

Optimal Control

Pascal Michailat

1. Preliminary Results	1
1.1. L'Hopital's rule	1
1.2. Leibniz's rule	1
2. Consumption-Saving Problem	2
2.1. Description of the Problem	2
2.2. Optimal-Control Approach	2
3. Solution with Present-Value Hamiltonian	3
3.1. State and Control Variables	3
3.2. Present-Value Hamiltonian	3
3.3. Optimality Conditions	3
3.4. Transversality Condition	4
3.5. Euler Equation	4
3.6. CRRA Utility	5
4. Solution with Current-Value Hamiltonian	6
4.1. Current-Value Hamiltonian	6
4.2. Optimality Conditions	6
4.3. Transversality Condition	7
4.4. Euler Equation	7
4.5. Comparison of the Two Solutions	8
5. Theory of Optimal Control	9
5.1. General Optimization Problem	9
5.2. Present-Value Hamiltonian	9

5.3. Present-Value Optimality Conditions	9
5.4. Current-Value Hamiltonian	10
5.5. Current-Value Optimality Conditions	10
6. Heuristic Derivation of the Maximum Principle	11
6.1. Value Function for the Discretized Optimization Problem	11
6.2. Bellman Equation	11
6.3. Hamilton-Jacobi-Bellman Equation	12
6.4. Derivation of the Optimality Conditions	13
7. Applications of the Hamilton-Jacobi-Bellman Equation	15

1. Preliminary Results

Before presenting the techniques of optimal control, we review two mathematical results: L'Hopital's rule and Leibniz's rule. These results are used in section 6 for the derivations of some of the results.

1.1. L'Hopital's rule

L'Hopital's rule states that if

$$\lim_{x \rightarrow x_0} f(x) = \lim_{x \rightarrow x_0} g(x) = 0,$$

then

$$\lim_{x \rightarrow x_0} \frac{f(x)}{g(x)} = \lim_{x \rightarrow x_0} \frac{f'(x)}{g'(x)}.$$

1.2. Leibniz's rule

Leibniz's rule states that if

$$I(z) \equiv \int_{a(z)}^{b(z)} f(x, z) dx$$

where x is the integration variable, z is a parameter and $f(x, z)$ is assumed to have a continuous derivative $\partial f(x, z)/\partial z$ in the interval $[a(z), b(z)]$, then the effect of change in z on the integral is given by

$$\frac{dI}{dz}(z) = \int_{a(z)}^{b(z)} \frac{\partial f}{\partial z}(x, z) dx + \frac{db}{dz}(z) \cdot f(b(z), z) - \frac{da}{dz}(z) \cdot f(a(z), z).$$

2. Consumption-Saving Problem

This section considers the same consumption-saving problem as in the notes on dynamic programming. The only difference is that the problem is now set in continuous time.

2.1. Description of the Problem

Taking initial wealth a_0 as given, the problem is to choose the consumption path $\{c(t)\}_{t \geq 0}$ to maximize the lifetime utility

$$(1) \quad \int_0^{\infty} e^{-\rho \cdot t} \cdot u(c(t)) dt,$$

subject to the law of motion

$$(2) \quad \dot{a}(t) = r \cdot a(t) - c(t).$$

The parameter $r > 0$ is the constant interest rate at which wealth is invested. The parameter $\rho > 0$ is the discount factor. The notation \dot{a} denotes the derivative of wealth a with respect to time t :

$$\dot{a}(t) \equiv \frac{\partial a(t)}{\partial t}.$$

2.2. Optimal-Control Approach

To solve this problem, it is inconvenient to use the Lagrangian technique or dynamic programming technique because they are designed for discrete-time optimization problems. Instead, we use a technique called *optimal control*. We will see in section 6 that optimal control is related both to the Lagrangian technique and the dynamic programming technique. But optimal control is designed for continuous-time optimization problems.

There are two ways to apply the optimal-control approach: one by forming a *present-value Hamiltonian*, the other by forming a *current-value Hamiltonian*. These two ways are roughly equivalent, but using the current-value Hamiltonian is usually simpler.

3. Solution with Present-Value Hamiltonian

We start by solving the consumption-saving problem with the present-value Hamiltonian.

3.1. State and Control Variables

First, we identify the *state variables* and *control variables*. A control variable can be adjusted at any time t whereas the evolution of a state variable follows a law of motion such as (2). Here, the control variable is consumption $c(t)$ and the state variable is wealth $a(t)$.

3.2. Present-Value Hamiltonian

Second, we write down the present-value Hamiltonian

$$\mathcal{H}(t) = e^{-\rho \cdot t} \cdot u(c(t)) + \lambda(t) \cdot (r \cdot a(t) - c(t)).$$

To form the Hamiltonian, we introduce a new variable $\lambda(t)$, which we call the *costate variable* associated with the state variable $a(t)$. In general, we introduce as many costate variables as there are state variables. As a consequence, we introduce as many costate variables as there are laws of motion, because each state variable is associated with a law of motion.

The costate variable is analogous to a Lagrange multiplier, except that it is associated to a law of motion instead of a static constraint. But just like a Lagrange multiplier, the costate variable measures the value from marginally relaxing the constraint at the optimum.

3.3. Optimality Conditions

Next we write down the two optimality conditions, which are derived from the Maximum Principle:

$$(3) \quad \frac{\partial \mathcal{H}(t)}{\partial c(t)} = 0$$
$$(4) \quad \frac{\partial \mathcal{H}(t)}{\partial a(t)} = -\dot{\lambda}(t).$$

The optimality conditions can be rewritten as

$$(5) \quad 0 = e^{-\rho \cdot t} \cdot u'(c(t)) - \lambda(t)$$

$$(6) \quad -\dot{\lambda}(t) = \lambda(t) \cdot r.$$

3.4. Transversality Condition

Then we impose a *transversality condition*:

$$(7) \quad \lim_{t \rightarrow +\infty} \lambda(t) \cdot a(t) = 0.$$

The transversality condition describes what must be satisfied at the end of the time horizon at the optimum. Here it says that at the optimum, at the end of time, either there is no wealth left ($a(t) = 0$) or wealth has no value ($\lambda(t) \cdot a(t) = 0$).¹ This clearly has to be the case at the optimum: if there was some wealth left and wealth had value, then all the wealth that is left over should be consumed. The idea behind the transversality condition is that if there is any flexibility at the end of time, then the marginal benefit from exploiting that flexibility must be zero at the optimum.

3.5. Euler Equation

We solve explicitly for the optimal consumption path by eliminating the costate variable $\lambda(t)$ using (5) and (6).

To eliminate $\lambda(t)$, we first take log of (5):

$$-\rho \cdot t + \ln(u'(c(t))) = \ln(\lambda(t)).$$

We then take time derivatives in this equation:

$$\rho + \left[\frac{-u''(c(t)) \cdot c(t)}{u'(c(t))} \right] \cdot \left[\frac{\dot{c}(t)}{c(t)} \right] = -\frac{\dot{\lambda}(t)}{\lambda(t)}.$$

Equation (6) can be rewritten as

$$-\frac{\dot{\lambda}(t)}{\lambda(t)} = r.$$

¹Recall that the costate variable $\lambda(t)$ measures the marginal value of wealth at the optimum.

Combining these two equations, we obtain the Euler equation for optimal consumption:

$$(8) \quad \frac{\dot{c}(t)}{c(t)} \cdot \left[\frac{-u''(c(t)) \cdot c(t)}{u'(c(t))} \right] = r - \rho.$$

The term

$$\frac{-u''(c(t)) \cdot c(t)}{u'(c(t))}$$

measures relative risk aversion. The coefficient of relative risk aversion also corresponds to the inverse of the intertemporal elasticity of substitution.

3.6. CRRA Utility

Consider the following utility function:

$$u(c) = \frac{c^{1-\gamma} - 1}{1-\gamma}.$$

This utility function is known as Constant Relative Risk Aversion (CRRA) utility. It is characterized by a constant coefficient of relative risk aversion γ as

$$\frac{-u''(c) \cdot c}{u'(c)} = \gamma.$$

With CRRA utility, the Euler equation simplifies to

$$\frac{\dot{c}(t)}{c(t)} = \frac{r - \rho}{\gamma}.$$

4. Solution with Current-Value Hamiltonian

Next we solve the consumption-saving problem with a current-value Hamiltonian.

4.1. Current-Value Hamiltonian

The present-value Hamiltonian \mathcal{H} depends on t because of the discounting $e^{-\rho \cdot t}$, which might create some difficulties in deriving and analyzing solutions to the problem. Multiplying \mathcal{H} by $e^{\rho \cdot t}$ addresses this problem.

We denote the resulting Hamiltonian \mathcal{H}^* as current-value Hamiltonian. The current-value Hamiltonian is given by

$$\mathcal{H}^*(t) \equiv e^{\rho \cdot t} \cdot \mathcal{H}(t) = u(c(t)) + e^{\rho \cdot t} \cdot \lambda(t) \cdot (r \cdot a(t) - c(t)).$$

We define a new costate variable $q(t)$ as

$$q(t) \equiv e^{\rho \cdot t} \cdot \lambda(t).$$

The current-value Hamiltonian becomes

$$\mathcal{H}^*(t) = u(c(t)) + q(t) \cdot (r \cdot a(t) - c(t)).$$

This is the expression of the current-value Hamiltonian that we use in practice.

4.2. Optimality Conditions

The optimality conditions are slightly different. The optimality conditions become

$$\begin{aligned} \frac{\partial \mathcal{H}^*(t)}{\partial c(t)} &= 0 \\ \frac{\partial \mathcal{H}^*(t)}{\partial a(t)} &= \rho \cdot q(t) - \dot{q}(t). \end{aligned}$$

Compared to the optimality conditions (3) and (4) with the present-value Hamiltonian, there is an extra term $+\rho \cdot q(t)$ in the second condition. The extra term arises because the costate variable $q(t)$ is defined differently from the costate variable $\lambda(t)$ that we used in the present-value Hamiltonian.

The optimality conditions can be rewritten as

$$(9) \quad 0 = u'(c(t)) - q(t)$$

$$(10) \quad \rho \cdot q(t) - \dot{q}(t) = q(t) \cdot r.$$

4.3. Transversality Condition

The transversality condition becomes

$$\lim_{t \rightarrow +\infty} e^{-\rho \cdot t} \cdot q(t) \cdot a(t) = 0.$$

Compared to the transversality condition (7) with the present-value Hamiltonian, there is an extra factor $e^{-\rho \cdot t}$ in this condition. The extra factor arises because the costate variable $q(t)$ is defined differently from the costate variable $\lambda(t)$ that we used in the present-value Hamiltonian. But the interpretation of the transversality condition remains the same: at the optimum, at the end of time, any leftover wealth must be consumed, except if the present-discounted marginal value of wealth is zero at the end of time.

4.4. Euler Equation

By combining equations (9) and (10), we obtain exactly the same condition (8) as with the present-value Hamiltonian. We take the log of equation (9) to obtain

$$\ln(u'(c(t))) = \ln(q(t)).$$

We then take time derivatives in this equation:

$$\left[\frac{-u''(c(t)) \cdot c(t)}{u'(c(t))} \right] \cdot \left[\frac{\dot{c}(t)}{c(t)} \right] = -\frac{\dot{q}(t)}{q(t)}.$$

Equation (10) can be rewritten

$$-\frac{\dot{q}(t)}{q(t)} = r - \rho.$$

Combining these two equations, we obtain the same Euler equation for optimal consumption as the one we obtained with the present-value Hamiltonian:

$$\frac{\dot{c}(t)}{c(t)} \cdot \left[\frac{-u''(c(t)) \cdot c(t)}{u'(c(t))} \right] = r - \rho.$$

4.5. Comparison of the Two Solutions

The two approaches—the approach with the present-value Hamiltonian and the approach with the current-value Hamiltonian—are equivalent. They lead to the same Euler equation. However, it is often more convenient to work with the current-value Hamiltonian.

5. Theory of Optimal Control

Optimal control can be used to solve any continuous-time optimization problem. This section provides optimality conditions for a general optimization problem.

5.1. General Optimization Problem

The general problem is to choose $\{c(t)\}_{t \geq 0}$ to maximize

$$(11) \quad \int_0^{\infty} e^{-\rho \cdot t} \cdot u(a(t), c(t))$$

given the constraint that for all t

$$(12) \quad \dot{a}(t) = g(a(t), c(t)),$$

and taking a_0 as given. The parameter $\rho > 0$ is the discount rate. The functions u and g are concave and twice differentiable.

5.2. Present-Value Hamiltonian

An important result in optimal control theory is the Maximum Principle. It is due to Pontryagin. In addition to the control and state variables, we introduce a costate variable $\lambda(t)$ associated with the state variable. The costate variable measures the shadow price of the associated state variable. The costate variable enters the optimal control problem through the present-value Hamiltonian, defined as

$$(13) \quad \mathcal{H}(t) = e^{-\rho \cdot t} \cdot u(a(t), c(t)) + \lambda(t) \cdot g(a(t), c(t)).$$

5.3. Present-Value Optimality Conditions

The Maximum Principle gives necessary conditions for optimality. There are three conditions. The first two conditions are

$$(14) \quad \frac{\partial \mathcal{H}(t)}{\partial c(t)} = 0$$

$$(15) \quad \frac{\partial \mathcal{H}(t)}{\partial a(t)} = -\dot{\lambda}(t).$$

Condition (14) implies that the Hamiltonian must be maximized with respect to the control variable at any point in time. Condition (15) says that the marginal change of the Hamiltonian associated with a unit change of the state variable is equal to minus the rate of change of the costate variable.

The optimal solution must also satisfy a third condition, which we call transversality condition:

$$(16) \quad \lim_{t \rightarrow \infty} \lambda(t) \cdot a(t) = 0.$$

The transversality condition implies the product of costate and state must be converging to zero as time goes to infinity.

5.4. Current-Value Hamiltonian

We can reformulate the results from the Maximum Principle with the current-value Hamiltonian which is often easier to manipulate. The current-value Hamiltonian is defined as

$$\mathcal{H}^*(t) = u(a(t), c(t)) + q(t) \cdot g(a(t), c(t)),$$

where $q(t)$ is the costate variable associated with the state variable $a(t)$.

5.5. Current-Value Optimality Conditions

With the current-value Hamiltonian, the three necessary conditions (14), (15), and (16) for optimality become

$$\begin{aligned} \frac{\partial \mathcal{H}^*(t)}{\partial c(t)} &= 0 \\ \frac{\partial \mathcal{H}^*(t)}{\partial a(t)} &= \rho \cdot q(t) - \dot{q}(t) \\ \lim_{t \rightarrow \infty} e^{-\rho \cdot t} \cdot q(t) \cdot a(t) &= 0. \end{aligned}$$

6. Heuristic Derivation of the Maximum Principle

In this section, we provide an heuristic derivation of the necessary conditions for optimality provided by the Maximum Principle. One way to derive the optimality conditions (14) and (15) is to apply informally the results from dynamic programming. Formally, many of the claims below are imprecise, but they will serve the purpose of providing intuition for the Maximum Principle.

6.1. Value Function for the Discretized Optimization Problem

We begin by defining the value function of the problem, which is the maximized value of the objective function as a function of the state variable $a(t)$ and time t :

$$V(a(t), t) = \max_{\{c(s)\}_{s \geq t}} \int_t^\infty e^{-\rho \cdot (s-t)} \cdot u(a(s), c(s)) ds,$$

where the maximization is subject for all $s \geq t$ to the law of motion of the state variable

$$(17) \quad \dot{a}(s) = g(a(s), c(s)).$$

6.2. Bellman Equation

Since the problem has a recursive structure, we can apply the Principle of Optimality and write the value function as the solution to a Bellman equation:

$$V(a(t), t) = \max_{\{c(s)\}_{t \leq s \leq t+\Delta t}} \left\{ \int_t^{t+\Delta t} e^{-\rho \cdot (s-t)} \cdot u(a(s), c(s)) ds + e^{-\rho \cdot \Delta t} \cdot V(a(t+\Delta t), t+\Delta t) \right\},$$

where the maximization is subject for all $t \leq s \leq t + \Delta t$ to (17).

Subtract $V(a(t), t)$ from both side and divide by Δt :

$$(18) \quad 0 = \max_{\{c(s)\}_{t \leq s \leq t+\Delta t}} \left[\frac{\int_t^{t+\Delta t} e^{-\rho \cdot (s-t)} \cdot u(a(s), c(s)) ds}{\Delta t} + \frac{e^{-\rho \cdot \Delta t} \cdot V(a(t+\Delta t), t+\Delta t) - V(a(t), t)}{\Delta t} \right],$$

where the maximization is subject for all $t \leq s \leq t + \Delta t$ to (17).

6.3. Hamilton-Jacobi-Bellman Equation

We now take the limit of (18) as $\Delta t \rightarrow 0$ to obtain the Hamilton-Jacobi-Bellman equation.

Limit of First Term. We start with the first term in the curly brackets. Since numerator and denominator of the first term approach zero as $\Delta t \rightarrow 0$, we apply L'Hopital's rule. The derivative of the denominator with respect to Δt is 1. We apply Leibniz's rule to determine the derivative with respect to Δt of the integral in the numerator. Leibniz's rule tells us that the derivative of the integral with respect to Δt is

$$e^{-\rho \cdot \Delta t} \cdot u(a(t + \Delta t), c(t + \Delta t)).$$

Therefore, the limit as $\Delta t \rightarrow 0$ for the first term in the bracket is

$$(19) \quad u(a(t), c(t)).$$

Limit of Second Term. We move on to the second term. Since

$$\lim_{\Delta t \rightarrow 0} e^{-\rho \cdot \Delta t} = 1$$

and

$$\lim_{\Delta t \rightarrow 0} V(a(t + \Delta t), t + \Delta t) = V(a(t), t),$$

both numerator and denominator approach zero as $\Delta t \rightarrow 0$. Therefore we apply L'Hopital's rule. The derivative of the denominator with respect to Δt is 1. The derivative of the numerator with respect to Δt is

$$\begin{aligned} & -\rho \cdot e^{-\rho \cdot \Delta t} \cdot V(a(t + \Delta t), t + \Delta t) + e^{-\rho \cdot \Delta t} \cdot \frac{\partial V}{\partial a}(a(t + \Delta t), t + \Delta t) \cdot \dot{a}(t + \Delta t) \\ & + e^{-\rho \cdot \Delta t} \cdot \frac{\partial V}{\partial t}(a(t + \Delta t), t + \Delta t). \end{aligned}$$

We have the following limits:

$$\begin{aligned} \lim_{\Delta t \rightarrow 0} \rho \cdot e^{-\rho \cdot \Delta t} \cdot V(a(t + \Delta t), t + \Delta t) &= \rho \cdot V(a(t), t) \\ \lim_{\Delta t \rightarrow 0} e^{-\rho \cdot \Delta t} \cdot \frac{\partial V}{\partial t}(a(t + \Delta t), t + \Delta t) &= \frac{\partial V}{\partial t}(a(t), t) \\ \lim_{\Delta t \rightarrow 0} e^{-\rho \cdot \Delta t} \cdot \frac{\partial V}{\partial a}(a(t + \Delta t), t + \Delta t) \cdot \dot{a}(t + \Delta t) &= \frac{\partial V}{\partial a}(a(t), t) \cdot \dot{a}(t) = \frac{\partial V}{\partial a}(a(t), t) \cdot g(a(t), c(t)), \end{aligned}$$

where the last equality results from the law of motion (12) of state variable $a(t)$.

Therefore, the limit as $\Delta t \rightarrow 0$ for the first term in the bracket is

$$(20) \quad -\rho \cdot V(a(t), t) + \frac{\partial V}{\partial a}(a(t), t) \cdot g(a(t), c(t)) + \frac{\partial V}{\partial t}(a(t), t).$$

Derivation of the Hamilton-Jacobi-Bellman Equation. Combining equations (18), (19), and (20), we obtain a version of the Bellman equation for the continuous-time optimization problem. This equation is called the *Hamilton-Jacobi-Bellman equation*. The equation is

$$(21) \quad \rho V(a(t), t) = \max_{c(t)} \left[u(a(t), c(t)) + \frac{\partial V}{\partial a}(a(t), t) \cdot g(a(t), c(t)) + \frac{\partial V}{\partial t}(a(t), t) \right],$$

where $a(t)$ is given. We define

$$\lambda(t) \equiv e^{-\rho \cdot t} \cdot \frac{\partial V}{\partial a}(a(t), t).$$

We can rewrite the Hamilton-Jacobi-Bellman equation as

$$(22) \quad \rho V(a(t), t) = \max_{c(t)} \left[u(a(t), c(t)) + e^{\rho \cdot t} \cdot \lambda(t) \cdot g(a(t), c(t)) + \frac{\partial V}{\partial t}(a(t), t) \right].$$

6.4. Derivation of the Optimality Conditions

Taking the first-order condition with respect to $c(t)$ in the Hamilton-Jacobi-Bellman equation (22) implies

$$\frac{\partial u}{\partial c(t)}(a(t), c(t)) + e^{\rho \cdot t} \cdot \lambda(t) \cdot \frac{\partial g}{\partial c(t)}(a(t), c(t)) = 0.$$

Furthermore, the envelope theorem implies

$$\rho \frac{\partial V}{\partial a(t)}(a(t), t) = \frac{\partial u}{\partial a(t)}(a(t), c(t)) + e^{\rho \cdot t} \cdot \lambda(t) \cdot \frac{\partial g}{\partial a(t)}(a(t), c(t)) + \frac{\partial^2 V}{\partial t \partial a(t)}(a(t), t).$$

The last two equations become equivalent to optimality conditions (14) and (15). This is the case because, using the definition of the present-value Hamiltonian, equation (14) can be written

$$\frac{\partial H(t)}{\partial c(t)}(a(t), c(t)) = 0$$

and equation (15) can be written

$$\frac{\partial u}{\partial a(t)}(a(t), c(t)) + e^{\rho \cdot t} \cdot \lambda(t) \cdot \frac{\partial g}{\partial a(t)}(a(t), c(t)) = -e^{\rho \cdot t} \cdot \frac{\partial \lambda(t)}{\partial t},$$

which implies

$$\frac{\partial H(t)}{\partial a(t)}(a(t), c(t)) = -\frac{\partial \lambda(t)}{\partial t}.$$

Note that we consider that the Hamiltonian is a function of a_t , t , and λ_t , and we only take the partial derivative with respect to a_t , thus keeping λ_t constant.

7. Applications of the Hamilton-Jacobi-Bellman Equation

The Hamilton-Jacobi-Bellman equation (21) is an optimality condition that equates flow costs with flow benefits. In practice, we write it down without going through all the algebra relating to Δt .

This equation is commonly used in macroeconomics. For instance, it is frequently used in search-and-matching models of the labor market. In a search-and-matching model, a vacant job costs c per unit time and becomes occupied according to a Poisson process with arrival rate q . In the labor market, the occupied job yields net returns $p - w$, where p is real output and w is the cost of labor. The job runs a risk λ of being destroyed.

Let V be the value of the vacant job and J be the value of occupied job. Let r be the discount factor. In steady state, the Hamilton-Jacobi-Bellman equations are

$$\begin{aligned}r \cdot V &= -c + q \cdot (J - V), \\r \cdot J &= p - w + \lambda \cdot (0 - J) = p - w - \lambda \cdot J.\end{aligned}$$

There is no maximization on the right-hand-side in this particular example. These equations simply describe the relationship between the equilibrium values V , J , and w .