Dynamic Optimization and Optimal Control

Mark Dean+

Lecture Notes for Fall 2014 PhD Class - Brown University

1 Introduction

To finish off the course, we are going to take a laughably quick look at optimization problems in dynamic settings. We will start by looking at the case in which time is discrete (sometimes called dynamic programming), then if there is time look at the case where time is continuous (optimal control).

2 Dynamic Programming

We are interested in recursive methods for solving dynamic optimization problems. While we are not going to have time to go through all the necessary proofs along the way, I will attempt to point you in the direction of more detailed source material for the parts that we do not cover.

We are going to begin by illustrating recursive methods in the case of a finite horizon dynamic programming problem, and then move on to the infinite horizon case.

2.1 The Finite Horizon Case

2.1.1 The Dynamic Programming Problem

The environment that we are going to think of is one that consists of a sequence of time periods, indexed $1, \ldots, T < \infty$. We are going to think about the problem of someone who is choosing a sequence of control variables, $c_t \in C \subset \mathbb{R}$, one for each period (in a standard consumption problem, this represents how much one consumes in each period).

We assume that the maximizing choice in each period may depend on a set of parameters that describe the environment of the decision maker, but also any information that she has at the time of deciding. The state at time t contains all the information that is available to the decider and relevant for decisions today and in the future.

In each period we describe the current state of the economy by $x_t \in X \subset \mathbb{R}^n$ (possibly a vector). The objective function in a dynamic problem is typically the discounted sequence of instantaneous functions $u : X \times C \to \mathbb{R}$, which we will assume is continuous and bounded (in the consumption problem, this would just be the utility of consumption in a given period, but note that period utility could depend on the state variable as well). With discount factor β, the objective function is

$$\sum_{t=0}^{T} \beta^t u(x_t, c_t).$$

In addition, we will allow for uncertainty in the model given by a random variable $z_t \in Z \subset \mathbb{R}^m$, which is assumed to follow a stationary Markov process. This means the distribution of the random variable at time $t+1$ depends only on its value at time t, so that $\Pr(z_{t+1} \leq z | z_t, z_{t-1}, \ldots) = \Pr(z_{t+1} \leq z | z_t)$. We will denote the stochastic process with $Q(z', z) = \Pr\left(z_{t+1} \leq z' | z_t = z\right)$. In the standard consumption problem z might be for example income. We will also assume that the value of z_t is known to the individual at time t, so (z_t, x_t) is the complete description of the state of the economy at time t.

We assume that all the relevant information at time t is contained in x_t and z_t. This requirement means "finding the state is an art" (Tom Sargent, 2003). In a completely standard (deterministic) consumption savings problem, the state of the system is just the current asset level. However, in more complicated problems, finding a usable state variable may be very difficult, or impossible. In particular, since the state space consists of all possible combinations of the different values of each parameter, the state space grows quickly with the number of dimensions. This problem of a rapidly expanding state space is called the 'curse of dimensionality'. For example, in models where markets are incomplete and agents are heterogenous, it may be that the state needs to include the wealth level of every person in the economy, making the optimization problem intractable.

In addition to the objective function and the state, we need to describe the constraints on the choices that the individual can make. We will assume that there is a stationary constraint correspondence, denoted by $\Gamma : X \times Z \rightrightarrows C$, which describes the choice set for a given state. Note that the assumption that Γ is stationary is restrictive: it means for example that there may be credit constraints but we cannot require that the agent cannot die in debt. Finally, we need a law of motion for the state variable, $f : X \times Z \times C \to X$. This function tells us the state in time $t+1$, given x_t, z_t and c_t. In the standard consumption savings problem, this would be an equation that describes how wealth changes (i.e. wealth next period is wealth this period, minus consumption, plus income, all multiplied by the rate of return).

The last missing piece is a specification of the agent's choice. In the consumption problem, she chooses consumption levels for each period. We may also allow her to condition the level of consumption at time t on the state variable at time t. If the problem was deterministic, this wouldn't matter very much, but as we are allowing for a stochastic element to the problem, this is an important distinction. Here the consumer is really choosing a sequence of **decision rules**, or conditional consumption rules. Let $g_t : X \times Z \to C$ be a decision rule for time t. We call the set $\pi_T = (g_0, g_1, \ldots, g_T)$ a **policy** for T periods. We say that a policy is feasible if, for every t, $g_t(x_t, z_t) \in \Gamma(x_t, z_t)$. We say that a policy is stationary if the decision rule does not depend on t.

We can now describe the expected present value of a policy $\pi_{(T)}$ given the initial state variables x_0 and z_0. Expectations are taken with respect to the distribution $Q(z', z)$, and the state variable is assumed to follow the law of motion:

$$W_{(T)}\left(x_0, z_0, \pi_T\right) = E_0\left[\sum_{t=0}^{T} \beta^t u\left(x_t, g_t\left(x_t, z_t\right)\right)\right].$$

We can now state the dynamic programming problem:

$$\begin{aligned} \max_{\pi_{(T)}} W_{(T)}\left(x_0, z_0, \pi_{(T)}\right) \;\; s.t. & \\ x_{t+1} &= f\left(x_t, z_t, g_t\left(x_t, z_t\right)\right) \\ g_t\left(x_t, z_t\right) &\in \Gamma\left(x_t, z_t\right) \\ x_0, z_0, Q(z', z) \text{ given.} & \end{aligned}$$

This is a complex problem, but luckily the recursive approach will allow us to solve for the policy functions g_t one at a time. First we show that a solution to this problem exists. In order to do so, we need some additional assumptions:

Assumptions We will assume that the problem has the following structure:

1. For every $x \in X$ and $z \in Z$, the constraint set $\Gamma(x, z)$ is non-empty and compact,
2. Γ is continuous,
3. f is continuous, and
4. Q satisfies some restriction to ensure that E_0 is bounded and continuous in x_t[1].

Now we can apply the Theorem of the Maximum that we have discussed previously.

Theorem 1 (Theorem of the Maximum) *Let $X \subset \mathbb{R}^l$ and $Y \subset \mathbb{R}^m$. Suppose $u : X \times C \to \mathbb{R}$ is a continuous function, and $\Gamma : X \to C$ is a compact-valued, continuous correspondence. Then $U(x) = \max_{c \in \Gamma(x)} u(x, c)$ is well defined and continuous, and $c^*(x) = \arg\max_{y \in \Gamma(x)} u(x, y)$ is well defined, non-empty, compact-valued and upper hemicontinuous.*

Using the Theorem of the Maximum, we prove the following statement:

Theorem 2 *Under the stated assumptions, the dynamic programming problem has a solution, the optimal policy π_T^*. The value function $V_{(T)}(x_0, z_0) = W_{(T)}\left(x_0, z_0, \pi_{(T)}^*\right)$ is continuous in x_0.*

Proof. *We will prove this iteratively. If $T = 0$, the statement follows directly from the theorem of the maximum. Now let $\pi_{(T)} = \{g_t^*(x_t, z_t)\}_{t=0}^T$ be the optimal policy from period 0 to T. Define*

$$V_{(T)}(x_0, z_0) = E_0 \sum_{t=0}^{T} \beta^t u\left(x_t, g_t^*(x_t, z_t)\right)$$

as the value of the optimal policy for T periods given initial state values (x_0, z_0). Assume that $V_{(T)}$ is continuous in x_0. Now note that

$$\begin{aligned}
\max_{\pi_{(T+1)}} W_{(T+1)}\left(x_0, z_0, \pi_{(T+1)}\right) &= \max_{\pi_{(T+1)}} E_0 \sum_{t=0}^{T+1} \beta^t u\left(x_t, g_t(x_t, z_t)\right) \\
&= \max_{\pi_{(T+1)}} E_0 \left\{ u(x_0, c_0) + \sum_{t=1}^{T+1} \beta^t u\left(x_t, g_t(x_t, z_t)\right) \right\} \\
&= \max_{\pi_{(T+1)}} E_0 \left\{ u(x_0, c_0) + E_1 \sum_{t=1}^{T+1} \beta^t u\left(x_t, g_t(x_t, z_t)\right) \right\} \\
&= \max_{c_0} E_0 \left\{ u(x_0, c_0) + \max_{\pi_{(T)}} E_1 \sum_{t=1}^{T+1} \beta^t u\left(x_t, g_t(x_t, z_t)\right) \right\},
\end{aligned}$$

[1] One example is the Feller property, see page 220 of Stokey and Lucas

where $x_1 = f(x_0, z_0, c_0)$ and the distribution of z_1 is determined by Q and z_0. Therefore

$$\max_{\pi_{(T+1)}} W_{(T+1)}(x_0, z_0, \pi_{T+1}) = \max_{c_0} \left\{ u(x_0, c_0) + \beta E_0 \left[V_{(T)}(f(x_0, z_0, c_0), z_1) \right] \right\}$$

These steps follows from the law of iterated expectations, and from the recursive structure of the problem (i.e the choice of control variable at time s can only affect the state variable at times after s). But this new problem is now a problem of maximizing a continuous function of c_0 on a compact valued continuous correspondence (as $V_{(T)}(x, z)$ is continuous in x by assumption, $f(x, z, c)$ is continuous, $u(x, c)$ is continuous, and the expectation operator is continuous). Thus, by the Theorem of the Maximum $V_{(T+1)}(x_0, z_0)$ is well defined and continuous. ■

A corollary to this statement (that we will not prove) is as follows:

Corollary 1 *If $\Gamma(x_t, z_t)$ is convex and $u(\cdot)$ and $f(\cdot)$ are strictly concave in c_t, then $g_t(x_t, z_t)$ is also continuous.*

Notice that during this proof, we re-wrote the problem is a **recursive** way. That is, we wrote the value today as a function of the (optimal) choice today plus the value tomorrow, assuming that the optimal policy would be carried out from that point onwards (this is the function $V_{(T)}$). Generalizing this result to the case in which we have s periods left to go, inserting the state equation into next period's value function, and using the definition of conditional expectation, we arrive at *Bellman's equation* of dynamic programming with a finite horizon (named after Richard Bellman (1956)):

$$V_{(s)}(x, z) = \max_{c \in \Gamma(x,z)} \left\{ u(x, c) + \beta \int_Z V_{(s-1)} \left(f(x, z, c), z' \right) dQ(z', z) \right\} \tag{1}$$

where x and z denote more precisely x_{T-s} and z_{T-s} respectively, and z' denotes z_{T-s+1}.

Bellman's equation is useful because it reduces the **choice of a sequence** of decision rules to a **sequence of choices** for the decision rules. It is sufficient to solve the problem in (1) sequentially $T + 1$ times, as shown in the next section. Hence a dynamic problem is reduced to a sequence of static problems. A consequence of this result is the so-called *Bellman's principle of optimality* which states that if the sequence of functions $\pi_T^* = \{g_0^*, g_1^*, ..., g_T^*\}$ is an optimal policy, then to maximize the objective function $W_{(T-s)}\left(x_s, z_s, \pi_{(T-s)}\right)$ after s periods, the decision rules $\left\{g_s^*, g_{s+1}^*, ..., g_T^*\right\}$ are still optimal. In other words, as time advances, there is no incentive to depart from the original plan. Policies with this property are also said to be *time-consistent*.

Time consistency depends on the recursive structure of the problem and does not apply to more general settings. David Laibson became famous by suggesting that people do not discount exponentially. Instead, he proposed that people suffer from "present bias" and put extra weight on consumption that occurs in the current period. Thus, their utility function is given by

$$E_0 \left\{ u(x_0, c_0) + \delta \sum_{t=1}^{T} \beta^t u(x_t, c_t) \right\}$$

where $0 < \delta < 1$ is the additional discounting for consumption that does not occur today. Assuming that there exists an optimal policy $c_t = g_t^* (x_t, z_t)$, note that

$$V_{(T)} (x_0, z_0) = E_0 \left\{ u(x_0, g_0^* (x_0, z_0)) + \delta \sum_{t=1}^{T} \beta^t u(x_t, g_t^* (x_t, z_t)) \right\}$$

is *not* the same as

$$E_0 \left\{ u(x_0, g_0^* (x_0, z_0)) + \delta \beta E_0 \left[V_{(T-1)} (f(x_0, z_0, g_0^* (x_0, z_0)), z_1) \right] \right\}.$$

The reason is that the objective function from period $t = 1$ onwards looks different viewed from period zero and from period one:

$$\delta \sum_{t=1}^{T} \beta^t u (x_t, c_t)$$

$$u(x_1, c_1) + \delta \sum_{t=2}^{T} \beta^t u(x_t, c_t).$$

This implies both that the problem does not have a recursive structure, and that optimal plans made at period 0 may no longer be optimal in period 1.

2.1.2 Backward Induction

If the problem we are considering is actually recursive, we can apply backward induction to solve it.

1. Start from the last period T , with 0 periods to go. Then the problem is static and reads:

   $$V_{(0)}(x_T, z_T) = \max_{c_T \in \Gamma(x_T, z_T)} u (x_T, c_T)$$

   which yields the optimal choice $g_T^* (x_T, z_T)$ for the final value of x_T and the final realization of z_T. Hence,given a given specification of $u(\cdot)$, we have an explicit functional form for $V_{(0)}(x_T, z_T)$.

2. We can now go back by one period and use the constraint $x_T = f(x_{T-1}, z_{T-1}, c_{T-1})$ to write:

$$\begin{aligned} V_{(1)}(x_{T-1}, z_{T-1}) &= \max_{c_{T-1}\in\Gamma(x_{T-1},z_{T-1})} \{u\left(x_{T-1}, c_{T-1}\right) + \\ &\quad +\beta \int_Z V_{(0)}\left(f(x_{T-1}, z_{T-1}, c_{T-1}), z_T\right) dQ(z_T, z_{T-1})\Big\} \end{aligned}$$

which allows to solve for $g^*_{T-1}\left(x_{T-1}, z_{T-1}\right)$ and obtain $V_{(1)}(x_{T-1}, z_{T-1})$ explicitly.

3. We continue until time 0 and collect the sequence of decision rules in each t. This way we can construct the optimal policy contingent on x_0 and any realization of $\{z_t\}_{t=0}^T$.

As a simple example, consider the following 'cake eating' problem:

$$\begin{aligned} &\max_{\{c_t\}_{t=0}^T} \sum_{t=0}^T \beta^t \ln(c_t) \\ \text{subject to } x_{t+1} &= (1-\delta)x_t - c_t \\ c_t &\geq 0 \\ x_{t+1} &\geq 0 \\ &x_0 \text{ given} \end{aligned}$$

You should check that this satisfies our assumptions (note that we can reformulate the constraints as $\Gamma(x) = [0, x]$).

The Bellman equation for this problem is given by

$$V_{(s)}\left(x\right) = \max_{c\in\Gamma(x,z)} \left\{\ln(c) + \beta V_{(s-1)}\left((1-\delta)x - c\right)\right\}$$

For the final period, we know that the optimal policy must be to consume whatever is left, yielding

$$\begin{aligned} g_T(x) &= x \\ V_{(0)}(x) &= \ln(x). \end{aligned}$$

Thus, going back one period, we have the static problem

$$V_{(1)}(x) = \max_{c\in\Gamma(x,z)} \left\{\ln(c) + \beta V_{(0)}((1-\delta)x - c)\right\} = \max_{c\in\Gamma(x,z)} \left\{\ln(c) + \beta \ln\left((1-\delta)x - c\right)\right\}.$$

Taking first order conditions gives us

$$\frac{1}{c} = \frac{\beta}{((1-\delta)x - c)}$$

so

$$
\begin{aligned}
(1-\delta)x - c &= c\beta \\
\Rightarrow c &= \frac{(1-\delta)}{1+\beta}x
\end{aligned}
$$

and so

$$
\begin{aligned}
g_{T-1}(x) &= \frac{(1-\delta)}{1+\beta}x \\
V_{(1)}(x) &= \ln\left(\frac{(1-\delta)}{1+\beta}x\right) + \beta\ln\left((1-\delta)x - \frac{(1-\delta)}{1+\beta}x\right) \\
&= \ln\left(\frac{(1-\delta)}{1+\beta}x\right) + \beta\ln\left(\beta\frac{(1-\delta)}{1+\beta}x\right) \\
&= (1+\beta)\ln\left(\frac{(1-\delta)}{1+\beta}x\right) + \beta\ln\beta
\end{aligned}
$$

This process can then be iterated to solve the dynamic programming problem.

2.2 The Infinite Horizon Case

We will now drop the assumption that the problem is finitely lived, and think about the case in which $T = \infty$. The bad news is that (obviously), we can no longer solve these problems by backward induction. The good news is that this where recursive dynamic programming really comes in to its own. Note that the sequence problem is now

$$
\begin{aligned}
&\max_{\pi} W(x_0, z_0, \pi)\ s.t. \\
&\qquad x_{t+1} = f(x_t, z_t, g_t(x_t, z_t)) \\
&\qquad g_t(x_t, z_t) \in \Gamma(x_t, z_t) \\
&x_0, z_0, Q(z', z) \text{ given.}
\end{aligned}
$$

Intuitively, the equivalent of the recursive functional equation should be

$$
V(x_0, z_0) = \max_{c} u(x_0, z_0, c) + \beta E_0 V(f(x_0, z_{0,}c), z_1)
$$

It should be the case that the problem is not only recursive, it is also stationary. After all, in the infinite horizon case, the problem starting tomorrow looks exactly as it does starting today, as long as the initial state is the same. Thus, we should be able to drop the time index from policy and value functions. In fact, this is the case (under a couple of regularity conditions). This is the **principle of optimality**, first formulated by Richard Bellman. However, it does require proving (I'm going to omit the proof, but if you are interested it is in Section 4.1 of Stokey and Lucas).

Does such a value function exist? Luckily we have already got the machinery to prove this is the case. Consider a mapping T that maps a set of functions $V : X \times Z \to \mathbb{R}$ into itself. Let T be given by

$$T(V)(x,z) = \max_{c} u(x,c) + \beta E(V(f(x,z,c),z_1)).$$

We want to find a fixed point of this mapping, i.e. $V = T(V)$. Recall that in previous lectures we have shown the following result:

Theorem 3 (Banach Fixed Point Theorem) *Let X be a complete metric space, and f be a contraction on X. Then there exists a unique x^* such that $f(x^*) = x^*$.*

Now we combine this with

Lemma 1 (Blackwell's Contraction Lemma) *Let S be a non-empty set, and F a non-empty subset of $\mathcal{B}(S)$ that is closed under addition of positive constant functions. Assume that T is an increasing self-map on F. If there exists a $0 < \delta < 1$ such that*

$$T(f+\alpha) \leq T(f) + \delta\alpha \text{ for all } (f,\alpha) \in \mathcal{B}(S) \times \mathbb{R}_+$$

then T is a contraction. (Note that we define $f \leq g$ if $f(x) \leq g(x)\ \forall\ x \in S$).

It is true (but you should check), that our assumptions are enough to guarantee that the conditions of Blackwell's Lemma hold for T, and that T maps the space of continuous, bounded functions (which is complete) into itself. Thus, we know that T has a fixed point, and we can characterize V as a fixed point of a contraction mapping. One advantage is that this provides us with a straightforward way of finding V. Recall that the the proof of Banach simply used the iteration of the contraction, starting from any x_0, by setting $x_1 = f(x_0)$ and $x_2 = f(x_1)$ and so on.

The fixed point is just the limit of this sequence. V is therefore the limit of

$$\begin{aligned} V_{(1)}(x,z) &= \max_{c\in\Gamma(x,z)}\left\{u(x,c)+\beta\int_Z V_{(0)}\left(f(x,z,c),z'\right)dQ(z',z)\right\} \\ V_{(2)}(x,z) &= \max_{c\in\Gamma(x,z)}\left\{u(x,c)+\beta\int_Z V_{(1)}\left(f(x,z,c),z'\right)dQ(z',z)\right\} \\ &\vdots \end{aligned}$$

Thus, we know that V exists and is unique, and that we can think of the value function of the infinite horizon problem as the limit of a sequence of finite horizon problems. V can be found computationally by iterating as above. Moreover, the contraction guarantees that V preserves some of the properties T, such as monotonicity and continuity.

2.2.1 Characterization of the Policy Function

Associated with the stationary value function, there is a stationary policy function $g : X \times Z \to C$ that solves the relevant maximization problem

$$V(x,z) = \max_{c\in\Gamma(x,z)}\left\{u(x,c)+\beta\int_Z V\left[f(x,z,c),z'\right]dQ(z',z)\right\}.$$

Assuming that the mapping between states $f(x,z,c)$ is invertible, we can write $c = \gamma(x,x',z)$ and treat the problem as if the decision maker chooses x', the state next period:

$$V(x,z) = \max_{x'}\left\{u\left[x,\gamma(x,x',z)\right]+\beta\int_Z V\left(x',z'\right)dQ(z',z)\right\}$$

If we knew that V was differentiable in x, then we could characterize the solution using the FOC

$$u_c(x,c)\,\gamma_2(x,x',z)+\beta\int_Z V_x\left(x''\right)dQ(z',z)=0.$$

How do we know that V is differentiable? The following result comes in handy:

Theorem 4 (The Envelope Theorem) *Let V be a concave function defined on the set X, let $x_0 \in int(X)$, and let $N(x_0)$ be a neighborhood of x_0. If there is a concave differentiable function $\Omega : N(x_0) \to \Re$ such that $\Omega(x) \leq V(x)$, $\forall x \in N(x_0)$ with the equality holding at x_0, then V is differentiable at x_0 and $V_x(x_0) = \Omega_x(x_0)$.*

Proof. *See Stokey-Lucas, p. 85.* ■

We can use this result to show that the value function is differentiable, and to find the derivative. To see this, let x'_0 be the value of x' that solves the optimization problem at (x_0, z_0). Then let

$$\Omega(x,z) = u\left[x, \gamma(x, x'_0, z)\right] + \beta \int_Z V\left(x'_0, z'\right) dQ(z', z),$$

that is we are holding x'_0 fixed but allow x and z to vary. By definition, $\Omega(x_0, z_0) = V(x_0, z_0)$. Moreover, as $V(x,z)$ is the maximized value of $u\left[x, \gamma((x, x', z)\right] + \beta \int_Z V\left(x', z'\right) dQ(z', z)$, it must also be the case that $\Omega(x,z) \leq V(x,z)\ \forall\ x$. Thus, assuming that u is concave and differentiable, we get

$$V_x(x,z) = \Omega_x(x,z) = u_c(x,c)\gamma_1(x,x',z) + u_x(x,c)$$

This gives us the Euler equation

$$u_c(x,c)\,\gamma_2(x,x',z) + \beta \int_Z u_x(x',c') + u_c(x',c')\gamma_1(x',x'',z'',z) = 0$$

To understand this, it is worth thinking about the standard consumption problem (let's do in the deterministic case). Here the problem is

$$\begin{aligned} & \max \sum_{t=0}^{\infty} \beta^t u(c_t) \\ \text{subject to } w_{t+1} &= (1-\delta)w_t - c_t + I_t \end{aligned}$$

where w_t is wealth at time t and I_t is income. Thus, $\gamma(w_t, w_{t+1}) = w_{t+1} - (1-\delta)w_t - I_t$, and therefore $\gamma_2(w_t, w_{t+1}) = 1$ and $\gamma_1(w_t, w_{t+1}) = -(1-\delta)$. The Euler equation gives us

$$\begin{aligned} u'(w_{t+1} - (1-\delta)w_t - I) &= \beta(1-\delta)u'(w_{t+2} - (1-\delta)w_{t+1} - I) \\ u'(c_t) &= \beta(1-\delta)u'(c_{t+1}) \end{aligned}$$

This is a nonlinear difference equation in the state variable. To be able to fully characterize the optimal dynamic path of the state $\{x_t^*\}_{t=0}^{\infty}$, we need two boundary conditions:

- initial condition: x_0, z_0 given
- transversality condition (TVC): $\lim_{t\to\infty} \beta^t u_x(x_t^*, c_t^*) x_t^* = 0$

The latter condition requires that the present discounted value of the state at time t along the optimal path tends to zero as t goes to infinity. Hence, either the state variable is small enough in the limit, or its marginal value is small enough.

2.2.2 Sufficiency of the Euler Equation and TVC

It is possible to prove that if a sequence of states $\{x_t^*\}_{t=0}^{\infty}$ satisfies the Euler equation and the TVC, then it is optimal for the dynamic programming problem. To do so, we will show that the difference D between the objective function evaluated at $\{x_t^*\}$ and at $\{x_t\}$, any alternative feasible sequence of states, is non-negative. In this particular proof, we abstract from the stochastic nature of the problem.[2]

First, we slightly rewrite the utility function and therefore the Euler equation as:

$$\begin{aligned} u_c(x,c)\gamma_2(x,x') &+ \beta u_x(x',c') + \beta u_c(x',c')\gamma_1(x',x'') = 0 \\ \hat{u}_2(x,x') &+ \beta \hat{u}_1(x',x'') = 0 \end{aligned}$$

by replacing $u(x,\gamma(x,x'))$ with $\hat{u}(x,x')$.

We assume concavity of the return function $\hat{u}$, and choose x such that $\hat{u}_1(x,x') \geq 0$ and $x \geq 0$. Concavity of f means that for any x, x', $f(x) - f(x') \geq f'(x)(x - x')$. Thus

$$\begin{aligned} D &= \lim_{T\to\infty} \sum_{t=0}^{T} \beta^t \left[\hat{u}(x_t^*, x_{t+1}^*) - \hat{u}(x_t, x_{t+1})\right] \\ &\geq \lim_{T\to\infty} \sum_{t=0}^{T} \beta^t \left[\hat{u}_1(x_t^*, x_{t+1}^*)(x_t^* - x_t) + \hat{u}_2(x_t^*, x_{t+1}^*)(x_{t+1}^* - x_{t+1})\right] \end{aligned}$$

Since $x_0^* = x_0$ from the initial condition, we can rewrite the sum above as:

$$\begin{aligned} D \geq \lim_{T\to\infty} \quad & \left\{ \sum_{t=0}^{T-1} \beta^t \left[\hat{u}_2(x_t^*, x_{t+1}^*) + \beta \hat{u}_1(x_{t+1}^*, x_{t+2}^*)\right] (x_{t+1}^* - x_{t+1}) \right. \\ & \left. + \beta^T u_2(x_T^*, x_{T+1}^*)(x_{T+1}^* - x_{T+1}) \right\} \end{aligned}$$

The Euler equation implies that each term of the sum is zero. Moreover, using the Euler equation to substitute the last term, we obtain:

$$D \geq \lim_{T\to\infty} \beta^T u_2(x_T^*, x_{T+1}^*)(x_{T+1}^* - x_{T+1}) = -\lim_{T\to\infty} \beta^{T+1} u_1(x_{T+1}^*, x_{T+2}^*)(x_{T+1}^* - x_{T+1})$$

Since $u_1 \geq 0$ and $x_{T+1} \geq 0$, then:

$$D \geq -\lim_{T\to\infty} \beta^{T+1} u_1(x_{T+1}^*, x_{T+2}^*)(x_{T+1}^* - x_{T+1}) \geq -\lim_{T\to\infty} \beta^{T+1} u_1(x_{T+1}^*, x_{T+2}^*) x_{T+1}^*$$

Using the TVC, we can finally establish that $D \geq 0$.

[2] This proof is from Stokey-Lucas, pp. 98-99.

2.3 The Stochastic Growth Model

Consider the following optimal growth problem, (stated as a social planner problem)

$$\max_{\{c_t, i_t\}} E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t)\right]$$

$$s.t.$$

$$k_{t+1} = (1-\delta)\, k_t + i_t$$

$$f(k_t, z_t) = c_t + i_t$$

$$k_0, z_0,\ Q(z_{t+1}, z_t) \text{ given}$$

Let's write the recursive formulation of the problem, using the DP techniques:

$$V(k,z) = \max_{k'} \left\{ u\left(f(k,z) - k' + (1-\delta)k\right) + \beta \int V(k', z') dQ(z', z) \right\}$$

The FOC w.r.t. k' and the Envelope condition for this problem are respectively:

- (FOC): $u'\left(f(k,z) - k' + (1-\delta)k\right) = \beta \int V_1(k', z') dQ(z', z)$
- (ENV): $V_1(k,z) = u_k(\cdot) = u'(\cdot)\left[f_1(k,z) + (1-\delta)\right]$

Therefore, we obtain the Euler Equation:

$$u'\left(f(k,z) - k' + (1-\delta)k\right) = \beta \int u'\left(f(k',z') - k'' + (1-\delta)k'\right)\left[1 + f_1(k',z') - \delta\right] dQ$$

which, together with the TVC

$$\lim_{t\to\infty} \beta^t u_1(\cdot)\, k_t = 0$$

and the initial condition of the economy (k_0, z_0) fully characterizea the solution of the problem.

The Euler equation can be given an intuitive interpretation. It states that the marginal utility of consumption this period, along the optimal path, must be equal to the discounted expected marginal utility of consumption next period accounting for the fact that postponing consumption by one period will increase consumption tomorrow by the rate of return on capital net of depreciation.

3 Optimal Control

We will now move onto the case where is time is continuous, rather than discrete.

3.1 The General Control Problem

Consider a dynamic system evolving over time, and an agent, the planner, who has the task to control the behavior of this system. Time is continuous and indexed by t, with $t \in [0,T]$. The interval $[0,T]$ is called the planning horizon. We start by assuming that T is finite, later we'll relax this assumption. Once again, we will summarize the state of the system with a set of state variables, denoted as $s(t) \in S$, with $s : [0,T] \to \mathbb{R}^n$. The planner has knowledge of the initial and final state of the system, hence we can define the **boundary conditions** $s(0) = s_0$, $s(T) = s_T$.

The state variables can be affected through a set of **control variables** (or instruments) $c(t) \in C$, $c : [0,T] \to \mathbb{R}^m$ $m \leq n$. By choosing in every instant a value of each control variable, the planner can modify the paths of the state variables and the dynamic behavior of the system. The planner knows the relationship between the actions taken and the evolution of the states, which are summarized by a **law of motion** of the states, a differential equation of the type

$$\frac{ds(t)}{dt} \equiv \dot{s}(t) = g\left(s(t), c(t), t\right), \text{ with } g : S \times C \times [0,T] \to \mathbb{R}^n.$$

Once the values are chosen for the control variables at a given instant t, the rates of change of the state variables at t are obtained by the law of motion, and given the value of the state at t (which is predetermined), the future values of the state variables are determined. Often the planner is limited in its choice of the controls by a **set of constraints** $f\left(s(t), c(t), t\right) \geq 0$ which defines the set of admissable controls.

In order to choose properly the controls, the planner needs an **objective function**, i.e. a function describing the value of the system for any given path of states and controls, which we denote as $V = \int_0^T u\left(s(t), c(t), t\right) dt$. The function $u(\cdot)$ is called **instantaneous return (or utility) function**, and the function V, **value function** as it summarizes the value of any given path of controls and states.

We now enrich the problem with some additional structure, through a couple of key assumptions:

Assumption 1: $u(\cdot,\cdot,\cdot)$, $g(\cdot,\cdot,\cdot)$, $f(\cdot,\cdot,\cdot)$, $s(\cdot), c(\cdot)$ *are continuous and differentiable*

Assumption 2: $w(s(t),t) = \{c(t) \in C \ s.t. \ f(s(t),c(t),t) \geq 0\}$ *is compact and convex.*

We are now ready to formulate the general control problem **(CP)** faced by the planner:

$$
\begin{aligned}
&\max_{\{c(t)\}_{t\in[0,T]}} V = \int_0^T u(s(t),c(t),t)\,dt \\
&\qquad s.t. \\
&\dot{s}(t) = g(s(t),c(t),t) \\
&s(t) \in S,\ c(t) \in C \\
&f(s(t),c(t),t) \geq 0 \\
&s(0) = s_0,\ s(T) = s_T
\end{aligned}
\tag{CP}
$$

The **solution** of the problem (CP) above is the optimal path of the control variable (and the state, given the initial conditions and the law of motion) for each instant in the time interval considered. The salient difference with a static optimization problem is obvious: the solution is not just a singe optimal value for each choice variable, but an optimal time path for each control variable, i.e. a function of time $\{c^*(t)\}$. Note also that for the moment this is a deterministic, rather than a stochastic problem.

Example 1 (Lifetime Consumption Alloctaion) *Denote by $c(t)$ the flow of consumption at time t and assume that households obtain utility only by consuming, through the instantaneous utility function $u(c(t))$. Each household knows that he will live for T periods. He starts his life with no wealth, i.e. $w(0) = 0$ and does not want to leave behind him any bequest, hence $w(T) = 0$. Every period he receives an endowment y that he can either consume or save for future consumption. Savings are remunerated at rate r. He will therefore face the following problem:*

$$
\begin{aligned}
&\max_{\{c(t)\}_{t\in[0,T]}} \int_0^T u(c(t))dt \\
&\qquad s.t. \\
&\dot{w}(t) = s(t) + rw(t) \\
&y = c(t) + s(t) \\
&w(0) = 0,\ w(T) = 0.
\end{aligned}
$$

The simplest solution method for problems in the class of (CP) is the **Maximum Principle**, developed by the Russian mathematician Pontryagin in the late 50's.

4 The Maximum Principle

4.1 A Derivation of the Maximum Principle from the Lagrangean

It is useful to derive the Maximum Principle from the Lagrangean associated to (CP) since we are more acquainted with using the Lagrangean in static optimization. We simplify the exposition in three ways. First, we ignore the feasibiliy constraint f, or equivalently, we assume that we can substitute the set of constraints of f into the transition equation, thus the transition equation g incorporates the constraint f as well. Second, we assume that $m = n = 1$. Third, to simplify the notation, we denote the dynamic path for the variable x, i.e. $\{x(t)\}_{t\in[0,T]}$ more simply as $\{x(t)\}$ and we omit the subscript.

The Lagrangean corresponding to problem (CP), can therefore be written as:

$$\mathcal{L} = \int_0^T \left\{ u\left(s(t), c(t), t\right) + \lambda(t) \ \left[g\left(s(t), c(t), t\right) - \dot{s}(t) \right] \right\} dt \tag{L}$$

It is clear that the Lagrangean is obtained by summing up (under the integral sign) all the static Lagrangeans for each time t. Note that $\{\lambda(t)\}$ are the dynamic Lagrange multipliers (or costates), which are no longer constants, but functions of time.

Now, we need to characterize the maximum wrt to $\{c(t)\}$. Integrate by parts the last component of (L), $\int_0^T \lambda(t)\dot{s}(t)dt$, and obtain:

$$\mathcal{L} = \int_0^T \left[u\left(s(t), c(t), t\right) + \lambda(t) g\left(s(t), c(t), t\right) \right] dt - \left[\lambda(T)s(T) - \lambda(t)s(0) \right] + \int_0^T \dot{\lambda}(t)s(t)dt \tag{2}$$

Define now the **Hamiltonian** as $H(s, c, \lambda, t) = u(s, c, t) + \lambda(t)g(s, c, t)$ (where we have omitted the argument t for the sake of simpler notation) and rewrite (2) as:

$$\mathcal{L} = \int_0^T \left[H\left(s(t), c(t), \lambda(t), t\right) + \dot{\lambda}(t)s(t) \right] dt - \left[\lambda(T)s(T) - \lambda(0)s(0) \right].$$

Now consider now an arbitrary perturbation of $\{c(t)\}$ which will imply a change in $\{s(t)\}$ through the law of motion $g(\cdot)$. The implied change in the Lagrangean is

$$\Delta\mathcal{L} = \int_0^T \frac{\partial H(t)}{\partial c(t)} \Delta c(t) + \left[\frac{\partial H(t)}{\partial s(t)} + \dot{\lambda}(t) \right] \Delta s(t)$$

It follows that $\Delta\mathcal{L} = 0$ (so that we are on a maximum) for all perturbation paths of $c(t)$ and $s(t)$ only if at every instant t:

$$\frac{\partial H(t)}{\partial c(t)} = 0 \qquad \text{and} \qquad \frac{\partial H(t)}{\partial s(t)} = -\dot{\lambda}(t) \tag{3}$$

We have therefore obtained a set of conditions, (**??**) and (3), which together with the boundary conditions are **necessary** for the solution. This set of conditions constitutes the **Maximum Principle (MP)** that we restate more formally below:

Maximum Principle: An optimal solution to the program (CP) above is a triplet $\{s^*(t), c^*(t), \lambda^*(t)\}$ that, given $H(s, c, \lambda, t) = u(s, c, t) + \lambda g(s, c, t)$, must satisfy at each date t:

(i) $\dfrac{\partial H(s^*, c^*, \lambda^*, t)}{\partial c} = 0.$

(ii) $\dfrac{\partial H(s^*, c^*, \lambda^*, t)}{\partial s} = -\dot{\lambda}^*(t).$

(iii) $\dfrac{\partial H(s^*, c^*, \lambda^*, t)}{\partial \lambda} = \dot{s}^*(t).$(i.e $g(s, c, t) = \dot{s}^*(t)$)

(iv) boundary conditions $s(0) = s_0$, $s(t) = s_T$.

This is the great advantage of the MP: we were able to split the dynamic problem (CP) involving optimization over an entire interval of time into a sequence of simpler problems of static optimization that look all the same, so it is enough to maximize one Hamiltonian for a generic time t and we are done!

4.2 Sufficiency of the Maximum Principle

Once we add some concavity restrictions to the objective function and the constraint set, we are ready to prove the sufficiency of the conditions of the MP . .

Theorem 1: Suppose that the triplet $\{s^*, c^*, \lambda^*\}$ satisfies the conditions of the MP, and suppose the Hamiltonian evaluated at $\lambda = \lambda^*$ is jointly (strictly) concave in $\{s, c\}$, then $\{s^*, c^*, \lambda^*\}$ is the (unique) solution to the problem (CP).

Proof. Call V^* the value of the program for the candidate optimal path of the control and the states which satisfy the MP, (s^*, c^*). Call V the value corresponding to any other arbitrary *admissable* path of (s, c). We want to show that under the conditions of Theorem 1,

$$V^* - V = \int_0^T (u^* - u)dt \geq 0 \tag{4}$$

where the "strict" inequality holds under strict concavity. Defining $H^* = H(s^*, c^*, \lambda^*)$ and $H = H(s, c, \lambda^*)$, we obtain:

$$V^* - V = \int_0^T (u^* - u)dt = \int_0^T \left[\left(H^* - \lambda^* \dot{s}^* \right) - \left(H - \lambda^* \dot{s} \right) \right] dt.$$

Integrating by parts, one obtains:

$$\begin{aligned} V^* - V &= \int_0^T \left[\left(H^* + \dot{\lambda}^* s^* \right) - \left(H + \dot{\lambda}^* s \right) \right] dt - [\lambda^*(T)s^*(T) - \lambda^*(0)s^*(0)] \\ &\quad + [\lambda^*(T)s(T) - \lambda^*(0)s(0)] \\ &= \int_0^T \left[(H^* - H) + \dot{\lambda}^* (s^* - s) \right] dt \end{aligned}$$

where the last equality follows from that any admissable path must have the same initial and terminal conditions as the optimal path. From the above definition of concavity, it follows that:

$$\begin{aligned} \int_0^T \left[(H^* - H) + \dot{\lambda}^* (s^* - s) \right] dt &\geq \int_0^T \left[(s^* - s)H_s^* + (c^* - c)H_c^* + \dot{\lambda}^* (s^* - s) \right] dt \\ &= \int_0^T \left\{ (s^* - s) \left[H_s^* + \dot{\lambda}^* \right] + (c^* - c)H_c^* \right\} dt = 0 \end{aligned}$$

where the last equality follows from that all " $*$ " variables satisfy the conditions of the MP by assumption. Thus $V^* - V \geq 0$. ■

Remark: One can ensure the sufficiency of the MP from more primitive assumptions on u and g. In particular, it is immediate that if u is concave in (s, c), g is concave (convex) in (s, c), and $\lambda^* \geq 0$ $(\lambda^* \leq 0)$, then H evaluated at λ^* is jointly concave in (s, c).

Remark: We could have assumed that $c(t)$ is *piecewise-continuous*, i.e. we can have the control variables jumping a finite number of times and the whole derivation of the MP would still hold true. However, state variables cannot jump, i.e. must be continuous.

5 Generalizations of the Maximum Principle

5.1 Discounting

In several economic applications the return function is written with a discount factor $e^{-\rho t}$ with $\rho > 0$, to capture the fact that the future is given lower weight than the present in intertemporal decisions. The discounted control problem (DCP) can be written as:

$$\begin{aligned} &\max_{\{c(t)\}} \int_0^T e^{-\rho t} u(s(t), c(t)) dt \\ &\quad s.t. \\ &\dot{s}(t) = g(s(t), c(t)), \\ &s(0) = s_0, \;\; s(T) = s_T. \end{aligned} \tag{DCP}$$

It is clear that the only difference with (CP) is in the return function. From the usual definition of the (discounted) Hamiltonian, $H(s, c, \lambda) = e^{-\rho t} u(s(t), c(t)) + \lambda(t) g(s(t), c(t))$, we can apply the MP to obtain the familiar set of conditions:

$$\begin{aligned} \frac{\partial H}{\partial c} &= 0 \Rightarrow e^{-\rho t} \frac{\partial u}{\partial c} + \lambda(t) \frac{\partial g}{\partial c} = 0 \\ \frac{\partial H}{\partial s} &= -\dot{\lambda} \Rightarrow \left[e^{-\rho t} \frac{\partial u}{\partial s} + \lambda(t) \frac{\partial g}{\partial s} \right] = -\dot{\lambda}(t) \\ \frac{\partial H}{\partial \lambda} &= \dot{s} \Rightarrow g(s(t), c(t)) = \dot{s}(t) \end{aligned}$$

First notice that the third condition (the law of motion of the state) does not change, but the first two are modified. We define now the **current value multiplier** $\mu(t)$ as:

$$\mu(t) = e^{\rho t} \lambda(t) \Leftrightarrow \lambda(t) = e^{-\rho t} \mu(t),$$

and, through it, the **current value Hamiltonian** as:

$$H^c(s, c, \lambda) = u\left(s(t), c(t)\right) + \mu(t) g\left[s(t), c(t)\right].$$

Notice that: $e^{-\rho t}H^c(s,c,\lambda) = H(s,c,\lambda)$, so if we discount the current value H^c at rate ρ, we obtain the discounted Hamiltonian. Now, we aim at expressing the first two conditions of the MP in terms of H^c. From the first condition:

$$\frac{\partial H}{\partial c} = 0 \Rightarrow e^{-\rho t}\frac{\partial u}{\partial c} + e^{-\rho t}\mu(t)\frac{\partial g}{\partial c} = 0 \Rightarrow e^{-\rho t}\left[\frac{\partial u}{\partial c} + \mu(t)\frac{\partial g}{\partial c}\right] = 0 \Rightarrow \frac{\partial H^c}{\partial c} = 0$$

Therefore,

$$\frac{\partial H}{\partial c} = 0 \Leftrightarrow \frac{\partial H^c}{\partial c} = 0.$$

>From the second condition:

$$\begin{aligned}\frac{\partial H}{\partial s} = -\dot{\lambda}(t) \quad & \Leftrightarrow e^{-\rho t}\frac{\partial u}{\partial s} + e^{-\rho t}\mu(t)\frac{\partial g}{\partial s} = -e^{-\rho t}\dot{\mu}(t) + \rho e^{-\rho t}\mu(t) \\ & \Leftrightarrow \frac{\partial u}{\partial s} + \mu(t)\frac{\partial g}{\partial s} = -\dot{\mu}(t) + \rho\mu(t) \\ & \Leftrightarrow \frac{\partial H^c}{\partial s} - \rho\mu(t) = -\dot{\mu}(t)\end{aligned}$$

Therefore,

$$\frac{\partial H}{\partial s} = -\dot{\lambda}(t) \Leftrightarrow \frac{\partial H^c}{\partial s} = \rho\mu(t) - \dot{\mu}(t)$$

We can now state the:

Maximum Principle for the Discounted Case: An optimal solution to the program (DCP) above is a triplet $\{s^*(t), c^*(t), \mu^*(t)\}$ that given $H^c(s,c,\lambda,t) = u(s,c,t) + \mu g(s,c,t)$ must satisfy at each date t:

(i) $\dfrac{\partial H^c(s^*, c^*, \mu^*, t)}{\partial c} = 0.$

(ii) $\dfrac{\partial H^c(s^*, c^*, \mu^*, t)}{\partial s} = \rho\mu(t) - \dot{\mu}^*(t).$

(iii) $\dfrac{\partial H^c(s^*, c^*, \mu^*, t)}{\partial \mu} = \dot{s}^*(t).$

(iv) boundary conditions $s(0) = s_0$, $s(t) = s_T$.

5.2 Infinite Horizon

Consider the (DCP) without an explicit terminal condition for $s(T)$, but just with the additional constraint $s(T) \geq 0$. It is easy to prove, using the proof of necessity of the MP based on the perturbation argument, that the additional necessary condition of the MP is

$$e^{-\rho T}\mu(T)s(T) = 0.$$

This condition, called **Transversality Condition** (TVC), is very reminiscent of the complementary slackness conditions of static optimization. The meaning of this condition is that as we approach the end of the planning horizon, the optimal solution requires that either $s(T)$ is equal to zero (so nothing is waisted, as $s(T)$ gives utility) or $s(T)$ is positive, but nothing is waisted anyway because its shadow value $\mu(T)$ is zero (in the next section we explain why the costate can be interpreted as a shadow value of the state). With $T \to \infty$, the TVC becomes:

$$\lim_{t\to\infty} \lambda(t)s(t) = 0.$$

In the infinite horizon case, the TVC above replaces the terminal condition, whereas all other conditions of the MP remain equal.

Remark: The strangest things can happen when $t \to \infty$. See Barro, Sala-i-Martin, page 505, and the article by P. Michel on *Econometrica* in 1982 for a more formal treatment.

6 An Economic Interpretation of Optimal Control Theory

The MP can be seen as the dynamic generalization of the Lagrange method, so it should not come as a surprise that the dynamic Lagrange multipliers $\lambda(t)$ have a similar economic interpretation as **shadow prices**.

Consider the Lagrangean associated to the problem (CP) evaluated along the optimal path $\{s^*(t), c^*(t)\}$, for any arbitrary function $\lambda(t)$ of time:

$$\mathcal{L} = \int_0^T \left[u\left(s^*(t), c^*(t)\right) + \lambda(t)g\left(s^*(t), c^*(t)\right) - \lambda(t)\dot{s}^*(t)\right] dt$$

Using the definition of Hamiltonian and integrating by parts, we obtain:

$$
\begin{aligned}
\mathcal{L} &= \int_0^T \left[H\left(s^*, c^*, \lambda\right) + \dot{\lambda}(t) s^*(t) \right] dt - \left[\lambda(T) s_T - \lambda(0) s_0\right] \\
\frac{d\mathcal{L}}{ds_0} &= \int_0^T \left[H_s \frac{\partial s^*}{\partial s_0} + H_c \frac{\partial c^*}{\partial s_0} + H_\lambda \frac{\partial \lambda}{\partial s_0} + \dot{\lambda}(t) \frac{\partial s^*}{\partial s_0} + s^*(t) \frac{\partial \dot{\lambda}}{\partial s_0} \right] dt \\
&\quad - \frac{\partial \lambda(T)}{\partial s_0} s_T + \frac{\partial \lambda(0)}{\partial s_0} s_0 + \lambda(0)
\end{aligned}
$$

Since $\lambda(t)$ is an arbitrary function of time, we have that $\frac{\partial \lambda}{\partial s_0} = 0$ and $\frac{\partial \dot{\lambda}}{\partial s_0} = 0$. Now, if we evaluate this derivative at $\{\lambda^*(t)\}$:

$$
\frac{d\mathcal{L}^*}{ds_0} = \int_0^T \left\{ \left[H_s^* + \dot{\lambda}^*(t) \right] \frac{\partial s^*}{\partial s_0} + H_c^* \frac{\partial c^*}{\partial s_0} \right\} dt + \lambda^*(0) = \lambda^*(0)
$$

Where the last equality follows from that $\{s^*(t), c^*(t), \lambda^*(t)\}$ satisfy the conditions of the MP when the Hamiltonian is evaluated at λ^*. We conclude that $\lambda^*(0)$ is the marginal increase in the total value of the program due to a marginal increase in s_0, the initial value of the state.

Since the optimal control problem can be always divided into two parts at any time $t \geq 0$, then more in general, we can state that:

$$
\frac{d\mathcal{L}^*}{ds(t)} = \lambda^*(t),
$$

hence $\lambda^*(t)$ can be given the interpretation of shadow price of $s^*(t)$. A useful application of this result, due to Dorfman (1969), is in neoclassical capital theory.

Example 2: Neoclassical Capital Theory

Consider the problem of a firm which maximizes the sum of future profits $\pi(t)$ in the time interval $[0, T]$. Profits depend positively on the stock of capital $k(t)$ (as more can be produced and sold with more capital), and negatively on the flow of investments $x(t)$, as new investment is costly, so $\pi_k > 0$ and $\pi_x < 0$. Capital (the state) and investment (the control) are related by the law of motion $\dot{k}(t) = g\left(k(t), x(t)\right)$, for example $g(k, x) = x - \delta k$, and in general $g_k < 0$ and $g_x > 0$.

The problem for the firm is to choose optimally the time path of investment, i.e.

$$\max_{\{x(t)\}} \int_0^T \pi\left(k(t), x(t)\right) dt$$

$$s.t.$$

$$\dot{k}(t) = g\left(k(t), x(t)\right)$$

$$k(0) = k_0, \quad k(T) \geq 0$$

The intertemporal trade off involved in the investment choice is that investing today reduces the current profit, but increases the future profits by increasing future capital k.

To use the MP for this problem, we need to define the Hamiltonian: $H = \pi(k, x) + \lambda g(k, x)$. The 3 conditions of the MP (besides the law of motion) are:

$$(1)\ \ \frac{\partial H^*}{\partial x} = 0 \Rightarrow -\frac{\partial \pi^*}{\partial x} = \lambda^*(t)\frac{\partial g^*}{\partial x}.$$

Interpretation: Along the optimal path, the marginal decrease in profits due to an additional unit of investment is equal to the value (computed at the shadow price of capital) of the additional capital generated by the investment.

$$(2)\ \ \frac{\partial H^*}{\partial x} = -\dot{\lambda}^* \Rightarrow \frac{\partial \pi^*}{\partial k} + \lambda^*(t)\frac{\partial g^*}{\partial k} = -\dot{\lambda}^*(t).$$

Interpretation: $\dot{\lambda}^*(t)$ is the change in the value of capital in a small interval of time, or also the loss firms incur if the production of k is postponed for a short time. That condition above then says that along the optimal path each unit of capital good must decrease in value at the same rate at which it contributes to current profit and the value of cumulated capital.

$$(3)\ \ \text{TVC:} \quad \lambda(T)k(T) = 0.$$

Interpretation: In the limit the value of capital must be zero, either because $k = 0$ or because the marginal value of an additional unit is zero, i.e. $\lambda = 0$.

Example 1B: Optimal Lifetime Consumption

Consider the example described earlier, where a household has to choose optimally consumption $c(t)$ and savings $s(t)$ at every instant along a finite lifetime, given that every period he receives labour income y, and capital income $rw(t)$, where $w(t)$ is the cumulated wealth. Preferences are of the log family and there is no discounting. Initial wealth is zero, and final wealth must be zero. The problem of the household is

$$\max_{\{c(t)\}} \int_0^T \ln(c(t))dt$$
$$s.t.$$

$$\dot{w}(t) = s(t) + rw(t)$$

$$y = c(t) + s(t)$$

$$w(0) = 0, \qquad w(T) = 0.$$

Substitute out $s(t)$ from the two constraints, to end up with the **intertemporal budget constraint**:

$$\dot{w}(t) + c(t) = y + rw(t).$$

The state variable is $w(t)$,and the control variable is $c(t)$. The Hamiltonian for this problem is:

$$H = \ln(c(t)) + \lambda(t)\left[rw(t) + y - c(t)\right].$$

We can now apply the Maximum Principle:

$$H_c = 0 \Rightarrow \frac{1}{c(t)} = \lambda(t) \tag{5}$$

Interpretation: the utility of the marginal unit of consumption equals the shadow value of that same unit of the good if saved and added to the capital stock.

$$H_w = -\dot{\lambda} \Rightarrow \lambda(t)r = -\dot{\lambda}(t) \tag{6}$$

Interpretation: the return on the marginal unit of capital equals its loss of value over time $\frac{\dot{\lambda}(t)}{\lambda(t)}$.

$$H_\lambda = \dot{w} \Rightarrow \dot{w}(t) = rw(t) + y - c(t)$$

Differentiating (5) with respect to time[3] and using (6), one obtains:

$$\frac{\dot{c}(t)}{c(t)} = r$$

[3]Taking logs of (5):

$$\ln c(t) = -\ln \lambda(t)$$

Solving this simple first order differential equation, we obtain the optimal path of consumption up to a constant:

$$c(t) = \overline{c}e^{rt}. \tag{7}$$

Notice that for $t = 0$, $c(0) = \overline{c}$ so we know that the constant is nothing but the initial value of consumption at time zero. How do we solve for $c(0)$? We need to use the transition equation, which we still haven't used. Substituting (7) into the law of motion, we obtain:

$$\dot{w}(t) = rw(t) + y - c(0)e^{rt}.$$

If we premultiply the budget constraint above by e^{-rt}, we obtain:

$$\dot{w}(t)e^{-rt} - w(t)re^{-rt} = ye^{-rt} - c(0).$$

Where the left hand side (LHS) is clearly $\dfrac{\partial w(t)e^{-rt}}{\partial t}$. Hence, integrating out we obtain:

$$w(t)e^{-rt} = \overline{w} - y\frac{e^{-rt}}{r} - c(0)t. \tag{8}$$

Now, we can derive the expressions of the two constants of integrations $\{\overline{w}, c(0)\}$ from the boundary conditions.

For $t = 0$ (8) becomes: $0 = \overline{w} - y/r \Rightarrow \overline{w} = y/r$. Substituting this result into (7):

$$w(t)e^{-rt} = \left(\frac{1 - e^{-rt}}{r}\right) y - c(0)t.$$

For $t = T$ (8) becomes: $c(0) = \dfrac{y}{T}\left(\dfrac{1 - e^{-rT}}{r}\right)$. Hence, when we substitute $c(0)$ into (7) we obtain

$$c^*(t) = \frac{y}{T}\left(\frac{1 - e^{-rT}}{r}\right) e^{rt}, \tag{9}$$

which is the final expression for optimal consumption c^*.

Notice that consumption is increasing over time, so in the initial period of his life the household will save to guarantee his future consumption. Then, as the end of the horizon T becomes closer, he

---

and differentiating wrt time

$$\frac{1}{c(t)}\frac{\partial c(t)}{\partial t} = -\frac{1}{\lambda(t)}\frac{\partial \lambda(t)}{\partial t} \Leftrightarrow \frac{\dot{c}(t)}{c(t)} = -\frac{\dot{\lambda}(t)}{\lambda(t)}.$$

starts consuming more than y, hence he will be decumulating wealth. This is the typical life-cycle pattern: first saving, then dissaving.

It is interesting to notice that one can get the same result for $c(0)$ from the lifetime budget constraint:

$$\int_0^T ye^{-rt}dt = \int_0^T c(t)e^{-rt}dt$$

stating that the discounted value of income must equal the discounted value of consumption expenditures. The right hand side (RHS) above, using (7), can be rewritten as $\int_0^T c(0)dt = c(0)T$. It follows that:

$$c(0)T = y\int_0^T e^{-rt}dt = y\left[-\frac{e^{-rt}}{r}\right]_0^T = y\left[\frac{1-e^{-rt}}{r}\right]$$

which yields the same expression for $c(0)$.