PDE in Finance, Spring 2008,
`http://www.math.nyu.edu/faculty/goodman/teaching/PDEfin/index.html`
Written by Jonathan Goodman. Please report errors, sources of confusion, and typos on the class Blackboard site.
Last updated April 15, 2008.

Section 8: Optimal exercise and calculus of variations.

1 Early exercise

An American style option is the right to do a transaction at a specified price at any time. The quantitative finance definition of *option* is any cash flow that depends on the values of some set of economic variables. By definition, an American style option is one in which the cash flow depends on decisions made by the holder. Naturally, the holder will try to make the decisions in an optimal way. This makes pricing of American style options a part of stochastic optimization.

The simplest American style option problem involves a diffusion process, $X(t)$, that represents the relevant part of the economy, and a payout function $V(x)$. At time t the holder has the option of exercising or holding the option. Exercise gives the holder a cash flow, or payout, $V(x)$. We may and do assume that $V(x) \geq 0$ for all x. A value $V(x) < 0$ is indistinguishable from $V(x) = 0$ since we would not exercise the option at any x with $V(x) \leq 0$. Therefore, the payout $V(x)$ and $V_+(x)$ ($V_+ = \max(V, 0)$ is the non-negative part of V) are equivalent from a cash flow point of view.

An exercise strategy is a stopping time, τ, with respect to the diffusion, X. Recall the definition of a stopping time from stochastic calculus. We write A_t for the event that $\tau \leq t$. Then τ is a stopping time if $A_t \in \mathcal{F}_t$ for all t. This says that the decision of whether to stop at time t is made on the basis of values of $X(s)$ for $s \leq t$, but not on values for $s > t$. A stopping time is the same as a decision strategy.

(In principle, we could consider decision strategies that involve tossing coins unrelated to X. For example, we could say that each time a stock reaches a certain level, we toss a coin and exercise with probability p. In this case, $\mathcal{F}_t$ would not be the filtration generated by X, but could be larger. The real assumption is that X should be a stochastic process with respect to $\mathcal{F}_t$. We do not worry about this distinction because it turns out that optimal decision strategies do not involve coin tossing.)

One of the conclusions of the theory of dynamic programming is that the solution to an optimal control problem has the form of feedback control. That is, decisions made at time t depend only on the state at time t. Here, the only decision is whether to exercise or not. Therefore, there will be some values of $X(t)$ that call for exercise at time t and the rest that call for *continuation* (not exercise). The exercise set (or *early exercise* set) will be written $\mathcal{E}(t)$. The

continuation set is C(t). Clearly, $\mathcal{E}(t) \cap \mathcal{C}(t)$ is empty, and $\mathcal{E}(t) \cup \mathcal{C}(t)$ is the whole state space. The stopping time is $\tau = \min \{ t \mid X(t) \in \mathcal{E}(t) \}$. We write $\mathcal{C}$ for the trajectory in set space $\{\mathcal{C}(t)\}$, just as we write X for the trajectory $\{X(t)\}$. As a convention, we suppose that the common boundary of $\mathcal{C}$ and $\mathcal{E}$ (defined in the same way as $\mathcal{C}$) belongs to $\mathcal{E}$. We write Γ for this common boundary. This implies that $\mathcal{E}$ is a closed set and $\mathcal{C}$ is open. (Ignore this if you don't know what closed and open mean in this context.) We put Γ into $\mathcal{E}$ so that the minimum in the definition of τ is taken, i.e., that we don't need to say inf instead of min in the definition of τ. This is possible because the trajectories $X(t)$ are continuous functions of t. The mathematical framework is more technical for jump diffusions.

The boundary of $\mathcal{E}$, which is Γ, is the *decision boundary*. As we just said, choosing an optimal early exercise strategy is the same as choosing a decision boundary. In some cases we can find analytical solutions. More commonly, we use numerical techniques that find the decision boundary and the (optimal) value function at the same time. The binomial or trinomial tree method for American style options is an example of such a numerical technique, though more sophisticated and effective techniques exist. As for other problems, the curse of dimensionality prevents us from computing the value function for high dimensional problems. In fact, the curse of dimensionality prevents us even from representing a general surface, Γ, even if we did not have to compute the value function.

There are several possible approaches to problems that face the curse of dimensionality. One is heuristics such as the Longstaff and Schwartz approach to American options. Several people have proposed Monte Carlo methods that in principle converge to the solution of the American option problem. Even the most successful of these, by Brodie and Glasserman, have total computation time that is exponential in the number of time steps (e.g. 50^n in the simplest case), so they too are not practical for high accuracy computation. Remarkably, they provide useful approximations even with as few as 4 or 5 time steps. Another approach related to the previous two is to specify a parametric form of the early exercise decision boundary, say, using neural nets or other ideas from machine learning. We then can use Monte Carlo optimization (what machine learning people call *learning*) to find the optimal parameters. This approach is only as accurate as the heuristic decision boundary family you use. Finally, there are perturbation techniques that provide answers when some of the components of X have little uncertainty.

The simplest early exercise problem, in the spirit of earlier problems we have treated would be to choose τ to optimize discounted expected utility. This is not the Black Scholes theory for American options and does not lead to the Black Scholes PDE. We will discuss below the reasoning that does lead to Black Scholes. But for now, we consider the utility of the payout is $U(V(X(\tau)))$ (count parentheses please). It turns out here that there is no distinction between the mathematical problems of maximizing expected cash flow and maximizing the utility of the expected cash flow. Indeed, optimizing U is the same as optimizing $\widetilde{U} = U - u_0$ for any u_0. We take advantage of this to take $u_0 = U(0)$. Since

$U(v)$ is an increasing function of v, this implies that $\widetilde{U}(V(x)) > 0$ if and only if $V(x) > 0$. Therefore, if we define $\widetilde{V}(x) = \widetilde{U}(V(x))$, then we may as well regard $\widetilde{V}$ as the payout. In other words, we may as well pretend we are paid in *utils* (Utils are units of utility. It's an economists' joke that emphasizes the uncertain foundations of utility theory.) rather than in units of a particular currency. Therefore, in the rest of the rest of this section, we discuss optimizing expected cash flow, because the mathematical structure is the same as the more proper one using utility theory.

The value function is the solution to the optimization problem starting at x at time t:

$$f(x,t) \;=\; \max_{\tau} E_{x,t}\left[\, e^{-r(\tau - t)}\, V(X(\tau))\,\right] \;. \tag{1}$$

For $x \in \mathcal{C}$, you do nothing but watch $X(t)$. Therefore, this f satisfies the backward equation (assuming $dX = a(X)dt + b(X)dW$)

$$\partial_t f \;+\; \frac{1}{2} b b^t_{jk}(x) \partial_{x_j} \partial_{x_k} f \;+\; a_j(x) \partial_{x_j} f \;+\; rf \;=\; 0 \;, \tag{2}$$

Of course,

$$f(x,t) \;=\; V(x) \quad , \qquad \text{for } \; x \in \mathcal{E}(t) \;. \tag{3}$$

This serves as a boundary condition for the PDE (2), but in that case it needs to be applied only on the boundary of $\mathcal{C}$, which is Γ.

The final condition is that

$$u(x,T) = V(x) \quad \text{for all } x \;. \tag{4}$$

We say that x is *in the money* and write $x \in \mathcal{P}$ if $V(x) > 0$. If $x \in \mathcal{P}$ and the option is expiring right now, you have the choice between a positive cash flow or nothing. You choose to exercise and get the positive. If $x \notin \mathcal{P}$, you get zero whether or not you exercise the option. This implies that $u(x,T) = 0$ if $V(x) = 0$. If $t < T$, we should have $\Gamma(t) \subset \mathcal{P}$. Otherwise we would be exercising the option when it is worthless. Only at time T is $\Gamma(T)$ the boundary of $\mathcal{P}$ rather than properly inside of $\mathcal{P}$.

With all this PDE stuff, the stochastic optimal control problem (1) is equivalent to maximizing f subject to the constraints (2), (3), and (4). This means optimizing over Γ. For any Γ, we can use (2), (3), and (4) to find the corresponding f. The optimal Γ is that which produces the largest f. This is a *free boundary* problem because the unknowns are not only the PDE solution but also the PDE boundary, Γ.

To find the optimal Γ we need optimality conditions. It turns out that these optimality conditions take the form of an extra boundary condition that must be satisfied on Γ. This is typical of other free boundary conditions. The two boundary conditions together determine both f and Γ. We derive the extra boundary conditions below.

2 Calculus of variations

The *calculus of variations* is a set of mathematical ideas for finding an optimal function according to some criteria. In general, if $u(x)$ is a function and $F(u)$ is some quantity determined by u, then the calculus of variations is a technique that tries to find the function, u, that optimized F. The goal is to calculate the gradient $\nabla_u F(u)$ and set it to zero. The problem is that in then infinite dimensional space of all functions it is not clear exactly what $\nabla_u F(u)$ means. The calculus of variations is a way to identify the gradient. Setting it to zero then turns into some kind of equation or equation system that we try to solve. A function that depends on another function, such as $F(u)$ often is called a *functional*. The analogue of the gradient is the *first variation*, written δF (or $\delta F/\delta u$ or $\delta_u F$).

In our application the unknown function u is replaced by the unknown surface, Γ. The functional we want to optimize is $f(x,t)$. It turns out in this problem that the optimality conditions do not depend on the values of x and t. Before treating this example, we treat some simpler examples to explain the basic method.

The first variation is a generalization of the idea of the gradient of a function of many variables. When you want to generalize something to a more general context, you ask yourself: "what is this thing really doing", then figure out a way to do that in the more general situation. For $x \in R^n$, the gradient, $\nabla_x F$ tells you how much F changes if we make a small change in x. In notation sometimes used in this situation, we pick another $\dot{x} \in R^n$ and ask how F changes when you move from x in the direction $\dot{x}$. (Note, there is no time variable here and $\dot{x}$ is not the derivative of x with respect to time.) More precisely, we have the formula from the chain rule:

$$F(x + \epsilon\dot{x}) \; - \; F(x) \; = \; \epsilon \nabla F(x) \cdot \dot{x} \; + \; O\left(\epsilon^2\right) \; . \tag{5}$$

Put another way, the definition of $\nabla_x F(x)$ is that it is the vector so that, for every $\dot{x}$,

$$\frac{d}{d\epsilon} F(x + \epsilon\dot{x}) \Big|_{\epsilon=0} \; = \; \nabla F(x) \cdot \dot{x} \; . \tag{6}$$

In R^n, the optimality criterion is that $\nabla_x F(x) = 0$. The reason is that if $\nabla_x F(x) \neq 0$ then there is some $\dot{x}$ with $\nabla_x F(x) \cdot \dot{x} > 0$. Then (5) tells us that if we take $\epsilon > 0$ but sufficiently small, then $F(x + \epsilon\dot{x}) > F(x)$, so x was not optimal. We want to do the same for functionals of functions: define $\delta_u F$ so that (6) holds. From this we conclude that the optimality condition is $\delta_u F(u) = 0$. This turns out to be a differential equation for u.

A simple example is finding a function $u(x)$ defined for $0 \leq x \leq u$ that minimizes

$$F(u) \; = \; \int_0^1 u'(x)^2 \, dx \; , \tag{7}$$

subject to the constraints (or boundary conditions) $u(0) = a$, and $u(1) = b$. The first variation is a function, $v(x)$ so that for "all" functions $\dot{u}(x)$ that respect

the boundary conditions,

$$\frac{d}{d\epsilon} F(u + \epsilon \dot{u}) \bigg|_{\epsilon=0} \;=\; \int_0^1 v(x) \dot{u}(x)\, dx \;. \tag{8}$$

If this holds, we write $v = \delta_u F(u)$. Respecting the boundary conditions means that $(u + \epsilon \dot{u}(0) = a$ for all ϵ. This is the same as $\dot{u}(0) = 0$. Clearly also $\dot{u}(1) = 0$. Thus, finding the first variation is the same as finding a v so that (8) holds for all $\dot{u}$ with $\dot{u}(0) = 0$ and $\dot{u}(1) = 0$.

For the functional (7), and for most others, the calculus of variations boils down to integration by parts. Clearly, $\partial_x \left(u + \epsilon \dot{u} \right) = \partial_x u + \epsilon \partial_x \dot{u}$. Therefore

$$F(u + \epsilon \dot{u}) \;=\; F(u) \;+\; 2\epsilon \int_0^1 u'(x)\, \dot{u}'(x)\, dx \;+\; O\left(\epsilon^2\right).$$

The integral on the right of (8) involves $\dot{u}$, not $\dot{u}'$. We come to this by integration by parts with respect to x. The boundary term is $u'(1)\dot{u}(1) - u'(0)\dot{u}(0)$. The boundary conditions for $\dot{u}$ make these boundary terms vanish. What remains is

$$\frac{d}{d\epsilon} F(u + \epsilon \dot{u}) \bigg|_{\epsilon=0} \;=\; -2 \int u''(x)\, \dot{u}(x)\, dx \;.$$

This implies that (using (8)) the first variation of the functional (7) is

$$\delta_u \int_0^2 u'(x)^2\, dx \;=\; -2u''(x) \;. \tag{9}$$

The function that connects $u(0) = a$ to $u(1) = b$ with the minimum of (7) has second derivative equal to zero, i.e. is linear.

3 Optimality condition for free boundaries

The goal of this section is to figure out how to calculate $\delta_\Gamma f(x,t)$, or at least set it to zero. Right now, that expression is not well defined. We must give it a sense similar to (6). The main technical point is to find a definition of a family of boundary surfaces, Γ_ϵ with $\Gamma_0 = \Gamma$ so that the functional derivative may be defined and calculated. We choose the following approach.

For every point $p \in \Gamma$, let $n(p)$ be the unit normal pointing toward $\mathcal{C}$. A point on Γ_ϵ will be a point $p \in \Gamma$ transported by an amount $\epsilon \dot{\gamma}(p)$ in the direction $n(p)$. More precisely, suppose there is a function $\dot{\gamma}(p)$ defined for all $p \in \Gamma$. For sufficiently small ϵ, we may define the perturbed Γ_ϵ as the set of all points $p \in \Gamma$ moved by a distance $\epsilon \dot{\gamma}(p)$ in the direction $n(p)$:

$$\Gamma_\epsilon \;=\; \{\, p \;+\; \epsilon \dot{\gamma}(p) n(p) \mid p \in \Gamma \,\}$$

If $\dot{\gamma}(p) > 0$ then that point of Γ_ϵ is inside $\mathcal{C}$.

For sufficiently small ϵ, there is a perturbed value function $f(x,t,\epsilon)$ that satisfies the PDE, (2), the final condition, (4), and the boundary condition

$f(x,t,\epsilon) = V(x)$ for all $x \in \Gamma_\epsilon$. We want to characterize the perturbed value function

$$\dot{f}(x,t) \;=\; \frac{d}{d\epsilon} f(x,t,\epsilon)\Big|_{\epsilon=0} \quad .$$

More precisely, we want to find conditions under which $\dot{f} \equiv 0$ for all x and t. The fact that this is possible is responsible for the possibility of optimizing all the values $f(x,t)$ at the same time, rather than having a different optimal Γ for each x and t.

To characterize $\dot{f}$, we need a PDE, final values, and boundary values. The first two are easy. We simply differentiate both sides of (2) with respect to ϵ and then set $\epsilon = 0$. Since differentiations with respect to different variables commute and since none of the coefficients a, b, and r depend on ϵ, we see that $\dot{f}$ satisfies the same backward equation,

$$\partial_t \dot{f} \;+\; \frac{1}{2} b b^t_{jk}(x) \partial_{x_j} \partial_{x_k} \dot{f} \;+\; a_j(x) \partial_{x_j} f \;+\; r\dot{f} \;=\; 0 \,.$$

The final condition $f(x,T,\epsilon) = V(x)$ also does not depend on ϵ. This implies (differentiating both sides with respect to ϵ) that $\dot{f}(x,T) = 0$. So far it looks like $\dot{f} \equiv 0$ is possible. But this cannot be true all the time because not all Γ are optimal.

The interesting condition is the boundary condition satisfied by $\dot{f}$. We write this as

$$f(p + \epsilon\dot{\gamma}(p)n(p), t, \epsilon) \;=\; V(p + \epsilon\dot{\gamma}(p)n(p)) \,. \tag{10}$$

This condition is supposed to hold for all $p \in \Gamma_0$ because Γ_ϵ is the set of all points of the form $p + \epsilon\dot{\gamma}(p)n(p)$ for $p \in \Gamma_0$. We find the boundary condition by differentiating both sides with respect to ϵ and then setting $\epsilon = 0$. On the right side we get the directional derivative:

$$\Big(\nabla_x V(p) \cdot n(p) \Big) \, \dot{\gamma}(p) \,.$$

The left, f depends on ϵ in two ways because both the first and third arguments involve ϵ. In the chain rule for partial differentiation there are terms corresponding to each way f depends on ϵ:

$$\frac{d}{d\epsilon} f(x,t,\epsilon)\Big|_{\epsilon=0} \;=\; \Big(\nabla_x f(p,t,0) V(p) \cdot n(p) \Big) \, \dot{\gamma}(p) \;+\; \dot{f}(x,t) \,.$$

Setting them equal using (10) gives the identity

$$\Big(\nabla_x f(p,t,0) V(p) \cdot n(p) \Big) \, \dot{\gamma}(p) \;+\; \dot{f}(x,t) \;=\; \Big(\nabla_x \cdot n(p) \Big) \, \dot{\gamma}(p) \,.$$

Re-ordering this gives

$$\dot{f}(x,t) \;=\; \Big(\nabla_x \big(f(p,t,0) - V(p) \big) \cdot n(p) \Big) \, \dot{\gamma}(p) \,. \tag{11}$$

This is the key relation.

The relation (11) tells us that if we want $\dot{f} \equiv 0$ on Γ for all $\dot{\gamma}$, we must set the coefficient of $\dot{\gamma}$ to zero. That is:

$$n(p) \cdot \Big(\nabla_x f(p,t,0) \; - \; \nabla_x V(p) \Big) \; = \; 0 \; . \tag{12}$$

The optimal Γ is the one that allows f to satisfy the PDE (2), the final conditions (4), and both the boundary conditions (3) and (12). The extra condition (12) is the *smooth pasting* condition. Together with the boundary condition (3), it implies that not only is $f(x,t)$ equal to $V(x)$ on Γ, but the slope of f matches the slope of V on Γ as well. Both f and $\nabla_x f$ are continuous at points $p \in \Gamma$.