Stochastic optimal control

Recap

Itô process Let $b : [0,\infty) \times \Omega \to \mathbb{R}^n$ be measurable and adapted [i.e. $b(t,\cdot)$ is $\mathcal{F}_t$ mble $\forall t \geq 0$]
Let $\sigma : [0,\infty) \times \Omega \to \mathbb{R}^{n\times m}$ be progressively measurable [i.e. $\sigma|_{[0,t]\times\Omega}$ is $\mathcal{B}[0,t] \times \mathcal{F}_t$ mble $\forall t \geq 0$.

Suppose $\mathbb{P}\left(\int_0^t \sum_{i,j} \sigma_{ij}^2(s) + |b_i(s)|\, ds < \infty \;\; \forall t \geq 0\right) = 1$

Let B be a m-dimensional Brownian motion and ~~X_0 be a $\mathcal{F}_0^B$ mble r.v., then~~ $x \in \mathbb{R}^n$

$$X_t = \cancel{X_0}\, x + \int_0^t b(s)\, ds + \int_0^t \sigma(s)\, dB_s$$

is called Itô process.

Itô diffusion Let $b : \mathbb{R}^n \to \mathbb{R}^n$, $\sigma : \mathbb{R}^n \to \mathbb{R}^{n\times m}$ such that

$$\|b(x) - b(y)\| + \left(\sum_{i,j} |\sigma_{ij}(x) - \sigma_{ij}(y)|^2\right)^{1/2} \leq D\, \|x-y\|$$

$\forall x, y \in \mathbb{R}^n$ and some $D > 0$. Let $s \geq 0$, $x \in \mathbb{R}^n$. Then

$$X_t = x + \int_s^t b(X_r)\, dr + \int_s^t \sigma(X_r)\, dB_r, \quad t \geq s,$$

is called Itô diffusion.

A solution to this equation is unique up to indistinguishability, i.e. two solutions X and $\hat{X}$ satisfy
$$\mathbb{P}_{x,s}(X_t = \hat{X}_t \quad \forall t \in [s,\infty)) = 1.$$
Moreover, X is adapted to $(\mathcal{F}_t^B)_{t\geq s}$.

The uniqueness implies two important properties of Itô diffusions

1. Time homogeneous
$$\begin{aligned} X_{s+t} &= x + \int_s^{s+t} b(X_r)\, dr + \int_s^{s+t} \sigma(X_r)\, dB_r \\ &= x + \int_0^t b(X_{r+s})\, dr + \int_0^t \sigma(X_{s+r})\, d(B_{r+s} - B_s) \end{aligned}$$
$$\underset{\text{uniqueness}}{\Longrightarrow} (X_{s+t} : t \geq 0 \,;\, \mathbb{P}_{x,s}) = (X_t : t \geq 0 \,;\, \mathbb{P}_{x,0})$$

2. Markov property
Let $\Phi : \mathbb{R}^n \to \mathbb{R}$ be mble and bounded, τ a $(\mathcal{F}_t^B)_t$ stopping time, $s \geq 0$, $x \in \mathbb{R}^n$ and $\mathbb{P}_{x,s}(\tau < \infty) = 1$.

Then $\mathbb{E}_{x,s}[\Phi(X_{\tau+h}) \mid \mathcal{F}_\tau^B] = \mathbb{E}_{X_\tau,0}[\Phi(X_h)] \quad \forall h \geq 0$
$\mathbb{P}_{x,s}$ a.s.

Proof idea in case $\tau = t \geq 0$ constant

$$X_{t+h} = X_t + \int_t^{t+h} b(X_r)\,dr + \int_t^{t+h} \sigma(X_r)\,d(B_r - B_t)$$

The solution to this eq is unique and mble w.r.t
$\sigma(B_r - B_t\ ;\ r \geq t)$ what is independent of $\mathcal{F}_t^B$

$$\Rightarrow \mathbb{E}_{x,s}[\Phi(X_{t+h}) \mid \mathcal{F}_t^B] = \mathbb{E}_{X_t,t}[\Phi(X_{t+h})]$$

$$\underset{\text{homogeneous}}{=} \mathbb{E}_{X_t,0}[\Phi(X_h)]$$

We denote $\forall x \in \mathbb{R}^n,\ \Phi \in C^2(\mathbb{R}^n)$

$$L\Phi(x) = \sum_{i=1}^n b_i(x) \frac{\partial \Phi}{\partial x_i}(x) + \sum_{i,j=1}^n \frac{1}{2}\left(\sigma\sigma^T(x)\right)_{ij} \frac{\partial^2 \Phi}{\partial x_i \partial x_j}(x)$$

the generalized generator of X.

By Itô's formula,

$$\mathbb{E}_{x,0}[\Phi(X_\tau)] = \Phi(x) + \mathbb{E}_{x,0}\left[\int_0^\tau L\Phi(X_s)\,ds\right] \quad \boxed{\text{Dynkins formula}}$$

for all $\Phi \in C^2_b$ and $\mathcal{F}_t^B$-stopping times τ for which

$$\mathbb{E}_{x,0}\left[\int_0^\tau \sigma_{ij}(X_t) \frac{\partial \Phi}{\partial x_i}(X_t)\,dB_j\right] = 0 \quad \forall i,j.$$

Examples
~~τ is the exit time~~
$\Phi \in C_0^2(\mathbb{R}^2)$, $\mathbb{E}_{x,0}[\tau] < \infty$ (see last week)
$\Phi \in C^2(\mathbb{R}^2)$ and τ exit time of a bounded domain $\mathbb{E}_{x,0}[\tau] < \infty$.
the continuity of σ_{ij} and $\frac{\partial \Phi}{\partial x_i}$ and the dominated convergence theorem.

Problem statement - Stochastic Control

Let $U \subseteq \mathbb{R}^k$ be a measurable set
$b: \mathbb{R} \times \mathbb{R}^n \times U \to \mathbb{R}^n$, $\sigma: \mathbb{R} \times \mathbb{R}^n \times U \to \mathbb{R}^{n \times m}$ measurable functions and B a m-dimensional Brownian motion.

Suppose there is an U-valued, $(\mathcal{F}_t^B)_t$ adapted process $(u_t)_{t \geq 0}$ and an $\mathbb{R}^n$-valued Itô process such that

$$X_t = x + \int_s^t b(r, X_r, u_r)\,dr + \int_s^t \sigma(r, X_r, u_r)\,dB_r \quad \forall t \geq s$$

There are more assumptions needed to make the terms well-defined / guarantee existence but we will state the problem under the simple assumption that the process exists.

Let $f: \mathbb{R}\times\mathbb{R}^n\times U \to \mathbb{R}$, $g: \mathbb{R}\times\mathbb{R}^n \to \mathbb{R}$ be continuous functions and $G \subseteq \mathbb{R}\times\mathbb{R}^n$ be a ~~bounded~~ domain.

$\hat{J}_G = \inf\{r > s \quad (r, X_r) \notin G\}$ the first exit time from G.

Suppose

(1.4) $\mathbb{E}_{x,s}\left[\int_s^{\hat{J}_G} |f(r, X_r, u_r)|\,dr + |g(\hat{J}_G, X_{\hat{J}_G})|\, 1_{\hat{J}_G < \infty}\right] < \infty$ $\forall x \in \mathbb{R}^n, s \geq 0$.

We define

control

$J^u(x,s) = \mathbb{E}_{x,s}\left[\int_s^{\hat{J}_G} f(r, X_r, u_r)\,dr + g(\hat{J}_G, X_{\hat{J}_G})\, 1_{\hat{J}_G < \infty}\right]$

performance fctn — profit rate function — bequest function = inheritance

i.e. accrued interest + final returns if we sell

Let $\mathcal{A}$ be a family of allowed controls, ie

- X exists
- u adapted and U valued
- (1.4) holds

$\Big\}\ \forall u \in \mathcal{A}$.

For all $y \in G$ we are looking for $u^* = u^{*y} \in \mathcal{A}$ such that

$$\Phi(y) := \sup_{u \in \mathcal{A}} J^u(y) = J^{u^*}(y)$$

optimal control.

admissible controls

optimal performance function or value function.

Before we look at ways to solve these problems, we try to "remove" the time dependence from coefficients b and σ to get closer to the setup of Itô diffusions where we have many results already.

~~Write $Y_t = \begin{pmatrix} s+t \\ X_{s+t}\end{pmatrix} = \begin{pmatrix} s \\ X_{s+t}\end{pmatrix}$~~

Note that

$$Y_t := \begin{pmatrix} s+t \\ X_{s+t}\end{pmatrix} = \begin{pmatrix} s \\ x\end{pmatrix} + \int_s^{s+t} \begin{pmatrix} 1 \\ b(r, X_r, u_r)\end{pmatrix} dr + \int_s^{s+t} \begin{pmatrix} 0 \cdots 0 \\ \sigma(r, X_r, u_r)\end{pmatrix} dB_r$$

$$= \begin{pmatrix} s \\ x\end{pmatrix} + \int_0^t \begin{pmatrix} 1 \\ b(Y_r, \tilde{u}_r)\end{pmatrix} dr + \int_0^t \begin{pmatrix} 0 \cdots 0 \\ \sigma(Y_r, \tilde{u}_r)\end{pmatrix} d(B_{r+s} - B_s)$$

with $\tilde{u}_r = u_{r+s}$

④

$\Rightarrow$ Y solves an equation with a different BM but non-time dependent coefficients

For the optimal control problem we obtain

$$\hat{\tau}_G - s = \inf\{t > s \; (t, X_t) \notin G\} - s$$
$$= \inf\{t > 0 \; Y_t \notin G\} =: \tau_G$$

$$\Rightarrow J^u(x,s) = \mathbb{E}_{x,s}\left[\int_s^{\hat{\tau}_G} f(r, X_r, u_r)\,dr + g(\hat{\tau}_G, X_{\hat{\tau}_G})\,1_{\hat{\tau}_G < \infty}\right]$$

$$= \mathbb{E}_{x,s}\left[\int_0^{\hat{\tau}_G - s} f(Y_r, u_{r+s})\,dr + g(\hat{\tau}_G - s + s, X_{\hat{\tau}_G - s + s})\,1_{\hat{\tau}_G < \infty}\right]$$

$$\underset{y=\binom{s}{x}}{=} \mathbb{E}_y\left[\int_0^{\tau_G} f(Y_r, \tilde{u}_r)\,dr + g(Y_{\tau_G})\,1_{\tau_G < \infty}\right]$$

Markov controls: Todays lecture is mainly looking at Markov controls.

Suppose $u: \mathbb{R}^{n+1} \to U$ is a function with $u_t = u(t, X_t)$. Then

$$\tilde{u}_t = u_{s+t} = u(Y_t).$$

u is called Markov control if $y \mapsto b(y, u(y))$ and $y \mapsto \sigma(y, u(y))$ are Lipschitz continuous on $\mathbb{R}^{n+1}$ and (1.4) holds. The generalized generator is [use $b_1 = 1$, $\sigma_{1j} = 0$]

$$L^u \phi(y) = \frac{\partial \phi}{\partial s}(y) + \sum_{i=1}^n b_i(y, u(y)) \frac{\partial \phi}{\partial x_i} + \sum_{i,j=1}^n a_{ij}(y, u(y)) \frac{\partial^2 \phi}{\partial x_i \partial x_j}(y) \quad \forall \phi \in C^2, \quad a = \tfrac{1}{2}\sigma\sigma^T$$

Thm The Hamilton-Jacobi-Bellman (HJB) equation (I)

Let $\mathcal{A} = \{u : u \text{ Markov control}\}$.

Suppose $\Phi \in C^2(G) \cap C(\bar{G})$ satisfies

$$\mathbb{E}_y\left[|\Phi(Y_\alpha)| + \int_0^\alpha |L^v \Phi(Y_t)|\,dt\right] < \infty$$

Interpret v in L^v as a constant function

for all bounded stopping times $\alpha \le \tau_G$, $y \in G$, $v \in U$.

Moreover, suppose that an optimal Markov control u^* exists and that ∂G is regular for $(Y_t^{(u^*)})_{t \ge 0}$.

Then
$$\sup_{v \in U}\left\{ \cancel{f^v(y)}\; f(y,v) + L^v \Phi(y)\right\} = 0 \quad \forall y \in G$$

$$\Phi(y) = g(y) \quad \forall y \in \partial G$$

and
$$f(y, u^*(y)) + (L^{u^*}\Phi)(y) = 0 \quad \forall y \in G.$$

Note: Regular means

$$\mathbb{P}_y^{u^*}[\tau_G = 0] = 1 \quad \forall \; y \in \partial G.$$

Proof

Let $\alpha \le \mathcal{T}_G$ be a bounded stopping time and $u \in \mathcal{A}$ a Markov control.
Then

$$\mathbb{E}_y[\, J^u(Y_\alpha)] \underset{\text{Def}}{=} \mathbb{E}_y\Big[\, \mathbb{E}_{Y_\alpha}\Big[\int_0^{\mathcal{T}_G} f(Y_r, u(Y_r))\,dr + g(Y_{\mathcal{T}_G})\,1_{\mathcal{T}_G<\infty}\Big]\Big]$$

$$\underset{\substack{\text{Strong}\\\text{Markov}}}{=} \mathbb{E}_y\Big[\, \mathbb{E}_y\Big[\int_0^{\mathcal{T}_G-\alpha} f(Y_{r+\alpha}, u(Y_{r+\alpha}))\,dr + g(Y_{\mathcal{T}_G})\,1_{\mathcal{T}_G<\infty}\Big]\,\mathcal{F}_\alpha\Big]$$

time shift by $\alpha \Rightarrow \mathcal{T}_G \rightsquigarrow \mathcal{T}_G - \alpha$, $Y_r \rightsquigarrow Y_{r+\alpha}$

$$\underset{\text{Tower}}{=} \mathbb{E}_y\Big[\int_0^{\mathcal{T}_G} f(Y_r, u(Y_r))\,dr + g(Y_{\mathcal{T}_G})\,1_{\mathcal{T}_G<\infty} - \int_0^{\alpha} f(Y_r, u(Y_r))\,dr\Big]$$

$$= J^u(y) - \mathbb{E}_y\Big[\int_0^{\alpha} f(Y_r, u(Y_r))\,dr\Big]$$

$$\Rightarrow J^u(y) = \mathbb{E}_y\Big[\int_0^{\alpha} f(Y_r, u(Y_r))\,dr\Big] + \mathbb{E}_y[\, J^u(Y_\alpha)] \quad \triangleleft \quad (2.6)$$

<u>Dynamic programming principle</u>

Note that we needed only that u is a Markov control and none of the other assumptions.

Let $W = \{(r,z) \in G : r < t_1\}$ for a fixed $s < t_1 < \infty$

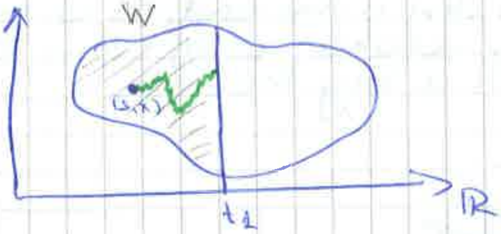

Note that Y_t always moves to the right because it's first component is $s+t$.

Let $\alpha = \inf\{t \ge 0 \;\; Y_t \notin W\} \le t_1 - s$

Let
$$u(r,z) = \begin{cases} v & \text{if } (r,z) \in W \\ u^*(r,z) & \text{if } (r,z) \in G \setminus W \end{cases}$$

Where $v \in U$ is arbitrary. Then

$$\Phi(Y_\alpha) \underset{\substack{\uparrow\\ u^* \text{ optimal}}}{=} J^{u^*}(Y_\alpha) \underset{\substack{\uparrow\\ J^{u^*}(Y_\alpha) \text{ depends only on } u \text{ outside } W}}{=} J^u(Y_\alpha)$$

$$\Rightarrow \Phi(y) \ge J^u(y) \underset{(2.6)}{=} \mathbb{E}_y\Big[\int_0^{\alpha} f(Y_r, u(Y_r))\,dr\Big] + \mathbb{E}_y[\,\Phi(Y_\alpha)]$$

Since $\Phi \in C^2$ ~~and~~ Dynkins formula gives

We want to use this but conditions are not satisfied. Approximate ...

$$\mathbb{E}_y[\,\Phi(Y_\alpha)] = \Phi(y) + \mathbb{E}_y\Big[\int_0^{\alpha} (L^u \Phi)(Y_r)\,dr\Big]$$

... bounded sol and use dom'd convergence

⑥

$$\underset{\text{Def } u}{\Rightarrow} \quad \mathbb{E}_y\left[\int_0^\alpha f(Y_r, v) + (L^v \Phi)(Y_r)\, dr\right] \le 0$$

$$\Rightarrow \quad \frac{1}{\mathbb{E}_y[\alpha]} \mathbb{E}_y\left[\int_0^\alpha f(Y_r, v) + (L^v \Phi)(Y_r)\, dr\right] \le 0 \quad \forall t_1 > s.$$

Taking $t_1 \downarrow s$ the continuity of the involved functions yields

$$f(y,v) + (L^v \Phi)(y) \le 0 \quad \forall y \in G, v \in U.$$

It remains to show the boundary condition and that the supremum is attained.

For the boundary condition:

Since u^* is optimal, we have

$$\Phi(y) = J^{u^*}(y) = \mathbb{E}_y\left[\int_0^{\tau_G} f(Y_s, u^*(Y_s))\, ds + g(Y_{\tau_G}) 1_{\tau_G < \infty}\right]$$

~~If $y \in \partial G$~~ If $y \in \partial G$, then $\mathbb{P}_y(\tau_G = 0) = 1$ (since ∂G is regular) and therefore

$$\Phi(y) = 0 + g(y) \quad \checkmark$$

The identity $\quad L^{u^*}\Phi(y) = -f(y, u^*(y)) \quad \forall y \in G$

says that Φ solves the Dirichlet Poisson problem. Maybe I say sth about this in the end. It is not too hard to show that any solution is of this form but that Φ really solves it is a big theorem.

$\blacksquare$

~~Thm~~ ~~The HJB(II) equation.~~

Interpretation of HJB(I): When an optimal control exists and everything is nice, then its value at point y is the maximizer of

$$v \mapsto f(y,v) + (L^v \Phi)(y)$$

$\Rightarrow$ can solve real valued optimization problem instead of stochastic control problem.

But we know only that being the maximizer is necessary. Is it also sufficient?

Thm The HJB(II) equation

Let $\Phi \in C^2(G)\cap C(\bar{G})$ such that $\forall v\in U$

$$f(y,v) + (L^v\Phi)(y) \le 0 \quad \forall y\in G \quad \text{and}$$

$$\lim_{t\to\tau_G} \Phi(Y_t) = g(Y_{\tau_G})\, 1_{\tau_G<\infty} \quad P_y\text{-a.s.}$$

$\{\Phi^-(Y_\tau) : \tau$ stopping time $\tau\le\tau_G\}$ uniformly P_y integrable for all Markov controls u and $y\in G$.

Then $\Phi(y) \ge J^u(y)$ $\forall$ Markov controls u and all $y\in G$.

Moreover, if $\forall y\in G$ $\exists$ $u_0(y)$ such that

$$f(y,u_0(y)) + L^{u_0(y)}\Phi(y) = 0$$

and

$\{\Phi(Y_\tau) : \tau$ stopping time $\tau\le\tau_G\}$ uniformly P_y integrable $\forall u\in\mathcal{A}$, $y\in G$, and if u_0 is admissible, ~~then~~

then ~~u is a Markov control such that~~ $\Phi(y) = J^{u_0}(y)$

and ~~if u_0 is admissible,~~ $\Phi(y) = \bar{\Phi}(y)$.

Proof Let u be a Markov control and set

$$T_R = \min\{R, \tau_G, \inf\{t>0 : |Y_t|\ge R\}\} \quad \forall R<\infty.$$

By Dynkin's formula

$$E_y[\Phi(Y_{T_R})] = \Phi(y) + E_y\left[\int_0^{T_R} (L^u\Phi)(Y_r)\,dr\right]$$

$$\underset{\text{assumption}}{\le} \Phi(y) - E_y\left[\int_0^{T_R} f(Y_r, u(Y_r))\,dr\right]$$

$$\Rightarrow \Phi(y) \ge \liminf_{R\to\infty} E_y\left[\Phi(Y_{T_R}) + \int_0^{T_R} f(Y_r,u(Y_r))\,dr\right]$$

$$\underset{\substack{\text{Assumptions}\\ \text{incl (1.4)}}}{\ge} E_y\left[g(Y_{\tau_G})\,1_{\tau_G<\infty} + \int_0^{\tau_G} f(Y_r,u(Y_r))\,dr\right] = J^u(y)$$

In the second case all inequalities become identities ∎

Intuition:...

The HJB equations provide conditions which allow to guess and verify. But they are only for Markov controls. The next Theorem says that this is not too restrictive.

Thm 2.3

$\Phi_M(y) = \sup\{ J^u(y) : u \in A_M\}$ A_M = Markov controls

$\Phi_a(y) = \sup\{ J^u(y) : u \in A_G\}$ A_G = all allowed controls

Suppose there exists an optimal Markov control u_0, ie $\Phi_M(y) = J^{u_0}(y)$ $\forall y \in G$, such that all boundary points of G are regular w.r.t $(Y_t^{u_0})_t$ and that

Φ_M is a bounded function in $C^2(G) \cap C(\bar{G})$ with

$E_y[|\Phi_M(Y_\alpha)| + \int_0^\alpha |L^u \Phi_M(Y_t)| dt] < \infty$

$\forall$ bounded stopping times $\alpha \le \tau_G$, all adapted controls u and all $y \in G$.

Then

$$\Phi_M(y) = \Phi_a(y) \qquad \forall y \in G.$$

Proof

Let Φ be a bounded fctn in $C^2(G) \cap C(\bar{G})$ with

$$E_y[|\Phi(Y_\alpha)| + \int_0^\alpha |L^u \Phi(Y_t)| dt] < \infty$$

for all bounded stopping times $\alpha \le \tau_G$, adapted controls u and all $y \in G$.

Moreover, assume

$$f(y,v) + (L^v \Phi)(y) \le 0 \quad \forall y \in G, v \in U$$

and $\quad \Phi(y) = g(y) \quad \forall y \in \partial G.$

Notice that Φ_M satisfies these condition by HJB (I) and the assumptions.

Let Y be the Itô process associated to an adapted control $\tilde{u}$.

Then Dynkins formula yields

$$E_y[\Phi(Y_{T_R})] = \Phi(y) + E_y[\int_0^{T_R} (L^{\tilde{u}} \Phi)(Y_t) dt]$$

$$\underset{\text{assumption}}{\le} \Phi(y) - E_y[\int_0^{T_R} f(Y_t, \tilde{u}_t) dt]$$

$$\Rightarrow \Phi(y) \ge \liminf_{R \to \infty} E_y[\Phi(Y_{T_R}) + \int_0^{T_R} f(Y_t, \tilde{u}_t) dt]$$

$$\underset{\Phi \text{ bounded}}{=} E_y[\Phi(Y_{\tau_G}) 1_{\tau_G < \infty} + \int_0^{\tau_G} f(Y_t, \tilde{u}_t) dt] = J^{u}(y). \quad \text{choosing } \Phi = \Phi_M.$$

Remark
We can apply the theory to the minimum problem

$$\Psi(y) = \inf_{u\in\mathcal{A}} J^u(y) = -\sup_{u\in\mathcal{A}} -J^u(y)$$

$$= -\sup_{u\in\mathcal{A}} \mathbb{E}_y\Big[\int_0^{\tau_G} -f(Y_t, u_t)\,dt - g(Y_{\tau_G})1_{\tau_G<\infty}\Big]$$

$\Rightarrow$ $-\Psi(y)$ is the optimal performance function for profit rate function $-f$ and bequest function $-g$
The HJB eq turns into

$$\inf_{v\in U}\{f(y,v) + L^v\Psi(y)\} = 0 \quad \forall y\in G$$

Example

Recall from the first lecture

$$X_t = x + \int_0^t (1-u_s)\,ds + \int_0^t \sigma\,dB_s \qquad m = n = 1.$$

$$\Rightarrow \quad b(r,x,u) = 1-u \qquad \sigma(r,x,u) = \sigma, \quad U = [0,1]$$

$$L^u\Phi(y) = \frac{\partial\Phi}{\partial s}(y) + (1-u)\frac{\partial\Phi}{\partial x}(y) + \frac{1}{2}\sigma^2\frac{\partial^2\Phi}{\partial x^2}.$$

$$J^u(x,s) = \mathbb{E}_{x,s}\Big[\int_s^{\hat\tau_G} e^{-rt}u_t\,dt\Big]$$

$$\Rightarrow \quad f(y,u) = e^{-rs}u \qquad g = 0$$

$G = \mathbb{R}\times(0,\infty)$ $\quad \hat\tau_G =$ the ~~first~~ time of bankruptcy

HJB eq:

$$f(y,v) + L^v\Phi(y) = e^{-rs}v + \frac{\partial\Phi}{\partial s}(y) + (1-v)\frac{\partial\Phi}{\partial x} + \frac{1}{2}\sigma^2\frac{\partial^2\Phi}{\partial x^2}$$

what leads to the same eq as we solved in the example there $\rightsquigarrow$ eq (3).