Simulation

1977

University of California, Davis
Department of Agricultural Economics

CONTROL MODEL SPECIFICATION AND SOLUTION TECHNIQUES*

Richard Howitt**

*Invited paper for presentation at the AAEA-WAEA meeting in San Diego, California, July 31-August 3, 1977.

** Assistant Professor of Agricultural Economics and Assistant Agricultural Economist in the Experiment Station and on the Giannini Foundation, University of California, Davis.

Richard Howitt - Stochastic Control Symposium Notes

CONTROL MODEL SPECIFICATION AND SOLUTION TECHNIQUES

Introduction

The early interest and development of dynamic programming in agricultural economics has not been generally adopted by the mainstream of quantitative agricultural economics, as have the essentially static optimal policy analysis techniques of linear and quadratic programming, econometric multiplier analysis, and suboptimal simulation. Why? The most likely answer is that the problems with the most quantitative appeal during this decade were simply not solvable by the numerical search dynamic programming techniques currently available. The emphasis in quantitative modeling was, and still is, on multi-variable systems and their stochastic properties. Except for cases of remarkable aggregation or simple physical systems, the aptly named "curse of dimensionality" prevented the achievement of solutions even under deterministic assumptions.

My main point today is that for the class of passive adaptive models and policies, the curse, although not entirely lifted has lost most of its effect due to alternative specifications and advances in solution techniques. Currently, the specification of the control problem which an applied researcher would face in deriving optimal policy actions from a conventionally estimated simultaneous system is readily solvable as an L.Q.G. problem. Given the solution complexities and inconclusive advantages of the active adaptive models surveyed by Rausser, the passive adaptive models appear to be a satisfactory compromise between model complexity and the truly optimal solution. Indeed, from the results of the illustrative analytical example presented by Rausser in section 3, one can conclude that investment in passive experimental information is more fruitful where there is either considerable doubt or precise knowledge about the controlled system. Under these conditions, the passive adaptive approach to control will have a similar performance as the active closed loop approaches.

Initially, I will discuss two classes of solution techniques and some examples of their use with which I am familiar. However, as cautionary points, there are two significant problem areas that economists will encounter in the application of passive adaptive controls. First, the need for inequality constraints on states and controls for all initial trajectory models and some tracking models. Second, the problems encountered in specifying objective and penalty functions as cardinal measures of benefits.

I. Two Solution Approaches

Polak in a survey of computational methods in optimal control [1973] characterized all methods utilizing Bellman's "principle of optimality" to separate the control horizon as Feedback Solutions. Referring to the conventional form of dynamic programming with its search over the numerical solutions at feasible modes, he concludes:

> "In spite of these attempts to make dynamic programming a practical algorithm, it cannot be considered to be particularly successful in this role, ... The main use of dynamic programming has been as a conceptual tool, particularly, before the Pontryagin maximum principle was well understood" [1973].

This unfortunate dimensionality result can be traced to the prevalent analytical intractability of the partial difference equation of the intertemporal rate of change of the objective functional (Bellman's equation $-\frac{\partial J^*}{\partial t}$).

Linear Quadratic Gaussian Solutions

Application of Pontryagin's maximum principle to the dynamic programming two point boundary value problem by Kalman [1960] and others yields an analytic solution to the dynamic programming problem under specific conditions. The analytic solution is obtained by the introduction of the costate or "cost to go" matrix; which determines the functional relationship between the vector of states and costates at any time. Under given conditions, the costate matrix is expressed as the solution to a discrete matrix Riccati equation that can be solved recursively backwards.

As the name suggests

(1) the control system must be linear

(2) the error terms on the equation of motion must have Gaussian (Normal) distributions.

This control problem specification has many Convenient Properties

i. It is not restricted by dimensionality, but only by the computer's matrix inversion capacity, and is cheap to run average sized econometric models and planning horizons.

ii. It has a convenient economic structure for many models in that

(a) the linear and Gaussian form are the way in which the reduced forms of many simultaneous equation models are calculated.

(b) the quadratic form of objective function has had defendable justification as a direct measure of well being (Takayama and Judge [1964], Goreaux, Manne, et al. [1973]). As an approximation to more complex objective functions, the quadratic form has been shown to be comparatively robust (Zellner and Giesel [1970]).

iii. The L.Q.G. approach yields analytical solutions for both the Certainty Equivalent case and the Stationary Stochastic Case (Chow [1975] and Aoki [1976]).

iv. Under the certainty equivalent specification, the separation principle (Joseph and Tou [1961]) allows independent solution of the optimal policy solution and the optimal conditional state estimator. This estimator is often obtained from a Kalman filter which, symmetrically, is also solved by the L.Q.G. procedure.

v. Where the system equations estimated are not linear, the L.Q.G. solution can be used to track a target trajectory with linearized equations of motion. The quadratic form of the criterion function minimizes the errors due to the linearization (Athans [1972]).

vi. The optimal solution yields the decision relevant (Marschak [1971]) value of passive information on the initial condition estimates, the equations of motion and the observation error, if specified.

Practical Use

Use of the L.Q.G. approach at Davis in the last four years has been over initial trajectory, tracking and filtering problems (Dixon [1976], Howitt [1975]) applied to optimal natural resource policy and externality policy. The routines have been written to be problem specific using subroutines from a system developed by White and Lee [1971]. A more general L.Q.G. package is under development.

A model much used in published stochastic control results is by Chow [1967]. The L.Q.G. routine on a Burroughs 6700 ran this nine equation model for a ten period horizon using 5.0 seconds C.P.U. at $1-30 for the certainty equivalent solution, and 6.0 seconds C.P.U. at $1-65 for the stationary stochastic solution.

Nonlinear Programming Solutions to Control Problems

The deterministic nonlinear control problem can be formulated as a nonlinear programming problem by stacking the equations of motion for each time period as constraints (Pindyk [1973]). Understandably this leads a constraint set of large dimensionality. A more elegant and efficient approach is advanced by Canon, Cullum, and Polak [1970], Aoki [1971], Fair [1974] and others, which solve the nonlinear difference equations for their transition matrix. A state vector of endogenous variables at any time period (t) can thus be expressed as a function of the initial state conditions and the trajectory of controls to (t). Since the equation of motion constraints are now nested via the transition matrix in the objective function, the optimal solution, assuming no state or control constraints, may be obtained by unconstrained gradient approaches.

Algorithms

The optimization literature abounds with alternative numerical nonlinear optimization routines. My practical experience has so far been limited to the variable Metric Davidon Fletcher Powell algorithm that combines some of the advantages of both gradient and Newton-Raphson approaches (National Bureau of Economic Research [1976]). Comparisons by Fair [1974] indicate the D.F.P. algorithm to have good properties for large models.

Tests with the same Chow [1967] model on the D.F.P. algorithm on an IBM 360 solve the certainty equivalent solution in 11 seconds C.P.U. and $ 2.50 .

Convenient Properties

(i) Direct numerical solution of nonlinear problems without linearization.

(ii) Derivation of the initial target trajectory for the stochastic linearized L.Q.G. method to track.

(iii) The ability to efficiently handle inequality constraints on states and controls through additional constraints in the Pindyk formulation, or penalty functions on the Canon Cullum Polak formulation.

(iv) An open loop feedback approach to the solution of the stochastic nonlinear problem has been suggested by Fair [1974]. Essentially it involves recomputation of a deterministic nonlinear control problem at each time period using updated certainty equivalent estimates.

Constraints on State and Control Variables

In the scramble to adopt alluring control models from engineering applications to macroeconomic, and more recently microeconomic problems, the thorny problem of inequality constraints on state and control variables has largely been ignored. The two text books currently published on economic applications of stochastic control Chow [1975] and Aoki [1976] make only passing reference to the problem. An explanation of this gap in development, may be that the analytic stochastic models originally developed in the engineering literature were largely tracking models. The objective functions on these models were usually concerned only with the terminal trajectory state such as a moon landing or obliteration of a city by a missile (Polak [1973]). For these problems, fluctuations of the trajectory are of no consequence if the final state is achieved in a minimum time or with minimum fuel. Initial applications of this type of model to macroeconomic teaching problems invariably operated without the need for control or state constraints. Microeconomic applications, however, routinely require derivation of the optimal initial trajectory of states; and for market pricing, inventory, investment and resource allocation decisions this initial trajectory is the crux of the control problem.

The microeconomist is faced with three principle types of inequality constraints.

(a) Feasible region constraints on states and controls, the simplest example being nonnegativity constraints on market prices and production quantities.

(b) Institutional constraints on the range over which policy control actions can be set. Most control specifications define the controls as belonging to a feasible set, but rarely say how the solution procedure keeps them there.

(c) Constraints on the range of state variable values over which the linear state system estimates are viable approximations. Where the equations of motion are linearized by a Taylor series expansion, significant departures of the state variable values from those about which the system was linearized will lead to errors. For equations of motion that are directly estimated in the linear form, the estimates may be only valid over a given range.

The problem in imposing state or control constraints on the class of analytical feedback solution techniques, is that global optima in controls are needed to solve for the optimal cost to go relation. Using Bellman's principle, these relations have to be solved backwards in time or "offline;" while the values of states or controls which may or may not be binding, are solved forwards in real time. Clearly, if a state or control becomes constrained, when the marginal value of previous actions is predicated on an unconstrained solution, the prior decisions are suboptimal. To paraphrase - "If I had known that gas rationing was going to be imposed, I wouldn't have bought a Dodge Charger."

The degree of suboptimality of the constrained solution will depend on (1) the length of the control horizon; (2) the timing of the binding constraints; and (3) the degree of control constraint.

Methods of Inequality Constraints

1. Analytical Feedback Solutions. An "ad hoc" suboptimal method is to simply institute a subroutine in the "forward" loop that solves for optimal state and control values. The subroutine checks the values against the constraint boundaries at each time period, and substitutes in the boundary value if the constraint is binding. The degree of suboptimality of this method is specific to different solutions, but could be ascertained by numerical comparison with programming solution approaches.

2. Programming Solutions. Introduction of inequality constraints in the Pindyk [1973] programming formulation is a simple addition to the large set of equality constraints. In the Canon, Cullum, Polak approach, inequality constraints are implemented by penalty functions, which preserves the convenient feature of the unconstrained gradient search. However, this method can lead to constraint violations if the penalty functions are inappropriately chosen since the penalty function form is usually quadratic. Kim, Goreaux, and Kendrick [1975] report on the successful use of gradient modification to satisfy constraints.

Users of ad hoc constraints on feedlock models should consider testing the degree of suboptimality of their results with deterministic programming solutions.

The Economic Meaning of Objective and Penalty Functions

The origins of analytic stochastic control in physical tracking problems has often influenced the specification of the objective function for initial trajectory problems. Athans [1972] in his survey of the L.Q.G. approach has a section on the "Selection of Weighting Matrices." The microeconomist, for whom the value of the objective function signifies a cardinal level of well-being does not enjoy such selection alternatives. The basis of microeconomic objectives are likely to be in demand, yield or damage functions; and model specifications must satisfy economic as well as optimization constraints. Principle points are:

(i) Symmetry in quadratic objective functions can lead to absurd results in initial trajectory models if state variables are unconstrained. In some cases this can be avoided by normalizing the quadratic function such that the minimum (maximum) is never approached. Tracking models have similar problems in that the costs of over-achieving are rarely the same as under-achieving. Athans [1972] has suggested a procedure for deriving weights based on the second derivative of the Hamiltonian along the initial trajectory. Gradient solution approaches do not encounter similar problems due to the ability to specify asymmetric or truncated objective functions; and constraints on state variables.

(ii) Terminal Period Values. The effect of changes in the terminal period objective function value depends on the length of the planning horizon. For empirical microeconomic models the policy horizon will probably be short, depending on institutions and politics. The terminal period of a policy model does not imply the termination or scrapping of the investment, as is the usual engineering interpretation of the terminal value function. For this type of medium horizon policy model, a satisfactory compromise is achieved by specifying the terminal value on the same basis as the other time periods. This decision implicitly assumes that an alternative policy planning model will post-date the current one before the terminal time is reached.

(iii) Penalty Functions are suggested by Fair [1974], Chow [1975] and others as a method of implementing inequality constraints. Of the two alternative specifications, the "exterior method" promises the least chance of distorting the cardinal value of the objective function. However, where the linear approximation of the function is poor, or its weighting low, the intrusion of penalty function values into the optimal objective function occurs. Penalty functions have the advantage of exactly representing nonlinear constraints if they are required.

(iv) Policy Transaction Costs. An advantage of the objective function generally specified for control models is that it specifically includes the cost of controls. That is, the social transaction costs of implementing policy actions. These costs may be hard to empiricise, but where the state related parameters measure consumer's surplus, the transaction costs of policy actions should also be included to yield an optimal social policy.

Conclusions

The curse of dimensionality which has prevented the routine use of optimal dynamic policy analysis for simultaneous equation models is being lifted. If a researcher can specify an objective function or policy weightings on the endogenous variables in a linear or quadratic form, they can obtain optimal policies under certainty equivalent and stationary stochastic information structures quickly and cheaply. The main problems I foresee in the policy use of stochastic control in microeconomics are in constraining states and controls, and specifying meaningful objective functions.

Currently, a promising compromise of solution techniques is to use a gradient programming approach to solve the deterministic and constrained initial trajectory. Then use the L.Q.G. stationary stochastic approach to examine the optimal policy under exogenous stochastic conditions, and the resulting passive information values.

bf
7-29-77

FIGURE 1

Simulation Versus Control Results

Chow's [1967] Model is simulated to give the target endogenous variable values - YD - (change in income) is the only value plotted here to show the precision of tracking

```
   200.    220.    240.    260.    280.    300.    320.    340.    360.    380.    400.    420.
     !=======!=======!=======!=======!=======!=======!=======!=======!=======!=======!=======!
1950 !  2    !       !       !       !       !       !       !       !       !       !       !
     !       !       !       2       !       !       !       !       !       !       !       !
     !       !       !       !       !       !       !AB     !       !       !       !       !
     !       !       !       !     2 !       !       !       !       !       !       !       !
     !       !       !       !     AB!       !       !       !       !       !       !       !
1955 ! - - - ! - - - ! - - - ! - - - ! - - - ! - - -2! - - - ! - - - ! - - - ! - - - ! - - - !
     !       !       !       !       !     2 !       !       !       !       !       !       !
     !       !       !       !       !       !       !      2!       !       !       !       !
     !       !       !       !       !       !       !       !       !      2!       !       !
     !       !       !       !       !       !       !       !       !       !      2!       !
     !=======!=======!=======!=======!=======!=======!=======!=======!=======!=======!=======!
   200.    220.    240.    260.    280.    300.    320.    340.    360.    380.    400.    420.

*************************************LEGEND*************************************************

TIME BOUNDS: 1950 TO 1959

SYMBOL SCALE NAME
  A     #1    RES1_YD  - Control Results
  B     #1    YD_CHOW  - Simulation Results

********************************************************************************************
```

FIGURE 2

Policy variable G (Change in Government Spending) - control action is lower or equal to simulated action.

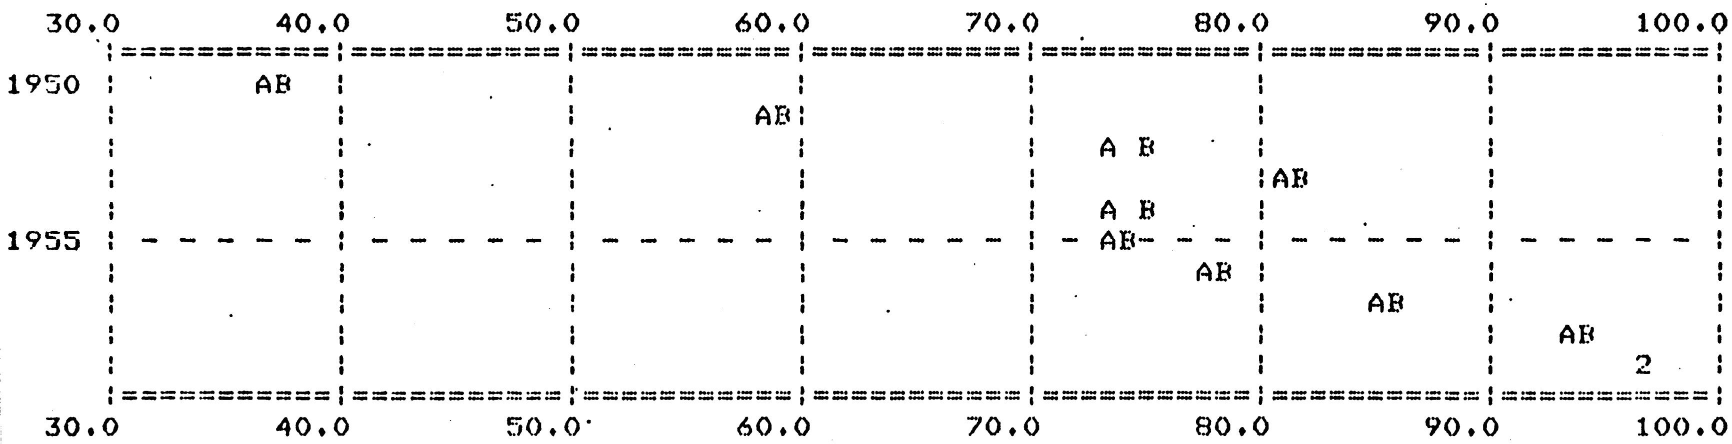

************************************LEGEND************************************

TIME BOUNDS: 1950 TO 1959

SYMBOL SCALE NAME
A #1 RES1_G - Control Policy Action
B #1 G_CHOW - Simulation Policy Action

****