OPTIMAL CONTROL WITHOUT SOLVING THE BELLMAN EQUATION

Gregory C. Chow
Princeton University

Econometric Research Program
Research Memorandum No. 364

March 1992

Abstract

This paper recommends against solving the Bellman partial differential equation for the value function in optimal control problems involving stochastic differential or difference equations. It recommends solving for the vector Lagrange multiplier associated with a first-order condition for maximum. The method is preferable to Bellman's in exploiting this first-order condition and in solving only algebraic equations in the control variable and Lagrange multiplier and its derivatives rather than a functional equation. The solution requires no global approximation of the value function and is exact for continuous-time models and nearly exact for discrete-time models.

Acknowledgments

The author would like to thank Giuseppe Bertola, John Campbell, Harald Uhlig, an anonymous referee, the Editor and participants of the Econometric Research Program Seminar at Princeton University for helpful comments on an earlier draft.

Econometric Research Program
Princeton University
203 Fisher Hall
Princeton, NJ 08544-1021, U.S.A.

OPTIMAL CONTROL WITHOUT SOLVING THE BELLMAN EQUATION

Gregory C. Chow

Since the 1950's the method of dynamic programming suggested by Richard Bellman (1957) has been considered the main tool for solving optimal control problems where the dynamics is governed by a system of stochastic differential or difference equations. The approach has been to solve for the value function in a partial differential equation known as the Bellman Equation. This paper suggests that solving such a partial differential equation is unnecessary and for most applications unwise. To maximize a concave differentiable function subject to differentiable constraints a standard method is to use Lagrange multipliers and to solve the first-order conditions for maximum. Solving for the value function is contrary to the spirit of this method and fails to exploit the first-order conditions. Given the value of the vector x(t) of state variables, one needs only to find the value of the vector u(t) of control variables to maximize a multiperiod objective function, but to find the value u(t) it is unnecessary to know the value function, which gives the maximum values of the objective function for all conceivable states x(t). We will show that the method of Lagrnge multipliers can be extended to solve dynamic and stochastic optimization problems as well as static maximization problems. In the case of a consumer maximizing a differentiable utility function of quantities of consumption goods subject to budget constraints, the Bellman approach amounts to abandoning the algebraic method of Lagrange multipliers and attempting to solve a partial differential equation for the indirect utility function.

In Chow (1992), the use of Lagrange multipliers without solving for the value function is suggested for solving stochastic control problems in discrete time. This paper is concerned also with continuous-time models in the form of a system of stochastic differential equations. Like the previous paper, it exploits the first-order conditions as suggested by the method of Lagrange multipliers without seeking the value function. Furthermore, by solving algebraic equations rather than functional equations globally it provides a simpler and more accurate method for exploiting the first-order conditions in computing the optimal control function. Section 1 reviews the method of dynamic programming for solving a stochastic control problem in continuous time and points out the main characteristic of our approach. Section 2 provides a method to implement our approach in obtaining the optimal control function and the associated value function without solving Bellman's partial differential equation. Section 3 applies the same method for solving an analogous control problem in discrete time. Section 4 concludes by pointing out the essence of our approach and the weakness of the Bellman approach.

1. Dynamic Programming for a Continuous-Time Model

The problem is

$$\max_u E_t \int_t^{\infty} e^{-\beta(\tau-t)}\, r\Big(x(\tau),u(\tau)\Big)\, d\tau = V(x(t)) \tag{1}$$

subject to

$$dx = f(x,u)dt + S(x,u)dw \tag{2}$$

where $x(t)$ is a $p\times 1$ vector of state variables, $u(t)$ is a $q\times 1$ vector of control variables, β is a discount rate, E_t is the conditional expectation operator given information at time t which includes $x(t)$, $w(t)$ is a vector Wiener process with covariance matrix $\mathrm{cov}(dw)=\Phi dt$, $r(x,u)$ is a differentiable and concave utility function, both f and S in the stochastic differential equation (2) are differentiable, the covariance matrix of Sdw is $S\Phi S'dt\equiv\Sigma dt$, and $V(x)$ is the value function defined by (1).

By considering the problem of finding the optimal control $u(t)$ in a small time interval from t to t+dt, one can write

$$V(x(t)) = \max_u E_t\left[r\left(x(t),u(t)\right)dt + e^{-\beta dt} V\left(x(t+dt)\right)\right]$$

$$= \max_u \left\{ r\left(x(t),u(t)\right)dt + E_t(1-\beta dt)\left[V(x(t)) + dV(x(t))\right]\right\} \qquad (3)$$

implying

$$\beta V(x(t)) = \max_u \left\{ r(x,u) + \frac{1}{dt} E_t dV(x(t))\right\} \qquad (4)$$

where dV(x(t)) is defined as V(x(t+dt)) - V(x(t)) and can be evaluated by Ito's lemma given the stochastic differential equation (2) for dx

$$dV(x) = \left[f' \frac{\partial V}{\partial x} + \frac{1}{2} \operatorname{tr}\left(\frac{\partial^2 V}{\partial x \partial x'} \cdot \Sigma \right)\right] dt + \frac{\partial V}{\partial x'} S dw \qquad (5)$$

Using (5) to evaluate $E_t dV$ and substituting into (4) one obtains the Bellman equation

$$\beta V(x) = \max_u \left\{ r(x,u) + f' \frac{\partial V}{\partial x} + \frac{1}{2} \operatorname{tr}\left(\frac{\partial^2 V}{\partial x \partial x'} \cdot \Sigma \right)\right\} \qquad (6)$$

A standard approach to solve this optimal control problem using dynamic programming consists of two steps. First, assuming the value function V(x) to be known, find the optimum u by solving the maximization problem of (6). For differentiable functions, one may use the first-order condition obtained by differentiation with respect to u_i.

$$\frac{\partial r}{\partial u_i} + \frac{\partial f'}{\partial u_i} \cdot \frac{\partial V}{\partial x} + \frac{1}{2} \operatorname{tr}\left(\frac{\partial^2 V}{\partial x \partial x'} \cdot \frac{\partial \Sigma}{\partial u_i} \right) = 0 \qquad i = 1,\ldots,q \qquad (7)$$

Denote the solution to the q equations by the vector $\hat{u}' = (\hat{u}_1,\ldots,\hat{u}_q)$. If there are side conditions $g(x,u) \le 0$, they can be imposed at this stage. Second, given the solution $\hat{u}(x)$ from (7), solve for the value function V(x) using (6), namely, solve the following partial differential equation for V(x)

$$\beta V(x) = r\left(x,\hat{u}(x)\right) + f'\left(x,\hat{u}(x)\right) \frac{\partial V}{\partial x} + \frac{1}{2} \operatorname{tr}\left[\frac{\partial^2 V}{\partial x \partial x'} \Sigma \left(x,\hat{u}(x)\right)\right] \qquad (8)$$

The approach suggested in this paper does not seek the value function V(x) in obtaining u(x). Denoting $\frac{\partial V}{\partial x}$ by the vector $\lambda(x)$, we rewrite (7) as

$$\frac{\partial r}{\partial u_i} + \frac{\partial f'}{\partial u_i}\lambda + \frac{1}{2}\,\mathrm{tr}\left(\frac{\partial \lambda}{\partial x'}\cdot\frac{\partial \Sigma}{\partial u_i}\right) = 0 \qquad i = 1,\ldots,q \tag{9}$$

First, we obtain from (9) a solution $\hat{u}$ which can be considered a function of x, λ and the matrix $\frac{\partial \lambda}{\partial x'} = \frac{\partial \lambda'}{\partial x}$. Second, given $\hat{u}$ and assuming the existence of the third derivative of the value function we differentiate (8) with respect to x_i to obtain

$$\beta\lambda_i = \frac{\partial}{\partial x_i} r(x,\hat{u}) + \frac{\partial}{\partial x_i} f'(x,\hat{u})\lambda + \frac{\partial \lambda'}{\partial x_i} f(x,\hat{u})$$

$$+ \frac{1}{2}\,\mathrm{tr}\left[\frac{\partial}{\partial x_i}\left(\frac{\partial \lambda}{\partial x'}\right)\cdot\Sigma\right] + \frac{1}{2}\,\mathrm{tr}\left[\frac{\partial \lambda}{\partial x'}\,\frac{\partial}{\partial x_i}\,\Sigma(x,\hat{u})\right] \qquad i = 1,\ldots,p \tag{10}$$

It is suggested that equations (9) and (10) be used to obtain $\hat{u}$ without solving the partial differential equation (8) for the value function V.

2. An Algebraic Method for Finding the Optimal Control Function

Using (9) and (10) the problem is to find an optimal control function $\hat{u}(x)$. For any value x, one seeks a method to compute the value $\hat{u}$ of the optimal control which satisfies (9) and (10). These equations involve the variables u, x, λ, $\frac{\partial \lambda}{\partial x'}$, and $\frac{\partial}{\partial x_i}\left[\frac{\partial \lambda}{\partial x'}\right]$, i = 1,...,p. To find $\hat{u}$ given x, we first assume the two sets of derivatives $\frac{\partial \lambda}{\partial x'}$ and $\frac{\partial}{\partial x_i}\left[\frac{\partial \lambda}{\partial x'}\right]$, i = 1,...,p, to be given, and solve (9) and (10) for $\hat{u}$ and λ. In the second stage, given $\hat{u}$ and λ, we will find the two sets of derivatives satisfying (9) and (10).

In the first stage, we treat all derivatives of λ as given and find $\hat{u}$ and λ by solving (9) and (10). In the second stage, we evaluate the derivatives of λ, given $\hat{u}$ and λ. To do so, we take the total differential of (9) and (10) treating u, λ and x as variables and the derivatives of λ as fixed. Noting r, f and Σ to be functions of x and u, we obtain the differential of (9) as

$$\sum_j^q \left\{ \frac{\partial^2 r}{\partial u_i \partial u_j} + \frac{\partial^2 f'}{\partial u_i \partial u_j}\lambda + \frac{1}{2}\,\mathrm{tr}\left[\frac{\partial \lambda}{\partial x'} \cdot \frac{\partial^2 \Sigma}{\partial u_i \partial u_j} \right]\right\} du_j + \sum_j^p \frac{\partial f_j}{\partial u_i} d\lambda_j$$

$$+ \sum_j^p \left\{ \frac{\partial^2 r}{\partial u_i \partial x_j} + \frac{\partial^2 f'}{\partial u_i \partial x_j}\lambda + \frac{\partial f'}{\partial u_i}\cdot\frac{\partial \lambda}{\partial x_j} + \frac{1}{2}\,\mathrm{tr}\left[\frac{\partial \lambda}{\partial x'} \cdot \frac{\partial^2 \Sigma}{\partial u_i \partial x_j} \right]\right\} dx_j = 0 \qquad i = 1,\ldots,q \tag{11}$$

The differential of (10) is (with δ_{ij} denoting the Kronecker delta)

$$\sum_j^q \left\{ \frac{\partial^2 r}{\partial x_i \partial u_j} + \frac{\partial^2 f'}{\partial x_i \partial u_j}\lambda + \frac{\partial \lambda'}{\partial x_i}\frac{\partial f}{\partial u_j} + \frac{1}{2}\,\mathrm{tr}\left[\frac{\partial}{\partial x_i}\left(\frac{\partial \lambda}{\partial x'}\right)\frac{\partial \Sigma}{\partial u_j} + \frac{\partial \lambda}{\partial x'}\cdot\frac{\partial^2 \Sigma}{\partial x_i \partial u_j} \right]\right\} du_j$$

$$+ \sum_j^p \left[\frac{\partial f_j}{\partial x_i} - \delta_{ij}\beta \right] d\lambda_j + \sum_j^p \left\{ \frac{\partial^2 r}{\partial x_i \partial x_j} + \frac{\partial^2 f'}{\partial x_i \partial x_j}\lambda + \frac{\partial f'}{\partial x_i}\cdot\frac{\partial \lambda}{\partial x_j} + \frac{\partial \lambda'}{\partial x_i}\frac{\partial f}{\partial x_j}\right.$$

$$\left. + \frac{1}{2}\,\mathrm{tr}\left[\frac{\partial}{\partial x_i}\left(\frac{\partial \lambda}{\partial x'}\right)\frac{\partial \Sigma}{\partial x_j} + \frac{\partial \lambda}{\partial x'}\cdot\frac{\partial^2 \Sigma}{\partial x_i \partial x_j} \right]\right\} dx_j \qquad i = 1,\ldots,p \tag{12}$$

Equations (11) and (12) can be written as

$$A \begin{bmatrix} du \\ d\lambda \end{bmatrix} + B dx = 0 \tag{13}$$

where A is a q+p by q+p matrix and B is a q+p by p matrix. The solution of (13) is

$$\begin{bmatrix} du \\ d\lambda \end{bmatrix} = -A^{-1} B dx \tag{14}$$

To evaluate $\frac{\partial}{\partial x_i}\left[\frac{\partial \lambda}{\partial x'}\right]$, $i = 1,\ldots,p$, we differentiate $-A^{-1}B$ partially with respect to x_i. Differentiating the identity $AA^{-1}=I$ partially with respect to x_i yields

$$\frac{\partial A}{\partial x_i} A^{-1} + A \frac{\partial A^{-1}}{\partial x_i} = 0 \; ; \qquad \frac{\partial A^{-1}}{\partial x_i} = -A^{-1}\frac{\partial A}{\partial x_i} A^{-1}$$

Applying the above identity and treating the derivatives of λ as fixed, one obtains

$$\frac{\partial}{\partial x_i}(-A^{-1}B) = -A^{-1}\frac{\partial B}{\partial x_i} + A^{-1}\frac{\partial A}{\partial x_i}A^{-1}B \tag{15}$$

The first q rows of (15) are the second partials of $\hat{u}$ with respect to x and x_i while the last p rows are the second partials of λ.

We repeat this two-stage procedure iteratively until $\hat{u}$, λ, $\partial\lambda/\partial x'$ and $\frac{\partial}{\partial x_i}(\partial\lambda/\partial x')$, $i = 1,\ldots,p$, satisfy (9) and (10). If any u_i is required to be nonnegative, then the i^{th} equation in (9) will be replaced by an inequality <0 for $u_i=0$, with the equality holding for $u_i>0$ as in standard treatment of the method of Lagrange multipliers with nonnegative decision variables (see Dixit, 1990, p. 28). In the case of $u_i=0$, du_i disappears from equations (11), (12) and (13). This method using equations (8)-(15) can also provide analytical solutions to $\hat{u}$, λ, $\partial\lambda/\partial x'$, $\partial^2\lambda/\partial x_i\partial x'$ and V(x) if analytical solutions exist. For example, I have solved the problem of Abel (1983) analytically by this method.

For the first stage a good initial value for the matrix $\frac{\partial}{\partial x_i}(\partial\lambda/\partial x')$ of second partials of λ is the zero matrix, which is correct if the value function is locally quadratic near x so that $\lambda = \partial V/\partial x$ is linear. To find a good initial value for $\partial\lambda/\partial x'$, consider the nonstochastic control problem by letting $S = \Sigma = 0$. The Bellman equation is

$$\beta V(x) = \max_u \left\{ r(x,u) + f'\frac{\partial V}{\partial x} \right\} \tag{16}$$

Equation (9) for u and equation (10) for λ become respectively

$$\frac{\partial r}{\partial u_i} + \frac{\partial f'}{\partial u_i}\lambda = 0 \qquad i = 1,\ldots,q \tag{17}$$

$$\beta\lambda_i = \frac{\partial}{\partial x_i} r(x,\hat{u}) + \frac{\partial}{\partial x_i} f'(x,\hat{u})\lambda + \frac{\partial\lambda'}{\partial x_i} f(x,\hat{u}) \qquad i = 1,\ldots,p \tag{18}$$

As the second partials of λ are not involved, we can apply the two-stage iterative method to find $\hat{u}$, λ and $\partial\lambda/\partial x'$, using 0 for the initial value of $\partial\lambda/\partial x'$ in the first stage. This is the solution for the deterministic control problem. It provides a good initial value for $\hat{u}$, λ, and $\partial\lambda/\partial x'$ required in the first stage of our iterative method for solving the stochastic control problem.

The above iterative method computes an optimal value $\hat{u}$ of the control variable for any value x of the state variable. The solution is exact in the sense that the first-order conditions are satisfied except for rounding errors. No approximation of the value function is required. In fact, after $\hat{u}$, λ and $\partial\lambda/\partial x'$ are known for a given x, we can substitute them into (8) to evaluate V(x). Thus the value function is computed exactly point by point.

When the optimal control function $\hat{u}(x)$ can be so computed, we can study the dynamics of the system under optimal control by substituting $\hat{u}(x)$ for u in equation (2), yielding

$$dx = f(x,\hat{u}(x))dt + S(x,\hat{u}(x))dw \tag{19}$$

The dynamics of the costate variables λ can be readily obtained by using Ito's lemma and (10). By Ito's lemma,

$$d\lambda_i = \left\{ \frac{\partial\lambda_i}{\partial x'} f + \frac{1}{2}\, \mathrm{tr} \left[\frac{\partial}{\partial x_i} \left(\frac{\partial\lambda}{\partial x'}\right) \cdot \Sigma \right] \right\} dt + \frac{\partial\lambda_i}{\partial x'} Sdw \qquad (i = 1,\ldots,p) \tag{20}$$

where $\partial\lambda_i/\partial x = \partial\lambda/\partial x_i$. To eliminate the second partials of λ, we replace the sum of the two terms in curly brackets by the four remaining terms of (10) to obtain

$$d\lambda_i = -\left\{\frac{\partial}{\partial x_i} r(x,\hat{u}(x)) + \frac{\partial}{\partial x_i} f'(x,\hat{u}(x))\lambda + \frac{1}{2}\, \mathrm{tr}\left[\frac{\partial\lambda}{\partial x'} \frac{\partial}{\partial x_i} \Sigma(x,\hat{u}(x))\right] - \beta\lambda_i\right\}dt$$

$$+ \frac{\partial\lambda_i}{\partial x'} Sdw \qquad (i = 1,\ldots,p) \tag{21}$$

In the literature, equations (10) and (21) are known. For example, a version of (10) can be found in Benveniste and Scheinkman (1979) and is referred to by Sargent (1987, p. 21). A discrete-time version of (10) can be found in Chow (1975, pp. 158 and 281). A version of (21) can be found in Malliaris and Brock (1982, p. 112). However, equation (10) has not been treated as an essential component in obtaining the optimal control function $\hat{u}(x)$. The standard procedure using dynamic programming has been to solve the partial differential equation (8) for the value function V(x). This paper suggests that $\hat{u}(x)$ can be obtained directly by exploiting only the first-order conditions (9) and (10) without having to solve for the value function. Our computations yield $\hat{u}(x)$, V(x), the shadow price vector $\lambda(x)$, and the matrices of first and second partials of $\hat{u}$ and λ as given in stage two.

In this section, I have used the value function of dynamic programming to give an exposition of our method as most readers are familiar with dynamic programming. In the next section on discrete-time models, I will abandon the concepts of value function and the principle of optimality of dynamic programming completely and derive our method by using Lagrange multipliers to begin with. The first-order conditions (9) and (10) will become apparent. The reader will recognize that the stochastic nature of the problem does not present much additional complication as long as one can differentiate under the expectation operator. Using our method one does not encounter the curse of dimensionality which often occurs in solving the Bellman equation for the value function in dynamic programming.

3. Solution for Discrete-Time Models

For discrete-time models, the problem analogous to (1) and (2) is

$$\max_{\{u_t\}_{t=0}^{\infty}} E_0\left[\sum_{t=0}^{\infty} \beta^t r(x_t,u_t) \right] \tag{22}$$

subject to

$$x_{t+1} = f(x_t,u_t) + \varepsilon_{t+1} \tag{23}$$

where ε_{t+1} is an i.i.d. random vector with mean zero and covariance matrix Σ. Chow (1992) solves this problem by introducing the p×1 vector λ_t of Lagrange multipliers and setting to zero the derivatives of the Lagrangean expression

$$\mathscr{L} = E_0\left[\sum_{t=0}^{\infty}\left\{\beta^t r(x_t,u_t) - \beta^{t+1}\lambda'_{t+1}\left[x_{t+1} - f(x_t,u_t) - \varepsilon_{t+1}\right]\right\}\right] \tag{24}$$

with respect to u_t and x_t (t=0,1,2,...). The first-order conditions analogous to (9) and (10) are

$$\frac{\partial}{\partial u_t} r(x_t,u_t) + \beta \frac{\partial}{\partial u_t} f'(x_t,u_t)E_t\lambda_{t+1} = 0 \tag{25}$$

$$\lambda_t = \frac{\partial}{\partial x_t} r(x_t,u_t) + \beta \frac{\partial}{\partial x_t} f'(x_t,u_t)E_t\lambda_{t+1} \tag{26}$$

To justify the above method of solution, four observations can be made. First, if the problem were nonstochastic, i.e. if E_0 were absent and ε_{t+1} were constants, the use of Lagrange multipliers is justified since variables in different time periods are simply treated as different variables and the constraint $x_{t+1}-f(x_t,u_t)-\varepsilon_{t+1}=0$ for each period required a separate (vector) multiplier $\beta^{t+1}\lambda_{t+1}$, the scaling factor β^{t+1} being harmless but convenient. Second, if the problem were stochastic but unconstrained, the procedure is also justified because the expectation to be maximized is a function of the variables u_t, x_t and λ_t, and first-order conditions can be obtained by differentiation with respect to these variables, with the order of differentiation and taking expectation interchanged under suitable regularity conditions. Third, the method of Lagrange multipliers is to convert a constrained maximization problem to an unconstrained one by introducing the additional variables λ_t as is done above. Fourth, note that the problem is not to choose $u_0,u_1,\ldots$ all at once in an open-loop policy, but to choose u_t sequentially given the information x_t at time t in a closed-loop policy.

Since x_t is in the information set when u_t is to be determined, the expectations in equations (25) and (26) for the determination of u_t and λ_t at period t are E_t and not E_0.

Chow (1992) suggests using equation (26) instead of the Bellman equation to obtain an optimal control function $\hat{u}(x_t)$, but still recommends some global approximation to the Lagrange function $\lambda(x_t)$ in the process of obtaining the optimal control function. The present paper recommends doing away with such a global approximation as it is a poor strategy for finding $\hat{u}_t$ and λ_t satisfying the first-order conditions (25) and (26) given a specific x_t. Only a local approximation of $\lambda(x_t)$ near x_t is required. This is accomplished by a second-order Taylor expansion of $\lambda_{t+1}=\lambda(x_{t+1})$ in the evaluation of $E_t\lambda_{t+1}$.

$$\lambda_{i,t+1} = \lambda_i(x_{t+1}) = \lambda_i(x_t) + \frac{\partial \lambda_i}{\partial x'} (x_{t+1}-x_t) + \frac{1}{2}\, \mathrm{tr}\left[\frac{\partial \lambda_i}{\partial x \partial x'} (x_{t+1}-x_t)(x_{t+1}-x_t)' \right]$$

$$= \lambda_{it} + \frac{\partial \lambda_i}{\partial x'} \left[h(x_t,u_t) + \varepsilon_t \right] + \frac{1}{2}\, \mathrm{tr}\left[\frac{\partial \lambda_i}{\partial x \partial x'} (h+\varepsilon_{t+1})(h+\varepsilon_{t+1})' \right] \qquad (27)$$

where we have defined

$$h(x_t,u_t) = f(x_t,u_t) - x_t \qquad (28)$$

Hence

$$E_t\lambda_{i,t+1} = \lambda_{it} + \frac{\partial \lambda_i}{\partial x'} h(x_t,u_t) + \frac{1}{2}\, \mathrm{tr}\left[\frac{\partial \lambda_i}{\partial x \partial x'} (hh'+\Sigma) \right] \qquad (29)$$

Substitution for $E_t\lambda_{i,t+1}$ in (25) and (26), with the time subscript t omitted, gives

$$\frac{\partial}{\partial u_i} r(x,u) + \beta \frac{\partial}{\partial u_i} \sum_j f_j(x,u) \left\{ \lambda_j + \frac{\partial \lambda_j}{\partial x'} h(x,u) + \frac{1}{2}\, \mathrm{tr}\left[\frac{\partial^2 \lambda_j}{\partial x \partial x'} (hh'+\Sigma) \right] \right\}$$

$$i = 1,\ldots,q \qquad (30)$$

$$\lambda_i = \frac{\partial}{\partial x_i} r(x,u) + \beta \frac{\partial}{\partial x_i} \sum_j f_j(x,u)\left\{\lambda_j + \frac{\partial \lambda_j}{\partial x'} h(x,u) + \frac{1}{2} \operatorname{tr}\left[\frac{\partial^2 \lambda_j}{\partial x \partial x'} (hh'+\Sigma)\right]\right\}$$

$$i = 1,\ldots,p \tag{31}$$

Applying an analogous two-stage procedure as specified in section 2, one can solve (30) and (31) for $\hat{u}$, λ, $\frac{\partial \lambda}{\partial x'}$, and $\frac{\partial^2 \lambda_i}{\partial x \partial x'} = \frac{\partial^2 \lambda}{\partial x' \partial x_i}$, $i = 1,\ldots,p$ for any given x. I have applied this procedure to solve the optimal control problem associated with the baseline real business cycle model discussed in Chow (1992). With consumption and labor supply as two control variables and technology (Solow residual) and capital stock as two state variables, equations (30) and (31) are four equations for u_1, u_2, λ_1 and λ_2 which involve $\lambda_{ij} = \partial\lambda_i/\partial x_j$ $(i,j = 1,2)$ if we use only a first-order Taylor expansion for λ_i. Taking total differentials of (30) and (31) gives 4 equations in du_1, du_2, $d\lambda_1$, $d\lambda_2$, dx_1 and dx_2 written as (13). The second and third rows of the 4×2 matrix $-A^{-1}B$ provide estimates of λ_{ij} $(i,j = 1,2)$ to be used in the second stage of our iterative procedure. The procedure has been found to converge rapidly.

This procedure has an advantage over the procedure recommended by Chow (1992) in not having to approximate the (vector) function $\lambda(x)$ globally in finding a global optimal control function $\hat{u}=\hat{u}(x)$. We employ only a local approximation of $\lambda(x_{t+1})$ around x_t for which the value of optimal control $\hat{u}(x_t)$ has to be computed for a given x_t, and not a global approximation of λ for all x_t. A quadratic approximation to λ amounts to a cubic approximation to the value function V. Hence it is better to solve for λ using the first-order conditions than for V using the Bellman equation. Furthermore, this paper avoids the possibly large errors introduced in a global approximation to λ or V in finding the value $\hat{u}$ of optimal control associated with a particular x. By using a quadratic approximation for λ locally for a given x in evaluating a particular point $\hat{u}(x)$ of the optimal control function, one allows the control function, the global

Lagrange function $\lambda(x)$ and the value function $V(x)$ to take almost any form. In fact, even a linear approximation for λ locally, by dropping the second-order terms of (30) and (31), can yield a global $\lambda(x)$ better than a quadratic or a cubic global approximation. To sum up, seeking a global approximation for the value function in computing a global control function is a poor strategy in two important respects. A better function to seek is the Lagrange function. Much error is introduced by forcing a global approximation of λ or V in computing the value for the optimal control $\hat{u}$ associated with any specific value x of the state, as it occurs in practice when the method of dynamic programming is applied.

4. Conclusion - Why Not To Seek the Value Function Globally

The logic of seeking the value function in dynamic programming becomes apparent if one considers the problem of maximizing $r(x,u)$ with respect to vector u subject to a vector constraint $x = f(u)$. By the method of Lagrange multipliers, one differentiates the Lagrange expression $r-\lambda'(x-f)$ with respect to u, x and λ to obtain three first-order conditions

$$\frac{\partial r}{\partial u} + \frac{\partial f'}{\partial u}\lambda = 0 \; ; \qquad \frac{\partial r}{\partial x} - \lambda = 0 \; ; \qquad x - f = 0$$

which provide three equations for the variables u, x and λ. However, if one ignores the second first-order condition and solves the first first-order condition for $\hat{u}$ as a function of x and λ, one may substitute the result in the function r to be maximized to yield

$$V(x) = r(x,\hat{u}(x,\lambda))$$

This is the value function. It satisfies the above partial differential equation with $\lambda=\partial V/\partial x$. Bellman recommends that we solve this partial differential equation for the value function V. This paper suggests that we return the problem to one of solving algebraic equations by including the second first-order condition and including λ as a variable. It is obviously much better than ignoring this condition and solving a partial differential equation for the value function globally, except for very special examples. In fact, after $\hat{u}$ is obtained for a given x, we can substitute it in r to evaluate the value function at x. Thus both the control function and the value function can be obtained exactly point by point.

References

Abel, Andrew B., 1983, Optimal investment uncer uncertainty, American Economic Review 73, 228-233.

Bellman, Richard, 1957, Dynamic programming (Princeton University Press, Princeton).

Benveniste, Lawrence and Jose Scheinkman, 1979, On the differentiability of the value function in dynamic models of economics, Econometrica 47, 727-732.

Chow, Gregory C., 1975, Analysis and Control of Dynamic Economic Systems (John Wiley & Sons, Inc., New York).

Chow, Gregory C., 1992, Dynamic optimization without dynamic programming, Economic Modelling 9, 3-9.

Dixit, Avinash K., 1990, Optimization in Economic Theory (Oxford University Press, Oxford).

Malliaris, A.G. and W.A. Brock, 1982, Stochastic methods in economics and finance (Elsevier Science Publishers B.V., Rotterdam).

Sargent, Thomas J., 1987, Dynamic macroeconomic theory (Harvard University Press, Cambridge).