OPTIMAL CONTROL OF GENERAL EQUILIBRIUM MODELS

by

John Derpanopoulos*

Invited Paper for the American Agricultural Economics Association
Summer Meeting, Reno, Nevada, July 27-30, 1986.

OPTIMAL CONTROL OF GENERAL EQUILIBRIUM MODELS

Introduction

This paper discusses a methodology for applying optimal control to dynamic general equilibrium models (GEMs). Given a welfare criterion, optimal government policies over a fixed horizon can be ascertained using the classic quadratic linear approximation technique whereby the model is linearized and solved using an extension of Johansen's method. Intraperiod updates of the models' state variables, as well as inter-period updates, are performed automatically by means of a prespecified default value. This reduces the approximation error which can be further reduced and made even arbitrarily small by suitable reductions in the default value.

The use of a self-updating Johansen system rather than a more formal solution algorithm greatly economizes on the computation time required for final convergence which, even so, entails some difficulties arising from bang-bang behavior. For this reason, welfare functions monotonic in consumption or income were found very difficult to deal with regardless of terminal constraints, thus leading to the employment of outright quadratic or "target" objective functions. Similarly, the use of more than one policy instrument as control variable made final convergence more difficult which, in the case of a single control, is achievable within 8 to 12 iterations.

Solution Techniques

In their most generalized form, GEMs can be expressed as

(1) $$G_t(Y_t, X_t, Z_t) = 0$$

where

G_t = vector of valued functions representing the equations of the model (at time t)

Y_t = vector of endogenous variables (including prices)

X_t = vector of policy or control instruments

Z_t = vector of exogenous and lagged endogenous variables.

In obtaining a first-order approximation of this system, the total derivative of (1) is

(2) $$G_t^Y\, dY_t + G_t^X\, dX_t + G_t^Z\, dZ_t = 0 \quad \text{or}$$

$$\tilde{A}_t\, dY_t = \tilde{B}_t\, dX_t + \tilde{C}_t\, dZ_t$$

and solving for dY_t

$$dY_t = (\tilde{A}_t)^{-1}\, (\tilde{B}_t\, dX_t + \tilde{C}\, dZ_t)$$

where

$$\tilde{A}_t' = -G_t^Y$$

$$\tilde{B}_t' = G_t^X$$

$$\tilde{C}_t' = G_t^Z.$$

Alternatively, one can rewrite (2) in Johansen form as

(3) $$A_t\, y_t = B_t\, x_t + C_t\, z_t$$

where small letters denote rates of change, i. e.,

$$y_t = \left[\frac{1}{Y_t}\right] dY_t \qquad \text{and} \qquad A_t = \tilde{A}_t \left[Y_t\right]$$

$$x_t = \left[\frac{1}{X_t}\right] dX_t \qquad\qquad B_t = \tilde{B}_t \left[X_t\right]$$

$$z_t = \left[\frac{1}{Z_t}\right] dZ_t \qquad\qquad C_t = \tilde{C}_t \left[Z_t\right]$$

with brackets denoting square-diagonal matrices.

The advantage of the rate of change set up over a simple first-order Taylor expansion lies in the resultant matrices of derivatives A_t, B_t, and C_t. In (3), the elements of these matrices are either constant (pertaining to log linear equations) or simple functions of Y_t, X_t, and Z_t, respectively, such that one can perform simple updates as follows:[1]

$$A_{t+1} = A_t[1 + y_t]$$

$$B_{t+1} = B_t[1 + x_t]$$

$$C_{t+1} = C_t[1 + z_t].$$

This is in contrast to Johansen, who treats the above matrices as constant. Furthermore, one can perform intraperiod updates for the sake of higher accuracy. Since it is not practically feasible to increase the number of such updates ad infinitum (something that would duplicate the actual nonlinear system), we opted for some sort of consistency rule that would indicate the proper number of "period segmentations." Dixon et al., who use a segmented version of Johansen's technique for purposes of policy simulation, link this number to the percentage change of the exogenous variables in each period. This, however, lacks consistency since a given change in exogenous variables will have varying effects according to the state of the economy.

The rule which we used (Derpanopoulos) calls for an update whenever the largest absolute value deviation amongst all the variables exceeds a certain default value, M. Assume that the largest such deviation is 2.3M. Then the original vectors of control and exogenous changes (x_t and z_t) are applied in three segments before solving for the complete y_t in (3). First, (1/2.3) x_t and (1/2.3) z_t are used to solve for part of y_t followed by an update of A_t, B_t, and C_t; then the process is repeated a second time, leaving a remainder of (.3/2.3) of exogenous and policy variable changes to be applied the third (and final) time.

Different default values were tried with a 12.5 percent default value proving sufficiently accurate and at the same time economical for the purpose of control. The average error resulting from this operation was in the order of 0.4 percent with individual errors varying from 0 to 1.57 percent in the worst case.

Optimal Control

Having set up an efficient and economical framework for solving GEMs, one can proceed to derive optimal government policies over a fixed time horizon for given preference orderings as expressed in a social utility function, W, where

(4) $$W = \sum_{t=1}^{T} \left(\frac{1}{1 + \rho_t}\right) W_t \; (Y_t, X_t),$$

and where W = total discounted utility and ρ_t = discount rate for period t. The optimization algorithm used is the quadratic linear approximation technique (Athans) whereby a second-order expansion of the welfare function (4), together with a first-order expansion of the system equations (1), are performed around a base path for the economy. This path could be a simple projection of the current state of affairs together with anticipated levels for policy and exogenous variables over the horizon of the planning exercise. Thus, given $(X_t^0, Z_t^0)_{t=1}^{T}$, one could, through the use of the solution algorithm discussed above, solve for $(Y_t^0)_{t=1}^{T}$. In doing so, one has also obtained the necessary first-order expansions around the base path from equation (3) with the derivative matrices A_t, B_t, and C_t evaluated at each interval of the base path for t = 1, ..., T. One, therefore, is left with only the task of performing a quadratic approximation of (4) around the base path, i.e.,

$$\underset{(Y_1^0, \; \ldots, \; Y_T^0)}{dW_t} = \left(\frac{1}{1 + \rho_t}\right) \left\{ W_t^Y \, dY_t + dW_t^X \, dX_t \right.$$

$$\left. + \frac{1}{2}(dY_t \;\; dX_t) \begin{pmatrix} W_t^{YY} & W_t^{YX} \\ W_t^{XY} & W_t^{XX} \end{pmatrix} \begin{pmatrix} dY_t \\ dX_t \end{pmatrix} \right\}.$$

Making the latter compatible with the Johansen framework, the control problem can thus be expressed as:

$$\underset{(x_1, \ldots, x_T)}{\text{Maximize}} \quad dW = \sum_{t=1}^{T} \left(\frac{1}{1+\rho_t}\right) W_t^y \, y_t + W_t^x \, x_t \tag{5}$$

$$+ \frac{1}{2} (y_t \; x_t) \begin{pmatrix} W_t^{yy} & W_t^{yx} \\ W_t^{xy} & W_t^{xx} \end{pmatrix} \begin{pmatrix} y_t \\ x_t \end{pmatrix} \Bigg\}$$

subject to

$$y_t = A_t^{-1}(B_t \, x_t + C_t \, z_t)$$

where

$$y_t = \frac{Y_t - Y_t^0}{Y_t^0}$$

$$x_t = \frac{X_t - X_t^0}{X_t^0}$$

$$z_t = \frac{Z_t - Z_t^0}{Z_t^0}$$

$$W_t^y = W_t^Y \, [Y_t^0]$$

$$W_t^x = W_t^X \, [X_t^0]$$

$$W_t^{yy} = [Y_t^0] \, W_t^{YY} \, [Y_t^0]$$

$$W_t^{xx} = [X_t^0] \, W_t^{XX} \, [X_t^0]$$

$$W_t^{xy} = [X_t^0] \, W_t^{XY} \, [Y_t^0] \; .$$

This problem can be approached using dynamic programming where one solves for a set of closed-loop feedback rules $(G_t, g_t)_{t=1}^{T}$. These indicate the optimal deviation in the control variables for each period, x_t, conditional on the changes in the lagged endogenous and exogenous variables of that period, z_t.

Working backward from the optimality conditions of the last period,

$$x_T = G_T z_T + g_T$$

where

$$G_T = \phi_T \{B_T \left(A_T^{-1}\right)' W_T^{yy} A_T C_T\}$$

(6) $$g_T = \phi_T \{\left(W_T^{x}\right)' + B_T\left(A_T^{-1}\right)' W_t^{y}\}$$

$$\phi_T = \{B_T\left(A_T^{-1}\right)' W_T^{yy} A_T^{-1} B_T + W_T^{xx}\}^{-1},$$

one must next solve the Riccati equations for the previous period,

$$h_{T-1} = W_{T-1}^{y} + \left\{g_T'\left(A_{T-2}\right)^{-1} W_T^{yy} + W_T^{y}\right\}\left\{A_T^{-1}(B_T G_T + C_T)\right\} + W_T^{x} G_T,$$

$$H_{T-1} = W_{T-1}^{yy} + \left\{A_{T-1}^{-1}(B_T G_T + C_T)\right\}' W_T^{yy}\left\{A_{T-1}^{-1}(B_T G_T + C_T)\right\} + G_T' W_T^{xx} G_T.$$

H_{T-1} and $h_{T=1}$ can be used in place of W_T^{y} and W_T^{yy}, respectively, in (6) to compute G_{T-1} and g_{T-1} which can then be used for H_{T-2} and h_{T-2} and so on until one finally arrives at the optimal feedback rules for the first period. This optimality, however, pertains to the linear quadratic version (or approximation) of the actual problem. It is customary practice to use the derived feedback rules and the associated control deviations as indicative of the direction toward which welfare is increased and not of the actual location of the optimum. This leads to the use of step sizes whereby a certain fraction of the indicated change in control variables is actually enforced at each time period. Thus, if the deviation indicated for period t is

$$x_t = G_t z_t + g_t = \xi,$$

then, at the simulation stage,

$$x_t = a_t \xi$$

where a_t is the step size for period t. Under the new deviations of the policy variables, simulation of the GEM for the T periods will give rise to a new path for the economy,

$$\left\{Y_t^1, X_t^1\right\}_{t=1}^{T},$$

where

$$Y_t^1 = Y_t^0 [1 + y_t]$$

$$X_t^1 = X_t^0 [1 + x_t] \qquad 1 \leq t \leq T.$$

Normally, this new path should be closer to the optimal than the first (i. e., should increase the value of the objective function). It is not unusual, however, if, instead, it gives rise to a decrease in welfare. This will be the case if the step size is too large leading to "overshooting." Unfortunately, there is no rule for establishing the optimum step size although conditions do exist for placing upward bounds on the step size which guarantee convergence (Ortega and Rheinbolt;

Kailath). Such bounds are usually on the conservative side and, in practice, much larger step sizes can do the job more efficiently. We discovered that step sizes of either 1/4 or 1/8 worked quite well--at least for the early iterations. As one approaches an optimum, these step sizes have to be reduced for the sake of smoother convergence.

Finally, for convergence we required that the welfare function increase by less than 1 percent for two consecutive iterations. In the majority of cases, such a situation was achieved within 8 to 12 iterations.[2]

The Welfare Criterion

Regarding the functional form of the welfare function, one can distinguish between two general types: monotonic functions and quadratic functions. Monotonic functions are logically more sound, adhering to principles of nonsatiation with respect to welfare-bearing arguments (whether these are consumption, income, capital stock, etc.). In practice, we encountered considerable difficulty in inplementing this type of welfare criterion. Attempts were made in this direction following the traditional method of second-order approximation of these functions. Under this approach, the idea is to guide the system in the direction of the steepest ascent, i. e., the direction which increases welfare most rapidly. One problem is that gradients are not constant from iteration to iteration but change depending on the point of approximation. This results in a somewhat erratic control policy which inhibits smooth convergence to an optimum (to the extent that the latter is not well defined since it changes with each iteration).

A more serious problem with monotonic welfare functions is that optimal solutions are not guaranteed to be as appealing as the principle of monotonicity upon which they are built. Bang-bang solutions were often encountered, characterized by greatly reduced income in certain periods for the sake of higher income in other periods. Under monotonic welfare functions, where the goal is to maximize total discounted income (or consumption), bang-bang solutions will arise from the dynamics of the particular model.

GEMs generally have few lagged dependent variables. In our model, these were the capital stocks of each sector, the price level, and the rate of money increase. All remaining variables which did not appear in lagged form were associated with unitary eigenvalues. This, in itself, is sufficient to give the system oscillatory tendencies which, in certain cases, can diverge in an outward direction with time. Optimal control will exploit this characteristic unless there is a stabilizing force exerted by the welfare function. Monotonic functions do not have such a stabilizing effect since the goal of maximum income or consumption might be achievable through alternations of high and low income periods.

These difficulties led us to employ quadratic or target-oriented welfare functions. These functions have the chief shortcoming of being symmetric around their particular targets and, hence, of penalizing equally outcomes which are on different sides of a target. One way to get around this difficulty is to ensure that the targets are never superseded. There is a trade-off, however, between the infeasibility of the targets on the one hand and the smoothness of the results on the other. By setting the targets too far, one will induce the same sort of bang-bang behavior that is characteristic of monotonic-function solutions. In other words, given that targets in such a case cannot all be satisfied simultaneously, optimal policy will tend to satisfy certain of these targets at the expense of others.

Once again, we can attribute this to two reasons. On the one hand, distant targets are assocated with steep gradients which induce sharp and erratic behavior. On the other, there is the potential for system oscillation (mentioned earlier) which is exploited by the optimization process.[3]

Consequently, in our experiments, we operated with targets which were slightly beyond reach. By this, we mean growth rates of relevant variables which are roughly between 1/2 and 1 percentage point above what can be actually achieved.

Footnotes

*John Derpanopoulos is Chief Economist for Sarasin International Securities, London, England.

[1]In CES-type functions where the elasticity of substitution (ρ) is not equal to unity, the relevant matrix of elements must undergo additional updating of the sort $\alpha_{ij(t+1)} = \alpha_{ij(t)} (1 + y_{it})^{-\rho-1}$.

[2]Regarding computation time, an average iteration required approximately 30 seconds of CPU time on a PDP 11/780 for a 53 x 53 dimensional problem. If several within-period approximation updates were involved, computation time obviously exceeded this.

[3]As one would expect, oscillations are greatly reduced with a reduction in returns to scale. Other factors which influence oscillations in a negative (reduced) direction are the share of labor in value added and the speed at which the effective capital stock is affected by investment. Surprisingly, oscillations increase as one lowers the elasticity of substitution in production.

References

Athans, Michael. "The Discrete Time Linear-Quadratic-Gaussian Stochastic Control Problem." Annals of Economic and Social Measurement. 1(1972):449-92.

Derpanopoulos, John. "Optimal Macroeconomic Policies: The Case of India." Ph.D. thesis, University of California, Berkeley, 1983.

Dixon, Peter B., B. P. Parmenter, John Sutton, and D. P. Vincent. ORANI: A Multisectoral Model of the Australian Economy. Amsterdam: North-Holland Publishing Company, 1982.

Johansen, Leif. A Multi-Sectoral Study of Economic Growth. Amsterdam: North-Holland Publishing Company, 1964.

Kailath, Thomas. Linear Systems. Englewood Cliffs, N. J.: Prentice-Hall, Inc., 1980.

Ortega, James, and W. C. Reinboldt. Iterative Solution of Nonlinear Equations in Several Variables. New York: Academic Press, 1970.