Optimal control theory with applications to resource and environmental economics

Michael Hoel, March 16, 2015

This note gives a brief, non-rigorous sketch of basic optimal control theory, which is a useful tool in several simple economic problems, such as those in resource and environmental economics.

Consider the dynamic optimization problem

$$\max \int_0^\infty e^{-rt} f(x(t), S(t), t) dt \tag{1}$$

subject to

$$\dot{S}(t) = g(x(t), S(t), t) \tag{2}$$

$$S(0) = S_0 \text{ historically given} \tag{3}$$

$$S(t) \geq 0 \text{ for all } t \tag{4}$$

where f and g are continuous and differentiable functions (and in many cases concave in (x, S)), and r is an exogenous positive discount rate. The variable $S(t)$ is a *stock* variable, also called a *state* variable, and can only change gradually over time as given by (2). The variable $x(t)$, on the other hand, is a variable that the decision maker chooses at any time. It is often called a *control* variable. In many economic problems the variable $x(t)$ will be constrained to be non-negative.

Remark 1: In the problem above there is only one control variable and one state variable. It is straightforward to generalize to many control and state variables, and the number of control variables need not be equal to the number of state variables.

Remark 2: The constraint (4) is more general than it might seem, as we often can reformulate the problem so we get this type of constraint. Assume e.g. that the constraint was $S(t) \leq \bar{S}$. We can then reformulate the problem by defining $Z(t) = \bar{S} - S(t)$, implying that $Z(t) \geq 0$. In this case the dynamic equation (2) must be replaced by $\dot{Z} = -g(x(t), \bar{S} - Z(t), t)$ and $S(t)$ in (1) must be replaced by $\bar{S} - Z(t)$.

The current value Hamiltonian

The current value Hamiltonian H is defined as

$$H(x, S, \lambda, t) = f(x, S, t) + \lambda g(x, S, t)$$

where $\lambda(t)$ is continuous and differentiable. The variable $\lambda(t)$ is often called a *co-state variable*. This variable will be non-negative in all problems where "more of the state variable" is "good". More precisely: The derivative of the maximized integral in (1) with respect to S_0 is equal to $\lambda(0)$. For this reason $\lambda(t)$ is also often called the *shadow price* of the state variable $S(t)$.

Conditions for an optimal solution

A solution to the problem (1)-(4) is a time path of the control variable $x(t)$ and an associated time path for the state variable $S(t)$. For optimal paths, there exist a differentiable function $\lambda(t)$ and a piecewise continuous function $\gamma(t)$ such that the following equations must hold for all t:

$$\frac{\partial H(x(t), S(t), \lambda(t), t)}{\partial x} = 0 \tag{5}$$

$$\dot{\lambda}(t) = r\lambda(t) - \frac{\partial H(x(t), S(t), \lambda(t), t)}{\partial S} - \gamma(t) \tag{6}$$

$$\gamma(t) \geq 0 \text{ and } \gamma(t)S(t) = 0 \tag{7}$$

$$Lim_{t \to \infty} e^{-rt} \lambda(t) S(t) = 0 \tag{8}$$

Remark 3: If $x(t)$ is constrained to be non-negative, (5) must be replaced by $\frac{\partial H}{\partial x} \leq 0$ and $\frac{\partial H}{\partial x} x(t) = 0$.

Remark 4: If we know from the problem that $S(t) > 0$ for all t, we can forget about $\gamma(t)$, since it always will be zero.

Remark 5: Condition (8) is a transversality condition. Transversality conditions are simple in problems with finite horizons, but considerably more complicated for problems with an infinite horizon (like our problem). The condition (8) holds for all problems where $\lambda(t) \geq 0$, and also for most problems in economics even if $\lambda(t) < 0$.

Remark 6: If f and g are concave in (x, S) and $\lambda(t) \geq 0$, the conditions (5)-(8) are sufficient for an optimal solution. If we can find a time path for $x(t)$ and for $S(t)$ satisfying (5)-(8) in this case, we thus know that the time paths $(x(t), S(t))$ are optimal.

Remark 7: As mentioned in Remark 1, it is straightforward to generalize to many control and state variables. If there are n state variables, there are also n co-state variables $(\lambda_1, ...\lambda_n)$, n Lagrangian multipliers $(\gamma_1 ... \gamma_n)$, and n differential equations of each of the types (2) and (6).

Example 1: The optimal use of a non-renewable resource

Consider the dynamic optimization problem

$$\max \int_0^\infty e^{-rt} u((x(t))dt \tag{9}$$

subject to

$$\dot{S}(t) = -x(t) \tag{10}$$

$$S(0) = S_0 \text{ historically given initial resource stock} \tag{11}$$

$$x(t) \geq 0 \text{ for all } t \tag{12}$$

$$S(t) \geq 0 \text{ for all } t \tag{13}$$

where $u(0) = 0$, $u' > 0$, $u'' < 0$ and $u'(0) = b$. The Hamiltonian in this case is

$$H(x, S, \lambda) = u(x) - \lambda x$$

and the conditions (5)-(8) are now

$$\frac{\partial H}{\partial x} = u'(x(t)) - \lambda(t) \leq 0 \text{ and } \left[u'(x(t)) - \lambda(t)\right] x(t) = 0 \tag{14}$$

$$\dot{\lambda}(t) = r\lambda(t) - \gamma(t) \tag{15}$$

$$\gamma(t) \geq 0 \text{ and } \gamma(t)S(t) = 0 \tag{16}$$

$$Lim_{t\to\infty} e^{-rt}\lambda(t)S(t) = 0 \quad (17)$$

As long as $S(t) > 0$ we have $\gamma(t) = 0$ implying from (15) that

$$\dot{\lambda}(t) = r\lambda(t) \text{ or } \lambda(t) = \lambda(0)e^{rt} \quad (18)$$

It follows from (14) and (18) that $\dot{x}(t) \leq 0$. For $x(t) > 0$ we have

$$u'(x(t) = \lambda(0)e^{rt}$$

giving a declining $x(t)$. At some time T , $\lambda(0)e^{rT} = b$, giving $x(T) = 0$ since $u'(0) = b$. The resource stock must reach 0 at T: $S(T) < 0$ would violate the condition $S(t) \geq 0$ for all t, while $S(T) > 0$ would violate the transversality condition (17).

To conclude: Optimal resource extraction declines gradually over time, making the marginal utility u' rise over time at the rate r. Extraction eventually reaches zero; this occurs simultaneously with the resource stock being completely depleted.

Example 2: Optimal climate policy for a given carbon budget

Same as above, but let $u(x)$ be benefit of emitting carbon (x), i.e. of using fossil fuels. Moreover, let S_0 be the total amount of carbon emissions in the future (from date $t = 0$) that are consistent with a political goal of total temperature increase (see Allen et. al, 2009). The model describes how u' must develop over time. Users of carbon set u' equal to the carbon tax. Hence we can conclude that the optimal carbon tax must rise at interest rate. Moreover, the *level* of this carbon tax is higher the lower is S_0, i.e. the lower temperature increase we accept.

Reference:

Allen et. al 2009:

Warming caused by cumulative carbon emissions towards the trillionth tonne

http://www.nature.com/nature/journal/v458/n7242/full/nature08019.html

Quotation:

the relationship between cumulative emissions and peak warming is remarkably insensitive to the emission pathway (timing of emissions or peak emission rate). Hence policy targets based on limiting cumulative emissions of carbon dioxide are likely to be more robust to scientific uncertainty than emission-rate or concentration targets. Total anthropogenic emissions of one trillion tonnes of carbon (3.67 trillion tonnes of CO2), about half of which has already been emitted since industrialization began, results in a most likely peak carbon-dioxide induced warming of 2^{o}C above pre-industrial temperatures, with a 5–95% confidence interval of 1.3–3.9^{0}C.

Example 3: Stock pollution with an environmental damage cost function.

Example 2 described a stock pollution problem. Instead of an absolute limit to the stock of pollution, assume now that there at time t is an environmental cost $D(S(t), t)$ of a stock $S(t)$ of the pollutant in the environment, and that $D_S > 0$ and $D_{SS} \geq 0$. Notice that the relationship between S and D may vary over time; if e.g. income growth implies an increased willingness to pay for avoiding environmental damage we will have $D_t > 0$. The sign and size of D_t is, however, of no importance for the derivations below. (We could also assume that u depends directly on t, this would not affect the analysis below.)

Assume that the development of the stock S depends on the flow x as follows:

$$\dot{S}(t) = x(t) - \delta S(t) \tag{19}$$

where $\delta \geq 0$. The constraints (10)-(12) remain valid, and it follows from (12) and (19) that $S(t)$ can never become negative, so we need not explicitly include the constraint (13).

The optimization problem is now

$$\max \int_0^\infty e^{-rt} \left[u((x(t)) - D(S(t), t) \right] dt \tag{20}$$

and the Hamiltonian is in this case

$$H(x,S,\lambda,t) = u(x) - D(S,t) + \lambda\left[x - \delta S\right]$$

The conditions (5)-(8) are now

$$\frac{\partial H}{\partial x} = u'(x(t)) + \lambda(t) \leq 0 \text{ and } \left[u'(x(t)) + \lambda(t)\right] x(t) = 0 \qquad (21)$$

$$\dot{\lambda}(t) = (r+\delta)\lambda(t) + D_S(S(t),t) \qquad (22)$$

$$Lim_{t\to\infty} e^{-rt}\lambda(t)S(t) = 0 \qquad (23)$$

Assume that the problem has properties implying $x(t) > 0$ for all t. Then $S(t)$ must also always be positive, so that the transversality condition implies $Lim_{t\to\infty} e^{-rt}\lambda(t) = 0$.

It is useful to define $q(t) = -\lambda(t)$. Since $x(t)$ is always positive we can rewrite (21) and (22) as

$$u'(x(t)) = q(t) \qquad (24)$$

$$\dot{q}(t) = (r+\delta)q(t) - D_S(S(t)) \qquad (25)$$

We may interpret $q(t)$ as the optimal emission tax. If we know how this tax evolves over time we can deduce from (24) how emissions $x(t)$ will evolve over time. To solve for $q(t)$ we first define

$$\mu(t) = e^{-(r+\delta)t} q(t) \qquad (26)$$

implying

$$\dot{\mu}(t) = -(r+\delta)e^{-(r+\delta)t}q(t) + e^{-(r+\delta)t}\dot{q}(t)$$

Inserting from (25) gives

$$\dot{\mu}(t) = -e^{-(r+\delta)t} D_S(S(t),t) \qquad (27)$$

By definition, we have for any $T > t$ that

$$\mu(T) - \mu(t) = \int_t^T \dot{\mu}(\tau) d\tau$$

Letting $T \to \infty$ and inserting (27) gives

$$\mu(t) = Lim_{T\to\infty}\mu(T) + \int_t^\infty e^{-(r+\delta)\tau} D_S(S(\tau), \tau) d\tau \qquad (28)$$

Since $Lim_{T\to\infty} e^{-rT}\lambda(T) = 0$, it follows that $Lim_{t\to\infty} e^{-rT}\left[-e^{(r+\delta)T}\mu(T)\right] = -Lim_{t\to\infty} e^{\delta T}\mu(T) = 0$, which can only hold if $Lim_{t\to\infty}\mu(T) = 0$. Hence, it follows from (26) and (28) that

$$q(t) = \int_t^\infty e^{-(r+\delta)(\tau - t)} D_S(S(\tau), \tau) d\tau \qquad (29)$$

This equation for the optimal emission tax has an obvious interpretation: The amount of 1 unit of emissions at time t remaining in the atmosphere at $\tau(> t)$ is $e^{-\delta(\tau-t)}$. To get from the additional stocks at τ to additional damages at τ we must multiply the additional stock at τ by the marginal damage at τ, which is $D_S(S(\tau), \tau)$, giving a damage equal to $e^{-\delta(\tau-t)} D_S(S(\tau), \tau)$ for 1 unit emissions at t. The total additional damage caused by 1 unit of emissions at time t is the discounted sum of additional damages at all dates from t to infinity caused by the additional stocks from t to infinity. The marginal damage of 1 additional unit of emissions at t is thus given by the expression above.