Dynamic Optimization in Continuous-Time Economic Models
(A Guide for the Perplexed)

Maurice Obstfeld*

University of California at Berkeley

First Draft: April 1992

*I thank the National Science Foundation for research support.

I. Introduction

The assumption that economic activity takes place continuously is a convenient abstraction in many applications. In others, such as the study of financial-market equilibrium, the assumption of continuous trading corresponds closely to reality. Regardless of motivation, continuous-time modeling allows application of a powerful mathematical tool, the theory of optimal dynamic control.

The basic idea of optimal control theory is easy to grasp--indeed it follows from elementary principles similar to those that underlie standard static optimization problems. The purpose of these notes is twofold. First, I present intuitive derivations of the first-order necessary conditions that characterize the solutions of basic continuous-time optimization problems. Second, I show why very similar conditions apply in deterministic and stochastic environments alike. [1]

A simple unified treatment of continuous-time deterministic and stochastic optimization requires some restrictions on the form that economic uncertainty takes. The stochastic models I discuss below will assume that uncertainty evolves continuously according to a type of process known as an Itô (or Gaussian

---

[1] When the optimization is done over a finite time horizon, the usual second-order sufficient conditions generalize immediately. (These second-order conditions will be valid in all problems examined here.) When the horizon is infinite, however, some additional "terminal" conditions are needed to ensure optimality. I make only passing reference to these conditions below, even though I always assume (for simplicity) that horizons are infinite. Detailed treatment of such technical questions can be found in some of the later references.

diffusion) process. Once mainly the province of finance theorists, Itô processes have recently been applied to interesting and otherwise intractable problems in other areas of economics, for example, exchange-rate dynamics, the theory of the firm, and endogenous growth theory. Below, I therefore include a brief and heuristic introduction to continuous-time stochastic processes, including the one fundamental tool needed for this type of analysis, Itô's chain rule for stochastic differentials. Don't be intimidated: the intuition behind Itô's Lemma is not hard to grasp, and the mileage one gets out of it thereafter truly is amazing.

II. Deterministic Optimization in Continuous Time

The basic problem to be examined takes the form: Maximize

$$(1) \qquad \int_0^\infty e^{-\delta t}U[c(t),k(t)]dt$$

subject to

$$(2) \qquad \dot{k}(t) = G[c(t),k(t)], \qquad k(0) \text{ given,}$$

where U(c,k) is a strictly concave function and G(c,k) is concave. In practice there may be some additional inequality constraints on c and/or k; for example, if c stands for consumption, c must be nonnegative. While I will not deal in any detail with such constraints, they are straightforward to incorporate.[2] In general, c and k can be vectors, but I will concentrate on the notationally simpler scalar case. I call c the *control* variable for the optimization problem and k the *state* variable. You should think of the control variable as a flow--for example, consumption per unit time--and the state variable as a stock--for example, the stock of capital, measured in units of consumption.

The problem set out above has a special structure that we can exploit in describing a solution. In the above problem, planning starts at time $t = 0$. Since no exogenous variables enter (1) or (2), the maximized value of (1) depends *only* on $k(0)$, the predetermined initial value of the state variable. In other words, the problem is *stationary*, i.e., it does not change in form with the passage of time.[3] Let's denote this maximized value by $J[k(0)]$, and call $J(k)$ the *value function* for the problem. If $\{c^*(t)\}_{t=0}^{\infty}$ stands for the associated optimal path of the control and $\{k^*(t)\}_{t=0}^{\infty}$ for that of the state,[4] then by definition,

---

[2] The best reference work on economic applications of optimal control is still Kenneth J. Arrow and Mordecai Kurz, *Public Investment, the Rate of Return, and Optimal Fiscal Policy* (Baltimore: Johns Hopkins University Press, 1970).

[3] Nonstationary problems often can be handled by methods analogous to those discussed below, but they require additional notation to keep track of the exogenous factors that are changing.

[4] According to (2), these are related by

$$k^*(t) = \int_0^t G[c^*(s),k^*(s)]ds + k(0).$$

$$J[k(0)] = \int_0^\infty e^{-\delta t} U[c^*(t),k^*(t)]dt.$$

The nice structure of this problem relates to the following property. Suppose that the optimal plan has been followed until a time $T > 0$, so that $k(T)$ is equal to the value $k^*(T)$ given in the last footnote. Imagine a *new* decision maker who maximizes the discounted flow of utility from time $t = T$ onward,

$$(3) \qquad \int_T^\infty e^{-\delta(t-T)} U[c(t),k(t)]dt,$$

subject to (2), but with the intial value of k given by $k(T) = k^*(T)$. *Then the optimal program determined by this new decision maker will coincide with the continuation, from time T onward, of the optimal program determined at time 0, given* $k(0)$*.* You should construct a proof of this fundamental result, which is intimately related to the notion of "dynamic consistency."

You should also convince yourself of a key implication of this result, that

$$(4) \qquad J[k(0)] = \int_0^T e^{-\delta t} U[c^*(t),k^*(t)]dt + e^{-\delta T} J[k^*(T)],$$

where $J[k^*(T)]$ denotes the maximized value of (3) given that $k(T) = k^*(T)$ and (2) is respected. Equation (4) implies that we can think of our original, $t = 0$, problem as the finite-horizon problem of maximizing

$$\int_0^T e^{-\delta t} U[c(t),k(t)]dt + e^{-\delta T} J[k(T)]$$

subject to the constraint that (2) holds for $0 \leq t \leq T$. Of course, in practice it may not be so easy to determine the correct functional form for J(k), as we shall see below!

Nonetheless, this way of formulating our problem--which is known as *Bellman's principle of dynamic programming*--leads directly to a characterization of the optimum. Because this characterization is derived most conveniently by starting in discrete time, I first set up a discrete-time analogue of our basic maximization problem and then proceed to the limit of continuous time.

Let's imagine that time is carved up into discrete intervals of length h. A decision on the control variable c, which is a flow, sets c at some fixed level per unit time over an entire period of duration h. Furthemore, we assume that changes in k, rather than accruing continuously with time, are "credited" only at the very end of a period, like monthly interest on a bank account. We thus consider the problem: Maximize

$$(5) \quad \sum_{t=0}^{\infty} e^{-\delta th} U[c(t),k(t)]h$$

subject to

(6) $k(t+h) - k(t) = hG[c(t),k(t)], \qquad k(0)$ given.

Above, c(t) is the fixed rate of consumption over period t while k(t) is the given level of k that prevails from the very end of period $t - 1$ until the very end of t. In (5) [resp. (6)] I have multiplied U(c,k) [resp. G(c,k)] by h because the cumulative flow of utility [resp. change in k] over a period is the product of a fixed instantaneous *rate* of flow [resp. *rate* of change] and the period's length.

Bellman's principle provides a simple approach to the preceding problem. It states that the problem's value function is given by

$$(7) \quad J[k(t)] = \max_{c(t)} \left\{ U[c(t),k(t)]h + e^{-\delta h} J[k(t+h)] \right\},$$

subject to (6), for any initial k(t). It implies, in particular, that optimal c*(t) must be chosen to maximize the term in braces. By taking functional relationship (7) to the limit as $h \to 0$, we will find a way to characterize the continuous-time optimum. [5]

We will make four changes in (7) to get it into useful form. First, subtract J[k(t)] from both sides. Second, impose the

---

[5] All of this presupposes that a well-defined value function *exists*--something which in general requires justification. (See the extended example in this section for a concrete case.) I have also not proven that the value function, when it does exist, is *differentiable*. We know that it will be for the type of problem under study here, so I'll feel free to use the value function's first derivative whenever I need it below. With somewhat less justification, I'll also use its second and third derivatives.

constraint (6) by substituting for k(t+h), k(t) + hG[c(t),k(t)]. Third, replace $e^{-\delta h}$ by its power-series representation, $1 - \delta h + (\delta^2 h^2)/2 - (\delta^3 h^3)/6 + \ldots$. Finally, divide the whole thing by h. The result is

$$(8) \quad 0 = \max_c \left\{U(c,k) - [\delta - (\delta^2 h/2) + \ldots]J[k + hG(c,k)] + \{J[k + hG(c,k)] - J(k)\}/h\right\},$$

where implicitly all variables are dated t. Notice that

$$\frac{J[k + hG(c,k)] - J(k)}{h} = \frac{\{J[k + hG(c,k)] - J(k)\}G(c,k)}{G(c,k)h}.$$

It follows that as $h \to 0$, the left-hand side above approaches $J'(k)G(c,k)$. Accordingly, we have proved the following

PROPOSITION II.1. *At each moment, the control* c* *optimal for maximizing (1) subject to (2) satisfies the Bellman equation*

$$(9) \quad 0 = U(c^*,k) + J'(k)G(c^*,k) - \delta J(k) = \max_c \left\{U(c,k) + J'(k)G(c,k) - \delta J(k)\right\}.$$

This is a very simple and elegant formula. What is its interpretation? As an intermediate step in interpreting (9), define the *Hamiltonian* for this maximization problem as

(10) $\mathcal{H}(c,k) \equiv U(c,k) + J'(k)G(c,k).$

In this model, the intertemporal tradeoff involves a choice between higher current c and higher future k. If c is consumption and k wealth, for example, the model is one in which the utility from consuming now must continuously be traded off against the utility value of savings. The Hamiltonian $\mathcal{H}(c,k)$ can be thought of as a measure of the flow value, in current utility terms, of the consumption-savings combination implied by the consumption choice c, given the predetermined value of k. The Hamiltonian solves the problem of "pricing" saving in terms of current utility by multiplying the flow of saving, $G(c,k) = \dot{k}$, by $J'(k)$, the effect of an increment to wealth on total lifetime utility. A corollary of this observation is that $J'(k)$ has a natural interpretation as the shadow price (or marginal current utility) of wealth. More generally, leaving our particular example aside, $J'(k)$ is the shadow price one should associate with the state variable k.

This brings us back to the Bellman equation, equation (9). Let c^* be the value of c that maximizes $\mathcal{H}(c,k)$, given k, which is arbitrarily predetermined and therefore might not have been chosen optimally.[6] Then (9) states that

(11) $\mathcal{H}(c^*,k) = \max_{c} \{\mathcal{H}(c,k)\} = \delta J(k).$

---

[6]It is important to understand clearly that at a given point in time t, k(t) is not an object of choice (which is why we call it a state variable). Variable c(t) can be chosen freely at time t (which is why it is called a control variable), but its level influences the *change* in k(t) over the next infinitesimal time interval, $k(t + dt) - k(t)$, not the current value k(t).

In words, the maximized Hamiltonian is a fraction δ of an optimal plan's total lifetime value. Equivalently, the instantaneous value flow from following an optimal plan divided by the plan's total value--i.e., the plan's rate of return--must equal the rate of time preference, δ. Notice that if we were to measure the current payout of the plan by $U(c^*,k)$ alone, we would err by not taking proper account of the value of the current increase in k. This would be like leaving investment out of our measure of GNP! The Hamiltonian solves this accounting problem by valuing the increment to k using the shadow price $J'(k)$.

To understood the implications of (9) for optimal consumption we must go ahead and perform the maximization that it specifies (which amounts to maximizing the Hamiltonian). As a by-product, we obtain the Pontryagin necessary conditions for optimal control.

Maximizing the term in braces in (9) over c, we get [7]

$$(12)\quad U_c(c^*,k) = -G_c(c^*,k)J'(k).$$

The reason this condition is necessary is easy to grasp. At each moment the decision maker can decide to "consume" a bit more, but at some cost in terms of the value of current "savings." A unit of additional consumption yields a marginal payoff of $U_c(c^*,k)$, but at the same time, savings change by $G_c(c^*,k)$. The utility

---

[7] I assume interior solutions throughout.

value of a marginal *fall* in savings thus is $-G_c(c^*,k)J'(k)$; and if the planner is indeed at an optimum, it must be that this marginal cost just equals the marginal current utility benefit from higher consumption. In other words, unless (12) holds, there will be an incentive to change c from c*, meaning that c* cannot be optimal.

Let's get some further insight by exploiting again the recursive structure of the problem. It is easy to see from (12) that for any date t the optimal level of the control, $c^*(t)$, depends only on the inherited state $k(t)$ (regardless of whether $k(t)$ was chosen optimally in the past). Let's write this functional relationship between optimal c and k as $c^* = c(k)$, and assume that $c(k)$ is differentiable. (For example, if c is consumption and k total wealth, $c(k)$ will be the household's consumption function.) Functions like $c(k)$ will be called *optimal policy functions*, or more simply, policy functions. Because $c(k)$ is defined as the solution to (9), it automatically satisfies

$$0 = U[c(k),k] + J'(k)G[c(k),k] - \delta J(k).$$

Equation (12) informs us as to the optimal relation between c and k at a point in time. To learn about the implied optimal behavior of consumption over time, let's differentiate the preceding equation with respect to k:

$$0 = [U_c(c^*,k) + J'(k)G_c(c^*,k)]c'(k) + U_k(c^*,k)$$
$$+ [G_k(c^*,k) - \delta]J'(k) + J''(k)G(c^*,k).$$

The expression above, far from being a hopeless quagmire, is actually just what we're looking for. Notice first that the left-hand term multiplying $c'(k)$ drops out entirely thanks to (12): another example of the envelope theorem. This leaves us with the rest,

$$(13)\quad U_k(c^*,k) + J'(k)[G_k(c^*,k) - \delta] + J''(k)G(c^*,k) = 0.$$

Even the preceding simplified expression probably isn't totally reassuring. Do not despair, however. A familiar economic interpretation is again fortunately available.

We saw earlier that $J'(k)$ could be usefully thought of as the shadow price of the state variable k. If we think of k as an asset stock (capital, foreign bonds, whatever), this shadow price corresponds to an asset price. Furthermore, we know that under perfect foresight, asset prices adjust so as to equate the asset's total instantaneous rate of return to some required or benchmark rate of return, which in the present context can only be the time-preference rate, δ. As an aid to clear thinking, let's introduce a new variable, λ, to represent the shadow price $J'(k)$ of the asset k:

$$\lambda \equiv J'(k).$$

Our next step will be to rewrite (13) in terms of λ. The key observation allowing us to do this concerns the last term on the right-hand side of (13), $J''(k)G(c^*,k)$. The chain rule of calculus implies that

$$J''(k)G(c^*,k) = \frac{dJ'(k)}{dk} \times \frac{dk}{dt} = \frac{d\lambda}{dk} \times \frac{dk}{dt} = \frac{d\lambda}{dt} = \dot{\lambda};$$

and with this fact in hand, it is only a matter of substitution to express (13) in the form

$$\frac{U_k + \lambda G_k + \dot{\lambda}}{\lambda} = \delta. \tag{14}$$

This is just the asset-pricing equation promised in the last paragraph.

Can you see why this last assertion is true? To understand it, let's decompose the total return to holding a unit of stock k into "dividends" and "capital gains." The "dividend" is the sum of two parts, the direct effect of an extra unit of k on utility, U_k, and its effect on the rate of increase of k, λG_k. (We must multiply G_k by the shadow price λ in order to express the *physical* effect of k on $\dot{k}$ in the same terms as U_k, that is, in terms of utility.) The "capital gain" is just the increase in the price of k, $\dot{\lambda}$. The sum of dividend and capital gain, divided by the asset price λ, is just the rate of return on k, which, by

(14) must equal δ along an optimal path.

Example

Let's step back for a moment from this abstract setting to consolidate what we've learned through an example. Consider the standard problem of a consumer who maximizes $\int_0^\infty e^{-\delta t}U[c(t)]dt$ subject to $\dot{k} = f(k) - c$ (where c is consumption, k capital, and f(k) the production function). Now $U_k = 0$, $G(c,k) = f(k) - c$, $G_c = -1$, and $G_k = f'(k)$. In this setting, (14) becomes the statement that the rate of time preference should equal the marginal product of capital plus the rate of accrual of utility capital gains,

$$\delta = f'(k) + \frac{\dot{\lambda}}{\lambda}.$$

Condition (12) becomes $U'(c) = \lambda$. Since this last equality implies that $\dot{\lambda} = U''(c)\dot{c}$, we can express the optimal dynamics of c and k as a nonlinear differential-equation system:

$$(15) \quad \dot{c} = -\frac{U'(c)}{U''(c)}\left[f'(k) - \delta\right], \qquad \dot{k} = f(k) - c.$$

You can see the phase diagram for this system in figure 1. (Be sure you can derive it yourself! The diagram assumes that $\lim_{k\to\infty} f'(k) = 0$, so that a steady-state capital stock exists.) The diagram makes clear that, given k, *any* positive initial c initiates a path along which the two preceding differential equations for c and k are respected. *But not all of these paths are optimal, since the differential equations specify conditions that are merely necessary, but not sufficient, for optimality.*

Indeed, only one path will be optimal in general: we can write the associated policy function as as $c^* = c(k)$ (it is graphed in figure 1). For given k, paths with initial consumption levels exceeding c(k) imply that k becomes negative after a finite time interval. Since a negative capital stock is nonsensical, such paths are not even feasible, let alone optimal. Paths with initial consumption levels below c(k) imply that k gets to be too large, in the sense that the individual could raise lifetime utility by eating some capital and never replacing it. These "overaccumulation" paths violate a sort of terminal condition stating that the present value of the capital stock should converge to zero along an optimal path. I shall not take the time to discuss such terminal conditions here.

If we take

$$U(c) = \frac{c^{1-(1/\varepsilon)} - 1}{1 - (1/\varepsilon)}, \qquad f(k) = rk,$$

where ε and r are positive constants. we can actually find an algebraic formula for the policy function c(k).

Let's conjecture that optimal consumption is proportional to wealth, that is, that $c(k) = \eta k$ for some constant η to be determined. If this conjecture is right, the capital stock k will follow $\dot{k} = (r - \eta)k$, or, equivalently,

$$\frac{\dot{k}}{k} = r - \eta.$$

This expression gives us the key clue for finding η. If $c = \eta k$, as we've guessed, then also

$$\frac{\dot{c}}{c} = r - \eta.$$

But necessary condition (15) requires that $\frac{\dot{c}}{c} = \varepsilon(r - \delta)$, which contradicts the last equation unless

(16) $\eta = (1 - \varepsilon)r + \varepsilon\delta.$

Thus, $c(k) = [(1 - \varepsilon)r + \varepsilon\delta]k$ is the optimal policy function. In the case of log utility ($\varepsilon = 1$), we simply have $\eta = \delta$. We get the same simple result if it so happens that r and δ are equal.

Equation (16) has a nice interpretation. In Milton Friedman's permanent-income model, where $\delta = r$, people consume the annuity value of wealth, so that $\eta = r = \delta$. This rule results in a *level* consumption path. When $\delta \neq r$, however, the optimal consumption path will be tilted, with consumption rising over time if $r > \delta$ and falling over time if $r < \delta$. By writing

(16) as

$$\eta = r - \varepsilon(r - \delta)$$

we can see these two effects at work. Why is the deviation from the Friedman permanent-income path proportional to ε? Recall that ε, the elasticity of intertemporal substitution, measures an individual's willingness to substitute consumption today for consumption in the future. If ε is high and $r > \delta$, for example, people will be quite willing to forgo present consumption to take advantage of the relatively high rate of return to saving; and the larger is ε, *certeris paribus*, the lower will be η. Alert readers will have noticed a major problem with all this. If $r > \delta$ and ε is sufficiently large, η, and hence "optimal" consumption, will be *negative*. How can this be? Where has our analysis gone wrong?

The answer is that when $\eta < 0$, *no optimum consumption plan exists!* After all, nothing we've done demonstrates existence: our arguments merely indicate some properties that an optimum, *if* one exists, will need to have.

No optimal consumption path exists when $\eta < 0$ for the following reason. Because optimal consumption growth necessarily satisfies $\dot{c}/c = \varepsilon(r - \delta)$, and $\varepsilon(r - \delta) > r$ in this case, optimal consumption would have to grow at a rate exceeding the rate of return on capital, r. Since capital growth obeys $\dot{k}/k = r - (c/k)$, however, and $c \geq 0$, the growth rate of capital, and hence that of output, is *at most* r. With consumption positive and growing at 3 percent per year, say, but with capital growing at a lower yearly rate, consumption would eventually grow to be greater than total output--an impossibility in a closed economy. So the proposed path for consumption is not feasible. This means that no feasible path--other than the obviously suboptimal path with $c(t) = 0$, for all t--satisfies the necessary conditions for optimality. Hence, no feasible path is optimal: no optimal path exists.

Let's take our analysis a step further to see how the value function J(k) looks. Observe first that at any time t, $k(t) = k(0)e^{(r-\eta)t} = k(0)e^{\varepsilon(r-\delta)t}$, where k(0) is the starting capital stock and η is given by (16). Evidently, the value function at t = 0 is just

$$J[k(0)] = \left[1 - \frac{1}{\varepsilon}\right]^{-1}\left\{ \int_0^\infty e^{-\delta t}[\eta k(t)]^{1-(1/\varepsilon)}dt - \frac{1}{\delta} \right\}$$

$$= \left[1 - \frac{1}{\varepsilon}\right]^{-1}\left\{ \int_0^\infty e^{-\delta t}[\eta k(0)e^{\varepsilon(r-\delta)t}]^{1-(1/\varepsilon)}dt - \frac{1}{\delta} \right\}$$

$$= \left[1 - \frac{1}{\varepsilon}\right]^{-1}\left\{\frac{[\eta k(0)]^{1-(1/\varepsilon)}}{\delta - (\varepsilon - 1)(r - \delta)} - \frac{1}{\delta}\right\} .$$

So the value function J(k) has the same general form as the utility function, but with k in place of c. This is not the last time we'll encounter this property. Alert readers will again notice that to carry out the final step of the last calculation, I had to assume that the integral in braces above is convergent, that is, that $\delta - (\varepsilon - 1)(r - \delta) > 0$. Notice, however, that $\delta - (\varepsilon - 1)(r - \delta) = r - \varepsilon(r - \delta) = \eta$, so the calculation is valid provided an optimal consumption program exists. If one doesn't, the value function clearly doesn't exist either: we can't specify the maximized value of a function that doesn't attain a maximum. This counterexample should serve as a warning against blithely assuming that all problems have well-defined solutions and value functions.

Return now to the theoretical development. We have seen how to solve continuous-time determinstic problems using Bellman's method of dynamic programming, which is based on the value function $J(k)$. We have also seen how to interpret the derivative of the value function, $J'(k)$, as a sort of shadow asset price, denoted by λ. The last order of business is to show that we have actually proved a simple form of Pontryagin's *Maximum Principle*:[8]

PROPOSITION II.2. (Maximum Principle) *Let* $c^*(t)$ *solve the problem of maximizing (1) subject to (2). Then there exist variables* $\lambda(t)$*--called costate variables--such that the Hamiltonian*

---

[8]First derived in L.S. Pontryagin et al., *The Mathematical Theory of Optimal Processes* (New York and London: Interscience Publishers, 1962).

$$\mathcal{H}[c,k(t),\lambda(t)] \equiv U[c,k(t)] + \lambda(t)G[c,k(t)]$$

is maximized at $c = c^*(t)$ *given* $\lambda(t)$ *and* $k(t)$*; that is,*

(17) $$\frac{\partial\mathcal{H}}{\partial c}(c^*,k,\lambda) = U_c(c^*,k) + \lambda G_c(c^*,k) = 0$$

at all times (assuming, as always, an interior solution). Furthermore, the costate variable obeys the differential equation

(18) $$\dot{\lambda} = \lambda\delta - \frac{\partial\mathcal{H}}{\partial k}(c^*,k,\lambda) = \lambda\delta - [U_k(c^*,k) + \lambda G_k(c^*,k)]$$

for $\dot{k} = G(c^*,k)$ *and* $k(0)$ *given*. [9]

---

[9]You should note that if we integrate differential-equation (18), we get the *general* solution

$$\lambda(t) = \int_t^\infty e^{-\delta(s-t)}\frac{\partial\mathcal{H}}{\partial k}[c^*(s),k(s),\lambda(s)]ds + Ae^{\delta t},$$

where A is an arbitrary constant. [To check this claim, just differentiate the foregoing expression with respect to t: if the integral in the expression is I(t), we find that $\dot{\lambda} = \delta I - (\partial\mathcal{H}/\partial k) + \delta Ae^{\delta t} = \delta\lambda - (\partial\mathcal{H}/\partial k)$.] I referred in the prior example to an additional terminal condition requiring the present value of the capital stock to converge to zero along an optimal path. Since $\lambda(t)$ is the price of capital at time t, this terminal condition usually requires that $\lim_{t\to\infty} e^{-\delta t}\lambda(t) = 0$, or that A = 0 in the

You can verify that if we identify the costate variable with the derivative of the value function, $J'(k)$, the Pontryagin necessary conditions are satisfied by our earlier dynamic-programming solution. In particular, (17) coincides with (12) and (18) coincides with (14). So we have shown, in a simple stationary setting, why the Maximum Principle "works." The principle is actually more broadly applicable than you might guess from the foregoing discussion--it easily handles nonstationary environments, side constraints, etc. *And* it has a stochastic analogue, to which I now turn. [10]

---

solution above. The *particular* solution that remains equates the shadow price of a unit of capital to the discounted stream of its shadow "marginal products," where the latter are measured by partial derivatives of the flow of value, $\mathcal{H}$, with respect to k.

[10] For more details and complications on the deterministic Maximum Principle, see Arrow and Kurz, *op. cit.*

III. Stochastic Optimization in Continuous Time

The optimization principles set forth above extend directly to the stochastic case. The main difference is that to do continuous-time analysis, we will have to think about the right way to model and analyze uncertainty that evolves *continuously* with time. To understand the elements of continuous-time stochastic processes requires a bit of investment, but there is a large payoff in terms of the analytic simplicity that results.

Let's get our bearings by looking first at a discrete-time stochastic model.[11] Imagine now that the decision maker maximizes the von Neumann-Morgenstern expected-utility indicator

$$(19)\quad \mathbf{E}_0 \sum_{t=0}^{\infty} e^{-\delta th} U[c(t),k(t)]h,$$

where $\mathbf{E}_t X$ is the expected value of random variable X conditional on all information available up to (and including) time t.[12] Maximization is to be carried out subject to the constraint that

$$(20)\quad k(t+h) - k(t) = G[c(t),k(t),\theta(t+h),h], \qquad k(0) \text{ given},$$

[11]An encyclopedic reference on discrete-time dynamic programming and its applications in economics is Nancy L. Stokey and Robert E. Lucas, Jr. (with Edward C. Prescott), *Recursive Methods in Economic Dynamics* (Cambridge, Mass.: Harvard University Press, 1989). The volume pays special attention to the foundations of stochastic models.

[12]Preferences less restrictive than those delimited by the von Neumann-Morgenstern axioms have been proposed, and can be handled by methods analogous to those sketched below.

where $\{\theta(t)\}_{t=-\infty}^{\infty}$ is a sequence of exogenous random variables with a known joint distribution, and such that only realizations up to and including $\theta(t)$ are known at time t. For simplicity I will assume that the θ process is *first-order Markov*, that is, that the joint distribution of $\{\theta(t+h), \theta(t+2h), ...\}$ conditional on $\{\theta(t), \theta(t-h), ...\}$ depends *only* on $\theta(t)$. For example, the AR(1) process $\theta(t) = \rho\theta(t-h) + \upsilon(t)$, where $\upsilon(t)$ is distributed independently of past θ's, has this first-order Markov property.

Constraint (20) differs from its deterministic version, (6), in that the time interval h appears as an argument of the transition function, but not necessarily as a multiplicative factor. Thus, (20) is somewhat more general than (6). The need for this generality arises because $\theta(t+h)$ is meant to be "proportional" to h in a sense that will become clearer as we proceed.

Criterion (19) reflects inherent uncertainty in the realizations of c(t) and k(t) for t > 0. Unlike in the deterministic case, the object of individual choice is *not* a single path for the control variable c. Rather, it is a sequence of *contingency plans* for c. Now it becomes really essential to think in terms of a policy function mapping the "state" of the program to the optimal level of the control variable. The optimal policy function giving c*(t) will not be a function of the state variable k(t) alone, as it was in the last section; rather, it will depend on k(t) *and* $\theta(t)$, because $\theta(t)$ (thanks to the first-order Markov assumption) is the piece of current information that helps forecast the future realizations $\theta(t+h)$, $\theta(t+2h)$, etc. Since k(t) and $\theta(t)$ evolve stochastically, writing $c^*(t) = c[k(t);\theta(t)]$ makes it clear that from the perspective of any time before t, c*(t) will be a random variable, albeit one that depends in a very particular way on the realized values of k(t) and $\theta(t)$.

Bellman’s principle continues to apply, however. To implement it, let us write the value function--again defined as the maximized value of (19)--as $J[k(0);\theta(0)]$. Notice that $\theta(0)$ enters the value function for the same reason that $\theta(t)$ influences c*(t). If θ is a positive shock to capital productivity (for example), with θ positively serially correlated, then a higher current value of θ leads us to forecast higher θ’s for the future. This higher expected path for θ both raises raises expected lifetime utility and influences the optimal consumption choice.

In the present setting we write the Bellman equation as

$$(21)\quad J[k(t);\theta(t)] = \max_{c(t)} \left\{ U[c(t),k(t)]h + e^{-\delta h}\mathbf{E}_t J[k(t+h);\theta(t+h)] \right\},$$

where the maximization is done subject to (20). The rationale for this equation basically is the same as before. The contingent rules for $\{c(s)\}_{s=t+1}^{\infty}$ that maximize $\mathbf{E}_t \sum_{s=t}^{\infty} e^{-\delta sh} U[c(s),k(s)]h$ subject to (20), given k(t) and the optimal choice c*(t), will also maximize

$$\mathbf{E}_t \sum_{s=t+1}^{\infty} e^{-\delta sh} U[c(s),k(s)]h$$ subject to (20), given the probability distribution for k(t+h) induced by c*(t).

Equation (21) is the stochastic analogue of (7) for the case of first-order Markovian uncertainty. The equation is immediately useful for discrete-time analysis: just use (20) to eliminate k(t+h) from (21) and differentiate away. But our concern here is with continuous-time analysis. We would like to proceed as before, letting the market interval h go to zero in (21) and, hopefully, deriving some nice expression analogous to (9). Alas, life is not so easy. If you try to take the route just described, you will end up with an expression that looks like the expected value of

$$\frac{J[k(t+h);\theta(t+h)] - J[k(t);\theta(t)]}{h}.$$

This quotient need not, however, converge (as $h \to 0$) to a well-defined random variable. One way to appreciate the contrast between the present setup and the usual setup of the calculus is as follows. Because $J[k(t);\theta(t)]$ is a random variable, a plot of its realizations against time--a *sample path*--is unlikely to be differentiable. Even after time is carved up into very small intervals, the position of the sample path will change abruptly from period to period as new realizations occur. Thus, expressions like the quotient displayed above may have no well-defined limiting behavior as $h \to 0$. To proceed further we need a new mathematical theory that allows us to analyze infinitesimal changes in random variables. The stochastic calculus is designed to accomplish precisely this goal.

Stochastic Calculus

Let $X(t)$ be a random variable whose change between periods $t - 1$ and t, $\Delta X(t) = X(t) - X(t - 1)$, has mean μ and variance σ^2. To simplify matters I'll assume that $\Delta X(t)$ is normally distributed, although this is not at all necessary for the argument. [13]

We are interested in the case where $\Delta X(t)$, the change in random variable X over the period of length 1 between $t - 1$ and t, can be viewed as a sum (or integral) of very small (in the limit, infinitesimal) random changes. We would also like each of these changes, no matter how small, to have a normal distribution. Our method, as in the usual calculus, is to divide the time interval $[t - 1, t]$ into small segments. But we need to be sure that no matter how finely we do the subdivision, $\Delta X(t)$, the sum of the smaller changes, remains $N(\mu,\sigma^2)$.

To begin, carve up the interval $[t - 1, t]$ into n disjoint subintervals, each of length $h = 1/n$. For every $i \in \{1,2,...,n\}$,

---

[13] For a simplified yet rigorous exposition of these matters,

let $\nu(i)$ be a N(0,1) random variable with $\mathbf{E}\nu(i)\nu(j) = 0$ for $i \neq j$. Suppose that $\Delta X(t)$ can be written as

$$(22)\quad \Delta X(t) = \sum_{i=1}^{n} \mu h + \sigma h^{1/2}\nu(i)$$

Then since nh = 1, (22) is consistent with our initial hypothesis that $\mathbf{E}\Delta X(t) = \mu$ and $\mathbf{V}\Delta X(t) = \sigma^2$. For example,

$$\mathbf{V}\Delta X(t) = \sigma^2 \sum_{i=1}^{n} \sum_{j=1}^{n} \mathbf{E}\nu(i)\nu(j)/n = \sum_{i=1}^{n} \mathbf{E}\nu(i)^2/n = \sigma^2.$$

Equation (22) expresses the finite change $\Delta X(t)$ as the sum of tiny independent normal increments of the form $\mu h + \sigma h^{1/2}\nu$. It is customary to denote the limit of such an increment as $h \to 0$ by $\mu dt + \sigma dz$, where for any instant τ, $dz(\tau) = \lim_{h\to 0} h^{1/2}\nu(\tau)$. When this limit is well-defined, we say that X(t) follows the Gaussian *diffusion process*

$$(23)\quad dX(t) = \mu dt + \sigma dz(t),$$

which means, in notation that is suggestive but that I will not attempt to define rigorously, that

$$X(t) = X(\tau) + \mu(t-\tau) + \sigma\int_{\tau}^{t} dz(s) = X(\tau) + \mu(t-\tau) + \sigma[z(t)-z(\tau)]$$

for all $\tau \le t$. [14]

---

[14]Again, see Merton, *op. cit.*, for a more rigorous treatment. To make all this more plausible, you may want to write (22) (for our

Think of X(t) as following a continuous-time random walk with a predictable rate of drift μ and an instantaneous rate of variance (variance per unit of time) σ^2. When $\sigma = 0$, we are back in the deterministic case and are therefore allowed to assert that X(t) has time derivative μ: $dX(t)/dt = \mu$. But when $\sigma > 0$, X(t) has sample paths that are differentiable nowhere. So we use a notation, (23), that does not require us to "divide" random differences by dt. Because we are looking at arbitrarily small increments over arbitrarily small time intervals, however, the sample paths of X(t) are continuous.

Now that we have a sense of what (23) means, I point out that this process can be generalized while maintaining a Markovian setup in which today's X summarizes all information useful for forecasting future X's. For example, the process

$$(24)\quad dX = \mu(X,t)dt + \sigma(X,t)dz.$$

---

earlier case with $\tau = t - 1$) as

$$\Delta X(t) - \mu = \sum_{i=1}^{n} \nu(i)/\sqrt{n} \ ,$$

where $n = 1/h$ is the number of increments in $[t - 1, t]$. We know from the central-limit theorem that as $n \to \infty$, the right-hand side above is likely to approach a limiting normal distribution even if the $\nu(i)$'s aren't normal (so my assumptions above were stronger than necessary). Obviously, also, $X(t) - X(t - h)$ will be normally distributed with variance $h\sigma^2$ no matter how small h is. But $X(t) - X(t - h)$ *divided by* h therefore explodes as $h \to 0$ (its variance is σ^2/h). This is why the sample paths of diffusion processes are not differentiable in the usual sense.

allows the drift and variability of dX to be functions of the level of X(t) itself, which is known at time t, and of time.

There is a further set of results we'll need before tackling the one major theorem of stochastic analysis applied below, Itô's chain rule. We need to know the rules for multiplying stochastic differentials. We're familiar, from the usual differential calculus, with the idea that quantities of order dt are important, whereas quantities of order dt^m, m > 1, are not. For example, in calculating the derivative of the function y^2, we compute h^{-1} times the limit of $(y + h)^2 - y^2 = 2yh + h^2$ as $h \to 0$. The derivative is simply 2y, because the term h^2 goes to zero even after division by h. The same principle will apply in stochastic calculus. Terms of order greater than h are discarded. In particular $dt^2 = \lim_{h\to\infty} h^2$ will be set to zero, just as always.

What about something like the product dzdt? Since this is the limit of $h^{3/2}\nu$ as $h \to \infty$, it shrinks faster than h and accordingly will be reckoned at zero:

(25) $dzdt = 0.$

Finally, consider $dz^2 = \lim_{h\to\infty} h\nu^2$. This is of order h, and thus does not disappear as h gets very small. But the *variance* of this term can be shown to be $2h^2$, which *is* zero asymptotically.[15]

---

[15] To prove this, note that because ν is N(0,1), $\mathbf{V}h\nu^2 = \mathbf{E}(h\nu^2 -$

By Chebyshev's inequality, $h\nu^2$ thus converges in probability to its expected value, h, as $h \to 0$, and so we write

(26) $dz^2 = dt.$

Let's turn now to Itô's famous lemma. Suppose that the random variable X(t) follows a diffusion process such as (24). The basic idea of Itô's Lemma is to help us compute the stochastic differential of the random variable f[X(t)], where $f(\cdot)$ is a differentiable function. If $\sigma(X,t) \equiv 0$, then the chain rule of ordinary calculus gives us the answer: the change in f(X) over an infinitesimal time interval is given by $df(X) = f'(X)dX = f'(X)\mu(X,t)dt$. If $\sigma(X,t) \not\equiv 0$ but $f(\cdot)$ is linear, say $f(X) = aX$ for some constant a, then the answer is also quite obvious: in this special case, $df(X) = f'(X)dX = a\mu(X,t)dt + a\sigma(X,t)dz$.

Even if $f(\cdot)$ is nonlinear, however, there is often a simple answer to the question we've posed:

Itô's Lemma. Let X(t) follow a diffusion process, and let $f:\mathbb{R} \to \mathbb{R}$ be twice continuously differentiable. The stochastic differential of f(X) is

(27) $df(X) = f'(X)dX + \frac{1}{2}f''(X)dX^2.$

$h)^2 = \mathbf{E}(h^2\nu^4 - 2h^2\nu^2 + h^2) = 3h^2 - 2h^2 + h^2 = 2h^2.$

Comment. If X follows the diffusion process (24), then, using rules (25) and (26) to compute dX^2 in (27), we get

$$(28)\quad df(X) = [\mu(x,t)f'(X) + \frac{\sigma(X,t)^2}{2}f''(X)]dt + \sigma(X,t)f'(X)dz.$$

You'll notice that (28) differs from the "naive" chain rule only in modifying the expected drift in f(X) by a term that depends on the curvature of $f(\cdot)$. If $f''(X) > 0$ so that $f(\cdot)$ is strictly convex, for example, (28) asserts that $\mathbf{E}_t df(X) = \mathbf{E}_t f[X(t+dt)] - f[X(t)]$ is greater than $f'(X)\mu(X,t)dt = f'(X)\mathbf{E}_t dX = f[\mathbf{E}_t X(t+dt)] - f[X(t)]$. But anyone who remembers Jensen's Inequality knows that $\mathbf{E}_t f[X(t+dt)] \geq f[\mathbf{E}_t X(t+dt)]$ for convex $f(\cdot)$, and that the opposite inequality holds for concave $f(\cdot)$. So Itô's Lemma should not come as a surprise.[16]

---

[16]In case you don't remember Jensen's Inequality, here's a quick

Motivation for Itô's Lemma. The proof of Itô's Lemma is quite subtle, so a heuristic motivation of this key result will have to suffice.[17] Once again I'll rely on a limit argument. For an interval length h, Taylor's theorem[18] implies that

$$f[X(t+h)] - f[X(t)] = f'[X(t)][X(t+h) - X(t)]$$
$$+ \frac{1}{2}f''\{X(t) + \xi(h)[X(t+h) - X(t)]\}[X(t+h) - X(t)]^2,$$

where $\xi(h) \in [0,1]$. It may look "obvious" to you that this converges to (27) as $h \to 0$. Beware. It turns out to be quite a chore to ensure that the right-hand side of this expression is well behaved as $h \to 0$, largely because of the complicated dependence of the term $f''\{X(t) + \xi(h)[X(t+h) - X(t)]\}$ on h. Fortunately, as $h \to 0$, the randomness in this term does disappear quickly enough that we can safely equate it to $f''[X(t)]$ in the limit. The result is (27). It should now be clear how one would

---

sketch of a proof. Recall that a convex function has the property that $\gamma f(X_1) + (1-\gamma)f(X_2) \geq f[\gamma X_1 + (1-\gamma)X_2]$ $\forall\gamma \in [0,1]$. It is easy to extend this to the proposition that $\sum_i \pi_i f(X_i) \geq f(\sum_i \pi_i X_i)$ for $(\pi_1,...,\pi_n)$ in the unit simplex. (Try it.) So for finite discrete probability distributions we're done. (Obviously concave functions work the same way, with the inequalities reversed.) Now consider the case in which the random variable X has an arbitrary continuous density function $\pi(X)$. We can approximate $\mathbf{E}f(X)$ by sums of the form $\sum_i f(X_i)\pi(X_i)h$, each of which must be at least as great as $f[\sum_i X_i\pi(X_i)h]$ if we choose the

[17]For Taylor's theorem with remainder, see any good calculus text.

motivate a multivariate version of Itô's Lemma using the multivariate Taylor expansion.

The preceding digression on stochastic calculus has equipped us to answer the question raised at the outset: What is the continuous-time analogue of (21), the stochastic Bellman equation?

To make matters as simple as possible, in analogy with section II's time-stationary setup, I'll assume that $\theta(t+h) = X(t+h) - X(t)$, where $X(t)$ follows the simple diffusion process (23), $dX = r dt + \sigma dz$, for constant r and σ. Under this assumption $\mathbf{E}_t\theta(t+h) = rh$ always, so knowledge of $\theta(t)$ gives us no information about future values of θ. Thus the value function depends on the state variable k alone. Now (21) becomes

$$(29) \quad J[k(t)] = \max_{c(t)} \left\{ U[c(t),k(t)]h + e^{-\delta h}\mathbf{E}_t J[k(t+h)] \right\}.$$

Let's carry on by adapting the last section's strategy of subtracting $J[k(t)]$ from both sides of (21) and replacing $e^{-\delta h}$ by $1 - \delta h$. (We now know we can safely ignore the terms in h^m for $m \geq 2$.) The result is

$$0 = \max_{c(t)} \left\{ U[c(t),k(t)]h + \mathbf{E}_t J[k(t+h)] - J[k(t)] - \delta \mathbf{E}_t J[k(t+h)]h \right\}.$$

Now let $h \to 0$. According to (20), $dk = G(c,k,dX,dt)$, and I assume that this transition equation defines a diffusion process for k. Itô's Lemma then tells us that

$$(30)\quad dJ(k) = J'(k)dk + \frac{1}{2}J''(k)dk^2,$$

Thus as $h \to 0$, $\mathbf{E}_tJ[k(t+h)] - J[k(t)] \to J'[k(t)]\mathbf{E}_tdk(t) + \frac{1}{2}J''[k(t)]\mathbf{E}_tdk(t)^2$. Furthermore, as $h \to 0$, $\mathbf{E}_tJ[k(t+h)] \to J[k(t)]$. So we end up with the following:

PROPOSITION III.1. (Continuous-Time Stochastic Bellman Equation) *Consider the problem of maximizing* $\mathbf{E}_0\int_0^\infty e^{-\delta t}U(c,k)dt$ *subject to a diffusion process for* k *controlled by* c, *and given* k(0). *At each moment, the optimal control* c* *satisfies the Bellman equation*

$$(31)\quad 0 = U(c^*,k)dt + J'(k)\mathbf{E}_tG(c^*,k,dX,dt) + \frac{1}{2}J''(k)\mathbf{E}_tG(c^*,k,dX,dt)^2 - \delta J(k)dt$$
$$= \max_{c(t)}\left\{U(c,k)dt + J'(k)\mathbf{E}_tdk + \frac{1}{2}J''(k)\mathbf{E}_tdk^2 - \delta J(k)dt\right\}.$$

Equation (31) is to be compared with equation (9), given in Proposition II.1. Indeed, the interpretation of Proposition III.1 is quite similar to that of Proposition II.1. Define the stochastic Hamiltonian [in analogy to (10)] as

$$(32)\quad \mathcal{H}(c,k) \equiv U(c,k) + J'(k)\frac{\mathbf{E}_tdk}{dt} + \frac{1}{2}J''(k)\frac{\mathbf{E}_tdk^2}{dt}.$$

The Hamiltonian has the same interpretation as (10), but with a stochastic twist. The effect of a given level of "savings" on next period's "capital stock" now is uncertain. Thus the Hamiltonian measures the *expected* flow value, in current utility terms, of the consumption-savings combination implied by the consumption choice c, given the predetermined (and known) value of k. The analogy will be clearer if you use (30) to write (32) as [18]

$$\mathcal{H}(c,k) = U(c,k) + \frac{\mathbf{E}_t dJ(k)}{dt},$$

and if you use the ordinary chain rule to write the deterministic Hamiltonian (10) as $U(c,k) + J'(k)\dot{k} = U(c,k) + dJ(k)/dt$.

The stochastic Bellman equation therefore implies the same rule as in the deterministic case, but in an expected-value sense. Once again, optimal consumption c* satisfies (11),

$$\mathcal{H}(c^*,k) = \max_c \{\mathcal{H}(c,k)\} = \delta J(k).$$

Rather than proceeding exactly as in our deterministic analysis, I will sacrifice generality for clarity and adopt a specific (but widely used) functional form for the continuous-

---

[19]The notation in (32) and in the next line below is common. Since $\mathbf{E}_t dk$, for example, is deterministic, $(\mathbf{E}_t dk)/dt$ can be viewed as the expected *rate* of change in k. Since diffusion processes aren't differentiable, $\mathbf{E}_t(dk/dt)$ is in contrast a nonsensical expression.

time version of (20), $dk = G(c,k,dX,dt)$. I will assume the linear transition equation

$$(33)\quad dk = kdX - cdt = (rk - c)dt + \sigma kdz$$

(since $dX = rdt + \sigma dz$). What form does (31) now assume? To see this we have to calculate $\mathbf{E}_t dk$ and $\mathbf{E}_t dk^2$. It is clear from (33) that $\mathbf{E}_t dk = (rk - c)dt$. Invoking (25) and (26), and recalling that $dt^2 = 0$, we see that $dk^2 = \mathbf{E}_t dk^2 = k^2 dX^2 - 2ckdXdt + c^2 dt^2 = \sigma^2 k^2 dt$. We thus conclude that c* must solve

$$(34)\quad 0 = \max_{c(t)} \left\{ U(c,k) + J'(k)(rk - c) + \frac{1}{2}J''(k)k^2\sigma^2 - \delta J(k) \right\}.$$

In principle this equation is no harder to analyze than was (9): the two are identical [if $G(c,k) = rk - c$] aside from the additional second derivative term in (34), due to Itô's Lemma. So we proceed as before, starting off by maximizing the Hamiltonian.

Since k is predetermined and known at each moment, the necessary condition for c* to maximize the right hand of (34) is

$$(35)\quad U_C(c^*,k) = J'(k),$$

which is the same as (12) because I've assumed here that $G_C = -1$.

We can also define the optimal policy function $c^* = c(k)$, just as before. By definition c(k) satisfies the equation

(36) $$0 = U[c(k),k] + J'(k)[rk - c(k)] + \frac{1}{2}J''(k)k^2\sigma^2 - \delta J(k).$$

One would hope to understand better the implied dynamics of c by differentiating with respect to the state variable. The result is

(37) $$U_k(c^*,k) + J'(k)(r - \delta) + J''(k)k\sigma^2 + J''(k)(rk - c^*) + \frac{1}{2}J'''(k)k^2\sigma^2 = 0,$$

where I've already applied the envelope condition (35).

It is tempting to give up in the face of all these second and third derivatives; but it is nonetheless possible to interpret (37) in familiar economic terms. Let's again define the shadow price of k, λ, by

$$\lambda \equiv J'(k).$$

This shadow price is known at time t, but its change over the interval from t to t + dt is stochastic. Equation (37) differs from (13) only by taking this randomness into account; and by writing (37) in terms of λ, we can see precisely how this is done.

To do so we need two observations. First, Itô's Lemma discloses the stochastic differential of λ to be

$$(38)\quad d\lambda = dJ'(k) = J''(k)(kdX - cdt) + \frac{1}{2}J'''(k)k^2\sigma^2 dt$$

(verify this), so that

$$(39)\quad \frac{\mathbf{E}_t d\lambda}{dt} = J''(k)(rk - c) + \frac{1}{2}J'''(k)k^2\sigma^2.$$

Second, the term $J''(k)k\sigma^2$ in (37) can be expressed as

$$(40)\quad J''(k)k\sigma^2 = -J'(k)R(k)\sigma^2,$$

where $R(k) \equiv -J''(k)k/J'(k)$ should be interpreted as a coefficient of relative risk aversion.

Using (39) and (40), rewrite (37) in terms of $\lambda = J'(k)$ as

$$U_k(c^*,k) + \lambda[r - R(k)\sigma^2 - \delta] + \frac{\mathbf{E}_t d\lambda}{dt},$$

or, in analogy to (14), as

$$(41)\quad \frac{U_k + \lambda[r - R(k)\sigma^2/2] + [(\mathbf{E}_t d\lambda)/dt - \lambda R(k)\sigma^2/2]}{\lambda} = \delta,$$

To compare (41) with (14), notice that under the linear transition equation (33), r corresponds to the expected value of G_k; we adjust this expectation downward for risk by subtracting the product of the risk-aversion coefficient and $\sigma^2/2$. An identical risk adjustment is made to the expected "capital gains" term, $(\mathbf{E}_t d\lambda)/dt$. Otherwise, the equation is the same as (14), and has a corresponding "efficient asset price" interpretation.

Example

An individual maximizes the expected discounted utility of consumption, $\mathbf{E}_0\int_0^\infty e^{-\delta t}U(c)dt$, subject to a stochastic capital accumulation constraint that looks like (33):

$$dk = rkdt + \sigma kdz - cdt, \quad k(0) \text{ given.}$$

What is the meaning of this savings constraint? Capital has a mean marginal product of r, but its realized marginal product fluctuates around r according to a white-noise process with instantaneous variance σ^2. The flow utility function is

$$U(c) = \frac{c^{1-(1/\varepsilon)} - 1}{1 - (1/\varepsilon)},$$

as in the second part of the last section's example.

To solve the problem I'll make the same guess as before, that the optimal consumption policy function is $c(k) = \eta k$ for an appropriate η. As will be shown below--and as was the case in a deterministic setting--the value function J(k) is a linear function of $k^{1-(1/\varepsilon)}$, making the risk aversion coefficient R(k) defined after (40) a constant, $R \equiv 1/\varepsilon$. For now I will assume this, leaving the justification until the end.

How can we compute η in the policy function $c(k) = \eta k$? The argument parallels our earlier discussion of the nonstochastic case, which you may wish to review at this point.

Start by thinking about the implications of the postulated policy function for the dynamics of capital. If $c(k) = \eta k$, then

$$dk = rkdt + \sigma kdz - c(k)dt = (r - \eta)kdt + \sigma kdz.$$

But as optimal c is proportional to k,

$$dc = (r - \eta)cdt + \sigma cdz.$$

Above we defined λ as $J'(k)$; but first-order condition (35) implies that $\lambda = U'(c) = c^{-1/\varepsilon}$. Application of Itô's Lemma to $\lambda = c^{-1/\varepsilon}$ leads to

$$d\lambda = -\left(\frac{1}{\varepsilon}\right)c^{-1-(1/\varepsilon)}dc + \left(\frac{1}{2}\right)\left(\frac{1}{\varepsilon}\right)\left(1 + \frac{1}{\varepsilon}\right)c^{-2-(1/\varepsilon)}dc^2.$$

Because we've already established that $\mathbf{E}_t dc = (r - \eta)cdt$ and that $dc^2 = \sigma^2 c^2 dt$, we infer from the equation above that

$$\frac{E_t d\lambda}{dt} = \frac{c^{-(1/\varepsilon)}}{\varepsilon}\left[\eta - r + \left(\frac{1}{2}\right)\left(1 + \frac{1}{\varepsilon}\right)\sigma^2\right]$$

But there is an alternative way of describing the dynamics of λ: equation (41) can be written here as

$$\frac{\mathbf{E}_t d\lambda}{dt} = \lambda[\delta - (r - R\sigma^2)] = c^{-1/\varepsilon}[\delta - (r - \sigma^2/\varepsilon)].$$

So we have derived two potentially different equations for $(\mathbf{E}_t d\lambda)/dt$; clearly the two are mutually consistent if and only if

$$[\delta - (r - \sigma^2/\varepsilon)] = \left(\frac{1}{\varepsilon}\right)\left[\eta - r + \left(\frac{1}{2}\right)\left(1 + \frac{1}{\varepsilon}\right)\sigma^2\right],$$

or, solving for η, if and only if

$$\eta = r - \varepsilon(r - \delta) + \frac{(\varepsilon - 1)}{2\varepsilon}\sigma^2.$$

The implied consumption rule is similar to the one that arose in the nonstochastic example analyzed earlier, but it corrects for the unpredictable component of the return to capital. (Notice that we again obtain $\eta = \delta$ if $\varepsilon = 1$.) The analogy with (16) will be clearest if the rule is written as

$$(42) \quad \eta = (1 - \varepsilon)(r - \frac{1}{2}R\sigma^2) + \varepsilon\delta.$$

In (42), η appears as the weighted average of the time-preference rate and a *risk-adjusted* expected return on investment.

Problems still arise if $\eta \leq 0$. In these cases an optimum fails to exist, for reasons essentially the same as those discussed in section II's example.

As a final exercise let's calculate the value function J(k) and confirm the assumption about its form on which I've based my analysis of the optimal consumption policy function. In the process we'll learn some more about the importance of Itô's Lemma. One way to approach this task is to calculate the (random) path for k under an optimal consumption plan, observe that the optimal contingency rule for consumption is $c = \eta k$, and then use this formula to compute the optimal (random) consumption path and lifetime expected utility. Indeed, we took a very similar tack in the deterministic case. So we start by asking what the optimal transition equation for the capital stock, $dk = (r - \eta)kdt + \sigma dz$, implies for the *level* of k. [Throughout the following discussion, you should understand that η is as specified by (42).]

Observe first that the optimal capital-stock transition equation can be written as

$$dk/k = (r - \eta)dt + \sigma dz.$$

A crucial warning. You might think that dk/k is the same thing as dlog(k), as in the ordinary calculus. *If* this were true, we would conclude that the capital stock follows the stochastic process

$$\log[k(t)] = \log[k(0)] + (r - \eta)t + \sigma\int_0^t dz(s),$$

or, equivalently, that

$$k(t) = k(0)e^{(r-\eta)t + \sigma[z(t)-z(0)]}.$$

But this is incorrect. Itô's Lemma tells us that $d\log(k) = (dk/k) - \frac{1}{2}\sigma^2 dt = (r - \eta - \frac{1}{2}\sigma^2)dt + \sigma dz$. [The reason for this divergence is Jensen's Inequality--log(•) is a strictly concave function.] It follows that the formula for k(t) below is the right one:

(43) $$k(t) = k(0)e^{(r-\eta-\sigma^2/2)t + \sigma[z(t)-z(0)]}.$$

At an optimum, k(t) will be conditionally lognormally distributed, with an expected growth rate of $r - \eta$: $\mathbf{E}_0 k(t)/k(0) = e^{(r-\eta)t}$.[20]

As a result of (43), the value function at t = 0 is

---

[20] If X is a normal random variable with mean μ and variance σ^2, e^X is said to be lognormally distributed. The key fact about lognormals that is used repeatedly is that when X is normal,

$$\mathbf{E}e^X = e^{\mu+\sigma^2/2}.$$

For a proof, see any good statistics text.

$$J[k(0)] = \left[1 - \frac{1}{\varepsilon}\right]^{-1}\mathbf{E}_0\left\{\int_0^\infty e^{-\delta t}[\eta k(t)]^{1-(1/\varepsilon)}dt - \frac{1}{\delta}\right\}$$

$$= \left[1 - \frac{1}{\varepsilon}\right]^{-1}\left\{\int_0^\infty e^{-\delta t}\mathbf{E}_0\left[\eta k(0)e^{(r-\eta-\sigma^2/2)t+\sigma[z(t)-z(0)]}\right]^{1-(1/\varepsilon)}dt - \frac{1}{\delta}\right\}$$

$$= \left[1 - \frac{1}{\varepsilon}\right]^{-1}\left\{\left[\eta k(0)\right]^{1-(1/\varepsilon)}\int_0^\infty e^{-\delta t}e^{[1-(1/\varepsilon)](r-\eta-\sigma^2/2\varepsilon)t}dt - \frac{1}{\delta}\right\}$$

$$= \left[1 - \frac{1}{\varepsilon}\right]^{-1}\left\{\frac{[\eta k(0)]^{1-(1/\varepsilon)}}{\delta - (\varepsilon - 1)(r - R\sigma^2/2 - \delta)} - \frac{1}{\delta}\right\}.$$

You'll recognize the final product above as the same formula for $J[k(0)]$ that we encountered on p. 16 above, with the sole amendment that the risk-adjusted expected return $r - R\sigma^2/2$ replaces r everywhere [including in η; recall (42)].[21] Because $\delta - (\varepsilon - 1)(r - R\sigma^2/2 - \delta) = \eta$, $\eta > 0$ ensures convergence of the integral defining $J(k)$. Finally, $J(k)$ is a linear function of $k^{1-(1/\varepsilon)}$, as claimed earlier.

There is another, more direct way to find the value

---

[21]To move from the second to the third equality above, I used the fact that the normal random variable $[1 - (1/\varepsilon)]\sigma[z(t) - z(0]$ has mean zero and variance $[1 - (1/\varepsilon)]^2\sigma^2 t$ conditional on $t = 0$ information.

function, one that also applies in the deterministic case. [Had we known the value function in advance, we could have used (35) to compute the consumption function without trial-and-error guesses.] By (35), the optimal control must satisfy

$$c(k) = J'(k)^{-\varepsilon}.$$

Thus by (34),

$$0 = \frac{[J'(k)]^{1-\varepsilon}}{1 - (1/\varepsilon)} + J'(k)[rk - J'(k)^{-\varepsilon}] + \frac{1}{2}J''(k)k^2\sigma^2 - \delta J(k).$$

This is just an ordinary second-order differential equation which in principle can be solved for the variable J(k). You may wish to verify that the value function J(k) we derived above is indeed a solution. To do the nonstochastic case, simply set $\sigma^2 = 0$.

The similarities between this example and its deterministic analogue are striking. They are not always so direct. Nonetheless, it is noteworthy that for the linear state transition equation considered above, there exists a stochastic version of Pontryagin's Maximum Principle. One could attack the problem in full generality,[22] but as my goal here is the more modest one of illustrating the basic idea, I will spare you this.

---

[22] As does Jean-Michel Bismut, "Growth and the Optimal Intertemporal Allocation of Risks," *Journal of Economic Theory* 10 (April 1975): 239-257.

PROPOSITION III.2. (Stochastic Maximum Principle) *Let* $c^*(t)$ *solve the problem of maximizing*

$$\mathbf{E}_0\int_0^\infty e^{-\delta(s-t)}U[c(s),k(s)]ds$$

subject to the transition equation

$$dk(t) = rk(t)dt + \sigma k(t)dz(t) - c(t)dt, \qquad k(0) \text{ given},$$

where $z(t)$ *is a standard Gaussian diffusion. Then there exist costate variables* $\lambda(t)$ *such that if* $\zeta(t)$ *is the instantaneous conditional covariance of* $\lambda(t)$ *and* $z(t)$*, the* risk-adjusted *Hamiltonian*

$$\tilde{\mathcal{H}}[c,k(t),\lambda(t),\zeta(t)] \equiv U[c,k(t)] + \lambda(t)[rk(t) - c] + \zeta(t)\sigma k(t)$$

is maximized at $c = c^*(t)$ *given* $\lambda(t)$, $\zeta(t)$, *and* $k(t)$*; that is,*

$$(44) \quad \frac{\partial\tilde{\mathcal{H}}}{\partial c}(c^*,k,\lambda,\zeta) = U_c(c^*,k) - \lambda = 0$$

at all times (assuming an interior solution). Furthermore, the costate variable obeys the stochastic differential equation

$$(45) \quad \begin{aligned} d\lambda &= \lambda\delta dt - \frac{\partial\tilde{\mathcal{H}}}{\partial k}(c^*,k,\lambda,\zeta)dt + \zeta dz \\ &= \lambda\delta dt - [U_k(c^*,k) + \lambda r + \zeta\sigma]dt + \zeta dz \end{aligned}$$

for $dk = rkdt - c^*dt + \sigma kdz$ *and* $k(0)$ *given*

To understand how this proposition follows from our earlier discussion, observe first that because λ will again equal $J'(k)$, the instantaneous conditional covariance of $\lambda(t)$ and $z(t)$ can be seen from (25), (26), and (38) to be

$$(46) \quad \zeta = (\mathbf{E}_t d\lambda dz)/dt = J''(k)\sigma k.$$

Thus, with reference to the definition (32) of the *unadjusted* stochastic Hamiltonian, given here by

$$\mathcal{H}(c,k) = U(c,k) + J'(k)(rk - c) + \frac{1}{2}J''(k)\sigma^2k^2,$$

we have

$$\tilde{\mathcal{H}}(c,k,\lambda,\zeta) = \mathcal{H}(c,k) + \frac{1}{2}J''(k)\sigma^2k^2 = \mathcal{H}(c,k) - \lambda R(k)\sigma^2k/2,$$

where $R(k)$ is the relative risk-aversion coefficient defined above. Accordingly, we can interpret $\tilde{\mathcal{H}}$ as the expected instantaneous flow of value *minus* a premium that measures the riskiness of the stock of capital currently held.

With (46) in hand it is easy to check the prescriptions of the Stochastic Maximum Principle against the results we've already derived through other arguments. Clearly (44) corresponds directly to (35). Likewise, if you multiply (37) by dt and combine the result with (38), you will retrieve (45).

IV. Conclusion

These notes have offered intuitive motivation for the basic optimization principles economists use to solve deterministic and stochastic continuous-time models. My emphasis throughout has been on the Bellman principle of dynamic programming, which offers a unified approach to all types of problems. The Maximum Principle of optimal control theory follows from Bellman's approach in a straightforward manner.

I have only been able to scratch the surface of the topic. Methods like those described above generalize to much more complex environments, and have applications much richer than those I worked through for you. The only way to gain a true understanding of these tools is through "hands on" learning: you must apply them yourself in a variety of situations. As I noted at the outset, abundant applications exist in many areas of economics. I hope these notes make this fascinating body of research more approachable.