Annals of Economic and Social Measurement, 1/4, 1972

OPTIMAL STABILIZATION POLICIES VIA DETERMINISTIC CONTROL

BY ROBERT S. PINDYCK

Economic stabilization policy is defined in terms of a linear-quadratic tracking problem. The implications of alternative cost functionals and non-linear models for planning are then discussed in terms of computational simplicity and improved performance of models.

1. INTRODUCTION

The development of macro-econometric models over the past two decades has provided a vehicle for studying the simultaneous time-dependent relationships between economic variables and their response over the short term to policy instruments, such as tax rates and the money supply. The more recent availability of computational algorithms for the efficient solution of sets of simultaneous difference equations has made the computer simulation of econometric models a particularly useful way to determine and compare the dynamic effects of different economic stabilization policies, and to test policies and weed out those whose economic effects would be undesirable. Although it is an extremely useful tool for the planning and analysis of stabilization policies, simulation does not provide a direct means of obtaining a policy that is optimal with respect to a fixed set of objectives.

In recent years, there has been an interest in optimal control theory as a possible tool for economic planning. Given an econometric model that one is willing to accept as a reasonable and fixed representation (at least over the short term) of the economy, and given a cost (objective) functional that represents the goals and objectives of economic stabilization, then the design of a stabilization policy can easily be thought of as a problem in deterministic optimal control.

Of course, econometric models are not really deterministic systems. Each equation has an implicit additive error term associated with it, and every estimated coefficient is itself a random variable. Most of our work, however, is done under the guidance of simplification, and so we might choose to ignore the stochastic properties of the model or else make the necessary simplifying assumptions about them that would allow us to invoke "certainty equivalence" [9] in obtaining a solution. Certainly, a deterministic treatment of the optimal stabilization problem is simpler in many respects to a stochastic treatment. To judge its usefulness and adequacy, we will have to examine some of the results which are now becoming available.

Most, though not all, of the recent work in applying deterministic optimal control to economic stabilization policy assumes that the econometric model is either linear or else has been linearized, and it is not always clear as to how appropriate this assumption is. The cost functionals that have been used have generally been quadratic or quadratic-linear in structure. Such a formulation is perhaps somewhat restrictive but results in an optimization problem that is mathematically tractable.

As an example of the use of deterministic control in the formulation of optimal stabilization policies, we will discuss the treatment of stabilization policy as a linear-quadratic tracking problem in optimal control. We will then briefly discuss alternatives to the linear-quadratic formulation in the context of some other recent applications of deterministic control to stabilization policy. Finally, we will make some remarks about the use of deterministic optimal control as a practical tool for the planning and analysis of stabilization policies.

2. Economic Stabilization Policy as a Linear-Quadratic Tracking Problem

Economic stabilization policy can be approached as a deterministic optimal control problem that involves tracking as closely as possible nominal state and nominal policy trajectories, subject to a quadratic cost functional and the constraint of a linear system. This formulation is actually quite general, and allows for penalization for variations in, as well as the levels of, the state variables and control variables.

The deterministic system is of the form

$$x_{i+1} - x_i = Ax_i + Bu_i + Cz_i \tag{1}$$

with a given initial condition

$$x_0 = \xi \tag{2}$$

Here x_i is the n-dimensional state vector at time i, u_i, the r-dimensional control vector at time i, and z_i, an s-dimensional vector representing, at time i, s exogenous variables which are known for all i but cannot be controlled by the policy planner. A, P, and C are known $n \times n$, $n \times r$, and $n \times s$ matrices. Note that n, the number of state variables, will generally be larger than the number of endogenous variables since the structural form of the model will usually contain difference equations of order greater than one.

We then define $\hat{x}_i$ and $\hat{u}_i$ as the nominal (ideal) state and control vectors that we would like to track, and we assume that they have been specified for the entire planning period. The nominal time paths for variables such as GNP and investment, for example, would probably grow at some steady rate, while that for unemployment might drop and then remain low for the remainder of the planning period. The control variables themselves cannot be manipulated in any way whatsoever, but must also stay close to a set of nominal or "ideal" time paths. For example, it is probably undesirable for government spending or the money supply to increase by 100 percent in one year and decrease by 200 percent in the next year. Manipulating policy variables has real costs associated with it, and these costs must be embodied in the cost functional

The quadratic cost functional then is given by:

$$J = 1/2 \sum_{i=0}^{N} \{(x_i - \hat{x}_i)'Q(x_i - \hat{x}_i) + (u_i - \hat{u}_i)'R(u_i - \hat{u}_i)\} \tag{3}$$

where Q is an $n \times n$ positive semi-definite matrix, and R is an $r \times r$ positive definite matrix. The diagonal elements of Q, some of which may be zero, give the relative costs for deviating from the nominal path of each state variable—for example, the cost of deviating from nominal GNP relative to the cost for deviating from nominal unemployment. The diagonal elements of R (all of which must be non-zero), give the relative costs for deviating from the nominal paths of the control variables; for example, we would expect it to be more costly to manipulate the tax rate than to manipulate the money supply. Finally, the comparative magnitudes of Q and R give the relative costs of control versus the objectives of control.

The optimal control problem is then to find a control sequence $\{u_i^*, i = 0, 1, \ldots, N-1\}$ such that u_i^* and the resulting x_i^* satisfy equations (1) and (2), and the cost functional (3) is minimized.

The solution to this problem, which is described elsewhere [5, 6], provides a convenient method of obtaining optimal policies that are computationally compatible even with reasonably large econometric models. The optimal closed-loop feedback control is linear, i.e. of the form

$$u_i^* = F_i x_i^* + G_i \tag{4}$$

This control law tends to be "self-correcting," i.e. if random shocks are introduced into the system so that at different times x_i moves away from its optimal value x_i^*, the resulting optimal control (in the following period) will force the state variables towards their optimal paths.

As an example of the application of this method, this author [5] calculated several optimal stabilization policies using a 28-state variable (ten equations) quarterly econometric model [7]. The optimal policies were based on different cost functions (i.e. different Q and R matrices) designed to provide insight into the trade-offs inherent in policy formulation in the context of the model. The results demonstrate the usefulness of the approach as a tool both for policy planning and for the analysis and better understanding of a model's dynamic behavior.

3. Alternative Cost Functionals and Non-Linear Models

The quadratic cost functional has become a familiar in economic optimization problems. Besides having the nice property of yielding linear decision rules when applied to the constraints of a linear system, it is, as Theil [10] and others have argued, a very reasonable way to model the costs of deviating from desired objectives. It is restrictive however, in that it is symmetrical, i.e. overshooting a policy target incurs the same cost as undershooting the target.

Attempts have been made to solve dynamic optimization problems with more general cost functionals. In a recent application to economic stabilization policy, Friedman [2] has extended Theil's [10] specification and solution of a linear-quadratic optimization problem by working with a cost functional that is piece-wise quadratic. For each endogenous variable and each policy variable, the range of possible values is divided into three regions; values within the middle region are assigned zero cost, but values within the two extreme regions are penalized quadratically but asymetrically (e.g. overshooting a target might cost less than undershooting). The algorithm devised to solve this problem calculates optimal policies through an iterative process which solves Theil's standard problem along the different pieces of the cost functional.

The decision as to whether this or any other deviation from the standard quadratic cost functional is worth the resulting computational expense depends in part on the range of values that the endogenous and exogenous variables are likely to take on. Unemployment and inflation will probably always be higher than is desired, and GNP growth lower than desired, so there is some question as to whether the symmetry of the quadratic cost is itself a serious limitation. There may be functional forms other than the quadratic which are more representative of actual social costs. However, the definition of society's economic goals and preferences in any parametrizable functional form is probably a much more complex problem than is the attainment of those goals, but as long as one is willing to assume that the definition is possible, a quadratic specification does not seem too unreasonable, particularly in view of its analytical tractability.

Probably a more serious restriction than the quadratic cost is that of a linear model. Most econometric models are at least quasi-linear in structure, but sometimes the more interesting aspects of their dynamic behavior arise from the non-linearities. Livesey [4] has recently approached the optimal stabilization problem with a continuous-time 15 state-variable non-linear model of the U.K. The control variables included the interest rate, the growth rate of government expenditures, and the rates of change of three tax rates, and the quadratic cost functional penalized for unemployment and an adverse trade balance, while assigning a positive utility to the terminal capital stock.

Livesey solved this stabilization problem computationally, using a conjugate gradient method [3], but large amounts of computation time are typically involved in the iterative solution of a non-linear optimal control problem. Even though his model was fairly small, the large number of iterations required made it too costly to reach the true optimum (i.e. to allow the solution algorithm to iterate until convergence) or to repeat the optimization for several alternative cost functionals. Livesey's results raise the fundamental question that has been encountered again and again in engineering application of optimal control, namely, is the non-linear optimization worth all of the computational difficulty that it entails. This question is particularly important in economic stabilization where the specification of the cost functional is so arbitrary, thus making it desirable to test different objectives by computing several different optimal policies. The experience in engineering has been that often the closed-loop control for a linear model can be applied adequately to the control of a physical system that is non-linear. We have had less experience with control theory in economics, but we can expect that the adequacy or inadequacy of linear or linearized models will depend on how much of the dynamic behavior of the economic system is determined by the non-linearities in its structure. Our analytical tools for dealing with the dynamics of non-linear systems are meager and so we may have to look at computational results to get a better feeling for how much we can rely on linear optimal control as a means of obtaining stabilization policies. As an example, it would be interesting for comparison to solve Livesey's optimization problem using a linearized version of the model with the same cost functional.

4. Deterministic Control as a Tool for Planning

As we mentioned before, computer simulation of econometric models has lately become accepted as a useful tool for the planning and analysis of short-term stabilization policies, and this acceptance is at least in part due to the availability of efficient computational algorithms and the resulting ease by which numerical simulation results can be obtained. An economist can take a model of almost any size and, by simulating it over and over again, experiment with different values of policy parameters and different time-paths for policy and other exogenous variables.

If optimal control is ever to gain the acceptance that simulation has as a practical tool for policy planning and analysis, it is imperative that it yield solutions that are computationally tractable. An economist should be able to get numerical solutions easily so that he can experiment, much as he would with simulation, with different values for the parameters in his cost functional or different time-paths for non-policy exogenous variables. This should be an important consideration when translating stabilization policy into an optimal control problem, and the specification of the optimal control problem should be such that its solution will make it possible to obtain efficiently and easily computational results for policies using models of reasonable size.

The linear-quadratic specification is robust in its applicability to the stabilization problem, and has the special advantage of being computationally tractable. Whether deviations from the linear-quadratic specification are worth the added computational expense depends partly on how big that expense is, but also on exactly how the properties of the resulting optimal policies depend on the non-linearities or alternative cost functionals that one might want to introduce. For now, even within the context of the linear-quadratic formulation, the dynamic properties of optimal stabilization policies are not well understood. Sengupta [8] for example has used simple multiplier-accelerator models to show that the stability characteristics (e.g. the possible presence of oscillatory behavior) of the optimal policy depends on the lag structure of the model, as well as possible constraints on the control. When dealing with non-linear models and more complicated cost functionals, it is all the more difficult to get an analytical understanding of the dynamics of the optimal policy and to assess the loss (the degree of sub-optimality) involved in a linearization of the model.

The points raised above also apply to the choice between a deterministic versus a stochastic approach to finding optimal policies. The econometric model is a stochastic system, but solving a stochastic optimal control problem that takes into account both the implicit additive error terms and the statistical properties of the estimated coefficients can be extremely difficult, especially when one wants to obtain computational results for a large model. In simplifying the problem and using a deterministic treatment, we rely on past engineering experience with the linear-quadratic deterministic control of non-linear stochastic systems. The self-correcting nature of the linear control law (equation 4) that results seems to have provided rather satisfactory results—satisfactory not only because they can be computed easily, but also because they do not seem to be that sub-optimal in their performance.

Massachusetts Institute of Technology

References

[1] Athans, M. and Falb, P., *Optimal Control—An Introduction to the Theory and Its Applications*, McGraw-Hill, 1966.

[2] Friedman, B. M., "Methods in Optimization for Economic Stabilization Policy," Economics Ph.D. Thesis, Harvard University, 1971.

[3] Lasdon, L. S., Mitter, S. K., and Waren, A. D., "The Conjugate Gradient Method for Optimal Control Problems," *IEEE Transactions on Automatic Control*, April, 1967.

[4] Livesey, D. A., "Optimizing Short-Term Economic Policy," *Economic Journal*, September, 1971.

[5] Pindyck, R. S., *Optimal Planning for Economic Stabilization*, North-Holland, forthcoming.

[6] Pindyck, R. S., "An Application of the Linear-Quadratic Tracking Problem to Economic Stabilization Policy," *IEEE Transactions on Automatic Control*, June, 1972.

[7] Pindyck, R. S., "A Small Quarterly Model of the U.S. Economy," April, 1970, unpublished.

[8] Sengupta, J. K., "Optimal Stabilization Policy with a Quadratic Criterion Function," *Review of Economic Studies*, January, 1970.

[9] Theil, H., "A Note on Certainty Equivalence in Dynamic Planning," *Econometrica*, April, 1957.

[10] Theil, H., *Optimal Decision Rules for Government and Industry*, North-Holland, 1964.