EVALUATION OF MACROECONOMIC POLICIES BY STOCHASTIC CONTROL TECHNIQUES

by

Gregory C. Chow

Econometric Research Program
Research Memorandum No. 205
December 1976

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

Evaluation of Macroeconomic Policies by Stochastic Control Techniques

Gregory C. Chow [1]

1. Introduction

The purpose of this note is to provide a method, based on the theory of optimal control for stochastic systems, which can be used to evaluate the macroeconomic policy prevailing during any historical period.

A reasonable evaluation of economic performance requires the following four elements. First, some objective or loss function of the relevant economic variables has to be agreed upon. Second, some model governing the economy has to be postulated. Such a model is necessary for the assertion that an alternative course of action could have led, with a reasonable degree of confidence, to a different and better result. Third, in a world of uncertainty, an action should be judged not by the actual outcome which is partly the result of chance (or luck!), but by the probability distribution (often summarized by its mathematical expectation) of the utility or loss resulting from the action. Fourth, if it is recognized that the consequences of decisions are extended to many periods, the action of one period has to be judged not only by the probability distribution of the outcome for that period resulting from it as compared with the probability distribution resulting from an alternative action, but by the distributions of all relevant outcomes in future periods. These four essential elements for policy evaluation will be explicitly taken into account in the proposed method.

Not all of these four elements are generally recognized and incorporated by economists attempting to evaluate government policies. For example, in the interesting paper by Kmenta and Smith [7] , historical fiscal and monetary policies in the quarters from 1954.2 to 1963.4 were evaluated using a (very simple) loss function and a quarterly macroeconometric model, thus incorporating the first two elements listed above. Uncertainly was partly accounted for by evaluating the change in the only target variable, real GNP, due to the changes in the policy variables, net of the random disturbances in the model. However, the multiperiod nature of the decision process was not adequately taken into account since a policy was judged by its effect only on current real GNP without considering its effect on real GNP in the future.

A recent study by Fair [4] on the use of optimal control techniques to measure economic performance does incorporate the above four elements. The main source of differences between the approach proposed here and Fair's approach is that ours is grounded on a theory of optimal control for stochastic systems whereas Fair's is based mainly on open-loop control for deterministic systems. As a consequence, our theory is more comprehensive and the computations required are much simpler. The contrast with Fair's approach will be briefly discussed in section 3, after we have set out our theory in section 2. Section 4 sets forth the limitations of our approach and the directions for improvement and refinement. Section 5 provides an illustrative evaluation.

2. A Theory for Policy Evaluation

The theoretical framework on which we base our proposal for policy evaluation is the theory of optimal control for stochastic systems. See [1, Chapters 8 and 12]. Our proposed method follows directly from the stochastic control theory based on the method of dynamic programming. We will restate the main feature of this theory very briefly under simplifying assumptions, develop the method of policy evaluation from it, and discuss limitations and extensions of the method later on.

Initially we make two simplifying assumptions. First, a quadratic loss function for T periods

$$(2.1)\quad W = \sum_{t=1}^{T} (y_t - a_t)' K_t (y_t - a_t) = \sum_{t=1}^{T} (y_t' K_t y_t - 2y_t' k_t + a_t' K_t a_t)$$

is assumed to measure the performance of a vector of economic variables y_t. Second, the econometric model used is linear, time-dependent, having known coefficients, A_t, b_t and C_t, and a serially independent additive randon disturbance u_t :

$$(2.2) \qquad y_t = A_t\, y_{t-1} + C_t x_t + b_t + u_t$$

where x_t is a vector of control variables which may be incorporated in y_t as a subvector if necessary.

Recall that, using the method of dynamic programming to find the optimal strategy [1, Section 8.1], we first minimize the expected loss for only the last period T with respect to x_T and obtain a linear feedback control equation $\hat{x}_T = G_T\, y_{T-1} + g_T$. We then minimize the sum of the expected losses for the last two periods with respect to x_{T-1}, assuming that the last period policy x_T shall be optimal, i.e., substituting the minimum expected loss for period T into the minimand. Continuing the process backward in time, we finally minimize the sum of the expected losses for all T periods with respect to x_1 of the first period, assuming that x_2, ..., x_T shall be optimal. This sum, after all the future minimum expected losses from period 2 onward have been duly inserted, is the expectation of a quadratic function of the economic variables y_1 for the first period only:

$$(2.3) \qquad V_1 = E(y_1' H_1 y_1 - 2_1' h_1 + c_1)$$

where the coefficients H_1, h_1 and c_1 can be calculated by standard formulas [1, p. 179]. Using (2.2) to substitute for y_1 in (2.3) and taking expectations, we have

$$(2.4) \qquad V_1 = x_1'C_1'H_1C_1x_1 + 2x_1'C_1'(H_1A_1y_o + H_1b_1 - h_1)$$
$$+ (A_1y_o + b_1)'H_1(A_1y_o + b_1) + E\, u_1'H_1u_1 - 2(A_1y_o + b_1)'h_1 + c_1 .$$

Thus, assuming that the policies from period 2 onward shall be optimal, the expected multiperiod loss is a quadratic function of x_1, as given by (2.4). The optimal first-period policy $\hat{x}_1$ is obtained by minimizing (2.4) with respect to x_1, yielding the associated minimum expected multiperiod loss $\hat{V}_1 = V_1(\hat{x}_1)$.

To evaluate any policy x_1 adopted in the first period, we propose to use the difference $V_1(x_1) - V_1(\hat{x}_1)$. To evaluate a sequence of policies $x_1, x_2, \ldots, x_N$ $(N < T)$, we propose to use the sum of the differences

$$(2.5) \qquad \sum_{t=1}^{N} [V_t(x_t) - V_t(\hat{x}_t)].$$

Note that the function $V_2(x_2)$ measures the total expected loss from period 2 to period T if x_2 is adopted for period 2 and optimal policies shall be followed from period 3 onward.

The rationale of the proposed measure is as follows. At the beginning of period 1, if the optimal feedback policies $\hat{x}_t = G_ty_t + g_t$ are followed sequentially for $t = 1, 2, \ldots, T$, the minimum expected loss for all T periods is $V_1(\hat{x}_1)$. However, if a nonoptimal policy x_1 is chosen for the first period, but optimal policies are chosen from period 2 on, the total expected loss will be $V_1(x_1)$. The difference $V_1(x_1) - V_1(\hat{x}_1)$ measures the extra expected loss which is attributable to the nonoptimal policy x_1. At the beginning of period 2, again if the optimal feedback policies are followed until period T, the minimum expected loss is $V_2(\hat{x}_2)$. Since $V_2(x_2)$ is given by (2.4) with all subscripts increased by 1, it is a function of y_1. Whatever policy x_1 was chosen for the first period, it has affected the initial condition y_1 which in turn influences the total expected loss from period 2 on. But bygone is bygone. The best that one can do now is to choose $\hat{x}_2 = G_2 y_1 + g_2$. If x_2 is chosen instead, the extra expected loss due to the nonoptimal policy in period 2 is $V_2(x_2) - V_2(\hat{x}_2)$. This argument applies to many future periods, provided that the number of periods N chosen for policy evaluation is smaller than the planning horizon T used in the optimal control calculations.

The proposed measure can be applied to a nonlinear econometric model. Since a truly optimal solution to control a nonlinear model with random disturbances is not available, we employ the approximately optimal solution of Chow [1, Chapter 12] and [2]. This amounts to finding the optimal path for the nonstochastic control problem formulated by setting the random disturbances equal to their expected values, and then linearizing the model around this path to yield a linear model with time-dependent coefficients as given in (2.2). We then follow the procedure recommended above.

3. Comparison with Fair's Approach

We basically agree with the logic of Fair [4] in measuring the economic performance of several political administrations in the United States. However, there are several differences to be explained below.

Fair's measure for one four-year administration equals (1) the expected loss during the four-year interval using the policy actually adopted, minus (2) the expected loss during the four-year interval using an optimal policy, plus (3) the expected loss in the next four-year period if the next administration behaved optimally given that the present administration did not, minus (4) the expected loss in the next four-year period if both administrations had behaved optimally. The difference between the components (3) and (4) is intended to measure the damage done by the administration being evaluated to the next administration.

To see the similarity between Fair's measure and ours, consider the case where $T = 32$ quarters and $N = 1$. That is, the performance of only the first quarter is to be evaluated, or, in Fair's terms, the administration lasted only for the first quarter. The sum of his terms (1) and (3) would measure the expected loss for all T periods if the actual x_1 is used for period 1 and optimal policies will be followed from period 2 on. Our $V_1(x_1)$ is intended to measure the same. The sum of Fair's terms (2) and (4) would measure the expected loss for all T periods if optimal policies were followed throughout. So does our $V_1(\hat{x}_1)$.

From the viewpoint of theory in the general case of policy evaluation for many quarters, there are three differences between our approaches. First, when Fair measures the expected loss for many periods resulting from the actual historical policy, he uses the open-loop approach. In other words, the decision maker is assumed to have decided on the actual values of the policy instruments at the beginning of the first quarter of his administration. This is the assumption on which Fair's first term is evaluated. According to our theory, the policy maker was free at the beginning of each quarter to choose his policy. When the policy for quarter 2 was chosen, for example, the actual economic data y_1 were known to the decision maker and were used to compute $V_2(x_2;y_1)$ in our measure $V_2(x_2;y_1) - V_2(\hat{x}_2;y_1)$ of economic performance for quarter 2, which is the first period in the optimal control problem involving T-1 periods. We believe that it is more realistic to assume that actual decisions were made sequentially quarter by quarter and not made once for all four years at the beginning of each administration. This difference would disappear if we were to evaluate an administration lasting only for one quarter.

Second, when Fair measures the expected loss for many periods resulting from the optimal policy, his expectation is conditional on the particular realization of the random residuals in the equations as they occured in history.

By contrast, our $V_1(\hat{x}_1)$ is not conditional on the historical observations on the equation residuals. To evaluate the multiperiod expected loss from an optimal feedback policy by stochastic simulation, one would have to use repeated random drawings of the residuals through time, and, for each set of residuals, perform the same calculations which Fair applies to the set of historical residuals. The resulting losses would have to be averaged to obtain the required expectation. It should be noted that Fair also follows the feedback approach in evaluating the optimal policy since he makes the optimal policy of period two dependent on the residuals occuring in period 1. However, his expected loss is based on only one observation on the random residuals.

Third, the theory of section 2 provides a measure of economic performance for each quarter (if a quarterly model is used) without having to break up the evaluation into four-year or whatever time intervals. Fair's measure can be applied to any time interval as desired, but the computations would have to be redone when the time interval changes.

From the viewpoint of computations, the method of section 2 is much simpler than Fair's. Once we have obtained the (approximately) optimal control solution for the T period problem in the form of feedback control equations, the functions $V_t(x_t;y_{t-1})$ are available. Our measure merely requires the evaluation of $V_t(x_t)$ and $V_t(\hat{x}_t)$. The computations are extremely trivial using any modern computer. On the other hand, the expected values defined in the four components of Fair's measure all require extensive stochastic simulations. The cost of performing some of these stochastic simulations was thought to be so high that Fair did not actually compute the measure as he first defined it, but rather used an approximation based on open-loop control of a deterministic model. The computer program described in Chow [2] can handle a nonlinear model of eighty simultaneous equations. We have used it to control the

University of Michigan Quarterly Econometric Model of 61 simultaneous equations for 17 quarters. The (approximately) optimal feedback solution based on linearization as described above costs about \$54 on the IBM 360-91 computer at Princeton University; it takes three rounds of linearizations to converge, with each round costing about \$18. A model 1.5 times as large would cost not more than $(1.5)^2$ or 2.25 times as much. An illustrative calculation using the Michigan model will be provided in section 5.

The evaluation of mathematical expectations by stochastic simulations always involves sampling errors which our analytical method using $V_t(x_t)$ avoids. However, our method is not truly optimal. Our feedback control equations $\hat{x}_t = G_t y_{t-1} + g_t$ are not truly optimal because they are based on linearization about the optimal path derived from the deterministic control problem and the actual path will most probably deviate from this path. A different point used to perform the linearization would yield different coefficients A_t, C_t and b_t for the model, and thus G_t and g_t for the feedback control equation, but the differences will not be large if the model is not highly nonlinear or if the actual path is not very far away from the above optimal path. To obtain an (approximately) optimal policy for period t, Fair minimizes the expectation (calculated by stochastic simulations) of a multiperiod loss function resulting from open-loop, rather than feedback, policies. His solution, besides incurring sampling errors due to stochastic simulations, is not optimal either, since open-loop policies are not optimal when random disturbances are present in an econometric model.

4. Limitations and Extensions

While we believe that the method of section 2 can fruitfully be applied to the operating nonlinear econometric models for the purpose of economic policy evaluation, we mention several obvious extensions of the method when some assumptions are relaxed. First, the loss function may not be quadratic and the control variables may be subject to inequality constraints. To deal with these problems partially, we propose to incorporate the above assumptions in deriving the optimal path for the deterministic control problem formulated by setting the random disturbances equal to their expected values, and then linearize the model about this path to obtain approximately optimal linear feedback control equations. Methods to solve the deterministic control problem required are given in Friedman [5], Chow [1, p. 284] and Gordon and Jorgenson [6]. Second, to deal with uncertainty in the parameters of a nonlinear econometric model, we can use the method of Chow [3] by deriving an approximate probability distribution for the coefficients A_t, b_t and C_t in the linearized reduced form from the assumed distribution of the parameters in the original structural equations, and obtaining the expectation $V_t(x_t)$ from the former distribution.

A computationally more difficult problem is to deal with the possible revision of the estimates of the parameters of the econometric model through time. To anticipate future revision of the parameter estimates in deriving an optimal policy for the current period, i.e., to allow for active learning, is too difficult, and probably not worth the trouble. To allow for passive revision of the parameter estimates as time passes, which Fair also mentions in his paper [4], we would have to use a new set of parameter estimates for each period t in computing $V_t(x_t)$. This means that, for each period t, a different multiperiod control problem from t to T has to be solved, each with a new set of coefficients A_s, b_s and C_s ($s = t, \ldots, T$) which result from the corresponding set of estimates for the structural parameters. One may question the value of the extra computations involved, but to solve a multiperiod control problem (from t to T) for each period t to evaluate

$V_t(x_t)$ may be worthwile in order to allow for the possible lack of information available to the decision maker concerning the future values of the exogenous variables outside his control. Our method incorporates the combined effect of these exogenous variables in the intercept b_t of (2.2). Rather than using the actual values of these variables after they have occurred, which the policy maker at the time of decision did not know, one may use the official forecasts at the time, and the resulting b_t, to calculate $V_t(x_t)$. Uncertainty in b_t can be treated by using a probability distribution, as pointed out in the last paragraph. The forecasts of the exogenous variables may turn out to be very poor, but the proposed method of policy evaluation is valid if the policy makers are not to blame for utilizing their imperfect knowledge to behave optimally.

5. An Illustrative Evaluation

To illustrate the method of section 2, the Michigan Quarterly Econometric Model [7] has been used to calculate the function $V_1(x_1)$. The optimal control problem solved is the one posted by the participants of the NSF-NBER Econometric Model Comparison Seminar directed by Lawrence Klein. That is, the number of periods is 17, covering the quarters from 1971.1 to 1975.1. The objective is to minimize the loss function

$$\text{(5.1)} \quad \sum_{t=1}^{17} \left[(\dot{p}_t - \alpha_t)^2 + .75(u_t - 4.0)^2 + .75(\text{GNP Gap}_t - 0)^2 + (TB_t - 0)^2 + .1(UR\$_t - \gamma_t)^2 \right]$$

where $\dot{p}$ is the annual rate of inflation measured by the GNP deflator, $\alpha_t = 3.0$ for $t=1, \ldots, 12$, $\alpha_t = 7.0$ for $t=13, \ldots, 17$, u is the unemployment rate, GNP Gap is the percentage deviation of GNP in 1958 dollars

(GNP58) from capacity output, TB is trade balance as a percentage of GNP in current dollars, UR$ is unborrowed reserves in billions of current dollars, and γ_t represents a smooth expansionary path for UR$. The policy variables set up for our computations are unborrowed reserves UR$ and nondefense government purchases of goods and services GFO$ both in billions of current dollars. The UR$ term in the loss function serves to prevent erratic behavior of the monetary instrument. The estimated residuals in the structural equations were added back to the intercepts, at the suggestion of the Seminar participants, resulting in a deterministic control problem, although our method can handle a stochastic control problem as well.

The quadratic function $V_1(x_1) = x_1'Qx_1 + 2x_1'q + d$ is given by the following matrix Q and vector q

$$(5.2) \qquad Q = \begin{bmatrix} .103{,}013 & .000{,}999 \\ .000{,}999 & .008{,}614 \end{bmatrix} \qquad q = \begin{bmatrix} -2.7954 \\ -0.7011 \end{bmatrix}$$

The vector of control variables in period 1 which minimizes V_1 is

$$(5.3) \qquad \hat{x}_1 = [\ 26.3770 \qquad 78.3233\ ]$$

We will evaluate (5.2) at the historical values of the control variables which are 29.5 and 24.1 respectively. Note that the second figure differs a great deal from the optimum value 78.3 given in (5.3). At these historical values, the multiperiod loss V_1 given by (5.2) is 26.00 higher than its minimum.

To examine the figure 26.00 more closely, we will compute the difference between the first-period losses of the two policies. The figure 26.00 is composed of this difference and the remainder which measures the extra loss from period 2 to T attributable to the historical (nonoptimal) policy as it affects the initial economic condition at the end of period one. The optimum solution values and the historical solution values of the variables entering the loss function (5.1) for period one (1971.1) are

	$\dot{p}_1$	u_1	GNP Gap_1	TB_1	$UR\$_1$
Optimum	6.07	5.08	.48	-.182	26.4
Historical	4.69	5.95	5.01	.278	29.5

The first-period losses resulting from these two solutions are respectively 11.47 and 24.61, the difference being 13.14. (Note that the target value for $UR\$_1$ was set at its historical value 29.5, accounting for a contribution of $.1(26.4 - 29.5)^2 = .961$ to the first-period loss of the optimum policy). The main contribution to the first-period loss of the historical policy is from the GNP Gap , equal to $.75(5.01)^2 = 18.83$. Although the loss function (5.1) weighs the inflation term by 1 and the unemployment term by only .75, it penalizes low output quite heavily through the GNP Gap. In short, 13.14 or about half of the extra multiperiod loss 26.00 is allotted to the difference between the first-period losses. The assumption is that, no matter whether the first-period policy is optimal or not, the policies from period 2 on shall be optimally chosen. The nonoptimal policy in period 1 adds 13.14 to the loss in period 1 itself, as computed from the two sets of solution values of the variables included in the loss function shown above. It adds an almost equal amount to the total loss from period 2 to period 17, assuming the policies in these periods to be optimal. The 26.00 figure can also be compared with the total loss of 248.03 for all 17 periods if optimal policies had been followed throughout. The same analysis can be applied to period 2 and other future periods, but since our time horizon T is only 17, the number of quarters N to be evaluated should be much smaller than 17 to avoid problems associated with the terminal conditions.

Footnote

1. I would like to thank Ray Fair for comments on section 3 and to acknowledge financial support from the National Science Foundation in the preparation of this paper.

References

[1] Chow, G.C., Analysis and Control of Dynamic Economic Systems, (New York: John Wiley and Sons, 1975).

[2] __________, "An Approach to the Feedback Control of Nonlinear Econometric Systems," Annals of Economic and Social Measurement, V (Summer, 1976), 297-309.

[3] __________, "The Control of Nonlinear Econometric Systems with Unknown Parameters," Econometrica, XLIV (July, 1976), 685-695.

[4] Fair, R.C., "The Use of Optimal Control Techniques to Measure Economic Performance," Cowles Foundation Discussion Paper No. 420 (Yale University, January, 1976).

[5] Friedman, B.M., Methods in Optimization for Economic Stabilization Policy (Amsterdam: North-Holland Publishing Company, 1973).

[6] Gordon, R.H. and D. Jorgenson, "The Investment Tax Credit and Counter-Cyclical Policy," in Otto Eckstein (ed.), Parameters and Policies in the U.S. Economy (Amsterdam: North-Holland Publishing Company, 1976), 275-314.

[7] Hymans, S.H. and H. Shapiro, "The Michigan Quarterly Econometric Model of the U.S. Economy," The Economic Outlook for 1973, The University of Michigan, Ann Arbor, Michigan, 1973.

[8] Kmenta, J. and P.E. Smith, "Autonomous Expenditures Versus Money Supply: An Application of Dynamic Multipliers," The Review of Economics and Statistics, LV (August, 1973), 299-307.