Rules Rather than Discretion: The Inconsistency of Optimal Plans
Author(s): Finn E. Kydland and Edward C. Prescott
Source: *The Journal of Political Economy*, Vol. 85, No. 3 (Jun., 1977), pp. 473-492
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/1830193
Accessed: 28/09/2008 06:29

Rules Rather than Discretion: The Inconsistency of Optimal Plans

Finn E. Kydland

Norwegian School of Economics and Business Administration

Edward C. Prescott

Carnegie-Mellon University

Even if there is an agreed-upon, fixed social objective function and policymakers know the timing and magnitude of the effects of their actions, discretionary policy, namely, the selection of that decision which is best, given the current situation and a correct evaluation of the end-of-period position, does not result in the social objective function being maximized. The reason for this apparent paradox is that economic planning is not a game against nature but, rather, a game against rational economic agents. We conclude that there is *no* way control theory can be made applicable to economic planning when expectations are rational.

I. Introduction

Optimal control theory is a powerful and useful technique for analyzing dynamic systems. At each point in time, the decision selected is best, given the current situation and given that decisions will be similarly selected in the future. Many have proposed its application to dynamic economic planning. The thesis of this essay is that it is not the appropriate tool for economic planning even when there is a well-defined and agreed-upon, fixed social objective function.

We find that a discretionary policy for which policymakers select the

We would like to thank Walter Dolde, Leif Johansen, Robert E. Lucas, Jr., Christopher A. Sims, and Neil Wallace, who all provided comments on an earlier draft. We also would like to acknowledge the support of the Guggenheim Foundation, National Science Foundation, and the Bank of Norway.

[*Journal of Political Economy*, 1977, vol. 85, no. 3]

best action, given the current situation, will not typically result in the social objective function being maximized. Rather, by relying on some policy rules, economic performance can be improved. In effect this is an argument for rules rather than discretion, but, unlike Friedman's (1948) argument, it does not depend upon ignorance of the timing and magnitude of the effects of policy.

The reasons for this nonintuitive result are as follows: optimal control theory is an appropriate planning device for situations in which current outcomes and the movement of the system's state depend only upon current and past policy decisions and upon the current state. But, we argue, this is unlikely to be the case for dynamic economic systems. Current decisions of economic agents depend in part upon their expectations of future policy actions. Only if these expectations were invariant to the future policy plan selected would optimal control theory be appropriate. In situations in which the structure is well understood, agents will surely surmise the way policy will be selected in the future. Changes in the social objective function reflected in, say, a change of administration do have an immediate effect upon agents' expectations of future policies and affect their current decisions. This is inconsistent with the assumptions of optimal control theory. This is not to say that agents can forecast future policies perfectly. All that is needed for our argument is that agents have *some* knowledge of how policymakers' decisions will change as a result of changing economic conditions. For example, agents may expect tax rates to be lowered in recessions and increased in booms.

The paradox also arises in situations in which the underlying economic structure is not well understood, which is surely now the case for aggregate economic analyses. Standard practice is to estimate an econometric model and then, at least informally, to use optimal-control-theory techniques to determine policy. But as Lucas (1976) has argued, since optimal decision rules vary systematically with changes in the structure of series relevant to the decision maker, any change in policy will alter the structure of these rules. Thus changes in policy induce changes in structure, which in turn necessitate reestimation and future changes in policy, and so on. We found for some not implausible structures that this iterative procedure does not converge, and, instead, stabilization efforts have the perverse effect of contributing to economic instability. For most examples, however, it did converge, and the resulting policy was consistent but suboptimal. It was consistent in the sense that at each point in time the policy selected was best, given the current situation. In effect the policymaker is failing to take into account the effect of his policy rule upon the optimal decison rules of the economic agents.

In this paper, we first define consistent policy and explain for the two-period problem why the consistent policy is suboptimal. The implications of the analysis are then considered for patent policy and

flood-control problems for which consistent policy procedures are not seriously considered. Then, for the aggregate demand management problem, it is shown that the application of optimal control theory is equally absurd, at least if expectations are rational. Doing what is best, given the current situation, results in an excessive level of inflation, but unemployment is no lower than it would be if inflation (possibly deflation or price stability) were at the socially optimal rate. Consistency for infinite-period recursive economic structures is then considered. In equilibrium, optimizing agents follow rules which specify current decisions as a function of the current state.[1] Methods are developed for computing these equilibrium decision rules for certain specialized structures. The methods are used to evaluate alternative investment-tax-credit policies designed both to stabilize and to yield optimal taxation. Among the policies evaluated is the suboptimal consistent policy. Within the class of feedback policy rules, we found that the optimal one depended upon the initial conditions. Thus it was not optimal to continue with the initial policy in subsequent periods; that is, the optimal policy was inconsistent.

II. Consistent Policy

Let $\pi = (\pi_1, \pi_2, \ldots, \pi_T)$ be a sequence of policies for periods 1 to T (which may be infinite) and $x = (x_1, x_2, \ldots, x_T)$ be the corresponding sequence for economic agents' decisions. An agreed-upon social objective function

$$S(x_1, \ldots, x_T, \pi_1, \ldots, \pi_T) \tag{1}$$

is assumed to exist.[2] Further, agents' decisions in period t depend upon all policy decisions and their past decisions as follows:

$$x_t = X_t(x_1, \ldots, x_{t-1}, \pi_1, \ldots, \pi_T), \qquad t = 1, \ldots, T. \tag{2}$$

In such a framework an optimal policy, if it exists, is that feasible π which maximizes (1) subject to constraints (2). The concept of consistency is less obvious and is defined as follows:

> *Definition:* A policy π is *consistent* if, for each time period t, π_t maximizes (1), taking as given previous decisions, $x_1, \ldots, x_{t-1}$, and that future policy decisions (π_s for $s > t$) are similarly selected.

[1] The original objective of this research was to demonstrate the applicability of optimal control methods in a rational-expectations world. We recognized the nonoptimality of the consistent solution obtained by using control-theory techniques, but initially considered this a minor problem. Further thought, in large part motivated by C. A. Sims's criticism of our initial analyses, led us to the radical conclusions of this essay.

[2] Uncertainty is not the central issue of this essay. As with Arrow-Debreu state-preference theory, one need only define the decision elements to be functions contingent upon observables to incorporate uncertainty as is done for the stabilization example in Sec. V.

The inconsistency of the optimal plan is easily demonstrated by a two-period example. For $T = 2$, π_2 is selected so as to maximize

$$S(x_1, x_2, \pi_1, \pi_2), \tag{3}$$

subject to

$$x_1 = X_1(\pi_1, \pi_2)$$

and

$$x_2 = X_2(x_1, \pi_1, \pi_2). \tag{4}$$

For a plan to be consistent, π_2 must maximize (3), given the past decisions π_1, x_1, and constraint (4). Assuming differentiability and an interior solution, then necessarily

$$\frac{\partial S}{\partial x_2}\frac{\partial X_2}{\partial \pi_2} + \frac{\partial S}{\partial \pi_2} = 0.$$

The consistent policy ignores the effects of π_2 upon x_1. For the optimal decision rule, the first-order condition is

$$\frac{\partial S}{\partial x_2}\frac{\partial X_2}{\partial \pi_2} + \frac{\partial S}{\partial \pi_2} + \frac{\partial X_1}{\partial \pi_2}\left[\frac{\partial S}{\partial x_1} + \frac{\partial S}{\partial x_2}\frac{\partial X_2}{\partial x_1}\right] = 0.$$

Only if either the effect of π_2 upon x_1 is zero (i.e., $\partial X_1/\partial \pi_2 = 0$) or the effect of changes in x_1 upon S both directly and indirectly through x_2 is zero (i.e., $[\partial S/\partial x_1 + \partial S/\partial x_2\, \partial X_2/\partial x_1] = 0$) would the consistent policy be optimal.

Pollak (1968) resolved a planning inconsistency which arose because different generations had different preference orderings by assuming at each stage that the policy selected was best (relative to that generation's preferences), given the policies which will be followed in the future. For the T-period problem, the π_T is determined which, conditional upon previous decisions π_t and x_t, is best:

$$\pi_T = \Pi_T(\pi_1, \ldots, \pi_{T-1}, x_1, \ldots, x_{T-1}).$$

Once the functional relationship Π_T is known, the determination of the best policy rule $\pi_{T-1} = \Pi_{T-1}(\pi_1, \ldots, \pi_{T-2}, x_1, \ldots, x_{T-2})$ can be determined, and in general the consistent policy

$$\pi_t = \Pi_t(\pi_1, \ldots, \pi_{t-1}, x_1, \ldots, x_{t-1})$$

can be determined once future policy rules are known. With such a procedure, the policy decision at each stage is optimal, given the rules

for future policy selection.[3] But as the simple example illustrated, this procedure is suboptimal.

Two examples follow:

The issues are obvious in many well-known problems of public policy. For example, suppose the socially desirable outcome is not to have houses built in a particular flood plain but, given that they are there, to take certain costly flood-control measures. If the government's policy were not to build the dams and levees needed for flood protection and agents knew this was the case, even if houses were built there, rational agents would not live in the flood plains. But the rational agent knows that, if he and others build houses there, the government will take the necessary flood-control measures. Consequently, in the absence of a law prohibiting the construction of houses in the flood plain, houses are built there, and the army corps of engineers subsequently builds the dams and levees.

A second example is patent policy. Given that resources have been allocated to inventive activity which resulted in a new product or process, the efficient policy is not to permit patent protection. For this example, few would seriously consider this optimal-control-theory solution as being reasonable. Rather, the question would be posed in terms of the optimal patent life (see, e.g., Nordhaus 1969), which takes into consideration both the incentive for inventive activity provided by patent protection and the loss in consumer surplus that results when someone realizes monopoly rents. In other words, economic theory is used to predict the effects of alternative policy rules, and one with good operating characteristics is selected.

III. The Inflation-Unemployment Example

The suboptimality of the consistent policy is not generally recognized for the aggregate demand management problem. The standard policy prescription is to select that policy which is best, given the current situation. This may seem reasonable, but for the structure considered, which we argue is a plausible abstraction of reality, such policy results in excessive rates of inflation without any reduction in unemployment. The policy of maintaining price stability is preferable.

[3] There are some subtle game-theoretic issues which have not been addressed here. Peleg and Yaari (1973) criticized Pollak's solution because sometimes it did not exist and proposed an alternative solution to the noncooperative intergeneration game. As explained by Kydland (1975*b*), in the language of dynamic games, Pollak used the feedback solution and Peleg and Yaari the open-loop solution. For policy selection, the policymaker is dominant, and for dominant-player games, the open-loop solution is inconsistent (see Kydland 1975*a*, 1975*b* for further details). That is why Peleg and Yaari's solution was not considered here.

The attempts of economists to rationalize the apparent trade-off between unemployment and inflation in modern theoretical terms have resulted in models with the following structure: unemployment (employment) is a decreasing (increasing) function of the discrepancy between actual and expected inflation rates. This example assumes such a relationship and that it is linear:

$$u_t = \lambda(x_t^e - x_t) + u^*, \tag{5}$$

where u_t is unemployment in period t, λ a positive constant, x_t the inflation rate, x_t^e the forecasted or expected inflation rate, and u^* the natural rate implied by these theories. As has been recently shown by Phelps and Taylor (1975), one need not rely upon imperfect information across firms about the "generality" of shock or imperfect foresight about the persistence of shock over time to obtain a similar relationship. They obtained one by assuming price rigidities, namely, that prices and wages are set prior to the realization of demand.

The crucial issue is what assumption to make concerning price expectations. The conventional approach is to assume that expectations depend in some mechanical ad hoc way upon past prices. If so, control theory would be an appropriate tool to determine the optimal path of unemployment and inflation. The policy decision in each period would consider both current outcomes and a proper evaluation of the terminal price expectations state variable. Such a treatment of expectations is difficult to justify either on a priori or empirical grounds. A change in administration which reflects a change in the relative costs society assigns to unemployment and inflation will have an immediate effect upon expectations—contrary to the implicit assumption of the proponents of control theory. Moreover, private agents or their agents have as much information about the economic structure as does the policymaker and some information concerning the implicit objective function which rationalizes policy selections. Therefore their forecasts of future policy behavior will be related to actual policy selection. This does not imply that policy is perfectly predicted, but then neither is the behavior of private agents. Just partial predictability of policy is sufficient to invalidate the use of optimal control theory.

For this example, we shall assume that the expectations are rational, so that the mathematical expectation of inflation equals the expected rate:

$$x_t^e = Ex_t.$$

Whether forecasts are rational is still open to debate. In Sargent (1973) the rational-expectations hypothesis is tested and accepted. He also explains why many other tests that rejected the hypothesis are invalid. He does not, however, comment on the Hirsch and Lovell (1969) test

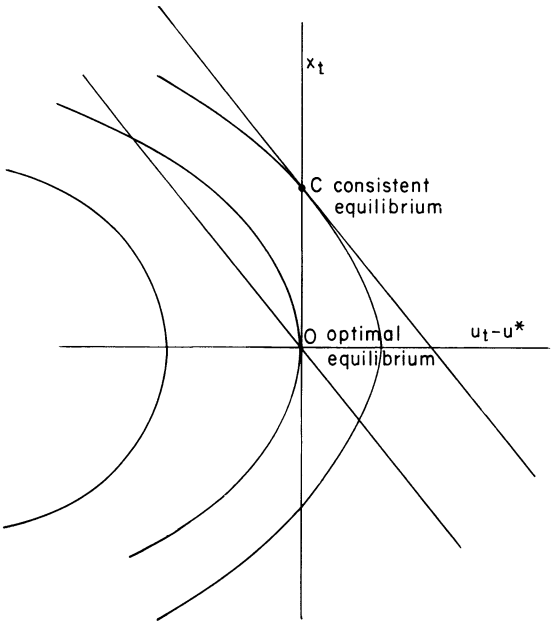

FIG. 1.—Consistent and optimal equilibrium

which used direct measures of expectations and found that forecast errors were systematically related to lagged sales, so we will do so. Responses to this finding are that there may be biases in their measurement of expectations, and these biases are related to lagged sales. This is not implausible, given the subtleness of the expectations concept and the imprecision of survey instruments. Further, even if there were a systematic forecast error in the *past*, now that the Hirsch and Lovell results are part of agents' information sets, future forecast errors should not be subject to such biases.

To complete the model, a theory of policy selection is needed. Here it is assumed that there is some social objective function which rationalizes policy choice:

$$S(x_t, u_t).$$

If the rationalization is not perfect, a random term must be introduced into the function. The consistent policy maximizes this function subject to the Phillips curve constraint (5).

Figure 1 depicts some Phillips curves and indifference curves. From (5) the Phillips curves are straight lines having slope $-\lambda^{-1}$ and intersecting the vertical axis at x_t^e. For a consistent equilibrium, the indifference curve must be tangent to a Phillips curve at a point along the vertical axis—as at point C. Only then are expectations rational and the policy selected

best, given the current situation. The indifference curves imply that the socially preferred inflation rate is zero, which seems consistent with the public's preferences. We of course recognize that inflation is a tax on reserves and currency, and a more informed public might prefer some positive or negative inflation rate. If so, x_t need only be interpreted as deviation from the optimal rate. The outcome of a consistent policy selection clearly is not optimal. If the policymakers were compelled to maintain price stability and did not have discretionary powers, the resulting equilibrium would have no higher unemployment than the consistent policy. The optimal equilibrium is point O, which lies on a higher indifference curve than the consistent-equilibrium point C.

It is perhaps worthwhile to relate our analysis to that of Taylor's (1975), in which he found that the optimal monetary policy was random in a rational-expectations world. Similar results would hold for our problem if uncertainty in the social objective function had been introduced. Both for his structure and for ours, the optimal policy is inconsistent, and consequently it is not optimal for the policymaker to continue with his original policy rules.

IV. Consistent Planning for the Infinite Horizon

The method of backward induction cannot be applied to infinite-period problems to determine a consistent policy because, unlike the finite-period problem, there is no final period with which to begin the induction. For recursive structures, however, the concept of consistency can be defined in terms of policy *rules*. Suppose that the economy at time t can be described by a vector of state variables y_t, a vector of policy variables π_t, a vector of decision variables x_t for the economic agents, and a vector of random shocks ε_t which are temporally independent. The movement over time of these variables is given by the system of equations

$$y_{t+1} = F(y_t, \pi_t, x_t, \varepsilon_t).$$

Let the feedback policy rule for future periods be

$$\pi_s = \Pi^f(y_s), \qquad s > t. \tag{6}$$

For certain structures, rational economic agents will in the future follow a rule of the form

$$x_s = d^f(y_s; \Pi^f). \tag{7}$$

It is important to note that changes in policy rule Π^f change the functional form of d^f, a point convincingly made by Lucas (1976) in his critique of current econometric policy-evaluation procedures. The decisions of agents in the current period will have the form

$$x_t = d^c(y_t, \pi_t; \Pi^f).$$

Again, it is important to note that expectations of future policy affect current decisions. For example, the effect of an increase in the investment tax credit will depend upon the expected future levels of the investment tax credit.

If, in addition, the social objective function is of the form

$$\sum_{s=t}^{\infty} \beta_\pi^s q(x_s, y_s, \pi_s), \qquad 0 < \beta_\pi < 1,$$

and the objective is to minimize its expected value, the optimal value for π_t will depend upon both y_t and Π^f, the policy rule which will be used in the future. In other words, the best policy rule for the current period $\Pi^c(y)$ is functionally related to the policy rule used in the future $\Pi^f(y)$, say

$$\Pi^c = g(\Pi^f).$$

A stationary policy rule Π is consistent if it is a fixed point of mapping g, for then it is best to use the same policy rule as the one expected to be used in the future.[4]

Suppose policymakers and agents do not have a clear understanding of the dynamic structure of the economy. Over time, agents will grope for and most likely converge to the equilibrium rules of forms (6) and (7). Policymakers taking the decision rules of agents as given, when evaluating alternative decisions, typically would consider the trade-off of current outcomes relative to the desirability or value of the end-of-period state. Assuming that their valuation of the terminal state is approximately correct, they will be selecting the approximately consistent policy, assuming also that agents have approximately rational expectations. Thus it seems likely that the current practice of selecting that policy which is best, given the current situation, is likely to converge to the consistent but *suboptimal* policy.[5]

It is hard to fault a policymaker acting consistently. The reason that such policies are suboptimal is not due to myopia. The effect of this decision upon the entire future is taken into consideration. Rather, the suboptimality arises because there is no mechanism to induce *future* policymakers to take into consideration the effect of their policy, via the expectations mechanism, upon *current* decisions of agents.

[4] This is the solution concept used by Phelps and Pollak (1968) for an infinite-period second-best growth problem when different generations had inconsistent preferences.

[5] Optimal policy refers to the best policy, assuming it exists, within a certain class of policies. Within the class of linear feedback rules $\Pi(y_t)$, we found that the best policy rule depended upon the initial condition. The most general class of decision policies are characterized by a sequence of probability measures indexed by the history $\{\Pi_t(x^t, \pi^t, y^t)\}$, with the superscripted variables denoting all previously observed values of the variables. It was necessary to consider probability distributions because for some games a randomized strategy will be optimal and not dominated by a deterministic one. For games against nature, only deterministic strategies need be considered.

V. The Investment-Tax-Credit Example

In this section an equilibrium framework is developed for evaluating a class of investment-tax-credit policies. The assumed technological structure is similar to the one used by Jorgenson (1963), though increasing costs associated with rapid adjustment in capacity are assumed. A firm uses k_t units of capital and n_t units of labor to produce an output $Ak_t^{\alpha}n_t^{(1-\alpha)}$. Output price is p_t, and the real wage is assumed to be a constant, say 1. Investment planned in period t and carried out that period and the next, x_t, does not become productive until period $t + 2$. The relationship between current productive capital, planned investment, and future productive capital is

$$k_{t+2} = x_t + (1 - \delta)k_{t+1},$$

where δ is the constant physical rate of depreciation. Investment costs associated with x_t, the new investment plans in period t, occur in both period t and period $t + 1$. This reflects the fact that time is required to expand capacity, and investment expenditures occur over the entire time interval. The fraction of the investment effort induced by plan x_t in the current period is ϕ, and the fraction induced in the subsequent period is $1 - \phi$. The investment rate in period t is then

$$z_t = \phi x_t + (1 - \phi)x_{t-1}.$$

Following Haavelmo (1960), Eisner and Strotz (1963), Lucas (1967), Gould (1968), and Treadway (1969), we assume that the investment expenditures are an increasing convex function of the rate of capital expansion z_t. In order to insure constant returns to scale in the long run, the function is assumed to have slope equal to the price q of capital goods at $z_t = \delta k_t$. Making the quadratic approximation, the investment expenditures in period t are then

$$i_t = qz_t + \gamma(z_t - \delta k_t)^2,$$

where γ is positive. Observe that i_t depends upon investment plans in both the current and previous periods and that, if x_t is constant over time and sufficient to maintain the capital stock, $i_t = qx_t$.

The gross cash inflow during period t is

$$p_t Ak_t^{\alpha}n_t^{(1-\alpha)} - p_t n_t - i_t. \tag{8}$$

In period t a tax rate θ is applied to sales less labor costs and depreciation. Letting Ψ be the fraction of the "true" depreciation being tax deductible, the tax bill is then

$$\theta(p_t Ak_t^{\alpha}n_t^{(1-\alpha)} - p_t n_t - \Psi\delta k_t). \tag{9}$$

Finally, an investment tax credit is offered at the value π_t, so there will be a tax offset of

$$\pi_t z_t. \tag{10}$$

The view is that the adjustment cost term reflects costs internal to the firm and therefore not eligible for the investment tax credit.[6]

The net cash inflow in period t is (8) less (9) plus (10). The objective of the firm is to maximize the expected present value of this net cash inflow stream. Maximizing each period's cash flow over the period's n_t, the objective function to be maximized becomes

$$E \sum_{t=0}^{\infty} \beta^t[(1 - \theta)p_t\alpha\lambda k_t + \theta\Psi\delta k_t - (q - \pi_t)z_t - \gamma(z_t - \delta k_t)^2],$$

where $\lambda = [1/(1 - \alpha)][A(1 - \alpha)]^{1/\alpha}$ is output per unit of capital and β is the discount factor.

The inverse aggregate demand function is assumed to be linear. Letting capital letters denote the aggregates of the corresponding variables for the individual firms, the inverse demand function is of the form

$$p_t = a_t - b\lambda K_t,$$

where b is a positive constant, a_t is a stochastic demand shift parameter, and K_t is the aggregate capital stock for the firms. We assume that a_t is subject to the first-order autoregressive process

$$a_{t+1} = \rho a_t + \varepsilon_t, \qquad -1 < \rho < 1, \tag{11}$$

where the ε_t are positive independent random variables with mean μ and variance σ_ε^2.

For the economy to be in equilibrium, the expected and actual distribution of the random elements must be equal. Here we are assuming rational expectations of Muth (1961) and of Lucas and Prescott (1971). Brock (1972) has characterized such expectations schemes as being self-fulfilling. We are implicitly assuming that the economy is Hicksian in the sense that a single consumer could have the implicit excess demand function. For such economies wealth effects net out, and our equilibrium yields the same allocation as the Arrow-Debreu state preference equilibrium upon extension of their analysis to infinite-dimensional space.

In the Appendix we develop direct methods for computing the competitive equilibrium, given the policy rule.[7] Also discussed in the Appendix is the stability of the equilibrium.

[6] We are also implicitly assuming a per unit rather than a percentage tax credit.

[7] If policy does not depend upon agents' decisions, the competitive equilibrium is efficient and therefore maximizes the utility of the economy-wide consumer, a fact exploited in Lucas and Prescott (1971) to characterize the competitive equilibrium. For these examples, policy rules are of the feedback variety and depend upon past decisions of private economic agents. In effect this introduces an externality. Suppose, e.g., that future investment tax credits depend positively upon the magnitude of future capital stocks. If all agents invest less now, future capital stocks will be smaller and, consequently, future investment tax credits larger. Because of this externality, the competitive equilibrium will not in general maximize the utility of the economy-wide consumer, given the policy rule. This is why it was necessary to devise direct methods.

The policymakers are choosing values for π in each time period so as to minimize the value of some social preference function. We use a quadratic approximation of a function which includes the terms likely to carry any weight. The assumed form is

$$E \sum_{t=0}^{\infty} \beta_{\pi}^{t}[\omega_1(\lambda K_t - g_1)^2 + \omega_2(Z_t - g_2)^2 + \omega_3(\lambda K_t + Z_t - g_3)^2 + \omega_4(p_t - p_{t-1})^2 + \omega_5 \pi_t Z_t + \omega_6 \pi_t^2],$$

where each ω_i, $i = 1, \ldots, 6$, indicates the relative weight on each of the components. The terms g_1, g_2, and g_3 are targets for real output of the industry, for real investments, and for the total of the two, respectively. Reasons for including the tax credit in the loss function are that the amount paid may have to be collected elsewhere in the form of taxes and inefficiencies generally caused by such measures.

Our examples assume that a passive investment-tax-credit stabilization policy had been pursued in the past and that the function describing investment behavior was equilibrium, given this passive policy. They also assume that econometricians have estimated the investment relationship[8] and that the policymaker uses control theory to determine which policy rule is optimal under the incorrect assumption that the equilibrium investment function is invariant to the policy rule used. Subsequent to the implementation of this policy rule, the economy moves to the new equilibrium investment function. Econometricians revise their estimate of the investment function, arguing that there has been structural change, and the policymaker uses optimal control to determine a new policy rule. The change in policy induces still another change in the investment function, which in turn induces a change in the policy rule, once the shift in the investment function is recognized. This iterative process, we think, captures the essence of what is actually happening. We observe that econometricians are continually revising their estimates of the structure on the basis of which new policies are devised and are continually surprised to find that the structure has changed.

Our results can be summarized as follows:

Factor shares, the capital output ratio, tax rates, and capital consumption allowances were used to deduce not unreasonable values for parameters of technology and preferences with the exception of ϕ, the distributed lag coefficient of investment, and ρ, the autoregressive parameter of the

[8] We recognize that there may be problems involved in estimating the investment function because of perfect multicollinearity between π_t and the other independent variables. However, the model could be modified to permit random fluctuations in the price of capital. For estimation purposes, one would then, instead of using the investment tax credit, use the price of capital which is affected by the tax credit in a predictable way.

aggregate demand relation.[9] Only these parameters and the parameters of the objective function were subject to variation. Space constraints preclude more than a brief summary of the results.[10]

Typically the iterative process of the policy rule change inducing investment function change inducing policy rule change, etc., did converge. Given that it converges, the limiting policy rule is consistent in the sense described in Sections II and IV. In all cases for which it did converge, we searched for and found linear feedback policy rules which were superior to this consistent rule, typically by a substantial amount.

For one example ($\omega_1 = 2$, $\omega_2 = 4$, $\omega_3 = 1$, $\omega_4 = 10$, $\omega_5 = 20$, $\omega_6 = 10$, and $\rho = 0.6$), the application of optimal control initially improved the performance of the economy relative to the assumed objective function. For the first two iterations, the economy was subject to less fluctuation and fluctuated about a preferred point. After the third iteration, however, performance deteriorated, and the consistent policy to which the process converged was decidedly inferior to the passive policy for which the investment tax credit was not varied. The difference in performance corresponded roughly to the variables being 10 percent on average away from their targets.

For another example ($\omega_1 = 1$, $\omega_2 = 2$, $\omega_3 = 1$, $\omega_4 = 10$, $\omega_5 = 3$, $\omega_6 = 20$, and $\rho = -0.6$), the iterative process did not converge. Changes in the policy rule induced ever larger changes in the investment function. The variables fluctuated about their targeted values but fluctuated with increased amplitude with each iteration. This is a very disturbing result, for it indicates that current practice, if continued, could conceivably result in even greater fluctuations than are now being experienced.

There are two lessons we learned from the examples. First, the use of optimal control theory is hazardous and could very well increase economic fluctuations or even make a stable economy unstable. Second, even when it does work reasonably well, it can be improved upon by following some other simple feedback rule.

This is not an argument that economic fluctuations are either desirable or unavoidable. That our economy has experienced periods of reasonable stability is evidence that much of the fluctuation is avoidable. Rather, it is a plea for the use of economic theory to evaluate correctly the performance of a policy rule before it is implemented. We emphasize that

[9] The values used for the fixed parameters were $\delta = 0.1$, $\theta = 0.5$, $A = 1.15$, $\gamma = 0.03$, $\alpha = 0.28$, $b = 0.4$, $\lambda = 0.7$, $q = 1$, $\beta = \beta_\pi = 0.9$, $\phi = 0.3$, and $\sigma_\varepsilon = 0.03$. Given ρ, the parameter μ was chosen so as to give a mean of 2.5 for a_t. For instance, if $\rho = 0.6$, as in the first example, we get $\mu = 1$. In the examples discussed, we used $g_1 = 2.008$, $g_2 = 0.2837$, and $g_3 = 2.292$, which are the stationary levels of the corresponding target variables when the passive policy of $\pi_t = 0$ is used in every period.

[10] More details of the results of the numerical examples can be found in our original working paper, which is available upon request.

optimal control theory can *not* be made applicable to economic planning by taking into consideration the way changes in the policy rule change the behavioral equation of the model when expectations are rational.

VI. Discussion

The analysis has implications in other situations as well. Kydland (1975*a*) has explored the implications for a dynamic oligopoly problem with a dominant firm. Like the policymaker, the dominant firm takes into consideration the reaction of the other agents in selecting its decision. Precisely the same paradox arises.

The analysis also has implications for constitutional law. A majority group, say, the workers, who control the policy might rationally choose to have a constitution which limits their power, say, to expropriate the wealth of the capitalist class. Those with lower discount rates will save more if they know their wealth will not be expropriated in the future, thereby increasing the marginal product and therefore wage and lowering the rental price of capital, at least for most reasonable technological structures.

Still another area is the current energy situation. We suspect that rational agents are not making investments in new sources of oil in the anticipation that price controls will be instituted in the future. Currently there are those who propose to tax away "excessive" profits of the oil companies with the correct argument that this will not affect past decisions. But rational agents anticipate that such expropriations may be made in the future, and this expectation affects their current investment decisions, thereby reducing future supplies.

VII. Summary and Conclusions

We have argued that control theory is not the appropriate tool for dynamic economic planning. It is not the appropriate tool because current decisions of economic agents depend upon expected future policy, and these expectations are not invariant to the plans selected. We have shown that, if in each period the policy decision selected is the one which maximizes the sum of the value of current outcomes and the discounted valuation of the end-of-period state, the policy selected will be consistent but not optimal. This point is demonstrated for an investment-tax-credit policy example, using a rational-expectations equilibrium theory with costs of adjustment and distributed lags for expenditures. In fact, active stabilization effects did, for some distributed lag expenditure schedules, contribute to economic instability and even make a stable economy unstable.

The structures considered are far from a tested theory of economic fluctuations, something which is needed before policy evaluation is undertaken. The implication of this analysis is that, until we have such a theory, active stabilization may very well be dangerous and it is best that it not be attempted. Reliance on policies such as a constant growth in the money supply and constant tax rates constitute a safer course of action.

When we do have the prerequisite understanding of the business cycle, the implication of our analysis is that policymakers should follow rules rather than have discretion. The reason that they should not have discretion is not that they are stupid or evil but, rather, that discretion implies selecting the decision which is best, given the current situation. Such behavior either results in consistent but suboptimal planning or in economic instability.

If we are not to attempt to select policy optimally, how should it be selected? Our answer is, as Lucas (1976) proposed, that economic theory be used to evaluate alternative policy rules and that one with good operating characteristics be selected. In a democratic society, it is probably preferable that selected rules be simple and easily understood, so it is obvious when a policymaker deviates from the policy. There could be institutional arrangements which make it a difficult and time-consuming process to change the policy rules in all but emergency situations. One possible institutional arrangement is for Congress to legislate monetary and fiscal policy rules and these rules to become effective only after a 2-year delay. This would make discretionary policy all but impossible.

Appendix

Let y be the state variables and x the decision variables for the firm. There is a linear relationship between the next period's state variables, y_{t+1}, and the current x_t and y_t:

$$y_{t+1} = f(x_t, y_t). \tag{A1}$$

The movement of economy-wide state variables Y and aggregate (or per firm) decision variables X are described by the same linear function:

$$Y_{t+1} = F(X_t, Y_t). \tag{A2}$$

We also include a vector of autonomous shocks, W, which are subject to a first-order autoregressive process,

$$W_{t+1} = \Omega W_t + \eta_t, \tag{A3}$$

where Ω is a matrix of fixed coefficients and η a random vector with finite variances. In the vector W we may also include other variables on which decisions can depend and which may be common to the firm and the economy as a whole. An example would be the lagged price level of output.

The firm's objective is to maximize

$$E\left[\sum_{t=0}^{T} \beta^t R(x_t, y_t, X_t, Y_t, W_t, \pi_t)\right], \qquad 0 < \beta < 1,$$

where π_t is a vector of policy variables assumed to be given by a sequence of linear policy rules,

$$\pi_t = \Pi_t(Y_t, W_t), \quad t = 0, \ldots, T, \tag{A4}$$

which in equilibrium are correctly anticipated by the firm.

The cash-flow function R is quadratic. The decisions x_t are selected sequentially conditional on y_t, X_t, Y_t, and W_t. Let $v_t(y_t, Y_t, W_t)$ be the value of the firm at time t. The v_t functions satisfy the recursive relationship

$$v_t(y_t, Y_t, W_t) = R(x_t, y_t, X_t, Y_t, W_t, \pi_t) + \beta E[v_{t+1}(y_{t+1}, Y_{t+1}, W_{t+1})] \tag{A5}$$

subject to constraints (A1)–(A4) and one additional constraint. To explain this last constraint, note that, since x_t is chosen so as to maximize the valuation at time t, if v_t is quadratic and the right-hand side of (A5) concave in x_t, the x_t which maximizes the right-hand side of (A5), taking as given X_t, Y_t, W_t, and the motion of the economy-wide state variables, will be linearly related to y_t, X_t, Y_t, W_t, and π_t:

$$x_t = d_t(y_t, X_t, Y_t, W_t, \pi_t). \tag{A6}$$

In order for the economy to be in equilibrium, we have to impose the constraint that, when firms behave according to (A6), the aggregate or per firm X_t is indeed X_t. Therefore[11]

$$X_t = d_t(Y_t, X_t, Y_t, W_t, \pi_t),$$

which can be rewritten as

$$X_t = D_t(Y_t, W_t, \pi_t). \tag{A7}$$

As the constraints are all linear and the right-hand side of (A5) quadratic, the function v_t is quadratic, given that v_{t+1} is quadratic. The function v_{T+1} is the null function and therefore trivially quadratic, so by induction all the v_t are quadratic. The equilibrium per firm decision function for each time period t is given by (A7).

If the social objective function is quadratic of the form

$$E\left[\sum_{t=0}^{T} \beta_\pi^t S(X_t, Y_t, W_t, \pi_t)\right], \qquad 0 < \beta_\pi < 1,$$

the determination of the consistent policy is straightforward. Let

$$u_t(Y_t, W_t) = E\left[\sum_{i=t}^{T} \beta_\pi^{i-t} S(X_i, Y_i, W_i, \pi_i)\right],$$

given that policy is selected consistently and that the economy is competitive. Thus the function u_t gives the total expected value of the social objective function from period t throughout the rest of the horizon for the consistent policy. By backward induction

$$u_t(Y_t, W_t) = \min_{\pi_t} \{S(X_t, Y_t, W_t, \pi_t) + \beta_\pi E[u_{t+1}(Y_{t+1}, W_{t+1})]\},$$

[11] We think of a large corporation as being the aggregate of several small firms. Therefore the effect of an investment-tax-credit policy is proportional to size.

subject to (A2), (A3), and (A7). If u_{t+1} is quadratic, a quadratic function is being minimized subject to linear constraints. Therefore u_t must be quadratic if u_{t+1} is quadratic. As $u_{T+1} = 0$ and is thus trivially quadratic, all the u_t are quadratic by backward induction, and the consistent policy is a linear function of Y and W:

$$\pi_t = \Pi_t(Y_t, W_t).$$

It is perhaps worthwhile to make the connection between the structure just analyzed and the one described in Section V. The state vector is

$$y_t = (k_t, x_{t-1})',$$

and the linear equations governing its movement over time are

$$y_{t+1} = \begin{pmatrix} 1 - \delta & 1 \\ 0 & 0 \end{pmatrix} y_t + \begin{pmatrix} 0 \\ 1 \end{pmatrix} x_t.$$

The equations governing the economy-wide variables are the same. Furthermore,

$$W_{t+1} \equiv a_{t+1} = \rho a_t + \varepsilon_t \equiv \Omega W_t + \eta_t.$$

The revenue function R for the firm is

$$(1 - \theta)\alpha\lambda p_t k_t + \theta\Psi\delta k_t - (q - \pi_t)z_t - \gamma(z_t - \delta k_t)^2,$$

which has the assumed form, given that

$$p_t = a_t - b\lambda K_t$$

and

$$z_t = \phi x_t + (1 - \phi)x_{t-1}.$$

Finally, the social objective function S given by

$$\omega_1(\lambda K_t - g_1) + \omega_2(Z_t - g_2) + \omega_3(\lambda K_t + Z_t - g_3) + \omega_4(p_t - p_{t-1})^2 + \omega_5\pi_t Z_t + \omega_6\pi_t^2$$

also has the quadratic form, given the assumed definitions of the variables.

Computations for the Infinite-Period Problem

Equilibrium decision rules for agents were determined as the limit of first-period decision rules as the life of the economy went to infinity. There is an interesting and as yet unsolved problem as to the uniqueness of the equilibrium.[12] For these examples the equilibrium associated with a stationary policy rule did appear to be unique, for when we used the method of successive approximation in the value space (i.e., the v function in [A5]) the value function and therefore decision rules converged to the same limit for a number of initial approximations. For some unreasonable policy rules and finite T, there were no competitive equilibria.

Consistent solutions were computed in two different ways. The first determined the first-period consistent policy for T-period problems and the limit determined as T went to infinity. The second determined the nth approximation to the consistent-equilibrium investment function X^n, given the nth approximation to the consistent policy rule Π^n, using the methods described above. Optimal control theory was then used to determine the policy Π^{n+1} which would be

[12] Standard dynamic programming arguments such as those of Denardo (1967) could not be applied because there was not monotonicity of the mapping in the value space.

optimal if X^n were not to change as a result of the change in Π^n. Given initial linear feedback rule Π^0, sequences of linear rules $\{\Pi^n, X^n\}$ were obtained. When such sequences existed and converged, the limits constituted a consistent policy rule and the corresponding equilibrium investment function. In no case did we ever obtain two different consistent policies for the same structure, though both methods of successive approximations were used and a number of different starting values tried.

Stability of the Competitive Equilibrium

We also checked whether the computed competitive equilibria were stable, as follows: given the expected aggregate investment function (which implies expectations) at stage n,

$$X = G_n(y, W, \pi),$$

and given structural relations (A2) and (A3), one finds the optimal firm investment function

$$x = d_n(y, X, Y, W, \pi),$$

which in the aggregate becomes

$$X = D_n(Y, W, \pi).$$

Now let

$$G_{n+1}(Y, W, \pi) = G_n(Y, W, \pi) + \xi[D_n(Y, W, \pi) - G_n(Y, W, \pi)].$$

For the numerical examples in Section V for which we found competitive equilibria, this process converged for various initial aggregate investment functions G_0 for $\xi = 1$ and of course for smaller positive values of ξ as well.

References

Brock, W. A. "On Models of Expectations That Arise from Maximizing Behavior of Economic Agents over Time." *J. Econ. Theory* 5 (December 1972): 348–76.

Denardo, E. "Contraction Mappings in the Theory Underlying Dynamic Programming." *SIAM Review* 9 (April 1967): 165–77.

Eisner, R., and Strotz, R. "Determinants of Business Investment." In Commission on Money and Credit, *Impacts of Monetary Policy*. Englewood Cliffs, N.J.: Prentice-Hall, 1963.

Friedman, M. "A Monetary and Fiscal Framework for Economic Stabilization." *A.E.R.* 38 (June 1948): 245–64.

Gould, J. P. "Adjustment Costs in the Theory of Investment of the Firm." *Rev. Econ. Studies* 35 (January 1968): 47–55.

Haavelmo, T. *A Study of the Theory of Investment*. Chicago: Univ. Chicago Press, 1960.

Hirsch, A. A., and Lovell, M. C. *Sales Anticipations and Inventory Behavior*. New York: Wiley, 1969.

Jorgenson, D. W. "Capital Theory and Investment Behavior." *A.E.R.* (Proc.) 53 (May 1963): 247–59.

Kydland, F. "Equilibrium Solutions in Dynamic Dominant Player Models." Discussion Paper, Norwegian School of Economics and Business Administration, 1975. (*a*)

———. "Noncooperative and Dominant Player Solutions in Discrete Dynamic Games." *Internat. Econ. Rev.* 16 (June 1975): 321–35. (*b*)

Lucas, R. E., Jr. "Adjustment Costs and the Theory of Supply." *J.P.E.* 75, no. 4, pt. 1 (August 1967): 321–34.

———. "Econometric Policy Evaluation: A Critique." In *The Phillips Curve and Labor Markets*, edited by K. Brunner and A. H. Meltzer. Amsterdam: North-Holland, 1976.

Lucas, R. E., Jr., and Prescott, E. C. "Investment under Uncertainty." *Econometrica* 39 (September 1971): 659–81.

Muth, J. F. "Rational Expectations and the Theory of Price Movements." *Econometrica* 29 (July 1961): 315–35.

Nordhaus, W. D. *Invention Growth and Welfare: A Theoretical Treatment of Technological Change*. Cambridge, Mass.: M.I.T. Press, 1969.

Peleg, B., and Yaari, M. E. "On the Existence of a Consistent Course of Action When Tastes Are Changing." *Rev. Econ. Studies* 40 (July 1973): 391–401.

Phelps, E. S., and Pollak, R. A. "On Second-best National Saving and Game-Equilibrium Growth." *Rev. Econ. Studies* 35 (April 1968): 185–99.

Phelps, E. S., and Taylor, J. B. "Stabilizing Properties of Monetary Policies under Rational Price Expectations." Department of Economics Working Paper no. 75-7607, Columbia Univ., 1975.

Pollak, R. A. "Consistent Planning." *Rev. Econ. Studies* 35 (April 1968): 201–8.

Sargent, T. J. "Rational Expectations, the Real Rate of Interest, and the 'Natural' Rate of Unemployment." *Brookings Papers on Economic Activity*, no. 2 (1973), pp. 429–72.

Taylor, J. B. "Monetary Policy during a Transition to Rational Expectations." *J.P.E.* 83, no. 5 (October 1975): 1009–22.

Treadway, A. B. "On Rational Entrepreneurial Behavior and the Demand for Investment." *Rev. Econ. Studies* 36 (April 1969): 227–39.