Partial Differential Equations for Finance
G63.2706, Spring 2003
Mondays 7:10-9pm
Silver Center 207

Instructor: Robert V. Kohn. Office: 612 Warren Weaver Hall. Phone: 212-998-3217. Email: kohn@cims.nyu.edu. Web: www.math.nyu.edu/faculty/kohn. Office hours: Tuesday 4-5 and by appointment.

Teaching Assistant: Fan-fu Feng. Office: 1309. Phone: 212-998-3329.
Email: fffeng@cims.nyu.edu. Office hours: Tuesday 12-1 and Thursday 5-6. Room 613 has been reserved for the Thursday timeslot.

Special Dates: First lecture Jan. 27. No lecture Feb. 17 (Presidents' Day) and March 17 (spring break). Last lecture May 5. Final exam: May 12.

Content: An introduction to those aspects of partial differential equations and optimal control most relevant to finance. PDE's naturally associated to diffusion processes: the forward and backward Kolmogorov equations and their applications. Linear parabolic equations: fundamental solution, boundary value problems, maximum principle, transform methods. Dynamic programming and optimal control: Hamilton-Jacobi-Bellman equation, verification arguments, optimal stopping. Applications to finance will be distributed throughout the course, including: barrier, Asian, and other exotic options; pricing and hedging in incomplete markets; options on an underlying that can jump; American options, portfolio optimization, and other examples of optimal decision-making.

Prerequisites: Working knowledge of stochastic calculus, and some familiarity with financial models. The fall semester course Stochastic Calculus (G63.2902) is ideal; actually I'll assume somewhat less, roughly the material covered by Chapters 9, 10, and 11 of S. Neftci's book. See my handout "Stochastic Calculus Review" for a quick refresher.

In addition to stochastic calculus, students also need considerable scientific maturity, at a level typically obtained through an undergraduate math or science major.

Course requirements: There will be several homework sets, one every couple of weeks, probably 6 in all. Collaboration on homework is encouraged (homeworks are not exams) but registered students must write up and turn in their solutions individually. There will be one in-class final exam.

Lecture notes: Lecture notes and homework sets will be handed out, and also posted on my web-site as they become available.

Text: There is no textbook for this course – the right book just doesn't exist. Lectures will draw from many sources, including recent articles from the quantitative finance literature. See the separate handout "Library Reserve" for a list of some books that may be useful. I may from time to time make use of articles that are not easily downloaded electronically; if so then I'll place copies in the "Green Box" associated with my name, which you can request from the CIMS library staff.

Partial Differential Equations for Finance
G63.2706, Spring 2003
Reserve List

I'm providing two separate lists: first, some books that are more or less at the level of this class; then some that are more advanced than this class. I don't recommend buying *any* of these books for this class alone. But MS-level students of mathematical finance will probably find Merton and Wilmott-Howison-Dewynne well worth owning (both are inexpensive paperbacks).

Except for Dixit and Pindyck, all the books in both lists are on reserve in the CIMS library.

Books roughly at the level of this class.

- C.W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences*, Springer-Verlag 1985. No financial applications here – the book is aimed at applications in the physical sciences. But its heuristic, not-overly-rigorous style is a lot like this course, making it a useful reference for stochastic differential equations, backward and forward Kolmogorov equations, and their applications.

- R.C. Merton, *Continuous Time Finance*, Blackwell, 1992. Our discussions of several topics (stochastic optimal control, portfolio optimization, options on underlyings with jumps) will draw on Merton's classic work.

- S. Neftci, *An introduction to the mathematics of financial derivatives*, Academic Press, 1996. (A second edition was published recently, it contains everything from the first edition, plus some new material.) Chapters 9-11 cover roughly the stochastic calculus prerequisites to this class. There's relatively little about PDE's here.

- L.C. Evans, *Partial Differential Equations*, American Math Society, 1998. This is a standard graduate text on partial differential equations. But be warned: the parts relevant to this class – concerning the linear heat equation, and concerning Hamilton-Jacobi equations – is just a small part of the book.

- F. John, *Partial differential equations*, 4th edition, Springer-Verlag. Another standard graduate-level text on PDE's. The usual caveat: the part relevant to this class (on the linear heat equation) is just a small portion of the book.

- Jack Macki and Aaron Strauss, *Introduction to optimal control theory*, Springer-Verlag. Discusses deterministic optimal control (not stochastic optimal control), mainly with reference to applications in the physical sciences. Therefore it has only tangential relevance to this class.

- W. Strauss, *Partial Differential Equations; an Introduction*, John Wiley & Sons, 1992. This is a standard undergraduate text on partial differential equations. Same caveat as Evans and John: the part relevant to this class – concerning the linear heat equation – represents just a small portion of the book.

- P. Wilmott, S. Howison, and J. Dewynne, *The mathematics of financial derivatives: a student introduction*, Cambridge Univ Press, 1995. This book avoids almost all discussion of diffusion processes associated with option pricing, focusing instead as much as possible on the associated PDE's. Relatively easy to read; it goes much further than this class on numerical approximation schemes, American options, and some other PDE-related topics.

- P. Wilmott, *Derivatives*, John Wiley & Sons, 1998. Longer and deeper than the Wilmott-Howison-Dewynne book, but it has the same strong bias toward PDE (away from risk-neutral pricing). Just a few relatively short sections are directly relevant to this class.

Books more advanced than this class.

- Dimitri Bertsekas, *Dynamic programming: deterministic and stochastic models*, Prentice-Hall, 1987. No PDE here either, and very little finance. But the early chapters give some enlightening examples of discrete-time dynamic programming.

- A. Dixit and R. Pindyck, *Investment under Uncertainty*, Princeton University Press, 1994. Applies dynamic programming to financial decision-making. The sections on "mathematical background" may be helpful, though this book's focus (mainly economics) is quite different from ours (mainly option pricing and hedging). (Another book on dynamic programming with a focus on economics applications is A. Chiang, *Elements of Dynamic Optimization*, McGraw Hill, 1992.)

- W. Fleming and R. Rishel, *Deterministic and stochastic optimal control*, Springer-Verlag, 1975 (reprinted with corrections recently). A little out of date (e.g. predates viscosity solutions) but still a classic treatment of optimal control.

- I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*, Springer-Verlag, second edition, 1991. Harder to read than Oksendal, but includes some valuable topics not found there.

- R. Korn and E. Korn, *Option Pricing and Portfolio Optimization: Modern Methods of Financial Mathematics*, American Math. Society. I haven't actually seen this yet, so I can't comment on it.

- B.K. Oksendal, *Stochastic differential equations: an introduction with applications* (5th edition, paper) Springer-Verlag, 1998. A great book, well worth reading if you have the background. Some of my lectures on the backward and forward Kolmogorov equations, optimal stopping, etc. are watered-down versions of material from here.

- J. Michael Steele, *Stochastic Calculus and Financial Applications*, Springer-Verlag, 2001. The most accessible book in this list. Steele uses measure-theoretic probability, with rigorous proofs; but he always explains the main idea before addressing the nitty-gritty details – making the book delightful reading for those with sufficient background. The short chapter on diffusion equations (Chapter 11) is independent of the rest of the book and about at the level of this class.

PDE for Finance Notes, Spring 2003 – Section 1.
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course PDE for Finance, G63.2706.

Links between stochastic differential equations and PDE. A stochastic differential equation, together with its initial condition, determines a diffusion process. We can use it to define a deterministic function of space and time in two fundamentally different ways:

(a) by considering the expected value of some "payoff," as a function of the initial position and time; or

(b) by considering the probability of being in a certain state at a given time, given knowledge of the initial state and time.

Students of finance will be familiar with the Black-Scholes PDE, which amounts to an example of (a). Thus in studying topic (a) we will be exploring among other things the origin of the Black-Scholes PDE. The basic mathematical ideas here are the *backward Kolmogorov equation* and the *Feynman-Kac formula*.

Viewpoint (b) is different from (a), but not unrelated. It is in fact *dual* to viewpoint (a), in a sense that we will make precise. The evolving probability density solves a different PDE, the *forward Kolmogorov equation* – which is actually the adjoint of the backward Kolmogorov equation.

It is of interest to consider how and when a diffusion process crosses a barrier. This arises in thinking subjectively about stock prices (e.g. what is the probability that IBM will reach 200 at least once in the coming year?). It is also crucial for pricing barrier options. Probabilistically, thinking about barriers means considering *exit times*. On the PDE side this will lead us to consider *boundary value problems* for the backward and forward Kolmogorov equations.

A fairly accessible treatment of some of this material is given by Gardiner (Sections 5.1 and 5.2). Wilmott's Chapter 10 discusses exit times (but too briefly!). Parts of my notes draw from Fleming-Rishel and Oksendal, however the treatments there are much more general and sophisticated so not easy to read.

Our main tool will be Ito's formula, coupled with the fact that any Ito integral of the form $\int_a^b f\,dw$ has expected value zero. (Equivalently: $m(t) = \int_a^t f\,dw$ is a martingale.) Here w is Brownian motion and f is non-anticipating. The stochastic integral is defined as the limit of Ito sums $\sum_i f(t_i)(w(t_{i+1} - w(t_i))$ as $\Delta t \to 0$. The sum has expected value zero because each of its terms does: $E[f(t_i)(w(t_{i+1}) - w(t_i))] = E[f(t_i)]E[w(t_{i+1}) - w(t_i)] = 0$.

****

Expected values and the backward Kolmogorov equation. Here's the most basic version of the story. Suppose $y(t)$ solves the scalar stochastic differential equation

$$dy = f(y,s)ds + g(y,s)dw,$$

and let

$$u(x,t) = E_{y(t)=x}\left[\Phi(y(T))\right]$$

be the expected value of some payoff Φ at maturity time $T > t$, given that $y(t) = x$. Then u solves

$$u_t + f(x,t)u_x + \tfrac{1}{2}g^2(x,t)u_{xx} = 0 \text{ for } t < T, \text{ with } u(x,T) = \Phi(x). \tag{1}$$

The proof is easy: for any function $\phi(y,t)$, Ito's lemma gives

$$\begin{aligned} d(\phi(y(s),s)) &= \phi_y dy + \tfrac{1}{2}\phi_{yy}dydy + \phi_s ds \\ &= (\phi_s + f\phi_y + \tfrac{1}{2}g^2\phi_{yy})dt + g\phi_y dw. \end{aligned}$$

Choosing $\phi = u$, the solution of (1), we get

$$u(y(T),T) - u(y(t),t) = \int_t^T (u_t + fu_y + \tfrac{1}{2}g^2u_{yy})ds + \int_t^T gu_y dw.$$

Taking the expected value and using the PDE gives

$$E_{y(t)=x}\left[\Phi(y(T))\right] - u(x,t) = 0$$

which is precisely our assertion.

That was the simplest case. It can be jazzed up in many ways. We discuss some of them:

Vector-valued diffusion. Suppose y solves a vector-valued stochastic differential equation

$$dy_i = f_i(y,s)ds + \sum_j g_{ij}(y,s)dw_j,$$

where each component of w is an independent Brownian motion. Then

$$u(x,t) = E_{y(t)=x}\left[\Phi(y(T))\right]$$

solves

$$u_t + \mathcal{L}u = 0 \text{ for } t < T, \text{ with } u(x,T) = \Phi(x),$$

where $\mathcal{L}$ is the differential operator

$$\mathcal{L}u(x,t) = \sum_i f_i \frac{\partial u}{\partial x_i} + \tfrac{1}{2}\sum_{i,j,k} g_{ik}g_{jk}\frac{\partial^2 u}{\partial x_i \partial x_j}.$$

The justification is just as in the scalar case, using the multidimensional version of Ito's lemma. The operator $\mathcal{L}$ is called the "infinitesimal generator" of the diffusion process $y(s)$.

The Feynman-Kac formula. We discuss the scalar case first, for clarity. Consider as above the solution of

$$dy = f(y,s)dt + g(y,s)dw$$

but suppose we are interested in a suitably "discounted" final-time payoff of the form:

$$u(x,t) = E_{y(t)=x}\left[e^{-\int_t^T b(y(s),s)ds}\Phi(y(T))\right] \tag{2}$$

for some specified function $b(y)$. Then u solves

$$u_t + f(x,t)u_x + \tfrac{1}{2}g^2(x,t)u_{xx} - b(x,t)u = 0 \tag{3}$$

instead of (1). (Its final-time condition is unchanged: $u(x,T) = \Phi(x)$.) If you know some finance you'll recognize that when y is log-normal and b is the interest rate, (3) is precisely the Black-Scholes partial differential equation. Also: if $b(y(s),s)$ is the spot interest rate, then (3) with $\Phi = 1$ gives the time-t value of a zero-coupon bond with maturity T, given that the spot interest rate at time t is $b(x,t)$.

To explain (3), we must calculate the stochastic differential $d[z_1(s)\phi(y(s),s)]$ where $z_1(s) = e^{-\int_t^s b(y(r))dr}$. The multidimensional version of Ito's lemma gives

$$d[z_1(s)z_2(s)] = z_1 dz_2 + z_2 dz_1 + dz_1 dz_2.$$

We apply this with z_1 as defined above and $z_2(s) = \phi(y(s),s)$. Ito's lemma (or ordinary differentiation) gives

$$dz_1(s) = -z_1 b(y(s))ds$$

and we're already familiar with the fact that

$$\begin{aligned} dz_2(s) &= (\phi_s + f\phi_y + \tfrac{1}{2}g^2\phi_{yy})ds + g\phi_y dw \\ &= (\phi_s + \mathcal{L}\phi)ds + g\phi_y dw. \end{aligned}$$

Notice that $dz_1 dz_2 = 0$. Applying the above with $\phi = u$, the solution of the PDE (3), gives

$$\begin{aligned} d\left(e^{-\int_t^s b(y(r))dr}u(y(s),s)\right) &= z_1 dz_2 + z_2 dz_1 \\ &= z_1\left[(u_s + \mathcal{L}u)ds + gu_y dw\right] - z_1 ubds \\ &= z_1 gu_y dw. \end{aligned}$$

The right hand side has expected value 0, so

$$E_{y(t)=x}[z_1(T)z_2(T)] = z_1(t)z_2(t) = u(x,t)$$

as asserted.

A moment's thought reveals that vector-valued case is no different. The discounted expected payoff (2) solves the PDE

$$u_t + \mathcal{L}u - bu = 0$$

where $\mathcal{L}$ is the infinitesimal generator of the diffusion y.

Running payoff. Suppose we are interested in

$$u(x,t) = E_{y(t)=x}\left[\int_t^T \Psi(y(s),s)ds\right]$$

for some specified function Ψ. Then u solves

$$u_t + \mathcal{L}u + \Psi(x,t) = 0.$$

The final-time condition is $u(x,T) = 0$, since we have included no final-time term in the "payoff." The proof is hardly different from before: by Ito's lemma,

$$\begin{aligned} d[u(y(s),s)] &= (u_t + \mathcal{L}u)ds + \nabla u \cdot g \cdot dw \\ &= -\Psi(y(s),s)ds + \nabla u \cdot g \cdot dw. \end{aligned}$$

Integrating and taking the expectation gives

$$E_{y(t)=x}\left[u(y(T),T)\right] - u(x,t) = E_{y(t)=x}\left[-\int_t^T \Psi(y(s),s)ds\right].$$

This gives the desired assertion, since $u(y(T),T) = 0$.

In valuing options, "running payoffs" are relatively rare. However terms of this type will be very common later in the course, when discuss optimal control problems.

****

Boundary value problems and exit times. The preceding examples use stochastic integration from time t to a fixed time T, and they give PDE's that must be solved for all $x \in R^n$. It is also interesting to consider integration from time t to the first time y exits from some specified region. The resulting PDE must be solved on this region, with suitable boundary data.

Let D be a region in R^n. Suppose y is an R^n-valued diffusion solving

$$dy = f(y,s)ds + g(y,s)dw \text{ for } s > t, \text{ with } y(t) = x$$

with $x \in D$. Let

$$\begin{aligned} \tau(x) &= \text{the first time } y(s) \text{ exits from } D, \text{ if} \\ &\quad \text{prior to } T; \text{ otherwise } \tau(x) = T. \end{aligned}$$

This is an example of a *stopping time*. (Defining feature of a stopping time: the statement "$\tau(x) < t$" is $\mathcal{F}_t$-measurable; in other words, the decision whether to stop or not at time t depends only on knowledge of the process up to time t. This is clearly true of the exit time defined above.)

Here is the basic result: the function

$$u(x,t) = E_{y(t)=x}\left[\int_t^{\tau(x)} \Psi(y(s),s)ds + \Phi(y(\tau(x)),\tau(x))\right]$$

solves

$$u_t + \mathcal{L}u + \Psi = 0 \text{ for } x \in D$$

with boundary condition

$$u(x,t) = \Phi(x,t) \text{ for } x \in \partial D \tag{4}$$

and final-time condition

$$u(x,T) = \Phi(x,T) \text{ for all } x \in D. \tag{5}$$

The justification is entirely parallel to our earlier examples. The only change is that we integrate, in the final step, to the stopping time τ rather than the final time T. (This is permissible for any stopping time satisfying $E[\tau] < \infty$. The statement that $E\left[\int_t^\tau f\, dw\right] = 0$ when $E[\tau] < \infty$ is known as Dynkin's theorem.)

A subtlety is hiding here: the hypothesis that $E[\tau] < \infty$ is not a mere technicality. Rather, there are simple and interesting examples where it is false and $E\left[\int_t^\tau f\, dw\right] \neq 0$. One such example is related to the "gambler's ruin" paradox. Consider the standard Brownian process $w(s)$, starting at $w(0) = 0$. Let τ_* be the first time $w(s)$ reaches 1. Then $w(\tau_*) - w(0) = 1 - 0 = 1$ certainly does not have mean 0, so $E\left[\int_0^{\tau_*} dw\right] \neq 0$ in this case. This doesn't contradict Dynkin's theorem; it just shows that $E[\tau_*] = \infty$. To understand the situation better, consider τ_n= the time of first exit from $[-n, 1]$. You'll show on HW1 that $E[\tau_n] < \infty$ for each n, but $E[\tau_n] \to \infty$ as $n \to \infty$. The Brownian motion process eventually reaches 1 with probability one, but it may make extremely large negative excursions before doing so. Here's the coin-flipping version of this situation: consider a gambler who decides to bet by flipping coins and never quitting till he's ahead by a fixed amount. If there is no limit on the amount he is permitted to lose along the way, then he'll eventually win with probability one. But if there is a threshold of losses beyond which he must stop then there is a nonzero probability of ruin and his expected outcome is 0.

There's something slightly misleading about our notation in (4)-(5). We use the same notation Φ for both the boundary condition (4) and the final-time condition (5) because they come from the same term in the payoff: $\Phi(y(\tau), \tau)$ where τ is the time the curve $(y(s), s)$ exits from the cylinder $D \times [0, T]$. But Φ should be thought of as representing two *distinct* functions – one at the spatial boundary $\partial D \times [0, T]$, the other at the final time boundary $D \times \{T\}$ (see the figure). These two functions need have nothing to do with one

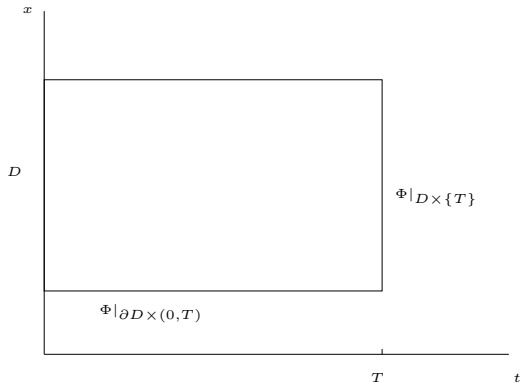

Figure 1: *Distinguishing between the two different parts of* Φ.

another. Often one is chosen to be zero, while the other is nontrivial. [A financial example:

when one values a barrier option using the risk-neutral expectation of the payoff, Φ is zero at the knock-out price, and it equals the payoff at the maturity time.]

Elliptic boundary-value problems. Now suppose f and g in the stochastic differential equation don't depend on t, and for $x \in D$ let

$$\tau(x) = \text{the first time } y(s) \text{ exits from } D.$$

(Unlike the previous example, we do not impose a final time T. Clearly this amounts to taking $T = \infty$ in the previous definition.) Suppose furthermore the process does eventually exit from D, (and more: assume $E[\tau(x)] < \infty$, for all $x \in D$). Then

$$u(x) = E_{y(0)=x}\left[\int_0^{\tau(x)} \Psi(y(s))ds + \Phi(y(\tau(x)))\right]$$

solves

$$\mathcal{L}u + \Psi = 0 \text{ for } x \in D,$$

with boundary condition

$$u = \Phi \text{ for } x \in \partial D.$$

The justification is again entirely parallel to our earlier examples.

****

Applications: some properties of the Brownian motion process. Let us use these results to deduce – by solving appropriate PDE's – some properties of the Brownian motion process. (This discussion is taken Oksendal's example 7.4.2. Related material on exit times will be explored in HW1.)

QUESTION 1. Consider n-dimensional Brownian motion starting at x. What is the mean time it takes to exit from a ball of radius R, for $R > |x|$? Answer: apply the last example with $f = 0$, $g =$ identity matrix, $\Psi = 1$, $\Phi = 0$. It tells us the mean exit time is the solution $u(x)$ of

$$\tfrac{1}{2}\Delta u + 1 = 0$$

in the ball $|x| < R$, with $u = 0$ at $|x| = R$. The (unique) solution is

$$u(x) = \frac{1}{n}(R^2 - |x|^2).$$

(To do this calculation we must know in advance that the expected exit time is finite. We'll justify this as Question 3 below.)

QUESTION 2. Consider the scalar lognormal process

$$dy = \mu y dt + \sigma y dw$$

with μ and σ constant. Starting from $y(0) = x$, what is the mean exit time from a specified interval (a, b) with $a < x < b$? Answer: the mean exit time $u(x)$ solves

$$\mu x u_x + \tfrac{1}{2}\sigma^2 x^2 u_{xx} + 1 = 0 \text{ for } a < x < b$$

with boundary conditions $u(a) = u(b) = 0$. The solution is

$$u(x) = \frac{1}{\frac{1}{2}\sigma^2 - \mu}\left(\log(x/a) - \frac{1 - (x/a)^{1-2\mu/\sigma^2}}{1 - (b/a)^{1-2\mu/\sigma^2}}\log(b/a)\right)$$

(readily verified by checking the equation and boundary conditions). This answer applies only if $\mu \neq \frac{1}{2}\sigma^2$. See HW1 for the case $\mu = \frac{1}{2}\sigma^2$.

QUESTION 3: Returning to the setting of Question 1, how do we know the mean exit time is finite? Answer: assume D is a bounded domain in R^n, and $y(s)$ is multidimensional Brownian motion starting at $x \in D$. Recall that by Ito's lemma, $t \to \phi(y(t))$ satisfies

$$d\phi = \nabla\phi \, dw + \tfrac{1}{2}\Delta\phi \, dt \tag{6}$$

for any function ϕ. Let's apply this with $\phi(y) = |y|^2$, integrating in time up to the stopping time

$$\tau_T(x) = \min\{\tau(x), T\} = \begin{cases} \text{first time } y(s) \text{ exits from } D & \text{if less than } T \\ T & \text{otherwise.} \end{cases}$$

We get

$$\begin{aligned} E\left[|y(\tau_T(x))|^2\right] - |x|^2 &= \frac{1}{2}\int_0^{\tau_T(x)} \Delta\phi(y(s))ds \\ &= nE\left[\tau_T(x)\right] \end{aligned} \tag{7}$$

since $\Delta\phi = 2n$. Now let $T \to \infty$. The left hand side of (7) stays finite, since we're considering a *bounded* domain, and by definition $y(\tau_T(x))$ is either in D or on the boundary of D. Thus we conclude that

$$\lim_{T\to\infty} E\left[\tau_T(x)\right] < \infty.$$

It follows (using the monotone convergence theorem, from real variables) that the exit time $\tau = \lim_{T\to\infty} \tau_T$ is almost surely finite, and $E[\tau] < \infty$, for any starting point $x \in D$.

QUESTION 4: Consider Brownian motion in R^n, starting at a point x with $|x| = b$. Given $r < b$, what is the probability that the path ever enters the ball of radius r centered at 0? Answer: for $n = 1, 2$ this probability is 1. (Interpretation: Brownian motion is "recurrent" in dimensions 1 and 2 – it comes arbitrarily close to any point, infinitely often, regardless of where it starts.) In higher dimensions the situation is different: in dimension $n \geq 3$ the probability of entering the ball of radius r is $(b/r)^{2-n}$. (Interpretation: Brownian motion is "transient" in dimension $n \geq 3$.)

Consider first the case $n \geq 3$. We use the stopping time $\tau_k =$ first exit time from the annulus

$$D_k = \{r < |x| < 2^k r\}.$$

Since D_k is bounded, $E[\tau_k] < \infty$ and we can integrate the stochastic differential equation (6) up to time τ_k. Let's do this with the special choice

$$\phi(y) = |y|^{2-n}.$$

This ϕ solves Laplace's equation $\Delta\phi = 0$ away from its singularity at $y = 0$. (The singularity does not bother us, since we only evaluate ϕ at points $y(s) \in D_k$ and 0 does not belong to D_k.) The analogue of (7) is

$$E\left[|y(\tau_k)|^{2-n}\right] - b^{2-n} = \frac{1}{2}\int_0^{\tau_k} \Delta\phi(y(s))ds = 0.$$

If p_k is the probability that y leaves the annulus D_k at radius r, and $q_k = 1 - p_k$ is the probability that it leaves the annulus at radius $2^k r$, we have

$$r^{2-n}p_k + (2^k r)^{2-n}q_k = b^{2-n}.$$

As $k \to \infty$ this gives $p_k \to (b/r)^{2-n}$, as asserted.

The case $n = 2$ is treated similarly, using

$$\phi(y) = \log y,$$

which solves $\Delta\phi = 0$ in the plane, away from $y = 0$. Arguing as before we get

$$p_k \log r + q_k \log(2^k r) = \log b.$$

As $k \to \infty$ this gives $q_k \to 0$. So $p_k \to 1$, as asserted.

The case $n = 1$ is similar to $n = 2$, using $\phi(y) = |y|$.

****

Another application: distribution of first arrivals. Consider a scalar diffusion whose drift and volatility are functions of y alone, independent of t:

$$dy = f(y(s))ds + g(y(s))dw.$$

The initial condition is $y(t) = x$. We are interested in the first arrival of $y(s)$ at a given threshold, say $y = 0$. Assume to fix ideas that $x > 0$.

What is the distribution of arrival times? Let the density of arrival times be $\rho(s)$. Its cumulative distribution function $\int_0^T \rho(s)\,ds$ is the probability that the first arrival occurs by time T. According to our discussion of the backward Kolmogorov equation, this is $u(x, 0)$ where u solves

$$u_t + fu_x + \tfrac{1}{2}g^2 u_{xx} = 0 \quad \text{for } x > 0, \quad 0 < t < T \tag{8}$$

with boundary condition

$$u = 1 \quad \text{at } x = 0$$

and final-time condition

$$u = 0 \quad \text{at } t = T.$$

Clearly u depends also on the final time T; let's make that dependence explicit by writing $u = u(x, t; T)$. Evidently $u(x, 0; T) = \int_0^T \rho(s)\, ds$, so by differentiation we get

$$\rho(s) = \frac{\partial u}{\partial T}(x, 0; s).$$

For special choices of f and g (e.g. Brownian motion with drift, or lognormal) the PDE (8) can be solved explicitly (we'll discuss how to do this later), yielding an explicit formula for the distribution of first arrivals.

Suppose the mean arrival time is finite. Then we know it should be given by $v(x)$ where $fv_x + \frac{1}{2}gv_{xx} = -1$ for $x > 0$ with $v = 0$ at $x = 0$. On the other hand, the mean arrival time is

$$\int_0^\infty s\rho(s)\, ds = \int_0^\infty s\partial_s u(x, 0; s)\, ds.$$

Are these apparently different expressions consistent? Yes indeed! To show this, we observe (writing u_s for $\partial_s u(x, 0; s)$, for simplicity) that

$$\int_0^\infty s u_s\, ds = -\int_0^\infty s(1-u)_s\, ds = \int_0^\infty (1-u)\, ds$$

by integration by parts, since $\int_0^\infty \partial_s[s(1-u)]\, ds = s(1-u)|_0^\infty = 0$. Moreover the function $v(x) = \int_0^\infty (1-u)\, ds$ clearly satisfies $v = 0$ at $x = 0$, and

$$fv_x + \tfrac{1}{2}g^2 v_{xx} = -\int_0^\infty f u_x + \tfrac{1}{2}g^2 u_{xx}\, ds.$$

But since f and g are independent of time, $u(x, t; T)$ depends on t and T only through $T - t$, so $\partial u/\partial T = -\partial u/\partial t$. Therefore, using the backward Kolmogorov equation,

$$-\int_0^\infty f u_x + \tfrac{1}{2}g^2 u_{xx}\, ds = \int_0^\infty u_t\, ds = -\int_0^\infty \partial_s u(x, 0; s)\, ds.$$

The last expression is clearly $u(x, 0; 0) - u(x, 0; \infty) = 0 - 1 = -1$. Thus v solves the anticipated PDE.

****

Transition probabilities and the forward Kolmogorov equation. We've shown that when the state evolves according to a stochastic differential equation

$$dy_i = f_i(y, s)ds + \sum_j g_{ij}(y, s)dw_j$$

the expected final position

$$u(x, t) = E_{y(t)=x}\left[\Phi(y(T))\right]$$

solves the backward Kolmogorov equation

$$u_t + \sum_i f_i \frac{\partial u}{\partial x_i} + \tfrac{1}{2} \sum_{i,j,k} g_{ik} g_{jk} \frac{\partial^2 u}{\partial x_i \partial x_j} = 0 \text{ for } t < T, \text{ with } u = \Phi \text{ at } t = T. \tag{9}$$

We can write the backward Kolmogorov equation as

$$u_t + \mathcal{L} u = 0 \tag{10}$$

with

$$\mathcal{L} u = \sum_i f_i \frac{\partial u}{\partial x_i} + \sum_{i,j} a_{ij} \frac{\partial^2 u}{\partial x_i \partial x_j}, \tag{11}$$

where $a_{ij} = \frac{1}{2} \sum_k g_{ik} g_{jk} = \frac{1}{2} (g g^T)_{ij}$.

The solution of the stochastic differential equation is a *Markov process*, so it has a well-defined *transition probability*

$p(z, s; x, t) =$ probability of being at z at time s, given that it started at x at time t.

More precisely: $p(\cdot, s; x, t)$ is the probability density of the state at time s, given that it started at x at time t. Of course p is only defined for $s > t$. To describe a Markov process, p must satisfy the Chapman-Kolmogorov equation

$$p(z, s; x, t) = \int_{R^n} p(z_1, s_1; x, t) p(z, s; z_1, s_1) \, dz_1$$

for any s_1 satisfying $t < s_1 < s$. Intuitively: the state can get from (x, t) to (z, s) by way of being at various intermediate states z_1 at a chosen intermediate time s_1. The Chapman-Kolmogorov equation calculates $p(z, s; x, t)$ by adding up (integrating) the probabilities of getting from (x, t) to (z, s) via (z_1, s_1), for all possible intermediate positions z_1.

How should we visualize p? Consider first the case when y is multidimensional Brownian motion. Then $p(\cdot, s; x, t)$ is the density of a Gaussian random variable with mean x and variance $s - t$. The graph of $z \to p(z, s; x, t)$ always has volume 1 below it (since p is a probability density); as $s \to \infty$ its maximum value tends to 0 (a Brownian particle diffuses further and further away, on average, as time increases); as $s \to t$ it becomes infinitely tall and thin (at time $s \approx t$ the Brownian particle is very close to its initial position x). The situation for a general stochastic differential equation is similar: p becomes infinitely tall and thin, concentrating at $z = x$, as $s \to t$; and if $g g^T > 0$ then the graph of p keeps spreading as $s \to \infty$. Of course in the general case p does not describe a Gaussian distribution, and there is no simple formula for the mean or variance – they are simply the mean and variance of $y(s)$.

If the stochastic differential equation does not involve time explicitly, then the transition probability depends only on the "elapsed time":

if $dy = f(y)dt + g(y)dw$ with f, g depending only on y, then $p(z, s; x, t) = p(z, s - t; x, 0)$.

If the stochastic differential equation does not involve space explicitly, then the transition probability depends only on the "relative position":

if $dy = f(t)dt + g(t)dw$ with f, g depending only on t, then $p(z, s; x, t) = p(z - x, s; 0, t)$.

The initial position of a Markov process need not be deterministic. Even if it is (e.g. if $y(0) = x$ is fixed), we may wish to consider a later time as the "initial time." The transition probability determines the evolution of the spatial distribution, no matter what its initial value: if $\rho_0(x)$ is the probability density of the state at time t then

$$\rho(z, s) = \int_{R^n} p(z, s; x, t)\rho_0(x)\, dx \tag{12}$$

gives the probability density (as a function of z) at any time $s > t$.

The crucial fact about the transition probability is this: it solves the *forward Kolmogorov equation* in s and z:

$$-p_s - \sum_i \frac{\partial}{\partial z_i}\left(f_i(z, s)p\right) + \tfrac{1}{2}\sum_{i,j,k} \frac{\partial^2}{\partial z_i \partial z_j}\left(g_{ik}(z, s)g_{jk}(z, s)p\right) = 0 \text{ for } s > t, \tag{13}$$

with initial condition

$$p = \delta_x(z) \text{ at } s = t.$$

We can write the forward Kolmogorov equation as

$$-p_s + \mathcal{L}^* p = 0 \tag{14}$$

with

$$\mathcal{L}^* p = -\sum_i \frac{\partial}{\partial z_i}\left(f_i p\right) + \sum_{i,j} \frac{\partial^2}{\partial z_i \partial z_j}\left(a_{ij} p\right). \tag{15}$$

Here $a_{ij} = \frac{1}{2}(gg^T)_{ij}$ just as before. The initial condition $p = \delta_x(z)$ encapsulates the fact, already noted, that the graph of $p(\cdot, s; x, t)$ becomes infinitely tall and thin at x as s decreases to t. The technical meaning is that

$$\int_{R^n} p(z, s; x, t)f(z)\, dz \to f(x) \text{ as } s \text{ decreases to } t \tag{16}$$

for any continuous f.

Recall that if the initial state distribution is ρ_0 then the evolving distribution is $\rho(z, s) = \int p(z, s; x, t)\rho_0(x)\, dx$. This function $\rho(z, s)$ automatically solves the forward equation (just bring the derivatives under the integral, and use that p solves it). The initial condition on p is just what we need to have $\rho(z, s) \to \rho_0(z)$ as $s \to t$. (Demonstration: multiply (16) by $\rho_0(x)$ and integrate in x to see that

$$\int \rho(z, s)f(z)\, dz = \int p(z, s; x, t)f(z)\rho_0(x)\, dzdx \to \int f(x)\rho_0(x)\, dx$$

as $s \to t$. Since this is true for every continuous f, we conclude that $\rho(z,s)$ converges [weakly] to $\rho_0(z)$ as $s \to t$.)

Please note that the forward Kolmogorov equation describes the probability distribution by solving an initial-value problem, while the backward Kolmogorov equation describes the expected final payoff by solving a final-value problem. Students familiar with pricing options via binomial trees will find this familiar. The stock prices at various nodes of a tree are determined by working forward in time; the option values at various nodes of a tree are determined by working backward in time.

Notice that the forward and backward Kolmogorov equations are, in general, completely different. There is one case, however, when they are closely related: for Brownian motion the forward equation starting at $t = 0$ is

$$p_s - \tfrac{1}{2}\Delta p = 0 \text{ for } s > 0$$

while the backward equation with final time T is

$$u_t + \tfrac{1}{2}\Delta u = 0 \text{ for } t < T.$$

In this special case the backward equation is simply the forward equation with time reversed. More careful statement: if $u(x,t)$ solves the backward equation then $\tilde{u}(z,s) = u(z, T-s)$ solves the forward equation, and conversely. This is an *accident,* associated with the the self-adjointness of the Laplacian. The situation is different even for Brownian motion with constant drift f: then the forward equation is $p_s + f \cdot \nabla p - \frac{1}{2}\Delta p = 0$, while the backward equation is $u_t + f \cdot \nabla u + \frac{1}{2}\Delta u = 0$, and the two are not equivalent under time-reversal.

Students with a background in physical modeling will be accustomed to equations of the form $v_t = \text{div } (a(x)\nabla v)$. Neither the forward nor the backward Kolmogorov equation has this form. Such equations are natural in physics, but not in problems from control theory and stochastic differential equations.

Application: steady-state distributions. The backward Kolmogorov equation comes up more often than the forward one in finance. But one important application involves the large-time behavior of a diffusion. If $\rho(z,s)$ is the probability density of a diffusion, then evidently $\rho_\infty(z) = \lim_{s\to\infty} \rho(z,s)$ represents (if it exists) the large-time statistics of the process. For Brownian motion $\rho_\infty = 0$, reflecting the fact that Brownian particles wander a lot. The situation is quite different however for the Ornstein-Uhlenbeck process $dy = -kyds + \sigma dw$. We expect y to remain near 0 due to the deterministic term $-ky$, which constantly pushes it toward 0. And indeed the steady-state distribution is

$$\rho_\infty(z) = Ce^{-kz^2/\sigma^2}$$

where C is chosen so that ρ_∞ has integral 1. (It's easy to check that this gives a steady-state solution of the forward Kolmogorov equation. In fact this gives the long-time asymptotics of a fairly general initial condition, but this is *not* so obvious.)

This application can be generalized. Consider the stochastic PDE $dy = -V'(y)dt + \sigma dw$. Its deterministic part pushes y toward a local minima of V. If V grows rapidly enough at

infinity (so the diffusion is succesfully confined, and does not wander off to infinity) then the long-time statistics are described by the steady-state distribution

$$\rho_\infty(z) = Ce^{-2V(z)/\sigma^2}.$$

Testing the plausibility of the forward equation. We will explain presently why the forward equation holds. But first let's get used to it by examining some consequences and checking some special cases. Let $\rho_0(x)$ be the probability density of the state at time 0, and consider

$$\rho(z,s) = \int p(z,s;x,0)\rho_0(x)\,dx$$

for $s > 0$. It gives the probability density of the state at time s.

Checking the integral. Since ρ is a probability density we expect that $\int \rho(z,s)\,dz = 1$ for all s. In fact, from the forward equation

$$\begin{aligned}\frac{d}{ds}\int \rho\,dz &= \int \rho_s\,dz \\ &= \int \mathcal{L}^*\rho\,dz \\ &= 0\end{aligned}$$

since each term of $\mathcal{L}^*\rho$ is a perfect derivative. (Here and below, we repeatedly integrate by parts, with no "boundary terms" at $\pm\infty$. We are implicitly assuming that ρ and its derivatives decay rapidly as $z \to \pm\infty$. This is true, provided the initial distribution ρ_0 has this property.)

If the stochastic differential equation has no drift then the expected position is independent of time. In general, $E[y(s)] - E[y(0)] = E\int_0^s f(y(r),r)\,dr$ since the expected value of the integral dw vanishes. Thus when $f = 0$ the expected position $E[y(s)]$ is constant. Let's prove this again using the forward equation:

$$\begin{aligned}\frac{d}{ds}(\text{expected position}) &= \frac{d}{ds}\int z\rho(z,s)\,dz \\ &= \int z\rho_s(z,s)\,dz \\ &= \int z\mathcal{L}^*\rho(z,s)\,dz \\ &= 0 \quad \text{when } f = 0.\end{aligned}$$

The last step is the result of integration by parts; for example, if y is scalar valued ($dy = g(y,t)dw$) we have

$$\begin{aligned}\int z\mathcal{L}^*\rho\,dz &= \tfrac{1}{2}\int z\left(g^2\rho\right)_{zz}\,dz \\ &= -\tfrac{1}{2}\int \left(g^2\rho\right)_z\,dz \\ &= 0.\end{aligned}$$

(As noted above, to justify the integrations by parts one must know that ρ vanishes rapidly enough at spatial infinity.)

The special case f = constant, $g = 0$. If $g = 0$ then we're studying a deterministic motion. If in addition f = constant then the solution is explicit and very simple: $y(t) = y(0) + ft$. Clearly

$$\text{Prob of being at } z \text{ at time } s = \text{Prob of being at } z - fs \text{ at time } 0,$$

whence

$$\rho(z, s) = \rho_0(z - fs).$$

In particular, $\rho_s + f \cdot \nabla\rho = 0$, which agrees with the forward equation (since f is constant).

Biting the bullet. Enough playing around; let's explain why the forward equation holds. The first main ingredient is the observation that

$$E_{y(t)=x}\left[\Phi(y(T))\right] = \int \Phi(z)p(z, T; x, t)\, dz. \tag{17}$$

We know how to determine the left hand side (by solving the backward equation, with final value Φ at $t = T$). This relation determines the integral of $p(\cdot, T; x, t)$ against any function Φ, for any value of x, t, T. This is a lot of information about p – in fact, it fully determines p. Our task is to make this algorithmic, i.e. to explain how p can actually be computed. (The answer, of course, will be to solve the forward equation in z and s.)

The second main ingredient is the relation between $\mathcal{L}$ and $\mathcal{L}^*$. Briefly: $\mathcal{L}^*$ is the *adjoint* of $\mathcal{L}$ in the L^2 inner product. Explaining this: recall from linear algebra that if A is a linear operator on an inner-product space, then its adjoint A^* is defined by

$$\langle Ax, y\rangle = \langle x, A^*y\rangle.$$

When working in R^n we can represent A by a matrix, and A^* is represented by the transpose A^T. The situation is similar here, but our inner product space consists of all (square-integrable, scalar-valued) functions on R^n, with inner product

$$\langle v, w\rangle = \int_{R^n} v(x)w(x)\, dx.$$

We claim that

$$\langle \mathcal{L}v, w\rangle = \langle v, \mathcal{L}^*w\rangle. \tag{18}$$

When y is scalar-valued our claim says that

$$\int_R \left(fv_x + \tfrac{1}{2}g^2v_{xx}\right) w\, dx = \int_R v\left(-(fw)_x + \tfrac{1}{2}(g^2w)_{xx}\right)\, dx.$$

This is a consequence of integration by parts. For example, the first term on the left equals the first term on the right since

$$\int_R [fw]v_x\, dx = -\int_R [fw]_x v\, dx.$$

The second term on each side matches similarly, integrating by parts twice. Notice that f and g can depend on time as well as space; it doesn't change the argument. The proof of (18) when y is vector valued is essentially the same as the scalar case.

The third main ingredient is hiding in our derivation of the backward equation. We know from this derivation that

$$E_{y(t)=x}\left[\phi(y(T),T)\right] - \phi(x,t) = E_{y(t)=x}\left[\int_t^T (\phi_s + \mathcal{L}\phi)(y(s),s)\,ds\right] \tag{19}$$

for any function $\phi(y,s)$. Our main use of this relation up to now was to choose ϕ so that the right hand side vanished, i.e. to choose ϕ to solve the backward equation. But we don't have to make such a restrictive choice: relation (19) holds for *any* ϕ.

Let's put these ingredients together. Rewriting (19) using the transition probabilities gives

$$\int_{R^n} \phi(z,T)p(z,T;x,t)\,dz - \phi(x,t) = \int_t^T \int_{R^n} (\phi_s + \mathcal{L}\phi)(z,s)p(z,s;x,t)\,dzds. \tag{20}$$

Using (18) and doing the obvious integration by parts in time, the right hand side becomes

$$\int_t^T \int_{R^n} -\phi p_s + \phi\mathcal{L}^* p\,dzds + \int_{R^n} \phi(z,s)p(z,s;x,t)\,dz\Big|_{s=t}^{s=T}. \tag{21}$$

This is true for *all* ϕ. Since the left hand side of (20) involves only the initial and final times (t and T) we conclude that

$$-p_s + \mathcal{L}^* p = 0.$$

Therefore (20)-(21) reduce to

$$\int_{R^n} \phi(z,t)p(z,t;x,t)\,dz = \phi(x,t)$$

for all ϕ, which is what we mean by the initial condition "$p = \delta_x$ when $s = t$". Done!

The argument is simple; but maybe it's hard to encompass. To recapitulate its essence, let's give a new proof (using the forward equation) of the fact (known via Ito calculus) that

$$u \text{ solves the backward equation} \Longrightarrow \frac{d}{ds}E\left[u(y(s),s)\right] = 0.$$

In fact: if $\rho(z,s)$ is the probability distribution of the state at time s,

$$\begin{aligned}
\frac{d}{ds}E\left[u(y(s),s)\right] &= \frac{d}{ds}\int u(z,s)\rho(z,s)\,dz \\
&= \int u_s\rho + u\rho_s\,dz \\
&= \int u_s\rho + u\mathcal{L}^*\rho\,dz \\
&= \int u_s\rho + (\mathcal{L}u)\rho\,dz \\
&= 0
\end{aligned}$$

using in the last step our hypothesis that u solves the backward equation.

****

Boundary value problems. The preceding discussion concerned the backward and forward Kolmogorov equations in all space. When working in a bounded domain, the boundary conditions for the forward Kolmogorov equation depend on what the random walk does when it reaches the boundary.

We discuss here just the case of most interest for financial applications: the *absorbing* boundary condition, i.e. a random walk that we track only till it hits the boundary for the first time. (After that time we think of the random walker as disappearing, i.e. being "absorbed" by the boundary.) The corresponding boundary condition for the forward Kolmogorov equation is that the probability density vanish there (since it represents the density of not-yet-absorbed walkers).

Let's explain briefly why this choice is right. Consider the backward Kolmogorov equation in a bounded domain, with boundary condition $u = 0$:

$$\begin{gathered} u_t + \mathcal{L}u = 0 \text{ for } x \in D, t < T \\ u(x,T) = \phi(x) \text{ at } t = T \\ u(x,t) = 0 \text{ for } x \in \partial D. \end{gathered}$$

We know that

$$u(x,t) = E_{y(t)=x}\left[\Phi(y(\tau),\tau)\right]$$

where $\tau = \tau(x)$ is the exit time from D (or T, if the path doesn't exit by time T) and

$$\Phi = 0 \text{ for } x \in \partial D;\ \Phi = \phi \text{ at the final time } T.$$

This formula for u can be written as

$$u(x,t) = \int_{R^n} \phi(z) q(z,T;x,t)\, dz$$

where

$$\begin{aligned} &q(z,s;x,t) = \text{probability that the diffusion arrives at } z \text{ at time } s, \\ &\text{starting from } x \text{ at time } t, \text{ without hitting } \partial D \text{ first.} \end{aligned}$$

Our assertion is that $q(z,s;x,t)$ *solves the forward Kolmogorov equation for* $z \in D$ *and* $s > t$*, with boundary condition* $q = 0$ *for* $z \in \partial D$*, and initial condition* $q = \delta_x$. The justification is very much like the argument given above for R^n.

One thing changes significantly when we work in a bounded domain: $\int_D q(z,s;x,t)\, dz < 1$. The reason is that q gives the probability of arriving at z at time s *without hitting the boundary first.* Thus

$$1 - \int_D q(z,s;x,t)\, dz = \text{prob of hitting } \partial D \text{ by time } s, \text{ starting from } x \text{ at time } t.$$

Evidently $\int q(z,s;x,t)\,dz$ is decreasing in time. Let's check this for Brownian motion, for which $q_s - \frac{1}{2}\Delta q = 0$. We have

$$\begin{aligned}\frac{d}{ds}\int_D q(z,s;x,t)\,dz &= \int_D q_s\,dz \\ &= \tfrac{1}{2}\int_D \Delta q\,dz \\ &= \tfrac{1}{2}\int_{\partial D} \frac{\partial q}{\partial n} \\ &\leq 0.\end{aligned}$$

The inequality in the last step is a consequence of the maximum principle (to be discussed in a future lecture): since $q = 0$ at ∂D and $q \geq 0$ in D we have $\partial q/\partial n \leq 0$ at ∂D, where n is the outward unit normal. (In fact $\partial q/\partial n < 0$; this is a "strong version" of the maximum principle.)

Application to the exit time distribution. We used the backward Kolmogorov equation to express the probability that a diffusion reaches a certain threshold before time T (see (8)). The forward Kolmogorov equation gives a very convenient alternative expression for the same quantity. Indeed, if ρ solves the forward Kolmogorov equation in the domain D of interest, with $\rho = 0$ at the boundary and $\rho = \delta_x$ at time 0, then $\int_D \rho(x,T)\,dx$ gives the probability of surviving till time T. So $1 - \int_D \rho(x,T)\,dx$ is the probability of hitting the boundary by time T, given that you started at x at time 0. When D is a half-space, this is an alternative expression for $u(x,0;T)$ defined by (8).

PDE for Finance Notes, Spring 2003 – Section 2.
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course PDE for Finance, G63.2706.

Solution formulas for the linear heat equation. Applications to barrier options. Section 1 was relatively abstract – we listed many PDE's but solved just a few of them. This section has the opposite character: we discuss explicit solution formulas for the linear heat equation – both the initial value problem in all space and the initial-boundary-value problem in a halfspace. This is no mere academic exercise, because the (constant-volatility, constant-interest-rate) Black-Scholes PDE can be reduced to the linear heat equation. As a result, our analysis provides all the essential ingredients for valuing barrier options. The PDE material here is fairly standard – most of it can be found in John or Evans for example. For the financial topics (reduction of Black-Scholes to the linear heat equations; valuation of barrier options) see Wilmott-Howison-Dewynne or Wilmott. For much more detail on barrier options see Taleb's book *Dynamic Hedging*.

****

The heat equation and the Black-Scholes PDE. We've seen that linear parabolic equations arise as *backward* Kolmogorov equations, determining the expected values of various payoffs. They also arise as *forward* Kolmogorov equations, determinining the probability distribution of the diffusing state. The simplest special cases are the backward and forward linear heat equations $u_t + \frac{1}{2}\sigma^2 \Delta u = 0$ and $p_s - \frac{1}{2}\sigma^2 \Delta p = 0$, which are the backward and forward Kolmogorov equations for $dy = \sigma dw$, i.e. for Brownian motion scaled by a factor of σ. From a PDE viewpoint the two equations are equivalent, since $v(t,x) = u(T-t,x)$ solves $v_t - \frac{1}{2}\sigma^2 \Delta v = 0$, and final-time data for u at $t = T$ determines initial-time data for v at $t = 0$.

This basic example has direct financial relevance, because the Black-Scholes PDE can be reduced to it by a simple change of variables. Indeed, the Black-Scholes PDE is

$$V_t + rsV_s + \tfrac{1}{2}\sigma^2 s^2 V_{ss} - rV = 0.$$

It is to be solved for $t < T$, with specified final-time data $V(s,T) = \Phi(s)$. (Don't be confused: in the last paragraph s was time, but here it is the "spatial variable" of the Black-Scholes PDE, i.e. the stock price.) We claim this is simply the standard heat equation $u_t = u_{xx}$ written in special variables. To see this, consider the preliminary change of variables $(s,t) \to (x,\tau)$ defined by

$$s = e^x, \quad \tau = \tfrac{1}{2}\sigma^2 (T-t),$$

and let $v(x,\tau) = V(s,t)$. An elementary calculation shows that the Black-Scholes equation becomes

$$v_\tau - v_{xx} + (1-k)v_x + kv = 0$$

with $k = r/(\frac{1}{2}\sigma^2)$. We've done the main part of the job: reduction to a constant-coefficient equation. For the rest, consider $u(x,t)$ defined by

$$v = e^{\alpha x + \beta \tau} u(x, \tau)$$

where α and β are constants. The equation for v becomes an equation for u, namely

$$(\beta u + u_\tau) - (\alpha^2 u + 2\alpha u_x + u_{xx}) + (1-k)(\alpha u + u_x) + ku = 0.$$

To get an equation without u or u_x we should set

$$\beta - \alpha^2 + (1-k)\alpha + k = 0, \quad -2\alpha + (1-k) = 0.$$

These equations are solved by

$$\alpha = \frac{1-k}{2}, \quad \beta = -\frac{(k+1)^2}{4}.$$

Thus,

$$u = e^{\frac{1}{2}(k-1)x + \frac{1}{4}(k+1)^2 \tau} v(x, \tau)$$

solves the linear heat equation $u_\tau = u_{xx}$ with initial condition $u(x,0) = e^{\frac{1}{2}(k-1)x}\Phi(e^x)$.

****

The initial-value problem. Consider the equation

$$f_t = \alpha \Delta f \quad \text{for } x \in R^n, t > 0 \tag{1}$$

with specified data $f(x,0) = f_0(x)$. Its solution is:

$$f(x,t) = (4\pi\alpha t)^{-n/2} \int e^{-|x-y|^2/4\alpha t} f_0(y)\, dy. \tag{2}$$

Why? Because (1) can be viewed as the forward Kolmogorov equation for $dy = \sigma dw$ when $\alpha = \frac{1}{2}\sigma^2$. Let's take $\alpha = \frac{1}{2}$ for simplicity, so $\sigma = 1$. The probability of a Brownian particle being at x at time t, given that it started at y at time 0, is $(2\pi t)^{-n/2} e^{-|x-y|^2/2t}$. If the initial probability distribution is $f_0(y)$ then the probability of being at x at time t is $(2\pi t)^{-n/2} \int e^{-|x-y|^2/2t} f_0(y)\, dy$, exactly our formula specialized to $\alpha = 1/2$.

We have, in effect, used our knowledge about Brownian motion to write down a specific solution of (1). To know it is the *only* solution, we must prove a uniqueness theorem. We'll address this in the next Section.

For what class of initial data f_0 is the solution formula (2) applicable? From our probabilistic argument it might appear that f_0 has to be a probability density (positive, with integral 1). In fact however there is no such limitation. It isn't even necessary that f_0 be integrable. To prove that the solution formula works much more generally, one must verify (a) that it solves the equation, and (b) that it has the desired initial value. The proof of (a) is

easy, by differentiating under the integral. The proof of (b) is more subtle. Most textbooks present it assuming f_0 is continuous, but the standard argument also works somewhat more generally, e.g. if f_0 is piecewise continuous.

There is however one requirement: the solution formula must make sense. This requires a modest restriction on the growth of f_0 at ∞, to make the integral on the right hand side of (2) converge. For example, if $f_0(x) = \exp(c|x|^2)$ with $c > 0$ then the integral diverges for $t > (4\alpha c)^{-1}$. The natural growth condition is thus

$$|f_0(x)| \leq Me^{c|x|^2} \tag{3}$$

as $|x| \to \infty$.

Solutions growing at spatial infinity are uncommon in physics but common in finance, where the heat equation arises by a logarithmic change of variables from the Black-Scholes PDE, as shown above. The payoff of a call is linear in the stock price s as $s \to \infty$. This leads under the change of variable $x = \log s$ to a choice of f_0 which behaves like e^{cx} as $x \to \infty$. Of course this lies well within what is permitted by (3).

What option payoffs are permitted by (3)? Since $x = \log s$, the payoff must grow no faster as $s \to \infty$ than $M \exp\left(c(\log s)^2\right)$. This condition is pretty generous: it permits payoffs growing like any power of s as $s \to \infty$, though it excludes growth like e^s.

Discontinuous initial conditions are relatively uncommon in physics, but common in finance. A digital option, for example, pays a specified value if the stock price at maturity is greater than a specified value, and nothing otherwise. This corresponds to a discontinuous choice of f_0. Notice that even if f_0 is discountinuous, the solution $f(x,t)$ is smooth for $t > 0$. This can be seen by differentiating under the integral in the solution formula.

We have reduced the Black-Scholes equation to the heat equation, and we have given an explicit solution formula for the heat equation. Unraveling all this gives an explicit solution for the Black-Scholes equation. Of course in the end this is just the familiar formula, giving the value of an option as the discounted risk-neutral expected payoff.

It may seem we haven't gained much. And indeed, for vanilla options we haven't. The PDE viewpoint is much more useful, however, for considering barrier options.

****

The initial-boundary value problem for a halfspace. Now consider

$$u_t = u_{xx} \quad \text{for } t > 0 \text{ and } x > x_0, \text{ with } u = g \text{ at } t = 0 \text{ and } u = \phi \text{ at } x = x_0. \tag{4}$$

Since this is a boundary-value problem, we must specify data both at the initial time $t = 0$ and at the spatial boundary $x = 0$. We arrived at this type of problem (with t replaced by $T - t$) in our discussion of the backward Kolmogorov equation when we considered a payoff defined at an exit time. The associated option-pricing problems involve barriers. If the option becomes worthless at when the stock price crosses the barrier then $\phi = 0$ (this

is a knock-out option). If the option turns into a different instrument when the stock price crosses the barrier then ϕ is the value of that instrument. (When $\phi = 0$, (4) can also be viewed as a forward Kolmogorov equation, for $\sqrt{2}$ times Brownian motion with an absorbing boundary condition at $x = 0$.)

Please notice that if u is to be continuous (up to the boundary, and up to the initial time) then the boundary and initial data must be *compatible*, in other words they must satisfy $g(0) = \phi(0)$. When the data are incompatible, the solution is bound to be singular near $x = t = 0$ even if g and ϕ are individually smooth. The incompatible case is directly relevant to finance. For example, the pricing problem for a down-and-out call has incompatible data if the strike price is below the barrier. Such options become difficult to hedge if, near the time of maturity, the stock price wanders near the barrier.

The pricing problem can be decomposed, by linearity, into two separate problems: $u = v+w$ where v solves

$$v_t = v_{xx} \quad \text{for } t > 0 \text{ and } x > x_0, \text{ with } v = g \text{ at } t = 0 \text{ and } v = 0 \text{ at } x = x_0 \tag{5}$$

(in other words: v solves the same PDE with the same initial data but boundary data 0) and w solves

$$w_t = w_{xx} \quad \text{for } t > 0 \text{ and } x > x_0, \text{ with } w = 0 \text{ at } t = 0 \text{ and } w = \phi \text{ at } x = x_0 \tag{6}$$

(in other words: w solves the same PDE with the same boundary data but initial condition 0). Careful though: in making this decomposition we run the risk of replacing a compatible problem with two incompatible ones. As we'll see below, the solution formula for (5) is very robust even in the incompatible case. The formula for (6) is however much less robust. So in practice our decomposition is most useful when $\phi(0) = 0$, so that the w-problem has compatible data. This restriction represents no real loss of generality: if $\phi(0) = c \neq 0$ then our decomposition can be used to represent $u - c$, which solves the same PDE with data $\phi - c$ and $g - c$.

The half-space problem with boundary condition 0. We concentrate for the moment on v. To obtain its solution formula, we consider the whole-space problem with the *odd reflection of* g as initial data. Remembering that $x_0 = 0$, this odd reflection is defined by

$$\tilde{g}(x) = \begin{cases} g(x) & \text{if } x > 0 \\ -g(-x) & \text{if } x < 0 \end{cases}$$

(see Figure 1). Notice that the odd reflection is continuous at 0 if $g(0) = 0$; otherwise it is discontinuous, taking values $\pm g(0)$ just to the right and left of 0.

Let $\tilde{v}(x, t)$ solve the whole-space initial-value problem with initial condition $\tilde{g}$. We claim

- $\tilde{v}$ is a smooth function of x and t for $t > 0$ (even if $g(0) \neq 0$);
- $\tilde{v}(x, t)$ is an odd function of x for all t, i.e. $\tilde{v}(x, t) = -\tilde{v}(-x, t)$.

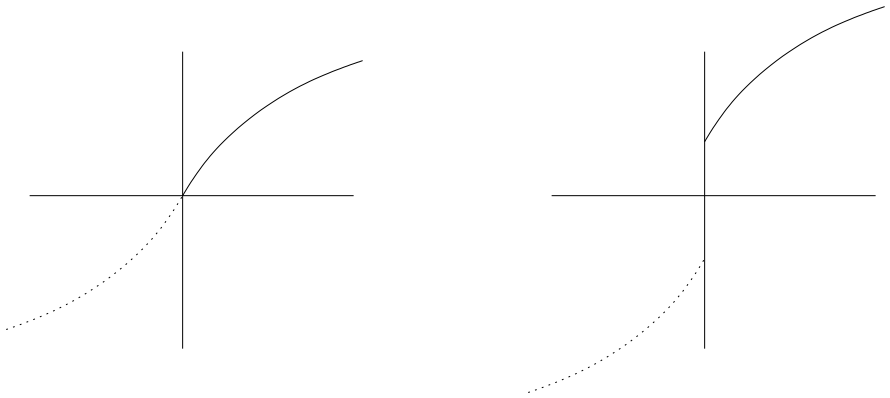

Figure 1: *Odd reflection. Note that the odd reflection is discontinuous at* 0 *if the original function doesn't vanish there.*

The first bullet follows from the smoothing property of the heat equation. The second bullet follows from the uniqueness of solutions to the heat equation, since $\tilde{v}(x,t)$ and $-\tilde{v}(-x,t)$ both solve the heat equation with the *same* initial data $\tilde{g}$. (Please accept this uniqueness result for now; we'll prove it in the next Section.)

We're essentially done. The oddness of $\tilde{v}$ gives $\tilde{v}(0,t) = -\tilde{v}(0,t)$, so $\tilde{v}(0,t) = 0$ for all $t > 0$. Thus

$$v(x,t) = \tilde{v}(x,t), \quad \text{restricted to } x > 0$$

is the desired solution to (5). Of course it can be expressed using (2): a formula encapsulating our solution procedure is

$$\begin{aligned} v(x,t) &= \int_0^\infty k(x-y,t)g(y)\,dy + \int_{-\infty}^0 k(x-y,t)(-g(-y))\,dy \\ &= \int_0^\infty [k(x-y,t) - k(x+y,t)]g(y)\,dy \end{aligned}$$

where $k(x,t)$ is the fundamental solution of the heat equation, given by

$$k(z,t) = \frac{1}{\sqrt{4\pi t}} e^{-z^2/4t}. \tag{7}$$

In other words

$$v(x,t) = \int_0^\infty G(x,y,t)g(y)\,dy$$

with

$$G(x,y,s) = k(x-y,t) - k(x+y,t). \tag{8}$$

Fortunately $G(x,y,s) = G(y,x,s)$ so we don't have to try to remember which variable (x or y) we put first. The function G is called the "Green's function" of the half-space problem. Based on our discussion of the forward Kolmogorov equation, we recognize $G(x,y,t)$ as giving the probability that a Brownian particle starting from y at time 0 reaches position

x at time t without first reaching the origin. (I'm being sloppy: the relevant random walk is not Brownian motion but $\sqrt{2}$ times Brownian motion.)

The half-space problem with initial condition 0. It remains to consider w, defined by (6). It solves the heat equation on the half-space, with initial value 0 and boundary value $\phi(t)$.

The solution w is given by

$$w(x,t) = \int_0^t \frac{\partial G}{\partial y}(x,0,t-s)\phi(s)\,ds \tag{9}$$

where $G(x,y,t)$ is the Green's function of the half-space problem given by (8). Using the formula derived earlier for G, this amounts to

$$w(x,t) = \int_0^t \frac{x}{(t-s)\sqrt{4\pi(t-s)}} e^{-x^2/4(t-s)}\phi(s)\,ds$$

Notice that the integral is quite singular near $x = t = 0$. That's why the w-problem is best applied to compatible data ($\phi(0) = 0$).

The justification of (9) is not difficult, but it's rather different from what we've done before. To represent the value of w at location x_0 and time t_0, consider the function ψ which solves the heat equation *backward in time* from time t_0, with final-time data concentrated at x_0 at time t_0. We mean ψ to be defined only for $x > 0$, with $\psi = 0$ at the spatial boundary $x = 0$. In formulas, our definition is

$$\psi_\tau + \psi_{yy} = 0 \quad \text{for } \tau < t_0 \text{ and } y > 0, \text{ with } \psi = \delta_{x_0} \text{ at } \tau = t_0 \text{ and } \psi = 0 \text{ at } y = 0.$$

A formula for ψ is readily available, since the change of variable $s = t_0 - \tau$ transforms the problem solved by ψ one considered earlier for v:

$$\psi(y,\tau) = G(x_0, y, t_0 - \tau). \tag{10}$$

What's behind our strange-looking choice or ψ? Two things. First, the choice of final-time data gives

$$w(x_0,t_0) = \int \psi(y,t_0)w(y,t_0)\,dy.$$

(The meaning of the statement "$\psi = \delta_{x_0}$ at time t_0" is precisely this.) Second, if w solves the heat equation forward in time and ψ solves it backward in time then

$$\begin{aligned}
\frac{d}{ds}\int_0^\infty \psi(y,s)w(y,s)\,dy &= \int_0^\infty \psi_s w + \psi w_s\,dy \\
&= \int_0^\infty -\psi_{yy}w + \psi w_{yy}\,dy \\
&= \int_0^\infty -(\psi_y w)_y + (\psi w_y)_y\,dy \\
&= (-\psi_y w + \psi w_y)|_0^\infty \ . \qquad (11)
\end{aligned}$$

(I've used here that the heat equation backward-in-time is the formal adjoint of the heat equation forward-in-time; you saw this before in the discussion of the forward Kolmogorov equation, which is always the formal adjoint of the backward Kolmogorov equation.) Because of our special choice of ψ the last formula simplifies: ψ and ψ_y decay rapidly enough at ∞ to kill the "boundary term at infinity," and the fact that $\psi = 0$ at $y = 0$ kills one of the two boundary terms at 0. Since $w(0, s) = \phi(s)$ what remains is

$$\frac{d}{ds}\int_0^\infty \psi(y,s)w(y,s)\,dy = \psi_y(0,s)\phi(s).$$

We're essentially done. Substitution of (10) in the above gives, after integration in s,

$$\int_0^\infty \psi(y,t_0)w(y,t_0)\,dy - \int_0^\infty \psi(y,0)w(y,0) = \int_0^{t_0} G_y(x_0,0,t_0-s)\phi(s)\,ds.$$

The first term on the left is just $w(x_0, t_0)$, by our choice of ψ, and the second term on the left vanishes since $w = 0$ at time 0, yielding precisely the desired solution formula (9).

****

Barrier options. The preceding analysis provides the main ideas needed for pricing any (European) barrier option. By definition, a barrier option is like a vanilla option except that it acquires or loses its value if the stock price goes above or below a specified barrier X:

An **up-and-out** option loses its value if the stock price crosses X from below prior to maturity.

A **down-and-out** option loses its value if the stock price crosses X from above prior to maturity.

An **up-and-in** option pays only if the stock price crosses X from below prior to maturity.

A **down-and-in** option pays off only if the stock price crosses X from above prior to maturity.

For example, a down-and-out call with strike price K, maturity T, and barrier X, pays $(s - K)_+$ if the stock price stays above X, and nothing if the stock price dips below X prior to maturity. The corresponding down-and-in call has no value until the stock price crosses the barrier X from above; if that ever happens then it behaves like a standard call thereafter.

To connect our PDE's with their financial applications, let's discuss in detail the case of a down-and-out barrier call with $X < K$. This case has compatible data (i.e. the payoff vanishes at $s = X$), making it easier than the case $X > K$. The value of this barrier call is

$$V(s,t) = C(s,t) - \left(\frac{s}{X}\right)^{(1-k)} C(X^2/s,t)$$

where $k = r/(\frac{1}{2}\sigma^2)$ and $C(s, t)$ is the value of the ordinary European call with strike K and maturity T. One can, of course, check by mere arithmetic that this formula solves the PDE

and the boundary conditions. But we prefer to show how this formula emerges from our understanding of the initial-boundary-value problem for the heat equation.

Recall that under the change of variables

$$s = e^x, \quad \tau = \tfrac{1}{2}\sigma^2(T-t), \quad V(s,t) = e^{\alpha x + \beta\tau} u(x,\tau)$$

with $\alpha = (1-k)/2$, $\beta = -(k+1)^2/4$, the Black-Scholes PDE becomes

$$u_\tau = u_{xx}.$$

Restricting $s > X$ is the same as restricting $x > \log X$, so the linear heat equation is to be solved for $x > \log X$, with $u = 0$ at $x = \log X$. Its initial value $u_0(x) = u(0,x)$, is obtained from the payoff of the call by change of variables: $u_0(x) = e^{-\alpha x}(e^x - K)_+$.

Since the boundary condition is now at $x = \log X$, we can impose the condition $u = 0$ at $x = \log X$ via odd reflection about $\log X$. That is, we look for a solution of $u_t = u_{xx}$ for all x, with the property that

$$u(x',t) = -u(x,t) \quad \text{when } x' \text{ is the reflection of } x \text{ about } \log X.$$

Such a solution must satisfy $u(\log X, t) = 0$, since the condition of odd symmetry gives $u(\log X, t) = -u(\log X, t)$.

Let's be more explicit about the condition of odd symmetry. Two points $x' < \log X < x$ are related by reflection about $\log X$ if $x - \log X = \log X - x'$, i.e. if $x' = 2\log X - x$. So a function $u(x,t)$ has odd symmetry about $\log X$ if it satisfies

$$u(2\log X - x, t) = -u(x,t) \quad \text{for all } x.$$

OK, the plan is clear: First (a) extend the initial data by odd symmetry about $x = \log X$; then (b) solve the linear heat equation for $t > 0$ and all x. Carrying out step (a): the desired initial data is $u_0(x) = e^{-\alpha x}(e^x - K)_+$ for $x > \log X$; moreover our assumption that $X < K$ assures that $u_0(x) = 0$ for $x \le \log X$. So the extended initial data is

$$f_0(x) = u_0(x) - u_0(2\log X - x).$$

Carrying out step (b) is of course trivial: the solution $f(x,t)$ is given by the convolution of f_0 with the fundamental solution, i.e. by (2).

To make the value of the option explicit without an orgy of substitution, we use the fact that our PDE's are linear. So the value $V(s,t)$ of our barrier option is the difference of two terms. The first corresponds under the change of variables to

$$f_1(x,t) = \text{solution of the whole-space heat equation with initial data } u_0(x),$$

i.e. the first term is the value $C(s,t)$ of the ordinary European call. (We note for later use the relation $f_1(x,\tau) = C(e^x,t)e^{-\alpha x - \beta\tau}$.) The second term corresponds under the change of variables to

$$\begin{aligned} f_2(x,t) &= \text{solution of the whole-space heat equation with initial data } u_0(2\log X - x) \\ &= f_1(2\log X - x), \end{aligned}$$

so it is

$$\begin{aligned} e^{\alpha x+\beta\tau} f_2(x,\tau) &= e^{\alpha x+\beta\tau} f_1(2\log X - x, t) \\ &= e^{\alpha x+\beta\tau} e^{-(\alpha[2\log X - x]+\beta\tau)} C(e^{2\log X - x}, t) \\ &= X^{-2\alpha} s^{2\alpha} C(X^2/s, t) = (s/X)^{1-k} C(X^2/s, t). \end{aligned}$$

The solution formula asserted above is precisely the difference of these two terms.

We close with a comment about the associated down-and-in call. (Remember: it has no value until the stock price crosses the barrier X from above; if this ever happens then it behaves like a standard call thereafter.) At first we seem to have to solve the Black-Scholes PDE with value of a vanilla call as its boundary condition. But actually there's no need for such hard work. Indeed, it's obvious that

$$\text{down-and-out call} + \text{down-and-in call} = \text{standard call}$$

since the two portfolios are equivalent. So the value of a down-and-in call is just the difference between the value of the standard call and the down-and-out call.

PDE for Finance Notes, Spring 2003 – Section 3.
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course PDE for Finance, G63.2706.

More about linear PDE's: the heat equation on an interval; uniqueness via the maximum principle. Section 2 gave the solution formula for the heat equation in a half-space – the essential tool for pricing barrier options. The natural next topic is to consider the heat equation in an interval – the essential tool for pricing double-barrier options. Then we turn to the question of uniqueness, something we've been assuming up to now. The proof uses the maximum principle – a basic and powerful tool for obtaining qualitative information about solutions of parabolic equations.

Our treatment of the heat equation in an interval is by separation of variables; this can be found in most basic PDE books. The recent article *Exotic Spectra* by Vadim Linetsky (Risk Magazine, April 2002) reviews this topic and explores its financial applications; a xerox copy is on reserve in the CIMS library green box for this class. Uniqueness via the maximum principle is also quite standard; students of finance may find it convenient to read this in Steele's book (the relevant section is in the middle of the book but is quite self-contained).

****

The heat equation in an interval. A double-barrier option has both an upper and lower barrier. Its value satisfies the Black-Scholes PDE with appropriate boundary data at the barriers. If the underlying is lognormal then this problem can be reduced, as shown in Section 2, to solving the linear heat equation on an interval with appropriate initial and boundary data. For simplicity we focus on the case when the interval is $0 < x < 1$. Thus we wish to solve:

$$u_t = u_{xx} \quad \text{for } t > 0 \text{ and } 0 < x < 1$$

with initial data $u = g$ at $t = 0$, and boundary data $u = \phi_0$ at $x = 0$, $u = \phi_1$ at $x = 1$.

Consider first the case when the boundary condition is zero, i.e. $\phi_0 = \phi_1 = 0$. We will use the following basic fact: any function $f(x)$ defined for $0 < x < 1$ and vanishing at the endpoints $x = 0, 1$ can be expanded as a Fourier sine series:

$$f(x) = \sum_{k=1}^{\infty} a_n \sin(n\pi x); \tag{1}$$

moreover the coefficients a_n are determined by f via

$$a_n = 2 \int_0^1 f(x) \sin(n\pi x)\, dx. \tag{2}$$

If you accept (1), the justification of (2) is easy. It follows from the fact that the functions $\{\sin(n\pi x)\}_{n=1}^{\infty}$ are orthogonal in the inner product $(f, g) = \int_0^1 f(x)g(x)\, dx$, and each has norm $1/\sqrt{2}$.

Now consider the solution of $u_t = u_{xx}$ with boundary data 0 and initial data $g(x)$. Applying (1) at each time, we have

$$u(x,t) = \sum_{k=1}^{\infty} a_n(t) \sin(n\pi x)$$

for some functions $a_n(t)$. To get initial data right, $a_n(0)$ must be the Fourier sine-series coefficients of g. To satisfy the PDE we must have

$$da_n/dt = -n^2\pi^2 a_n,$$

whence $a_n(t) = a_n(0)e^{-n^2\pi^2 t}$. Thus the solution is

$$u(x,t) = \sum_{n=1}^{\infty} g_n e^{-n^2\pi^2 t} \sin(n\pi x) \tag{3}$$

where

$$g_n = 2\int_0^1 g(x) \sin(n\pi x)\, dx. \tag{4}$$

Notice an important feature of this solution formula: the nth term decays like $e^{-n^2\pi^2 t}$. Thus terms with higher n (corresponding to higher-order modes of the initial data) decay faster. If you only need the solution to a certain accuracy at a fixed time T then you only need consider a limited number of modes – and the number of modes decreases as T increases! Thus valuing an option by this method gets *easier* as the maturity gets larger. This is opposite to the behavior of a time-stepping numerical scheme (such as finite-difference approximation).

It is natural to seek a solution formula in the form

$$u(x,t) = \int_0^1 G(x,y,t) g(y)\, dy \tag{5}$$

since $G(x,y,t)$ is then the probability that a random walker starting at y arrives at x at time t without first hitting the boundary. (Our random walker executes, as usual, $\sqrt{2}$ times Brownian motion.) This is just a matter of manipulation: combining (3) and (4) gives (5) with

$$G(x,y,t) = 2\sum_{n=1}^{\infty} e^{-n^2\pi^2 t} \sin(n\pi x) \sin(n\pi y). \tag{6}$$

One might think the solution formula just presented was limited to initial conditions with $g = 0$ at the boundary, i.e. "consistent" initial data. However this is not the case: our formula is correct even for inconsistent data. The reason is that *any* function on $[0,1]$ has a Fourier sine series – which equals the function almost everywhere (though perhaps not at the endpoints).

Let's turn now to the boundary value problem, i.e. the solution of $u_t - u_{xx} = 0$ with initial condition 0 and boundary data $u = \phi_0$ at $x = 0$, $u = \phi_1$ at $x = 1$. We did all the work already in Section 2: the argument given there shows that the solution is

$$u(x,t) = \int_0^t \frac{\partial G}{\partial y}(x,0,t-s)\phi_0(s)\, ds - \int_0^t \frac{\partial G}{\partial y}(x,1,t-s)\phi_1(s)\, ds$$

where G is given by (6). In particular, $\frac{\partial G}{\partial y}(x,0,t)$ is the probability that the random walker, starting from x, hits the boundary first at 0 and arrives there at time t. Similarly, $-\frac{\partial G}{\partial y}(x,1,t)$ is the probability that the random walker, starting from x, hits the boundary first at 1 and arrives there at time t.

The "separation of variables" solution method just presented is not limited to the constant-coefficient heat equation. It can be applied to any equation of the form

$$u_t - \tfrac{1}{2}a^2(x)u_{xx} - b(x)u_x + r(x)u = 0$$

on an interval, provided a is strictly positive on this interval. Thus it can be used to price double-barrier options on a fairly general class of underlyings (provided that the volatility is independent of time). The key is to recognize that Fourier sine series (and the associated decay rates $n^2\pi^2$) must be replaced by the eigenfunctions (and eigenvalues) of

$$-\tfrac{1}{2}a^2(x)v_{xx} - b(x)v_x + r(x)v = \lambda v$$

with $v = 0$ at the boundary. These eigenfunctions are orthogonal not in the L^2 inner product, but rather in a different one determined by the coefficients a and b. See Linetsky's article *Exotic Spectra* for details and applications.

****

Why does time have a preferred direction? This is a convenient time to ask: why is it so important that we solve $u_t - u_{xx} = 0$ forward rather than backward in time? The answer is that solving in the correct time direction is stable (moreover the solution is instantly smooth, regardless how rough the initial data may be) whereas solving in the wrong time direction is extremely unstable. Indeed, consider the separation of variables technique presented above. We can, in principle, apply it in the wrong time direction, since the ODE $da_n/dt = -n^2\pi^2 a_n$ is reversible. Thus the solution of $u_t - u_{xx} = 0$ for $t < T$ with final-time data $u = g$ at $t = T$ is

$$u(x,t) = \sum_{n=1}^{\infty} g_n e^{n^2\pi^2(T-t)} \sin(n\pi x)$$

But watch out! The nth mode grows extremely fast: as a constant times $e^{n^2\pi^2(T-t)}$ as $T-t$ increases. The tiniest high-frequency ripple in g will become huge as you proceed backward in time.

The whole-space problem behaves similarly: solving backward in time is extremely unstable. This is, in fact, a consequence of the fact that solving forward in time is so stable and smoothing.

Might there still be some interest in solving the heat equation the "wrong way" in time? Sure. This is the simplest example of "deblurring," a typical task in image enhancement.

Consider a photograph taken with an out-of-focus camera. Its image is (roughly speaking) the convolution of the true image with a Gaussian of known variance. Finding the original image amounts to backsolving the heat equation with the blurry photo as final-time data. Backsolving the heat equation is a typical example of an *ill-posed problem* – one whose answer depends in an unreasonably sensitive way on the data.

****

The maximum principle and uniqueness. Are our solution formulas the *only* solution of the heat equation with the specified initial and/or boundary conditions? By linearity, this amounts to asking whether $u = 0$ is the only solution with data 0? The answer is yes. We shall prove this using the maximum principle.

The maximum principle. This is an elementary yet far-reaching fact about solutions of linear parabolic equations. Here is the simplest version:

> Let D be a bounded domain. Suppose $f_t - \Delta f \leq 0$ for all $x \in D$ and $0 < t < T$. Then the maximum of f in the closed cylinder $\bar{D} \times [0, T]$ is achieved either at the "initial boundary" $t = 0$ or at the "spatial boundary" $x \in \partial D$.

Notice the asymmetry between the initial boundary $t = 0$ (where f can easily achieve its maximum) and the final boundary $t = T$ (where f does not achieve its maximum – except in the trivial case when f is constant). This asymmetry reflects once again the fact that time has a "preferred direction" when solving a parabolic PDE.

If $f_t - \Delta f$ were *strictly* negative, the principle just enunciated would be a calculus exercise. Indeed, f must achieve its maximum *somewhere* in the cylinder or on its boundary (we use here that D is bounded). Our task is to show this doesn't occur in the interior or at the "final boundary" $t = T$. At an interior maximum all first derivatives would vanish and $\partial^2 f/\partial x_i^2 \leq 0$ for each i; but then $f_t - \Delta f \geq 0$, contradicting the hypothesis that $f_t - \Delta f < 0$. At a final-time maximum (in the interior of D) all first derivatives in x would still vanish, and we would still have $\partial^2 f/\partial x_i^2 \leq 0$; we would only know $f_t \geq 0$, but this would still give $f_t - \Delta f \geq 0$, again contradicting the hypothesis of strict negativity.

If all we know is $f_t - \Delta f \leq 0$ then the preceding argument doesn't quite apply. But the fix is simple: we can apply it to $f_\epsilon(x, t) = f(x, t) - \epsilon t$ for any $\epsilon > 0$. As $\epsilon \to 0$ this gives the desired result.

There is an analogous minimum principle:

> Let D be a bounded domain. Suppose $f_t - \Delta f \geq 0$ for all $x \in D$ and $0 < t < T$. Then the minimum of f in the closed cylinder $\bar{D} \times [0, T]$ is achieved either at the "initial boundary" $t = 0$ or at the "spatial boundary" $x \in \partial D$.

It follows from the maximum principle applied to $-f$. In particular, if $f_t - \Delta f = 0$ in the cylinder then f assumes its maximum and minimum values at the spatial boundary or the initial boundary.

The proof just given for the maximum principle generalizes straightforwardly to more general linear parabolic equations, provided there is no zeroth-order term. For example: if $f_t - \sum_{i,j} \alpha_{ij}(x,t)\nabla^2_{ij} f - \sum_i \beta_i(x,t)\nabla_i f \leq 0$ then f achieves its maximum in $\bar{D} \times [0,T]$ at the initial or spatial boundary.

Uniqueness. Uniqueness of the initial-boundary-value problem in a bounded domain follows immediately from the maximum principle. Since the equation is linear, if there were two solutions with the same data then their difference would be a solution with data 0. So the main point is this:

> Suppose $f_t = \Delta f$ for $t > 0$ and $x \in D$. Assume moreover f has initial data 0 ($f(x,0) = 0$ for $x \in D$) and boundary data 0 ($f(x,t) = 0$ for $x \in \partial D$). Then $f(x,t) = 0$ for all $x \in D$, $t > 0$.

Indeed: the maximum and minimum of f are 0, by the maximum (and minimum) principles. So f is identically 0 in the cylinder.

To show uniqueness for the initial-value problem in all space one must work a bit harder. The problem is that we no longer have a spatial boundary – and we mean to allow solutions that grow at ∞, so the maximum of $f(x,t)$ over all $0 < t < T$ and $x \in R^n$ might well occur as $x \to \infty$. We already know, however, that it's natural to assume a growth condition of the form $|f(x,t) \leq Me^{c|x|^2}$ for some M and c. Subtracting two possible solutions, our task is thus to show the following:

> Suppose $f_t = \Delta f$ for $t > 0$ and $x \in R^n$. Assume moreover f has initial data 0 and $|f(x,t) \leq Me^{c|x|^2}$ for some M and c. Then $f(x,t) = 0$ for all $x \in R^n$, $t > 0$.

A brief simplification: we need only show that $f = 0$ for $0 < t \leq t_0$ for some $t_0 > 0$; then applying this statement k times gives $f = 0$ for $t \leq kt_0$ and we can let $k \to \infty$. Another simplification: we need only show $f \leq 0$; then applying this statement to $-f$ we conclude $f = 0$.

Here's the idea: we'll show $f \leq 0$ by applying the maximum principle not to f, but rather to

$$g(x,t) = f(x,t) - \frac{\delta}{(t_1 - t)^{n/2}} e^{\frac{|x|^2}{4(t_1 - t)}}.$$

for suitable choices of the constants δ and t_1. The space-time cylinder will be of the form $D \times [0, t_0]$ where D is a large ball and $t_0 < t_1$.

Step 1. Observe that $g_t - \Delta g = 0$. This can be checked by direct calculation. But a more conceptual reason is this: the term we've subtracted from f is a constant times the fundamental solution evaluated at ix and $t_1 - t$. The heat equation is invariant under this change of variables.

Step 2. Let D be a ball of radius r. We know from the maximum principle that the maximum of g on $D \times [0, t_0]$ is achieved at the initial boundary or spatial boundary. At the initial boundary clearly

$$g(x,0) < f(x,0) = 0.$$

At the spatial boundary we have $|x| = r$ so

$$\begin{aligned} g(x,t) &= f(x,t) - \frac{\delta}{(t_1-t)^{n/2}} e^{\frac{r^2}{4(t_1-t)}} \\ &\leq Me^{c|x|^2} - \frac{\delta}{(t_1-t)^{n/2}} e^{\frac{r^2}{4(t_1-t)}} \\ &\leq Me^{cr^2} - \frac{\delta}{t_1^{n/2}} e^{\frac{r^2}{4t_1}} \end{aligned}$$

We may choose t_1 so that $1/(4t_1) > c$. Then when r is large enough the second term dominates the first one, giving

$$g(x,t) \leq 0 \quad \text{at the spatial boundary } |x| = r.$$

We conclude from the maximum principle that $g(x,t) \leq 0$ on the entire space-time cylinder. This argument works for any sufficiently large r, so we have shown that

$$f(x,t) \leq \frac{\delta}{(t_1-t)^{n/2}} e^{\frac{|x|^2}{4(t_1-t)}}$$

for all $x \in R^n$ and all $t < t_1$. Restricting attention to $t < t_0$ for some fixed $t_0 < t_1$, we pass to the limit $\delta \to 0$ to deduce that $f \leq 0$ as desired. This completes the proof of uniqueness.

PDE for Finance Notes, Spring 2003 – Addendum to Section 3.
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course PDE for Finance, G63.2706.

Numerical solution by finite differences. Before leaving the linear heat equation, let's briefly discuss how it can be solved numerically. These notes consider only the most basic numerical scheme: explicit finite differences, following roughly the discussion in F. John's book. For more information (including more sophisticated schemes) see e.g. Chapter 8 of the "student guide" by Wilmott, Howison, and Dewynne.

But first, some corrections to recent handouts:

- Problem 5b on HW2 was wrong by a factor of 2: the first passage time density is $\frac{1}{2}\frac{\partial G}{\partial z}(z_0, 0, t)$. The extra factor of $\frac{1}{2}$ arises because the process under consideration is Brownian motion with drift, not $\sqrt{2}$ times Brownian motion with drift. The HW2 solution sheet also missed this point. (The solution sheet now posted on my web page has been corrected.)

- When the infinitesimal generator is self-adjoint, the Green's function is symmetric, i.e. $G(x, y, t) = G(y, x, t)$ (see Problem 2 on HW3). Otherwise it isn't (Brownian motion with drift is an example where G is not symmetric – as we know from HW2). So it is a good idea to maintain the convention that $G(x, y, t)$ is the probability of that a walker starting from x at time 0 arrives at y at time t. Section 3 was sloppy about this, in Eqn. (5) and the text just following it. I should have said "It is natural to seek a solution formula in the form

$$u(y, t) = \int_0^1 G(x, y, t) g(x)\, dy$$

since $G(x, y, t)$ is then the probability that a random walker starting at x arrives at y at time t without first hitting the boundary."

****

Numerical time-stepping is perhaps the most straightforward method for solving a linear parabolic PDE in a bounded domain. An in-depth treatment is beyond the scope of the present course. But let's spend a moment on the most basic example: an explicit finite-difference scheme for the linear heat equation $f_t = f_{xx}$ on the unit interval $0 < x < 1$ as our spatial interval. We suppose, as usual, that the value of f is specified at the boundary points $x = 0$ and $x = 1$.

If the timestep is Δt and the spatial length scale is Δx then the numerical f is defined at $(x, t) = (j\Delta x, k\Delta t)$. The explicit finite difference scheme determines f at time $(j + 1)\Delta t$ given f at time $j\Delta t$ by reading it off from

$$\frac{f((j+1)\Delta t, k\Delta x) - f(j\Delta t, k\Delta x)}{\Delta t} = \frac{f(j\Delta t, (k+1)\Delta x) - 2f(j\Delta t, k\Delta x) + f(j\Delta t, (k-1)\Delta x)}{(\Delta x)^2}.$$

Notice that we use the initial data to get started, and we use the boundary data when $k\Delta x$ is next to the boundary.

This method has the stability restriction

$$\Delta t < \frac{1}{2}(\Delta x)^2. \tag{1}$$

To see why, observe that the numerical scheme can be rewritten as

$$f((j+1)\Delta t, k\Delta x) = \tfrac{\Delta t}{(\Delta x)^2} f(j\Delta t, (k+1)\Delta x) + \tfrac{\Delta t}{(\Delta x)^2} f(j\Delta t, (k-1)\Delta x) + (1 - 2\tfrac{\Delta t}{(\Delta x)^2}) f(j\Delta t, k\Delta x).$$

If $1 - 2\frac{\Delta t}{(\Delta x)^2} > 0$ then the scheme has a discrete maximum principle: if $f \leq C$ initially and at the boundary then $f \leq C$ for all time; similarly if $f \geq C$ initially and at the boundary then $f \geq C$ for all time. The proof is easy, arguing inductively one timestep at a time. (If the stability restriction is violated then the scheme is unstable, and the discrete solution can grow exponentially.)

One can use this numerical scheme to prove existence (see e.g. John). But let's be less ambitious: let's just show that the numerical solution converges to the solution of the PDE as Δx and Δt tend to 0 while obeying the stability restriction (1). The main point is that the scheme is consistent, i.e.

$$\frac{g(t + \Delta t, x) - g(t, x)}{\Delta t} \to g_t \quad \text{as } \Delta t \to 0$$

and

$$\frac{g(t, x + \Delta x) - 2g(t, x) + g(t, x - \Delta x)}{(\Delta x)^2} \to g_{xx} \quad \text{as } \Delta x \to 0$$

if g is smooth enough. Let f be the numerical solution, g the PDE solution, and consider $h = f - g$ evaluated at gridpoints. Consistency gives

$$\begin{aligned} h((j+1)\Delta t, k\Delta x) &= \tfrac{\Delta t}{(\Delta x)^2} h(j\Delta t, (k+1)\Delta x) + \tfrac{\Delta t}{(\Delta x)^2} h(j\Delta t, (k-1)\Delta x) \\ &\quad + (1 - 2\tfrac{\Delta t}{(\Delta x)^2}) h(j\Delta t, k\Delta x) + \Delta t e(j\Delta t, k\Delta x) \end{aligned}$$

with $|e|$ uniformly small as Δx and Δt tend to zero. Stability – together with the fact that $h = 0$ initially and at the spatial boundary – gives

$$|h(j\Delta t, k\Delta x)| \leq j\Delta t \max |e|.$$

It follows that $h(t, x) \to 0$, uniformly for bounded $t = j\Delta t$, as Δt and Δx tend to 0.

The preceding argument captures, in this special case, a general fact about numerical schemes: that stability plus consistency implies convergence.

PDE for Finance Notes, Spring 2003 – Section 4

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, G63.2706.

Deterministic optimal control. We've been studying the solution of an SDE, modelling for example the price of an asset in the marketplace. One focus of our attention has been the backward Kolmogorov equation, whose solution $u(x,t)$ is the expected final-time payoff, if the system starts in state x at time t. The backward Kolmogorov equation is a linear PDE, which must be solved backward in time.

We're heading toward *stochastic control.* That theory considers SDE's over which we have some influence, modelling for example the value of a portfolio. One typical goal is to maximize the utility of final-time wealth. The task is now two-fold: (i) to identify an optimal strategy, and (ii) to evaluate the associated "value function" $u(x,t)$ – the optimal utility of final-time wealth, if the system starts in state x at time t. It solves the Hamilton-Jacobi-Bellman (HJB) equation – the analogue in this setting of the backward Kolmogorov equation. The HJB equation is usually *nonlinear*, due to the effect of our decision-making. Like the backward Kolmogorov equation, it must be solved backward in time. Underlying the derivation and solution of the HJB equation is the *dynamic programming principle* – a powerful scheme for solving optimization problems by gradually increasing the time-to-maturity (or a similar parameter).

Selected financial applications of stochastic control include: (a) optimizing the allocation of assets between distinct risky investment opportunities (e.g. a time-dependent version of CAPM); (b) optimizing the rate at which to spend income from an investment portfolio; (c) optimal hedging of an option on a non-traded underlying; and (d) pricing of American options (i.e. optimization of the "exercise rule"). All these problems involve a blend of (i) stochasticity and (ii) control. We are already well-versed in stochasticity. We now take a little time to become well-versed in control, by studying *deterministic* control problems. Then we'll combine the two ingredients (in Sections 5 and beyond) to address financially signficant examples.

The material covered in this section is "standard," however I don't know a really good place to read it. The book by Fleming and Rishel, *Deterministic and Stochastic Optimal Control*, covers everything here and much more – but it goes much deeper than the level of this class. Roughly the same comment applies to the book of Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models.* The charming and inexpensive book A. Dixit, *Optimization in Economic Theory* (Oxford Univ Press, 1990, paperback) covers some closely related material. Most basic engineering-oriented texts such as Macki & Strauss, *Introduction to Optimal Control Theory*, emphasize the Pontryagin Maximum Principle and de-emphasize the method of dynamic programming. By contrast, we will emphasize dynamic programming because it is the more useful approach for many of the financial applications. (There is a stochastic-control analogue of the Pontryagin Maximum Principle. It is the "martingale method," presented e.g. in Section 5.2 of Korn & Korn, *Option Pricing and Portfolio Optimization.* Perhaps we'll find time to discuss this toward the end of the semester.) Those with a lot of background in PDE's will enjoy reading Chapter 10

of Evans, *Partial Differential Equations*, where he discusses the relation between deterministic optimal control and the *viscosity solution* of the associated Hamilton-Jacobi-Bellman equation. However the analysis of viscosity solutions lies beyond the scope of this course.

Once we get to stochastic control it will be easier to suggest things to read. We'll be following, among other things, pieces of Merton's work; the corresponding articles – reprinted in R.C. Merton, *Continuous Time Finance* – are very rich and not too difficult.

This section is relatively long, because there's a lot to cover. We shall focus however on two key examples. *Example 1* is the deterministic analogue of Merton's classic example involving optimal investment and consumption. *Example 2* is a minimum-time problem whose HJB equation is the Eikonal equation $|\nabla u| = 1$. Example 2 has, I guess, no financial interpretation but its easy-to-visualize character makes it a convenient aid to understanding. General topics which we'll explain, specializing where convenient to these examples, include: finding the relevant *Hamilton-Jacobi-Bellman equation*; identifying the optimal *feedback law*; and proving optimality by means of a *verification argument*.

****

What is optimal control? A typical problem of optimal control is this: we have a system whose *state* at any time t is described by a vector $y = y(s) \in R^n$. The system evolves in time, and we have the ability to influence its evolution through a vector-valued *control* $\alpha(s) \in R^m$. The evolution of the system is determined by an *ordinary differential equation*

$$\dot{y}(s) = f(y(s), \alpha(s)), \quad y(0) = x, \tag{1}$$

and our goal is to choose the function $\alpha(s)$ for $0 < s < T$ so as to maximize some utility or minimize some cost, e.g.

$$\max \int_0^T h(y(s), \alpha(s))\, ds + g(y(T)). \tag{2}$$

The problem is determined by specifying the dynamics f, the initial state x, the final time T, the "running utility" h and the "final utility" g. The problem is solved by finding the optimal control $\alpha(s)$ for $0 < s < T$ and the value of the maximum.

The mathematical and engineering literature often focuses on minimizing some sort of cost; the economic literature on maximizing utility. The two problems are mathematically equivalent.

One needs some hypotheses on f to be sure the solution of the ODE defining $y(s)$ exists and is unique. We do not make these explicit since the goal of these notes is to summarize the main ideas without getting caught up in fine points. See Evans for a mathematically careful treatment. Another technical point: it's possible (even easy) to formulate optimal control problems that have no solution. If the utility is bounded above, then for any $\epsilon > 0$ there's certainly a control $\alpha_\epsilon(s)$ achieving a value within ϵ of optimal. But the controls α_ϵ might not converge to a meaningful control as $\epsilon \to 0$. Note however that even if an optimal control doesn't exist, the optimal *value* (the maximum utility) is still well-defined.

An optimal control problem is evidently a special type of optimization problem. What's special is that we're dealing with functions of time, and decisions that must be made as time proceeds. Often the optimal control is described by a *feedback law*. Such a law determines the optimal control $\alpha(s)$ as having the form $\alpha(s) = F(y(s), s)$ for some function F (the feedback law).

Example 1: Here is a simple example which already has financial interest. (It's a deterministic version of Merton's famous example of optimal investment and consumption; we'll do the version with investment in a few weeks). Consider an individual whose wealth today is x, and who will live exactly T years. His task is to plan the rate of *consumption* of wealth $\alpha(s)$ for $0 < s < T$. All wealth not yet consumed earns interest at a fixed rate r. The state equation is thus

$$\dot{y} = ry - \alpha, \quad y(0) = x. \tag{3}$$

The control is $\alpha(s) \geq 0$, and the state is constrained by $y(s) \geq 0$ (he cannot consume wealth he doesn't have). The goal is

$$\max \int_0^T e^{-\rho s} h(\alpha(s))\, ds$$

where ρ is the discount rate and $h(\alpha)$ is the utility of consumption. (The function h, which must be given as part of the formulation of the problem, should be monotonically increasing and concave. A typical choice is $h(\alpha) = \alpha^\gamma$ with $0 < \gamma < 1$.) We have, for simplicity, assigned no utility to final-time wealth (a bequest), so the solution will naturally have $y(T) = 0$. Our goal is not strictly of the form (2) due to the presence of discounting; well, we omitted discounting from (2) only for the sake of simplicity.

The state constraint $y(s) \geq 0$ is awkward to deal with. In practice it tells us that if the investor ever runs out of wealth (i.e. if $y(s)$ ever reaches 0) then $\alpha = 0$ and $y = 0$ thereafter. This state constraint can be avoided by reformulating the goal as

$$\max \int_0^\tau e^{-\rho s} h(\alpha(s))\, ds$$

where τ is the first time y reaches 0 if this occurs before T, or $\tau = T$ if y is positive for all $s < T$. With this goal we need not impose the state constraint $y(s) \geq 0$.

Control theory is related to – but much more general than – the one-dimensional calculus of variations. A typical calculus of variations problem is

$$\max_{y(s)} \int_0^T W(s, y(s), \dot{y})\, ds$$

subject, perhaps, to endpoint conditions on $y(0)$ and $y(T)$. The example just formulated can be expressed in this form,

$$\max_{y(s)} \int_0^T e^{-\rho s} h(ry - \dot{y})\, ds, \quad \text{subject to } y(0) = x,$$

except that we have additional constraints $ry(s) - \dot{y}(s) \geq 0$ and $y(s) \geq 0$ for all s.

We will shortly discuss the method of dynamic programming as a scheme for solving optimal control problems. The key to this method is to consider how the solution depends on the

initial time and *initial state* as parameters. Thus rather than start arbitrarily at time 0, it is better to introduce a variable initial time t. And it is fruitful to consider the *value function* $u(x,t)$, the optimal value achievable using initial time t and initial state x. In the context of our basic framework (1) this means changing the state equation to

$$\dot{y}(s) = f(y(s), \alpha(s)), \quad y(t) = x.$$

The control $\alpha(s)$ is now to be determined for $t < s < T$, and the value function is

$$u(x,t) = \max \int_t^T h(y(s), \alpha(s))\, ds + g(y(T)).$$

In the context of Example 1 it means changing the state equation to

$$\dot{y} = ry - \alpha, \quad y(t) = x,$$

and the objective to

$$u(x,t) = \max \int_t^T e^{-\rho s} h(\alpha(s))\, ds.$$

(Warning: with this definition $u(x,t)$ is the utility of consumption discounted to time 0. The utility of consumption discounted to time t is $e^{\rho t}u(x,t)$.)

We started by formulating the "typical" optimal control problem (1)-(2). Now let's discuss some of the many variations on this theme, to get a better sense of the scope of the subject. We repeat for clarity the state equation:

$$\dot{y}(s) = f(y(s), \alpha(s)) \text{ for } t < s < T \text{ with initial data } y(t) = x.$$

Sometimes we may wish to emphasize the dependence of $y(s)$ on the initial value x, the initial time t, and the choice of control $\alpha(s), t < s < T$; in this case we write $y = y_{x,t,\alpha}(s)$. The control is typically restricted to take values in some specified set A, independent of s:

$$\alpha(s) \in A \text{ for all } s;$$

the set A must be specified along with the dynamics f. Sometimes it is natural to impose *state constraints*, i.e. to require that the state $y(s)$ stay in some specified set Y:

$$y_{x,t,\alpha}(s) \in Y \text{ for all } s;$$

when present, this requirement restricts the set of admissible controls $\alpha(s)$. Our basic example (2) is known as a **finite horizon** problem; its value function is

$$u(x,t) = \max_\alpha \left\{ \int_t^T h(y_{x,t,\alpha}(s), \alpha(s))\, ds + g(y_{x,t,\alpha}(T)) \right\}. \tag{4}$$

For the analogous **infinite horizon** problem it is customary to set the starting time to be 0, so the value function depends only on the spatial variable x:

$$u(x) = \max_\alpha \int_0^\infty e^{-\rho s} h(y_{x,0,\alpha}(s), \alpha(s))\, ds. \tag{5}$$

Discounting is important for the infinite-horizon problem, since without it the integral defining u could easily be infinite. (As already noted in our example, it is also often natural to include discounting in a finite-horizon problem.)

The **minimum time** problem is a little bit different. It minimizes the time it takes $y(s)$ to travel from x to some target set $\mathcal{G}$. The value function is thus

$$u(x) = \min_{\alpha} \{\text{time at which } y_{x,0,\alpha}(s) \text{ first arrives in } \mathcal{G}\}. \tag{6}$$

The minimum time problem is somewhat singular: if, for some x, the solution starting at x cannot arrive in $\mathcal{G}$ (no matter what the control) then the value is undefined. The **discounted minimum time** problem avoids this problem: its value function is

$$u(x) = \min_{\alpha} \int_0^{\tau(x,\alpha)} e^{-s}\, ds \tag{7}$$

where $\tau(x,\alpha)$ is the time that $y_{x,0,\alpha}(s)$ first arrives in $\mathcal{G}$, or infinity if it never arrives. Notice that the integral can be evaluated: the quantity being minimized is $\int_0^{\tau(x,\alpha)} e^{-s}\, ds = 1 - e^{-\tau(x,\alpha)}$. So we're still minimizing the arrival time, but the value function is $1 - \exp(-\text{arrival time})$ instead of the arrival time itself.

Example 2. Here is a simple example of a minimum-time problem, with the great advantages that (a) we can easily visualize everything, and (b) we know the solution in advance. In its simplest form the problem is: given a point x in R^n, and a set $\mathcal{G}$ not containing x, find the distance from x to $\mathcal{G}$. We recognize this as a minimum time problem, by reformulating it in terms of paths travelled with speed ≤ 1. The state equation is

$$dy/ds = \alpha(s), \quad y(0) = x,$$

and the controls are restricted by

$$|\alpha(s)| \leq 1.$$

The minimum arrival time

$$u(x) = \min_{\alpha} \{\text{time of arrival at } \mathcal{G}\}$$

is of course the distance from x to $\mathcal{G}$, and the optimal strategy is to travel with constant velocity (and unit speed) toward the point in $\mathcal{G}$ that is closest to x. We remark that $u(x) = \text{dist}\,(x, \mathcal{G})$ solves the differential equation

$$|\nabla u| = 1$$

in its natural domain $\Omega = R^n - \mathcal{G}$, with boundary condition $u = 0$ at $\partial\Omega$. This is an example of a (time-independent) Hamilton-Jacobi equation. The solution is typically not smooth: consider for example the case when Ω is a circle or a square. The optimal control is determined by a feedback law ("wherever you are right now, proceed at unit speed toward the nearest point on the target $\mathcal{G}$"). The non-smoothness of u reflects the fact that the feedback law is discontinuous, with nonuniqueness where ∇u is discontinuous. There is clearly nothing pathological about this example: non-smooth value functions, and discontinuous feedback laws, are commonplace in deterministic optimal control.

****

Dynamic programming. There are basically two systematic approaches to solving optimal control problems: one known as the Pontryagin Maximum Principle, the other known as Dynamic Programming. The two approaches are fundamentally equivalent, though in specific problems one may be easier to apply than the other. We shall emphasize dynamic programming, because (a) it extends more easily to the random case (time-dependent decision-making to optimize expected utility), and (b) it extends the familiar financial procedure of valuing an option by working backward through a tree.

The essence of dynamic programming is pop psychology: "today is the first day of the rest of your life." More: every day is the first day of the future thereafter. How to use this insight? One way is to make it the basis of a numerical solution scheme. Another way is to use it to derive a PDE for $u(x,t)$. These two ideas are of course closely related: our numerical solution scheme is in fact a crude numerical scheme for solving the PDE.

Let's start with the numerical scheme, concentrating on the finite-horizon problem (4), and keeping space one-dimensional for simplicity. Our goal is to compute (approximately) the value function $u(x,t)$. Of course any numerical scheme must work in discrete space and time, so t is a multiple of Δt, and x is a multiple of Δx. It's also natural to consider that the controls are discretized: $\alpha(s)$ is piecewise constant with mesh Δt. Now work backward in time:

First Consider the problem with initial time $t = T$. In this case the dynamics is irrelevant. So are the control and the running utility . Whatever the value of x, the associated value function is $g(x)$. In other words: $u(x,T) = g(x)$.

Next Consider the problem with initial time $t = T - \Delta t$. Approximate the dynamics as

$$y(s + \Delta t) = y(s) + f(y(s), \alpha(s))\Delta t.$$

Since there is just one time interval between the initial time t and the final time $T = t + \Delta t$, and since the control is piecewise constant, the unknown is now just a single vector $\alpha = \alpha(t)$ (not a function). It is determined by optimization. We may approximate the objective integral by a sum (dropping terms of higher order in Δt), leading to

$$u(x, T - \Delta t) = \max_{\alpha} \{h(x,\alpha)\Delta t + g\left(x + f(x,\alpha)\Delta t\right)\}.$$

This must be evaluated for each x (i.e. every multiple of Δx), and the maximization over α must be done globally (we need the global optimum, not just a local optimum). For a real numerical scheme some further structure is needed here: we should solve a problem in a bounded spatial domain, and impose concavity hypotheses assuring that there are no local optima. For the present conceptual discussion let us ignore such practical issues and proceed. (One might worry that when the spatial dimension is greater than 1 this scheme is utterly impractical, since the number of grid points x to be considered at each time t is of order $(\Delta x)^{-n}$ in dimension n. This worry is well-founded: our scheme is impractical in higher dimensions. However

there are good numerical schemes for multidimensional problems. One option is to solve the Hamilton-Jacobi-Bellman equation we'll derive presently, using a suitable finite-difference scheme.) At the end of this step we have computed $u(\cdot, T - \Delta t)$ as a function of space.

Next Consider the problem with initial time $t = T - 2\Delta t$. For any initial state $x = y(t)$, the possible controls are now represented by a pair of vectors $\alpha(t), \alpha(t + \Delta t)$. However we can still solve the problem by considering just the current control $\alpha = \alpha(t)$, since the optimal choice of $\alpha(t + \Delta t)$ has already been determined in the course of evaluating $u(x, T - \Delta t)$. Making crucial use of the fact that the "running utility" is an integral in time, we may determine the optimal value $u(x, T - 2\Delta t)$ by solving

$$u(x, T - 2\Delta t) = \max_{\alpha} \{h(x, \alpha)\Delta t + u\left(x + f(x, \alpha)\Delta t, T - \Delta t\right)\}.$$

Here the unknown is just the control α to be used during the time interval from $T - 2\Delta t$ to $T - \Delta t$. The optimal α depends of course on x, and the optimization in α must be done for each choice of x separately. (Again, this is the conceptual but impractical version; numerical optimal control uses various workarounds to make it more practical.) At the end of this step we have computed $u(\cdot, T - 2\Delta t)$ as a function of space.

Continue The scheme continues, working backward time-step by time-step. Notice that for computing $u(x, T - (j + 1)\Delta t)$ we need only save the values of $u(x, T - j\Delta t)$. However if we wish to synthesize an optimal control starting at an arbitary point x and time $t = T - (j+1)\Delta t$ we must save much more information: namely the feedback law $\alpha = F(y, s)$, obtained in the course of calculating $u(y, s)$ for $s > t$. (This is the optimal initial-time-period value of the control, when the initial state is y and the initial time is s). This information permits us to synthesize the optimal control and solve the state equation at the same time: starting from x at time t, the state evolves by

$$y_\alpha(s + \Delta t) = y_\alpha(s)) + f(y_\alpha(s), \alpha(s))\Delta t$$

with $\alpha(s)$ determined by

$$\alpha(s) = F\left(y_\alpha(s), s\right).$$

We remark that a similar philosophy can be used in many other settings. One example is this standard scheme for computing the shortest path between two nodes of a graph. Pick one of the nodes (call it an endpoint). Find all nodes that lie distance 1 from it, then all points that lie distance 2 from it, etc. Stop when the other endpoint appears in the set you come up with.

Students of math finance will have noticed by now that dynamic programming looks a lot like the binomial-tree method for valuing a European or American option. The resemblance is no coincidence. The biggest difference is that for the European option no optimization need be done at any point in the calculation; for the American option the optimization is simple – over just two alternatives, to exercise or not to exercise. This is due to the completeness of the underlying market model. In a multiperiod market that's not complete, there *is* an

optimization to be done at each stage. We'll discuss an example of this type when we get to stochastic optimal control.

The discrete-time, discrete-space scheme described above can be viewed as a crude numerical scheme for solving the PDE satisfied by the value function. This is known as the **Hamilton-Jacobi-Bellman equation**. We shall derive it, in essence, by taking the formal limit $\Delta t \to 0$ in our numerical discussion. This viewpoint can be used for all the optimal control problems we've discussed (finite-horizon, infinite-horizon, least-time, with or without discounting) but to fix ideas we concentrate on the usual finite-horizon example

$$u(x,t) = \max_{\alpha} \left\{ \int_t^T h(y(s), \alpha(s))\, ds + g(y(T)) \right\}$$

where the controls are restricted by $\alpha(s) \in A$, and the state equation is

$$dy/ds = f(y(s), \alpha(s)) \text{ for } t < s < T \text{ and } y(t) = x.$$

(Space can be multidimensional here.) The Hamilton-Jacobi-Bellman equation in this case is

$$u_t + H(\nabla u, x) = 0 \quad \text{for } t < T \tag{8}$$

with

$$u(x,T) = g(x) \quad \text{at } t = T,$$

where H (the "Hamiltonian") is defined by

$$H(p,x) = \max_{a \in A} \{ f(x,a) \cdot p + h(x,a) \}. \tag{9}$$

(Note that p is a vector with the same dimensionality as x; a is a vector with the same dimensionality as α.)

To explain, we start with the **dynamic programming principle**, which was in fact the basis of our discrete scheme. It says:

$$u(x,t) = \max_{\alpha} \left\{ \int_t^{t'} h(y_{x,t,\alpha}(s), \alpha(s))\, ds + u(y_{x,t,\alpha}(t'), t') \right\} \tag{10}$$

whenever $t < t' < T$. The justification is easy, especially if we assume that an optimal control exists (this case captures the main idea; see Evans for a more careful proof, without this hypothesis). Suppose the optimal utility starting at x at time t is achieved by an optimal control $\alpha_{x,t}(s)$. Then the restriction of this control to any subinterval $t' < s < T$ must be optimal for its starting time t' and starting position $y_{x,t,\alpha}(t')$. Indeed, if it weren't then there would be a new control $\alpha'(s)$ which agreed with α for $t < s < t'$ but did better for $t' < s < T$. Since the utility is additive – the running utility is $\int_t^T h(y,\alpha)\, ds = \int_t^{t'} h(y,\alpha)\, ds + \int_{t'}^T h(y,\alpha)\, ds$ – this new control would be better for the entire time period, contradicting the optimality of α. Therefore in defining $u(x,t)$ as the optimal utility, we can restrict our attention to controls that are optimal from time t' on. This leads immediately to (10).

Now let us derive (heuristically) the Hamilton-Jacobi-Bellman equation. The basic idea is to apply the dynamic programming principle with $t' = t+\Delta t$ and let $\Delta t \to 0$. Our argument is heuristic because (i) we assume u is differentiable, and (ii) we assume the optimal control is adequately approximated by one that is constant for $t < s < t + \Delta t$. (Our goal, as usual, is to capture the central idea, referring to Evans for a more rigorous treatment.) Since Δt is small, the integral on the right hand side of (10) can be approximated by $h(x, a)\Delta t$, where $a \in A$ is the (constant) value of α for $t < s < t + \Delta t$. Using a similar approximation for the dynamics, the dynamic programming principle gives

$$u(x, t) \geq h(x, a)\Delta t + u(x + f(x, a)\Delta t, t + \Delta t) + \text{errors we wish to ignore}$$

with equality when a is chosen optimally. Using the first-order Taylor expansion of u this becomes

$$u(x, t) \geq h(x, a)\Delta t + u(x, t) + (\nabla u \cdot f(x, a) + u_t)\Delta t + \text{error terms}$$

with equality when a is optimal. In the limit $\Delta t \to 0$ this gives

$$0 = u_t + \max_{a \in A}\{\nabla u \cdot f(x, a) + h(x, a)\},$$

i.e. $u_t + H(\nabla u, x) = 0$ with H as asserted above. The final-time condition is obvious: if $t = T$ then the dynamics is irrelevant, and the total utility is just $g(x)$.

That was easy. Other classes of optimal control problems are treated similarly. Let's look at the minimum-time problem, where the state evolves by

$$dy/ds = f(y, \alpha), \quad y(t) = x,$$

and the controls are restricted by

$$\alpha(s) \in A \quad \text{for all } s$$

for some set A. The associated Hamilton-Jacobi-Bellman equation is

$$H(\nabla u, x) = -1 \quad \text{for } x \notin \mathcal{G}$$

with Hamiltonian

$$H(p, x) = \min_{a \in A}\{f(x, a) \cdot p\} = 0.$$

The boundary condition is

$$u = 0 \quad \text{for } x \in \mathcal{G}.$$

To see this, we argue essentially as before: the value function (the time it takes to arrive at $\mathcal{G}$) should satisfy

$$u(x) \leq \Delta t + u(x + f(x, a)\Delta t) + \text{error terms}$$

for any $a \in A$, with equality when a is optimal. Using Taylor expansion this becomes

$$u(x) \leq \Delta t + u(x) + \nabla u \cdot f(x, a)\Delta t + \text{error terms}.$$

Optimizing over a and letting $\Delta t \to 0$ we get

$$1 + \min_{a \in A}\{f(x,a) \cdot \nabla u\} = 0,$$

which is the desired equation.

Let us specialize this to Example 2. In that example the set A is the unit ball, and $f(y,\alpha) = \alpha$, so $H(p,x) = \min_{|a|\le 1} p \cdot a = -|p|$ and the Hamilton-Jacobi equation becomes $|\nabla u| = 1$, as expected.

****

Solutions of the Hamilton-Jacobi-Bellman equation are not unique (at least, not when we understand "solution" in the naive almost-everywhere sense). For example, there are *many* Lipschitz continuous solutions of $|\nabla u| = 1$ in a square, with $u = 0$ at the boundary. If one were smooth we might prefer it – however there is *no* smooth solution. So, is the HJB equation really of any use?

The answer is yes, it's very useful, for three rather distinct reasons. The first is obvious; the second is elementary but not obvious; the third is subtle, representing a major mathematical achievement of the past 20 years:

(a) In deriving the HJB equation, we deduced a relation between the optimal control and the value of ∇u: briefly, $\alpha(s)$ achieves the optimum in the definition of $H(p,x)$ with $p = \nabla u(y(s),s)$ and $x = y(s)$. Thus the derivation of the HJB equation tells us the relation between the value function and the optimal control. In many settings, this argument permits us to deduce a feedback law once we know the value function.

(b) The argument used for the HJB equation can often be reorganized to show that a conjectured formula for the value function is correct. This sort of argument is called a *verification theorem.*

(c) There is a more sophisticated notion of "solution" of a Hamilton-Jacobi equation, namely the notion of a *viscosity solution.* Viscosity solutions exist, are unique, and can be computed by suitable numerical schemes. Moreover the value function of a dynamic programming problem is automatically a viscosity solution of the associated HJB equation. (Chapter 10 of Evans' book gives an excellent introduction to the theoretical side of this topic. The book *Level Set Methods* by J. Sethian, Cambridge Univ Press, provides a readable introduction to the numerical side, concentrating on the special class of HJB equations associated with geometric evolution problems – closely connected with our minimum time example.)

Point (a) should be clear, and it will be illuminated further by various examples later on. Point (c) is an interesting story, but beyond the scope of the present discussion. Our present intention is to concentrate on point (b). We focus as usual on the setting of the finite-horizon problem. As usual, $u(x,t)$ denotes the value function (the maximal value achievable starting from state x at time t). Our plan is to develop schemes for proving upper and lower bounds

on u. If we do a really good job the upper and lower bounds will coalesce – in which case they will fully determine u.

There's always one type of bound that is easy. Since we're maximizing utility, these are the lower bounds. Any scheme for choosing the control – for example a conjectured feedback law specifying $\alpha(s)$ as a function of $y(s)$ – provides a lower bound $v(x,t) =$ the value achieved by this scheme. The inequality

$$v(x,t) \leq u(x,t)$$

is obvious, since u is the maximal value obtainable using *any* control – including the ones used to define v.

The **verification theorem** provides the other bound. In its most basic form – specialized to the present setting – it says the following. Suppose $w(x,t)$ is defined (and continuously differentiable) for $t < T$, and it solves the Hamilton-Jacobi equation (8) with $w = g$ at $t = T$. Then w is an upper bound for the value function:

$$u(x,t) \leq w(x,t).$$

To see why, consider any candidate control $\alpha(s)$ and the associated state $y = y_{x,\alpha}(s)$ starting from x at time t. The chain rule gives

$$\begin{aligned} \frac{d}{ds}w(y(s),s) &= w_s(y(s),s) + \nabla w(y(s),s)\cdot \dot{y}(s) \\ &= w_s(y(s),s) + \nabla w(y(s),s)\cdot f(y(s),\alpha(s)) \\ &\leq w_s + H(\nabla w, y) - h(y(s),\alpha(s)) \qquad (11) \\ &= -h(y(s),\alpha(s)), \end{aligned}$$

using for (11) the relation

$$H(p,y) = \max_{a\in A}\{f(y,a)\cdot p + h(y,a)\} \geq f(y,\alpha)\cdot p + h(y,\alpha)$$

with $y = y(s)$, $\alpha = \alpha(s)$, and $p = \nabla w(y(s),s)$. Now integrate in time from t to T:

$$w(y(T),T) - w(x,t) \leq -\int_t^T h(y(s),\alpha(s))\,ds.$$

Since $w(y(T),T) = g(y(T))$ this gives

$$g(y(T)) + \int_t^T h(y(s),\alpha(s))\,ds \leq w(x,t).$$

The preceding argument applies to any control $\alpha(s)$. Maximizing the left hand side over all admissible controls, we have

$$u(x,t) \leq w(x,t)$$

as asserted.

We presented the task of finding lower and upper bounds as though they were distinct, but of course they are actually closely correlated. A smooth solution w of the Hamilton-Jacobi

equation comes equipped with its own feedback law (as discussed in point (a) above). It is natural to consider the lower bound v obtained using the feedback law associated with w. I claim that this v is equal to w. To see this, follow the line of reasoning we used for the verification theorem, noticing that (11) holds with equality if α is determined by the feedback associated with w. Therefore integration gives

$$w(x,t) = g(y(T)) + \int_t^T h(y(s), \alpha(s))\, ds$$

and the right hand side is, by definition, $v(x,t)$. In conclusion: if w is a (continuously differentiable) solution of the HJB equation, satisfying the appropriate final-time condition too, then w *is in fact the value function* $u(x,t)$.

It sounds like a great scheme, and in many ways it is. There is however a small fly in the ointment. Sometimes the value function isn't continuously differentiable. (Consider, for example, the minimum time problem). In such a case our proof of the verification theorem remains OK for paths that avoid the locus of nonsmoothness – or cross it transversely. But there is a problem if the state should happen to hug the locus of nonsmoothness. Said more plainly: if $w(x,t)$ has discontinuous derivatives along some set Γ in space-time, and if a control makes $(y(s), s)$ move along Γ, then the first step in our verification argument

$$\frac{d}{ds} w(y(s), s) = w_s(y(s), s) + \nabla w(y(s), s) \cdot \dot{y}(s)$$

doesn't really make sense (for example, the right hand side is not well-defined). Typically this problem is overcome by using the fact that the verification argument has some extra freedom: it doesn't really require that w solve the HJB equation exactly. Rather, it requires only that w satisfy the *inequality* $w_t + H(\nabla w, t) \leq 0$.

To give an example where this extra freedom is useful consider our geometrical Example 2, with target $\mathcal{G}$ the complement of the unit square in R^2. The HJB equation is $|\nabla u| = 1$ in Ω=unit square, with $u = 0$ at $\partial\Omega$. The value function is defined as $u(x)$=minimum time of arrival to $\partial\Omega$ (among all paths with speed ≤ 1). Simple geometry tells us the solution is the distance function $\text{dist}\,(x, \partial\Omega)$, whose graph is a pyramid. We wish to give an entirely PDE proof of this fact.

One inequality is always easy. In this case it is the relation $u(x) \leq \text{dist}\,(x, \partial\Omega)$. This is clear, because the right hand side is associated with a specific control law (namely: travel straight toward the nearest point of the boundary, with unit speed). To get the other inequality, observe that if $w \leq 0$ at $\partial\Omega$ and $|\nabla w| \leq 1$ in Ω then

$$\begin{aligned} \frac{d}{ds} w(y(s)) &= \nabla w(y(s)) \cdot \dot{y}(s) \\ &= \nabla w(y(s)) \cdot \alpha(s) \\ &\geq -|\nabla w(y(s))| \geq -1. \end{aligned}$$

(Here $y(s)$ solves the state equation $\dot{y} = \alpha$, with initial condition $y(0) = x$ and any admissible control $|\alpha(s)| \leq 1$.) If τ is the time of arrival at $\partial\Omega$ then integration gives

$$w(y(\tau)) - w(x) \geq \int_0^\tau (-1)\, ds.$$

Since $w(y(\tau)) \leq 0$ we conclude that

$$w(x) \leq \tau.$$

Minimizing the right hand side over all admissible controls gives

$$w(x) \leq u(x).$$

We're essentially done. We cannot set w equal to the distance function itself, because this choice isn't smooth enough. However we can choose w to be a slightly smoothed-out version of the distance function minus a small constant. It's easy to see that we can approach the distance function from below by such functions w. Therefore (using these w's and passing to a limit)

$$\text{dist}\,(x, \partial\Omega) \leq u(x),$$

completing our PDE argument that the value function is in this case the distance function.

It's time for a more financial example. Let's give the **solution to Example 1** for a power-law utility. The state equation is

$$\dot{y} = ry - \alpha, \quad y(t) = x$$

where x is the initial wealth and α is the consumption rate, restricted by $\alpha \geq 0$ (an explicit constraint on the controls). We consider the problem of finding

$$u(x,t) = \max_{\alpha} \int_t^T e^{-\rho s} \alpha^\gamma(s)\, ds,$$

which amounts to the utility of consumption with the power-law utility function $h(\alpha) = \alpha^\gamma$. Utility functions should be concave so we assume $0 < \gamma < 1$.

First, before doing any real work, let us show that the value function has the form

$$u(x,t) = g(t)x^\gamma$$

for some function $g(t)$. It suffices for this purpose to show that the value function has the homogeneity property

$$u(\lambda x, t) = \lambda^\gamma u(x,t), \tag{12}$$

for then we can take $g(t) = u(1,t)$. To see (12), suppose $\alpha(s)$ is optimal for starting point x, and let $y_x(s)$ be the resulting trajectory. We may consider the control $\lambda\alpha(s)$ for the trajectory that starts at λx, and it is easy to see that the associated trajectory is $y_{\lambda x}(s) = \lambda y_x(s)$. Using the power-law form of the utility this comparison demonstrates that

$$u(\lambda x, t) \geq \lambda^\gamma u(x,t).$$

This relation with λ replaced by $1/\lambda$ and x replaced by λx gives

$$u(x,t) \geq \lambda^{-\gamma} u(\lambda x, t),$$

completing the proof of (12).

Now let's find the HJB equation. This is almost a matter of specializing the general calculation to the case at hand. But we didn't have a discount term before, so let's redo the argument to avoid any doubt. From the dynamic programming principle we have

$$u(x,t) \geq e^{-\rho t} a^\gamma \Delta t + u(x + (rx - a)\Delta t, t + \Delta t) + \text{error terms}$$

with equality when $a \geq 0$ is chosen optimally. Using the first-order Taylor expansion of u this becomes

$$u(x,t) \geq e^{-\rho t} a^\gamma \Delta t + u(x,t) + (u_x(rx - a) + u_t)\Delta t + \text{error terms}$$

with equality when a is optimal. In the limit $\Delta t \to 0$ this gives

$$u_t + \max_{a \geq 0} \{u_x(rx - a) + e^{-\rho t} a^\gamma\} = 0.$$

This is the desired HJB equation, to be solved for $t < T$. The final-time condition is of course $u = 0$ (since no utility is associated to final-time wealth).

It's obvious that $u_x > 0$. (This follows from the observation that $u(x,t) = g(t)x^\gamma$. Or it's easy to prove using the original problem formulation and a suitable comparison argument.) Therefore the optimal a is easy to find, by differentiation, and it is positive:

$$\gamma a^{\gamma - 1} = e^{\rho t} u_x.$$

This is the feedback law, determining the optimal control (once we know u_x). Remembering that $u(x,t) = g(t)x^\gamma$, we can write the feedback law as

$$a = \left[e^{\rho t} g(t)\right]^{1/(\gamma - 1)} x$$

To find g (and therefore u) we substitute $u = g(t)x^\gamma$ into the HJB equation. This leads, after some arithmetic and cancelling a common factor of x^γ from all terms, to

$$\frac{dg}{dt} + r\gamma g + (1 - \gamma) g (e^{\rho t} g)^{1/(\gamma - 1)}.$$

This equation (with the end condition $g(T) = 0$) is entirely equivalent to the original HJB equation. It looks ugly, however it is not difficult to solve. First, multiply each term by $e^{\rho t}$ to see that $G(t) = e^{\rho t} g(t)$ solves

$$G_t + (r\gamma - \rho)G + (1 - \gamma) G^{\gamma/(\gamma - 1)} = 0.$$

Next, multiply by $(1-\gamma)^{-1} G^{\gamma/(1-\gamma)}$ to see that $H(t) = G^{1/(1-\gamma)}$ satisfies the linear equation

$$H_t - \mu H + 1 = 0 \quad \text{with } \mu = \tfrac{\rho - r\gamma}{1 - \gamma}.$$

This is a linear equation! The solution satisfying $H(T) = 0$ is

$$H(t) = \lambda^{-1} \left(1 - e^{-\lambda(T-t)}\right).$$

Unraveling our changes of variables gives finally

$$g(t) = e^{-\rho t}\left[\frac{1-\gamma}{\rho - r\gamma}\left(1 - e^{-\frac{(\rho - r\gamma)(T-t)}{1-\gamma}}\right)\right]^{1-\gamma}.$$

We've solved the HJB equation. Have we actually found the value function and the optimal feedback (consumption) policy? Yes indeed. The verification theorem given above supplies the proof. (Well, it should be redone with discounting, and with the more precise formulation of the objective which integrates the utility only up to the first time τ when $y = 0$, if this occurs before T. These modifications require no really new ideas.) Nothing fancy is needed since $u(x, t)$ is smooth.

PDE for Finance Notes, Spring 2003 – Section 5

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, G63.2706, Spring 2000.

Stochastic optimal control. Stochastic optimal control is like deterministic optimal control except that (i) the equation of state is a *stochastic* differential equation, and (ii) the goal is to maximize or minimize the *expected* utility or cost. To see the structure of the theory in a simple, uncluttered way, we begin by examining what becomes of a standard deterministic utility maximization problem when the state equation is perturbed by a little noise. Then we present a finance classic: Merton's analysis of optimal consumption and investment, in the simplest meaningful case (a single risk-free asset and a risk-free account). Our treatment follows more or less Chapter VI of Fleming and Rishel. For those who wish to learn more, Merton's original paper is highly recommended; it is, in my view, easier to read than Fleming and Rishel. The reference is: "Optimal consumption and portfolio rules in a continuous-time model", *J. Economic Theory* 3, 1971, 373-413 (reprinted as Chapter 5 of his book *Continuous Time Finance*).

****

Perturbation of a deterministic problem by small noise. We've discussed at length the deterministic dynamic programming problem with state equation

$$dy/ds = f(y(s), \alpha(s)) \text{ for } t < s < T, \quad y(t) = x,$$

controls $\alpha(s) \in A$, and objective

$$\max_\alpha \left\{ \int_t^T h(y(s), \alpha(s))\, ds + g(y(T)) \right\}.$$

Its value function satisfies the HJB equation

$$u_t + H(\nabla u, x) = 0 \text{ for } t < T, \quad u(x,T) = g(x),$$

with Hamiltonian

$$H(p,x) = \max_{a \in A} \{ f(x,a) \cdot p + h(x,a) \}. \tag{1}$$

Let us show (heuristically) that when the state is perturbed by a little noise, the value function of resulting stochastic control problem solves the perturbed HJB equation

$$u_t + H(\nabla u, x) + \frac{1}{2}\epsilon^2 \Delta u = 0 \tag{2}$$

where H is still given by (1), and $\Delta u = \sum_i \frac{\partial^2 u}{\partial x_i^2}$.

Our phrase "perturbing the state by a little noise" means this: we replace the ODE governing the state by the stochastic differential equation (SDE)

$$dy = f(y, \alpha) ds + \epsilon dw,$$

keeping the initial condition $y(t) = x$. Here dw is a standard, vector-valued Brownian motion (each component w_i is a scalar-valued Brownian motion, and different components are independent).

The evolution of the state is now stochastic, hence so is the value of the utility. Our goal in the stochastic setting is to maximize the *expected* utility. The value function is thus

$$u(x,t) = \max_{\alpha} E_{y(t)=x}\left\{\int_t^T h(y(s),\alpha(s))\,ds + g(y(T))\right\}.$$

There is some subtlety to the question: what is the class of admissible controls? Of course we still restrict $\alpha(s) \in A$. But since the state is random, it's natural for the control to be random as well – however its value at time s should depend only on the past and present, not on the future (which is after all unknown to the controller). Such controls are called "non-anticipating." A simpler notion, sufficient for most purposes, is to restrict attention to *feedback* controls, i.e. to assume that $\alpha(s)$ is a deterministic function of s and $y(s)$. One can show (under suitable hypotheses, when the state equation is a stochastic differential equation) that these two different notions of "admissible control" lead to the same optimal value.

Courage. Let's look for the HJB by applying the usual heuristic argument, based on the principle of dynamic programming applied to a short time interval:

$$u(x,t) \approx \max_{a\in A}\left\{h(x,a)\Delta t + E_{y(t)=x}u(y(t+\Delta t), t+\Delta t)\right\}.$$

The term $h(x,a)\Delta t$ approximates $\int_t^{t+\Delta t} h(y(s),a)\,ds$, because we assume h is smooth and $y(s) = x +$ terms tending to 0 with Δt. Notice that $h(x,a)\Delta t$ is deterministic. The expression $u(y(t+\Delta t), t+\Delta t)$ is the optimal expected utility starting from time $t+\Delta t$ and spatial point $y(t+\Delta t)$. We must take its expected value, because $y(t+\Delta t)$ is random. (If you think carefully you'll see that the Markov property of the process $y(s)$ is being used here.)

We're almost in familiar territory. In the deterministic case the next step was to express $u(y(t+\Delta t), t+\Delta t)$ using the state equation and the Taylor expansion of u. Here we do something analogous: use Ito's lemma and the stochastic differential equation. Ito's lemma says the process $\phi(s) = u(y(s), s)$ satisfies

$$\begin{aligned} d\phi &= \frac{\partial u}{\partial s}ds + \sum_i \frac{\partial u}{\partial y_i}dy_i + \frac{1}{2}\sum_{i,j}\frac{\partial^2 u}{\partial y_i \partial y_j}dy_i dy_j \\ &= u_t(y(s),s)ds + \nabla u\cdot(f(y(s),\alpha(s))ds + \epsilon dw) + \frac{1}{2}\epsilon^2\Delta u\,ds. \end{aligned}$$

The real meaning of this statement is that

$$\begin{aligned} u(y(t'),t') - u(y(t),t) &= \int_t^{t'}[u_t(y(s),s) + \nabla u(y(s),s)\cdot(f(y(s),\alpha(s)) + \frac{1}{2}\epsilon^2\Delta u(y(s),s)]ds \\ &\quad + \int_t^{t'}\epsilon\nabla u(y(s),s)\cdot dw. \end{aligned}$$

The expected value of the second integral is 0, so

$$E_{y(t)=x}[u(y(t+\Delta t),t+\Delta t)] - u(x,t) \approx [u_t(x,t) + \nabla u(x,t)\cdot f(x,a) + \frac{1}{2}\epsilon^2 \Delta u(x,t)]\Delta t.$$

Assembling these ingredients, we have

$$u(x,t) \approx \max_{a\in A}\left\{h(x,a)\Delta t + u(x,t) + [u_t(x,t) + \nabla u(x,t)\cdot f(x,a) + \frac{1}{2}\epsilon^2\Delta u(x,t)]\Delta t\right\}.$$

This is almost identical to the relation we got in the deterministic case. The only difference is the new term $\frac{1}{2}\epsilon^2\Delta u(x,t)\Delta t$ on the right. It doesn't depend on a, so the optimal a is unchanged – it still maximizes $h(x,a) + f(x,a)\cdot\nabla u$ – and we conclude, as asserted, that u solves (2).

Before going to another topic, let's link this discussion to the notion of "viscosity solution." We noted in Section 4 that the solution of the deterministic HJB equation can be nonunique. (For example, our geometric Example 2 has the HJB equation $|\nabla u| = 1$ with boundary condition $u = 0$ at the target; it clearly has many almost-everywhere solutions, none of them smooth). We also mentioned in Section 4 that this difficulty can be resolved by working with the "viscosity solution." One characterization of the viscosity solution is this: it is the solution obtained by including a little noise in the problem formulation (with variance ϵ, as above), then taking the limit $\epsilon \to 0$.

****

Optimal portfolio selection and consumption. This is the simplest of a class of problems solved by Robert Merton in his paper "Optimal consumption and portfolio rules in a continuous-time model", *J. Economic Theory* 3, 1971, 373-413 (reprinted in his book *Continuous Time Finance.*) As you'll see, the math is almost the same as our Example 1 – though the finance is more interesting.

We consider a world with one risky asset and one risk-free asset. The risk-free asset grows at a constant risk-free rate r, i.e. its price per share satisfies $dp_1/dt = p_1 r$. The risky asset executes a geometric Brownian motion with constant drift $\mu > r$ and volatility σ, i.e. its price per share solves the stochastic differential equation $dp_2 = \mu p_2 dt + \sigma p_2 dw$.

The control problem is this: an investor starts with initial wealth x at time t. His control variables are

$$\begin{aligned} \alpha_1(s) &= \text{fraction of total wealth invested in the risky asset at time } s \\ \alpha_2(s) &= \text{rate of consumption at time } s. \end{aligned}$$

It is natural to restrict these controls by $0 \le \alpha_1(s) \le 1$ and $\alpha_2(s) \ge 0$. We ignore transaction costs. The state is the investor's total wealth y as a function of time; it solves

$$dy = (1-\alpha_1)yr dt + \alpha_1 y(\mu dt + \sigma dw) - \alpha_2 dt$$

so long as $y(s) > 0$. We denote by τ the first time $y(s) = 0$ if this occurs before time T, or $\tau = T$ (a fixed horizon time) otherwise. The investor seeks to maximize the discounted total utility of his consumption. We therefore consider the value function

$$u(x,t) = \max_{\alpha_1,\alpha_2} E_{y(t)=x} \int_t^\tau e^{-\rho s} h[\alpha_2(s)]\, ds$$

where $h[\cdot]$ is a specified utility function (monotone increasing and concave, with $h(0) = 0$). We shall specialize eventually to the power-law utility $h(\alpha_2) = \alpha_2^\gamma$ with $0 < \gamma < 1$. (We have chosen, as in Example 1, to work with the utility discounted to time 0. It is also possible, as in HW1, to work with the utility discounted to time t. The latter choice would give an autonomous HJB equation, i.e. time would not appear explicitly in the equation.)

We find the HJB equation by essentially the same method used above. The principle of dynamic programming applied on a short time interval gives:

$$u(x,t) \approx \max_{a_1,a_2} \left\{ e^{-\rho t} h(a_2)\Delta t + E_{y(t)=x} u(y(t+\Delta t), t+\Delta t) \right\}.$$

To evaluate the expectation term, we use Ito's lemma again. Using the state equation

$$dy = [(1-\alpha_1)yr + \alpha_1 y\mu - \alpha_2]dt + \alpha_1 y\sigma dw$$

and skipping straight to the conclusion, we have

$$u(y(t'),t') - u(y(t),t) = \int_t^{t'} [u_t + u_y[(1-\alpha_1)yr + \alpha_1 y\mu - \alpha_2] + \frac{1}{2}u_{yy}y^2\alpha_1^2\sigma^2]dt + \int_t^{t'} \alpha_1\sigma y u_y dw.$$

The expected value of the second integral is 0, so

$$E_{y(t)=x}[u(y(t+\Delta t), t+\Delta t)] - u(x,t) \approx [u_t + u_y[(1-\alpha_1)yr + \alpha_1 y\mu - \alpha_2 + \frac{1}{2}u_{yy}y^2\alpha_1^2\sigma^2]\Delta t.$$

Assembling these ingredients,

$$u(x,t) \approx \max_{a_1,a_2} \left\{ e^{-\rho t} h(a_2)\Delta t + u(x,t) + [u_t + u_x[(1-a_1)xr + a_1 x\mu - a_2] + \frac{1}{2}u_{xx}x^2 a_1^2\sigma^2]\Delta t \right\}.$$

Cleaning up, and taking the limit $\Delta t \to 0$, we get

$$u_t + \max_{a_1,a_2} \left\{ e^{-\rho t} h(a_2) + [(1-a_1)xr + a_1 x\mu - a_2]u_x + \frac{1}{2}x^2 a_1^2\sigma^2 u_{xx} \right\} = 0.$$

This is the relevant HJB equation. It is to be solved for $t < T$, with $u(x,T) = 0$ since we have associated no utility associated to final-time wealth.

That looks pretty horrible, but it isn't really so bad. First of all, if we constrain $a_1 = 0$ it reduces to the HJB equation from Example 1. (Well, it has to: if $a_1 = 0$ then all investment is in the risk-free asset, and the problem *is* Example 1.) So we charge ahead.

Let us assume $u_x > 0$ (practically obvious: larger initial wealth should produce larger total utility; what comparison argument would you use to prove it?). Let's also assume $u_{xx} < 0$ (not quite so obvious: this reflects the concavity of the utility function; it will be easy to

check it on our explicit solution at the end). Then the optimal a_1 (ignoring the constraint $0 \le a_1 \le 1$) is

$$a_1^* = -\frac{(\mu - r)u_x}{\sigma^2 x u_{xx}}$$

which is positive. We proceed, postponing till later the verification that $a_1^* \le 1$. The optimal a_2 satisfies

$$h'(a_2^*) = e^{\rho t} u_x;$$

we can be sure this a_2^* is positive by assuming that $h'(0) = \infty$.

To go further we now specialize to the power-law utility $h(a_2) = a_2^\gamma$ with $0 < \gamma < 1$. The same argument we used in the deterministic case shows that the solution must have the form

$$u(x,t) = g(t)x^\gamma.$$

The associated a_1^* and a_2^* are evidently

$$a_1^* = \frac{(\mu - r)}{\sigma^2(1-\gamma)}, \quad a_2^* = \left[e^{\rho t} g(t)\right]^{1/(\gamma-1)} x.$$

We assume henceforth that $\mu - r < \sigma^2(1-\gamma)$ so that $a_1^* < 1$. Substituting these values into the HJB equation gives, after some arithmetic,

$$\frac{dg}{dt} + \nu\gamma g + (1-\gamma) g (e^{\rho t} g)^{1/(\gamma-1)} = 0$$

with

$$\nu = r + \frac{(\mu - r)^2}{2\sigma^2(1-\gamma)}.$$

We must solve this with $g(T) = 0$. This is the same nonlinear equation we dealt with in Example 1 – with ν in place of r. So we can go straight to the answer: $u = g(t)x^\gamma$ with

$$g(t) = e^{-\rho t}\left[\frac{1-\gamma}{\rho - \nu\gamma}\left(1 - e^{-\frac{(\rho - \nu\gamma)(T-t)}{1-\gamma}}\right)\right]^{1-\gamma}.$$

It should not be surprising that we had to place some restrictions on the parameters to get this solution. When these restrictions fail, inequalities that previously didn't bother us become important (namely the restrictions $0 \le a_1 \le 1$, which prohibit borrowing and short-selling).

We have solved the HJB equation; but have we found the value function? The answer is yes, as we now show using a verification argument.

****

The verification argument. In the deterministic case we used a heuristic argument to derive the HJB equation, but then showed completely honestly that a (sufficiently smooth) solution of the HJB equation (satisfying appropriate boundary or final-time conditions) provides a bound on the value attainable by any control. A similar result holds in the stochastic setting.

Rather than give a general result at this time, let's focus on the example just completed (Merton's optimal selection and consumption problem). All the ideas required for the general case are already present here. Brief review of our task: the state equation is

$$dy = [(1-\alpha_1)yr + \alpha_1 y\mu - \alpha_2]dt + \alpha_1 y\sigma dw$$

which we shall write for simplicity as

$$dy = f(y, \alpha_1, \alpha_2)dt + \alpha_1 y\sigma dw.$$

The value function is

$$u(x,t) = \max_{\alpha} E_{y(t)=x} \int_t^\tau e^{-\rho s} h[\alpha_2(s)]\, ds$$

where τ is either the first time $y = 0$ (if this happens before time T) or $\tau = T$ (if y doesn't reach 0 before time T). The HJB equation is

$$v_t + \max_{a_1,a_2} \left\{ e^{-\rho t} h(a_2) + f(x, a_1, a_2) v_x + \frac{1}{2} x^2 a_1^2 \sigma^2 v_{xx} \right\} = 0$$

for $t < T$, with $v = 0$ at $t = T$. We didn't fuss over it before, but clearly v should also satisfy $v(0, s) = 0$ for all s. We write v instead of u, to reserve notation u for the optimal value. The goal of the verification argument is to show that $v \geq u$, i.e. to show that no control strategy can achieve an expected discounted utility better than v. Our argument will also show that the feedback strategy associated with the HJB calculation – namely

$$\alpha_1(s) = -\frac{(\mu - r)v_x}{\sigma^2 x v_{xx}}(y(s), s), \quad h'(\alpha_2)(s) = e^{\rho s} v_x(y(s), s) \tag{3}$$

does indeed achieve expected discounted value v; in other words $v \leq u$. This suffices of course to show $v = u$.

Consider any control $\tilde{\alpha}(s)$, and the associated state $\tilde{y}(s)$ starting from $\tilde{y}(t) = x$. Of course we assume $\tilde{\alpha}$ is non-anticipating, i.e. it depends only on knowledge of $\tilde{y}(s)$ in the present and past, not the future. (If this condition confuses you, just assume $\tilde{\alpha}$ is given by a feedback law, i.e. $\tilde{\alpha}(s) = F(y(s), s)$ for some deterministic function $F(y, s)$. Such controls are automatically non-anticipating.) We wish to show that

$$v(x,t) \geq E_{y(t)=x} \int_t^{\tilde{\tau}} e^{-\rho s} h[\tilde{\alpha}_2(s)]\, ds.$$

Consider $\phi(s) = v(\tilde{y}(s), s))$: by the Ito calculus it satisfies

$$\begin{aligned} d\phi &= v_s ds + v_y d\tilde{y} + \frac{1}{2} v_{yy} d\tilde{y} d\tilde{y} \\ &= v_s ds + v_y [f(\tilde{\alpha}, \tilde{y}) ds + \tilde{\alpha}_1(s)\tilde{y}(s)\sigma dw] + \frac{1}{2} v_{yy} \tilde{\alpha}_1^2(s) \tilde{y}^2(s) \sigma^2 ds. \end{aligned}$$

Therefore

$$v(\tilde{y}(t'),t') - v(\tilde{y}(t),t) = \int_t^{t'} [v_s + v_y f + \frac{1}{2} v_{yy} \tilde{y}^2 \tilde{\alpha}_1^2 \sigma^2] dt + \int_t^{t'} \sigma \tilde{\alpha}_1 \tilde{y} v_y dw$$

where each integrand is evaluated at $y = \tilde{y}(s)$, $\alpha = \tilde{\alpha}(s)$ at time s. The expected value of the second integral is 0 (here is where we use that α is nonanticipating; we will return to this when we discuss stochastic integrals). Thus taking the expectation, and using the initial condition:

$$E\left[v(\tilde{y}(t'),t')\right] - v(x,t) = E\left[\int_t^{t'} (v_s + v_y f + \frac{1}{2} v_{yy} \tilde{y}^2 \tilde{\alpha}_1^2 \sigma^2) dt\right].$$

Now from the definition of the Hamiltonian we have

$$v_t(\tilde{y}(s),s) + \left\{ e^{-\rho s} h(\tilde{\alpha}_2(s)) + f(\tilde{y}(s),\tilde{\alpha}(s)) v_y(\tilde{y}(s),s) + \frac{1}{2} \tilde{y}^2(s) \tilde{\alpha}_1^2(s) \sigma^2 v_{yy}(\tilde{y}(s),s) \right\} \leq 0. \tag{4}$$

Combining this with the preceding relation gives

$$E\left[v(\tilde{y}(t'),t')\right] - v(x,t) \leq -E\left[\int_t^{t'} e^{-\rho s} h(\tilde{\alpha}_2(s) ds\right].$$

Taking $t' = \tilde{\tau}$ and using the fact that $v(\tilde{y}(t'),t') = 0$, we conclude that

$$v(x,t) \geq E\left[\int_t^{\tilde{\tau}} e^{-\rho s} h(\tilde{\alpha}(s) ds\right].$$

Maximizing the right hand side over all $\tilde{\alpha}$ we conclude that

$$v \geq u$$

For the special feedback law associated with the HJB equation, which fixes the control α by (3), relation (4) becomes equality. This shows that $v \leq u$, since v is the value achieved by a specific control strategy and u is the maximum value over all possible strategies. Thus $v = u$. In summary: the function v, defined by solving the HJB equation with appropriate boundary and initial conditions, is in fact the value function of this stochastic control problem, and the control strategy (3) is indeed optimal.

Notice that this calculation rests on pretty much the same tools we used to derive the HJB: (a) the Ito calculus, to get a representation of $u(\tilde{y}(s),s)$, and (b) the fact that any integral "dw" has expected value 0.

PDE for Finance Notes, Spring 2003 – Section 6

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, G63.2706.

Optimal stopping and American options. Optimal stopping refers to a special class of stochastic control problems where the only decision to be made is "when to stop." The decision *when to sell an asset* is one such problem. The decision *when to exercise an American option* is another. Mathematically, such a problem involves optimizing the expected payoff over a suitable class of stopping times. The value function satisfies a "free boundary problem" for the backward Kolmogorov equation.

We shall concentrate on some simple yet representative examples which display the main ideas, namely: (a) a specific optimal stopping problem for Brownian motion; (b) when to sell a stock which undergoes log-normal price dynamics; and (c) the pricing of a perpetual American option. At the end we discuss how the same ideas apply to the pricing of an American option with a specified maturity. My discussion of (a) borrows from Raghu Varadhan's PDE for Finance notes; my discussion of (b) is similar to Oksendal's Examples 10.2.2 and 10.4.2; the discussion of (c) can be found in many places, e.g. Ingersoll's book, or Wilmott's.

****

Optimal stopping for 1D Brownian motion. Let $y(t)$ be 1D Brownian motion starting from $y(0) = x$. For any function f, we can consider the simple optimal stopping problem

$$u(x) = \max_{\tau} E_{y(0)=x} \left[e^{-\tau} f(y(\tau)) \right].$$

Here τ varies over all *stopping times*. We have set the discount rate to 1 for simplicity. We first discuss some general principles then obtain an explicit solution when $f(x) = x^2$.

What do we expect? The x-axis should be divided into two sets, one where it is best to *stop immediately*, the other where it is best to *stop later*. For x in the stop-immediately region the value function is $u(x) = f(x)$ and the optimal stopping time is $\tau = 0$. For x in the stop-later region the value function solves a PDE. Indeed, for Δt sufficiently small (and assuming the optimal stopping time is larger than Δt)

$$u(x) \approx e^{-\Delta t} E_{y(0)=x} \left[u(y(\Delta t)) \right].$$

By Ito's formula

$$E_{y(0)=x}[u(y(t)] = u(x) + \int_0^t \tfrac{1}{2} u_{xx}(y(s))\, ds.$$

Applying this with $t = \Delta t$ and approximating the integral by $\frac{1}{2} u_{xx}(x) \Delta t$ we conclude that $u(x) \approx e^{-\Delta t}(u(x) + \frac{1}{2} u_{xx} \Delta t)$. As $\Delta t \to 0$ this gives the PDE in the stop-later region:

$$\tfrac{1}{2} u_{xx} - u = 0.$$

The preceding considerations tell us a little more. Stopping immediately is always a candidate strategy; so is waiting. So for every x we have

$$u(x) \geq f(x)$$

since this is the value obtained by exercising immediately. Also, $u(x) \geq \lim_{\Delta t \to 0} e^{-\Delta t}(u(x) + \frac{1}{2}u_{xx}(x)\Delta t)$, since this is the value obtained by waiting a little. Evaluating the limit gives

$$\tfrac{1}{2}u_{xx} - u \leq 0.$$

These considerations restrict the location of the *free boundary* separating the stop-now and stop-later regions; in particular, we must have $\frac{1}{2}f_{xx} - f \leq 0$ everywhere in the stop-now region, since there we have both $u = f$ and $\frac{1}{2}u_{xx} - u \leq 0$.

To specify the free boundary fully, however, we need a more subtle condition, the *high-order contact condition*: the value function is C^1 at the free boundary. In other words the value of u_x at the free boundary is the same whether you approach it from the stop-immediately side (where $u_x = f_x$) or from the stop-later side (where u_x is determined by the PDE). In truth we used this property above, when we applied Ito's Lemma to $u(y(t))$ (the usual proof of Ito's Lemma assumes u is C^2, but a more careful argument shows that it applies even if u_{xx} is discountinuous across a point, provided u_x is continuous). The rationale behind the high-order contact condition is easiest to explain a little later, in the context of Example 1.

Example 1. Let us obtain an explicit solution when the payoff is $f(x) = x^2$. It is natural to guess that the free boundary is symmetric, i.e. it lies at $x = \pm a_*$ for some a. If so, then the optimal strategy is this: if $y(0) = x$ satisfies $|x| < a_*$, stop at the first time when $|y(t)| = a_*$; if on the other hand $y(0) = x$ has $|x| \geq a_*$ then stop immediately. We will find the value of a_* and prove this guess is right. Notice that we know $a_* \geq 1$, since $\frac{1}{2}f_{xx} - f = 1 - x^2 \leq 0$ everywhere the stop-immediately region.

Consider the strategy described above, with any choice of a. The value $u_a(x)$ associated with this strategy is easy to evaluate: by the argument used above (or remembering Section 1, i.e. using the Feynman-Kac formula with stopping), it satisfies $\frac{1}{2}u_a'' - u = 0$ on the interval $-a < x < a$ with boundary condition $u_a = a^2$ at $x = \pm a$. This can be solved explicitly: the general solution of $\frac{1}{2}v'' - v = 0$ is $v = c_1 e^{\sqrt{2}x} + c_2 e^{-\sqrt{2}x}$. Using the boundary conditions to determine c_1 and c_2 gives

$$u_a(x) = a^2 \frac{\cosh \sqrt{2}x}{\cosh \sqrt{2}a}$$

for $|x| \leq a$. We use the high-order contact condition to determine a_*. It is the choice of a for which $u_a'(\pm a) = f'(\pm a) = \pm 2a$. This amounts to

$$a^2 \sqrt{2} \frac{\sinh \sqrt{2}a}{\cosh \sqrt{2}a} = 2a$$

which simplifies to

$$\tanh \sqrt{2}a = \frac{\sqrt{2}}{a}.$$

This equation has two (symmetric) solutions, $\pm a_*$. Notice that since $|\tanh x| < 1$ we have $|a_*| > \sqrt{2} > 1$. The PDE defines u_a only for $|x| \le a$. For $|x| > a$ it is $u_a(x) = f(x) = x^2$, since this is the value associated with stopping immediately.

We promised to explain the high-order contact condition. Here is the logic behind it. If we assume the optimal strategy will be of the kind just considered for some a, then the value function must be

$$u(x) = \max_a u_a(x)$$

for all $|x| < |a_*|$. In particular, $\partial u_a(x)/\partial a = 0$ for all $|x| < |a_*|$. Taking $x = a$ and using chain rule, we get

$$\frac{\partial}{\partial a}[u_a(a)] = \left.\frac{\partial u_a(x)}{\partial a}\right|_{x=a} + \left.\frac{\partial u_a(x)}{\partial x}\right|_{x=a}$$

for $|a| < a_*$. The left hand side is $f'(a)$ for any a. In the limit as a approaches a_* the first term on the right is 0 and the second term on the right is $u'_{a_*}(a)$. Thus the high-order contact condition holds at the optimal a.

We noted as part of our general discussion that $u \ge f$ in the "stop-later" region. There is no further freedom – we have fully determined $u_{a_*}(x)$ – so this had better be satisfied, i.e. we want $u_{a_*} \ge f$ on the interval $[-a_*, a_*]$. Since the function is completely explicit, this relation is easily verified by direct calculation.

Let us finally prove our guess is right. The function $u = u_{a_*}$ is, by construction, the value function of an admissible strategy. We must show it is optimal, i.e. that for any stopping time τ

$$u_{a_*}(x) \ge E_{y(0)=x}\left[e^{-\tau} f(y(\tau)\right].$$

Applying Ito's formula to $\phi(t) = e^{-t}u_{a_*}(y(t))$ gives

$$d(e^{-t}u_{a_*}(y(t)) = e^{-t}u'_{a_*}dy + e^{-t}(\tfrac{1}{2}u''_{a_*} - u)dt$$

(we used here the fact that u_{a_*} is smooth away from $x = a_*$ and C^1 across $x = a_*$ so the use of Ito's lemma can be justifed). Integrating up to the stopping time and taking the expectation gives

$$E_{y(0)=x}\left[e^{-\tau}u_{a_*}(y(\tau)\right] - u_{a_*}(x) = E_{y(0)=x}\left[\int_0^\tau e^{-s}(\tfrac{1}{2}u''_{a_*} - u)(y(s))\,ds\right].$$

Since $\frac{1}{2}u''_{a_*} - u \le 0$ and $u_{a_*} \ge f$, this implies

$$E_{y(0)=x}\left[e^{-\tau} f(y(\tau)\right] - u_{a_*}(x) \le 0$$

which is the desired assertion.

****

Example 2: When to sell an asset. This problem is familiar to any investor: when to sell a stock you presently own? Keeping things simple (to permit a closed-form solution), we suppose the stock price executes geometric brownian motion

$$dy = \mu y ds + \sigma y dw$$

with constant μ and σ. Assume a fixed commission a is payable at the time of sale, and suppose the present value of future income is calculated using a constant discount rate r. Then the time-0 value realized by sale at time s is $e^{-rs}[y(s) - a]$. Our task is to choose the time of sale optimally, i.e. to find

$$u(x) = \max_{\tau} E_{y(0)=x} \left[e^{-r\tau}(y(\tau) - a) \right] \tag{1}$$

where the maximization is over all stopping times. This example differs from Example 1 in that (a) the underlying process is lognormal, and (b) the payoff is linear. The analysis is however parallel to that of Example 1.

It is natural to assume that $\mu < r$, and we shall do so. If $\mu > r$ then the maximum value of (1) is easily seen to be ∞; if $\mu = r$ then the maximum value (1) turns out to be xe^{-rt}. When $\mu \geq r$ there is no optimal stopping time – a sequence of better and better stopping times tends to ∞ instead of converging. (You will be asked to verify these assertions on HW5.)

Let's guess the form of the solution. Since the underlying is lognormal it stays positive. So the optimal strategy should be: sell when the underlying reaches a threshold h_*, with h_* depending only on the parameters of the problem, i.e. μ, σ, r, and a. The positive reals are divided into two regions: a "sell-later" region where $x < h_*$ and a "sell-now" region where $x > h_*$.

In the sell-now region clearly $u(x) = x - a$. In the sell-later region it satisfies the PDE

$$\tfrac{1}{2}\sigma^2 x^2 u_{xx} + \mu x u_x - ru = 0$$

with boundary condition $u(x) = x - a$ at $x = h_*$. Moreover we have the global inequalities

$$u(x) \geq x - a \quad \text{and} \quad \tfrac{1}{2}\sigma^2 x^2 + \mu x u_x - ru \leq 0$$

by the same arguments used earlier for Example 1.

To identify the optimal sales threshold h_*, we proceed as in Example 1. Consider any candidate threshold h. The associated value function u_h solves $\sigma^2 x^2 u_h'' + \mu x u_h' - r u_h = 0$ for $x < h$, with boundary condition $u_h = x - a$ at $x = h$. This can be solved explicitly. The general solution of $\sigma^2 x^2 \phi'' + \mu x \phi' - r\phi = 0$ is

$$\phi(x) = c_1 x^{\gamma_1} + c_2 x^{\gamma_2}$$

where c_1, c_2 are arbitrary constants and

$$\gamma_i = \sigma^{-2} \left[\tfrac{1}{2}\sigma^2 - \mu \pm \sqrt{(\mu - \tfrac{1}{2}\sigma^2)^2 + 2r\sigma^2} \right].$$

We label the exponents so that $\gamma_2 < 0 < \gamma_1$. To determine u_h we must specify c_1 and c_2. Since u_h should be bounded as $x \to 0$ we have $c_2 = 0$. The value of c_1 is determined by the boundary condition at $x = h$: evidently $c_1 = h^{-\gamma_1}(h - a)$. Thus the expected payoff using sales threshold h is

$$u_h(x) = \begin{cases} (h - a)\left(\frac{x}{h}\right)^{\gamma_1} & \text{if } x < h \\ (x - a) & \text{if } x > h. \end{cases}$$

In Example 1 we used the high-order contact condition to determine h_*, and we could do the same here. But for variety (and to gain intuition) let's maximize $u_h(x)$ over h instead. One verifies by direct calculation that the optimal h is

$$h_* = \frac{a\gamma_1}{\gamma_1 - 1}$$

(notice that $\gamma_1 > 1$ since $\mu < r$). Let's spend a moment visualizing the geometry underneath this optimization, which is shown in Figure 1. As an aid to visualization, suppose $\gamma_1 = 2$ (the general case is not fundamentally different, since $\gamma_1 > 1$). Then the graph of $x - a$ is a line, while the graph of $(h - a)(x/h)^2$ is a parabola. The two graphs meet when $x - a = (h - a)(x/h)^2$. This equation is quadratic in x, so it has two roots, $x = h$ and $x = ah/(h - a)$ — unless $h = 2a$, in which case the two roots coincide. The optimal choice $h = h_*$ is the one for which the roots coincide. Some consideration of the figure shows why: if $h < h_*$ then increasing h slightly raises the parabola and increases u_h; similarly if $h > h_*$ then decreasing h slightly raises the parabola and increases u^h.

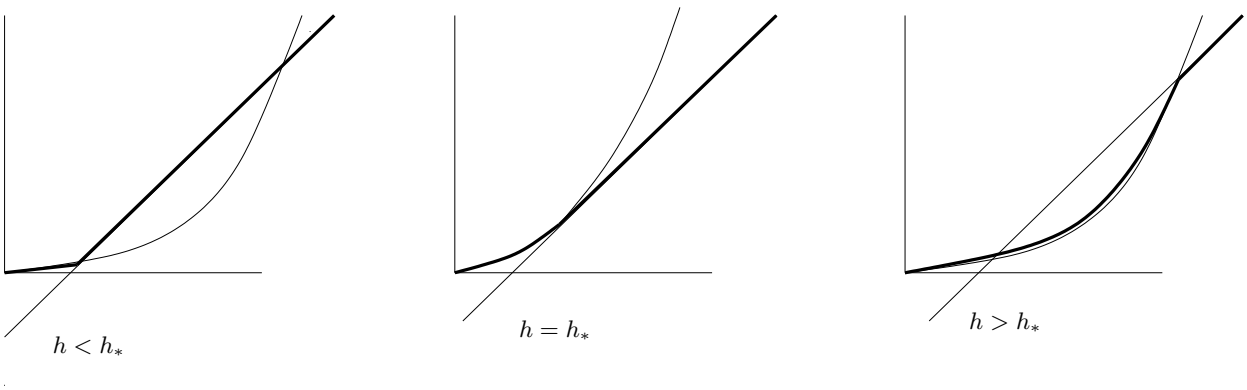

Figure 1: *Graph of* u_h.

Summing up (and returning to the general case, i.e. we no longer suppose $\gamma_1 = 2$): the optimal policy is to sell when the stock price reaches a certain threshold h_*, or immediately if the present price is greater than h_*; the value achieved by this policy is

$$u_{h_*}(x) = \max_h u_h(x,t) = \begin{cases} \left(\frac{\gamma_1 - 1}{a}\right)^{\gamma_1 - 1} \left(\frac{x}{\gamma_1}\right)^{\gamma_1} & \text{if } x < h_* \\ (x - a) & \text{if } x > h_*. \end{cases} \tag{2}$$

Our figure shows that the high-order-contact condition holds, i.e. u_{h_*} is C^1. In other words, while for general h the function u_h has a discontinuous derivative at h, the optimal h is also the choice that makes the derivative continuous there. This can of course be verified by direct calculation, and explained by the (actually quite general) argument presented in Example 1.

It remains to prove that our guess is right, i.e. to prove that this u_{h_*} achieves the optimal value among *all* sales strategies (stopping times). This is a verification argument, entirely parallel to that of Example 1; its details are left to the reader.

****

Example 3: the perpetual American put. An American option differs from a European one in the feature that it can be exercised at any time. Therefore the associated optimal stopping problem is to maximize the expected discounted value *at exercise*, over all possible exercise times. The decision whether to exercise or not should naturally depend only on present and past information, i.e. it must be given by a stopping time. Consider, to fix ideas, a put option with strike K (so the payoff is $(K-x)_+$), for a stock with lognormal dynamics $dy = \mu y ds + \sigma y dw$, and discount rate r. (For option pricing this should be the risk-neutral process not the subjective one. If the stock pays no dividends then $\mu = r$; if it pays continuous dividends at rate d then $\mu = r - d$.) To make maximum contact with the preceding two examples, we focus for now on a *perpetual* option, i.e. one that never matures. Then the holder decides his exercise strategy by solving the optimal control problem

$$u(x) = \max_{\tau} E_{y(0)=x}\left[e^{-r\tau}(K - y(\tau))_+\right]. \tag{3}$$

This problem differs from Example 2 only in having a different payoff. The method we used for Examples 1 and 2 works here too. Here is an outline of the solution:

- It is natural to guess that the optimal policy is determined by an exercise threshold h as follows: exercise immediately if the price is below h; continue to hold if the price is above h. Clearly we expect $h < K$ since it would be foolish to exercise when the option is worthless.

- For a given candidate value of h, we can easily evaluate the expected value associated with this strategy. It solves

  $$-ru_h + \mu x u_h' + \tfrac{1}{2}\sigma^2 x^2 u_h'' = 0 \quad \text{for } x > h$$

  and

  $$u_h(x) = (K - x) \quad \text{for } 0 < x \le h.$$

- To find u_h explicitly, recall that the general solution of the PDE was $c_1 x^{\gamma_1} + c_2 x^{\gamma_2}$ with $\gamma_2 < 0 < \gamma_1$ given by

  $$\gamma_i = \sigma^{-2}\left[\tfrac{1}{2}\sigma^2 - \mu \pm \sqrt{(\mu - \tfrac{1}{2}\sigma^2)^2 + 2r\sigma^2}\right].$$

  This time the relevant exponent is the negative one, γ_2, since it is clear that u_h should decay to 0 as $x \to \infty$. The constant c_2 is set by the boundary condition $u_h(h) = (K - h)$. Evidently

  $$u_h(x) = \begin{cases} (K-h)\left(\frac{x}{h}\right)^{\gamma_2} & \text{if } x > h \\ (K - x) & \text{if } x < h. \end{cases}$$

- The correct exercise threshold is obtained by either (i) imposing the high-order contact condition $u_h'(h) = -1$, or (ii) maximizing with respect to h. (The two procedures are equivalent, as shown above.) The optimal value is $h_* = \frac{K\gamma_2}{\gamma_2 - 1}$, which is less than K as expected.

- When $h = h_*$ the function $v = u_{h_*}$ satisfies satisfies

  (a) $v \geq (K - x)_+$ for all $x > 0$;

  (b) $\mathcal{L}v \leq 0$ for all $x > 0$

  (c) v is C^1 at $x = h_*$ and smooth everywhere else.

  (d) equality holds in (a) for $0 < x < h_*$ and in (b) for $x > h_*$

  where $\mathcal{L}v = -rv + \mu x v_x + \frac{1}{2}\sigma^2 x^2 v_{xx}$.

- Properties (a)-(d) imply, by the usual verification argument, that v is indeed optimal (i.e. no exercise policy can achieve a better discounted expected value).

****

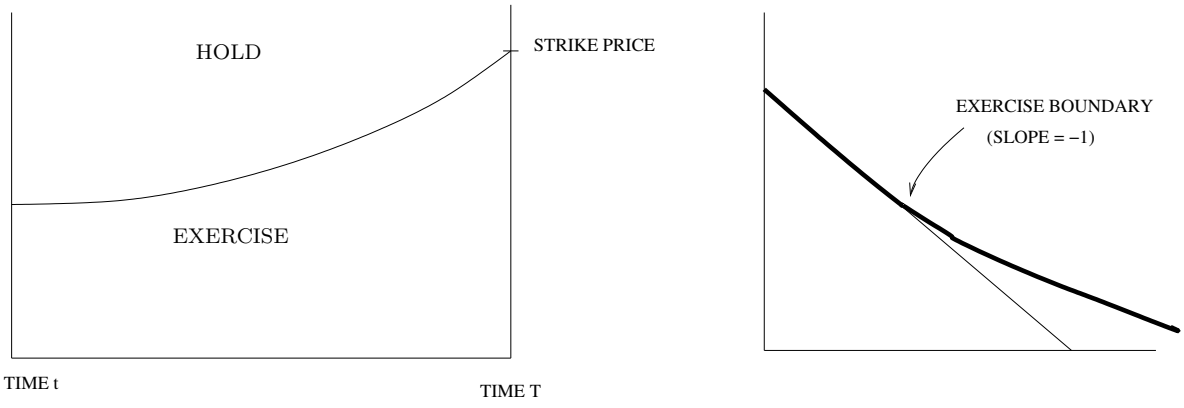

Figure 2: *The exercise boundary of an American option, and its value as a function of stock price at a given time t*

.

American options with finite maturity. What about American options with a specified maturity T? The same principles apply, though an explicit solution formula is no longer possible. The relevant optimal control problem is almost the same – the only difference is that the option must be exercised no later than time T. As a result the optimal value becomes a nontrivial function of the start time t:

$$u(x,t) = \max_{\tau \leq T} E_{y(t)=x}\left[e^{-r(\tau - t)}(K - y(\tau))_+\right].$$

The exercise threshold $h = h(t)$ is now a function of t: the associated policy is to exercise immediately if $x < h(t)$ and continue to hold if $x > h(t)$ (see Figure 2). It's clear, as before, that $h(t) < K$ for all t. Optimizing h is technically more difficult than in our previous examples because we must optimize over all *functions* $h(t)$. The most convenient characterization of the result is the associated variational inequality: the optimal exercise threshold $h(t)$ and the associated value function v satisfy

(a) $v \geq (K - x)_+$ for all $x > 0$ and all t;

(b) $v_t + \mathcal{L}v \leq 0$ for all $x > 0$ and all t;

(c) v is C^1 at $x = h(t)$ and smooth everywhere else.

(d) equality holds in (a) for $0 < x < h(t)$ and in (b) for $x > h(t)$

The proofs of (a) and (b) are elementary – using essentially the same ideas as in the Examples presented above. It is much more technical to prove that when h is optimized we get the high-order contact property (c); however the essential idea is the same as explained in Example 1. If you accept that (a)-(d) has a solution, its optimality is readily verified by the usual argument (modulo technicalities – mainly the validity of Ito's Lemma though v is not C^2 across the free boundary).

PDE for Finance Notes, Spring 2003 – Section 7

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, G63.2706.

About the final exam: Our exam is Monday May 12, 8-10pm, in the usual room Silver 207. Note the time-shift (8-10 not 7-9), intended to give students taking both Scientific Computing and PDE for Finance some breathing room. If this late start is a hardship for anyone, see me – it is possible by request to take the exam from 7-9pm instead of 8-10pm. The exam will be closed-book, but you may bring two sheets of notes (8.5×11, both sides, any font). The preparation such notes is an excellent study tool.

Addendum to Section 6: We solved various optimal stopping problems (Examples 1-3 in the Section 6 notes) by (a) guessing that the optimal strategy involved a certain threshold, then (b) choosing the threshold optimally, by direct maximization of the associated value or by applying the high-order-contact condition. We also gave (c) a verification argument, showing that the resulting value function was truly optimal – no stopping criterion could do better. There are two subtleties to the verification argument. One was noted – it requires applying Ito's formula to a function that's only piecewise smooth; this can be justified since the value function is C^1. The other subtlety was however entirely swept under the rug: we used the assertion that $E[\int_0^\tau g\,dw] = 0$. But this is clear in general only for *bounded* stopping times; we saw in HW1 that it can fail for an unbounded stopping time. How to get around this? Here's the answer: for any (possibly unbounded) stopping time τ, consider the truncated stopping times $\tau_k = \min\{\tau, k\}$. Clearly $\tau_k \to \tau$ as $k \to \infty$. Since τ_k is bounded, there's no problem applying the verification argument to it. In the context of Example 1 of Section 6, for example, this gives

$$E_{y(0)=x}\left[e^{-\tau_k} f(y(\tau_k)\right] \le u_{a_*}(x).$$

The limit $k \to \infty$ is handled by Fatou's lemma from real variables. It tells us that

$$E_{y(0)=x}\left[\liminf_{k\to\infty} e^{-\tau_k} f(y(\tau_k)\right] \le \liminf_{k\to\infty} E_{y(0)=x}\left[e^{-\tau_k} f(y(\tau_k)\right]$$

provided the payoff f is bounded below. In the present setting $\liminf_k e^{-\tau_k} f(y(\tau_k) = e^{-\tau} f(y(\tau))$, so these relations combine to give

$$E_{y(0)=x}\left[e^{-\tau} f(y(\tau))\right] \le u_{a_*}(x)$$

as desired.

Recommended reading: Merton and beyond. In Section 5, HW4, and HW5 you've been exposed to Merton's work applying dynamic programming to (i) portfolio optimization, and and (ii) the selection of consumption rates. Merton went much further than our treatment of course, and his articles are a pleasure to read; they are reprinted (with some updating) in Robert C. Merton, *Continuous Time Finance*, Blackwell, 1992, chapters 4 and 5. Research continues on closely related issues, for example: (a) the analogue of Merton's analysis in the presence of transaction costs [see e.g. M.H.A. Davis and A.R. Norman, *Portfolio selection with transaction costs*, Math. of Operations Research 15 (1990) 676-713]; and

(b) optimal pricing pricing and hedging of options when the market is incomplete, or the underlying is not a tradeable [see e.g. T. Zariphopoulou, *A solution approach to valuation with unhedgeable risks*, Finance and Stochastics 5 (2001) 61-82, and V. Henderson, *Valuation of claims on nontraded assets using utility maximization*, Mathematical Finance 12 (2002) 351-373].

****

Discrete-time dynamic programming. This section achieves two goals at once. One is to demonstrate the utility of discrete-time dynamic programming as a flexible tool for decision-making in the presence of uncertainty. The second is to introduce some more modern financially-relevant issues. To achieve these goals we shall discuss three specific examples: (1) optimal control of execution costs (following a paper by Bertsimas and Lo); (2) a discrete-time version of when to sell an asset (following Bertsekas' book); and (3) least-square replication of a European option (following a paper by Bertsimas, Kogan, and Lo).

In the context of this course it was natural to address continuous-time problems first, because we began the semester with stochastic differential equations and their relation to PDE's. Most courses on optimal control would however discuss the discrete-time setting first, because it is in many ways easier and more flexible. Indeed, continuous-time dynamic programming uses stochastic differential equations, Ito's formula, and the HJB equation. Discrete-time dynamic programming on the other hand uses little more than basic probability and the viewpoint of dynamic programming. Of course many problems have both discrete and continuous-time versions, and it is often enlightening to consider both (or compare the two). A general discussion of the discrete-time setting, with many examples, can be found in Dimitri Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Prentice-Hall, 1987 (on reserve), especially Chapter 2. Our approach here is different: we shall explain the method by presenting a few financially-relevant examples.

Example 1: Optimal control of execution costs. This example is taken from the recent article: Dimitris Bertsimas and Andrew Lo, *Optimal control of execution costs*, *J. Financial Markets* 1 (1998) 1-50. You can download a copy from the site www.sciencedirect.com (this works from the nyu.edu domain; to do it from outside NYU, see the Bobst Library databases web page for instructions how to set your browswer to use the NYU proxy server).

The problem is this: an investor wants to buy a large amount of some specific stock. If he buys it all at once he'll drive the price up, thereby paying much more than necessary. Better to buy part of the stock today, part tomorrow, part the day after tomorrow, etc. until the full amount is in hand. But how best to break it up?

Here's a primitive model. It's easy to criticize (we'll do this below), but it's a good starting point – and an especially transparent example of stochastic optimal control. Suppose the investor wants to buy $S_{\rm tot}$ shares of stock over a period of N days. His control variable is S_i, the number of shares bought on day i. Obviously we require $S_1 + \ldots + S_N = S_{\rm tot}$.

We need a model for the impact of the investor's purchases on the market. Here's where this model is truly primitive: we suppose that the price P_i the investor achieves on day i is related to the price P_{i-1} on day $i-1$ by

$$P_i = P_{i-1} + \theta S_i + \sigma e_i \tag{1}$$

where e_i is a Gaussian random variable with mean 0 and variance 1 (independent of S_i and P_{i-1}). Here θ and σ are fixed constants.

And we need a goal. Following Bertsimas and Lo we focus on minimizing the expected total cost:

$$\min E\left[\sum_{i=1}^{N} P_i S_i\right].$$

To set this up as a dynamic programming problem, we must identify the *state*. There is a bit of art here: the principle of dynamic programming requires that we be prepared to start the optimization at any day $i = N, N-1, N-2, \ldots$ and when $i = 1$ we get the problem at hand. Not so hard here: the state on day i is described by the most recent price P_{i-1} and the amount of stock yet to be purchased $W_i = S_{\text{tot}} - S_1 - \ldots - S_{i-1}$. The state equation is easy: P_i evolves as specified above, and W_i evolves by

$$W_{i+1} = W_i - S_i.$$

Dynamic programming finds the optimal control by starting at day N, and working backward one day at a time. The relation that permits us to work backward is the one-time-step version of the principle of dynamic programming. In this case it says:

$$V_i(P_{i-1}, W_i) = \min_s E\left[P_i s + V_{i+1}(P_i, W_{i+1})\right].$$

Here $V_i(P, W)$ is the value function:

$$\begin{aligned} V_i(P, W) &= \text{optimal expected cost of purchasing } W \text{ shares} \\ &\quad \text{starting on day } i, \text{ if the most recent price was } P. \end{aligned}$$

(The subscript i plays the role of time.)

To find the solution, we begin by finding $V_N(P, W)$. Since $i = N$ the investor has no choice but to buy the entire lot of W shares, and his price is $P_N = P + \theta W + e_N$, so his expected cost is

$$V_N(P, W) = E\left[(P + \theta W + \sigma e_N)W\right] = PW + \theta W^2.$$

Next let's find $V_{N-1}(P, W)$. The dynamic programming principle gives

$$\begin{aligned} V_{N-1}(P, W) &= \min_s E\left[(P + \theta s + \sigma e_{N-1})s + V_N(P + \theta s + \sigma e_{N-1}, W - s)\right] \\ &= \min_s E\left[(P + \theta s + \sigma e_{N-1})s + (P + \theta s + \sigma e_{N-1})(W - s) + \theta(W - s)^2\right] \\ &= \min_s \left[(P + \theta s)s + (P + \theta s)(W - s) + \theta(W - s)^2\right] \\ &= \min_s \left[W(P + \theta s) + \theta(W - s)^2\right]. \end{aligned}$$

The optimal s is $W/2$, giving value

$$V_{N-1}(P,W) = PW + \frac{3}{4}\theta W^2.$$

Thus: starting at day $N-1$ (so there are only 2 trading days) the investor should split his purchase in two equal parts, buying half the first day and half the second day. His impact on the market costs him, on average, an extra $\frac{3}{4}\theta W^2$ over the no-market-impact value PW.

Proceeding similarly for day $N-2$ etc., a pattern quickly becomes clear: starting at day $N-i$ with the goal of purchasing W shares, if the most recent price was P, the optimal trade on day i (the optimal s) is $W/(i+1)$, and the expected cost of all W shares is

$$V_{N-i}(P,W) = WP + \frac{i+2}{2(i+1)}\theta W^2.$$

This can be proved by induction. The inductive step is very similar to our calculation of V_{N-1}, and is left to the reader.

Notice the net effect of this calculation is extremely simple: no matter when he starts, the investor should divide his total goal W into equal parts – as many as there are trading days – and purchase one part each day. Taking $i = N-1$ we get the answer to our original question: if the most recent price is P and the goal is to buy S_{tot} over N days, then this optimal strategy leads to an expected total cost

$$V_1(P, S_{\text{tot}}) = PS_{\text{tot}} + \frac{\theta}{2}(1 + \frac{1}{N})S_{\text{tot}}^2.$$

There's something unusual about this conclusion. The investor's optimal strategy is not influenced by the random fluctuations of the prices. It's always the same, and can be fixed in advance. That's extremely unusual in stochastic control problems: the optimal control can usually be chosen as a *feedback control*, i.e. a deterministic function of the state – but since the state depends on the fluctuations, so does the control.

I warned you it was easy to criticize this model. Some comments:

1. The variance of the noise in the price model never entered our analysis. That's because our hypothetical investor is completely insensitive to risk – he cares only about the expected result, not about its variance. No real investor is like this.

2. The price law (1) is certainly wrong: it has the ith trade S_i increasing not just the ith price P_i but also *every subsequent price*. A better law would surely make the impact of trading *temporary*. Bertismas and Lo consider one such law, for which the problem still has a closed-form solution derived by methods similar to those used above.

The take-home message: Discrete-time stochastic dynamic programming is easy and fun. Of course a closed-form solution isn't always available. When there is no closed-form solution one must work backward in time *numerically*. The hardest part of the whole thing is keeping your indices straight, and remembering which information is known at time i, and which is random.

****

Example 2: When to sell an asset. This is an optimal stopping problem, analogous to Example 2 of Section 6. My discussion follows Section 2.4 of Bertsekas.

The problem is this: you have an asset (e.g. a house) you wish to sell. One offer arrives each week (yes, this example is oversimplified). The offers are independent draws from a single, known distribution. You must sell the house by the end of N weeks. If you sell it earlier, you'll invest the cash (risk-free), and its value will increase by factor $(1+r)$ each week. Your goal is to maximize the expected present value of the cash generated by the sale. We shall ignore transaction costs.

The control, of course, is the decision (taken each week) to sell or not to sell. The value function is

$$\begin{aligned} V_i(w) &= \text{expected present-value at week } i \text{ of current and future sales, if the house} \\ &\quad \text{is still unsold, the current week is } i, \text{ and the current offer is } w. \end{aligned}$$

We start as usual with the final time, $i = N$. If the house isn't already sold you have no choice but to sell it, realizing

$$V_N(w) = w.$$

The key to working backward is the principle of dynamic programming, which in this setting says:

$$V_i(w) = \max\left\{w, (1+r)^{-1}E[V_{i+1}(w')]\right\}.$$

Here w' is an independent trial from the specified distribution (the next week's offer); the first choice corresponds to the decision "sell now", the second choice to the decision "don't sell now".

The optimal decision in week i is easily seen to be:

$$\begin{aligned} \text{accept offer } w \quad &\text{if} \quad w \geq \alpha_i \\ \text{reject offer } w \quad &\text{if} \quad w \leq \alpha_i \end{aligned}$$

with

$$\alpha_i = (1+r)^{-1}E[V_{i+1}(w', n)].$$

To complete the solution to the problem we must find the sequence of real numbers $\alpha_0, \ldots \alpha_{N-1}$. Since

$$V_{i+1}(w,n) = \begin{cases} w & \text{if } w > \alpha_{i+1} \\ \alpha_{i+1} & \text{if } w \leq \alpha_{i+1} \end{cases}$$

we have

$$\begin{aligned} \alpha_i &= \frac{1}{1+r}\int_0^{\alpha_{i+1}} \alpha_{i+1}\, dP(w) + \frac{1}{1+r}\int_{\alpha_{i+1}}^{\infty} w\, dP(w) \\ &= \frac{1}{1+r}\alpha_{i+1}P(\alpha_{i+1}) + \frac{1}{1+r}\int_{\alpha_{i+1}}^{\infty} w\, dP(w) \end{aligned}$$

where $P(\lambda) = \text{prob}\{w < \lambda\}$ is the distribution function of w. This relation, with the initialization $\alpha_N = 0$, permits one to calculate the α's one by one (numerically). It can be shown that they are monotone in i: $\alpha_0 > \alpha_1 > \dots$ (see Bertsekas). This is natural, since early in the sales process it makes no sense to take a low offer, but later on it may be a good idea to avoid being forced to take a still lower one on week N. One can also show that after many steps of the recursion relation for α_i, the value of α_i approaches the fixed point α_* which solves

$$\alpha_* = \frac{1}{1+r}\alpha_* P(\alpha_*) + \frac{1}{1+r}\int_{\alpha_*}^{\infty} w\, dP(w).$$

Thus when the horizon is very far away, the optimal policy is to reject offers below α_* and accept offers above α_*.

Let's compare this discussion to Example 2 in Section 6. There we assumed lognormal dynamics (making the PDE easy to solve by hand) and considered the case where there was no deadline. Had we used some other SDE the PDE might have been harder to solve explicitly; had we imposed a deadline the value function would have become time-dependent and we would have been forced to solve for it numerically. Our discrete-time version includes both difficulties (an arbitrary probability distribution for offers, and a time-dependent value function). Therefore it must be solved numerically. But since time is intrinsically discrete, there are no technicalities such as discretization of the HJB equation.

****

Least-square replication of a European option. This discussion follows the paper D. Bertsimas, L. Logan, and A.W. Lo: *Hedging derivative securities and incomplete markets: an ϵ-arbitrage approach*, Operations Research 49 (2001) 372-397. (Downloadable from Andrew Lo's web site. Search Google for Andrew Lo MIT to get there.) The paper is quite rich; I focus for specificity on the simplest case, when the returns at different times are independent trials from a single probability distribution. However you'll see as we go along that this hypothesis isn't really being used; the method is actually much more general. (I'll comment on its scope at the end.)

Here's the problem. Consider a stock that can be traded only at discrete times $j\Delta t$, and suppose its price P_j at the jth time satisfies

$$P_j = P_{j-1}(1 + \phi_{j-1}) \tag{2}$$

where ϕ_{j-1} is chosen from a specified distribution, independent of j. (The discrete-time analogue of standard lognormal dynamics is obtained by taking $\log(1+\phi_j) = \mu\Delta t + \sigma\sqrt{\Delta t}e$ where e is Gaussian with mean 0 and variance 1.) You are an investment banker, and at time $j = 0$ you sell an option with maturity N and payout $F(P_N)$, receiving cash V_0 in payment. Your goal is to invest this cash wisely, trading in a self-financing way, so the value at the final time comes as close as possible to replicating the payout $F(P_N)$.

The *state* at time j is

$$\begin{aligned} V_j &= \text{the value of your portfolio at time } j, \text{ and} \\ P_j &= \text{the price at which trades can be made at time } j. \end{aligned}$$

We suppose that on each day, knowing V_j and P_j (but not the next day's price P_{j+1}) you make a decision how to rebalance your portfolio, buying or selling at price P_j till you hold θ_j units of stock and B_j units of cash. Thus θ_j is the *control*. Each trade must be "self-financing." To understand what this means, observe that going into the jth day your portfolio is worth

$$\theta_{j-1}P_j + B_{j-1}$$

while after rebalancing it is worth

$$\theta_j P_j + B_j.$$

For the trade to be self-financing these two expressions must be equal; this gives the restriction

$$P_j(\theta_j - \theta_{j-1}) + (B_j - B_{j-1}) = 0.$$

Since the value of your portfolio on the jth day is

$$V_j = \theta_j P_j + B_j,$$

the value changes from day to day by the law

$$V_j - V_{j-1} = \theta_{j-1}(P_j - P_{j-1}).$$

We interpret the goal of replicating the payout "as well as possible" in a least-square sense: your aim is to choose the θ_j's so as to

$$\text{minimize } E\left[(V_N - F(P_N))^2\right].$$

This time it is fairly obvious how to fit the problem into the dynamic programming framework. At time i the value function is

$$J_i(V,P) = \min_{\theta_i,\ldots,\theta_{N-1}} E_{V_i=V,\,P_i=P}\left[|V_N - F(P_N)|^2\right].$$

The final-time condition is

$$J_N(V,P) = |V - F(P)|^2$$

since on day N there is no decision to be made. The principle of dynamic programming gives

$$J_i(V,P) = \min_{\theta_i} E_{P_i=P}\left[J_{i+1}(V + \theta_i(P_{i+1} - P_i), P_{i+1})\right].$$

Now a small miracle happens (this is the advantage of the least-square formulation): *the value function is, at each time, a quadratic polynomial in V, with coefficients which are computable functions of P.* In fact:

Claim: The value functions have the form

$$J_i(V_i,P_i) = a_i(P_i)|V_i - b_i(P_i)|^2 + c_i(P_i)$$

and the optimal control θ_i is given by a feedback law that's linear in V_i:

$$\theta_i(V_i,P_i) = p_i(P_i) - V_i q_i(P_i).$$

The functions p_i, q_i, a_i, b_i, and c_i are determined inductively by the following explicit formulas:

$$\begin{aligned}
p_i(P_i) &= \frac{E\left[a_{i+1}(P_{i+1}) \cdot b_{i+1}(P_{i+1}) \cdot (P_{i+1} - P_i)\right]}{E\left[a_{i+1}(P_{i+1}) \cdot (P_{i+1} - P_i)^2\right]} \\
q_i(P_i) &= \frac{E\left[a_{i+1}(P_{i+1}) \cdot (P_{i+1} - P_i)\right]}{E\left[a_{i+1}(P_{i+1}) \cdot (P_{i+1} - P_i)^2\right]} \\
a_i(P_i) &= E\left[a_{i+1}(P_{i+1}) \cdot [1 - q_i(P_i)(P_{i+1} - P_i)]^2\right] \\
b_i(P_i) &= \frac{1}{a_i(P_i)} E\left[a_{i+1}(P_{i+1}) \cdot [b_{i+1}(P_{i+1}) - p_i(P_i)(P_{i+1} - P_i)] \cdot [1 - q_i(P_i)(P_{i+1} - P_i)]\right] \\
c_i(P_i) &= E\left[c_{i+1}(P_{i+1})\right] - a_i(P_i) \cdot b_i(P_i)^2 + E\left[a_{i+1}(P_{i+1}) \cdot [b_{i+1}(P_{i+1}) - p_i(P_i)(P_{i+1} - P_i)]^2\right]
\end{aligned}$$

where all expectations are over the uncertainties associated with passage from time i to $i + 1$. These relations can be solved backward in time, starting from time N, using the initialization

$$a_N(P_N) = 1, \quad b_N(P_N) = F(P_N), \quad c_N(P_N) = 0.$$

Play before work. Let's explore the impact of the claim.

Main consequence: The price you charged for the option – V_0 – never enters the analysis. But of course it's not irrelevant! If you charged V_0 for the option and the day-0 price was P_0, then your expected replication error is $J_0(V_0, P_0) = a_0(P_0)|V_0 - b_0(P_0)|^2 + c_0(P_0)$. The first term is always positive, so *the price that minimizes the replication error is* $V_0 = b_0(P_0)$.

Is $V_0 = b_0(P_0)$ necessarily the market price of the option? Not so fast! This would be so – by the usual absence-of-arbitrage argument – if $c_0(P_0)$ were zero, since in that case the payout is exactly replicatable. However in general $c_0(P_0)$ is positive. (It is clearly nonnegative, since the mean-square replication error is nonnegative no matter what the value of V_0. It is generally positive, due to market incompleteness: even the Black-Scholes marketplace is not complete in the discrete-time setting.) If c_0 is small then the price of the option should surely be close to b_0. However there is no logical reason why it should be exactly b_0. For example, the *sign* of the replication error $V_N - F(P_N)$ makes a great deal of difference to the seller (and to buyer) of the option, but it did not enter our discussion. Moreover there is no reason a specific investor should accept the *quadratic* replication error as the appropriate measure of risk.

What about the Black-Scholes marketplace, where the classical Black-Scholes-Merton analysis tells us how to price and hedge an option? That analysis is correct, of course, but it requires trading *continuously* in time. If you can trade only at discrete times $j\Delta t$ then the market is no longer complete and options are not exactly replicatable. If you use the optimal trading strategy determined in our Claim, your mean-square replication error will be *smaller than* the value obtained by using the continuous-time Black-Scholes hedging strategy (which sets $\theta_j = \partial V/\partial P$ evaluated at $P = P_j$ and $t = j\Delta t$, where V solves the Black-Scholes PDE). How much smaller? This isn't quite clear, at least not to me. The

paper by Bertsimas-Logan-Lo does show, however, that the discrete-time results converge to those of the continuous-time analysis as $\Delta t \to 0$.

OK, now work. We must justify the claim. Rather than do a formal induction, let us simply explain the first step: why the formulas are correct when $i = N - 1$. This has all the ideas of the general case, and the notation is slightly simpler since in this case $a_{i+1} = a_N = 1$. The principle of dynamic programming gives

$$J_{N-1}(V_{N-1}, P_{N-1}) = \min_{\theta_{N-1}} E\left[|V_{N-1} + \theta_{N-1}(P_N - P_{N-1}) - F(P_N)|^2\right].$$

Simplifying the notation, let us write the right hand side as

$$\min_{\theta} E\left[|V + \theta\delta P - F|^2\right], \tag{3}$$

bearing in mind that $\delta P = P_N - P_{N-1}$ and $F = F(P_N)$ are random variables, and V and θ are deterministic constants.

Identification of the optimal θ is essentially a task of linear algebra, since

$$\langle \xi, \eta \rangle = E[\xi\eta]$$

is an inner product on the vector space of random variables. We need to view the constant function V as a random variable; let us do this by writing it as $V\mathbf{1}$ where V is scalar and $\mathbf{1}$ is the random variable which always takes the value 1. Then (3) can be written as

$$\min_{\theta} \|V\mathbf{1} + \theta\delta P - F\|^2$$

where $\|\xi\|^2 = \langle \xi, \xi \rangle = E[\xi^2]$. Decomposing $\mathbf{1}$ and F into the parts parallel and orthogonal to δP, we have

$$\mathbf{1} = (\mathbf{1} - q\delta P) + q\delta P \quad \text{with } q = \langle \mathbf{1}, \delta P \rangle \|\delta P\|^{-2}$$

and

$$F = (F - p\delta P) + p\delta P \quad \text{with } p = \langle F, \delta P \rangle \|\delta P\|^{-2},$$

and

$$\begin{aligned} \|V\mathbf{1} + \theta\delta P - F\|^2 &= \|V(\mathbf{1} - q\delta P) - (F - p\delta P) + (\theta + Vq - p)\delta P\|^2 \\ &= \|V(\mathbf{1} - q\delta P) - (F - p\delta P)\|^2 + (\theta + Vq - p)^2 \|\delta P\|^2. \end{aligned}$$

The optimal θ makes the second term vanish: $\theta = p - Vq$, and the resulting value is

$$V^2 \|\mathbf{1} - q\delta P\|^2 - 2V\langle \mathbf{1} - q\delta P, F - p\delta P \rangle + \|F - p\delta P\|^2$$

This is, as expected, a quadratic polynomial in V, which can be written in the form $a(V - b)^2 + c$. Expressing p, q, a, b and c in the original probabilistic notation gives precisely the formulas of the Claim with $i = N - 1$. The general inductive step is entirely similar.

Let's close by discussing the scope of this method. The case considered above – returns that are independent and identically distributed at each time step – is already of real interest. It

includes the time discretization of the standard Black-Scholes marketplace, but it is much more general. For example, it can also be used to model a stock whose price process has jumps (see Section 3.3 of Bertsimas-Logan-Lo).

Moreover the framework is by no means restricted to the case of such simple price dynamics. All one really needs is that (i) the price is determined by a Markov process, and (ii) the payout at maturity depends on the final state of this process. Thus the same framework can be used for problems as diverse as:

- An exotic option whose payout is the maximum stock price between times 0 and N. Just replace the stock price process P_j with the process (P_j, M_j) = (price at time j, max price through time j), defined by

  $$P_j = P_{j-1}(1 + \phi_{j-1}), \quad M_j = \max\{P_j, M_{j-1}\}$$

  and replace the payout $F(P_N)$ by one of the form $F(M_N)$.

- Stochastic volatility. Just replace the stock price process P_j with a process (P_j, σ_j) = (price at time j, volatility at time j), with the time-discretization of your favorite stochastic-volatility model as the dynamics. The payout would, in this case, still have the form $F(P_N)$.

PDE for Finance Notes, Spring 2003 – Section 8

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, G63.2706.

MINOR CORRECTIONS AND CLARIFICATIONS ADDED 4/28/03.

Underlyings with jumps. We began the semester studying diffusions and the associated linear PDE's, namely the backward and forward Kolmogorov equations. This section asks: what becomes of that story when the underlying process has jumps? The answer, briefly, is that the backward and forward Kolmogorov equations become *integrodifferential* equations. In general they must be solved numerically, but in the constant-coefficient setting they can be solved using the *Fourier transform.* One application (though perhaps not the most important one) is to option pricing. But be careful: when the underlying can jump the market is not complete; so arbitrage alone does not determine the value of options. Why, then, should an option have a well-defined price? The answer, due to Merton, assumes that the extra randomness due to jumps is uncorrelated with the market – i.e. its β is zero. This means it can be made negligibile by diversification, and (by the Capital Asset Pricing Model) only the *average* effect of the jumps is important for pricing.

For basic issues involving processes with jumps and option pricing, I strongly recommend Merton's 1976 article "Option pricing when underlying stock returns are discontinuous" (reprinted as Chapter 9 of the book *Continuous Time Finance*). The Fourier Transform is discussed in many books; a fairly basic treatment can be found in H.F. Weinberger, *A First Course in Partial Differential Equations with Complex Variables and Transform Methods* (an inexpensive Dover paperback). Two current research papers on this topic are A. Lipton, *Assets with jumps*, *RISK*, Sept. 2002, 149-153; and A. Lewis, "A simple option formula for general jump-diffusion and other exponential Levy processes" (2001 preprint, available at http://optioncity.net/pubs/ExpLevy.pdf).

****

Jump-diffusion processes. The standard (constant-volatility) Black-Scholes model assumes that the logarithm of an asset price is normally distributed. In practice however the observed distributions are not normal – they have "fat tails," i.e. the probability of a very large positive or negative change is (though small) much larger than permitted by a Gaussian. The jump-diffusion model provides a plausible mechanism for explaining the fat tails and their consequences.

We work for simplicity in one space dimension. A diffusion solves $dy = f\,dt + g\,dw$. A jump-diffusion solves the same stochastic differential equation most of the time, but the solution occasionally jumps.

We need to specify the statistics of the jumps. We suppose the occurrence of a jump is a Poisson process with rate λ. This means the jumps are entirely independent of one another. Some characteristics of Poisson processes:

(a) The probability that a jump occurs during a short time interval of length Δt is $\lambda \Delta t + o(\Delta t)$.

(b) The probability of two or more jumps occuring during a short time interval of length Δt is negligible, i.e. $o(\Delta t)$.

(c) The probability of exactly n jumps occuring in a time interval of length t is $\frac{(\lambda t)^n}{n!}e^{-\lambda t}$.

(d) The mean waiting time for a jump is $1/\lambda$.

We also need to specify what happens when a jump occurs. Our assumption is that a jump takes y to $y + J$. The jump magnitudes are independent, identically distributed random variables. In other words, each time a jump occurs, its size J is selected by drawing from a pre-specified probability distribution.

This model is encapsulated by the equation

$$dy = f dt + g dw + J dN$$

where N counts the number of jumps that have occurred (so it takes integer values, starting at 0) and J represents the random jump magnitude. Ito's Lemma can be extended to this setting: if $v(x,t)$ is smooth enough and y is as above then $v(y(t),t)$ is again a jump-diffusion, with

$$d[v(y(t),t)] = (v_t + v_x f + \tfrac{1}{2}v_{xx}g^2)dt + v_x g dw + [v(y(t)+J,t) - v(y(t),t)]dN.$$

All the terms on the right are evaluated at $(y(t),t)$, as usual. In writing the jump term, we're trying to communicate that while the occurence of a jump in $v(y(t),t)$ is determined by N (i.e. by the presence of a jump in y) the size of the jump depends on $y(t)$ and the form of v.

Now consider the expected final-time payoff

$$u(x,t) = E_{y(t)=x}\left[w(y(T))\right]$$

where $w(x)$ is an arbitrary "payoff" (later it will be the payoff of an option). It is described as usual by a backward Kolmogorov equation

$$u_t + \mathcal{L}u = 0 \text{ for } t < T, \quad \text{with} \quad u(x,T) = w(x) \text{ at } t = T. \tag{1}$$

The operator $\mathcal{L}$ for our jump-diffusion process is

$$\mathcal{L}u = f u_x + \tfrac{1}{2}g^2 u_{xx} + \lambda E\left[u(x+J,t) - u(x,t)\right].$$

The expectation in the last term is over the probability distribution of jumps. The proof of (1) is no different from the one we gave in Section 1 for diffusions: let u solve (1), and apply Ito's formula. This gives

$$\begin{aligned} u(y(T),T) - u(x,t) &= \int_0^T (u_x g)(y(s),s)\, dw + \int_0^T (u_s + u_x f + \tfrac{1}{2}u_{xx}g^2)(y(s),s)\, ds \\ &\quad + \int_0^T [u(y(s)+J,s) - u(y(s),s)]dN. \end{aligned}$$

Now take the expectation. Only the jump term is unfamiliar; since the jump magnitudes are independent of the Poisson jump occurence process, we get

$$E\left([u(y(s)+J,s)-u(y(s),s)]dN\right) = E\left([u(y(s)+J,s)-u(y(s),s)]\right)\lambda ds.$$

Thus when u solves (1) we get

$$E[u(y(T),T)] - u(x,t) = 0.$$

This gives the result, since $u(y(T),T) = w(y(T))$ from the final-time condition on u.

A similar argument shows that the discounted final-time payoff

$$u(x,t) = E_{y(t)=x}\left[e^{-r(T-t)}w(y(T))\right]$$

solves

$$u_t + \mathcal{L}u - ru = 0 \text{ for } t < T, \quad \text{with} \quad u(x,T) = w(x) \text{ at } t = T,$$

using the same operator $\mathcal{L}$.

What about the probability distribution? We know from Section 1 that it solves the forward Kolmogorov equation,

$$p_s - \mathcal{L}^* p = 0 \text{ for } s > 0, \quad \text{with} \quad p(z,0) = p_0(z)$$

where p_0 is the initial probability distribution and $\mathcal{L}^*$ is the adjoint of $\mathcal{L}$. What is the adjoint of the new jump term? For any functions $\xi(z), \eta(z)$ we have

$$\int_{-\infty}^{\infty} E[[\xi(z+J) - \xi(z)]\,\eta(z)\,dz = \int_{-\infty}^{\infty} \xi(z) E[[\eta(z-J) - \xi(z)]\;dz =$$

since $\int_{-\infty}^{\infty} E[\xi(z+J)]\eta(z)\,dz = \int_{-\infty}^{\infty} \xi(z)E[\eta(z-J)]\,dz$. Thus

$$\mathcal{L}^* p = \tfrac{1}{2}(g^2 p)_{zz} - (fp)_z + \lambda E\left[p(z-J) - p(z)\right],$$

i.e. the probability distribution satisfies

$$p_s - \tfrac{1}{2}(g^2 p)_{zz} + (fp)_z - \lambda E\left[p(z-J,s) - p(z,s)\right] = 0.$$

****

Solution via Fourier transform. In general the backward and forward Kolmogorov equations must be solved numerically, but in the constant-coefficient setting we can do better. So we now focus on

$$dy = \mu dt + \sigma dw + J dN$$

i.e. we take f and g constant. To be specific let's focus on the *forward* Kolmogorov equation, which is now

$$p_s - \tfrac{1}{2}\sigma^2 p_{zz} + \mu p_z - \lambda E\left[p(z-J,s) - p(z,s)\right] = 0 \tag{2}$$

and let us solve it with initial condition $p(0) = \delta_{z=0}$. (This will give us the fundamental solution, i.e. $p(z,s)$ = probability of being at z at time s given that you started at 0 at time 0.)

Why is the Fourier transform useful in this setting? Basically, because the forward equation is a mess – nonlocal in space (due to the jump term) and involving derivatives too (the familiar terms). But when we take its Fourier transform in space we get a simple, easy-to-solve ODE. The result is a simple expression not for the probability distribution itself, but rather for its Fourier transform – what a probabilist would call the *characteristic function* of the distribution.

Most students in this class will be relatively unfamiliar with the Fourier transform. Here's a brief summary of what we'll use:

(a) Given a function $f(x)$, its Fourier transform $\hat{f}(\xi) = \mathcal{F}[f](\xi)$ is defined by

$$\hat{f}(\xi) = \int_{-\infty}^{\infty} f(x) e^{ix\xi}\, dx.$$

Notice that even if f is real-valued, $\hat{f}$ is typically complex-valued.

(b) Elementary manipulation reveals that the translated function $f(x-a)$ has Fourier transform

$$\mathcal{F}[f(x-a)](\xi) = e^{i\xi a}\hat{f}(\xi)$$

(c) Integration by parts reveals that the Fourier transform takes differentiation to multiplication:

$$\mathcal{F}[f_x](\xi) = -i\xi\hat{f}(\xi)$$

(d) It is less elementary to prove that the Fourier transform takes convolution to multiplication: if $h(x) = \int_{-\infty}^{\infty} f(x-y)g(y)\, dy$ then

$$\hat{h}(\xi) = \hat{f}(\xi)\hat{g}(\xi).$$

(e) Another less elementary fact is Plancherel's formula:

$$\int_{-\infty}^{\infty} \overline{f} g\, dx = \frac{1}{2\pi}\int_{-\infty}^{\infty} \overline{\mathcal{F}[f]}\mathcal{F}[g]\, d\xi$$

where $\overline{f}$ is the complex conjugate of f.

(f) The Fourier transform is invertible, and

$$f(x) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-ix\xi}\hat{f}(\xi)\, d\xi$$

(g) The Fourier transform of a Gaussian is again a Gaussian. More precisely, for a centered Gaussian with variance σ^2,

$$\mathcal{F}\left[\frac{1}{\sqrt{2\pi}\sigma}e^{-x^2/2\sigma^2}\right] = e^{-\sigma^2\xi^2/2}.$$

OK, let's use this tool to solve the forward equation (2). We apply the Fourier transform (in space) to each term of the equation. The first three terms are easy to handle using property (c). For the jump term, let the jump J have probability density ω, that

$$E\left[p(z-J)-p(z)\right] = \int [p(z-J)-p(z)]\omega(J)\,dJ = \int p(z-J)\omega(J)\,dJ - p(z).$$

By property (d) this the Fourier transform of this term is $(\hat{\omega}-1)\hat{p}$. Thus in the Fourier domain the equation becomes

$$\hat{p}_s + \tfrac{1}{2}\sigma^2\xi^2\hat{p} - i\xi\mu\hat{p} - \lambda(\hat{\omega}-1)\hat{p} = 0$$

with initial condition $\hat{p}(0) = \int e^{i\xi z}\delta_{z=0}\,dz = 1$. Writing the equation in the form

$$\hat{p}_s = K(\xi)\hat{p}$$

with

$$K(\xi) = -\tfrac{1}{2}\sigma^2\xi^2 + i\xi\mu + \lambda(\hat{\omega}(\xi)-1)$$

we recognize immediately that the solution is

$$\hat{p}(\xi,s) = e^{sK(\xi)}. \tag{3}$$

The probability density itself can of course be obtained by taking the inverse Fourier transform:

$$p(z,s) = \frac{1}{2\pi}\int_{-\infty}^{\infty} e^{-iz\xi+sK(\xi)}\,d\xi. \tag{4}$$

****

Option pricing. What use is this? Well, one can explore what the tails of the distribution look like, as you vary the jump rate λ and jump distribution ω. Or one can use it to price options. Recall that the value of an option should be the discounted expected payoff relative to the risk-neutral probability distribution. Assume for the moment that the under the risk-neutral distribution the stock price is e^y where y is a jump-diffusion as above. Then the time-0 value of an option with payoff w is

$$e^{-rT}E_{y(0)=\ln S_0}[w(e^y)]$$

if the time-0 stock price is S_0. If $w(S)$ vanishes for S near 0 and ∞ (corresponding to y very large negative or positive) then the option price is easy to express. Let $v(y) = w(e^y)$ be the payoff as a function of y, and let $x = \ln S_0$. By translation invariance, the probability density of y is $p(z-x,T)$ where p is given by (4). Therefore

$$\begin{aligned} E_{y(0)=\ln S_0}[w(e^y)] &= \int p(z-x,T)v(z)\,dz \\ &= \frac{1}{2\pi}\int \overline{\mathcal{F}[p(z-x,T)]}\mathcal{F}[v]\,d\xi \\ &= \frac{1}{2\pi}\int e^{-i\xi x}\hat{p}(-\xi,T)\hat{v}(\xi)d\xi. \end{aligned} \tag{5}$$

In the last step we used that $\mathcal{F}[p(z-x,T)] = e^{i\xi x}\hat{p}(\xi,T)$, and the fact that the complex conjugate of $\hat{p}(\xi,T) = \int e^{i\xi z}p(z,T)\,dz$ is $\int e^{-i\xi z}p(z,T)\,dz = \hat{p}(-\xi,T)$ since p is real. Equation (5) reduces the task of option pricing to the calculation of two Fourier transforms (those of ω and v) followed by a single integration (5).

The hypothesis that the payoff $w(S)$ vanishes near $S = 0$ and $S = \infty$ is inconvenient, because neither a put nor a call satisfies it. Fortunately this hypothesis can be dispensed with. Consider for example the call $w(S) = (S-K)_+$, for which $v(y) = (e^y - K)_+$. Its Fourier transform is not defined on the real axis, because the defining integral diverges as $y \to \infty$. But $e^{-\alpha y}v(y)$ decays near ∞ for $\alpha > 1$. So its Fourier transform in y is well-defined. This amounts to examining the Fourier transform of v along the line $\Im\xi = \alpha$ (here $\Im\xi$ is the imaginary part of ξ) since

$$\mathcal{F}[e^{-\alpha y}v(y)](\xi_1) = \int_{-\infty}^{\infty} e^{i(\xi_1+i\alpha)y}v(y)\,dy.$$

Fortunately, the Plancherel formula isn't restricted to integrating along the real axis in Fourier space; one can show that

$$\int_{-\infty}^{\infty} \overline{f}g\,dx = \frac{1}{2\pi}\int_{-\infty}^{\infty} \overline{\mathcal{F}[f]}(\xi_1+i\alpha)\mathcal{F}[g](\xi_1+i\alpha)\,d\xi_1$$

if the Fourier transforms of f and g exist (and are analytic) at $\Im\xi = \alpha$. Using this, an argument similar to that given above shows that

$$E_{y(0)=\ln S_0}[w(e^y)] = \int p(z-x,T)v(z)\,dz = \frac{1}{2\pi}\int e^{i(-\xi_1+i\alpha)x}\hat{p}(-\xi_1+i\alpha,T)\hat{v}(\xi_1+i\alpha)d\xi_1.$$

By the way, in the case of the call the Fourier transform of v is explicit and easy:

$$\hat{v}(\xi) = \int_{\ln K}^{\infty} e^{iy\xi}(e^y - K)\,dy = -\frac{K^{1+i\xi}}{\xi^2 - i\xi}$$

by elementary integration. Here $\xi = \xi_1 + i\alpha$ is any point in the complex plane such that the integral converges (this requires $\alpha > 1$).

How practical is this? Not very. European options are relatively liquid. Their prices are visible in the marketplace. Interpreting their prices using standard Black-Scholes (deducing an implied volatility from the option price) one obtains a result in contradiction to the model: the implied volatility depends on maturity and strike price (the dependence on strike price is often called the "volatility smile"). The introduction of jumps provides a plausible family of models that's better able to fit the market data. But it introduces headaches of modeling and calibration (e.g. how to choose the distribution of jumps?). If the goal were merely to price Europeans, there's no reason to bother – their prices are visible in the marketplace.

So why do this? Three reasons. One is the desire for a consistent theoretical framework. The second, more practical, is the need to hedge (not simply to price) options. The Delta predicted by a jump-diffusion model is different from that of the Black-Scholes framework; this will be clear in the next subsection. A third reason is the need to price and hedge

exotic options (e.g. barriers) which are less liquid. The basic idea: calibrate your jump-diffusion model using the market prices of European options, then use it to price and hedge barrier options. The difficulty of course is that pricing a barrier requires solving a *boundary-value* problem; thinking probabilistically, it requires knowing the probability distribution of hitting times. Such information can be obtained for some jump-diffusion processes. This topic is however beyond the scope of the present discussion (see e.g. the article by Lipton cited on page 1).

****

Hedging and the risk-neutral process. We assumed above that the time-0 value of an option with payoff $w(S)$ should be the discounted final-time payoff under "the risk-neutral dynamics." This is not obvious; indeed, it is a modeling hypothesis not consequence of arbitrage. I now explain the meaning and logic of this hypothesis, following Merton.

It is convenient to focus on the stock price itself not its logarithm. If (as above) J represents a typical jump of $\log S$ then the stock price leaps from S to $e^J S$, i.e. the jump in stock price is $(e^J - 1)S$. So (applying the jump-diffusion version of Ito to $S = e^y$, with $dy = \mu dt + \sigma dw$) the stock dynamics is

$$dS = (\mu + \tfrac{1}{2}\sigma^2)Sdt + \sigma S dw + (e^J - 1)SdN.$$

The associated risk-neutral process is determined by two considerations:

(a) it has the same volatility and jump statistics – i.e. it differs from the subjective process only by having a different drift; and

(b) under the risk-neutral process $e^{-rt}S$ is a martingale, i.e. $dS - rSdt$ has mean value 0.

We easily deduce that the risk-neutral process is

$$dS = (r - \lambda E[e^J - 1])Sdt + \sigma S dw + (e^J - 1)SdN. \tag{6}$$

Applying Ito once more, we see that under the risk-neutral dynamics $y = \log S$ satisfies

$$dy = (r - \tfrac{1}{2}\sigma^2 - \lambda E[e^J - 1])dt + \sigma dw + JdN$$

Thus the formalism developed in the preceding subsection can be used to price options; we need only set $\mu = r - \frac{1}{2}\sigma^2 - \lambda E[e^J - 1]$.

But is this right? And what are its implications for hedging? To explain, let's examine what becomes of the standard demonstration of the Black-Scholes PDE in the presence of jumps. Assume the option has a well-defined value $u(S(t), t)$ at time t. Suppose you try to hedge it by holding a long position in the option and a short position of Δ units of stock. Then over a short time interval the value of your position changes by

$$\begin{aligned} d[u(S(t),t)] - \Delta dS &= u_t dt + u_S([\mu + \tfrac{1}{2}\sigma^2]Sdt + \sigma S dw) + \tfrac{1}{2}u_{SS}\sigma^2 S^2 dt \\ &\quad + [u(e^J S(t), t) - u(S(t), t)]dN \\ &\quad - \Delta([\mu + \tfrac{1}{2}\sigma^2]Sdt + \sigma S dw) - \Delta(e^J - 1)SdN. \end{aligned}$$

There are two sources of randomness here – the Brownian motion dw and the jump process dN – but only one tradeable. So the market is incomplete, and there is no choice of Δ that makes this portfolio risk-free.

But consider the choice $\Delta = u_S(S(t), t)$. With this choice the randomness due to dw cancels, leaving only the uncertainty due to jumps:

$$\text{change in portfolio value} = (u_t + \tfrac{1}{2}\sigma^2 S^2 u_{SS})dt + \{[u(e^J S(t), t) - u(S(t), t)] - u_S(e^J S - S)\}dN.$$

To make progress, we must assume something about the statistics of the jumps. Merton's suggestion (still controversial) was to assume they are uncorrelated with the marketplace. The impact of such randomness can be eliminated by diversification. Put differently: according the the Capital Asset Pricing Model, for such an investment (whose β is zero) only the mean return is relevant to pricing. So the mean return on our hedge portfolio should be the risk-free rate:

$$(u_t + \tfrac{1}{2}\sigma^2 S^2 u_{SS})dt + \lambda E[u(e^J S(t), t) - u(S(t), t) - (e^J S - S)u_S]dt = r(u - Su_S)dt. \quad (7)$$

After rearrangement, this is precisely the backward Kolmogorov equation describing the discounted final-time payoff under the risk-neutral dynamics (6):

$$u_t + (r - \lambda E[e^J - 1])Su_S + \tfrac{1}{2}\sigma^2 S^2 u_{SS} - ru + \lambda E[u(e^J S, t) - u(S, t)] = 0.$$

A final remark about the experience of the investor who follows this hedge rule. If the option value is convex in S (as we expect for a call or a put) then the term in (7) associated with the jumps is positive:

$$E[u(e^J S(t), t) - u(S(t), t) - (e^J S - S)u_S] \geq 0.$$

So in the absence of jumps the value of the hedge portfolio (long the option, short $\Delta = u_S$ units of stock) rises a little slower than the risk-free rate. Without accounting for jumps, the investor who follows this hedge appears to be falling behind (relative to a cash investment at the risk-free rate). But due to convexity the net effect of the jumps is favorable – exactly favorable enough that the investor's long-term (mean) experience is risk-free.

PDE for Finance Notes, Spring 2003 – Section 9

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, G63.2706.

About the final exam: As previously announced, our exam is Monday May 12, 8-10pm, in the usual room Silver 207. Note the time-shift (8-10 not 7-9), intended to give students taking both Scientific Computing and PDE for Finance some breathing room. If this late start is a hardship for you, tell me – it is possible by request to take the exam 7-9pm instead of 8-10pm. The exam will be closed-book, but you may bring two sheets of notes (8.5×11, both sides, any font). The preparation such notes is an excellent study tool.

The exam covers the material in Sections 1-8 of the lecture notes, and Homeworks 1-6. Sole exception: there will be no exam questions using the Fourier Transform. This final Section 9 will *not* be on the exam.

A good exam question can be answered with very little calculation, provided you understand the relevant ideas. Most of the exam questions will be similar to (parts of) homework problems or examples discussed in the notes. The final lecture will be devoted to answering questions and reviewing what we've achieved this semester.

****

The martingale method for dynamic portfolio optimization. Sections 5-7 were devoted to stochastic control. We discussed the *value function* and the *principle of dynamic programming*. In the discrete-time setting dynamic programming gives an iterative scheme for finding the value function; in the continuous-time setting it leads to the Hamilton-Jacobi-Bellman PDE. Stochastic control is a powerful technique for optimal decision-making in the presence of uncertainty. In particular it places few restrictions on the sources of randomness, and it does not require special hypotheses such as market completeness.

In the continuous-time setting, a key application (due to Merton, around 1970) is *dynamic portfolio optimization*. We examined two versions of this problem: one optimizing the utility of consumption (Section 5), the other optimizing the utility of final-time wealth (Homework 5).

This section introduces an alternative approach to dynamic portfolio optimization. It is much more recent – the main papers were by Cox & Huang; Karatzas, Lehoczky, & Shreve; and Pliska, all in the mid-80's. A very clear, rather elementary account is given in R. Korn and E. Korn, *Option Pricing and Portfolio Optimization: Modern Methods of Financial Mathematics* (American Mathematical Society, 2001). My discussion is a simplified (i.e. watered-down) version of the one in Korn & Korn.

This alternative approach is called the "martingale method," for reasons that will become clear presently. It is closely linked to the modern understanding of option pricing via the discounted risk-neutral expected value. (Therefore this Section, unlike the rest of the course, requires some familiarity with continous-time finance.) The method is much less general than stochastic control; in particular, it requires that the market be complete. When it applies, however, it provides an entirely fresh viewpoint, quite different from Merton's. To

capture the main idea with a minimum of complexity, I shall (as in Section 5 and HW 5) consider just the case of a single risky asset. Moreover I shall focus on the case (like Problem 1 of HW5) where there is no consumption. So the goal is this: consider an investor who starts with wealth x at time 0. He can invest in a risk-free asset ("bond") which offers constant interest r, or a risky asset ("stock") whose price satisfies

$$dS = \mu S dt + \sigma S dw. \tag{1}$$

He expects to mix these, putting fraction $\theta(t)$ of his wealth in stock and the rest in the bond; the resulting wealth process satisfies

$$dX = [(1-\theta)r + \theta\mu]X dt + \theta\sigma X dw \tag{2}$$

with initial condition $X(0) = x$. His goal is to maximize his expected utility of final-time wealth:

$$\max_{\theta(t)} E\left[h(X(T)\right]$$

where h is his utility function and T is a specified time. Of course his investment decisions must be non-anticipating: $\theta(t)$ can depend only on knowledge of S up to time t.

Review of risk-neutral option pricing. Recall that the time-0 value of an option with payoff $f(S_T)$ is its discounted risk-neutral expected payoff:

$$\text{option value } = e^{-rT} E_{\text{RN}}[f(S_T)].$$

Moreover the risk-neutral process differs from (1) by having a different drift: it solves $dS = rSdt + \sigma S dw$. By Girsanov's theorem we can write the risk-neutral expected payoff using the subjective probability as

$$E_{\text{RN}}[f(S_T)] = E[e^{-z(T)} f(S_T)] \tag{3}$$

where

$$z(t) = \int_0^t \frac{\mu - r}{\sigma}\, dw + \tfrac{1}{2}\int_0^t \left(\frac{\mu - r}{\sigma}\right)^2 ds. \tag{4}$$

Clearly (3) can be written as

$$\text{option value } = E[H(T) f(S_T)]$$

with $H(t) = e^{-rt}e^{-z(t)}$. One verifies using Ito's formula and (4) that

$$dH = -rH dt - \frac{\mu - r}{\sigma} H dw \tag{5}$$

with initial condition $H(0) = 1$.

This option pricing formula doesn't come from thin air. It comes from the absence of arbitrage, together with the fact that the option payoff $f(S_T)$ can be replicated by a *hedge portfolio*, consisting of stock and bond in suitable weights. The option value is the value of the hedge portfolio at time 0, i.e. the initial capital needed to establish this (time-varying but self-financing) replicating portfolio.

The connection with portfolio optimization. What does this have to do with portfolio optimization? Plenty. The hedge portfolio represents a specific choice of weights $\theta(t)$. The option pricing formula tells us the condition under which a random final-time wealth of the form $X(T) = f(S(T))$ is achievable by a suitable trading strategy: the only restriction is that f satisfy $x = E[H(T)f(S_T)]$.

In portfolio optimization we are not mainly interested final-time wealths of the form $f(S_T)$. Rather, we are interested in those achievable by a suitable trading strategy (i.e. a non-anticipating choice of $\theta(t)$ for $0 \le t \le T$). The resulting random final-time wealth will, in general, be path-dependent; however it is certainly $\mathcal{F}_T$ measurable, i.e. it is determined by knowledge of the entire Brownian process $\{w(t)\}_{0 \le t \le T}$.

Our discussion of option pricing extends, however, to more general ($\mathcal{F}_T$-measureable) final-time wealths. The crucial question is: which (random, $\mathcal{F}_T$-measureable) final-time wealths B are associated with nonanticipating trading strategies $\theta(t)$ using initial capital x? The answer is simple: B has this property exactly if

$$x = E[H(T)B] \tag{6}$$

and in that case the associated wealth process $X(t)$ satisfies

$$H(t)X(t) = E\left[H(T)B \middle| \mathcal{F}_t\right]. \tag{7}$$

We'll prove just the easy direction: that if $B = X(T)$ for some trading strategy $\theta(t)$ then (6) and (7) hold. (See Korn & Korn for the converse.) Let's apply Ito's formula in the form

$$d(HX) = HdX + XdH + dHdX$$

(there is no factor of $1/2$ in front of the last term, because $f(h,x) = xh$ has $\partial^2 f/\partial h \partial x = \partial^2 f/\partial x \partial h = 1$). Substituting (2) and (5) gives

$$d(HX) = -HX(rdt + \frac{\mu - r}{\sigma}dw) + HX(\{(1-\theta)r + \theta\mu\}dt + \theta\sigma dw) - HX\frac{\mu - r}{\sigma}\theta\sigma dt.$$

The dt terms cancel, leaving only dw terms. So

$$H(T)X(T) - H(t)X(t) = \int_t^T [\text{stuff}]\, dw.$$

Setting $t = 0$ and taking the expectation of both sides gives

$$E[H(T)X(T)] = H(0)X(0) = x;$$

similarly, for any t we take the expectation of both sides conditioned on time-t data to get

$$E\left[H(T)X(T) \middle| \mathcal{F}_t\right] = H(t)X(t).$$

These are the desired assertions. (See e.g. Korn & Korn for the converse.)

The martingale approach to portfolio optimization. Relation (6) changes the task of dynamic portfolio optimization to a *static* optimization problem: the optimal final-time wealth B solves

$$\max_{E[H(T)B]=x} E[h(B)].$$

The solution is easy. To avoid getting confused let's pretend the list of possible final-time states was discrete, with state α having probability p_α, $1 \le \alpha \le N$. Then the random variable B would be characterized by its list of values $(B_1, \ldots, B_N)$ and our task would be to solve

$$\max_{\sum H(T)_\alpha B_\alpha p_\alpha = x} \sum h(B_\alpha) p_\alpha$$

for the optimal $B = (B_1, \ldots, B_N)$. This is easily achieved using the method of Lagrange multipliers. If λ is the Lagrange multiplier for the constraint then the solution satisfies $\partial L/\partial B_\alpha = 0$ and $\partial L/\partial \lambda = 0$ where

$$L(B, \lambda) = \sum h(B_\alpha) p_\alpha + \lambda \left(x - \sum H(T)_\alpha B_\alpha p_\alpha \right).$$

The derivative in λ recovers the constraint

$$\sum H(T)_\alpha B_\alpha p_\alpha = x$$

and the derivative with respect to B_α gives

$$h'(B_\alpha) = \lambda H(T)_\alpha$$

for each α. These $N+1$ equations determine the values of the $N+1$ unknowns B_α and λ.

The continuous case is no different: the optimal B satisfies

$$h'(B) = \lambda H(T) \tag{8}$$

as random variables. Since h is concave, h' is invertible, so we can solve (8) for B:

$$B = \{h'\}^{-1}(\lambda H(T))$$

where $\{h'\}^{-1}$ is the inverse function of h'. The value of λ is uniquely determined by the condition that this B satisfy

$$E[H(T)B] = x.$$

An example: the Merton problem with logarithmic utility. So far we have not assumed anything about the drift and volatility: they can be functions of time and stock price ($\mu = \mu(t, S)$ and $\sigma = \sigma(t, S)$). But to bring things down to earth let's consider a familiar example: the constant-drift, constant-volatility case, with utility $h(x) = \log x$. Notice that for this utility $h'(x) = 1/x$, so $\{h'\}^{-1}(y) = 1/y$. Also, since the drift and volatility are constant

$$H(t) = e^{-rt - z(t)}$$

with

$$z(t) = \frac{\mu - r}{\sigma} w(t) + \tfrac{1}{2} \left(\frac{\mu - r}{\sigma} \right)^2 t$$

while

$$S(t) = S_0 e^{(\mu - \frac{1}{2}\sigma^2)t + \sigma w(t)}.$$

Thus both $H(t)$ and $S(t)$ are determined by knowledge of t and $w(t)$. Put differently: in this case $H(T)$ is a function of $S(T)$.

Specializing (8) to our example gives $\frac{1}{B} = \lambda H(T)$, so

$$B = \frac{1}{\lambda H(T)}.$$

The value of λ is fixed by (6), which gives $x = E[H(T)B] = \frac{1}{\lambda}$ thus

$$\lambda = \frac{1}{x}.$$

The wealth at time T is determined by (7); it gives

$$H(t)X(t) = E\left[H(T)B\middle|\mathcal{F}_t\right] = E\left[\lambda^{-1}\middle|\mathcal{F}_t\right] = 1/\lambda$$

whence

$$X(t) = H^{-1}(t)x.$$

To implement this solution practically, what we really need is the weight $\theta(t)$ associated with the optimal policy. To find it, observe that by Ito's formula applied to (5),

$$d[H^{-1}x] = -H^{-2}xdH + H^{-3}xdHdH = H^{-1}x\left[rdt + \left(\frac{\mu - r}{\sigma}\right)^2 dt + \frac{\mu - r}{\sigma}dw\right].$$

This can be written as a wealth-process SDE, i.e. it has the form

$$dX = [(1-\theta)r + \theta\mu]Xdt + \theta\sigma Xdw$$

with $X(t) = H^{-1}x$ and

$$\theta = \frac{\mu - r}{\sigma^2}.$$

Thus the optimal policy is to keep fraction $\theta = (\mu - r)/\sigma^2$ of your wealth in the stock, and fraction $1 - \theta$ in the risk-free bond. This agrees, as it should, with the solution to Problem 1c of HW5.

The optimal portfolio is easy to describe but difficult to maintain, since it requires continuous-time trading. Indeed, each time $S(t)$ changes, some stock must be bought or sold so the value of the stock portfolio remains a constant proportion of the total wealth. This is of course impractical – one gets killed by transaction costs. In the constant-drift, constant-volatility case a convenient alternative presents itself. We saw that the optimal B is a function of $S(T)$ alone. So it can be expressed as the payoff of a suitable option, i.e. there is a function f (depending on T) such that $B = f(S(T))$. (It is not hard to find a formula for f, using the information given above.) If an option with payoff f were available in the marketplace, you could simply buy this option at time 0 and do no trading. (The optimal portfolio is identical to this option's hedge portfolio, which is in turn functionally equivalent to the option itself.) In practice the option is not available in the marketplace; even so, it can be approximated by a suitable combination of calls. This combination of options, held statically, reproduces (approximately) the optimal Merton payoff without incurring any transaction costs. For discussion of this observation and its consequences see M. Haugh and A. Lo, *Asset allocation and derivatives,* Quantitative Finance 1 (2001) 45-72.

PDE for Finance, Spring 2003 – Homework 1
Distributed 1/27/03, due 2/10/03.

1) Consider a diffusion $dy = f(y)ds + g(y)dw$, with initial condition $y(0) = x$. Suppose $u(x)$ solves the PDE

$$fu_x + \tfrac{1}{2}g^2 u_{xx} - q(x)u = 0 \quad \text{for } a < x < b, \text{ with } u = 1 \text{ at } x = a, b$$

for some function $q(x)$. Show that

$$u(x) = E_{y(0)=x}\left[e^{-\int_0^\tau q(y(s))ds}\right]$$

where τ is the exit time from $[a, b]$. (You should assume that $E[\tau] < \infty$.)

2) Consider the lognormal random walk

$$dy = \mu y dy + \sigma y dw$$

starting at $y(0) = x$. Assume $\mu \neq \frac{1}{2}\sigma^2$. The Section 1 notes examine the mean exit time from an interval $[a, b]$ where $0 < a < x < b$. There we used the PDE for the mean exit time

$$\mu x u_x + \tfrac{1}{2}\sigma^2 x^2 u_{xx} = -1 \quad \text{for } a < x < b \tag{1}$$

with boundary conditions $u(a) = u(b) = 0$ to derive an explicit formula for u.

(a) Show that the general solution of (1), without taking any boundary conditions into account, is

$$u = \frac{1}{\frac{1}{2}\sigma^2 - \mu}\log x + c_1 + c_2 x^\gamma$$

with $\gamma = 1 - 2\mu/\sigma^2$. Here c_1 and c_2 are arbitrary constants. [The formula given in the notes for the mean exit time is easy to deduce from this fact, by using the boundary conditions to solve for c_1 and c_2; however you need not do this calculation as part of your homework.]

(b) Argue as in the notes to show that the mean exit time from the interval (a, b) is finite. (Hint: mimic the argument used to answer Question 3, using $\phi(y) = \log y$.)

(c) Let p_a be the probability that the process exits at a, and $p_b = 1 - p_a$ the probability that it exits at b. Give an equation for p_a in terms of the barriers a, b and the initial value x. (Hint: mimic the argument used in the answer to Question 4, using $\phi(y) = y^\gamma$.) How does p_a behave in the limit $a \to 0$?

3) Consider a diffusion $dy = f(y)ds + g(y)dw$ starting at x at time 0, with $a < x < b$. Let τ be its exit time from the interval $[a, b]$, and assume $E[\tau] < \infty$.

(a) Let $u_a(x)$ be the probability it exits at a. Show that u_a solves $fu_x + \frac{1}{2}g_2 u_{xx} = 0$ with boundary conditions $u_a(a) = 1, u_a(b) = 0$.

(b) Apply this method to Problem 2(c). Is this approach fundamentally different from the one indicated by the hint above?

4) Consider once again a diffusion $dy = f(y)ds + g(y)dw$ starting at x at time 0. We know the mean arrival time to the boundary $v(x) = E[\tau]$ satisfies $fv_x + \frac{1}{2}g^2 v_{xx} = -1$ with $v = 0$ at $x = a, b$. Now consider the *second moment* of the arrival time $h(x) = E[\tau^2]$. Show that it satisfies $fh_x + \frac{1}{2}g^2 h_{xx} = -2v(x)$, with $h = 0$ at $x = a, b$.

5) Examine the analogues of Problem 2(a)–(c) when $\mu = \frac{1}{2}\sigma^2$. (Hint: notice that $xu_x + x^2 u_{xx} = u_{zz}$ with $z = \log x$.)

6) Let $w(t)$ be standard Brownian motion, starting at 0. Let τ_n be the first time w exits from the interval $[-n, 1]$, and let τ_∞ the the first time it reaches $w = 1$.

(a) Find the expected value of τ_n, and the probability that the path exits $[-n, 1]$ at $-n$.

(b) Verify by direct evaluation that $w(\tau_n)$ has mean value 0. (This must of course be true, since $E\left[\int_0^{\tau_n} dw = 0\right]$ by Dynkin's theorem.)

(c) Conclude from (a) that $E[\tau_\infty] = \infty$.

(d) Show that τ_∞ is almost-surely finite.

PDE for Finance, Spring 2003 – Homework 2
Distributed 2/10/03, due 2/24/03.

1) Consider the linear heat equation $u_t - u_{xx} = 0$ in one space dimension, with discontinuous initial data

$$u(x,0) = \begin{cases} 0 & \text{if } x < 0 \\ 1 & \text{if } x > 0. \end{cases}$$

(a) Show by evaluating the solution formula that

$$u(x,t) = N\left(\frac{x}{\sqrt{2t}}\right) \tag{1}$$

where N is the cumulative normal distribution

$$N(z) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{z} e^{-s^2/2}\, ds.$$

(b) Explore the solution by answering the following: what is $\max_x u_x(x,t)$ as a function of time? Where is it achieved? What is $\min_x u_x(x,t)$? For which x is $u_x > (1/10)\max_x u_x$? Sketch the graph of u_x as a function of x at a given time $t > 0$.

(c) Show that $v(x,t) = \int_{-\infty}^{x} u(z,t)\, dz$ solves $v_t - v_{xx} = 0$ with $v(x,0) = \max\{x, 0\}$. Deduce the qualitative behavior of $v(x,t)$ as a function of x for given t: how rapidly does v tend to 0 as $x \to -\infty$? What is the behavior of v as $x \to \infty$? What is the value of $v(0,t)$? Sketch the graph of $v(x,t)$ as a function of x for given $t > 0$.

2) We showed, in the Section 2 notes, that the solution of

$$w_t = w_{xx} \quad \text{for } t > 0 \text{ and } x > 0, \text{ with } w = 0 \text{ at } t = 0 \text{ and } w = \phi \text{ at } x = 0$$

is

$$w(x,t) = \int_0^t \frac{\partial G}{\partial y}(x, 0, t-s)\phi(s)\, ds \tag{2}$$

where $G(x,y,s)$ is the probability that a random walker, starting at x at time 0, reaches y at time s without first hitting the barrier at 0. (Here and throughout this problem set, the random walker solves $dy = \sqrt{2}dw$, i.e. it executes the scaled Brownian whose backward Kolmogorov equation is $u_t + u_{xx} = 0$.) Let's give an alternative demonstration of this fact, following the line of reasoning at the end of the Section 1 notes.

(a) Express, in terms of G, the probability that the random walker (starting at x at time 0) hits the barrier before time t. Differentiate in t to get the probability that it hits the barrier at time t. (This is known as the *first passage time* density).

(b) Use the forward Kolmogorov equation and integration by parts to show that the first passage time density is $\frac{\partial G}{\partial y}(x,0,t)$.

(c) Deduce the formula (2).

3) Give "solution formulas" for the following initial-boundary-value problems for the linear heat equation

$$w_t - w_{xx} = 0 \quad \text{for } t > 0 \text{ and } x > 0$$

with the specified initial and boundary conditions.

(a) $w_1 = 0$ at $x = 0$; $w_1 = 1$ at $t = 0$. Express your solution in terms of the function $u(x,t)$ defined in Problem 1.

(b) $w_2 = 0$ at $x = 0$; $w_2 = (x-K)_+$ at $t = 0$, with $K > 0$. Express your solution in terms of the function $v(x,t)$ defined in Problem 1(c).

(c) $w_3 = 0$ at $x = 0$; $w_3 = (x-K)_+$ at $t = 0$, with $K < 0$.

(d) $w_4 = 1$ at $x = 0$; $w_4 = 0$ at $t = 0$.

Interpret each as the expected payoff of a suitable barrier-type instrument, whose underlying executes the scaled Brownian motion $dy = \sqrt{2}dw$ with initial condition $y(0) = x$ and an absorbing barrier at 0. (Example: $w_1(x,T)$ is the expected payoff of an instrument which pays 1 at time T if the underlying has not yet hit the barrier and 0 otherwise.)

4) The Section 2 notes reduce the Black-Scholes PDE to the heat equation by brute-force algebraic substitution. This problem achieves the same reduction by a probabilistic route. Our starting point is the fact that

$$V(s,t) = e^{-r(T-t)} E_{y(t)=s} \left[\Phi(y(T)\right] \tag{3}$$

where $dy = rydt + \sigma ydw$.

(a) Consider $z = \frac{1}{\sigma} \log y$. By Ito's formula it satisfies $dz = \frac{1}{\sigma}(r - \frac{1}{2}\sigma^2)dt + dw$. Express the right hand side of (3) as a discounted expected value with respect to z process.

(b) The z process is Brownian motion with drift $\mu = \frac{1}{\sigma}(r - \frac{1}{2}\sigma^2)$. The Cameron-Martin-Girsanov theorem tells how to write an expected value relative to z as a weighted expected value relative to the standard Brownian motion w. Specifically:

$$E_{z(t)=x} \left[F(z(T))\right] = E_{w(t)=x} \left[e^{\mu w(T) - \frac{1}{2}\mu^2 (T-t)} F(w(T)) \right] \tag{4}$$

where left side is an expectation using the path-space measure associated with z, and the right hand side is an expectation using the path-space measure associated with Brownian motion. Apply this to get an expression for $V(s,t)$ whose right hand side involves an expected value relative to Brownian motion.

(c) An expected payoff relative to Brownian motion is described by the heat equation (more precisely by an equation of the form $u_t + \frac{1}{2}u_{xx} = 0$). Thus (b) expresses the solution of the Black-Scholes PDE in terms of a solution of the heat equation. Verify that this representation is the same as the one given in the Section 2 notes.

5) As noted in Problem 4(b), questions about Brownian motion with drift can often be answered using the Cameron-Martin-Girsanov theorem. But we can also study this process directly. Let's do so now, for the process $dz = \mu dt + dw$ with an absorbing barrier at $z = 0$.

(a) Suppose the process starts at $z_0 > 0$ at time 0. Let $G(z_0, z, t)$ be the probability that the random walker is at position z at time t (and has not yet hit the barrier). Show that

$$G(z_0, z, t) = \frac{1}{\sqrt{2\pi t}} e^{-|z-z_0-\mu t|^2/2t} - \frac{1}{\sqrt{2\pi t}} e^{-2\mu z_0} e^{-|z+z_0-\mu t|^2/2t}.$$

(Hint: just check that this G solves the relevant forward Kolmogorov equation, with the appropriate boundary and initial conditions.)

(b) Show that the first passage time density is

$$\frac{\partial G}{\partial z}(z_0, 0, t) = \frac{2z_0}{t\sqrt{2\pi t}} e^{-|z_0+\mu t|^2/2t}.$$

PDE for Finance, Spring 2003 – Homework 3
Distributed 3/3/03, due 3/24/03.

1) Consider the linear heat equation $u_t - u_{xx} = 0$ on the interval $0 < x < 1$, with boundary condition $u = 0$ at $x = 0, 1$ and initial condition $u = 1$.

(a) Interpret u as the value of a suitable double-barrier option.

(b) Express $u(t, x)$ as a Fourier sine series, as explained in Section 3.

(c) At time $t = 1/100$, how many terms of the series are required to give $u(t, x)$ within one percent accuracy?

2) Consider the SDE $dy = f(y)dt + g(y)dw$. Let $G(x, y, t)$ be the fundamental solution of the forward Kolmogorov PDE, i.e. the probability that a walker starting at x at time 0 is at y at time t. Show that if the infinitesimal generator is self-adjoint, i.e.

$$-(fu)_x + \tfrac{1}{2}(g^2 u)_{xx} = f u_x + \tfrac{1}{2} g^2 u_{xx},$$

then the fundamental solution is symmetric, i.e. $G(x, y, t) = G(y, x, t)$.

3) Consider the stochastic differential equation $dy = f(y, s)ds + g(y, s)dw$, and the associated backward and forward Kolmogorov equations

$$u_t + f(x, t)u_x + \tfrac{1}{2} g^2(x, t)u_{xx} = 0 \quad \text{for } t < T, \text{ with } u = \Phi \text{ at } t = T$$

and

$$\rho_s + (f(z, s)\rho)_z - \tfrac{1}{2}(g^2(z, s)\rho)_{zz} = 0 \quad \text{for } s > 0, \text{ with } \rho(z) = \rho_0(z) \text{ at } s = 0.$$

Recall that $u(x, t)$ is the expected value (starting from x at time t) of payoff $\Phi(y(T))$, whereas $\rho(z, s)$ is the probability distribution of the diffusing state $y(s)$ (if the initial distribution is ρ_0).

(a) The solution of the backward equation has the following property: if $m = \min_z \Phi(z)$ and $M = \max_z \Phi(z)$ then $m \le u(x, t) \le M$ for all $t < T$. Give two distinct justifications: one using the maximum principle for the PDE, the other using the probabilistic interpretation.

(b) The solution of the forward equation does *not* in general have the same property; in particular, $\max_z \rho(z, s)$ can be larger than the maximum of ρ_0. Explain why not, by considering the example $dy = -yds$. (Intuition: $y(s)$ moves toward the origin; in fact, $y(s) = e^{-s} y_0$. Viewing $y(s)$ as the position of a moving particle, we see that particles tend to collect at the origin no matter where they start. So $\rho(z, s)$ should be increasingly concentrated at $z = 0$.) Show that the solution in this case is $\rho(z, s) = e^s \rho_0(e^s z)$. This counterexample has $g = 0$; can you also give a counterexample using $dy = -yds + \epsilon dw$?

4) The solution of the forward Kolmogorov equation is a probability density, so we expect it to be nonnegative (assuming the initial condition $\rho_0(z)$ is everywhere nonnegative). In light of Problem 2b it's natural to worry whether the PDE has this property. Let's show that it does.

(a) Consider the initial-boundary-value problem

$$w_t = a(x,t)w_{xx} + b(x,t)w_x + c(x,t)w$$

with x in the interval $(0,1)$ and $0 < t < T$. We assume as usual that $a(x,t) > 0$. Suppose furthermore that $c < 0$ for all x and t. Show that if $0 \le w \le M$ at the initial time and the spatial boundary then $0 \le w \le M$ for all x and t. (Hint: a positive maximum cannot be achieved in the interior or at the final boundary. Neither can a negative minimum.)

(b) Now consider the same PDE but with $\max_{x,t} c(x,t)$ positive. Suppose the initial and boundary data are nonnegative. Show that the solution w is nonnegative for all x and t. (Hint: apply part (a) not to w but rather to $\bar{w} = e^{-Ct}w$ with a suitable choice of C.)

(c) Consider the solution of the forward Kolmogorov equation in the interval, with $\rho = 0$ at the boundary. (It represents the probability of arriving at z at time s without hitting the boundary first.) Show using part (b) that $\rho(z,s) \ge 0$ for all s and z.

[Comment: statements analogous to (a)-(c) are valid for the initial-value problem as well, when we solve for all $x \in R$ rather than for x in a bounded domain. The justification takes a little extra work however, and it requires some hypothesis on the growth of the solution at ∞.]

5) Consider the solution of

$$u_t + au_{xx} = 0 \quad \text{for } t < T, \text{ with } u = \Phi \text{ at } t = T$$

where a is a positive constant. Recall that in the stochastic interpretation, a is $\frac{1}{2}g^2$ where g represents volatility. Let's use the maximum principle to understand qualitatively how the solution depends on volatility.

(a) Show that if $\Phi_{xx} \ge 0$ for all x then $u_{xx} \ge 0$ for all x and t. (Hint: differentiate the PDE.)

(b) Suppose $\bar{u}$ solves the analogous equation with a replaced by $\bar{a} > a$, using the same final-time data Φ. We continue to assume that $\Phi_{xx} \ge 0$. Show that $\bar{u} \ge u$ for all x and t. (Hint: $w = \bar{u} - u$ solves $w_t + \bar{a}w_{xx} = f$ with $f = (a - \bar{a})u_{xx} \le 0$.)

6) Consider the standard finite difference scheme

$$\frac{u((m+1)\Delta t, n\Delta x) - u(m\Delta t, n\Delta x)}{\Delta t} = \frac{u(m\Delta t, (n+1)\Delta x) - 2u(m\Delta t, n\Delta x) + u(m\Delta t, (n-1)\Delta x)}{(\Delta x)^2} \tag{1}$$

for solving $u_t - u_{xx} = 0$. The stability restriction $\Delta t < \frac{1}{2}\Delta x^2$ leaves a lot of freedom in the choice of Δx and Δt. Show that

$$\Delta t = \frac{1}{6}\Delta x^2$$

is special, in the sense that the numerical scheme (1) has errors of order Δx^4 rather than Δx^2. In other words: when u is the exact solution of the PDE, the left and right sides of (1) differ by a term of order Δx^4. [Comment: the argument sketched in the Section 3 Addendum shows that if u solves the PDE and v solves the finite difference scheme then $|u - v|$ is of order Δx^2 in general, but it is smaller – of order Δx^4 – when $\Delta t = \frac{1}{6}\Delta x^2$.]

PDE for Finance, Spring 2003 – Homework 4
Distributed 3/10/03, due 3/31/03.

Problems 1 – 4 concern deterministic optimal control (Section 4 material); problems 5 – 7 concern stochastic control (Section 5 material). Warning: this problem set is longer than usual (mainly because Problems 1 – 4, though not especially difficult, are fairly laborious.)

1) Consider the finite-horizon utility maximization problem with discount rate ρ. The dynamical law is thus

$$dy/ds = f(y(s), \alpha(s)), \quad y(t) = x,$$

and the optimal utility discounted to time 0 is

$$u(x,t) = \max_{\alpha \in A} \left\{ \int_t^T e^{-\rho s} h(y(s), \alpha(s))\, ds + e^{-\rho T} g(y(T)) \right\}.$$

It is often more convenient to consider, instead of u, the optimal utility discounted to time t; this is

$$v(x,t) = e^{\rho t} u(x,t) = \max_{\alpha \in A} \left\{ \int_t^T e^{-\rho(s-t)} h(y(s), \alpha(s))\, ds + e^{-\rho(T-t)} g(y(T)) \right\}.$$

(a) Show (by a heuristic argument similar to those in the Section 4 notes) that v satisfies

$$v_t - \rho v + H(x, \nabla v) = 0$$

with Hamiltonian

$$H(x,p) = \max_{a \in A} \{ f(x,a) \cdot p + h(x,a) \}$$

and final-time data

$$v(x,T) = g(x).$$

(Notice that the PDE for v is autonomous, i.e. there is no explicit dependence on time.)

(b) Now consider the analogous infinite-horizon problem, with the same equation of state, and value function

$$\bar{v}(x,t) = \max_{\alpha \in A} \int_t^\infty e^{-\rho(s-t)} h(y(s), \alpha(s))\, ds.$$

Show (by an elementary comparison argument) that $\bar{v}$ is independent of t, i.e. $\bar{v} = \bar{v}(x)$ is a function of x alone. Conclude using part (a) that if $\bar{v}$ is finite, it solves the stationary PDE

$$-\rho \bar{v} + H(x, \nabla \bar{v}) = 0.$$

2) Recall Example 1 of the Section 4 notes: the state equation is $dy/ds = ry - \alpha$ with $y(t) = x$, and the value function is

$$u(x,t) = \max_{\alpha \geq 0} \int_t^\tau e^{-\rho s} h(\alpha(s))\, ds$$

with $h(a) = a^\gamma$ for some $0 < \gamma < 1$, and

$$\tau = \begin{cases} \text{first time when } y = 0 \text{ if this occurs before time } T \\ T \text{ otherwise.} \end{cases}$$

(a) We obtained a formula for $u(x,t)$ in the Section 4 notes, however our formula doesn't make sense when $\rho - r\gamma = 0$. Find the correct formula in that case.

(b) Let's examine the infinite-horizon-limit $T \to \infty$. Following the lead of Problem 1 let us concentrate on $v(x,t) = e^{\rho t} u(x,t) =$ optimal utility discounted to time t. Show that

$$\bar{v}(x) = \lim_{T \to \infty} v(x,t) = \begin{cases} G_\infty x^\gamma & \text{if } \rho - r\gamma > 0 \\ \infty & \text{if } \rho - r\gamma \leq 0 \end{cases}$$

with $G_\infty = [(1-\gamma)/(\rho - r\gamma)]^{1-\gamma}$.

(c) Use the stationary PDE of Problem 1(b) (specialized to this example) to obtain the same result.

(d) What is the optimal consumption strategy, for the infinite-horizon version of this problem?

3) Consider the analogue of Example 1 with the power-law utility replaced by the logarithm: $h(a) = \ln a$. To avoid confusion let us write u_γ for the value function obtained in the notes using $h(a) = a^\gamma$, and $u_{\log}$ for the value function obtained using $h(a) = \ln a$. Recall that $u_\gamma(x,t) = g_\gamma(t) x^\gamma$ with

$$g_\gamma(t) = e^{-\rho t} \left[\frac{1-\gamma}{\rho - r\gamma} \left(1 - e^{-\frac{(\rho - r\gamma)(T-t)}{1-\gamma}} \right) \right]^{1-\gamma}.$$

(a) Show, by a direct comparison argument, that

$$u_{\log}(\lambda x, t) = u_{\log}(x,t) + \frac{1}{\rho} e^{-\rho t} (1 - e^{-\rho(T-t)}) \ln \lambda$$

for any $\lambda > 0$. Use this to conclude that

$$u_{\log}(x,t) = g_0(t) \ln x + g_1(t)$$

where $g_0(t) = \frac{1}{\rho} e^{-\rho t}(1 - e^{-\rho(T-t)})$ and g_1 is an as-yet unspecified function of t alone.

(b) Pursue the following scheme for finding g_1: Consider the utility $h = \frac{1}{\gamma}(a^\gamma - 1)$. Express its value function u_h in terms of u_γ. Now take the limit $\gamma \to 0$. Show this gives a result of the expected form, with

$$g_0(t) = g_\gamma(t)|_{\gamma=0}$$

and

$$g_1(t) = \frac{dg_\gamma}{d\gamma}(t)\Big|_{\gamma=0}.$$

(This leads to an explicit formula for g_1 but it's messy; I'm not asking you to write it down.)

(c) Indicate how g_0 and g_1 could alternatively have been found by solving appropriate PDE's. (Hint: find the HJB equation associated with $h(a) = \ln a$, and show that the ansatz $u_{\log} = g_0(t) \ln x + g_1(t)$ leads to differential equations that determine g_0 and g_1.)

4) Our Example 1 considers an investor who receives interest (at constant rate r) but no wages. Let's consider what happens if the investor also receives wages at constant rate w. The equation of state becomes

$$dy/ds = ry + w - \alpha \quad \text{with } y(t) = x,$$

and the value function is

$$u(x,t) = \max_{\alpha \geq 0} \int_t^T e^{-\rho s} h(\alpha(s))\, ds$$

with $h(a) = a^\gamma$ for some $0 < \gamma < 1$. Since the investor earns wages, we now permit $y(s) < 0$, however we insist that the final-time wealth be nonnegative ($y(T) \geq 0$).

(a) Which pairs (x,t) are acceptable? The strategy that maximizes $y(T)$ is clearly to consume nothing ($\alpha(s) = 0$ for all $t < s < T$). Show this results in $y(T) \geq 0$ exactly if

$$x + \phi(t)w \geq 0$$

where

$$\phi(t) = \frac{1}{r}\left(1 - e^{-r(T-t)}\right).$$

Notice for future reference that ϕ solves $\phi' - r\phi + 1 = 0$ with $\phi(T) = 0$.

(b) Find the HJB equation that $u(x,t)$ should satisfy in its natural domain $\{(x,t) : x + \phi(t)w \geq 0\}$. Specify the boundary conditions when $t = T$ and where $x + \phi w = 0$.

(c) Substitute into this HJB equation the ansatz

$$v(x,t) = e^{-\rho t} G(t)(x + \phi(t)w)^\gamma.$$

Show v is a solution when G solves the familiar equation

$$G_t + (r\gamma - \rho)G + (1-\gamma)G^{\gamma/(\gamma-1)} = 0$$

(the same equation we solved in Example 1). Deduce a formula for v.

(d) In view of (a), a more careful definition of the value function for this control problem is

$$u(x,t) = \max_{\alpha \geq 0} \int_t^\tau e^{-\rho s} h(\alpha(s))\, ds$$

where

$$\tau = \begin{cases} \text{first time when } y(s) + \phi(s)w = 0 \text{ if this occurs before time } T \\ T \text{ otherwise.} \end{cases}$$

Use a verification argument to prove that the function v obtained in (c) is indeed the value function u defined this way.

5) Our geometric Example 2 gave $|\nabla u| = 1$ in D (with $u = 0$ at ∂D) as the HJB equation associated with starting at a point x in some domain D, traveling with speed at most 1, and arriving at ∂D as quickly as possible. Let's consider what becomes of this problem when we introduce a little noise. The state equation becomes

$$dy = \alpha(s)ds + \epsilon dw, \quad y(0) = x,$$

where $\alpha(s)$ is a (non-anticipating) control satisfying $|\alpha(s)| \leq 1$, y takes values in R^n, and each component of w is an independent Brownian motion. Let $\tau_{x,\alpha}$ denote the arrival time:

$$\tau_{x,\alpha} = \text{time when } y(s) \text{ first hits } \partial D,$$

which is of course random. The goal is now to minimize the *expected* arrival time at ∂D, so the value function is

$$u(x) = \min_{|\alpha(s)| \leq 1} E_{y(0)=x} \left\{ \tau_{x,\alpha} \right\}.$$

(a) Show, using an argument similar to that in the Section 5 notes, that u solves the PDE

$$1 - |\nabla u| + \tfrac{1}{2}\epsilon^2 \Delta u = 0 \quad \text{in } D$$

with boundary condition $u = 0$ at ∂D.

(b) Your answer to (a) should suggest a specific feedback strategy for determining $\alpha(s)$ in terms of $y(s)$. What is it?

6) Let's solve the differential equation from the last problem explicitly, for the special case when $D = [-1, 1]$:

$$\begin{aligned} 1 - |u_x| + \tfrac{1}{2}\epsilon^2 u_{xx} &= 0 \quad \text{for } -1 < x < 1 \\ u &= 0 \quad \text{at } x = \pm 1. \end{aligned}$$

(a) Assuming that the solution u is unique, show it satisfies $u(x) = u(-x)$. Conclude that $u_x = 0$ and $u_{xx} < 0$ at $x = 0$. Thus u has a maximum at $x = 0$.

(b) Notice that $v = u_x$ solves $1 - |v| + \delta v_x = 0$ with $\delta = \frac{1}{2}\epsilon^2$. Show that

$$\begin{aligned} v &= -1 + e^{-x/\delta} \qquad \text{for } 0 < x < 1 \\ v &= +1 - e^{x/\delta} \qquad \text{for } -1 < x < 0. \end{aligned}$$

Integrate once to find a formula for u.

(c) Verify that as $\epsilon \to 0$, this solution approaches $1 - |x|$.

[Comment: the assumption of uniqueness in part (a) is convenient, but it can be avoided. Outline of how to do this: observe that any critical point of u must be a local maximum (since $u_x = 0$ implies $u_{xx} < 0$). Therefore u has just one critical point, say x_0, which is a maximum. Get a formula for u by arguing as in (b). Then use the boundary condition to see that x_0 had to be 0.]

7) Let's consider what becomes of Merton's optimal investment and consumption problem if there are two risky assets: one whose price satisfies $dp_2 = \mu_2 p_2 dt + \sigma_2 p_2 dw_2$ and another whose price satisfies $dp_3 = \mu_3 p_3 dt + \sigma_3 p_3 dw_3$. To keep things simple let's suppose w_2 and w_3 are independent Brownian motions. It is natural to assume $\mu_2 > r$ and $\mu_3 > r$ where r is the risk-free rate. (Why?) Let $\alpha_2(s)$ and $\alpha_3(s)$ be the proportions of the investor's total wealth invested in the risky assets at time s, so that $1 - \alpha_2 - \alpha_3$ is the proportion of wealth invested risk-free. Let β be the rate of consumption. Then the investor's wealth satisfies

$$dy = (1 - \alpha_2 - \alpha_3) y r ds + \alpha_2 y(\mu_2 ds + \sigma_2 dw_2) + \alpha_3 y(\mu_3 ds + \sigma_3 dw_3) - \beta ds.$$

(Be sure you understand this; but you need not explain it on your solution sheet.) Use the power-law utility: the value function is thus

$$u(x,t) = \max_{\alpha_2, \alpha_3, \beta} E_{y(t)=x} \left[\int_t^\tau e^{-\rho s} \beta^\gamma(s)\, ds \right]$$

where τ is the first time $y(s) = 0$ if this occurs, or $\tau = T$ otherwise.

(a) Derive the HJB equation.

(b) What is the optimal investment policy (the optimal choice of α_2 and α_3)? What restriction do you need on the parameters to be sure $\alpha_2 > 0$, $\alpha_3 > 0$, and $\alpha_2 + \alpha_3 < 1$?

(c) Find a formula for $u(x,t)$. [Hint: the nonlinear equation you have to solve is not really different from the one considered in Section 5.]

PDE for Finance, Spring 2003 – Homework 5
Distributed 4/7/03, due 4/21/03.

Problem 1 is a classic example (due to Merton) of optimal asset allocation. Problems 2-4 reinforce our discussion of optimal stopping and American options. Problem 5 displays the power of dynamic programming for solving a different type of optimal stopping problem (one that's intrinsically discrete).

1) Consider the following asset-allocation problem. Two investment opportunities are available. One is risk-free, earning (constant) interest r. The other is lognormal, with (constant) drift μ and volatility σ, i.e. it satisfies $dp = \mu p ds + \sigma p dw$. You start at time t by investing wealth x. Your control is the weighting of your portofolio between these two assets, i.e.

$$\alpha(s) = \text{fraction of wealth invested in the risky asset at time } s$$

subject to $0 \leq \alpha \leq 1$. You never withdraw from or add to the portfolio, and you have a fixed horizon T. Your goal is to maximize the utility of your portfolio value at time T; in other words, your value function is

$$u(x,t) = \max_{\alpha(s)} E_{y(t)=x}\left[h(y(T))\right]$$

where $y(s)$ is the value of the portfolio at time s.

(a) Find the HJB equation satisfied by u.

(b) Find the solution – and the optimal investment strategy – if your utility is $h(y) = y^\gamma$ with $0 < \gamma < 1$.

(c) Find the solution – and the optimal investment strategy – if your utility is $h(y) = \log y$.

2) Example 2 of the Section 6 notes discusses when to sell a stock. The goal proposed in the notes was to maximize the discounted wealth realized by the sale, i.e.

$$\max_\tau E_{y(0)=x}\left[e^{-r\tau}(x-a)\right]$$

A different goal would be to maximize the discounted *utility* of wealth realized by the sale, i.e.

$$\max_\tau E_{y(0)=x}\left[e^{-r\tau}h(x-a)\right]$$

where h is your utility.

(a) Consider the utility $h(y) = y^\gamma$ with $0 < \gamma < 1$. (This is concave only for $y > 0$, but that's OK – it would clearly be foolish to sell at a price that realizes a loss.) Find the value function and the optimal strategy.

(b) The example in the notes corresponds to $\gamma = 1$. Using $\gamma < 1$ corresponds to introducing risk-averseness, and decreasing γ corresponds to increasing the risk-averseness. How is this reflected in the γ-dependence of the optimal strategy?

3) In Example 2 of the Section 6 notes we assumed $\mu < r$. Let's explore what happens when $\mu \geq r$. All other conventions of Example 2 remain in effect: the asset price satisfies $dy = \mu y dt + \sigma y dw$ and the value function is $u(x) = \max_\tau E_{y(0)=x}[e^{-r\tau}(y(\tau) - a)]$.

(a) Show that if $\mu > r$ then $u = \infty$.

(b) Show that if $\mu = r$ then $u(x) = x$.

(Hint: consider the value associated with sales threshold h, as $h \to \infty$.)

4) For a lognormal underlying with continuous dividend yield d, the risk-neutral process is $dy = (r - d)y dt + \sigma y dw$. The value of a perpetual American call with strike K is thus

$$u(x) = \max_\tau E_{y(0)=x}\left[e^{-r\tau}(y(\tau) - K)_+\right]$$

where r is the risk-free rate.

(a) Find the value of this option, and the optimal exercise rule, for $d > 0$.

(b) Show that as $d \to 0$ the value approaches $u(x) = x$.

5) [from Dimitri Bertsekas, *Dynamic Programming: Deterministic and Stochastic Models*, Chapter 2, problem 19]. A driver is looking for a parking place on the way to his destination. Each parking place is free with probability p, independent of whether other parking spaces are free or not. The driver cannot observe whether a parking place is free until he reaches it. If he parks k places from his destination, he incurs a cost k. If he reaches the destination without having parked the cost is C.

(a) Let F_k be the minimal expected cost if he is k parking places from his destination. Show that

$$\begin{aligned} F_0 &= C \\ F_k &= p\min[k, F_{k-1}] + qF_{k-1}, \quad k = 1, 2, \ldots \end{aligned}$$

where $q = 1 - p$.

(b) Show that an optimal policy is of the form: never park if $k \geq k^*$, but take the first free place if $k < k^*$, where k is the number of parking places from the destination, and k^* is the smallest integer i satisfying $q^{i-1} < (pC + q)^{-1}$.

PDE for Finance, Spring 2003 – Homework 6
Distributed 4/21/03, due 5/5/03. *Includes a hint for problem 4.*

1) This problem develops a continuous-time analogue of the simple Bertsimas & Lo model of "Optimal control of execution costs" presented in the Section 7 notes. The state is (w, p), where w is the number of shares yet to be purchased and p is the current price per share. The control $\alpha(s)$ is the rate at which shares are purchased. The state equation is:

$$\begin{aligned} dw &= -\alpha\, ds \text{ for } t < s < T, \quad w(t) = w_0 \\ dp &= \theta\alpha\, ds + \sigma dz \text{ for } t < s < T, \quad p(t) = p_0 \end{aligned}$$

where dz is Brownian motion and θ, σ are fixed constants. The goal is to minimize, among (nonanticipating) controls $\alpha(s)$, the expected cost

$$E\left\{\int_t^T [p(s)\alpha(s) + \theta\alpha^2(s)]\, ds + [p(T)w(T) + \theta w^2(T)]\right\}.$$

The optimal expected cost is the value function $u(w_0, p_0, t)$.

(a) Show that the HJB equation for u is

$$u_t + H(u_w, u_p, p) + \frac{\sigma^2}{2} u_{pp} = 0$$

for $t < T$, with Hamiltonian

$$H(u_w, u_p, p) = -\frac{1}{4\theta}(p + \theta u_p - u_w)^2.$$

The final value is of course

$$u(w, p, T) = pw + \theta w^2.$$

(b) Look for a solution of the form $u(w, p, t) = pw + g(t)w^2$. Show that g solves

$$\dot{g} = \frac{1}{4\theta}(\theta - 2g)^2$$

for $t < T$, with $g(T) = \theta$. Notice that u does not depend on σ, i.e. setting $\sigma = 0$ gives the same value function.

(c) Solve for g. (Hint: start by rewriting the equation for g, "putting all the g's on the left and all the t's on the right".)

(d) Show by direct examination of your solution that the optimal $\alpha(s)$ is constant.

(Food for thought: what happens if one takes the running cost to be $\int_t^T p(s)\alpha(s)\, ds$ instead of $\int_t^T p(s)\alpha(s) + \theta\alpha^2(s)\, ds$?)

2) The Section 7 notes discuss work by Bertsimas, Kogan, and Lo involving least-square replication of a European option. The analysis there assumes all trades are *self-financing*, so the value of the portfolio at consecutive times is related by

$$V_j - V_{j-1} = \theta_{j-1}(P_j - P_{j-1}).$$

Let's consider what happens if trades are permitted to be non-self-financing. This means we introduce an additional control g_j, the amount of cash added to (if $g_j > 0$) or removed from (if $g_j < 0$) the portfolio at time j, and the portfolio values now satisfy

$$V_j - V_{j-1} = \theta_{j-1}(P_j - P_{j-1}) + g_{j-1}.$$

It is natural to add a quadratic expression involving the g's to the objective: now we seek to minimize

$$E\left[(V_N - F(P_N))^2 + \alpha \sum_{j=0}^{N-1} g_j^2\right]$$

where α is a positive constant. The associated value function is

$$J_i(V,P) = \min_{\theta_i, g_i; \ldots; \theta_{N-1}, g_{N-1}} E_{V_i=V,\, P_i=P}\left[(V_N - F(P_N))^2 + \alpha \sum_{j=i}^{N-1} g_j^2\right].$$

The claim enunciated in the Section 7 notes remains true in this modified setting: J_i can be expressed as a quadratic polynomial

$$J_i(V_i, P_i) = \bar{a}_i(P_i)|V_i - \bar{b}_i(P_i)|^2 + \bar{c}_i(P_i)$$

where $\bar{a}_i, \bar{b}_i$, and $\bar{c}_i$ are suitably-defined functions which can be constructed inductively. Demonstrate this assertion in the special case $i = N-1$, and explain how $\bar{a}_{N-1}, \bar{b}_{N-1}, \bar{c}_{N-1}$ are related to the functions $a_{N-1}, b_{N-1}, c_{N-1}$ of the Section 7 notes.

3) Consider scaled Brownian motion with drift, $dy = \mu dt + \sigma dw$, starting at $y(0) = 0$. The solution is of course $y = \mu t + \sigma w(t)$, so its probability distribution at time t is Gaussian with mean μt and variance $\sigma^2 t$. Show that solution $\hat{p}(\xi, t)$ obtained by Fourier transform in the Section 8 notes is consistent with this result.

4) Consider scaled Brownian motion with drift and jumps: $dy = \mu dt + \sigma dw + JdN$, starting at $y(0) = 0$. Assume the jump occurences are Poisson with rate λ, and the jump magnitudes J are Gaussian with mean 0 and variance δ^2. Find the probability distribution of the process y at time t. (*Hint*: don't try to use the Fourier transform. Instead observe that you know, for any n, the probability that n jumps will occur before time t; and after conditioning on the number of jumps, the distribution of y is a Gaussian whose mean and variance are easy to determine. Assemble these ingredients to give the density of y as an infinite sum.)

PDE for Finance Final Exam Questions
Spring 2003 – Professor Kohn

1) Give a probabilistic interpretation for the solution of each PDE. (You must justify your answer to receive full credit.)

(a) $u_t + f(x)u_x + \frac{1}{2}g^2(x)u_{xx} = 0$ for $t < T$ and $x \in R$, with final-time condition $u(x,T) = \phi(x)$.

(b) $f(x)u_x + \frac{1}{2}g^2(x)u_{xx} = -1$ on the interval $a < x < b$, with $u = 0$ at the boundary points $x = a, b$.

2) This problem concerns the explicit solution formulas for the linear heat equation in a half-space and a bounded interval.

(a) The solution of

$$u_t = u_{xx} \quad \text{for } t > 0 \text{ and } x > 1, \text{ with } u = (x-1)^3 \text{ at } t = 0 \text{ and } u = 0 \text{ at } x = 1$$

can be expressed as $u(x,t) = \frac{1}{\sqrt{4\pi t}} \int e^{-|x-y|^2/4t}\phi(y)\,dy$. What is $\phi(y)$?

(b) The solution of

$$u_t = u_{xx} \quad \text{for } t > 0 \text{ and } 0 < x < 1, \text{ with } u = g(x) \text{ at } t = 0 \text{ and } u = 0 \text{ at } x = 0, 1$$

can be expressed as $u(x,t) = \sum_{n=1}^{\infty} a_n(t)\sin(n\pi x)$. Find $a_n(t)$ in terms of g.

3) This problem concerns the arrival time at the boundary, for a random walker solving $dy = f\,dt + g\,dw$ on the interval $[a,b]$.

(a) Let $G(x,y,t)$ be the probability, starting from x at time 0, of being at y at time t without having yet hit the boundary. What version of the forward Kolmogorov equation does G solve?

(b) Express, as an integral involving G_t, the "first passage time density to the boundary," i.e. the probability that the process, starting from $a < x < b$, first hits the boundary at time t.

(c) Using your answers to (a) and (b) and some further manipulation, show that

$$\text{first passage time density to the boundary} = -\tfrac{1}{2}\,\frac{\partial}{\partial y}(g^2 G(x,y,t))\Big|_{y=b} + \tfrac{1}{2}\,\frac{\partial}{\partial y}(g^2 G(x,y,t)\Big|_{y=a}.$$

4) Consider the following version of the Merton asset allocation problem:

- There is a risk-free asset, whose price satisfies $dp_1 = rp_1 ds$.
- There are two risky assets, whose prices p_2 and p_3 satisfy $dp_i = \mu_i p_i ds + \sigma_i p_i dw_i$ for $i = 2, 3$. We assume for simplicity that w_2 and w_3 are independent Brownian motions.
- Your controls are $\alpha_i(s) =$ the fraction of your wealth invested in asset i at time s, $i = 1, 2, 3$; note that $\alpha_1 + \alpha_2 + \alpha_3 = 1$.
- There is no consumption, and your goal is to optimize your expected utility of wealth at a predetermined time T. Your utility function is h.

Answer the following:

(a) What stochastic differential equation describes the evolution of your total wealth?

(b) Define an appropriate value function $u(x,t)$.

(c) Specify the Hamilton-Jacobi-Bellman equation and final-time condition u should satisfy.

(d) How does the value function determine the optimal asset allocations α_i?

5) In pricing a perpetual American put, we considered an underlying satisfying $dy = \mu y ds + \sigma y dw$ and the goal was to evaluate $\max_\tau E_{y(0)=x}\left[e^{-r\tau}(K - y(\tau))_+\right]$. Show that if v is a differentiable function with $v \geq (K-x)_+$ and $-rv + \mu x v_x + \frac{1}{2}\sigma^2 x^2 v_{xx} \leq 0$ for all x then v gives an upper bound:

$$E_{y(0)=x}\left[e^{-r\tau}(K - y(\tau))_+\right] \leq v(x)$$

for any bounded, nonanticipating stopping time τ.

6) This is a variant of the Bertsimas-Kogan-Lo least-squares-replication problem.considered in Section 7. It differs from the version in the notes in two ways: (i) the underlying has stochastic volatility; and (ii) the goal is not least-square replication but rather maximizing the utility of final-time wealth.

The underlying is a stock which can be traded at discrete times $i\Delta t$. Its price P_i and volatility σ_i at the ith time satisfy

$$\begin{aligned} \sigma_{i+1} &= \sigma_i + f(\sigma_i)\Delta t + g(\sigma_i)\phi_i\sqrt{\Delta t} \\ P_{i+1} &= P_i + \sigma_i P_i \psi_i \sqrt{\Delta t} \end{aligned}$$

where f and g are specified functions and ψ_i, ϕ_i are independent standard Gaussians (with mean 0 and variance 1).

You have sold an option on this stock with payoff $F(P_N)$, receiving cash V_0 in payment. Your goal is to invest this cash wisely, trading in a self-financing way, to maximize the expected utility of your final-time wealth $E[h(V_N - F(P_N))]$. Here h is your utility.

(a) Set this up as a discrete-time optimal control problem. What are the state variables? What is the control? Define an appropriate value function (call it J_i) at time $i\Delta t$. Be sure to specify the arguments of J_i, i.e. the variables it depends on.

(b) What is the value of J_N?

(c) Give a recursion relation that specifies J_i in terms of J_{i+1} for $i < N$.

7) Consider scaled Brownian motion with jumps: $dy = \sigma dw + J dN$, starting at $y(0) = x$. Assume the jump occurences are Poisson with rate λ, and the jumps have mean 0 and variance δ^2.

(a) Find $E[y^2(T)]$. (Hint: for a Poisson process with rate λ, the expected number of arrivals by time T is λT.)

(b) What backward Kolmogorov equation does part (a) solve?

PDE for Finance Notes, Spring 2003 – Stochastic Calculus Review

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use in connection with the NYU course PDE for Finance, G63.2706.

These notes provide a quick review of stochastic calculus. They explain somewhat more than you need to know for the course PDE for Finance. Still, if you have trouble understanding these notes you probably don't have sufficient background for PDE for Finance. To check whether you have a working knowledge of Ito's formula, try the exercises at the end.

The material presented here is mostly covered in the book by Neftci, at about the same (not-fully-rigorous) level. Deeper treatments of these topics can be found in many places, for example Richard Durett's books *Probability: theory and examples* and *Stochastic calculus: a practical introduction*, or J. Michael Steele's book *Stochastic calculus and financial applications*.

Brownian motion. Brownian motion $w(t)$ is the stochastic process with the following properties:

- For $s < t$ the increment $w(t) - w(s)$ is Gaussian with mean zero and variance $E[(w(t) - w(s))^2] = t - s$. Moreover the increments associated with disjoint intervals are independent.
- Its sample paths are continuous, i.e. the function $t \mapsto w(t)$ is (almost surely) continuous.
- It starts at 0, in other words $w(0) = 0$.

This process is unique (up to a suitable notion of equivalence). One "construction" of Brownian motion obtains it as the limit of discrete-time random walks; students of finance who have considered the continuous-time limit of a binomial lattice have seen something very similar.

The sample paths of Brownian motion, though continuous, are non-differentiable. Here is an argument that proves a little less but captures the main point. Given any interval (a, b), divide it into subintervals by $a = t_1 < t_2 \ldots < t_N = b$. Clearly

$$\sum_{i=1}^{N-1} |w(t_{i+1}) - w(t_i)|^2 \le \max_i |w(t_{i+1}) - w(t_i)| \cdot \sum_{i=1}^{N-1} |w(t_{i+1}) - w(t_i)|.$$

As $N \to \infty$, the left hand side has expected value $b-a$ (independent of N). The first term on the right tends to zero (almost surely) by continuity. So the second term on the right must tend to infinity (almost surely). Thus the sample paths of w have unbounded total variation on any interval. One can show, in fact, that $|w(t)-w(s)|$ is of order $\sqrt{|t-s| \log\log 1/|t-s|}$ as $|t-s| \to 0$.

It's easy to construct, for any constant $\sigma > 0$, a process whose increments are mean-value-zero, independent, and variance $\sigma^2|t-s|$: just use $\sigma w(t)$. The vector-valued version of this construction is more interesting. We say $w(t) = (w_1, \ldots, w_n)$ is an R^n-valued Brownian motion if its components are *independent* scalar Brownian motions. Thus

$E[(w(t) - w(s))_i (w(t) - w(s))_j]$ equals 0 if $i \neq j$ and $|t - s|$ if $i = j$. Given such w, we can obtain a process with correlated increments by taking linear combinations, i.e. by considering $z(t) = Aw(t)$ where A is a (constant) matrix. Its covariance is $E[(z(t) - z(s))_i (z(t) - z(s))_j] = (AA^T)_{ij}|t - s|$. If the desired variance σ is a function of state and time (deterministic, or random but nonanticipating) then construction of the associated process requires solving the stochastic differential equation $dx = \sigma dw$ (to be discussed below). That's the scalar case; the vector-valued situation is similar: to construct a process with independent, mean-value-zero increments with specified covariance Σ we have only to set $A = \sqrt{\Sigma}$ (the unique nonnegative, symmetric square root of Σ) and solve $dx = Adw$.

Filtrations and conditional expectations. It is important, in discussing stochastic processes, to remember that at time t one knows (with certainty) only the past and the present, not the future. This is important for understanding the term "martingale." It will also be crucial later in the class when we discuss optimal decision-making.

The meaningful *statements* about a Brownian motion (or any stochastic process, for that matter) are statements about its values at various times. Here is an example of a statement: "$-3 < w(.5) < -2$ and $w(1.4) > 3$". Here is another: "$\max_{0 \leq t \leq 1} |w(t)| < 3$". A statement is either true or false for a given sample path; it has a certain probability of being true. We denote by $\mathcal{F}_t$ the set of all statements about w that involve only the values of w up to time t. Obviously $\mathcal{F}_s \subset \mathcal{F}_t$ if $s < t$. These $\mathcal{F}_t$'s are called the *filtration* associated with w.

We can also consider *functions* of a Brownian path. When we take the expected value of some expression involving Brownian motion we are doing this. Here are some examples of functions: $f[w] = w(.5) - w(1)^2$; $g[w] = \max_{0 \leq t \leq 1} |w(t)|$. Notice that both these examples are determined entirely by time-1 information (jargon: f and g are $\mathcal{F}_1$-measureable). It's often important to discuss the expected value of some uncertain quantity given the information available at time t. For example, we may wish to know the expected value of $\max_{0 \leq t \leq 1} |w(t)|$ given knowledge of w only up to time .5. This is a *conditional expectation*, sometimes written $E_t[g] = E[g|\mathcal{F}_t]$ (in this case t would be .5). We shall define it in a moment via orthogonal projection. This definition is easy but not so intuitive. After giving it, we'll explain why the definition captures the desired intuition.

Let V be the vector space of all functions $g[w]$, endowed with the inner product $\langle f, g \rangle = E[fg]$. It has subspaces

$$V_t = \text{space of functions whose values are determined by time-t information.}$$

The conditional expectation is defined by orthogonal projection:

$$E_t[g] = \text{orthogonal projection of } g \text{ onto } V_t.$$

The standard linear-algebra definition of orthogonal projection characterizes $E_t[g]$ as the unique element of V_t such that

$$\langle E_t[g], f \rangle = \langle g, f \rangle \text{ for all } f \in V_t.$$

Rewriting this in terms of expectations: $E_t[g]$ is the unique function in V_t such that

$$E\left[E_t[g]f\right] = E[gf] \text{ for all } f \in V_t.$$

All the key properties of conditional expectation follow easily from this definition. Example: "tower property"

$$s < t \Longrightarrow E_s\left[E_t[f]\right] = E_s[f]$$

since projecting first to V_t then to $V_s \subset V_t$ is the same as projecting directly to V_s. Another fact: E_0 is the ordinary expectation operator E. Indeed, V_0 is one-dimensional (its elements are functions of a single point $w(0) = 0$, i.e. it consists of those functions that aren't random at all). From the definition of orthogonal projection we have

$$E_0[g] \in V_0 \text{ and } E\left[E_0[g]f\right] = E[gf] \text{ for all } f \in V_0.$$

But when f is in V_0 it is deterministic, so $E[gf] = fE[g]$. Similarly $E\left[E_0[g]f\right] = fE_0[g]$. Thus $E_0[g] = E[g]$.

To see that this matches our intuition, i.e. that E_t is properly interpreted as "the expected value based on future randomness, given all information available at time t", let's consider the simplest possible discrete-time analogue. Consider a 2-stage coin-flipping process which obtains at each stage heads (probability p) or tails (probability $q = 1 - p$). We visualize it using a (nonrecombinant) binomial tree, numbering the states as shown in Figure 1.

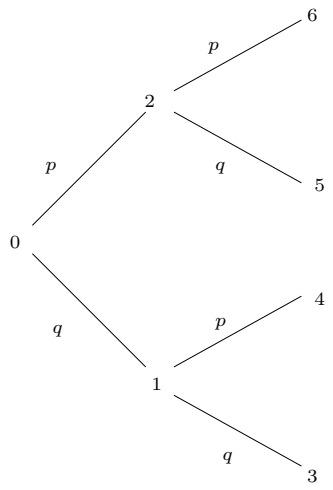

Figure 1: *Binomial tree for visualizing conditional probabilities*

The space V_2 is 4-dimensional; its functions are determined by the full history, i.e. they can be viewed as functions of the time-2 nodes (numbered $3, 4, 5, 6$ in the figure). The space V_1 is two-dimensional; its functions are determined by just the first flip. Its elements can be viewed as functions of the time-1 nodes (numbered $1, 2$ in the figure); or, equivalently, they are functions $f \in V_2$ such that $f(3) = f(4)$ and $f(5) = f(6)$. (Such an function can be

viewed as a function of the time-1 nodes by setting $f(1) = f(3) = f(4)$ and $f(2) = f(5) = f(6)$. The "expected value of g given time-1 information" intuitively has values

$$\tilde{E}_1[g](1) = pg(4) + qg(3), \quad \tilde{E}_1[g](2) = pg(6) + qg(5).$$

To check that this agrees with our prior definition, we must verify that $\langle f, \tilde{E}_1[g]\rangle = \langle f, g\rangle$ for all $f \in V_1$. In other words we must check that

$$E\left[\tilde{E}_1[g]f\right] = E\left[gf\right] \tag{1}$$

whenever $f(2) = f(5) = f(6)$ and $f(1) = f(3) = f(4)$. The left hand side is

$$q\tilde{E}_1[g](1)f(1) + p\tilde{E}_1[g](2)f(2)$$

while the right hand side is

$$q^2 f(3)g(3) + pqf(4)g(4) + pqf(5)g(5) + p^2 f(6)g(6)$$

which can be rewritten (since $f(1) = f(3) = f(4)$ and $f(2) = f(5) = f(6)$) as

$$q(qg(3) + pg(4))f(1) + p(qg(5) + pg(6))f(2).$$

The formula given above for $\tilde{E}_1[g]$ is precisely the one that makes (1) correct.

A stochastic process $x(t)$ is "adapted" to $\mathcal{F}_t$ if its values up to and including time t are determined by the statements in $\mathcal{F}_t$. (The stochastic processes obtained from Brownian motion by solving stochastic differential equations automatically have this property.) Such a stochastic process is called a *martingale* if $E_s[x(t)] = x(s)$ for $s < t$. An equivalent statement: $E_s[x(t) - x(s)] = 0$ for $s < t$. Intuitively: given current information, there's no point betting on the future of the process; it's equally likely to go up or down. (That's not quite right; it confuses the mean and the median. The correct statement is that the expected future value, based on present information, is exactly the present value.)

A stochastic process $f(t)$ is called *nonanticipating* if its value at time t depend onlys on information available at time t, i.e. if $f(t)$ is adapted to $\mathcal{F}_t$. An example is $f(t) = F(t, w(t))$ for any (deterministic) function $F : R^2 \to R$. But this isn't the only type of example – for example $f(t) = \int_0^t w(s)\, ds$ is also nonanticipating.

Stochastic integrals. We are interested in stochastic differential equations of the type

$$dy = f(y, s)ds + g(y, s)dw, \quad y(t) = x.$$

(Pretty much everything we'll say extends straightforwardly to SDE's of the form $dy = fds + gdw$ with f and g random but nonanticipating.) The stochastic differential equation is really shorthand for the associated integral equation

$$y(b) = x + \int_t^b f(y(s), s)ds + \int_t^b g(y(s), s)dw. \tag{2}$$

To understand what this means we must understand the two integrals on the right.

The first one is relatively easy. If y is continuous in s then

$$\int_t^b f(y(s), s)ds$$

makes perfect sense as a Riemann integral. (All the processes considered in this review do have y continuous in s, so this hypothesis is OK. The case when f is random but nonanticipating is more subtle; the successful treatment is similar to the one of $\int g dw$ explained below.)

The second "stochastic" integral is more subtle. The proper interpretation is this: for any random but nonanticipating integrand $g(s)$,

$$\int_a^b g dw = \lim_{\Delta t \to 0} \sum_{i=1}^{N-1} g(t_i)[w(t_{i+1}) - w(t_i)] \tag{3}$$

with the notation $a = t_1 < t_2 < \ldots < t_N = b$ and $\Delta t = \max_i |t_{i+1} - t_i|$. (We may, but we don't have to, choose the t_i's equally spaced.) The important point is that we evaluate g at the beginning of the increment. We'll show presently that making the opposite choice

$$\sum_{i=1}^{N-1} g(t_{i+1})[w(t_{i+1}) - w(t_i)]$$

would give a *different* answer. Thus the stochastic integral is not a Riemann integral, but something different.

A key property of the stochastic integral is immediately clear: since g is nonanticipating,

$$E_a \int_a^b g dw = 0 \tag{4}$$

because each term in the sum has

$$E_{t_i} \left[g(t_i) \left(w(t_{i+1}) - w(t_i) \right) \right] = 0$$

(since $w(t_{i+1}) - w(t_i)$ is independent of all time-t_i information, hence independent of $g(t_i)$). Therefore by the tower property $E_a \left[g(t_i)[w(t_{i+1}) - w(t_i)] \right] = 0$, and summing gives (4). Remembering the definition of a martingale, (4) says the solution of a stochastic differential equation of the form $dy = g dw$ (with no dt term on the right) is a martingale.

What kind of limit do we mean in (3)? The mean-square kind. If a sequence of functions $\phi_n(x)$ is defined for $x \in (0, 1)$, one says $\phi = \lim_{n \to \infty} \phi_n$ in the mean-square sense if $\int_0^1 |\phi_n(x) - \phi(x)|^2 \, dx \to 0$. The situation for the stochastic integral is similar, except the integral is replaced by expectation:

$$E \left[\left(\int_a^b g dw - \sum_{i=1}^{N-1} g(t_i)[w(t_{i+1}) - w(t_i)] \right)^2 \right] \to 0.$$

We won't prove the existence of this limit in any generality. Instead let's do a simple example – which displays many of the essential ideas of the general case. Specifically: let's show that

$$\int_a^b w\,dw = \tfrac{1}{2}w^2(b) - \tfrac{1}{2}w^2(a) - (b-a)/2.$$

Notice that this is *different* from the formula you might have expected based on elementary calculus ($wdw \neq \frac{1}{2}w^2$). The calculus rule is based on Chain Rule, whereas in the stochastic setting we must use Ito's formula – as we'll explain presently. If I skip too many details, you'll find a slower treatment in Neftci (in the first edition the relevant pages are 179-184).

According to the definition, $\int_a^b w\,dw$ is the limit of

$$\sum_{i=1}^{N-1} w(t_i)(w(t_{i+1}) - w(t_i)).$$

A bit of manipulation shows that this is exactly equal to

$$\tfrac{1}{2}w^2(b) - \tfrac{1}{2}w^2(a) - \tfrac{1}{2}\sum_{i=1}^{N-1}(w(t_{i+1}) - w(t_i))^2,$$

so our assertion is equivalent to the statement

$$\lim_{\Delta t\to 0}\sum_{i=1}^{N-1}(w(t_{i+1}) - w(t_i))^2 = b - a. \tag{5}$$

In the Ito calculus we sometimes write "$dw \times dw = dt$;" when we do, it's basically shorthand for (5). Notice that each term $(w(t_{i+1}) - w(t_i))^2$ is random (the square of a Gaussian random variable with mean 0 and variance $t_{i+1} - t_i$). But in the limit the sum is deterministic, by a sort of law of large numbers. If you believe it's deterministic then the value is clear, since $S_N = \sum_{i=1}^{N-1}(w(t_{i+1}) - w(t_i))^2$ has expected value $b - a$ for any N.

To prove (5) in the mean-square sense, we must show that

$$E\left[(S_N - (b-a))^2\right] \to 0$$

as $N \to \infty$. Expanding the square, this is equivalent to

$$E\left[S_N^2 - (b-a)^2\right] \to 0.$$

Now,

$$\begin{aligned} E\left[S_N^2\right] &= E\left[\sum_{i=1}^{N-1}(w(t_{i+1}) - w(t_i))^2 \sum_{j=1}^{N-1}(w(t_{j+1}) - w(t_j))^2\right] \\ &= E\left[\sum_{i,j=1}^{N-1}(w(t_{i+1}) - w(t_i))^2(w(t_{j+1}) - w(t_j))^2\right]. \end{aligned}$$

The last term is easy to evaluate, using the properties of Brownian motion:

$$E\left[(w(t_{i+1}) - w(t_i))^2 (w(t_{j+1}) - w(t_j))^2\right] = (t_{i+1} - t_i)(t_{j+1} - t_j)$$

when $i \neq j$, and

$$E\left[(w(t_{i+1}) - w(t_i))^4\right] = 3(t_{i+1} - t_i)^2.$$

(The latter follows from the fact that $w(t_{i+1}) - w(t_i)$ is Gaussian with mean 0 and variance $t_{i+1} - t_i$.) We deduce after some manipulation that

$$\begin{aligned} E\left[S_N^2 - (b-a)^2\right] &= 2\sum_{i=1}^{N-1} (t_{i+1} - t_i)^2 \\ &\leq 2(\max_i |t_{i+1} - t_i|)(b-a) \end{aligned}$$

which does indeed tend to 0 as $\max_i |t_{i+1} - t_i| \to 0$.

We now confirm a statement made earlier, that the stochastic integral just defined is different from

$$\lim_{\Delta t \to 0} \sum_{i=1}^{N-1} w(t_{i+1})[w(t_{i+1}) - w(t_i)]. \tag{6}$$

Indeed, we have

$$\sum_{i=1}^{N-1} w(t_{i+1})[w(t_{i+1}) - w(t_i)] - \sum_{i=1}^{N-1} w(t_i)[w(t_{i+1}) - w(t_i)] = \sum_{i=1}^{N-1} [w(t_{i+1}) - w(t_i)]^2$$

which tends in the limit (we proved above) to $b-a$. Thus the alternative (wrong) definition (6) equals $\frac{1}{2}w^2(b) - \frac{1}{2}w^2(a) + \frac{1}{2}(b-a)$. If we had used this definition, the stochastic integral would not have been a martingale.

Stochastic differential equations. It's important to say that stochastic differential equations have solutions, under reasonable conditions on the form of the equation. Moreover the resulting stochastic process $y(s)$ has continuous sample paths (y is a continuous function of s). For the purposes of the course PDE for Finance you do *not* need to know how to prove this.

The Ito calculus. If $y(s)$ solves a stochastic differential equation, it's natural to seek a stochastic differential equation for $\phi(s, y(s))$ where ϕ is any smooth function. If y solved an ordinary differential equation we would obtain the answer using chain rule. When y solves a stochastic differential equation we must use the Ito calculus instead. It replaces the chain rule.

Let's first review the situation for ordinary differential equations. Suppose $dy/dt = f(y,t)$ with initial condition $y(0) = x$. It is a convenient mnemonic to write the equation in the form

$$dy = f(y,t)dt.$$

This reminds us that the solution is well approximated by its finite difference approximation $y(t_{i+1}) - y(t_i) = f(y(t_i), t_i)(t_{i+1} - t_i)$. Let us write

$$\Delta y = f(y,t)\Delta t$$

as an abbreviation for the finite difference approximation. (In this discussion Δ is always an increment, never the Laplacian.) The ODE satisfied by $z(t) = \phi(y(t))$ is, by chain rule, $dz/dt = \phi'(y(t))dy/dt$. The mnemonic for this is

$$d\phi = \phi' dy.$$

It reminds us of the proof, which boils down to the fact that (by Taylor expansion)

$$\Delta\phi = \phi'(y)\Delta y + \text{error of order } |\Delta y|^2.$$

In the limit as the time step tends to 0 we can ignore the error term, because $|\Delta y|^2 \le C|\Delta t|^2$ and the sum of these terms is of order $\max_i |t_{i+1} - t_i|$.

OK, now the stochastic case. Suppose y solves

$$dy = f(y,t)dt + g(y,t)dw.$$

Ito's lemma, in its simplest form, says that if ϕ is smooth then $z = \phi(y)$ satisfies the stochastic differential equation

$$dz = \phi'(y)dy + \tfrac{1}{2}\phi''(y)g^2 dt = \phi'(y)g dw + \left[\phi'(y)f + \tfrac{1}{2}\phi''(y)g^2\right] dt.$$

Here is a heuristic justification: carrying the Taylor expansion of $\phi(y)$ to second order gives

$$\begin{aligned} \Delta\phi &= \phi(y(t_{i+1})) - \phi(y(t_i)) \\ &= \phi'(y(t_i))[y(t_{i+1}) - y(t_i)] + \tfrac{1}{2}\phi''(y(t_i))[y(t_{i+1}) - y(t_i)]^2 + \text{error of order } |\Delta y|^3. \end{aligned}$$

So far we haven't cheated. It's tempting to write the last expression as

$$\phi'(y)(g\Delta w + f\Delta t) + \tfrac{1}{2}\phi''(y)g^2(\Delta w)^2 + \text{errors of order } |\Delta y|^3 + |\Delta w||\Delta t| + |\Delta t|^2$$

where $\phi'(y) = \phi'(y(t_i))$, $g = g(y(t_i), t_i)$, $\Delta w = w(t_{i+1}) - w(t_i)$, etc. (In other words: it's tempting to substitute $\Delta y = f\Delta t + g\Delta w$.) That's not quite right: in truth $\Delta y = y(t_{i+1}) - y(t_i)$ is given by a stochastic integral from t_i to t_{i+1}, and our cheat pretends that the integrand is constant over this time interval. But fixing this cheat is a technicality – much as it is in the deterministic setting – so let's proceed as if the last formula were accurate. I claim that the error terms are negligible in the limit $\Delta t \to 0$. This is easy to see for the $|\Delta t|^2$ terms, since

$$\sum_i (t_{i+1} - t_i)^2 \le \max_i |t_{i+1} - t_i| \sum_i |t_{i+1} - t_i|$$

A similar argument works for the $|\Delta t||\Delta w|$ terms. The $|\Delta y|^3$ term is a bit more subtle; we'll return to it presently. Accepting this, we have

$$\Delta\phi \approx \phi'(y)(g\Delta w + f\Delta t) + \tfrac{1}{2}\phi''(y)g^2(\Delta w)^2.$$

Now comes the essence of the matter: we can replace $(\Delta w)^2$ by Δt. A more careful statement of this assertion: if $a = t_1 < t_2 < \ldots < t_N = b$ then

$$\lim_{\Delta t \to 0} \sum_{i=1}^{N-1} h(t_i)[w(t_{i+1}) - w(t_i)]^2 = \int_a^b h(t)dt \tag{7}$$

if h is non-anticipating. Notice: we don't claim that $h(\Delta w)^2$ is literally equal to $h\Delta t$ for any single time interval, no matter how small. Rather, we claim that once the contributions of *many* time intervals are combined, the fluctuations of w cancel out and the result is an integral dt. We proved (7) in the case $h = 1$; the general case is more technical, of course, but the ideas are similar.

We skipped over why the $|\Delta y|^3$ error terms can be ignored. The reason is that they're controlled by

$$\max_i |y(t_{i+1}) - y(t_i)| \sum_i |y(t_{i+1}) - y(t_i)|^2.$$

The argument above shows that the sum is finite. Since $y(t)$ is continuous, $\max_i |y(t_{i+1}) - y(t_i)|$ tends to zero. So this term is negligible.

The same logic applies more generally, when w is a vector-valued Brownian motion, y is vector-valued, and ϕ is a function of time as well as y. The only new element (aside from some matrix algebra) is that the quadratic terms in Δw are now of the form

$$\Delta w_j \Delta w_k = [w_j(t_{i+1}) - w_j(t_i)][w_k(t_{i+1}) - w_k(t_i)].$$

An argument very much like the proof of (5) shows that

$$\lim_{\Delta t \to 0} \sum_{i=1}^{N-1} [w_j(t_{i+1}) - w_j(t_i)][w_k(t_{i+1}) - w_k(t_i)] = \begin{cases} 0 & \text{if } j \neq k \\ b - a & \text{if } j = k, \end{cases}$$

which justifies (at the same heuristic level as our scalar treatment) the rule that $\Delta w_j \Delta w_k$ should be replaced by dt when $j = k$, and 0 when $j \neq k$.

Examples. Knowing the preceding arguments is less important than having some facility with the actual application of Ito's lemma. Here are some basic examples.

Calculating $\int w dw$. Above we showed, by brute force calculation based on the definition of a stochastic integral, that

$$\int_a^b w\, dw = \tfrac{1}{2}(w^2(b) - w^2(a)) - \tfrac{1}{2}(b - a).$$

Ito's lemma gives a much easier proof of the same result: applying it to $\phi(w) = w^2$ gives

$$d(w^2) = 2wdw + dwdw = 2wdw + dt$$

which means that $w^2(b) - w^2(a) = 2\int_a^b w\, dw + (b - a)$.

Log-normal dynamics. Suppose

$$dy = \mu(t)ydt + \sigma(t)ydw \tag{8}$$

where $\mu(t)$ and $\sigma(t)$ are (deterministic) functions of time. What stochastic differential equation describes $\log y$? Ito's lemma gives

$$\begin{aligned} d(\log y) &= y^{-1}dy - \tfrac{1}{2}y^{-2}dydy \\ &= \mu(t)dt + \sigma(t)dw - \tfrac{1}{2}\sigma^2(t)dt. \end{aligned}$$

Remembering that $y(t) = e^{\log y(t)}$, we see that

$$y(t_1) = y(t_0)e^{\int_{t_0}^{t_1}(\mu-\sigma^2/2)ds + \int_{t_0}^{t_1}\sigma dw}.$$

In particular, if μ and σ are constant in time we get

$$y(t_1) = y(t_0)e^{(\mu-\sigma^2/2)(t_1-t_0)+\sigma(w(t_1)-w(t_0))}.$$

Stochastic stability. Consider once more the solution of (8). It's natural to expect that if μ is negative and σ is not too large then y should tend (in some average sense) to 0. This can be seen directly from the solution formula just derived. But an alternative, instructive approach is to consider the second moment $\rho(t) = E[y^2(t)]$. From Ito's formula,

$$d(y^2) = 2ydy + dydy = 2y(\mu ydt + \sigma ydw) + \sigma^2 y^2 dt.$$

Taking the expectation, we find that

$$E[y^2(t_1)] - E[y^2(t_0)] = \int_{t_0}^{t_1}(2\mu + \sigma)E[y^2]ds$$

or in other words

$$d\rho/dt = (2\mu + \sigma)\rho.$$

Thus $\rho = E[y^2]$ can be calculated by solving this deterministic ODE. If the solution tends to 0 as $t \to \infty$ then we conclude that y tends to zero in the mean-square sense. When μ and σ are constant this happens exactly when $2\mu + \sigma < 0$. When they are functions of time, the condition $2\mu(t) + \sigma(t) \le -c$ is sufficient (with $c > 0$) since it gives $d\rho/dt \le -c\rho$.

An example related to Girsanov's theorem. For any (deterministic or non-anticipating) function $\gamma(s)$,

$$E\left[e^{\int_a^b \gamma(s)dw - \frac{1}{2}\int_a^b \gamma^2(s)ds}\right] = 1.$$

In fact, this is the expected value of $e^{z(b)}$, where

$$dz = -\tfrac{1}{2}\gamma^2(t)dt + \gamma(t)dw, \quad z(a) = 0.$$

Ito's lemma gives

$$d(e^z) = e^z dz + \tfrac{1}{2} e^z dz dz = e^z \gamma dw.$$

So

$$e^{z(b)} - e^{z(a)} = \int_a^b e^z \gamma dw.$$

The right hand side has expected value zero, so

$$E[e^{z(b)}] = E[e^{z(a)}] = 1.$$

Notice the close relation with the previous example "lognormal dynamics": all we've really done is identify the conditions under which $\mu = 0$ in (8).

[Comment for those who know about risk-neutral pricing: this example is used in the discussion of Girsanov's theorem, which gives the relation between the "subjective" and "risk-neutral" price processes. The expression

$$e^{\int_a^b \gamma(s) dw - \frac{1}{2}\int_a^b \gamma^2(s) ds}$$

is the Radon-Nikodym derivative relating the associated measures on path space. The fact that it has expected value 1 reflects the fact that both measures are probability measures.]

The Ornstein-Uhlenbeck process. You should have learned in calculus that the deterministic differential equation $dy/dt + Ay = f$ can be solved explicitly when A is constant. Just multiply by e^{At} to see that $d(e^{At}y)/dt = e^{At}f$ then integrate both sides in time. So it's natural to expect that linear stochastic differential equations can also be solved explicitly. We focus on one important example: the "Ornstein-Uhlenbeck process," which solves

$$dy = -cydt + \sigma dw, \quad y(0) = x$$

with c and σ constant. (This is *not* a special case of (8), because the dw term is not proportional to y.) Ito's lemma gives

$$d(e^{ct}y) = ce^{ct}ydt + e^{ct}dy = e^{ct}\sigma dw$$

so

$$e^{ct}y(t) - x = \sigma \int_0^t e^{cs} dw,$$

or in other words

$$y(t) = e^{-ct}x + \sigma \int_0^t e^{c(s-t)} dw(s).$$

Now observe that $y(t)$ is a Gaussian random variable – because when we approximate the stochastic integral as a sum, the sum is a linear combination of Gaussian random variables. (We use here that a sum of Gaussian random variables is Gaussian; also that a limit of Gaussian random variables is Gaussian.) So $y(t)$ is entirely described by its mean and variance. They are easy to calculate: the mean is

$$E[y(t)] = e^{-ct}x$$

since the "dw" integral has expected value 0. To calculate the variance let us accept for a moment the formula

$$E\left[\left(\int_a^b g(y(s),s)\,dw\right)^2\right] = \int_a^b E[g^2(y(s),s)]\,ds. \tag{9}$$

Using this, the variance of the Ornstein-Uhlenbeck process is easily determined:

$$\begin{aligned} E\left[(y(t)-E[y(t)])^2\right] &= \sigma^2 E\left[\left(\int_0^t e^{c(s-t)}dw(s)\right)^2\right] \\ &= \sigma^2\int_0^t e^{2c(s-t)}ds \\ &= \sigma^2\frac{1-e^{-2ct}}{2c}. \end{aligned}$$

The justification of the formula (9) is easy. Just approximate the stochastic integral as a sum. The square of the stochastic integral is approximately

$$\left(\sum_{i=1}^{N-1} g(y(s_i),s_i)[w(s_{i+1})-w(s_i)]\right)\left(\sum_{j=1}^{N-1} g(y(s_j),s_j)[w(s_{j+1})-w(s_j)]\right)$$
$$= \sum_{i,j=1}^{N-1} g(y(s_i),s_i)g(y(s_j),s_j)[w(s_{i+1})-w(s_i)][w(s_{j+1})-w(s_j)] \quad .$$

For $i \neq j$ the expected value of the i,jth term is 0 (for example, if $i<j$ then $[w(s_{j+1})-w(s_j)]$ has mean value 0 and is independent of $g(y(s_i),s_i)$, $g(y(s_j),s_j)$, and $[w(s_{i+1})-w(s_i)]$). For $i=j$ the expected value of the i,jth term is $E[g^2(y(s_i),s_i)][s_{i+1}-s_i]$. So the expected value of the squared stochastic integral is approximately

$$\sum_{i=1}^{N-1} E[g^2(y(s_i),s_i)][s_{i+1}-s_i],$$

and passing to the limit $\Delta s \to 0$ gives the desired assertion.

We close this example with a brief discussion of the relevance of the Ornstein-Uhlenbeck process. One of the simplest interest-rate models in common use is that of Vasicek, which supposes that the (short-term) interest rate $r(t)$ satisfies

$$dr = a(b-r)dt + \sigma dw$$

with a, b, and σ constant. Interpretation: r tends to revert to some long-term average value b, but noise keeps perturbing it away from this value. Clearly $y = r-b$ is an Ornstein-Uhlenbeck process, since $dy = -aydt + \sigma dw$. Notice that $r(t)$ has a positive probability of being negative (since it is a Gaussian random variable); this is a reminder that the Vasicek model is not very realistic. Even so, its exact solution formulas provide helpful intuition.

Historically, the Ornstein-Uhlenbeck process was introduced by physicists Ornstein and Uhlenbeck, who believed that a diffusing particle had brownian *acceleration* not brownian velocity. Their idea was that the position $x(t)$ of the particle at time t should satisfy

$$\begin{aligned} dx &= vdt \\ \epsilon dv &= -vdt + dw \end{aligned}$$

with $\epsilon > 0$ small. As $\epsilon \to 0$, the resulting $x_\epsilon(t)$ converges to a brownian motion process. Formally: when $\epsilon = 0$ we recover $0 = -vdt + dw$ so that $dx = (dw/dt)dt = dw$. Honestly: we claim that $|x_\epsilon(t) - w(t)|$ converges to 0 (uniformly in t) as $\epsilon \to 0$. In fact, writing the equations for the Ornstein-Uhlenbeck process as

$$\begin{aligned} dx_\epsilon &= v_\epsilon dt \\ dw &= v_\epsilon dt + \epsilon dv_\epsilon \end{aligned}$$

then subtracting, we see that

$$d(x_\epsilon - w) = \epsilon dv_\epsilon.$$

Now use our explicit solution formula for the Ornstein Uhlenbeck process to represent v_ϵ in terms of stochastic integrals, ultimately concluding that $\epsilon v_\epsilon(t) \to 0$ as $\epsilon \to 0$. (Details left to the reader.)

Some Exercises. The material in these notes goes beyond what we'll actually use in PDE for Finance. But if you have trouble with the following problems, you probably aren't prepared for the class.

Problem 1. Let $s(t) = s(0)e^{\mu t + \sigma w(t)}$, where μ and σ are constant and w is Brownian motion.

(a) Show, using Ito's formula, that $ds = (\mu + \frac{1}{2}\sigma^2)sdt + \sigma s dw$.

(b) Conclude that $E[s(t)] - E[s(0)] = (\mu + \frac{1}{2}\sigma^2)\int_0^t E[s(\tau)]d\tau$, where E denotes expected value.

(c) Conclude that $E[s(t)] = s(0)e^{(\mu + \frac{1}{2}\sigma^2)t}$.

[Comment: taking $t = 1$, this gives a proof of the lemma – more commonly proved by direct integration – that if X is Gaussian with mean μ and standard deviation σ then $E[e^X] = e^{\mu + \sigma^2/2}$.]

Problem 2. Let's use Ito's formula to calculate the moments of a Gaussian. Set

$$\beta_k(t) = E[w^k(t)],$$

where $w(t)$ is Brownian motion (with $w(0) = 0$).

(a) Show, using Ito's formula, that for $k = 2, 3, \ldots$,

$$\beta_k(t) = \frac{1}{2}k(k-1)\int_0^t \beta_{k-2}(s)\, ds.$$

(b) Deduce that $E[w^4(t)] = 3t^2$. What is $E[w^6(t)]$?

[Comment: the moments of w can also be calculated by direct integration, since $w(t)$ is Gaussian with mean 0 and variance 1.]

Problem 3. Let $\mu(t)$ and $\sigma(t)$ be deterministic functions of time. Consider the associated lognormal process, i.e. the solution of

$$ds = \mu(t)sdt + \sigma(t)sdw$$

with initial condition $s(0) = s_0$. Show that $\log s(t)$ is Gaussian, with mean

$$E[\log s(t)] = \log s_0 + \int_0^t [\mu(\tau) - \tfrac{1}{2}\sigma^2(\tau)]\, d\tau$$

and variance

$$E[(\log s(t))^2] - (E[\log s(t)])^2 = \int_0^t \sigma^2(\tau)\, d\tau.$$