

ON THE SOLUTION OF OPTIMAL CONTROL PROBLEMS
AS MAXIMIZATION PROBLEMS

Ray C. Fair

Econometric Research Program
Research Memorandum No. 146
February 1973

Econometric Research Program
PRINCETON UNIVERSITY
207 Dickinson Hall
Princeton, New Jersey

ON THE SOLUTION OF OPTIMAL CONTROL PROBLEMS
AS MAXIMIZATION PROBLEMS

Ray C. Fair*

1. Introduction

There appears to be among many economists the view that the computation of optimal controls for moderate- to large-scale nonlinear econometric models is not feasible. Pindyck [19], for example, has questioned whether "non-linear optimization [is] worth all of the computational difficulty that it entails."¹ Shupp [24] has stated that "the size and complexity of these models preclude formal optimization."² Holbrook [13] has been led to consider only quadratic objective functions and to propose a method of obtaining approximate solutions to optimal control problems for nonlinear models.³ The results presented in this paper indicate that this view is not correct, even for models of up to 100 or 200 equations. The results suggest that it is feasible to compute optimal controls for most econometric models encountered in practice.

Historically, optimal control problems have been formulated in continuous time and have been looked upon as problems in choosing functions of time to maximize an objective function. Fairly advanced mathematical techniques are required to solve these problems. For discrete-time models, however, which include virtually all large-scale econometric models, optimal control problems can also be looked upon as problems in choosing parameters to maximize an objective function. The number of parameters to be determined is equal to the number of

*I would like to thank Gregory C. Chow, Kenneth D. Garbade, Stephen M. Goldfeld, and Richard E. Quandt for many helpful comments.

¹Pindyck [19], p. 388.

²Shupp [24], p. 94.

³Holbrook's method is based on obtaining numerically a linear approximation to the relationship between the value of each control variable for each period and the value of each target variable for each period. It is unclear as to how close a solution obtained using Holbrook's method is likely to be to the true *optimum solution*.

control variables times the number of time periods chosen for the problem. From this perspective, optimal control problems are straightforward maximization problems, and in attempting to solve problems in this way, one can take advantage of the recent advances that have been made in computational algorithms for maximizing nonlinear functions of parameters. This approach, of treating optimal control problems as problems of maximizing a nonlinear function of parameters, is the approach taken in this paper.

2. The General Method of Solution

Assume that the model under consideration is deterministic⁴ and has g equations. Write each equation for each period of time as

$$(1) \quad f_{it}(y_t, z_t, x_t, \alpha_{it}) = 0 ; \quad \begin{array}{l} i = 1, \dots, g; \\ t = 1, \dots, T; \end{array}$$

where y_t is a vector of observations for period t on the g endogenous variables in the model, z_t is a vector of observations for period t on the noncontrol, predetermined variables in the model, x_t is a vector of observations for period t on the control variables in the model, and α_{it} is a vector of non-zero parameters that are included in equation i for period t . The t subscripts in α_{it} and f_{it} allow for the possibility that some parameters and some functional forms are changing over time.⁵ Lagged endogenous variables are included in the z_t vector. T is the total number of periods to be considered in the control problem.

The model in (1) is assumed to be such that, for each t , given values for

⁴Stochastic models are discussed in Section 7.

⁵It is assumed throughout this paper that the values of α_{it} and the values of the exogenous variables in the z_t vector are known with certainty.

z_t , x_t , and α_{it} ($i = 1, \dots, g$) one can solve numerically for y_t . In practice, most large-scale econometric models are solved each period by some version of the Seidel method.⁶ Further, one can frequently isolate each component of the y_t vector on one side of one equation, which greatly aids in the solution of the model. If the model is solved for more than one period and if the solution is meant to be a dynamic one, then the solution values of the endogenous variables for previous periods are used, when appropriate, as values for the lagged endogenous variables in the z_t vector. For linear models, of course, values of y_t are merely obtained from reduced form equations.

For a time horizon of T periods, the objective function, h , is taken to be a function of y_t , z_t , and x_t ($t = 1, \dots, T$):

$$(2) \quad W = h(y_1, \dots, y_T; z_1, \dots, z_T; x_1, \dots, x_T),$$

where W , a scalar, is the value of the objective function corresponding to values of y_t , z_t , and x_t ($t = 1, \dots, T$).

The optimal control problem for this discrete-time, deterministic model is to choose values of x_1, \dots, x_T so as to maximize W subject to the equation-constraints in (1). The givens of the problem are the value of each α_{it} , the values for each period of the purely exogenous variables, and initial values for the lagged endogenous variables. Assume that x_t is of dimension k , so that there are kT control values to determine. Let x be a kT -component vector denoting these values: $x = (x_1, \dots, x_T)$. Now, for each value of x , one can compute a value of W by first solving the model in (1) for y_1, \dots, y_T and then

⁶See, for example, Fromm and Klein [10], pp. 373-382.

using these values along with the values for z_1, \dots, z_T and x to compute W in (2). The optimal control problem can thus be looked upon as a problem in choosing parameters (the elements of x) to maximize a nonlinear function. The function in the optimal control case is h in combination with the f_{it} . In general it will not be possible to express y_t explicitly in terms of z_t , x_t , and α_{it} , so that in general it will not be possible to write W in (2) explicitly as a function of z_t , x_t , and α_{it} ($t = 1, \dots, T$). Nevertheless, given values for z_t and α_{it} ($t = 1, \dots, T$), values of W can be obtained numerically for different values of x .

There are many algorithms available for maximizing (or minimizing) nonlinear functions of parameters. Since W cannot in general be written as an explicit function of x , it will in general be difficult to obtain analytically the partial derivatives of h with respect to the elements of x . Therefore, in attempting to solve optimal control problems by treating them as problems in maximizing a nonlinear function of parameters, one will usually be required either to use algorithms that do not require derivatives or else to compute derivatives numerically. Both approaches have been followed for the results in Sections 4 and 5.

Algorithms that do not require derivatives and algorithms for which derivatives are obtained numerically spend most of their time doing function evaluations. For the results in Sections 4 and 5, over 75 percent of the time was spent doing function evaluations for all algorithms tried except in two cases, where the figures were 52 and 53 percent. One function evaluation in the present context corresponds to the solution of a g -equation model for T periods (plus the rather trivial computation, once y_1, \dots, y_T are determined, of W in (2)). It is therefore quite important to solve a model in the most efficient way possible, since for one solution of the optimal control problem a

model will usually be solved hundreds or thousands of times. Some suggestions are presented in Section 6 for efficient ways of solving models.

Much of the engineering literature on optimal control is concerned with continuous-time models and so is not of direct concern here. Polak [20], however, does present a good discussion of the discrete optimal control problem in engineering.⁷ The discrete-time model considered by Polak differs from the standard econometric model considered in this paper in that his model is already in reduced form. In the notation of this paper, each component of y_t would be written as an explicit function of z_t , x_t , and α_{it} for Polak's model. The fact that the derivatives of y_t with respect to z_t and x_t can be directly obtained for Polak's model allows Polak to obtain fairly easily the derivatives of the objective function with respect to the values of the control variables. Polak also reports that the time horizon for the problems he is considering may be as large as 1000 periods,⁸ which is much larger than the time horizon for most problems in economics, where the horizon is likely to be much less than even 100 periods. The discrete optimal control problem in economics is thus on the one hand easier than the corresponding problem in engineering in that the time horizon appears to be much smaller and on the other hand more difficult in that analytic derivatives of the objective function with respect to the values of the control variables are not easy to obtain because of the non-reduced-form nature of most econometric models.

⁷See especially pp. 66-71. See also Athans [1] for a discussion of the linear-quadratic-gaussian stochastic control problem for discrete-time models.

⁸Polak [20], p. 67. Polak does not, however, report on any actual solutions of problems of this sort in his book.

3. The Computational Algorithms Used

Three basic algorithms were used for the results in Sections 4 and 5. The first is the 1964 algorithm of Powell [21], which does not require any derivatives. The second is a gradient algorithm, which requires first derivatives. The third is the quadratic hill-climbing algorithm of Goldfeld, Quandt, and Trotter [12], which requires both first and second derivatives. The gradient algorithm that was used in this study is a member of the class of algorithms considered by Huang [15].⁹ The algorithms within this class basically differ from each other in how the approximation to the inverse of the matrix of second partial derivatives is updated after each iteration. One member of this class is the well-known DFP variable metric algorithm.¹⁰ Some results using the DFP algorithm are reported below, but the main gradient algorithm that was used in this study is the one that updates by means of the "rank one correction formula."¹¹ This algorithm appeared to give the best results. Some results using one other member of the class of algorithms considered by Huang are also reported below.¹² All three of the gradient algorithms considered in this study use linear searches on each iteration.

All of the computer programs were compiled in FORTRAN-H and were run on an IBM 360-91 computer at Princeton University.¹³ All derivatives for the gradient and quadratic hill-climbing algorithms were computed numerically. For the gradient

⁹See Powell [23] for an excellent summary of Huang's theory.

¹⁰See Davidon [7] and Fletcher and Powell [9].

¹¹See Powell [23], p. 41.

¹²See Powell [23], equations (31) and (32), p. 41, for a presentation of this algorithm.

¹³The Powell and quadratic-hill-climbing algorithms were programmed by Stephen M. Goldfeld and Richard E. Quandt. The three gradient algorithms were programmed by Thomas Russell.

algorithms the derivatives were computed in two ways. For one set of runs derivatives were obtained for each iteration by computing two function evaluations per parameter, each parameter being perturbed by equal amounts around the value available from the previous iteration. For the other set of runs derivatives were obtained for each iteration by computing only one function evaluation per parameter. The percentage amount by which parameters were perturbed (.01 percent) was not varied from iteration to iteration.¹⁴ Stewart [25] has proposed a more sophisticated way of computing numeric derivatives when using gradient algorithms, but his method was not tried in this study. For the quadratic hill-climbing algorithm first derivatives were always obtained by computing two function evaluations per parameter, as these computations had to be made anyway to obtain the own second derivatives, but the cross partial derivatives were computed in two ways. For one set of runs the cross partial derivatives were obtained by computing four extra function evaluations per set of two parameters, and for the other set of runs the derivatives were obtained by computing only one extra function evaluation per set of two parameters.¹⁵ The reason two methods were used to obtain derivatives for the gradient and quadratic hill-climbing algorithms --one more expensive but likely to be more accurate and one less expensive but likely to be less accurate--was to see how sensitive the results were to the

¹⁴ Let $f(a,b)$ be a function of two parameters. Then the formulas that were used to obtain the partial derivative of f with respect to a for the two runs are $(f(a+\epsilon,b) - f(a-\epsilon,b))/2\epsilon$ and $(f(a+\epsilon,b) - f(a,b))/\epsilon$, where $\epsilon = 0.0001a$ or 0.000001 , whichever is larger. For all of the runs the problems were set up so that the solution values of the parameters would be between about 0.1 and 10.0.

¹⁵ Using the notation in footnote 14, the formula used for the own second derivatives is $(f(a+\epsilon,b) - 2f(a,b) + f(a-\epsilon,b))/\epsilon^2$. The two formulas used for the cross partial derivatives are $(f(a+\epsilon,b+\eta) - f(a-\epsilon,b+\eta) - f(a+\epsilon,b-\eta) + f(a-\epsilon,b-\eta))/4\epsilon\eta$ and $(f(a+\epsilon,b+\eta) - f(a,b+\eta) - f(a+\epsilon,b) + f(a,b))/\epsilon\eta$, where $\eta = 0.0001b$ or 0.000001 , whichever is larger. In the second formula, values for $f(a,b+\eta)$ and $f(a+\epsilon,b)$ are available from the own-second-derivative calculations.

way in which the derivatives were obtained. Box, Davies, and Swann [5], for example, report that their experience is that "gradient methods employing numerical differentiation are, (with the exception of Stewart, 1967) usually inferior to the best direct search methods, and therefore not recommended."¹⁶ The results in this study do not confirm this view.

In the programs, the algorithms were taken to have converged when the absolute value of the difference between the value of each parameter on successive iterations was within a prescribed tolerance level. The Powell algorithm was generally more sensitive to the particular tolerance level used than were the gradient and quadratic hill-climbing algorithms, and for the results in Section 4 two sets of runs were obtained using the Powell algorithm, corresponding to two different tolerance levels.

Studies that have been done comparing different computational algorithms have tended to limit the size of the problems considered to 20 parameters or less. This is true, for example, of the comparisons in Bard [3], Box [4], Goldfeld and Quandt [11], Kowalik and Osborne [16], Murtagh and Sargent [17], Pearson [18], and Stewart [25]. Powell [22] reports that the DFP algorithm using analytic derivatives has been successful for problems of size 100 and that his 1964 algorithm and the DFP algorithm using numeric derivatives in the manner proposed by Stewart have solved problems of size 20.¹⁷ Wolfe [26] states that the upper limit to the size of problems that can be solved in which derivatives can be calculated analytically is around 100. For problems in which derivatives cannot be calculated, Wolfe's diagram indicates that the upper limit is about 10.¹⁸ The results reported below

¹⁶Box, Davies, and Swann [5], p. 32.

¹⁷Powell [22], p. 95.

¹⁸Wolfe [36], pp. xi-xii. It should be noted, however, that it is not clear from Wolfe's notes whether for these particular figures Wolfe is also including problems in which there are inequality constraints.

indicate that the upper limit to the size of problems that can be solved when derivatives are not calculated analytically is much larger than 10 or 20. The largest problem solved below was of size 239, and a number of problems between size 59 and 100 were solved. In fact, one of the main reasons why the method proposed in this paper appears feasible for most econometric models is the ease in which algorithms appear to be able to solve large problems even when analytic derivatives are not calculated.

4. An Example Using a Linear Model with a Quadratic Objective Function.

The method proposed in Section 2 was first used to solve one of the optimal control problems solved by Chow [6] for his nine-equation, linear econometric model. The model has two control variables. Chow solved various 10-period optimal control problems corresponding to different quadratic objective functions (to be minimized). The problem chosen to solve in this study is the second problem in Table 3 of Chow [6]. Two control variables and ten periods means that there are 20 parameters to be determined. The initial values for the 20 parameters were chosen to be zero, although in practice one could obviously choose better initial values than these. The results of solving this problem are presented in the first row of Table 1. Two runs for the Powell algorithm are reported, one which used a tolerance level of .0005 and one which used a tolerance level of .00001. Two runs each for the gradient and quadratic hill-climbing algorithm are also reported, corresponding to the two ways of computing derivatives. The latter two algorithms used a tolerance level of .00001.

Powell's no-derivative algorithm required 1687 function evaluations to attain the optimum using a tolerance level of .0005 and 2633 function evaluations using a tolerance level of .00001. The value of the objective function at the stopping point was smaller for the smaller tolerance level, but only by a very small amount.

TABLE 1
RESULTS FOR CHOW MODEL

	Powell										Gradient										Hill-Climbing												
	No. of Control Vars.	No. of Periods	No. of Parameters	Starting Point	Time per Funct. (sec.)	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct. (10) ⁹	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct. (10) ⁹	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct. (10) ⁹	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct. (10) ⁹	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct. (10) ⁹	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct. (10) ⁹				
1.	2	10	20	0's	.0018	3.27	1687	14	.12494578	1.33	614	21	.12494570	3.21	929	4	.12494558	3.21	929	4	.12494558	3.21	929	4	.12494558	3.21	929	4	.12494558	3.21	929	4	.12494558
2.	2	20	40	0's	.0034	17.38	4752	20	.47827735	6.44	1575	32	.47827742	22.21	3450	4	.47827668	22.21	3450	4	.47827668	22.21	3450	4	.47827668	22.21	3450	4	.47827668	22.21	3450	4	.47827668
3.	2	30	60	0's	.0050	41.78	7820	22	1.4724707	17.81	2945	42	1.4724725	17.81	2945	42	1.4724725	17.81	2945	42	1.4724725	17.81	2945	42	1.4724725	17.81	2945	42	1.4724725	17.81	2945	42	1.4724725
4.	2	40	80	0's	.0067	81.29	10960	23	4.2400571	34.27	4432	49	4.2400653	34.27	4432	49	4.2400653	34.27	4432	49	4.2400653	34.27	4432	49	4.2400653	34.27	4432	49	4.2400653	34.27	4432	49	4.2400653
5.	2	50	100	0's	.0085	129.94	14883	25	11.875411	57.54	5854	53	11.875442	57.54	5854	53	11.875442	57.54	5854	53	11.875442	57.54	5854	53	11.875442	57.54	5854	53	11.875442	57.54	5854	53	11.875442
6.	2	40	80	a	.0067	42.60	6253	15	4.2400580	12.04	1396	16	4.2400573	12.04	1396	16	4.2400573	12.04	1396	16	4.2400573	12.04	1396	16	4.2400573	12.04	1396	16	4.2400573	12.04	1396	16	4.2400573
*7.	2	30	60	0's	.0050	42.60	6253	15	4.2400580	20.22	2842	17	4.2400560	20.22	2842	17	4.2400560	20.22	2842	17	4.2400560	20.22	2842	17	4.2400560	20.22	2842	17	4.2400560	20.22	2842	17	4.2400560
**8.	2	30	60	0's	.0050	42.60	6253	15	4.2400580	26.14	4499	64	1.4724720	26.14	4499	64	1.4724720	26.14	4499	64	1.4724720	26.14	4499	64	1.4724720	26.14	4499	64	1.4724720	26.14	4499	64	1.4724720
	2	30	60	0's	.0050	42.60	6253	15	4.2400580	40.28	7455	57	1.4726138	40.28	7455	57	1.4726138	40.28	7455	57	1.4726138	40.28	7455	57	1.4726138	40.28	7455	57	1.4726138	40.28	7455	57	1.4726138
	2	30	60	0's	.0050	42.60	6253	15	4.2400580	27.69	4836	68	1.4724713	27.69	4836	68	1.4724713	27.69	4836	68	1.4724713	27.69	4836	68	1.4724713	27.69	4836	68	1.4724713	27.69	4836	68	1.4724713
	2	30	60	0's	.0050	42.60	6253	15	4.2400580	47.57	8790	67	1.4724700	47.57	8790	67	1.4724700	47.57	8790	67	1.4724700	47.57	8790	67	1.4724700	47.57	8790	67	1.4724700	47.57	8790	67	1.4724700

Notes: Powell = Powell [21].

Gradient = Gradient algorithm that uses rank one correction formula. See Powell [23], p. 41.

Hill-Climbing = Goldfeld, Quandt, and Trotter [12].

nc = not computed.

^aAnswer to 3 plus trend values otherwise.

*Gradient = DFP variable metric algorithm. See Powell [23], p. 37.

**Gradient = Algorithm presented in eqs. (31) and (32) in Powell [23], p. 41.

First run for Powell for each problem used tolerance level of .0005. Second run used .00001.

First run for Gradient for each problem used less expensive way of obtaining first derivatives. Second run used more expensive way.

First run for Hill-Climbing for each problem used less expensive way of obtaining cross partial derivatives. Second run used more expensive way.

Tolerance level for Gradient and Hill-Climbing was .0000

Chow Model: 9 equations; linear; lags of first order; welfare function is quadratic.

The corresponding parameter values for the two runs agreed to three significant digits, with the largest difference being .00015 (.70272 versus .70287). The gradient algorithm required 614 function evaluations to attain the optimum using one function evaluation per derivative per parameter and 1033 function evaluations using two. The value of the objective function at the stopping point was smaller for the second run, but again by only a very small amount. The corresponding parameter values for these two runs also agreed to three significant digits. The quadratic hill-climbing algorithm required 929 function evaluations to attain the optimum using one function evaluation per cross derivative and 3209 function evaluations using four. For these two runs the values of the objective function at the stopping point were the same. The time per function evaluation for the Chow-model, 10-period problem was .0018 of a second. The optimum obtained for this problem was the same as Chow had obtained.

The optimal control problem for the Chow Model was next made progressively larger by increasing the time horizon. The largest problem considered was a time horizon of 50 periods, which meant that there were 100 parameters to estimate. The results for 40, 60, 80, and 100 parameters are presented in rows 2 through 5 in Table 1 respectively. For the various problems the gradient algorithm clearly dominated Powell's in terms of speed of convergence. The use of the smaller tolerance level for the Powell algorithm increased the number of function evaluations considerably, and the values of the objective functions at the stopping points were only slightly larger for the larger tolerance level. Likewise, for the gradient algorithm the values of the objective functions at the stopping points were only slightly larger for the runs using one function evaluation per derivative. For the quadratic hill-climbing algorithm no accuracy at all was lost using one function evaluation per cross derivative. The quadratic hill-climbing

algorithm was not tried after 40 parameters, although the use of the algorithm for problems of, say, size 100 is not completely out of the question. Using the less expensive way of obtaining cross derivatives, it requires $0.5N^2 + 1.5N$ function evaluations to compute the vector of first derivatives and the matrix of second partial derivatives per iteration (where N is the number of parameters). If 4 iterations are required to attain convergence, then roughly 20600 function evaluations would be required to solve the 100-parameter problem.

Adding extra periods for the Chow model in general had little effect on the optimal parameter values of previous periods, so that, for example, the answer to the 60-parameter problem was close to the answer to the 80- or 100-parameter problem for the first 60 parameters. In view of this, the answer to smaller problems should be a good starting point for larger problems, and so to test this, the answer to the 60-parameter problem was used as a starting point for the first 60 parameters of the 80-parameter problem. Starting points for the other 20 parameters were obtained by letting the values of the two control variables grow by 6 and 5 percent respectively, these figures being obtained by observing how the control variables were growing in the answer to the 60-parameter problem. The results of this test are presented in row 6 of Table 1. For the gradient algorithm the number of function evaluations was cut by about a factor of 3 (from 4432 to 1396 and from 8517 to 2842), a substantial savings. For the Powell algorithm the number of function evaluations was cut from 10960 to 6253 using the larger tolerance level and from 16219 to 6253 using the smaller tolerance level. In both cases for the Powell algorithm, a slightly smaller value of the objective function was obtained by starting the parameter values from zero.

As a final test using the Chow model, two other gradient algorithms were tried for the 60-parameter problem. The results are reported in rows 7 and 8 of Table 1. Neither algorithm worked as well as the rank one algorithm. The

DFP algorithm required about 155⁴ more function evaluations than did the rank-one algorithm for the run using one function evaluation per derivative. For the run using two function evaluations per derivative, the DFP algorithm did not quite attain the optimum.

5. An Example Using a Nonlinear Model with a Non-Quadratic Objective Function.

The method of Section 2 was next used to solve a more complicated optimal control problem. The model used was the Fair model [8], less the monthly housing starts sector. The model used consists of 19 equations, is nonlinear, has lags of up to eighth order, and was estimated under the assumption of first-order serial correlation of most of the error terms.¹⁹ The initial period was taken to be 1962III, and the horizon for the various runs was either 10, 20, 25, or 60 quarters. The number of control variables was varied between one and four. Government spending was always taken to be a control variable. The other three variables that were sometimes used as control variables were the level of consumer sentiment, plant and equipment investment expectations, and nonfarm quarterly housing starts. These latter three variables are clearly not variables under the direct control of the government, but for purposes of illustrating the method of solution, there is no harm in treating them as if they were. The objective function was deliberately chosen to be non-quadratic in the variables of the model. The objective function (to be minimized) was:

$$\begin{aligned}
 W = \sum_{t=1}^T \{ & 10(g_{PD_t})^2 + 10/UR_t - .030/2 + (\frac{CD_t}{GNP_t} - .094)^2 \\
 & + (\frac{CN_t}{GNP_t} - .275)^2 + (\frac{CS_t}{GNP_t} - .257)^2 \\
 & + (\frac{IP_t}{GNP_t} - .101)^2 + (\frac{IH_t}{GNP_t} - .038)^2 \} ,
 \end{aligned}$$

¹⁹The coefficients were taken from Table 11-4 in [8].

where g_{PD_t} is the rate of growth (at an annual rate) of the private output deflator, UR_t is the unemployment rate, and the five ratios are the ratios of durable consumption, non-durable consumption, service consumption, plant and equipment investment, and housing investment to gross national product respectively. The slashes around $UR_t - .030$ denote the fact that $|UR_t - .030|$ was taken to be equal to $UR_t - .030$ if $UR_t \geq .030$ and zero otherwise. In other words, welfare was not improved for an unemployment rate below .030, but it was not decreased either, as a straight quadratic function would imply. The objective function is non-quadratic in this respect, as well as in targeting ratios of the various components of GNP to GNP itself. The rate of inflation and the unemployment rate were weighted ten times more heavily in the objective function than were the ratios. It should be noted that the welfare function is not differentiable at $UR_t = .030$. In the present case, however, the optimum values of UR_t were always greater than .030, and the lack of differentiability at $UR_t = .030$ did not appear to be a problem for the algorithms for which numeric derivatives had to be computed. In general, if the lack of differentiability of either the model or the welfare function appears to be important (as it might be, for example, for models in which capacity ceilings play an important role), then algorithms that do not require the computation of derivatives may be better choices than those that do.

The results for the various runs using the Fair model are presented in Table 2. The second control variable, the level of consumer sentiment, does not enter the model currently, but only with lags of one or more periods, so when this variable was used as a control variable, the number of values of this variable to be determined was one less than the number of periods. Except for lines 7 and 8, historic values were used as starting points for the values of the control variables. Again, two runs each for the gradient and quadratic hill-climbing algorithms are reported, corresponding to the two ways of computing

TABLE 2
RESULTS FOR FAIR MODEL

15.	No. of Control Vars.		Powell				Gradient				Hill-Climbing				
	No. of Periods	No. of Parameters	Starting Point	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct.	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct.	Time (sec.)	No. of Funct. Evals.	No. of Iterations	Value of Obj. Funct.
1.	1	10	hist. val.	8.40	1123	12	.076917048	1.35	163	9	.076917048	2.33	269	4	.076917047
			val.					1.80	225	8	.076917048	6.43	809	4	.076917047
2.	1	20	hist. val.	32.67	2152	12	.16427133	4.78	307	11	.16427053	15.45	929	4	.16427052
			val.					8.71	573	12	.16427052	52.16	3209	4	.16427052
3.	2	20	hist. val.	48.16	3142	10	.16380396	20.83	1371	29	.16373695	63.26	3284	4	.16373693
			val.					40.39	2756	32	.16373693	198.52	12176	4	.16373693
4.	3	20	hist. val.	84.36	5580	10	.16290711	35.82	2344	35	.16288105	nc	nc	nc	.20080975
			val.					73.84	4787	38	.16288104				.20080975
5.	4	20	hist. val.	156.20	10513	15	.16250254	113.82	7314	84	.16234170				.20080975
			val.					197.26	12807	77	.16234169				.20080975
6.	4	25	hist. val.	231.40	12822	14	.20624995	204.47	10181	95	.20511032				.30350850
			val.					384.73	19631	95	.20511030				.30350850
7.	4	60	a	.0439	nc	nc		189.51	25661	104	.58885958				.80730797
									nc	nc					.80730797
8.	4	20	b	.0148	2615	4	.16249898	110.60	7047	81	.16234170				.16288104
								196.45	12793	77	.16234169				.16288104
9.	3	20	hist. val.	.0148	-	-		35.12	2305	34	.16288106				.20080975
			val.					73.98	4821	38	.16288104				.20080975
10.	3	20	hist. val.	.0148	-	-		35.53	2336	35	.16288105				.20080975
			val.					72.55	4777	38	.16288104				.20080975

Notes: See Table 1. Tolerance level for Powell was .000005. Tolerance level for Gradient and Hill-Climbing was .00005.
 *Gradient = DFP variable metric algorithm.
 **Gradient = Algorithm presented in eqs. (31) and (32) in Powell [23], p. 41. Answer to 6 plus trend values otherwise.
 Fair Model: 19 equations; nonlinear; lags of up to eighth order; first-order serial correlation of the error terms in most of the equations; welfare function is not quadratic.

derivatives. The tolerance level used for these two algorithms was .00005.

The tolerance level used for the Powell algorithms was .000005.

From the results in Table 2, it can be seen that the gradient algorithm worked better than Powell's. The number of function evaluations was usually less for the gradient algorithm, and for the problems of greater than 20 parameters, the Powell algorithm did not quite attain the optima that the gradient algorithm did. For the 39- through 99-parameter problems, the largest differences between the parameter values computed by the Powell algorithm and the corresponding parameter values computed by the gradient algorithm were 26, 8, 34, and 88 percent respectively. An even smaller tolerance level was tried for some of the runs using the Powell algorithm (.0000001 versus .000005) to see if this resulted in a smaller value of the objective function, but the results were not improved using the smaller tolerance levels. For the gradient algorithm the use of the less expensive way of obtaining derivatives resulted in virtually no loss in accuracy for any of the runs. For the quadratic hill-climbing algorithm the use of the less expensive way of computing cross partial derivatives resulted in no loss in accuracy at all and, of course, substantial savings on cost. For the problem of 4 control variables and 25 periods (99 parameters), the gradient algorithm using the less expensive way of computing derivatives required 10181 function evaluations and took about 3.4 minutes to attain the optimum.

When the 79-parameter problem was started from the answer to the 59-parameter problem plus historical values otherwise (line 8), the speed of convergence was only slightly increased for the gradient algorithm. The number of function evaluations fell from 7314 to 7047 for the one run and from 12807 to 12793 for the other. The number of function evaluations fell substantially for the Powell algorithm, but the optimum was still not attained.

When the other two gradient algorithms were tried for the 59-parameter problem (lines 9 and 10), the results were virtually the same as for the rank one algorithm. For this problem there is nothing to choose among the three algorithms.

The largest problem tried for the Fair model was 4 control variables and 60 periods (1962III - 1977II) for a total of 239 parameters. The answer to the 99-parameter problem was used as a starting point plus historical or extrapolated values otherwise. Only the gradient algorithm using the less expensive way of obtaining derivatives was tried for this problem. The program was allowed to run for approximately 20 minutes. At the end of 20 minutes and 104 iterations, the value of the objective function was changing only in the eighth decimal place between iterations and the largest difference between any corresponding parameter values on the last two / iterations was .0007. The value of the objective function at the starting point was .80730797 and the value after 104 iterations was .58885958. The starting point turned out to be fairly far away from the stopping point, with unemployment rates of about 7 percent near the end of the horizon compared with the stopping-point values of around 5 percent. The stopping-point values for the 239-parameter problem appeared to be in line with what would be expected from observing the answers to the smaller problems. The Powell algorithm was started from the values attained by the gradient algorithm on the 53rd iteration (an objective-function value of .58890611) to see if it would go anywhere. A tolerance level of .000005 was used. The algorithm went one iteration, lowered the objective function to .58890571, and stopped (the convergence criterion having been met for all parameters), a clear failure in view of the value obtained by the gradient algorithm. One other result is also of interest to note here. The gradient algorithm was also started from the values attained on the 53rd iteration. A tolerance level of .000005 was used. The algorithm went one

iteration, lowered the objective function to .58890575, and stopped (the convergence criterion having been met), also a clear failure. By starting the gradient algorithm over on the 53rd iteration, one lost the approximation to the inverse of the matrix of second partial derivatives that had been developed over 53 iterations, which in the present case was obviously quite important. A similar result occurred when experimenting with the 99-parameter problem. These results suggest that if one contemplates having to restart the gradient algorithm for one reason or another (like running out of time on the computer), one ought to save the latest approximation to the inverse of the matrix of second partial derivatives to be used when the algorithm is restarted. The results also suggest, oddly enough, that when using the gradient algorithm one ought not to start the algorithm too close to the (presumed) optimum for fear that the algorithm will get stuck before it has a chance to build up a good approximation to the inverse of the matrix of second partial derivatives.

The answers to the problems for the Fair model were characterized by a large value of government spending in the first period (compared with the historical value) and large values near the end of the time horizon. In the model employment responds faster to government spending than does the price level, and so the relatively large values of government spending for the last few periods of the horizon are taking advantage of this fact and lowering the unemployment rate without having too much effect on the price level.²⁰ The large value of spending in the first period is apparently designed to lower the

²⁰To avoid undesirable end-point effects in practice, one can always extend the horizon a few periods beyond the actual horizon of interest. For the Fair model it appeared that the horizon should be lengthened by about 5 quarters. Because of the end-point effects, the last few answers to the 99-parameter problem for each control variable were not used as starting points for the 239-parameter problem.

unemployment rate quickly from its relatively high historic level. Excluding beginning and ending effects, the particular objective function used resulted in an unemployment rate of about 5.0 percent and an annual rate of inflation of about 2.2 percent. The IP_t/GNP_t and IH_t/GNP_t ratios were met almost exactly when plant and equipment investment expectations and housing starts were used as control variables, as would be expected. The three consumption ratios were not met as exactly when consumer sentiment was used as a control variable since in this case there was, in effect, only one main control variable influencing three ratios.

In Table 3 are presented estimates for each run in Tables 1 and 2 of the percentage of time that was spent doing function evaluations. The estimates were obtained by multiplying the time per function evaluation by the number of function evaluations and dividing this figure by the total time for the job. For the Fair model abnormal exits sometimes occurred from the function-evaluation program (before all of the computations were performed), which causes some of the percentages for the Fair model in Table 3 to be too high. Abnormal exits occur when parameter values imply that the logarithm of a negative number should be taken. The estimates in Table 3 are also subject to error for reasons that have to do with the way that computation time in the computer is estimated. In general, the percentages are quite high in Table 3, indicating the importance of writing efficient programs for evaluating functions.

6. An Evaluation of the Practical Usefulness of the Method.

The results in Sections 4 and 5 are very encouraging as to the feasibility of using the method proposed in Section 2 even for large-scale models. For a 20-period problem the 19-equation Fair model takes .0148 of a second per function evaluation on the IBM 360-91 computer. The Fair model can be solved without the use of the Seidel method since the nonlinear part of the model is recursive.

TABLE 3

ESTIMATES OF PERCENTAGE OF TIME SPENT DOING FUNCTION EVALUATIONS

From Table 1:

Row	powell		Gradient		Hill-climbing	
	(1)	(2)	(1)	(2)	(1)	(2)
1.	93	97	83	95	52	82
2.	93	95	83	93	53	79
3.	94	90	83	91	-	-
4.	90	92	87	90	-	-
5.	97	95	86	91	-	-
6.	98	98	78	94	-	-
7.	-	-	86	93	-	-
8.	-	-	87	92	-	-

From Table 2:

Row	Powell		Gradient		Hill-climbing	
	(1)	(2)	(1)	(2)	(1)	(2)
1.	96		87	90	83	91
2.	97		95	97	89	91
3.	97		97	101	77	91
4.	98		97	96	-	-
5.	100		95	96	-	-
6.	100		90	92	-	-
7.	-		95	-	-	-
8.	99		94	96	-	-
9.	-		97	96	-	-
10.	-		97	97	-	-

If a 100-equation model could be solved in the same way, it should take only about 5 times longer to solve this model than it takes to solve the Fair model since the number of computations per equation is not likely to vary much from model to model. Econometric models tend to be larger because of more equations and not because of more variables per equation. If the Seidel method must be used to solve a model and if for each iteration for each period the entire model must be passed through, then the cost per solution of the model is increased in proportion to the number of iterations that are required to solve the model each period. If, for example, it takes 5 iterations to solve a 100-equation model each period, it should take about 25 times longer to solve this model than it takes to solve the Fair model. Since algorithms that do not require derivatives or for which derivatives are computed numerically spend most of their time doing function evaluations, the total time that it takes to solve a control problem for a 100-equation model that requires 5 iterations per solution of the model should be about 25 times greater for the same problem than the corresponding time in Table 2 for the Fair model. A 20-period problem with one control variable should thus take about 2.0 minutes using the gradient algorithm and the less expensive way of obtaining derivatives (25×4.78 seconds). A 20-period problem with two control variables should take about 8.7 minutes (25×20.83 seconds). The problem of 4 control variables and 25 periods should take about 85.2 minutes (25×204.47 seconds).

Although the times just mentioned are not completely out of the range of practicality, it is possible that in practice the times can be substantially cut down. First, good starting points can be quite important, and significant time may be saved by first solving a small problem (say one control variable), using the answer to this problem as a starting point for a somewhat larger problem (say two control variables), and so on, building up to the largest problem that one wants to consider. Also, once one has solved a particular optimal control problem

once, the answer to this problem may be a good starting point for a slightly different problem (say a slight change in the objective function). In other words, it may not be too costly to experiment with different objective functions or a slightly different specification of the model once one solution to a particular problem has been obtained. It may also be the case that from a welfare point of view or from the point of view of feasibility one wants to keep the control variables within certain bounds. This can be done by including control variables in the objective function and penalizing deviations of the values of the control variables from target values. If this is done, one has a natural starting point for the control variables--the target values--and this may significantly increase the speed of convergence of the algorithm being used, in addition perhaps to decreasing the likelihood that the algorithm goes to a local but not the global optimum.

A second way in which much time might be saved by models that need to be solved by the Seidel method is by choosing good initial values of the endogenous variables to begin the solution of the model each period. Since most algorithms perturb the parameters (in the present case, the values of the control variables) only a slight amount between function evaluations, particularly when derivatives are being computed, a good choice for the initial values of the endogenous variables is likely to be the solution values obtained in the process of computing the previous function evaluation. It is possible that this choice can cut the number of iterations needed per solution of the model per period to two or three, which would greatly save on cost.

A third way in which time can be saved is to write the program that does function evaluations in such a way that no computations are performed other than those that are absolutely needed in going from values of the control variables to the value of the objective function. For example, any sets of calculations using exogenous variables that are not changed as a result of changes in the

values of the control variables should not be done in the function-evaluation program, but only once before the solution of the optimal control problem begins. This kind of efficient programming was not done for the results in Tables 1 and 2.

If for a 100-equation model one could, by following the above suggestions, cut the number of iterations using the Seidel method to an average of 2.5 and could further cut the time per function evaluation by 25 percent, then the times quoted above (2.0, 8.7, and 85.2 minutes) would be cut to 0.75, 3.3, and 32.0 minutes respectively. These times may be further cut by a factor of 2 or more by better choices of initial parameter values than those used for the results in Table 2.²¹

In terms of the size of the problem that the method proposed in this paper can handle, there is an obvious tradeoff between the size of the model, the number of control variables, and the length of the decision horizon. It is hard to establish any precise rules as to what problems are practical to solve and what are not because no two models and problems are the same. Furthermore, for some problems one algorithm may work best and for others another may work

²¹ Albert Ando has communicated to the author a "conservative" estimate that for the solution of the 200-equation FMP model it takes about .00500 of a second per iteration per period on an IBM 370-165 computer. This figure compares with .00072 for the solution of the 19-equation Fair model (divide .0072 in Table 2 by 10). Since the FMP model has 10.5 times more equations than the Fair model, one would expect the time per iteration per period to be about 10.5 times greater for the FMP model. The figure supplied by Ando indicates that the time is only 6.9 times greater. Ando's results thus suggest that the times cited in the text above may be too conservative. It should also be noted that Ando's results are for a program that was not written with optimal control problems in mind.

The FMP model usually takes between 10 and 15 iterations to solve per period using the Seidel method. However, the values used as initial values for the endogenous variables are the solution values of the previous quarter, and, as suggested above, in an optimal-control context one should be able to do much better than this.

best. Each person must to some extent determine for oneself through experimentation the practical limits to the size of problems that one can solve. Nevertheless, the results in this study can give some indication of the likely cost of various problems. One important question in this regard is how rapidly the number of function evaluations increases as the number of parameters to be estimated increases. From the results in Tables 1 and 2 one can compute the extra number of function evaluations required per additional parameter ($\Delta FE/\Delta N$, where FE is the number of function evaluations and N is the number of parameters) and observe how this quantity varies as the total number of parameters varies. These computations are presented in Table 4. For the quadratic hill-climbing algorithm, $\Delta FE/\Delta N$ clearly increases as N increases since the number of function evaluations required to compute first and second derivatives per iteration increases as the square of N . From the results for the Chow model there is only a slight tendency for $\Delta FE/\Delta N$ to increase as N increases for the Powell and gradient algorithms. From the results for the Fair model there is somewhat more of a tendency in this direction for the two algorithms, but this tendency is far from being uniform. In general, the results in Table 4 indicate that there is only a slight tendency for $\Delta FE/\Delta N$ to increase as N increases for the Powell and gradient algorithms.

The time required per function evaluation should be roughly proportional to the number of periods times the number of equations in the model times the number of Seidel iterations required to solve the model. The time required to solve a control problem is roughly equal to the time required per function evaluation times the number of function evaluations. If the number of function evaluations varies only in proportion to the number of parameters ($\Delta FE/\Delta N$ not increasing as N increases), then the time required to solve a control problem should be roughly proportional to the square of the number of periods times the number of control variables times the number of equations times the

TABLE 4
VALUES OF $\Delta FE/\Delta N$

From Table 1:

N	ΔN	Powell		Gradient		Hill-climbing	
		(1)	(2)	(1)	(2)	(1)	(2)
20	20	84.4	141.6	30.7	51.7	46.5	160.5
40	20	153.3	216.8	48.1	91.1	126.1	480.1
60	20	153.4	214.5	68.5	130.6	-	-
80	20	157.0	238.1	74.4	152.6	-	-
100	20	196.2	236.1	71.1	132.0	-	-

From Table 2:

N	ΔN	Powell		Gradient		Hill-climbing	
		(1)	(2)	(1)	(2)	(1)	(2)
10	10	112.3		16.3	22.5	26.9	80.9
20	10	102.9		14.4	34.8	66.0	240.0
39	19	52.1		56.0	114.9	123.9	471.9
59	20	121.9		48.7	101.6	-	-
79	20	246.7		248.5	401.0	-	-
99	20	115.5		143.4	341.2	-	-
239	140			110.6 ^a	-	-	-

N = number of parameters

FE = number of function evaluations

^aThe 239-parameter run was started from a more accurate point than the others and was terminated at a tolerance level of only .0007 versus .00005 for the others.

number of seidel iterations. In this case, if the number of seidel iterations required to solve a model does not increase as the number of equations of the model increases, then the time required to solve a control problem should increase only in proportion to the increase in the number of equations. Otherwise, the time will increase more than in proportion to the increase in the number of equations.²² The time required to solve a control problem is proportional to the square of the number of periods because an increase in the number of periods increases both the number of parameters and the time required per function evaluation. If the number of function evaluations increases more than in proportion to the number of parameters, then the time required to solve a control problem will increase more than in proportion to the increase in the square of the number of periods times the number of control variables.

Baring further results, some tentative conclusions can be drawn from the results in this study as to the size of problems that it appears feasible to solve using the method discussed in section 2. For models of about 20 equations, it appears quite practical to solve problems in which the product of the number of control variables and the number of periods is greater than 100. For models of about 100 equations, a product of 100 is probably within the range of practicality. For models of about 200 equations, a product of 60 may be close to the limit of practicality.²³ The use of good starting points

22

If the objective function to be maximized becomes less well behaved as the number of equations increases, this should also cause the time required to solve a control problem to increase more than in proportion to the increase in the number of equations. Without further experimentation using other models it is not clear how sensitive the shape of the objective function is likely to be to the number of equations in the model.

23

Holbrook, using his method, obtained an approximate solution to a control problem involving 4 control variables and 8 periods for a model of 261 equations. Since in the present context this involves determining only 32 parameters, it seems quite likely that this problem could be solved by the method proposed in this paper, even though the model is quite large.

and efficient programming may, of course, greatly extend even these limits. Since most econometric models do not exceed 200 equations and since the number of control variables in any one model can usually be kept under, say, 5 without seriously restricting the problem, the method considered in this paper should be able to handle most problems of interest to policy makers who use econometric models in their decision-making process. It should also be noted that the method considered in this paper requires relatively little human effort. All one has to do is write a program to solve the model and compute the objective function. No derivatives are required, no analytic approximations have to be made, and the model does not have to be set up in any special form.

The results in Tables 1 and 2 indicate that the gradient algorithm using the less expensive way of obtaining derivatives is the most efficient. Slightly more accuracy may be obtained by using the more expensive way of obtaining derivatives or by using the quadratic hill-climbing algorithm, but in general this increased accuracy is not likely to be worth the cost. For the quadratic hill-climbing algorithm no accuracy was gained using the more expensive way of computing cross partial derivatives, and so this way is not recommended. The Powell algorithm was generally more expensive than the gradient algorithm, and for the Fair model it had a tendency to get close to but not quite to the optimum. The results in the two tables do, of course, indicate that quite large problems can be solved even when derivatives are obtained numerically. In practice, it may be desirable, after having attained an answer from one algorithm, to start another algorithm from this answer to be more certain that the true optimum has been attained. The quadratic hill-climbing algorithm, while being the most expensive for large problems, is likely to be the most robust to attaining the true optimum.

7. Stochastic Models

In the case of a linear model with additive error terms and a quadratic objective function it is well known that solving the deterministic control problem derived by setting the error terms to their expected values will provide the optimal first-period control values for the stochastic, closed-loop, feedback control problem. Therefore, if one solves the deterministic control problem each period, after observations on the state of the system for the previous period become available, one will over time make the same decisions regarding the current values of the control variables (i.e., the values of the control variables that the decision maker actually sets) as would be made by one who had solved the stochastic, closed-loop, feedback control problem explicitly in terms of feedback equations. To this extent, feedback equations need not be obtained, and one can concentrate on solving deterministic control problems as considered in the previous sections of this paper.²⁴ For most economic applications sufficient time is usually available to recompute the entire sequence of optimal controls each period.

For nonlinear models the first-period certainty-equivalence property does not hold. One procedure that might be followed in this situation is merely to treat the nonlinear-model case in the same way as one would treat the linear-model case, i.e., setting error terms to their expected values, and solve the deterministic control problem each period. This procedure is probably the one most often used in practice for solving nonlinear models,

²⁴

Knowledge of feedback equations for a particular model may aid one in understanding the dynamic properties and other characteristics of the model, and for this reason it may be useful to compute feedback equations even though they are not actually needed for the solution of the optimal control problem.

although Howrey and Kelejian [14] have shown that solving a nonlinear model by setting the error terms equal to their expected values is not equivalent to solving the reduced-form equations of the model.

For a nonlinear model the mean values of the endogenous variables can be obtained by means of stochastic simulation. A number of drawings from the joint probability distribution of the error terms can be taken, and for each drawing one can obtain by solving the model a set of values for the endogenous variables. The mean value for each endogenous variable can then be computed as the average of the values obtained from solving the model for the various drawings. Using the procedure of stochastic simulation, it may be possible for relatively small problems to obtain optimal open-loop controls for nonlinear, stochastic models in a manner similar to that done above for nonlinear, deterministic models. Say the aim were to maximize the expected value of the objective function. For each choice of control values, one could compute by means of stochastic simulation the mean value of W . The computed mean value of W would be the value returned to the maximization algorithm, and the algorithm would be used in the usual way in an attempt to find that set of control values for which the mean value of W were at a maximum. Each function evaluation in the stochastic case would correspond to an entire stochastic simulation. If, for example, 50 drawings from the joint probability distribution of the error terms were needed to obtain an adequate approximation to the expected value of W , then approximately 50 times more time would be needed per function evaluation for the stochastic problem than for the deterministic problem. Even though the cost is high for the stochastic problem, it may be feasible for small problems to carry out the above suggestion. If one did carry out the above suggestion and found the optimum and if one recomputed the entire sequence of optimal controls each period, one would over time

make the same decisions regarding the current values of the control variables as would be made by one who had solved the stochastic, open-loop, feedback control problem explicitly in terms of feedback equations.

For the control problem for nonlinear, stochastic models, Athans [1], [2] has suggested first solving the deterministic control problem (the deterministic problem being obtained by setting the error terms equal to their expected values) and then linearizing around the deterministic-control paths to obtain linear feedback equations around the paths. The aim is over time to keep the actual paths close to the deterministic-control paths. While Athans' suggestion may be useful for engineering applications, where reoptimization each period may not be feasible, the suggestion is likely to be of less use for economic applications. If sufficient time is available to reoptimize each period, then it is much more straightforward just to solve the deterministic control problem each period.²⁵ The results in this paper certainly indicate that it is feasible to reoptimize each period when, say, the period is a month or a quarter. The procedure of reoptimizing each period is also somewhat more appealing on intuitive grounds than Athans' procedure. If stochastic simulation is ruled out, then both procedures are based on the incorrect practice of setting error terms equal to their expected values. If one follows Athans' procedure, however, further approximations have to be made that do not have to be made if one reoptimizes each period.

²⁵These remarks should not be interpreted as meaning that Athans would necessarily disagree with them. For example, Athans [1], p. 449, has stated: "It should be stressed that trends in stochastic control research by engineers has been greatly influenced by two factors
 (a) a need to minimize on-line computations, and
 (b) the requirements in many aerospace applications that the control system be realized by analog hardware.

In economic applications these requirements are not present, since the time period between decisions does allow for extensive digital computer calculations. Thus, one does have the luxury of examining more sophisticated decision and control algorithms, which however have increased computational requirements."

REFERENCES

- [1] Athans, Michael, "The Discrete Time Linear-Quadratic-Gaussian Stochastic Control Problem," Annals of Economic and Social Measurement, I (October 1972), 449-491.
- [2] Athans, Michael, "The Role and Use of the Linear-Quadratic-Gaussian Problem in Control System Design," IEEE Transactions on Automatic Control, AC-16 (December 1971), 529-552.
- [3] Bard, Yonathan, "Comparison of Gradient Methods for the Solution of Non-Linear Parameter Estimation Problems," SIAM Numerical Analysis, VII (March 1970), 157-186.
- [4] Box, M.J., "A Comparison of Several Current Optimization Methods, and the Use of Transformations in Constrained Problems," Computer Journal, IX (May 1966), 67-77.
- [5] Box, M.J., D. Davies, and W.H. Swann, Non-Linear Optimization Techniques (Edinburgh: Oliver and Boyd Ltd., 1969).
- [6] Chow, Gregory C., "How Much Could be Gained by Optimal Stochastic control Policies," Annals of Economic and Social Measurement, I (October 1972), 391-406.
- [7] Davidon, W.C., "Variable Metric Method for Minimization," A.E.C. Research and Development Report ANL-5990 (Revised), 1959.
- [8] Fair, Ray C., A Short-Run Forecasting Model of the United States Economy. (Lexington: D.C. Heath and Co., 1970).
- [9] Fletcher, R., and M.J.D. Powell, "A Rapidly Convergent Descent Method for Minimization," Computer Journal, VI (July 1963), 163-168.
- [10] Fromm, Gary, and Lawrence R. Klein, "Solutions of the Complete System," in Duesenberry, James S., Gary Fromm, Lawrence R. Klein, and Edwin Kuh, The Brookings Model: Some Further Results (Chicago: Rand McNally & Co., 1969), 362-421.
- [11] Goldfeld, Stephen M., and Richard E. Quandt, Nonlinear Methods in Econometrics (Amsterdam: North-Holland Publishing Co., 1972).
- [12] Goldfeld, Stephen M., Richard E. Quandt, and Hale F. Trotter, "Maximization by Quadratic-Hill-Climbing," Econometrica, XXXIV (July 1966), 541-551.
- [13] Holbrook, Robert S., "Optimal Policy Choice Under a Nonlinear Constraint," mimeographed.
- [14] Howrey, E. Philip, and H.H. Kelejian, "Simulation Versus Analytical Solutions: The Case of Econometric Models," Chapter 12 in Naylor, Thomas H., Computer Simulation Experiments with Models of Economic Systems. (New York: John Wiley & Sons, 1971).
- [15] Huang, H.Y., "Unified Approach to Quadratically Convergent Algorithms for Function Minimization," Journal of Optimization Theory and Applications, V (June 1970), 405-423.

- [16] Kowalik, J., and M.R. Osborne, Methods for Unconstrained Optimization Problems (New York: Elsevier Publishing Co., 1968).
- [17] Murtagh, B.A., and R.W.H. Sargent, "Computational Experience with Quadratically Convergent Minimization Methods," Computer Journal, XIII (May 1970), 185-194.
- [18] Pearson, J.D., "On Variable Metric Methods of Minimization," Computer Journal, XI (May 1969), 171-178.
- [19] Pindyck, Robert S., "Optimal Stabilization Policies via Deterministic Control," Annals of Economic and Social Measurement, I (October 1972), 385-389.
- [20] Polak, E., Computational Methods in Optimization (New York: Academic Press, 1971).
- [21] Powell, M.J.D., "An Efficient Method for Finding the Minimum of a Function of Several Variables Without Calculating Derivatives," Computer Journal, VII (July 1964), 155-162.
- [22] Powell, M.J.D., "A Survey of Numerical Methods for Unconstrained Optimization," SIAM Review, XII (January 1970), 79-97.
- [23] Powell, M.J.D., "Recent Advances in Unconstrained Optimization," Mathematical Programming, I (October 1971), 26-57.
- [24] Shupp, Franklin R., "Uncertainty and Stabilization for a Nonlinear Model," The Quarterly Journal of Economics, LXXXVI (February 1972), 94-110.
- [25] Stewart, G.W. III, "A Modification of Davidon's Minimization Method to Accept Difference Approximations of Derivatives," Journal of the Association of Computing Machinery, XIV (January 1967), 72-83.
- [26] Wolfe, Philip, A review of some notes of Philip Wolfe in Fletcher, R., ed., Optimization (London: Academic Press, 1969), xi-xv.