Lectures on Mean Field Games: II. Calculus over Wasserstein Space, Control of McKean-Vlasov Dynamics, and the Master Equation

René Carmona

Department of Operations Research & Financial Engineering
PACM
Princeton University

Minerva Lectures, Columbia U. October 2016

THE ANALYTIC (PDE) APPROACH TO MFGS

For **fixed** $\boldsymbol{\mu} = (\mu_t)_t$, the **value function**

$$V^{\boldsymbol{\mu}}(t,x) = \inf_{(\alpha_s)_{t \le s \le T}} \mathbb{E}\Big[\int_t^T f(s, X_s, \mu_s, \alpha_s)ds + g(X_T, \mu_T)\Big|X_t = x\Big]$$

solves a **HJB (backward)** equation

$$\partial_t V^{\boldsymbol{\mu}}(t,x) + \inf_{\alpha}[b(t,x,\mu_t,\alpha)\cdot\partial_x V^{\boldsymbol{\mu}}(t,x) + f(t,x,\mu_t,\alpha)]$$
$$\frac{1}{2}\text{trace}[\sigma(t,x)^\dagger\sigma(t,x)\partial^2_{xx}V^{\boldsymbol{\mu}}(t,x)] = 0$$

with terminal condition $V^{\boldsymbol{\mu}}(T,x) = g(x,\mu_T)$

The fixed point step is implemented by requiring that $t \to \mu_t$ solves the **(forward)** Kolmogorov equation

$$\partial_t\mu_t = \mu_t\mathcal{L}_t^\dagger$$

This is also a **nonlinear PDE** because μ_t appears in b ......

System of **strongly coupled nonlinear PDEs**! Time goes in **both directions**

HJB Equation from Itô's Formula

Classical Optimal Control set-up (μ fixed)

Dynamic Programming Principle

$$t \hookrightarrow V^{\mu}(t, X_t) \text{ is a } \textbf{martingale} \text{ when } (X_t)_{0 \le t \le T} \text{ is } \textbf{optimal}$$

Classical Itô formula to compute:

$$d_t V^{\mu}(t, X_t)$$

when $(t, x) \hookrightarrow V^{\mu}(t, x)$ is **smooth** and

$$dX_t = b(t, X_t, \hat{\alpha}_t)dt + \sigma(t, X_t, \hat{\alpha}_t)dW_t$$

is optimal to

- set the drift to 0
- get HJB

MFG Counterpart

- MFG **is not** an optimization problem per-se
- Optimal control arguments (for μ fixed) affected by fixed point step
- What is the effect of last step substitution $\mu_t = \mathbb{P}_{X_t}$?
- In **equilibrium**, do we still have:
  - **Dynamic Programming Principle**?
  - **Martingale property** of

$$t \hookrightarrow V^{\mu}(t, X_t)$$

- What would be the right **Itô formula** to compute:

$$d_t V^{\mu}(t, X_t)$$

when

$$dX_t = b(t, X_t, \hat{\alpha}_t)dt + \sigma(t, X_t, \hat{\alpha}_t)dW_t$$

is optimal and $\mu_t = \mathbb{P}_{X_t}$?

More Reasons to Differentiate Functions of Measures

Back to the **N-player games** (with reduced or distributed controls):

$$dX_t^i = b\big(t, X_t^i, \overline{\mu}_t^N, \phi(t, X_t^i, \overline{\mu}_t^N)\big)\,dt + \sigma\big(t, X_t^i, \overline{\mu}_t^N, \phi(t, X_t^i, \overline{\mu}_t^N)\big)\,dW_t^i\,, \quad t \in [0, T],$$

Propagation of Chaos

- $X_t^1, \cdots, X_t^k, \cdots$ become independent in the limit $N \to \infty$
- $\mathbf{X}^i = (X_t^i)_{0 \le t \le T} \Longrightarrow \mathbf{X} = (X_t)_{0 \le t \le T}$ solution of the McKean–Vlasov equation:

  $$dX_t = b\big(t, X_t, \mathbb{P}_{X_t}, \phi(t, X_t, \mathbb{P}_{X_t})\big)\,dt + \sigma\big(t, X_t, \mathbb{P}_{X_t}, \phi(t, X_t, \mathbb{P}_{X_t})\big)\,dW_t, \quad t \in [0, T],$$

  where $\mathbf{W} = (W_t)_{0 \le t \le T}$ is a standard Wiener process.

Expected Costs:

$$J^i(\phi) = \mathbb{E}\left[\int_0^T f\big(t, X_t^i, \overline{\mu}_t^N, \phi(t, X_t^i, \overline{\mu}_t^N)\big)\,dt + g\big(X_T^i, \overline{\mu}_T^N\big)\right],$$

converge to:

$$J(\phi) = \mathbb{E}\left[\int_0^T f\big(t, X_t, \mathbb{P}_{X_t}, \phi(t, X_t, \mathbb{P}_{X_t})\big)\,dt + g\big(X_T, \mathbb{P}_{X_T}\big)\right].$$

Optimization after the limit: Control of McKean-Vlasov equations !

Taking Stock

SDE State Dynamics for N players	*Optimization* $\longrightarrow$	Nash Equilibrium for N players
$\downarrow$ Fixed Point $\lim_{N\to\infty}$		$\downarrow$ Fixed Point $\lim_{N\to\infty}$
McKean Vlasov Dynamics	*Optimization* $\longrightarrow$	**Mean Field Game?** **Controlled McKean-Vlasov SDE?**

Is the above diagram commutative?

Controlled McKean-Vlasov SDEs

$$\inf_{\boldsymbol{\alpha}=(\alpha_t)_{0\le t\le T}} \mathbb{E}\left[\int_0^T f(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dt + g(X_T, \mathbb{P}_{X_T})\right]$$

under dynamical constraint $dX_t = b(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dt + \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dW_t$.

- State $(X_t, \mathbb{P}_{X_t})$ infinite dimensional
- State trajectory $t \mapsto (X_t, \mu_t)$ is a very thin submanifold due to constraint $\mu_t = \mathbb{P}_{X_t}$
- Open loop form: $\boldsymbol{\alpha} = (\alpha_t)_{0\le t\le T}$ adapted
- Closed loop form: $\alpha_t = \phi(t, X_t, \mathbb{P}_{X_t})$

Whether we use

- Infinite dimensional HJB equation
- Pontryagin stochastic maximum principle with Hamiltonian

$$H(t, x, \mu, y, z, \alpha) = b(t, x, \mu, \alpha)\cdot y + \sigma(t, x, \mu, \alpha)\cdot z + f(t, x, \mu, \alpha)$$

and introduce the adjoint equations,

WE NEED TO DIFFERENTIATE FUNCTIONS OF MEASURES !

Differentiability of Functions of Measures

$\mathcal{M}(\mathbb{R}^d)$ space of **signed** (finite) measures on $\mathbb{R}^d$

- Banach space (dual of a space of continuous functions)
- Classical differential calculus available
- If

$$\mathcal{M}(\mathbb{R}^d) \ni m \hookrightarrow \phi(m) \in \mathbb{R}$$

  "ϕ ***is differentiable***" has a meaning
- For $m_0 \in \mathcal{M}(\mathbb{R}^d)$ one can define

$$\frac{\delta\phi(m_0)}{\delta m}(\cdot)$$

  as a function on $\mathbb{R}^d$ in **Fréchet** or **Gâteaux** sense

Bensoussan-Frehe-Yam alternative is to work only with measures with **densities** and view ϕ as a function on $L^1(\mathbb{R}^d, dx)$!

Measures appearing in MFG theory are probability distributions of random variables !!!

Wasserstein space

$$\mathcal{P}_2(\mathbb{R}^d) = \left\{ \mu \in \mathcal{P}(\mathbb{R}^d);\ \int_{\mathbb{R}^d} |x|^2 d\mu(x) < \infty \right\}$$

Metric space for the 2-Wasserstein distance

$$W_2(\mu, \nu) = \inf_{\pi \in \Pi(\mu,\nu)} \left[\int_{\mathbb{R}^d \times \mathbb{R}^d} |x-y|^2 \pi(dx, dy) \right]^{1/2}$$

where $\Pi(\mu, \nu)$ is the set of probability measures coupling μ and ν.
Topological properties of Wasserstein space well understood as following statements are equivalents

- $\mu^N \longrightarrow \mu$ in Wasserstein space
- $\mu^N \longrightarrow \mu$ weakly and $\int |x|^2 \mu^N(dx) \longrightarrow \int |x|^2 \mu(dx)$

Glivenko-Cantelli in Wasserstein Space

$X^1, X^2, \cdots$, i.i.d. random variables in $\mathbb{R}^d$ with common distribution μ s.t.

$$M_q(\mu) = \int_{\mathbb{R}^d} |x|^q \mu(dx) < \infty.$$

If $q = 2$,

$$\mathbb{P}\left[\lim_{N\to\infty} W_2(\overline{\mu}^N, \mu) = 0\right] = 1.$$

where $\overline{\mu}^N = \frac{1}{N}\sum_{i=1}^N \delta_{X^i}$ is a (random) empirical measure. Standard LLN !

Crucial Estimate: **Glivenko-Cantelli** If $q > 4$ for each dimension $d \geq 1$, $\exists C = C(d, q, M_q(\mu))$ s.t. for all $N \geq 1$:

$$\mathbb{E}\left[W_2(\overline{\mu}^N, \mu)^2\right] \leq C \begin{cases} N^{-1/2}, & \text{if } d < 4, \\ N^{-1/2}\log N, & \text{if } d = 4, \\ N^{-2/d}, & \text{if } d > 4. \end{cases} \tag{1}$$

Differential Calculus on Wasserstein Space

What does it mean "ϕ ***is differentiable***" or "ϕ ***is convex***" for

$$\mathcal{P}_2(\mathbb{R}^d) \ni \mu \hookrightarrow \phi(\mu) \in \mathbb{R}$$

Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$ is a **metric space** for W_2

- **Optimal transportation** (**Monge-Ampere-Kantorovich**)
- Curve length and shortest paths (**geodesics**)
- Notion of **convex function** on $\mathcal{P}_2(\mathbb{R}^d)$
- **Tangent spaces** and **differential geometry** on $\mathcal{P}_2(\mathbb{R}^d)$.
- **Differential calculus** on Wasserstein space

Brenier, Benamou, Ambrosio, Gigli, Otto, Caffarelli, Villani, Carlier, ...

DIFFERENTIABILITY IN THE SENSE OF P.L.LIONS

If $\mathcal{P}_2(\mathbb{R}^d) \ni \mu \hookrightarrow \phi(\mu) \in \mathbb{R}$ is "***differentiable***" on **Wasserstein space** what about

$$\mathbb{R}^{dN} \ni (x^1, \cdots, x^N) \mapsto u(x^1, \cdots, x^N) = \phi\left(\frac{1}{N}\sum_{j=1}^{N} \delta_{x^j}\right) \ ?$$

How does $\partial\phi(\mu)$ relate to $\partial_{x^i} u(x^1, \cdots, x^N)$?

Lions' Solution

- **Lift** ϕ up to $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ into $\tilde{\phi}$ defined by $\tilde{\phi}(X) = \phi(\tilde{\mathbb{P}}_X)$
- Use Fréchet differentials on **flat space** L^2

Definition of L-differentiability

ϕ is differentiable at μ_0 if $\tilde{\phi}$ is Fréchet differentiable at X_0 s.t. $\tilde{\mathbb{P}}_{X_0} = \mu_0$

- Check definition is **intrinsic**

Properties of L-differentials

- $\partial\phi(\mu_0) = D\tilde{\phi}(X_0) \in L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$
- The distribution of the random variable $\partial\phi(\mu_0)$ depends only on μ_0,
  NOT ON THE RANDOM VARIABLE X_0 used to represent it
- $\exists \xi : \mathbb{R}^d \mapsto \mathbb{R}^d$ uniquely defined μ_0 a.e. such that $\partial\phi(\mu_0) = D\tilde{\phi}(X_0) = \xi(X_0)$
- we use $\partial\phi(\mu_0)(\cdot) = \xi$

Examples

$$\phi(\mu) = \int_{\mathbb{R}^d} h(x)\mu(dx) \quad \Longrightarrow \quad \partial\phi(\mu)(\cdot) = \partial h(\cdot)$$

$$\phi(\mu) = \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} h(x-y)\mu(dx)\mu(dy) \quad \Longrightarrow \quad \partial\phi(\mu)(\cdot) = [2\partial h(\cdot) * \mu](\cdot)$$

$$\phi(\mu) = \int_{\mathbb{R}^d} \varphi(x,\mu)\mu(dx) \quad \Longrightarrow \quad \partial\phi(\mu)(\cdot) = \partial_x\varphi(\cdot,\mu) + \int_{\mathbb{R}^d} \partial_\mu\varphi(x',\mu)(\cdot)\mu(dx')$$

Two More Examples

Assume $\phi : \mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}$ is L-differentiable and define

$$\phi^N : \mathbb{R}^d \times \cdots \times \mathbb{R}^d \ni (x^1, \cdots, x^N) \hookrightarrow \phi^N(x^1, \cdots, x^N) = \phi\left(\frac{1}{N}\sum_{i=1}^N \delta_{x^i}\right)$$

$$\partial_{x^i}\phi^N(x^1, \cdots, x^N) = \frac{1}{N}\partial_\mu \phi\left(\frac{1}{N}\sum_{i=1}^N \delta_{x^i}\right)(x_i)$$

Assume $\phi : \mathcal{M}_2(\mathbb{R}^d) \mapsto \mathbb{R}$ has a linear functional derivative (at least in a neighborhood of $\mathcal{P}_2(\mathbb{R}^d)$ and that $\mathbb{R}^d \ni x \mapsto [\delta\phi/\delta m](m)(x)$ is differentiable and the derivative

$$\mathcal{M}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (m, x) \mapsto \partial_x \left[\frac{\delta \phi}{\delta m}\right](m)(x) \in \mathbb{R}^d$$

is jointly continuous in (m, x) and is of linear growth in x, then ϕ is L-differentiable and

$$\partial_\mu \phi(\mu)(\,\cdot\,) = \partial_x \frac{\delta \phi}{\delta m}(\mu)(\,\cdot\,), \quad \mu \in \mathcal{P}_2(\mathbb{R}^d).$$

Convex Functions of Measures

$\phi : \mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}$ is said to be **L-convex** if

$$\forall \mu, \mu' \quad \phi(\mu') - \phi(\mu) - \mathbb{E}[\partial_\mu \phi(\mu)(X) \cdot (X' - X)] \geq 0,$$

whenever $\mathbb{P}_X = \mu$ and $\mathbb{P}_{X'} = \mu'$.

Example1

$$\mu \mapsto \phi(\mu) = g\Big(\int_{\mathbb{R}^d} \zeta(x) d\mu(x) \Big),$$

- for $g : \mathbb{R} \to \mathbb{R}$ is non-decreasing convex differentiable
- and $\zeta : \mathbb{R}^d \to \mathbb{R}$ convex differentiable with derivative of at most of linear growth

Example2

$$\mu \mapsto \phi(\mu) = \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} g(x, x') d\mu(x) d\mu(x')$$

- If $g : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$ is convex differentiable (∂g linear growth)

A sobering counter-example. If $\mu_0 \in \mathcal{P}_2(E)$ is fixed, the square distance function

$$\mathcal{P}_2(E) \ni \mu \to W_2(\mu_0, \mu)^2 \in \mathbb{R}$$

may not be convex or even L-differentiable!

BACK TO THE CONTROL OF MCKEAN-VLASOV EQUATIONS

$$\inf_{\boldsymbol{\alpha}=(\alpha_t)_{0\le t\le T}} \mathbb{E}\left[\int_0^T f\big(t, X_t, \mathbb{P}_{X_t}, \alpha_t\big)\,dt + g\big(X_T, \mathbb{P}_{X_T}\big)\right]$$

under the dynamical constraint

$$dX_t = b\big(t, X_t, \mathbb{P}_{X_t}, \alpha_t\big)\,dt + \sigma\big(t, X_t, \mathbb{P}_{X_t}, \alpha_t\big)\,dW_t.$$

EXAMPLE: POTENTIAL MEAN FIELD GAMES

Start with **Mean Field Game** à la **Lasry-Lions**

$$\inf_{\boldsymbol{\alpha}=(\alpha_t)_{0\le t\le T},\ dX_t=\alpha_t dt+\sigma dW_t} \mathbb{E}\Big[\int_0^T \big[\frac{1}{2}|\alpha_t|^2 + f(t, X_t, \mu_t)\big]\,dt + g(X_T, \mu_T)\Big]$$

s.t. f and g are differentiable w.r.t. x and there exist differentiable functions F and G

$$\partial_x f(t,x,\mu) = \partial_\mu F(t,\mu)(x) \quad \text{and} \quad \partial_x g(x,\mu) = \partial_\mu G(\mu)(x)$$

Solving this **MFG** is equivalent to solving the **central planner** optimization problem

$$\inf_{\boldsymbol{\alpha}=(\alpha_t)_{0\le t\le T},\ dX_t=\alpha_t dt+\sigma dW_t} \mathbb{E}\Big[\int_0^T \big[\frac{1}{2}|\alpha_t|^2 + F(t, \mathbb{P}_{X_t})\big]\,dt + G(\mathbb{P}_{X_T})\Big]$$

Special case of **McKean-Vlasov optimal control**

The Adjoint Equations

Lifted Hamiltonian

$$\tilde{H}(t, x, \tilde{X}, y, \alpha) = H(t, x, \mu, y, \alpha)$$

for any random variable $\tilde{X}$ with distribution μ.

Given an admissible control $\boldsymbol{\alpha} = (\alpha_t)_{0\le t\le T}$ and the corresponding controlled state process $\mathbf{X}^\alpha = (X_t^\alpha)_{0\le t\le T}$, any couple $(Y_t, Z_t)_{0\le t\le T}$ satisfying:

$$\begin{cases} dY_t = -\partial_x H(t, X_t^\alpha, \mathbb{P}_{X_t^\alpha}, Y_t, \alpha_t)dt + Z_t dW_t \\ \qquad\qquad\qquad -\tilde{\mathbb{E}}[\partial_\mu H(t, \tilde{X}_t, X, \tilde{Y}_t, \tilde{\alpha}_t)]|_{X=X_t^\alpha} dt \\ Y_T = \partial_x g(X_T^\alpha, \mathbb{P}_{X_T^\alpha}) + \tilde{\mathbb{E}}[\partial_\mu g(x, \tilde{X}_t)]|_{x=X_T^\alpha} \end{cases}$$

where $(\tilde{\alpha}, \tilde{X}, \tilde{Y}, \tilde{Z})$ is an independent copy of (α, X^α, Y, Z), is called a set of **adjoint processes**

BSDE of Mean Field type according to Buckhdan-Li-Peng !!!

Extra terms in red are the ONLY difference between MFG and Control of McKean-Vlasov dynamics !!!

Pontryagin Maximum Principle (Sufficiency)

Assume

1. Coefficients continuously differentiable with bounded derivatives;
2. Terminal cost function g is convex;
3. $\boldsymbol{\alpha} = (\alpha_t)_{0 \le t \le T}$ admissible control, $\mathbf{X} = (X_t)_{0 \le t \le T}$ corresponding dynamics, $(\mathbf{Y}, \mathbf{Z}) = (Y_t, Z_t)_{0 \le t \le T}$ adjoint processes and

$$(x, \mu, \alpha) \hookrightarrow H(t, x, \mu, Y_t, Z_t, \alpha)$$

is $dt \otimes d\mathbb{P}$ a.e. **convex**,

then, if moreover

$$H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) = \inf_{\alpha \in A} H(t, X_t, \mathbb{P}_{X_t}, Y_t, \alpha), \quad \text{a.s.}$$

Then α is an optimal control, i.e.

$$J(\boldsymbol{\alpha}) = \inf_{\boldsymbol{\beta} \in \mathbb{A}} J(\boldsymbol{\beta}).$$

Particular Case: Scalar Interactions

$$b(t,x,\mu,\alpha) = \tilde{b}(t,x,\langle\psi,\mu\rangle,\alpha) \quad \sigma(t,x,\mu,\alpha) = \tilde{\sigma}(t,x,\langle\phi,\mu\rangle,\alpha)$$
$$f(t,x,\mu,\alpha) = \tilde{f}(t,x,\langle\gamma,\mu\rangle,\alpha) \quad g(x,\mu) = \tilde{g}(x,\langle\zeta,\mu\rangle)$$

- ψ, ϕ, γ and ζ **differentiable** with at most quadratic growth at ∞,
- $\tilde{b}$, $\tilde{\sigma}$ and $\tilde{f}$ **differentiable** in $(x,r) \in \mathbb{R}^d \times \mathbb{R}$ for $t,\alpha)$ fixed
- $\tilde{g}$ **differentiable** in $(x,r) \in \mathbb{R}^d \times \mathbb{R}$.

Recall that the adjoint process satisfies

$$Y_T = \partial_x g(X_T,\mathbb{P}_{X_T}) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T,\mathbb{P}_{\tilde{X}_T})(X_T)].$$

but since

$$\partial_\mu g(x,\mu)(x') = \partial_r \tilde{g}\big(x,\langle\zeta,\mu\rangle\big)\partial\zeta(x'),$$

the terminal condition reads

$$Y_T = \partial_x \tilde{g}\big(X_T,\mathbb{E}[\zeta(X_T)]\big) + \tilde{\mathbb{E}}\big[\partial_r \tilde{g}\big(\tilde{X}_T,\mathbb{E}[\zeta(X_T)]\big)\big]\partial\zeta(X_T)$$

Convexity in μ follows convexity of $\tilde{g}$

Scalar Interactions (cont.)

$$H(t,x,\mu,y,z,\alpha)=\tilde{b}(t,x,\langle\psi,\mu\rangle,\alpha)\cdot y+\tilde{\sigma}(t,x,\langle\phi,\mu\rangle,\alpha)\cdot z+\tilde{f}(t,x,\langle\gamma,\mu\rangle,\alpha).$$

$\partial_\mu H(t,x,\mu,y,z,\alpha)$ can be identified wih

$$\begin{aligned}\partial_\mu H(t,x,\mu,y,z,\alpha)(x') = &\big[\partial_r\tilde{b}(t,x,\langle\psi,\mu\rangle,\alpha)\cdot y\big]\partial\psi(x')\\ &+\big[\partial_r\tilde{\sigma}(t,x,\langle\phi,\mu\rangle,\alpha)\cdot z\big]\partial\phi(x')\\ &+\partial_r\tilde{f}(t,x,\langle\gamma,\mu\rangle,\alpha)\,\partial\gamma(x')\end{aligned}$$

and the adjoint equation rewrites:

$$\begin{aligned}dY_t = &-\Big\{\partial_x\tilde{b}(t,X_t,\mathbb{E}[\psi(X_t)],\alpha_t)\cdot Y_t+\partial_x\tilde{\sigma}(t,X_t,\mathbb{E}[\phi(X_t)],\alpha_t)\cdot Z_t\\ &+\partial_x\tilde{f}(t,X_t,\mathbb{E}[\gamma(X_t)],\alpha_t)\Big\}dt+Z_t dW_t\\ &-\Big\{\tilde{\mathbb{E}}\big[\partial_r\tilde{b}(t,\tilde{X}_t,\mathbb{E}[\psi(\tilde{X}_t)],\tilde{\alpha}_t)\cdot\tilde{Y}_t\big]\partial\psi(X_t)+\tilde{\mathbb{E}}\big[\partial_r\tilde{\sigma}(t,\tilde{X}_t,\mathbb{E}[\phi(\tilde{X}_t)],\tilde{\alpha}_t)\cdot\tilde{Z}_t\big]\partial\phi(X_t)\\ &+\tilde{\mathbb{E}}\big[\partial_r\tilde{f}((t,\tilde{X}_t,\mathbb{E}[\gamma(\tilde{X}_t)],\tilde{\alpha}_t)\big]\partial\gamma(X_t)\Big\}dt\end{aligned}$$

Anderson - Djehiche

Solution of the McKV Control Problem

Assume

- $b(t, x, \mu, \alpha) = b_0(t) \int_{\mathbb{R}^d} x d\mu(x) + b_1(t)x + b_2(t)\alpha$
  with b_0, b_1 and b_2 is $\mathbb{R}^{d\times d}$-valued and are bounded.
- f and g as in MFG problem.

Thn there exists a solution $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = (X_t, Y_t, Z_t)_{0\le t\le T}$ of the McKean-Vlasov FBSDE

$$\begin{cases} dX_t = b_0(t)\mathbb{E}(X_t)dt + b_1(t)X_t dt + b_2(t)\hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t)dt + \sigma dW_t, \\ dY_t = -\partial_x H\big(t, X_t, \mathbb{P}_{X_t}, Y_t, \hat{\alpha}_t\big)\, dt \\ \qquad\qquad - \mathbb{E}\big[\partial_\mu \tilde{H}\big(t, \tilde{X}_t, X_t, \tilde{Y}_t, \tilde{\hat{\alpha}}_t\big)\big]\, dt + Z_t dW_t. \end{cases}$$

with $Y_t = u(t, X_t, \mathbb{P}_{X_t})$ for a function

$$u : [0, T] \times \mathbb{R}^d \times \mathcal{P}_1(\mathbb{R}^d) \ni (t, x, \mu) \mapsto u(t, x, \mu)$$

uniformly of Lip-1 and with linear growth in x.

A Finite Player Approximate Equilibrium

For N independent Brownian motions $(W^1, \ldots, W^N)$ and for a square integrable exchangeable process $\beta = (\beta^1, \ldots, \beta^N)$, consider the system

$$dX_t^i = \frac{1}{N} b_0(t) \sum_{j=1}^{N} X_t^j + b_1(t) X_t^i + b_2(t)\beta_t^i + \sigma dW_t^i, \quad X_0^i = \xi_0^i,$$

and define the common cost

$$J^N(\beta) = \mathbb{E}\left[\int_0^T f(s, X_s^i, \bar{\mu}_s^N, \beta_s^i)\, ds + g(X_T^1, \bar{\mu}_T^N)\right], \quad \text{with } \bar{\mu}_t^N = \frac{1}{N}\sum_{i=1}^{N} \delta_{X_t^i}.$$

Then, there exists a sequence $(\epsilon_N)_{N\geq 1}$, $\epsilon_N \searrow 0$, s.t. **for all** $\boldsymbol{\beta} = (\boldsymbol{\beta}^1, \ldots, \boldsymbol{\beta}^N)$,

$$J^N(\boldsymbol{\beta}) \geq J^N(\boldsymbol{\alpha}) - \epsilon_N,$$

where, $\boldsymbol{\alpha} = (\boldsymbol{\alpha}^1, \cdots, \boldsymbol{\alpha}^N)$ with

$$\alpha_t^i = \hat{\alpha}(s, \tilde{X}_t^i, u(t, \tilde{X}_t^i), \mathbb{P}_{X_t})$$

where X and u are from the solution to the **controlled McKean Vlasov problem**, and $(\tilde{X}^1, \ldots, \tilde{X}^N)$ is the state of the system controlled by α, i.e.

$$d\tilde{X}_t^i = \frac{1}{N}\sum_{j=1}^{N} b_0(t)\tilde{X}_t^j + b_1(t)\tilde{X}_t^i + b_2(t)\hat{\alpha}(s, \tilde{X}_s^i, u(s, \tilde{X}_s^i), \mathbb{P}_{X_s}) + \sigma dW_t^i, \quad \tilde{X}_0^i = \xi_0^i.$$

Application #2: Chain Rule

Assume

$$dX_t = b_t dt + \sigma_t dW_t, \quad X_0 \in L^2(\Omega, \mathcal{F}, \mathbb{P}),$$

where

- $\mathbf{W} = (W_t)_{t \geq 0}$ is a $\mathbb{F}$-Brownian motion with values in $\mathbb{R}^d$
- $(b_t)_{t \geq 0}$ and $(\sigma_t)_{t \geq 0}$ are $\mathbb{F}$-progressive processes in $\mathbb{R}^d$ and $\mathbb{R}^{d \times d}$
- Assume

$$\forall T > 0, \quad \mathbb{E}\left[\int_0^T \left(|b_t|^2 + |\sigma_t|^4\right) dt\right] < +\infty.$$

Then for any $t \geq 0$, if $\mu_t = \mathbb{P}_{X_t}$, and $a_t = \sigma_t \sigma_t^\dagger$ then:

$$u(\mu_t) = u(\mu_0) + \int_0^t \mathbb{E}\big[\partial_\mu u(\mu_s)(X_s) \cdot b_s\big] ds + \frac{1}{2} \int_0^t \mathbb{E}\big[\partial_v \big(\partial_\mu u(\mu_s)\big)(X_s) \cdot a_s\big] ds.$$

Control of McKean-Vlasov SDEs: Verification Theorem

Problem: if $f : \mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}$, minimize

$$J(\boldsymbol{\alpha}) = \int_0^T f\left(\mathbb{P}_{X_t^{\boldsymbol{\alpha}}}\right) dt + \mathbb{E}\left[\int_0^T \frac{1}{2}|\alpha_t|^2 dt\right]$$

under the constraint:

$$dX_t^{\boldsymbol{\alpha}} = \alpha_t dt + dW_t, \quad 0 \le t \le T,$$

Verification Argument: Assume $u : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ is $\mathcal{C}^{1,2}$, and satisfies

$$\partial_t u(t,\mu) - \frac{1}{2}\int_{\mathbb{R}^d} |\partial_\mu u(t,\mu)(v)|^2 d\mu(v) + \frac{1}{2}\mathrm{trace}\left[\int_{\mathbb{R}^d} \partial_v \partial_\mu u(t,\mu)(v) d\mu(v)\right] + f(\mu) = 0,$$

then, the McKean-Vlasov SDE

$$d\hat{X}_t = -\partial_\mu u(t, \mathbb{P}_{\hat{X}_t})(\hat{X}_t) dt + dW_t, \quad 0 \le t \le T,$$

has a unique solution $(\hat{X}_t)_{0 \le t \le T}$ satisfying $\mathbb{E}[\sup_{0 \le t \le T} |\hat{X}_t|^2] < \infty$ which is the unique optimal path since $\hat{\alpha}_t = -\partial_\mu u(t, \mathbb{P}_{\hat{X}_t})(\hat{X}_t)$ minimizes the cost:

$$J(\hat{\boldsymbol{\alpha}}) = \inf_{\boldsymbol{\alpha} \in \mathbb{A}} J(\boldsymbol{\alpha}).$$

PROOF (SKETCH OF)

For a generic admissible control $\boldsymbol{\alpha} = (\alpha_t)_{0\le t\le T}$, set $X_t^{\boldsymbol{\alpha}} = X_0 + \int_0^T \alpha_s ds + W_t$ and apply the chain rule:

$$
\begin{aligned}
& du(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}}) \\
&= \Big[\partial_t u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}}) + \mathbb{E}\big[\partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}}) \cdot \alpha_t \big] + \frac{1}{2}\mathbb{E}\big[\text{trace}\big[\partial_v \partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}}) \big] \big] \Big] dt \\
&= \Big[-f(\mathbb{P}_{X_t^{\boldsymbol{\alpha}}}) + \frac{1}{2}\mathbb{E}\big[|\partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}})|^2 \big] + \mathbb{E}\big[\partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}}) \cdot \alpha_t \big] \Big] dt \\
&= \Big[-f(\mathbb{P}_{X_t^{\boldsymbol{\alpha}}}) - \frac{1}{2}\mathbb{E}[|\alpha_t|^2] + \frac{1}{2}\mathbb{E}\big[|\alpha_t + \partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}})|^2 \big] \Big] dt
\end{aligned}
$$

where we used the PDE satisfied by u, and identified a *perfect square*. Integrate both sides and get:

$$
J(\boldsymbol{\alpha}) = u(0, \mathbb{P}_{X_0}) + \frac{1}{2}\mathbb{E}\Big[\int_0^T \big[|\alpha_t + \partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}})|^2 \big] dt \Big],
$$

which shows that $\alpha_t = -\partial_\mu u(t, \mathbb{P}_{X_t^{\boldsymbol{\alpha}}})(X_t^{\boldsymbol{\alpha}})$ is **optimal**.

Joint Chain Ruile

- If u is smooth
- If $d\xi_t = \eta_t dt + \gamma_t dW_t$
- If $dX_t = b_t dt + \sigma_t dW_t$ and $\mu_t = \mathbb{P}_{X_t}$

$$
\begin{aligned}
u(t,\xi_t,\mu_t) &= u(0,\xi_0,\mu_0) + \int_0^t \partial_x u(s,\xi_s,\mu_s)\cdot(\gamma_s dW_s) \\
&+ \int_0^t \Big(\partial_t u(s,\xi_s,\mu_s) + \partial_x u(s,\xi_s,\mu_s)\cdot \eta_s + \frac{1}{2}\mathrm{trace}\big[\partial^2_{xx} u(s,\xi_s,\mu_s)\gamma_s\gamma_s^\dagger\big]\Big) ds \\
&+ \int_0^t \tilde{\mathbb{E}}\big[\partial_\mu u(s,\xi_s,\mu_s)(\tilde{X}_s)\cdot \tilde{b}_s\big] ds + \frac{1}{2}\int_0^t \tilde{\mathbb{E}}\big[\mathrm{trace}\big(\partial_v\big[\partial_\mu u(s,\xi_s,\mu_s)\big](\tilde{X}_s)\tilde{\sigma}_s\tilde{\sigma}_s^\dagger\big)\big] ds
\end{aligned}
$$

where the process $(\tilde{X}_t, \tilde{b}_t, \tilde{\sigma}_t)_{0\le t\le T}$ is an **independent copy** of the process $(X_t, b_t, \sigma_t)_{0\le t\le T}$, on a different probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$

Deriving the Master Equation

If $(t, x, \mu) \hookrightarrow \mathcal{U}(t, x, \mu)$ is the master field

$$\left(\mathcal{U}(t, X_t, \mu_t) - \int_0^t f\big(s, X_s, \mu_s, \hat{\alpha}(s, X_s, \mu_s, Y_s)\big)\, ds\right)_{0 \leq t \leq T}$$

is a martingale whenever $(X_t, Y_t, Z_t)_{0 \leq t \leq T}$ is the solution of the forward-backward system characterizing the optimal path under the flow of measures $(\mu_t)_{0 \leq t \leq T}$. So if we compute its Itô differential, the drift must be 0

An Example of Derivation

$$dX_t = b(t, X_t, \mu_t, \alpha_t)dt + dW_t$$
$$H(t, x, \mu, y, \alpha) = b(t, x, \mu, \alpha) \cdot y + f(t, x, \mu, \alpha)$$
$$\hat{\alpha}(t, x, \mu, y) = \arg\inf_{\alpha} H(t, x, \mu, y, \alpha)$$

Itô's Formula with $\mu_t = \mathbb{P}_{X_t}$

(set $\hat{\alpha}_t = \hat{\alpha}(t, X_t, \mu_t, \partial U(t, X_t, \mu_t))$ and $b_t = b(t, X_t, \mu_t, \hat{\alpha}_t)$)

$$d\mathcal{U}(t, X_t, \mu_t) =$$
$$\left(\partial_t \mathcal{U}(t, X_t, \mu_t) + b_t \cdot \partial_x \mathcal{U}(t, X_t, \mu_t) + \frac{1}{2}\mathrm{trace}[\partial^2_{xx}\mathcal{U}(t, X_t, \mu_t)] + f(t, x, \mu, \hat{\alpha}_t)\right)dt$$
$$+ \mathbb{E}\left[b_t \cdot \partial_\mu \mathcal{U}(t, X_t, \mu_t)(X_t) + \frac{1}{2}\partial_v \partial_\mu \mathcal{U}(t, X_t, \mu_t)](X_t)]\right] dt + \partial_x \mathcal{U}(t, X_t, \mu_t) dW_t$$

THE ACTUAL MASTER EQUATION

$$
\begin{aligned}
&\partial_t \mathcal{U}(t,x,\mu) + b\big(t,x,\mu,\hat{\alpha}(t,x,\mu,\partial\mathcal{U}(t,x,\mu))\big)\cdot\partial_x\mathcal{U}(t,x,\mu) \\
&\quad + \frac{1}{2}\mathrm{trace}\Big[\partial^2_{xx}\mathcal{U}(t,x,\mu)\Big] + f\big(t,x,\mu,\hat{\alpha}(t,x,\mu,\partial\mathcal{U}(t,x,\mu))\big) \\
&\quad + \int_{\mathbb{R}^d}\Big[b\big(t,x',\mu,\hat{\alpha}(t,x,\mu,\partial\mathcal{U}(t,x,\mu))\big)\cdot\partial_\mu\mathcal{U}(t,x,\mu)(x') \\
&\qquad\qquad + \frac{1}{2}\mathrm{trace}\Big(\partial_v\partial_\mu\mathcal{U}(t,x,\mu)(x')\Big)\Big]d\mu(x') = 0,
\end{aligned}
$$

for $(t,x,\mu)\in[0,T]\times\mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d)$, with the **terminal** condition $V(T,x,\mu)=g(x,\mu)$.