PDE for Finance Notes – 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 1999.

Viscosity solutions. This section discusses "viscosity solutions" of Hamilton-Jacobi-Bellman equations. We mainly follow chapter 10 of Evans. Of course we barely scratch the surface. Good places to learn more are (a) Evans; (b) the book *Viscosity solutions and optimal control*, R. J. Elliott, Pitman Research Notes in Mathematics 165, 1987 (not on reserve); and the chapter by Crandall in *Viscosity solutions and applications* (on reserve). Sources (b) and (c) are strictly more difficult than Evans, and are recommended only for people with substantial background in PDE.

As announced in class last week, Mr. Xiang is now holding office hours on Mondays (4-5) as well as Thursdays (5-7). Moreover he has arranged to use Room 613 when a classroom setting is convenient. So check both 613 and his office (606) when looking for him during office hours.

Some students found the discussion of verification theorems in Section 1 of the notes too brief. Don't be confused by all the inequalities. The main point is that if u is smooth and u satisfies the HJB equation (including appropriate boundary or final-time conditions) then – by an elementary argument – u bounds the payoff obtained using any control. Examining the argument that proves this statement, you always find that you don't really need u to be a *solution* of the HJB equation – what you need is some inequality. This fact is useful because the argument requires u to be smooth (i.e. continuously differentiable), whereas there may be no smooth *solution* to the HJB equation. Note that the Section 1 Addendum includes an example of a verification theorem (complete with proof). It appears on pages 3-4; start reading at the heading "The Hamilton-Jacobi-Bellman equation". By the way, the Green box contains a copy of the article by Shreve, Soner, and Xu cited in those notes.

****

Let's review where we stand. We introduced the method of dynamic programming as a scheme for approaching certain optimal control problems. We explained that the value function solves a first-order differential equation, the Hamilton-Jacobi-Bellman equation associated to the control problem. Then we observed that merely "solving" the HJB equation in a naive sense doesn't describe the value function uniquely. We nevertheless made good use of the HJB equation – or a related inequality – in our discussion of verification theorems. These ideas date mainly from the late 1950's and 1960's.

The theory of viscosity solutions is newer: it dates from the early 1980's. Key points:

(a) The value function of a dynamic programming problem is automatically a viscosity solution;

(b) Viscosity solutions are unique.

In short, when understood this way the Hamilton-Jacobi-Bellman equation provides a complete characterization of the value function. Beyond its elegance, this theory provides a framework for designing good numerical schemes – including methods that do not have a direct dynamic-programming interpretation. (References concerning numerical schemes: P. Souganides, "Approximation schemes for viscosity solutions of Hamilton-Jacobi equations," J. Diff. Eqns 59, 1985, 1-43; also P.-L. Lions & P. Souganides, "Convergence of MUSCL and filtered schemes for scalar conservation laws and Hamilton-Jacobi equations", Numer. Math. 69, 1995, 441-470.)

The viscosity solution of an initial value problem. Most of the mathematical literature (including Evans) focuses on initial value problems. So let's start there, though in the end we're more interested in final value problems. If our goal is to solve

$$u_t + H(Du, x) = 0 \text{ for } t > 0, \; u(x, 0) = g(x),$$

a plausible approach is to solve the "regularized" equation

$$u^\epsilon_t + H(Du^\epsilon, x) - \epsilon\Delta u^\epsilon = 0 \text{ for } t > 0, \; u^\epsilon(x, 0) = g(x),$$

for $\epsilon > 0$, then pass to the limit $\epsilon \to 0$. This is the right idea – it gives the solution we want – but it is not a convenient *definition* of the viscosity solution. Instead, the definition captures a property of $u = \lim_{\epsilon\to 0} u^\epsilon$ that can be stated without making reference to u^ϵ. This property is the following:

- If $\phi(x, t)$ is smooth and $u - \phi$ has a local maximum at (x_0, t_0) then $\phi_t(x_0, t_0) + H(D\phi(x_0, t_0), x_0) \le 0$.
- If $\phi(x, t)$ is smooth and $u - \phi$ has a local minimum at (x_0, t_0) then $\phi_t(x_0, t_0) + H(D\phi(x_0, t_0), x_0) \ge 0$.

A *viscosity solution of the initial value problem* is a continuous function $u(x, t)$ satisfying these conditions and $u = g$ at $t = 0$.

To explain the definition, we now show that $u = \lim_{\epsilon\to 0}$ has this property. We'll assume (1) u^ϵ is differentiable when $\epsilon > 0$, and (2) $u^\epsilon - \phi$ has a local maximum or minimum at some point (x_ϵ, t_ϵ) near (x_0, t_0). (See Evans, pp. 541-2, for justification of the second assumption.) For the first bullet, suppose $u - \phi$ has a local maximum at (x_ϵ, t_ϵ), and observe from calculus that $D(u^\epsilon - \phi)(x_\epsilon, t_\epsilon) = 0$ and $\Delta(u^\epsilon - \phi)(x_\epsilon, t_\epsilon) \le 0$. Therefore at (x_ϵ, t_ϵ) we have

$$\begin{aligned} 0 &= u^\epsilon_t + H(Du^\epsilon, x_\epsilon) - \epsilon\Delta u^\epsilon \\ &= \phi_t + H(D\phi, x_\epsilon) - \epsilon\Delta u^\epsilon \\ &\ge \phi_t + H(D\phi, x_\epsilon) - \epsilon\Delta\phi. \end{aligned}$$

Passing to the limit $\epsilon \to 0$ gives $0 \ge \phi_t(x_0, t_0) + H(D\phi(x_0, t_0), x_0)$, as expected. For the second bullet the calculation is the same, except that $u^\epsilon - \phi$ has a local minimum, so $\Delta(u^\epsilon - \phi) \ge 0$ and the inequality is reversed.

Why the name "viscosity solution"? This comes from a link to conservation laws, and where viscosity plays a regularizing role similar to our $\epsilon\Delta u^\epsilon$. The link to conservation laws is easy: in one space dimension, the equation $u_t + H(u_x) = 0$ can be rewritten as $v_t + (H(v))_x = 0$ with $v = u_x$. This is a "scalar conservation law." The associated regularized equation is $v_t^\epsilon + (H(v^\epsilon))_x = \epsilon v_{xx}^\epsilon$. To see why the parameter ϵ is called "viscosity," compare this with the slightly more complicated (but still simple) conservation law of 1D elasticity:

$$\begin{aligned} e_t &= v_x \\ v_t &= (f(e))_x + \epsilon v_{xx} \end{aligned}$$

in which $e(x,t)$ is strain, $v(x,t)$ is velocity, and $f(e)$ is the inviscid stress, and ϵv_x is the viscous stress. The first equation expresses the fact that $e = u_x$ and $v = u_t$ where $u(x,t)$ is displacement; the second equation expresses Newton's law in the form $u_{tt} = (\text{total stress})_x$.

The viscosity solution of a final value problem. In dynamic programming the value function satisfies a final-value problem not an initial-value problem. The appropriate regularization is different:

$$u_t^\epsilon + H(Du^\epsilon, x) + \epsilon\Delta u^\epsilon = 0 \text{ for } t < T, \ u^\epsilon(x,T) = g(x),$$

with $\epsilon > 0$. (We'll explain why later this semester; briefly the linear heat equation $u_t - \Delta u = 0$ is well-behaved for initial-value problems but badly-behaved for final-value problems; $u_t + \Delta u = 0$ is equally well-behaved for final-value problems but badly-behaved for initial-value problems.) Arguing as above, but remembering the modified regularization, we find that $u = \lim_{\epsilon\to 0} u^\epsilon$ satisfies

- If $\phi(x,t)$ is smooth and $u - \phi$ has a local maximum at (x_0, t_0) then $\phi_t(x_0,t_0) + H(D\phi(x_0,t_0), x_0) \geq 0$.
- If $\phi(x,t)$ is smooth and $u - \phi$ has a local minimum at (x_0, t_0) then $\phi_t(x_0,t_0) + H(D\phi(x_0,t_0), x_0) \leq 0$.

A *viscosity solution of the final value problem* is a continuous function $u(x,t)$ satisfying these conditions and $u = g$ at $t = T$.

In the dynamic programming setting the regularization has a very natural interpretation that has nothing to do with conservation laws or physical viscosity. Adding $\epsilon\Delta u$ to the equation corresponds to introducing a bit of randomness in the dynamics (making the state equation a stochastic differential equation rather than an ODE). We'll be explaining this later in the semester.

Section 1 considered, as an example, the final value problem

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

We saw that the associated dynamic programming problem had value function $u(x,t) \equiv 0$ but there were lots of other "solutions" such as

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

Let's verify that this v is *not* a viscosity solution. In fact, consider $\phi(x,t) = \frac{1}{2}(T-t)$. Then $v \le \phi$, and $v - \phi$ has a local maximum at $(0,t)$ for any $t < T$. If v were a viscosity solution we would have $\phi_t + \frac{1}{2}|\phi_x|^2 \ge 0$ at $x = 0$. This is false, since $\phi_t = -1$ and $\phi_x = 0$.

The viscosity solution of a time-independent equation. For the equation $H(Du, x) = 0$ we get two notions of viscosity solution, depending on whether we perturb by $\epsilon\Delta u$ or $-\epsilon\Delta u$. They give different solutions! The Eikonal equation $|\nabla u| = 1$ provides an interesting example. If we perturb it as

$$|\nabla u^\epsilon| - 1 - \epsilon\Delta u^\epsilon = 0$$

we get

$$u - \phi \text{ has a local max} \Rightarrow |\nabla\phi| - 1 \le 0,$$
$$u - \phi \text{ has a local min} \Rightarrow |\nabla\phi| - 1 \ge 0.$$

If instead we perturb it as

$$|\nabla u^\epsilon| - 1 + \epsilon\Delta u^\epsilon = 0$$

we get

$$u - \phi \text{ has a local max} \Rightarrow |\nabla\phi| - 1 \ge 0,$$
$$u - \phi \text{ has a local min} \Rightarrow |\nabla\phi| - 1 \le 0.$$

Suppose we are interested in the minimum time problem, with target $E = \{|x| \ge 1\}$. Then the value function is defined in the unit ball, and $u(x) = \text{dist}(x, E) = 1 - |x|$. In which sense (if either) is this a viscosity solution? Answer: it is a viscosity solution in the first sense, but not the second one.

To see u is not a viscosity solution in the second sense, we focus on $x = 0$ (where u is singular), and consider $\phi \equiv 1$. Clearly $u - \phi$ has a local max; the second notion of viscosity solution would require $|\nabla\phi| \ge 1$ whereas actually $\nabla\phi = 0$.

To see that u is a viscosity solution in the first sense is slightly harder, because we must prove a statement about all ϕ and all x_0 rather than simply giving a counterexample. Observe that we need only focus on $x_0 = 0$. (Explanation: if $u - \phi$ has a local max or min at $x_0 \ne 0$ then since u is smooth near x_0 we have $\nabla\phi(x_0) = \nabla u(x_0) = 1$). We may also assume $\phi(0) = u(0)$. (Explanation: if not, replace ϕ by $\phi + \text{constant}$.) There does not exist a smooth ϕ such that $\phi(0) = u(0)$ and $\phi(x) \le u(x)$ near $x = 0$. There are plenty of smooth ϕ such that $\phi(0) = u(0)$ and $\phi \ge u$ near $x = 0$ – but one easily verifies, using Taylor's theorem, that any such ϕ satisfies $|\nabla\phi(0)| \le 1$. Thus u passes the test.

A similar argument shows that $|x| - 1$ is a viscosity solution in the second sense. We see here an important feature of viscosity solutions: one of the two definitions permits convex singularities but not concave ones, the other vice versa.

The value function is automatically a viscosity solution of the HJB equation. Let us explain this in the context of the finite-horizon utility maximization problem discussed in Section 1:

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

where the control $\alpha(s) \in A$, and the state satisfies

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

The associated HJB equation is

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

with

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

where H is defined by

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

We showed (heuristically) in Section 1 that *if* u is smooth then it solves the HJB equation. We show now (essentially honestly, without assuming differentiability) that u is a viscosity solution of the HJB equation:

(a) If $\phi(x, t)$ is smooth and $u - \phi$ has a local maximum at (x_0, t_0) then $\phi_t(x_0, t_0) + H(\nabla\phi(x_0, t_0), x_0) \geq 0$.

(b) If $\phi(x, t)$ is smooth and $u - \phi$ has a local minimum at (x_0, t_0) then $\phi_t(x_0, t_0) + H(\nabla\phi(x_0, t_0), x_0) \leq 0$.

We assume in the following that $f(x, a)$, $h(x, a)$, $g(x)$, and $u(x, t)$ are continuous functions. (See Evans for conditions on f, h, g that ensure continuity of u.)

Consider (b) first. Suppose $u - \phi$ has a local minimum at (x_0, t_0). We wish to show that $\phi_t + H(\nabla\phi, x) \leq 0$ at (x_0, t_0). So let's assume the opposite and seek a contradiction:

$$\phi_t + \max_a\{f(x, a) \cdot \nabla\phi + h(x, a)\} \geq \delta > 0 \tag{1}$$

at (x_0, t_0). If the maximum is achieved at a_0, then this assumption is

$$\phi_t(x_0, t_0) + f(x_0, a_0) \cdot \nabla\phi(x_0, t_0) + h(x_0, a_0) \geq \delta > 0.$$

By continuity a similar relation holds for all (x, t) near (x_0, t_0):

$$\phi_t(x, t) + f(x, a_0) \cdot \nabla\phi(x, t) + h(x, a_0) \geq \delta/2 \tag{2}$$

for (x, t) in some neighborhood $\mathcal{N}$ of (x_0, t_0). We may suppose, choosing $\mathcal{N}$ smaller if necessary, that $u - \phi$ has its minimum in $\mathcal{N}$ at (x_0, t_0).

Consider the following strategy, starting from x_0 at time t_0: set $\alpha(s) = a_0$ for an initial short time interval $t_0 < s < t_1$, then proceed optimally thereafter. The optimal strategy starting from x_0 at time t_0 can do no worse, so

$$u(x_0, t_0) \geq \int_{t_0}^{t_1} h(y(s), a_0)\, ds + u(y(t_1), t_1) \tag{3}$$

where $y(s)$ is the associated solution of the state equation. By taking t_1 close to t_0 we may arrange that $(y(s), s)$ stays in the region $\mathcal{N}$ where (2) holds for all s, $t_0 \leq s \leq t_1$.

Equation (3) gives a relation between $u(x_0, t_0)$ and $u(y(t_1), t_1)$. But we also have another relation between the same quantities: since $u - \phi$ is minimized in $\mathcal{N}$ at (x_0, t_0),

$$u(x_0, t_0) - \phi(x_0, t_0) \leq u(y(t_1), t_1) - \phi(y(t_1), t_1),$$

or in other words

$$u(x_0, t_0) \leq \phi(x_0, t_0) - \phi(y(t_1), t_1) + u(y(t_1), t_1). \tag{4}$$

The ϕ terms on the right can be estimated by the same sort of argument used in a verification theorem:

$$\begin{aligned} \phi(x_0, t_0) - \phi(y(t_1), t_1) &= -\int_{t_0}^{t_1} \frac{d}{ds}\phi(y(s), s)\, ds \\ &= \int_{t_0}^{t_1} -\phi_s(y(s), s) - \nabla\phi(y(s), s) \cdot f(y(s), a_0) \\ &\leq \int_{t_0}^{t_1} h(y(s), a_0)\, ds - \frac{1}{2}\delta(t_1 - t_0). \end{aligned}$$

Combining this with (4) gives

$$u(x_0, t_0) \leq \int_{t_0}^{t_1} h(y(s), a_0)\, ds + u(y(t_1), t_1) - \frac{1}{2}\delta(t_1 - t_0).$$

This contradicts (3), since $\delta(t_1 - t_0) > 0$. So our assumption (1) was wrong. Thus u satisfies condition (b) in the definition of a viscosity solution.

Now consider the other condition (a). Suppose $u - \phi$ has a local maximum at (x_0, t_0). We must show that $\phi_t + H(\nabla\phi, x) \geq 0$ at (x_0, t_0). So let's assume the opposite and seek a contradiction:

$$\phi_t + \max_a \{f(x, a) \cdot \nabla\phi + h(x, a)\} \leq -\delta < 0 \tag{5}$$

at (x_0, t_0). This means that

$$\phi_t(x_0, t_0) + f(x_0, a) \cdot \nabla\phi(x_0, t_0) + h(x_0, a) \leq -\delta$$

for all $a \in A$. By continuity a similar relation holds for all (x, t) near (x_0, t_0):

$$\phi_t(x, t) + f(x, a) \cdot \nabla\phi(x, t) + h(x, a) \leq -\delta/2 \tag{6}$$

for all $a \in A$ and all (x, t) in a neighborhood $\mathcal{N}$ of (x_0, t_0). We may suppose, choosing $\mathcal{N}$ smaller if necessary, that $u - \phi$ has its maximum in $\mathcal{N}$ at (x_0, t_0).

Consider the optimal strategy $\alpha(s)$ and the associated state $y(s)$. Choose t_1 so that $y(s)$ stays in $\mathcal{N}$ for $t_0 \leq s \leq t_1$. We have

$$u(x_0, t_0) = \int_{t_0}^{t_1} h(y(s), \alpha(s))\, ds + u(y(t_1), t_1) \tag{7}$$

from the principle of dynamic programming. This is a relation between $u(x_0, t_0)$ and $u(y(t_1), t_1)$. We also have another relation between these quantities: since $u - \phi$ is maximized in $\mathcal{N}$ at (x_0, t_0),

$$u(x_0, t_0) - \phi(x_0, t_0) \geq u(y(t_1), t_1) - \phi(y(t_1), t_1),$$

or in other words

$$u(x_0, t_0) \geq \phi(x_0, t_0) - \phi(y(t_1), t_1) + u(y(t_1), t_1). \tag{8}$$

The ϕ terms on the right can be estimated much as before:

$$\begin{aligned}
\phi(x_0, t_0) - \phi(y(t_1), t_1) &= -\int_{t_0}^{t_1} \frac{d}{ds}\phi(y(s), s)\, ds \\
&= \int_{t_0}^{t_1} -\phi_s(y(s), s) - \nabla\phi(y(s), s) \cdot f(y(s), \alpha(s)) \\
&\geq \int_{t_0}^{t_1} h(y(s), \alpha(s))\, ds + \frac{1}{2}\delta(t_1 - t_0).
\end{aligned}$$

Combining this with (8) gives

$$u(x_0, t_0) \geq \int_{t_0}^{t_1} h(y(s), a_0)\, ds + u(y(t_1), t_1) + \frac{1}{2}\delta(t_1 - t_0).$$

This contradicts (7). So our assumption (5) was wrong. Thus u satisfies condition (a) in the definition of a viscosity solution.

(We cheated slightly in the treatment of (a), when we assumed the existence of an optimal control. But once you know the structure of the argument it is easy to eliminate this assumption. We didn't really need that $\alpha(s)$ was optimal; all we needed was that it gave a value within $\frac{1}{2}\delta(t_1 - t_0)$ of optimal.)

Viscosity solutions are unique. This theorem is more difficult, but the essential idea is simple and elegant. The following discussion aims to capture this idea without getting lost in technicalities or details. See Evans for a rigorous treatment, with carefully stated hypotheses. (Evans considers the initial-value problem, here we address the final-value problem, but the differences are trivial.)

Our goal is to show that the problem

$$u_t + H(\nabla u, x) = 0 \text{ for } t < T,\ x \in R^n,\ u(x, 0) = g(x)$$

has at most one bounded, Lipschitz continuous viscosity solution. It suffices to show that if u and $\tilde{u}$ are two different solutions then $u \leq \tilde{u}$. (Then reversing the roles of u and $\tilde{u}$ we also have $\tilde{u} \leq u$ so in fact $u = \tilde{u}$. We argue by contradiction: let us assume that for some x, t,

$$u(x, t) - \tilde{u}(x, t)) \geq \delta > 0 \tag{9}$$

and seek a contradiction.

Obviously we must use the definition of a viscosity solution. But how to choose ϕ and (x_0, t_0)? Here's the answer: consider the function

$$\Sigma(x, y, t, s) = u(x, t) - \tilde{u}(y, s) + \lambda(t + s) - \frac{1}{\epsilon^2}(|x - y|^2 + |t - s|^2) - \epsilon(|x|^2 + |y|^2) \qquad (10)$$

with λ and ϵ positive but small, and let (x_0, y_0, t_0, s_0) be the point where it achieves its maximum over all $x, y \in R^n$ and $t, s < T$. Then

- $u(x, t) - \phi(x, t)$ is maximized at (x_0, t_0) if we choose

$$\phi(x, t) = \tilde{u}(y_0, s_0) - \lambda(t + s_0) + \frac{1}{\epsilon^2}(|x - y_0|^2 + |t - s_0|^2) + \epsilon(|x|^2 + |y_0|^2)$$

  since $u(x, t) - \phi(x, t) = \Sigma(x, y_0, t, s_0)$; also

- $\tilde{u}(y, s) - \tilde{\phi}(y, s)$ is minimized at (y_0, s_0) if we choose

$$\tilde{\phi}(y, s) = u(x_0, t_0) + \lambda(t_0 + s) - \frac{1}{\epsilon^2}(|x_0 - y|^2 + |t_0 - s|^2) - \epsilon(|x_0|^2 + |y|^2)$$

  since $\tilde{u}(y, s) - \tilde{\phi}(y, s) = -\Sigma(x_0, y, t_0, s)$.

We thus have two opportunities to use the definition of a viscosity solution: one for u and one for $\tilde{u}$. The first gives

$$\phi_t + H(\nabla\phi, x) \geq 0 \quad \text{at } x_0, t_0$$

and the second gives

$$\tilde{\phi}_t + H(\nabla\tilde{\phi}, y) \leq 0 \quad \text{at } y_0, s_0.$$

Doing the differentiations, these become:

$$-\lambda + \frac{2(t_0 - s_0)}{\epsilon^2} + H(\frac{2}{\epsilon^2}(x_0 - y_0) + 2\epsilon x_0, x_0) \geq 0 \qquad (11)$$

and

$$\lambda + \frac{2(t_0 - s_0)}{\epsilon^2} + H(\frac{2}{\epsilon^2}(x_0 - y_0) - 2\epsilon y_0, y_0) \leq 0. \qquad (12)$$

We have not yet used our assumption (9), nor the smallness of λ and ϵ. Consider the formula for Σ. The term involving λ and the term $\epsilon(|x|^2 + |y|^2)$ are present to be sure Σ achieves its maximum at a finite point (x_0, y_0, t_0, s_0), but they are relatively unimportant in understanding the argument. The important terms are the others,

$$u(x, t) - \tilde{u}(y, s) - \frac{1}{\epsilon^2}(|x - y|^2 + |t - s|^2),$$

so for a heuristic understanding let's suppose (x_0, y_0, t_0, s_0) maximize this expression instead. We claim that under hypothesis (9) and assuming $u, \tilde{u}$ are Lipschitz continuous:

(a) the maximal value of this expression is at least δ;

(b) The optimal x_0, y_0, t_0, s_0 satisfy $|x_0 - y_0| \leq C\epsilon^2$, $|t_0 - s_0| \leq C\epsilon^2$.

(c) The optimal t_0 and s_0 are not at T.

Point (a) is easy: we can achieve value δ by setting $y = x$ and $s = t$ optimizing over x, t. Point (b) isn't so hard either: the optimal choices must do better than than $(x, y, t, s) = (x_0, x_0, t_0, t_0)$,

$$u(x_0, t_0) - \tilde{u}(y_0, s_0) - \frac{1}{\epsilon^2}(|x_0 - y_0|^2 + |t_0 - s_0|^2) \geq u(x_0, t_0) - \tilde{u}(x_0, t_0)$$

which gives

$$\frac{1}{\epsilon^2}(|x_0 - y_0|^2 + |t_0 - s_0|^2) \leq \tilde{u}(x_0, t_0) - \tilde{u}(y_0, s_0) \leq C(|x_0 - y_0| + |t_0 - s_0|)$$

using the Lipschitz continuity of $\tilde{u}$ in the last step. This gives (b). Point (c) follows easily from (a), (b), the fact that both solutions have the same final data, and the Lipschitz continuity of u and $\tilde{u}$.

OK, now we're ready to finish. Of course to make some combined use of (11) and (12) we must make some regularity assumption on H. A convenient hypothesis is that $H(p, x)$ is Lipschitz continuous in either variable with the other held fixed:

$$|H(p, x) - H(q, x)| \leq C|p - q| \quad \text{and} \quad |H(p, x) - H(p, y)| \leq C|x - y|.$$

Subtracting (12) from (11) gives

$$-2\lambda + H(\frac{2}{\epsilon^2}(x_0 - y_0) + 2\epsilon x_0, x_0) - H(\frac{2}{\epsilon^2}(x_0 - y_0) - 2\epsilon y_0, y_0) \geq 0,$$

which implies under our hypotheses on H

$$-2\lambda \leq -C\epsilon(|x_0| + |y_0|) - C|x_0 - y_0|.$$

To finish, we hold λ fixed and positive, and send $\epsilon \to 0$. (We suppressed the role of λ above; a more careful analysis shows that while λ must be small, it doesn't have to tend to 0 with ϵ.) The right hand side tends to 0 (the second term tends to 0 by point (b); the first one tends to 0 since $\epsilon(|x_0|^2 + |y_0|^2) \leq \max \Sigma \leq \max u + \max \tilde{u}$). In the limit we get $-2\lambda \leq 0$, which is the desired contradiction.

Notice that in the above, we used just *half* the definition of a viscosity solution for u, and the other half for $\tilde{u}$. Naming these halves, we've actually shown that if u is a viscosity *subsolution* and v is a viscosity *supersolution* then $u \leq v$. This suggests the key idea in the PDE proof that viscosity solutions exist: one shows that the largest subsolution is a viscosity solution. (So is the smallest supersolution, and by uniqueness they are the same.)

Comments. The importance of the theory of viscosity solutions goes far beyond dynamic programming. It applies equally well to Hamiltonians that aren't convex in ∇u (which have interpretations involving two-player games, see Bardi) and to second-order equations (which have interpretations involving stochastic control, as we'll see soon). For problems in bounded domains, the handling of boundary conditions is subtle (you've seen an example of such subtlety in Homework 1, the example involving $|\nabla u| = 1$ for $x \notin E$, $u = g$ at ∂E.)