6. Applications of optimal control to natural resource problems
AGEC 642 - 2024

I. The model of Hotelling 1931

Hotelling's 1931 article, "The Economics of Exhaustible Resources" is a classic that provides very important intuition that applies not only to natural resources but any form of depletable asset. Hotelling does not use the methodology of optimal control (since it wasn't discovered yet), but this methodology is easily applicable to the problem.

A. The General Hotelling model

Hotelling considers the problem of a depletable resource (like oil or minerals) and how might it be optimally used over time. *What are the state and control variables of such a problem?*

Let x_t be the stock of the resource remaining at time t and let z_t be the rate at which the stock is being depleted. For simplicity, first assume that extraction costs are zero and that the market is perfectly competitive. In this case, the representative owner of the resource will receive $p_t z_t$ from the extraction of z_t in period t and this will be pure profit or, more accurately, *quasi-rents*.

Definitions (taken from http://www.bized.ac.uk/, which appears to no longer exist)

Economic rent: A surplus paid to any factor of production over its supply price. Economic rent is the difference between what a factor of production is earning (its return) and what it would need to be earning to keep it in its present use. It is, in other words, the amount a factor is earning over and above what it could be earning in its next best alternative use (its transfer earnings).

Quasi-rent: Short-term economic rent arising from a temporary inelasticity of supply.

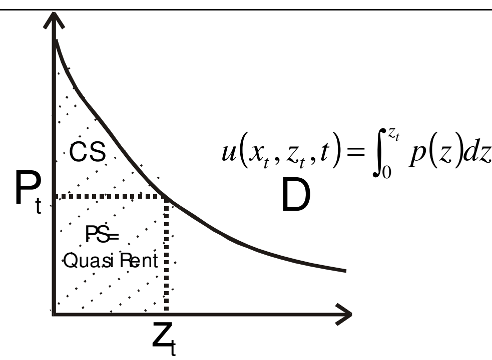

We consider the problem of a social planner who wants to maximize the present value of consumer surplus plus rents (= producer surplus in this case). CS + PS at any instant in time is equal to the area under the inverse demand curve, i.e., $u(x_t, z_t, t) = \int_0^{z_t} p(z)\,dz$, where $p(z)$ is the inverse demand curve for extractions of the resource. Notice that the marginal value of extraction up to z_t is obtained using Leibniz Rule: $p(z_t)$.

The problem is constrained by the fact that the original supply of the resource is finite, $x(t=0)=x_0$, and any extraction of the resource will reduce the available stock, $\dot{x}=-z$. We know that in any period $x_t \geq 0$ and simple intuition assures us that $x_T=0$. *Do you see why* $x_T=0$*?*

A formal statement of the planner's problem is:

$$\max_{z_t} \int_0^T e^{-rt} u(x_t, z_t, t)\, dt = \max_{z_t} \int_0^T e^{-rt} \left[\int_0^{z_t} p(z)\, dz \right] dt \text{ s.t.}$$

$\dot{x}_t = -z_t$

$x(t=0)=x_0$

$x_t \geq 0$

The current-value Hamiltonian of this problem is, therefore,

$H_c = u(\cdot) + \mu_t(-z_t)$

and the maximization criteria are:

1. $H_z=0$: $\quad u'(\cdot) - \mu_t = 0 \Rightarrow p(z_t) - \mu_t = 0$

2. $H_x = -\dot{\mu}_t + r\mu_t$: $\quad 0 = -\dot{\mu}_t + r\mu_t$

3. $H_\mu = \dot{x}_t$: $\quad \dot{x}_t = -z_t$

The transversality condition in this case is found by the terminal point condition,
4. $x_T=0$ since economic intuition tells us that it would be inefficient to leave a positive amount of the resource in the ground at the end of the planning horizon.

Looking at 1 and using the intuition developed by Dorfman, we see that at every moment in time the marginal benefit of extraction in t, $p(z_t)$, must be equal to the marginal cost in terms of foregone future net benefits, μ_t. Said another way, at every time t it is optimal to keep extracting up to the point where the marginal value of extraction, $p(z_t)$, is equal to the marginal value of leaving the resource in the ground, μ_t.

From 2 we see that since $\dot{\mu}_t = r\mu_t$, which means that $\dot{\mu}_t / \mu_t = r$; μ_t grows at the rate r.

This exponential growth in μ_t holds true in any discounted dynamic optimization problem in which neither the benefit function nor the state equation depend on the state variable. To see this, consider a Hamiltonian that takes the form $$H_c = f(z_t) + \mu g(z_t).$$ In this case, the FOC w.r.t. the state variable is $$\partial H_c / \partial x_t = 0 = r\mu_t - \dot{\mu}_t \Rightarrow \dot{\mu}_t / \mu_t = r.$$

For the Hotelling problem, solving the differential equation, $\dot{\mu}_t / \mu_t = r$, we find that $\mu_t = \mu_0 e^{rt}$ and, using 1, $p(z_t) = \mu_0 e^{rt}$.

<u>This is important.</u> It shows that the optimal price will grow at the discount rate, and this is true regardless of the demand function (as long as we have an interior solution). [Note that in this example the marginal extraction cost is set at zero so that the price is equal to the marginal quasi-rents earned by the producer. More generally, the marginal quasi-rents would be equal to price minus marginal cost, and this would grow at the rate of interest.] Resource economists refer to a price path along these lines with terms like "Hotelling-type prices" or "follows the Hotelling rule."

Another thing that is interesting in this model is that the value of μ_t rises at the rate of interest is that it means that the present-value co-state variable, $\lambda_t = e^{-rt}\mu_t$, would be constant. That means that the marginal increment to the objective function (the whole integral) of a unit of the resource stock, when viewed from the perspective of t=0, is constant – the planner is completely indifferent between receiving a marginal unit of the resource at time 0 and the instant before T, as long as the change is known a t=0.

There's some nice intuition for the result for λ_t. Since we know that the benefits are a function of the stream of utility that can be obtained over the planning horizon, whether you get the marginal increment to x at the beginning or at the end, it has the same effect on the total amount of the resource that is available to use over the interval from 0 to T. Hence, λ_t should be constant over the entire planning horizon.

B. The Hotelling Model with a Specific Functional Form

If we want to proceed further to obtain a specific analytical result, it is necessary to define a particular functional form for our demand equation. Suppose that $p(z)=e^{-\gamma z}$ so that the inverse demand curve looks like the figure above.

Hence, from 1, H_z=0⇒ $e^{-\gamma z_t} = \mu_t$, or $e^{-\gamma z_t} = \mu_t$ so that,

$$-\gamma z_t = \ln(\mu_t)$$

or

5 $$z_t = -\frac{\ln(\mu_t)}{\gamma}$$

The last main step to solve the problem is to find the value of μ_t as a function of parameters of the model and to do this we must use the transversality condition.

Note that $x_t = x_0 + \int_{\tau=0}^{t} \dot{x}_\tau d\tau$. Hence, from our transversality condition, 4,

$$x_T = 0 \Rightarrow x_0 = -\int_{\tau=0}^{T} \dot{x}_\tau d\tau.$$

From 3 and 5, this can be rewritten $\int_0^T z_t dt = x_0$ or $\int_0^T \left(-\frac{\ln(\mu_t)}{\gamma}\right) dt = x_0$ and, using $\mu_t = \mu_0 e^{rt}$, $x_0 = \int_0^T \left(-\frac{\ln(\mu_0 e^{rt})}{\gamma}\right) dt = \int_0^T \left(-\frac{\ln(\mu_0)+rt}{\gamma}\right) dt$.

Evaluating the integral leads to

$$\frac{1}{\gamma}\left(-\ln(\mu_0)t - \frac{r}{2}t^2\right)_0^T = x_0$$

$$\left(-\ln(\mu_0) - \frac{r}{2}T\right)\frac{T}{\gamma} = x_0$$

$$-\ln(\mu_0) = \frac{\gamma}{T}x_0 + \frac{r}{2}T.$$

Hence, we now know that $\mu_0 = e^{\left(-\frac{\gamma}{T}x_0 - \frac{r}{2}T\right)}$ and $\mu_t = e^{rt}\mu_0 = e^{rt} \cdot e^{\left(-\frac{\gamma}{T}x_0 - \frac{r}{2}T\right)}$.

We can then solve explicitly for z_t by substituting for μ_t into 5, yielding

$$z_t = -\frac{\ln(\mu_t)}{\gamma}$$

$$z_t = -\frac{\ln(\mu_0)+rt}{\gamma} = -[\ln \mu_0 + rt]/\gamma = -\left[\left(-\frac{\gamma}{T}x_0 - \frac{r}{2}T\right) + rt\right]\Big/\gamma$$

6. $z_t = \frac{x_0}{T} + \frac{r}{2\gamma}T - \frac{r}{\gamma}t \quad = \quad \frac{x_0}{T} + \frac{r}{\gamma}\left(\frac{T}{2} - t\right)$

To verify that this is correct, check the integral of this, from 0 to T

$$\int_0^T z_t dt = \frac{x_0}{T}T + \frac{r}{2\gamma}T^2 - \frac{r}{2\gamma}T^2 = x_0.$$

Homework tip: Notice that at this point z_t, is a function of only known parameters: x_0, which we know because we can observe this before starting the problem, T, which is assumed to be known in advance, and γ, which is a known parameter. When solving an optimal control problem, if you still have unknown variables in your result, e.g., μ, then you are not finished. Frequently, you need to use the transversality condition at this point

Looking at 6, we see that the rate of consumption at any point in time is determined by two parts: a constant portion of the total stock, $\frac{x_0}{T}$, plus a portion that declines linearly over time $\frac{r}{\gamma}\left(\frac{T}{2} - t\right)$. This second portion is greater than zero until $t = \frac{T}{2}$, and is then less than zero for the remainder of the period.

We can now reconsider whether our assumption that this is a simple fixed end point problem with x_T=0 is correct. Note that, that $z_T < 0$ if

7. $T > \sqrt{2\gamma x_0 / r}$.

If this inequality is satisfied, i.e., if the time horizon is long enough, then along the optimal path that results from choices made following 6, z_t will become negative, implying that the resource stock is being rebuilt as we approach T. This means that x_t is negative over some range, violating the constraint $x_t \geq 0$.

Hence, if 7 holds, the solution violates the constraint, and we will need to re-solve the problem explicitly noting the constraint on x_t. We will evaluate such constrained problems in Lecture 13.

C. Some variations on the theme and other results

Hotelling's analysis certainly doesn't end here.

Q: Consider again the question, "What would happen if we used the present-value instead of the current-value Hamiltonian?"
A: As mentioned above, the present value co-state variable, λ_t, would be constant over time. What's the economic interpretation of λ?

Q: What if there are costs to extraction $c(z_t)$ so that the planner's problem is to maximize the area under the demand curve minus the area under the marginal cost curve?
A: First recognize that if we define $\tilde{u}(\cdot) = \int_0^{z_t} p(z) - c'(z, x_t)\, dz$, where c' is the marginal cost function, then the general results will be exactly the same as in the original case after substituting "marginal quasi rents" for "price". That is, along the optimal path the marginal surplus will rise at the rate of interest. Obviously getting a nice clean closed-for solution z^* will be more difficult the more complex and realistic you make $c(\cdot)$, but the economic intuition does not change. This economic principle is a central to a wide body of economic analysis.

Q: Would the social optimum be achieved in a competitive market?
A: First, assuming that both consumers and producers are interested in maximizing the present value of their respective welfare, then we've maximized total surplus, i.e., it is a Pareto Efficient outcome. So, we can then ask, *Do the assumptions of the 2nd Welfare Theorem hold? If they do, then what does that tell us about the social optimum?* If these hold, then for a Pareto efficient there exists a price vector for which any Pareto efficient allocation will be a competitive equilibrium. Finding the Pareto optimal allocation also gives a competitive equilibrium. Hence, our findings are not only normative, but more importantly, they're positive, i.e., a prediction of what choices would actually occur in a perfectly competitive economy.

Now, let's look at this question a little more intuitively. We know that one of the basic results is that the price (or marginal quasi rents) grows at the rate of interest along with the shadow value of the resource stock, μ_t? Is this likely to occur in a competitive economy as well? In the words of Hotelling, "it is a matter of indifference to the owner of a mine whether he receives for a unit of his product a price p_0 now or a price p_0e^{rt} after time t" (p. 140). That is, price takers will look at the future and decide to extract today, or a unit tomorrow at a higher price. The price must increase by at least the rate of interest in this simple model because, if not, the market would face a glut today. If the price rose faster than the rate of interest, then the owners would choose to extract none today. Assuming that the inverse-demand curve is downward sloping, supply and demand can be equal only if each individual is completely indifferent as to when he or she extracts which also explains the constancy of λ.

This also gets at an important difference between *profit* and *rents*. We all know that in a perfectly competitive economy with free entry, profits are pushed to zero -- so why do the holders of the resource still make money in this case? Because there is not free entry. The total resource endowment is fixed at x_0. An owner of a portion of that stock is able to make *resource rents* because he or she has access to a restricted profitable input. Further, the owner is able to exploit the tradeoffs between current and future use to make economic gains. This is what is meant by *Hotelling rents*.

II. Hartwick's model of national accounting and the general interpretation of the Hamiltonian

Hartwick (1990) has a very nice presentation of the Hamiltonian's intuitive appeal as a measure of welfare in a growth economy. The analogies to microeconomic problems will be considered at the end of this section, so read that first if you want to be convinced that this is relevant to you. Hartwick's 1990 paper builds on Weitzman (1976) and is a generalization of Hartwick's more often cited paper from 1977.

A. *The general case*

We'll first present the general case and then look at some of Hartwick's particulars. Consider the problem of optimal growth in an economy maximizing

$$\int_0^\infty U(C)e^{-\rho t}dt$$

subject to a state equation for a malleable capital stock, x_0, that can either be consumed or saved for next period

$$\dot{x}_0 = g_0(\mathbf{x},\mathbf{z}) - C$$

and n additional state equations for the n other assets in the economy (e.g., infrastructure, human capital, environmental quality, etc.).

$$\dot{x}_i = g_i(\mathbf{x},\mathbf{z}),\ i=1,\ldots,n.$$

> Please excuse the possibly confusing notation. Here the subscript is an index of the good and the time subscript is suppressed.

Let $\mathbf{z}$ represent a vector of control variables and C be the numeraire choice variable (think consumption). The vector of state variables is denoted $\mathbf{x}$. The variables x_0 and C differ from the other state and choice variables because they enter into the equation in a strictly linear way and are not included in any of the state equations other than that for x_0.

The general current-value Hamiltonian of this optimization problem is

$$H_c = U\left(C\right) + \mu_0\left(g_0\left(\mathbf{x},\mathbf{z}\right) - C\right) + \sum_{j=1}^{n} \mu_j g_j\left(\mathbf{x},\mathbf{z}\right).$$[1]

This is our first exposure to the problem of optimal control with multiple state and control variables, but the maximization conditions are the simple analogues of the single variable case:

$$\frac{\partial H}{\partial C} = \frac{\partial H}{\partial z_i} = 0 \text{ for all } i$$ [or in general, maximize H with respect C and all the z_i's]

$$\frac{\partial H}{\partial x_j} = \rho\mu_j - \dot{\mu}_j \text{ for all } j$$

$$\frac{\partial H}{\partial \mu_j} = \dot{x}_j \text{ for all } j$$

Given the specification of utility, $\frac{\partial H}{\partial C} = U' - \mu_0 = 0 \Rightarrow \mu_0 = U'$.

(Remember, μ_0 is the co-state variable on the numeraire good, <u>not</u> the co-state variable at t=0.)

Similar to the approach used by Dorfman, Hartwick uses a linear approximation of current utility, $U(C) \approx U' \cdot C$, and, if we measure consumption in terms of dollars, U' is the marginal utility of income. Using this approximation, the Hamiltonian can be rewritten

$$H = U' \cdot C + \mu_0\left(g_0\left(\mathbf{x},\mathbf{z}\right) - C\right) + \sum_{j=1}^{n} \mu_j g_j\left(\mathbf{x},\mathbf{z}\right).$$

Dividing both sides by the marginal utility of consumption and remembering that $\mu_0 = U'$, he obtains the relationship

$$\frac{H}{U'} = C + \dot{x}_0 + \sum_{j=1}^{n} \frac{\mu_j}{\mu_0} \dot{x}_j$$

If you look at the RHS of this equation, you will see that this is equivalent to net national product (NNP) in a closed economy without government. NNP is equal to the value of goods and services (C) plus the net change in the value of the assets of the economy, $\left(\dot{x}_0 + \sum_{j=1}^{n} \frac{\mu_j}{\mu_0} \dot{x}_j\right)$.

The first lesson from this model is a general one and, as we will discuss below, it carries over quite nicely to microeconomic problems: maximizing the current-value Hamiltonian is equivalent to maximizing NNP. So, a policy that seeks to maximize NNP is consistent with the social planner's objective of maximizing the present value of utility. The

---

[1] Again, to write more concisely, from this point further we will write H instead of H_c, which we typically use in these notes.

current-value Hamiltonian, therefore, is a measure of the benefits to the decision maker that correctly takes into account both current benefits and costs, and future consequences.

Second, Hartwick's model can provide us with formal guidance on the appropriate shadow prices to be used when measuring changes in an economy's assets. Specifically, if the j^{th} asset changes in quantity, the monetary units that should be used to account for this change is $\frac{\mu_j}{\mu_0}$. In the next section couple of sections, we consider cases where this framework can be helpful in identifying the right prices for resource stock.

B. The case of a non-renewable resource

Consider an economy in which there are three state variables.

- First, there's the fungible capital stock, x_0 which we will now call *K*.
- Second, there's a nonrenewable resource or mine, *S* which falls as the resource is extracted, *R*, and grows when there are discoveries, *D*. Extractions, *R* are used in the production function $F(\cdot)$ but there is a cost to extraction, $f(R,S)$.
- Discovery costs rise over time as a function of cumulative discoveries so that the marginal cost of finding more of the resource increases over time. The total cost of discovery in a period is $v(D)$, linearly approximated as $v_D \cdot D$ with v_D changing over time.[2]
- Hartwick also includes labor, *L*. However, since the economy is always assumed to be at full employment and the growth rate of labor is exogenous, labor can be treated as an intermediate variable and can, therefore, be largely ignored.

The three state equations are, therefore,

Capital stock: $\dot{K} = F(K,L,R) - C - f(R,S) - v_D D$

Resource Stock: $\dot{S} = -R + D$

Discovery Cost: $\dot{v}_D = g(D)$

and the resulting current value Hamiltonian is

$$H = U(C) + \mu_K\left[F(K,L,R) - C - f(R,S) - v_D D\right] + \mu_S\left[-R + D\right] + \mu_D g(D)$$

The FOCs w.r.t. the choice variables are:

$H_C=0$: $U' = \mu_K$

$H_R=0$: $\mu_K\left[F_R - f_R\right] - \mu_S = 0$.

$H_D=0$: $-\mu_K v_D + \mu_S + \mu_D g' = 0$.

A linear approximation of the current-value Hamiltonian can be written

$$H = U'C + \mu_K\dot{K} + \mu_S\left[-R + D\right] + \mu_D g'D$$

Dividing by $U'=\mu_k$, we get

[2] This is a refinement of the specification in Hartwick (1990) as proposed Hamilton (1994).

$$\frac{H}{U'} = C + \dot{K} - \frac{\mu_S}{\mu_K} R + \frac{\mu_S}{\mu_K} D + \frac{\mu_D}{\mu_K} g'D$$

Using the H_R and H_D conditions, it follows that $\frac{\mu_S}{\mu_K} = \left[F_R - f_R\right]$ and

$$\mu_D = \frac{\mu_K v_D}{g'} - \frac{\mu_S}{g'} \text{ or } \mu_D = \frac{\mu_K v_D}{g'} - \frac{\mu_K\left[F_R - f_R\right]}{g'}$$

Hence, the linear approximation of the Hamiltonian can be rewritten

$$\frac{H}{U'} = C + \dot{K} - \frac{\mu_K\left[F_R - f_R\right]}{\mu_K} R + \frac{\mu_K\left[F_R - f_R\right]}{\mu_K} D + \left(\frac{\mu_K v_D}{\mu_K g'} - \frac{\mu_K\left[F_R - f_R\right]}{\mu_K g'}\right) g'D,$$

which, cancelling terms, can be simplified to

8 $$\frac{H}{U'} = C + \dot{K} - \left[F_R - f_R\right] R + v_D D.$$

We know that in a competitive economy, the price paid for the resource would equal F_R (resources are paid their marginal value product). Hence, to arrive at NNP current 'Hotelling Rents' from extractions, namely $\left[F_R - f_R\right]R$, should be netted out of GNP, and discoveries, priced at the marginal cost of discovery, should be added back in.[3]

So, we see that there are appropriate prices that could be used to adjust NNP to take into account resource extraction and discoveries. Is this common practice in national accounting? No! The depreciation of natural resource assets is ignored in the system of national accounts. This results in a misrepresentation of national welfare. One reason for this is the ability to actually implement the necessary accounting practice. Hartwick elaborates, "The principal problem of implementing the accounting rule above is in obtaining *marginal* extraction costs for minerals extracted."

Note that equation 8 can be rewritten by expressing $\dot{K}$ more completely:

$$\frac{H}{U'} = C + \left[F\left(K,L,R\right) - C - f\left(R,S\right) - v_D D\right] - \left[F_R - f_R\right] R + v_D D$$

9 $$\frac{H}{U'} = \left[F\left(K,L,R\right) - f\left(R,S\right) - v_D D\right] - \left[F_R - f_R\right] R + v_D D.$$

Notice that consumption, C, does not appear in 9. This makes sense. Consumption does not determine the value of the economy; that is determined by net output in the economy, $F\left(K,L,R\right) - C - f\left(R,S\right) - v_D D$.

Equation 9 can be simplified further:

$$\frac{H}{U'} = \left[F\left(K,L,R\right) - f\left(R,S\right)\right] - \left[F_R - f_R\right] R.$$

Discoveries, $v_D D$, cancel out because they count as an increase in capital stock in 8, but also decrease the value of net production so interestingly discoveries cancel out since

[3] This result differs from that presented in Hamilton (1994). I have not attempted to unravel where the difference comes from. Help is welcome to improve these notes.

their marginal value to our wealth is exactly equal to the marginal cost to current consumption.

C. Implications beyond the realm of national income accounting

If you're not interested in the national income accounts or environmental and natural resource economics, the above discussion may seem academic. However, clearly, the correct measurement of income is not an academic pursuit limited to the national income accounts.

Paraphrasing Hicks' (1939, *Value and Capital*), **income is defined as the maximum amount that an individual can consume in a week without diminishing his or her ability to consume next week.** Clearly, just as for a national account, farmers and managers also need to be aware of the distinction between investment, capital consumption, and true income. Hartwick's Hamiltonian formulation of NNP, therefore, with its useful presentation of the correct prices for use in the calculation of income, might readily be applied to a host of microeconomic problems of concern to applied economists.

III. References

Hartwick, John M. 1977. Intergenerational Equity and the Investing of Rents from Exhaustible Resources. *American Economic Review* 67(5):972-74.

Hartwick, John M. "Natural Resources, National Accounting and Economic Depreciation." *Journal of Public Economics* 43(December 1990):291-304.

Hamilton, Kirk. 1994. Green Adjustments to GDP. *Resources Policy* 20(3):155-68.

Hicks, John Richard. *Value and capital; an inquiry into some fundamental principles of economic theory*. Oxford: The Clarendon press, 1939.

Hotelling, Harold. 1931. The Economics of Exhaustible Resources. *The Journal of Political Economy* 39(2):137-175.

Weitzman, Martin L. "On the Welfare Significance of National Product in a Dynamic Economy." *Quarterly Journal of Economics* 90(Feb 1976):156-62.

IV. Readings for next class

For fun, read Dreyfus, S. 2002. Richard Bellman on the Birth of Dynamic Programming. Operations Research 50(1):48-51. (link to on-line version from notes page). This is just light reading that will give you a glimpse of how the principles of DP arose.

You should also watch the video that works through the simple example in the notes.