4. End points and transversality conditions

AGEC 637 - Summer 2009

In this lecture we consider a variety of alternative ending conditions for continuous-time dynamic optimization problems. For example, it might be that the state variable, x, must equal zero at the terminal time T, i.e., $x_T=0$, or it might be that it must be less than some function of t, $x_T \le \phi(T)$. We also consider problems where the ending time is flexible or $T\to\infty$. In the process, we will provide a more formal development of Pontryagin's maximum principle.

I. First, a word on salvage value

First, let's more completely introduce the notion of salvage value. Salvage value refers to the benefits (or costs) that occur after you exit the problem, say $S(x_T, T)$. For example, operating a car is certainly a dynamic problem and there is typically some value (perhaps negative) to your vehicle when you're finally finished with it. Similarly, farm production problems might be thought of as a dynamic optimization problem in which there are costs during the growing season, followed by a salvage value at harvest time. The general optimization problem with salvage value becomes

$$\max_z \int_0^T F(t,x,z)\,dt \quad +S(x_T,T) \quad \text{s.t.}$$
$$\dot{x}_t = f(t,x,z)$$
$$x_0 = x_0$$

Incorporating salvage value into a problem can change it in rather important ways. For the time being, however, we will assume that $S(\cdot) = 0$.

II. **Transversality conditions for a variety of ending points**

(Based on Chiang pp. 181-184)

A. Vertical or Free-endpoint problems

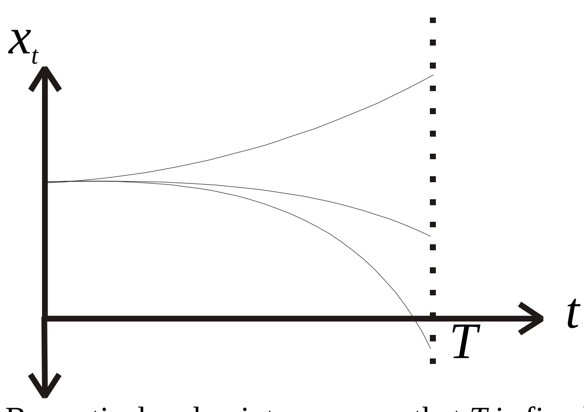

By vertical end point, we mean that T is fixed and x_T can take on any value. This would be appropriate if you are managing an asset or set of assets over a fixed horizon and it doesn't matter what condition the assets are in when it's all over. This case we have considered previously. When looked at from the perspective of the beginning of the planning horizon, the value that t takes on at T is free and, moreover, it has no effect on what happens in the future. So it is a fully free variable and we would maximize V over x_T. Hence, it follows that the shadow price of x_T must equal zero, giving us our transversality condition, $\lambda_T = 0$.

We will now confirm this intuition by deriving the transversality condition for this particular problem and at the same time giving a more formal presentation of Pontryagin's maximum principle.
The objective function is

$$V \equiv \int_0^T F(t,x,z) \quad dt$$

now, setting up an equation as a Lagrangian with the state-equation constraint, we have

$$L = \int_0^T \left[F(t,x,z) + \lambda_t \left(f(t,x,z) - \dot{x}_t \right) \right] dt .$$

We put the constraint inside the integral because it must hold at every point in time. Note that the shadow price variable, λ_t, is actually not a single variable, but is instead defined at every point in time in the interval 0 to T. Since the state equation must be satisfied at each point in time, at the optimum, it follows that $\lambda_t \left(f(t,x,z) - \dot{x}_t \right) = 0$ at each instant t, so that the value of L must equal the value of V. Hence, we might write instead

$$V = \int_0^T \left[F(t,x,z) + \lambda_t \left(f(t,x,z) - \dot{x}_t \right) \right] dt$$

or

$$V = \int_0^T \left[\underbrace{\left\{ F(t,x,z) + \lambda_t f(t,x,z) \right\}} - \lambda_t \dot{x}_t \right] dt$$

$$V = \int_0^T \left[\overbrace{H(t,x,z,\lambda)} - \lambda_t \dot{x}_t \right] dt$$

It will be useful to reformulate the last term, $\lambda_t \dot{x}_t$, by integrating by parts:

$$\int udv = vu - \int vdu$$

with $\lambda = u$ and $x = v$, so that $dv = \dot{x}$, we get

$$-\int_0^T \lambda_t \dot{x}_t dt = -\left[\lambda_t x_t \right]_0^T + \int_0^T \dot{\lambda}_t x_t dt$$

$$= \int_0^T \dot{\lambda}_t x_t dt + \lambda_0 x_0 - \lambda_T x_T$$

so, we can rewrite V as

1. $V = \int_0^T \left[H(t,x,z,\lambda) + \dot{\lambda}_t x_t\right]dt + \lambda_0 x_0 - \lambda_T x_T$

> **Derivation of the maximum conditions** (Based on Chiang chapter 7)
>
> From 1, we can easily derive the first two conditions of the maximum principal. Assuming an interior solution and twice-differentiability, a necessary condition for an optimum is that the first derivatives of choice variables are equal to zero.
>
> First consider our choice variable, z_t. At each point in time it must be that $\partial V/\partial z_t = 0$. This reduces to $\partial H/\partial z = 0$, which is the first of the conditions stated without proof in lecture 3.
>
> Next, for all t>0, x_t is also a choice variable in 1, so it must also hold that $\partial V/\partial x_t = 0$. This reduces to if $-H_x = \dot{\lambda}$, which is the second of the conditions stated in lecture 3.
>
> Finally, the FOC with respect to λ_t is more directly derived from the Lagrangian above. $\partial L/\partial \lambda_t = f(t,x,z) - \dot{x}_t$, so this implies that $\partial L/\partial \lambda_t = 0 \Rightarrow \dot{x}_t = f(t,x,z)$.

If the terminal condition is that x_T can take on any value, then it must be that the marginal value of a change in x_T must equal to zero, i.e., $\partial V/\partial x_T=0$. Hence, the first-order condition with respect to x_T is

$$\frac{\partial V}{\partial x_T} = \int_0^T \left[H_t \frac{\partial t}{\partial x_T} + H_x \frac{\partial x_t}{\partial x_T} + H_z \frac{\partial z_t}{\partial x_T} + H_\lambda \frac{\partial \lambda_t}{\partial x_T} + \dot{\lambda}_t \frac{\partial x_t}{\partial x_T} + x_t \frac{\partial \dot{\lambda}_t}{\partial x_T} \right] dt \quad - \lambda_T = 0$$

Several terms in this derivative must equal zero. First, clearly it holds that $\partial t/\partial x_T = 0$ so $H_t \frac{\partial t}{\partial x_T} = 0$.

Second, as stated above when we converted from L to V, λ_t will have no effect on V as long as the constraint is satisfied, i.e., as long as the state equation is satisfied. Hence, the terms that involve $\frac{\partial V}{\partial \lambda_t}$ or $\frac{\partial V}{\partial \dot{\lambda}_t}$ can be ignored. Hence,

$$\frac{\partial V}{\partial x_T} = \int_0^T \left[\cancel{H_t \frac{\partial t}{\partial x_T}} + H_x \frac{\partial x_t}{\partial x_T} + H_z \frac{\partial z_t}{\partial x_T} + \cancel{H_\lambda \frac{\partial \lambda_t}{\partial x_T}} + \dot{\lambda}_t \frac{\partial x_t}{\partial x_T} + \cancel{x_t \frac{\partial \dot{\lambda}_t}{\partial x_T}} \right] dt \quad - \lambda_T = 0$$

or

$$\frac{\partial V}{\partial x_T} = \int_0^T \left[H_x \frac{\partial x_t}{\partial x_T} + H_z \frac{\partial z_t}{\partial x_T} + \dot{\lambda}_t \frac{\partial x_t}{\partial x_T} \right] dt - \lambda_T = 0$$

$$\frac{\partial V}{\partial x_T} = \int_0^T \left[\left(H_x + \dot{\lambda}_t\right) \frac{\partial x_t}{\partial x_T} + H_z \frac{\partial z_t}{\partial x_T} \right] dt \quad - \lambda_T = 0$$

As we derived above, the maximum principle requires that $H_x = -\dot{\lambda}_t$ and H_z=0, so both of the terms inside the integral equal zero at the optimum. Hence, we are left with

$$\frac{\partial V}{\partial x_T} = -\lambda_T = 0 .$$

The minus sign on the LHS is there because it reflects the marginal cost of leaving a marginal unit of the stock at time T. In general, we can show that λ_t is the value of an additional unit of the stock at time t. Setting this FOC equal to zero, we obtain the transversality condition, λ_T=0.

This confirms our intuition that since we're attempting to maximize V over our planning horizon, from the perspective of the beginning of that horizon x_T is a variable to be chosen, it must hold that λ_T, the marginal value of an additional unit of x_T, must equal zero. Note that this is the marginal value to V, i.e., to the sum of all benefits over time for 0 to T, not the value to the benefit function, $F(\cdot)$. Although an additional unit may add value if it arrived at time T, i.e., $\partial F(\cdot)/\partial x_T > 0$, the costs that are necessary for that marginal unit of x to arrive at T must exactly balance the marginal benefit.

B. Horizontal terminal line or fixed-endpoint problem

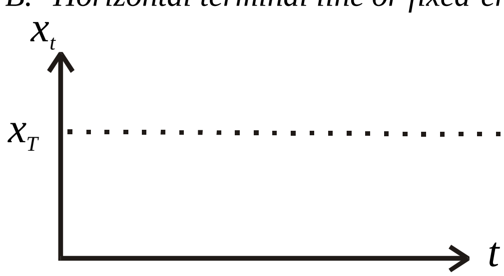

In this case there is no fixed endpoint, but the ending state variables must have a given level. For example, you can keep an asset as long as you wish, but at the end of your use it must be in a certain state. Again, we will use equation 1:

$$V = \int_0^T \left[H(t, x, z, \lambda) + \dot{\lambda}_t x_t\right]dt + \lambda_0 x_0 - \lambda_T x_T .$$

Now, if we have the right terminal time, it must be the case that $\partial V/\partial T$=0, for otherwise it would certainly be the case that the optimal thing to do would be increase T if it were positive and decrease if it were negative. (Note that this is a necessary, but not sufficient condition -- for the sufficient condition we'll have to wait until we introduce an infinite horizon framework). Evaluating this derivative (remember Leibniz's rule in PS 1), we get

$$\frac{\partial V}{\partial T} = \left[H\left(T, x_T, z_T, \lambda_T\right) + \dot{\lambda}_T x_T\right] - \left(\dot{\lambda}_T x_T + \lambda_T \dot{x}_T\right) = 0$$

The second and third terms cancel and, since we are restricted to have x_T equal to a specific value, it follows that $\dot{x}_T = 0$. Hence, the condition reduces to $H(T,x_T,z_T,\lambda_T)=0$, i.e., $H=F(T,x_T,z_T)+\lambda_T(f(T, x_T,z_T))=0$

C. Fixed Terminal Point

In this case both x_T and T are fixed. Such would be the case if you're managing the asset and, at the end of a fixed amount of time you have to have the asset in a specified condition. A simple case: you rent a car for 3 days and at the end of that time the gas tank has to have 5 gallons in it. There's nothing complicated about the transversality condition here, it is satisfied by the constraints on T and x_T , i.e. x_3=5.

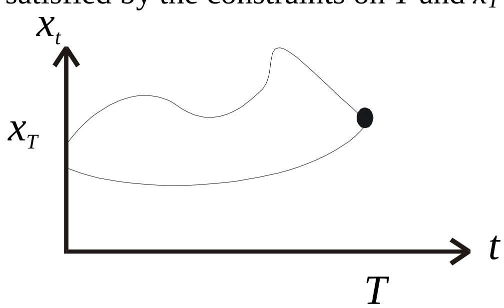

When added to the other optimum criteria, this transversality equation gives you enough equations to solve the system and identify the optimal path.

D. Terminal Curve

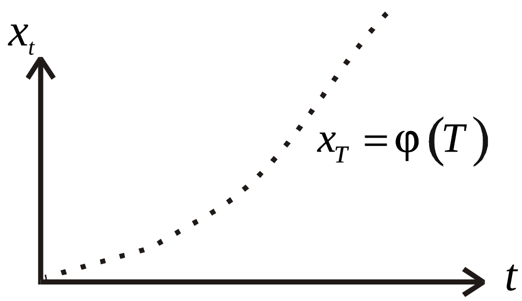

In this case the terminal condition is a function, $x_T = \varphi(T)$. Again, we use

1 $\quad V = \int_0^T \left[H(t,x,z,\lambda) + \dot{\lambda}_t x_t \right] dt \;\; + \lambda_0 x_0 - \lambda_T x_T$.

Taking the derivative with respect to T and substituting in $\dot{x}_T = \phi'(T)$

$$\frac{\partial V}{\partial T} = H(T, x_T, z_T, \lambda_T) \quad + \dot{\lambda}_T x_T - \dot{\lambda}_T x_T - \lambda_T \phi'(T) = 0$$

which can be simplified to the transversality condition,

$$\frac{\partial V}{\partial T} = H(T, x_T, z_T, \lambda_T) - \lambda_T \phi'(T) = 0$$

E. Truncated Vertical Terminal Line

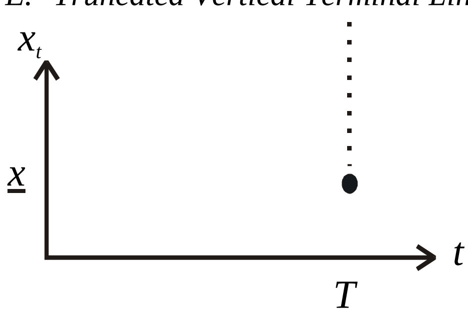

In this case the terminal time is fixed, but x_T can only take on a set of values, e.g. $x_T \geq \underline{x}$. This would hold, for example, in a situation where you are using a stock of inputs that must be used before you reach time T and $x_T \geq 0$. You can use the input from 0 to T, but x_t can never be negative.

For such problems there are two possible transversality conditions. If $x_T > \underline{x}$, then the transversality condition $\lambda_T = 0$ applies. On the other hand, if the optimal path is to reach the constraint on x, then the terminal condition would be $x_T = \underline{x}$. In general, the Kuhn-Tucker specification is what we want. That is, our maximization objective is the same, but we now have an inequality constraint, i.e., we're seeking to maximize

$$V = \int_0^T \left[H(t, x, z, \lambda) + \dot{\lambda}_t x_t \right] dt + \lambda_0 x_0 - \lambda_T x_T \quad \text{s.t. } x_T \geq \underline{x}.$$

The Kuhn-Tucker conditions for the optimum then are:

$\lambda_T \geq 0$, $x_T \geq \underline{x}$, and $(x_T - \underline{x})\lambda_T = 0$

where the last of these is the complementary slackness condition of the Kuhn-Tucker conditions.

As a practical matter, rather than burying the problem in calculus and algebra, I suggest that you would typically take a guess, Is x_T going to be greater than $\underline{x}$? If you think it is, then solve, the problem first using $\lambda_T = 0$. If your solution leads to $x_T \geq \underline{x}$, you're done. If not, substitute in $x_T = \underline{x}$ and solve again. This will usually work. *When would this approach not work?*

F. Truncated Horizontal Terminal Line

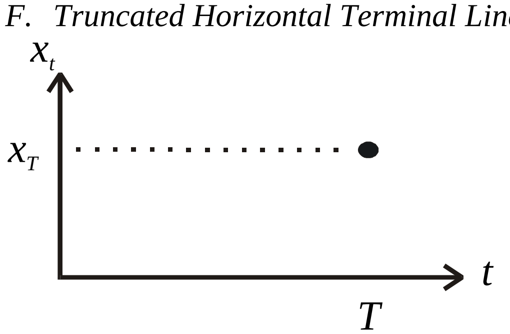

In this case the time is flexible up to a point, e.g., $T \leq T_{max}$, but the state is fixed at a given level, say x_T is fixed. Again there are two possibilities, $T = T_{max}$ or $T < T_{max}$. Using the horizontal terminal line results from above, the transversality condition takes on a form similar to the Kuhn-Tucker conditions above,

$T \leq T_{max}$, $H(T, x_T, z_T, \lambda_T) \geq 0$, and $(T - T_{max})H_T = 0$.

G. Vertical terminal line with salvage value

Let's now re-introduce the salvage value function $S(T, x_T)$ in the context of a problem with a defined length but freedom as to the terminal stock. Recall the transversality condition of the problem without a salvage value was simply that $\lambda_T = 0$, that is, that at the end of the period the marginal benefit of an additional unit of the stock contributes nothing to the objective. Similar intuition holds here. In this case we want the possible benefit of a change in the stock to be equal to the cost associated with not leaving that increment for salvage value. The transversality condition can easily be derived by noting that the optimization criterion can now be written:

$$1' \quad V = \int_0^T \left[H(t, x, z, \lambda) + \dot{\lambda}_t x_t \right] dt + \lambda_0 x_0 - \lambda_T x_T + S(T, x_T)$$

Hence, our transversality condition becomes:

$$\lambda_T = \frac{\partial S(T, x_T)}{\partial x_T}.$$

Note that the addition of the salvage value does not affect the Hamiltonian, nor will it affect the first 3 of the criteria that must be satisfied. *What would be the transversality condition for a horizontal end-point problem with a salvage value?*

H. An important caveat

Most of the results above will not hold exactly if there are additional constraints on the problem or if there is a salvage value. However, you should be able to derive similar transversality conditions equation 1 and similar logic.

III. Infinite horizon problems

It is frequently the case (I would argue, usually the case) that the true problem that we are interested in is of infinite horizon. The optimality conditions for an infinite horizon problem are identical to those of a finite horizon problem with the exception of the transversality condition. Hence, in solving the problem the most important change is how we deal with the need for the transversality conditions. [Obviously, in infinite horizon problems the mnemonic of *transversing* to the other side doesn't really work because there is no "other side" to which we might transverse.]

A. Fixed finite x

If we have a value of *x* to which we must arrive, i.e., $x_\infty \equiv \lim_{t\to\infty} x_t = k$, then the problem is identical to the horizontal terminal line case considered above.

B. Flexible x_T

Recall from above that for the finite horizon problem we used equation 1:

$$V = \int_0^T \left[H(t, x, z, \lambda) + \dot{\lambda}_t x_t \right] dt + \lambda_0 x_0 - \lambda_T x_T .$$

In the infinite horizon case this equation is rewritten:

$$V = \int_0^\infty \left[H(t, x, z, \lambda) + \dot{\lambda}_t x_t \right] dt + \lambda_0 x_0 - \lim_{t\to\infty} \lambda_t x_t$$

and, for problem in which x_∞ is free, the condition analogous to the transversality condition in the finite horizon case is $\lim_{t\to\infty} \lambda_t = 0$. Note that if our objective is to maximize the present-value of benefits, this means that the **present value** of the marginal value of an additional unit of *x* must go to zero as *t* goes to infinity. Hence, the current value (at time *t*) of an additional unit of *x* must either be finite or grow at a rate slower than *r* so that the discount factor, e^{-rt}, pushes the present value to zero.

One way that we frequently present the results of infinite horizon problems is to evaluate the equilibrium where $\dot{\lambda} = \dot{x} = 0$. Using these equations (and evaluating convergence and stability via a phase diagram) we can then solve the problem.

IV. Summary

The central idea behind all transversality conditions is that if there is any flexibility at the end of the time horizon, then the marginal benefit from taking advantage of that flexibility must be zero at the optimum. You can apply this general principal to problems with more than one variable, to problems with constraints and, as we have seen, to problems with a salvage value.

V. Reading for next class

Dorfman, Robert. 1969. An Economic Interpretation of Optimal Control Theory. *American Economic Review* 59(5):817-31.