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 $h \rightarrow 0$) 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 $\tau \leq t$. [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) $dX = \mu(X,t)dt + \sigma(X,t)dz.$

---

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(·). If $f''(X) > 0$ so that f(·) 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(·), and that the opposite inequality holds for concave f(·). 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), $dX = rdt + \sigma dz$, 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 $h \to 0$, $\mathbf{E}_t J[k(t+h)] - J[k(t)] \to J'[k(t)]\mathbf{E}_t dk(t) + \frac{1}{2}J''[k(t)]\mathbf{E}_t dk(t)^2$. Furthermore, as $h \to 0$, $\mathbf{E}_t J[k(t+h)] \to J[k(t)]$. 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*

$$(31)\quad 0 = U(c^*,k)dt + J'(k)\mathbf{E}_t G(c^*,k,dX,dt) + \frac{1}{2}J''(k)\mathbf{E}_t G(c^*,k,dX,dt)^2 - \delta J(k)dt$$

$$= \max_{c(t)}\left\{U(c,k)dt + J'(k)\mathbf{E}_t dk + \frac{1}{2}J''(k)\mathbf{E}_t dk^2 - \delta J(k)dt\right\}.$$

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

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

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

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

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

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

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

(verify this), so that

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

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

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

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

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

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

Problems still arise if $\eta \le 0$. 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) $$\frac{\partial\tilde{\mathcal{H}}}{\partial c}(c^*,k,\lambda,\zeta) = U_c(c^*,k) - \lambda = 0$$

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

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

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

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.