MA2 – Advanced Macroeconomics, Spring 2000

Dr. Gianluca Violante

Notes on Discrete Time Stochastic Dynamic Programming

1. The Finite Horizon Case

Time is discrete and indexed by $t = 0, 1, ..., T < \infty$. The environment is stochastic. The uncertainty is introduced through z_t, an exogenous random variable (or shock) which follows a Markov process with transition function (or cumulative conditional distribution) $Q(z', z) = \Pr(z_{t+1} \leq z' | z_t = z)$ and with z_0 (initial value for the shock) given. We assume that z_t (but not z_{t+1}) is known at time t, so clearly z_t is a state of the economy.

The objective function is defined by the expected sum of instantaneous returns $u(x_t, c_t)$ discounted by the factor $\beta < 1$. This return function $u(\cdot)$ is a continuous and bounded function of the state variable x_t and of the control variable c_t. We assume $x_t \in X \subset \Re$ and $c_t \in C(x_t, z_t) \subset \Re$, hence we allow the feasible choice set C to depend on the pair of states of the economy (x_t, z_t).

The law of motion of the state x is: $x_{t+1} = f(x_t, z_t, c_t)$, with x_0 given. The action c_t in each period will depend only on the current states through the (possibly) time-varying function $g_t : X \times Z \to C$, i.e. $c_t = g_t(x_t, z_t), \forall t$. Given the time-separability of the objective function and the assumptions on the Markov process and the law of motion $f(\cdot)$, the pair (x_t, z_t) completely describe the state of the economy at any time t. The function g_t is referred to as decision rule. The sequence of decision rules $\pi_T = (g_0, g_1, ..., g_T)$ is called policy. A policy is feasible if each generic element g_t belongs to $C(x_t, z_t)$. It is said to be stationary if it does not depend on time, i.e. $g_t(x_t, z_t) = g(x_t, z_t), \forall t$.

Define now the expected discounted present value of following a given feasible policy π_T from the initial time 0 until the final time T as:

$$W_T(x_0, z_0, \pi_T) = E_0 \sum_{t=0}^{T} \beta^t u(x_t, g_t(x_t, z_t))$$

where it is understood that the expectation is taken with respect to $Q(z', z)$, that x_0, z_0 are given and that x_{t+1} follows $f(x_t, z_t, g_t(x_t, z_t))$. The dynamic programming (DP) problem is to choose π_T^* that maximizes W_T by solving:

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

We will abstract from most of the properties we should assume on Q to establish the main results. For all practical purposes it is sufficient to know that Q must satisfy the **Feller property**, a restriction which guarantees that the expectation function E is bounded and continuous in x_t.[1]

We want to know when the DP problem above has a solution. The following theorem does so.

Theorem of the Maximum: If the constraint set $C(x_t, z_t)$ is nonempty, compact and continuous, $u(\cdot)$ is continuous and bounded, $f(\cdot)$ is continuous, and Q has the Feller property, then there exists a solution to the problem above, called optimal policy π_T^* and the **value function** $V_T(x_0, z_0) = W_T(x_0, z_0, \pi_T^*)$ is also continuous.

Proof. See Stokey-Lucas, p. 62.

Rmk: Notice that the value function is the expected discounted present value of the optimal plan, i.e.

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

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

Given the existence of a solution, we can therefore write:

$$V_T(x_0, z_0) = \max_{\pi_T} E_0 \left\{ u(x_0, c_0) + \sum_{t=1}^{T} \beta^t u(x_t, c_t) \right\}.$$

[1] A careful explanation of these properties with an introduction to Markov Chains can be found in Stokey-Lucas, chapters 7,8,9.

By the law of iterated expectations: $E_0(x_1) = E_0(E_1(x_1))$, hence:

$$V_T(x_0, z_0) = \max_{\pi_T} E_0 \left\{ u(x_0, c_0) + E_1 \sum_{t=1}^{T} \beta^t u(x_t, c_t) \right\}.$$

Now let's cascade the max operator:

$$V_T(x_0, z_0) = \max_{c_0} E_0 \left\{ u(x_0, c_0) + \max_{\pi_{T-1}} E_1 \sum_{t=1}^{T} \beta^t u(x_t, c_t) \right\},$$

where $\pi_{T-1} = \{c_1, c_2, ..., c_T\}$. We can perform this step because the recursive structure of the problem guarantees that c_s will affect the dynamics of the state for t 1 s, but not for $t < s$. With some simple algebra on the discount factor, we obtain:

$$V_T(x_0, z_0) = \max_{c_0} E_0 \left\{ u(x_0, c_0) + \beta \max_{\pi_{T-1}} \underbrace{E_1 \sum_{t=1}^{T} \beta^{t-1} u(x_t, c_t)}_{W_{T-1}(x_1, z_1)} \right\}.$$

Hence, if we use the definition of V_{T-1} (the expected present value of the optimal policy with $T-1$ periods left) and pass through the expectation operator, we reach:

$$V_T(x_0, z_0) = \max_{c_0} \{ u(x_0, c_0) + \beta E_0 [V_{T-1}(x_1, z_1)] \}.$$

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 finite horizon (named after Richard Bellman (1956)):

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

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 reduces the choice of a sequence of decision rules to a sequence of choices for the control variable. 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 the optimal policy that maximizes $W_T(x_0, z_0, \pi_T)$, then if we consider the remainder of the objective function after s periods $W_{T-s}(x_s, z_s, \pi_{T-s})$, then the functions $\left\{g_s^*, g_{s+1}^*, ..., g_T^*\right\}$ which were optimal for the original problem, are still the optimal ones. Thus, 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.

1.1 Backward Induction

This is an algorithm to solve finite horizon DP problems.

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

$$V_0(x_T, z_T) = \max_{c_T \in C(x_T, z_T)} u(x_T, c_T)$$

   which yields the optimal choice $g_T^*(x_T, z_T)$ depending on the final value for x_T and the final realization of z_T. Hence, given a specification of $u(\cdot)$, we have an explicit functional form for $V_0(x_T, z_T)$.

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

$$\begin{aligned} V_1(x_{T-1}, z_{T-1}) = \quad & \max_{c_{T-1} \in C} \{u(x_{T-1}, c_{T-1}) + \\ & + \beta \int_Z V_0(f(x_{T-1}, z_{T-1,} c_{T-1}), z_T) \, dQ(z_T, z_{T-1})\right\} \end{aligned}$$

   which allows to solve once again for $g_{T-1}^*(x_{T-1}, z_{T-1})$ and to obtain $V_1(x_{T-1}, z_{T-1})$ explicitly.

3. We continue until time 0 and collect the sequence of decision rules into the optimal policy vector.

4. Given the initial conditions at time 0, we can reconstruct the whole optimal path for the state and the control, contingent on any realization of $\{z_t\}_{t=0}^T$.

2. The Infinite Horizon Case

When the horizon is infinite, i.e. $T \rightarrow \infty$, we cannot proceed with the backward induction algorithm to solve the problem, as there is no last period in which to start. The DP problem (conditional on the initial state) is the same at each point since we always have an infinite number of periods left to go, hence the environment is stationary. It follows that the value function $V(x, z)$ will be time invariant as well. In the infinite horizon case, we can write the Bellman's equation as:

$$V(x,z) = \max_{c \in C(x,z)} \left\{ u(x,c) + \beta \int_Z V(f(x,z,c), z') dQ(z', z) \right\} \tag{2}$$

where we have dropped the subindexes in the value function. The solution to this problem will be a stationary (time-invariant) policy function $c^* = g^*(x, z)$.

It will be useful to think of the above equation as a functional equation, i.e. an equation where the unknown is not a constant but a function $\Psi(x, z)$ belonging to some functional space B. The equation will be of the type:

$$T(\Psi) = \max_{c \in C(x,z)} \left\{ u(x,c) + \beta \int_Z \Psi(x', z') dQ(z', z) \right\} \tag{3}$$

where T is a mapping from B into B. Our objective is to find the solution to that equation, i.e. a function V belonging to the same functional space B that satisfies the fixed point property $V = T(V)$ displayed by the Bellman equation (2). We also want to try to establish whether equation (2) can be thought of as the limit of the DP problem with a finite horizon displayed in (1) as s approaches infinity.

To establish the existence of a solution and to characterize it, we therefore need to introduce some basic elements of functional analysis.[2]

2.1 Some Useful Mathematical Concepts

Definition: Metric Space. A metric space is a set S together with a metric or distance function $d : S \times S \rightarrow \Re_+$ satisfying the following conditions $\forall x, y, z \in S$:

1. $d(x, y) = 0 \; iff \; x = y$

[2] Most of them can be found in Chapters 3 and 4 of Stokey-Lucas (1989), and in the Appendix of Sargent (1987). For an introduction to the main concepts of real analysis, see Kolmogorov-Fomin (1970).

2. $d(x,y) = d(y,x)$
3. $d(x,z) \leq d(x,y) + d(y,z)$

Ex. 1: The Euclidean space $\Re^n$ equipped with the Euclidean metric is a metric space (PROVE IT).

Ex. 2: The space of continuous and bounded real functions defined over some set X, defined as:

$$\mathsf{B} \equiv \{\Psi : X \rightarrow \Re \mid \Psi \in \mathsf{C} \text{ and } m < \Psi(x) < M, \forall x \text{ for some } m, M \in \Re\}$$

equipped with the sup metric $d_\infty = d(\Psi_1, \Psi_2) = \sup_x |\Psi_1(x) - \Psi_2(x)|$ (where the sup operator is the supremum over the domain of the functions) is a metric space (PROVE IT).

Definition: Norm. If the point zero belongs to the metric space, then we can define the norm of y as the distance between the point y in the metric space and zero, and denote it as $||y||$.

Rmk: For instance the sup norm would be defined as $||\Psi||_\infty = \sup_x |\Psi(x)|$

Definition: Cauchy sequence. A sequence $\{y_n\}$ in a metric space (S,d) is called a Cauchy sequence if $\forall \varepsilon > 0, \exists\ N(\varepsilon)$ s.t. $d(y_n, y_m) < \varepsilon$ if $n, m \geq N(\varepsilon)$.

Rmk: From the definition it follows that a sequence is Cauchy if $\lim_{n,m\rightarrow\infty} d(y_n, y_m) = 0$. Hence, all convergent sequences are Cauchy sequences, but not the other way around (PROVE IT). An example is the sequence $y_n = 1/n$ defined on the set of positive real numbers which converges to zero, but zero does not belong to that metric space. The problem is that the metric space has a hole. We define now the set of metric spaces without holes.

Definition: Complete metric space. A metric space (S,d) is complete if every Cauchy sequence in (S,d) converges to a point in the space. This space is also called Banach space.

Rmk: The metric space (B, d_∞) —where B is the set of continuous and bounded functions and d_∞ is the sup metric– is a complete metric space (PROVE IT).

Definition: Operator. A function T mapping a metric space into itself is called an operator.

Definition: Contraction Mapping. Let (S, d) be a complete metric space and $T : (S, d) \rightarrow (S, d)$ the associated operator. Then T is said to be a contraction with modulus β if there is a number $\beta \in (0, 1)$ s.t.

$$\forall (y_1, y_2) \in (S, d)\,, \;\; d\left(T(y_1), T(y_2)\right) \leq \beta d\left(y_1, y_2\right)$$

Evidently, a contraction "contracts" any two points so that their images are closer together than the two points were. An example of contraction is the linear function $T : \Re \rightarrow \Re$ defined by $T(y) = a + by$ for $|b| < 1$ (PROVE IT).

We are now ready to state a theorem which is central in DP analysis.

Banach Fixed Point Theorem (Contraction Mapping Theorem): Let (S, d) be a complete metric space and let T be a contraction mapping with modulus β. Then, it follows that:

1. T has a unique fixed point y^* in S
2. For any $y^0 \in S$, the sequence $y^{n+1} = T(y^n)$ started at y^0 converges to y^* in the metric d.

Proof. See Sargent, pp. 343-344.

This is a very useful theorem because it ensures the existence and uniqueness of a fixed point under very general conditions and provides through point 2. an algorithm to compute the fixed point, by simple iteration. We now state a set of sufficient conditions to recognize a contraction mapping when dealing with particular metric spaces.

Blackwell Sufficient Conditions: Let T be an operator on a metric space (S, d_∞) where S is a space of functions with domain X and d_∞ is the sup metric. Then T is a contraction mapping with modulus β if it has the following properties:

- (monotonicity) For any pair of functions $(\psi_1, \psi_2) \in S$, $\psi_1(x)$ 1 $\psi_2(x)$ for all $x \Rightarrow T(\psi_1)$ 1 $T(\psi_2)$

- (discounting) For any function $\psi(x) \in S(X)$, positive real numbers $c > 0$, and $\beta \in (0,1)$, it is true that for any x on which ψ is defined: $T(\psi + c) \leq T(\psi) + \beta c$.

Proof. See Sargent, p. 345.

2.2 The value function in Bellman's equation as a fixed point of a contraction

Consider the functional equation in (3):

$$T\Psi\left(x,z\right) = \max_{c \in C(x,z)} \left\{u(x,c) + \beta E\left[\Psi(f(x,z,c),z')\right]\right\} \tag{4}$$

which defines a mapping T with domain equal to the space of functions Ψ. Let us restrict our attention to the space of functions Ψ which are continuous and bounded, call it B. We can establish two results:

1. Given the assumption on the boundedness of u and Ψ, it is immediate to show that $T\Psi$ is bounded as well.
2. By the Theorem of the Maximum, we obtain that if Ψ is continuous, then $T\Psi$ is continuous as well (just interpret Ψ as V_{s-1} and $T\Psi$ as V_s).

Now equip the space B with the sup norm d_∞. We know that (B,d_∞) is a metric space. Hence the mapping T defined above maps this metric space into itself, and it is an operator (call it the Bellman's operator). Being (B,d_∞) a complete metric space (as we have established earlier), we can use Blackwell's sufficient conditions to check if T is a contraction mapping. Let's check both conditions:

- (monotonicity) Take $\Psi_1 \geq \Psi_2$. Then:

$$\begin{aligned} T\Psi_1 &= \max_c \left\{u(x,c) + \beta E\Psi_1(f(x,z,c),z')\right\} \\ &\geq \max_c \left\{u(x,c) + \beta E\Psi_2(f(x,z,c),z')\right\} \\ &= T\Psi_2 \end{aligned}$$

- (discounting) For any function ψ, positive real numbers $c > 0$, and $\beta \in (0,1)$

$$
\begin{aligned}
T(\Psi + a) &= u(x, c^*) + \beta E\left[\Psi(f(x, z, c^*), z') + a\right] \\
&= u(x, c^*) + \beta E\left[\Psi(f(x, z, c^*), z')\right] + \beta a \\
&= T\Psi + \beta a
\end{aligned}
$$

It follows that T is a contraction mapping, hence it satisfies Banach fixed point theorem.

Rmk: We could have proved it by applying directly the definition of contraction mapping (PROVE IT).

We can therefore use the contraction mapping theorem in characterizing the solution of the DP problem in infinite horizon. This is extremely useful because we can interpret the value function of the infinite horizon DP problem in (2) as the fixed point of a contraction mapping. In particular:

- We know that the infinite horizon Bellman's equation

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

  has a solution V and this solution is unique under general conditions.

- From point 2. of the contraction mapping theorem, we can interpret the infinite horizon value function as the limit as $T \to \infty$ of a finite horizon DP problem.
- Point 2. of the theorem also gives us a useful algorithm to compute the fixed point, as we can start iterating on (3) from any initial (continuous and bounded) function Ψ_0, and we are certain to converge to the solution V.
- Contractions have the feature of guaranteeing that the fixed point V (weakly) preserves the properties of the functions Ψ on which we iterate, such as monotonicity, continuity and concavity.

2.3 Characterization of the Policy Function

In the solution of Bellman's equation:

$$
\begin{aligned}
&V(x, z) = \max_{c \in C(x,z)} \left\{ u(c) + \beta \textstyle\int_Z V(x', z') dQ(z', z) \right\} \\
&s.t. \quad x' = f(x, z, c)
\end{aligned}
$$

the fixed point value function V is associated to an optimal policy function $c^* = g^*(x, z)$. We now characterize the optimal policy function with the usual methods of calculus, i.e. by differentiation, like in the static case or in the optimal control problem in continuous time.

Suppose we can invert the function f to obtain $c = \varphi(x', x, z)$. Then we substitute out the control variable in the Bellman's equation and the choice becomes on the state next period x'. Hence, the problem becomes:

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

If we knew that V were differentiable, by taking the FOC with respect to x' we would obtain:

$$u_1(x', x, z) - \beta \int_Z V_1(x', z') dQ(z', z) = 0 \tag{5}$$

While we are free to make assumptions on u and f, given these assumptions the differentiability of V must be established. The following result shown by Bienveniste and Scheinkman (1979) is useful in this regard.

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 $V_1(x_0) = \Omega_1(x_0)$.

Proof. See Stokey-Lucas, p. 85.

Let's apply this result to the Bellman's equation. Define:

$$\Omega(x, z) = u(x, g^*(x_0, z)) + \beta \int_Z V(f(x_0, z, g^*(x_0, z)), z') dQ(z', z)$$

It follows that $\Omega(x_0, z) = V(x_0, z)$ and $\forall x \in N(x_0)$:

$$\Omega(x, z) \leq \max_{c \in C(x,z)} \left\{ u(x, c) + \beta \int_Z V(f(x, z, c), z') dQ(z', z) \right\} = V(x, z)$$

since $g^*(x_0, z)$ is not the optimal policy function for $x \neq x_0$. Assuming that $u(\cdot)$ is concave and differentiable implies that $\Omega(x, z)$ is concave and differentiable as well, since the integral element is just a constant. We can therefore apply the envelope theorem to obtain that:

$$V_1(x,z) = \Omega_1(x,z) = u_x(x, g^*(x,z))$$

which tells us that the derivative of the value function is the partial derivative of the return function with respect to the state variable evaluated at the optimal value for the control.

We can use this result to substitute out the derivative of the value function from (5) to obtain:

$$u_1(x', x, z) - \beta \int u_1(x', c') dQ(z', z) = 0$$

If we substitute out next period control through $c' = \varphi(x'', x', z')$, we obtain the Euler equation:

$$u_1(x', x, z) - \beta \int u_1(x', x'', z') dQ(z', z)$$

To stress that this equation is satisfied by the optimal sequence of states $\{x_t^*\}_{t=0}^{\infty}$, we can re-express it with time subscripts as:

$$u_1(x_{t+1}^*, x_t^*, z_t) - \beta \int u_1(x_{t+1}^*, x_{t+2}^*, z_{t+1}) dQ(z_{t+1}, z_t) \tag{6}$$

This is a second order 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. This conditions are:

- 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.

2.4 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 DP 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.[3]

- First, let's assume concavity and differentiability of the return function u. Then from concavity it follows that:

$$\begin{aligned} D &= \lim_{T\to\infty} \sum_{t=0}^{T} \beta^t \left[u(x_t^*, x_{t+1}^*) - u(x_t, x_{t+1})\right] \\ &\geq \lim_{T\to\infty} \sum_{t=0}^{T} \beta^t \left[u_1(x_t^*, x_{t+1}^*)(x_t^* - x_t) + u_2(x_t^*, x_{t+1}^*)(x_{t+1}^* - x_{t+1})\right] \end{aligned}$$

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

$$\begin{aligned} D \geq \lim_{T\to\infty} \quad & \left\{\sum_{t=0}^{T-1} \beta^t \left[u_2(x_t^*, x_{t+1}^*) + \beta u_1(x_{t+1}^*, x_{t+2}^*)\right] (x_{t+1}^* - x_{t+1}) \right. \\ & \left. + \beta^T u_2(x_T^*, x_{T+1}^*)(x_{T+1}^* - x_{T+1})\right\} \end{aligned}$$

- Since the expression in the square brackets $[\cdot]$ satisfies the Euler equation by assumption, each term of the sum is zero. Using the Euler equation (6) 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 $u_1 > 0$ and $x_{T+1} > 0$, then:

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

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

3. The Stochastic Growth Model

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

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

$$\max_{\{c_t,i_t\}} E_0\left[\sum_{t=0}^{\infty}\beta^t u(c_t)\right]$$
$$s.t.$$
$$k_{t+1} = (1-\delta)\,k_t + i_t$$
$$f\,(k_t, z_t) = c_t + i_t$$
$$k_0, z_0,\ Q(z_{t+1}, z_t) \text{ given}$$

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

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

The FOC and the Envelope condition for this problem are respectively:

- (FOC): $u_1\,(f(k,z) - k' + (1-\delta)k) = \beta\int V_1(k',z')dQ(z',z)$
- (ENV): $V_1(k,z) = u_k\,(\cdot) = u_1\,(\cdot)\,[f_1(k,z) + (1-\delta)]$

Therefore, we obtain the Euler Equation:

$$u_1\,(f(k,z) - k' + (1-\delta)k) = \beta\int u_1\,(f(k',z') - k'' + (1-\delta)k')\,[1 + f_1(k',z') - \delta]\,dQ$$

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 characterize 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.

4. References

[1] Bertsekas, D. (1976), Dynamic Programming and Stochastic Control, New York: Academic Press.

[2] Harris M. (1987), Dynamic Economic Analysis, Oxford, UK: Oxford University Press.

[3] Kolmogorov A.N., F.V. Fomin (1970), *Introductory Real Analysis*, New York: Dover Publications, Inc.

[4] Sargent T. (1987), *Dynamic Macroeconomic Theory*, Cambridge, MA: Harvard University Press.

[5] Stokey N., R.E. Lucas and E. Prescott (1989), *Recursive Methods in Economic Dynamics*, Cambridge, MA: Harvard University Press.