PDE for Finance Notes – 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, Spring 1999.

Applications to finance: optimal stopping and option pricing. These notes discuss several topics, each a central application of stochastic pde methods to finance:

(a) **When to sell an asset.** We've seen that the expected final-time payoff, at a fixed maturity time, is determined by solving a linear PDE. But what if the final time is not fixed, but is rather a "stopping time" to be chosen optimally? We considered some deterministic problems of this type in Section 3 and HW3. We'll do a continuous-time example here. It leads to a *free boundary problem* for the backward Kolmogorov equation.

(b) **European options, with and without transaction costs**. The theory developed by Black-Scholes-Merton determines a price for any European option by identifying an equivalent replicating portfolio. However this theory ignores transaction costs, and it requires the market to be complete. We'll discuss how it can be embedded in a more general framework that permits taking transaction costs into account and can be used for almost any market model, whether or not it is complete.

(c) **Risk-neutral pricing**. We've noticed that the Black-Scholes PDE describes the expected payoff of a European option – provided we take the "expected payoff" using a suitable "risk-neutral" stochastic differential equation for the underlying asset prices. We'll explain this, discussing (well, really just stating) Girsanov's theorem in the process.

(d) **American options**. When viewed via risk-neutral pricing, pricing an American option amounts to an optimal stopping problem. The price is therefore described by a free boundary problem a lot like that of topic (a).

Concerning (a) we follow Oksendal, Examples 10.2.2 and 10.4.2. Concerning (b) we follow the chapter by Soner in *Viscosity Solutions and Applications,* Lecture Notes in Math 1660, Springer-Verlag, section 4. Discussions of topic (c) and (d) can be found in many places, including Baxter-Rennie, Neftci, and Oksendal; our treatment is therefore very abbreviated.

The methods under discussion here are applicable in considerable generality – for markets with many sources of randomness (described by vector-valued solutions of stochastic differential equations), and for non-constant interest rates, drifts, and volatilities. However for simplicity we'll focus on the simplest case when the interest rate, drift, and volatility are constant.
Since today is the last lecture and this material is not reinforced by homework, the topics covered here will not be on the final exam.

****

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. The decision to sell may depend on the stock price, and in principle on all information about the stock price history – but not on knowledge of the future. Thus the sales time τ is random but non-anticipating, i.e. it is a *stopping time*. We plan to use the method of dynamic programming, so it is natural to formulate the problem with an arbitrary initial time and initial state (but with the objective always discounted to time 0). Our goal is thus to find

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

where the maximization is over all stopping times.

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. (Exercise: prove the assertions in this paragraph.)

Our plan is a lot like the one we used for other optimal control problems: we shall guess, using a combination of rigorous and heuristic arguments, the optimal stopping rule. Then we'll prove our guess is right by a suitable verification argument.

We naturally expect that

$$u(x,t) \geq e^{-rt}(x - a)$$

since one possible strategy is to sell immediately. Moreover it is optimal to sell immediately (at time t) exactly if $u = e^{-rt}(x - a)$. By the principle of dynamic programming, we should consider, at each moment $s > t$, whether to sell immediately or hold longer. Thus the optimal stopping rule should have the form

$$\text{sell when } (y(s), s) \text{ leaves the set } H$$

where H is the "hold" region

$$H = \{(x,t) \ : \ u(x,t) > e^{-rt}(x - a)\}.$$

We claim that H is independent of t, and it really has the form

$$H = \{(x,t) \ : \ (0 < x < h)\} \tag{2}$$

for some "selling threshold" h. To see why H is independent of t, observe that

$$\begin{aligned} u(x,t) &= e^{-rt} \max_\tau E_{y(t)=x}\left[e^{-r(\tau-t)}(y(\tau)-a)\right] \\ &= e^{-rt}\tilde{u}(x) \end{aligned}$$

where $\tilde{u}$ is the optimal payoff discounted to the starting time (which is therefore independent of the starting time). Thus $u > e^{-rt}(x-a)$ exactly if $\tilde{u} > x-a$. So the decision whether to "sell immediately" or "hold longer" depends only on the initial stock price x. It's natural to expect that this dependence is through a *sales threshold* h, i.e. that the set where $\tilde{u} > x-a$ is an interval. Rather than prove this now, consider it a guess to be verified later.

If h were known then $u(x,t)$ would be fully determined as the solution to an *exit-time* problem similar to those discussed in Section 5. There we considered

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

where τ was the exit time from a domain D. We saw that v solves

$$v_t + \mathcal{L}v + \Psi = 0 \text{ for } x \in D, \quad v = \Phi \text{ for } x \in \partial D$$

where $\mathcal{L}$ is the infinitesimal generator of the SDE. (The discussion in Section 5 had a fixed maturity time T, and set $\tau = T$ if the exit time was greater than T. However if the expected exit time is finite then we can pass to the limit $T \to \infty$. This is equivalent to solving an elliptic boundary value problem for $\tilde{u}$, as will be clear presently.)

The special case of interest here is $D = (0,h)$, $\Phi(y,s) = e^{-rs}(y-a)$, $\Psi = 0$. Writing $u^h(x,t)$ for the expected final cost, we deduce that

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

and

$$u^h(x,t) = e^{-rt}(x-a) \quad \text{at } x = h.$$

(We do not impose a boundary condition at $x = 0$ because geometric Brownian motion never reaches 0.)

Let's find u^h explicitly. We showed above that $u(x,t) = e^{-rt}\tilde{u}(x)$, and the same argument shows that $u^h(x,t) = e^{-rt}\tilde{u}^h(x)$. The PDE for $\tilde{u}^h$ is evidently

$$-r\tilde{u}^h + \mu x \tilde{u}_x^h + \tfrac{1}{2}\sigma^2 x^2 \tilde{u}_{xx}^h = 0 \quad \text{for } 0 < x < h$$

with

$$\tilde{u}^h = (x-a) \quad \text{at } x = h.$$

The general solution of $-r\phi + \mu x\phi_x + \frac{1}{2}\sigma^2 x^2 \phi_{xx} = 0$ 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 $\tilde{u}^h$ we must specify C_1 and C_2. Since $\tilde{u}^h$ 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

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

Any sales threshold is permitted, of course, so we should optimize over h. One verifies by direct calculation that the optimal threshold is

$$h_{\text{opt}} = \frac{a\gamma_1}{\gamma_1 - 1}$$

(notice that $\gamma_1 > 1$ since $\mu < r$). Summing up: the optimal policy is to sell when the stock price reaches h_{opt}, or immediately if the present price is greater than h_{opt}; the value achieved by this policy is

$$u(x,t) = \max_h u^h(x,t) = \begin{cases} e^{-rt}\left(\frac{\gamma_1 - 1}{a}\right)^{\gamma_1 - 1}\left(\frac{x}{\gamma_1}\right)^{\gamma_1} & \text{if } x < h_{\text{opt}} \\ e^{-rt}(x-a) & \text{if } x > h_{\text{opt}}. \end{cases} \tag{3}$$

One can verify by direct calculation that u 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 is not an accident: it is a general feature of optimal stopping problems.

OK, we have surely found the optimal policy. But we did it by making some guesses. The proof that our answer is right requires a verification argument. Our prior verification arguments showed that the solution of a suitable Hamilton-Jacobi-Bellman equation gave a one-sided bound on the value achieved by any strategy. We shall do something similar here, but the HJB equation is replaced by a *variational inequality.*

Claim. Let $\mathcal{L}$ be the generator of y ($\mathcal{L}\phi = \mu x \phi_x + \frac{1}{2}\sigma^2 x^2 \phi_{xx}$). Suppose there exists a function $v(x,t)$ and constant x_0 such that

(a) $v(x,t) \ge e^{-rt}(x-a)$ for all $x > 0$ and all t;

(b) $v_t + \mathcal{L}v \le 0$ for all $x > 0$ and all t;

(c) v is C^1 at $x = x_0$ and smooth everywhere else.

(d) equality holds in (b) for $0 < x < x_0$ and in (a) for $x > x_0$;

Then for any stopping time τ we have

$$v(x,t) \ge E_{y(t)=x}\left[e^{-r\tau}(y(\tau) - a)\right]$$

for all x, t.

Explanation. We argue as in Section 5: for any sufficiently differentiable function $\phi(x,t)$,

$$d[\phi(y(s),s)] = (\phi_s + \mathcal{L}\phi)ds + \text{a term involving } dw.$$

Taking $\phi = v(x,t)$, integrating from time t to the stopping time τ, then taking the expected value, we get

$$\begin{aligned} E_{y(t)=x}\left[v(y(\tau),\tau)\right] - v(x,t) &= E_{y(t)=x} \int_t^\tau (v_s + \mathcal{L}v)(y(s),s)\, ds \\ &\leq 0 \end{aligned} \tag{4}$$

using (b). Therefore

$$\begin{aligned} v(x,t) &\geq E_{y(t)=x}\left[v(y(\tau),\tau)\right] \\ &\geq E_{y(t)=x}\left[e^{-r\tau}(y(\tau)-a)\right] \end{aligned} \tag{5}$$

using (a). Done!

We have glossed over some technical points (for example, if τ is unbounded we should use this argument on $\tau_k = \min\{\tau, k\}$ then let $k \to \infty$). More interesting: we appear not to have made use of (c) and (d). But notice that we've applied Ito's lemma to a function v that isn't smooth. If v were not C^1 it would be hopeless: you can't do second-order Taylor expansion on a function that's not at least piecewise C^2. Under (c) and (d) the situation is not so bad: the terms appearing in Ito's lemma are uniformly bounded, though v_{xx} is discontinuous at x_0. The discontinuity of v_{xx} turns out to be a minor matter – one can justify this application of Ito's lemma (see Oksendal Theorem 10.4.1).

Examining the argument, one sees that it proves more than was stated in the claim. Namely: if

$$\tau^* = \begin{cases} t & \text{if } x \geq x_0 \\ \text{first time } y(s) \text{ reaches } x_0 & \text{if } x < x_0 \end{cases}$$

then equality holds in both (4) and (5), so $v(x,t)$ is in fact the optimal value and τ^* is the optimal stopping time.

The point of all this, of course, is that the function $u = \max_h u^h$ satisfies conditions (a)-(d) with $x_0 = h_{\text{opt}}$. The only not-entirely-obvious point is that (a) holds for $x > x_0$, i.e. that

$$\phi = e^{-rt}(x-a) \text{ satisfies } v_t + \mathcal{L}v \leq 0 \text{ for } x > x_0.$$

This reduces to the assertion $x_0 \geq (ra)/(r-\mu)$, which follows from the explicit formula for $x_0 = h_{\text{opt}}$.

The solution of the general optimal stopping problem is similar. The value function is described by a variational inequality similar to (a)-(d). There is a "free boundary": on one side equality holds in (the analogue of) (a), on the other side equality holds in (the analogue of) (b). The location of the free boundary determines the optimal stopping policy.

****

European options without transaction costs. This topic will be familiar to students of mathematical finance. We continue to concentrate on the simplest case: a European option on a single stock executing a geometric Brownian motion in a constant-interest-rate environment. We're really interested in two assets this time: a risk-free money-market account, which earns interest at rate r,

$$dy_0 = ry_0 dt,$$

and the stock, which evolves by

$$dy_1 = \mu y_1 dt + \sigma y_1 dw. \tag{6}$$

A *European option* with payoff g and maturity T provides its holder with a cash payment of $g(y_1(T))$ at the maturity date T.

The insight of Black-Scholes-Merton was that the payoff can be *replicated* by a *self-financing portfolio* managed according to a certain *trading strategy.* Since the self-financing portfolio is functionally equivalent to the option, the time-0 value of the option must be the same as the initial endowment of the replicating portfolio.

What is a portfolio? Answer: a non-anticipating vector-valued function $\theta(t) = (\theta_0(t), \theta_1(t))$, giving the amount of each asset held. Its value at time t is

$$V(t) = \theta(t) \cdot y(t) = \theta_0(t)y_0(t) + \theta_1(t)y_1(t).$$

We say it is self-financing if

$$dV = \theta_0 dy_0 + \theta_1 dy_1.$$

This is a stochastic differential equation; it means the stochastic integral of the left hand side equals the stochastic integral of the right hand side. Remembering the meaning of the stochastic integral, this means

$$V(s') - V(s) = \lim_{\Delta t \to 0} \sum_i \theta_0(t_i)[y_0(t_{i+1}) - y_0(t_i)] + \theta_1(t_i)[y_1(t_{i+1}) - y_1(t_i)]$$

where $s = t_0 < t_1 < \ldots < t_N = s'$. Briefly: the value of the portofolio changes — but only because the values of the stock and money market accounts change during time intervals during which the portfolio is held constant. The investor who maintains such a portfolio is buying or selling stock every time θ_0 and θ_0 change – but in doing so he is neither putting new money into the portfolio nor taking money out; the cost of any stock purchases and the proceeds of any sales come from and go to the money market account, not the investor's pocket.

We claim that if $u(x, t)$ solves the Black-Scholes partial differential equation

$$u_t + rxu_x + \tfrac{1}{2}\sigma^2 x^2 u_{xx} - ru = 0 \text{ for } t < T, \text{ with final value } u(x, T) = g(x)$$

then there is a trading strategy determining a self-financing portfolio with initial value $V(0) = u(y_1(0), 0)$ whose final value is $g(y_1(T))$, no matter what the stock price does. This claim shows that the time-0 value of the European option with payoff g and maturity T is $u(y_1(0), 0)$.

To justify our claim we need merely specify the trading strategy and check that it works. Here's the trading strategy: we set

$$\theta_1(t) = \frac{\partial u}{\partial x}(y_1(t), t),$$

and we choose $\theta_0(t)$ so that

$$u(y_1(t), t) = \theta_0(t) y_0(t) + \theta_1(t) y_1(t).$$

This amounts of course to setting

$$\theta_0 = \frac{1}{y_0} \left. (u - x u_x) \right|_{x = y_1(t)} .$$

This strategy clearly has value $V(t) = u(y_1(t), t)$ at every time. To see it is self-financing, notice first that

$$dV = d[u(y_1(t), t)] = [u_t + \mu x u_x + \tfrac{1}{2}\sigma^2 x^2 u_{xx}]dt + [\sigma x u_x]dw$$

where the terms in brackets on the right are evaluated at $x = y_1(t)$. On the other hand

$$\theta_1 dy_1 = [\mu x u_x]dt + [\sigma x u_x]dw$$

evaluated at $x = y_1(t)$, and

$$\theta_0 dy_0 = [u - x u_x] r dt.$$

One verifies using the PDE for u that $dV = \theta_0 dy_0 + \theta_1 dy_1$. Thus the portfolio is self-financing, as asserted. Its final value is $u(y_1(T), T) = g(y_1(T))$ since $u = g$ at $t = T$. This completes the proof of the claim.

If you never saw this argument before it undoubtedly looks mysterious. Slower, more motivated treatments can be found in many sources, for example Neftci or Dewynne-Howison-Wilmott.

****

European options in the presence of transaction costs. The Black-Scholes-Merton theory rests on a miracle: the fact that any possible payoff g has a replicating portfolio. This property, known as *completeness* of the market, extends to more complicated market models – with nonconstant drift and volatility, for example – provided the number of stocks under discussion is equal to the number of distinct sources of randomness. However the Black-Scholes-Merton theory ignores transaction costs, and completeness is immediately lost when we take such costs into account. Completeness can also be lost in other ways, e.g. when the market has two distinct sources of randomness but just one risky asset that can be purchased (example: a stock executing geometric Brownian motion and paying dividends at a random rate).

What might a model with transaction costs look like? Section 1 presented a portfolio optimization model with transaction costs, involving two risk-free assets (a high-yield account and a money-market account). We can introduce modified state variables

- $Y_0(t)$ = value of money-market holdings at time t
- $Y_1(t)$ = value of stock holdings at time t

and control variables

- $\alpha(t) \geq 0$ is the rate at funds are being removed from the money-market account to buy stock
- $\beta(t) \geq 0$ is rate at which stock is being sold, with the proceeds going into the money-market account.

If transaction costs are proportional to the transactions themselves then the state variables evolve by

$$\begin{aligned} dY_0 &= rY_0dt - \alpha dt + [1-\lambda]\beta dt \\ dY_1 &= \mu Y_1 dt + \sigma Y_1 dw - \beta dt + (1-\lambda)\alpha dt \end{aligned}$$

where λ gives the proportion of any transaction that's lost to transaction costs.

Hodges and Neuberger proposed a scheme for pricing of options in such a marketplace in a recent paper *Review of Future Markets* 8, 1989, 222-239. The theory was explored further by Davis, Panas, and Zariphopoulou in *SIAM J. Control Optim.* 31, 1993, 470-493, and by these and other authors in subsequent publications, and it remains an area of active research. Rather than attempt to develop the theory in any generality, we explain (following Soner) what it looks like when specialized to the transaction-free, Black-Scholes setting.

We could return to the Black-Scholes notation used above (holdings $(\theta_0(t), \theta_1(t))$ and prices $(y_0(t), y_1(t))$) but it's more traditional to work instead with the investor's total wealth $V(t) = \theta_0(t)y_0(t) + \theta_1(t)y_1(t)$ as the state, and the proportion π_0 of his wealth in money-market as our control. (In terms of our prior notation, $\theta_0 y_0 = \pi_0 V$.) The proportion of his wealth in stock is naturally $\pi_1 = 1 - \pi_0$. We do *not* restrict π_0 and π_1 to be between 0 and 1, since borrowing and short-selling are permitted. The stochastic differential equation satisfied by V is

$$dV = \pi_0 V r dt + \pi_1 V(\mu dt + \sigma dw).$$

To simplify notation slightly (well, really out of laziness) we set $r = 0$ for the rest of this discussion, and we relabel $\pi_1 = \pi$. The state equation is thus

$$dV = \pi V(\mu dt + \sigma dw)$$

Suppose an agent (an investor, or a bank) has initial endowment z and is considering selling a call option with payoff g and maturity T. Suppose moreover the agent's utility of time-T wealth is a specified function h. And suppose the agent doesn't know portfolios are replicatable (they wouldn't be, when transaction costs are included). Then the agent may reason as follows.

First, suppose it does not issue an option, but just manages its present wealth optimally till time T. This means maximizing the final-time utility

$$E_{V(t)=z}\left[h(V(T))\right]$$

over all (non-anticipating) choices of the control π. The associated value function

$$u^{(1)}(z,t) = \max_{\pi} E_{V(t)=z}\left[h(V(T))\right]$$

solves the HJB equation

$$u_t^{(1)} + \max_{\pi}\{\mathcal{L}^{\pi} u^{(1)}\} = 0$$

with final value $u^{(1)} = h$ at $t = T$. Here $\mathcal{L}^{\pi}$ is the infinitesimal generator with control π:

$$\mathcal{L}^{\pi}\phi(z,t) = \pi\mu z\phi_z + \tfrac{1}{2}\pi^2\sigma^2 z^2\phi_{zz}.$$

Next, suppose the agent were to sell the option. Still denoting by z its intial endowment (which now includes any proceeds from selling the option), the agent's policy is naturally again to manage its wealth optimally till time T. At maturity it has to pay the obligation $g(y_1(T))$, so the expected final-time utility under this scenario is

$$E_{V(t)=z,y_1(t)=x}\left[h[V(T) - g(y_1(T))]\right]$$

and the stochastic control problem is

$$u^{(2)}(z,x,t) = \max_{\pi} E_{V(t)=z,y_1(t)=x}\left[h[V(T) - g(y_1(T))]\right].$$

The value function $u^{(2)}$ satisfies the HJB equation:

$$u_t^{(2)} + \max_{\pi}\{\mathcal{L}^{\pi} u^{(2)} + \sigma^2\pi z x u_{zx}^{(2)}\} + \mathcal{B}u^{(2)} = 0$$

with final value $u^{(2)}(z,x,T) = h(z - g(x))$, and with

$$\mathcal{B}\phi(z,x,t) = \mu x\phi_x + \tfrac{1}{2}\sigma^2 x^2\phi_{xx}.$$

OK. Now the proposal of Neuberger and Hodges. They propose that if the time-t stock price is x, the agent would sell the option for price e only if

$$u^{(2)}(z+e,x,t) \geq u^{(1)}(z,t).$$

In other words, this agent's sales price $e_*(z,x,t)$ is determined by

$$u^{(2)}(z+e_*,x,t) = u^{(1)}(z,t).$$

Notice that the sales price e_* depends on the agent's initial wealth z as well as on the stock price x. The marketplace price of the option should be obtained by optimizing over z, bearing in mind that no agent desires negative final-time utility: it equals $e_*(z_*, x, t)$ where z_* is such that $u^{(1)}(z_*, t) = 0$.

The preceding framework can be used in any market, with or without transaction costs, whether or not it is complete. To show the framework is reasonable, we make the following

Claim: In the Black-Scholes setting (with no transaction costs), the price $e_*(z, x, t)$ is equal to the Black-Scholes value $u(x, t)$ for any z.

Explanation. One proof – the obvious one, but by far the more tedious one – is to use the Hamilton-Jacobi-Bellman differential equations. Our claim asserts that

$$u^{(2)}(z + u(x, t), t) = u^{(1)}(z, t).$$

Both sides clearly have the same final value, and a tedious but straightforward calculation (using the HJB equation characterizing $u^{(2)}$, the Black-Scholes equation characterizing u, and chain rule) shows that $u^{(2)}(z + u(x, t), t)$ satisfies the HJB equation characterizing $u^{(1)}$. It follows (from uniqueness of solutions of the HJB equations) that the two expressions determine the same function. (See Soner for details.)

A better proof – easier and more conceptual – uses the key feature of the Black-Scholes world: the existence of replicating portfolios. To see that $u^{(2)}(z + e, x, t) \geq u^{(1)}(z, t)$ when $e = u(x, t)$, consider any control (trading strategy) for the definition of $u^{(1)}(z, t)$. It starts with endowment z at time t and reaches a certain expected final utility of wealth. We associate with it a trading strategy for the definition of $u^{(2)}(z+e, x, t)$: the initial endowment is $z+e$, and the portfolio consists of two sub-portfolios – one associated with $u^{(1)}$ with initial endowment z, the other the Black-Scholes portfolio replicating g with initial endowment e. Comparing the expected final payoffs shows that $u^{(2)}(z+e, x, t) \geq u^{(1)}(z, t)$. The argument proving the reverse inequality is essentially the same: given a trading strategy for the definition of $u^{(2)}(z + e, x, t)$, we associate with it a trading strategy for $u^{(1)}(z, t)$ consisting of two sub-portfolios: one associated with $u^{(2)}$ with initial endowment $z + e$, the other the Black-Scholes portfolio replicating $-g$ with initial endowment $-e$. Comparing the expected final payoffs shows that $u^{(2)}(z + e, x, t) \leq u^{(1)}(z, t)$.

****

Risk-neutral pricing. Our derivation of the Black-Scholes differential equation was based on identifying a replicating portfolio. That argument is efficient, and it captures the essence of what makes Black-Scholes work (by specifying the optimal trading strategy as well as the value). However it does not explain why the Black-Scholes PDE is the backward Kolmogorov equation (well, really the PDE from the Feynman-Kac formula) for the expected discounted final payoff

$$u(x, t) = E^{\mathrm{RN}}_{y(t)=x}\left[e^{-r(T-t)} g(y(T))\right]$$

associated with the *different* stochastic PDE

$$dy = rydt + \sigma ydw. \tag{7}$$

The superscript RN on E^{RN} means "risk-neutral", and it refers to the fact that we the expected value is taken relative to the stochastic differential equation (7) rather than relative to the original equation $dy_1 = \mu y_1 dt + \sigma y_1 dw$.

This story really deserves an entire section of notes of its own. But here are the bare essentials, specialized to the constant-drift, constant-volatility setting. Rather than work with y_0 and y_1 solving (6), it is convenient to work with the *normalized* prices

$$\bar{y}_0(t) = 1, \quad \bar{y}_1(t) = y_1(t)/y_0(t).$$

This amounts to a change of units: we choose to measure the stock price in inflated dollars (using the risk-free rate r as the inflation rate) rather than in some other more arbitary units. The state equation (6) gives

$$d\bar{y}_1 = (\mu - r)\bar{y}_1 dt + \sigma \bar{y}_1 dw. \tag{8}$$

Martingales measures and the absence of arbitrage. A key requirement of any market model is that it not admit arbitrage. Informally: there should be no investment strategy that (a) requires no initial investment, (b) never loses money, and (c) makes money with positive probability. More formally: there should be no self-financing portfolio such that (a) $V(0) = 0$, (b) $V(T) \geq 0$ almost surely, (c) $V(T) > 0$ with positive probability, and (d) $V(t) \geq -K$ almost surely, for some constant K. (This last condition is a technical but important one; see Oksendal for discussion of it.)

If $\mu = r$ then we can easily show the absence of arbitrage using ordinary expectation. Indeed, when $\mu = r$ we see from (8) that $\bar{y}_1$ is a martingale (its stochastic differential equation has only dw terms). The self-financing condition is

$$dV = \bar{\theta}_1 d\bar{y}_1 = \bar{\theta}_1 \bar{y}_1 \sigma dw$$

so

$$E[V(T)] = V(0).$$

Thus if $V(0) = 0$ and $V(T) \geq 0$ we must have $V(T) = 0$ with probability 1.

If $\mu \neq r$ we can use a similar strategy to rule out arbitrage, but we cannot use the ordinary expectation. Instead we must use a *weighted* expectation. Let dP be the measure on path space that's implicit in our expectation notation:

$$E[f] = \int_{\text{paths}} f \, dP.$$

By a weighted expectation we mean an expression of same form but with a weighted measure $dQ = MdP$:

$$E^Q[f] = \int_{\text{paths}} f \, dQ = \int_{\text{paths}} fM \, dP.$$

We wish to restrict our attention to Q's that are themselves probability measures on path space; the associated weights M satisfy $M \geq 0$ and

$$E[M] = \int_{\text{paths}} M\, dP = E^Q[1] = 1.$$

Finally we want the sets of measure 0 under P and Q to be the same, so we need $M > 0$ almost surely.

Suppose we can find a weighted expectation operator E^Q with all these properties, and one more crucial one: that $\bar{y}_1$ is a martingale relative to E^Q, i.e.

$$E^Q\left[\int_0^T \alpha(s) d\bar{y}_1\right] = 0$$

for any nonanticipating α. Then we can prove the absence of arbitrage exactly as before. The value of any self-financing portfolio satisfies

$$E^Q[V(T)] = E^Q[V(0)] = V(0).$$

So if $V(0) = 0$ and $V(T) \geq 0$ we must have $V(T) = 0$ with probability 1.

This may seem a complicated method for proving absence of arbitrage. But actually it's quite natural, and in fact universal. The weighted expectation operators under discussion are called "martingale measures" or "risk-neutral probabilities." A market without arbitrage has at least one. If the market is complete it has only one. The proof of this statement in a general, continuous-time setting is pretty technical. (Oksendal has a good treatment.) The discrete-time version is an application of the duality theorem of linear programming. The one-period version is in Section 3 of my Math Finance I lecture notes (available from my web page).

The point of the last paragraph was to suggest that the M and Q we need should indeed exist. For the Black-Scholes market we can say more: one can give an explicit formula for M using Girsanov's theorem. We'll get to this in a moment.

Martingale measures and option pricing. We've shown that every contingent claim in the Black-Scholes market is replicatable. In other words, for any payoff function g there's a self-financing portfolio (with initial value $u = u(x,t)$, where x is the present stock price $y_1(t)$, and with final value $V(T) = g(y_1(T))$ almost surely). We claim now that the initial endowment has a simple expression in terms of the risk-neutral measure, namely

$$u(x,t) = E^Q\left[e^{-r(T-t)} g(y_1(T))\right]$$

To explain this, we observe first that if $\theta = (\theta_0, \theta_1)$ determines a self-financing portfolio in the y variables with value $V(t) = \theta(t) \cdot y(t)$, then it also gives a self-financing portfolio in the $\bar{y}$ variables with value $\bar{V}(t) = \theta(t) \cdot \bar{y}(t) = y_0^{-1}(t)V(t)$. Indeed,

$$\begin{aligned} d\bar{V} &= d(y_0^{-1}V) \\ &= -V y_0^{-2} dy_0 + y_0^{-1} dV \\ &= -(\theta_0 y_0 + \theta_1 y_1) y_0^{-2} dy_0 + y_0^{-1}(\theta_0 dy_0 + \theta_1 dy_1) \\ &= \theta_0 d\bar{y}_1 \end{aligned}$$

using $dy_0 = ry_0 dt$ and (8) for the last step. Applying this observation to the Black-Scholes replicating portfolio, we get

$$E^Q\left[\bar{V}(T)\right] = E^Q\left[\bar{V}(t)\right]$$

or equivalently

$$E^Q\left[y_0^{-1}(T)V(T)\right] = E^Q\left[y_0^{-1}(t)V(t)\right]$$

or again

$$E^Q\left[e^{-r(T-t)}g(y_1(T))\right] = V(t) = u(x,t),$$

which is the desired relation.

Girsanov's theorem and the risk-neutral dynamics. The remaining task is to explain why taking expectations relative to the risk-neutral probability Q is equivalent to simply changing the stochastic differential equation from (6) to (7). The essential content of Girsanov's theorem is this: if $\alpha(s)$ is random but non-anticipating, we may choose

$$M(t) = e^{\left(\int_0^t \alpha(s)\,dw - \frac{1}{2}\int_0^t \alpha^2(s)\,ds\right)}.$$

It is a martingale, with $E[M_t] = 1$ for all t, so the associated measure Q is a probability measure. Moreover the process $\tilde{w}$ defined by

$$d\tilde{w} = -\alpha ds + dw$$

is a Brownian motion *relative to the probability* Q. In particular its increments have mean 0 and variance Δs – these means and variances being taken relative to the probability measure Q rather than the original measure P. (There is a technical condition, known as the Novikoff condition: α must be sufficiently well-behaved that $E\left[\exp(\int_0^t \alpha^2\,ds)\right] < \infty$. This presents no problem in financial applications provided the volatility is bounded below.)

The essential meaning of this result is this: changing the measure on path space (from dP to $dQ = MdP$ with M given by Girsanov's theorem) has the effect of changing the ds term in the stochastic differential equation, since it replaces dw with $d\tilde{w} + \alpha ds$.

There's no need to compute further. We saw that the "ordinary" expectation is a martingale measure for a special choice of drift – namely when $\mu = r$. Girsanov's theorem tells us we can always arrange that the drift have this form – by changing the probability measure. The Q that achieves this special drift is the risk-neutral measure. (A little algebra determines the value of α required to change the drift, and Girsanov gives the associated M. But we don't really need M to compute the value of an option. Better to use the risk-neutral stochastic differential equation directly.)

****

American options. An American option differs from a European one in the feature that it can be exercised at any time. If the time of exercise were known, its value would be given by the discounted expected final value (using the risk-neutral probability, or equivalently solving the risk-neutral stochastic differential equation!). Since the time of exercise isn't known it must be determined optimally. We thus get a stopping time problem very much like the one considered at the beginning of these notes:

$$u(x,t) = \max_{\tau} E_{y(t)=x} \left[e^{-r(T-t)} g(y(T)) \right]$$

where y solves the risk-neutral SDE

$$dy = rydt + \sigma ydw.$$

The value is again determined by a variational inequality, whose "free boundary" determines the price at which the option should be exercised at any given time. The main difference from our earlier example is that in this setting the free boundary depends on both space and time.