A Dynamic Optimization Primer

John C. Bluedorn
for 1st Year M.Phil. students,
University of Oxford 2003-2004

I see a certain order in the universe and math is one way of making it visible.
– May Sarton (1912-1995)

Dynamic optimization is about making optimal plans through time or space. In what follows, I will present some results which I found particularly helpful in understanding dynamic optimization theory and the relationships amongst various solution methods. Essentially, what follows is a brief distillation of basic dynamic optimization theory, as applied to economics.[1] My own sources (from where I learned dynamic optimization methods) include Kamien & Schwartz (1991), Romer (2001), and Obstfeld & Rogoff (1996), and I recommend any of them for an introduction to these methods.

The three solution methods for dynamic problems I will look at include: the calculus of variations, optimal control theory, and dynamic programming.

In solving a dynamic optimization problem, there are four main questions which must be answered. These are: (1) does an optimal solution exist?; (2) is the optimal solution unique?; (3) what are the properties or characteristics which the optimal solution must have?; and (4) can we recover the optimal solution from these properties? Just as Fox Mulder on the X-Files assumes that "the truth is out there," I will assume that "the optimal solution is out there and is unique." The existence of an optimal solution to a dynamic problem is hard to show, and so I set it aside to concentrate on the necessary conditions which the optimal solution must fulfill (see Kamien & Schwartz (1991) for more discussion on this issue). Furthermore, I assume that the second order conditions for a maximum are satisfied.[2]

I draw heavily on the wonderful work of Kamien & Schwartz (1991), working through some of their derivations step-by-step in this primer.[3] A couple of quick notes on notation: partial and cross-partial derivatives of functions are denoted by subscripts of the argument(s) against which the derivative is taken; if the values at which a function is evaluated are clearly understood (e.g., evaluated at the optimal solution), then I may drop them in the notation, to avoid cluttering the page.

[1] Dynamic optimization theory is useful in solving many problems. In economics, most of these problems involve making optimal plans through time. Another possible use of dynamic optimization theory though is for making optimal plans through space. In fact, as related in Kamien & Schwartz (1991), the original motivation for creating dynamic optimization theory was to enable solution of the famous brachistochrone problem of John Bernoulli in 1696. The problem involves finding the optimal path (shortest time) through space for an object in a uniform gravitational field moving between two points.

[2] Under the calculus of variations for a maximum to be found, it is sufficient for the objective function $F(\cdot)$ to be concave in $x(t)$ and $x'(t)$. Note how this is analogous to the calculation of a maximum for a stationary function. A similar condition holds in optimal control theory.

[3] I really like Kamien & Schwartz (1991)! It presents a more comprehensive introduction to dynamic optimization methods than many introductory texts.

1 Calculus of Variations

Assume that we have the following optimization problem:

$$\max_{x(t)} \int_{t_0}^{t_1} F\left(t, x\left(t\right), x'\left(t\right)\right) dt$$
$$\text{s.t.} \quad x\left(t_0\right) = x_0 \text{ and } x\left(t_1\right) = x_1.$$

The initial and terminal condition on the function $x(t)$ are necessarily to pin down the solution we are looking for from the general class of functions which will fulfill the transition equation (or Euler equation) which we will find. Notice how our objective function $F(\cdot)$ may depend on time directly (through t), and indirectly, through our choice of $x(t)$ and implicitly of $x'(t)$, the time derivative of $x(t)$. Note though that these are all treated as distinct arguments.

The class of functions $x(t)$ which are admissible as possible optimal solutions are continuously differentiable functions defined over $[t_0, t_1]$ satifying the initial and terminal conditions.

To find the Euler equation, which will describe how the optimal solution evolves over time, we are going to use some really neat tricks. Suppose that the optimal solution to the maximization problem above exists and is denoted $x^*(t)$. Let $x(t)$ denote some other admissible function. Further, let the difference between this admissible function and the optimal solution be denoted:

$$h(t) = x(t) - x^*(t).$$

Given that both $x(t)$ and $x^*(t)$ are admissible, then we know that $h(t_0) = 0$ and $h(t_1) = 0$. Notice that we can rearrange our definition of $h(t)$ to get an alternative way to express $x(t)$. Specifically, we know that $x(t) = x^*(t) + h(t)$, which is admissible (by assumption).

Given that $x(t)$ is admissible, we know then that the function:

$$y(t) = x^*(t) + ah(t), \text{ for some constant } a,$$

is also admissible. The function $y(t)$ satisfies all of our criteria for admissibility listed above. Now, define the objective function evaluated at $y(t)$ as:

$$\begin{aligned} g(a) &= \int_{t_0}^{t_1} F\left(t, y(t), y'(t)\right) dt \\ &= \int_{t_0}^{t_1} F\left(t, x^*(t) + ah(t), x'^*(t) + ah'(t)\right) dt. \end{aligned}$$

By the optimality of $x^*(t)$, we know that $g(a)$ achieves its maximum at $a = 0$. This then implies that the first derivative of g with respect to its argument a must be zero at $a = 0$. Hence, we know that:

$$g'(0) = 0.$$

To find the derivative of g, we need to apply several differentiation rules from calculus, including the chain rule and Leibniz's Rule for differentiating an integral. We know that:

$$\begin{aligned} g'(a) &= F(t_1, y(t_1), y'(t_1)) \frac{dt_1}{da} - F(t_0, y(t_0), y'(t_0)) \frac{dt_0}{da} \\ &\quad + \int_{t_0}^{t_1} \frac{\partial F(t, x^*(t) + ah(t), x'^*(t) + ah'(t))}{\partial a} dt \\ &= \int_{t_0}^{t_1} \frac{\partial F(t, x^*(t) + ah(t), x'^*(t) + ah'(t))}{\partial a}, \end{aligned}$$

since we know that the endpoints (t_1 and t_0) do not depend on the parameter a (which means that $\frac{dt_1}{da} = \frac{dt_0}{da} = 0$). To get the derivative of F with respect to a, we apply the chain rule and find that:

$$\begin{aligned} g'(a) &= \int_{t_0}^{t_1} \left[F_t(t, y(t), y'(t)) \frac{dt}{da} + F_x(t, y(t), y'(t)) \frac{dy(t)}{da} + F_{x'}(t, y(t), y'(t)) \frac{dy'(t)}{da} \right] dt \\ &= \int_{t_0}^{t_1} \left[F_x(t, y(t), y'(t)) h(t) + F_{x'}(t, y(t), y'(t)) h'(t) \right] dt. \end{aligned}$$

Note that I have used the subscripts for the partial derivatives from the original maximization problem. Evaluated at $a = 0$ and applying our first order condition for optimality, we know that:

$$g'(0) = \int_{t_0}^{t_1} F_x(t, x^*(t), x'^*(t)) h(t) dt + \int_{t_0}^{t_1} F_{x'}(t, x^*(t), x'^*(t)) h'(t) dt = 0.$$

This expression is referred to as the *first variation* (this is where the calculus of variations gets its name!).[4] The above condition is a bit hard to interpret, so let's simplify it a bit.

We can integrate the second integral by parts, and get something a bit friendlier.[5] Integration by parts implies that:

$$\begin{aligned} \int_{t_0}^{t_1} F_{x'}(t, x^*(t), x'^*(t)) h'(t) dt &= F_{x'} h|_{t_0}^{t_1} - \int_{t_0}^{t_1} h \frac{dF_{x'}}{dt} dt \\ &= -\int_{t_0}^{t_1} h \frac{dF_{x'}}{dt} dt, \text{ since } h(t_1) = h(t_0) = 0. \end{aligned}$$

[4]To see why it is the calculus of *variations* (plural), it helps to know that $g''(0)$ is referred to as the *second variation*. In fact, to be at a maximum, the second variation must be nonpositive for all admissible functions $h(t)$ (this is a necessary condition).

[5]Recall that integration by parts merely comes from the product rule:

$$\begin{aligned} d(uv) &= udv + vdu \\ \Rightarrow uv &= \int udv + \int vdu \\ \Rightarrow \int udv &= uv - \int vdu. \end{aligned}$$

The integration sign here indicates that you should take the antiderivative.

We can now see that:

$$\begin{aligned} g'(0) &= \int_{t_0}^{t_1} F_x h dt - \int_{t_0}^{t_1} h \frac{dF_{x'}}{dt} dt = 0 \\ &\Rightarrow \int_{t_0}^{t_1} \left[F_x - \frac{dF_{x'}}{dt} \right] h dt = 0. \end{aligned}$$

Note that this will always hold given that $x^*(t)$ is optimal and that $h(t)$ is derived from an admissible function. It therefore must also hold whenever the term multiplying $h(t)$ is zero. Hence, we can see that a necessary condition, which describes the Euler equation (or the optimal transition equation) is that:

$$F_x(t, x^*(t), x'^*(t)) = \frac{dF_{x'}(t, x^*(t), x'^*(t))}{dt}.$$

In fact, the above equation must *always* hold. The following lemma confirms this.

Lemma 1 *Suppose that $g(t)$ is a given, continuous function defined over $[t_0, t_1]$. If*

$$\int_{t_0}^{t_1} g(t) h(t) dt = 0$$

for every continuous function $h(t)$ defined on $[t_0, t_1]$ and satisfying our admissability conditions (it's zero at the endpoints), then $g(t) = 0 \forall t \in [t_0, t_1]$.

PROOF: Suppose that the above conclusion is not true. Hence, $g(t)$ is nonzero for some t over the interval $[t_0, t_1]$. Suppose that $g(t) > 0$, for some $t \in [t_0, t_1]$. Then, since $g(t)$ is continuous, $g(t)$ must be positive over some interval $[a, b]$ within $[t_0, t_1]$. Now, we will take a particular $h(t)$ which is admissible under the conditions given in Lemma ## (it is continuous, defined over the interval $[t_0, t_1]$, and is zero at the endpoints). Let:

$$h(t) = \begin{cases} (t-a)(b-t), \text{ for } t \in [a, b] \\ 0, \text{ everywhere else.} \end{cases}.$$

Hence, we know that:

$$\int_{t_0}^{t_1} g(t) h(t) dt = \int_a^b g(t)(t-a)(b-t) dt > 0, \text{ since the integrand is positive.}$$

However, we can see that this clearly contradicts the hypothesis that $\int_{t_0}^{t_1} g(t) h(t) dt = 0$ for every admissible function $h(t)$, which is one of the assumptions of the lemma. Hence, we know that our assertion that $g(t) > 0$ for some $t \in [t_0, t_1]$ is false under the assumptions of the lemma. Similarly, we can show that a contradiction arises in the case where the assertion is that $g(t) < 0$ for some $t \in [t_0, t_1]$. Hence, the lemma must be true. ■

Does the calculus of variations approach give us a means to actually solve for the optimal solution $x^*(t)$, and not just the Euler equation? It does. Note that:

$$\frac{dF_{x'}(t, x^*(t), x'^*(t))}{dt} = F_{x't} + F_{x'x}x' + F_{x'x'}x''.$$

Hence, the Euler equation can be rewritten as:

$$F_x = F_{x't} + F_{x'x}x' + F_{x'x'}x'',$$

where all terms are evaluated at the optimal solution. In this form, we can see that the extremal (optimal function) $x(t)$ can actually be solved from the second order differential equation defined by the Euler equation (with the terminal conditions)

There is actually a clever mneumonic device to recall all the necessary conditions for an optimal solution for a calculus of variations problem. In fact, we will see that the mneumonic device hints at a link between optimal control theory to the calculus of variations (beyond the fact that they address similar problems). Define the function $\lambda(t) = F_{x'}(t, x, x')$. Given that $F_{x'x'} \neq 0$, then we know that x' can be expressed as a function of t, x, and λ by the implicit function theorem. We can do a nice little trick. Define the function:

$$H(t, x, \lambda) = -F(t, x, x') + \lambda x'.$$

For a typical utility maximization problem in economics, λ will be the shadow price of wealth or the marginal utility of wealth (the value in utility terms of an extra unit of wealth). The total differential of $H(\cdot)$ is then:

$$\begin{aligned} dH &= -F_t dt - F_x dx - F_{x'}dx' + \lambda' x' + \lambda dx' \\ &= -F_t dt - F_x dx - F_{x'}dx' + \frac{dF_{x'}}{dt}x' + F_{x'}dx' \\ &= -F_t dt - F_x dx + \lambda' x' \\ &= \frac{\partial H}{\partial t}dt + \frac{\partial H}{\partial x}dx + \frac{\partial H}{\partial \lambda}d\lambda, \end{aligned}$$

using the definition of λ and of total differentiation. Then, we can see that:

$$\frac{\partial H}{\partial x} = -F_x, \ \frac{\partial H}{\partial \lambda} = x', \text{ and } d\lambda = \lambda'.$$

Given that the function $x^*(t)$ satisfies the Euler equation we found earlier, we know that the following must be true:

$$\begin{gathered} \lambda' = F_x = -\frac{\partial H}{\partial x} \\ \Rightarrow \frac{\partial H}{\partial x} = -\lambda' \\ \text{and } \frac{\partial H}{\partial \lambda} = x'. \end{gathered}$$

These two first order differential equations are another way to state the Euler equation (because we remember that $\lambda = F_{x'}$ by definition).

EXAMPLE: We can solve a standard consumption/investment problem in continuous time now. Let the instantaneous utility function be denoted $u(c(t))$, which is concave by assumption. Further, suppose that future consumption is discounted exponentially at the constant rate β. Let saving and borrowing occur through the holding of a single asset, capital, which yields a riskless rate of return r each instant (holdings can be positive or negative; negative holdings imply that an individual must pay a rate r each instant). There is an exogenous income stream, given by $y(t)$. Then, the lifetime utility maximization problem for the individual is given by:

$$\begin{gathered}\max_{c(t)} \int_0^T e^{-\beta t} u(c(t))\, dt \\ \text{s.t.} \quad c(t) + k'(t) = y(t) + rk(t), \\ k(0) = k_0 \text{ and } k(T) = k_T.\end{gathered}$$

The budget constraint here says that consumption equals the difference between total income $(y(t) + rk(t))$ and any change in capital holdings $(k'(t))$. Initially, it doesn't appear that the above problem is amenable to solution by the calculus of variations. However, a little inspection reveals that we can merely substitute the budget constraint into the objective function, and recover a classic calculus of variations problem. It is:

$$\max_{k(t),k'(t)} \int_0^T e^{-\beta t} u(y(t) + rk(t) - k'(t))\, dt.$$

Recall that the Euler equation is then:

$$\begin{aligned} F_x &= \frac{dF_{x'}}{dt} \Rightarrow \\ e^{-\beta t} u'(y(t) + rk(t) - k'(t))\, r &= \frac{d\left[e^{-\beta t} u'(y(t) + rk(t) - k'(t))(-1)\right]}{dt} \\ e^{-\beta t} u'(c(t))\, r &= -\beta e^{-\beta t} u'(c(t))(-1) + e^{-\beta t} u''(c(t))\, c'(t)(-1) \\ r &= \beta - \frac{u''(c(t))}{u'(c(t))} c'(t) \\ -\frac{u''(c(t))\, c(t)\, c'(t)}{u'(c(t))\, c(t)} &= r - \beta \\ \frac{c'(t)}{c(t)} &= \frac{1}{\sigma(t)}(r - \beta). \end{aligned}$$

The last line is relatively easy to interpret. Essentially, the growth rate of consumption along the optimal path should equal a proportion of the difference between the rate of return on capital r and the rate of time preference (or time discount) β. The exact proportion is given by the reciprocal of the coefficient of relative risk aversion $\sigma(t) = -\frac{u''(c(t))c(t)}{u'(c(t))}$. If the instantaneous utility function is $u(c(t)) = \ln(c(t))$, then we know that $\sigma(t) = 1$ for all t. In the case of log utility then, we can see that:

$$\frac{c'(t)}{c(t)} = (r - \beta).$$

The consumption function can readily be solved from this first order linear differential equation, given the initial and terminal conditions on capital asset holdings. According to the

Euler equation, consumption will grow over the individual's lifetime if they are more patient than the market ($\beta < r$), while consumption will fall over the individual's lifetime if they are more impatient than the market ($\beta > r$). Consumption is flat if the rate of time preference and the rate of return on capital are identical.

2 Optimal Control Theory

Pontryagin's maximum principle (developed in the 1950s) of optimal control applies to any problems addressable by the calculus of variations. However, optimal control is somewhat more flexible than the calculus of variations. For example, constraints on the derivatives of functions can be more readily applied.

In an optimal control problem, we have two basic kinds of variables: *state* variables, which follow some first order differential equation, and *control* variables, for which values can be optimally chosen in a piecewise, continuous fashion. An example of a state variable is wealth, or really any asset subject to accumulation (something that follows a differential equation). An example of a control variable is consumption. Consider the following problem:

$$\max \int_{t_0}^{t_1} f\left(t, x\left(t\right), u\left(t\right)\right) dt$$
$$\text{s.t. } \; x'\left(t\right) = g\left(t, x\left(t\right), u\left(t\right)\right)$$
$$x\left(t_0\right) = x_0, \text{ while } x\left(t_1\right) \text{ is free.}$$

Here, the objective function f and the constraint function g are assumed to be continuously differentiable functions (members of C^1 at least). Notice that if we let $g\left(t, x\left(t\right), u\left(t\right)\right) = u\left(t\right)$, we recover the calculus of variations problem presented in section 1.

As before, we will assume that an optimum exists. Note that a choice of an optimal control variable $u^*\left(t\right)$, defines the optimal state variable $x^*\left(t\right)$ since $x\left(t\right)$ is specified at the endpoints and its evolution is given by a differential equation which depends on $u\left(t\right)$. Hence, we can focus on finding the optimal control alone. Consider the control function given by $u\left(t\right) = u^*\left(t\right) + ah\left(t\right)$, where $u^*\left(t\right)$ is the optimal control, $h\left(t\right)$ is some fixed function, and a is a parameter. Let $y\left(t, a\right)$ denote the state variable implied by the differential equation $x' = g$ and $u\left(t\right)$. We can see that $y\left(t, 0\right) = x^*\left(t\right)$. Further, we will constrain y so that the initial condition is satisfied for any a. Thus, $y\left(t_0, a\right) = x_0$.

The objective function for the optimal control problem can be rewritten as:

$$\int_{t_0}^{t_1} f\left(t, x\left(t\right), u\left(t\right)\right) dt = \int_{t_0}^{t_1} \left[f\left(t, x\left(t\right), u\left(t\right)\right) + \lambda\left(t\right) g\left(t, x\left(t\right), u\left(t\right)\right) - \lambda\left(t\right) x'\left(t\right)\right] dt,$$

where the sum of the coefficients of $\lambda\left(t\right)$ (which are the function g and x') must be zero since the constraint is assumed to hold. No restrictions have yet been placed on the multiplier function, $\lambda\left(t\right)$. It could be any old C^1 function. Now, integrate the last term by parts, to get:

$$-\int_{t_0}^{t_1} \lambda\left(t\right) x'\left(t\right) dt = -\lambda\left(t_1\right) x\left(t_1\right) + \lambda\left(t_0\right) x\left(t_0\right) + \int_{t_0}^{t_1} x\left(t\right) \lambda'\left(t\right) dt.$$

By substituting this expression into the earlier equation, we have:

$$\int_{t_0}^{t_1} f(t, x(t), u(t))\, dt =$$
$$\int_{t_0}^{t_1} \left[f(t, x(t), u(t)) + \lambda(t) g(t, x(t), u(t)) + \lambda'(t) x(t)\right] dt - \lambda(t_1) x(t_1) + \lambda(t_0) x(t_0).$$

Now, define the objective function evaluated at the proposed control function to be:

$$\begin{aligned} J(a) &= \int_{t_0}^{t_1} f(t, y(t,a), u^*(t) + ah(t))\, dt \\ &= \int_{t_0}^{t_1} \left[f(t, y(t,a), u^*(t) + ah(t)) + \lambda(t) g(t, y(t,a), u^*(t) + ah(t)) + \lambda'(t) y(t,a)\right] dt \\ &\quad -\lambda(t_1) y(t_1, a) + \lambda(t_0) y(t_0, a). \end{aligned}$$

We can see then that for $J(a)$ to achieve its maximum, $a = 0$. Further, from the first order conditions for an optimum, we know that $J'(0) = 0$ (it's the same condition for either a maximum or a minimum). So, let's differentiate our expression for $J(a)$ with respect to a and evaluate it at $a = 0$. Evaluating all terms at $a = 0$ and noting that $y(t_0, 0) = x_0$, a constant, we get:

$$\begin{aligned} J'(0) &= \int_{t_0}^{t_1} \left[f_x y_a + f_u h + \lambda g_x y_a + \lambda g_u h + \lambda' y_a\right] dt - \lambda(t_1) y_a(t_1, 0) + \lambda(t_0) y_a(t_0, 0) \\ &= \int_{t_0}^{t_1} \left[(f_x + \lambda g_x + \lambda') y_a + (f_u + \lambda g_u) h\right] dt - \lambda(t_1) y_a(t_1, 0). \end{aligned}$$

Now, to get any farther in understanding the solution, we need to restrict the functional form of $\lambda(t)$. Specifically, we will assume that $\lambda(t)$ behaves in such a way that we will not need to worry about how the implicit function $y(t,a)$ changes with a. Essentially, we're choosing a particular set of possible $\lambda(t)$ to make our life easier. Hence, let $\lambda(t)$ follow the differential equation:

$$\lambda(t) = -\left[f_x(t, x^*, u^*) + \lambda g_x(t, x^*, u^*)\right], \text{ where } \lambda(t_1) = 0.$$

This enables us to get rid of the first and last terms in $J'(0)$.

Accordingly, we can see that $J'(0) = 0$ requires that:

$$\int_{t_0}^{t_1} \left[f_u(t, x^*(t), u^*(t)) + \lambda(t) g_u(t, x^*(t), u^*(t))\right] h(t)\, dt = 0, \text{ for any arbitrary function } h(t).$$

Now, we're going to be clever and see what happens when we choose $h(t) = f_u(t, x^*(t), u^*(t)) + \lambda(t) g_u(t, x^*(t), u^*(t))$, since h can be any function. Then, we know that:

$$\int_{t_0}^{t_1} \left[f_u(t, x^*(t), u^*(t)) + \lambda(t) g_u(t, x^*(t), u^*(t))\right]^2 dt = 0.$$

However, this means that the optimal solution requires that:

$$f_u\left(t, x^*\left(t\right), u^*\left(t\right)\right) + \lambda\left(t\right) g_u\left(t, x^*\left(t\right), u^*\left(t\right)\right) = 0, \text{ for all } t \in \left[t_0, t_1\right].$$

The optimal solution to the problem thus must satisfy the following conditions:

$$\begin{gathered} x'\left(t\right) = g\left(t, x\left(t\right), u\left(t\right)\right) \text{ and } x\left(t_0\right) = x_0, \\ \lambda'\left(t\right) = -\left[f_x\left(t, x\left(t\right), u\left(t\right)\right) + \lambda g_x\left(t, x\left(t\right), u\left(t\right)\right)\right] \text{ and } \lambda\left(t_1\right) = 0, \\ f_u\left(t, x\left(t\right), u\left(t\right)\right) + \lambda\left(t\right) g_u\left(t, x\left(t\right), u\left(t\right)\right) = 0, \text{ for all } t \in \left[t_0, t_1\right]. \end{gathered}$$

The first equation is the state equation, the second equation is the costate or multiplier equation, and the last equation is the optimality condition.

How to remember all of these conditions? Just as we used the Hamiltonian as a mneumonic device to generate the conditions for an optimal solution to a calculus of variations problem, we can employ the Hamiltonian for optimal control problems (given that all calculus of variations problems can be solved via optimal control methods, this is not surprising). Define the Hamiltonian for the optimal control problem to be:

$$H\left(t, x\left(t\right), u\left(t\right), \lambda\left(t\right)\right) = f\left(t, x\left(t\right), u\left(t\right)\right) + \lambda\left(t\right) g\left(t, x\left(t\right), u\left(t\right)\right).$$

Notice that we can now recover all of the conditions for an optimal solution by taking the appropriate derivatives of the Hamiltonian. Specifically, we can see that:

$$\begin{aligned} \frac{\partial H}{\partial u} &= f_u + \lambda g_u = 0 \text{ gives us the optimality condition;} \\ -\frac{\partial H}{\partial x} &= -\left[f_x + \lambda g_x\right] = \lambda'\left(t\right) \text{ gives us the costate equation;} \\ \frac{\partial H}{\partial \lambda} &= g = x'\left(t\right) \text{ gives us the state equation.} \end{aligned}$$

Of course, we must also require that $x\left(t_0\right) = x_0$ and $\lambda\left(t_1\right) = 0$ to complete the solution. To ensure that we are at a maximum solution, we further require that $H_{uu}\left(t, x^*\left(t\right), u^*\left(t\right), \lambda\left(t\right)\right) \leq 0$ (for a minimum, it's that $H_{uu} \geq 0$).

Try out the utility maximization problem from the calculus of variations section using the methods of optimal control. You should recover the same results. Notice that since the objective function of the problem contains an exponential discount factor, the resulting Hamiltonian is sometimes called the *present value* Hamiltonian (because everything is discounted to the present effectively). The *current value* Hamiltonian is formed by multiplying the present value Hamiltonian by the reciprocal of the exponential discount factor (so everything is valued in current terms). Hence, $\widetilde{H} = e^{\beta t} H$, in the case of the utility maximization problem, where $\widetilde{H}$ denotes the current value Hamiltonian, while H denotes the standard or present value Hamiltonian. Then the new multiplier (or costate variable) is:

$$\begin{aligned} m\left(t\right) &= e^{\beta t} \lambda\left(t\right) \Rightarrow \\ m'\left(t\right) &= e^{\beta t}\left(\beta\right) \lambda\left(t\right) + e^{\beta t} \lambda'\left(t\right) \\ m'\left(t\right) &= \beta m\left(t\right) - e^{\beta t} \frac{\partial H}{\partial x}. \end{aligned}$$

Note that $\frac{\partial \widetilde{H}}{\partial x} = e^{\beta t} \frac{\partial H}{\partial x} \Rightarrow \frac{\partial H}{\partial x} = e^{-\beta t} \frac{\partial \widetilde{H}}{\partial x}$. Hence, we can see that:

$$-\frac{\partial \widetilde{H}}{\partial x} = m'(t) - \beta m(t) \text{ and that}$$
$$\frac{\partial H}{\partial u} = \frac{\partial \left(e^{-\beta t} curl H \right)}{\partial u} = e^{-\beta t} \frac{\partial \widetilde{H}}{\partial u} = 0 \Rightarrow \frac{\partial \widetilde{H}}{\partial u} = 0.$$

The above conditions characterize the optimum when we choose to use the current value Hamiltonian.

3 Dynamic Programming Approach

The dynamic programming approach to dynamic optimization problems employs *Bellman's principle of optimality* to solve these problems (whether discrete or continuous).[6] Essentially, the principle of optimality states that the optimal path for the control variable (the choice variable) in a problem will be the same whether we solve the problem over the entire time interval of interest, or solve for future periods as a function of the initial conditions given by past optimal solutions. We can break the problem up into two bits, solving only for the current optimal path, taking the fact that the future path will also be optimal (with an initial condition of today's choice) as a given.

Consider the following dynamic optimization problem:

$$\max \int_0^T f(t, x(t), u(t))\, dt$$
$$\text{s.t.} \quad x'(t) = g(t, x(t), u(t)),\ x(0) = a.$$

Define $J(t_0, x(t_0))$ to be the maximum value of this function starting at time t_0, subject to the constraint on the motion of $x(t)$. Hence, we have:

$$\begin{aligned} J(t_0, x(t_0)) &= \max_{u(t)} \int_{t_0}^T f(t, x(t), u(t))\, dt \\ \text{s.t.} \quad x'(t) &= g(t, x(t), u(t)),\ x(t_0) = x_0. \end{aligned}$$

Now, we can break up the above integral into two pieces. We then have:

$$J(t_0, x(t_0)) = \max_{u(t)} \int_{t_0}^{t_0+\Delta t} f(t, x(t), u(t))\, dt + \int_{t_0+\Delta t}^T f(t, x(t), u(t))\, dt,$$

where Δt is a small positive increment of time. Throughout, we're assuming that the maximization is done subject to the appropriate constraints. Here is where we'll apply the Bellman's principle of optimality, by arguing that the solution to the second integral will be optimal given our choices regarding the first integral. Notice that the two integrals are linked by the behavior of the state variable $x(t)$ (the second integral takes its starting state

[6] Dynamic programming methods are probably more commonly applied to discrete time problems, though they also work in continuous time.

variable value from the problem represented by the first integral). The problem can then be rewritten as:

$$J(t_0, x(t_0)) = \max_{u(t) \text{ for } t\in[t_0,t_0+\Delta t]} \left[\int_{t_0}^{t_0+\Delta t} f(t, x(t), u(t))\, dt \right.$$
$$\left. + \max_{u(t) \text{ for } t\in[t_0+\Delta t,T]} \int_{t_0+\Delta t}^{T} f(t, x(t), u(t))\, dt \right]$$
$$\text{s.t.} \quad x'(t) = g(t, x(t), u(t)),\ x(t_0) = x_0, \text{ and } x(t_0 + \Delta t) = x_0 + \Delta x.$$

Hence, we have:

$$J(t_0, x(t_0)) = \max_{u(t) \text{ for } t\in[t_0,t_0+\Delta t]} \left[\int_{t_0}^{t_0+\Delta t} f(t, x(t), u(t))\, dt + J(t_0 + \Delta t, x_0 + \Delta x)\right]$$
$$\text{s.t.} \quad x'(t) = g(t, x(t), u(t)),\ x(t_0) = x_0.$$

Now, let's approximate the value of the integral in $J(t_0, x_0)$. For small values of Δt, we'll have that:

$$\int_{t_0}^{t_0+\Delta t} f(t, x(t), u(t))\, dt \approx f(t, x(t), u(t))\, \Delta t.$$

Further, we'll assume that the value function J is at least twice continuously differentiable (it's a member of C^2). Then, we can take a Taylor approximation of J around (t_0, x_0). We can then see that:

$$J(t_0 + \Delta t, x_0 + \Delta x) = J(t_0, x_0) + J_t(t_0, x_0)(t_0 + \Delta t - t_0) + J_x(t_0, x_0)(x_0 + \Delta x - x_0) + h.o.t,$$

where $h.o.t.$ stands for the higher order terms in the Taylor approximation. Substituting these results into the value function, we have:

$$J(t_0, x_0) = \max_{u(t) \text{ for } t\in[t_0,t_0+\Delta t]} [f(t, x(t), u(t))\, \Delta t + J(t_0, x_0) + J_t(t_0, x_0)(t_0 + \Delta t - t_0)$$
$$+ J_x(t_0, x_0)(x_0 + \Delta x - x_0) + h.o.t]$$

By subtracting $J(t_0, x_0)$ from both sides (and suppressing t as an argument when it's clear), we have that:

$$0 = \max_u [f(t, x, u)\, \Delta t + J_t(t_0, x_0)\, \Delta t + J_x(t_0, x_0)\, \Delta x + h.o.t].$$

Now, divide through by Δt and let $\Delta t \to 0$, to get:

$$\begin{aligned}
\frac{0}{\Delta t} &= \max_u \left[f(t, x, u)\frac{\Delta t}{\Delta t} + J_t(t_0, x_0)\frac{\Delta t}{\Delta t} + J_x(t_0, x_0)\frac{\Delta x}{\Delta t} + h.o.t\right] \\
0 &= \max_u \left[f(t, x, u) + J_t(t_0, x_0) + J_x(t_0, x_0)\frac{\Delta x}{\Delta t} + h.o.t\right] \Rightarrow \\
0 &= \max_u \left[f(t, x, u) + J_t(t, x) + J_x(t, x)\, x' + h.o.t\right].
\end{aligned}$$

The zero subscripts are no longer necessary (since with $\Delta t \to 0$, we are at the initial condition!). Hence, we can see that:

$$\begin{aligned} 0 &= \max_u \left[f(t,x,u) + J_t(t,x) + J_x(t,x)\, g(t,x,u) + h.o.t \right] \\ -J_t(t,x) &= \max_u \left[f(t,x,u) + J_x(t,x)\, g(t,x,u) + h.o.t \right], \end{aligned}$$

since $J_t(t,x)$ at the initial condition doesn't depend on u. This is the fundamental partial differential equation which the value function J must obey. It is called the *Hamilton-Jacobi-Bellman equation*, sometimes shortened to just the *Bellman equation*. First, you solve the maximization problem on the right-hand side of the equation. Then, with that solution in hand, you solve the resulting partial differential equation for J.

Notice how the expression to on the right-hand side of the Bellman equation looks a lot like the standard Hamiltonian, except that $\lambda(t)$ is replaced by $J_x(t,x)$. In fact, this is quite apt. Recall that the multiplier function $\lambda(t)$ gives us the marginal valuation of the state variable, which is exactly $J_x(t,x)$ in this case. Hence, the right-hand side is the standard Hamiltonian problem from optimal control! The equation of motion for $\lambda(t)$ can also be derived from the Bellman equation. Suppose that the optimal control (chosen in terms of t, x, and J_x is given by u^*. Ignoring the higher order terms in the maximization problem, we know that:

$$f_u(t,x,u^*) + J_x(t,x)\, g_u(t,x,u^*) = 0$$

Since the Bellman equation must hold for any function x, we know that it must hold for small changes in the function x. Hence, we can see that:

$$\begin{gathered} -J_t(t,x) = f(t,x,u^*) + J_x(t,x)\, g(t,x,u^*) \Rightarrow \\ -J_{tx}(t,x) = f_x(t,x,u^*) + J_{xx}(t,x)\, g(t,x,u^*) + J_x(t,x)\, g_x(t,x,u^*) \\ -J_{tx}(t,x) - J_{xx}(t,x)\, g(t,x,u^*) = f_x(t,x,u^*) + J_x(t,x)\, g_x(t,x,u^*). \end{gathered}$$

Notice that the total derivative with respect to t of J_x is given by:

$$\begin{aligned} \frac{d(J_x(t,x))}{dt} &= J_{xt}(t,x) + J_{xx}(t,x)\, x'(t) \\ &= J_{xt}(t,x) + J_{xx}(t,x)\, g(t,x,u). \end{aligned}$$

Let $\lambda(t) = J_x(t,x(t))$. Then, we have that:

$$-\frac{d\lambda(t)}{dt} = f_x(t,x,u^*) + \lambda(t)\, g_x(t,x,u^*).$$

Thus, the optimal conditions from optimal control can be generated by the dynamic programming approach! It's all connected. You could try to solve the partial differential equation for the value function directly, but partial differential equations are notoriously difficult to solve. A typical method would be to postulate a particular solution for the value function, and then see if it satisfies the partial differential equation.

The dynamic programming approach also generates our familiar asset-value like equations of motion. Suppose that we have an infinite time horizon problem with an exponential discount factor, where the objective function and state equation do not depend on time directly:

$$\max_{u(t)} \int_0^\infty e^{-\beta t} f\left(x\left(t\right), u\left(t\right)\right) dt$$
$$\text{s.t.} \quad x'\left(t\right) = g\left(x\left(t\right), u\left(t\right)\right),\ x\left(0\right) = \bar{x}.$$

This implies that:

$$\begin{aligned} J\left(t_0, x_0\right) &= \max_{u(t)} \int_{t_0}^\infty e^{-\beta t} f\left(x\left(t\right), u\left(t\right)\right) dt \\ &= \max_{u(t)} \int_{t_0}^\infty e^{-\beta t_0} e^{-\beta(t-t_0)} f\left(x\left(t\right), u\left(t\right)\right) dt \\ &= e^{-\beta t_0} \max_{u(t)} \int_{t_0}^\infty e^{-\beta(t-t_0)} f\left(x\left(t\right), u\left(t\right)\right) dt. \end{aligned}$$

Notice how the maximization problem now depends only upon elapsed time $(t - t_0)$ and the initial state (x_0) and not upon the date (t_0). In this case, the maximization problem piece only depends upon the initial state, and not upon the initial date.[7] Hence, we can write:

$$V\left(x\left(t_0\right)\right) = V\left(x_0\right) = \max_{u(t)} \int_{t_0}^\infty e^{-\beta(t-t_0)} f\left(x\left(t\right), u\left(t\right)\right) dt.$$

We can then see that the following must be true:

$$\begin{aligned} J\left(t, x\right) &= e^{-\beta t} V\left(x\right) \\ J_t\left(t, x\right) &= -\beta e^{-\beta t} V\left(x\right) \\ J_x\left(t, x\right) &= e^{-\beta t} V'\left(x\right). \end{aligned}$$

This implies that the Hamilton-Jacobi-Bellman equation for this problem can be written as:

$$\begin{aligned} \beta e^{-\beta t} V\left(x\right) &= \max_u \left[e^{-\beta t} f\left(x, u\right) + e^{-\beta t} V'\left(x\right) g\left(x, u\right)\right] \\ \beta V\left(x\right) &= \max_u \left[f\left(x, u\right) + V'\left(x\right) g\left(x, u\right)\right] \end{aligned}$$

Suppose that V represents the value of holding some asset. Hence, the path for the state variable x must be chosen so that the instantaneous return on holding the asset (βV) is equal to the maximum instantaneous dividend paid by the asset (f) plus the maximum instantaneous capital gain of the asset we can achieve as we change the state variable $(V'g)$.

4 Final Notes

Although all of the above examples and derivations employ only a single state variable and a single control variable, the optimization method is easily generalizable to as many state and control variables as required by a problem. The necessary conditions are then defined in matrix form, employing the Jacobian. Note that the optimal control approach requires the introduction of as many co-state variables as there are state variables.

[7] The infinite horizon helps us here. The problem for $V(\cdot)$ looks the same regardless of the particular starting time at which it is evaluated.

References

Kamien, M. I. & Schwartz, N. L. (1991), *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*, Vol. 31 of *Advanced Textbooks in Economics*, second edn, Elsevier B.V., Amsterdam.

Obstfeld, M. & Rogoff, K. (1996), *Foundations of International Macroeconomics*, first edn, MIT Press, Cambridge, Massachusetts.

Romer, D. (2001), *Advanced Macroeconomics*, second edn, McGraw-Hill, New York, New York.