Calculus of Variations, Lectures 5+6, 2/26/2013 + 3/5/2013

New topic : optimal control + HJ eqns

We're still focusing on 1D var'l pbms, but the viewpt is a bit different from Lecture 4, leading to additional types of applns (eg economics + finance), and deep linkages to pde (through the "value function", Hamilton-Jacobi eqns, and the theory of viscosity solutions of HJ eqns)

There are lots of books on optimal control, eg

- LM Hocking, Optimal Control: An Introduction to the Theory with Applications, Oxford Univ Press 1991 (lots of examples, not much HJ)
- A K Dixit, Optimization in Economic Theory, Oxford Univ Press 1990 (chaps 10+11, emphasizes economic applns)
- L.C. Evans' book on pde, chap 10 (explains why value fn is a viscosity soln of the HJ eqn)

See also my Spr 2011 PDE for Finance notes, Section 4 and HW 4 (which are relatively close to the treatment here).

Typical examples of opt'l control problems

A) $\min \int_0^T h(y(s), \alpha(s))\,ds + g(y(T))$

where the maximization is over "controls" $\alpha(s)$, which determine evolution of the "state" via an ODE

$$\dot{y}(s) = f(y(s), \alpha(s)), \quad y(0) = y_0$$

Typical engineering appl'n: send a spacecraft to the moon. Then $\vec{y}$ = (position, velocity), $\vec{\alpha}(s)$ controls firing of rockets, $\dot{y} = f(y, \alpha)$ is eqn of Newtonian mechanics, T = desired arrival time (treated here as fixed), $g(y(T))$ favors desired arrival location with velocity near 0, and h = fuel consumption.

Typical economics appl'n: max rather than min; $\alpha(s)$ controls investment policy and/or consumption of resources; h + g are utilities assoc to consumption + final-time wealth

Link to problems considered recently: var'l pbm

$$\min_{u(0)=0} \int_0^T (u_x^2-1)^2 + u^2 \, dx$$

can easily be put in this form:

$$\min \int_0^T (\alpha^2-1)^2 + y^2 \, dt$$

where $\alpha(t) \in \mathbb{R}$ is the control
$dy/dt = \alpha(t)$, $y(0)=0$ is the "state eqn"
$h(y,\alpha) = (\alpha^2-1)^2 + y^2$ is the "running cost"
$g = 0$ (no final-time cost)

As we've noticed before, the nonconvexity in α leads to nonexistence of a minimizer, though the min value is perfectly well-defined (it is 0 in this case).

B) min arrival time

min { time at which $y(s)$ reaches some target set Γ }

where $y(s)$ solves an ode
$\dot{y} = f(y(s), \alpha(s))$, $y(0) = y_0$
and the optimization is over the "control" $\alpha(s)$.

This is nearly a special case of (A) (setting $T=\infty$ and $h = \begin{cases} 1 & \text{if target has not been reached} \\ 0 & \text{at target} \end{cases}$)

but it's different enough to deserve special treatment.

Special case that's easy to visualize: given $D \subset \mathbb{R}^n$ and $x \in D$, consider

$$\min \{ \text{time to exit } D, \text{ starting from } x \text{ and travelling at velocity} \le 1 \}.$$

Evidently state eqn is $\dot{y} = \alpha$, where $\alpha(s) \in \mathbb{R}^n$ must satisfy $|\alpha| \le 1$. Opt'l path is of course straight line to nearest pt on ∂D, &

$$\text{min value} = \text{dist}(x, \partial D)$$

(well-defined, though opt'l path may not be unique).

Our goals are

(1) a scheme for guessing the form of the solution, and proving that the guess

is correct

(2) a far-from-obvious link between optimal control (which intrinsically involves problems in *one* variable, typically "time") and pde (namely Hamilton-Jacobi eqns).

The two goals will be achieved together (They are interdependent).

For problem class A) the trick is to study the dependence of the optimal *value* on the initial *position* and *time*; so define

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

where $\dot{y}(s) = f(s,\alpha(s))$ for $t<s<T$

with $y(t) = x$

We'll derive a pde for u. The main tool is the *dynamic programming principle*

$$u(x,t) = \min_{\substack{\alpha(s)\in A \\ t<s<t'}} \left\{ \int_t^{t'} h(y(s),\alpha(s))\,ds + u(y_{x,t,\alpha}(t'),t') \right\}$$

Interpretation: the optimal strategy must do something between t + t'; starting from t', it should solve the same problem with a new starting time + position.

We can derive a pde for u (formally, ie assuming more differentiability than might really be true) by applying this with $t' = t + \Delta t$ in the limit $\Delta t \to 0$. Here is the formal argument:

- let's guess that over $t < s < t + \Delta t$ the optimal $\alpha(s)$ is (more or less) constant. Then

$$u(x,t) \approx \min_{a \in A} \left\{ h(x,a)\Delta t + u(x + f(t,a)\Delta t,\ t + \Delta t) \right\}$$

dropping corrections of order $(\Delta t)^2$

- let's assume u is differentiable, and expand via Taylor series

$$\cancel{u(x,t)} \simeq \min_{a \in A} \left\{ h(x,a)\Delta t + \cancel{u(x,t)} + \nabla u \cdot f(t,a)\,\Delta t + u_t \Delta t \right\}$$

- Cancelling the Δt terms, we get

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

ie a pde of the form $u_t + H(t,x,\nabla u) = 0$. It is to be solved for $t<T$, $x \in \mathbb{R}^n$, with final-time data

$$u(x,T) = g(x)$$

(since if starting time $= T$ then $u(x,t) = g(x)$ from the very defn of u). Note that the H we get this way is concave in ∇u (being the min of linear fns of ∇u).

If we had started with a max problem instead of a min problem, same calcn would have given

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

ie eqn of form $u_t + H(t,x,\nabla u) = 0$ with H concave in ∇u.

Example: The Hopf-Lax formula for $u_t + H(\nabla u) = 0$ with $H(\vec{p})$ convex. We have only to write

$$H(\vec{p}) = \max_{\vec{a}} \langle \vec{a}, \vec{p} \rangle + h(\vec{a})$$

(evidently, $h = -H^*$ where H^* is the Fenchel transform we defined a couple of lectures ago):

to guess that the relevant soln of the pde with $u = g$ at $t = T$ is

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

using eqn of state

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

Fact: given any choice of $y(t)$, the best path is the one with constant α. This follows from concavity of h, and Jensen's ineq

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

Since avg velocity depends only on endpts, ie

$$\frac{1}{T-t}\int_t^T \frac{dy}{ds}\, ds = \frac{1}{T-t}\left[y(T) - y(t)\right]$$

we arrive at the "Hopf-Lax solution formula"

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

which reduces solving the pde $u_t + H(\nabla u) = 0$ $(t < T)$ with $u = g$ at $t = T$ to a 1-D optimization at each time t + spatial pt x.

For problem class B, ie min arrival time problems, the situation is similar: if

$$u(x) = \min_{\alpha(s)} \{ \text{time when } y(s) \text{ reaches target } \Gamma \}$$

using eqn of state

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

[note that starting time is now fixed!] then dyn progr prin says

$$u(x) = \min_{\substack{\alpha(s) \in A \\ 0 < s < t'}} \{ u(y_{\alpha,x}(t')) + t' \}$$

Arguing as before (taking $t' = \Delta t \to 0$ and using Taylor expansion) we get

$$u(x) \approx \min_{a \in A} \{ u(x + f(x,a)\Delta t) + \Delta t \}$$

$$\simeq \min_{a \in A} \{ u(x) + (f(x,a) \cdot \nabla u)\Delta t + \Delta t \}.$$

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

an eqn of form $H(\nabla u) + 1 = 0$ (with H concave), to be solved for $x \notin \Gamma$, with bc $u = 0$ at Γ.

Example: in the special case where state eqn is

$$\dot{y} = \alpha(s)$$

and speed is $|\alpha(s)| \le 1$ we know optimal path is straight + goes toward pt of Γ closest to x, so that

$$u(x) = dist(x, \Gamma).$$

Assoc HJ eqn is

$$\min_{|a|\le 1} \{ \nabla u \cdot a \} + 1 = 0$$

ie

$$-|\nabla u| + 1 = 0$$

ie the eikonal eqn

$$|\nabla u| = 1 \quad \text{off } \Gamma$$
$$u = 0 \quad \text{at } \Gamma$$

Our discn thus far has ignored some very important issues:

① The soln $u(x,t)$ we want may not be

differentiable (calling into question the derivation of the pde). For example, the eikonal eqn

$$|\nabla u| = 1 \quad \text{in } D \subset \mathbb{R}^n$$
$$u = 0 \quad \text{at } \partial D$$

has <u>no</u> C^1 soln

(2) The pde may have many ae solns; for example

$$|u_x| = 1 \text{ on } [-1,1], \quad u=0 \text{ at endpts}$$

has lots of solns 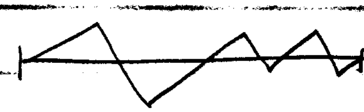

How to know which one we want?

(3) Our real goal was to <u>solve</u> problems; can the pde be used for this (either by hand or numerically)?

Sketch of answers:

To (1) and (2): there's a notion of a <u>viscosity</u>

solution of the pde. Viscosity solns are unique, and they are the special (ae) soln of the pde that gives the value fn. (This is explained in Evans' chap 10)

To (3): The derivation of the pde gives us a pretty good idea how the control should be related to ∇u (it should achieve the optzn that determined $H(x, \nabla u)$).

Once we have a conjectured solution, we can often harness the argument that derived the pde to prove that it's optimal (using what is sometimes called a "verification argument").

Before starting (3), let's work an example: find

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

where $0<q<1$, $\rho>0$, and where the state eqn is

$$\frac{dy}{dt} = ry - a \quad , \quad y(0) = x$$

and the control + state must satisfy $a(s) \geq 0$, $y(s) \geq 0$. (Here $r>0$ is a constant interest rate.)

<u>Interpretation</u>: an investor has initial wealth x at time t, and plans his consumption $a(s)$ to maximize his discounted "utility" up to a fixed final time T. (We use the power law utility a^γ, $0<\gamma<1$, because it makes the HJB solvable by separation of variables.)

<u>Step 0</u>: Let's show the value fn must have the form

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

for some $g(t)$. Suff't to show

$$(*) \qquad u(\lambda x, t) = \lambda^\gamma u(x,t)$$

for all $\lambda > 0$ (Then $g(t) = u(1,t)$). To see $(*)$, consider control $\lambda a(s)$ for problem starting from λx, where $a(s)$ is opt'l choice starting from x. Assoc soln of state eqn is $y_{\lambda x}(s) = \lambda y_x(s)$. Using form of utility we conclude that

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

Same reln with λ replaced by λ^{-1} gives

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

Together, these give (*).

<u>Step 1</u>: Find HJB eqn. Almost a special case of calcn done before – except now we have a discount term $e^{-\rho(s-t)}$. Arguing as before: formally

$$u(x,t) \simeq \max_{\alpha\geq 0}\left\{\alpha^{\gamma}\Delta t + e^{-\rho\Delta t}u(x+(rx-\alpha)\Delta t, t+\Delta t)\right\}$$

$$\simeq \max_{\alpha\geq 0}\left\{\alpha^{\gamma}\Delta t + (1-\rho\Delta t)\left(u(x,t)+u_t\Delta t + (rx-\alpha)u_x\Delta t\right)\right\}$$

$$\simeq u(x,t) + \Delta t \max_{\alpha\geq 0}\left\{\alpha^{\gamma} - \rho u + u_t + (rx-\alpha)u_x\right\}$$

so as $\Delta t \to 0$ we get

$$(**) \qquad u_t + \max_{\alpha\geq 0}\left\{\alpha^{\gamma} + (rx-\alpha)u_x\right\} - \rho u = 0$$

<u>Step 2</u> Optimal consumption policy is easy to find. Clearly $u_x > 0$ (clear from step 0), so opt'l α is positive, namely

$$\alpha = \left(\frac{1}{\gamma}u_x\right)^{\frac{1}{\gamma-1}}$$

Recalling that $u = g(t)\, x^{\gamma}$ we get

$$\alpha(t) = g(t)^{\frac{1}{\gamma-1}} \cdot x$$

To find $g(t)$ we substitute this into the pde + do some arithmetic:

$$g_t x^{\gamma} - \rho g x^{\gamma} + \left(g^{\frac{\gamma}{\gamma-1}}(1-\gamma) + r\gamma\, g \right) x^{\gamma} = 0$$

ie

$$\frac{dg}{dt} + (r\gamma - \rho)\, g(t) + (1-\gamma)\, g(t)^{\frac{\gamma}{\gamma-1}} = 0 .$$

Multiplying by $(1-\gamma)^{-1} g^{\frac{\gamma}{1-\gamma}}$, we find that $H(t) = g(t)^{\frac{1}{1-\gamma}}$ satisfies the linear eqn

$$(***) \qquad H_t - \mu H + 1 = 0 \qquad \text{with } \mu = \frac{\rho - r\gamma}{1-\gamma}$$

Step 3 We forgot to note the final-time condn, which in this pblm is $u(x,T) = 0$ (since there is no final-time term in the optimization).

Due to its simple form, we can easily solve (***) with $H = 0$ at $t = T$. Soln is

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

This determines H, whence $g(t) = H^{1-\gamma}$ and $u = g(t)x^{\gamma}$.

But: our derivation of the HJB eqn was formal. So, is the soln just found really the optimal value, i.e. is it really

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

Answer is yes, by the following verification argument. Let $u(x,t)$ be the optimal value, and $\tilde{u}(x,t)$ = conjectured optimal value assoc our explicit soln. Then

(A) $u(x,t) \geq \tilde{u}(x,t)$ because $\tilde{u}(x,t)$ is the value assoc with a particular consumption plan, namely the one we found in step 2. (This can be checked directly, but we'll also see why it's true in part B.)

(B) to show $u(x,t) \leq \tilde{u}(x,t)$ let's calculate

$$\frac{d}{dt} \tilde{u}(y_{x,a}(t), t) \quad \text{where } y_{x,a} \text{ solves state}$$

eqn for any (fixed) policy $a(t)$. We get

$$\frac{d}{dt}\tilde{u}(y(t),t) = \tilde{u}_t + \nabla\tilde{u}\cdot(ry-a).$$

$$\leq \rho\tilde{u} - a^{\gamma}$$

using that $\tilde{u}$ solves the HJB eqn (**) in the last step. (Note that this calcn has = in the last step when $a(t)$ is the optimal policy found in step 2.)

So

$$\frac{d}{dt}\, e^{-\rho t}\tilde{u}(y(t),t) \leq -e^{-\rho t}a^{\gamma}(t)$$

Integrate + use that $\tilde{u}(y(T),T)=0$ to get

$$-e^{-\rho t}\tilde{u}(x,t) \leq -\int_t^T e^{-\rho s}a^{\gamma}(s)\,ds.$$

$$\Rightarrow \quad \tilde{u}(x,t) \geq \int_t^T e^{-\rho(s-t)}a^{\gamma}(s)\,ds$$

Maximizing RHS over all choices of $a(s)$ we see that

$$\tilde{u}(x,t) \geq u(x,t)$$

as desired.

Similar arg't shows more generally that if soln of HJB eqn is C^1 then it is indeed the optimal control.

Alas, soln of HJB eqn is often not C^1.
Simple example: recall from pg 10 that "min time pblm" assoc "travel, starting from x, with max speed 1, until you arrive in a target set Γ" has eikonal eqn

$$|\nabla u| = 1 \text{ off } \Gamma, \qquad u = 0 \text{ at } \Gamma$$

as its HJB eqn. This has many ae solns, but none that are smooth, if eg $\Gamma = \partial D$.

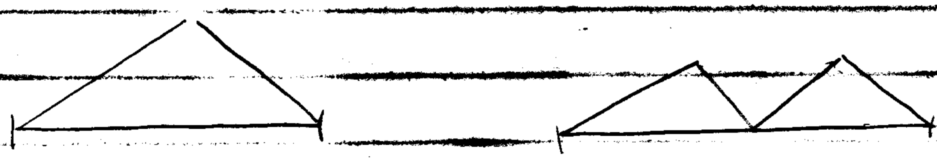

$D = [0,1]$ $\qquad\qquad$ $D = [0,1]$.

two ae solns of $|u_x| = 1$ in D, $u = 0$ at ∂D

Can verification argt still be used? Yes as follows. The goal would be to give a verifn-style pf that

$$\tilde{u}(x) = \text{dist}(x, \partial D)$$

equals

$$u(x) = \min_{|\alpha| \le 1} \{ \text{arrival time to } \partial D \}$$

where the state eqn is $\dot{y}(t) = \alpha(t)$, $y(0) = x$.

Obviously $\tilde{u}(x) \geq u(x)$ since we know how to achieve $\tilde{u}$ (namely: travel at const speed 1 toward nearest bdry pt).

To see $\tilde{u}(x) \leq u(x)$ we first argue formally as before (pretending $\tilde{u}$ is diff'ble): for any $\alpha(t)$ st $|\alpha(t)| \leq 1$,

$$\frac{d}{dt}\tilde{u}(y(t)) = \nabla\tilde{u} \cdot \frac{dy}{dt} = \nabla\tilde{u} \cdot \alpha(t)$$

$$\geq \min_{|\alpha(t)| \leq 1} \nabla\tilde{u} \cdot \alpha(t) = -1.$$

whence if arrival occurs at time τ then

$$\cancel{\tilde{u}(y(\tau))}^{\,0} - \tilde{u}(y(0)) \geq \int_0^\tau -1$$

$$\tilde{u}(x) \leq \tau.$$

Optimizing over all $\alpha(t) \Rightarrow \tilde{u}(x) \leq u(x)$.

To make this honest, observe that we pretended $\tilde{u}$ was smooth (which isn't true) but we only used the HJB eqn as an <u>inequality</u>

$$\min_{|\alpha| \leq 1} \nabla\tilde{u} \cdot \alpha = -|\nabla\tilde{u}| \underset{\substack{\uparrow \\ \text{all we used.}}}{\geq} -1$$

So our arg't shows (quite honestly) that

if w is C^1, $w=0$ at ∂D, and $|\nabla w| \le 1$
then $w(x) \le u(x)$.

Apply this not to $\tilde{u} = \text{dist}(x, \partial D)$ but rather to a "smoothed out" approx'n

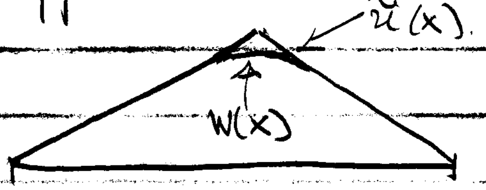

As $w \uparrow \tilde{u}$ we conclude (honestly) that $\tilde{u}(x) \le u(x)$, as desired.

Generalization of this: assertion that a singular sol'n of HJB eq'n is the opt'l value can often be achieved by verif'n arg't applied to a smooth approx'n.

I have barely mentioned viscosity sol'ns of HJ eq'ns. It would take us too far afield, and Evans' treatment is excellent. But briefly: situation is a bit like study of shock waves (eg Burgers' eq'n).

a) Though HJB eq'n may have no smooth sol'n + ~~many~~ a.e. sol'ns, there's a special one (called the "viscosity soln," though artificial viscosity is not the most convenient analytical tool here).

b) value fn of an opt'l control pbm is always the viscosity soln! (Thus: no need for any verifn arg't, if we can manage to find the viscosity soln.)

Suggested exercises

(1)(a) In our example involving optimal consumption (pp 5.12 - 5.15) we got an explicit soln of the HJB eqn, but the formula doesn't make sense if $\rho - r\gamma = 0$. What is the solution in that case?

(b) Show that $u_\infty(x)$ = limit of $u(x,t)$ as the final-time horizon $T \to \infty$ is

$$u_\infty(x) = \begin{cases} G_\infty x^\gamma & \text{if } \rho - r\gamma > 0 \\ \infty & \text{if } \rho - r\gamma < 0 \end{cases}$$

(c) What is the optimal consumption strategy, in the limit $T \to \infty$?

(2) Consider the analogous example when the goal is to maximize

$$\int_t^T e^{-\rho(s-t)} \ln a(s)\, ds$$

and let $u(x,t)$ be the associated value function.

a) Show that for any $\lambda > 0$,
$$u(\lambda x, t) = u(x,t) + \frac{1}{\rho} \ln \lambda \cdot \left(1 - e^{-\rho(T-t)}\right).$$

b) Using (a), conclude that
$$u(x,t) = g_0(t) \ln x + g_1(t)$$
for some functions g_0 and g_1.

c) What ODE's and final-time conditions should g_0 + g_1 solve? (The ODE's can be solved explicitly; g_0 is pretty simple but g_1 is a little messier.)

(3) We discussed a "minimum travel time" problem whose value function u solves $|\nabla u| = 1$ in D and $u = 0$ at ∂D

a) Find a related problem whose value function solves $|\nabla u| = 1$ in D and $u = g$ at ∂D, where g is a specified function.

b) Consider the 2D case, ie let D be a domain in $\mathbb{R}^2$ (assume ∂D is smooth). Describe the optimal controls + paths, if g is smooth + its derivative w.r.t. to arc length has $|g'|<1$.

c) What changes if $|g'|>1$ on some part of ∂D?

(4) This problem is a special case of the "linear-quadratic regulator" widely used in engineering applications. The state is $y(s) \in \mathbb{R}^n$ + the control is $\alpha(s) \in \mathbb{R}^n$ (with no ptwise restriction). The state eqn is

$$\frac{dy}{ds} = Ay + \alpha \quad , \quad y(t) = x$$

where A is a given (constant) matrix. The goal is to find

$$u(x,t) = \min_{\alpha(s)} \int_t^T |y(s)|^2 + |\alpha(s)|^2 \, ds \; + \; |y(T)|^2$$

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

a) Find the HJB eqn. Explain why we should expect the relation $\alpha(s) = -\frac{1}{2}\nabla u(y(s))$ to hold along optimal trajectories.

b) Since the problem is quadratic, it is natural to guess that

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

where $K(t)$ is a symmetric-matrix-valued function. Show that u solves the HJB eqn iff

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

with $K(T) = I$ (the $n\times n$ identity matrix). [Hint: two quadratic forms agree exactly if the assoc symmetric matrices agree.]

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

(5) We showed (pp 5.18-5.20) how a verification argt can be used to show that $\tilde{u}(x) = \mathrm{dist}(x, \partial D)$ is the value function of a simple "min travel time" optimal control problem.

In 1D with $D = [-1,1]$ and $\bar{u}(x)$ as shown

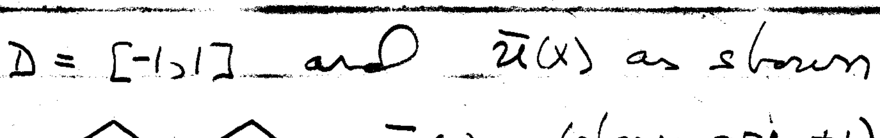

we could try to use a similar argument to show that $\bar{u}(x)$ is the value function of this problem.

Of course we must fail (since $\bar{u}(x) \neq \text{dist}(x, \partial D)$) even though $|\bar{u}_x| = 1$ a.e. What goes wrong?