

PDE for Finance Notes – Section 6

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. For use only in connection with the NYU course PDE for Finance, MATH-GA 2706. Prepared in 2003, minor updates made in 2011 and 2014.

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 notes Raghu Varadhan prepared when he taught PDE for Finance; my treatment of (b) is a dumbed-down version of one in Oksendal’s book “Stochastic Differential Equations;” the perpetual American put is discussed in many places. My main suggestion for further reading in this area is however F-R Chang’s excellent book “Stochastic Optimization in Continuous Time” (Cambridge Univ Press), which is on reserve in the CIMS library and available electronically through Bobcat. It includes many examples of stochastic control, some close to the ones considered here and others quite different.

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

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

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)

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

By Ito’s formula

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

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 $\Delta t \rightarrow 0$ this gives the PDE in the stop-later region:

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

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

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

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 discontinuous 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 $\frac{1}{2}u_a'' - u_a = 0$ 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 $|x| \leq a$. 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 $u \geq f$ 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 $u_{a_*} \geq f$ 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

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

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

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

(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 justified). Integrating up to the stopping time and taking the expectation gives

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

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

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

which is the desired assertion.

The preceding argument isn't quite complete. We used the assertion that $E[\int_0^\tau g dw] = 0$; but this is clear in general only for *bounded* stopping times (recall from HW1 that it can fail for an unbounded stopping time). How to know no stopping time – not even an unbounded one – can achieve a better result than u_{a^*} ? Here's the answer: for any (possibly unbounded) stopping time τ , consider the truncated stopping times $\tau_k = \min\{\tau, k\}$. Clearly $\tau_k \rightarrow \tau$ as $k \rightarrow \infty$. Since τ_k is bounded, the argument presented above applies to it, so

$$E_{y(0)=x} [e^{-\tau_k} f(y(\tau_k))] \leq u_{a^*}(x).$$

The limit $k \rightarrow \infty$ is handled by Fatou's lemma from real variables. It tells us that

$$E_{y(0)=x} \left[\liminf_{k \rightarrow \infty} e^{-\tau_k} f(y(\tau_k)) \right] \leq \liminf_{k \rightarrow \infty} E_{y(0)=x} [e^{-\tau_k} f(y(\tau_k))]$$

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

$$E_{y(0)=x} [e^{-\tau} f(y(\tau))] \leq u_{a^*}(x)$$

as desired.

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

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

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 x . 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

$$\frac{1}{2}\sigma^2 x^2 u_{xx} + \mu x u_x - r u = 0$$

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

$$u(x) \geq x - a \quad \text{and} \quad \frac{1}{2}\sigma^2 x^2 u_{xx} + \mu x u_x - r u \leq 0$$

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 $\frac{1}{2}\sigma^2 x^2 u_h'' + \mu x u_h' - r u_h = 0$ for $x < h$, with boundary condition $u_h = x - a$ at $x = h$. This can be solved explicitly. The general solution of $\frac{1}{2}\sigma^2 x^2 \phi'' + \mu x \phi' - r \phi = 0$ is

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

where c_1, c_2 are arbitrary constants and

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

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 $x \rightarrow 0$ 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

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

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

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

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

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

and

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

- 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

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

This time the relevant exponent is the negative one, γ_2 , since it is clear that u_h should decay to 0 as $x \rightarrow \infty$. 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
 - (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 $\mathcal{L}v = -rv + \mu x v_x + \frac{1}{2} \sigma^2 x^2 v_{xx}$.




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

- 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).

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).