Introduction to Stochastic Control Theory and Economic Systems

Michael Athans*
Gregory C. Chow**

ECONOMETRIC RESEARCH PROGRAM
Research Memorandum No. 143
June 1972

*Massachusetts Institute of Technology

**Research supported by NSF Grant GS 32003X. This paper will appear in Annals of Economic and Social Measurement, Vol. 1, No. 4 (October 1972).

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

Introduction to Stochastic Control Theory and Economic Systems

by

Michael Athans and Gregory C. Chow*

A workshop on "Stochastic Control Theory and Economic Systems" was held on May 5th and 6th, 1972, at Princeton University. It was under the sponsorship of the National Bureau of Economic Research, with financial support from the National Science Foundation and the International Business Machines Corporation. Approximately 40 economists and 25 control theorists attended. A list of participants is included in Appendix A of this introduction. Nine invited papers were presented, six by economists and three by control scientists. These nine papers, either in their entirety or some abstract form after revisions, have been collected in this volume.

To introduce these papers it would be helpful to point out the nature and the objective of the workshop. Early in January 1972, the planning committee of this workshop, consisting of Michael Athans, Gregory C. Chow, Edwin Kuh, and M. Ishaq Nadiri, met to discuss the general plans for the workshop.[1] We felt that the development of stochastic dynamic economics using tools related to optimal stochastic control had reached such a point that it would be extremely useful to bring together research workers from both the economics and control

---

professions to report on current research work, to exchange ideas, and to evaluate the prospects for future developments. In addition, recent trends in control theory research deal with the fundamental understanding of large scale systems and decentralized decision-making; clearly, such problems are common in economic systems. Hence, exchange of ideas in this area was also judged to be of importance. Michael Athans and Gregory C. Chow were asked to serve as co-chairmen of this workshop to invite interested participants and to select the papers to be presented, with Athans responsible for inviting the control scientists and Chow responsible for inviting the economists,[2] while Kuh and Nadiri were constantly providing ideas and suggestions.

We thought that it would be useful to have the economists present their papers first, and in the second day of the conference, to have the control theorists present three state-of-the-art papers in their field which would be of interest to the group. There were definitely other interesting and relevant papers by economists than the six actually selected, and now published in this volume. One of the criteria for selecting these six papers was their continuity and coherence, aspects that would facilitate understanding by the participants of the workshop, and would provide convenient focal points for discussion and exchanges by the participants and, it is to be hoped, also by readers of this volume.

The first three papers, besides possible substantive contributions that they may make to the analysis of economic policy, contain discussions

---

and expositions of some basic tools and ideas of stochastic control for readers who are not already familiar with the subject.[3] The first paper, by Robert S. Pindyck, applies deterministic control theory to study the optimal time paths for the policy variables, using a linear econometric model of the United States economy that the author has constructed. The welfare function used is quadratic. It penalizes the performance of an economic variable (if it is selected in the welfare function) by the sum of squares of its deviations from the target values during the discrete time periods (quarters in this case) of a finite planning horizon. The derivation of the optimal time paths in feedback form is based on deterministic control theory - deterministic because the random disturbances in the econometric model were ignored (treated as having values zero) and because the parameters in the linear system are assumed to be given constants not subject to uncertainty. Pindyck's analysis illustrates what target paths, especially for the rates of unemployment and of inflation, are feasible under the assumptions of his econometric model. In other words, if one specifies a target for unemployment of two percent during the planning period, the optimal solution may give unemployment rates of between three to four percent approximately, but also an inflation rate of, say, over five percent per year. Thus, his analysis provides empirical measures of the trade-off between unemployment and inflation, a relationship well-known in the economic literature as the Phillips' curve, named after the control engineer and economist, A.W. Phillips.[4]

The second paper, by Gregory C. Chow, contains an elementary exposition of the simplest case of stochastic control theory. Unlike Pindyck's paper, this paper incorporates the effects of random disturbances in the econometric model, and the welfare function becomes the mathematical expectation of a weighted sum of squares of deviations of the variables under control (now stochastic time series) from their target paths. A main substantive problem of this paper is to measure the difference between the welfare cost, as defined by the above mathematical expectation, for an optimal policy and the cost for the best policy that maintains a constant rate of growth for each policy variable. The latter policy is a deterministic policy in the sense that it can be specified at the beginning of the planning period without observing future occurrences of the economic time series in question as generated by the random process. Hence, the difference in welfare costs can be decomposed into the difference between the optimal policy using stochastic control and the best deterministic policy, plus the difference between the best deterministic policy and the suboptimal deterministic policy which is required to have each policy variable follow a constant rate of change. The econometric model used is a simple macro model constructed by Chow using annual data of the United States for 1931-1940 and 1948-1963. The results, for different welfare functions chosen and for both variables in their levels and in first differences, show that the stochastic part of the gain, or difference, is much larger than the deterministic part, suggesting the importance of incorporating the random disturbances in an econometric model for the purpose of measuring welfare gains from an optimal stochastic control policy.

The third paper, by Stanley Fischer and J. Phillip Cooper, studies the relationship between the lag structure and the effectiveness of certain stabilization policies. The models used are linear and rather simple, but they do include stochastic disturbances, and some parameters are assumed to be random in parts of this study. The policy rules studied are simple, proportional and derivative feedback rules, rather than fully optimal feedback policies. Effectiveness of each policy is measured mainly by the variances around stable paths. The optimal parameters in the proportional feedback equations were obtained by numerical techniques. The relationships of these optimal parameters to the mean length and variability of the lag structure were investigated. This study concludes that, in general, the longer are the lags, the more active should stabilization policy be, under the assumption that the parameters in the lag structure are given constants. Another finding is the obvious deterioration in performance of the feedback rules when the lagged effects of the policy instrument were both long and variable. It is interesting to note that this type of problem has been under general theoretical investigation in the control literature; it is commonly refered to as the "stochastic stability problem."

The fourth paper, which is the first in the second group of papers presented by economists,[5] is a study of stabilization policy using the recursive model of the Federal Reserve Bank of St. Louis and is contributed by H. Woods Bowman and Anne Marie Laporte. A recursive model has the property that, although one dependent variable may be explained by another dependent variable in a stochastic difference equation, the latter dependent variable is not explained by the former in another stochastic equation. Recursiveness simplifies the mathematical derivation of optimal control policies which are confined to one-period policies in this paper, but the analysis allows for randomness in the parameters of the recursive system.[6] This paper is of interest because it employs an econometric model which has received quite a bit of attention in the study of monetary and fiscal policies, and because it attempts to measure how much more conservative a policy should be - conservative being to react less to the changing economic data of recent past - when the parameters of the model are subject to uncertainty.

While the Bowman-Laporte paper treats a planning horizon of only one period, the next paper by Edward C. Prescott treats planning problems for many periods but only when one parameter, namely the slope of a simple linear model explaining the current dependent variable by one control variable, is considered to be uncertain. Although he used a simple model, Prescott was able to study the trade-off

---

between the control purposes of steering the dependent variable to target in the present and of increasing the accuracy of the estimated parameters in order to further the objective of control in the future. This study contains both analytical and simulation results. It was pointed out in the discussion that a more interesting, and still very simple, model would result from introducing the lagged dependent variable into the system. A participant, Karl Astrom of Sweden, discussed some of his work on this and related problems.

The next paper, by Elizabeth Chase MacRae, is not only multiperiod in formulation but also allows for randomness in the parameters of a linear model. The minimization of expected quadratic welfare under such circumstances is a well-known problem in stochastic control theory. A completely analytical solution that can be numerically implemented is not yet available, and this paper provides one of the approximate solutions. In this approximation, the state variables (which may include both dependent and control variables) at a certain future period, which are of course random, are replaced by their expected values in a certain stage of the calculations. This paper assumes that only the random coefficients in a linear system are unknown, but the residual disturbance of the system has a known variance. Therefore, using the Bayesian rule, the posterior distribution of the random coefficients in each period, which is proportional to the product of the likelihood and the prior distribution of the same parameters a period earlier, both of which are normal, will be multivariate normal for each period. The equations given by MacRae are readily interpretable in terms of price of information and value of estimating, and can be implemented, but how good the approximation is remains an open question. Again, it is interesting to note that this type of problem has received quite a bit of attention in the control literature dealing with adaptive control; the specific mathematical technique goes under the name of "open-loop-feedback-optimal" control.

The first state-of-the-art paper is by Michal Athans. As the author would call it, it is a "bread and butter" paper in the implementation of control theory by its engineering practitioners. When dealing with a possibly nonlinear stochastic model, when the welfare function may not be quadratic, and when the parameters of the system may be random, the author suggests solving the problem of optimal control in three stages. In the first, all the random disturbances or even parameters are replaced by their expected values and the converted, deterministic problem will be solved by whatever means available, such as Pontryagin's minimum principle or nonlinear programming. The second stage deals with the construction of estimates of the state variables, using Kalman-Bucy filtering techniques, on the basis of noisy (uncertain) measurements. The third stage consists of solving the linear-quadratic-Gaussian problem of steering the deviations of the variables from the optimal paths as determined by stage one to zero, after the dynamic system is linearized around the above optimal path to yield a linear system with time-varying, but non-random, parameters, utilizing the Kalman filter estimates. This appears to be a reasonable first approximation to optimal control when the model is non-linear, when the cost may not be quadratic, and when the parameters may be random. A second approximation, as suggested by Charles Holt, would be to replace the above deterministic optimal path by the mathematical expectation of that path, the reason for the suggestion being that the expectation of a nonlinear function of random variables, which is what is involved, is not equal to the function of the expectations. While it was argued by others that such expectations might not be easily computed, Gregory Chow suggested that Monte Carlo techniques be used to generate samples of random sequences (in place of the expectations of the parameters) to solve the deterministic problem and that these experiments be replicated in order to obtain an estimate of the mean path to be employed in the second stage.

The second paper of the same session, by Hans S. Witsenhausen, is concerned with the separation of estimation and control for discrete time systems. One important difference between the basic model of the control engineer and that of the economist is that the former assumes that the state variables are subject to measurement errors (the structure of the errors partially known, at least) whereas the economist, more often than not, ignores the possible errors of measurement. This is not the place to discuss this important issue except to point out its existence, and to say that perhaps greater interest by economists in errors of measurement in the context of control may be desirable; in fairness to the economists, though, the incorporation of measurement errors might often be more difficult in economic applications where knowledge of the error structure cannot be assumed.

The term "estimation" as used in the control literature refers not to the estimation of parameters as in classical econometrics or classical statistics, but to the estimation of the value of a random vector of state variables which cannot be observed directly. If the state variables are not directly observed, and therefore the problem of its estimation exists, one might choose to separate the solution to the optimal control problem, if such a solution exists, into two parts, one to estimate the state variables and the second to apply certain control rules to the estimated values of the state variables. Under the assumptions that the model is linear with known parameters but unknown additive Gaussian random disturbances and that the welfare function is quadratic, it is well-known that the optimal control solution can be separated into two parts, the first being the estimation of the unknown state by its conditional expectation through the Kalman-Bucy filter and the second being the application of linear feedback control to these expected values. This may be considered an example of the <u>certainty equivalence principle</u>

---

in which the optimal solution is obtained by replacing certain random and unknown values by their conditional expectations. It is also well-known that if the parameters are random, then the replacement of random values by their expectations in the optimal control equations would no longer yield an optimal solution. Witsenhausen's paper defines clearly the different aspects of the separation problem and states some sufficient conditions under which different aspects of separation will exist. In his oral presentation, Witsenhausen further illuminated the difference between the problem of statistical decision, in which the state of the world is assumed to exist though unknown, and the optimal control problem in which the purpose is to change the state in a certain optimal way.

The last invited paper, by Pravin Varaiaya, discusses trends in the theory of decision-making in large scale systems. It is a summary of some of the literature dealing with the theoretical aspects of decision-making in large systems. It attempts to survey three related lines of development. The first is team theory, contributed to a large extent by Jacob Marschak and Roy Radner. In contrast with statistical decision theory, team theory brings in the information structure, namely, a mapping from the state of the world to the observations by any one of the agents of an organization, as a subject for analysis, whereas in statistical decision theory, the information structure is usually assumed to be known, and the only remaining task is to determine the optimal decision functions relating observations to actions. The second line is competitive equilibrium theory, which is mentioned only briefly, and the third is research in organization form arranged in a hierarchy. The last is exemplified by some contributions from works in operations research, and by the formulation of an optimal control problem using a linear system in which there are different agents each of whom may employ certain feedback control rules on observations of that part of the state of the system which is available to him. The last obviously is related to the problem of decentralized decision-making for the purpose of controlling a large system.[7] This survey paper is interesting because it brings together a few important and related areas of research, and it is somewhat broader in scope than most of the previous papers.

Having introduced the nine papers of this volume, we wish to point out that there are other interesting research works dealing with stochastic control and economic systems which were briefly reported by their authors in an open session held in the afternoon of May 6th.[8] David Kendrick has been applying the tools of stochastic control to the problem of stabilization of the international cocoa market. Benjamin Friedman reported briefly on his research in extending the approach to economic policy of H. Theil to the case where the welfare function may not be quadratic but is approximated by several quadratic segments. Alfred L. Norman discussed the technique of generating a number of optimal paths for a model as a means of checking the reasonableness of the estimates of its parameters. Franklin R. Shupp described briefly

---

his experience in the Federal Reserve Board on how monetary policies are actually made; he also mentioned briefly his work on stabilization using a nonlinear model.[9] Since many people in the control group asked how econometric models are built and how good they are, Ray C. Fair responded with brief remarks based on his experience with the Fair Model.[10]

The question of how useful stochastic control theory is for economics was subject to lively discussions and comments at the workshop. Without reporting the discussion in full detail, we would like to mention that, in the pessimistic extreme, L. A. Zadeh argued that, because of the imprecise nature of economic modeling and the making of economic decisions, it would be less useful to deal with quantitative systems and decisions as we have formulated them in the economics literature, than to use the "fuzzy sets" which he recommends.[11] Zadeh admitted, however, that his formulation of fuzzy sets does not include the problem of estimating either the state or the structure of an economic system in order to acquire knowledge for the purpose of control. On the more optimistic side were the comments by Charles Holt. According to Holt, most of the economic decisions actually made consist of examining various actions together with their associated outcomes, and choosing that action which would generate the most desirable outcome. This primitive way of decision-making could be improved upon by using an econometric model for the purpose of examining various decision rules and their associated

---

consequences in terms of the time paths of the variables (deterministic or random) so generated. The second improvement, according to Holt, would be to specify some welfare functions to generate certain optimal decision rules, not necessarily because we believe that certain welfare functions can be agreed upon, but because the ad hoc decision rules examined by the previous method might not be optimal for any reasonable welfare functions, and because some better rules could be discovered by the optimal control approach which would otherwise remain unnoticed by the preceding method.

The tools of optimal control can be used to study the dynamic responses of the system, and how good the performances are, under various assumptions concerning the welfare function and the parameters of the model. It is our opinion that, while there are many problems in theory and in computations that are shared by both the control theorists and the economists, the uses of control by the engineer would likely be different from the uses by an economist. The former may actually apply the control rules to an operating system, but the latter, at least in the present state of econometric knowledge, is more likely to use the tools to study qualitatively the dynamic characteristics of an economy as they may respond to certain policy rules. It is hoped that the papers of this volume will contribute to the understanding of the ideas and problems involved, and thus contribute to better knowledge of the economic structure and to better economic policy.

FOOTNOTES:

* We would like to acknowledge financial support from the National Science Foundation in the preparation of this paper, and to thank M. Ishaq Nadiri and Richard E. Quandt for editorial comments on an earlier draft.

[1] We would like to use this opportunity to thank Neville Beharie of the National Bureau of Economic Research for having kept the workshop running, and to thank Grace B. Lilley of the Econometric Research Program of Princeton University for assisting in local arrangements as well as other aspects of the running of the workshop.

[2] In this connection, Chow received helpful suggestions from David Kendrick and would like to express his appreciation.

[3] Michael C. Lovell served as chairman of the first session in which these three papers were presented.

[4] A.W. Phillips, "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," Economica, Vol. 25 (1958), pp. 283-299.

[5] George G. Judge served as chairman of this session.

[6] If the parameters of the reduced-form of a fully simultaneous system are regarded as random, and if one-period optimal policies are desired, the method indicated at the end of the paper by Chow in this volume can be applied. In fact, Chow's method is also applicable

to finding multiperiod optimal control policies, provided that one is willing to ignore the effect of observations during the control process on the uncertainty (or posterior density) of the parameters. The papers by Prescott and by MacRae to be introduced in the following paragraphs do not ignore this effect.

[7] For a survey of some of these problems using deterministic control theory in continuous time, the reader may refer to Edwin Burmeister and A. Rodney Dobell, "Guidance and Optimal Control of Free-Market Economies: A New Interpretation," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-2, No. 1 (January 1972), pp. 9-15.

[8] Michael D. Intriligator served as chairman in the open session.

[9] See Franklin R. Shupp, "Uncertainty and Stabilization for a Nonlinear Model," Quarterly Journal of Economics, Vol. LXXXVI, No. 1 (February 1972), pp. 94-110.

[10] Ray C. Fair, A Short-Run Forecasting Model of the United States Economy (Lexington, Mass: D. C. Heath and Company, 1971). Quarterly forecasts based on this model have been released since July, 1970.

[11] Zadeh offered any interested participant of the workshop copies of his reprints, including, among others, L. A. Zadeh, "Fuzzy Sets," Information and Control, Vol. 8, No. 3 (June 1965), pp. 338-353; R. E. Bellman and L. A. Zadeh, "Decision Making in a Fuzzy Environment," Management Science, Vol. 17, No. 4 (December 1970),

B-141 - B164; S.S.L. Chang and L.A. Zadeh, "On Fuzzy Mapping and Control," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-2, No. 1 (January, 1972), pp. 30-34.

APPENDIX A: List of Participants (omitted)