Math Material for Algo- & HFT, Book by Cartea-Jaimungal-Penalva

User 1006 • 3 Dec 2024

This is just mathematical background needed to understand the book.It contains the Math Appendix and Ch5 (the Math tools chapter). I will collect substantive results from the book in a separate post.

For further reference:

1. Bertsekas-Shreve 1978: Stochastic Optimal Control, The Discrete Time Case (fairly advanced due to its generality and abstractness, to my surprise)
2. Yong-Zhou 1999: SMP & HJB
3. Fleming-Soner 2006: Controlled Markov Processes and Viscosity Solutions
4. Oksendal-Sulem 2007: Applied Stochastic Control of Jump Diffusions
5. Pham 2010: Cont-time Stochastic Control and Optimization with Fin Applications
6. Touzi 2013: Optimal Stochastic Control, Stochastic Target Problems, and Backard SDE (advanced)

Appendix: Stochastic Calculus

definition of Brownian Motion;

definition of stochastic (Ito) integral. Shorter way:

1. for adapted, L2 proecss g, define
   $I_t = \int_0^t g_s dW_s = \lim_{\|\Pi_k\|\to 0} \sum_{m=1}^{n_k} g_{t_{m-1}}(W_{t_m} - W_{t_{m-1}})$
2. this I_t is called an **Ito process**, and often written as $dI_t = g_t dW_t$. Can show that this SP (I_t) is a martingale.
3. More generally, Ito process can be written as $I_t = \int_0^t \mu_s ds + \int_0^t \sigma_s dW_s$

4. generally, we just need μ_t and σ_t to be adapted and satisfy certain integrability conditions. But in the special case where $\mu_t, \sigma_t = \mu(t, I_t), \sigma(t, I_t)$, the equation $dI_t = \mu_t dt + \sigma_t dW_t$ is called **SDE**. But namewise, book also mentioned that **Ito processes are stochastic processes satisfying SDEs with Brownian noise terms**.

definition of stochastic (Ito) integral. Rigorous way:

1. define Ito integral for simple functions
2. prove that any $g \in L^2$ can be approximated, in $L2$, by sequence of simple functions
3. define Ito integral for g as the limiting value of the Ito integral of the sequence of simple functions

Ito isometry: for adapted $g \in L^2$, $E[(\int_0^T g_s dW_s)^2] = E[\int_0^T g_s^2 ds]$

infinitesimal generator: the generalization of derivative of a function, to make it applicable to stochastic process. $\mathcal{L}_t f(x) \equiv \lim_{h \downarrow 0} \frac{E[f(X_{t+h})|X_t = x] - f(x)}{h}$

This is the generator of an Ito process satisfying a certain SDE, e.g., $dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t$.

Jump Processes

Poisson process
$N = (N_t)_{t \in [0,T]}$, valued in $\mathbb{Z}_+$, with intensity param λ, is a SP s.t.:

1. $N_0 = 0$, a.s.
2. $N_t - N_0$ has Poisson distro with param λt: $P(N_t - N_0 = n) = e^{-\lambda t} \frac{(\lambda t)^n}{n!}$
3. has independent increments: $N_t - N_s$ is independent of $N_v - N_u$
4. has stationary increments: $N_{s+t} - N_s \stackrel{d}{=} N_t$

Classic result 1: time between successive jumps of N are independent, and exponentially distributed.

compensated Poisson process:
$\hat{N}$ with $\hat{N}_t \equiv N_t - \lambda t$. Note that this is a martingale.

as with BM, we can define stochastic integrals wrt compensated Poisson processes in a way that the resulting integral process is a martingale.

let g be adapted, define stochastic integral $Y = (Y_t)_{t \in [0,T]}$ of g wrt $\hat{N}$ by:

$Y_t = \int_0^t g_{s^-} d\hat{N}_s = \sum_{k=1}^{N_t} g_{\tau_k^-} - \int_0^t g_s \lambda ds$, where the τ_i's are jump times.

1. need g_{s^-}, not g_s, to make integral a martingale.

2. alt. def: replace first term with $\sum_{s\in[0,t]}^{N_t} g_{s^-}\Delta N_s$, where $\Delta N_s \equiv N_s - N_{s^-}$, which in this case is either 0 or 1. **sum over a continuum of s, what is the formal def?**

Ito formula for Poisson process

recall we can write $dY_t = -\lambda g_t dt + g_{t^-}\Delta N_t$

Ito's formula for such process Y is:
suppose $Z = (Z_t)_{t\in[0,T]}$ satisfies $Z_t = f(t, Y_t)$ for differentiable f. Then:

$$dZ_t = (\partial_t f(t, Y_t) - \lambda g_t \partial_y f(t, Y_t))dt + [f(t, Y_{t^-} + g_{t^-}) - f(t, Y_{t^-})]dN_t$$

or written in compensated poisson process.

we also see from this (compensated version) formula that the generator of the process Y is $\mathcal{L}_t^Y f(y) = \lambda([f(y + g_t) - f(y)] - g_t\partial_y f(y))$

jump diffusion $Y_t = \int_0^t f_s ds + \int_0^t g_s dW_s + \int_0^t h_{s^-} d\hat{N}_s$

Ito formula for jump diffusion for the above Y, let $Z = (Z_t)$ be defined by $Z_t = \ell(t, Y_t)$, then:

$$dZ_t = (\partial_t + f_t\partial_y + \tfrac{1}{2}g_t^2\partial_{yy} - \lambda h_t\partial_y)\ell(t, Y_t)dt + [\ell(t, Y_{t^-} + h_{t^-}) - \ell(t, Y_{t^-})]dN_t$$

again, common to write it using compensated poisson process $d\hat{N}_t$.

compound Poisson process $J = (J_t)_{t\in[0,T]}$ is built out of:

1. a Poisson process N with intensity λ
2. a collection of iid RVs $\{\epsilon_1, \epsilon_2, ...\}$, with common distro F

$J_t \equiv \sum_{k=1}^{N_t} \epsilon_i$. The process jumps when Poisson event arrives, but the jump size is drawn from F.

We can show, as before:

1. $\hat{J}$ defined by $\hat{J}_t = J_t - E[\epsilon]\lambda t$ is a martingale
2. we can define stochastic integral wrt compound Poisson too.
   $\int_0^t g_{s^-} d\hat{J}_s = \sum_{s\le t} g_{s^-}\Delta J_s - \int_0^t g_s \lambda E[\epsilon] ds$
3. note that $\Delta J_t = J_t - J_{t^-} = \epsilon_{N_s}\Delta N_s$
4. a Ito's formula for $Z_t \equiv \ell(t, Y_t)$ where $dY_t = f_t dt + g_t dW_t + h_{t^-} d\hat{J}_t$

Doubly Stochastic Poisson Processes (Cox process)

these are jump processes which has stochastic intensity

1. given counting process N, we want its intensity process $\lambda = (\lambda_t)_{t\in[0,T]}$ be stochastic
2. the approach is to give a way to compute the probability that an event arrives at t, given info we have at time s: to define $P(N_t - N_s = n|\mathcal{F}_s)$, where $\mathcal{F}$ is the natural filtration generated by (N, λ).
3. $P(N_t - N_s = n|\mathcal{F}_s \vee \sigma(\{\lambda_u\}_{u\in[s,t]})) = e^{-\int_s^t \lambda_u du}\frac{(\int_s^t \lambda_u du)^n}{n!}$
4. this means: $P(N_t - N_s = n|\mathcal{F}_s) = E[e^{-\int_s^t \lambda_u du}\frac{(\int_s^t \lambda_u du)^n}{n!}|\mathcal{F}_s]$
5. the driver of the intensity process can be diverse, leading to Feller/OU/ Hawkes processes.
6. as before, can define its compensated version $\hat{N}_t = N_t - \int_0^t \lambda_s ds$ which is a martingale; can define stochastic integral wrt the compensated doubly stochastic Poisson process, and can derive a Ito's formula for such integral processes, and from which we can derive an expression for the generator of the **joint process** (N, λ).

Feynman-Kac

certain linear PDEs are related to SDEs.

Let X be an Ito process satisfying: $dX_t = \mu(t, X_t)dt + \sigma(t, X_t)dW_t$.

The generator of X is then $\mathcal{L}_t^X$ where $\mathcal{L}_t^X f = \mu(t, x)\partial_x f + \frac{1}{2}\sigma^2(t, x)\partial_{xx} f$

Now suppose we try to solve PDE:
$\partial_t f(t, x) + \mathcal{L}_t^X f(t, x) + g(t, x) = \gamma(t, x)f(t, x)$, with terminal condition
$f(T, x) = h(x)$,

then we have a probabilistic representation of solution $f(t, x)$:

$f(t, x) = E_{t,x}[\int_t^T e^{-\int_t^s \gamma(u,X_u)du} g(s, X_s)ds + e^{-\int_t^T \gamma(u,X_u)du} h(X_T)]$

(note there is a typo in the book)

Consider the simplest example: $\begin{cases} \partial_t h(t, x) + \frac{1}{2}\partial_{xx} h(t, x) = 0 \\ h(T, x) = H(x) \end{cases}$

1. Now introduce a BM X, and define $f_t \equiv E[H(X_T)|\mathcal{F}_t]$
2. (f_t) is a martingale, and Markov: $f_t = g(t, X_t)$ for some g

3. use Ito's lemma to write out $g(t+h, X_{t+h}) - g(t, X_t)$
4. devide the above by h and take limit, we get $0 = \partial_t g(t,x) + \frac{1}{2}\partial_{xx} g(t,x)$
5. by definition, $g(T,x) = H(x)$. Thus, $g(t,x)$ satisfies the PDE. Recall $g(t,x) \equiv E[H(X_T)|\mathcal{F}_t]$.

Ch5 Stochastic Optimal Control and Stopping

A few motivating examples (just to be familiar with notation)

Merton Problem

value function: $H(S,x) = \sup_{\pi \in \mathcal{A}_{0,T}} E_{S,x}[U(X_T^\pi)]$, where:

1. at t, place π_t dollars in risky asset
2. wealth level is X_t

state dynamics follow:

1. $\frac{dS_t}{S_t} = (\mu - r)S_t dt + \sigma S_t dW_t$
2. $dX_t^\pi = (\pi_t(\mu - r) + rX_t^\pi)dt + \pi_t \sigma dW_t$

$\mathcal{A}_{t,T}$ is the admissible set, the set of $\mathcal{F}$-predictable, self-financing strategies satisfying $\int_t^T \pi_s^2 ds < \infty$. (to prevent doubling strategies)

Optimal Liquidation

$$H(x,S,q) = \sup_{\nu \in \mathcal{A}_{0,T}} E[X_T^\nu + Q_T^\nu(S_T^\nu - \alpha Q_T^\nu) - \phi \int_0^T (Q_s^\nu)^2 ds]$$

state dynamics follow:

1. $dQ_t^\nu = -\nu_t dt$ (note the sign)
2. $dS_t^\nu = -g(\nu_t)dt + \sigma dW_t$
3. $\hat{S}_t^\nu = S_t^\nu - h(\nu_t)$
4. $dX_t^\nu = \nu_t \hat{S}_t^\nu dt$

$\mathcal{A}_{t,T}$ is the set of $\mathcal{F}$-predictable, non-negative bounded strategies (excluding repurchasing of shares, and keep liquidation rate finite)

optimal Limit Order placement identical value function expression, just change the ν to δ, which means that agent posts a LO at $S_t + \delta_t$ when current stock price is S_t.

state dynamics:

1. M_t denotes market orders

2. $S_t = S_0 + \sigma W_t$
3. $dX_t^\delta = (S_t + \delta_t)(-dQ_t^\delta)$
4. $dQ_t^\delta = -\mathbf{1}_{U(P(\delta_t))} dM_t$

book mentioned uniform distribution U, but I don't quite understand it. The logic is straightforward though: the probability of your limit order getting executed is a decreasing function of δ_t.

below the book gave 3 types of problems:

1. control for diffusion proecsses
2. control for counting processes
3. countrol of stopping times

The derivation of DPP & HJB may seem pedantic, but one has to see such arguments at some point. I find the book's exposition at a nice balance of rigor and accessibility (sacrificing a bit generality, compared with other treatment say in Nisio's book) therefore I will not skip the derivations and only provide a "cookbook". Due to the similarities, I will only copy the derivation of DPP&HJB for control of diffusion processes.

general control problem for diffusion processes

problem statement:

$$H(\boldsymbol{x}) = \sup_{\boldsymbol{u} \in \mathcal{A}_{0,T}} E[g(\boldsymbol{X}_T^{\boldsymbol{u}}) + \int_0^T F(s, \boldsymbol{X}_s^{\boldsymbol{u}}, \boldsymbol{u}_s) ds]$$

where:

1. $d\boldsymbol{X}_t^u = \boldsymbol{\mu}(t, \boldsymbol{X}_t^{\boldsymbol{u}}, \boldsymbol{u}_t) dt + \boldsymbol{\sigma}(t, \boldsymbol{X}_t^{\boldsymbol{u}}, \boldsymbol{u}_t) d\boldsymbol{W}_t$
2. $\boldsymbol{X}_0^{\boldsymbol{u}} = \boldsymbol{x}$

Here $\mathcal{A}$ is the set of $\mathcal{F}$-**predictable** processes s.t. the state dynamics admits a strong solution. Also assume some nice properties of $\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t$ such as Lipschitz continuity.

note that predictability is necessary since otherwise the agent may be able to peek into the future to optimize her strategy.

we embed optimization problem into a larger class of problems indexed by time, but equal to the original problem at $t = 0$.

performance criterion (associated with u)
$H^{\boldsymbol{u}}(\boldsymbol{x}) \equiv E_{t,\boldsymbol{x}}[G(\boldsymbol{X}_T^{\boldsymbol{u}}) + \int_t^T F(s, \boldsymbol{X}_s^{\boldsymbol{u}}, \boldsymbol{u}_s) ds]$

value function $H(t,\boldsymbol{x}) \equiv \sup_{\boldsymbol{u}\in\mathcal{A}_{t,T}} H^{\boldsymbol{u}}(t,\boldsymbol{x})$

First, we establish DPP:

$\forall(t,\boldsymbol{x}) \in [0,T]\times\mathbb{R}^n$, and $\forall$ stopping times $\tau \le T$, we have:

$H(t,\boldsymbol{x}) = \sup_{\boldsymbol{u}\in\mathcal{A}} E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$

We prove this by showing two-sided inequality.

$H(t,\boldsymbol{x}) \le \sup_{\boldsymbol{u}\in\mathcal{A}} E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$

1. $H^{\boldsymbol{u}}(t,\boldsymbol{x}) = E_{t,\boldsymbol{x}}[G(\boldsymbol{X}^{\boldsymbol{u}}_T) + \int_\tau^T F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$
2. using LIE:
   $H^{\boldsymbol{u}}(t,\boldsymbol{x}) = E_{t,\boldsymbol{x}}[E_{\tau,\boldsymbol{X}^{\boldsymbol{u}}_\tau}[G(\boldsymbol{X}^{\boldsymbol{u}}_T) + \int_\tau^T F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds] + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$
3. $H^{\boldsymbol{u}}(t,\boldsymbol{x}) = E_{t,\boldsymbol{x}}[H^{\boldsymbol{u}}(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$
4. we know that $\forall \boldsymbol{u}, H(t,\boldsymbol{x}) \ge H^{\boldsymbol{u}}(t,\boldsymbol{x})$
5. $H^{\boldsymbol{u}}(t,\boldsymbol{x}) \le E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$
6. taking supremum over $\boldsymbol{u}\in\mathcal{A}$ on the RHS, and then taking supremum on the LHS, we finally get:
7. $H(t,\boldsymbol{x}) \le \sup_{\boldsymbol{u}\in\mathcal{A}} E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$

Now, we show the reverse inequality.

$H(t,\boldsymbol{x}) \ge \sup_{\boldsymbol{u}\in\mathcal{A}} E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$

1. **assuming the value function is continuous in the space of controls**, we pick a control $\boldsymbol{v}^\epsilon \in \mathcal{A}$ such that it is almost perfect:
   $H(t,\boldsymbol{x}) \ge H^{\boldsymbol{v}^\epsilon}(t,\boldsymbol{x}) \ge H(t,\boldsymbol{x}) - \epsilon$
2. Take an arbitrary control $\boldsymbol{u}$, modify our almost-optimal control:
   $\tilde{\boldsymbol{v}}^\epsilon = \boldsymbol{u}_t \mathbf{1}_{t\le\tau} + \boldsymbol{v}^\epsilon \mathbf{1}_{t>\tau}$
3. Note that this modified control is almost-optimal after τ, but suboptimal on $[t,\tau]$. Anyway, we have $H(t,\boldsymbol{x}) \ge H^{\tilde{\boldsymbol{v}}^\epsilon}(t,\boldsymbol{x})$
4. by LIE: $H^{\tilde{\boldsymbol{v}}^\epsilon}(t,\boldsymbol{x}) = E_{t,\boldsymbol{x}}[H^{\tilde{\boldsymbol{v}}^\epsilon}(\tau, \boldsymbol{X}^{\tilde{\boldsymbol{v}}^\epsilon}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\tilde{\boldsymbol{v}}^\epsilon}_s, \tilde{\boldsymbol{v}}^\epsilon_s)ds]$
5. the RHS above is equal to: $E_{t,\boldsymbol{x}}[H^{\boldsymbol{v}^\epsilon}(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds]$
6. the RHS above satisfies inequality:
   $RHS \ge E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}^{\boldsymbol{u}}_\tau) + \int_t^\tau F(s, \boldsymbol{X}^{\boldsymbol{u}}_s, \boldsymbol{u}_s)ds] - \epsilon$

7. let $\epsilon \downarrow 0$, and take the supremum of RHS, we arrive at the desired inequality.

the DPP is really a sequence of equations. An even more powerful equation can be found by looking at its infinitesimal version – DPE (HJB).

two key ideas in deriving DPE:

1. let τ be small: specifically, let it be the minimum between the time it takes for $\boldsymbol{X}_t^{\boldsymbol{u}}$ to exit a ball of radius ϵ around the starting position, and a fixed small time h. **note that if we let $h \downarrow 0$, we would eventually have $\tau = t + h$, since as h shrinks, it is less and less likely that $\boldsymbol{X}$ will exit the ball first**.

2. **assuming enough regularity of value function**, write the value function using Ito's lemma:

$$H(\tau, \boldsymbol{X}_\tau^{\boldsymbol{u}}) = H(t, \boldsymbol{x}) + \int_t^\tau (\partial_t + \mathcal{L}_s^{\boldsymbol{u}}) H(s, \boldsymbol{X}_s^{\boldsymbol{u}}) ds + \int_t^\tau \boldsymbol{\mathcal{D}}_x H(s, \boldsymbol{X}_s^{\boldsymbol{u}})' \boldsymbol{\sigma}_s^{\boldsymbol{u}} dW_s$$

Note that here $\boldsymbol{u}$ is an **arbitrary** control, and $\mathcal{L}_t^{\boldsymbol{u}}$ is the infinitesimal generator of $\boldsymbol{X}_s^{\boldsymbol{u}}$. Note also that as an example, $\mathcal{L}_t^u = \mu(t, x, u)\partial_x + \frac{1}{2}\sigma^2(t, x, u)\partial_{xx}$, so it is about the local behavior: note that the u in μ, σ^2 is an action, not the whole (time-indexed) strategy.

Now, we derive HJB: $\partial_t H(t, \boldsymbol{x}) + \sup_{\boldsymbol{u} \in \mathcal{A}} (\mathcal{L}_t^{\boldsymbol{u}} H(t, \boldsymbol{x}) + F(t, \boldsymbol{x}, \boldsymbol{u})) = 0$

We prove this by showing two-sided inequality.

$\partial_t H(t, \boldsymbol{x}) + \sup_{\boldsymbol{u} \in \mathcal{A}} (\mathcal{L}_t^{\boldsymbol{u}} H(t, \boldsymbol{x}) + F(t, \boldsymbol{x}, \boldsymbol{u})) \leq 0$

1. take arbitrary $\boldsymbol{v} \in \mathcal{A}$ **such that it is CONSTANT over** $[t, \tau]$

2. as shown before (DPP),
   $H(t, \boldsymbol{x}) \geq \sup_{\boldsymbol{u} \in \mathcal{A}} E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}_\tau^{\boldsymbol{u}}) + \int_t^\tau F(s, \boldsymbol{X}_s^{\boldsymbol{u}}, \boldsymbol{u}_s) ds]$

3. $RHS \geq E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}_\tau^{\boldsymbol{v}}) + \int_t^\tau F(s, \boldsymbol{X}_s^{\boldsymbol{v}}, \boldsymbol{v}) ds]$

4. $RHS = E_{t,\boldsymbol{x}}[H(t, \boldsymbol{x}) + \int_t^\tau (\partial_t + \mathcal{L}_s^{\boldsymbol{v}}) H(s, \boldsymbol{X}_s^{\boldsymbol{v}}) ds + \int_t^\tau \mathcal{D}_x H(s, \boldsymbol{X}_s^{\boldsymbol{v}})' \boldsymbol{\sigma}_s^{\boldsymbol{v}} dW_s + \int_t^\tau F(s, \boldsymbol{X}_s^{\boldsymbol{v}}, \boldsymbol{v}) ds]$
   by our choice of τ, can show the stochastic integral is indeed a martingale, therefore we have:

6. $H(t, \boldsymbol{x}) \geq E_{t,\boldsymbol{x}}[H(t, \boldsymbol{x}) + \int_t^\tau (\partial_t + \mathcal{L}_s^{\boldsymbol{v}}) H(s, \boldsymbol{X}_s^{\boldsymbol{v}}) ds + \int_t^\tau F(s, \boldsymbol{X}_s^{\boldsymbol{v}}, \boldsymbol{v}) ds]$

7. Now we let $h \downarrow 0$, so that $\tau = t + h, a.s.$

8. $0 \geq \lim_{h \downarrow 0} E_{t,\boldsymbol{x}}[\frac{1}{h} \int_t^\tau \{(\partial_t + \mathcal{L}_s^{\boldsymbol{v}}) H(s, \boldsymbol{X}_s^{\boldsymbol{v}}) + F(s, \boldsymbol{X}_s^{\boldsymbol{v}}, \boldsymbol{v})\} ds]$

9. RHS is equal to: $(\partial_t + \mathcal{L}_t^{\boldsymbol{v}})H(t,\boldsymbol{x}) + F(t,\boldsymbol{x},\boldsymbol{v})$, where we used the mean value theorem

10. this this inequality holds for arbitrary $v \in \mathcal{A}$, take the supremum we have: $\partial_t H(t,\boldsymbol{x}) + \sup_{\boldsymbol{u}\in\mathcal{A}} (\mathcal{L}_t^{\boldsymbol{u}} H(t,\boldsymbol{x}) + F(t,\boldsymbol{x},\boldsymbol{u})) \leq 0$

note that the $\boldsymbol{v}$ here is used a bit loosely here. Sometimes it denotes the whole strategy, sometimes it denotes the constant action applied in $[t,\tau]$.

Now we show the reverse inequality, by showing that for the optimal control $\boldsymbol{u}^*$, we have $\partial_t H(t,\boldsymbol{x}) + \sup_{\boldsymbol{u}\in\mathcal{A}} (\mathcal{L}_t^{\boldsymbol{u}} H(t,\boldsymbol{x}) + F(t,\boldsymbol{x},\boldsymbol{u})) = 0$

1. by LIE, $H(t,\boldsymbol{x}) = E_{t,\boldsymbol{x}}[H(\tau, \boldsymbol{X}_\tau^{\boldsymbol{u}^*}) + \int_t^\tau F(s, \boldsymbol{X}_s^{\boldsymbol{u}^*}, \boldsymbol{u}^*)ds]$

2. apply Ito's lemma as before, writing $H(\tau, \boldsymbol{X}_\tau^{\boldsymbol{u}^*})$ in terms of $H(t,\boldsymbol{x})$, we will find the desired equality.

Combined these two parts, we arrive at DPE (HJB):

$$\begin{cases} \partial_t H(t,\boldsymbol{x}) + \sup_{\boldsymbol{u}\in\mathcal{A}} (\mathcal{L}_t^{\boldsymbol{u}} H(t,\boldsymbol{x}) + F(t,\boldsymbol{x},\boldsymbol{u})) = 0 \\ H(T,\boldsymbol{x}) = G(\boldsymbol{x}) \end{cases}$$

Note that optimal control in HJB is an action, and can be written in feedback form in terms of the value function. Substituting this optimal control back into HJB we get non-linear PDEs.

Often we define (maximized) Hamiltonian as:

$$\mathfrak{H}(t,x,\boldsymbol{\mathcal{D}}_x H, \boldsymbol{\mathcal{D}}_{xx}^2 H) = \sup_{\boldsymbol{u}\in\mathcal{A}} (\mathcal{L}_t^{\boldsymbol{u}} H(t,x) + F(t,x,u))$$

Note that some define Hamiltonian with generic costates, and optimality is associated with those costates being equal to partial derivatives of value function.

DPE provides a necessary condition for optimality. We use verification theorems to prove sufficiency. Basically, it says that if you can find a solution to DPE, and demonstrate that it is a classical solution (once differentiable in time and twice differentiable in state vars), and the resulting control is admissible, then the solution is indeed the value function, and the resulting control is indeed the optimal Markov control. Under some more technical assumptions, can show that the optimal control is indeed Markov, and therefore we have found not just the optimal Markov control but the optimal $\mathcal{F}-$predictable control.

(not so general) control problem for counting processes

just a special case to build intuition: agent control the frequency of jumps of a counting process N.

performance criteria

$H^u(n) \equiv E[G(N_T^u) + \int_0^T F(s, N_s^u)ds]$

value function $H(n) = \sup_{u \in \mathcal{A}_{0,T}} E[G(N_T^u) + \int_0^T F(s, N_s^u, u_s)ds]$, where:

1. $u = (u_t)_{t\in[0,T]}$ is control process
2. $(N_t^u)_{t\in[0,T]}$ is a controlled doubly stochastic Poisson process, starting at $N_{0^-} = n$, with intensity $\lambda_t^u = \lambda(t, N_{t^-}^u, u_t)$
3. as a result, $\hat{N}_t^u = N_t - \int_0^t \lambda_s^u ds$ is a martingale

DPP

$H(t,n) = \sup_{u\in\mathcal{A}} E_{t,n}[H(\tau, N_\tau^u) + \int_t^\tau F(s, N_s^u, u_s)ds]$

DPE

$$\begin{cases} \partial_t H(t,n) + \sup_{\boldsymbol{u}\in\mathcal{A}_t} (\mathcal{L}_t^{\boldsymbol{u}} H(t,n) + F(t,n,u)) = 0 \\ H(T,n) = G(n) \end{cases}$$

Note that $\mathcal{L}_t^u H(t,n) = \lambda(t,n,u)[H(t,n+1) - H(t,n)]$

Thus, if we plug in the infinitesimal form, we get:

$$\sup_{\boldsymbol{u}\in\mathcal{A}_t} \lambda(s,n,u)[H(t,n+1) - H(t,n)] + F(t,n,u)$$

if $F = 0$, then optimal control is to make λ as large as possible or as small as possible depending on the sign of $[...]$, i.e., bang-bang controls.

Two ways to break this uninteresting feature:

1. $F \neq 0$
2. add another SP driven by the counting process, thus this new SP is controlled indirectly, and have this new SP affect performance

Similar DPE for the problem with jump-diffusions.

Stopping Problems

performance criterion $H^\tau(t, \boldsymbol{x}) = E_{t,\boldsymbol{x}}[G(\boldsymbol{X}_\tau)]$

where $\boldsymbol{X} = (\boldsymbol{X}_t)_{t\in[0,T]}$ is a jump diffusion following:

$d\boldsymbol{X}_t = \boldsymbol{\mu}(t, \boldsymbol{X}_t)dt + \boldsymbol{\sigma}(t, \boldsymbol{X}_t)d\boldsymbol{W}_t + \boldsymbol{\gamma}(t, \boldsymbol{X}_t)d\boldsymbol{N}_t$

where $\boldsymbol{N}$ is a multi-dim counting process with intensities $\boldsymbol{\lambda}(t, \boldsymbol{X}_t)$.

value function $H(t, \boldsymbol{x}) = \sup_{\tau \in \mathcal{T}_{[t,T]}} H^\tau(t, \boldsymbol{x})$

The difficult problem is to characterize the (boundary) of the stopping region. Again we have our DPP and DPE.

DPP $H(t, \boldsymbol{x}) = \sup_{\tau \in \mathcal{T}_{[t,T]}} E_{t,\boldsymbol{x}}[G(\boldsymbol{X}_\tau)\mathbf{1}_{\tau<\theta} + H(\theta, \boldsymbol{X}_\theta)\mathbf{1}_{\tau \geq \theta}]$, for all stopping times $\theta \leq T$.

intuition: if τ^* occurs prior to θ, then agent's value function is just the reward at τ^*. If not, then at θ agent receives the value function evaluated at the current state.

DPE H solves the variational inequality:

$\max\{\partial_t H + \mathcal{L}_t H, G - H\} = 0$, on $[0, T] \times \mathbb{R}^m$.

The proof follows Touzi 2013, which is nice as it is related to viscosity solutions (needed when value function itself is not smooth enough to differentiate)

Interpretation of this HJBQVI:

1. in the continuation region, $h_t \equiv H(t, \boldsymbol{X}_t)$ is a martingale
2. in the stopping region, if you don't stop, the linear operator tries to render the value function negative. But we pin it to the reward (constant). therefore, it is again a martingale.
3. so the SP h_t corresponding to the flow of the value function is a martingale on the entire $[0, T] \times \mathbb{R}^m$.

Finally, for combined stopping and control problems, we have DPE of the same format, but now we need to also do optimization in the continuation region, therefore the DPE reads:

$\max\{\partial_t H + \sup_{\boldsymbol{u} \in \mathcal{A}_t} \mathcal{L}_t^{\boldsymbol{u}} H, G - H\} = 0$

You see now that in the continuation region, the value function needs to satisfy a general non-linear HJB, instead of a linear PDE.