PDE for Finance Notes, Spring 2000 – Section 2

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.

Special note: My normal Thursday 4-5 office hours are cancelled Feb 17 due to a travel commitment.

More on deterministic optimal control. This is the second of two sections on deterministic optimal control. Topics: Pontryagin's maximum principle; an example involving transaction costs; the Hopf-Lax formula (a reorganized and extended version of the last bit of Section 1). The discussion of Hopf-Lax touches ever-so-briefly on the concept of viscosity solution. The section closes with some suggestions for further reading.

****

Pontryagin's maximum principle. Many books on optimal control hardly mention dynamic programming and the HJB equation at all. Instead they use an alternative (ultimately equivalent) approach known as Pontryagin's maximum principle.

For those who already know some PDE – in particular, who know what it means to solve a first-order PDE by the method of characteristics – it's easy to explain what's going on. Pontryagin's maximum principle simply identifies the characteristics of the HJB equation, i.e. it solves the HJB equation by the method of characteristics.

For those who know nothing of PDE, but who know something about nonlinear optimization and first-order optimality conditions – in particular, who know the method of Lagrange multipliers for solving a constrained optimization – it's also easy to explain what's going on. The optimal control problem is an (infinite-dimensional) nonlinear optimization problem. Pontryagin's maximum principle basically states its first-order optimality conditions (obtained by the method of Lagrange multipliers). [I'm simplifying matters a bit: the preceding statement is accurate when the problem is sufficiently convex so that optimal controls exist; the Pontryagin maximum principle also addresses what happens when optimal controls don't exist, as in the final problem of HW1. But that part of the theory is beyond the scope of this course.]

Actually you don't need prior knowledge of either PDE or nonlinear optimization to understand this topic. The following (formal) derivation depends exclusively on the saddle-point principle that $\min_x \max_y F(x,y) = \max_y \min_x F(x,y)$ when F is convex in x and concave in y. (The special case when F is linear in x and y can be used to develop the theory of duality in linear programming; Peter Lax's new book on Linear Algebra takes this approach. When F is nonlinear one needs some further technical conditions; they rarely fail in practice so we shall not fuss over them.)

OK, let's start. We focus as usual on the finite-horizon problem with equation of state

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

and value function

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

Our treatment is similar to the one in chapter 10 of A.K. Dixit, *Optimization in Economic Theory*, Oxford University Press. (A xerox copy of this chapter is in the green box reserve in the CIMS library.)

The equation of state determines $y(s)$ for all s, once the control $\alpha(s)$ is given. However for the present purposes we'd like to view *both* $y(s)$ and $\alpha(s)$ as unknowns, and the equation of state as a constraint coupling them. Thus we think of the problem as:

$$u(x,t) = \max_{\substack{dy/ds = f(y,s),\ y(t)=x \\ \alpha \in A}} \left\{ \int_t^T h(y(s), \alpha(s))\, ds + g(y(T)) \right\}.$$

A standard means of handling constrained optimization is the method of Lagrange multipliers. In the present setting the constraint is an *equation*, valid for all s; so the Lagrange multiplier is a *function* of s, call it $\pi(s)$. The preceding definition of u can be written as

$$u(x,t) = \max_{\substack{y(t)=x \\ \alpha \in A}} \min_{\pi(s)} \left\{ \int_t^T \pi(s) \cdot [f(y(s), \alpha(s)) - \tfrac{dy}{ds}]\, ds + \int_t^T h(y(s), \alpha(s))\, ds + g(y(T)) \right\}$$

since the minimum over π is $-\infty$ unless $dy/ds = f(y(s), \alpha(s))$ for all s. (Note: if y takes values in R^n then so does π, and the expression $\pi \cdot [f - dy/ds]$ is the inner product of the two n-vectors π and $f - dy/ds$.)

Let's assume it's correct to interchange the max and the min. Integrating by parts as well, we get

$$u(x,t) = \min_{\pi(s)} \max_{\substack{y(t)=x \\ \alpha \in A}} \left\{ \int_t^T \left[y \cdot \tfrac{d\pi}{ds} + f(y, \alpha) \cdot \pi + h(y, \alpha) \right] ds + g(y(T)) + x \cdot \pi(t) - y(T) \cdot \pi(T) \right\}.$$

Now maximize over α. For any given $\pi(s)$ and $y(s)$, the best choice of $\alpha(s)$ is the one that maximizes $f \cdot \pi + h$. Therefore

$$u(x,t) = \min_{\pi(s)} \max_{y(t)=x} \left\{ \int_t^T \left[y \cdot \tfrac{d\pi}{ds} + H(\pi, y) \right] ds + g(y(T)) + x \cdot \pi(t) - y(T) \cdot \pi(T) \right\} \tag{1}$$

with

$$H(p,x) = \max_{a \in A} \left[p \cdot f(x,a) + h(x,a) \right]. \tag{2}$$

The connection with the HJB equation is now evident: the HJB equation for this problem is $u_t + H(\nabla u, x) = 0$ with the *same* function H. Since the optimization (1) involves $H(\pi, y)$ and the HJB equation involves $H(\nabla u, x)$, it's natural to guess that the optimal π is equal to $\nabla u(s)$. This is indeed the case. (The rigorous proof of this fact involves verification that the Pontryagin maximum principle, stated below, does give the characteristics of the HJB equation. We shall not attempt such an argument here.)

The formula (2) has two important consequences. The first is

$$\frac{\partial H}{\partial p} = f(x, a_*) \quad \text{where } a_* = a_*(p, x) \text{ is the optimal } a \text{ in the definition of } H(p, x). \tag{3}$$

The explanation is simplest when the control is unconstrained (i.e. there is no constraint $\alpha(s) \in A$). Applying the chain rule to the relation $H(p, x) = p \cdot f(x, a_*(p, x)) + h(x, a_*(p, x))$ gives

$$\frac{\partial H}{\partial p} = f(x, a_*) + \frac{\partial}{\partial a}\left[p \cdot f(x, a) + h(x, a)\right]\Big|_{a=a_*} \frac{\partial a_*}{\partial p}.$$

But the second term is zero, because a_* is optimal, and this shows (3). When the control is constrained to stay in some set A the argument must be modified a bit, but the idea is the same. The second fact is that if we choose the control at time s to be $a_*(\pi(s), y(s))$ then the state $y(s)$ satisfies

$$dy/ds = \nabla_p H(\pi(s), y(s)). \tag{4}$$

This is an immediate consequence of (3), using the equation of state $dy/ds = f$.

We return now to the min-max representation of u, equation (1). The inner maximization is over all functions $y(s)$ such that $y(t) = x$. They need not satisfy any differential equation or correspond to any control. So $y(s)$ is an independent variable at each time $s > t$, and the optimization over y can be done separately for each s. The expression $y \cdot d\pi/ds + H(\pi, y)$ is maximized when its derivatives with respect to each y_i are zero. We thus conclude that

$$d\pi/ds = -\frac{\partial H}{\partial y}. \tag{5}$$

A similar argument at the final time gives

$$\pi(T) = \nabla g(y(T)) \quad \text{at the final time } T. \tag{6}$$

Collecting the information in (3)–(6), we have shown **Pontryagin's maximum principle**:

(a) At each time s, the control $\alpha(s)$ should be the value of $a \in A$ that maximizes $\pi \cdot f(y, a) + h(y, a)$ with $\pi = \pi(s)$ and $y = y(s)$.

(b) The evolution of $y(s)$ and $\pi(s)$ are governed by the differential equations

$$\begin{aligned} dy/ds &= \nabla_\pi H(\pi(s), y(s)) \\ d\pi/ds &= -\nabla_y H(\pi(s), y(s)) \end{aligned}$$

for $t < s < T$.

(c) The initial condition for $y(s)$ is known from the initial formulation of the problem, $y(t) = x$. The initial condition for $\pi(s)$ is not given to us explicitly; rather it is determined implicitly by the final-time condition $\pi(T) = \nabla g(y(T))$.

The ODE system in (b) has the same form as that of Hamiltonian mechanics. That's why H is called the Hamiltonian.

The advantage of this approach is that it is in some sense more *local* than dynamic programming. Rather than calculate $u(x,t)$ everywhere in space and time, we need only solve certain ODE's along a single trajectory $y = y(s)$, $\pi = \pi(s)$ for $t < s < T$. Doing so, we get the optimal control, the associated path, and even the gradient of the value function $\nabla u(y(s)) = \pi(s)$ along that path.

In the end, however, the method is less local than it looks. This is due to the lack of an initial condition for $\pi(t)$. In its place we have a final-time condition for $\pi(T)$. To determine the optimal trajectory and control strategy one must somehow determine the special value of $\pi(t)$ that makes the solution of the requisite ODE's (the Hamiltonian system specified in (b) above) satisfy the final-time condition $\pi(T) = \nabla g(y(T))$ at time T.

Let's bring this down to earth by considering how it applies to Example 1. To keep matters as simple as possible we restrict attention to the power law utility, and we don't discount consumption. The equation of state is then

$$dy/ds = ry - \alpha$$

and the goal is

$$\max_{\alpha} \int_t^T [\alpha(s)]^\gamma \, ds.$$

The Hamiltonian is

$$H(p,x) = \max_a \{p(rx - a) + a^\gamma\}.$$

The optimal choice is $a = (p/\gamma)^{1/(\gamma-1)}$, resulting in

$$H(p,x) = prx - c_\gamma p^{\gamma/(\gamma-1)}$$

with $c_\gamma = (1/\gamma)^{1/(\gamma-1)} - (1/\gamma)^{\gamma/(\gamma-1)}$. Therefore the evolution of π and y satisfy

$$dy/ds = \partial H/\partial p = ry(s) - (\pi(s)/\gamma)^{1/(\gamma-1)}$$

(the algebra is simplest if you make use of (3)) and

$$d\pi/ds = -\partial H/\partial x = -r\pi(s).$$

Solving for $\pi(s)$ is easy: evidently

$$\pi(s) = \pi(t)e^{-r(s-t)}.$$

To finish the problem using this approach we should:

(a) Substitute this $\pi(s)$ into the equation for y and solve the resulting ODE, using the initial condition $y(t) = x$. This gives a formula for $y(T)$ in which $\pi(t)$ is a parameter.

(b) Use the condition $y(T) = 0$ to determine $\pi(t)$.

This calculation is clearly possible, though a bit laborious to do by hand. (A good symbolic-integration package such as Matlab, Maple, or Mathematica can help here.) Another example, close to problem 4 of HW1, is worked out in detail in Dixit.

Notice that while a complete solution by this method is somewhat laborious, some valuable qualitative information comes easily, even without determining the exact value of $\pi(t)$. In fact, Pontryagin's maximum principle tells us the qualitative time-dependence of the optimal consumption rate: since $\pi(s)$ decays exponentially at rate r, the optimal consumption rate increases exponentially in time at rate $r/(1-\gamma)$, i.e. $\alpha(s) = Ce^{rs/(1-\gamma)}$ for some constant C. This situation is fairly typical: the Pontryagin maximum principle is a good source of qualitative information, though some additional more global argument is usually needed to determine the precise solution associated with a given initial state x.

Suggestions for further reading on Pontryagin's maximum principle and its applications: I've already noted Chapter 10 of Dixit. The book by Macki and Strauss emphasizes this viewpoint, but takes most of its examples from the physical sciences. I recently learned of a good book – actually better – at about the same level as Macki and Strauss. It is *Optimal Control: An Introduction to the Theory with Applications* by Leslie M. Hocking, Oxford University Press, 1991 (Bobst has one copy, not on reserve; there's a paperback edition that's not too expensive). Like many others, this book emphasizes Pontryagin's maximum principle; its strength is that it presents a wide variety of applications including economics, biology, and other areas of science. Dixit recommends the paper *Dynamic limit pricing: optimal pricing under threat of entry* by D.W. Gaskins, Jr., J. Econ. Theory 3 (1971) 306-322 and he's right: it gives a nice – and relatively elementary – economic application involving competition and pricing. I've put a copy of Gaskin's article in the green box reserve.

****

An example involving transaction costs. Our examples were chosen carefully. Example 1 is the simplest problem I know involving optimization of investment and consumption; we shall soon be discussion similar stochastic problems (first analyzed by Merton). Example 2 has the advantage of being intuitive and easy to visualize, and it shows that the value function need not be smooth.

A major application of optimal control to mathematical finance is the analysis of models which take into account transaction costs. The standard Black-Scholes approach to option pricing assumes continuous-in-time hedging, so it specifically ignores transaction costs. Merton's analysis of optimal investment and consumption also ignores transaction costs. But it's relatively easy to formulate approaches to these problems that account for transaction costs. Analyzing the behavior of the resulting control problems is less easy: this is an area of current research. The book on reserve by Capuzzo Dolcetta and Lions has an article by Soner which gives a good summary of work in this direction. The following example reveals some key ideas in a deterministic setting. The same problem with a general (rather than power-law) utility function is analyzed in *Optimal investment and consumption with* *two bonds and transaction costs*, by S.E. Shreve, H.M. Soner, and G.-L. Xu, Mathematical Finance 1 (1991) 53-84 (I've put a copy in the green box, but be warned that it's not easy reading).

Example 3. Consider an investor who can choose between two different investment opportunities:

- a risk-free money-market account, paying constant interest r, and
- a risk-free high-yield account, paying constant interest $R > r$.

The investor can move money between the two accounts, but in doing so he incurs a transaction fee proportional to the amount of money moved:

- when moving funds from money-market to high-yield, X dollars in money-market becomes $(1 - \mu)X$ dollars in high-yield (μX is the transaction cost);
- when moving funds from high-yield to money-market, Y dollars in high-yield becomes $(1 - \mu)Y$ dollars in money-market (μY is the transaction cost).

The investor can remove money from the portfolio only by taking it out of the money market account; there is no transaction fee associated with such consumption. The investor can take short positions in either account: when the money-market balance is negative it accrues interest charges at rate r; when the high-yield balance is negative it accrues interest charges at the higher rate R. However we impose a *solvency constraint*: liquidation of the whole portfolio into money market should never leave the investor in debt. If X is the money-market position and Y is the high-yield position, this solvency constraint says:

- if $Y \geq 0$ then $X + (1 - \mu)Y \geq 0$. (Liquidation involves turning Y dollars in high-yield into $(1 - \mu)Y$ dollars in money-market; this must be sufficient to pay off any money-market debt).
- if $Y \leq 0$ then $X + Y/(1 - \mu) \geq 0$. (Liquidation involves paying off the high-yield debt by removing $|Y|/(1 - \mu)$ dollars from money-market; the resulting money-market balance must not be negative.)

The investor's goal is to maximize the discounted utility of his total future consumption. To keep things simple we shall focus on the power-law utility $h(c) = \frac{1}{p}c^p$ with $0 < p < 1$, and we take the discount rate to be $\rho = 1$.

We can formulate this as a control problem as follows. The *state* is an R^2-valued function of time, $(X(t), Y(t))$, where

- $X(t)$ = money-market position at time t,
- $Y(t)$ = high-yield position at time t.

The solvency condition restricts X and Y to the wedge-shaped "solvency region" shown in Figure 1. The control is an R^3-valued function of time, $(\alpha(t), \beta(t), \gamma(t))$, where

- $\alpha(t) \geq 0$ is the rate at which money is being moved from money-market to high-yield at time t,

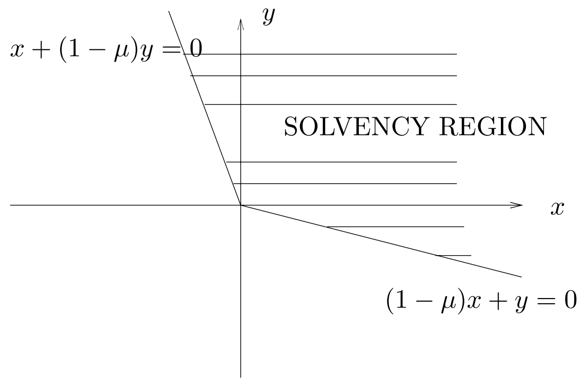

Figure 1: *The solvency region.*

- $\beta(t) \geq 0$ is the rate at which money is being moved from high-yield to money-market at time t,
- $\gamma(t) \geq 0$ is the consumption rate at time t.

The state equation is thus

$$\begin{aligned} dX/dt &= rX - \alpha + (1-\mu)\beta - \gamma \\ dY/dt &= RY + (1-\mu)\alpha - \beta, \end{aligned}$$

with initial conditions $X(0) = x, Y(0) = y$, and the value function is

$$u(x, y) = \max_{\alpha,\beta,\gamma} \int_0^\infty \frac{1}{p}\gamma^p(s)e^{-s}\, ds.$$

This is an infinite-horizon problem; the value function $u(x, y)$ gives the (discounted) present utility of all future consumption, if the investor's present position is (x, y) and he behaves optimally (and never dies).

Some comments about the formulation: We permit the investor to "move" money from money-market to high-yield even when $X < 0$; this amounts to borrowing money at the money-market rate to purchase the high-yield investment. Similarly he can "move" money from high-yield to money-market even when $Y < 0$. Of course the transaction costs should be smaller than the transactions themselves, so we require $0 < \mu < 1$. If the investor moves a certain amount of money all at once at time t_* then α or β is formally infinite (like a "delta-function") at t_* and $X(t), Y(t)$ are discontinuous; such policies can be approximated, of course, by more regular ones for which $X(t)$ and $Y(t)$ remain continuous but change rapidly.

Analysis of Example 3. The main purpose of introducing this example was to show how transaction costs can be included in the optimization of investment and consumption. The analysis of this example gets somewhat technical. Let us do as much as we can without getting bogged down in tedious detail.

Homogeneity. The value function satisfies

$$u(\lambda x, \lambda y) = \lambda^p u(x, y)$$

for any $\lambda > 0$. The proof is exactly like the argument given in Section 1 for Example 1. This result is special to the power-law utility, and the hypothesis that the cost of a transaction is proportional to its size.

The $x > 0$ solvency boundary. It's easy to see that if the initial portfolio (x, y) lies on the $x > 0$ part of the insolvency boundary then the only possible policy is to liquidate immediately, transfering funds from money-market to pay off the short high-yield position, effectively jumping to $(0, 0)$. Thus $u = 0$ on this part of the solvency boundary. In fact: our hypothesis is $(1 - \mu)x + y = 0$, and combining the state equations gives

$$\frac{d}{ds}[(1-\mu)X + Y] = r[(1-\mu)X + Y] + (R - r)Y + [(1-\mu)^2 - 1]\beta - \gamma.$$

The right hand side is nonpositive at $s = 0$, but it cannot be negative because this would make the trajectory leave the solvency region. This forces $\beta = \gamma = 0$, but more: it forces the liquidation to take place immediately, since otherwise the $(R - r)Y$ term would push the solution out of the solvency region.

Restrictions on parameters. We expect u to be infinite if the interest rates (r and R) are large enough compared to the discount rate for utility (which we have set to 1). This lesson is clear from Problem 2 of HW1.

To derive a restriction on r and R, also gain some insight concerning the solution, consider what happens if the initial portfolio is on the $y > 0$ part of the solvency boundary. Then a plausible strategy is to set $\alpha = \beta = 0$, and choose $\gamma(s)$ in such as way as to keep $(X(s), Y(s))$ on the the solvency boundary. To find γ, observe that our strategy implies $X(s) + (1 - \mu)Y(s) = 0$, $Y(s) = ye^{Rs}$, and $\dot{X} = rX - \gamma$. A little algebra gives

$$\gamma(s) = (R - r)(1 - \mu)ye^{Rs}$$

and this gives discounted utility of consumption

$$\frac{1}{p}\int_0^\infty (R - r)^p (1-\mu)^p y^p e^{(Rp-1)s}\, ds = \frac{(1-\mu)^p (R-r)^p}{p(1 - Rp)} y^p$$

provided $pR < 1$. The integral would be infinite if $pR \geq 1$. Thus to have an everywhere-finite value function we must require

$$pR < 1. \tag{7}$$

The strategy just considered is plausible, but it is not the only possibility[1]. Another way to stay on the $y > 0$ solvency boundary is to transfer funds continuously from the money

[1] I suggest skipping the following discussion; it's really a technicality. Go to the paragraph on the HJB equation.

market account into the high yield account (really: borrow money at the money-market rate and use it to invest at the high-yield rate). Suppose this is done using

$$\alpha(s) = \frac{1 - pR}{1 - \mu} c_0 Y(s)$$

for some $c_0 > 0$. Then a bit of calculation gives

$$Y(s) = ye^{c_1 s} \quad \text{with } c_1 = R + (1 - pR)c_0,$$

and the relations $X(s) + (1 - \mu)Y(s) = 0$, $\dot{X} = rX - \gamma - \alpha$ determine that

$$\gamma(s) = \left\{ (1 - \mu)(R - r) - \frac{2\mu - \mu^2}{1 - \mu}(1 - pR)c_0 \right\} ye^{c_1 s}.$$

One can look for the condition that the payoff be finite, and if it is finite then one can optimize c_0. Let's look instead for the condition that this policy be no better than the one considered a moment ago, i.e. the condition that the optimal c_0 be $c_0 = 0$. For any c_0 the discounted utility of consumption is

$$\frac{1}{p} y^p \frac{\left\{ (1 - \mu)(R - r) - \frac{2\mu - \mu^2}{1 - \mu}(1 - pR)c_0 \right\}^p}{(1 - Rp)(1 - c_0 p)}$$

provided $c_0 p < 1$. The condition that the optimal c_0 be 0 is

$$(1 - \mu)(R - r) < \frac{2\mu - \mu^2}{1 - \mu}(1 - pR), \tag{8}$$

which is stronger than (7).

The Hamilton-Jacobi-Bellman equation. The HJB equation is easy to derive. We revert to the generic Section 1 notation: the state equation is $dy/ds = f(y, \alpha)$ with $y(0) = x$, and the value function is $u(x) = \max_\alpha \int_0^\infty e^{-s} h(y, \alpha)\, ds$. Arguing in our usual heuristic way:

$$u(x) \geq h(x, a)\Delta t + e^{-\Delta t} u(x + f(x, a)\Delta t).$$

Taylor expanding u, taking the limit $\Delta t \to 0$, and optimizing in a gives

$$0 = \max_a \{h + f \cdot \nabla u\} - u.$$

(We have essentially just done problem 1(a) of HW1.)

Now we change back to the notation of our example to implement this. It says
$H(u_x, u_y, x, y) - u = 0$ with

$$H = \max_{\alpha,\beta,\gamma>0} \left\{ \frac{1}{p}\gamma^p + (rx - \alpha + (1 - \mu)\beta - \gamma)u_x + (Ry + (1 - \mu)\alpha - \beta)u_y \right\}.$$

Finiteness of H requires

$$-u_x + (1 - \mu)u_y \leq 0 \quad \text{with strict negativity implying } \alpha = 0;$$

$$(1-\mu)u_x - u_y \le 0 \quad \text{with strict negativity implying } \beta = 0\text{; and}$$
$$u_x \ge 0.$$

The optimal γ satisfies $\gamma^{p-1} = u_x$, whence $\frac{1}{p}\gamma^p - \gamma u_x = \frac{1-p}{p}u_x^{p/(p-1)}$. With this substitution, and assuming the finiteness conditions above,

$$H = \frac{1-p}{p}u_x^{p/(p-1)} + rxu_x + Ryu_y.$$

Pontryagin maximum principle. What does the maximum principle tell us about optimal trajectories? The discussion given earlier must be modified slightly to take into account the discounting of future utility. This is easy: if the utility is $\int_0^\infty e^{-\rho s}h(y,\alpha)\,ds$ one can repeat the argument used earlier with $\pi(s)$ replaced by $\pi(s)e^{-\rho s}$ to see that

$$\frac{dy}{ds} = \nabla_\pi H(\pi, y), \quad \frac{d\pi}{ds} = -\nabla_y H + \rho\pi.$$

In the present setting $\rho = 1$ and the Hamiltonian is

$$H = H(\pi_1, \pi_2, x, y) = \frac{1-p}{p}{\pi_1}^{p/(p-1)} + rx\pi_1 + Ry\pi_2$$

so we get

$$\frac{dX}{ds} = rX - {\pi_1}^{1/(p-1)}, \quad \frac{dY}{ds} = RY$$

and

$$\frac{d\pi_1}{ds} = (1-r)\pi_1, \quad \frac{d\pi_2}{ds} = (1-R)\pi_2$$

wherever u is smooth with $-u_x + (1-\mu)u_y < 0$ and $(1-\mu)u_x - u_y < 0$. The solutions correspond to a very simple strategy: do no transfers between accounts, and consume at the specified rate from the money-market account.

A proposal for the optimal strategy. Based on the preceding considerations, we expect that an optimal strategy can have three components: (i) an initial, instantaneous transfer if necessary to put sufficient funds in the high-yield account; (ii) otherwise do no transactions (thus incurring no transaction fees) while solvent; (iii) transact as necessary to stay solvent, once you reach the solvency boundary. This strongly suggests that (if the high-yield rate is high enough, i.e. if (8) holds) we have the following picture:

- The solvency region is divided into two parts by a ray $x = h_0 y$ (see Figure 2).
- To the right of this ray the optimal policy is to move immediately to this ray by transfering funds from money-market to high-yield. Consequently, to the right of this ray u is constant along each line $(1-\mu)x + y =$ constant. This wedge might be called the "transfer from money-market region".
- To the left of this ray the optimal policy is to do no transactions, and to consume at rate $\gamma = u_x^{1/(p-1)}$. This wedge might be called the "no transaction region". The portfolio will approach and eventually hit (in finite time) the $y > 0$ part of the solvency boundary.

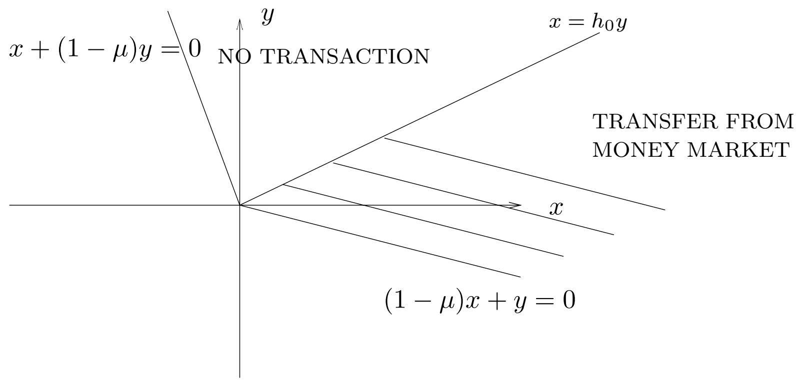

Figure 2: *Different regimes.*

- Once the portfolio is on the solvency boundary the optimal policy is to stay on this boundary, doing no transactions. The associated consumption law and the value of u along this boundary were described above in the paragraph concerning restrictions on the parameters.

This proposal is of course fully consistent with what we obtained from the Pontryagin maximum principle. Notice however that the trajectories we obtained from the maximum principle only describe the evolution in the "no transaction" sector. That's because we assumed, in our discussion of the maximum principle, that $-u_x + (1-\mu)u_y < 0$ and $\alpha = 0$. In the "transfer from money market" sector we have a different situation: α is infinite and $-u_x + (1-\mu)u_y = 0$.

The value of h_0 can be found by considering the strategy just proposed and optimizing the total discounted utility with respect to h_0. Or we can use the homogeneity property to obtain the result more or less analytically. The homogeneity property implies that

$$u(x, y) = y^p \phi(x/y)$$

for some function $\phi(\xi)$. The last bullet above determines u on the $y > 0$ part of the solvency boundary, so it determines $\phi(\mu - 1)$:

$$\phi(\mu - 1) = \frac{(1-\mu)^p (R-r)^p}{p(1-Rp)}. \tag{9}$$

The third bullet says that in the no-transaction region u solves $-u + \frac{1-p}{p} u_x^{p/(p-1)} + rxu_x + Ryu_y = 0$. Since

$$u_x = y^{p-1}\phi'(x/y) \quad u_y = y^{p-1}\left[p\phi(x/y) - (x/y)\phi'(x/y)\right]$$

we deduce a first-order differential equation for $\phi(\xi)$:

$$-\phi + \frac{1-p}{p}(\phi')^{p/(p-1)} + r\xi\phi'(\xi) + Rp\phi - R\xi\phi'(\xi) = 0. \tag{10}$$

Remember we need $(1-\mu)u_y - u_x \le 0$; in terms of ϕ this restriction reads

$$(1-\mu)\left[p\phi(\xi) - \xi\phi'(\xi)\right] - \phi'(\xi) \le 0. \tag{11}$$

The parameter restriction assures that this is true "initially", at $\xi = \mu - 1$. We may treat (9) and (10) as an initial-value problem for ϕ, but we must stop when (11) reaches 0. This determines a critical value of ξ, which we named named h_0 in the first bullet. For $\xi > h_0$ the value of ϕ is determined by extending u to the "transfer from money-market" region so it is constant on lines $(1-\mu)x + y =$ constant, i.e. so that $(1-\mu)u_y - u_x = 0$.

We have not used the restriction $(1-\mu)u_x - u_y \le 0$. This is obvious in the "transfer from money market" region, since there we have $(1-\mu)u_y - u_x = 0$. It should hold in the "no transaction" region too, and I'm sure it does (from Shreve, Soner, and Xu), but I don't yet have a simple argument. In terms of ϕ this amounts to the assertion that the function ϕ constructed above satisfies

$$[(1-\mu)+\xi]\phi'(\xi) - p\phi(\xi) \le 0 \quad \text{for } \mu - 1 \le \xi \le h_0.$$

(Please tell me if you see how to finish this detail.)

Verification. This proposal is certainly plausible, but how might one show it is optimal? There's only one game in town – a verification argument similar to that in Section 1. Implementing it is a bit technical, because of the singular character of the optimal control (which calls for an instantaneous, large transaction if the initial portfolio falls in the "transfer from money-market" sector). But the essential idea is no different from that of Section 1. We do not attempt to give the details here.

****

The Hopf-Lax solution formula for $u_t + H(\nabla u) = 0$. (This discussion repeats, with much reorganization and some elaboration, the final bit of Section 1.) We turn now to a different topic. It can be viewed in two different (essentially equivalent) ways:

(a) A discussion of the special class of finite-horizon control problems with equation of state

$$dy/ds = \alpha(s), \tag{12}$$

with initial value $y(t) = x$ and value function

$$u(x,t) = \max_{\alpha} \left[\int_t^T h(\alpha(s))\, ds + g(y(T)) \right] \tag{13}$$

when $h = h(\alpha)$ is a *concave* function of α alone.

(b) A discussion of the final-value problem for a Hamilton-Jacobi equation

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

when $H = H(\nabla u)$ is a *convex* function of ∇u alone. We shall derive a sort of solution formula, representing the unique "dynamic programming" solution of this equation.

As usual, we introduce this topic with a specific purpose in mind. Applying it to analyze the time-dependent Hamilton-Jacobi equation $u_t + \frac{1}{2}|u_x|^2 = 0$, we shall see that almost-everywhere solutions are far from unique, and dynamic-programming solutions can have discontinuous gradients. This example is more or less the dynamical analogue of our Example 2 (the minimum time problem whose HJB equation was $|\nabla u| = 1$). It lays the groundwork for discussing the effect of small stochastic perturbations and the notion of a viscosity solution.

To begin let us focus on viewpoint (a). The key observation is that for the special equation of state (12), and a concave running utility $h(\alpha)$, *the optimal control must be constant*, yielding a constant-velocity path in state space. In fact, since $dy/ds = \alpha$, the control is the precisely the velocity of the path $y(s)$. The concavity of h gives

$$h[\text{average velocity}] \geq \text{average of } h[\text{velocity}].$$

By the fundamental theorem of calculus, the average velocity of a path depends only on its endpoints, since

$$\frac{1}{T-t}\int_t^T \frac{dy}{ds}\,ds = \frac{1}{T-t}(y(T)-y(t)).$$

It follows that replacing any path by one with the same endpoints and constant velocity can only improve the utility. Indeed, for any path $y(s)$ with velocity $dy/ds = \alpha(s)$,

$$\begin{aligned}\int_t^T h(\alpha(s))\,ds &= (T-t)\cdot\text{average of } h[\text{velocity}]\\ &\leq (T-t)\cdot h[\text{average velocity}]\\ &\leq (T-t)\cdot h\left[(y(T)-y(t))/(T-t)\right]\end{aligned}$$

and the last expression is the running utility of the constant-velocity path with the same endpoints. We deduce a sort of "solution formula:"

$$u(x,t) = \max_z \left\{(T-t)h\left(\frac{z-x}{T-t}\right) + g(z)\right\}. \tag{14}$$

Here z represents the state at time T, the only remaining unknown.

Of course u solves the HJB equation associated with this control problem:

$$u_t + H(\nabla u) = 0 \quad \text{for } t < T \text{ with } u = g \text{ at } t = T \tag{15}$$

with Hamiltonian

$$H(p) = \max_p \left[a\cdot p + h(a)\right]. \tag{16}$$

This H is automatically convex, since (16) specifies it as the maximum of a family of linear functions. The HJB equation can have many almost-everywhere solutions (as we shall show in a moment). The value of the associated optimal control problem is, however, unique; let us call it the "dynamic programming" solution of the HJB equation. In this setting, the expression (14) for u is known as the *Hopf-Lax solution formula* of the HJB equation (15).

Example 4. Let's bring this down to earth by considering the specific example $h(a) = -\frac{1}{2}|a|^2$. The associated Hamiltonian is $H(p) = \frac{1}{2}|p|^2$, so the HJB equation is $u_t + \frac{1}{2}|\nabla u|^2 = 0$.

The following discussion extends easily to $x \in R^n$, but the issues that interest us arise already in one space dimension, so we focus on that case. The HJB equation is thus

$$u_t + \frac{1}{2}u_x^2 = 0, \quad u(x,T) = g(x), \tag{17}$$

and the solution formula is

$$u(x,t) = \max_z \left\{ g(z) - \frac{|z-x|^2}{2(T-t)} \right\}.$$

The observations we wish to make are three-fold:

(i) The HJB equation has infinitely many almost-everywhere solutions, only one of which agrees with the solution formula.

(ii) The dynamic programming solution of the HJB equation can have discontinuous derivatives: in fact, the graph of u as a function of x can have sharp valleys.

(iii) The graph of the dynamic programming solution cannot have sharp peaks.

(iv) The HJB equation is formally invariant under changing u to $-u$. However the dynamic programming solution with final-time data $-g$ is not, in general, $-u$.

Concerning (a): consider the case $g = 0$. Then the dynamic programming solution is clearly $u(x,t) = 0$ for all $t < T$. However the PDE has lots of other solutions: for example the function

$$u(x,t) = \begin{cases} \frac{1}{2}(T-t) - |x| & \text{if } |x| \le \frac{1}{2}(T-t) \\ 0 & \text{otherwise.} \end{cases}$$

is a solution (see figure 3).

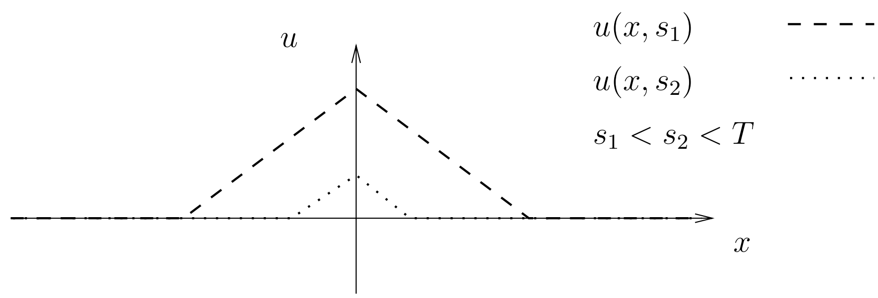

Figure 3: *Nonuniqueness of almost-everywhere solutions.*

There is nothing special about $x = 0$: a similar "bump" could emerge from any point on the x axis. Thus there are infinitely many almost-everywhere solutions of (17) when $g = 0$. This should not be a great surprise: we already noted the analogous nonuniqueness for the eikonal equation $|\nabla u| = 1$, the HJB equation associated with our geometrical minimum-time problem.

Concerning (b): consider the case $g(y) = |y|$. Our solution formula for the dynamic programming solution is readily evaluated; the optimal z is $z = x + (T-t)x/|x|$, leading to

$$u(x,t) = \frac{T-t}{2} + |x|,$$

which is clearly not smooth (see figure 4).

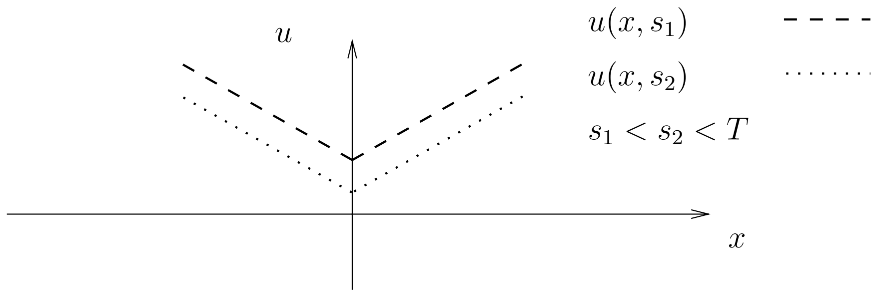

Figure 4: *The solution with final data* $g(y) = |y|$.

Toward (c) and (d): consider what happens if we change the sign of g. When $g(y) = -|y|$ the solution formula is again easy to evaluate, but now the best z is $z = x - (T-t)x/|x|$ if $|x| \geq (T-t)$, and $z = 0$ if not. This leads to

$$u(x,t) = \begin{cases} (T-t)/2 - |x| & \text{if } |x| \geq (T-t) \\ -|x|^2/2(T-t) & \text{otherwise,} \end{cases}$$

(see figure 5), which is clearly smooth for $t < T$, and clearly quite different from the negative of the solution obtained when $g = |y|$.

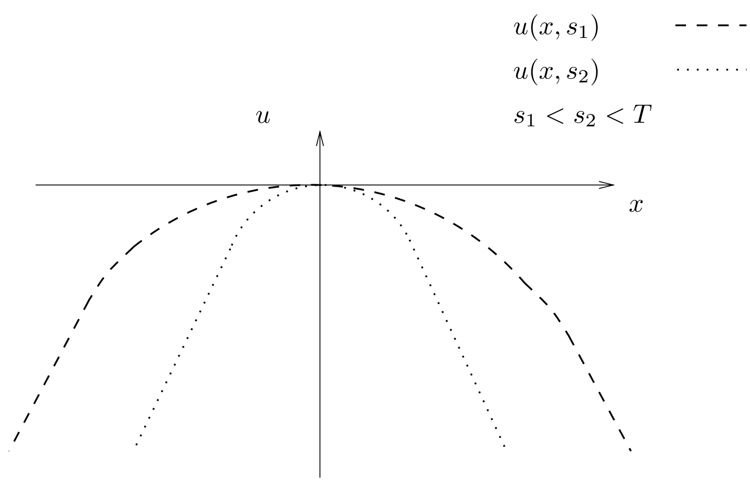

Figure 5: *The solution with final data* $g(y) = -|y|$.

Our assertion (c) is rather bold. Our example supports it, but of course it is just an example. Let us go a bit deeper, since (c) provides good intuition about the role of *viscosity solutions*. One definition – not the best, but historically the first definition – of the viscosity solution of a HJB equation $u_t + H(\nabla u) = 0$ with $u = g$ at $t = T$ is $u = \lim_{\epsilon\to 0} u^\epsilon$, where $u^\epsilon(x,t)$ solves the singularly perturbed problem

$$u^\epsilon_t + H(\nabla u^\epsilon) + \epsilon \Delta u^\epsilon = 0 \quad \text{for } t < T \text{ with } u^\epsilon = g \text{ at } t = T$$

with $\epsilon > 0$. (As shall discuss very soon, u^ϵ represents the value function of a *stochastic control* problem obtained by introducing a bit of noise in the equation of state.) Let us accept that this viscosity solution is equal to the dynamic programming solution. Let us also accept that u^ϵ is smooth when $\epsilon > 0$. For our 1D example $u_t + \frac{1}{2}u_x^2 = 0$ we expect u_x to change sign across a peak or valley (so u_x^2 is the same on both sides) and we expect $u_t = -\frac{1}{2}u_x^2$ to be negative. (This is of course consistent with all our examples.) Consider what would happen if the graph of u as a function of x had a sharp peak. Then u^ϵ, the solution of $u^\epsilon_t + \frac{1}{2}(u^\epsilon_x)^2 + \epsilon u^\epsilon_{xx} = 0$, would have associated rounded peak. At the top of this rounded peak we would have $u^\epsilon_x = 0$ and $u^\epsilon_{xx} \leq 0$, whence $u^\epsilon_t = -\epsilon u^\epsilon_{xx} \geq 0$. Passing to the limit $\epsilon \to 0$ we conclude (somewhat formally) that $u_t \geq 0$ at the top of the peak of u. But we argued a moment ago that $u_t < 0$. Having reached a contradiction, we conclude that there should be no peaks.

There can, of course, be valleys. Our argument does nothing to rule them out. In fact since $u^\epsilon_{xx} \geq 0$ at the bottom of a valley an argument parallel to the above simply says $u_t \leq 0$ at any valley, confirming something we already knew.

We are done with the example. It remains to discuss point (b) above, namely the use of the Hopf-Lax solution formula to define the (dynamic programming) solution of a general Hamilton Jacobi equation $u_t + H(\nabla u) = 0$ when H is convex. The idea is simple: given a convex H, there is a unique concave h such that

$$H(p) = \max_a \{a \cdot p + h(a)\},$$

namely

$$h(a) = \min_p \{H(p) - a \cdot p\}. \tag{18}$$

This h is indeed concave, since it is defined as a minimum of linear functions. To justify our assertion we must show that when H is convex and h is given by (18), the resulting Hamiltonian

$$\max_a \{a \cdot p + h(a)\}$$

is equal to $H(p)$. The definition of h shows that

$$h(a) \leq H(p) - a \cdot p$$

for all p, with equality when p is chosen optimally (say, $p = p_*(a)$). It is a fact of convex analysis that $a \to p_*(a)$ is invertible; in other words, each p arises as $p_*(a)$ for some a. Therefore we can rewrite the preceding relation as

$$H(p) \geq a \cdot p + h(a)$$

for all a and p, with equality when a is chosen optimally (depending on p). Maximization over a gives the desired relation.

The calculation we just did is is best understood in a more general context, as a fact about Fenchel transforms. For any function $F(z)$ defined for $z \in R^n$, its Fenchel transform is defined as

$$F^*(w) = \max_z \{w \cdot z - F(z)\}.$$

Here are two basic facts about the Fenchel transform:

- The double Fenchel transform $(F^*)^*(z)$ is the *convexification* of F, i.e. its graph is the convex hull of the graph of F.
- If F is convex then its double Fenchel transform $(F^*)^*$ is equal to F.

The proof of these facts follows the argument sketched briefly above. The fact that when H is convex the choice (18) leads to a HJB with Hamiltonian H is just an application of the second bullet, since (18) is equivalent to $h = -H^*$.

In conclusion: if H is convex, we can define a "dynamic programming" solution of the final-value problem for $u_t + H(\nabla u) = 0$. This solution is the value function of the finite-horizon dynamic programming problem with running utility (18), and it is characterized by the Hopf-Lax formula (14)

****

Suggestions for further study. We shall be turning to stochastic optimal control. Of course there's lots more we could have done in the deterministic setting. Important topics we've barely touched include:

Viscosity solutions. Chapter 10 of Evans is by far the best place to start. Key things to watch for: (a) the modern definition of viscosity solution; (b) smooth solutions of the HJB equation are automatically viscosity solutions; (c) the value function of a dynamic programming problem is automatically a viscosity solution of the HJB equation; (d) viscosity solutions are unique.

Numerical approach by direct solution of a discretized control problem. We started our discussion of dynamic programming this way, but we did not attempt to give a practical numerical algorithm. Some schemes based on Runge-Kutta time-discretization are analyzed in W.W. Hager, *Rate of convergence for discrete approximations to unconstrained control problems*, SIAM J. Numer. Anal. 13 (1976) 449-471. A recent survey is A.L. Dontchev, *Discrete approximations in optimal control*, in Nonsmooth Analysis and Geometric Methods in Deterministic Optimal Control (B.S. Mordukhovich and H.J. Sussman, eds.), Springer-Verlag (1996) 59-80.

Numerical approach by discretization of the HJB equation. This method gets the value function directly, without any explicit attention to the optimal control. (Of course we usually want the optimal control; it can generally be expressed in terms of ∇u once that has been determined). The best starting place I know is J.A. Sethian, Level Set Methods and Fast Marching Methods, Cambridge University Press (2nd edn, 1999). This book is not really about dynamic programming; even so, Chapter 6 provides a fine introduction to numerical methods for Hamilton-Jacobi equations. The rest of the book is great fun to read too.