Continuous time optimal control - the basics and applications

The general continuous time constrained optimal control problem can be written in the following form

$$\max_{\{u_s|x_s\}_{s=t_0}^{t_1}} \left\{ F = \int_{t_0}^{t_1} f\left(t, x_t, u_t\right) dt \right\} \tag{1}$$

$$s.t.$$

$$\frac{dx_t}{dt} \equiv \dot{x}_t = g\left(t, x_t, u_t\right) \tag{2}$$

$$x\left(t_0\right) = x_0, x\left(t_1\right) \text{ is free.} \tag{3}$$

For simplicity in the reminder of the text assume that f and g are continuously differentiable functions of time and u is piecewise continuous function of time. The constraint (2) is the law of motion of the state variable x_t which is predetermined at the beginning of each "period" (e.g., capital). Meanwhile u_t is the control variable (e.g., the amount of consumption) which, given the value of x_t and its law of motion (2), we choose in order to maximize F (e.g., the lifetime utility of household). The solution of this problem is the optimal path of state and control variables, (x_t^*, u_t^*). This path should be feasible. In other words, it should satisfy the law of motion or the dynamic constraint (2) and the initial condition (3). Moreover, given the definition of control and state variables $u_t^* = u_t^*\left(x_t\right)$. The values of $x\left(t_1\right)$ and $u\left(t_1\right)$ satisfy maximization problem (i.e., these values are a choice).

Digression: *When we consider a household's intertemporal problem we - usually - have*

$$f\left(t, x_t, u_t\right) = e^{-\rho t} \tilde{f}\left(u_t\left(x_t\right)\right),$$

where $e^{-\rho t}$ is the discounting function and ρ is the discount rate, $\tilde{f}$ is the instantaneous (one period) utility from consumption $c_t\left(\equiv u_t\right)$ - erroneously we use the letter u for $\tilde{f}$. The optimal consumption c_t^*, in turn, is function of capital $k_t\left(\equiv x_t\right)$. Meanwhile, the constraint (2) represents the accumulation rule of assets/capital - in such case we basically solve the optimal consumption and saving paths, where the latter determines the optimal path of capital.

The rigorous approach to solving the problem is through "Lagrangian." Let q_t be the Lagrange multiplier of the constraint (2). The optimal problem written in terms of Lagrangian is the following.

$$\max_{\{u_s|x_s\}_{s=t_0}^{t_1}} \left\{ L = \int_{t_0}^{t_1} \left\{ f\left(t, x_t, u_t\right) + q_t\left[g\left(t, x_t, u_t\right) - \dot{x}_t\right]\right\} dt \right\}.$$

Integrate the last term by parts

$$-\int_{t_0}^{t_1} q_t \dot{x}_t dt = -\int_{t_0}^{t_1} q_t dx_t = -q_{t_1} x_{t_1} + q_{t_0} x_{t_0} + \int_{t_0}^{t_1} x_t \dot{q}_t dt$$

and rewrite the L

$$\max_{\{u_s|x_s\}_{s=t_0}^{t_1}}\left\{L=\int_{t_0}^{t_1}\left[f\left(t,x_t,u_t\right)+q_t g\left(t,x_t,u_t\right)+\dot{q}_t x_t\right]dt-q\left(t_1\right)x\left(t_1\right)+q\left(t_0\right)x\left(t_0\right)\right\}.$$

Necessary conditions

Let u_t^* be the optimal control function. Construct a family of "comparison" controls $u_t^*+\alpha h_t$, where h_t is some function and α is a real number. Denote $y\left(t,\alpha\right)$ the path of the state variable generated by the control $u_t^*+\alpha h_t$. Assume that $y\left(t,\alpha\right)$ is differentiable in arguments and $y\left(t_0,\alpha\right)=x\left(t_0\right)$ for any α [i.e., the optimal path x_t^* and $y\left(t,\alpha\right)$ start from the same point]. Notice that $y\left(t,0\right)=x_t^*$.

With this comparison controls the value of the Lagrangian L is

$$\begin{aligned}L\left(\alpha\right) &= \int_{t_0}^{t_1}\left[f\left(t,y\left(t,\alpha\right),u_t^*+\alpha h_t\right)+q_t g\left(t,y\left(t,\alpha\right),u_t^*+\alpha h_t\right)+\dot{q}_t y\left(t,\alpha\right)\right]dt\\ &\quad -q\left(t_1\right)y\left(t_1,\alpha\right)+q\left(t_0\right)x_0.\end{aligned}$$

Further, for simplicity let $L\left(\alpha\right)$ have one and interior maximum and let it be differentiable. Consider the following first order condition with slight abuse of previous notation

$$\begin{aligned}0 &= \left.\frac{dL\left(\alpha\right)}{d\alpha}\right|_{\alpha=0}\equiv L_\alpha'\left(0\right) \qquad (4)\\ &= \int_{t_0}^{t_1}\left(f_x' y_\alpha'+q_t g_x' y_\alpha'+\dot{q}_t y_\alpha'+f_u' h_t+q_t g_u' h_t\right)dt\\ &\quad -q\left(t_1\right)y_\alpha'\left(t_1,0\right).\end{aligned}$$

Apparently, the exact value of the RHS of this expression depends on q_t. It depends also on h_t and the way h_t influences the path of the state variable $y\left(t,\alpha\right)$. Meanwhile, the condition (4) should hold for any h_t (thus any y_α'). Therefore, one should select q_t (and $\dot{q}_t$) so that it eliminates the influence of h_t - note that we are basically deriving the envelope condition which states that the gradient of the maximand at the optimal point is orthogonal Select

$$\dot{q}_t = -\left[f_x'\left(t,x^*,u^*\right)+q_t g_x'\left(t,x^*,u^*\right)\right] \qquad (5)$$

$$q\left(t_1\right) = 0. \qquad (6)$$

Under such choice,

$$\begin{aligned}0 &= \tilde{L}'\left(0\right) \qquad (7)\\ &= \int_{t_0}^{t_1}\left[f_u'\left(t,x^*,u^*\right)+q_t g_u'\left(t,x^*,u^*\right)\right]h_t dt,\end{aligned}$$

which should hold for any h_t. Therefore, it should hold also for $h_t = f'_u(t, x^*, u^*) + q_t g'_u(t, x^*, u^*)$, which means that

$$\int_{t_0}^{t_1} \left[f'_u(t, x^*, u^*) + q_t g'_u(t, x^*, u^*)\right]^2 h_t dt = 0. \tag{8}$$

This in turn implies that

$$f'_u(t, x^*, u^*) + q_t g'_u(t, x^*, u^*) = 0. \tag{9}$$

The equations (5), (6), and (9) are the necessary conditions for optimality. Together with (2) and (3) they determine the optimal path of control and state variables (x_t^*, u_t^*).

A simple way for deriving the necessary conditions

Form a Hamiltonian:

$$H(t, x_t, u_t, q_t) \equiv f(t, x_t, u_t) + q_t g(t, x_t, u_t),$$

where q_t is the costate variable and is part of the solution to the optimal problem. The necessary conditions are obtained as:

$$\frac{\partial H}{\partial u} = 0, \tag{10}$$

$$-\frac{\partial H}{\partial x} = \dot{q}, \tag{11}$$

$$\frac{\partial H}{\partial q} = \dot{x}. \tag{12}$$

Notice that (10) is the same as (9), (11) is the same as (5), and (12) is (2). In addition, one gets an obvious condition $x(t_0) = x_0$ and $q(t_1) = 0$. The latter plays the role of transversality condition (TVC) in terms of finite time problem.

Digression: *The TVC requires that in a dynamically optimal path the choices are made in a way that ensures that at the end of the time horizon the state variable (e.g., capital) has no value and therefore the constraint is not binding. In economic terms, one wants the value of capital in terms of utility to be zero at the planning horizon. If its value is positive then at the end of the time the choice leaves a positive value of capital that gives no utility, which is against the optimality.*

In economic terms, the costate variable measures the shadow value of the associated state variable. Hence, it captures the gains (value) in the optimal control problem that stem from marginally increasing the state variable.

Sufficient conditions

In order the necessary conditions to be also sufficient we need further conditions.

- the functions f and g are concave in both arguments
- the optimal trajectories of x, u, and q satisfy the necessary conditions

- x_t and q_t are continuous functions with $q_t \geq 0$ for all t and if g is nonlinear in x or u, or both.

In order to prove the sufficiency define $f^* \equiv f(t, x^*, u^*)$ and $g^* \equiv g(t, x^*, u^*)$ and

$$D \equiv \int_{t_0}^{t_1} (f^* - f)\, dt.$$

Given that we are solving for a maximum we need to show that

$$D \geq 0.$$

Since f is concave

$$f^* - f \geq f_x^{*\prime}(x^* - x) + f_u^{*\prime}(u^* - u).$$

Therefore,

$$\begin{aligned} D &\geq \int_{t_0}^{t_1} \left[f_x^{*\prime}(x^* - x) + f_u^{*\prime}(u^* - u)\right] dt \qquad (13) \\ &= \int_{t_0}^{t_1} \left[(x^* - x)(-q g_x^* - \dot{q}) + (u^* - u)(-q g_u^*)\right] dt. \end{aligned}$$

Notice that

$$\begin{aligned} \int_{t_0}^{t_1} -\dot{q}(x^* - x)\, dt &= -\int_{t_0}^{t_1} (x^* - x)\, dq = -(x^* - x)\, q|_{t_0}^{t_1} + \int_{t_0}^{t_1} (g^* - g)\, q dt \\ &= \int_{t_0}^{t_1} (g^* - g)\, q dt \end{aligned}$$

since $x^*(t_0) = x(t_0)$ and $q(t_1) = 0$. Therefore, (13) can be written as

$$D \geq \int_{t_0}^{t_1} \left[(g^* - g) - g_x^*(x^* - x) - g_u^*(u^* - u)\right] q dt \geq 0.$$

The latter integral is greater or equal to zero since $q \geq 0$ and g is a concave function of x and u. This shows that the necessary conditions together with concavity of f and g and non-negativity of q are also sufficient conditions.

Infinite horizon discounted problem

A usual economic problem is written as

$$\max_{\{u_s|x_s\}_{s=0}^{\infty}} \left\{ U = \int_0^{\infty} \tilde{f}(x_t, u_t)\, e^{-\rho t} dt \right\} \tag{14}$$

$$s.t.$$

$$\dot{x}_t = g(t, x_t, u_t) \tag{15}$$

$$x(0) = x_0 \tag{16}$$

Notice that while $\tilde{f}$ - the instantaneous utility - is at time t the costate involves the value of changing the state from x_t incrementally over time, i.e., to $t + dt$. Thus the costate (and the Hamiltonian) has to take this into account. The current value Hamiltonian (discount factor = 1) is

$$\begin{aligned} H^C &= e^{\rho t} H = \tilde{f}(x_t, u_t) + \tilde{q}_t g(t, x_t, u_t), \\ \tilde{q}_t &= q_t e^{\rho t}. \end{aligned}$$

While the present value Hamiltonian (discount factor $= e^{-\rho t}$) is then

$$H^P = \tilde{f}(x_t, u_t)\, e^{-\rho t} + q_t g(t, x_t, u_t).$$

The necessary conditions for optimality are

$$\begin{aligned} & H_u^C = \tilde{f}_u(x_t, u_t) + \tilde{q}_t g_u(t, x_t, u_t) = 0, \\ & \frac{d}{dt}\tilde{q}_t = \rho \tilde{q}_t - H_x^C = \rho \tilde{q}_t - \left[\tilde{f}_x(x_t, u_t) + \tilde{q}_t g_x(t, x_t, u_t) \right], \\ & \lim_{t \to \infty} e^{-\rho t} q_t x_t = 0, \end{aligned}$$

where the last condition is the TVC for infinite horizon optimal problem. It states that the value of state variable in terms of utility should be zero in the limit, $t = \infty$.

Many states and controls

There could be many state and control variables - the numbers do not need to coincide. For more than one state simply one adds extra costate variables (multiplying the RHS of the dynamic constraints) to the Hamiltonian. For more than one control, one needs to derive one optimal condition for each control variable.

Continuous time Bellman equation (Hamilton-Jacobi-Bellman equation)

This section is for those who are familiar with recursive dynamic programming in discrete time. It illustrates the analogy between continuous time necessary conditions and the conditions derived for discrete time. Here I consider only the discounted problem, though all the logic can be applied for the more general case.

With a slight abuse of notation define the maximized value of the objective function as a function of the initial state x_t and initial time t [it's sufficient since $u_t = u(x_t)$].

$$V(t, x_t) = \max_{\{u_s : \dot{x}_s = g(x_s, u_s) | x_s\}_{s=t}^{\infty}} \left\{ \int_t^{\infty} \tilde{f}(x_s, u_s) e^{-\rho(s-t)} ds \right\}.$$

This can be rewritten in recursive form in the following way:

$$V(t, x_t) = \max_{\{u_s : \dot{x}_s = g(x_s, u_s) | x_s\}_{s=t}^{\infty}} \left\{ \int_t^{t+\Delta t} \tilde{f}(x_s, u_s) e^{-\rho(s-t)} ds + e^{-\rho \Delta t} V(t + \Delta t, x_{t+\Delta t}) \right\},$$

for any Δt.

Subtract from both sides $V(t, x_t)$ and divide by Δt.

$$0 = \max_{\{u_s : \dot{x}_s = g(x_s, u_s) | x_s\}_{s=t}^{\infty}} \left\{ \frac{1}{\Delta t} \int_t^{t+\Delta t} \tilde{f}(x_s, u_s) e^{-\rho(s-t)} ds + \frac{e^{-\rho \Delta t} V(t + \Delta t, x_{t+\Delta t}) - V(t, x_t)}{\Delta t} \right\}.$$

Take the limit $\Delta t \to 0$ (i.e., continuous time). By L'Hopital's rule

$$\lim_{\Delta t \to 0} \frac{1}{\Delta t} \int_t^{t+\Delta t} \tilde{f}(x_s, u_s) e^{-\rho(s-t)} ds = \tilde{f}(x_t, u_t).$$

Meanwhile, apply the definition of differential in order to get

$$\lim_{\Delta t \to 0} \frac{e^{-\rho \Delta t} V(t + \Delta t, x_{t+\Delta t}) - V(t, x_t)}{\Delta t} =$$

$$\lim_{\Delta t \to 0} \left[\frac{\left(e^{-\rho \Delta t} - 1\right) V(t + \Delta t, x_{t+\Delta t})}{\Delta t} + \frac{V(t + \Delta t, x_{t+\Delta t}) - V(t, x_{t+\Delta t})}{\Delta t} + \frac{V(t, x_{t+\Delta t}) - V(t, x_t)}{\Delta t} \right]$$

$$-\rho V(t, x_t) + \dot{V}(t, x_t) + V'_x(t, x_t) \dot{x}_t.$$

In sum this means that

$$\rho V(t, x_t) = \max_{u_t | x_t} \left\{ \tilde{f}(x_t, u_t) + V'_x(t, x_t) g(x_t, u_t) + \dot{V}(t, x_t) \right\}, \tag{17}$$

which is the Hamilton-Jacobi-Bellman equation. The second term in RHS captures the value gains from marginal change in the state variable, while the third term stands for the gains over time. The maximization gives the FOC:

$$\tilde{f}_u + V'_x g_u = 0,$$

which is the necessary condition for optimality, $H_u^C = 0$, where $V_x' = \tilde{q}_t$. This shows how the costate captures the effect of the change of the state on the objective function in current value terms. It also shows that $\tilde{q}_t$ depends on dynamic decisions.

The envelope condition is

$$\rho V_x' = \tilde{f}_x + V_x' g_x + V_{xx}'' g + \dot{V}_x'.$$

This is the necessary condition which describes the dynamics of the costate variable $\frac{d}{dt}\tilde{q}_t$ given that $\dot{x}_t = g\left(x_t, u_t^*\right)$.