Dynamic Optimization

In economics, we frequently encounter problems in which we seek to maximize an expression that is in the form of an integral of some basic quantity, like utility or profits, over an interval of time. Most commonly, these expressions are in the form of discounted present values. There are three basic solution strategies for the continuous time version of such problems; the calculus of variations, continuous time dynamic programming, and the maximum principle. For most economic examples, the preferred alternative is the maximum principle, or optimal control.

Optimal Control

To utilize optimal control, we need to specify the problem in terms of a state variable, $y(t)$ and a control variable, $u(t)$. The most basic form of the optimal control problem is characterized as:

$$Max \int_0^T F(t,y,u)dt$$

Subject to $\dot{y} = f(t,y,u)$, $y(0) = y_0$ and $y(T) = y_T$

This integral expression is suggested by the particular problem we wish to solve. The variable, $u(t)$, the control variable, is to be set in such a way as to maximize the objective function. We will later introduce some possible restrictions on the path of $u(t)$, but for now we leave any such restrictions as unspecified. The trajectory of $y(t)$, the state variable, must conform to a law of motion in which the first time derivative of the state path with respect to time is a given function, f, of the state variable and the control variable. Note that the function, f, may also be time-varying. For now, we take $u(t)$ and $y(t)$ as scalar-values functions. We will later allow multiple control variables and multiple state variables. We are given an explicit starting value, y_0, and an explicit terminal value, y_T, of the state variable, but the initial value of the control variable may be freely chosen. For now, we also specify a given ending time, T, for the planning horizon.

Our solution technique requires that the control path, $u(t)$, be piecewise continuous, although it may have a discrete number of jumps. The state path, $y(t)$, must be continuous, but it need only be piecewise differentiable. That is, it may have a discrete number of sharp corners. We may also specify a required control set, $\mathcal{U}$, such that $u(t) \in \mathcal{U}$. The set, $\mathcal{U}$, must be closed and convex. We now derive conditions which characterize the optimal path of the control variable. As a prelude define:

$$L \equiv \int_0^T [F(t,y,u) + \lambda(t)(f(t,y,u) - \dot{y})]dt,$$

where $\lambda(t)$ is an as of yet unspecified time function. Clearly, we require that the solution path for the state variable, $y^*(t)$, be such that:

$$\dot{y}^* = f(t, y^*(t), u^*(t))$$

Rearrange L to the form:

$$L = \int_0^T [F(t,y,u) + \lambda(t)f(t,y,u) - \lambda\dot{y}]dt$$

The variable, $\lambda(t)$, is denoted as a costate variable. The equation above obviously holds for any arbitrarily selected time-path, $\lambda(t)$. Define $H(t,y(t),u(t),\lambda(t)) \equiv F(t,y,u) + \lambda(t)f(t,y,u)$. We will refer to H as the Hamiltonian for the problem. By definition,

$$L = \int_0^T [H(t,y,u,\lambda) - \lambda\dot{y}]dt$$

Integration by parts yields:

$$\int_0^T \lambda\dot{y}dt + \int_0^T \dot{\lambda}ydt = \lambda y|_0^T$$

Therefore:

$$L(y(t),u(t),T) = \int_0^T [H(t,y,u,\lambda) + \dot{\lambda}y]dt + \lambda(0)y(0) - \lambda(T)y(T)$$

Now we introduce perturbed paths, $y(t) + \Delta y(t)$ and $u(t) + \Delta u(t)$. We hope to characterize conditions on $u(t)$ and $y(t)$ that are true for the optimal paths, $u^*(t)$ and $y^*(t)$, but which are not true for alternative path, like the paths $y(t) + \Delta y(t)$ and $u(t) + \Delta u(t)$. The paths, $\Delta y(t)$ and $\Delta u(t)$, are, for the most part, arbitrarily selected. However, clearly $\hat{y} \equiv y(t) + \Delta y(t)$ and $\hat{u}(t) \equiv u(t) + \Delta u(t)$, must be related according to $\frac{d\hat{y}}{dt} = f(t,\hat{y},\hat{u})$. The path, $\lambda(t)$ remains unperturbed. The most simple form of the optimal control problem also specifies a fixed ending value for the state path, $y_T = y(T)$.

Next introduce the difference term:

$$\Delta L = L(y(t) + \Delta y(t), u(t) + \Delta u(t)) - L(y(t), u(t))$$

If $y(t)$ and $u(t)$ comprise an optimal path, $\Delta L \leq 0$, with $\Delta L = 0$, at an interior maximum. We now rewrite ΔL in the form:

$$\begin{aligned}\Delta L = {} & \int_0^T [H(t, y + \Delta y, u + \Delta u, \lambda) + \dot{\lambda}(y + \Delta y)]dt \\ & + \lambda(0)y(0) + \lambda(0)\Delta y(0) - \lambda(T)y(T) - \lambda(T)\Delta y(T) \\ & - \int_0^T [H(t,y,u,\lambda) + \dot{\lambda}y]dt - \lambda(0)y(0) + \lambda(T)y(T)) \leq 0\end{aligned}$$

Recall that the paths $\Delta y(t)$ and $\Delta u(t)$ may be arbitrarily selected. For the problem at hand, it is useful to select paths such that $\Delta y(0) = 0$ and $\Delta y(T) = 0$. Expanding the integral in terms of Δy, and Δu, and collecting terms, we obtain:

$$\Delta L = \int_0^T \left[\frac{\partial H}{\partial u}\Delta u(t) + \left(\frac{\partial H}{\partial y} + \dot{\lambda} \right)\Delta y(t) \right] dt$$

This inequality holds for the optimal path, $u^*(t)$, and the implied path, $y^*(t)$, for arbitrary $\Delta u(t)$ and the implied $\Delta y(t)$. Therefore this condition allows us to distinguish between the optimal path and other, suboptimal paths. Because $\Delta L = 0$ must hold for all of the infinitely-dimensional $\Delta u(t)$ (and the implied $\Delta y(t)$) paths, each term in $\Delta L = 0$ must separately equal zero for every point in time on $[0, T]$. The inequality will be a strict equality for an interior solution. We therefore separately require:

$$\frac{\partial H}{\partial u} = 0 \text{ and } \frac{\partial H}{\partial y} = -\dot{\lambda},\ \forall t \subset [0, T]$$

To summarize, for an optimal time path, $u^*(t)$, it must be true that:

$$\frac{\partial H}{\partial u} = 0,\ \ \frac{\partial H}{\partial y} = -\dot{\lambda},\ \ \dot{y} = f(t, y, u)$$

$$y(0) = y_0 \text{ and } y(T) = y_T$$

So far, we have limited the nature of the problem in such a way that it may only be solved if we a fixed initial condition, $y(0) = y_0$, and a fixed ending value, $y(T) = y_T$. We have also specified a fixed, finite time horizon, T. We now seek solutions to a more general class of problems. First, consider the case in which the ending value, $y(T)$, may be freely chosen. In this case, the term in the ΔL equation, $\lambda(T)\Delta y(T)$, may no longer vanish for any arbitrary $\Delta y(t)$ path. To insure that this term equals zero, we require that $\lambda(T) = 0$.

Another special case is that in which the ending value is not explicitly given, but where the ending value must be greater than or equal to some specified value, $y_{\min}(T)$. Recall the term, $-\lambda(T)\Delta y(T)$, in the ΔL equation. There are two possibilities. First, suppose that $y(T) \geq y_{\min}(T)$ binds as a strict equality, $y(T) = y_{\min}(T)$. In this case, it must be the case that we could increase L by decreasing $y(T)$ below $y_{\min}(T)$. That is, $\Delta y(T) < 0$, must increase L. To hypothetically increase the term, $-\lambda(T)\Delta y(T)$, in ΔL, for $\Delta y(T) < 0$ requires that $\lambda(T) > 0$. The second possibility is that the constraint does not bind as a strict equality. In this case, for $\Delta L = 0$ for arbitrary $\Delta y(T)$, the term, $-\lambda(T)\Delta y(T)$, in ΔL can only be zero for arbitrary $\Delta y(T)$, if $\lambda(T) = 0$. We may therefore summarize the appropriate additional, complementary slackness condition as:

$$\lambda(T) \geq 0,\ (y(T) - y_{\min}(T)) \ \geq 0,\ \lambda(T)(y(T) - y_{\min}(T)) \ = 0$$

Either $y(T) = y_{\min}(T)$ binds and $\lambda(T)$ may be freely chosen, or $\lambda(T) = 0$ binds and $y(T)$ may be chosen freely. Alternatively, if we had specified a maximum value for $y(T)$, $y_{\max}(T)$, analogous reasoning would derive the alternative complementary slackness condition:

$$\lambda(T) \leq 0,\ (y(T) - y_{\max}(T)) \ \leq 0,\ \lambda(T)(y(T) - y_{\max}(T)) \ = 0$$

We may finally consider cases in which the terminal time, T, may also be chosen as part of the maximization procedure. In this case, we have a perturbed endpoint, ΔT, in addition to a perturbed terminal value, $\Delta y(T)$. We keep $\lambda(t)$ unperturbed. We also set $\Delta y(0) = \Delta u(0) = 0$. Clearly, there is an implied relationship between $y(t) + \Delta y(t)$ and $u(t) + \Delta u(t)$, since $\dot{y} = f(t, y, u)$. There may also be an implied relationship between ΔT and $\Delta y(T)$, due to a terminal surface in the form of a relationship between the terminal time, T, and the terminal state, $y(T)$. Denote the terminal surface by the implicit function, $\Phi(y, t) = 0$. A sample terminal surface is depicted below.

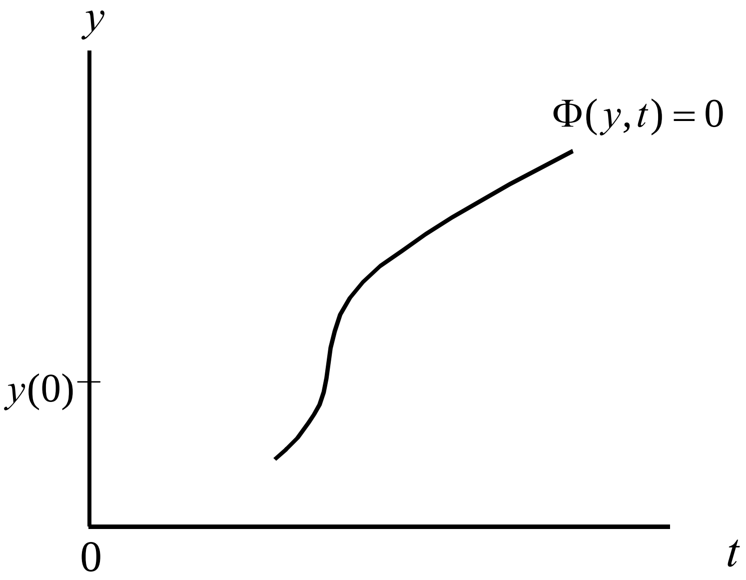

The implied relationship between $\Delta y(T)$ and ΔT can be depicted as:

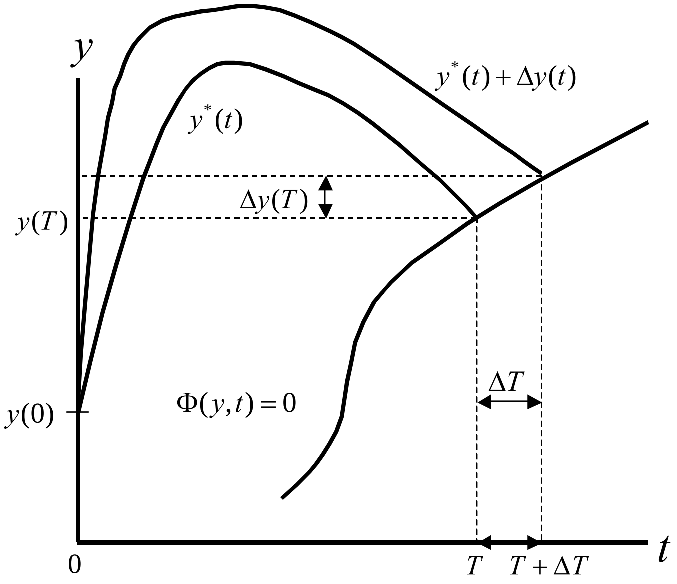

Now consider the amended difference term:

$$\Delta L = L(y(t)+\Delta y(t), u(t)+\Delta u(t), T+\Delta T) - L(y(t), u(t), T)$$

If $y(t)$ and $u(t)$ comprise an optimal path, $\Delta L \leq 0$, with $\Delta L = 0$, at an interior maximum. We now rewrite ΔL in the form:

$$\begin{aligned}\Delta L = & \int_0^{T+\Delta T} [H(t, y+\Delta y, u+\Delta u, \lambda) + \dot{\lambda}(y+\Delta y)]dt \\ & + \lambda(0)y(0) - \lambda(T)y(T) - \Delta(\lambda y)|_T \\ & - \int_0^T [H(t,y,u,\lambda) + \dot{\lambda} y]dt - \lambda(0)y(0) + \lambda(T)y(T)) \leq 0\end{aligned}$$

Expanding the integral in terms of Δy, Δu, and ΔT, and collecting terms, we obtain:

$$\Delta L = \int_0^T \left[\frac{\partial H}{\partial u}\Delta u(t) + \left(\frac{\partial H}{\partial y} + \dot{\lambda} \right) \Delta y(t) \right] dt$$

$$+ [H(t, y(t), u(t), \lambda(t)) + \dot{\lambda} y] \,\Big|_{t=T} \Delta T - \Delta(\lambda y)|_{t=T} \leq 0$$

This inequality holds for the optimal path, $u^*(t)$, and the implied path $y^*(t)$, for arbitrary $\Delta u(t)$. Therefore this condition allows us to distinguish between the optimal path and other, suboptimal paths. The inequality will be a strict equality for an interior solution. As in the simpler cases, we separately require:

$$\frac{\partial H}{\partial u} = 0, \ \frac{\partial H}{\partial y} = -\dot{\lambda}, \ \dot{y} = f(t, y, u)$$

When the terminal surface is vertical (T is fixed), $\Delta T = 0$, and the remaining terms in ΔL are given by:

$$-\Delta(\lambda y)|_{t=T} = -\lambda(T)\Delta y(T) \leq 0$$

More generally, we require that:

$$[H(t,y(t),u(t),\lambda(t)) \; + \dot{\lambda}y]\,\big|_{t=T}\Delta T - \Delta(\lambda y)|_{t=T}$$

$$= [H(t,y(t),u(t),\lambda(t)) \; + \dot{\lambda}y]\,\big|_{t=T}\Delta T - [\lambda\Delta y + y\Delta\lambda]|_{t=T} \leq 0$$

Recall that $y(T)$ and T are governed by the terminal surface constraint, $\Phi(y,t) = 0$, and so:

$$[\lambda\Delta y + y\Delta\lambda]|_{t=T} = \left[\lambda\left(\left(\frac{dy}{dt}\right)_{\Phi=0}\Delta T\right) + y\dot{\lambda}\Delta T\right]_{t=T}$$

The transversality terms may therefore be rewritten as:

$$[H(t,y(t),u(t),\lambda(t)) \; + \dot{\lambda}y]\,\big|_{t=T}\Delta T - \left[\lambda\left(\left(\frac{dy}{dt}\right)_{\Phi=0}\Delta T\right) + y\dot{\lambda}\Delta T\right]_{t=T}$$

$$= \left[H(t,y(t),u(t),\lambda(t)) \; - \lambda\left(\frac{dy}{dt}\right)_{\Phi=0}\right]\bigg|_{t=T}\Delta T \leq 0$$

For variable T, the additional transversality condition is given by:

$$H(T,y(T),u(T)) \; - \lambda(T)\left(\frac{dy}{dt}\right)_{\Phi=0} = 0, \text{ for free } T, \text{ and } \Phi(y(T),T) = 0$$

If we also require that $y(T) \geq y_{\min}$, we must add a complementary slackness condition. Recall that ΔL includes the term, $-[\lambda\Delta y]|_{t=T}$. When the inequality constraint binds as a strict equality, $\Delta y(T) \geq 0$, and so we would only benefit from relaxing this constraint (by setting $\Delta y(T) < 0$) by reducing $(y - y_{\min})|_{t=T}$ below zero) when $\lambda(T) > 0$. The CS condition is therefore given by:

$$\lambda(T) \geq 0,\ (y(T) - y_{\min}) \geq 0,\ \lambda(T)(y(T) - y_{\min}) = 0$$

If we are constrained by a requirement that $T \leq T_{\max}$, we may derive the additional complementary slackness condition:

$$H(T) - \lambda(T)\left(\frac{dy}{dt}\right)_{\Phi=0} \geq 0,\ (T - T_{\max}) \leq 0,$$

$$\left(H(T) - \lambda(T)\left(\frac{dy}{dt}\right)_{\Phi=0}\right)(T - T_{\max}) = 0.$$

To motivate this last condition, note that for $\left[H(t,y(t),u(t),\lambda(t)) \; - \lambda\left(\frac{dy}{dt}\right)_{\Phi=0}\right]\Big|_{t=T} > 0$, we could potentially increase L ($\Delta L > 0$) if we were allowed an extension of time beyond $T_{\max}$, by setting $\Delta T > 0$.

We may also modify the Maximum Principle to allow for corner solutions for the control path, $u(t)$. Although the proof is quite lengthy, the condition for the optimal $u(t)$ may more generally be characterized as:

$$H(t,y^*,u^*,\lambda^*) \geq H(t,y^*,u,\lambda^*),\ \forall t \in [0,T]$$

At corners, typically $\frac{\partial H}{\partial u} \neq 0$. Consider the specific control constraint, $u(t) \geq c$. If $u(t) \geq c$ binds as an equality, then $\frac{\partial H}{\partial u} < 0$. Hypothetically, H could be increased by reducing u below c. Alternatively, consider the specific control constraint, $u(t) \leq c$. If $u(t) \leq c$ binds as an equality, then $\frac{\partial H}{\partial u} > 0$. Hypothetically, H could be increased by increasing u above c.

The most general form of such a constraint may be written as:

$$g(t,y,u) \leq c$$

A constraint of this form may be formalized as a Kuhn-Tucker condition. The analysis precisely follows that of a static constrained optimization problem. Amend the Hamiltonian with a term of the form:

$$\theta(t)[c - g(t,y,u)]$$

We therefore construct a Lagrangian of the form:

$$\mathcal{L} = H(t,y,u,\lambda) + \theta(t)[c - g(t,y,u)]$$

First-order conditions include:

$$\frac{\partial \mathcal{L}}{\partial u} = 0,\ \frac{\partial \mathcal{L}}{\partial y} = -\dot{\lambda},\ \dot{y} = \frac{\partial \mathcal{L}}{\partial \lambda}$$

In addition to any other transversality conditions, we add the Kuhn-Tucker complementary slackness condition:

$$\frac{\partial \mathcal{L}}{\partial \theta} = c - g(t,y,u) \geq 0,\ \ \theta(t) \geq 0,\ \ \theta(t)[c - g(t,y,u)] = 0$$

Also note that

$$\frac{\partial \mathcal{L}}{\partial u} = \frac{\partial H}{\partial u} - \theta(t)\frac{\partial g(t,y,u)}{\partial u} = 0$$

Therefore, whenever $\theta(t) > 0$, $sgn\left(\frac{\partial H}{\partial u}\right) = sgn\left(\frac{\partial g(t,y,u)}{\partial u}\right)$.

Infinite Horizon Optimal Control

If we assume that T is given, and then take the limit as T approaches infinity, it is possible that, for some allowable control paths, the objective integral may no longer be finite. If we assume optimized objective functional converges, the results of the solution to the finite-horizon also govern the solution to the infinite-horizon problem. We must simply take limits as $T \to \infty$. We continue to require:

$$H(t,y^*,u^*,\lambda^*) \geq H(t,y^*,u,\lambda^*), t \in [0,\infty)$$

The revised condition for $\lim_{t\to\infty} y(t)$ free is given by:

$$\lim_{t\to+\infty} \lambda(t) = 0$$

For $\lim_{t\to\infty} y(t)$ free, $\lim_{t\to\infty} y(t) \geq y_{\min}$, the CS condition is:

$$\lim_{t\to+\infty} \lambda(t) \geq 0, \lim_{t\to+\infty} (y(t) - y_{\min}) \geq 0, \lim_{t\to+\infty} \lambda(t)(y(t) - y_{\min}) = 0$$

Implicitly, T is free, and so:

$$\lim_{t\to+\infty} H(t) = 0$$

also applies.

Current Value Hamiltonian

We often encounter problems in which the objective function F depends on t only through $e^{-\rho t}$. In such cases:

$$F(t,y,u) = G(y,u)e^{-\rho t}$$

First write the problem in the form:

$$\underset{u(t)}{Max} \int_0^T G(y,u)e^{-\rho t}dt$$

subject to: $\dot{y} = f(t,y,u)$

The Hamiltonian for this problem is given by:

$$H = G(y,u)e^{-\rho t} + \lambda(t)f(t,y,u)$$

Now define Current Value Hamiltonian:

$$\hat{H} \equiv He^{\rho t} = G(y,u) + \lambda(t)e^{\rho t}f(t,y,u)$$

Next define Current Value Costate Variable:

$$q(t) = \lambda(t)e^{\rho t}$$

Original problem requires:

$$\dot{\lambda}(t) = -\frac{\partial H}{\partial y}$$

From the definition of $\hat{H}$: $\frac{\partial \hat{H}}{\partial y} = \frac{\partial H}{\partial y}e^{\rho t} \Rightarrow \frac{\partial H}{\partial y} = \frac{\partial \hat{H}}{\partial y}e^{-\rho t}$

From the definition of q: $\dot{\lambda}(t) = \dot{q}(t)e^{-\rho t} - \rho q(t)e^{-\rho t}$. Therefore:

$$\dot{q}(t)e^{-\rho t} - \rho q(t)e^{-\rho t} = -\frac{\partial \hat{H}}{\partial y}e^{-\rho t}$$

Multiply by $e^{\rho t}$ to obtain:

$$\dot{q}(t) = -\frac{\partial \hat{H}}{\partial y} + \rho q(t)$$

Trivially, $\frac{\partial \hat{H}}{\partial u} = 0 \Leftrightarrow \frac{\partial H}{\partial u} = 0$.

Transversality Conditions

The rewritten condition for $y(T)$ free is given by:

$$q(T)e^{-\rho T} = 0$$

If we allow $y(T)$ to be free and impose $y(T) \geq y_{\min}$, the appropriate complementary slackness condition is given by:

$$q(T)e^{-\rho T} \geq 0, \ \ (y_T - y_{\min}) \geq 0, \ \ q(T)e^{-\rho T}(y(T) - y_{\min}) = 0$$

For the infinite-horizon case, we take limits as T approaches infinity. Therefore:

$$\lim_{t \to +\infty} q(t)e^{-\rho t} = 0$$

or

$$\lim_{t \to +\infty} q(t)e^{-\rho t} \geq 0, \ \lim_{t \to +\infty} (y(t) - y_{\min}) \geq 0, \ \ \lim_{t \to +\infty} q(t)e^{-\rho t}(y(t) - y_{\min}) = 0$$

We also require:

$$\lim_{t \to +\infty} \hat{H}e^{-\rho t} = 0$$