

A Brief Introduction to Optimal Control, Fall 2009

Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences. Prepared for my Fall 2009 class on the Calculus of Variations.

Topics considered here include: examples of optimal control problems; dynamic programming and the Hamilton-Jacobi-Bellman equation; verification theorems; the Pontryagin Maximum Principle. The examples include many with an economic flavor, but others too (including the Hopf-Lax solution formula for $u_t + H(Du) = 0$ with H convex). There's much more here than we'll have time to do in lecture.

These notes are an abridged version of material I prepared some years ago for the class "PDE for Finance." So they do not go deeply into the associated PDE issues such as viscosity solutions. Evans' PDE book is highly recommended for further reading in that direction. A brief treatment emphasizing economic applications (same level, more or less, as these notes) can be found in chapters 10 and 11 of A.K. Dixit, *Optimization in Economic Theory* (Oxford Univ Press, 1990). The book by Jost and Li-Jost has chapters on Hamilton-Jacobi equations and the Pontryagin Maximum Principle; they're worth reading, but quite different from my approach and rather terse, with relatively few examples.

What is optimal control? A typical problem of optimal control is this: we have a system whose *state* at any time t is described by a vector $y = y(s) \in R^n$. The system evolves in time, and we have the ability to influence its evolution through a vector-valued *control* $\alpha(s) \in R^m$. The evolution of the system is determined by an *ordinary differential equation*

$$\dot{y}(s) = f(y(s), \alpha(s)), \quad y(0) = x, \quad (1)$$

and our goal is to choose the function $\alpha(s)$ for $0 < s < T$ so as to maximize some utility or minimize some cost, e.g.

$$\max \int_0^T h(y(s), \alpha(s)) ds + g(y(T)). \quad (2)$$

The problem is determined by specifying the dynamics f , the initial state x , the final time T , the "running utility" h and the "final utility" g . The problem is solved by finding the optimal control $\alpha(s)$ for $0 < s < T$ and the value of the maximum.

The mathematical and engineering literature often focuses on minimizing some sort of cost; the economic literature on maximizing utility. The two problems are mathematically equivalent.

One needs some hypotheses on f to be sure the solution of the ODE defining $y(s)$ exists and is unique. We do not make these explicit since the goal of these notes is to summarize the main ideas without getting caught up in fine points. See Evans for a mathematically careful treatment. Another technical point: it's possible (even easy) to formulate optimal control problems that have no solution. If the utility is bounded above, then for any $\epsilon > 0$ there's certainly a control $\alpha_\epsilon(s)$ achieving a value within ϵ of optimal. But the controls α_ϵ

might not converge to a meaningful control as $\epsilon \rightarrow 0$. Note however that even if an optimal control doesn't exist, the optimal *value* (the maximum utility) is still well-defined.

An optimal control problem is evidently a special type of optimization problem. What's special is that we're dealing with functions of time, and decisions that must be made as time proceeds. Often the optimal control is described by a *feedback law*. Such a law determines the optimal control $\alpha(s)$ as having the form $\alpha(s) = F(y(s), s)$ for some function F (the feedback law).

Example 1: Here is a simple example which already has financial interest. (It's a deterministic version of Merton's famous example of optimal investment and consumption; we'll do the version with investment in a few weeks). Consider an individual whose wealth today is x , and who will live exactly T years. His task is to plan the rate of *consumption* of wealth $\alpha(s)$ for $0 < s < T$. All wealth not yet consumed earns interest at a fixed rate r . The state equation is thus

$$\dot{y} = ry - \alpha, \quad y(0) = x. \quad (3)$$

The control is $\alpha(s) \geq 0$, and the state is constrained by $y(s) \geq 0$ (he cannot consume wealth he doesn't have). The goal is

$$\max \int_0^T e^{-\rho s} h(\alpha(s)) ds$$

where ρ is the discount rate and $h(\alpha)$ is the utility of consumption. (The function h , which must be given as part of the formulation of the problem, should be monotonically increasing and concave. A typical choice is $h(\alpha) = \alpha^\gamma$ with $0 < \gamma < 1$.) We have, for simplicity, assigned no utility to final-time wealth (a bequest), so the solution will naturally have $y(T) = 0$. Our goal is not strictly of the form (2) due to the presence of discounting; well, we omitted discounting from (2) only for the sake of simplicity.

The state constraint $y(s) \geq 0$ is awkward to deal with. In practice it tells us that if the investor ever runs out of wealth (i.e. if $y(s)$ ever reaches 0) then $\alpha = 0$ and $y = 0$ thereafter. This state constraint can be avoided by reformulating the goal as

$$\max \int_0^\tau e^{-\rho s} h(\alpha(s)) ds$$

where τ is the first time y reaches 0 if this occurs before T , or $\tau = T$ if y is positive for all $s < T$. With this goal we need not impose the state constraint $y(s) \geq 0$.

Control theory is related to – but much more general than – the one-dimensional calculus of variations. A typical calculus of variations problem is

$$\max_{y(s)} \int_0^T W(s, y(s), \dot{y}) ds$$

subject, perhaps, to endpoint conditions on $y(0)$ and $y(T)$. The example just formulated can be expressed in this form,

$$\max_{y(s)} \int_0^T h(ry - \dot{y}) ds, \quad \text{subject to } y(0) = x,$$

except that we have additional constraints $ry(s) - \dot{y}(s) \geq 0$ and $y(s) \geq 0$ for all s .

We will shortly discuss the method of dynamic programming as a scheme for solving optimal control problems. The key to this method is to consider how the solution depends on the *initial time* and *initial state* as parameters. Thus rather than start arbitrarily at time 0, it is better to introduce a variable initial time t . And it is fruitful to consider the *value function* $u(x, t)$, the optimal value achievable using initial time t and initial state x . In the context of our basic framework (1) this means changing the state equation to

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

The control $\alpha(s)$ is now to be determined for $t < s < T$, and the value function is

$$u(x, t) = \max \int_t^T h(y(s), \alpha(s)) ds + g(y(T)).$$

In the context of Example 1 it means changing the state equation to

$$\dot{y} = ry - \alpha, \quad y(t) = x,$$

and the objective to

$$u(x, t) = \max \int_t^T e^{-\rho s} h(\alpha(s)) ds.$$

(Warning: with this definition $u(x, t)$ is the utility of consumption discounted to time 0. The utility of consumption discounted to time t is $e^{\rho t} u(x, t)$.)

We started by formulating the “typical” optimal control problem (1)-(2). Now let’s discuss some of the many variations on this theme, to get a better sense of the scope of the subject. We repeat for clarity the state equation:

$$\dot{y}(s) = f(y(s), \alpha(s)) \text{ for } t < s < T \text{ with initial data } y(t) = x.$$

Sometimes we may wish to emphasize the dependence of $y(s)$ on the initial value x , the initial time t , and the choice of control $\alpha(s)$, $t < s < T$; in this case we write $y = y_{x,t,\alpha}(s)$. The control is typically restricted to take values in some specified set A , independent of s :

$$\alpha(s) \in A \text{ for all } s;$$

the set A must be specified along with the dynamics f . Sometimes it is natural to impose *state constraints*, i.e. to require that the state $y(s)$ stay in some specified set Y :

$$y_{x,t,\alpha}(s) \in Y \text{ for all } s;$$

when present, this requirement restricts the set of admissible controls $\alpha(s)$. Our basic example (2) is known as a **finite horizon** problem; its value function is

$$u(x, t) = \max_{\alpha} \left\{ \int_t^T h(y_{x,t,\alpha}(s), \alpha(s)) ds + g(y_{x,t,\alpha}(T)) \right\}. \quad (4)$$

For the analogous **infinite horizon** problem it is customary to set the starting time to be 0, so the value function depends only on the spatial variable x :

$$u(x) = \max_{\alpha} \int_0^{\infty} e^{-\rho s} h(y_{x,0,\alpha}(s), \alpha(s)) ds. \quad (5)$$

Discounting is important for the infinite-horizon problem, since without it the integral defining u could easily be infinite. (As already noted in our example, it is also often natural to include discounting in a finite-horizon problem.)

The **minimum time** problem is a little bit different. It minimizes the time it takes $y(s)$ to travel from x to some target set \mathcal{T} . The value function is thus

$$u(x) = \min_{\alpha} \{\text{time at which } y_{x,0,\alpha}(s) \text{ first arrives in } \mathcal{T}\}. \quad (6)$$

The minimum time problem is somewhat singular: if, for some x , the solution starting at x cannot arrive in \mathcal{T} (no matter what the control) then the value is undefined. The **discounted minimum time** problem avoids this problem: its value function is

$$u(x) = \min_{\alpha} \int_0^{\tau(x,\alpha)} e^{-s} ds \quad (7)$$

where $\tau(x,\alpha)$ is the time that $y_{x,0,\alpha}(s)$ first arrives in \mathcal{T} , or infinity if it never arrives. Notice that the integral can be evaluated: the quantity being minimized is $\int_0^{\tau(x,\alpha)} e^{-s} ds = 1 - e^{-\tau(x,\alpha)}$. So we're still minimizing the arrival time, but the value function is $1 - \exp(-\text{arrival time})$ instead of the arrival time itself.

Example 2. Here is a simple example of a minimum-time problem, with the great advantages that (a) we can easily visualize everything, and (b) we know the solution in advance. In its simplest form the problem is: given a point x in R^n , and a set \mathcal{T} not containing x , find the distance from x to \mathcal{T} . We recognize this as a minimum time problem, by reformulating it in terms of paths travelled with speed ≤ 1 . The state equation is

$$dy/ds = \alpha(s), \quad y(0) = x,$$

and the controls are restricted by

$$|\alpha(s)| \leq 1.$$

The minimum arrival time

$$u(x) = \min_{\alpha} \{\text{time of arrival at } \mathcal{T}\}$$

is of course the distance from x to \mathcal{T} , and the optimal strategy is to travel with constant velocity (and unit speed) toward the point in \mathcal{T} that is closest to x . We remark that $u(x) = \text{dist}(x, \mathcal{T})$ solves the differential equation

$$|\nabla u| = 1$$

in its natural domain $\Omega = R^n - \mathcal{T}$, with boundary condition $u = 0$ at $\partial\Omega$. This is an example of a (time-independent) Hamilton-Jacobi equation. The solution is typically not

smooth: consider for example the case when Ω is a circle or a square. The optimal control is determined by a feedback law (“wherever you are right now, proceed at unit speed toward the nearest point on the target \mathcal{T} ”). The non-smoothness of u reflects the fact that the feedback law is discontinuous, with nonuniqueness where ∇u is discontinuous. There is clearly nothing pathological about this example: non-smooth value functions, and discontinuous feedback laws, are commonplace in deterministic optimal control.

Dynamic programming. There are basically two systematic approaches to solving optimal control problems: one known as the Pontryagin Maximum Principle, the other known as Dynamic Programming. The two approaches are fundamentally equivalent, though in specific problems one may be easier to apply than the other. We’ll discuss the Pontryagin Maximum Principle later.

The essence of dynamic programming is pop psychology: “today is the first day of the rest of your life.” More: every day is the first day of the future thereafter. How to use this insight? One way is to make it the basis of a numerical solution scheme. Another way is to use it to derive a PDE for $u(x, t)$. These two ideas are of course closely related: our numerical solution scheme is in fact a crude numerical scheme for solving the PDE.

Let’s start with the numerical scheme, concentrating on the finite-horizon problem (4), and keeping space one-dimensional for simplicity. Our goal is to compute (approximately) the value function $u(x, t)$. Of course any numerical scheme must work in discrete space and time, so t is a multiple of Δt , and x is a multiple of Δx . It’s also natural to consider that the controls are discretized: $\alpha(s)$ is piecewise constant with mesh Δt . Now work backward in time:

First Consider the problem with initial time $t = T$. In this case the dynamics is irrelevant. So are the control and the running utility. Whatever the value of x , the associated value function is $g(x)$. In other words: $u(x, T) = g(x)$.

Next Consider the problem with initial time $t = T - \Delta t$. Approximate the dynamics as

$$y(s + \Delta t) = y(s) + f(y(s), \alpha(s))\Delta t.$$

Since there is just one time interval between the initial time t and the final time $T = t + \Delta t$, and since the control is piecewise constant, the unknown is now just a single vector $\alpha = \alpha(t)$ (not a function). It is determined by optimization. We may approximate the objective integral by a sum (dropping terms of higher order in Δt), leading to

$$u(x, T - \Delta t) = \max_{\alpha} \{h(x, \alpha)\Delta t + g(x + f(x, \alpha)\Delta t)\}.$$

This must be evaluated for each x (i.e. every multiple of Δx), and the maximization over α must be done globally (we need the global optimum, not just a local optimum). For a real numerical scheme some further structure is needed here: we should solve a problem in a bounded spatial domain, and impose concavity hypotheses assuring that there are no local optima. For the present conceptual discussion let us

ignore such practical issues and proceed. (One might worry that when the spatial dimension is greater than 1 this scheme is utterly impractical, since the number of grid points x to be considered at each time t is of order $(\Delta x)^{-n}$ in dimension n . This worry is well-founded: our scheme is impractical in higher dimensions. However there are good numerical schemes for multidimensional problems. One option is to solve the Hamilton-Jacobi-Bellman equation we'll derive presently, using a suitable finite-difference scheme.) At the end of this step we have computed $u(\cdot, T - \Delta t)$ as a function of space.

Next Consider the problem with initial time $t = T - 2\Delta t$. For any initial state $x = y(t)$, the possible controls are now represented by a pair of vectors $\alpha(t), \alpha(t + \Delta t)$. However we can still solve the problem by considering just the current control $\alpha = \alpha(t)$, since the optimal choice of $\alpha(t + \Delta t)$ has already been determined in the course of evaluating $u(x, T - \Delta t)$. Making crucial use of the fact that the “running utility” is an integral in time, we may determine the optimal value $u(x, T - 2\Delta t)$ by solving

$$u(x, T - 2\Delta t) = \max_{\alpha} \{h(x, \alpha)\Delta t + u(x + f(x, \alpha)\Delta t, T - \Delta t)\}.$$

Here the unknown is just the control α to be used during the time interval from $T - 2\Delta t$ to $T - \Delta t$. The optimal α depends of course on x , and the optimization in α must be done for each choice of x separately. (Again, this is the conceptual but impractical version; numerical optimal control uses various workarounds to make it more practical.) At the end of this step we have computed $u(\cdot, T - 2\Delta t)$ as a function of space.

Continue The scheme continues, working backward time-step by time-step. Notice that for computing $u(x, T - (j + 1)\Delta t)$ we need only save the values of $u(x, T - j\Delta t)$. However if we wish to synthesize an optimal control starting at an arbitrary point x and time $t = T - (j + 1)\Delta t$ we must save much more information: namely the feedback law $\alpha = F(y, s)$, obtained in the course of calculating $u(y, s)$ for $s > t$. (This is the optimal initial-time-period value of the control, when the initial state is y and the initial time is s). This information permits us to synthesize the optimal control and solve the state equation at the same time: starting from x at time t , the state evolves by

$$y_{\alpha}(s + \Delta t) = y_{\alpha}(s) + f(y_{\alpha}(s), \alpha(s))\Delta t$$

with $\alpha(s)$ determined by

$$\alpha(s) = F(y_{\alpha}(s), s).$$

We remark that a similar philosophy can be used in many other settings. One example is this standard scheme for computing the shortest path between two nodes of a graph. Pick one of the nodes (call it an endpoint). Find all nodes that lie distance 1 from it, then all points that lie distance 2 from it, etc. Stop when the other endpoint appears in the set you come up with.

The discrete-time, discrete-space scheme described above can be viewed as a crude numerical scheme for solving the PDE satisfied by the value function. This is known as the

Hamilton-Jacobi-Bellman equation. We shall derive it, in essence, by taking the formal limit $\Delta t \rightarrow 0$ in our numerical discussion. This viewpoint can be used for all the optimal control problems we've discussed (finite-horizon, infinite-horizon, least-time, with or without discounting) but to fix ideas we concentrate on the usual finite-horizon example

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

where the controls are restricted by $\alpha(s) \in A$, and the state equation is

$$dy/ds = f(y(s), \alpha(s)) \text{ for } t < s < T \text{ and } y(t) = x.$$

(Space can be multidimensional here.) The Hamilton-Jacobi-Bellman equation in this case is

$$u_t + H(\nabla u, x) = 0 \quad \text{for } t < T \tag{8}$$

with

$$u(x, T) = g(x) \quad \text{at } t = T,$$

where H (the ‘‘Hamiltonian’’) is defined by

$$H(p, x) = \max_{a \in A} \{ f(x, a) \cdot p + h(x, a) \}. \tag{9}$$

(Note that p is a vector with the same dimensionality as x ; a is a vector with the same dimensionality as α .)

To explain, we start with the **dynamic programming principle**, which was in fact the basis of our discrete scheme. It says:

$$u(x, t) = \max_{\alpha} \left\{ \int_t^{t'} h(y_{x,t,\alpha}(s), \alpha(s)) ds + u(y_{x,t,\alpha}(t'), t') \right\} \tag{10}$$

whenever $t < t' < T$. The justification is easy, especially if we assume that an optimal control exists (this case captures the main idea; see Evans for a more careful proof, without this hypothesis). Suppose the optimal utility starting at x at time t is achieved by an optimal control $\alpha_{x,t}(s)$. Then the restriction of this control to any subinterval $t' < s < T$ must be optimal for its starting time t' and starting position $y_{x,t,\alpha}(t')$. Indeed, if it weren't then there would be a new control $\alpha'(s)$ which agreed with α for $t < s < t'$ but did better for $t' < s < T$. Since the utility is additive – the running utility is $\int_t^T h(y, \alpha) ds = \int_t^{t'} h(y, \alpha) ds + \int_{t'}^T h(y, \alpha) ds$ – this new control would be better for the entire time period, contradicting the optimality of α . Therefore in defining $u(x, t)$ as the optimal utility, we can restrict our attention to controls that are optimal from time t' on. This leads immediately to (10).

Now let us derive (heuristically) the Hamilton-Jacobi-Bellman equation. The basic idea is to apply the dynamic programming principle with $t' = t + \Delta t$ and let $\Delta t \rightarrow 0$. Our argument is heuristic because (i) we assume u is differentiable, and (ii) we assume the optimal control is adequately approximated by one that is constant for $t < s < t + \Delta t$. (Our goal, as usual, is to capture the central idea, referring to Evans for a more rigorous treatment.) Since Δt is

small, the integral on the right hand side of (10) can be approximated by $h(x, a)\Delta t$, where $a \in A$ is the (constant) value of α for $t < s < t + \Delta t$. Using a similar approximation for the dynamics, the dynamic programming principle gives

$$u(x, t) \geq h(x, a)\Delta t + u(x + f(x, a)\Delta t, t + \Delta t) + \text{errors we wish to ignore}$$

with equality when a is chosen optimally. Using the first-order Taylor expansion of u this becomes

$$u(x, t) \geq h(x, a)\Delta t + u(x, t) + (\nabla u \cdot f(x, a) + u_t)\Delta t + \text{error terms}$$

with equality when a is optimal. In the limit $\Delta t \rightarrow 0$ this gives

$$0 = u_t + \max_{a \in A} \{ \nabla u \cdot f(x, a) + h(x, a) \},$$

i.e. $u_t + H(\nabla u, x) = 0$ with H as asserted above. The final-time condition is obvious: if $t = T$ then the dynamics is irrelevant, and the total utility is just $g(x)$.

That was easy. Other classes of optimal control problems are treated similarly. Let's look at the minimum-time problem, where the state evolves by

$$dy/ds = f(y, \alpha), \quad y(t) = x,$$

and the controls are restricted by

$$\alpha(s) \in A \quad \text{for all } s$$

for some set A . The associated Hamilton-Jacobi-Bellman equation is

$$H(\nabla u, x) = -1 \quad \text{for } x \notin \mathcal{T}$$

with Hamiltonian

$$H(p, x) = \min_{a \in A} \{ f(x, a) \cdot p \} = 0.$$

The boundary condition is

$$u = 0 \quad \text{for } x \in \mathcal{T}.$$

To see this, we argue essentially as before: the value function (the time it takes to arrive at \mathcal{T}) should satisfy

$$u(x) \leq \Delta t + u(x + f(x, a)\Delta t) + \text{error terms}$$

for any $a \in A$, with equality when a is optimal. Using Taylor expansion this becomes

$$u(x) \leq \Delta t + u(x) + \nabla u \cdot f(x, a)\Delta t + \text{error terms}.$$

Optimizing over a and letting $\Delta t \rightarrow 0$ we get

$$1 + \min_{a \in A} \{ f(x, a) \cdot \nabla u \} = 0,$$

which is the desired equation.

Let us specialize this to Example 2. In that example the set A is the unit ball, and $f(y, \alpha) = \alpha$, so $H(p, x) = \min_{|a| \leq 1} p \cdot a = -|p|$ and the Hamilton-Jacobi equation becomes $|\nabla u| = 1$, as expected.

Solutions of the Hamilton-Jacobi-Bellman equation are not unique (at least, not when we understand “solution” in the naive almost-everywhere sense). For example, there are *many* Lipschitz continuous solutions of $|\nabla u| = 1$ in a square, with $u = 0$ at the boundary. If one were smooth we might prefer it – however there is *no* smooth solution. So, is the HJB equation really of any use?

The answer is yes, it’s very useful, for three rather distinct reasons. The first is obvious; the second is elementary but not obvious; the third is subtle, representing a major mathematical achievement that’s surprisingly recent (early 1980’s).

- (a) In deriving the HJB equation, we deduced a relation between the optimal control and the value of ∇u : briefly, $\alpha(s)$ achieves the optimum in the definition of $H(p, x)$ with $p = \nabla u(y(s), s)$ and $x = y(s)$. Thus the derivation of the HJB equation tells us the relation between the value function and the optimal control. In many settings, this argument permits us to deduce a feedback law once we know the value function.
- (b) The argument used for the HJB equation can often be reorganized to show that a conjectured formula for the value function is correct. This sort of argument is called a *verification theorem*.
- (c) There is a more sophisticated notion of “solution” of a Hamilton-Jacobi equation, namely the notion of a *viscosity solution*. Viscosity solutions exist, are unique, and can be computed by suitable numerical schemes. Moreover the value function of a dynamic programming problem is automatically a viscosity solution of the associated HJB equation. (Chapter 10 of Evans’ book gives an excellent introduction to the theoretical side of this topic. The book *Level Set Methods* by J. Sethian, Cambridge Univ Press, provides a readable introduction to the numerical side, concentrating on the special class of HJB equations associated with geometric evolution problems – closely connected with our minimum time example.)

Point (a) should be clear, and it will be illuminated further by various examples later on. Point (c) is an interesting story, but beyond the scope of the present discussion. Our present intention is to concentrate on point (b). We focus as usual on the setting of the finite-horizon problem. As usual, $u(x, t)$ denotes the value function (the maximal value achievable starting from state x at time t). Our plan is to develop schemes for proving upper and lower bounds on u . If we do a really good job the upper and lower bounds will coalesce – in which case they will fully determine u .

There’s always one type of bound that is easy. Since we’re maximizing utility, these are the lower bounds. Any scheme for choosing the control – for example a conjectured feedback law specifying $\alpha(s)$ as a function of $y(s)$ – provides a lower bound $v(x, t) =$ the value achieved by this scheme. The inequality

$$v(x, t) \leq u(x, t)$$

is obvious, since u is the maximal value obtainable using *any* control – including the ones used to define v .

The **verification theorem** provides the other bound. In its most basic form – specialized to the present setting – it says the following. Suppose $w(x, t)$ is defined (and continuously differentiable) for $t < T$, and it solves the Hamilton-Jacobi equation (8) with $w = g$ at $t = T$. Then w is an upper bound for the value function:

$$u(x, t) \leq w(x, t).$$

To see why, consider any candidate control $\alpha(s)$ and the associated state $y = y_{x,\alpha}(s)$ starting from x at time t . The chain rule gives

$$\begin{aligned} \frac{d}{ds}w(y(s), s) &= w_s(y(s), s) + \nabla w(y(s), s) \cdot \dot{y}(s) \\ &= w_s(y(s), s) + \nabla w(y(s), s) \cdot f(y(s), \alpha(s)) \\ &\leq w_s + H(\nabla w, y) - h(y(s), \alpha(s)) \\ &= -h(y(s), \alpha(s)), \end{aligned} \tag{11}$$

using for (11) the relation

$$H(p, y) = \max_{a \in A} \{f(y, a) \cdot p + h(y, a)\} \geq f(y, \alpha) \cdot p + h(y, \alpha)$$

with $y = y(s)$, $\alpha = \alpha(s)$, and $p = \nabla w(y(s), s)$. Now integrate in time from t to T :

$$w(y(T), T) - w(x, t) \leq - \int_t^T h(y(s), \alpha(s)) ds.$$

Since $w(y(T), T) = g(y(T))$ this gives

$$g(y(T)) + \int_t^T h(y(s), \alpha(s)) ds \leq w(x, t).$$

The preceding argument applies to any control $\alpha(s)$. Maximizing the left hand side over all admissible controls, we have

$$u(x, t) \leq w(x, t)$$

as asserted.

We presented the task of finding lower and upper bounds as though they were distinct, but of course they are actually closely correlated. A smooth solution w of the Hamilton-Jacobi equation comes equipped with its own feedback law (as discussed in point (a) above). It is natural to consider the lower bound v obtained using the feedback law associated with w . I claim that this v is equal to w . To see this, follow the line of reasoning we used for the verification theorem, noticing that (11) holds with equality if α is determined by the feedback associated with w . Therefore integration gives

$$w(x, t) = g(y(T)) + \int_t^T h(y(s), \alpha(s)) ds$$

and the right hand side is, by definition, $v(x, t)$. In conclusion: if w is a (continuously differentiable) solution of the HJB equation, satisfying the appropriate final-time condition too, then w is in fact the value function $u(x, t)$.

It sounds like a great scheme, and in many ways it is. There is however a small fly in the ointment. Sometimes the value function isn't continuously differentiable. (Consider, for example, the minimum time problem). In such a case our proof of the verification theorem remains OK for paths that avoid the locus of nonsmoothness – or cross it transversely. But there is a problem if the state should happen to hug the locus of nonsmoothness. Said more plainly: if $w(x, t)$ has discontinuous derivatives along some set Γ in space-time, and if a control makes $(y(s), s)$ move along Γ , then the first step in our verification argument

$$\frac{d}{ds}w(y(s), s) = w_s(y(s), s) + \nabla w(y(s), s) \cdot \dot{y}(s)$$

doesn't really make sense (for example, the right hand side is not well-defined). Typically this problem is overcome by using the fact that the verification argument has some extra freedom: it doesn't really require that w solve the HJB equation exactly. Rather, it requires only that w satisfy the *inequality* $w_t + H(\nabla w, t) \leq 0$.

To give an example where this extra freedom is useful consider our geometrical Example 2, with target \mathcal{T} the complement of the unit square in R^2 . The HJB equation is $|\nabla u| = 1$ in Ω =unit square, with $u = 0$ at $\partial\Omega$. The value function is defined as $u(x)$ =minimum time of arrival to $\partial\Omega$ (among all paths with speed ≤ 1). Simple geometry tells us the solution is the distance function $\text{dist}(x, \partial\Omega)$, whose graph is a pyramid. We wish to give an entirely PDE proof of this fact.

One inequality is always easy. In this case it is the relation $u(x) \leq \text{dist}(x, \partial\Omega)$. This is clear, because the right hand side is associated with a specific control law (namely: travel straight toward the nearest point of the boundary, with unit speed). To get the other inequality, observe that if $w \leq 0$ at $\partial\Omega$ and $|\nabla w| \leq 1$ in Ω then

$$\begin{aligned} \frac{d}{ds}w(y(s)) &= \nabla w(y(s)) \cdot \dot{y}(s) \\ &= \nabla w(y(s)) \cdot \alpha(s) \\ &\geq -|\nabla w(y(s))| \geq -1. \end{aligned}$$

(Here $y(s)$ solves the state equation $\dot{y} = \alpha$, with initial condition $y(0) = x$ and any admissible control $|\alpha(s)| \leq 1$.) If τ is the time of arrival at $\partial\Omega$ then integration gives

$$w(y(\tau)) - w(x) \geq \int_0^\tau (-1) ds.$$

Since $w(y(\tau)) \leq 0$ we conclude that

$$w(x) \leq \tau.$$

Minimizing the right hand side over all admissible controls gives

$$w(x) \leq u(x).$$

We're essentially done. We cannot set w equal to the distance function itself, because this choice isn't smooth enough. However we can choose w to be a slightly smoothed-out version of the distance function minus a small constant. It's easy to see that we can approach the

distance function from below by such functions w . Therefore (using these w 's and passing to a limit)

$$\text{dist}(x, \partial\Omega) \leq u(x),$$

completing our PDE argument that the value function is in this case the distance function.

As another example, let's give the **solution to Example 1** for a power-law utility. The state equation is

$$\dot{y} = ry - \alpha, \quad y(t) = x$$

where x is the initial wealth and α is the consumption rate, restricted by $\alpha \geq 0$ (an explicit constraint on the controls). We consider the problem of finding

$$u(x, t) = \max_{\alpha} \int_t^T e^{-\rho s} \alpha^\gamma(s) ds,$$

which amounts to the utility of consumption with the power-law utility function $h(\alpha) = \alpha^\gamma$. Utility functions should be concave so we assume $0 < \gamma < 1$.

First, before doing any real work, let us show that the value function has the form

$$u(x, t) = g(t)x^\gamma$$

for some function $g(t)$. It suffices for this purpose to show that the value function has the homogeneity property

$$u(\lambda x, t) = \lambda^\gamma u(x, t), \tag{12}$$

for then we can take $g(t) = u(1, t)$. To see (12), suppose $\alpha(s)$ is optimal for starting point x , and let $y_x(s)$ be the resulting trajectory. We may consider the control $\lambda\alpha(s)$ for the trajectory that starts at λx , and it is easy to see that the associated trajectory is $y_{\lambda x}(s) = \lambda y_x(s)$. Using the power-law form of the utility this comparison demonstrates that

$$u(\lambda x, t) \geq \lambda^\gamma u(x, t).$$

This relation with λ replaced by $1/\lambda$ and x replaced by λx gives

$$u(x, t) \geq \lambda^{-\gamma} u(\lambda x, t),$$

completing the proof of (12).

Now let's find the HJB equation. This is almost a matter of specializing the general calculation to the case at hand. But we didn't have a discount term before, so let's redo the argument to avoid any doubt. From the dynamic programming principle we have

$$u(x, t) \geq e^{-\rho t} a^\gamma \Delta t + u(x + (rx - a)\Delta t, t + \Delta t) + \text{error terms}$$

with equality when $a \geq 0$ is chosen optimally. Using the first-order Taylor expansion of u this becomes

$$u(x, t) \geq e^{-\rho t} a^\gamma \Delta t + u(x, t) + (u_x(rx - a) + u_t)\Delta t + \text{error terms}$$

with equality when a is optimal. In the limit $\Delta t \rightarrow 0$ this gives

$$u_t + \max_{a \geq 0} \{u_x(rx - a) + e^{-\rho t} a^\gamma\} = 0.$$

This is the desired HJB equation, to be solved for $t < T$. The final-time condition is of course $u = 0$ (since no utility is associated to final-time wealth).

It's obvious that $u_x > 0$. (This follows from the observation that $u(x, t) = g(t)x^\gamma$. Or it's easy to prove using the original problem formulation and a suitable comparison argument.) Therefore the optimal a is easy to find, by differentiation, and it is positive:

$$\gamma a^{\gamma-1} = e^{\rho t} u_x.$$

This is the feedback law, determining the optimal control (once we know u_x). Remembering that $u(x, t) = g(t)x^\gamma$, we can write the feedback law as

$$a = \left[e^{\rho t} g(t) \right]^{1/(\gamma-1)} x$$

To find g (and therefore u) we substitute $u = g(t)x^\gamma$ into the HJB equation. This leads, after some arithmetic and cancelling a common factor of x^γ from all terms, to

$$\frac{dg}{dt} + r\gamma g + (1 - \gamma)g(e^{\rho t} g)^{1/(\gamma-1)}.$$

This equation (with the end condition $g(T) = 0$) is entirely equivalent to the original HJB equation. It looks ugly, however it is not difficult to solve. First, multiply each term by $e^{\rho t}$ to see that $G(t) = e^{\rho t} g(t)$ solves

$$G_t + (r\gamma - \rho)G + (1 - \gamma)G^{\gamma/(\gamma-1)} = 0.$$

Next, multiply by $(1 - \gamma)^{-1} G^{\gamma/(1-\gamma)}$ to see that $H(t) = G^{1/(1-\gamma)}$ satisfies the linear equation

$$H_t - \mu H + 1 = 0 \quad \text{with } \mu = \frac{\rho - r\gamma}{1 - \gamma}.$$

This is a linear equation! The solution satisfying $H(T) = 0$ is

$$H(t) = \lambda^{-1} \left(1 - e^{-\lambda(T-t)} \right).$$

Unraveling our changes of variables gives finally

$$g(t) = e^{-\rho t} \left[\frac{1 - \gamma}{\rho - r\gamma} \left(1 - e^{-\frac{(\rho - r\gamma)(T-t)}{1-\gamma}} \right) \right]^{1-\gamma}.$$

We've solved the HJB equation. Have we actually found the value function and the optimal feedback (consumption) policy? Yes indeed. The verification theorem given above supplies the proof. (Well, it should be redone with discounting, and with the more precise formulation of the objective which integrates the utility only up to the first time τ when $y = 0$, if this occurs before T . These modifications require no really new ideas.) Nothing fancy is needed since $u(x, t)$ is smooth.

The **Hopf-Lax solution formula** provides another class of examples. Consider what becomes of the finite-horizon problem when f (the right hand side of the state equation) and h (the running utility) are independent of x . The state equation is thus

$$dy/ds = f(\alpha(s)) \quad \text{with } y(t) = x$$

and the value function is

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

The HJB equation has the form $u_t + H(\nabla u) = 0$ for $t < T$, with final-time condition $u = g$ at $t = T$ and with Hamiltonian

$$H(p) = \max_{a \in A} \{f(a) \cdot p + h(a)\}.$$

Notice that H is always a *convex* function of p , since the definition expresses it as a maximum of linear functions of p .

The function H does not determine f and h – different f 's and h 's can lead to the same HJB equation. But for any convex H there's an especially simple choice of an associated f and h , namely

$$f(a) = a, \quad h(a) = \min_p \{H(p) - a \cdot p\}. \tag{13}$$

(Comments on the latter formula: (a) h is concave, as it should be, since it is a minimum of linear functions; (b) $-h = \max_p \{a \cdot p - H(p)\}$ is the ‘‘Fenchel transform’’ of H). To justify this assertion we must show that when H is convex and f, h are given by (13), the resulting Hamiltonian

$$\max_{a \in A} \{f(a) \cdot p + h(a)\}$$

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 (depending on a). We can rewrite this 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 preceding calculation 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

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

(We can effectively restrict the maximization to $z \in A$ by choosing $F(z) = \infty$ when $z \notin A$.) 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 (13) leads to a HJB with Hamiltonian H is just an application of the second bullet.

Summarizing the above: if H is convex, we can define a “dynamic programming” solution of the HJB equation $u_t + H(\nabla u) = 0$ (with $u = g$ at $t = T$) as the solution of the finite-horizon dynamic programming problem associated with (13). This optimal control problem is easy to solve more or less explicitly. The key observation is that when $f(a) = a$ and $h(a)$ is concave, the optimal control is constant and the associated trajectory is a constant-velocity path. In fact, $f(a) = a$ means the control is the *velocity* of the path $y(s)$. The concavity of h gives

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

Notice moreover that the average velocity of a path depends only on its endpoints, since

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

Thus replacing any path by one with the same endpoints and constant velocity can only improve the utility. We thus arrive at the **Hopf-Lax solution formula**: when $f(a) = a$ and $h(a)$ is concave, the solution of the associated dynamic programming problem (with final-time utility g) is

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

Here z represents the state at time T – the only remaining unknown – and $(z-x)/(T-t)$ is the velocity of the associated path starting at x at time t and ending at z at time T . We view (14) as specifying the (dynamic programming) solution of the associated HJB equation $u_t + H(\nabla u) = 0$ with final-time condition $u = g$ at $t = T$.

Let’s bring this down to earth by considering the specific example: $f(a) = a$, $H(p) = \frac{1}{2}|p|^2$, and $h(p) = -\frac{1}{2}|a|^2$. Then the Hamilton-Jacobi-Bellman equation is

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

and the solution formula is

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

An important fact is immediately evident: the Hamilton-Jacobi-Bellman equation has many (almost-everywhere) solutions, only one of which agrees with the solution formula. For example, suppose $g = 0$. Then the solution formula gives $u(x, t) = 0$, which does solve the

Hamilton-Jacobi equation. However the PDE has lots of other solutions: for example the function

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

This example is easy to generalize, yielding infinitely many “solutions” of the PDE, all equal to 0 at $t = T$. 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 a geometrical minimum-time problem.

Another point to note: the Hamilton-Jacobi equation is *nonlinear*. If u_1 solves it with final data g_1 and u_2 solves it with final data g_2 , we should not expect $u_1 + u_2$ to solve it with final data $g_1 + g_2$. When $H(p) = H(-p)$, for example when $H(p) = |p|^2/2$, one might imagine that if u is the correct solution with final data g then $-u$ is the correct solution with final data $-g$. But even this is false, at least if we understand “correct solution” as “dynamic programming solution” in the sense developed above. Indeed, when $H(p) = |p|^2/2$ and $g(y) = |y|$ the solution formula gives

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

(the optimal z is $z = x + (T - t)x/|x|$), but when the final-time data is changed to $g(y) = -|y|$ the solution formula gives

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

(the optimal z is $z = x - (T - t)x/|x|$ if $|x| \geq (T - t)$, $z = 0$ otherwise).

We turn now to a discussion of **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. (Example: Leslie Hocking, *Optimal Control: An Introduction to the Theory with Applications*, Oxford University Press 1991 – a really good book, with lots of examples.)

From a strictly PDE viewpoint, the Pontryagin Maximum Principle is closely related to the method of characteristics. Roughly speaking: the principle identifies the characteristics of the HJB equation (and a bit more: it identifies the characteristic that’s directly relevant to the specific optimization problem under consideration). In short: it solves the HJB equation by the method of characteristics.

From the optimization viewpoint, the Pontryagin Maximum Principle provides optimality conditions – in much the same way that the convex dual gave us a substitute for the Euler-Lagrange equations in connection with convex variational problems. [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. That part of the theory is closely connected with *relaxation*, which we'll discuss (for nonconvex variational problems) soon. But it lies beyond the scope of these notes.]

The following (formal) derivation relies 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 same principle we used to understand linear programming (see Peter Lax's Linear Algebra book for a similar but fully rigorous analysis) and convex duality (see the book by Ekeland and Temam for rigorous analysis in that setting). A treatment of the Pontryagin Maximum Principle pretty close to the one given here is given by Dixit's book.

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

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

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:

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

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

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

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

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

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

$$u(x, t) = \min_{\pi(s)} \max_{y(t)=x} \left\{ \int_t^T \left[y \cdot \frac{d\pi}{ds} + H(\pi, y) \right] ds + g(y(T)) + x \cdot \pi(t) - y(T) \cdot \pi(T) \right\} \quad (15)$$

with

$$H(p, x) = \max_{a \in A} [p \cdot f(x, a) + h(x, a)]. \quad (16)$$

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 (15) 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. (I leave the justification as an exercise: it suffices to check that the Pontryagin maximum principle, stated below, identifies the characteristics of the HJB equation.)

The formula (16) has two important consequences. The first is

$$\frac{\partial H}{\partial p} = f(x, a_*) \quad \text{where } a_* = a_*(p, x) \text{ is the optimal } a \text{ in the definition of } H(p, x). \quad (17)$$

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

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

But the second term is zero, because a_* is optimal, and this shows (17). 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

$$dy/ds = \nabla_p H(\pi(s), y(s)). \quad (18)$$

This is an immediate consequence of (17), using the equation of state $dy/ds = f$.

We return now to the min-max representation of u , equation (15). 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

$$d\pi/ds = -\frac{\partial H}{\partial y}. \quad (19)$$

A similar argument at the final time gives

$$\pi(T) = \nabla g(y(T)) \quad \text{at the final time } T. \quad (20)$$

Collecting the information in (17)–(20), 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

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

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 (17)) 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 at the end of these notes, 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 .

Suggested exercises. More than usual in fact (don't feel you must do them all).

1. Consider the finite-horizon utility maximization problem with discount rate ρ . The dynamical law is thus

$$dy/ds = f(y(s), \alpha(s)), \quad y(t) = x,$$

and the optimal utility discounted to time 0 is

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

It is often more convenient to consider, instead of u , the optimal utility discounted to time t ; this is

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

- (a) Show (by a heuristic argument similar to those in the notes) that v satisfies

$$v_t - \rho v + H(x, \nabla v) = 0$$

with Hamiltonian

$$H(x, p) = \max_{a \in A} \{ f(x, a) \cdot p + h(x, a) \}$$

and final-time data

$$v(x, T) = g(x).$$

(Notice that the PDE for v is autonomous, i.e. there is no explicit dependence on time.)

- (b) Now consider the analogous infinite-horizon problem, with the same equation of state, and value function

$$\bar{v}(x, t) = \max_{\alpha \in A} \int_t^{\infty} e^{-\rho(s-t)} h(y(s), \alpha(s)) ds.$$

Show (by an elementary comparison argument) that \bar{v} is independent of t , i.e. $\bar{v} = \bar{v}(x)$ is a function of x alone. Conclude using part (a) that if \bar{v} is finite, it solves the stationary PDE

$$-\rho\bar{v} + H(x, \nabla\bar{v}) = 0.$$

2. Recall Example 1 of the notes: the state equation is $dy/ds = ry - \alpha$ with $y(t) = x$, and the value function is

$$u(x, t) = \max_{\alpha \geq 0} \int_t^{\tau} e^{-\rho s} h(\alpha(s)) ds$$

with $h(a) = a^\gamma$ for some $0 < \gamma < 1$, and

$$\tau = \begin{cases} \text{first time when } y = 0 \text{ if this occurs before time } T \\ T \text{ otherwise.} \end{cases}$$

- (a) We obtained a formula for $u(x, t)$ in the notes, however our formula doesn't make sense when $\rho - r\gamma = 0$. Find the correct formula in that case.
- (b) Let's examine the infinite-horizon-limit $T \rightarrow \infty$. Following the lead of Problem 1 let us concentrate on $v(x, t) = e^{\rho t} u(x, t) =$ optimal utility discounted to time t . Show that

$$\bar{v}(x) = \lim_{T \rightarrow \infty} v(x, t) = \begin{cases} G_\infty x^\gamma & \text{if } \rho - r\gamma > 0 \\ \infty & \text{if } \rho - r\gamma \leq 0 \end{cases}$$

with $G_\infty = [(1 - \gamma)/(\rho - r\gamma)]^{1-\gamma}$.

- (c) Use the stationary PDE of Problem 1(b) (specialized to this example) to obtain the same result.
- (d) What is the optimal consumption strategy, for the infinite-horizon version of this problem?
3. Consider the analogue of Example 1 with the power-law utility replaced by the logarithm: $h(a) = \ln a$. To avoid confusion let us write u_γ for the value function obtained in the notes using $h(a) = a^\gamma$, and u_{\log} for the value function obtained using $h(a) = \ln a$. Recall that $u_\gamma(x, t) = g_\gamma(t)x^\gamma$ with

$$g_\gamma(t) = e^{-\rho t} \left[\frac{1 - \gamma}{\rho - r\gamma} \left(1 - e^{-\frac{(\rho - r\gamma)(T-t)}{1-\gamma}} \right) \right]^{1-\gamma}.$$

- (a) Show, by a direct comparison argument, that

$$u_{\log}(\lambda x, t) = u_{\log}(x, t) + \frac{1}{\rho} e^{-\rho t} (1 - e^{-\rho(T-t)}) \ln \lambda$$

for any $\lambda > 0$. Use this to conclude that

$$u_{\log}(x, t) = g_0(t) \ln x + g_1(t)$$

where $g_0(t) = \frac{1}{\rho} e^{-\rho t} (1 - e^{-\rho(T-t)})$ and g_1 is an as-yet unspecified function of t alone.

- (b) Pursue the following scheme for finding g_1 : Consider the utility $h = \frac{1}{\gamma}(a^\gamma - 1)$. Express its value function u_h in terms of u_γ . Now take the limit $\gamma \rightarrow 0$. Show this gives a result of the expected form, with

$$g_0(t) = g_\gamma(t)|_{\gamma=0}$$

and

$$g_1(t) = \left. \frac{dg_\gamma}{d\gamma}(t) \right|_{\gamma=0}.$$

(This leads to an explicit formula for g_1 but it's messy; I'm not asking you to write it down.)

- (c) Indicate how g_0 and g_1 could alternatively have been found by solving appropriate PDE's. (Hint: find the HJB equation associated with $h(a) = \ln a$, and show that the ansatz $u_{\log} = g_0(t) \ln x + g_1(t)$ leads to differential equations that determine g_0 and g_1 .)
4. Our Example 1 considers an investor who receives interest (at constant rate r) but no wages. Let's consider what happens if the investor also receives wages at constant rate w . The equation of state becomes

$$dy/ds = ry + w - \alpha \quad \text{with } y(t) = x,$$

and the value function is

$$u(x, t) = \max_{\alpha \geq 0} \int_t^T e^{-\rho s} h(\alpha(s)) ds$$

with $h(a) = a^\gamma$ for some $0 < \gamma < 1$. Since the investor earns wages, we now permit $y(s) < 0$, however we insist that the final-time wealth be nonnegative ($y(T) \geq 0$).

- (a) Which pairs (x, t) are acceptable? The strategy that maximizes $y(T)$ is clearly to consume nothing ($\alpha(s) = 0$ for all $t < s < T$). Show this results in $y(T) \geq 0$ exactly if

$$x + \phi(t)w \geq 0$$

where

$$\phi(t) = \frac{1}{r} (1 - e^{-r(T-t)}).$$

Notice for future reference that ϕ solves $\phi' - r\phi + 1 = 0$ with $\phi(T) = 0$.

- (b) Find the HJB equation that $u(x, t)$ should satisfy in its natural domain $\{(x, t) : x + \phi(t)w \geq 0\}$. Specify the boundary conditions when $t = T$ and where $x + \phi w = 0$.
- (c) Substitute into this HJB equation the ansatz

$$v(x, t) = e^{-\rho t} G(t) (x + \phi(t)w)^\gamma.$$

Show v is a solution when G solves the familiar equation

$$G_t + (r\gamma - \rho)G + (1 - \gamma)G^{\gamma/(\gamma-1)} = 0$$

(the same equation we solved in Example 1). Deduce a formula for v .

- (d) In view of (a), a more careful definition of the value function for this control problem is

$$u(x, t) = \max_{\alpha \geq 0} \int_t^\tau e^{-\rho s} h(\alpha(s)) ds$$

where

$$\tau = \begin{cases} \text{first time when } y(s) + \phi(s)w = 0 \text{ if this occurs before time } T \\ T \text{ otherwise.} \end{cases}$$

Use a verification argument to prove that the function v obtained in (c) is indeed the value function u defined this way.

5. (An example of nonexistence of an optimal control.) Consider the following control problem: the state is $y(s) \in R$ with $y(t) = x$; the control is $\alpha(s) \in R$; the dynamics is $dy/dt = \alpha$; and the goal is

$$\text{minimize } \int_t^T y^2(s) + (\alpha^2(s) - 1)^2.$$

The value function $u(x, t)$ is the value of this minimum.

- (a) Show that when $x = 0$ and $t < T$, the value is $u(0, t) = 0$.
- (b) Show that when $x = 0$ and $t < T$ there is no optimal control $\alpha(s)$.

[The focus on $x = 0$ is only because this case is most transparent; nonexistence occurs for other (x, t) as well. Food for thought: What is the Hamilton-Jacobi-Bellman equation? Is there a modified goal leading to the same Hamiltonian and value function, but for which optimal controls exist?]

6. The Hopf-Lax solution formula solves the finite-horizon problem with state equation $dy/ds = \alpha$ and value function

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

with h concave. The key step was to show that an optimal trajectory has constant velocity. Give an alternative justification of this fact using Pontryagin's maximum principle.

7. Early in these notes we discussed a “minimum travel time” problem whose value function u solves $|\nabla u| = 1$ for $x \notin E$ with $u = 0$ at ∂E .
- Find a related dynamic programming problem whose value function (if smooth) should solve $|\nabla u| = 1$ for $x \notin E$ with $u = g$ at ∂E , where g is a specified function on ∂E .
 - Consider the 2D case, with E a planar region with smooth boundary ∂E . Describe the optimal controls and paths, if g is smooth and its derivative (with respect to arc-length) on E satisfies $|g'| < 1$.
 - What changes if $|g'|$ is bigger than 1 on some part of ∂E ?
8. This problem is a special case of the “linear-quadratic regulator” widely used in engineering applications. The state is $y(s) \in R^n$, and the control is $\alpha(s) \in R^n$. There is no pointwise restriction on the values of $\alpha(s)$. The evolution law is

$$dy/ds = Ay + \alpha, \quad y(t) = x,$$

for some constant matrix A , and the goal is to minimize

$$\int_t^T |y(s)|^2 + |\alpha(s)|^2 ds + |y(T)|^2.$$

(In words: we prefer $y = 0$ along the trajectory and at the final time, but we also prefer not to use too much control.)

- What is the associated Hamilton-Jacobi-Bellman equation? Explain why we should expect the relation $\alpha(s) = -\frac{1}{2}\nabla u(y(s))$ to hold along optimal trajectories.
- Since the problem is quadratic, it’s natural to guess that the value function $u(x, t)$ takes the form

$$u(x, t) = \langle K(t)x, x \rangle$$

for some symmetric $n \times n$ matrix-valued function $K(t)$. Show that this u solves the Hamilton-Jacobi-Bellman equation exactly if

$$\frac{dK}{dt} = K^2 - I - (K^T A + A^T K) \text{ for } t < T, \quad K(T) = I$$

where I is the $n \times n$ identity matrix. (Hint: two quadratic forms agree exactly if the associated symmetric matrices agree.)

- Show by a suitable verification argument that this u is indeed the value function of the control problem.