Bank of England

Discussion Papers

Technical Series

No 18

Optimal control of stochastic non-linear models

by

S G Hall

I R Harnett

M J Stephenson

October 1988

No 18

Optimal control of stochastic non-linear models

by

S G Hall

I R Harnett

M J Stephenson

October 1988

The object of this Technical Series of Discussion Papers is to give wider circulation to econometric research work predominantly in connection with revising and updating the various Bank models, and to invite comment upon it; any comments should be sent to the authors at the address given below.

This paper was presented at the ESRC Macroeconomic Modelling Seminar held at the University of Warwick in July 1988. The views expressed are those of the authors and do not necessarily represent those of the Bank of England.

Issued by the Economics Division, Bank of England, London, EC2R 8AH to which requests for individual copies and applications for mailing list facilities should be addressed; envelopes should be marked for the attention of the Bulletin Group.

ISBN 0 903312 96 4
ISSN 0263-6123

OPTIMAL CONTROL OF STOCHASTIC NON-LINEAR MODELS

Introduction

The use of large macroeconometric models has become central to the policy work of many national and international agencies. The models are not only complex in terms of their size but are also invariably non-linear. Using such a model in policy formulation leads inevitably towards specifying a set of objectives and trying to meet these objectives as closely as possible. The natural formal framework in which to characterise this procedure is clearly that of optimal control (see for example Fair (1984) which summarises a great deal of work in this area or Chow (1975)). Econometric models are however by their very nature stochastic and when they are also non-linear the solution to the deterministic control problem may differ substantially from the solution to the stochastic problem.

It is now almost ten years since the Report of the Committee on Policy Optimisation (1978) (the Ball Report) appeared. During that period the use of optimal control techniques has become more widespread and a considerable research effort has been made towards overcoming some of the difficulties identified in the Ball report. In particular, three main objections were made to the use of control techniques: the uncertainty which exists as to the correct specification of the economic model, the problem of identifying a suitable objective function, and the difficulty of formulating policy when the economy is made up of rational agents who alter their behaviour as the policy changes. The problem raised by rational agents has been the subject of an enormous body of literature over recent years - indeed it has even given rise to whole new areas of debate, such as the time inconsistency and credibility debates. There has been less obvious progress as to the correct specification of appropriate objective functions, but this is perhaps more an area where experience is needed rather than new techniques and theoretical insights, and experience of conducting optimal control exercises has grown enormously. In our view the main remaining objection to the use of optimal control in practical policy formulation now lies in the uncertainty over the true structure of the economy. It seems unlikely that economists will reach a consensus as the the true model of the economy in the near future, and even if such a consensus were reached there would still be some uncertainty attached to the model in the form of stochastic parameters and error terms. There is therefore a need for a practical procedure which would allow optimal control techniques to be extended to stochastic non-linear models.

This problem has been addressed by Chow (1976) from a theoretical standpoint, and he outlines an algorithm which calculates optimal control rules for stochastic non-linear models. The Chow algorithm, in essence, works by iterating over a number of linearisations of the stochastic model, using standard dynamic control theory to optimise the stochastic linearised model at each iteration. The key feature of the algorithm is that it is the stochastic model which is linearised, not the deterministic model. To linearise a large stochastic model once would be enormously difficult and to include this as part of an iteration procedure would be an order of magnitude more complex, and as far as we know the Chow (1976) algorithm has never been implemented in its full form. A few applications exist of stochastic optimal control of fairly small models (eg Bray (1975)) but this work has generally proceeded by linearising the deterministic model rather than the full stochastic model. These applications then tend to produce solutions close to the deterministic solution (as we would expect) - indeed if they were performed using fixed parameters and only error term uncertainty this algorithm converges on the deterministic solution.

There have been two other areas of development in applied work which contribute towards the development of optimal control techniques for stochastic non-linear models, although they both offer only a partial answer to the problem. Some studies have been undertaken which allow across model optimisation exercises to be carried out (see Becker et al (1986) or Hall and Henry (1988) for examples). This approach recognises that we are uncertain as to which of a range of models may truly represent the real world and so optimisation is carried out subject to a range of models weighted together. This approach is useful when we are faced with a small range of discrete alternatives but it is not a practical way of dealing with general uncertainty: indeed it cannot even allow for the stochastic terms in a single non-linear model. The other approach is to derive simple, supposedly robust, feedback rules for use in policy formulation (see for example Currie and Levine (1985), Taylor (1985), Vines, Maciejowski and Meade (1983) and Edison, Miller and Williamson (1988)). These simple rules are not generally optimal feedback rules, although their parameters may sometimes be chosen using an optimal control technique. This means that the use of a simple rule must inevitably lead to a loss in overall economic performance relative to the full optimal policy. The claim of robustness, which is often made, has no strong theoretical foundation. The full optimal feedback rule is not robust to changes in model structure, and there is no reason why a simple rule should be any more robust except that it may be so inefficient that a changing model structure has no real effect on it.

There is therefore a need for a practical algorithm to calculate optimal control solutions for a non-linear stochastic model directly. The purpose of the paper is to propose a new technique which combines the approach of stochastic simulation with standard optimal control techniques to produce an algorithm which allows such calculations. Section 2 of the paper will outline the algorithm and contrast it with an existing approach. Section 3 will present an illustration of the approach using the large forecasting model of the Bank of England.

2 The Algorithm

Suppose we define the solution of a non-linear model to be

$$f_i\ (Y,\ X,\ \Omega,\ A,\ e) = 0 \qquad i = 1,\ N \tag{1}$$

where Y is a vector of N endogenous variables, X is a vector of M exogenous variables, Ω represents the variance covariance matrix of the parameters A and the error terms e. If we define Y^* to be the solution to (1) subject to the stochastic parameters and error term then we may define

$$E\ (f_i\ (Y^*,\ X,\ \Omega,\ A,\ e)) = 0 \qquad i = 1,\ N \tag{2}$$

and Y^* will be the mathematical expectation of Y.

The deterministic solution to the model may be defined as $\hat{Y}$ where,

$$f_i\ (\hat{Y},\ X,\ 0,\ E(A),\ 0) = 0 \qquad i = 1,\ N \tag{3}$$

That is the variance covariance matrix of the parameters are set to zero and the error terms take their mean value, which is assumed to be zero without loss of generality. We know that, when the model is non-linear

$$\hat{Y} \neq Y^*$$

We may extend this framework to include optimal control by splitting the X vector into two sections: Z, a vector of exogenous variables and u, a vector of control variables. We then only need to specify a suitable objective function which is to be minimised: for the purposes of exposition we will use a conventional quadratic objective function. So let

$$E(J) = E\left(\sum_{i=1}^{n} A_i (Y_i - \bar{Y}_i)^2\right) \tag{4}$$

where $\bar{Y}_i$ is the desired value for variable Y_i and (4) is to be minimised subject to the model.

$$f(Y, Z, u, \Omega, A, e) = 0 \tag{5}$$

with respect to the control variables u. Again without loss of generality, we assume a one period time horizon so as to simplify the notation, the multi-period extension is trivial.

Now we may rewrite (4) in the following way

$$E(J) = \sum_{i=1}^{N} A_i\, E(Y_i^2 + \bar{Y}_i^2 - 2\, Y_i\, \bar{Y}_i) \tag{6}$$

$$= \sum_{i=1}^{N} A_i\, (E(Y_i^2) + \bar{Y}_i^2 - 2\, \bar{Y}\, E(Y_i)) \tag{7}$$

and given that $E(Y_i)^2 = E(Y_i)\, E(Y_i) + VAR(Y_i)$

$$E(J) = \sum_{i=1}^{N} A_i\, (E(Y_i)\, E(Y_i) + VAR(Y_i) + \bar{Y}_i^2 - 2\, \bar{Y}_i E(Y_i)) \tag{8}$$

and we may define $E(Y_i) = \hat{Y}_i + E(d_i)$, the expected value of Y_i equals the deterministic model solution $\hat{Y}_i$ plus the expected deviation of the deterministic value from the mean value $E(d_i)$. Substituting this into (8) gives

$$E(J) = \sum_{i=1}^{N} A_i\, [\hat{Y}_i\, \hat{Y}_i + E(d_i)\, E(d_i) + 2\, \hat{Y}_i\, E(d_i) + VAR(Y_i) + \bar{Y}_i^2 - 2\, \bar{Y}_i\, \hat{Y}_i$$

$$- 2\, \bar{Y}_i\, E(d_i)] \tag{9}$$

The advantage of (9) over (4) is that the stochastic elements of the solution have been isolated in the terms $VAR(Y_i)$ and $E(d_i)$ and we are able to provide numerical estimates for both of these terms through the use of stochastic simulation.

Stochastic simulation is a numerical technique which allow estimates to be made of the density function of the endogenous variables of large non-linear models. We will not discuss the details of this technique here, a general discussion of the approach may be found in Hall and Henry (1988) or Hall (1986). The main point of interest here is that stochastic simulation provides numerical estimates of both $VAR(Y_i)$ and $E(d_i)$. This suggests a solution algorithm which involves iterating between conventional optimal control exercises and stochastic simulation exercises to produce a solution to the full stochastic optimal control problem. This might be done by using the following step by step algorithm.

1 Calculate the optimal solution to the deterministic problem given by (4) and (3), let the solution be u^*.

2 Perform a set of stochastic simulation around the base given by u^* to produce estimates of $VAR(Y_i)$ and d_i $i=1, N$.

3 Using these estimates of d_i and $VAR(Y_i)$ we can now minimise (9) subject to (3) to produce a new optimal solution u'. If u' is within a convergence criteria of u^* ($|u' - u^*| < EPS$) then stop, if the convergence criteria is not met then set $u^* = u'$ and return to step 2.

This algorithm will at convergence, still entail a small degree of approximation although this will be much less than the usual method of producing a linear approximation to the non-linear model. The conventional procedure of linearising the deterministic model, discussed in Kendrick (1981) would involve producing a linear approximation to the model and then appealing to the certainty equivalence theorem to solve the resulting quadratic-linear model deterministicly. The problem with this approach is that when the objective function is quadratic and the parameters are known this procedure simply reproduces the deterministic solution.

We can see the source of the above approximation by noting that in general $VAR(Y_i)$ and d_i are both functions of the control variables u. Without loss of generality we may simplify the notation by considering an example with only one control variable (u) and one state variable Y. Then following the notation in (9) we may define

$$\hat{Y} = f(u) \tag{10}$$

$$VAR(Y) = g(u) \tag{11}$$

$$E(d) = h(u) \tag{12}$$

These terms may then be substituted into (9) to give

$$E(J) = f(u).f(u) + h(u).h(u) + 2\ f(u).h(u) + g(u) + \bar{Y}^2 - 2\bar{Y}.f(u)$$

$$- 2\bar{Y}.h(u) \tag{13}$$

This is now an unconstrained function in u which will be minimised when the following FOC is met.

$$2f(u)f' + 2h(u)\ h' + 2\ f(u)h' + 2\ h(u)f' + g' - 2\bar{Y}f' - 2\bar{Y}h' = 0 \tag{14}$$

In the algorithm given above during the calculation of the optimal solution the partial derivatives g′ and h′ are set to zero so the solution which is calculated will be characterised by

$$2f(u)\ f' + 2h(u)\ f' - 2\bar{Y}f' = 0 \tag{15}$$

The standard technique of linearising the model would also set h(u)=0 and so this term would also be lost in the approximation. It must be appreciated at this point that h(u), the deviation between the deterministic value of Y and its expected value, is of a quite different order of magnitude to g′ and h′, the derivatives of the deviation and the variance with respect to u. For most model applications g′ and h′ are likely to be so small that ignoring them is a reasonable approximation to make. However, if it is felt that a particular model is so non-linear that this is a damaging assumption then it is possible to reduce this level of approximation by estimating simple linear approximation for g(u) and h(u): two sets of stochastic simulation could be performed for different levels of u and a simple linear function for g(u) and h(u) could be calculated. Under normal circumstances however the main effect of the stochastic parts of the model will be captured by the term h(u).

Finally it is perhaps worth noting that the well known certainty equivalence theorem can be demonstrated via equation (9) and (14). Certainty equivalence states that if the objective function is quadratic and the model is linear with normally distributed error terms then the optimal control trajectory for the stochastic problem is identical to the solution to the deterministic problem. This can be seen as when the model is linear with normally distributed error terms then $h(u) = g' = h' = 0$ and so (14) reduces to

$$2f(u)\, f' - 2\bar{Y}\, f' = 0 \tag{16}$$

which is identical to the FOC for the deterministic model.

Feedback Rules and Minimising Variances

For many forms of applied policy making the possibility of stochastic optimal control is itself not a complete answer to the problem of uncertainty. The reason for this is that for many purposes we may not simply want to minimise the expected value of some objective function but we may also wish to add an effect related to the variance of the outcome relative to some baseline. As an example, suppose we wish to target a particular level of GDP and have a choice of either fiscal policy (in the form of government expenditure) or monetary policy (in the form of interest rates) as control variables. This is a one target two instrument problem and conventional deterministic control techniques will not be able to choose a single policy mix. But if we also want to minimise the variance of GDP around its desired level then a unique choice of policy instruments can generally be made. Another example where the variance is an important consideration is in the calculation of optimal feedback rules. Again in this case we might well not only want to hit a given target for the state variables but we might also want to minimise the variance around the target level.

We can again use information from stochastic simulation to construct an objective function which will make allowance for this effect. In particular Hall (1986b) or Fair (1980) gives an algorithm which evaluates the density function of a model's simulation properties. That is to say, if we change a variable u then the resulting change in Y will have a certain degree of uncertainty attached to it, as a simple approximation we may assume that $\Delta Y_i \sim (\delta_{ij}\Delta u_j, (\sigma_{ij}\Delta u_j)^2)$ where δ_{ij} is the expected value of the change in Y_i and σ_{ij} is the standard error of this change.

We could then augment the standard objective function (4) by terms in the variance of Y associated with different control variables. This would have the following general form

$$E(J) = E\left(\sum_{i=1}^{N} A_i (Y_i - \bar{Y}_i)^2\right) + B \sum_{i=1}^{N} \sum_{j=1}^{v} (\sigma_{ij}(u_j - \bar{u}_j))^2 \qquad (17)$$

The first term is the standard one from equation (4) the second term represents the variance in the state variables caused by the active use of the control variables relative to some base line value $\bar{u}$. B is a weighting parameter which evaluates the importance of minimising the variance relative to the target levels of the state variables.

If we were using this to calculate simple feedback rules then the control variables would be parameters of the feedback rule. In this case an obvious value for $\bar{u}$ would be zero, ie not operating a feedback mechanism.

The advantage of an objective function of this type is obvious: in standard control exercises we often have a selection of control variables which may be used which vary enormously in the degree of uncertainty over their effect. An obvious example here is direct government expenditure and interest rate policies. In most models the size and the sign of fiscal policies is fairly well defined: interest rates on the other hand work through the foreign exchange sectors and have an enormous degree of uncertainty attached to them. A conventional deterministic control exercise may well exploit the interest rate effects of the model to a great extent making no allowance for the degree of uncertainty attached to these effects. An objective function of the form of (17) would give much more weight to the more certain policy instruments and would not therefore tend to exploit the less well determined control variable.

3 An Application

To illustrate the technique described above, we give a simple optimal control example using the Bank of England's quarterly forecasting model, which is a large non-linear econometric model which involves around 600 variables. In this example, the objective function was constructed so as to penalise deviations of the current account of the balance of payments from zero, using government current expenditure on goods and services as the control variable. This simplifies both the standard objective function (4) and the extended version (9) by setting $\bar{Y}_i$ to zero for all i. Thus, the final three terms in (9) are set to zero.

The version of the Bank model employed is that deposited with the ESRC Macroeconomic Modelling Bureau in January 1988, and described in Patterson et al (1987), updated by Patterson and Harnett (1988).

The time period used in the problem is 1988 Q1 to 1990 Q4, using the database supplied to the Bureau with the model deposit. The stochastic simulations performed in stage 2 of the procedure are run over 500 replications, using the method of McCarthy (1972) to generate the random shocks to the residuals.

Chart 1 shows the outcome for the target variable. After some variation in the first two quarters, the original optimal control solution u* settles down to track the target value of 0 quite closely. However, the mean of the stochastic simulation around this base E(u*) deviates quite markedly from the target value in both the early and late quarters of the solution period.

The effect of including d_i and Var (Y_i) in the objective function for the second round optimal control solution u' is to produce an 'offset' for the deviation between u* and E(u*), so that the solution path deviates from the target value in the early and late quarters in the opposite direction to E(u*). When a stochastic simulation is performed around u', the mean E(u') track the target solution well throughout the solution period. No further rounds of the procedure were found to be necessry. Table 1 shows root mean squared errors for the four solution paths.

Table 1

Solution	u*	E(u*)	u'	E(u')
RMSE	19.78	70.77	78.45	11.23

Chart 2 shows the time paths for the control variable (in deviations from base). As can be seen, the amendments to this path to generate u' as opposed to u* are quite small in this case compared to the overall deviation from base.

This example illustrates two important points; the deviation caused by the non-linearities in the model is quite important in the sense that while the deterministic optimal trajectory follows the desired path quite closely its expected value does not. Second it is clear that the order of approximation involved in the stochastic control algorithm itself is actually negligible relative to the convergence criteria of the optimal control algorithm.

CHART 1: OUTCOMES FOR CURRENT BALANCE

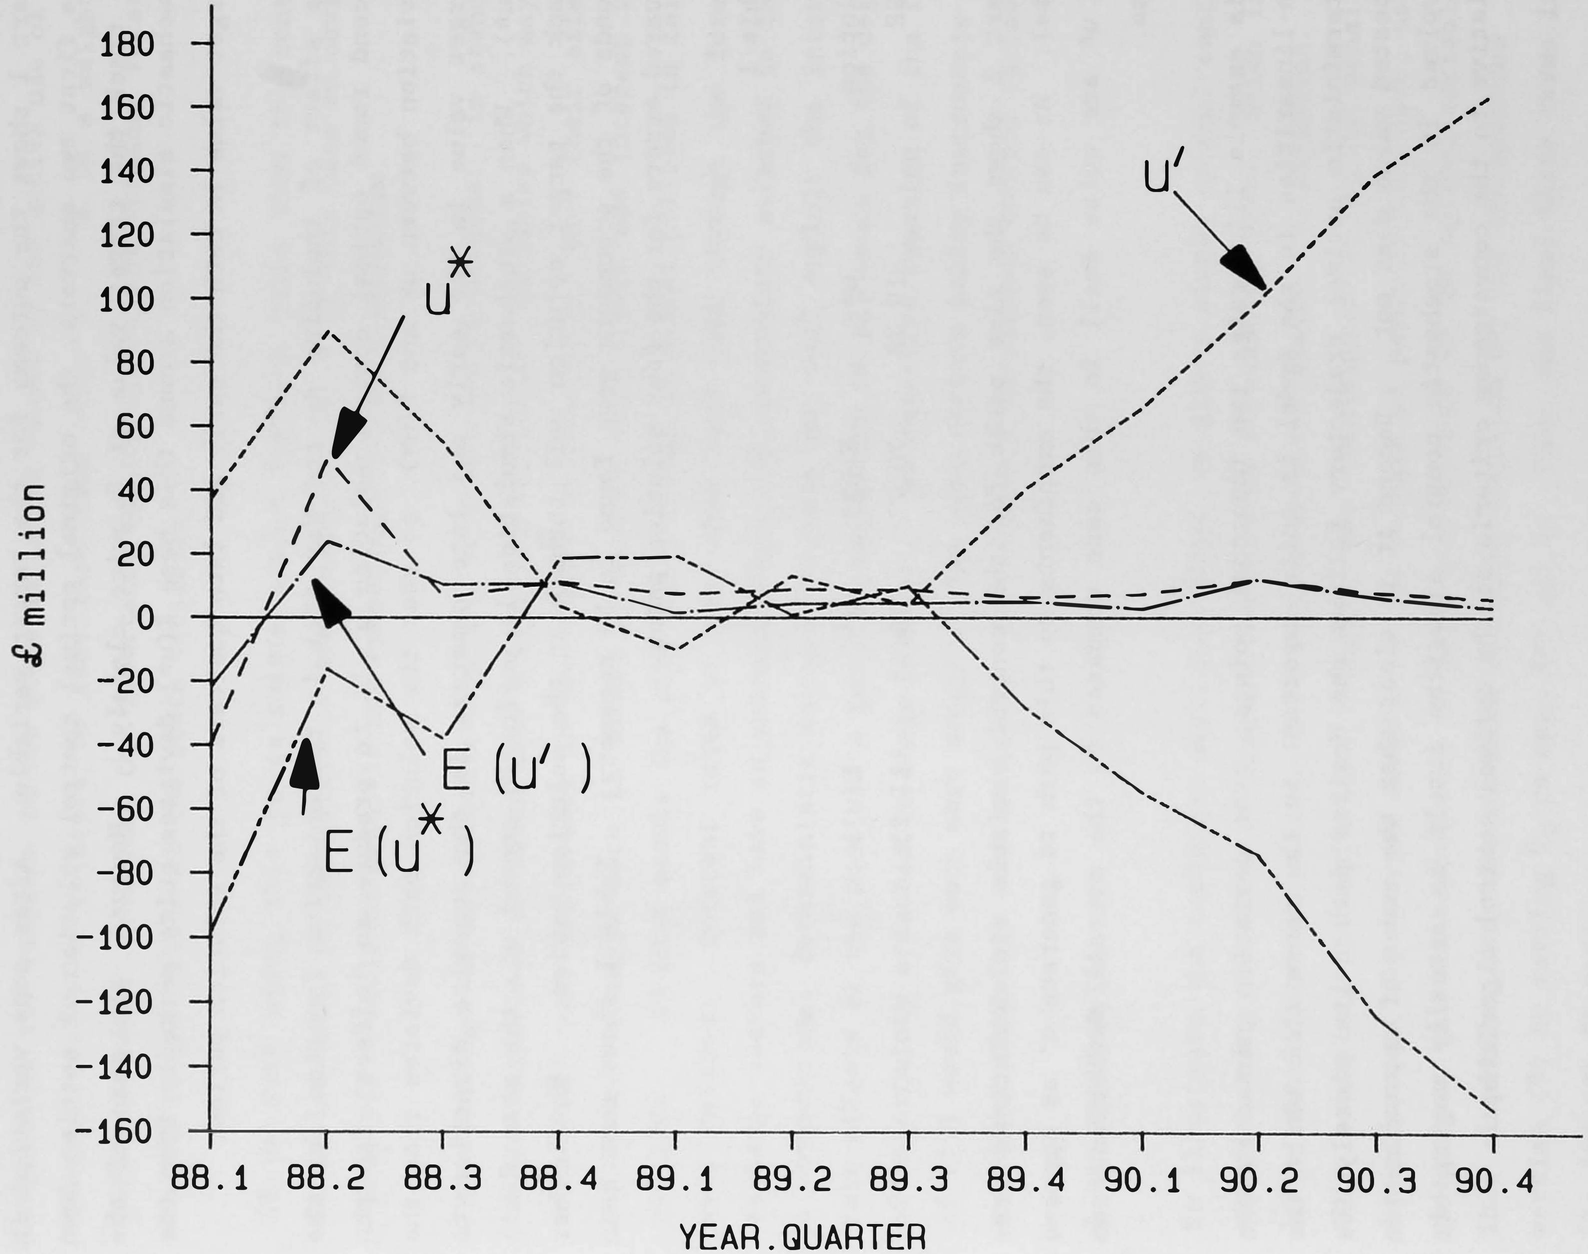

CHART 2: CONTROL VARIABLE OUTCOMES (DEVIATIONS FROM BASE)

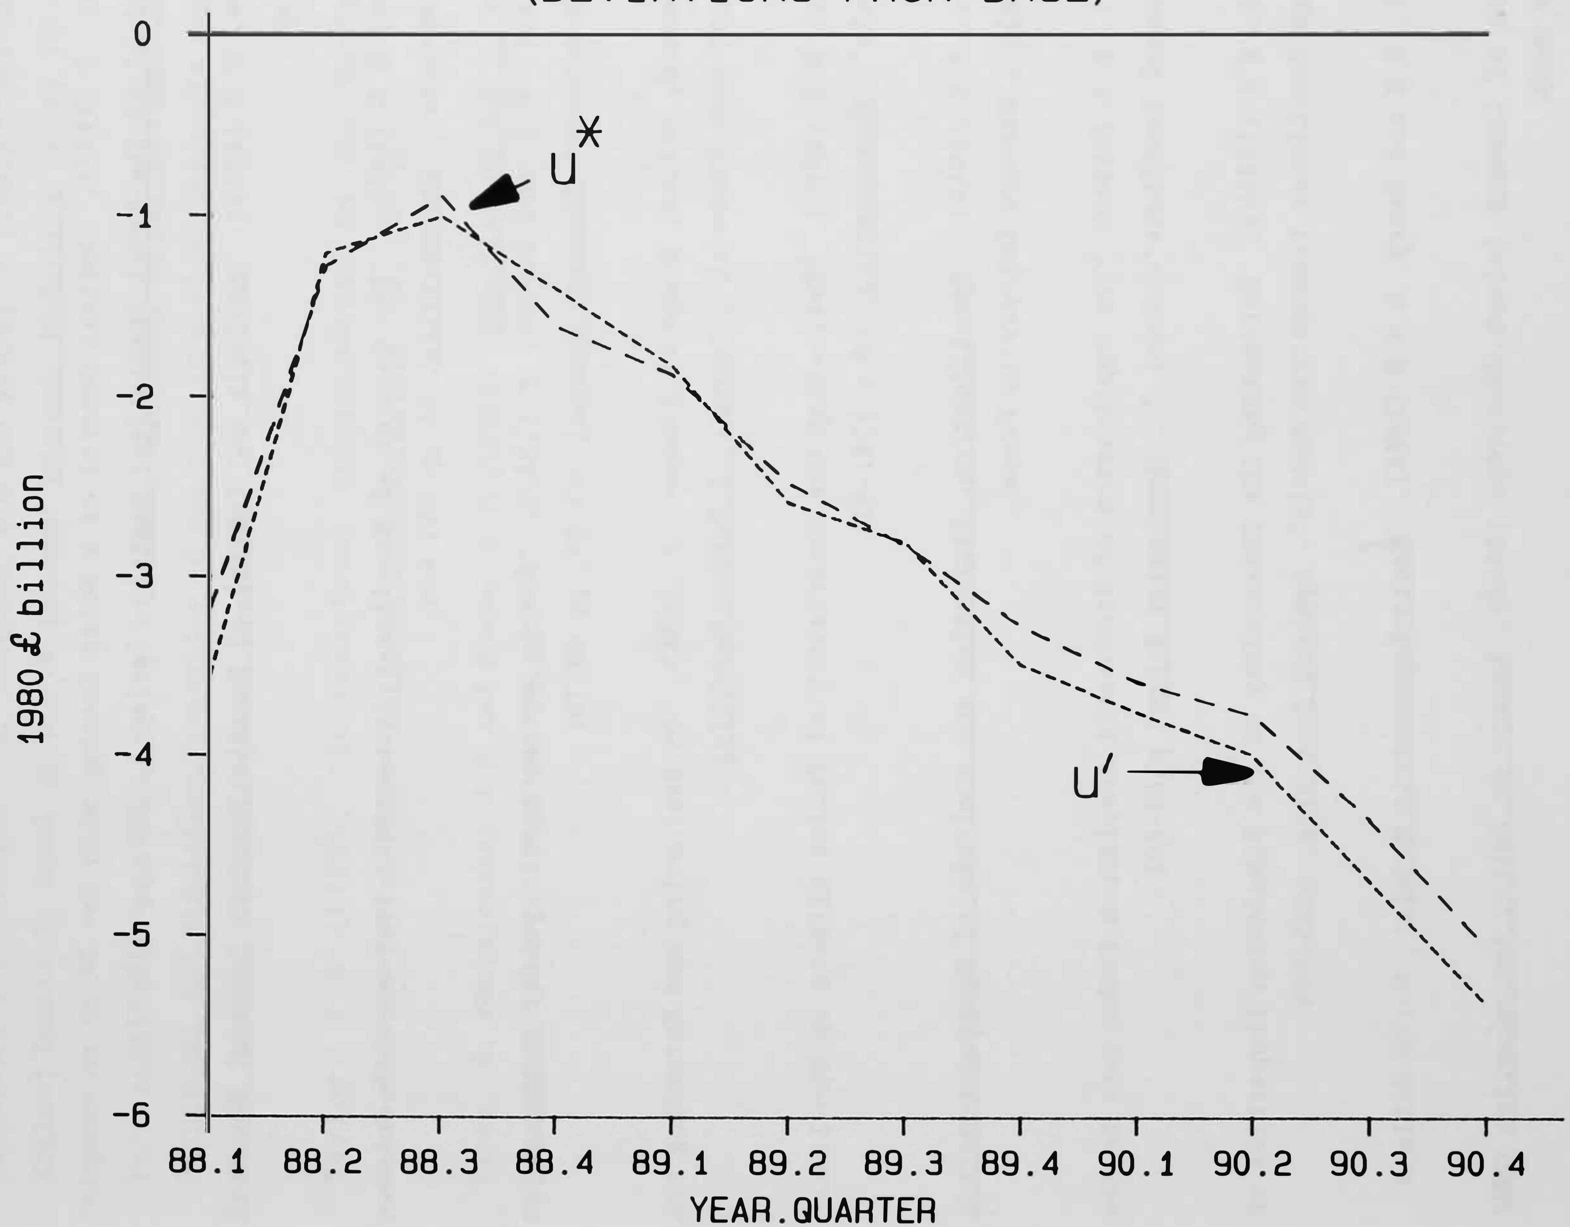

References

Becker, R G, Dwolatzky, B, Karakitsos, E and Rustem, B (1986), 'The Simultaneous Use of Rival Models in Policy Optimisation' Economic Journal 96 No 382 p425-448.

Bray, J (1975), 'Optimal Control of a Noisy Economy with the UK as an example' Journal of the Royal Statistical Society, Series A, 138, pp 339-373.

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

Chow, G C (1976), 'The Control of Non-linear Econometric Systems with Unknown Parameters', Econometrica, 44 pp 685-695.

Currie, D A and Levine, P (1985), 'Simple macroeconomic policy rules in an open economy' 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 (1980), 'Estimating the Uncertainty of Policy Effects in Non-linear Models', Econometrica, 48 p 1381-91.

Fair, R C (1984), 'Specification, Estimation and Analysis of Macroeconometric Models', Harvard University Press.

Hall, S G (1986), 'The Application of Stochastic Simulation Techniques to the National Institute's Model 7', Manchester School p180-201.

Hall, S G (1986b), 'Estimating the Uncertainty of the Simulation Properties of Large Non-Linear Econometric Models', Applied Economics, p985-993.

Hall, S G and Henry, S G B (1988), 'Macro-Econometric Models', North Holland.

House of Commons Select Committee (1978), Report on Policy Optimisation CMND 7148, HMSO.

Kendrick, D (1981), 'Stochastic Control for Economic Models', McGraw-Hill New York.

McCarthy, M D (1972), 'Some Notes on the Generation of Pseudo Structural Errors for Use in Stochastic Simulation Studies' in Hickman (1972) (ed) '<u>Econometric Models of Cyclical Behaviour</u>' Columbia University Press.

Patterson, K D, Harnett, I R, Robinson, G N and Ryding, J (1987), 'The Bank of England Quarterly Model of the UK Economy' <u>Economic Modelling</u>, 4, pp 398-629.

Patterson, K D and Harnett, I R (1988), <u>An Analysis of Changes in the Structural and Simulation Properties of the Bank of England Quarterly Model of the UK Economy</u>, Bank of England, mimeo.

Taylor, J B (1985), 'International Co-ordination in the Design of Macroeconomic Policy Rules' <u>European Economic Review</u>, 28, p53-82.

Vines, D, Maciejowski, J M and Meade, J E (1983), '<u>Stagflation Vol 2, Demand Management</u>' George Allen and Unwin.

Bank of England Discussion Papers

	Title	Author
1-5, 8, 11-14, 16-17, 19-22	*These papers are now out of print, but photocopies can be obtained from University Microfilms International (see below).*	
6	'Real' national saving and its sectoral composition	C T Taylor A R Threadgold
7	The direction of causality between the exchangte rate, prices and money	C A Enoch
9	The sterling/dollar rate in the floating rate period: the role of money, prices and intervention	I D Saville
10	Bank lending and the money supply	B J Moore A R Threadgold
15	Influences on the profitability of twenty-two industrial sectors	N P Williams
18	Two studies of commodity price behaviour: Interrelationships between commodity prices Short-run pricing behaviour in commodity markets	 Mrs J L Hedges C A Enoch
23	A model of the building society sector	J B Wilcox
24	The importance of interest rates in five macroeocnomic models	W W Easton
25	The effects of stamp duty on equity transactions and prices in the UK Stock Exchange	Mrs P D Jackson A T O'Donnell
26	An empirical model of company short-term financial decisions: evidence from company accounts data	Ms G Chowdhury C J Green D K Miles
27	Employment creation in the US and UK: an econometric comparison	I M Michael R A Urwin
28	An empirical model of companies' debt and dividend decisions: evidence from company accounts data	Ms G Chowdhury D M Miles
29	Expectations, risk and uncertainty in the foreign exchange market: some results based on survey data	M P Taylor
30	A model of UK non-oil ICCs' direct investment	E J Pentecost
31	What has the European Monetary System achieved?	M P Taylor M J Artis
32	The demographics of housing demand; household formations and the growth of owner-occupation	M J Dicks
33	Measuring the risk of financial institutions' portfolios: some suggestions for alternative techniques using stock prices	S G F Hall D K Miles
34	An error correction model of US consumption expenditure	I R Harnett

Papers presented to the Panel of Academic Consultants[(a)]

	Title	Author
8	International monetary arrangements the limits to planning*	P M Oppenheimer
9	Institutions in the financial markets: questions, and some tentative answers*	M V Posner
10	The arguments for and against protectionism*	M Fg Scott The Hon W A H Godley
14	The usefulness of macroeconomic models*	Prof W H Buiter T F Cripps Prof Angus Deaton Prof A P L Minford M V Posner
15	Factors underlying the recent recession*	G D N Worswick Dr A Budd
17	Why do forecasts differ?*	Prof M J Artis
19	Bank lending, monetary control and funding policy*	Prof A D Bain
20	The economics of pension arrangements*	Prof Harold Rose J A Kay
22	Monetary trends in the United Kingdom	Prof A J Brown Prof D F Hendry and N R Ericsson
23	The UK economic recovery in the 1930s	G D N Worswick P N Sedgwick Prof Michael Beenstock Dr Forrest Capie Prof Brian Griffiths
24	Employment, real wages and unemployment in the United Kingdom*	Prof J R Sargent Sir Bryan Hopkin

Technical Series

	Title	Author
1-11	*These papers are now out of print, but photocopies can be obtained from University Microfilms International (see below).*	
12	The development of expectations generating schemes which are asymptotically rational	K D Patterson
13	The arch model as applied to the study of international asset market volatility	R R Dickens
14	Modelling the UK economy in a stock-flow consistent manner	E P Davis
15	International comparison of asset market volatility: a further application of the ARCH model	R R Dickens
16	A three sector model of earning behaviour	D J Mackie
17	Integrated balance sheet and flow accounts for insurance companies and pension funds	Raymond Crossley
18	Optimal control of stochastic non-linear models	S G Hall I R Harnett M J Stephenson

* These papers are no longer available from the Bank, but photocopies can be obtained from University Microfilms International, at White Swan House, Godstone, Surrey RH9 8LW.

(a) Other papers in this series were not distributed.