Fall Semester '05-'06
Akila Weerapana

Lecture 24: Dynamic Optimization in Continuous Time

I. INTRODUCTION

- In the last few lectures we used the Bellman equation to solve discrete time dynamic optimization problems faced by consumers and firms. Today's lecture looks at techniques for solving dynamic optimization problems in continuous time. The techniques for solving continuous time problems came from a Russian mathematician named Pontryagin in the 1950s. Since that body of knowledge became accessible to economists, there has been a boom in dynamic models particularly in macroeconomics.
- We will first sketch out what the solution in continuous time looks like using the Bellman formulation we have been using. We will then adopt an easier technique to get the same equations in the continuous time case.
- Before we explore the solution techniques, be aware that the distinction between continuous time and discrete time problem manifests itself in both the control variable and the state variable.
- In a continuous time dynamic optimization problem, the agent is optimizing over an integral function of the control variable with the evolution of the state variable being described by a differential equation. In contrast, in the discrete time case we were minimizing a sum of functions of the control variable with the evolution of the state variable being given by a difference equation.

II. CONTINUOUS TIME OPTIMIZATION DECISIONS

- A continuous time maximization problem is typically of the form

$$\max_{x_t} \int_0^T \left[e^{-\rho t} f(x_t, A_t)\right] dt$$

subject to the constraint $\dot{x}_t = g(x_t, A_t)$.

- In this case, x is the choice variable, A is the state variable and f, g are arbitrary functions that describe the objective function and the evolution of the state variable respectively.
- ρ here is the discount rate, it is a non-negative number, with $\rho = 0$ denoting the case where we don't discount the future at all and the limiting case of $\rho \to \infty$ being where where we ignore the future completely.

Drawing Parallels to the Discrete Time Case

- Let's draw some parallels to the discrete case here. We can set up an approximate version of the above problem in discrete time as

$$\max_{x_t} \sum_{t=0}^{T} \left(\frac{1}{1+\rho}\right)^t f(x_t, A_t)$$

subject to the constraint $A_{t+1} = A_t + g(x_t, A_t)$. Note that the transition equation is slightly different here than what we typically worked with, simply so that we can draw a parallel between the evolution of the state variable in the case of discrete time, $A_{t+1} - A_t = g(x_t, A_t)$, and the evolution of the state variable in the case of continuous time, $\dot{A}_t = g(x_t, A_t)$.

- We also used the discount rate ρ here instead of the discount factor β that we usually use in the discrete time case. This does not really change anything, remember that $\beta = \frac{1}{1+\rho}$. When we don't discount the future at all we have $\rho = 0$ which corresponds to $\beta = 1$. When we discount the future completely we have $\rho \to \infty$, equivalent to $\beta = 0$.

- The value function in discrete time will be

$$V_t(A_t) = \max_{x_t} f(x_t, A_t) + \left(\frac{1}{1+\rho}\right) V_{t+1}(A_{t+1}) \text{ where } A_{t+1} = A_t + g(x_t, A_t)$$

- The FOC and the envelope conditions are

$$\begin{aligned} 0 &= f_x(x_t, A_t) + \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1}) g_x(x_t, A_t) \\ V'_t(A_t) &= f_A(x_t, A_t) + \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1})[1 + g_A(x_t, A_t)] \end{aligned}$$

- If we rearrange the envelope condition around a little (no particular reason why you would do so, just were you to do so) you would get

$$\begin{aligned} V'_t(A_t) - \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1}) &= f_A(x_t, A_t) + \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1})[g_A(x_t, A_t)] \\ \frac{(1+\rho)V'_t(A_t) - V'_{t+1}(A_{t+1})}{1+\rho} &= f_A(x_t, A_t) + \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1})[g_A(x_t, A_t)] \\ \frac{V'_t(A_t) - V'_{t+1}(A_{t+1})}{1+\rho} - \frac{\rho V'_t(A_t)}{1+\rho} &= f_A(x_t, A_t) + \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1})[g_A(x_t, A_t)] \end{aligned}$$

- The above rearrangement of the envelope condition is, of course, completely unnecessary to the solution in the discrete time case. However, what you will soon see is that the solution technique for the continuous time case generates two equations that are identical to the FOC and the rearranged envelope condition above.

Solving the Continuous Time Problem Using The Hamiltonian

- In order to find the solution to the continuous time optimization problem, we first define a special function known as the **Hamiltonian function**. We define the Hamiltonian as

$$H_t = f(x_t, A_t) + \mu_t \dot{A}_t$$

- μ_t is known as a co-state variable. Essentially, μ_t is a Lagrange multiplier term that basically tells us how much the inter-temporal objective function increases when we add one more unit of the state variable at time t.
- The Hamiltonian function has an economic interpretation. H_t basically describes the total impact on the inter-temporal objective function of choosing a particular level of the choice variable at time t.
- Picking a particular value of the choice variable has both a direct and indirect effect on the value of the inter-temporal objective function over the specified time period. The direct impact is captured by the $f(x_t, A_t)$ term.
- In addition, the choice variable also has an indirect effect on the value of the inter-temporal objective function over the specified time period by changing the level of the state variable, in the amount $\dot{A}_t$.
- Since μ_t tells us how much the objective function can be increased by one more unit of the state variable at the margin, it follows that $\mu_t \dot{A}_t$ summarizes the total indirect impact on the objective function.
- In other words, H_t summarizes the TOTAL impact of picking x_t on the inter-temporal objective function.
- The conditions for an optimum come next. Since H_t summarizes the total impact of choosing a given level of the choice variable on the inter-temporal utility function, it follows that the optimal solution is characterized by a FOC of the form $\frac{\partial H_t}{\partial x_t} = 0$.
- As in the Bellman equation, there is another, less intuitive condition involving the state variable. In the continuous time case this equation is that

$$\frac{\partial H_t}{\partial A_t} = \rho \mu_t - \dot{\mu}_t$$

- This is called the co-state equation. The Russian mathematician Lev Pontryagin derived these two conditions as part of what is known as the **Maximum Principle** characterizing the solution to a dynamic optimization problem.
- The final two constraints are of course the initial condition: $A_0 = \underline{A}$ and the terminal condition, $A_T = \overline{A}$. If the horizon is infinite, the terminal condition becomes $\lim_{T \to \infty} A_T \mu_T = 0$. This basically states that in the limit either there has to be no units of the state variable left over ($A_T = 0$) or if there are units left over, then their value has to be zero in terms of maximizing utility ($\mu_T = 0$).

- To summarize, given a general function of the form $\max_{x_t} \int_0^T \left[e^{-\rho t} f(x_t, A_t)\right]$ subject to the constraint $\dot{A}_t = g(x_t, A_t)$ we define the Hamiltonian as $H_t = f(x_t, A_t) + \mu_t \dot{A}_t$ and the conditions for a solution to be

$$\frac{\partial H_t}{\partial x_t} = 0 \quad \Rightarrow \quad f_x(x_t, A_t) + \mu_t g_x(x_t, A_t) = 0$$

$$\frac{\partial H_t}{\partial A_t} = \rho\mu_t - \dot{\mu}_t \quad \Rightarrow \quad f_A(x_t, A_t) + \mu_t g_A(x_t, A_t) = \rho\mu_t - \dot{\mu}_t$$

- This is combined with the initial condition $A_0 = \underline{A}$ and either a terminal condition $A_T = \overline{A}$ or a transversality condition $\lim_{T\to\infty} \mu_T A_T = 0$ to form a system of **differential** equations that characterize the solution to the problem.
- Comparing the FOC for the Hamiltonian case given above with the FOC for the Bellman case given earlier, we can see that they are equivalent if

$$\mu_t = \left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1})$$

- Why does this make sense? Well, we defined $V_t(A_t)$ as the optimized value of the lifetime objective function so $\left(\frac{1}{1+\rho}\right) V'_{t+1}(A_{t+1})$ is the change in the (discounted) maximized value of the objective function when we add 1 more unit of the state variable for next period, i.e. the equivalent of μ_t.
- Using the above equivalency, we can rewrite the envelope condition as

$$[\mu_{t-1} - \mu_t] + \rho\mu_{t-1} = f_A(x_t, A_t) + \mu_t g_A(x_t, A_t)$$

- If we move from discrete time to continuous time, where the difference between two time periods becomes infinitesimally small, the above envelope condition becomes

$$-\dot{\mu_{t-1}} + \rho\mu_{t-1} = f_A(x_t, A_t) + \mu_t g_A(x_t, A_t)$$

- When time periods are not as distinct as in the continuous time case, the above can be thought of as being analogous to

$$-\dot{\mu}_t + \rho\mu_t = f_A(x_t, A_t) + \mu_t g_A(x_t, A_t)$$

- This last condition is Pontryagin's co-state equation. So even though you did not have intuition about where that equation came from at the time I presented it, hopefully this exercise has shown that it at least parallels the Bellman equation.

Side Issue: Present Value Hamiltonians vs. Current Value Hamiltonians

- The following discussion is only for information purposes for those of you who may need to refer back to these notes later on in life.

- The Hamiltonian we presented was in current value terms, i.e. it represented the impact of making a decision in terms of utility at time t.
- Sometimes, the Hamiltonian is presented in **present-value** or discounted terms as

$$H_t = e^{-\rho t}\left[f(x_t, A_t) + \mu_t \dot{A}\right] \equiv \left[e^{-\rho t} f(x_t, A_t) + \lambda_t \dot{A}\right]$$

where $\lambda_t = e^{-\rho t}\mu_t$ is the Lagrange multiplier in time 0 present value (discounted to time 0) terms.

- In this case, the conditions for optimization are

$$\begin{aligned}
\frac{\partial H_t}{\partial x_t} &= 0 \Rightarrow e^{-\rho t} f_x(x_t, A_t) + \lambda_t g_x(x_t, A_t) = 0 \\
\frac{\partial H_t}{\partial A_t} &= -\dot{\lambda}_t \Rightarrow e^{-\rho t} f_A(x_t, A_t) + \lambda_t g_A(x_t, A_t) = -\dot{\lambda}_t
\end{aligned}$$

- Given the definition of $\lambda_t = e^{-\rho t}\mu_t$, you can calculate $\dot{\lambda}_t = e^{-\rho t}\dot{\mu}_t - \rho e^{-\rho t}\mu_t$ and show that these conditions are identical to the ones we derived

III. APPLICATIONS OF HAMILTONIANS

Cake Eating Problem

- Let's again begin with the "cake-eating" problem. Define Φ_t to be the size of a cake at time t. The problem is

$$\max_{C_t} \int_0^T \left[e^{-\rho t} U(C_t)\right] dt$$

subject to the constraint $\dot{\Phi}_t = -C_t$ and Φ_0, Φ_T are given

- The choice variable here is consumption, C_t, while the state variable is the size of the cake, Φ_t.
- We will define the Hamiltonian as

$$H_t = U(C_t) + \mu_t[\dot{\Phi}_t] \equiv U(C_t) + \mu_t[-C_t]$$

where μ_t is known as the co-state variable.

- The FOCs, according to the maximum principle are

$$\begin{aligned}
\frac{\partial H_t}{\partial C_t} = 0 \quad &\Rightarrow \quad U'(C_t) = \mu_t \\
\frac{\partial H_t}{\partial \Phi_t} = \rho\mu_t - \dot{\mu}_t \quad &\Rightarrow \quad 0 = \rho\mu_t - \dot{\mu}_t
\end{aligned}$$

- Combining

$$\begin{aligned} \rho U'(C_t) &= U''(C_t)\dot{C}_t \\ \Rightarrow \dot{C}_t &= \rho \left[\frac{U'(C_t)}{U''(C_t)}\right] \end{aligned}$$

- This is the Euler equation for the model. Recall that in the discrete case, the Euler equation states that at the optimum choices, we cannot gain utility by making a feasible switch of consumption from one period to the next. Since in continuous time we don't have distinct periods, we have to modify the underlying intuition a little.

- We can do this by rearranging the above equation as

$$\frac{\dot{C}_t U''(C_t)}{U'(C_t)} = \rho \Rightarrow \frac{d\ln(U'(C_t))}{dt} = \rho$$

  This equation states that the optimal path of consumption is such that the growth rate of marginal utility over time is equal to the discount rate.

- This is exactly analogous to the Euler equation for the cake-eating problem where we had the condition

$$U'(C_t) = \beta U'(C_{t+1}) \Rightarrow \frac{U'(C_{t+1}) - U'(C_t)}{U'(C_t)} = \frac{1-\beta}{\beta}$$

  which given $\beta = \frac{1}{1+\rho}$ collapses down to

$$\frac{U'(C_{t+1}) - U'(C_t)}{U'(C_t)} = \rho$$

- In other words, the interpretation of the Euler equation as saying that the growth rate of marginal utility over time is equal to the discount rate is exactly analogous to the interpretation that feasible reallocations of consumption yield no added utility gain.

- Given the utility function $U(C_t) = 2\sqrt{C_t}$ we have $U' = \frac{1}{\sqrt{C_t}}$ and $U'' = \frac{-1}{2C_t^{3/2}}$. We can rewrite the above as

$$\frac{\rho}{\sqrt{C_t}} = \frac{-1}{2C_t\sqrt{C_t}}\dot{C}_t$$

- Rearranging, we have $\dot{C}_t = -2\rho_t C_t \Rightarrow \frac{\dot{C}_t}{C_t} = -2\rho_t$

- This equation, combined with the equation describing the evolution of the size of the cake $\dot{\Phi}_t = -C_t$, give a system of differential equations in C and Φ which can be solved given the initial and terminal conditions $\Phi_0 = \bar{\Phi}$ and $\Phi_T = 0$.

Utility Maximization

- A continuous time analogy to the discrete time consumer optimization problem we solved is

$$\max_{C_t} \int_0^T \left[e^{-\rho t} U(C_t) dt\right]$$

subject to the constraint $\dot{A}_t = rA_t + Y_t - C_t$.

- Also suppose we are given some initial value of A, lets call that $A_0 = \underline{A}$ and a terminal value of A, $A_T = \overline{A}$.
- We can define the Hamiltonian as

$$H_t = U(C_t) + \mu_t \dot{A}_t \equiv U(C_t) + \mu_t[rA_t + Y_t - C_t]$$

- The first order conditions are

$$\begin{aligned}
\frac{\partial H}{\partial C_t} = 0 \qquad & \Rightarrow U'(C_t) - \mu_t = 0 \\
\frac{\partial H}{\partial A_t} = \rho\mu_t - \dot{\mu}_t \qquad & \Rightarrow r\mu_t = \rho\mu_t - \dot{\mu}_t \\
A_0 &= \underline{A} \\
A_T &= \overline{A}
\end{aligned}$$

- Combining the first two equations gives us the Euler equation for consumption in continuous time.

$$\begin{aligned}
U'(C_t) - \mu_t = 0 \quad & \Rightarrow \quad \mu_t = U'(C_t) \\
r\mu_t = \rho\mu_t - \dot{\mu}_t \quad & \Rightarrow \quad r\left[U'(C_t)\right] = \rho U'(C_t) - U''(C_t)\dot{C}_t
\end{aligned}$$

- This simplifies to the following Euler equation

$$U''(C_t)\dot{C}_t = (\rho - r)U'(C_t)$$

- Given the utility function $U(C_t) = 2\sqrt{C_t}$ we have $U' = \frac{1}{\sqrt{C_t}}$ and $U'' = \frac{-1}{2C_t^{3/2}}$.
- We can rearrange the above equation so that it reads

$$\frac{-\dot{C}}{C_t\sqrt{C_t}} = 2(\rho - r)\frac{1}{\sqrt{C_t}} \Rightarrow \frac{\dot{C}}{C_t} = 2(r - \rho)$$

- The growth rate of consumption is related to the difference between the rate at which we discount the future and the interest rate as well as to the elasticity of substitution. If the interest rate and the rate at which we discount the future are identical, i.e. $r - \rho$ then consumption will remain constant over time (analogous to the $\beta(1 + r) = 1$ case in discrete time optimization).
- If on the other hand, the interest rate exceeds the rate at which we discount the future, i.e. $r > \rho$ then consumption will be rising over time (analogous to the $\beta(1+r) > 1$ case in discrete time optimization).
- Finally, if the interest rate is less than the rate at which we discount the future, i.e. $r < \rho$ then consumption will be falling over time (analogous to the $\beta(1 + r) < 1$ case in discrete time optimization).

- Note that the Euler equation $\frac{\dot{C}}{C} = 2(r - \rho)$ along with the flow budget constraint $\dot{A}_t = rA_t + Y_t - C_t$ and the initial condition $A_0 = \underline{A}$ and terminal condition $A_T = \overline{A}$ constitute a system of differential equations that make up the solution to the problem.