensome and unhappy.

Moreover, the extent of any such commitment may be interpreted differently by the market and even by different members of the monetary authority itself (Thornton and Wheelock, 2000). It appeared, from a distance, to be a problem for the FOMC in the summer of 1999 that its statements about bias were at times misinterpreted by the market, requiring subsequent speeches by the Chairman to try to rectify.¹¹ Some observers also say that the fact that the Fed is not required to announce the bias at the end of each meeting has led the markets to interpret excessively the importance of the decision to announce bias. Economists say the statement has almost no utility as a predictive device since it is likely to give an indication of the next move in interest rates only 30 to 50 percent of the time. It was for reasons such as this that the MPC responded to the Treasury Select Committee of the House of Commons that it had no wish to publish a statement about future bias.¹²

A bias statement is, I would suggest, much less binding on future decisions than a quantified projection for a future path of interest rates. Unforeseen events (shocks) are likely to make the MPC depart from any such prefigured path. This may well be seen by markets as reneging on a prior commitment, and lead to accusations of misleading the market. It would be extraordinarily difficult to decide amongst ourselves, or to explain to the markets, just what extent of commitment was involved in following a rule-based path. This was less of a problem for the Reserve Bank of New Zealand since they projected a Monetary Conditions Index (MCI), which specified a complex combination of future expected developments, without implying any specific future path either for interest rates or for the exchange rate itself.¹³ The authors of the IMF working paper, Martijn and Samiei, are aware of this problem, but do not, in my view, give it sufficient weight; they assume this difficulty away when they state (p. 16), “Obviously, the Bank would

have to make it clear it was not committing itself to a particular path, so that without loss of credibility it might revise its projection at a later date as new information becomes available.”¹⁴

By contrast, a constant interest rate assumption is generally perceived as involving no forward commitment whatsoever. So it imposes no constraints on the MPC’s future decisions, leaving the MPC to respond with maximum flexibility to unforeseen events as they occur.

There is a clear division of views about the degree of forward signaling of, and market

¹¹ This was underlined, even as the first draft of this paper was being written, by a report that the “Fed [was] to examine policy bias announcements” in the *Financial Times*, December 14, 1999, p.11, by S. Fidler in Washington who reported as follows: “A policy step taken by the Federal Reserve earlier this year aimed at increasing the openness about its decision-making has confused financial markets, the central bank has admitted...Despite the aim to increase transparency about the Fed’s activities, and therefore to calm market volatility, the bias statement is widely acknowledged to have had the opposite effect. Critics have suggested that this is because of widespread confusion about what the bias statement means, confusion that appears to be shared by some Fed officials.”


¹² In any case the arguments in the minutes of the MPC meetings and the recorded individual votes give a considered view of the current inclinations of the MPC. For example the minutes of the April 2000 meeting (¶ 29-32) report the differing views of the members of the MPC. For example ¶ 31 states that “On one view, it would be better not to raise rates this month. The news over the month as a whole was inconclusive, with falls in manufacturing production, retail sales and consumer confidence, and oil prices, but with the determinants of domestic demand, both private and public, remaining robust. Against that background other factors were also important. First, the extent of price pressures stemming from the labour market required further analysis, not least in disentangling the effects of bonuses and other elements of wage drift from that of settlements. Second, the analysis undertaken for the *Inflation Report*, and the opportunity this provided to set out the Committee’s thinking in detail, were valid reasons not to move this month unless there was a strong case to do so. Third, for some, the volatility in equity markets introduced a possible downside risk which might mean that any increase in rates this month would need to be reversed soon afterwards. Finally, the imbalances in the economy, manifested in another fall in manufacturing production, seemed to have worsened. An increase might exacerbate these imbalances and it was possible that the weakness in some sectors might feed through into the rest of the economy. With inflation running below target, and expected to continue to do so for a while, there was no pressing reason to raise rates straightaway. For these members, no change in the repo rate was needed this month, although for some of them it was more likely than not that there would need to be an increase in rates in due course.”

¹³ While the Reserve Bank of New Zealand has, subsequently, abandoned the use of MCI, it has replaced this with a quantitative indication of the expected future path of short-term interest rates. Whether financial markets will feel aggrieved when, and if, interest rates depart from this prefigured path has yet to be seen.


¹⁴ In the next sentence they state that “The Bank could also include outside projections of inflation in the *Inflation Report*.” We already do, and the *Inflation Report* is the responsibility of the MPC, not the Bank alone; coverage of the two differs.

Figure 4


**Chart A
Interest Rates**



**Chart B
% Growth Output**



**Chart C
% Level of Inflation**



preparation for, future interest rate changes that (independent) central banks regard as suitable. The MPC certainly wants the market to be able to predict its future moves, and I believe that its members were not unhappy by a report that a computer had been programmed on an in-sample basis successfully to do so. (The program's out-of-sample subsequent predictive ability was less good!) Indeed the MPC wants monetary policy to be seen as "boring." But the MPC wants the basis of such predictability to be a considered assessment of its "reaction function" based on published accounts of members' actual votes and detailed minutes of the reasons for such votes, rather than more speculative comments in advance of the actual decision-making meetings of how (individual) members might vote. Given the large weight placed on individual accountability in the MPC and the inherent unpredictability of economic "shocks," any attempt to foreshadow in individual comments the prospective future decisions of the MPC as a whole would be seen as potentially embarrassing hostages to fortune.

In contrast, the Fed and ECB appear to consider the preparation of the market for future interest rate changes, for example by public statements by the Chairman, to be often desirable. This may reflect in part a tendency of the Fed and the ECB to have a more collegiate, consensual bias to decision-making, whereas the MPC places more weight on individual responsibility. It would be an interesting exercise to try to ascertain whether central banks differ significantly in their propensity to signal future interest rate moves, and, if so, what might be the causes and consequences of such differences. But that exercise has not yet been done.

How one might define "a signal" would itself be difficult to determine. In the absence of such an exercise, further discussion of "signaling" would take me beyond the self-imposed limits of this paper.

What Are the Consequential Incentives for Behavior?

Publishing the forecast for inflation, together with the assumption of constant interest rates from the forecast date to the horizon, has the effect of inducing the MPC to act now to adjust expected inflation at the appropriate horizon into line with the target. If the MPC could choose a time-varying path for interest rates, it might be more inclined to

defer interest rate changes until it was possible to learn more about evolving conditions.

There is a developing debate in the academic literature between those who suggest gradual changes in interest rates (for example, because of uncertainty about key parameters in the model [Brainard uncertainty], about possible errors in the data, and about opportunities for learning) and those who argue for more aggressive and preemptive action (e.g., the need for robust policy responses and minimizing potentially large errors). It is my conjecture that the choice of a constant versus a time-varying interest rate path has similar implications. The assumption of a constant interest rate path (plus a published forecast) may be thought to induce the MPC into early preemptive action. Whether or not “gradualism” or “preemptive” action is generally preferable is too large a subject to enter into here. My conjecture is that our procedure tilts us in the direction of early, preemptive actions, and my personal viewpoint is that this is desirable.

The working paper from the IMF is less than appreciative, or fair, in this respect. The authors do not seem to realize that publishing the forecast in this way forces us to act *now* on interest rates in order to drive inflation to its target level 18/24 months hence. Instead they question whether the forecasts are “credible.” Thus they state:

Since the MPC took over the job of monetary policy decision making, the analysis and the inflation forecasts reported in the Inflation Report cannot be treated as those of an independent reviewer: the body that makes interest rate decisions also assesses these decisions. As a result, it would be difficult to envisage a situation where the forecasts suggest that the two-year ahead inflation, on which public scrutiny tends to focus, would (in probabilistic terms) be missed, because then the report would in fact be questioning the committee’s own policy decisions. Indeed, since the Bank became independent, two-year ahead inflation has always been around the target. While this is possible logically, it raises questions as to the credibility of the Bank’s inflation forecast. (p. 15)

In so far as this suggests that, rather than adjust interest rates to hit the target, we manipulate the forecast to appear to be hitting the target, it is a

slur which members of the MPC would find offensive. Moreover in the large scale exercise involved in such a forecast, involving many Bank staff as well as all the MPC, such cosmetic manipulation could not get done without notice and huge risk of becoming publicly known.

One syndrome that I sometimes seem to perceive (both in myself and occasionally also in others) is the belief that the quantum of uncertainty (fog) will be reduced by waiting (e.g., to make a decision). Current shocks, implausible data points, and unintelligible anomalies are always very clear. But the *mean* expectation of future shocks is zero; hence, some of us may tend to assume implicitly that they will actually turn out to be zero. If that were so, uncertainty would indeed reduce over time. But subject to a qualification below, the distribution of shocks is *constant* over time; so, the best expectation of the future quantum of uncertainty is that it will be *constant*. If, as I believe, uncertainty myopia exists, then a procedure forcing an early decision may well be desirable.¹⁵

When I was a member of the MPC I thought that I was trying, at each forecast round, to set the level of interest rates so that, without the need for future rate changes, prospective (forecast) inflation would on average equal the target at the policy horizon. This was, I thought, what the exercise was supposed to be. But, if so, one might have expected short-term interest rates to follow, at least to a reasonable approximation and from quarter to quarter (*not* month to month), a random walk. But, in practice, changes in interest rates appear to have been just as serially correlated and gradual thus far under the MPC as in previous regimes and in other countries. If, however, we recognize that the appropriate periodicity is quarterly, then there have been only 12 observations of the MPC period, which is short for any formal testing of whether there was any significant shift in the extent of serial correlation.

Nevertheless my (casual) observation suggests there was not. Why? My preliminary hypothesis is that, during this period at least, there was some

¹⁵ There may be some seasonal patterns in uncertainty. There is some tendency in the U.K. for large negatively correlated movements in consumption in December and January, so assessment of consumption trends may be better left till late winter. Again, wage settlements and bonus payments tend to be concentrated in the winter and spring, so there is more uncertainty in December-June than in July-November. There are other instances of time-varying uncertainty relating to holidays, budget dates, etc. More research could be done to ascertain whether such seasonal patterns in uncertainty really do exist; it would be a nice simple exercise.

Table 1

Error Made in Forecasting

Date of:			
Forecast	Outturn	Output	Inflation
November 97	Q1:98	-0.58	0.08
February 98	Q2:98	0.64	0.32
May 98	Q3:98	0.42	0.20
August 98	Q4:98	-0.58	-0.03
November 98	Q1:99	-0.10	-0.03
February 99	Q2:99	0.67	-0.23
May 99	Q3:99	0.91	-0.12
August 99	Q4:99	1.05	-0.03
November 99	Q1:00	-0.18	0.09

NOTE: Positive values imply an outcome greater than forecast.

serial correlation in the forecasting process itself, i.e., that, whenever inflationary pressures were on the upward tack, or alternatively the downward tack, the forecasters initially underestimated the extent of such pressures. If so, the forecasters were playing “catch up,” and that led the MPC into a series of auto-correlated sequential decisions, although the MPC was not itself consciously trying to be “gradual” or inertial.

There is some plausibility in this hypothesis. By definition forecasters cannot predict shocks, such as that which hit in the late summer of 1998 or the sharper than expected recovery in 1999. Uncertainty, herd instinct among forecasters, and extrapolation from the past are all likely to make the expected (forecast) path smoother than the actual path. How far the serial correlation in interest rates derives from serial correlation in forecast errors, rather than from a conscious decision-making process, has not (yet) been properly researched, and I did not have the time/capacity to do it for this paper.

There would be considerable difficulties in doing this exercise in any case. If you take, say, the forecast *two years* hence and compare it with the *ex post* outcome, then, in the interval, both monetary and fiscal domestic policies will have changed, often sharply, from those that conditioned the original forecast. Should one do a counter-factual by trying to adjust the outcome for the intervening policy changes? But those policy changes will, in turn, have partly (largely) been in response to

shocks that the original forecasters could not have foreseen. On the other hand, if you take a very short-term forecast, say, over the coming quarter, there is usually sufficient inertia in the system that it is not a “model” forecast at all. Anyhow, in the one simple exercise that I asked to be done, the emphasis has been put, rightly or wrongly, on the *two-quarter-ahead* error in the forecast.

This is much more difficult to do for output than for the RPIX series, because the latter never gets revised, and the MPC puts weight on the latest actual figure. For output, the data get continuously revised, and the latest datum point at the time of the forecast is, itself, provisional. Moreover the quarterly output series are very noisy, more so than the annualized RPIX series. Even so, I asked the Bank staff to work out the MPC’s forecast errors for output and inflation two quarters hence. Given the short period, noise in the system, etc., the results can be at most suggestive. They are noted in Table 1. Note that there is some slight suggestion of inflation predictions being revised upward in the earlier part of the period *despite* a series of interest rate increases; then, this same pattern reverses in the middle of the period. This is consistent with my hypothesis that “gradualism”—the auto-correlation of interest rate movements—was, at least during this period, partly a function of the forecasting process and *not* of the decision-making process. It would, however, take much more careful empirical work to discover whether this was also the case in other regimes and countries. I just do not know, but I think it an exercise worth exploring.

IV. ON SKEWS AND ASSET PRICES

When the Bank was preparing its inflation forecast as background for the Chancellor’s decision on interest rates, the Bank could consider risks to its forecast in an unconditional way and to the best of my belief did so. But once the Bank obtained operational independence, the situation changed, since the question of what risks it would take into account would affect the decision itself. Let me first outline why it would affect the decision and then go on to argue that one would want the current decision to be affected in different ways, in some cases not at all, depending on the particular characteristics of each kind of risk.

When the probability distribution is symmetric, all three measures of central tendency (mode, median, and mean) are the same. In such cases the variance of the distribution, the extent of

uncertainty in the forecast, tends to play virtually no role in the interest rate decision. Discussion about forecast variances in the MPC were usually brief and uncontroversial; in this respect there was a kind of certainty equivalence. Pure variance was an academic issue.

This was not so for risk, because when we (as MPC members) think of risk, we are generally concerned with asymmetric possible outcomes. Asymmetry implies skew, and skew drives the measures of central tendency apart. The question of which measure of central tendency one should focus upon depends on one's own individual loss function.

Most forecasts, including that of the Bank, concentrate on the most likely outcome, i.e., the mode, and usually consider "the risks to the forecast" at the tail end of the exercise. But the interest rate decision should be taken only on the basis of the modal outcome if the occurrence of *any* deviation from the actual modal forecast, however small, is taken to be catastrophic. The usual academic assumption is that the loss function is quadratic in deviations from the desired outcome, in which case the focus should be on the mean forecast (not the mode). In my own case, I could never see why a 2 percent deviation from desired outcome was four times as bad as a 1 percent deviation, rather than just twice as bad. So I tended to concentrate on the median forecast, while others of my colleagues focused on the mean.¹⁶ But in either case the existence of a skew would affect our decision on the appropriate interest rate. Unlike uncertainty and variance, skew and risk mapped directly into the interest rate decision. So, discussions about risks were policy-relevant and not just academic.

But should all perceived risks affect policy now? I would argue not. Let me take some examples. Assume that there has been an ongoing "bubble" in housing prices, which have been driven far out of line with (perceived) fundamentals. The inflationary effect of this is, I assume, offset by other deflationary forces (say higher productivity) so the modal forecast is in line with the target. The risk is that the bubble will burst, and when it does this will be deflationary. So, if you take account of this risk in the forecast, the mean is driven below the mode. So the implication is that interest rates should now be cut!

This strikes me as a nonsensical response. The contrary argument is that faced with such an asset price bubble, interest rates should be (temporarily)

raised in order to check the misalignment and to prevent it from getting further out of hand and presumably collapsing with an even bigger detonation in the future.

This is *not*, alas, a hypothetical problem. Housing prices in 1988-90 in the U.K., equity prices in the U.S., and exchange rates in the U.K. all raise aspects of this question of how to deal with risk and skew. The MPC was criticized in the press after its May 2000 inflation forecast for not including the biggest risk of all, that of a sharp fall back in the exchange rate (notably against the euro), in its specific quantitative forecast. But if it had, that by itself would have implied that interest rates should have been tightened immediately.

An astute reader will by now have noted that this subject relates back to the earlier issue of conditioning on a constant future interest rate assumption. The point is that there are some categories of risk which should not involve countervailing action now, but only if and when they actually occur. In the meantime there is a case for nudging rates in the opposite direction, if that should help to reduce the misalignment and the risk and scale of the subsequent hard-landing collapse. This latter case was made in Cecchetti et al. (2000); it remains highly contentious. It is not the purpose of this paper to take sides on that (and anyhow I am ambivalent). All I want to argue is that there exists a sizeable category of risks whose future potential effects should be ignored in setting interest rates today.

Note that, so long as the inflation forecast continues to be conditioned on future unchanged interest rates, quantitative account of such risks are better omitted from the fan charts of output and inflation. The Bank was actually right to do just this in its May 2000 forecast. If *all* risks are to be included in the inflation forecast unconditionally, then the quid pro quo would have to be some discussion of future, non-constant paths for interest rates. As earlier described, this has some serious drawbacks and disadvantages. But the choice seems clear enough to me. Either condition on constant interest rates and exclude a certain category of

¹⁶ Commentators might expect the nature of the (individual member's) loss function to be affected by the 1 percent deviation from target, acting as a trigger for writing a letter of explanation and intent to the Chancellor. I can only say that in my own case it did not. Perhaps this was because inflation remained during these three years so close to the target, both currently and at the forecast horizon, that the possibility of having to write such a letter remained more hypothetical than immediate.

risk from the (formal) forecast or include all risks unconditionally and describe how the future path of interest rates might change over time in response.

What are the characteristics of the risks that should *not* lead to countervailing interest rate adjustments now? The first, and perhaps most important, is that the event of the risk coming into effect should be clearly observable. Examples are political or natural disasters (war in the Middle East, earthquakes in Tokyo, hurricanes in Washington, Liberal Democrats getting elected in the U.K. [only joking], etc.), and sharp changes in asset or commodity prices (exchange rates, oil, equity and housing prices, perhaps certain food prices).

Examples where the event is not observable, observable only with great difficulty, or observable too long after the event to be useful are changes in the non-accelerating inflation rate of unemployment, in future (trends in) productivity, in competitive pressures on margins (e.g., from the Internet), on international competitiveness, etc., etc. In so far as there are thought to be asymmetric risks attached to these, you need to take them into account now because you will never find a time in the future when you can perceive with any greater clarity or certainty whether or not they have been occurring.

A second, main characteristic of the risks that policy should *not* try to preempt is that they are low-probability events with a high payoff, if they should transpire at all.¹⁷ Fortunately the list of clearly observable risks and low-probability/high-payoff risks overlap rather closely.

Let me again take some examples. You should not run a systemically mildly inflationary policy because there is a non-zero risk of a 1929 (or a Japanese) collapse in asset prices. Nor should you run a systematically mildly deflationary policy because there is a non-zero chance of an exchange rate collapse or of entering an inflationary war.¹⁸ Similarly the distinction between observable and not-observable risks is not so easy to draw.

As far as I know the academic economics profession has not yet provided much help on the question of how to respond to asymmetric risk. Error terms, whether related to additive or multiplicative uncertainty, are universally taken to be symmetrically distributed. Standard practice is to run a forecast as if it was the mean of a symmetric distribution and then to add an asymmetric risk as an ad hoc afterthought.

Help on how better to treat asymmetric risks would be much appreciated. Until such assistance arrives, I stand by my view that some such risks, i.e., clearly observable, low probability/high payoff risks should *not* induce current countervailing action. This conclusion has, as already noted, some implications for the conduct and public presentation of the forecast process itself.

V. CONCLUSION

It will certainly be desirable to do further research on what either an optimal control or an assorted rule-based model would imply in the context of U.K. forecasting systems. But that should initially be for internal use only. Whether or not such further research might lead the MPC to change its forecasting assumption about a constant interest rate path is far from clear. Until such time, for the reasons set out above, there seem to me to be good grounds for maintaining present procedures in this respect. Nevertheless I would see some merit in exploring the presentation of an alternative, but equally arbitrary, set of scenarios, as in the charts in Section III (i.e., Figure 4).

Unless some such move is made away from the constant assumed time path of interest rates, those risks which should only evoke a policy response in the (unlikely) event of them occurring will have (to continue) to be excluded from any quantified inclusion in the inflation forecast. Only if the time path of interest rates can be presented as “conditional” can risks be included in an “unconditional” way.

A proper research exercise should be commissioned to examine whether the auto-correlation observed in interest rates has been due to the forecasting process rather than to the decision-making process.

In so far as the assumption of a constant time-path for interest rates was taken as an effective

¹⁷ Low-probability, low-payoff risks will get ignored.

¹⁸ Considerations of this kind make me somewhat unhappy also about the arguments of those who argue for “robust” policies (e.g., Sargent, 1999, and Stock, 1999). They suggest setting policy so as to minimize the likelihood of having a really bad outcome (a form of minimax policy). Since the worst outcome, from the point of view of a central bank, is to find oneself powerless to prevent worsening deviations from target, their policy prescription is to be more aggressive. But if I give a (very) low subjective probability, say 2 percent, to such an outcome (especially if I can observe it should it happen), does it make sense to aim consciously for a poor outcome (98 percent of the time) in order to offset a real disaster (2 percent of the time)?

constraint, rather than just a conditioning assumption, it would have the effect of inducing the monetary authority to choose a probably suboptimal and often noncredible path. But if correctly interpreted as just a conditioning assumption, it simplifies the decision-making process, making it easier to understand and to explain publicly; provides the greatest flexibility as well as minimum forward constraint on the MPC's future actions; and encourages the MPC to act in a robust, preemptive fashion.

REFERENCES

- Bank of England. Minutes of the August 1997 Monetary Policy Committee Meeting, Inflation Report. November 1997.
- _____. Inflation Report. February 2000.
- _____. Minutes of the April 2000 Monetary Policy Committee Meeting, Inflation Report. May 2000.
- Barro, R. and Gordon, D. "Rules, Discretion and Reputation in a Model of Monetary Policy." *Journal of Monetary Economics*, July 1983, 12(1), pp. 101-21.
- Batini, N. and Haldane, A. "Forward Looking Rules for Monetary Policy." Working Paper 91, Bank of England, 1999.
- Batini, N. and Nelson, E. "Optimal Horizons for Inflation Targeting." Working Paper 119, Bank of England, 2000.
- Cecchetti, Stephen G.; Genberg, Hans; Lipsky, John and Wadhvani, Sushil. "Asset Prices and Central Bank Policy." *Geneva Reports on the World Economy*. Vol. 2. ICMB and CEPR, 2000.
- Cukierman, A. *Central Bank Strategy, Credibility, and Independence: Theory and Evidence*. Cambridge, MA: MIT Press, 1992.
- Flemming, J.S. "Monetary Policy, Adequacy, Design and Presentation." London Business School and Oxford Economic Forecasts, *Autumn Economic Outlook*, 1999, pp. 9-13.
- Friedman, M. "A Monetary and Fiscal Framework for Economic Stability." *American Economic Review*, 1948, 38(3), pp. 245-64.
- King, M. "The Inflation Target Five Years On." London School of Economics Lecture on the 10th Anniversary of the LSE Financial Markets Group, 29 October 1997.
- Martijn, J.K. and Samiei, H. "Central Bank Independence and the Conduct of Monetary Policy in the United Kingdom." Working Paper 99/170, International Monetary Fund, 1999.
- Monetary Policy Committee, Bank of England. "The Transmission Mechanism of Monetary Policy." London: Bank of England, 1999.
- Rudebusch, G. and Svensson, L. "Eurosystem Monetary Targeting: Lessons from U.S. Data." Working Paper 7179, National Bureau of Economic Research, 1999a.
- _____ and _____. "Policy Rules for Inflation Targeting," in John B. Taylor, ed., *Monetary Policy Rules*. Chicago: University of Chicago Press, 1999b, pp. 203-46.
- Sargent, T.J. "Comment" on Laurence Ball's "Policy Rules for Open Economics," in John B. Taylor, ed., *Monetary Policy Rules*. Chicago: University of Chicago Press, 1999, pp. 144-54.
- Stock, J.H. "Comment" on Rudebusch and Svensson's "Policy Rules for Inflation Targeting," in John B. Taylor, ed., *Monetary Policy Rules*. Chicago: University of Chicago Press, 1999, pp. 253-59.
- Svensson, L. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." Working Paper 56, Bank of England, 1996.
- Thornton, D.L. and Wheelock, D.C. "A History of the Asymmetric Policy Directive." Federal Reserve Bank of St. Louis *Review*, September/October 2000, 82(5), pp.1-16.
- Williams, J.C. "Simple Rules for Monetary Policy." Working Paper, Board of Governors of the Federal Reserve System, 1997.

