14. Stochastic Optimal Control
AGEC 642 –2024

To Ngoc Nguyen is the co-author of these notes.
Andrew McMartin provided a helpful derivation of equation 1

These notes are based largely on Section 22, "Stochastic Optimal Control" in Kamien and Schwartz (1991). The other principal sources are Xepapadeas (1997), chapters 3 and section 7.7 and Malliaris and Brock (1982) Chapter 1. Equation numbers are taken from the original sources

I. Stochastic processes in continuous time

In discrete time it is easy to imagine that from t to t+1 there is a random shock so that $x_{t+1}=x_t+f(\cdot)+\varepsilon_t$. But what happens to ε as the time that elapses from t to t+1 goes to zero? In continuous time problems, the very notion of a stochastic process is a little hard to understand.

Randomness that continues no matter how short the time increment was first studied in the physical sciences when Robert Brown noticed the random movements of pollen particles floating in water.[1] Its mathematical formalization is attributed to the mathematician Norbert Wiener. Hence, such random processes are typically referred to as either Brownian Motion or a Wiener Process.

Brownian motion can most easily be thought of as a random walk in continuous time. Let's start with the discrete time example in which $x_{t+\Delta}=x_t+\varepsilon_t$. Assume that the distribution of the random parameter ε is symmetric, constant over time, and has a mean of zero. Figure 1 presents a series of simulated paths in which the size of Δ decreases from 5 periods to one-hundredth of a period. As you can see, the paths do not become smooth, but are instead fractal in nature – no matter how close you look they still jump around in a seemingly jittery fashion with sharp peaks and troughs.

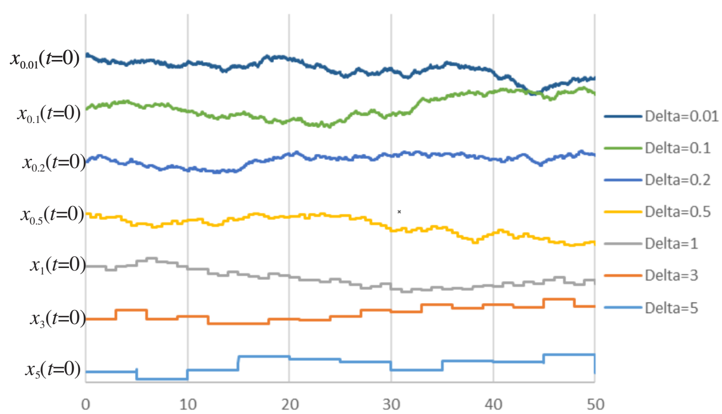

Figure 1: Random paths of $x_{t+\Delta}=x_t+\varepsilon_t$ for levels of Δ from 5 to 0.01 and the variance of ε declining to ensure that the variance over a single period is constant

[1] According to Wikipedia, Jan Ingenhousz observed a similar process over 40 years prior to Brown. However, Brown is typically attributed with this discovery (e.g., Bryson, Bill, 2003. *A Short History of Nearly Everything*, Broadway Books.)

Figure 2 presents a collection of simulated random walks. Since this is a random walk process, $E_t(x_{t+j})=x_t$ no matter how far out in the future you forecast. However, the variance of x_{t+j} increases in j, $var(x_{t+10})>var(x_{t+5})$.

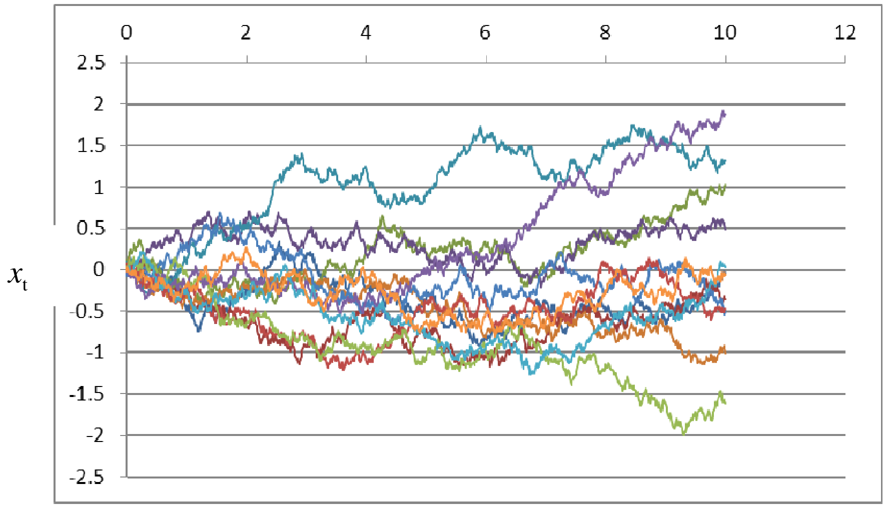

Figure 2: Random walks starting at zero with ε~uniform(–0.5,+0.5)

The random shock, which we will refer to as Z, is the Brownian motion drift. It can be expressed as the limit as the time between one shock and the next goes to zero as presented in Figure 1. If ε_k is the k^{th} independently and identically distributed shock with mean zero and variance 1, then

$Z_t=\lim_{\Delta\to 0}\sqrt{\Delta}\sum_{1\le k\le t/\Delta}\varepsilon_k$ is a Standard Wiener Process.[2]

Wiener processes have a number of interesting characteristics. First, the expectation is zero so that if $x_{t+\Delta t}=x_t+Z_{\Delta t}$, the $E_t(x_{t+k})=x_t$. Second, the variance of the random variable increases linearly in t, i.e., and $\text{var}_t(x_{t+\Delta t})=k\cdot\Delta t$. The increasing variance can be seen in Figure 2. Together, these characteristics mean that the changes in the variable over non-overlapping periods, e.g., $x(t_1)–x(t_0)$, $x(t_2)–x(t_1)$, $x(t_3)–x(t_2)$, …, are independently distributed with means equal to zero and variances proportional to the time lag considered, i.e., $t_1–t_0$, $t_2–t_1$, etc. If $(t_2–t_0)=2\cdot(t_1–t_0)$ then the variance of $[x(t_2)–x(t_0)]$ will be exactly double the variance of $[x(t_1)–x(t_0)]$.

A. Stochastic differential equations

Now let's formalize and generalize following Kamien and Schwartz (KS hereafter). A standard deterministic state equation would take the form

$$\dot{x}=\frac{\partial x}{\partial t}=g\left(t,x,u\right)$$

[2] The use of z is standard notation. Obviously it is important to not confuse it with z_t which we use to define the choice variable.

where u is the control variable and x is the state variable.[3] This can be written as a difference equation, $dx = g(t,x,u)dt$, which holds exactly for $\lim dt \to 0$.

When the state equation is stochastic, the changes in x are a function not only of time but also of the random shocks that occur. From 0 to t, the change in x would have a deterministic part, $\int g(t,x,u)dt$, and a stochastic part, $\sigma(\cdot)Z_t$, where $\sigma(\cdot)$ is a "volatility parameter" that scales the cumulative standard Wiener process, Z_t, which evolves like the lines in Figure 2. Hence, we can write $x_t = x_0 + \int g(t,x,u)dt + \sigma(\cdot)Z_t$. It follows that

$$E(x_t) = x_0 + \int_0^t g(\cdot)dt,$$

$$\text{var}(x_t) = \sigma(\cdot)^2 \text{var}(z) = \sigma(\cdot)^2 t$$

$$x_t \sim N\left(E(x_t), \sigma(\cdot)^2 t\right)$$

Hence, taking the total differential of the standard diffusion yields

$$dx = g(t,x,u)dt + \sigma(t,x,u)dz \tag{1}$$

where z is the standard Wiener process and dz is "the increment of a stochastic process." That is, z is a uniform Brownian motion term, which is scaled by σ for the particular problem that we're working with.

So, the evolution of x is governed not only by choices, u, and time, t, but also by the stochastic process, z. Figure 3 shows an example of Brownian motion characterized by (1), in which an increment in x, i.e., dx, includes 2 components: a deterministic linear drift, $g(.)dt$, and a random diffusion with standard error $\sigma(.)$. If $g(\cdot) = 0$ and $\sigma(.) = 1$ the stochastic process x becomes the standard Wiener process z.

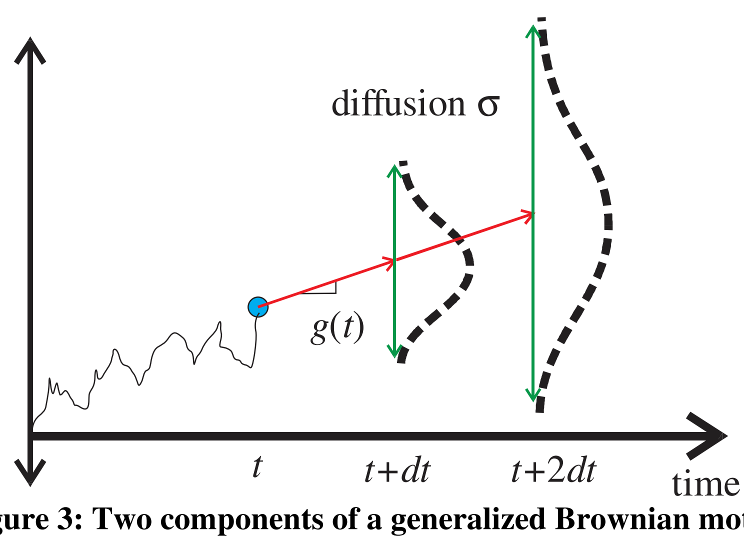

Figure 3: Two components of a generalized Brownian motion

[3] Apologies for the change in notation here. The use of u as a control variable is pretty common in the optimal control literature, so while it is terribly confusing for economists, we have to get used to it.

Recall that the variance of z increases in t, that z is independently distributed over time, and time, which is purely deterministic. Hence, for a standard Wiener process three key equalities are satisfied:

$$(dz)^2 = dt, \quad dzdt = dtdz = 0 \quad \text{and} \quad (dt)^2 = 0. \tag{2}$$

The first of these terms is particularly important, it means the distribution of $(z_t - z_s)$ depends only on the difference in time, $t-s$: the greater the difference in time between t and s, the greater the variance. The second equalities indicate that time and z are independent, and the third equality means that time is deterministic.

The mathematical usefulness of the equalities in (2) can be seen if we consider a second-order Taylor series expansion for a stochastic process. Suppose $y=F(t,z)$ is a random variable that is defined by a deterministic part, i.e., is a function of t, and a random part, z, where z is a Wiener process. A Taylor series expansion of dy would be

$$dy = F_t dt + F_z dz + \frac{1}{2}\left(F_{tt} dt^2 + F_{zz} dz^2\right) + F_{tz} dtdz + \text{higher order terms},$$

where $dt^2 = (dt)^2$ and $dz^2 = (dz)^2$. Using the fact that $dzdt = dt^2 = dtdz = 0$ and ignoring the higher order terms, we obtain

$$dy = F_t dt + F_z dz + \frac{1}{2}\left(\cancel{F_{tt} dt^2} + F_{zz} \underbrace{dz^2}_{(=dt)}\right) + \cancel{F_{tz} dtdz}$$

$$dy = \left(F_t + \frac{1}{2} F_{zz}\right) dt + F_z dz. \tag{3}$$

More generally, if $y=F(t,x)$ where $dx = g(t,x,u)dt + \sigma(t,x,u)dz$, then

$$dy = F_t dt + F_x dx + \frac{1}{2}\left(F_{tt} dt^2 + F_{xx} dx^2\right) + F_{tx} dtdx + \text{higher order terms}$$

$$dy \approx F_t dt + F_x \left(\underline{g(\cdot)dt + \sigma(\cdot)dz}\right) + \frac{1}{2}\left(F_{tt} dt^2 + F_{xx}\left(\underline{g(\cdot)dt + \sigma(\cdot)dz}\right)^2\right)$$
$$+ F_{tx} dt\left(\underline{g(\cdot)dt + \sigma(\cdot)dz}\right)$$

where the underlined pieces are equal to dx. Expanding and cancelling,

$$dy \approx F_t dt + F_x g(\cdot)dt + F_x \sigma(\cdot)dz + \frac{1}{2}\left(\cancel{F_{tt} dt^2} + F_{xx}\left(\cancel{g(\cdot)^2 dt^2} + \cancel{g(\cdot)\sigma(\cdot)dzdt} + \sigma(\cdot)^2 dz^2\right)\right)$$
$$+ \cancel{F_{tx} g(\cdot) dt^2} + \cancel{F_{tx} \sigma(\cdot) dzdt}$$

so

$$dy \approx F_t dt + F_x g(\cdot)dt + F_x \sigma(\cdot)dz + \frac{1}{2} F_{xx} \sigma^2(\cdot) dz^2$$

$$dy \approx F_t dt + F_x g(t,x,u)dt + F_x \sigma(t,x,u)dz + {F_{xx}\sigma^2(t,x,u)dt}/{2}.$$

This brings us to the differential equation that we can actually use:

$$dy = \left[F_t + F_x g(t,x,u) + \frac{1}{2} F_{xx} \sigma^2(t,x,u)\right] dt + F_x \sigma(t,x,u) dz, \tag{6}$$

in which we drop the ≈ symbol by assuming that dt is sufficiently small.

This expression can be extended to multiple state variables, $x_1,\ldots, x_n$, with correlations σ_{ij} as indicated in equation (7) and (8):

$$dx_i = g_i\left(t,x\right)dt + \sum_{j=1}^{n}\sigma_{ij}\left(t,x,u\right)dz_j,\ i=1,\ldots,n, \tag{7}$$

$$dy = \sum_{i=1}^{n}\frac{\partial F}{\partial x_i}dx_i + \frac{\partial F}{\partial t}dt + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}\frac{\partial^2 F}{\partial x_i \partial x_j}dx_i dx_j\ . \tag{8}$$

Equations (3) and (6) are special cases of what is known as *Itô's Lemma*, which is presented in a general form in equation (8).

B. Solving stochastic differential equations

Example #1 (KS, 266): Recall that if $\frac{\dot{y}}{y}=1$, then $y=C\cdot e^t$ where C is the constant of integration (which we will suppress from here forward). Or, using a difference equation presentation, $dy = ydt \Rightarrow y = e^t$. However, if $dy=ydz$ and z is a Wiener process, then $y = e^{z-t/2}$.

Why does stochastic integration work like this? Reproducing (3) from above, recall that if $y=F(t,z)$, then

$$dy = \left(F_t + \frac{1}{2}F_{zz}\right)dt + F_z dz\ . \tag{3}$$

So, we need a function $y=F(t,z)$ such that dy satisfies (3) and $dy=ydz$. One way that this equality will be satisfied is if the first term in (3) equals zero, i.e., $F_t = -\frac{F_{zz}}{2}$, and $F_z=y$. Using a carefully chosen "guess" that $F\left(t,z\right) = e^{z-t/2} = e^z e^{-t/2}$, we see that $F_z = F_{zz} = e^{z-t/2} = F\left(\cdot\right)$, and $F_t = -\frac{1}{2}e^{z-t/2} = -F\left(\cdot\right)/2$. Hence, if $F\left(t,z\right) = e^{z-t/2}$, then $dy = \left(F_t + \frac{1}{2}F_{zz}\right)dt + F_z dz = \left(-\frac{1}{2}e^{z-t/2} + \frac{1}{2}e^{z-t/2}\right)dt + e^{z-t/2}dz = 0 + e^{z-t/2}dz = ydz$. Hence, $y = e^{z-t/2}$ is the solution to the differential equation, $dy=ydz$ when z a Wiener process.

Example #2 (KS, 266): Let's look at another case: $dy = aydt + bydz$. Again, we start with

$$dy = \left(F_t + \frac{1}{2}F_{zz}\right)dt + F_z dz\ . \tag{3}$$

In this case

(i) $ay = F_t + \frac{1}{2}F_{zz}$ and (ii) $by = F_z$.

We use another carefully chosen "guess," $F(t,z)=y_0 e^{\left(a-\frac{b^2}{2}\right)t+bz}$ and then confirm that this works. In this case, we see that $F_z = by = bF(\cdot)$, so $F_{zz} = b\left(bF(\cdot)\right)=b^2F(\cdot)$. This can be substituted into (i) and simplified to obtain $F_t = \left(a-\frac{b^2}{2}\right)F(\cdot)$.

So we can rewrite the first term on the RHS of (3),
$\left(F_t + \frac{1}{2}F_{zz}\right)dt = \left(\left(a-\frac{b^2}{2}\right)F(\cdot)+\frac{1}{2}b^2F(\cdot)\right)dt = aF(\cdot)dt$ and the second term is

$F_z dz = bF(\cdot)dz$. Hence, if $y = F(t,z) = y_0 e^{\left(a-\frac{b^2}{2}\right)t+bz}$ then $dy = aydt + bydz$. So, this is the solution to the stochastic differential equation.

One other example taken from K&S has been placed in an appendix at the end of these notes.

II. Stochastic Control

With the background above (yes, it was just background!), we now turn to a stochastic optimization problem:

$$J(t_0,x_0)=\max_{u_t} E\left\{\int_{t=t_0}^{T} f(t,x,u)dt+\phi\left(x(T),T\right)\right\} \text{ s.t. } dx = g(t,x,u)dt+\sigma(t,x,u)dz$$

in which:

- $J(t_0,x_0)$ is the value function, a measure of the value of being at time t with a state variable x,
- x is our state variable and u is our control variable (for economists this is a confusing notational change, but it this is common to use u in this way in the control literature),
- $f(\cdot)$ is the benefit function
- $\phi(\cdot)$ is the salvage value, and
- the state equation, dx, is composed of two parts, a deterministic part $g(\cdot)$, and a stochastic component with possibly time-varying standard deviation, $\sigma(\cdot)$.

In a form reminiscent of the approach used by Dorfman (1969, see Lecture 5 notes), we can approximate $J(\cdot)$ by holding u constant for an interval from t to $t+\Delta t$ and then assuming maximizing behavior beyond that time:

$$J(t,x)\approx \max_u E\left[f(t,x,u)\Delta t + J(t+\Delta t, x+\Delta x)\right]. \tag{14}$$

Using a second-order Taylor series approximation and assuming twice-differentiability of J, we obtain

$$J(t+\Delta t,x+\Delta x)\approx J(t,x)+J_t(t,x)\Delta t+J_x(t,x)\Delta x+$$
$$(1/2)\left[J_{xx}(t,x)(\Delta x)^2+2\Delta x\Delta t J_{xt}(t,x)+J_{tt}(t,x)(\Delta t)^2\right]$$

Since we will let Δt go to zero, we can drop the $(\Delta t)^2$ and $(\Delta x\Delta t)$ terms, leaving us with equation (15):

$$J\left(t+\Delta t, x+\Delta x\right) \approx J\left(t,x\right)+J_t\left(t,x\right)\Delta t + J_x\left(t,x\right)\Delta x + \left(1/2\right)J_{xx}\left(t,x\right)\left(\Delta x\right)^2 . \tag{15}$$

Since $dx = g\left(\cdot\right)dt + \sigma\left(\cdot\right)dz$, we can rewrite

$$\left(\Delta x\right)^2 = \left(g\left(\cdot\right)\right)^2\left(\Delta t\right)^2 + \sigma^2\left(\Delta z\right)^2 + 2g\left(\cdot\right)\sigma\Delta t\Delta z .$$

Then, using the equalities from above,

$$(dz)^2 = dt, \quad dzdt = dtdz = 0 \quad \text{and} \quad (dt)^2 = 0, \tag{2}$$

we see that

$$\left(\Delta x\right)^2 = \cancel{\left(g\left(\cdot\right)\right)^2\left(\Delta t\right)^2} + \sigma^2 dt + \cancel{2g\left(\cdot\right)\sigma\Delta t\Delta z} = \sigma^2\Delta t .$$

Hence, (14) can be rewritten

$$J\left(t,x\right) \approx \max_{u_t} Ef\left(\cdot\right)\Delta t + J\left(t,x\right) + J_t\left(\cdot\right)\Delta t + J_x\left(\cdot\right)\left(g\left(\cdot\right)\Delta t + \sigma\Delta z\right) + \frac{1}{2}J_{xx}\left(\cdot\right)\sigma^2\Delta t .$$

Using the fact that $E(\Delta z)=0$ and canceling Δt we can simplify this as follows:

$$\cancel{J\left(t,x\right)} \approx \max_{u_t} Ef\left(\cdot\right)\Delta t + \cancel{J\left(t,x\right)} + J_t\left(\cdot\right)\Delta t + J_x\left(\cdot\right)\left(g\Delta t + \sigma\Delta z\right) + \frac{1}{2}J_{xx}\left(\cdot\right)\sigma^2\Delta t$$

$$0 \approx \max_{u_t} Ef\left(\cdot\right)\Delta t + \quad + J_t\left(\cdot\right)\Delta t + J_x\left(\cdot\right)\left(g\Delta t + \cancel{\sigma\Delta z}\right) + \frac{1}{2}J_{xx}\left(\cdot\right)\sigma^2\Delta t$$

$$0 \approx \max_{u_t} Ef\left(\cdot\right)\cancel{\Delta t} + \quad + J_t\left(\cdot\right)\cancel{\Delta t} + J_x\left(\cdot\right)\left(g\cancel{\Delta t} + \quad\right) + \frac{1}{2}J_{xx}\left(\cdot\right)\sigma^2\cancel{\Delta t}$$

$$0 \approx \max_{u_t} Ef\left(\cdot\right) + \quad + J_t\left(\cdot\right) \quad + J_x\left(\cdot\right)\left(g \quad + \quad\right) + \frac{1}{2}J_{xx}\left(\cdot\right)\sigma^2$$

$$-J_t\left(\cdot\right) \approx \max_u Ef\left(\cdot\right) + J_x\left(\cdot\right)g\left(\cdot\right) + \frac{1}{2}J_{xx}\left(\cdot\right)\sigma^2 . \tag{18}$$

As KS say, "This is the basic condition for the stochastic optimal control problem." It is also the stochastic version of the Hamilton-Jacobi-Bellman equation.

As shown by Xepapadeas (1997), from (18) we can obtain the stochastic version of a present-value Hamiltonian:

$$H = Ef\left(\cdot\right) + \mu g\left(\cdot\right) + \frac{1}{2}\mu_x\sigma^2 \qquad (7.51.1)^4$$

where $\mu = J_x$ and $\mu_x = J_{xx} = \partial^2 J\left(\cdot\right)/\partial x^2$.

As in deterministic problems, there are three canonical conditions for the optimum:

$$u^* = \arg\max_u H\left(x,u,\mu,\mu_x\right)$$

$$d\mu = -\frac{\partial H}{\partial x}dt + \sigma\left(\cdot\right)\mu_x dz \tag{7.51.2}$$

[4] If σ^2=0 so that uncertainty does not matter in this problem, then 7.51.1 would become the standard Hamiltonian.

where $\mu_x = \frac{\partial \mu}{\partial x}$. The boundary condition requires that

$$J(T, x_T) = \phi(x_T, T),$$

or, without a boundary condition (vertical terminal problems), the transversality condition is

$$\mu_T = 0. \tag{7.51.4}$$

If the problem is infinite horizon with standard discounting, i.e.

$$V(x_0) = \max_u E \int_{t=0}^{\infty} e^{-rt} f(x,u)\,dt \quad \text{s.t.} \quad dx = g(.)dt + \sigma(\cdot)dz$$

Then $J(t,x_t) = e^{-rt}V(x_t)$ and $-J_t(t,x_t) = -\left(-re^{-rt}V(x_t)\right)$ or, at $t=0$ $-J_t = rV(x_0)$ and (18) becomes

$$rV(x_0) \approx \max_u Ef(\cdot) + V_x(\cdot)g(\cdot) + \frac{1}{2}J_{xx}(\cdot)\sigma^2. \tag{30}$$

The Hamiltonian specification of the problem is just as in (7.51.1). The key difference is in the second optimization conditions, which now takes the form:

$$d\mu = \left(r\mu - \frac{\partial H}{\partial x}\right)dt + \sigma(\cdot)\mu_x dz$$

Noting that V_x is the shadow price on the state variable at time t, which we usually refer to as μ, then the right-hand side of (30) is the same as the Hamiltonian except for the additional term, $\frac{1}{2}\mu_x\sigma^2$.

III. Applications of stochastic control

A. A simple mathematical example

Let's start with the simple example from KS page 268. In this case the benefit function is $-\left(ax^2 + bu^2\right)$ and the stochastic state equation is $dx = udt + \sigma xdz$ with σ constant over time. In this case (18) takes the form

$$-J_t = \min_u \left[e^{-rt}\left(ax^2 + bu^2\right) + J_x u + \frac{1}{2}\sigma^2 x^2 J_{xx} \right].$$

This leads to the FOC

$$e^{-rt}\cdot 2bu + J_x = 0, \text{ or}$$

$$u = {-J_x e^{rt}}/{2b}, \tag{22}$$

which can be substituted into the objective function to obtain

$$-e^{rt}J_t = \left(ax^2 + b\left({-J_x e^{rt}}/{2b}\right)^2\right) + J_x\left({-J_x e^{rt}}/{2b}\right)e^{rt} + \frac{1}{2}\sigma^2 x^2 J_{xx}e^{rt}$$

$$-e^{rt}J_t = ax^2 + {J_x^{\,2} e^{2rt}}/{4b} - {J_x^{\,2} e^{rt}}/{2b} + \frac{1}{2}\sigma^2 x^2 J_{xx}e^{rt}$$

$$-e^{rt}J_t = ax^2 - \frac{J_x{}^2 e^{2rt}}{4b} + \frac{1}{2}\sigma^2 x^2 J_{xx} e^{rt} \tag{23}$$

This differential equation ($J_t = \partial J/\partial t$) can be integrated to obtain the solution. "Guessing" that it takes the form $J(t,x) = e^{-rt}Ax^2$, we can obtain the partial derivatives:

$$J_t = -re^{-rt}Ax^2, \quad J_x = 2e^{-rt}Ax, \quad J_{xx} = 2e^{-rt}A.$$

Substituting these into (23) we can rewrite that equation:

$$rAx^2 = ax^2 - \frac{A^2x^2}{b} + \sigma^2 x^2 A$$

$$rA = a - \frac{A^2}{b} + \sigma^2 A \tag{23b}$$

$$\frac{A^2}{b} + \left(r - \sigma^2\right)A - a = 0. \tag{25}$$

Returning to (22), it is possible to rewrite *u*:

$$u = \frac{-J_x e^{rt}}{2b} = \frac{2e^{-rt}Axe^{rt}}{2b} = -\frac{Ax}{b} \tag{27}$$

Finally, using the quadratic equation to solve 23b, we can find specific form for *A*:

$$\begin{aligned} &\frac{A^2}{b} + \left(r - \sigma^2\right)A - a = 0 \Rightarrow \\ &A^2 + b\left(r - \sigma^2\right)A - ab = 0 \Rightarrow \\ &A = \frac{-b\left(r-\sigma^2\right) \pm \sqrt{\left(b\left(r-\sigma^2\right)\right)^2 + 4ab}}{2} \end{aligned}.$$

Hence, since from (27) we know that *u* is a function of *A*, it follows that the optimal value of *u* will depend upon the degree of uncertainty, i.e., the value of σ^2.

B. An economic example: Xepapadeas (1997)

Xepapadeas (1997) constructs a relatively simple presentation of a stochastic control problem. The problem he considers is optimal management of greenhouse gases in which there is uncertainty surrounding the uptake by oceans. That is

$$dS(t) = \left[\sum_i e_i(t) - bS(t)\right]dt + \sigma S dz,\ S(0)=S_0, \tag{3.63}$$

where *S* is the stock of CO_2, e_i is emissions from the i^{th} source, and *z* is a Wiener process. Emissions lead to benefits, $B_i(e_i)$, and the stock leads to damages, $D(S)$.

Suppressing the time subscripts, the optimization problem solved by the planner is

$$\max_{\{e_i(t)\geq 0\}} E\int_0^\infty e^{-rt}\left[\sum_i B_i(e_i) - D(S)\right]dt \text{ s.t. (3.63).}$$

The stochastic Hamiltonian is

$$H = \left[\sum_i B_i(e_i) - D(S)\right] + \lambda\left(\sum_i e_i - bS\right) + \frac{1}{2}\sigma^2 S\lambda_S.$$

Note that $\lambda = V_S$, $\lambda_S = V_{SS}$ and that this is a <u>current value</u> Hamiltonian.

The optimality conditions for this problem are:

$$\frac{\partial H}{\partial e_i} = \frac{\partial B_i(e_i)}{\partial e_i} + \lambda = 0 \Rightarrow \frac{\partial B_i(e_i)}{\partial e_i} = -\lambda \text{, and} \tag{3.64}$$

$$d\lambda = \left(r\lambda - \frac{\partial H}{\partial S} \right) dt + \sigma(\cdot)\lambda_S dz = \left(r\lambda - \left(-D'(S) - \lambda b + \frac{1}{2}\sigma^2 \lambda_S \right) \right) dt + \sigma(\cdot)\lambda_S dz .$$

The second of these conditions can be simplified to

$$d\lambda = \left((r+b)\lambda + D'(S) - \frac{1}{2}\sigma^2\lambda_S \right) dt + \sigma(\cdot)\lambda_S dz . \tag{3.65}$$

The transversality condition for this infinite horizon planning problem is $\lim_{t\to\infty} E_0\left(e^{-rt}\lambda(t)S^*(t)\right) = 0$, i.e., from the perspective of t=0, the value of stocks in period t will go to zero as $t\to\infty$.

The first-order condition w.r.t. the choice variable, (3.64), is intuitive: at the optimum the marginal benefit of emissions from each source must be equal and they all must equal the marginal value (in terms of future net expected value) of the costs, i.e., MB=MC. We can easily see how this might be implemented through a tax on emissions equal to λ. The second condition, (3.65), tells us that the path of λ is stochastic. Hence, while we know the optimal level of emissions in t, we only know the distribution over that in the future.

Xepapadeas shows that for any level of S, the tax rate will be higher as uncertainty increases. He shows this by looking at the expected change in λ and S so that the terms dz drop out of (3.63) and (3.65),

$$\frac{\partial E(\lambda)}{\partial t} = (r+b)E\lambda + E\left(D'(S)\right) - \frac{1}{2}\sigma^2\lambda_S \text{ and } \frac{\partial E(S)}{\partial t} = \sum_i e_i^*(E\lambda) - bES .$$

If it is assumed that D is convex, then the isoclines in S, $E(\lambda)$ space, i.e. the points where $\partial E\lambda/\partial t = 0$ (referred to as λ_m by Xepapadeas) are below the equivalent $\dot{\lambda} = 0$ locus in S,λ space for the deterministic model (Figure 3.1 in the book reproduced below). When interpreting the figure, remember that λ<0 because it reflects the marginal *benefit* of additional S, which is negative. Hence, when there is uncertainty the certainty equivalent equilibrium level of S is lower than in the deterministic model. The risk aversion of the policy maker is captured in the term $-\frac{1}{2}\sigma\lambda_S$.

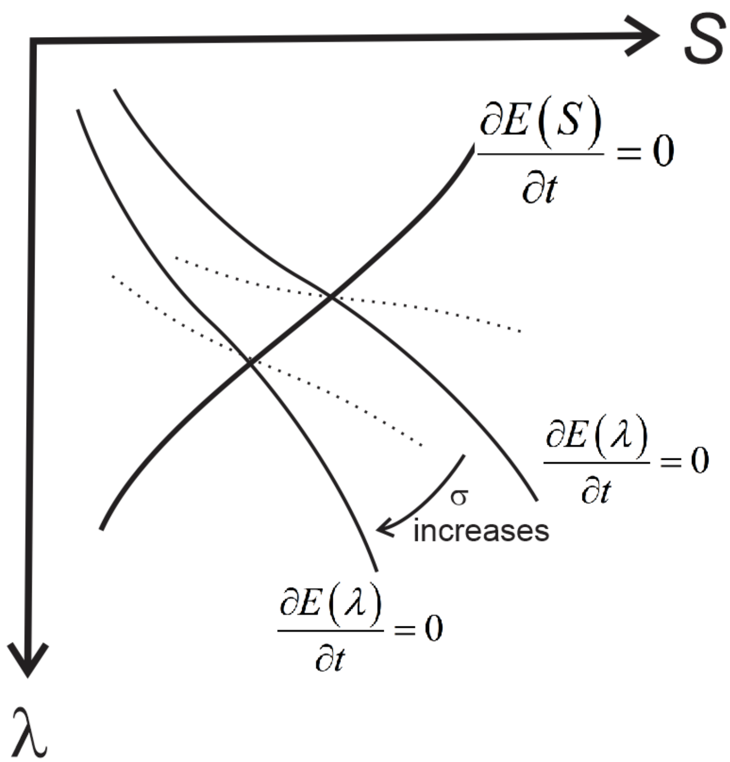

Adapted from Xepapadeas (1997) p. 76

The policy implications of this are relatively straightforward –lower levels of missions, e_i, are required for any level of S_t, meaning that the tax rate under uncertainty would be higher.

C. An economic example: McDonald and Hanf (1992)[5]

For a second economic example, I turn to one of the earliest applications of stochastic control in a natural resource application, McDonald and Hanf's 1992 analysis of a shrimp management problem. In their model there are three stochastic variables (though only 2 seem really important). Price evolves according to the equation

$$dp = \alpha p dt + \sigma_p p dz_p , \qquad \text{(MH 1)}$$

where z_p is a standard Wiener process and σ_p is the (constant) standard error term. In addition, the change in the stock is stochastic

$$ds = s(G - M)dt + s\sigma_s dz_s , \qquad \text{(MH 2)}$$

where G is the growth rate and M is the mortality rate and $\sigma_s dz_s$ is the random shock to the stock. The fishermen get benefits equal to price times harvest $p_t \cdot q$, where harvests are a function of the stock, s, and effort, h. The fisherman's costs are proportional to effort, $k \cdot h$. This leads to the optimization problem:

---

[5] I apologize for yet another example drawn from the area of environmental and resource economics. If you know of a relatively straight-forward example from other areas of economics, please share it with me and I will attempt to include it.

$$J\left(p,s,t\right)=\max_{h} E_t\int_0^T\left[p_t q_t\left(h_t,s_t\right)-kh_t\left(s_t,p_t\right)\right]e^{-rt}dt$$

subject to (MH 1), (MH 2), and

$$q_t=\rho s_t h_t\left(1-e^{-\rho h_t-\mu}\right)\Big/\rho h_t+\mu$$

where ρ is the "catchability coefficient" that controls what percentage of the stock can be harvested by a given level of effort and μ is the natural mortality rate of the shrimp.

The Hamilton-Jacobi-Bellman equation for the problem, therefore, can be written

$$\begin{aligned}-J_t=\max_{h}\Big\{&\left[p_t q_t-kh_t\right]e^{-rt}+J_p\alpha p+J_s\left(G-M\right)s\\&+\frac{1}{2}\left(J_{pp}\sigma_p{}^2p^2\right)+\frac{1}{2}\left(J_{ss}\sigma_s{}^2s^2\right)\Big\}\end{aligned}\qquad\text{(MH 6)}$$

McDonald and Hanf take the first-order condition of (MH 6) and simplify (using substantial algebra and Ito's lemma as shown in their appendix) to reach

$$\begin{aligned}dh=\left(d^2q/dh^2\right)^{-1}\Big\{&r\left(dq/dh-kp^{-1}+\sigma_p^2\left[kp^{-1}-\frac{1}{2}\left(\partial^3q/\partial h^3\right)\delta_1^2h^2\right]\right)\\&-\frac{1}{2}\left(\partial^3q/\partial h^3\right)\delta_2^2\sigma_s{}^2h^2-\alpha kp^{-1}\Big\}dt+\delta_1h\sigma_pdz_p+\delta_2h\sigma_sdz_s.\end{aligned}\qquad\text{(MH 18)}$$

To obtain some intuition behind this expression, first consider the case where there is no uncertainty, $\sigma_s=\sigma_p=0$, and there is no deterministic drift in prices, $\alpha=0$. In this case (MH 18) reduces to

$$\dot{h}=r\left(\frac{\partial q}{\partial h}-\frac{k}{p}\right)\Big/\left(d^2q/dh^2\right).$$

Hence, the rate of adjustment in the harvest rate is related to the curvature of the catch function; at the extreme, if q is linear in h, then $\dot{h}=\infty$, i.e., adjustment to the equilibrium is instantaneous. Further, the equilibrium and adjustment to the equilibrium is defined by the numerator:

$$\dot{h}=0\Leftrightarrow\frac{\partial q}{\partial h}=\frac{k}{p}\Leftrightarrow p\frac{\partial q}{\partial h}=k\ ,$$

i.e., where marginal benefit of effort, $p\dfrac{\partial q}{\partial h}$ equals the marginal cost.[6]

Equation (MH 18), therefore, defines the optimal adjustment in effort over time. The other terms, α, σ_s and σ_p are adjustments to how effort would tend to adjust over time, but do not affect the overall economic story.

---

[6] There may be a sign error somewhere in MH's expression, because as shown above, $\dot{h}<0$ if $p\dfrac{\partial q}{\partial h}>k$ when $\partial^2q/\partial h^2<0$, which is what one would typically expect.

Equations (MH 1) and (MH 2) define the stochastic evolution of the stock and the price. These three equations "make up a recursive system of equations for which econometric estimation and inference are possible" (MH p. 44). That estimation is then carried out in the rest of their paper.

IV. References

Kamien, Morton I. and Schwartz, Nancy Lou. 1991. *Dynamic Optimization : The Calculus of Variations and Optimal Control in Economics and Management*. New York, N.Y. : Elsevier.

Malliaris, A.G. and W.A. Brock. 1982. *Stochastic Methods in Economics and Finance*, Amsterdam: North-Holland.

McDonald, A David, and Claus-Hennig Hanf. 1992. Bio-economic Stability of the North Sea Shrimp Stock with Endogenous Fishing Effort. *Journal of Environmental Economics and Management* 22(1):38-56.

North-Holland Publishing Company, Amsterdam-New York-Oxford. McDonald, A David, and Claus-Hennig Hanf. 1992. Bio-economic Stability of the North Sea Shrimp Stock with Endogenous Fishing Effort. *Journal of Environmental Economics and Management* 22(1):38-56.

Xepapadeas, Anastasios. 1997. *Advanced Principles in Environmental Policy*. Northampton, Mass.: Edward Elgar.

Appendix

Example #3 (KS, 267): As a final example, we look at a problem in which the state equation is formulated as proportional change in the state variable, i.e.

$\frac{dy}{y} = cdt + b^2 dt - adt - bdz$. In this case, the trick to solving the differential equation involves defining two new variables, *P* and *Q*, $y = F(P, Q, t)$ where $dP/P = adt + bdz$ and $dQ/Q = cdt$ so that we can write $\frac{dx}{x} = \frac{dQ}{Q} - \frac{dP}{P}$, which, by integrating both sides we can obtain $\ln(x) = \ln(Q) - \ln(P) \Rightarrow x = Q/P$. Given this specification, we now apply the more general form of Itô's Lemma, (8) for two random variables,

$$dy = F_1 dx_1 + F_2 dx_2 + F_t dt + \frac{1}{2}\left(F_{11} dx_1{}^2 + F_{22} dx_2{}^2 + 2F_{12} dx_1 dx_2\right)$$

which for $y = F(P, Q, t)$ becomes

$$dy = F_P dP + F_Q dQ + \frac{1}{2}\left(F_{PP} dp^2 + F_{QQ} dx_Q{}^2 + 2F_{PQ} dP dQ\right) + F_t dt$$

Since $y=F(P,Q,t)=Q/P$, $F_Q=\frac{1}{P}$, $F_{QQ}=0$, $F_P=-\frac{Q}{P^2}$, $F_{PP}=2\frac{Q}{P^3}$, and $F_{PQ}=-\frac{1}{P^2}$, so

$$dy=-\frac{Q}{P^2}dP+\frac{dQ}{P}+\left(\frac{Q}{P^3}dp^2-\frac{dPdQ}{P^2}\right)$$

$$\frac{dy}{y}=\frac{P}{Q}\left(-\frac{Q}{P^2}dP+\frac{dQ}{P}+\left(\frac{Q}{P^3}dP^2-\frac{dPdQ}{P^2}\right)\right)$$

$$\frac{dy}{y}=-\frac{dP}{P}+\frac{dQ}{Q}+\left(\frac{dP}{P}\right)^2-\frac{dQ}{Q}\frac{dP}{P}$$

Now, using the definitions of dQ/Q and dP/P, we get

$$\frac{dy}{y}=-(adt+bdz)+cdt+(adt+bdz)^2-cdt(adt+bdz)$$

noting that $dtdz=dt^2=0$, we can simplify

$\frac{dy}{y}=-adt-bdz+cdt+b^2dz^2$. Then, finally recalling the $dz^2=dt$,

$$\frac{dy}{y}=(c+b^2-a)dt-bdz.$$

What KS do not show us is the solution. We now know that $y=Q/P$, but what are Q and P? If $dQ/Q=cdt$, then, integrating both sides we get $\ln(Q)=ct\Rightarrow Q=e^{ct}$. P is a bit trickier, $dP/P=adt+bdz$ or $dP=aPdt+bPdz$. This is of the form of Example #2, so $P=P_0e^{\left(a-\frac{b^2}{2}\right)t+bz}$.

Hence $y=\frac{Q}{P}=\frac{e^{ct}}{P_0e^{\left(a-\frac{b^2}{2}\right)t+bz}}=\frac{1}{P_0}e^{\left(c-a+\frac{b^2}{2}\right)t}e^{-bz}$.