CIM

Centre for International Macroeconomics

Discussion Paper No. 2000.01

MODELLING ECONOMIC POLICY RESPONSES WITH AN APPLICATION TO THE G3

BY

Michael Chui
(Bank of England)

Guglielmo Maria Caporale
(South Bank University

Stephen G. Hall
(Imperial College)

and

S.G.B. Henry
(Centre for International Macroeconomics, Oxford University)

May 2000

MODELLING ECONOMIC POLICY RESPONSES WITH AN APPLICATION TO THE G3

BY

Michael Chui
(Bank of England)

Guglielmo Maria Caporale
(South Bank University

Stephen G. Hall
(Imperial College)

and

S.G.B. Henry
(Centre for International Macroeconomics, Oxford University)

JEL Classification: E21, E22, E23

Keywords: Policy, Optimal Control, Feedback

Financial support from ESRC grant No. L116251003, "Macromodelling and Policy Analysis in a Changing World" is gratefully acknowledged. The usual disclaimer applies.

CONTENTS

1 INTRODUCTION

Within the framework of the large macroeconometric models there are two overriding reasons for modelling the process of policy response to economic events; The first is the obvious importance of this topic to the policy formation process itself. The second is the need to provide a basic model closure so that the model is a sensible forecasting tool.

The use of models in policy formulation has been one of the prime uses of these models since they were first developed. This can vary from the traditional technique of simply specifying exogenous fiscal and monetary policy and then using the model to forecast conditional on these assumptions to a complex range of optimal control techniques or endogenous feedback rules. The underlying objective is always the same; to help to specify a 'better' set of economic policies through the use of a fully specified model of the economy. Even in the most basic approach of exogenous economic policy variables, we should still think of policy as being an endogenous response, although it is not formally modelled as such. The model is used to investigate the effects of a particular, given, set of policies. These policies may be respecified many times by the model user until a satisfactory result for the economy emerges. In more complex modelling exercises this process is made formal by actually specifying the mechanism the policy maker uses to respond to economic events.

The model closure aspect of policy formulation is primarily a technical issue of interest to modellers. It is simply that the policy formulation process is so important to the economic properties of a model that it has become increasingly obvious that it is almost meaningless to consider a models properties in isolation from the assumption made about policy formulation. As models have become more sophisticated the `old' exogenous economic policy assumption has become increasingly untenable. For example a model which includes rational expectations will often simply not solve under the assumption of fixed policy settings, and indeed we know that often it should not. The particular specification of the policy response also may have enormous effects on a models response to any given shock and so it has become increasingly obvious that any model comparison exercise must be done on the basis of similar economic response assumptions (e.g. Bryant et al 1988, 1993) otherwise the exercise simply reduces to the finding that different policies give rise to different effects and nothing can be said about the actual differences between the models.

There are a range of formal approaches to the modelling of policy responses. This chapter will survey the approaches which have been used in the literature and explain the motivation and workings of the new algorithm which will be used later in this book. The modelling of policy response may be characterised in a number of ways depending on the issue being highlighted. It may be done in terms of specifying an objective function and a set of policy instruments or in terms of specifying an equation (feedback rule) which links the policy variables directly to the models outcomes. If we are considering the latter case we can consider equations which derive their parameters in a formal optimal way or equations which have some other method of choosing the parameters. We might consider only the deterministic model or we might focus on the stochastic model and the consider whether we are concerned with just executed outcomes or if we are also concerned with minimising the variance of the outcome. Finally we might consider the simple case of a single policy maker or we might set the problem within a game framework where two, or more, players compete in a policy game on the basis of a range of bargaining arrangements such as Nash stackelberg or co-operative solutions.

Our aim is to argue that in recent years there has been a general move towards the use of explicit feedback rules which are either optimal (in the sense that their parameters are chosen to minimse some criteria) or are designed to minimise the variance of the outcome of the economy but which do not do this optimally. This has come about mainly because of the numerical complexity of choosing optimal parameters with respect to a stochastic model. We will conclude by outlining a simple technique which allows the use of optimal parameter selection on a stochastic model and thus opens up the possibility of full optimal feedback rules for policy analysis even in a multi player game setting.

2 OPTIMAL CONTROL AND POLICY RULES

The basic underpinning of all policy work is the optimal control framework where we specify an objective function and an economic system and choose the optimal setting for the policy instruments in the light of the these two elements. Even the specification of exogenous policy variables may be interpreted in this way, with the optimisation being done informally `off model'. Optimal control has been carried out on large models for a considerable time, early examples are Chow(1975) and Bray(????), a survey of techniques used may be found in Fair(1984) or Hall and Henry(1988). The basic approach is to think of a macromodel in its most general form as a mapping from the known information set $X_t=x_0....x_T$, $Y_t=y_0....y_{t-1}$ onto the future endogenous variables Y_{t+i}, i=0...T, an expression for a general macromodel would be;

$$y\ (y_{t+i}\ , X, Y)\ =\ 0 \qquad \mathbf{1}$$

the solution to which can also be written as a function of current information and splitting the exogenous variables into *n* policy variables U, and the other exogenous variables X*:

$$y_{t+i}\ =\ w(U, X^*, Y) \qquad i = 1.....T \qquad \mathbf{2}$$

A policy rule might then be chosen to minimise any given cost function, such as the standard Quadratic function given below:

$$\text{Min } C\ =\ \sum_{i=0}^{T}\ (Y_{t+i}\ -\ \overline{Y}_{t+i}\ )^2 \qquad \mathbf{3}$$

where the optimal policy rule will be that which satisfies:

$$\sum_{i=0}^{T} 2(Y - \overline{Y})\ w_{u_{jk}}.(U, X^*, Y)\ =\ 0 \qquad \begin{matrix} j = 1.....n \\ k = 1.....T \end{matrix} \qquad \mathbf{4}$$

where

$$w_{u_{jk}}(U, X^*, Y)\ =\ \frac{dw(X,\ U)}{d\ u_{jk}} \qquad \mathbf{5}$$

In the case of a large non-linear model it will not generally be possible to find an explicit analytical solution to (4) but there are numerical techniques which are now well established for solving (4). This is called an open loop control procedure as the process does not allow us to take account of a change in the initial conditions or the exogenous variables except by recalculating the complete solution to (4). The alternative to this is to specify the solution to (4) in the form of an equation, this is then called the closed loop form of the solution and can, in general be written as,

$$u_{jk} = f(X,Y) \qquad \mathbf{6}$$

If the function Φ can be written down then this forms a closed loop feedback, in that it would be possible to calculate the appropriate change in u_j for any change in x_t. This is only possible for linear models. For the non-linear case one has to solve the optimisation problem numerically and hence derive open loop trajectories for the policy instruments, ie. a given value for each u_j in each t. If the future values of X change then the optimisation problem has to be resolved and a new trajectory calculated.

It is clear from equation 6 that the fully optimal rule makes use of the entire state vector of the model including all future values. This rule is therefore likely to be quite complex and, as already noted, it is not generally possible to solve this equation explicitly.

There are also a number of conceptual reasons why people have argued against the use of such a complex solution rule;

Firstly the function Φ exploits the full information about the structure of the model. What tends to happen in practice on an empirical macromodel, is that the rule will exploit the dynamic characteristics of the system being controlled or else may find some facet of the model that does not truly reflect the real world. This might simply be some odd quirk in the model, an odd non-linearity, a corner solution or even an extreme assumption such as rational expectations which the optimal policy rule is able to exploit. Optimal policy rules therefore tend to be highly model specific, as is demonstrated by Bray et al (1995). The optimal policy description may simply be erroneous or else be dependant on the accuracy of the model specification, as was first identified in the Ball Report. In particular, it will be little help if there is uncertainty about either the underlying structure of the economy or at the very least the rate of dynamic adjustment in the economy. In general, policy conclusions stemming from specific elements of a model in which we have little faith should obviously be avoided.

The second problem, also identified by the Ball Report related to the Lucus critique. Lucus (1976) was the first to raise doubts about the usefulness of macroeconomic models for policy making when economic agents formed expectations which where forward looking. The mere announcement of a future change in policy could therefore alter agents behaviour. The problem is that having changed agents expectations and hence their behaviour, the incentive to carry through the announced policy might evaporate. Put formally, the presence of forward looking expectations in the model will mean that the derivative of the model solution to future policy changes will not in general be zero, ie.

$$\frac{dY_t}{du_{jt+i}} \neq 0 \qquad \mathbf{7}$$

This means that a policy for period $t+i$ will be optimal for current period t, in which it is derived but may no longer be optimal when the future period $t+i$ actually arrives. This is a time inconsistent policy.

As Kydland and Prescott (1977) point out, this raises serious problems regarding the credibility of the optimal policy in the eyes of the private sector because there may be an incentive to renege on any preannounced policy. They regard the time inconsistency property as a fundamental problem in the use of optimal control methods for macroeconomic policy design. In the concluding paragraphs of this seminal paper Kydland and Prescott argue that instead of attempting to select a policy optimally,

> *"it is preferable that selected rules be simple and easily understood, so that it is obvious when a policy maker deviates from the policy. There should be institutional arrangements which make it a difficult and time-consuming process to change the policy rules in all but emergency situations."*

To some extent some developments in the literature have attempted to circumnavigate this problem. Barro and Gordon (1983) examine whether reputational considerations can restore credibility for policy makers and hence avoid the inferior outcome of the time consistency constraint. They assume that policy makers suffer a loss of reputation if they renege on their earlier commitments. With this "punishment" mechanism in place, Barro and Gordon show that credible and sustainable policies superior to the time consistent policy can exist.

This seems to ignore the main message in Kydland and Prescott, in that *building* credibility will require some process of monitoring the authorities actions. This is going to be very difficult if the optimal control rule is very complex. Of course in reality, as we have discussed above, there is not going to be a (feedback) "rule", rather, since the economy is non-linear, there is going to be a stated open loop trajectory for all the policy instruments. This would have to be updated with each new piece of economic data. Monitoring the authorities actions would therefore consist of everyone having access to the same model and data sources that the authorities were using to calculate the optimal policy. While this might prove costly and would not be particularly easy to communicate in the media for example, it is not completely out of the question. Ultimately therefore, it again comes down to the pertinence of the model.

The alternative which many authors have used to the full optimal feedback rule because of these disadvantages is to design rules that exploit only the information which we believe to be useful and to de-emphasise the less reliable elements of the model's structure. Thus simple feedback rules are generally a restricted form of the full optimal control solution which limit the amount of information drawn from the structure of the model to those areas which are of special relevance to the policy question at hand (see for example, Vines et. al. 1983; Currie and Levine, 1985; Taylor, 1985 and Edison et. al. 1988). By implicitly excluding much of the model, simple rules are supposedly robust to uncertainty. Furthermore, if articulated publicly, they meet Kydland and Prescott's criteria of being simple and easy to interpret, and therefore useful when it comes to monitoring the authorities.

Simple feedback rules were developed in the engineering literature and later applied to economic systems. Phillips (1954 & 1957) discusses the relative merits of proportional, integral, and derivative control. Proportional control for example, is where a control variable is set in proportion to how far a target variable is from some desired value, ie.

$$u_{t\,t} = b_2\,(Y_t - \overline{Y}_t) \qquad \mathbf{8}$$

Blake and Westaway (1994) point out that a proportional rule with respect with inflation will leave the price level indeterminate and hence will not adequately close a macromodel. In the language of the 1950s, this is because "complete error correction is not obtained" with such a rule "since the correcting action continues only because the error exists" (Phillips, 1954, p.298).

Integral control, on the other hand is where the policy instrument is adjusted in some proportion to the *sum* of the targets differences from its base or desired values:

$$u_t = b_3 \sum_{t=0}^{n} (Y_t - \overline{Y_t}) \qquad \mathbf{9}$$

Taking first differences this is equivalent to:

$$\Delta u_t = b_3\,(Y_t - \overline{Y}_t) \qquad \mathbf{10}$$

Technically, an integral control rule may be sufficient to ensure model closure. But this is true of only a limited class of models.

The performance of the final rule can be further enhanced by the inclusion of derivative control, where the control variable is set according to the rate at which the target variables are accelerating. This gives

$$u_t = b_1\,(\Delta Y_t - \Delta\overline{Y_t}) \qquad \mathbf{11}$$

This may be thought of as an attempt to deal with the slugishness of feedback rules by trying to make the rule react as rapidly as possible to any change in circumstances. Combining equations (8), (10) and (11), gives us the full standard closed loop feedback rule that includes elements of proportional, integral and derivative control;

$$\Delta u_t = b_1 \Delta^2 (Y_t - \overline{Y_t}) + b_2 \Delta (Y_t - \overline{Y}_t) + b_3 (Y_t - \overline{Y}_t) \qquad \mathbf{12}$$

Such a rule still only reflects a very small part of the optimal closed loop feedback rule given above, in particular the rule is simplified in two very important respects, only a very small part of the information set is utilised and only past or current values of the model variables affect the settings of the policy variables. This second point is particularly important as it means that the rule will normally only react to events rather sluggishly and this can give rise to slow and unstable policy responses. When we think how policy is actually conducted it is obvious that policy makers go to great lengths to anticipate future events so that they may react in anticipation of future problems. It is however possible to go beyond the standard framework of (12) to incorporate this effect by adding a forward looking component to the feedback rule following Hall and Nixon(1996) in the following specification,

$$u_t = au_{t-1} + (1-a)u_{t+1} + b_1\Delta^2(Y_t - \overline{Y_t}) + b_2\Delta(Y_t - \overline{Y}_t) + b_3(Y_t - \overline{Y}_t) \qquad \mathbf{13}$$

When a is close to 0 the rule will react to events a long way into the future while if a is close to 1 it will only react as changes in Y take place.

3. OPTIMAL, ESTIMATED AND CALIBRATED SIMPLE RULES

Within the context of the simple rule (13) there is an important distinction based on how the parameters of the rule are chosen. When such rules were first introduced into economics by Phillips in 1954 the parameters were selected in a largely informal way, mainly by trial and error. Later work was undertaken to parameterise the simple rules in a more formal way and so optimal simple rules were developed. This is done quite simply be specifying an objective function such as (3) and then treating the parameters of the control rule (13) as the control variables which are to be determined so that (3) is minimised. If the model were linear and all state variable appeared in the control rule this would simply be a direct way of calculating the optimal feedback rule. Generally the model is not linear and only a very restricted set of feedback variables are considered and so the form of the rule represents an overall constraint on the optimisation. It is hoped that this constraint rules out unreasonable policies while the use of formal optimisation gives the choice of parameters a proper foundation in the structure of the model. Of course, it is not necessarily the case that a particular simple rule will behave well and part of the skill of a modeller is in knowing the form of simple rule which will work satisfactorily in any given model.

Directly estimated rules

In work of this sort, what might be termed descriptive methods are used to identify policy rules. The question thus posed is whether over the past the policy actions of the central bank (or the fiscal authority) can be represented by an estimated rule linking instruments to targets. Examples include Clarida et al (1998) applied to G6 and E3 (UK, France and Italy), and Walton and Massone (1999), for the UK. Clarida et al provides the most fully worked out version, and we use this as representative. Their monetary reaction function starts from the form used by Taylor linking interest rates to discrepancies between inflation and its target, and the output gap,

$$r_t^* = a + b\,[\Pi_{t+n}] + dEx_t \qquad (14)$$

where r is the nominal short rate, (r* its target), Π the rate of inflation and x the output gap (y-y*; where y is real output). Assuming interest rate smoothing and rational expectations gives the equation in a form suitable for estimation, i.e.

$$r_t = (1-r)a + (1-r)\,b\,\Pi_{t+n} + (1-r)d\;x_t + r r_{t-1} + e_t \tag{15}$$

where the interest rate smoothing is given by the equation

$$r_t = (1-r)\,r_t^* + r\;r_{t-1} + v_t \tag{16}$$

with v_t white noise, and the error term e_t in the estimated equation above is a combination of “rational” forecast and equation error,

$$e_t = -(1-r)[\,b\;(\Pi_{t-n} - \mathrm{E}(\Pi_{t+n}) + d\;(x_t - \mathrm{E}(x_t)] + v_t \tag{17}$$

Estimated by GMM for the G3, and for what the authors describe as the E3 (namely the UK, France and Italy) using monthly data for the sample 1979M4-1993M12, the equations generally fit well.

We argue that there are two fundamental problems with this approach. The first is the problem of the econometric identification of the reaction function, the second is the economic interpretation of this function once it is derived. We will discuss these two in turn.

The econometric problems of identifying a reaction function where none of the variables are weakly exogenous are profound. The authors themselves interpret their equations as characterising monetary policy over the period, showing that there was a concerted move towards inflation targeting, albeit what the authors call “soft-hearted” targeting, i.e inflation targets with some stabilisation element in policy too (Svensson (1998) refers to this as “flexible” inflation targeting). That is, the response to a rise in expected inflation is to push up nominal rates by a sufficient amount to increase real interest rates. In this case $b > 1$ and the authorities move the real rate to stabilise inflation and output. Where $b < 1$ then the authorities do not move nominal rates by enough to stop real rates from declining, so increases in both inflation and output are possible. The results for the G3 show that the baseline specification given by equation (15) above works best. In other words, the addition of additional variables like money aggregates or exchange rates does not add significantly to the explanatory power of the equation. Also adding lagged inflation does not significantly improve the equation, and the authors interpret this as confirming the forward looking specification used in the model. In sum, the estimated equations are advanced as a plausible description of how central banks have conducted policy. There is also the suggestion that the results may be interpreted as showing what policies were actually desirable.(see Clarida et.al (1998) p1037). We discuss such an optimal interpretation of these equations below. Even interpreting the results as a description of what determined policy actions of the authorities is highly problematic. Fitting econometric equations to instruments and objectives and interpreting the result as an actual reaction function is almost certainly inappropriate. As there is a fundamental identification problem involved in exercises of this sort: there are at least two relations between these variables – the ”true “ policy reaction function (which we assert is not what the authors identify with their equations), and the relationships of the economy itself. The fitted equations combine these two in some unknown way. The authors argue that the benefit of their “weakly restricted” version of the reaction function is that it is sensible for a wide range of different macroeconomic frameworks (models). It is hard to know how we would establish this. Estimating by allowing for the regressor variables to be stochastic, as the authors do, does not deal with this issue. A Full Information method is required as a means of identifying both the responses of the economy to policy and other exogenous shocks on the one hand, and the policy responses to developments in the economy on the other. Furthermore, there is almost certainly structural change affecting both of these basic relations – the model and the policy reaction equations – in different ways. Thus we would strongly suspect that equations of this sort, although apparently well fitting, will exhibit structural instability. This suggests a simple way to test the validity of the Clarida et al approach, is to test its structural stability. Our argument suggests these estimates must be structurally or parameter unstable.

The second problem we identify is that even assuming the estimation has correctly identified the authorities reaction function, the interpretation of this equation is profoundly difficult. The best that can be said is that it represents what the authorities were actually doing. But were they behaving correctly, or optimally, or were they in fact following a completely erroneous set of policies. Even if we can assert that the authorities have been operating a good set of policies are they based on a particular form of co-operative structure? How would the policy have changed if the form of co-operation changed? All these questions are completely unanswerable from the perspective of this methodology and hence represent a sever limitation on its use for policy analysis.

4. GAMES

Policy analysis is often carried out as if the government was the only agent in the economy who is able to exercise any degree of discretion. Even in a closed economy this is not obviously the case, the monetary and fiscal policy makers may not always co-operate perfectly and many sectors of the economy may be able to make there own well informed decisions. When we consider an open economy then there is obviously a need to investigate the form of co-operation between policy makers in different countries and so we need to draw on the game theory literature to structure our analysis of how co-operation may take place and how to investigate the consequences of these structures.

In this section we will give a brief account of the main ideas behind game theory. A more complete exposition may be found in Intriligator (1971). Two primary sources of particular importance are Luce and Raiffa (1957) and von Neumann and Morgernstern (1944).

In the conventional discussion and use of control theory and optimisation given above, the implicit assumption is made that there is one decision maker whose preferences are represented by a single objective function. Game theory is the extension of conventional optimisation theory to the case of multiple decision makers. The generalisation to two ore more decision makers, or players, considerably complicates the problem, as one player's welfare not only depends on his decision and the equations of the system, but also on the decision made by the other player. In its most general form the multiplayer game will not generally yield a unique stable solution and so much of game theory has involved defining particular types of games which are tractable. Games can be classified by the nature of the payoff function; a zero sum game, a constant difference game of a non-zero sum game. They may be classified by the number of players, a two-player game or an n-player game. They may also be classified by the number of strategies each player may adopt, which may be either finite or infinite. Finally they may be classified by the amount of negotiation before play is initiated, we may consider co-operative games if coalitions between players are allowed, or non-co-operative games where no coalition are allowed.

In general the use of co-operative solutions really amounts to a reduction in the number of independent objective functions being maximised. So if we have a two player game a non co-operative game will involve each player independently maximising a utility function. The co-operative game will involve the formulation of a joint utility function (so there is only a single function in this case). Co-operative games, or games which involve subgroups of co-operating players can then be treated by standard optimisation techniques and the presence of co-operative behaviour does not complicate the game structure. The real increases in technical complexity then come about when we consider non-co-operative games.

The most widely explored game is the two-person zero sum game; in this game there are only two players and the set-up of the game is such that each player is competing for a larger share of a total but fixed payout. Even in such a simple game, the possibility arises that no unique solution exists, and locating a solution, assuming one exists, may be analytically very difficult.

The analysis may be extended to allow for uncertainty in the response of one player to the other, to allow for non-zero sum situations and to allow for co-operation between groups of players when there are three or more players. We will not pursue these elaboration's here as the essential point which we require is evident, that even in a simple game the assumption made about the other player's behaviour is critical.

Suppose we consider two individuals who have a respective set of decision variables x^1, x^2 and each individual welfare is a function of his own decision and the decision of the other individual, i.e.

$$U_1(x^1,\ x^2) \text{ and } U_2(x^1,\ x^2) \qquad \textbf{(18)}$$

We are dealing with a non-co-operative game so that each individual selfishly maximises his own welfare regardless of the effect on the other individual, subject only to any constraints which we assume have already been substituted into the objective function. The richness of the game arises from the inclusion of terms in the other player's discretionary variables in the first player's utility function. These terms are not, of course, exogenous to the actions of the player, as the other player may alter his behaviour in response to a move by the first player. Because each player's strategy depends on the strategy of the other player, we cannot use standard optimisation techniques to solve such a problem. In fact, the normal numerical solution procedure is an iterative one but we will not pursue that here. We can however characterise various types of solution. The general solution is given when both players optimise subject to the optimal strategy of the other player this is given by the Nash(1951) solution. It may be characterised by a pair of points $(x^1_N,\ x^2_N)$. which have the property that

$$U_1(x^1_N,\ x^2_N) \geq U_1(x^1{}_i, x^2{}_i)$$

$$U_2(x^1{}_N, x^2{}_N) \geq U_2(x^1{}_i, x^2{}_i) \qquad \textbf{(19)}$$

This states that given that player 2 implements x^2_N player 1 prefers x^1_N to any other permissible choice of x^1 open to him, and similarly given that player 1 implements x^1_N player 2 prefers x^2_N to any other permissible choice of x_2.

To solve a problem for the full Nash solution is numerically difficult even assuming that the solution exists, and often my not be analytically tractable. So two restricted forms of solution have been evolved which do not have the full optimising consistency of the Nash solution but which have the advantage of being much easier to solve. The first of these is the Cournot solution. This is defined as the solution which occurs when each player forms some expectation about the behaviour of the other player, and optimises subject to that expectation. In general this solution will differ from the Nash solution if the expectation about the other player's action differs from his finally chosen action. This solution may be characterised by a pair of points $(x^1_c,\ x^2_c)$. such that

$$\begin{aligned} U_1(x^1_c,\ x^2_o) &\geq U_1(x^1_i, x^2_w) \\ U_2(x^1_o,\ x^2_c) &\geq U_2(x^1_w, x^2_c) \end{aligned} \tag{20}$$

where x^1_o, x^2_o is player 2 and player 1's expectation of the other player's action respectively.

The second form of restricted solution is the Stackelberg game. In this game there is a clear leader and follower. The leader announces some action and the follower optimises subject to that announcement. The leader then optimises subject to the optimal behaviour of the follower. This may be characterised as a pair of points $(x^1_s,\ x^2_s)$ such that

$$\begin{aligned} U_1(x^1_s,\ x^2_s) &\geq U_1(x^1_s, x^2_s) \\ U_2(x^1_o,\ x^2_s) &\geq U_2(x^1_w, x^2_s) \end{aligned} \tag{21}$$

where player 1 is the leader and x^1_o is the announced policy. Player 2 plans on the basis of x^1_o even though it is not an optimal policy for player 1.

When we are dealing with large models all these basic forms of game solution may be solved numerically, although the computational burden of solving for the full Nash solution may be high. The Cournot solution is the easiest to compute as this simply involves standard optimisation conditional on the announced plans of each agent. The stackelberg game is also quite simple to solve as agent 2 carries out a simple optimisation conditional on agent 1's announced plan and then agent 1 carries out another standard optimisation conditional on the outcome of agent 2's optimisation. The Nash game requires an iterated solution procedure which basically consists of a series of standard optimisations, each one carried out conditional on the other agents optimal policies from the last round. Eventually, if the procedure converges each agent will be behaving optimally conditional on the other agents optimal behavior.

Within this structure the players in the game may be optimising an unconstrained objective function or they may be optimising the parameters of a simple feedback rule as discussed above.

5 STOCHASTIC MODEL ANALYSIS

Econometric models, are by there very nature subject to uncertainties. So it is often inappropriate to think only of the point forecast given by a conventional deterministic model solution. We may wish to consider two important consequences of the stochastic nature of models. First the deterministic forecast of a model is not generally the mean forecast of the model. This is true simply because if we have a stochastic process e, with mean zero and variance s^{2}, and we apply some non-linear transformation to it to give a variable z=f(e) then $E(z) \neq f(E(e))$ so whenever we are dealing with non-linear models the expected value of the outcome will be different from our model solution given by the deterministic model. Second we may wish to choose a set of policies which are deliberately designed to reduce the uncertainty of the actual outcome. That is we may wish to give some weight to reducing the variance of the models outcome when we design the policy regime.

The general technique for analysing stochastic models is called stochastic simulations and it is surveyed fully in Hall and Henry(1988). The basic approach used is to draw artificial random numbers from the a set of distributions which are given by the uncertain elements in the model (the error terms and the stochastic parameters). The model is then solved with these random terms added. The process is then repeated many times and the solution values are collected and there distribution is found. For a large enough sample this distribution will converge on the true distribution of the endogenous variables. In our simple illustration above we would draw a random value for e, e_i, then solve for z_i=f(e_i), this would be repeated many times then the average value of all the z's would be calculated and its variance and this would approximate the true distribution of z. This can clearly be a computationally expensive task but it has been widely used by modellers.

A number of approaches have been adopted to the analysis of policy setting within a stochastic regime which broadly parallel the analysis of deterministic models given above. First, optimal control algorithms have been developed for dealing with the correctly specified expected value of the model and also for allowing an optimal policy rule to be chosen which also minimises the variance of a particular outcome. Second, a great deal of work has been undertaken to investigate the effect of simple rules on the stochastic model and to select rules on the basis of minimising the resultant uncertainty.

The proposal of Hall and Stephenson(1990) generated an algorithm which allwed stochastic optimal control to be carried out where the primary focus was on the expected value of the level of the variables in the model (rather than their variance). The notion at the heart of the algorithm was simply to use stochastic simulation to derive a measure of the bias between the expected value of the endogenous variables and the deterministic value. Thus we use stochastic simulation to evaluste d_y in the following expression,

$$E(y_t) = \widehat{Y}_t + d_{yt} \qquad (22)$$

The conventional quadratic objective function can then be re-witten in the following way,

$$\begin{aligned} &E(\sum \Theta_y (y - y^*)^2) \\ &= \sum \Theta_y (\hat{y} + d_y - y^*)^2 + VAR(y) \end{aligned} \qquad (23)$$

Then assuming the VAR(y) is a constant this becomes a standard optimisation problem. Given that y^d is not constant we can iterate between optimal control runs and stochastic simulations but this is a small cost as typically this process converges very quickly. In effect this trick allows us to correct the deterministic controll solution for the deviation between the deterministic and stochastic solutions. It does not however allow any weight to be given to the uncertainty in the model as the variance is assumed fixed.

Rustem(1993) addressed the minimum variance problem by adding a term to the standard objective function which allowed the variance to be approximated. The idea here was to evaluate a matrix of partial derivatives,

$$J_e = \frac{dC}{de} \tag{24}$$

That is the matrix which shows the change in the cost function as the stochastic terms vary, then we simply add a term to the objective function such as,

$$\Omega J_e \tag{25}$$

Which is, in effect the variance of the objective function, and this can be weighted along with the rest of the function to give a robust term. The problem with this approach is that the matrix of partial derivatives is calculated in an iterative fashion between optimal control runs and so the optimal policy is not carried out on the basis of a changing matrix as the control variables change. Hence this technique tends to understate the posibility to minimise the variance. of the cost function.

In a non optimal setting late 1980's and early 90's saw a number of studies which evaluated a range of policy rules or policy options by using stochastic simulation to investigate which rule gave rise to the smallest variance in outcomes. Bryant et al (1989) included a number of studies looking at world models which focused on the stochastic nature of models and selected policy rules on the basis of minimum variance outcomes. Subsequently Brant et al (1992) extended this comparative work to consider a wider range of models. The work of Taylor(1993) is based solely around choosing between a range of policy rules based on a minimum variance criteria, this work has had considerable impact on the policy making community. A common problem with all this work is however that the parameters of the rule are selected in an arbitrary way, a range of different rules are selected to represent different policy regimes. Stochastic simulations are carried out to investigate the size of the resulting variances of the endogenous variables and then the `best' rule is selected as the one which gives the smallest variance in outcomes. However the parameters of each rule are simply chosen on an `ad hoc' basis and it is possible that different parameters in the same rules could have produced a different ordering of results or at the very least performed in a very different way.

Clearly to objectively evaluate the relative performance of a number of different rules we would want each of them to be performing at there most efficient level. So each should have the parameters of the rule selected so as to minimise a common loss function. This would ensure comparability between the rules. This however raises a further question as to what type of game is being played out between and how the optimisation should be carried out when there is more than one policy maker being considered. Ultimately this leads to the conclusion that what we should be doing is to select the parameters of the rule by optimal control so as to minimise the variance of the outcomes based on a range of solutions assumptions such as Nash or stackelberg games. Taken at face value this would imply that we should be carrying out iterating optimisations over the full stochastic model, this would be hugely expensive in computer time (a single stochastic simulation may take a day of computer time). In the next section we will outline a solution procedure which will allow us to calculate the same optimal solutions under a very weak set of assumption which is much more computationally efficient and therefore opens up this ideal form of model solution as a practical possibility for the first time.

6 OPTIMAL RULES TO MINIMISE THE VARIANCE OF ECONOMIC VARIABLES IN A GAME CONTEXT.

Given the preceding discussion, the type of analysis we propose as an advance over current applications and their limitations is to choose the parameters of a set of rules so as to minimise the variance of selected variables in the economy when it is subject to a particular set of stochastic shocks. Moreover, this often needs to be done allowing for possible strategic interaction between different policy makers, so the analysis has to allow for game playing which will involve successive optimisations over a number of players to achieve a range of different forms of solution, eg Nash, stackelberg, co-operative etc.. We therefore specify the problem, in compact notation as,

$$\min \ \operatorname{var}\ (C) = \operatorname{var}\left(\sum_{t=1}^{T} \Phi_t \sum_{i=1}^{n} \Theta_i Y_{it}\right) \tag{26}$$

$$Y_{it} = g(e,u)$$

where e is a matrix of k stochastic terms over the T periods of the model solution which have a given covariance matrix Ω_t, Y is the vector of endogenous variables in the model, Φ and Θ are weights in the cost function and u is a vector of control variables which in our case are the parameters of a control rule. In a policy game each player would have an objective function of this form.

The computational burden of this form of problem is considerable; to evaluate the variance alone needs a stochastic simulation involving thousands of conventional model solutions. This kind of solution would have to be calculated many hundreds of time during a conventional numerical optimisation. It seems that, for this reason alone, researchers have not pursued this approach to policy formulation. The innovation we propose is a simplification of the problem which will yield an identical solution for most forms of nonlinearity which are observed in the large macro models. The idea here is based on the notion that any monotonic transformation of the cost function will yield an identical solution for the control variables. So if we minimise the variance of the cost function (V(.)) with respect to a set of variables u then we will have exactly the same solution for u if we minimised a monotonic transformation of V (e.g. log (V) or V^2). We use these propositions to substantially reduce the computational problem in minimising V(.), using a special transformation based on two elements: the first is the technique of anti-thetic errors used in stochastic simulation, the second constructs a minimum set of replications which exactly reproduce the covariance matrix of the stochastic process.

Anti-thetic Errors

Anti-thetic errors simply mean that instead of drawing a sequence of completely random sets of shocks, the sets of shocks are chosen in symmetric pairs so that two replications from a stochastic simulation represent an exactly symmetric pair, in terms of the shocks being applied to the model. This technique increases the efficiency of stochastic simulations enormously but even one pair gives a lot of information. For example if the model is linear then the resulting average of the endogenous variables from the two solutions will be identical to the deterministic model solution, hence any divergence from the deterministic solution is an absolute sign of non-linearity

Minimum Set of Replications

For the moment lets assume that we are dealing with a single stochastic error term. In that case the following objective function would be our monotonic transformation of (26).

$$\min \; C^* = \sum_{t=1}^{t} \Phi_t \sum_{i=1}^{n} \Theta_i (\, | g(e,u) - g(0,u) | + | g(-e,u) - g(0,u) | \,) \qquad (27)$$

This objective function minimises the absolute deviation from the no shock solution after applying an arbitrary size shock to the model. The antithetic errors are represented by the two terms with plus the shock and minus the shock. Our claim is that there is a monotonic transformation between this objective function and (26). Hence the resulting optimal u will also be the solution to (26).

If we were dealing with a single error this would obviously be sufficient to give the solution we require. However, there is a further complication when the vector of errors is larger than a single scalar. The problem is that any single draw of the error vector cannot be representative of the whole distribution of errors, so it cannot represent the covariance matrix. A scalar error can have a value equal to its standard error but a vector cannot have both variances and covariance's equal to the full covariance matrix. This point can be seen by considering the bivariate case. Let the covariance matrix be,

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad (28)$$

Now any single pair of shocks cannot give both the variance and covariance's simultaneously. For example (1,1) has unit variances for both errors but a unit covariance, (1,0) would have a zero covariance but the variance on e_2 would also be zero. In fact, in this case it takes two sets of shocks to exactly replicate the covariance matrix. The required shocks are (1,-1), (1,1), which have unit variances for both errors and zero covariance. The anti-thetic pair corresponding to this would be (-1,1), (-1,-1). So if we were interested in solving the problem for a vector of two stochastic shocks we could do this by evaluating

$$\min \; C^* = \sum_{j=1}^{k} \; (\sum_{t=1}^{t} \Phi_t \sum_{i=1}^{n} \Theta_i (\, | g(e^j,u) - g(0,u) | + | g(-e^j,u) - g(0,u) | \,) \,) \qquad (29)$$

where k=2 and where the two vectors of shocks (e^j) are given as above. So in this case, instead of carrying out many thousand replications to estimate the variance of c, (VAR(C)) we can achieve the same object by calculating C* based on only four model solutions. This clearly brings the possibility of using optimal control within the bounds of computational feasibility, even in a game context.

The above case is an example of how the proposed procedure would work for a case of two shocks. In the general procedure we chose a set of k vectors of shocks such that

$$\Omega = \sum_{i=1}^{k} e_i e_i \tag{30}$$

This will generally involve approximately n=k sets of shocks where n is the number of stochastic elements in the model being examined. The reason why this is only approximate is that the relationship is different for an even and odd number of shocks. The above formulae gives an exact determination of the shocks when n is odd but when it is even we need some extra conditions to uniquely determine the shocks. In the bivariate case above, for example, there are actually an infinite number of pairs which would give the required covariance matrix. This can be seen by writing out the problem in full.

$$\begin{aligned} \Omega_{11} &= e_{11}^2 + e_{21}^2 \\ \Omega_{12} &= e_{11}e_{12} + e_{21}e_{22} \\ \Omega_{22} &= e_{12}^2 + e_{22}^2 \end{aligned} \tag{31}$$

this yields 3 equations in four unknowns and so we need to impose an extra condition to uniquely determine the shocks, we propose simply setting $e_{11}^2 = \Omega_{11}$ as the extra required restriction. For an odd number of shocks we exactly determine the k vectors of errors.

The following table gives the relationship between n, the dimension of the covariance matrix, k the minimum number of sets of shocks and r, the number of extra sets of restrictions required.

N	K	R
1	1	0
2	2	1
3	3	3
4	4	6

5	5	10
6	6	15
7	7	21

So in general, given the extra effect of the antithetic errors, we will need approximately twice number of replications as the dimension of the covariance matrix. If we wish to calculate an optimal policy rule for a countries monetary policy given shocks to both the exchange rate in that country and shocks to the exchange rate in two other country's we would therefore need six model solutions to evaluate the objective function we need to maximise.

The Monotonic Transformation

This proposed technique will not always give exactly the same answer as (26) above, it is possible that for a sufficiently non-linear model the mapping between (26) and (29) would cease to be monotonic and hence they would have different solutions. However our argument is that this would require and extremly perverse and unusual form of non-linearity to be present which is not typical of any macroeconomic model.

The essence of the monotonicity assumption is that if we have any two sets of control variables, u^1 and u^2, such that

$$C*(u^1) > C*(u^2) \tag{32}$$

that is, a deviation in C from its deterministic value is larger for the set of control variables u^1 than u^2. Then monotonicity between the two objective functions means that

$$\mathrm{var}(C(u^1)) > \mathrm{var}(C(u^2)) \tag{33}$$

This simply amounts to the assumption that if one set of control produces larger deviations in the model variables from their deterministic values then it will also lead to a larger variance. In our view it is almost inconceivable to think of an economic model where this would not be true.

7. The Policy Co-Ordination Exercises

In these exercises we initiate each solution with a shock to government expenditure in the US of approximately 1\% of GNP for the entire solution period from 1984--1994. Four quite separate forms of solutions are then compared. In the first, the US react optimally in terms of its own but monetary policy (interest rates) remain fixed in Germany and Japan. The same exercise is repeated for each country, giving a ``national'' optimal rule for each. In the second, every country uses its ``national'' rule from the first case to optimise. The third case is where each country optimises in the light of, and the knowledge of optimal behaviour in each other country ---A Nash solution. In the last case we have a fully cooperative solution. In what follows we refer to the first as single country optimising, the second as multicountry I (where each of the country assumes no policy reaction from the other). The third is multicountry II (Nash) and the last is multicountry III which is a fully co-operative solution. All solutions are for the period 1984--1994.

Single country optimising

In this alternative, each national authority optimises the weights of its PID monetary rule, in order to minimise the (proportional, integral and derivative) deviations of inflation from its base following the shock to the US fiscal deficit. But in these exercises, in each country, policy actions are governed by the national monetary rule, and there is no policy reaction from the other countries. There are {\em consequences}\/ for each country which flow from the actions of the others nevertheless. These take the form of the first of the spillover effects noted in our introductory section. That is, these are orthodox trade quantity and trade price effects affecting, in this case, Germany and Japan, following the US fiscal expansion. These operate through net trade and the real exchange rate. But as monetary policy in the US is tightened to counteract the inflationary effects of the fiscal stimulus there, we assume that interest rates (and hence real interest rates) do not rise elsewhere. Case (ii) introduces this further effect as explained in the introduction, and these interest rate changes will then exert additional effects upon wages, prices and employment in the medium term, and hence the inflation-unemployment choices for the authorities. How important are these latter interest rate effects? We can consider the evidence on this in comparison between single country optimising and multicountry I.

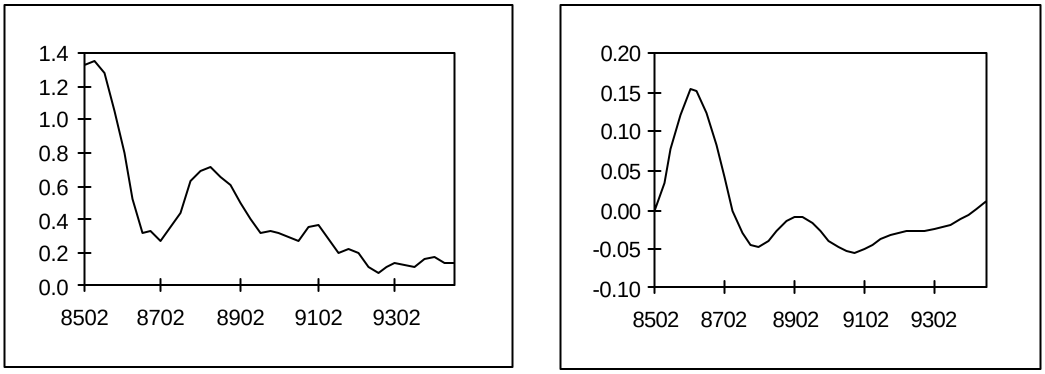

Figure 1. Single Country Optimising: Output and Inflation Effects on US of US Fiscal Shock.}

Figures 1a and 1b show the output and inflation effects of the US fiscal shock accompanied by optimal single country monetary policy response in that country. Output growth increases by over 1% initially, but reduces to around 0.5\% over the next 4 years as monetary is tightened. (Figure 1a). As Figure1b shows, the policy correction is successful in reducing the inflationary impulse, and by the end of the simulation the rate of inflation has reached its base value.

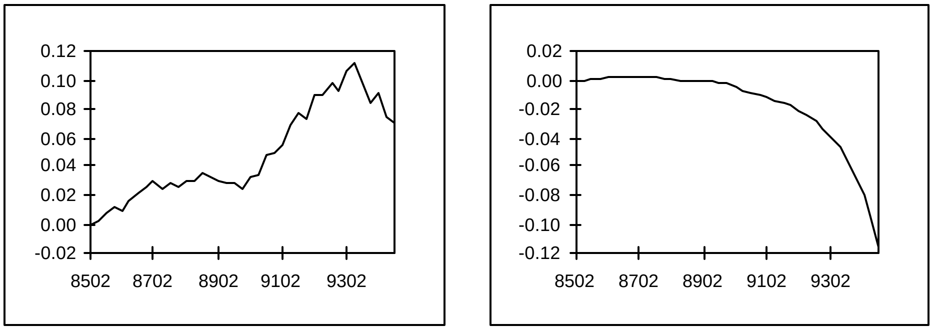

Figure 2. Single Country Optimising: Output and Inflation Effects on Japan of US Fiscal Shock.}

If we reverse the roles next, and let Japan's monetary policy react optimally to the US shock we get the effects shown in Figures2a and 2b. Output growth picks up steadily, but by small amounts, to reach 0.1% higher after about 8 years (Figure 2a) before falling. This stimulus operates through familiar net trade effects. There is a very small inflation effect from this. However, as Figure2b shows, the authorities reduce inflation by a small amount before base (0.1\% after 10 years), entailing a small gain in output with slightly lower inflation by the end of the simulation. A similar pattern emerges when Germany optimises monetary policy following the US expansion, with similar orders of magnitude but with some differences in timing. The differences overall are not significant enough to warrant separate treatment though. (Figures for Germany are therefore not included.) The optimal weights obtained in this set of single country exercises are then used in the next exercise, which begins the multicountry analysis proper.

Multicountry I

We are now in a position to analyse in a preliminary way the optimal responses to the US fiscal shock on a proper multicountry basis. In this next exercise, all countries responds together, each country according to its own optimal monetary policy rule derived from the simple country optimising exercise above. It is a limited form of multicountry response: although each country follows a (national) optimal rule, it assumes there will be no policy reaction in the other countries. this is an incorrect assumption to make, and we explore the effects of relaxing it in (iii) below. However, the present exercise does introduce further forms of spillover compared with the traditional case (which came in (i) above). Firstly, there are effects between interest rates across countries due to the workings of interest arbitrage. Second there are policy induced effects on interest rates, as each national authority seeks to offset the inflation consequences of the US fiscal expansion, using its own monetary policy rule. For both reasons, there will be inflation and unemployment effects due to the effects of changing interest rates on expenditures, including investment, and thence the capital stock.

Even though each country is (in this limited sense) making an optimal response to the US shock, the effects of it on Germany and Japan are striking. Figure 3 shows the effects on growth, Figure 4 the effects on inflation. Growth in the US expands more than in the previous case . Inflation rises there too, but by only a small amount (Figure 4), peaking at about 0.5\% above base after 1.5 years. The repercussions in the other countries are profound, especially in Japan. Inflation picks up markedly over the first two years, and remainsstubbornly high for a further 4 years, at something under 2% above base. The source is the appreciating dollar, and rising import price inflation in the two other countries. In consequences, monetary policy has to be tightened very sharply in both Japan and Germany.

This inflationary effect is compounded initially by the expansionary effect of the fiscal stimulus in the US on Japanese and German growth. After 2 years in the case of Germany, and 3.5 years in Japan, the strongly corrective monetary policy reduces growth. It proves difficult to reduce inflation in Japan [because $\dotsb$] and growth there is reduced substantially over most of the simulation period in the effort to contain the inflationary effects of the increased US deficit

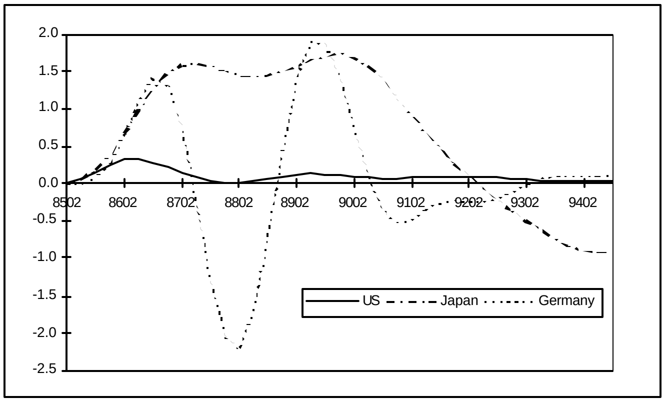

Figure 3.MultiCountry I: Inflation Effects of US Fiscal Shock.}

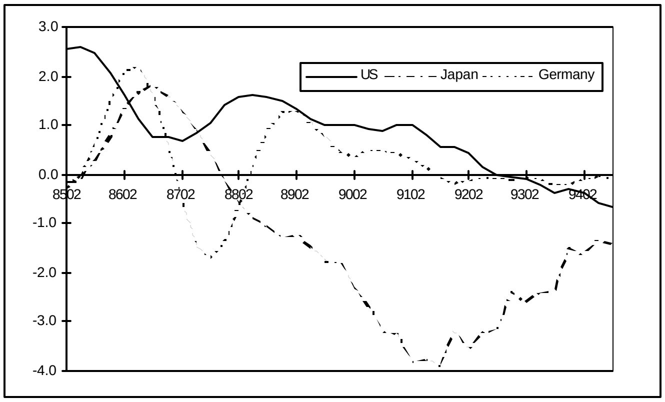

Figure 4. MultiCountry I: Inflation Effects of US Fiscal Shock.}

Although this exercise is obviously limited---it assumes that each country assumes the others will not react to its own policy changes, incorrectly---it indicates that the spillover effects of uncoordinated fiscal expansion can be very substantial indeed. Why does this finding differ so much from the typical finding of limited spillovers? There are two parts to the answer to this. The first is that the transmission mechanisms included in our exercise are more elaborate than normally used. In particular the emphasis we place upon the medium term effects of interest rate changes, the capital stock and the supply side gives added potency to the international transmission of fiscal shocks which themselves impinge upon interest rates (via orthodox crowding out {\em and}\/ because the fiscal shock stimulates monetary responses through the monetary rule). The second is that monetary policy is an optimal policy aimed to squeeze inflation shocks out. In practice monetary reactions to inflation changes have not proved so severe.

Multicountry II

One of the limiting assumptions in the previous exercise is now dropped, and we proceed to implement a full Nash solution on the optimisation. Allowing for each country to optimise, given that it assumes (correctly) that each of the other does the same, has evident consequences for the outcomes following the US fiscal shock. Figures5 and 6 show the growth and inflation differences from base in this regime. As compared with the previous exercise, the growth effects---with one exception---are more constrained: the expansion in the US is less initially, and more stable; while Japan's experience is also much less severe, although it again has the same sort of prolonged growth recession as in the previous case, the fall in growth being about half that of the previous case at its worst. For both countries, the major gain in this exercise is on inflation. Unlike the Multicountry I, inflation in Japan is reasonably well contained, rising between 0.3--0.5% until the end of 1990, but is effectively squeezed out thereafter (Figure 6}). Inflation in the US is broadly the same as in the earlier case. Hence, this case may be characterised as showing that better inflation can be achieved with smaller output losses when adopting Nash-type optimal strategic policies compared with single country optimising.

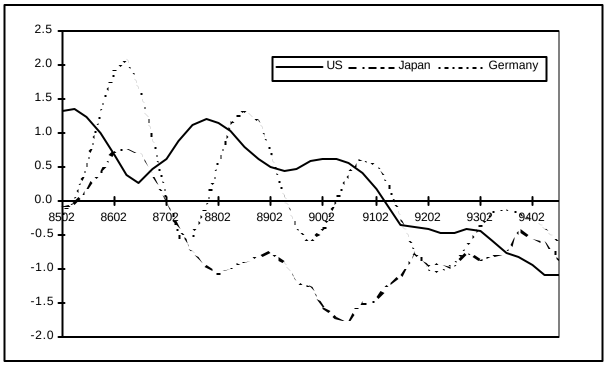

Figure 5. MultiCountry II:Output Effects of US Fiscal Shock.

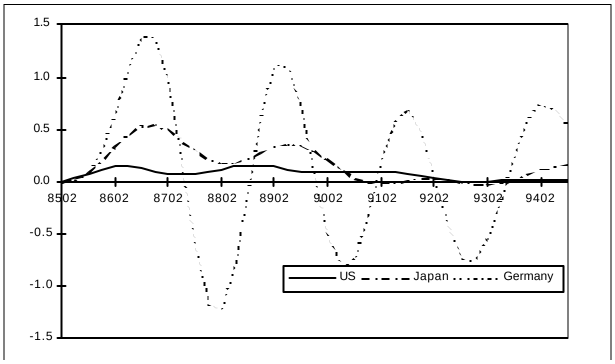

Figure 6. MultiCountry II: Inflation Effects of US Fiscal Shock.

[The exception to this is Germany, which appears to be the loser in this exercise. There is marked cyclicality in output responses and inflation, though being stabilised, is so very slowly, and again with marked oscillations.]

Multicountry III

Once a fully cooperative international policy regime is instituted the situation is transformed, showing substantial gains over the full Nash solution. Figures 7 and 8 give the growth and inflation differences from base for this case.

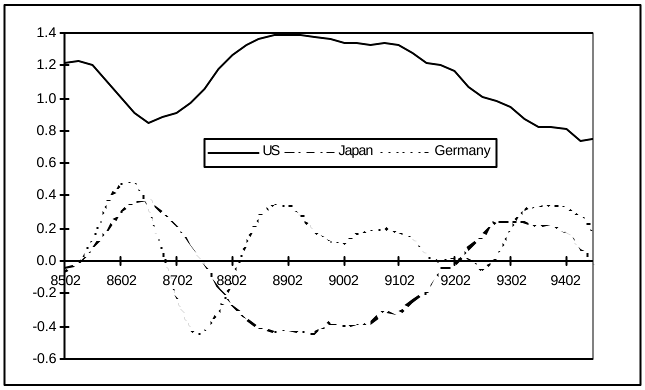

Figure 7. MultiCountry III: Output Effects of US Fiscal Shock.}

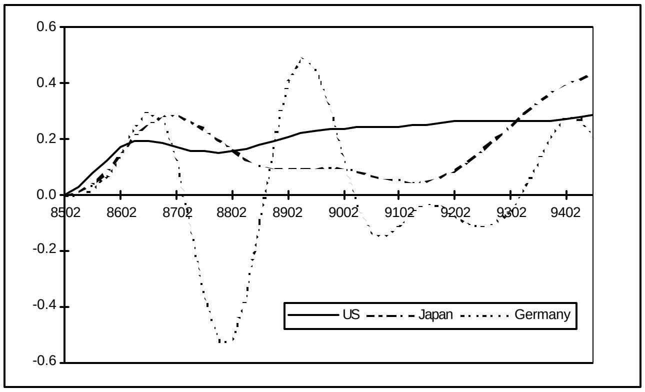

Figure 8. MultiCountry III: Inflation Effects of US Fiscal Shock.}

The most conspicuous effect is upon US growth which now is positive throughout and much more stable. Although not as high as initially in case (ii) where there is no policy reaction at all from other country, in this case, there is a positive increment to growth throughout the period (and by the end of 1994 it is still 0.75\% above base). similarly, in the other two countries, the adverse effects on growth are minimised in this regime. The adverse effects on growth in Japan are much shorter lived than in both of the non-cooperative exercise, and are much less severe. Germany also has a short-lived fall in growth compared with base, but positive effects there after. (Figure 7).

[Inflation is effectively contained in each country after some cycling in Germany. At the end, each has inflation some 0.3--0.4% above base. (Figure 8)].

7 CONCLUSION

In this chapter we have outlined the conventional approaches to the modelling of policy formulation. Traditionally this has fallen into one of two groups, fully optimal responses or the use of some simple feedback rule to represent the policy makers behaviour. These two approaches may be carried out either for a single agent or in a game setting between a number of policy makers. There are a number of conceptual reasons for favouring the use of simple rules and modern economics tends to favour the idea of policy stabilisation rather than trying to change the permanent behaviour of the economy. This leads us towards designing optimal policy rules which minimise the variance of the economy when it is subject to shocks. Computationally this has been very difficult with existing algorithms. We have proposed an algorithm here which makes this a feasible computational task even in a full game context.

.

REFERENCES

Allen C, Hall S G and Nixon J (1994) "The New London Business School Model of the UK Economy", *London Business School, Centre for Economic Forecasting Discussion Paper*, No.18-94.

Barro R J and Gordon (1983) "Rules, discretion and reputation in a model of monetary policy", *Journal of Monetary Economics*, Vol. 22, pp. 101-21.

Becker R G, Dwolatzky B, Karakitsos E and Rustem B (1986) "The simultaneous use of rival models in policy optimisation", *The Economic Journal*, Vol. 96, pp. 425-48.

Blake, A P and Westaway P (1994) "Targeting Inflation with Nominal Interest Rates", *National Institute of Economic and Social Research, Discussion Paper*, No. 70.

Bray J, Hall S, Kuleshov A, Nixon J and Westaway P (1995) "The Interfaces Between Policy Markers, Markets and Modellers", *The Economic Journal*, Vol. 105, No. 431, pp. 989-1000.

Bryant R C, Henderson D W, Holtham G, Hooper P, and Symanski S,(eds.)`(1988) `Empirical Macroeconomics for Interdependent Economies' Washington DC, Brookings Institute.

Bryant R C, Hooper P, and Mann CL,(eds.)(1993)`Evaluating Policy Regimes: New Research in Empirical Macroeconomics' Washington DC, Brookings Institute.

Budd A, Christodoulakis N, Holly S and Levine P (1989) "Stabilisation policy in Britain", in *Policy Making with Macroeconomic Models*, Andrew Britton (ED), Gower, Aldershot.

Chow G C (1975) *Analysis and Control of Dynamic Economic Systems*, Wiley, New York.

Church K B, Mitchell P N, Smith P N and Wallis R F (1995) "Comparative properties of models of the UK economy", *National Institute Economic Review, No. 153*, August.

Clarida, R., Gali, J., and Gerlter, M., (1998), "Monetary Policy Rules in Practice: Some International Evidence", "European Economic Review, 42. p.1033-1067.

Currie D A and Levine P (1985) "Simple macroeconomic policy rules in an open economy", *The Economic Journal*, Vol.85, pp. 60-70.

Edison H, Miller M and Williamson J (1988) "On evaluating and extending the target zone proposal", *Journal of Policy Modelling*.

Fair R C (1984) *Specification, Estimation and Analysis of Macroeconomic Models*, Harvard University Press, Cambridge, MA.

Garratt A and Hall S G (1993) "A Proposed Framework for Monetary Policy", *London Business School, Centre for Economic Forecasting Discussion Paper*, No. 9-93.

Gawthrop P and Vines D (1989) "The robustness of simple policy rules compared with optimal policy rules: an example", in *Dynamic Modelling and Control of National Economies*, Vol. 2, International Federation of Automatic Control.

Haldane A (1995) "Inflation Targets", *Bank of England Quarterly Bulletin*, Vol. 35, No. 3, August, pp. 250-259.

Hall S G and Garratt A (1992a) "Model Consistent Learning: the Sterling Deutschmark Rate in the London Business School Model", *London Business School, Centre for Economic Forecasting Discussion Paper*, No. 2-92.

Hall S G and Garratt A (1992b) "Expectations and Learning in Economic Models" *Economic Outlook* Vol. 16, No. 5, pp. 52-53.

Hall S G and Henry S G B (1988) *Macroeconomic Modelling*, North Holland, Amsterdam.

Hall S G and Stephenson M J (1990) "Optimal control of stochastic non-linear models", in Henry S G B and Patterson (eds), *Economic Modelling at the Bank of England*, Chapman and Hall, London.

H M Treasury (1993) *Financial Statement and Budget Report 1994-95*, November, London, HMSO.

Intriligator M.D.(1971) `Mathematical Optimisation and Economic Theory' Prentice Hall, London.

Kydland F E and Prescott E C (1977) "Rules rather than discretion: the inconsistency of optimal plans", *Journal of Political Economy*, pp. 473-92.

Leiderman L and Svensson L E O (1995) "*Inflation Targets*", Centre for Economic Policy Research, London.

Lucas R E (1976) "Econometric policy evaluation: a critique" in *The Phillips Curve and Labour Markets*, K Brunner and A H Meltzer (eds), North Holland, Amsterdam, pp 19-46.

Luce R.D. and Raiffa H. (1957) `Games and Decisions' John Wiley and Sons, New York.

McCallum B T (1988) "Robustness Properties of a Rule for Monetary Policy", Carnegie Rochester Conference Series, No. 29, Autumn, pp. 173-203.

Nash J.F.(1951)`Non-Co-operative Games' Annals of Mathematics, p286-95, 54.

Phillips A W (1954) "Stabilisation Policy in a Closed Economy", *The Economic Journal*, Vol. 64, pp. 290-323.

Phillips A W (1957) "Stabilisation Policy and the Time-form of Lagged Responses"; *Economic Journal*, Vol. 67, pp. 265-277.

Rustem B(1993) `Stochastic and Robust Control of Nonlinear Economic Systems’ European Journal of Operations Research.

Smith P N and Wallis K F "Policy Simulations and Long-Run Sustainability in Forward Looking Macroeconometric Models", *ESRC Macromodelling Bureau Discussion Paper* No. 34, University of Warwick.

Stock J H and Watson M W (1989) "New Indexes of Leading and Coincident Economic Indicators", *National Bureau of Economic Research Macroeconomics Annual*, pp 351-394.

Stock J H and Watson M W (1991) "A probability model of coincident economic indicators" in K Lahiri and G H Moore (eds), *Leading Economic Indicators: New approaches and forecasting records*, CUP, New York, pp. 63-85.

Svensson (1998), “Open Economy Inflation Targetting”, NBER Working Paper No. 6545. Cambridge, Mass.

Taylor J B (1985) International coordination in the design of macroeconomic policy rules", *European Economic Review*, Vol. 28, pp. 53-82.

Vines D, Maciejowski J M and Meade J E (1983) *Stagflation, Vol 2: Demand Management*, George Allen and Unwin, London.

Von Neumann J. and Morgenstern O. (1944)`Theory of Games and Economic Behaviour' Princeton University Press, Princeton NJ.

Walton and Massone (1999), “MPC Still Following Taylor Rule”, Goldman Sachs UK Weekly Analyst, Issue No. 316, 5th February.

Wood G E (1994) "A Time to Act or to Bide Your Time?", *Economic Affairs*, Vol.15, No.1.

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