The Lagrange Method for Dynamic Optimization
An Alternative to Dynamic Programming

Gregory C Chow

Outline of Presentation:

Brief Explanation of the Lagrange Method

Solution to a Standard Problem in Discrete Time

Numeral Solution by Linear Approximation of Lagrange Function

An Example from Economic Growth

An Example in Optimal Consumption and Portfolio Selection

Stochastic Models in Continuous Time

Solution by the Lagrange Method – Reduction to Pontryagin's

Maximum Principle as a Special Case

An Example in Optimal Consumption and Portfolio Selection

Characteristics of the Lagrange Method

Further Applications

References:

1. *Dynamic Economics: Optimization by the Lagrange Method*. Oxford, 1997.
2. "Equity Premium and Consumption Sensitivity …," *Journal of Economic Dynamics and Control,* vol. 26, no. 9, 2002
3. "Duplicating Contingent Claims by the Lagrange Method," *Pacific Economic Review,* vol. 4, no. 3, 1999.

BRIEF EXPLANATION OF THE LAGRANGE METHOD FOR DYNAMIC OPTIMIZATION

1. Start with the constrained maximization problem of max r(x,u) subject to the constraint x=f(u) using the Lagrange expression L = r(x,u) - λ[x-f(u)]. First-order conditions for max are obtained by setting to zero the derivatives of L with respect to x, u and λ. The resulting equations are solved for the three variables.

2. Easy to generalize this procedure to the case of many periods when the objective function is a weight sum of r(x(t), u(t)) over time t and the constraints are x(t+1)=f(x(t), u(t)) where we call x the state variable and u the control variable. We can set up the same Lagrange expression L using a multiplier λ(t+1) associated with the constraint x(t+1)-f(x(t), u(t))=0, t=1, …. Express optimum u(t) as a "policy function" g(x(t)).

3. Also straight-forward to generalize the above to the stochastic case with x(t+1)= f(x(t), u(t), ε(t)), ε(t) stochastic. We now have an expectation operator before the objective function and before the sum of all the products of λ(t+1) and the constraints. The first order conditions can still be obtained by differentiation after the summation sign.

(2.4), the choice is between (2.10) and (2.7). Equation (2.7) could be obtained by differentiating (2.10) with respect to x_1. To obtain the optimal control $u_1 = g_1(x_1)$, (2.7) is preferable because by both methods knowledge of V_2 is unnecessary in obtaining $g_1(x_1)$ from (2.9) but $\partial V_2/\partial x_2$ or λ_2 is required. For each period t, to seek the value function V_{t+1}, either analytically or numerically, and then differentiate it to find the optimal u_t is an uneconomical method, as I pointed out in section 1.2. After the optimization problem is solved by the Lagrange method, one can obtain the value function by substituting the optimal control function into the dynamic model to evaluate the objective function or by integrating the Lagrange function.

2.3 Solution of a Standard Dynamic Optimization Problem

A standard dynamic optimization problem is

$$\max_{\{u_t\}_{t=0}^{\infty}} E_0\left[\sum_{t=0}^{\infty} \beta^t r(x_t, u_t)\right], \tag{2.12}$$

subject to

$$x_{t+1} = f(x_t, u_t) + \varepsilon_{t+1}, \tag{2.13}$$

in which E_0 is expectation given information at time 0, and ε_{t+1} is an independent and identically distributed (i.i.d.) random vector with mean zero and covariance matrix Σ. As in section 2.1, this problem is solved by introducing the $p \times 1$ vector λ_t of Lagrange multipliers and setting to zero the derivatives of the Lagrangean expression

$$\mathscr{L} = E_0\left[\sum_{t=0}^{\infty}\left\{\beta^t r(x_t, u_t) - \beta^{t+1}\lambda'_{t+1}\left[x_{t+1} - f(x_t, u_t) - \varepsilon_{t+1}\right]\right\}\right], \tag{2.14}$$

with respect to u_t ($t = 0, 1 \ldots$) and x_t ($t = 1, 2 \ldots$). The first-order conditions analogous to (2.4) and (2.7) are

$$\frac{\partial}{\partial u_t} r(x_t, u_t) + \beta \frac{\partial}{\partial u_t} f'(x_t, u_t) E_t \lambda_{t+1} = 0, \tag{2.15}$$

$$\lambda_t = \frac{\partial}{\partial x_t} r(x_t, u_t) + \beta \frac{\partial}{\partial x_t} f'(x_t, u_t) E_t \lambda_{t+1}. \tag{2.16}$$

To justify this method of solution, four observations can be made. First, if the problem were nonstochastic, that is, if E_0 were absent and ε_{t+1} were constants, the use of Lagrange multipliers is justified because 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 requires a separate (vector) multiplier $\beta^{t+1}\lambda_{t+1}$. The factor β^{t+1} is used to discount the marginal value λ_{t+1} of x_{t+1} dated at $t + 1$. 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. Because x_t is in the information set when u_t is to be determined, the expectations in equations (2.15) and (2.16) for the determination of u_t and λ_t at period t are E_t and not E_0. As in section 2.1, I suggest using equation (2.16) instead of the Bellman equation (2.11), with E_t inserted before V_{t+1}, to obtain the optimal control function $u_t = g_t(x_t)$. In the next section, I present a well-known numerical method to solve (2.15) and (2.16) for a pair of functions $u_t = g(x_t)$ and $\lambda_t = \lambda(x_t)$ when the functions g and λ are time invariant in an infinite horizon dynamic optimization problem.

2.4 Numerical Solution by Linear Approximations for λ and *g*

To solve the first-order conditions (2.15) and (2.16) for u and λ, one can approximate the function $\lambda(x)$ by a linear function

$$\lambda(x) = Hx + h \tag{2.17}$$

One can also approximate $\partial r/\partial x$, $\partial r/\partial u$, and f by linear functions

$$\frac{\partial}{\partial x} r(x, u) = K_{11}x + K_{12}u + k_1 \tag{2.18}$$

$$\frac{\partial}{\partial u} r(x, u) = K_{21}x + K_{22}u + k_2 \tag{2.19}$$

$$f(x, u) = Ax + Cu + b \tag{2.20}$$

The linear functions are obtained by first-order Taylor expansion about some $(\bar{x}, \bar{u})$. The values $(\bar{x}, \bar{u})$ might be the sample means of the state and control variables in an econometric application. They may be the steady-state values of these variables under optimal control for the nonstochastic version of the optimal problem obtained by setting $\varepsilon_t = 0$. These values are obtained by solving equations (2.13), (2.15), and (2.16) as algebraic equations for x, u and λ omitting all time subscripts and the expectation sign (see problem 5).

In the first step of the solution procedure, the right-hand sides of (2.19) for $\partial r/\partial u$ and of $\lambda_{t+1} = Hx_{t+1} + h$ are substituted into (2.15) to yield

$$\partial r/\partial u + \beta C' E_t \lambda_{t+1} = K_{22}u + K_{21}x + k_2 + \beta C'H(Ax + Cu + b) + \beta C'h = 0. \tag{2.21}$$

Solving (2.21) for u gives $u = Gx + g$, in which

$$G = -(K_{22} + \beta C'HC)^{-1}(K_{21} + \beta C'HA), \tag{2.22}$$

$$g = -(K_{22} + \beta C'HC)^{-1}[k_2 + \beta C'(Hb + h)]. \tag{2.23}$$

In the next step, observe

$$\begin{aligned} E_t\lambda_{t+1} &= H(Ax + Cu + b) + h \\ &= H[(A + CG)x + Cg + b] + h, \end{aligned} \tag{2.24}$$

which is linear in x. Substituting (2.24) for $E_t\lambda_{t+1}$ and (2.18) for $\partial r/\partial x$ in (2.16),

$$\lambda_t = \lambda(x) = Hx + h = K_{11}x + K_{12}(Gx + g) + k_1 + \beta A'\{H[(A + CG)x + Cg + b] + h\}.$$

Equating coefficients on both sides gives

$$H = K_{11} + K_{12}G + \beta A'H(A + CG), \tag{2.25}$$

$$h = (K_{12} + \beta A'HC)g + k_1 + \beta A'(Hb + h). \tag{2.26}$$

Equations (2.22) and (2.25) are used to solve for G and H. Given H and G, (2.23) and (2.26) are used to solve for g and h. Equations (2.22) and (2.25) can be combined to yield the following equation in H alone:

$$H = K_{11} + \beta A'HA - (K_{12} + \beta A'HC)(K_{22} + \beta C'HC)^{-1}(K_{21} + \beta C'HA). \tag{2.27}$$

Equation (2.27) is known as the matrix Riccati equation. It can be solved iteratively. Given H, one can find G by using (2.22).

The method just described is an example of the *method of undetermined coefficients*. To solve two functional equations (2.15) and (2.16) for $u = g(x)$ and $\lambda = \lambda(x)$, assume both to be linear with coefficients (G, g, H, h). Substitute the linear functions into the equations. Equate the coefficients on both sides to solve for the unknown coefficients. The above method yields the same answer as given by the solution to the well-known linear quadratic optimal control problem. The latter problem assumes the function f to be linear and the function r to be quadratic and, hence, $\partial r/\partial x$ and $\partial r/\partial u$ to be linear. Under these assumptions the optimal control function g and the Lagrange function λ will be linear, and the value function V will be quadratic. When f is nonlinear and r is not quadratic, one may linearize f, $\partial r/\partial x$, and $\partial r/\partial u$ about a certain point $(\bar{x}, \bar{u})$ and solve equations (2.22), (2.23), (2.25), and (2.26) to find a linear control function $Gx + g$ together with a linear Lagrange function $\lambda = Hx + h$.

The first-order conditions (2.15) and (2.16) may have no solution when linear approximations are used for λ and g. Assuming that solutions for linear λ and g exist and that (2.27) can be solved for H and, hence, (2.22) can be solved for G, one finds an approximately optimum control function $u = g + Gx$. If a time-invariant linear control function $g + Gx$ is found, the steady-state value for x, if it exists, can be found by solving

$$x = f(x,\ g + Gx)$$

Note that a time-invariant optimal control function does not necessarily produce a steady state for x_t, depending on the dynamic model specified by $f(x, u)$. For example, if there are two explosive roots in the 2×2 matrix A in a linear f,

Economic Growth

3.1 The Brock-Mirman Growth Model

Consider a very simple optimal growth model

$$\max_{\{c_t\}} \sum_{t=0}^{\infty} E_t \beta^t \ln c_t \tag{3.1}$$

subject to

$$k_{t+1} = k_t^{\alpha} z_t - c_t \tag{3.2}$$

in which consumption c_t is the control variable, and capital stock k_t and random shock z_t are state variables. By using dynamic programming, set up the Bellman equation

$$V(x_t) = \max_{u_t} \{ r(x_t,\ u_t) + \beta E_t V(x_{t+1}) \} \tag{3.3}$$

for this problem and solve for the value function. Assuming, in favor of Bellman, that I conjecture correctly that the value function takes the form

$$V(k_t,\ z_t) = a + b \ln k_t + c \ln z_t, \tag{3.4}$$

where a, b, and c are three parameters to be determined. Evaluating $V_{t+1} = V(k_{t+1}, z_{t+1})$ by using the dynamic constraint and taking its conditional expectation E_t give

$$V(k_{t+1},\ z_{t+1}) = a + b \ln(k_t^{\alpha} z_t - c_t) + c \ln z_{t+1}$$

$$E_t V(k_{t+1},\ z_{t+1}) = a + b \ln(k_t^{\alpha} z_t - c_t) + c E_t \ln z_{t+1},$$

where $E_t \ln z_{t+1}$ is assumed to be zero. Maximizing the expression in curly brackets in (3.3) by differentiation with respect to c_t, one obtains

$$\frac{\partial \{\quad\}}{\partial c_t} = c_t^{-1} - \beta b (k_t^{\alpha} z_t - c_t)^{-1} = 0$$

yielding the optimal control function

$$c_t = (1+\beta b)^{-1} k_t^{\alpha} z_t.$$

Substituting this optimal c_t into the expression in curly brackets, one obtains

$$\max_{c_t}\{\quad\} = -\ln(1+\beta b) + \beta a + \beta b \ln\left(\frac{\beta b}{1+\beta b}\right) + \alpha(1+\beta b)\ln k_t + (1+\beta b)\ln z_t, \tag{3.5}$$

which, by the Bellman equation (3.3), equals the value function (3.4). By equating coefficients of (3.4) and (3.5), one solves for the unknown parameters of the value function and obtains, after some algebraic manipulations,

$$\begin{aligned} a &= [1-\beta]^{-1}\left[\ln(1-\alpha\beta) + \beta\alpha(1-\alpha\beta)^{-1}\ln(\alpha\beta)\right] \\ b &= \alpha(1-\alpha\beta)^{-1} \\ c &= (1-\alpha\beta)^{-1}. \end{aligned} \tag{3.6}$$

The method of Lagrange multipliers is simpler than dynamic programming because it does not seek the value function. By this method, solve the first-order conditions given by (2.15) and (2.16) of chapter 2, or

$$c_t^{-1} - \beta E_t \lambda_{t+1} = 0 \tag{3.7}$$

$$\lambda_t = \beta\alpha k_t^{\alpha-1} z_t E_t \lambda_{t+1} \tag{3.8}$$

If the time horizon is T and $k_{T+1} = 0$, the optimal c_T is $k_T^{\alpha} z_T$, which suggests the function $c_t = d k_t^{\alpha} z_t$, d being a parameter to be determined. By using this conjecture for c_t and combining conditions (3.7) and (3.8), one obtains $\lambda_t = d^{-1}\alpha k_t^{-1}$. By using this to evaluate $\lambda_{t+1} = d^{-1}\alpha(k_t^{\alpha} z_t - d k_t^{\alpha} z_t)^{-1}$ on the right side of (3.8) and equating coefficients on both sides of equation (3.8), one obtains $d = 1 - \alpha\beta$. Thus, easily obtained are the optimal control function for c_t and the Lagrangean function λ_t, namely,

$$\begin{aligned} c_t &= (1-\alpha\beta) k_t^{\alpha} z_t \\ \lambda_t &= (1-\alpha\beta)^{-1} \alpha k_t^{-1}. \end{aligned} \tag{3.9}$$

Only one unknown parameter $d = 1 - \alpha\beta$ is required, rather than three in the case of solving the Bellman equation (3.3). Note that $\lambda_t = \partial V/\partial k_t$.

This example demonstrates that, as pointed out in sections 1.2 and 2.2, to obtain the optimal control function, it is unnecessary to seek the value function, because the first-order condition (3.7) involves only $\lambda_{t+1} = \partial V(x_{t+1})/\partial x_{t+1}$ and not $V(x_{t+1})$ itself. The method of Lagrange multipliers is simpler in this example also because only the partial derivative of V with respect to the state variable k_t subject to constraint is required.

Characteristics of the Lagrange Method

1. No Bellman equation is required because the value function is not used in deriving the optimal policy.

2. Since L evaluated at the optimal policy equals the value function, and λ is the partial of L with respect to y, it is the partial of the value function with respect to y.

3. In my book *Dynamic Economics*, I have shown that in many examples the Lagrange method gives a simpler (algebraically and/or computationally) solution than dynamic programming. The main reason is that dynamic programming seeks the value function which contains more information than is required – it asks us to do extra work. To see this point, why not apply dynamic programming to solve a deterministic control problem in continuous time – in the latter case the Lagrange method is reduced to the Maximum Principle which is widely used instead of dynamic programming.

4. Dynamic programming fails if the model consists of expectations of future variables. M. Woodford, "Optimal monetary inertia:" x= output gap (deviation of log real output from trend minus "natural rate" of output; r= deviation of interest rate (control variable) from a steady state value; r^n= natural rate of interest; π= inflation rate. Model consists of two equations for x_t and π_t in which $E_t x_{t+1}$ and $E_t \pi_{t+1}$ appear. The Lagrange method is applicable, but dynamic programming is not.

Optimization for Stochastic Models in Continuous Time

We start with an intuitive explanation first. Consider a small time interval dt and treat the stochastic constraint as $dx = x(t+dt)-x(t) = f(x(t), u(t))dt + Sdz(t)$, where z(t) is a Brownian motion. Its increment $dz(t)=z(t+dt)-z(t)$ has a variance proportional to dt. Since successive increments are statistically independent, the variance of the increment for three unit time intervals is the sum of the three variances, or three times the variance for one unit time interval. dz is of order $\sqrt{dt}$, much smaller than dt when dt is small and can be ignored when we deal with terms of order dt. We replace the sum in the Lagrange expression over the constraints at all time t by an appropriate integral as in the above expression for L in our problem. Then proceed to differentiate L with respect to the state and control variables at each t and set the derivatives equal to zero. This is the method. One has to be more precise in defining the integral in this Lagrange expression L. To make the definition consistent with stochastic calculus, all one needs to do is to replace $\lambda(t+dt)$ by the sum $\lambda(t)+d\lambda$ and to break up the integral into the two parts, with the second involving the quadratic variation $[d\lambda, dz]$.

Merton (1969): To determine consumption c and fraction w_i of total wealth W to be invested in asset i with mean rate of return α_i and instantaneous standard deviation s_i.

$$L = \int E_t\{e^{-\beta t}u(c)dt - e^{-\beta(t+dt)}\lambda(t+dt)\,[dW-(W\Sigma_i w_i\alpha_i - c)dt - W\Sigma_i w_i s_i dz_i]\}$$

Problems and Economic Applications

Problems on the Lagrange method. Chapter 2, Problems 1-5, 9.

Applications to models of economic growth. Chapter 3, Problems 5 (growth model based on human capital and fertility), 6 (growth model based on technology) and 7 (growth model based on research and development).

Theories of market equilibrium. Economic examples drawn from Stokey and Lucus, *Recurve Methods in Dynamic Economics.* Chapter 4.

Business Cycles. Chapter 5.

Dynamic Games, Chapter 6.

Applications to financial economics. Chapter 7.

Models of Investment. Chapter 8.

Numerical Methods. Chapter 9.

Samuelson, Paul A., "Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics*, **51** (1969), 239-246.

Assume that there are two assets in the financial market: a stock and a bond. Let $R_t = \left[R_{1,t} \quad R_{2,t}\right]^T$ be a vector of returns to these two assets, meaning that one dollar invested in asset *i* at the beginning of period *t* will result in $R_{i,t}$ dollars at the end of period *t*. The covariance matrix of R_t will be denoted by $\Sigma = (\sigma_{ij})_{2\times 2}$. The consumer-investor is assumed to construct a self-financing portfolio with the two assets and consume C_t during period *t*. Let w_t be the proportion of wealth invested in the stock. Following Samuelson (1969), the beginning-of-period value Z_t of such a portfolio is governed by

$$Z_{t+1} = (Z_t - C_t)\left[w_t \quad 1 - w_t\right]R_t \tag{1}$$

$$L = \sum_{t=0}^{\infty} E_t \beta^t \left\{u(C_t) - \beta\lambda_{t+1}\left[Z_{t+1} - (Z_t - C_t)[w_t \quad 1 - w_t]\left(R_t\right)\right]\right\} \tag{3}$$

First-order conditions are obtained by setting to zero the partial derivatives of *L with* respect to the control variables C_t and w_t and the state variable Z_t. .

.