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 2000.

Reminder concerning the final: The exam will be Tuesday May 9, at the usual class time. It will be "closed-book" (no books, no lecture notes), however you may bring two sheets of your own notes (8.5×11, both sides, write as small as you like). You are responsible for material in Sections 1-6 of the lecture notes, and in Homeworks 1-6. See a separate handout for further discussion of what to expect. (The material in these Section 7 notes will not be on the exam. But it is important.)

Correction to the HW5 solution sheet: The solution given for 4(d) ("what happens when $\mu = \frac{1}{2}\sigma^2$?") was incorrect. Here's a sketch of the right answer. Observe first that $xu_x + x^2 u_{xx} = u_{zz}$ with $z = \log x$. So the general solution of $\mu(xu_x + x^2 u_{xx}) = -1$ is

$$\frac{-1}{2\mu} z^2 + c_1 z + c_2 = \frac{-1}{2\mu} (\log x)^2 + c_1 \log x + c_2.$$

The proof that the expected exit time is finite proceeds by the usual argument, using Ito's lemma applied to $\phi(x) = (\log x)^2$. To get the exit probabilities, use Ito applied to $\phi(x) = \log x$. This gives $\log x = E_{y(0)=x}[\log(y(\tau))] = p_a \log a + p_b \log b$. Using the fact that $p_b = 1 - p_a$ it follows easily that

$$p_a = \frac{\log b - \log x}{\log b - \log a}, \quad p_b = \frac{\log x - \log a}{\log b - \log a}.$$

****

Optimal stopping. 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.

As usual, we shall focus on a simple yet representative example which displays the main ideas, namely: when to sell a stock which undergoes log-normal price dynamics. Our treatment follows Oksendal, Examples 10.2.2 and 10.4.2. After completing this example, we shall discuss how similar ideas apply to the pricing of American options.

Note that we discussed some discrete-time optimal stopping problems earlier in the semester (when to sell an asset; when to park your car). Our goal here is to understand some analogous continuous-time problems.

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

sell when $(y(s), s)$ 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

$$u^h_t + \mu x u^h_x + \tfrac{1}{2}\sigma^2 x^2 u^h_{xx} = 0 \quad \text{for } 0 < x < h$$

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

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

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 $\mu < r$). It is important to 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 $h = h_{\text{opt}}$ is the one for which the roots coincide. Some consideration of the figure shows why: if $h < h_{\text{opt}}$ then increasing h slightly raises the parabola and increases u_h; similarly if $h > h_{\text{opt}}$ then decreasing h slightly raises the parabola and increases u^h.

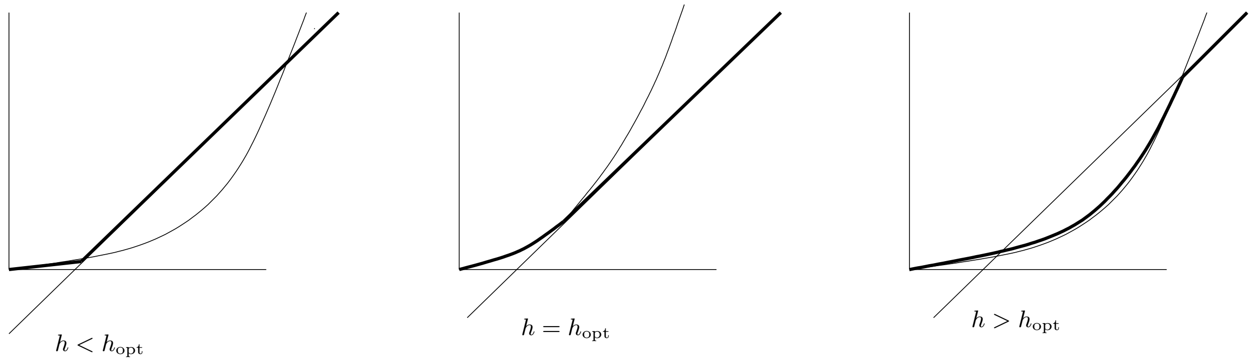

Figure 1: *Graph of u^h.*

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_{opt}, or immediately if the present price is greater than h_{opt}; the value achieved by this policy is

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

Our figure shows – and it can be verified 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) $v(x,t) \geq e^{-rt}(x-a)$ 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 $x = x_0$ 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

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

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

$$v = e^{-rt}(x-a) \text{ satisfies } v_t + \mathcal{L}v \leq 0 \text{ for } x > x_0. \tag{6}$$

This can be verified by direct arithmetic: it reduces to the assertion $x_0 \geq (ra)/(r-\mu)$, which follows from the explicit formula for $x_0 = h_{\text{opt}}$.

The validity of (6) is of course no accident; here's a heuristic explanation why it had to be true. We believe the optimal policy is to sell immediately if the current price is greater than h_{opt}. So if $y(t) = x > h_{\text{opt}}$ we think it would be a mistake to hold the asset a little longer. Apply the calculation that led to (4) using $\phi = e^{-rt}(x-a)$ and $\tau = t + \Delta t$, with Δt chosen small enough that $y(t+\Delta t)$ is still greater than h_{opt} to see that

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

We think the left hand side is negative; so we think right hand side is negative. When Δt is small the right hand side equals

$$\Delta t(\phi_t + \mathcal{L}\phi)(x,t) + \text{higher order terms in } \Delta t.$$

So we conclude that $\phi_t + \mathcal{L}\phi < 0$ for $x > h_{\text{opt}}$. In short: the correctness of our conjectured strategy *implies* the validity of the inequality (a).

The special, simple form of the problem considered above permitted us to find an explicit solution. For more general stopping problems it is typically not possible to solve the PDE explicitly. In such settings, conditions (a)-(d) of the verification argument (suitably modified for the problem of interest) provide the most useful representation of the solution. They describe the value function as the solution of a "variational inequality" The argument that the variational inequality is sufficient for optimality is always easy (modulo technicalities), just as it was above. The proof that the variational inequality has a solution is more difficult, unless the solution can be constructed more or less explicitly as we did above.

****

American options. 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. Then we are interested in

$$\tilde{u}(x) = \max_{\tau} E_{y(t)=x}\left[e^{-r(T-t)}(K-y(\tau))_+\right]. \tag{7}$$

(We call this $\tilde{u}$ rather than u to emphasize the link to our earlier discussion: the definition of $\tilde{u}$ discounts the payoff to the initial time t rather than to time 0. It is easy to see that the right hand side is independent of t.) If you've studied continuous-time finance you probably know that when $\mu = r$, (7) gives the price of a *perpetual American put* in the Black-Scholes marketplace. (It is a "perpetual" put because it never expires, i.e. it has no finite maturity time.) Anyway it makes perfectly good sense to try and evaluate the function $\tilde{u}$ defined this way using the same tools we applied above. This can be done. Here's an outline of the answer:

- It's natural to guess that the optimal policy is determined by an exercise threshold h as follows: exercise immediately 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

$$-r\tilde{u}^h + \mu x \tilde{u}^h_x + \tfrac{1}{2}\sigma^2 x^2 \tilde{u}^h_{xx} = 0 \quad \text{for } x > h$$

and

$$\tilde{u}^h(x,t) = (K-x) \quad \text{for } 0 < x \le h.$$

- To find $\tilde{u}^h$ 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's clear that $\tilde{u}^h$ should decay to 0 as $x \to \infty$. The constant C_2 is set by the boundary condition $\tilde{u}^h(h) = (K-h)$. Evidently

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

- The correct exercise threshold is obtained by maximizing with respect to h. The optimal value is $h_{\text{opt}} = \frac{K\gamma_2}{\gamma_2 - 1}$, which is less than K as expected.

- When $h = h_{\text{opt}}$ the function $v = \tilde{u}^h$ satisfies

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

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

(c) v is C^1 at $x = h_{\text{opt}}$ and smooth everywhere else.

(d) equality holds in (a) for $0 < x < h_{\text{opt}}$ and in (b) for $x > h_{\text{opt}}$

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).

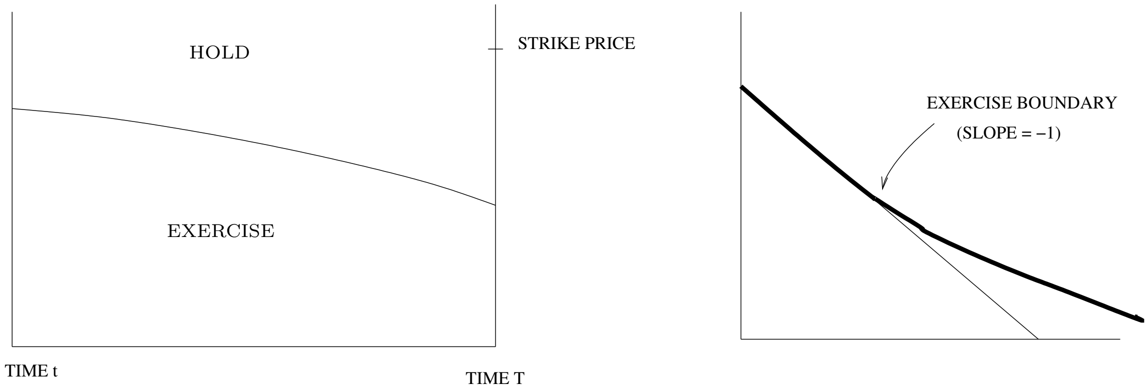

Figure 2: *The exercise boundary of an American option, and its value as a function of stock price at a given time t*

What about American options with a specified maturity time 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:

$$\tilde{u}(x,t) = \max_{\tau \leq T} E_{y(t)=x}\left[e^{-r(T-t)}(K - y(\tau))_+\right].$$

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 $x > h(t)$

If you accept that (a)-(d) has a solution, its optimality is readily verified by the usual argument (modulo the usual technicalities).