13. Optimal control with constraints and MRAP problems

AGEC 642 - 2024

We now return to an optimal control approach to dynamic optimization. This means that our problem will be characterized by continuous time and will be deterministic.

It is usually the case that we are not *Free to Choose*.[1] The choice set faced by decision makers is almost always constrained in some way and the nature of the constraint frequently changes over time. For example, a binding budget constraint or production function might determine the options that are available to the decision maker at any point in time. When this is true we will need to reformulate the simple Hamiltonian problem to take account of the constraints. Fortunately, in many cases, economic intuition will tell us that the constraint will not bind (except, for example, at $t=T$), in which case our life is much simplified. We consider here cases where we're not so lucky, where the constraints cannot be ruled out ex ante.

We will assume throughout that a feasible solution exists to the problem. Obviously, this is something that needs to be confirmed before proceeding to waste a lot of time trying to solve an infeasible problem.

In this lecture, we cover constrained optimal control problems rather quickly looking at the important conceptual issues. For technical details, I refer you to Kamien & Schwartz, which has chapters on constrained optimal control problems. We then go on to consider a class of problems where the constraints play a particularly central role in the solution.

I. A refresher on constrained optimization

Before we consider the economics of constrained optimal control problems, let's review the economics of inequality-constrained optimization problems and the Saddle Point Theorem.

Consider a simple inequality constrained optimization problem:

$$\max_z u(z) \quad \text{s.t. } h(z) \geq c$$

and, for better intuition, assume that $h(z)$ is monotonically decreasing with $\lim_{z\to-\infty} h(z) = +\infty$ and $\lim_{z\to+\infty} h(z) = -\infty$.

We know that this can be solved using a Lagrangian, but why? And what is the correct specification of the Lagrangian? For the Lagrange multiplier to have important economic meaning, it is critical that the constraint be specified correctly. The correct answer is

$$L(z,\lambda) = u(z) + \lambda\left(h(z) - c\right) \qquad \text{with } \lambda \geq 0.$$

A feature of the Lagrangian that many students are unaware of is that if z^* and λ^* are the optimal values of z and λ, then they are the solution to the optimization problem

$$L(z^*,\lambda^*) = \max_z \min_\lambda u(z) + \lambda\left(h(z) - c\right).$$

[1] This is an obtuse reference to the first popular press book on economics I ever read, *Free to Choose* by Milton and Rose Friedman.

This means that

1. $L(z^*,\lambda) \ge L(z^*,\lambda^*) \ge L(z,\lambda^*)$ for all z and $\lambda \ge 0$.

This is the Karush–Kuhn–Tucker or Saddle Point Theorem and, at the optimum, $\lambda^*\left(h(z^*)-c\right)=0$.[2]

Consider what the inequalities in 1 mean. Suppose that z^* and λ^* are optimal. Furthermore, suppose that at $h(z^*)-c>0$, i.e., the constraint does not bind. In this case, the second inequality in 1, $u(z^*)+\lambda^*\left(h(z^*)-c\right)>u(z)+\lambda^*\left(h(z)-c\right)$ for all z, is intuitive – if z^* is optimal, then this inequality will hold for all alternative values of z if λ^*=0. If $\lambda^* \neq 0$, will be values of z then such that $L(z^*,\lambda^*)<L(z,\lambda^*)$, which would be a violation of 1. Hence, we know that λ^*=0.

Now consider the first inequality,
$u(z^*)+\lambda\left(h(z^*)-c\right)\ge u(z^*)+\lambda^*\left(h(z^*)-c\right)$ for all $\lambda \ge 0$.
Since $u(z^*)$ appears on both sides, this can be simplified to
$\lambda\left(h(z^*)-c\right)>\lambda^*\left(h(z^*)-c\right)$ for all $\lambda \ge 0$.
Since $h(z^*)-c>0$ by assumption, the left-hand side is linearly increasing in λ and the only value λ^* such that this inequality will hold for all $\lambda \ge 0$ is λ^*=0. This would not be true if we had mistakenly reversed the constraint in the Lagrangian, i.e., if we had mistakenly written $\left(c-h(z^*)\right)$ instead of $\left(h(z^*)-c\right)$.

Now consider a case in which the constraint binds, i.e., $h(z^*)=c$. In this case we know that $\lambda^* \neq 0$. Why? And why is $\lambda^*>0$? Again, let's look at the second inequality:
$u(z^*)+\lambda^*\left(h(z^*)-c\right)\ge u(z)+\lambda^*\left(h(z)-c\right)$ for all z.
Since the constraint binds and because we've assumed that $h(\cdot)$ is downward sloping, it means that there is a value $z'>z^*$ such that $u(z')>u(z^*)$, but $h(z')<c$. In order for the constraint to bind we know that λ^* must be positive and big enough so that when z is pushed away from z^* in the direction of z', the gain in utility is less than $\lambda^*\left(h(z)-c\right)$ so that inequality no longer holds. Hence, at the optimum, the marginal value of increasing z, $\partial u(z)/\partial z$, must be less than or equal to the marginal change in the second term, $\lambda^* \partial h(z)/\partial z$. This helps us understand why λ^* is the shadow price.

[2] We will assume that certain regularity conditions hold, some of which are required for the Saddle Point Theorem to hold (i.e., that it is a convex programming problem satisfying the Slater condition) and others for convenience (e.g., continuity and twice differentiability).

The first inequality for the case is easier since $h(z^*)-c=0$, so

$u(z^*)+\lambda(h(z^*)-c)\geq u(z^*)+\lambda^*(h(z^*)-c)$ reduces to $u(z^*)\geq u(z^*)$, which holds trivially.

This discussion should help refresh your memory about why the Lagrangian can be used to solve constrained optimization problems, why the Lagrange multiplier has strong economic meaning, and can help you correctly specify your Lagrangians to ensure that the multipliers have the correct economic meaning.

For maximization problems with a constraint that can be written $g(z)\geq k$, you always want to specify your Lagrangian as $L=f(z)+\lambda(g(z)-k)$.

II. Optimal control with intratemporal equality constraints

A. Theory

Consider a simple dynamic optimization problem

$$\max_z \int_0^T e^{-rt}u(z_t,x_t,t)\,dt \quad \text{s.t.}$$

$$\dot{x}_t = g(z_t,x_t,t)$$

$$h(z_t,x_t,t)=c$$

$$x(0)=x_0$$

In this case we cannot use the Hamiltonian alone, because this would not take account of the constraint, $h(z,x,t)=c$. Rather, we need to maximize the Hamiltonian subject to a constraint, so we use a Lagrangian[3] in which H_c is the objective function, i.e.,

$$L=H_c-\phi_t\left(h(z_t,x_t,t)-c\right)$$

$$=u(z_t,x_t,t)+\mu_t g(z_t,x_t,t)-\phi_t\left(h(z_t,x_t,t)-c\right).$$

Equivalently, you can think about embedding a Lagrangian, within a Hamiltonian, i.e., $H_c=u(z_t,x_t,t)-\phi_t\left(h(z_t,x_t,t)-c\right)+\mu_t g(z_t,x_t,t)$. We'll use the first notation here.

Assuming that everything is continuously differentiable and that concavity assumptions hold, the FOC's of this problem, then, are:

2. $\dfrac{\partial L}{\partial z_t}=0$

3. $\dfrac{\partial L}{\partial x_t}=r\mu_t-\dot{\mu}_t$

and, of course, the constraints must be satisfied:

[3] This Lagrangian is given a variety of names in the literature. Some call it an augmented Hamiltonian, some a Lagrangian, some just a Hamiltonian. As long as you are explicit about what you are talking about, you can pretty much use any of these terms.

$$\frac{\partial L}{\partial \mu_t} = \dot{x}_t$$

$$\frac{\partial L}{\partial \phi_t} = c - h(z_t, x_t, t) = 0.$$

Let's look at these in more detail. The FOC w.r.t. z is

2'. $\frac{\partial L}{\partial z_t} = \frac{\partial u}{\partial z_t} + \mu_t \frac{\partial g}{\partial z_t} - \phi_t \frac{\partial h}{\partial z_t} = 0$,

which can be rewritten

2". $\frac{\partial u}{\partial z_t} - \phi_t \frac{\partial h}{\partial z_t} = -\mu_t \frac{\partial g}{\partial z_t}$.

As Dorfman showed us, the FOC w.r.t. the control variable tells us that at the optimum we balance off the marginal current benefit and marginal future costs. In this case the RHS is the cost to future benefits of a marginal increase in z. The LHS, therefore, must indicate the benefit to current utility from marginal increments to z. If $\partial u/\partial z$>RHS, then this implies that there is a cost to the constraint and $\phi_t \,{}^{\partial h}\!/_{\partial z_t}$ is the cost to current utility of the intratemporal constraint, h. If $h(\cdot)$ were marginally relaxed, then z_t could be changed to push it closer to balancing with the contribution of a marginal unit of z_t in the future.

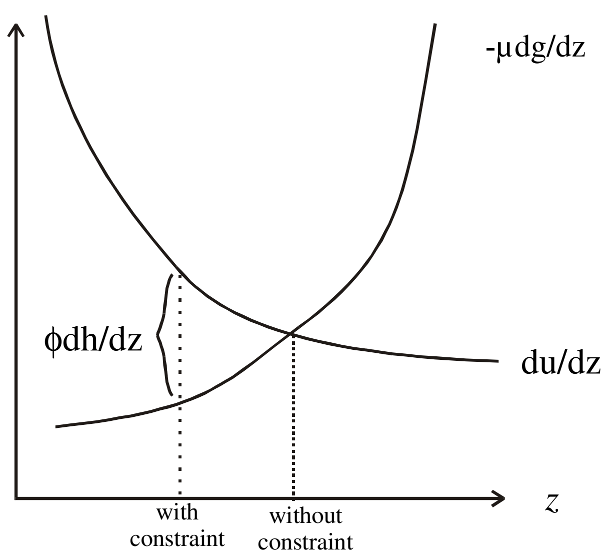

In principle, the problem can then be solved based on these equations. It is important to note that ϕ_t will typically change over time. *What is the economic significance of ϕ_t?*

B. Optimal control with multiple constraints

The extension to the case of multiple equality constraints, is easy; with n constraints the Lagrangian will take the form

$$L = u(z,x,t) + \lambda g(z,x,t) - \sum_{i=1}^{n} \phi_i \left(h_i(z,x,t) - c_i\right).$$

Obviously, there may not be a feasible solution unless some of the constraints do not bind or are redundant, especially if *n* is greater than the cardinality of *z*.

III. Optimal control with inequality constraints

A. Theory

Suppose now that the problem we face is one in which we have inequality constraints, $h_i(t, x, z) \geq c_i$, with $i=1,\ldots, n$, for *n* constraints and *x* and *z* are assumed to be vectors of the state and control variables respectively. For each $x_j \in x$, the state equation takes the form $\dot{x}_j = g_j(t, x, z)$.

As with standard constrained optimization problems, the Kuhn-Tucker conditions will yield a global maximum if any one of the Arrow-Hurwicz-Uzawa constraint qualifications is met (see Chiang p. 278). The way this is typically satisfied in most economic problems is for the h_i to be concave or linear in the control variables.

Assuming that the regularity conditions are satisfied, the optimum will be found using the Lagrangian specification in which a Hamiltonian, which takes the form

$$H(t,x,z,\lambda) = u(t,x,z) + \sum_{j=1}^{m} \lambda_{tj} g_j(t,x,z),$$

is embedded into the Lagrangian with the constraints,

$$L = H(t,x,z,\lambda) + \sum_{i=1}^{n} \phi_{ti}\left(h_i(t,x,z) - c_i\right)$$

$$L = u(t,x,z) + \sum_{j=1}^{m} \lambda_{tj} g_j(t,x,z) + \sum_{i=1}^{n} \phi_{ti}\left(h_i(t,x,z) - c_i\right).$$

The FOC's for this problem are:

$$\frac{\partial L}{\partial z_k} = 0 \Rightarrow \frac{\partial u}{\partial z_k} + \sum_{j=1}^{n} \lambda_{tj} \frac{\partial g_j}{\partial z_k} + \sum_{i=1}^{m} \phi_{ti} \frac{\partial h_i}{\partial z_k} = 0 \quad \text{for all } z_k \in z$$

$$\frac{\partial L}{\partial x_j} = -\dot{\lambda}_{tj} \text{ for all } j$$

$$\frac{\partial L}{\partial \lambda_j} = \dot{x}_j$$

and, for the constraints

$$\frac{\partial L}{\partial \phi_i} \geq 0 \Rightarrow h_i(x_t, z_t) \geq c_{ti}$$

with the complementary slackness conditions:

$$\phi_i \geq 0,\ h_i(x_t, z_t) \geq c_{ti} \text{ and } \phi_i \frac{\partial L}{\partial \phi_i} = 0 \text{ for all } i.$$

As with all such problems, the appropriate transversality conditions must be used and, if you use a current-value Hamiltonian, the necessary adjustments must be made. In the current value specification, the interpretation of both the co-state variable and the shadow price on the intratemporal constraint would in current-value terms, which is much more intuitive.

B. Example: Hotelling's optimal extraction problem

We return to Hotelling's problem from Lecture 6. The planner's problem is to maximize

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

$$\dot{x} = -z$$

$$x(0) = x_0,\ \ x_t \geq 0.$$

Economic intuition tells us that x_T=0. However, we found in lecture 6 that it is possible to find a solution in which x_t becomes negative and then, z_t is negative for a period to restore x_t so that x_T=0. However, by constraining x_T=0 and z_t≥0 for all t, we can indirectly ensure that x_t≥0 for all t. Hence, we replace the constraint x_t≥0 with z_t≥0.

The associated Lagrangian using a current-value Hamiltonian would be:

$$L = u(z_t) + \mu_t(-z_t) + \phi_t(z_t - 0).$$

We cover constraints on the state variable below

The associated maximization criteria are:

4. L_z=0: $\quad \dfrac{\partial u(z_t)}{\partial z_t} - \mu_t + \phi_t = 0$

5. $L_x = r\mu_t - \dot{\mu}_t$: $\quad 0 = r\mu_t - \dot{\mu}_t$

6. $L_\mu = \dot{x}$: $\quad \dot{x}_t = -z_t$

7. $\quad z_t \geq 0$

8. $\quad \phi_t \geq 0 \quad$ Kuhn Tucker Conditions

9. $\quad \phi_t \cdot z_t = 0$

The transversality condition is x_T=0.

Equation 5 can be rewritten, $\dot{\mu}_t/\mu_t = r$, i.e., μ_t grows at the rate r as we found in Lecture 6. Furthermore, solving this differential equation, we can write $\mu_t = \mu_0 e^{rt}$.

Recall from Lecture 6 that $\dfrac{\partial u(z_t)}{\partial z_t} = p(z_t)$, so 4 can be rewritten $p(z_t) = \mu_t - \phi_t$.

Using the assumed functional form for inverse demand curve, $p(z_t) = e^{-\gamma z_t}$, we obtain $e^{-\gamma z_t} = \mu_t - \phi_t = \mu_0 e^{rt} - \phi_t$. Taking logs, we get $-\gamma z = \ln(\mu_t - \phi_t)$, or

10. $\quad z_t = -\dfrac{\ln(\mu_t - \phi_t)}{\gamma}$.

Using the complementary slackness conditions, we know that if z_t>0 then ϕ_t=0 and if z_t=0, ϕ_t>0. The path can, therefore, be broken into two parts, the first part from 0 to T_1 during which z_t>0, and the second part, from T_1 to T, where z=0 and ϕ>0.

From 0 to T_1, when ϕ_t=0,

$$z_t = -\frac{\ln(\mu_t - 0)}{\gamma} = -\frac{\ln(\mu_0 e^{rt})}{\gamma} = -\frac{\ln(\mu_0) + rt}{\gamma} > 0.$$

Notice that since μ_t is positive and growing over this range, z_t is falling in value. The constraint on z binds at T_1, so that $\ln(\mu_{T_1}) = 0$ or $\mu_{T_1} = 1 = \mu_0 e^{rT_1}$, from which we also know that $\mu_0 = e^{-rT_1}$.

Over the period from T_1 to T, the constraint binds, z_t=0 and ϕ_t>0, so that $0 = -\frac{\ln(\mu_t - \phi_t)}{\gamma} \Rightarrow 0 = \ln(\mu_t - \phi_t)$ or $\mu_t - \phi_t = 1$. Using the fact that μ_t grows at the rate r, we can then write

11. $\phi_t = \mu_0 e^{rt} - 1$.

We know that since z_t=0 from T_1 onward, the transversality condition x_T=0 will only be satisfied if we exhaust the resource by T_1:

12. $\int_0^{T_1} z_t dt = x_0$ or $\int_0^{T_1}\left(-\frac{\ln(\mu_0) + rt}{\gamma}\right)dt = x_0$.

$$\int_0^{T_1}\left(-\frac{r}{\gamma}t\right)dt = -\frac{\ln(\mu_0)}{\gamma}T_t - \left[\frac{r}{2\gamma}t^2\right]_0^{T_1} = -\frac{\ln(\mu_0)}{\gamma}T_t - \frac{r}{2\gamma}T_1^2 = x_0$$

so that $-\frac{\gamma}{T_1}\left(\frac{r}{2\gamma}T_1^2 + x_0\right) = \ln(\mu_0)$ or $\mu_0 = \exp\left[-\frac{\gamma}{T_1}\left(\frac{r}{2\gamma}T_1^2 + x_0\right)\right]$. Finally, since $\mu_0 e^{rT_1} = 1$ we can write $\mu_0 = e^{-rT_1}$ so

$$\exp[-rT_1] = \exp\left[-\frac{\gamma}{T_1}\left(\frac{r}{2\gamma}T_1^2 + x_0\right)\right]$$

$$T_1^2 = \frac{\gamma}{r}\left(\frac{r}{2\gamma}T_1^2 + x_0\right)$$

$$T_1^2\left(1 - \frac{\gamma}{r}\frac{r}{2\gamma}\right) = \left(\frac{\gamma}{r}x_0\right)$$

$$T_1^2 = 2\left(\frac{\gamma}{r}x_0\right)$$

or

13. $T_1 = \sqrt{\frac{2\gamma}{r}x_0}$.

Hence, the resource will be exhausted by T_1 and the constraint on z is binding from T_1 onwards.

Now consider what happens for $t>T_1$. Recall that $\phi_t = \mu_0 e^{rt} - 1$, so ϕ_t grows geometrically from T_1 on so that $\mu_t - \phi_t$ always equals 1 and the meaning of ϕ_t is directly related to μ_t, i.e., the marginal value of the state variable if the constraint on x were relaxed.

IV. Constraints on the state space

A. Theory

Suppose now that we have constraints on the state variables which define a feasible range. This is common in economic problems. You may, for example, have limited storage space so that you cannot accumulate your inventory forever. Or, if you were dealing with a biological problem, you might be constrained to keep your stock of a species above a lower bound where reproduction begins to fail, and an upper bound where epidemics are common.

The approach to such problems is like that of the control problems.
Suppose we have an objective function

$\max \int_0^T u(t,x,z)dt$ s.t.

$\dot{x} = g(t,x,z),\ x(0) = x_0$ and

$h(t,x) \geq 0.$

The augmented Hamiltonian for this problem is

$$L = u(t,x,z) + \lambda g(t,x,z) + \phi h(t,x)$$

and the necessary conditions for optimality include the constraints plus

$$\frac{\partial L}{\partial z} = 0$$

$$\dot{\lambda} = -\frac{\partial L}{\partial x}$$

$\phi \geq 0$ and $\phi h = 0$

and the transversality condition.

Solving problems like this by hand, however, can be quite difficult, even for very simple problems. (See K&S p.232 if you want to convince yourself). (An alternative approach presented in Chiang (p. 300) is often easier and we follow this approach below). For much applied analysis, however, there may be no alternative to setting a computer to the problem to find a numerical solution.

V. Bang-bang OC problems

We now consider problems for which the optimal path does not involve a smooth approach to the steady state or gradual changes over time. Two important classes of such problems are known as "bang-bang" problems and most rapid approach problems. In such problems, the constraints play a central role in the solution.

A. Bang-bang example #1: A state variable constraint

Consider the following problem in which we seek to maximize discounted linear utility obtained from a nonrenewable stock (sometimes referred to as a cake-eating problem):

$$\max_z \int_0^T e^{-rt} z_t dt \quad \text{s.t.}$$

$$\dot{x} = -z$$

$$x(t) \geq 0$$

$$x(0) = x_0 > 0$$

What does intuition suggest about the solution to the problem? Will we want to consume the resource stock x gradually? Why or why not? Let's check our intuition.

Following the framework from above, we set up the Lagrangian by adding the constraint on the state variable to the Hamiltonian, i.e., $L=H+\phi$(constraint). Using the current-value specification, this gives us

$$L = z_t - \mu_t z_t + \phi_t x_t$$

The FOCs for the problem are:

(i) $\frac{\partial L}{\partial z} = 0$: $\quad 1 - \mu_t = 0$

(ii) $\frac{\partial L}{\partial x} = r\mu_t - \dot{\mu}_t$: $\quad \phi_t = r\mu_t - \dot{\mu}_t$

Because of the constraint, the complementary slackness condition must also hold:

(iii) $\phi_t x_t = 0$.

Equation i implies that μ_t=1. Since this holds no matter the value of t, we know that $\dot{\mu}_t = 0$ for all t. Conditions i and ii together indicate that

μ_t=1 and ϕ_t=r.

The second of these is most interesting. It shows us that ϕ_t, the Lagrange multiplier, is always positive. From the complementary slackness condition (iii), it follows that x_t must equal 0 <u>always</u>. But wait! We know x_0>0. However, at t=0, x_t is not variable – it is parametric to our problem. Since x_0 cannot be chosen, the requirement that x_t=0 only applies for t>0, i.e., at every instant except the immediate starting value.

So how big is z at zero? The first thought is that it must equal x_0 but this isn't quite right. To see this, suppose that we found that the constraint started to bind, not immediately, but after 10 seconds. To get the x to zero in 10 seconds, z per second would have to equal $x_0/10$. Now take the limit of this at the denominator goes to zero $\Rightarrow$ z goes to infinity. Hence, what happens is that for one instant there is a spike of z_t of infinite height and zero length that pushes x exactly to zero. This type of solution is known as a bang-bang problem because the state variable jumps discontinuously at a single point – BANG-BANG! Since, in the real world it's pretty difficult to push anything to infinity, we would typically interpret this solution as "consume it as fast as you can." This is formalized in the framework of most-rapid-approach path problems below.

B. Bang-Bang Example #2 (based on Kamien and Schwartz p. 205) A control variable constraint

Let x_t be a productive asset that generates output at the rate rx_t. This output can either be consumed or reinvested. The portion that is reinvested will be called z_t so $[1-z_t]$ is the portion that is consumed. We assume that the interest can be consumed, but the principal cannot be touched.[4] Our question is: What portion of the interest should be invested and what portion should be consumed over the interval $[0,T]$?

Formally, the problem is:

$$\max_z \int_0^T [1-z_t]rx_t dt \text{ s.t.}$$

$$\dot{x}_t = z_t r x_t,\ 0 \le z_t \le 1,\ x(0) = x_0$$

This time we have two constraints: $z_t \le 1$ and $z_t \ge 0$. Hence, our Lagrangian is

$$L = [1-z_t]rx_t + \lambda z_t r x_t + \phi_{1t}(1-z_t) + \phi_{2t} z_t,$$

where $[1-z_t]rx_t + \lambda z_t r x_t$ is the Hamiltonian part of the problem and the last two terms are the constraints.

The necessary conditions for an optimum are

14. $\dfrac{\partial L}{\partial z} = 0 \Leftrightarrow -rx_t + \lambda r x_t - \varphi_1 + \varphi_2 = 0$, and

15. $\dfrac{\partial L}{\partial x} = -\dot{\lambda}_t \Leftrightarrow -\dot{\lambda} = [1-z_t]r + \lambda z_t r$.

The transversality condition in this problem is $\lambda_T=0$ since x_T is unconstrained with the Kuhn-Tucker conditions,

KT_1: $\phi_1 \ge 0$ & $\phi_1(1-z_t)=0$, and

KT_2: $\phi_2 \ge 0$ & $\phi_2 z=0$.

From the KT_1, we know that if $\phi_1>0$, then the first constraint binds and $z_t=1$. Similarly, from KT_2, if $\phi_2>0$, then the second constraint binds and $z=0$. i.e.,

$\phi_1>0 \Rightarrow z = 1$ $\qquad$ $\phi_2>0 \Rightarrow z = 0.$

$\phi_1=0 \Leftarrow z < 1$ $\qquad$ $\phi_2=0 \Leftarrow z > 0.$

It is not possible for both ϕ_1 and ϕ_2 to be positive at the same time: if the z=1 constraint binds, then clearly the z=1 does not bind and vice versa.

The first FOC can be rewritten

$$(\lambda_t - 1)rx_t - \phi_1 + \phi_2 = 0.$$

[4] This problem is very similar to one looked at in Lecture 3. Comparing the two you'll see one key difference is that here utility is linear, while in lecture 3 utility was logarithmic.

We know that rx_t will always be positive since consumption of the capital stock is not allowed. Hence, we can see that three cases are possible:
1) if $\lambda=1 \Rightarrow \phi_1=0$ $\phi_2=0 \Rightarrow$ no constraint binds
2) if $\lambda>1 \Rightarrow \phi_1>0$ $\phi_2=0 \Rightarrow z_t=1$
3) if $\lambda<1 \Rightarrow \phi_1=0$ $\phi_2>0 \Rightarrow z_t=0$.

From the second FOC,
$\dot{\lambda} = -\{[1-z_t]r + \lambda_t z_t r\}$.
Since everything in the brackets is positive, the RHS of the equation is negative $\Rightarrow \lambda$ is always falling.

By the transversality condition, we know that eventually λ_t must reach 0 so that $\lambda_T=0$. Hence, eventually we'll reach case 3 where, $\lambda_t<1$ and $z_t=0$ and we consume all of our output. But when do we start consuming, right away or after x has grown for a while? We know from equation 2 that at $\lambda_t=1$ neither constraint binds.

- Suppose that at $t=n$ $\lambda_t=1$.
- For $t<n$ $\lambda_t>1$ and $z=1$.
- For $t>n$ $\lambda_t<1$ and $z=0$.

An important question then is when is n? We can figure this out by working backwards from $\lambda_T=0$. From the second FOC, we know that in the final period, (when $\lambda_t<1$) $z=0$, in which case
$\dot{\lambda} = -r$.
Solving this differential equation yields
$\lambda_t = -rt + A.$
Using the transversality condition,
$\lambda_T = -rT + A = 0$
$A = rT$
$\lambda_t = -rt + rT = r(T-t)$
Hence, substituting in $\lambda_n=1$ means $1 = -rn + rT = r(T-n)$. Simplifying,
$r(T-n) = 1$
$n = (rT-1)/r$
Hence, we find that the optimal strategy is to invest everything, $z_t=1$, from $t=0$ until $t = n = (rT-1)/r$. After $t=n$ consume all of the interest, $z_t=0$. If $(rT-1)/r < 0$ then it would be optimal to consume everything from the very outset.

For $(rT-1)/r>0$, we can graph the solution:

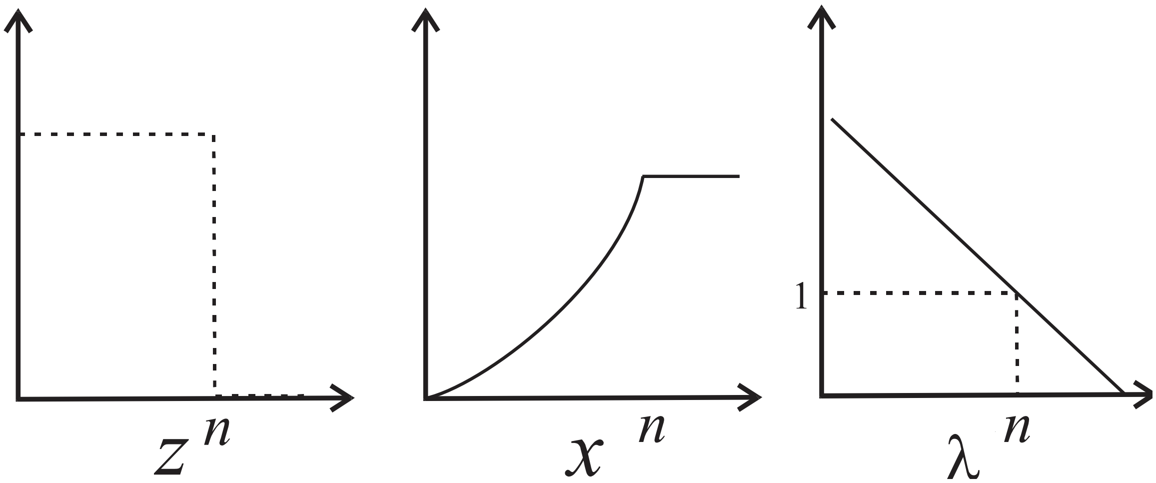

What would be the solution as $T\to\infty$? Does this make intuitive sense? What is it about the specification of the problem that makes it inconsistent with our economic intuition?

VI. Most Rapid Approach Path problems

Bang-bang problems fit into a general class of problems that are commonly found in economics: most-rapid-approach path problems (MRAP).[5] Here, the optimal policy is to get as quickly as possible to steady state where benefits are maximized. Consider the first example bang-bang example above. Wouldn't a solution in which we move toward the equilibrium as fast as possible rather than impossibly fast be more intuitively appealing?

A. MRAP example (Kamien & Schwartz p. 211)

A very simple firm generates output from its capital stock with the function $f(x_t)$ with the property that $\lim_{x\to 0} f'(x)=\infty$. The profit rate, therefore, is

$$\pi_t = p\cdot f(x_t) - c\cdot z_t$$

where x_t is the firm's capital stock, z_t is the rate at which it is investing per period, p and c are exogenously evolving unit price and unit cost respectively. The capital stock that starts with $x(0)=x_0$, depreciates at the rate b so that

$\dot{x}_t = z_t - bx_t$.

The firm's problem, therefore, is to maximize the present value of its profits,

$\int_0^\infty e^{-rt}\left[p\cdot f(x_t) - c\cdot z_t\right]dt$ subject to

$\dot{x}_t = z_t - bx_t$,

with three additional constraints:

i) $x(t)\geq 0$

ii) $z_t\geq 0$

iii) $p\cdot f(x_t) - c\cdot z_t \geq 0$

---

[5] Sometimes the term "bang-bang" is also used to describe MRAP problems.

Let's use economic intuition to help us decide if we need to explicitly include all the constraints in solving the problem.

- The constraint on x almost certainly does not need to be imposed because as long as f' gets big as $x\to 0$, the optimal solution will always avoid zero.
- The constraints on z, on the other hand might be relevant. But we'll start by assuming that neither constraint binds, and then see if we can figure out actual the solution based on the assumed interior solution or, if not, we'll need to use the Kuhn-Tucker specification. Note that if there does exist a steady state in x, then, as long as $b>0$, z must be greater than zero. Hence, we anticipate that much might be learned from the interior solution.
- Similarly, the profit constraint might also bind, but we would expect that in the long run, profits would be positive. So again, we start by solving for an interior solution, assuming $\pi>0$ where $\pi = p\cdot f\left(x_t\right)-c\cdot z_t$.

B. The interior solution

The current value Hamiltonian of the problem (assuming an interior solution w.r.t. z and x with $\pi>0$) is

$$H_c = p\cdot f\left(x_t\right)-c\cdot z_t+\mu_t\left(z_t-bx_t\right)$$

The necessary conditions for an interior solution are:

$$\frac{\partial H_c}{\partial z_t}=0 \qquad \Rightarrow \qquad -c+\mu_t=0$$

$$\frac{\partial H_c}{\partial x_t}=r\mu_t-\dot{\mu}_t \Rightarrow \qquad p\frac{\partial f\left(x_t\right)}{\partial x_t}-\mu_t b=r\mu_t-\dot{\mu}_t$$

Over any range where the constraints on z do not bind, therefore, we have

$c=\mu_t$

and, therefore, it must also hold that

$\dot{\mu}_t=\dot{c}=0$.

Substituting c for μ and rearranging, the second FOC becomes

16. $$p_t\frac{\partial f\left(x_t\right)}{\partial x_t}=\left(r+b\right)c,$$

which must hold over any interval where the constraints are not binding, i.e., $z>0$, and $p\cdot f\left(x_t\right)-c\cdot z_t>0$.

We see, therefore, that the optimum conditions tell us about the optimal level of x, say x^*. We can then use the state equation to find the value of z that maintains this relation.

Since c and p are constant, this means that the capital stock will be held at a constant level and 16 reduces to $\frac{pf'(x)}{r+b}=c$. This is known as the *modified golden rule*.

Let's think about this condition for a moment.

- In a static optimization problem, the optimal choice would be to choose x where the marginal product of increasing x is equal to the marginal cost, i.e., where $pf'=c$.
- In an infinite-horizon economy, if we could increase x at all points in time this would have a discounted present value of $\frac{pf'}{r}$. However, since the capital stock depreciates over time, this depreciation rate diminishes the present value of the gains that can be obtained from an increase in x today, hence the present value of the benefit of a marginal increase in x_t is $\frac{pf'}{r+b}$.

If p and c are not constant, but grow in a deterministic way (e.g., constant and equal to inflation) then we could de-trend the values and find a real steady state. If p and c both grow at a constant rate, say w, then there will be a unique and steady optimal value of x for all $z>0$.

Hence, our first-order conditions can be used to learn a lot about the nature of the situation when the constraints do not bind, i.e., when $z_t>0$ and $p\cdot f\left(x_t\right)>c\cdot z_t$. However, this is not the end of the story.

C. Corner solutions

All of the discussion above assumed that we are at an interior solution, where $0<z_t<p\cdot f\left(x_t\right)/c$. However, the interior solution only holds when the state variable x is at the point defined by equation 16; if the value of x is not at x^* at t=0, then it must be that we have a corner solution in which either z_t=0 or $p\cdot f\left(x_t\right)-c\cdot z_t=0$.

If $x_0>x^*$ then it will follow that z will equal zero until x_t depreciates to x^*. If $x_0< x^*$ then z will be as large as possible $\frac{p}{c}f\left(x_t\right)=z_t$ until x^* is reached.

Hence, economic intuition and a good understanding of the steady state can tell us where we want to get and how we're going to get there – in the most rapid approach possible.

D. Some theory and generalities regarding MRAP problems

A general statement of the conditions required for a MRAP result is presented by Wilen (1985, p. 64):

> Spence and Starrett show that for *any* problem whose augmented integrand (derived by substituting the dynamic constraint into the original integrand) can be written as
>
> $$\Pi_A\left(K,\dot{K}\right)=M\left(K\right)+N\left(K\right)\dot{K}$$
>
> the optimal solution reduces to one of simply reaching a steady state level $K=K^*$ as quickly as possible.

Where *K* is the state variable and by "integrand" they mean the objective function, profits in the case considered here.

How does this rule apply here? The integrand is $p_t f(x_t) - c_t z_t$. Using the state equation $bx_t + \dot{x}_t = z_t$, the integrand can be written
$p_t f(x_t) - c_t(bx_t + \dot{x}_t) = p_t f(x_t) - c_t bx_t - c_t \dot{x}_t$.
Converting this to the notation used by Wilen,
$M(K) = p_t f(x_t) - c_t bx_t$
and
$N(K)\dot{K} = c_t \dot{x}_t$.
Hence this problem fits into the general class of MRAP problems.

For a more intuitive understanding of why bang-bang and MRAP solutions arise, consider a general problem of the form

$$\max_{z_t} \int_0^T e^{-rt} f(x_t) z_t dt \qquad \text{s.t. } \dot{x}_t = g(x_t) z_t$$

so that both the benefit function and the state equation are linear in z_t. In this case, the Hamiltonian would be written

$$H_c = f(x_t) z_t + \mu_t \cdot g(x_t) z_t.$$

The optimization criterion remains: Maximize H_c with respect to z_t for all *t*. If $f(x_t) + \mu_t \cdot g(x_t) > 0$, then to maximize H_c we should set z_t at $+\infty$. If $f(x_t) + \mu_t \cdot g(x_t) < 0$, then z_t should be set at $-\infty$.[6] Hence, both the benefit function and the state equation are linear in z a bang-bang or MRAP solution will be obtained.

One lesson that can be obtained from this is that you need to be careful when specifying your model. While linear functions are nice to work with and frequently offer nice intuition, they frequently lead to corner solutions that are not intuitive, may not be easy to work with, and may lack the intuitive economic meaning that the model is set up to deliver.

VII. References

Chiang, Alpha C. 1991. *Elements of Dynamic Optimization*. McGraw Hill

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

Kamien, Morton I. and Schwartz, Nancy Lou. 1991. *Dynamic Optimization : The Calculus of Variations and Optimal Control in Economics and Management*. New York, N.Y. : Elsevier.

---

[6] I found this simple presentation in Rodriguez et al. (2011) very helpful, though I imagine that the presentation has been presented by others previously.

Rodriguez, Armando A. et al. 2011. "Confronting management challenges in highly uncertain natural resource systems: a robustness–vulnerability trade-off approach." *Environmental Modeling & Assessment* 16(1): 15-36.

Spence, Michael and David Starrett. 1975. Most Rapid Approach Paths in Accumulation Problems. *International Economic Review* 16(2):388-403.

Wilen, James E. 1985. Bioeconomics of Renewable Resource Use, In A.V. Kneese and J.L. Sweeney (eds.) *Handbook of Natural Resource and Energy Economics, vol. I.* New York: Elsevier Science Publishers B.V.