OPTIMAL CONTROL THEORY

PEDRO HEMSLEY

Preliminary and Incomplete - please do not cite

These short notes are based on Rochet's notes for a course in dynamic optimization[1,2]. They are no substitute for the original notes. They are no substitute for a course in optimization. I try to present the main results in optimal control theory along with some classic economic examples. All mistakes are mine. All relevant disclaimers not mentioned above nevertheless apply.

I didn't include anything on calculus of variations because it's just a particular case of optimal control theory. I didn't include anything on differential equations, which are necessary if you want to study dynamic optimization. And I skipped all proofs and most technical details.

Yet, I think it might help both those who need a quick reference, as those who don't know anything and need just a quick look before tackling a paper or following a class. The first and the second parts deal with finite horizon problems and are particularly relevant for contract theory. The third deals with infinite horizon and is the basic tool in macroeconomics. If you understand the first part, it will be enough in most applications. The final part presents some classic examples in contract theory.

Part 1. Basic Problem

The basic problem is to maximize a given function U over an interval $[t_0, t_1]$, including possibly a final value function A, under a set of restrictions.

Problem P:

$$\underset{x,v}{Max}\ \Psi(x,v) \equiv \int_{t_0}^{t_1} U(t, x(t), v(t))\,dt + A(x(t_1))$$

subject to:

$\dot{x}(t) = f(t, x(t), v(t))$ *whenever v is continuous*
$v(t)$ *admissable*[3]
$x(t_o) = x_0$ *given*

We call $x_t \in \mathbb{R}^n$ the state variable and v_t the control variable. We look for the optimal path of both over time (or over the relevant space). For example, in some Macro models, x_t is the stock of capital and v_t is consumption[4]. Notice that there's no restriction on the final value of the state variable x, only on the initial value.

If $f(t, x(t), v(t)) = v_t$, then we would be talking about calculus of variations: this is the case when the function U depends directly on the derivative $\dot{x}$. All theorems stated below would be valid.

We will never work directly with the integral form in the objective function. The basic tool that we will use to solve such problems is the Hamiltonian, defined below.

Definition 1. Hamiltonian

$$H(t, x_t, v_t, p_t) = U(t, x_t, v_t) + p_t f(t, x_t, v_t)$$

Maximized Hamiltonian

$$H^* = \underset{v}{Max}\ H$$

[1] Rochet, J-C. Dynamic Optimization in Continuous Time. Toulouse, 2007

[2] Comments are welcome: pedrohemsley@yahoo.com.br

[3] You won't need to worry about this in most basic examples.

[4] To help intuition, you may think of a consumer who lives from period t_0 to t_1 and wants to maximize his life utility plus a "legacy value" A (bequest, for example). He accumulates capital according to the function f and has an initial stock of capital (or wealth) x_0.

In these definitions, p_t is just the co-state variable, i.e., the (dynamic) multiplier. It plays the same role as the traditional Lagrange multiplier in static optimization problems, but now it must defined for each point of the interval $[t_0, t_1]$ - that is, it's a function. Notice that the terminal-value function $A(.)$ does not enter in the Hamiltonian.

The first result describes the necessary conditions for a pair (x, v) to solve the problem above.

Theorem 2. *Pontryaguin Maximum Principle*
if (x^*, v^*) solves P, then there is a function p such that the following conditions hold:
i) $v^* \in \underset{v}{argmax}\ H$
ii) $\dot{x} = \frac{\partial H^*}{\partial p}$ and $\dot{p} = -\frac{\partial H^*}{\partial x}$,
with boundary conditions $x(t_0) = x_0$ and $p(t_1) = \nabla A(x_1)$

Condition (i) is known as the maximum principle. In most problems, it just means $\frac{\partial H}{\partial v} = 0$. It plays the same role as setting the derivative of the Lagrangean equal to zero in static problems. Condition (ii) is the Jacobi-Hamilton system of (differential) equations[5,6]. Notice that the first one (the condition on $\dot{x}$) is just the law of motion of x, i.e., the original restriction. This will alway be the case. Notice also that we have $2N$ boundary conditions because $x \in \mathbb{R}^n$.

If you're interested only in the value of the control variable v, then maybe it will not be necessary to find x. This is often the case in growth models and contract theory models.

If $f(t, x(t), v(t)) = v$ and we are dealing with a problem of calculus of variations, then $\frac{\partial H}{\partial v} = \frac{\partial U}{\partial v} + p$ and conditions (i) and (ii) would read:

$$(i)\ \frac{\partial U}{\partial v} + p = 0$$

$$(ii)\ \dot{p} = -\frac{\partial H^*}{\partial x}$$

These two equations together imply $\frac{d}{dt}\left(\frac{\partial U}{\partial v}\right) = \frac{\partial H^*}{\partial x}$, the usual Euler equation.

We have two theorems that describe sufficient conditions. Essentially, the first one says that if the funcion $(x, v) \to H$ and the terminal funcion $x(t_1) \to A$ are concave, then the conditions above are sufficient. The second uses the optimized hamiltonian and states that if the functions $x \to H^*$ and $x(t_1) \to A$ are concave, then the conditions above are sufficient. The only difference is that in the second case we use the optimized hamiltonian, which doesn't depend on the control variable v. The formal statements are given below.

Theorem 3. *Mangasarian's sufficient conditions. Assume the conditions of Pontryaguin's theorem hold; the function $(x, v) \to H(t, x, p(t), v)$ is concave for all t; and A is concave. Then (x^*, v^*) is a solution to P.*

Notice that this function must be concave for all values of (x, v), and not only at the optimal point.

Theorem 4. *Arrow's sufficient conditions. Assume the conditions of Pontryaguin's theorem hold; the function $x \to H^*(t, p(t), v)$ is well-defined and concave for all t; and A is concave. Then (x^*, v^*) is a solution to P.*

Part 2. Some extensions

We can adapt Pontriaguin's Maximum Principle if the problem is not exactly the one described above[7].

First, consider that we have restrictions on the final value x_1. Remember it's a vector, and we may have restrictions on some dimensions. For simplicity, there's no terminal value function now ($A(x_1)$). Consider problem P':

$$\underset{x,v}{Max}\ \Psi(x, v) \equiv \int_{t_0}^{t_1} U(t, x(t), v(t))\, dt$$

subject to:

[5] This terminology is not always followed, though.

[6] Notation: $\nabla A(x_1) = \left(\frac{\partial A}{x_{11}}, ..., \frac{\partial A}{x_{N1}}\right)$

[7] This part should be skipped in a first reading.

$\dot{x}(t) = f(t, x(t), v(t))$ *whenever* v *is continuous*
$v(t)$ *admissable*
$x(t_o) = x_0$ *given*
$x_i(t_1) = x_{i1}, \forall i \in I,$ *given*
$x_j(t_1) \geq x_{j1}, \forall j \in J,$ *given*
$x_k(t_1)$ unrestricted $\forall k \in K$

in which I, J and K form a partition of $\mathbb{R}^N$ (remember that $x \in \mathbb{R}^N$).

The formal definition of the hamiltonian becomes a bit different:

$$H = p_0 U + pf$$

as we must now associate a co-state variable to the funcion U. The optimized hamiltonian is analogously rewritten. However, we can always normalize (p_0, p) such that $p_0 = 1$. This is not necessarily the adequate normalization in many problems, but to keep things simple here we will adopt it, so that the theorems given below will be a restricted version.

Theorem 5. *Pontryaguin Maximum Principle - version II.*
if (x^*, v^*) *solves* P'*, then there is a function* p *such that the following conditions hold:*
i) $v^* \in \underset{v}{argmax} H$
ii) $\dot{x} = \frac{\partial H^*}{\partial p}$ *and* $\dot{p} = -\frac{\partial H^*}{\partial x}$,
with boundary conditions $x(t_0) = x_0$; $x_i(t_1) = x_{i1}$ *if* $i \in I$; $p_i(t_1) \geq 0$ *and* $(p_i(t_1) \geq 0)(x_i(t_1) - x_{i1}) = 0$ *if* $i \in J$; *and* $p_i(t_1) = 0$ *if* $i \in K$.

Notice that the only thing that changes is the boundary conditions of the Jacobi-Hamilton equations: one must include "slackness" conditions in the same spirit as the KT theorem.

Sometimes you will see an inequality restriction on $\dot{x}$. If the objective function U does not depend directly on $\dot{x}$[8], it's easy to deal with it. The problem will read:

$$Max \int_{t_0}^{t_1} U(t, x(t))$$

$$\begin{cases} \dot{x}(t) \geq 0 \\ x(t_0) \geq 0 \end{cases}$$

Now, $\dot{x}$ will be the control variable, and x is as usual the state variable. The hamiltonian becomes:

$$H = U(t, x) + p\dot{x}$$

and the necessary conditions for a solution are:

$\frac{\partial H}{\partial x} = 0$

$$\begin{cases} \dot{p} = -\frac{\partial U}{\partial x} \\ p(t_1) = 0 \\ p(t_0)x(t_0) = 0 \end{cases} \quad \text{and} \quad \begin{cases} p(t) \leq 0 \\ x(t) \geq 0 \\ p(t)\dot{x}(t) = 0 \end{cases}$$

The first condition is the usual maximum principle. The H-J equation for $\dot{p}$ has an addition "slackness" boundary condition. Notice that there is no J-H equation for $\dot{x}$ (since there's none in the problem). The last conditions are not actually new if you're used to static optimization: the multiplier is non-zero only if the restriction is binding ($\dot{x} = 0$). This is simpler than it looks: it says that whenever the the restriction on $\dot{x}$ is not binding, we have the same solution as in the relaxed problem (i.e., ignoring $\dot{x} \geq 0$). If it's biding, then x is constant (zero derivative) and can be determined from the H-J equation.

Sufficient conditions are similar to the ones before and I think it's not that important to include them here. Take a look at Rochet's notes.

[8]This happens, for example, when you have bunching in adverse selection models.

Part 3. Infinite Horizon

We now turn to the infinite horizon problem, as you see all the time in macroeconomics. Essentially, the boundary conditions will change.

Problem P":

$\underset{x,v}{Max}\ \Psi(x,v) \equiv \int_{t_0}^{\infty} U(t,x(t),v(t))\,dt$

subject to:

$\dot{x}(t) = f(t,x(t),v(t))$ *whenever v is continuous*
$v(t)$ *admissable*
$x(t_o) = x_0$ *given*

Notice that it doesn't make sense to have a final value function $A(x_\infty)$, as the agent doesn't actually derive utility from it (which is good news). This problem makes sense only when the improper integral above converges to a finite value, as will be the case, for example, when $U(t,x(t),v(t)) = e^{-\delta t}u(v_t)$.

Theorem 6. *Pontryaguin Maximum Principle - version III*

if (x^*,v^*) *solves* P''*, then there is a function p such that the following conditions hold:*
i) $v^* \in \underset{v}{argmax}\ H$
ii) $\dot{x} = \frac{\partial H^*}{\partial p}$ and $\dot{p} = -\frac{\partial H^*}{\partial x}$,
with boundary conditions $x(t_0) = x_0$.

Notice that we do not have to impose anything on $\underset{t\to\infty}{lim}\,p$, but only because these conditions are necessary, not sufficient. The classic example for this case is the neoclassical growth model (Macro 1).

Part 4. Examples

1. Contract theory: classic example

We have one principal and one agent. The principal can sell a quantity $q \in \mathbb{R}$ to the agent for a total price t (transfer). They are characterized respectively by:

$$\pi = t - cq$$

$$u = v(\theta,q) - t$$

The agent has some private information θ that determines how much he's willing to pay for a given quantity. It's common knowledge that θ is distributed according to $F(\theta)$ on the support $[\theta_0,\theta_1]$ with associated density $f(\theta)$. The principal has all the bargaining power. He moves first and proposes a take-it-or-leave-it offer to the agent - that is, he offers a contract (t,q) to the agent, who can either accept or reject. Full commitment holds: if the agent accepts, the contract (t,q) is implemented (no kind of renegotiation). We will use the revelation principle and focus on contracts $(t(\theta),q(\theta))_\theta$.

The problem of the principal is to maximize his expected profit:

$$\begin{cases} \underset{\{q(\theta),t(\theta)\}_\theta}{Max} \int_{\theta_0}^{\theta_1} \{t(\theta) - cq(\theta)\} f(\theta)d\theta \\ v(\theta,q(\theta)) - t(\theta) \geq 0,\ \forall\theta \\ v(\theta,q(\theta)) - t(\theta) \geq v(\theta,q(\theta')) - t(\theta'),\ \forall\theta,\theta' \end{cases}$$

The interval $[\theta_0,\theta_1]$ plays the same role as $[t_0,t_1]$. This problem can be rewritten as follows[9]:

$$\begin{cases} \underset{\{q(\theta)\}_\theta}{Max} & \int_{\theta_0}^{\theta_1} \{v(\theta,q(\theta)) - cq(\theta) - u(\theta)\} f(\theta)d\theta \\ \dot{u}(\theta) = v_\theta(\theta,q(\theta)) & \forall\theta \\ \dot{q}(\theta) \geq 0 & \forall\theta \\ u(\theta_0) = 0 & \end{cases}$$

[9] For details, see Laffont & Martimort (2002).

Let's ignore the restriction $\dot{q}(\theta) \geq 0$ for now (just assume it's not binding[10]). Then we have a standard optimal control theory problem. The state variable is u; the control variable is q; and there's no final value function: $A(u(\theta_1)) = 0$ (keep in mind that this implies $A' = 0$)[11].

Write the hamiltonian, with co-state variable p:

$$H = (v(\theta, q) - cq - u)\, f(\theta) + p v_\theta(\theta, q)$$

By the maximum principle:

$$\frac{\partial H}{\partial q} = 0 \Rightarrow (v_q - c)\, f(\theta) + p v_{\theta q}(\theta, q) = 0 \tag{1.1}$$

The Jacobi-Hamilton equations are:

$$\frac{\partial H}{\partial p} = \dot{u} \Rightarrow \dot{u}(\theta) = v_\theta(\theta, q(\theta))$$

(as always, this is just the restriction)

$$\frac{\partial H}{\partial u} = -\dot{p} \Rightarrow -f(\theta) = -\dot{p} \Rightarrow f(\theta) = \dot{p}$$

The boundary conditions for this system are:

$$u(\theta_0) = 0$$

$$p(\theta_1) = 0$$

The last equation follows because $A' = 0$.

Notice first that the J-H "system" is actually made up of two independent equations that we can solve separately. Since we are not interested for the moment in the level of u at the optimum (i.e., the optimal path of the state variable) because it won't help me find the optimal q, I don't need to solve the first H-J equation.

Also, notice that the first boundary condition refers to θ_0, while the second is evaluated at θ_1 - this may be a little confusing if you're used to the "integration by parts" method.

Now we just need to solve the differential equation $\dot{p} = f(\theta)$ with boundary $p(\theta_1) = 0$. The solution is[12]

$$p(\theta) = F(\theta) - 1$$

If we plug this back into equation 1.1, we get:

$$v_q - \frac{(1 - F(\theta))}{f(\theta)} v_{q\theta} = c$$

at an interior solution. If the function $q(\theta)$ implicitly determined is non-decreasing, then this is the actual solution.

---

[10] It would still be an optimal control problem - with some adjustments - if we didn't ignore this restriction.

[11] Since u is unidimensional, $\nabla A \equiv A'$

[12] $\frac{dp(s)}{ds} = f(s) \Rightarrow dp(s) = f(s)ds \Rightarrow \int_\theta^{\theta_1} dp(s) = \int_\theta^{\theta_1} f(s)ds \Rightarrow p(\theta_1) - p(\theta) = F(\theta_1) - F(\theta) \Rightarrow -p(\theta) = 1 - F(\theta)$, in which the last equality follows from the boundary condition $p(\theta_1) = 0$ and the fact that the distribution function F equals one at the upper limit of the support.

2. Regulation: classic example

The principal is now the regulator who wants a firm (the agent) to exert cost-reducing effort. For details, see Laffont and Tirole (1993). The classic problem is:

$$\begin{cases} Max & \int_{\underline{\beta}}^{\bar{\beta}} \left[S-(1+\lambda)\left(\beta-e(\beta)+\psi(e(\beta))-\lambda U(\beta)\right]f(\beta)d\beta \\ st & \\ \dot{U}(\beta)= & -\psi'(e(\beta)) \\ \dot{e}(\beta)\leq & 1 \\ U(\overline{\beta})= & 0 \end{cases}$$

S is the (fixed) social value of the project; λ is the shadow cost of public funds; the firm's cost function is $C=\beta-e(\beta)$; and effort e costs $\psi(e)$ to the firm. The main objective of this exercise is to show that the same methods apply when you have β instead of θ. Again, we will ignore the monotonicity constraint $\dot{e}(\beta)\leq 1$. The control variable is e; the state variable is U. The hamiltonian is:

$$H=\left[S-(1+\lambda)\left(\beta-e(\beta)+\psi(e(\beta))-\lambda U(\beta)\right]f(\beta)+p\left[-\psi'(e(\beta))\right]$$

From the maximum principle, we get:

$$\frac{\partial H}{\partial e}=0\Rightarrow-(1+\lambda)\left(-1+\psi'\right)f(\beta)-p\psi''=0 \tag{2.1}$$

The J-H system is the following:

$$\dot{U}(\beta)=\frac{\partial H}{\partial p}=\dot{U}\Rightarrow-\psi'(e(\beta))$$

(guess what - this is just the restriction of the problem above)

$$\frac{\partial H}{\partial U}=-\dot{p}\Rightarrow-\lambda f(\beta)=-\dot{p}\Rightarrow\lambda f(\beta)=\dot{p}$$

Pay attention to the boundary conditions (this is tricky):

$U(\bar{\beta})=0$
$p(\underline{\beta})=0$

Notice that if we interpret $t_0=\underline{\beta}$ and $t_1=\bar{\beta}$ (that is, if we define "initial value" and "final value" just by looking at the limits of the integral), we would have $p(\overline{\beta})=0$. This would be totally wrong. The point here is that we are looking at the basic problem (problem P, at the beginning of these notes). So we have the initial value of the state variable ("x_0"). If you look at the statement of the problem we're solving now, you will see that we have the value of the state variable U evaluated at $\bar{\beta}$: $U(\bar{\beta})=0$. So, $\bar{\beta}$ is the "initial value" of this problem, even if it's actually the endpoint of the support of F! Then, the "final value" is $\underline{\beta}$. Yes, this is a bit confusing, but pay attention to it. This implies that the condition $p(t_1)=\nabla A(x_1)$ in problem $\overline{P}$ should be written here as $p(\underline{\beta})=0$ (remember that $A=0$).

Again, we don't need to solve the whole J-H system: we have two independent equations and we need only p in order to find the optimal control (that is, the only unknown in 2.1 is p). This amounts to solving $\lambda f(\beta)=\dot{p}(\beta)$ with boundary condition $p(\underline{\beta})=0$. Solving this, we get $p(\beta)=\lambda F(\beta)$. Plugging this into 2.1, you will find:

$$\psi'=1-\frac{\lambda}{1+\lambda}\frac{F(\beta)}{f(\beta)}\psi''$$

3. Public Finance

Let's solve the basic version of the optimal taxation problem (Mirrless, 1971[13]). We have the government (the principal) and workers/consumers (the agents). The government wants to maximize a function G that depends on the utility of each individual. Agents have private information about their productivity levels n, distributed according to $F(n)$ on support $[n_0, n_1]$. We assume the function G to be strictly increasing and concave, so that the government would like every agent to have the same consumption ("inequality aversion"). But this means that the most productive workers will consume less then they produce - therefore they have an incentive to hide their true type, and the government must take this into account. Also, total expenditure (private and public) cannot be larger than total production. The utility function of a typical agent is:

$$u(x,y) = x + v(1-y) = ny - T(ny) + v(1-y)$$

in which x stands for consumption, y for labor, n is the productivity of the agent and ny his realized output. Taxes $T(ny)$ depend on output, not directly on productivity. Government's expenditure is labeled E.

The problem reads as follows:

$$Max \int_{n_0}^{n_1} G\left[u(n)\right] f(n)dn$$

subject to:

$$\int_{n_0}^{n_1} \left[u(n) - v\left(1 - y(n)\right)\right] f(n) + E \leq \int_{n_0}^{n_1} ny(n)f(n)$$

$$\dot{u}(n) = y(n)\frac{v'\left[1-y(n)\right]}{n}$$

The first restriction is the budget constraint (BC): total consumption (the left-hand side - just look at the agent's utility function) cannot be greater than total production (the RHS). The second restriction is the usual incentive (truth-telling) constraint. Notice that there is no participation constraint: if you exist, you pay taxes.

Many people have trouble when they see an integral in the restriction, so for the moment let's focus on it. Assume initially there is no private information, so that the incentive constraint is irrelevant. Let p be the Lagrange multiplier of (BC)[14]. The lagrangean is:

$$L = \int_{n_0}^{n_1} G\left[u(n)\right] f(n)dn + p\left\{\int_{n_0}^{n_1} ny(n)f(n)dn - \int_{n_0}^{n_1}\left[u(n) - v\left(1-y(n)\right)\right] f(n)dn - E\right\}$$

Ignore the constant E (it doesn't affect the solution) and rewrite this expression as:

$$L = \int_{n_0}^{n_1} \left[G\left[u(n)\right] + \lambda p\left\{ny(n) - \left[u(n) - v\left(1-y(n)\right)\right]\right\}\right] f(n)dn$$

One can solve this problem by pointwise maximization - i.e., just ignore the integral sign. Just think of it as a summation over n. If you're not convinced, write the hamiltonian associated to this problem:

$$H = \left[G\left[u(n)\right] + \lambda\left\{ny(n) - \left[u(n) - v\left(1-y(n)\right)\right]\right\}\right] f(n)$$

There is no co-state variable, since the we assumed the restriction on the state variable u not to be binding. Then the J-H equations do not come into play, and the maximum principle $\frac{\partial H}{\partial y} = 0$ means that we're applying the usual first-order condition to the lagrangean, ignoring the integral sign.

To solve the complete problem, we just need to take IC into consideration (let $h(n)$ be the dynamic multiplier): The maximum principle implies:

$$-p\left\{n - v'\right\} f(n) = \frac{h(n)}{n}\left\{\left\{v' - y(n)v''\right\}\right\} \tag{3.1}$$

[13]Actually, I based myself on Diamond, 1998. The basic problem is the same.

[14]I use the original paper's notation, which is different from the one I've used so far. If you want, just take $p = \lambda$, $h(n) = p(n)$ and $n = \theta$. This probably will not affect the results.

The H-J equation for h is:

$$h'(n) = -\left\{G'(u) - p\right\} f(n)$$

With boundary condition $h(n_1) = 0$. I will ignore the J-H equation for the state variable u.

Solving this differential equation and plugging it back into 3.1, we get:

$$p(n - v')f = \frac{[v' - yv'']}{n} \left[\int_{n}^{n_1} (p - G')f(n)dn \right]$$

Notice that this is not a complete answer, since we haven't found the value of the static multiplier p - but this is beyond the scope of this notes. Keep in mind that $p > 0$ (the budget constraint is binding, since $G' > 0$), so most qualitative results of the basic taxation problem can be found in this last equation. The essential one is that when $n = n_1$, we have $v' = n_1$, which is the traditional "no distortion at the top" result: the marginal disutility of labor equals the individuals productivity.

Also, remember that the multiplier p is the shadow value of production[15] - that is, the value of an extra unit of y in the utility function of the government. Since preferences are quasi-linear, this extra unit will not affect the allocation of labor (no income effect). Then, on average, this should be equal to the marginal benefit G', i.e.,

$$p = \int_{n_0}^{n_1} G' f(n)dn$$

So the term in brackets on the RHS is equal to zero if $n = n_0$ and, again we have no distortion: $n_0 = v'\left(1 - y(n_0)\right)$. This "no distortion at the bottom" result is not totally general. It implies that the optimal scheme is "U-shaped": downwards distortion for every type but the highest and the lowest.

[15] Remeber the theory of the firm: when it minimizes cost subject to a given level of production, the multiplier is the marginal cost of production. Similar reasoning applies here.