Continuous Time Dynamic Programming

I. The Optimization Problem

We consider the problem of maximizing

$$E\left[\int_0^\infty U(C_t,K_t)e^{-\beta t}dt\right] \tag{1}$$

subject to

$$dK_t = \hat{K}(C_t,K_t)\ dt + \sigma_K(C_t,K_t)\ dW \tag{2}$$

and

$$h(C_t,K_t) \le 0\ , \tag{3}$$

where (2) and (3) hold for all t in $(0,\infty)$. In solving the problem we take K_0 as given and known, while the paths of C and thereby K_t for t>0 are subject to choice. We assume that C and K paths must be chosen so that decisions made at t depend only on information available at t. That is, our choice at t must be expressed as a mapping from what we will know at t, which is $\{C_s,K_s,W_s,\ s\le t\}$, to a real vector C_t of choices made at t. Of course (2) then determines K's behavior at t. We assume K and C are both vectors, so σ_K is a matrix.

II. Recursive Structure

The problem defined by (1)-(3) starts at time 0. We could as easily define the problem as starting at some arbitrary date t. Furthermore, for every initial value of K for which the problem has a solution, there will be some corresponding value of the objective function, V(K). That is, there is a function

$$V(K) = E\left[\int_t^\infty U(C^*_{t+s},K_{t+s})e^{-\beta s}ds\right]\ , \tag{4}$$

where the * indicates that C^* is chosen optimally. The function V is called the **value function**. Note that the optimal choice of C_t in optimizing (4) must depend only on K_t, not on other information available at t, since (4) itself and the constraints do not pend on the past except through K_t. We write $C_t^* = \gamma(K_t)$, and γ is the **policy function** solving the dynamic programming problem.

Now observe that (1) can be written as

$$E\left[\int_0^\infty U(C_t,K_t)e^{-\beta t}dt\right] = E\left[\int_0^T U(C_t,K_t)e^{-\beta t}dt\right] + e^{-\beta T}E\left[\int_T^\infty U(C_{T+s},K_{T+s})e^{-\beta s}ds\right] . \tag{5}$$

The last term on the right of (5) has the same structure as the original integral on the left of (5), except that it is shifted in time. Since our choices of C_{T+s} for $s>0$ have no effect on the value of the first integral on the right of (5), we can maximize the overall integral in two steps. First, maximize the second term on the right of (5), taking K_T as given. Then, taking account of how K_T affects what utility is obtainable from the second term on the right of (5), maximize over the (0,T) interval. This is **Bellman's principle** and can be stated more concisely as the assertion that the problem (1)-(3) is equivalent to the problem of maximizing

$$E\left[\int_0^T U(C_t,K_t)e^{-\beta t}dt + e^{-\beta T}V(K_T)\right] . \tag{6}$$

with respect to C and K subject to (2) and (3).

III. Conditions Determining an Optimum

Now assuming that, under the optimal choice of C, the objective function (1) has a well-defined finite value, we can define a new stochastic process

$$Z_t = E_t\left[\int_0^\infty U(\gamma(K_s),K_s)e^{-\beta s}ds\right] . \tag{7}$$

The process Z_t, being of the form E_tX, is by construction a martingale and therefore has

$$\hat{Z}_t = 0 \quad . \tag{8}$$

But now using (6) and Ito's lemma, we can form another expression for $\hat{Z}_t$, i.e.

$$\hat{Z}_t = e^{-\beta t}\left(U(\gamma(K),K) - \beta V(K) + D_K V\cdot\hat{K} + \frac{1}{2}\,\mathrm{tr}\left(\sigma_K' \; D_{KK}V \; \sigma_K\right)\right) = 0 \; . \tag{9}$$

Of course to form (9) we must assume that V is twice differentiable. This provides us with a necessary condition for V to be the value function and γ to be the optimal policy function, namely

$$U(\gamma(K),K) - \beta V(K) + D_K V\cdot\hat{K} + \frac{1}{2}\,\mathrm{tr}\left(\sigma_K' \; D_{KK}V \; \sigma_K\right) = 0 \; , \tag{10}$$

for all possible values of K.

Equation (10) by itself is not an equation we can solve for V or γ. To arrive at a soluble equation, we apply **Bellman's equation**, which is for this setup

$$\max_{h(C,K)\leq 0} \left\{ U - \beta V + D_K V\cdot\hat{K} + \frac{1}{2}\,\mathrm{tr}\left(\sigma_K' \; D_{KK}V \; \sigma_K\right)\right\} = 0. \tag{11}$$

If, for some V and γ, (11) and (12) hold for every possible value of K, then (subject to some regularity conditions) V is the value function and γ is the optimal policy. To see this, suppose ϕ is a stochastic process for C that makes it depend only on past information and, that, when it is used in (2) to generate K, satisfies (3). Let

$$\tilde{K}_0 = K_0 \tag{12}$$

and

$$d\tilde{K} = \hat{K}(\phi,\tilde{K})\, dt + \sigma_K(\phi,\tilde{K})dW \; . \tag{13}$$

Further let

$$Q_t = \int_0^t U(\phi_s,\tilde{K}_s)\; e^{-\beta s}\, ds + e^{-\beta t}V(\tilde{K}_t) \; . \tag{14}$$

Notice that $Q_0=V(K_0)$ and that

$$EQ_t \to E\left[\int_0^\infty U(\phi_s,\tilde{K}_s)\; e^{-\beta s}\; ds\right] . \tag{15}$$

In asserting (15) we are assuming an important regularity condition that actually fails in some versions of the first simple example we will consider below, the linear-quadratic permanent income model. The condition is that for any feasible K

$$E\left[e^{-\beta t}V(K_t)\right] \underset{t\to\infty}{\longrightarrow} 0 . \tag{16}$$

Now

$$\hat{Q} = e^{-\beta t}\left(U(\phi,\tilde{K}) - \beta V + D_K V\ \hat{K} + \frac{1}{2}\ \mathrm{tr}\left(\sigma_K\ D_{KK}\ \ \sigma_K V'\ \right)\right), \tag{17}$$

where the implicit C and K arguments in (17) are ϕ and $\tilde{K}$. Obviously if (11) holds for all possible values of K, $\hat{Q}\le 0$, all t. This, with (15), implies that ϕ as the choice for C yields a value of the objective function no greater than $V(K_0)$. Since this holds for any Φ, $V(K_0)$ is in fact the maximum attainable value of the objective function and is attained with the policy rule γ.

Note that while (11) is a standard, relatively easy to remember form for the Bellman equation, it suppresses, for the sake of simple notation, the important points that U, σ_K and $\hat{K}$ are both functions of C and K, while V depends on K alone.

If U, $\hat{K}$, and σ_K are differentiable in C, (11) implies the first-order condition

$$D_C U(C,K) + D_K V\cdot D_C\hat{K}(C,K) + .5\ \mathrm{tr}\left(D_{KK}V\ D_C\left[\sigma_K(C,K)\sigma_K'(C,K)\right]\right) = \mu\ D_C h , \tag{18}$$

where μ is a Kuhn-Tucker multiplier that vanishes for $h(C,K)<0$.

For the case of a one-dimensional K (18) takes the simpler form

$$D_C U(C,K) + V'\ D_C\hat{K}(C,K) + V''\ \sigma_K(C,K)\ D_C\sigma_K(C,K) = \mu\ D_C h . \tag{19}$$

Equation (18) or (19) can be solved, in principle, for C as a function of K, V_K, and V_{KK}. Substituting this back into (10) gives us a differential equation in V, V_K,

V_{KK}, and K alone. This is a second order differential equation (partial differential equation in the case of a non-scalar K) in V, and can in principle be solved, given appropriate boundary conditions.

Often in finding a solution it is helpful, instead of combining (10) with (18) to obtain a differential equation in V, to combine (18) with the derivative of (10) with respect to K. This is useful because often the resulting system can be reduced to a differential equation in γ, which is often more directly useful than V. Also, the derivative of (10) with respect to K has an interpretation as a differential equation in time (rather than K) that can sometimes be combined with (2) to obtain a system of differential equations with respect to time in C and K. This may allow us to proceed directly to solutions for or characterizations of C and K as functions of time, which may again be of more direct interest than either V or γ.

Differentiating (10) with respect to K gives

$$D_K U - \beta D_K V + D_K V\ D_K \hat{K} + .5\mathrm{tr}\left[D_{KK} V\ D_K\left(\sigma_K' \sigma_K\right)\right] \tag{20}$$
$$+ D_{KK} V\ \hat{K} + .5\ \mathrm{tr}\left(D_{KKK} V\ \sigma_K' \sigma_K\right)$$
$$+ D_C U\ \gamma' + D_K V\ D_C \hat{K}\ \gamma' + .5\mathrm{tr}\left[D_{KK} V\ D_C\left(\sigma_K' \sigma_K\right)\right]\ \gamma' = 0\ .$$

Note that when μ=0, (18) above implies that the last line of (20) (all the terms involving γ') is identically zero. When $\mu \neq 0$, the constraint $h(\gamma(K),K)=0$ holds, so

$$D_C h\ \gamma' = -D_K h\ . \tag{21}$$

Thus the terms on the last line of (20) always are equal to $-\mu D_K h$. This fact and the resulting simplification of (20) is what is known as the **envelope theorem**. It leaves us with

$$D_K U - \beta D_K V + D_K V\ D_K \hat{K} + .5\mathrm{tr}\left[D_{KK} V\ D_K\left(\sigma_K' \sigma_K\right)\right] \tag{22}$$
$$+ D_{KK} V\ \hat{K} + .5\ \mathrm{tr}\left(D_{KKK} V\ \sigma_K' \sigma_K\right) = -\mu\ D_K h\ .$$

Note that Ito's lemma implies that the last two terms on the left of (22) are exactly $D_{\hat{K}} V$, so that (22) can be written

$$-D_{\hat{K}}V = D_K U - \beta D_K V + D_K V \, D_K \hat{K} + .5\mathrm{tr}\left(D_{KK}V \, D_K\left(\sigma_K' \sigma_K\right)\right) + \mu \, D_K h \tag{23}$$

In fact, if we define $\lambda = D_K V$, we can rewrite (18) and (23) as

$$D_C U + \lambda \cdot D_C \hat{K} + .5 \, \mathrm{tr}\left(D_K \lambda \, D_C\left[\sigma_K \sigma_K'\right]\right) - \mu \, D_C h = 0 \tag{24}$$

$$-\hat{\lambda} = D_K U - \beta\lambda + \lambda \, D_K \hat{K} + .5\mathrm{tr}\left(D_K \lambda \, D_K\left(\sigma_K' \sigma_K\right)\right) + \mu \, D_K h \, . \tag{25}$$

In the deterministic case, where $\sigma_K \equiv 0$, or in any other case where σ_K is constant, (2), (24) and (25) form the usual Hamiltonian first-order conditions. While this system is still occasionally of some use in interpreting stochastic problems, it is not so directly useful because of the appearance of $D_K\lambda$ in the stochastic version. This prevents the system from being interpreted as a set of differential equations in the time paths of C, K and λ.

IV. The Linear-Quadratic Permanent Income Model

Let's apply what we've developed in the preceding sections to the continuous time version of the standard permanent income model. We consider the problem of maximizing

$$E\left[\int_0^\infty (C - .5C^2) e^{-\beta t} dt\right] \tag{26}$$

subject to

$$dA = (rA + \bar{Y} - C) \, dt + \sigma \, dW \, . \tag{27}$$

Specializing (24) and (25) to this case yields

$$1 - C = -\lambda \tag{28}$$

$$-\hat{\lambda} = (r - \beta)\lambda \, . \tag{29}$$

Combining (28) and (29) gives us

$$\hat{C} = (r - \beta)(1 - C) \tag{30}$$

which in the leading case of $r=\beta$ produces Hall's conclusion that consumption is a martingale. Rewriting (28) in terms of $\gamma(K)$ (to get back to (23) as applied to this problem) gives us

$$\gamma' \hat{A} + .5\, \gamma'' \sigma^2 = (r-\beta)(1-\gamma) \;, \tag{31}$$

an ordinary (nonlinear) second-order differential equation in γ.

Equation (31) has at least two very simple solutions. One is $\gamma(K)\equiv 1$. This is the policy of setting C at its satiation level forever. The other makes γ linear, so that $\gamma''=0$ and $\gamma=a+bA$ for some a and b. It is easy to check that (31) then implies

$$b\,(rA + \bar{Y} - a - bA) = (r - \beta)\,(1 - a - bA) \;. \tag{32}$$

From (32) we conclude that if $b\neq 0$, $b=2r-\beta$. In that case we can conclude further that $a=(2\beta/r-1)\bar{Y}+\beta/r-1$. In the special case $\beta=r$, this reduces to the rule $C=rA+\bar{Y}$.

These two solutions both display what is called **certainty equivalence**. That is, because the term in σ^2, the only one affected by the presence of uncertainty, disappears from (31) for these solutions, they are solutions also to the version of the problem that has $\sigma^2=0$. We could have found these solutions by ignoring the presence of uncertainty, and they would nonetheless have been correct when uncertainty was introduced. This result is obviously a special case. It always arises when U is quadratic in its arguments, $\hat{K}$ is linear in its arguments, and there is no side constraint h, the **linear-quadratic** case.

[A complete version of these notes would go on to discuss how the presence of two solutions that work for all K can be reconciled with the optimality principle and what to make of all the other, nonlinear solutions to (31).]