An Analysis of the Impact of Finite Horizons on Macroeconomic Control

by

Andrew P. Blake

and

Peter F. Westaway

June 1992

Macroeconomic control exercises which use large-scale non-linear consistent expectations models necessarily employ a finite time horizon since data bases must be constructed to simulate such models. Expectations past the terminal date must be proxied by terminal conditions, which can distort the solution. In contrast, much of the analytic literature has focused on the linear-quadratic infinite time horizon problem, which for the closed form solution is easy to derive and offers convenient interpretation. This paper examines the extent to which the necessary imposition of a finite time horizon distorts the optimal policy not only through the terminal condition but through the impact the finite horizon can have on expectations. This is carried out by comparing the results obtained from a large non-linear consistent expectations model (the National Institute UK model) with a linearised representation of the same model.

National Institute of Economic and Social Research
2 Dean Trench Street
Smith Square
London SW1P 3HE

An Analysis of the Impact of Finite Horizons on Macroeconomic Control

Andrew P. Blake and Peter F. Westaway†

1 Introduction

This paper is concerned with the practical application of optimal control techniques to the macroeconomic policy design problem using empirically based econometric models which embody forward-looking behaviour. In the context of UK models, such work can be divided into two broad approaches. One has involved the extension to rational expectations models of the well known non-linear optimal control algorithms[1] previously applied to backward-looking models, as in Holly and Zarrop (1983) and Westaway and Wren-Lewis (1992).[2] The alternative strategy has been to derive linearised versions of the non-linear models, thus allowing the more tractable properties of linear rational expectations models to be exploited, as in Christodoulakis *et al.* (1991) and Weale *et al.* (1989). Surprisingly, relatively little work has been done in attempting to reconcile these two approaches, in particular in comparing the policy implications which might emerge from applying the different techniques to the same problem. In this paper we explain why these differences might be important by emphasising the role of the finite horizon which we show can have large distortionary effects on the optimal control outcome.

Typically, macroeconomic models are large and non-linear; for example the National Institute model of the UK economy on which the empirical analysis in this paper is based contains just over 350 variables with the behavioral content embodied in around half of these (see NIESR, 1991). Inevitably, these models can only be solved using numerical techniques, for example using an extended Gauss-Seidel method (see Fisher and Hughes Hallet,

---

† The authors are grateful to Paul Fisher, members of the National Institute macromodelling team and seminar participants at the University of Warwick and at the 1991 meeting of the Society for Economic Dynamics and Control, in Capri, Italy for comments on earlier drafts of this paper. The macroeconomic programme at the Institute is financed by the ESRC.

[1] See Rüstem and Zarrop (1979) for example.

[2] These two applications differ in that Holly and Zarrop (1983) impose model consistent expectations via the optimisation algorithm itself, known as the penalty function method, while Westaway and Wren-Lewis (1992) adopt an iterative scheme to enforce expectational consistency which is independent of the optimisation.

so great.

The plan of the rest of this paper is as follows. Section 2 reviews the role of the terminal condition in finite horizon solution, illustrating the analytic results with both the non-linear and linearised model. Section 3 extends this discussion to consideration of the finite horizon control problem. In section 4, an empirical example is given using the National Institute model of the UK economy. The crucial role of instrument damping is illustrated for the case where the government is less forward-looking than the private sector. We also show that the non-linear control solution is close to the linear solution based on the local derivatives. Accordingly, in section 5, we carry out control exercises on a linearised version of the model (which we show in the appendix has remarkably similar properties to the non-linear model itself). We address the same issues as section 3, but now illustrate the implications of extending the time horizon and compare the results where possible with the infinite horizon analytic solution. Section 6 extends the analysis by considering the same issues under the alternative assumption that the reputational solution is not sustainable and only the 'discretionary' time consistent solution is available. Section 7 concludes by drawing out the implications for conventional approaches to empirically based policy analysis.

2 Finite horizons and terminal conditions

It is well known that the solution to a non-linear macroeconomic model requires the specification of a terminal condition if the model involves forward-looking expectations and is solved over a finite horizon. This arises because the solution will be affected by the expected values of variables beyond the end of the solution period. Conventionally, a terminal condition is imposed which assumes that the variable of which expectations are formed has reached its equilibrium path (Minford *et al.*, 1979). The equilibrium path can be fixed values for the expectations past the terminal date, but is more usually a constant (or zero) growth condition.

The question for practitioners is to set a terminal condition which does not distort the notional true solution by imposing such equilibrium behaviour in advance of the date at which this is appropriate. It is often argued that this can be done by setting the terminal

date of the simulation sufficiently far into the future so that a change in the terminal date (or indeed a change in the specification of the terminal condition itself) does not alter the solution over the period of interest.

We can illustrate these issues by considering a case where the correct terminal condition is available analytically but we impose an arbitrary one instead. Blake (1990) shows how terminal conditions affect the solution of linear rational expectations models. This helps us analyse the error in the initial period jump for variables where the terminal condition is enforced at too early a date. A simple linear model is given by

$$\begin{bmatrix} z_{t+1} \\ x^e_{t+1} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} z_t \\ x_t \end{bmatrix} \tag{1}$$

where z_t is a predetermined state variable and x_t a non-predetermined (expectational) variable which is free to take on any value at time t.[5]

Following Blanchard and Kahn (1980) diagonalise the transition matrix so that

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} = \begin{bmatrix} 1 & m_1 \\ 1 & m_2 \end{bmatrix}^{-1} \begin{bmatrix} \lambda_s & 0 \\ 0 & \lambda_u \end{bmatrix} \begin{bmatrix} 1 & m_1 \\ 1 & m_2 \end{bmatrix} \tag{2}$$

where λ_u is the unstable eigenvalue (greater than unity), λ_s is the stable eigenvalue (less than or equal to unity) and the matrix of left eigenvectors, partitioned conformably, has the first element normalised to unity. Setting

$$x_t = -(1/m_2)z_t = -nz_t \tag{3}$$

ensures that the unstable roots are cancelled. It then follows that the predetermined variable dynamics are given by

$$z_{t+1} = (a_{11} - a_{12}n)z_t = \lambda_s z_t. \tag{4}$$

At time T the *correct* terminal condition is the same, *i.e.* $x_T = -nz_T$.[6] Remember that for non-linear models this cannot be calculated analytically, requiring an *ad hoc* alternative

[5] This two variable model is a very simple one. We can augment the state vector to include further predetermined or even exogenous variables (e) with their own dynamic equation. The simplest is $e_{t+1} = \alpha e_t$. Setting α to unity gives step shocks, and to zero impulse shocks. Further expectational variables could be included. The results in this section generalise straightforwardly.

[6] We might refer to this as a *transversality condition* rather than a terminal condition in this context. For a saddlepath stable linear model many arbitrary terminal conditions will impose stable behaviour over a finite horizon, but no terminal condition other than the transversality condition will ensure stability over any time horizon.

terminal condition. To keep things simple, we assume a zero growth condition, $x_T = x_{T+1}$. In the context of our linear model, this can be expressed as the requirement that

$$x_T = \frac{a_{21}}{1 - a_{22}} z_T. \tag{5}$$

Using this as a terminal condition, the general result of Blake (1990) shows how the saddlepath condition needs to be modified. In this case, n in (3) should be replaced by

$$\hat{x}_t = -\frac{(1 - G_{T-t})}{(m_2 - G_{T-t}m_1)}\hat{z}_t = -\hat{n}_{(T-t)}\hat{z}_t \tag{6}$$

where

$$G_{T-t} = \left(\frac{1 - m_2(a_{21}/(1 - a_{22}))}{1 - m_1(a_{21}/(1 - a_{22}))}\right)\left(\frac{\lambda_s}{\lambda_u}\right)^{(T-t)}$$

and $\hat{x}_t$ and $\hat{z}_t$ are the values of the free and predetermined variables found using the terminal condition. Similarly, (4) is modified to

$$\hat{z}_{t+1} = (a_{11} - a_{12}\hat{n}_{(T-t)})\hat{z}_t \tag{7}$$

with the error in the initial jump is given by

$$\hat{x}_0 - x_0 = -\left(\frac{1 - G_T}{m_2 - G_T m_1} - \frac{1}{m_2}\right)z_0.$$

These expressions illustrate clearly how the distortion to the true analytic solution will be smaller, the closer G_k is to zero for any positive k. This will be achieved exactly if $a_{21}/(1 - a_{22}) = -1/m_2$ implying that the imposed terminal condition reflects the true behaviour of the system at that time.[7] Alternatively, and more usually, the terminal horizon can be chosen to be sufficiently long. This reduces the error in the earlier periods of the simulation since G_k tends to zero as k increases because it is discounted by the ratio of the stable to unstable roots, which must be less than unity.

Intuitively it is clear that if the predetermined variables evolve so that a shock is temporary ($\lambda_s < 1$), then the temporary effects of G_T will gradually die away until the original equilibrium path is reached, even if the unstable root λ_u is close to unity. Alternatively,

[7] Imposing this correct terminal condition will only be equivalent to the more usual rate of growth condition as the time horizon becomes large.

if any shock is permanent ($\lambda_s = 1$), so that z tends to a new steady state value, then the error on the initial jump is certain to disappear only when λ_u^{-T} tends to zero. For, say, an unstable root $\lambda_u = 1{\cdot}01$, T has to be very large.

The standard way to test whether the terminal condition is distorting the solution is to solve the model over different time horizons for the same shock and check whether the jump is invariant. This procedure is described for a range of large macro models in Wallis *et al.* (1986), and is carried out on a slightly older version of the NIESR model in Ireland and Wren-Lewis (1989). We briefly repeat this exercise here for the version of the model to be used in our empirical work.[8] The only forward-looking variable in our model is the exchange rate. The linearised model has an unstable root very close to the unit circle at around 1·01; this implies very forward-looking behaviour. We examine the sensitivity of the initial jump in this variable to an impulse shock to government spending. Figure 1a illustrates the results for 10 and 12 year simulation horizons. These results seem to confirm the conclusion of Ireland and Wren-Lewis (1989) that a simulation horizon of ten years is sufficient to avoid distortion of the results.

This has an important implication for the control problem. If we allow instrument movements over the entire simulation horizon then instruments could be (and we will argue below will be) moving sharply at the end of the simulation horizon. But for the terminal condition not to be affecting the simulation we need at least 10 years at the end when instruments constant. This implies a 'tail' is required to the optimisation where the instruments are constrained to a neutral setting, either a constant level or a steady growth rate, to allow the terminal condition effects to be minimised.

The conclusion that a 10 year tail is sufficient may be too optimistic since, as discussed above, it may not apply in the case of shocks where the steady state values change. This point is partly illustrated in figure 1b which shows how the initial jump in the exchange rate for a permanent shock to government spending does not seem to settle down as the simulation horizon is extended. This is certainly informative but fortunately we can more easily analyse the implications of changing the terminal date by working with a linearised version of the model. This is derived in the same way as in Weale *et al.* (1989), and

[8] Details of the modifications we made to the model are given in the appendix.

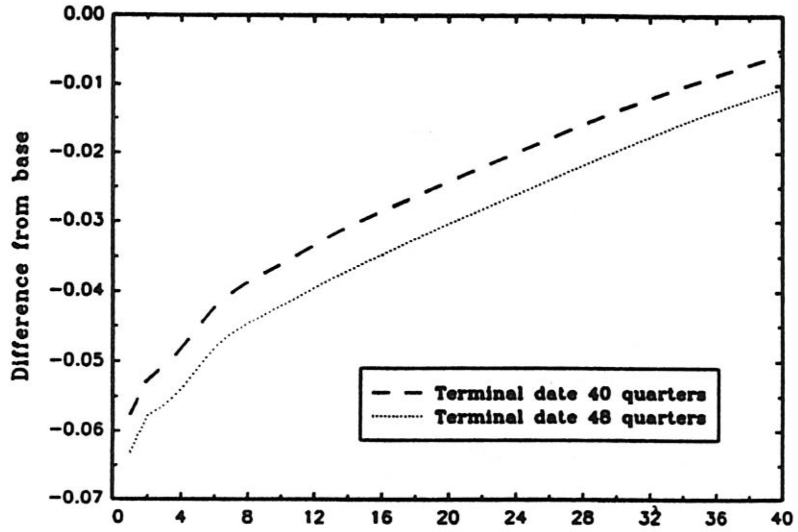

FIG. 1a. Exchange rate response to an impulse change to government spending.

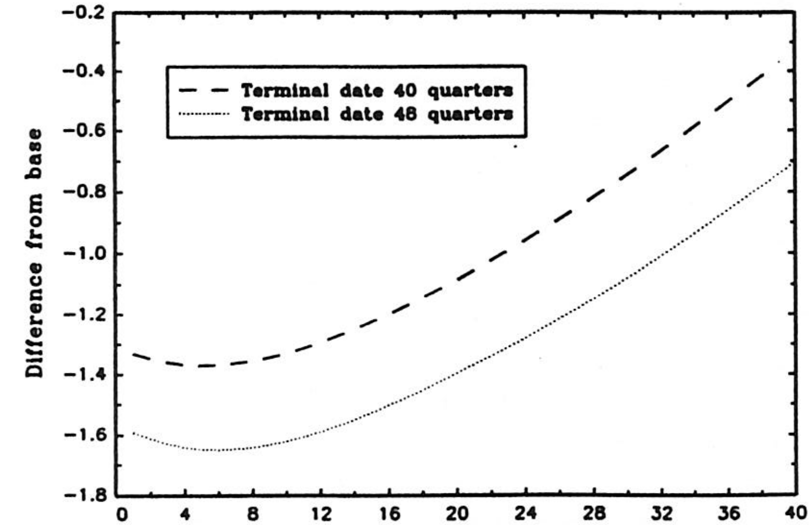

FIG. 1b. Exchange rate response to a permanent change to government spending.

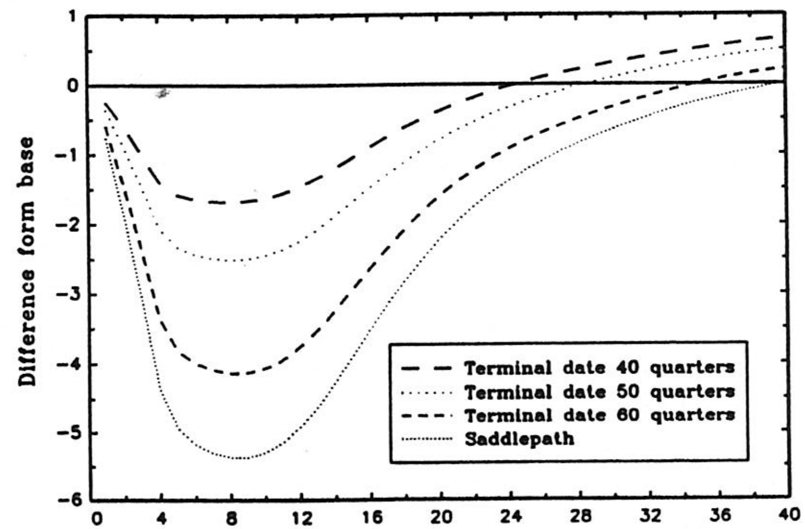

FIG. 2. Real exchange rate response to a permanent 0·5% increase in interest rates.

similar to the procedure described in Christodoulakis *et al.* (1991). For details of how the linearisation was obtained and a comparison with the non-linear model see the appendix. This allows us to solve the linearised model over an extended time horizon and impose the same terminal condition on the linear model as that used for the non-linear model.

Figure 2 illustrates the simulation properties of the linearised model, showing the response of the real exchange rate to a permanent increase in interest rates. Four simulations are reported, three where the terminal condition is imposed at 40, 50 and 60 quarters, and the infinite horizon solution found using the Blanchard-Kahn algorithm. Clearly, the terminal condition has a very marked effect on the initial jump. Indeed, even with a horizon of 100 quarters, it is still some distance away from the analytic solution; for this shock it requires a solution calculated over about 500 quarters before the two converge.

So far, both the analytic example and the empirical results strongly suggest that any simulation results from a model which is characterised by a high degree of forward-lookingness (*i.e.* with an unstable root close to the unit circle) will be substantially affected by the finite horizon through the mechanism of the terminal condition. We now turn to consideration of its influence on the control problem itself.

3 Finite horizons and optimal control

The finite nature of the terminal date also has implications for the optimal control problem, quite independent of the issues regarding the solution of forward-looking models described above. In fact, the difference between the finite and infinite time solutions is present even for purely backward-looking models. The problem manifests itself as sharp changes in the policy instruments near the end of the optimisation horizon. This well known result is easily explained for the backward-looking case. Because of the recursive nature of the instruments on the targets, the finite time control solution can be computed one period at a time taking the terminal period first and working backwards (see Bellman, 1957). Since policy instruments do not have anticipated effects, the policy setting in the last period will only be set with reference to the cost in that period. This may imply that it will be beneficial to manipulate policy in the final few periods if, say, an instrument movement involves short term benefits but long run costs which are incurred beyond the optimisation horizon. When the model is causal, it is possible to minimise the distortion to the results by setting the optimisation horizon sufficiently far in the future. This will not necessarily remove the irregular movement of policy instruments near the terminal date but it will render the chosen target variables unaffected over the initial periods, the immediate horizon of interest. In fact, it has been demonstrated that this finite horizon distortion will tend to zero in such cases subject to certain regularity conditions described in Flåm and Wets (1987); for forward-looking models, however, we explain how the finite horizon solution may not necessarily converge on an analytic infinite time solution, even for an arbitrarily long terminal horizon.

The appropriate solution techniques for the control problem applied to forward-looking models are also well known. For linear models, the infinite horizon analytic solution is described in Levine (1988), for example. Most practical applications, however, involve non-linear models solved over finite horizons. In these cases, the control solutions are based on a Newton-type algorithm which calculates local derivatives numerically and solves the model iteratively, assuming at each stage that the model is linear (see Rüstem and Zarrop,

1979). This algorithm applies equally to forward- and backward-looking models, although there have been relatively few applications of finite horizon optimal control on forward-looking non-linear models (see Westaway and Wren-Lewis (1992) for one example). Yet there are a number of interesting issues that arise when finite horizon control methods are applied to models with forward-looking expectations.

First, however, it is instructive to describe the policy problem and the rational expectations solution for the same simple analytic model employed in section 2. We now introduce policy instruments where their open-loop trajectory is taken as given; for now, this is convenient for expositional purposes since the instrument settings in the finite horizon control case can be interpreted as exogenous variables which are to be manipulated to derive the optimal outcome. Accordingly, the linear model (1) is modified to

$$\begin{bmatrix} z_{t+1} \\ x^e_{t+1} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} z_t \\ x_t \end{bmatrix} + \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} u_t \qquad (8)$$

where u_t is our policy instrument. The assumed policy problem is to minimise the following discounted quadratic loss function by choice of u

$$C = \frac{1}{2} \sum_{t=0}^{T} \rho^t (q z_t^2 + u_t^2) \qquad (9)$$

where T is the optimisation horizon, ρ is the discount factor, the target is to drive the predetermined state to zero and we include instrument costs, with a weight normalised to unity. More generally we would allow all variables to be vectors, we could include time-varying weights, or we could incorporate costs associated with the non-predetermined states and for non-zero targets, but the essence of the analysis would be unchanged. A necessary condition for the solution to this problem to be well defined is that the discounted sum of the target variables should be finite; this will be guaranteed if the rate of discounting is greater than the rate of growth of the target deviation which in our case is zero; this is sometimes referred to as Grinold's condition, as in Flåm and Wets (1987).

To find the rational solution we need to generalise (3) and (4). Expectational variables depend on future expected instrument values, and are calculated using a modification to (3)

$$x_t = -n z_t - (n b_1 + b_2) \sum_{i=0}^{\infty} \lambda_u^{-i-1} u_{t+i} \qquad (10)$$

where n and λ_u are as before. If the policy instruments are zero this collapses to the previous solution. Remembering that $\lambda_s = a_{11} - a_{12}n$ we modify (4) and the predetermined state evolves according to

$$z_{t+1} = \lambda_s z_t + b_1 u_t - ((a_{11} - \lambda_s)b_1 + a_{12}b_2) \sum_{i=0}^{\infty} \lambda_u^{-i-1} u_{t+i}. \tag{11}$$

Equations (10) and (11) demonstrate the significant difference between backward- and forward-looking models, that for the latter future instrument values affect the current state. It is apparent that the size of the unstable root is crucial in determining how much future instrument values affect current behaviour, as the inverse acts as a discount factor.

We have already noted in section 2 that the finite horizon control problem implies that the instrument must be set at some constant rate for the duration of the 'tail' up to the terminal date and implicitly beyond that date as well. Let f be the time up to which there is active control, and T be the terminal date. If u is set at some constant level u_∞ after $f \leq T$, then for $t < f$

$$x_t = -nz_t - (nb_1 + b_2) \sum_{i=0}^{f-t} \lambda_u^{-i-1} u_{t+i} - \left(\frac{nb_1 + b_2}{\lambda_u^{f-t}(\lambda_u - 1)} \right) u_\infty \tag{11a}$$

$$z_{t+1} = \lambda_s z_t + b_1 u_t - ((a_{11} - \lambda_s)b_1 + a_{12}b_2) \sum_{i=0}^{f-t} \lambda_u^{-i-1} u_{t+i} - \left(\frac{(a_{11} - \lambda_s)b_1 + a_{12}b_2}{\lambda_u^{f-t}(\lambda_u - 1)} \right) u_\infty \tag{11b}$$

and if $t \geq f$

$$x_t = -nz_t - \frac{nb_1 + b_2}{\lambda_u - 1} u_\infty \tag{12a}$$

$$z_{t+1} = \lambda_s z_t - \frac{(1 - a_{22})b_1 + a_{12}b_2}{\lambda_u - 1} u_\infty \tag{12b}$$

where we have used $a_{11} + a_{22} = \lambda_s + \lambda_u$ in deriving (12b). From (11) and (12), it is apparent that the last period value of the instrument exerts a significant influence throughout the entire optimisation because it implicitly carries on at that level for all periods past the terminal time. On these grounds, it is clear that the time horizon, f, at which the

instrument is set 'neutral' needs to be sufficient so that $\lambda_u^f(\lambda_u - 1)$ is a large number before the effects of the last period instrument can be safely assumed not to dominate the entire path of the target variables. The closer the unstable root to unity, the further away the tail needs to be.

The analysis so far, however, has only demonstrated the effect of a given change in the expected value of a future policy instrument on the current state. For the finite horizon control solution, we need to demonstrate that the effects of policy instruments in the distant future on the cost function itself should tend to zero as the optimisation horizon is extended. In fact, it is in this respect that there is a crucial difference between the backward- and forward-looking control problem. As with the backward-looking case, the standard control solution will deliver sharp changes to the policy instruments near the terminal date. But now, importantly, the effect of this distortion will not necessarily disappear as the terminal date is pushed further back. It is easy to demonstrate this for the simple univariate model described above. From equation (11), the effect of a change in the policy instrument at time N on the state variable at time j is given by

$$\frac{dz_j}{du_N} = -((a_{11} - \lambda_s)b_1 + a_{12}b_2)\lambda_u^{j-N-1} \left(\frac{(1 - r^j)}{(1 - r)} \right) \tag{12}$$

where $r = \lambda_s / \lambda_u$, from which we can derive the total effect of a unit change in the policy instrument on the cost function (9) as

$$\frac{dw}{du_N} = C(z) + C(u) \tag{13}$$

where the cost associated with the resultant movements in the state variables is given by

$$C(z) = \lambda_u^{-2N} \left(\frac{q((a_{11} - \lambda_s)b_1 + a_{12}b_2)^2}{(1 - r)^2 \lambda_u^2} \right) \sum_{t=1}^{T} \rho^t (1 - r^t)^2 \lambda_u^{-2t} \tag{14}$$

while the direct cost associated with the movement in the policy instrument is given by

$$C(u) = \rho^N \tag{15}$$

It is necessary for $C(u)$ to become larger than $C(z)$ as the movement in the policy instrument is extended further into the future (*i.e.* as N becomes larger); if not, instrument instability will result since it will always be possible to improve welfare by manipulating the policy instrument in the distant future to improve the overall value of the cost function since the costs of doing so will be increasingly outweighed by the benefits. Comparison of the expressions for $C(z)$ and $C(u)$ show that this instability problem will only be avoided if ρ/λ^2 tends to zero, that is if the discount factor is less than the squared inverse of the unstable root. This may be seen as a forward-looking extension to the conditions derived in Flåm and Wets (1987) for backward-looking models.

Once we extend this analysis to the multivariate framework, it will be the most forward-looking root of the model, that is the one closest to the unit circle, that must be compared with the discount factor. Indeed in our empirical work to be described below, we have found that small changes to the discount rate can make relatively small effects over the finite horizon but cause the infinite horizon solution to fail.

4 Finite horizon control on the non-linear model

In this section, we use the NIESR UK model (version 11.6, modified as described in the appendix), a large non-linear macroeconometric model of the UK economy to illustrate the issues raised in the last section. It is assumed that policymakers have two targets, GDP growth and price inflation, which they attempt to control using the level of short term interest rates and government expenditure (public authority consumption). Both instruments are constrained to have a five year tail out of a total optimisation horizon of 10 years (we should note that this tail is much too short since we showed in section 2 that a tail of 10 years was preferable even in the face of temporary shocks). It is assumed that the desired growth rate for GDP is at its base level while for inflation it is 3% below base levels (which approximates closely to a long run target of zero inflation). Deviations from these desired targets are weighted equally. A cost is also accorded to changes in policy instruments relative to base. Costs associated with deviations of the target variables are discounted by 0·99 a quarter while we examine the consequences of different assumptions regarding the discounting of instrument damping. Thus the objective function to be minimised is given by

$$C = \sum_{t=91Q1}^{00Q4} 0{\cdot}99^t[\dot{p}_t^2 + g_t^2 + k_t(\Delta r_t^2 + \gamma \Delta G_t^2)]$$

where $\dot{p}$ is the rate of price inflation (GDP deflator), g is the annual growth of real GDP (compromise measure) and Δr and ΔG are the one period changes in policy instruments. The base path used for the exercise is based on the National Institute February 1991 forecast.

Figures 3a–d illustrate the optimal control strategy which emerges when instrument movements are discounted at 0·99 a quarter in line with the target variables. As might be expected from the previous discussion, the results display the property that policy instruments are moved sharply in the last few periods (*i.e.* in the periods before the tail begins). More implausibly still, most of the improvement in inflation performance which occurs in the early years of the optimisation can be directly attributed to the appreciation in the exchange rate which is caused by the promise of the higher interest rates in the later

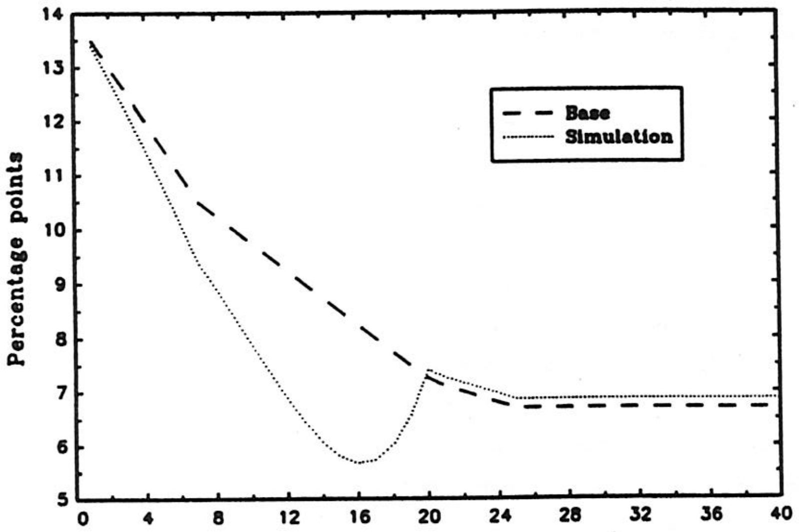

FIG. 3a. Treasury bill rate.

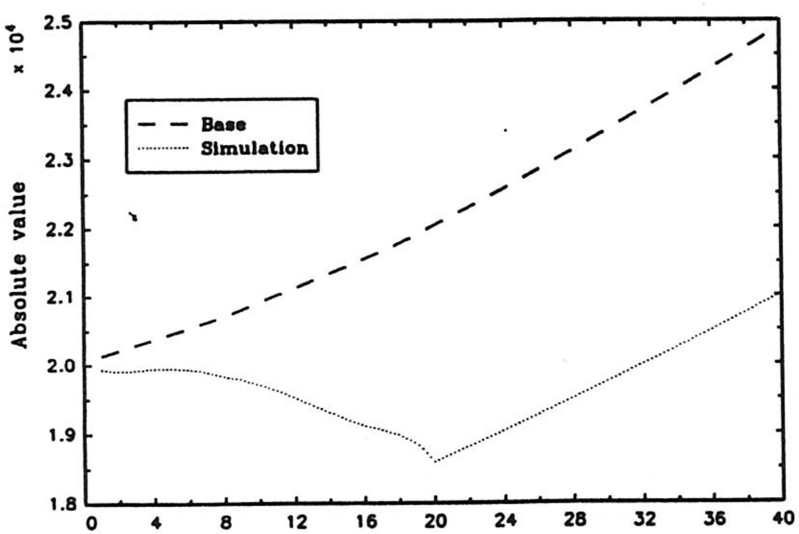

FIG. 3b. Public authority consumption.

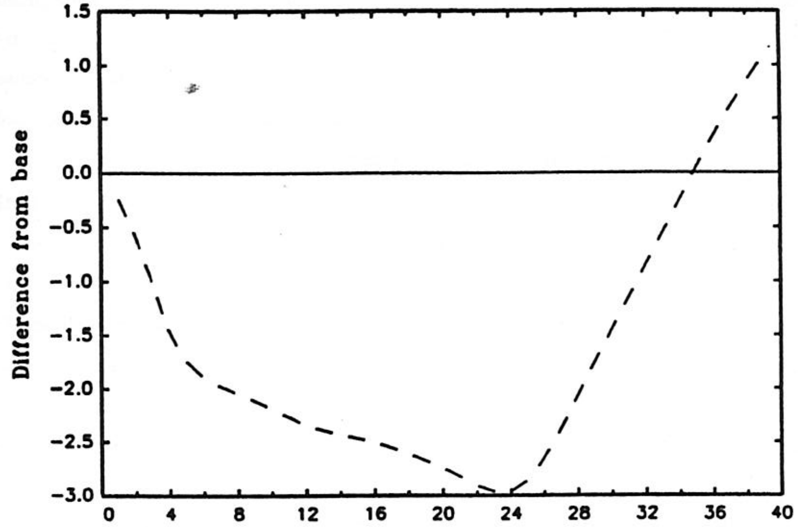

FIG. 3c. Annual price inflation.

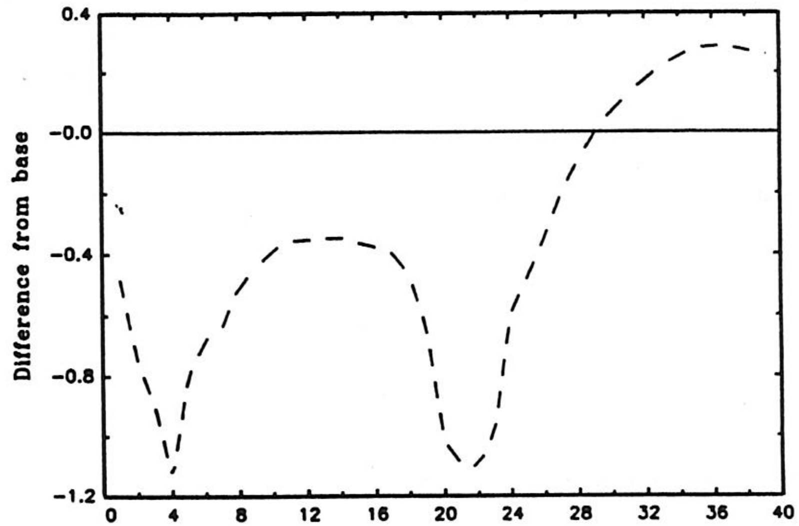

FIG. 3d. Annual GDP growth.

periods. Of course, this policy is a very stark example of the time inconsistency problem whereby policymakers would have a strong temptation to renege on this policy when the time came to raise interest rates. Yet this policy is implausible not simply because of the time inconsistency property but more importantly because such an announced policy setting would not be credible on the basis of observed past behaviour.

There are three possible responses to this problem. One is to derive the infinite time solution to the control problem instead. We defer this to the next section. An alternative is to address the issue of credibility more formally by deriving a time consistent solution; again, this is deferred to the discussion in section 6. Finally, the only practical choice if the model is genuinely non-linear is to build some mechanism into the optimisation problem which rules out such implausible instrument settings. One such mechanism would be to assume that the credibility of sharp changes in instruments diminishes the further into the future such changes are proposed. We attempt to capture this 'announcement credibility' mechanism by assuming that the weight on instrument damping increases through time; from the last section, we argued that the rate of increase of this damping had to be greater than the squared inverse of the dominant forward-looking root of the model in order to guarantee that the solution was well defined; this analytic result was borne out by our experience with our empirical application. Figures 3a–d illustrate the control outcomes for a discount rate of 1·25 on instrument damping (compared to 0·99 in the previous run). Clearly, the policy instrument paths now look much more plausible and the improvement in inflation which occurs is not dominated by the effect of the policy settings in the tail.

Yet the control solution is somewhat unsatisfactory since it is critically dependent on how much instruments are damped and leaves a number of crucial questions unanswered. Would the solution be very different if the optimisation horizon were lengthened? Would it converge on the infinite time solution? Since the non-linear model becomes very computationally expensive to solve over longer time horizons, it would be preferable to work with the linearised version of the model. A rough indication of the degree of linearity of the model is given in figures 4a–d which compare the control outcomes of the non-linear optimisation problem with that obtained if it is assumed that the local derivatives (used in the non-linear optimisation algorithm) represent the true model. This quasi-linearised version

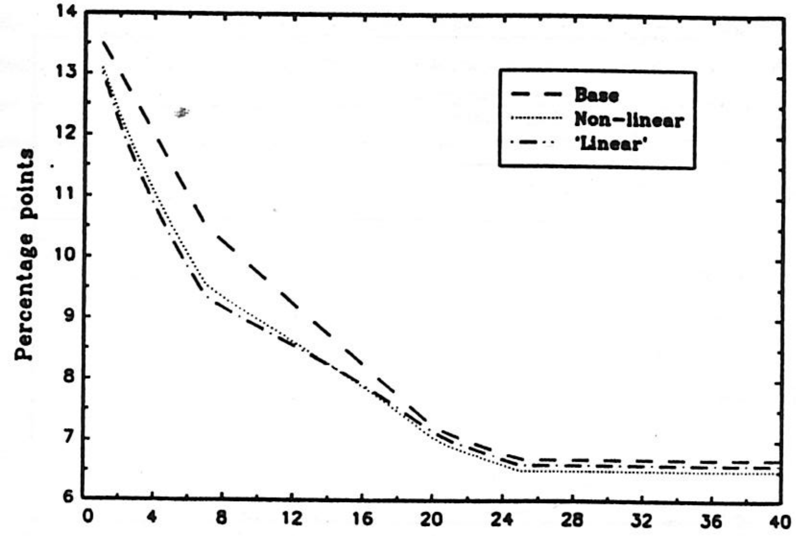

FIG. 4a. Treasury bill rate.

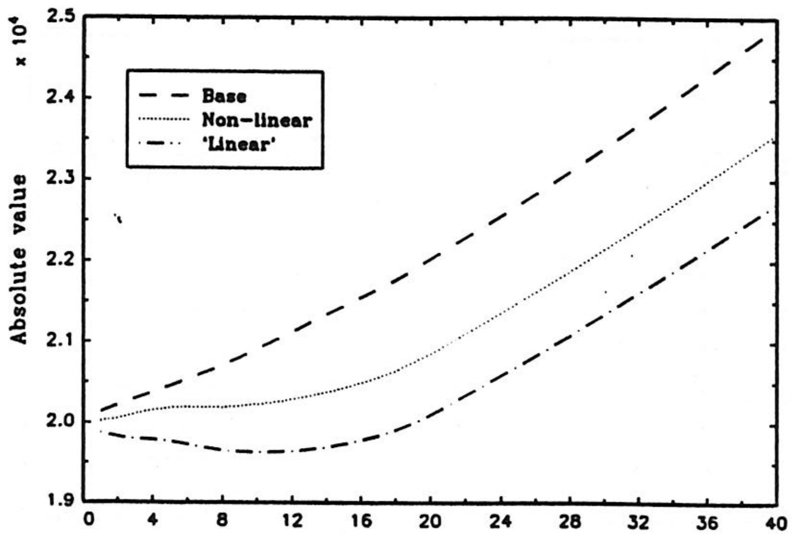

FIG. 4b. Public authority consumption.

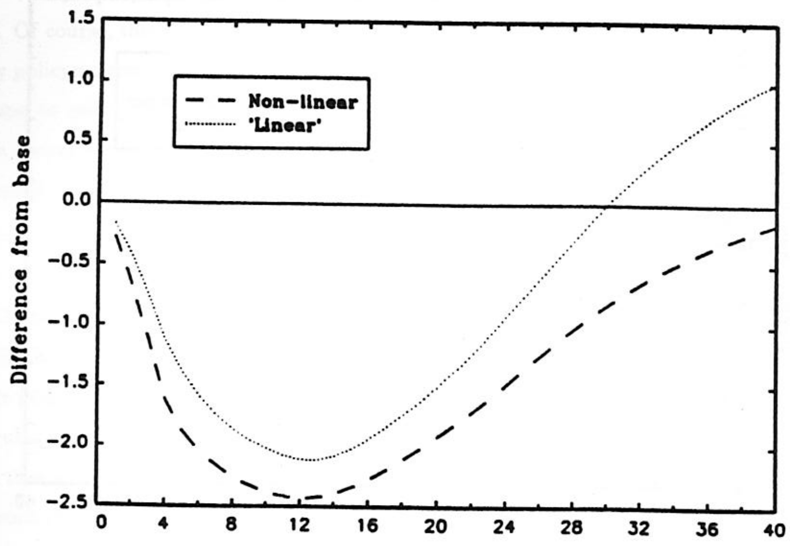

FIG. 4c. Annual price inflation.

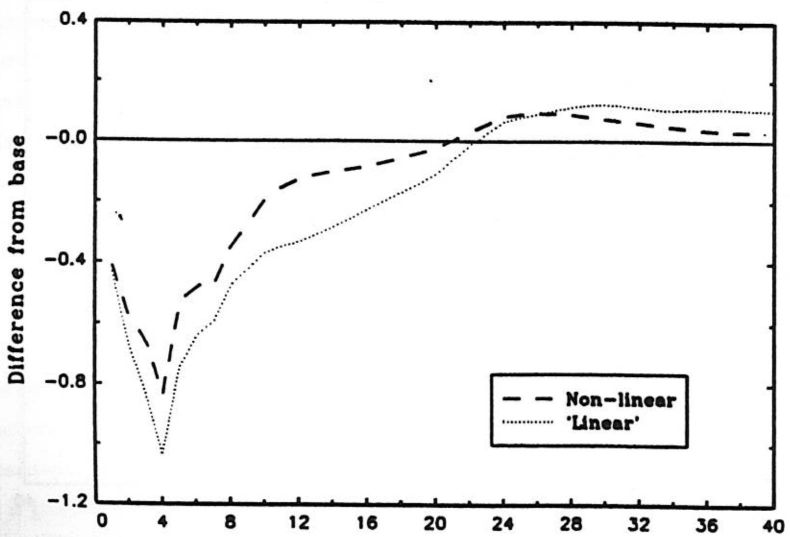

FIG. 4d. Annual GDP growth.

of the model shows that the responses of the targets to the instruments is approximately linear. Of course this type of approximate linearisation does not solve the problem of the finite horizon. For this the more sophisticated linearisation technique already described is needed. The results of adopting this linear approach are described in the next section.

5 Control using the linearised model

As described above, using the linear model offers a way of removing from the analysis of optimal policy the problems caused by the terminal condition interacting with the optimisation. We simply solve for expectations using the saddlepath property. The long run responses of the model embodied in the linear model are, of course, somewhat arbitrary but can be carefully imposed to be consistent with the expected long-run behaviour of the macroeconomic model.

In this section we concentrate on optimal time inconsistent control. There are two methods we use to calculate such policies. Firstly, the usual infinite horizon linear-quadratic solution. This solution is analytic, and described in, for example, Levine (1988).[9] Secondly, we calculated finite horizon optimal policies using directly analogous methods to the control exercises carried out on the non-linear model. Expectations are calculated using the Blanchard-Kahn method and the optimal policy calculated using the same method.[10] The linear problem solves in one iteration with exact derivatives. As before, instruments were allowed to vary over a fixed period, say 20 quarters, and then are set flat for a tail of 20 quarters with the cost function set to be minimised over the total 40 periods. The instruments, however, are set at these final levels permanently, as is implicit in the non-linear model. With the linear model we are able to extend the optimisation horizon to 200 quarters easily, and experiment with tails of arbitrary length.

In the notation of section 4 the cost function used with the linear model is

$$C = \sum_{t=0}^{T} 0{\cdot}99^t[(\dot{p} - \dot{p}^*)_t + g_t^2 + \Delta r_t^2 + \Delta G_t^2]$$

[9] The control problem does require numerical solution, however, as the solution to a matrix Ricatti equation must be found. All the results reported in this section and the next were calculated using a series of Gauss programs.

[10] The derivatives for the control problem were calculated using the appropriate multivariate generalisations of (11) and (12).

where variables are all as before and the weights have all been normalised to unity. A slight difference for the linear model is the inclusion of an explicit target variable for inflation, $\dot{p}^*$, with its own dynamic equation. T, the optimisation horizon, is varied between 40 (the shortest horizon we consider) and infinity. The target inflation rate is 1% below base and the target for growth is the base value.

The first example is where we fix the period of active control and vary the length of the tail. The results are illustrated in figures 5a and 5b. The period of active control is 40 quarters in all cases, but the tail is variously 10, 20 and 30 quarters. The departure from the previous section is that we can also show the infinite horizon solution.

The effects on inflation over the first 40 quarters are almost identical, with the shortest tailed run showing a marked upward movement towards the end. However, the instrument trajectories, illustrated in figure 5b, are very different. One point worth noting is that in general the qualitative finite horizon results are very similar to those that we obtained with the non-linear model (compare figure 5b and 3a). Just as in the non-linear model the impressive inflation control is achieved by promising increases in interest rates at the end of the optimisation (figure 5b) with the consequent impact on the profile of the exchange rate. This effect is nothing to do with terminal conditions, and everything to do with the finite horizon of the optimal control problem. The infinite horizon control shows a smooth interest rate path, which is clearly trending upwards towards the end.

Two points are to be noted. Firstly, a 30 period tail (for this control exercise) seems long enough for the tail length to be having no effect on the optimisation, a result we have confirmed by extending the tail further to almost no effect. Note that this is with the model solved using Blanchard-Kahn. Using terminal conditions, this might easily not be enough. A second, important, conclusion is that lengthening the tail does not bring us any closer necessarily to the infinite horizon solution. In fact, the model with a 10 period tail is 'closest' to the infinite time solution over first 40 periods, the length of the original control problem.

This is the type of effect we would hope to eliminate by extending the period of active control. Figure 6 shows the interest rate paths for 20, 40, 60 and 80 periods of active control, each with a 20 quarter tail, and the infinite horizon solution. After 40 quarters the

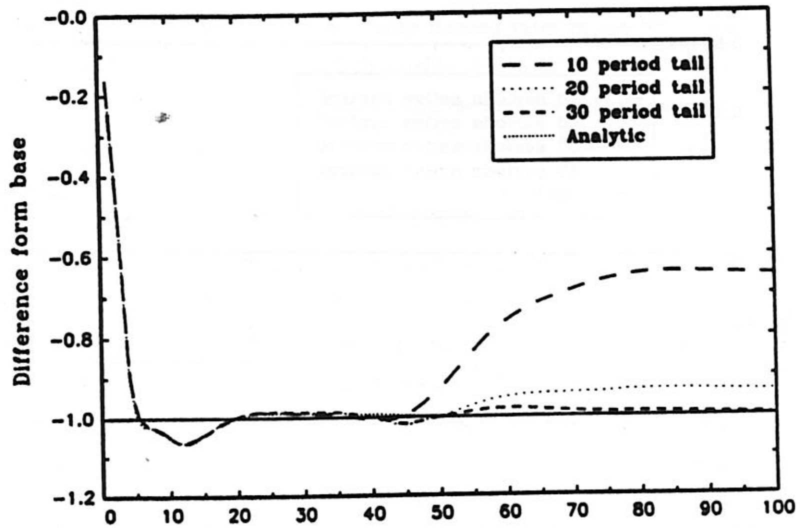

FIG. 5a. Inflation in the linear model with different tails.

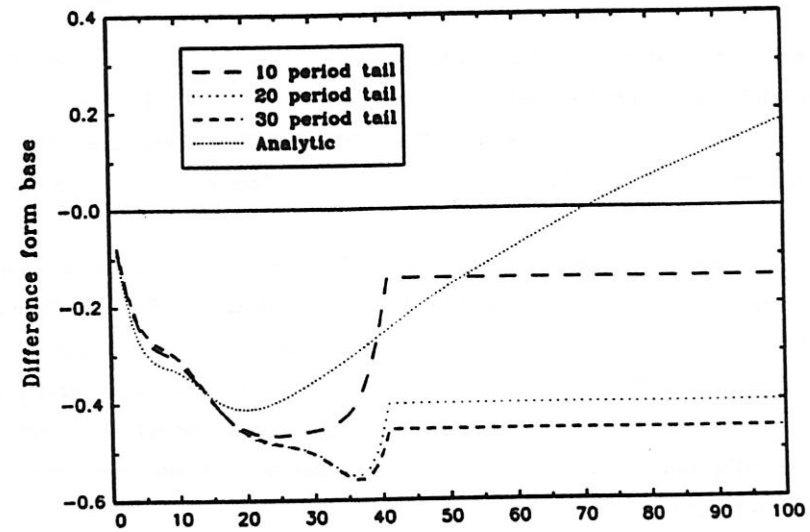

FIG. 5b. Treasury bill rate in the linear model with different tails.

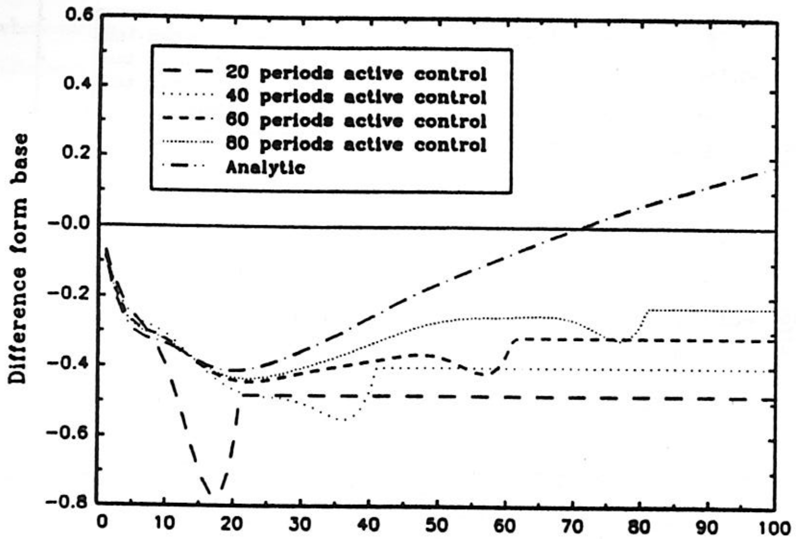

FIG. 6. Treasury bill rate in the linear with varying active control.

inflation rate (not illustrated) settles down to an indiscernibly different pattern. Similarly, growth is insufficiently different to reproduce. From the graphs it is clear that the finite horizon control is converging to the infinite but very slowly. Even after 160 periods of active control only the first 25 periods instrument values are the same as the infinite horizon solution.

Two further remarks can be made here. The effects illustrated are rather different if we reduce the discount factor slightly. We were unable to solve for the infinite horizon solution with a discount factor of 0·98. The finite horizon problem remains well defined and further simulations (not reported) seem to confirm that the forward-looking behaviour of agents is being exploited by the optimal policy as future interest rate movements get larger rather than smaller. The results are not as clear cut as the simple example in section 3 might suggest.

With further simulations, we were also able to confirm with the linear model that with an increasing weight on the instrument changes, at a rate of $1{\cdot}2^t$, the qualitative control results are unaffected by the finite horizon. Inflation gains are possible in the short run,

but without the ability to costlessly change interest rates in the future the spectacular control over long periods that were obtainable disappear. But this form of control is suboptimal, even if more credible. We can analyse time consistent (and therefore credible by construction) straightforwardly using the linear model. This we do in the next section.

6 Time consistent control using the linearised model

There have been several time consistent equilibria suggested for linear rational expectations models, with the best known the dynamic programming method suggested by Oudiz and Sachs (1985). This is technically a feedback Stackelberg equilibrium to a dynamic game, and an alternative feedback Nash equilibrium can be calculated in a similar way.[11] Such feedback equilibria are difficult (if not impossible) to calculate for non-linear models (see de Zeeuw and van der Ploeg, 1991) but can be calculated for the linearised model. An appealing feature of both feedback equilibria is that they are subgame (Markov) perfect, and as such are a complete description of the equilibrium strategies both on and off the equilibrium path.

Table 1

Welfare losses for the infinite horizon control problem using the linearised model

Control regime	*Welfare loss*
Time inconsistent	0·805
Feedback Stackelberg	0·839
Feedback Nash	0·849

The infinite horizon optimal time consistent policies are readily calculated and the welfare losses associated with both the time consistent and time inconsistent control regimes are shown in table 1. It is apparent that for the control problem we consider the restriction to time consistent policies does not seriously affect the welfare losses, with at most a 5%

[11] Blake (1992) gives iterative schemes to calculate these two equilibria for quite general linear rational expectations models.

reduction in welfare. The actual trajectories of the instruments and targets in all three regimes are so similar that we do not reproduce them here. However, we can conclude from this that the nature of the control problem is that it is not inherently time inconsistent. The rather incredible instrument movements required by the finite time optimisation make the optimal policy appear more time inconsistent than is actually the case.

Unfortunately, we are unable to calculate finite horizon feedback equilibria as the dynamic programming solutions only have a proper interpretation in the steady state. To evaluate the importance of finite horizons for time consistent control we must seek an alternative solution. We use the method proposed by Westaway (1989). This is an open loop time consistent solution and takes the form of a simple modification to the time inconsistent solution method (see Westaway, 1989, for details).

Figures 7a and 7b show the optimal trajectories for government spending for the time inconsistent and time consistent optimal policies, where the tail is always 20 quarters and the control horizon is from 100 to 180 periods for the consistent solution and 100 to 160 and the infinite horizon solution when the policy is time inconsistent. The control of the target variables is very similar in both cases. It is very noticeable that the consistent policy converges after 100 quarters, confirmed by the behaviour of interest rates (figure 8), in marked contrast to the time inconsistent path (figure 6). Both instruments settle down to near their long run values in all illustrated cases, a feature absent from the optimal inconsistent exercises. It is perhaps surprising that such similar control of targets is achievable with very different instrument trajectories.

These exercises show that a major advantage of linearising the model is that we can calculate a variety of time consistent equilibria to help us analyse the source of time inconsistency. As expected, we have been able to confirm that open loop time consistent equilibria seem to suffer much less from the problems associated with the finite horizon, although the horizon does need to be rather greater than usually available for non-linear control exercises.

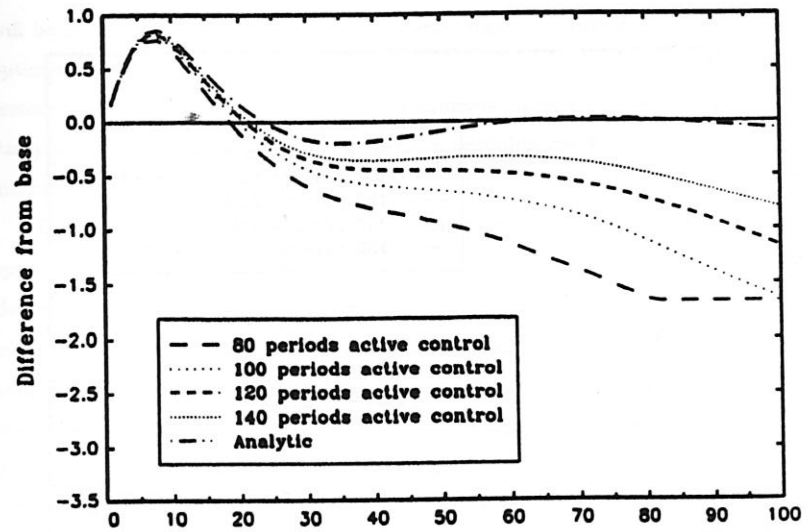

FIG. 7a. Government spending for time inconsistent control.

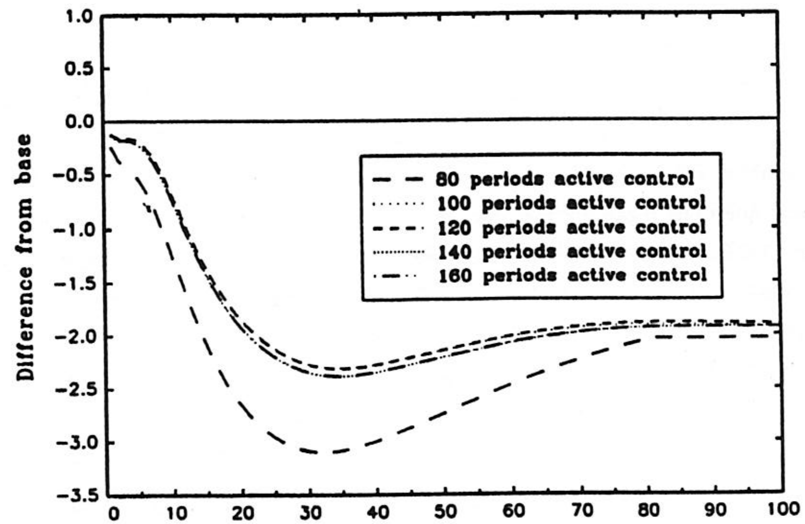

FIG. 7b. Government spending for time consistent control.

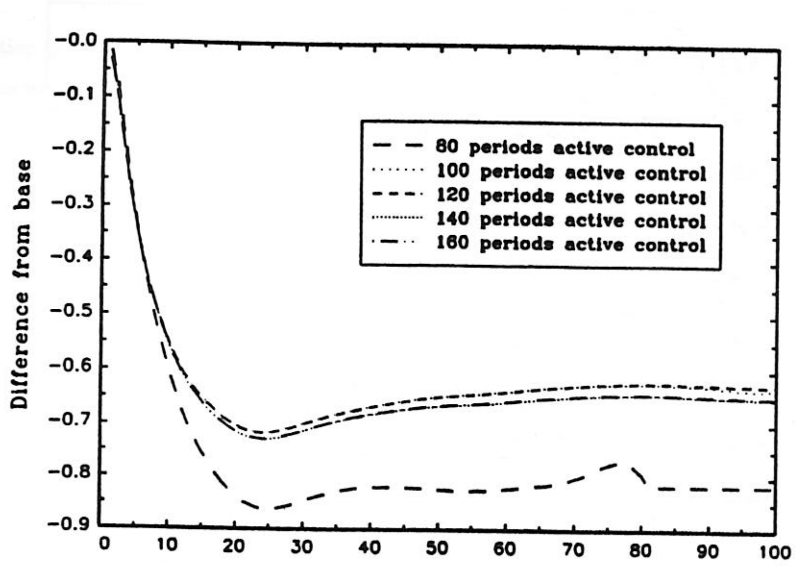

FIG. 8. Interest rates for time consistent control.

7 Conclusions

Many of the lessons to be drawn from this paper are as much salutary as definitive. Starting from the firmly held viewpoint that optimal control techniques ought to play a central role in guiding policymakers in their conduct of macroeconomic policy, we have argued that many questions regarding the practical use of these techniques are not only unresolved but also largely unconsidered. We have shown that attempts to use empirically based macroeconomic models which embody forward-looking behaviour present policy-makers with problems which do not arise in applications based on small analytic models. Central to these problems is the role of the finite horizon which must necessarily be assumed in work based on non-linear models. We have shown that this aspect of the problem distorts the implied policy solution in two ways, first because of the effect of the terminal condition on the solution (and the associated need to impose a tail on policy instruments), second because of the effect of the finite horizon on the control strategy itself. We have identified a class of problems, in particular when the government is less forward-looking than the private sector, where the finite horizon control solution can be calculated but where it

will be of dubious validity since it will not converge to an infinite horizon solution as the optimisation horizon is extended. We have shown how the convergence properties can be restored by assuming that the weight on instrument damping increases through time; this damping may be interpreted as capturing the declining credibility associated with future instrument movements.

Of course, whether or not this type of instrument damping is required to define the optimal control solution or not, the resulting solution will anyway be time inconsistent. However, consideration of credibility is central to the question of how empirically based models should be used in the design of macroeconomic policy. If, as is conventionally assumed, the time inconsistent policy is derived as the benchmark optimum that a government should strive for if it could build up sufficient reputation, then our finding here that this solution will not always be defined is an important one. On the other hand, if it is assumed that time consistent strategies are likely to be the only policies that are likely to remain sustainable, then we are still left with the problem that any finite horizon strategies derived on this basis are also likely to be distorted, although to a much lesser extent than with the time inconsistent cases.

Consequently, the results of this paper seem to point to the inevitable conclusion that meaningful policy analysis is best carried out on linearised representations of the original non-linear models so that the infinite horizon analytic results can be exploited. If this is true then it also follows that a much more rigorous scrutiny of the linearisation techniques that have been developed is required. This scrutiny would involve more thorough examination of the robustness of the techniques in capturing the behaviour of the underlying model as well as a more systematic examination of any important non-linearities in the model. Of course, once these analytic solutions are available for use on empirically based models, this facilitates a more rigorous analysis of the determination of the credibility of policymakers and the precommitment period associated with different policy regimes. These topics are the subject of our current research agenda.

APPENDIX

A.1 Modifications to the model

The model used in this exercise is a variant of the National Institute Domestic Econometric Model 11.6 (NIDEM 11.6). It was modified so that all except one of the expectational variables are solved in 'backwards mode'; for these variables expectations are generated by time series equations rather than set equal to model predictions. The exception to this is the exchange rate, where the usual model equation was retained. The non-linear control exercises were carried out on the model with no further changes.

This modification was made to simplify the linearisations. We describe why below. However, the modifications had little effect on the simulation properties of the model. Figures A1 and A2 show two key responses of NIDEM 11.6 in modified and unmodified form to a temporary public authority consumption shock, our fiscal instrument. These are the effective exchange rate and growth. Model properties are not seriously distorted by the changes, although growth is rather damped. The effect on inflation (not illustrated) was almost identical for the two different models.

A.2 Linearising the model

To linearise the model we used the procedure described by Maciejowski and Vines (1984), based on Kung (1978), modified to reproduce the dynamic multipliers from step changes to inputs. The output responses to chosen input variables are realised into a state-space representation by performing a singular value decomposition of the system Hankel matrix. For a detailed description of the method see Weale *et al.* (1989). The resulting linear model is of the form

$$z_{t+1} = Az_t + Bu_t \quad \text{(A.1a)}$$

$$y_t = Cz_t + Du_t \quad \text{(A.1b)}$$

where z_t is a vector of predetermined state variables, y_t the outputs and u_t the inputs. The predetermined states are analogous to principal components of the original model and have no direct economic interpretation (Moore, 1981). A major advantage of this method

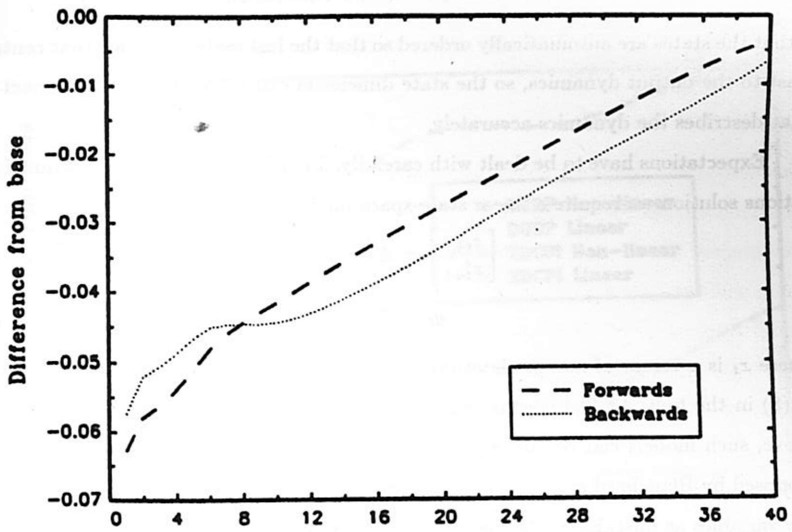

FIG. A1. Exchange rate response to a 1% one quarter increase in government spending.

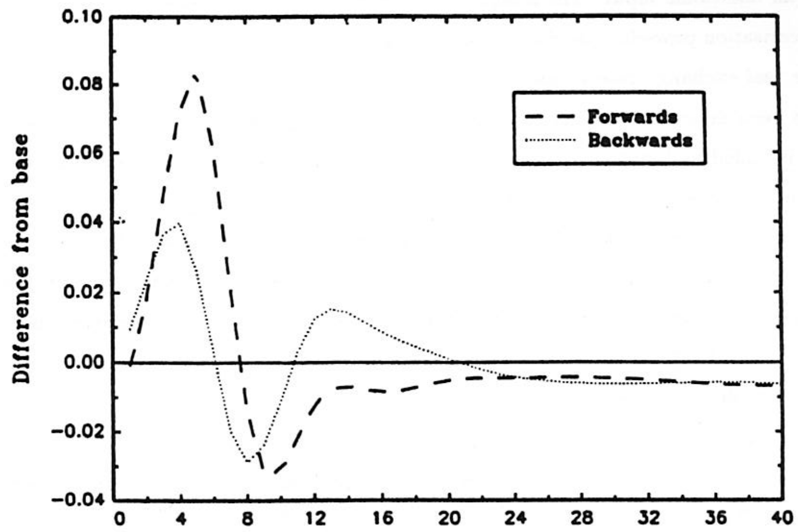

FIG. A2. Inflation response to a 1% one quarter increase in government spending.

is that the states are automatically ordered so that the last state is the one that contributes least to the output dynamics, so the state dimension can be reduced to a compact model that describes the dynamics accurately.

Expectations have to be dealt with carefully. In order to calculate the rational expectations solution we require a linear state-space model in the following form:

$$\begin{bmatrix} z_{t+1} \\ x^e_{t+1} \end{bmatrix} = A \begin{bmatrix} z_t \\ x_t \end{bmatrix} + Bu_t \quad \text{(A.2a)}$$

$$y_t = C \begin{bmatrix} z_t \\ x_t \end{bmatrix} + Du_t \quad \text{(A.2b)}$$

where x_t is a vector of non-predetermined variables. (A.2a) is the multivariate extension to (8) in the text. (A.2b) is sometimes known as an observation equation. As described above, such models can be solved for a rational expectations solution using the method proposed by Blanchard and Kahn (1980). If the model is saddlepath stable with as many free variables as unstable roots the solution is unique.

The linearisation was done with an exogenous ***real*** exchange rate, which was treated as an additional input. The model in this form has no forward-looking variables and the linearisation procedure used to reproduce the input-output behaviour of the model. With the real exchange rate as an input, we can include its equation ***after*** the linear model has been determined. This additionally means all the right hand side variables must to be included as inputs or outputs. As there are two targets and two instruments, three input variables and four outputs are required. The inputs were the Treasury bill rate, government spending and the real exchange rate (ρ). The outputs were inflation, expected one quarter inflation ($\dot{p}^e$), growth and the net overseas asset ratio (nar). All of the outputs are measured as deviations in levels from their base values and simulations run over ten years from 91Q1 to 00Q4. Simulation were also run over ten years from 93Q1 to check for base dependencies and non-linearities, with very similar results. All the responses were asymptotically stable, but in some cases extrapolated to a steady state. 59 states were found to be adequate to capture the dynamics of the responses exactly, with (A1) specifically

$$z_{t+1} = Az_t + B \begin{bmatrix} r_t \\ G_t \\ \rho_t \end{bmatrix} \quad \text{(A.3a)}$$

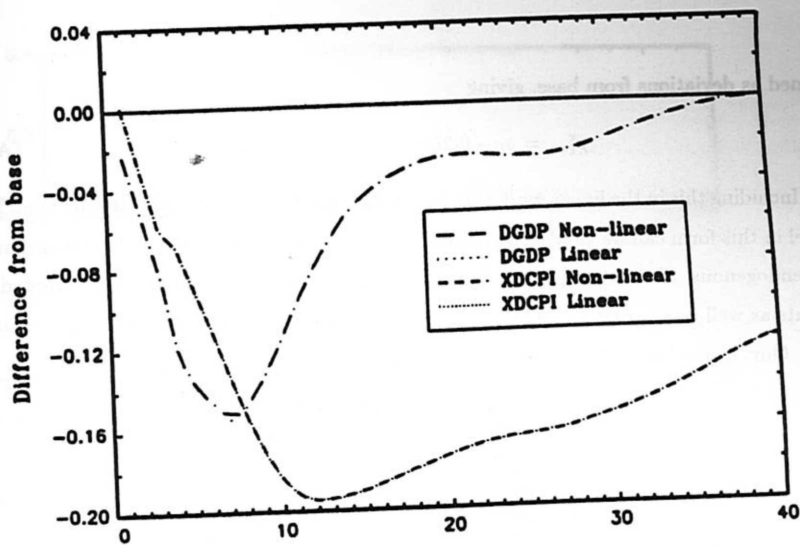

FIG. A3. Linear and non-linear model responses to 0·5% permanent increase in interest rates.

$$\begin{bmatrix} \dot{p}_t \\ g_t \\ nar_t \\ \dot{p}^e_t \end{bmatrix} = Cz_t + D \begin{bmatrix} r_t \\ G_t \\ \rho_t \end{bmatrix} \quad \text{(A.3b)}$$

where A, for example, is a 59×59 matrix.

Figure A3 shows two responses of the linear and the non-linear model measured as deviations from base for 40 quarters for a permanent 0·5% increase in interest rates. This level of accuracy is achieved for all responses. Linearisation error (with a fixed real exchange rate) is effectively zero.

A.3 Validating the linearisation

The linearisation procedure is very accurate in reproducing the input-output characteristics of the model, but does it mimic the rational expectations solution of the non-linear model? With the given A, B, C and D matrices the exchange rate can now be made to behave as a forward-looking variable. The exchange rate equation used on the non-linear model can be translated into a real exchange rate equation, taking account of the variables being defined as deviations from base, giving

$$\rho^e_{t+1} = \rho_t - 0{\cdot}25(r_t - \dot{p}^e_t) - 1{\cdot}12nar_t. \qquad (A.4)$$

Including this in the linear model obviously increased the number of states by one. The model in this form can be validated against the non-linear model solved with the exchange rate endogenous. For control purposes the changes in the instruments were included as outputs as well as a target value for inflation, requiring three more state variables, a total of 63. Our 'control model', the form (A2), is then:

$$\begin{bmatrix} z_{t+1} \\ \rho^e_{t+1} \end{bmatrix} = \hat{A}\begin{bmatrix} z_t \\ \rho_t \end{bmatrix} + \hat{B}\begin{bmatrix} r_t \\ G_t \end{bmatrix} \qquad (A.5a)$$

$$\begin{bmatrix} (\dot{p} - \dot{p}^*)_t \\ g_t \\ \Delta r_t \\ \Delta G_t \end{bmatrix} = \hat{C}\begin{bmatrix} z_t \\ \rho_t \end{bmatrix} + \hat{D}\begin{bmatrix} r_t \\ G_t \end{bmatrix} \qquad (A.5b)$$

where $\dot{p}^*_t$ is the target value of inflation with behavioral equation $\dot{p}^*_{t+1} = \dot{p}^*_t$ and $\Delta u_t = u_t - u_{t-1}$. With the model set up in this way the instruments are set to always drive the outputs to zero. The model is saddlepath stable open loop with one eigenvalue of 1·0103.

Validating the linear model with the exchange rate endogenised is the next step. As sections 3 and 4 above show, the rational solution is affected by the terminal condition used to solve the non-linear model. Therefore the linear model is solved using the same terminal condition as the non-linear model for comparison. This is that:

$$nar_T = 0{\cdot}25(nar_{T-1} + nar_{T-2} + nar_{T-3} + nar_{T-4}). \qquad (A.6)$$

Although we could use the method suggested by Blake (1990) it is simpler to solve the model on an unstable trajectory for an arbitrary initial value, using a simple golden search routine to locate the starting value for the exchange rate that satisfies the terminal condition.

The results of this exercise are illustrated in figures A4 and A5. Although the solutions are not identical in the way that the open loop (exogenous real exchange rate) simulations are, the linear model is again remarkably similar in simulation response to the non-linear model. We are confident that the linear model is an accurate representation of the non-linear model.

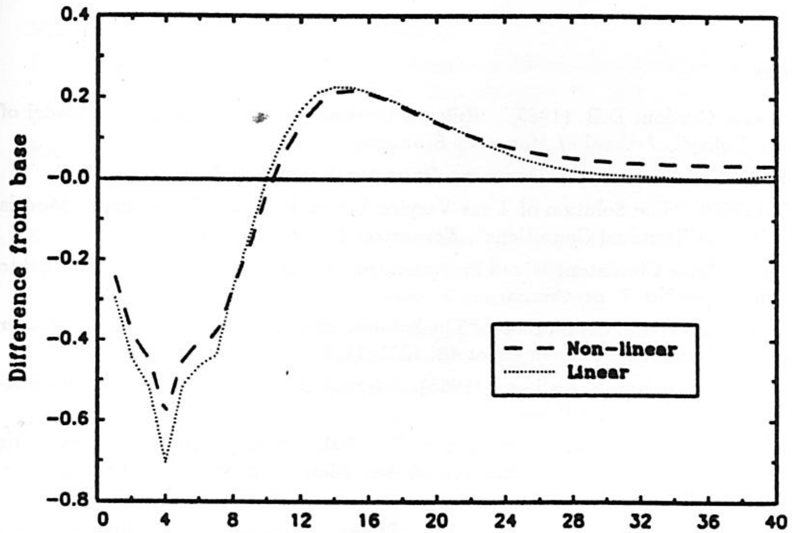

FIG. A4. Growth response to 0·5% permanent increase in interest rates.

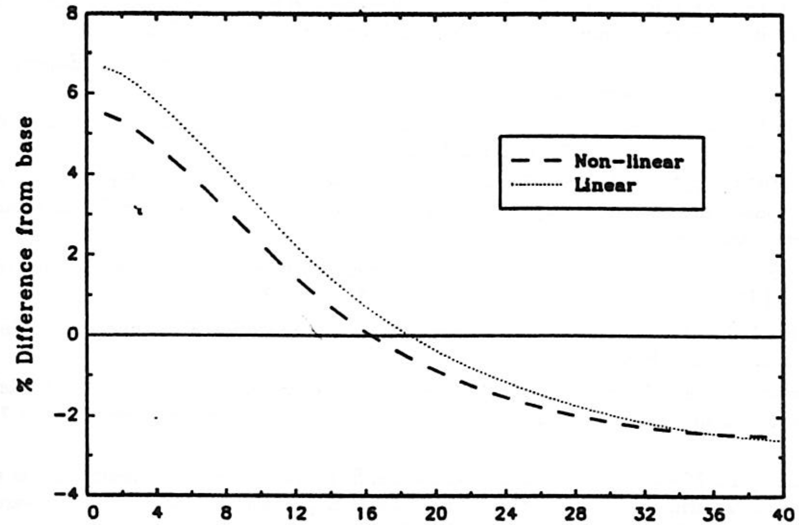

FIG. A5. Real exchange rate response to 0·5% permanent increase in interest rates.

References

Barro, R.J. and Gordon, D.B. (1983). "Rules, Discretion and Reputation in a Model of Monetary Policy", *Journal of Monetary Economics* **12**, 101–121.

Bellman, R. (1957). *Dynamic Programming*, Princeton University Press.

Blake, A.P. (1990). "The Solution of Time-Varying Linear Rational Expectations Models and the Role of Terminal Conditions", *Economics Letters* **33**, 265–269.

——— . (1992). "Time Consistent Mixed Precommitment Macropolicy", National Institute Discussion Paper No. 7, new series.

Blanchard, O.J. and Kahn, C.M. (1980). "The Solution of Linear Difference Models Under Rational Expectations", *Econometrica* **48**, 1305–1311.

Buiter, W.H. and Marston, R. (editors) (1985). *International Economic Policy Coordination*, Cambridge University Press.

Christodoulakis, N.C., Gaines, J. and Levine, P. (1991). "Macroeconomic Policy using Large Econometric Rational Expectations Models: Methodology and Application", *Oxford Economic Papers* **43**, 25–58.

Fisher, P.G. and Hughes Hallet, A.J. (1988). "Efficient Solution Techniques for Linear and Non-Linear Rational Expectations Models", *Journal of Economic Dynamics and Control* **12**, 635–657.

Flåm, S.D. and Wets, R.J-B. (1987). "Existence Results and Finite Horizon Approximates for Infinite Horizon Optimization Problems", *Econometrica* **55**, 1187–1209.

Holly, S. and Zarrop, M.B. (1983). "On Optimality and Time Consistency when Expectations are Rational", *European Economic Review* **20**, 23–40.

Ireland, J.M. and Wren-Lewis, S. (1989). "Terminal Dates and the Dynamic Properties of National Institute Model 11", *National Institute Economic Review* **128**, 91–93.

Kung, S. (1978). "A New Identification and Model Reduction Algorithm via Singular Value Decomposition", *Proceedings of the 12th Annual Asilomar Conference on Circuits and Systems*, Asilomar.

Kydland, F.E. and Prescott, E.C. (1977). "Rules versus Discretion: the Inconsistency of Optimal Plans", *Journal of Political Economy* **85**, 473–493.

Levine, P. (1988). "Does Time Inconsistency Matter?", Centre for Economic Policy Research, Discussion Paper No. 227.

Maciejowski, J.M. and Vines, D.A. (1984). "Decoupled Control of a Macroeconomic Model Using Frequency-Domain Methods", *Journal of Economic Dynamics and Control* **7**, 55–77.

Minford, A.P.L., Matthews, K.G.P. and Marwaha, S. (1979). "Terminal Conditions as a Means of Ensuring Unique Solutions for Rational Expectations Models with Forward Expectations", *Economics Letters* **4**, 117–120.

Moore, B.C. (1981). "Principal Component Analysis in Linear Systems: Controllability, Observability and Model Reduction", *IEEE Transactions on Automatic Control* **AC–26**, 17–32.

NIESR (1991). *National Institute Model 11.*

Oudiz, G. and Sachs, J. (1985). "International Policy Coordination in dynamic macroeconomic models", in Buiter, W.H. and Marston, R. (eds.), *International Economic Policy Coordination*, Cambridge University Press.

Rüstem, B. and Zarrop, M.B. (1979). "A Newton-type Method for the Optimisation and Control of Non-Linear Econometric Models", *Journal of Economic Dynamics and Control* **1**, 283–300.

Selten, R. (1978). "The Chain-store Paradox", *Theory and Decision* **9**, 280–312.

Wallis, K.E. (ed.), Andrews, M.J., Bell, D.N.F., Fisher, P.G. and Whitley, J.D. (1985). *Models of the UK Economy 2*, Oxford University Press.

Weale, M.R., Blake, A.P., Christodoulakis, N.C., Meade, J.E. and Vines, D.A. (1989). *Macroeconomic Policy: Inflation, Wealth and the Exchange Rate.* Unwin Hyman.

Westaway, P. (1989). "Partial Credibility: A Solution Technique for Econometric Models", National Institute Discussion Paper No. 157.

——— and Wren-Lewis, S. (1990). "Forecasting Government Policy: An Example of the Importance of Time Inconsistency", *International Journal of Forecasting* **6**, 401–405.

——— and ——— . (1992). "Is There a Case for the MTFS?", forthcoming in *The Manchester School*.

Zeeuw, A.J. de and van der Ploeg, F. (1991). "Difference Games and Policy Evaluation: A Conceptual Framework", *Oxford Economic Papers* **43**, 612-636.