Probabilistic Approach to Mean Field Games and the Control of McKean-Vlasov Dynamics

René Carmona

Bendheim Center for Finance
Department of Operations Research & Financial Engineering
Princeton University

Freddy's Festschrift, September 28, 2012

Joint work with

François Delarue

Université de Nice

- (with F. Delarue and A. Lachapelle) Control of McKean-Vlasov Dynamics versus Mean Field Games. *MAFE* (2012) (to appear).
- (with F. Delarue) Probabilistic Analysis of Mean Field Games. *submitted for publication*
- (with F. Delarue) Control of McKean Vlasov Dynamics *in preparation*
- (with F. Delarue) FBSDEs of McKean-Vlasov Type I. Existence *in preparation*

Simple Example from Systemic Risk (J.P. Fouque)

- Log-monetary reserves of N banks

$$X_t^{(i)}, i = 1, \ldots, N$$

- $W_t^{(i)}, i = 1, \ldots, N$ **independent Brownian motions**, $\sigma > 0$
- Model **borrowing and lending** through the drifts:

$$\begin{aligned} dX_t^{(i)} &= \frac{\alpha}{N} \sum_{j=1}^{N} (X_t^{(j)} - X_t^{(i)})\, dt + \sigma dW_t^{(i)} \\ &= \alpha(\overline{X}_t - X_t^{(i)})\, dt + \sigma dW_t^{(i)}\,, \quad i = 1, \ldots, N. \end{aligned}$$

- OU processes reverting to the **sample mean** $\overline{X}_t$ (rate $\alpha > 0$)
- $D < 0$ **default level**

Easy Conclusions

- Sample mean $\overline{X}_t$ is a BM a Brownian motion with vol. $\sigma/\sqrt{N}$
- Simulations “show" that **STABILITY** is created by increasing the rate α of borrowing and lending.
- Compute the loss distribution (how many firms fail)
- Large Deviations (Gaussian estimates) show that increasing α increases **SYSTEMIC RiSK**

Modified Model

New dynamics is

$$dX_t^i = \left[a(\overline{X}_t - X_t^i) + \alpha_t^i\right] dt + \sigma dW_t^i, \quad i = 1, \cdots, N$$

α^i is the control of bank i, and bank i tries to **minimize**

$$J^i(\alpha^1, \cdots, \alpha^N) = \mathbb{E}\left\{ \int_0^T \left[\frac{1}{2q}(\alpha_t^i)^2 - \alpha_t^i(\overline{X}_t - X_t^i) \right] dt \right\}$$

The regulator can choose the parameter $q > 0$ controlling the cost of borrowing and lending.

- If X_t^i is small (relative to the empirical mean $\overline{X}_t$) then bank i will want to borrow($\alpha_t^i > 0$)
- If X_t^i is large then bank i will want to lend ($\alpha_t^i < 0$)

Example of **Mean Field Game (MFG)** *à la Lasry - Lions*

APPROXIMATE NASH MFG-EQUILIBRIUM

◇ Banks act **independently** of each other

◇ Bank i chooses $\alpha_t^i = q(\overline{X}_t - X_t^i) - \eta_t X_t^i$

$$dX_t^i = \left[(a+q)(\overline{X}_t - X_t^i) - \eta_t X_t^i\right] dt + \sigma dW_t^i$$

for a deterministic function $t \hookrightarrow \eta_t$ solving a Ricatti equation.Therefore

$$d\overline{X}_t = -\eta_t \overline{X}_t dt + \frac{\sigma}{\sqrt{N}} d\overline{W}_t$$

where $\overline{W}_t = \frac{1}{\sqrt{N}} \sum_1^N dW_t^i$.

- **Note** that $\eta_t < 0$, and therefore $(\overline{X}_t)$ is a **repulsive OU**.
- Still **Gaussian** system, so similar **Large Deviation** estimates

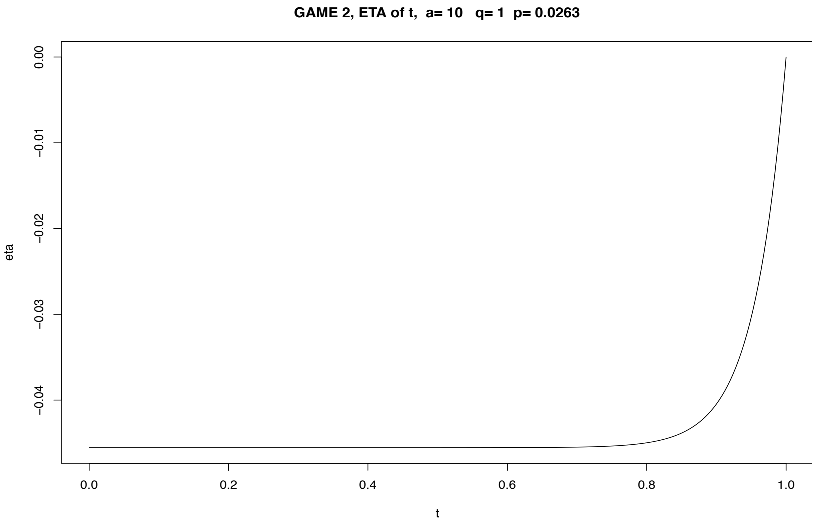

GAME 2, ETA of t, a= 10 q= 1 p= 0.0263
eta
t
0.00
-0.01
-0.02
-0.03
-0.04
0.0
0.2
0.4
0.6
0.8
1.0

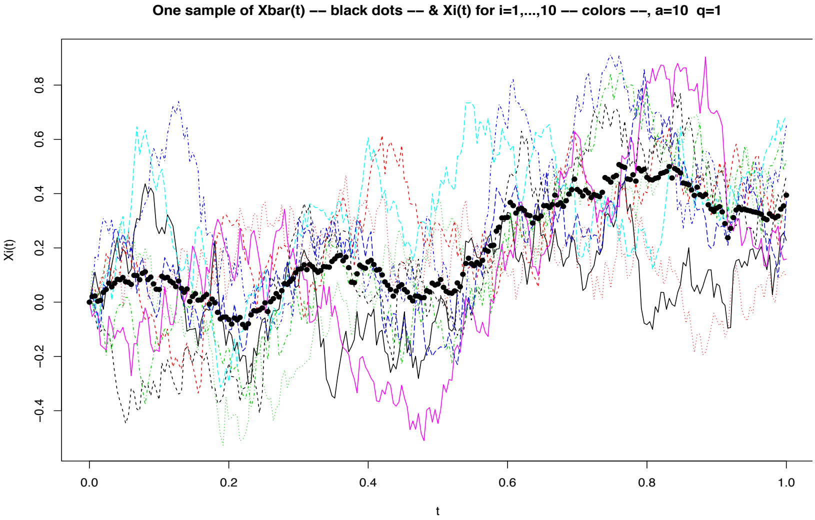
One sample of Xbar(t) -- black dots -- & Xi(t) for i=1,...,10 -- colors --, a=10 q=1
Xi(t)
-0.4
-0.2
0.0
0.2
0.4
0.6
0.8
0.0
0.2
0.4
0.6
0.8
1.0
t

One sample of Xbar(t) -- black dots -- & Xi(t) for i=1,...,10 -- colors --, a=100 q=1

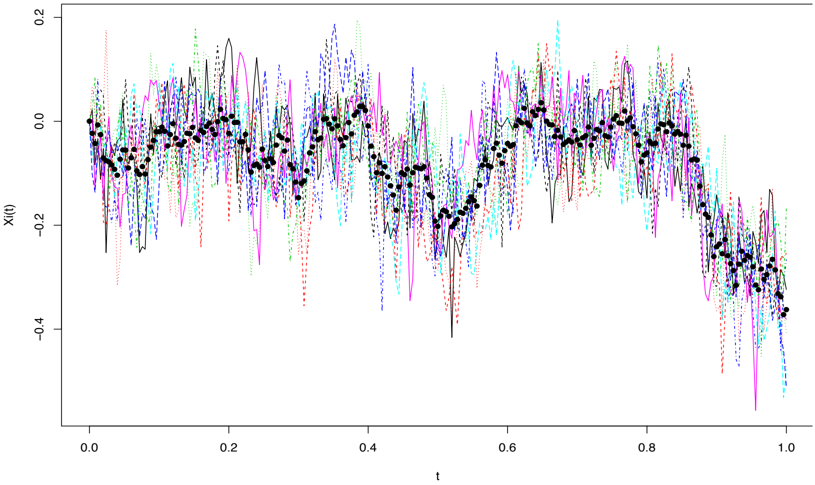

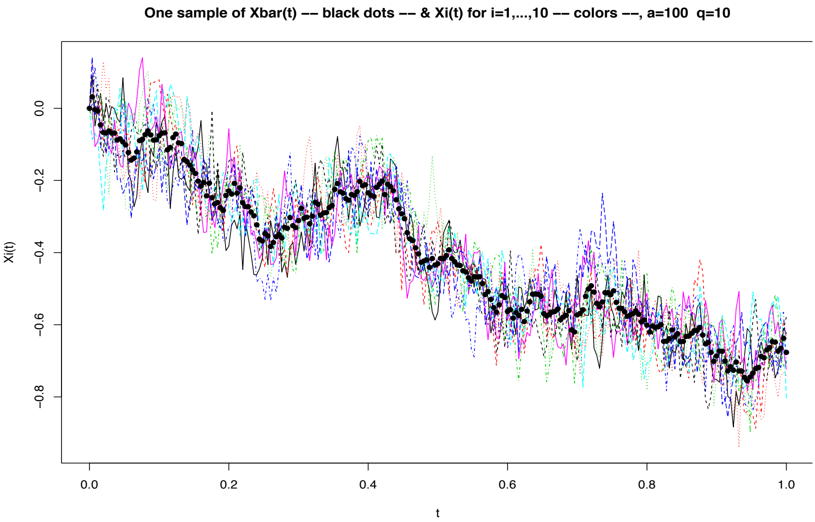

One sample of Xbar(t) -- black dots -- & Xi(t) for i=1,...,10 -- colors --, a=100 q=10
Xi(t)
0.0
−0.2
−0.4
−0.6
−0.8
0.0
0.2
0.4
0.6
0.8
1.0
t

High Dimensional Dynamical Systems

- Dynamic: equations will be ODEs, PDEs, SDEs, SPDEs, . . .
- High Dimensional (large populations)
- **Equilibrium Analysis**: Control or Game Theory
  - Each individual makes decisions based on
    - his/her current state
    - Risk / Reward expectation
  - Interaction with the rest of the population
    - Distribution of states in the entire population
  - **MEAN FIELD** interaction

The PDE Approach to MFGs

Motivation (**Lasry-Lions, Guéant, La Chapelle, . . .**)

$$u(t,x) = \sup_{(\alpha_s)_{t \le s \le T}, X_t = x} \mathbb{E}\left[\int_t^T e^{-\rho(s-t)}[g(m(s,X_s)) + h(|\alpha(s,X_s)|)]ds\right]$$

under constraint $dX_t = \alpha(t, X_t)dt + \sigma dW_t$

Formulation (given $m(0, \cdot)$ & $u(T, \cdot)$)

$$\partial_t u + \frac{\sigma^2}{2}\Delta u + H(\nabla u) - \rho u = -g(m) \qquad \text{(Hamilton-Jacobi-Bellman)}$$

$$\partial_t m + \nabla \cdot (mH'(\nabla u)) = \frac{\sigma^2}{2}\Delta m, \qquad \text{(Kolmogorov)}$$

where $m(t, \cdot)$ probability measure, $H(p) = \sup_a(ap - h(a))$.

Stationary Case

$$\frac{\sigma^2}{2}\Delta u + H(\nabla u) - \rho u = -g(m)$$

$$\nabla \cdot (mH'(\nabla u)) = \frac{\sigma^2}{2}\Delta m,$$

Probabilistic Approach

Disclaimer

"Mathematicians are like Frenchmen: whatever you say to them they translate into their own language and forthwith it is something entirely different."

Johann Wolfgang von Goethe

- Formulate the problem for (N **finite**) players with (**mean - field interactions**)
- Define the **Optimization Problem**: Cost / Reward functions
- Define the type of desired **equilibrium** (Nash, Pareto. centralized, etc)

Stochastic Differential Game or **Stochastic Control Problem**

- **Identify the limit** $N \to \infty$ to Lasry-Lions MFG problem, **or something else !**
- **Prove** that solution of $N = \infty$ (e.g. MFG problem) provides approximate (ϵ) equilibria for problems if (N) finite many players

Stochastic Differential Mean Field Games

First example of **private states** dynamics

$$dX_t^i = \frac{1}{N}\sum_{j=1}^{N} \tilde{b}(t, X_t^i, X_t^j, \alpha_t^i)dt + \sigma dW_t^i$$

Rewrite it as

$$dX_t^i = b(t, X_t^i, \overline{\mu}_t^N, \alpha_t^i)dt + \sigma dW_t^i$$

where

$$b(t, x, \mu, \alpha) = \int \tilde{b}(t, x, x', \alpha)\, d\mu(x')$$

and $\overline{\mu}_t^N$ is **empirical distribution** of the private states

$$\overline{\mu}_t^N = \frac{1}{N}\sum_{j=1}^{N} \delta_{X_t^j}.$$

linear interaction or *of order* 1

More General Mean Field Interactions

Quadratic interactions or *of order* 2

$$\frac{1}{N^2}\sum_{j,k=1}^{N}\tilde{b}(t, X_t^i, X_t^j, X_t^k, \alpha_t^i)dt + \sigma dW_t^i$$

rewritten as $b(t, X_t^i, \overline{\mu}_t^N, \alpha_t^i)$ with

$$b(t, x, \mu, \alpha) = \int \tilde{b}(t, x, x', x'', \alpha)\, d\mu(x')d\mu(x'').$$

Fully nonlinear Mean Field interaction

$$b : [0, T] \times \mathbb{R} \times \mathcal{P}(\mathbb{R}) \times A \hookrightarrow \mathbb{R}$$

Scalar Mean Field interaction

$$b(t, x, \mu, \alpha) = \int \tilde{b}(t, x, \langle\varphi, \mu\rangle, \alpha)$$

for some scalar function φ with $\langle\varphi, \mu\rangle = \int \varphi(x')d\mu(x')$

Optimization Problem

Simultaneous Minimization of

$$J^i(a) = \mathbb{E}\left\{\int_0^T f(t, X_t^i, \overline{\mu}_t^N, \alpha_t^i)dt + g(X_T, \overline{\mu}_T^N)\right\}, \quad i = 1, \cdots, N$$

under **constraints** of the form

$$dX_t^i = b(t, X_t^i, \overline{\mu}_t^N, \alpha_t^i)dt + \sigma dW_t^i, \quad i = 1, \cdots, N.$$

GOAL: search for **equilibriums**

Model Requirements

- Each player **cannot** on its own, influence **significantly** the global output of the game
- **Large** number of players ($N \to \infty$)
- **Closed loop** controls in **feedback** form

$$\alpha_t^i = \phi^i(t, (X_t^1, \cdots, X_t^N)), \qquad i = 1, \cdots, N.$$

- **Distributed** controls

$$\alpha_t^i = \phi^i(t, X_t^i), \qquad i = 1, \cdots, N.$$

- **Identical** feedback functions

$$\phi^1(t, \cdot\,) = \cdots = \phi^N(t, \cdot\,) = \phi(t, \cdot\,), \qquad 0 \leq t \leq T.$$

Touted Solution (Wishful Thinking)

- **Identify** a (distributed closed loop) **strategy** ϕ from **effective equations** (from stochastic optimization for large populations)
- Each player is assigned the same function ϕ
- At each time t, player i take action $\alpha_i = \phi(t, X_t^i)$

What is the resulting **population behavior**?

- Did we reach some form of equilibrium?
- If yes, what kind of equilibrium?

Nash Equilibrium: Optimizing First

$$\alpha_t^{1*} = \phi^{1*}(t, X_t^1), \cdots\cdots, \alpha_t^{N*} = \phi^{N*}(t, X_t^N)$$

is a **Nash equilibrium** means that for each player i, if we assume

$$\alpha^{-i*} \quad = \quad \alpha_t^{1*}, \cdots, \alpha^{i-1*}, \alpha_t^{i+1*}, \cdots, \alpha_t^{N*}$$

are **FIXED**, then:

$$\phi^{i*} = \arg\inf_{\phi} \mathbb{E}\left\{\int_0^T f(t, X_t^i, \overline{\mu}_t^N, \phi(t, X_t^i))dt + g(X_T, \overline{\mu}_T^N)\right\}.$$

When N is large small perturbations of ϕ

should not change empirical measure $\overline{\mu}_t^N$

So one could solve the optimization problem (approximate its solution)

FREEZING $(\overline{\mu}_t^N)_{0\leq t\leq T}$

Standard stochastic control problem (parameterized by $(\mu_t)_{0\leq t\leq T}$):
Once ϕ is found, μ_t should be the statistical distribution of the solution X_t !

Summary of the Lasry-Lions MFG Approach

1. Fix a deterministic function $[0, T] \ni t \hookrightarrow \mu_t \in \mathcal{P}(\mathbb{R})$;
2. Solve the standard stochastic control problem

$$\phi^* = \arg\inf_{\phi} \mathbb{E}\left\{\int_0^T f(t, X_t, \mu_t, \phi(t, X_t))dt + g(X_T, \mu_T)\right\}$$

   subject to

$$dX_t = b(t, X_t, \mu_t, \phi(t, X_t))dt + \sigma dW_t;$$

3. Determine the function $[0, T] \ni t \hookrightarrow \mu_t \in \mathcal{P}(\mathbb{R})$ so that

$$\forall t \in [0, T], \quad \mathbb{P}_{X_t} = \mu_t.$$

Once this is done,

$$\alpha_t^{j*} = \phi^*(t, X_t^j), \qquad j = 1, \cdots, N$$

form an **approximate Nash equilibrium** for the game.

MFG Adjoint Equations

Freeze $\mu = (\mu_t)_{0 \le t \le T}$, write Hamiltonian

$$H^{\mu_t}(t, x, y, \alpha) = b(t, x, \mu_t, \alpha) \cdot y + f(t, x, \mu_t, \alpha)$$

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

$$\begin{cases} dY_t = -\partial_x H^{\mu_t}(t, X_t^{\alpha}, Y_t, \alpha_t) dt + Z_t dW_t \\ Y_T = \partial_x g(X_T^{\alpha}, \mu_T) \end{cases}$$

is called a set of **adjoint processes**

Stochastic Minimum Principle (Pontryagin)

Determine

$$\hat{\alpha}^{\mu_t}(t,x,y) = \arg\inf_{\alpha\in A} H^{\mu_t}(t,x,y,\alpha)$$

Inject in **FORWARD** and **BACKWARD** dynamics and **SOLVE**

$$\begin{cases} dX_t = b(t, X_t, \mu_t, \hat{\alpha}^{\mu}(t, X_t, Y_t))dt + \sigma dW_t, & X_0 = x_0 \\ dY_t = -\partial_x H^{\mu_t}(t, X_t, Y_t, \hat{\alpha}^{\mu_t}(t, X_t, Y_t))dt + Z_t dW_t, & Y_T = \partial_x g(X_T, \mu_t) \end{cases}$$

Standard **FBSDE** (for each fixed $t \hookrightarrow \mu_t$)

Fixed Point Step

Solve the **fixed point problem**

$$(\mu_t)_{0\le t\le T} \quad \longrightarrow \quad (X_t)_{0\le t\le T} \quad \longrightarrow \quad (\mathbb{P}_{X_t})_{0\le t\le T}$$

Note: if we enforce $\mu_t = \mathbb{P}_{X_t}$ for all $0 \le t \le T$ in FBSDE we have

$$\begin{cases} dX_t = b(t, X_t, \mathbb{P}_{X_t}, \hat{\alpha}^{\mathbb{P}}_{X_t}(t, X_t, Y_t))dt + \sigma dW_t, & X_0 = x_0 \\ dY_t = -\partial_x H^{\mathbb{P}}_{X_t}(t, X^{\alpha}_t, Y_t, \hat{\alpha}^{\mathbb{P}}_{X_t}(t, X_t, Y_t))dt + Z_t dW_t, & Y_T = \partial_x g(X_T, \mathbb{P}_{X_T}) \end{cases}$$

FBSDE of McKean-Vlasov type !!!

Solvability of Forward Backard Systems

Existence of a solution of

$$\begin{cases} dX_t = b(t, X_t, Y_t, \mathbb{P}_{(X_t,Y_t)})dt + \sigma(t, X_t, Y_t, \mathbb{P}_{(X_t,Y_t)})dW_t \\ dY_t = -\Psi(t, X_t, Y_t, \mathbb{P}_{(X_t,Y_t)})dt + Z_t dW_t \\ X_0 = x, \; Y_T = g(X_T, \mathbb{P}_{X_T}) \end{cases}$$

if coefficients are **uniformly Lipschitz** and **bounded**

boundedness assumption **can be** relaxed

e.g. MFG and Controlled McKean-Vlasov models (later on in the talk)

Proof works for $\mathbb{P}_{(X_t,Y_t,Z_t)}$ instead of $\mathbb{P}_{(X_t,Y_t)}$

Assumptions

- Convex costs (f and g)
- Uncontrolled volatility ($\sigma(t, x, \mu, \alpha) \equiv \sigma > 0$)
- $b(t, x, \mu, \alpha) = b_0(t, \mu) + b_1(t)x + b_2(t)\alpha$ with bounded b_i's

Then

$$\hat{\alpha}(t, x, y, \mu) \in \arg\inf_{\alpha} H^{\mu}(t, x, y, \alpha)$$

is Lip-1 in (x, y, μ) uniformly in $t \in [0, T]$ and one can solve:

$$\begin{cases} dX_t = b(t, X_t, Y_t, \mathbb{P}_{X_t}, \hat{\alpha}(t, X_t, Y_t, \mathbb{P}_{X_t}))dt + \sigma dW_t \\ dY_t = -\partial_x f(t, X_t, Y_t, \mathbb{P}_{X_t}, \hat{\alpha}(t, X_t, Y_t, \mathbb{P}_{X_t}))dt - b_1(t) Y_t + Z_t dW_t \\ X_0 = x, \; Y_T = \partial_x g(X_T, \mathbb{P}_{X_T}) \end{cases}$$

and the solution is of the form

$$Y_t = u(t, X_t)$$

Back to the N-Player (Mean Field) Game

Assume:

$$dX_t^i = b(t, X_t^i, \overline{\mu}_t^N, \alpha_t^i)dt + \sigma dW_t^i, \qquad 0 \leq t \leq T, \quad 1 \leq i \leq N$$

where

$$\overline{\mu}_t^N = \frac{1}{N}\sum_{i=1}^{N} \delta_{X_t^i}.$$

Then the controls

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

form an ϵ-Nash equilibrium in the sense that for some $\epsilon_N \searrow 0$, for each $1 \leq i \leq N$

$$J(\hat{\alpha}_t^1, \cdots, \alpha_t^i, \cdots, \hat{\alpha}_t^N) \geq J(\hat{\alpha}_t^1, \cdots, \hat{\alpha}_t^i, \cdots, \hat{\alpha}_t^N) - \epsilon_N$$

Franchise Equilibrium

We say that $(t,x) \hookrightarrow \phi^*(t,x)$ gives a **franchise equilibrium** if

$$\phi^* = \arg\inf_{\phi} \mathbb{E}\left\{\int_0^T f(t, X_t^i, \overline{\mu}_t^N, \phi(t, X_t^i))dt + g(X_T, \overline{\mu}_T^N)\right\}.$$

where for each player $i \in \{1, \cdots, N\}$ we have $\alpha_t^i = \phi(t, X_t^i)$.

So when one player perturbs his/her ϕ

ALL players perturb their ϕ's in the same way!

So the streamlining procedure is

1. Take the limit $N \to \infty$ (i.e. solve the **fixed point problem**) **FIRST**
2. Solve the optimization problem **NEXT**

Taking the Limit $N \to \infty$ First

Propagation of Chaos
(**Mc Kean / Sznitmann / Jourdain-Méleard-Woyczinski**)

- Focus on N_0 (fixed) player in a large set ($N \to \infty$) of players
- Their private state processes X_t^j for $j = 1, \cdots, N_0$ become
  - (Asymptotically) **independent identically distributed**
  - (Asymptotically) **distributed** like the solution of (McKV)

$$dX_t = b(t, X_t, \mathbb{P}_{X_t}, \phi(t, X_t))dt + \sigma d\tilde{W}_t$$

The individual objective costs become

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

Control of McKean-Vlasov Dynamics

Stochastic optimization problem: minimize

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

over admissible control processes $\underline{\alpha} = (\alpha_t)_{0\le t\le T}$ subject to

$$dX_t = b(t, X_t, \mathbb{P}_{X_t}, \alpha_t)dt + \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)dW_t \qquad 0 \le t \le T,$$

Probabilistic approach based on **Pontryagin maximum principle**

Hamiltonian

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

(Informal) Natural Question

Is the diagram

SDE State Dynamics for N players	$\xrightarrow{\text{Optimization}}$	Nash Equilibrium for N players
$\downarrow N \to \infty$		$\downarrow N \to \infty$
McKean Vlasov Dynamics	$\xrightarrow{\text{Optimization}}$	Mean Field Game? Controlled McK-V Dynamics?

commutative?

DIFFERENTIABILITY AND CONVEXITY OF $\mu \hookrightarrow h(\mu)$

- Notions of differentiability for functions defined on spaces of measures from theory of optimal transportation, gradient flows, etc) studied by **Ambrosio, De Giorgi, Otto, Villani**, etc
- Tailored made notion (**Lions**' Collège de France Lectures, **Cardaliaguet**)

Lift a function $\mu \hookrightarrow h(\mu)$ to $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ into

$$X \hookrightarrow \tilde{h}(X) = h(\tilde{\mathbb{P}}_X)$$

and say

h is differentiable at μ if $\tilde{h}$ is Fréchet differentiable at X whenever $\tilde{\mathbb{P}}_X = \mu$.

A function g on $\mathbb{R}^d \times \mathcal{P}_1(\mathbb{R}^d)$ is said to be **convex** if for every (x, μ) and (x', μ') in $\mathbb{R}^d \times \mathcal{P}_1(\mathbb{R}^d)$ we have

$$g(x', \mu') - g(x, \mu) - \partial_x g(x, \mu) \cdot (x' - x) - \tilde{\mathbb{E}}[\partial_\mu g(x, \tilde{X}) \cdot (\tilde{X}' - \tilde{X})] \geq 0$$

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

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

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

where $(\tilde{\alpha}, \tilde{X}, \tilde{Y}, \tilde{Z})$ is an independent copy of $(\alpha, X^{\alpha}, Y, Z)$, 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 !!!

A Necessary Condition for Optimality

If $\underline{X} = X^{\underline{\alpha}}$ controlled McKean-Vlasov dynamics ($X_0 = x$), compute the **Gâteaux derivative of the cost functional** J at $\underline{\alpha}$ in the direction of $\underline{\beta}$ using dual processes and the variation process $\underline{V} = (V_t)_{0 \le t \le T}$ solution of the equation

$$dV_t = [\gamma_t V_t + \delta_t(\mathbb{P}_{(X_t,V_t)}) + \eta_t]dt + [\tilde{\gamma}_t V_t + \tilde{\delta}_t(\mathbb{P}_{(X_t,V_t)}) + \tilde{\eta}_t]dW_t$$

where the coefficients γ_t, δ_t, η_t, $\tilde{\gamma}_t$, $\tilde{\delta}_t$ and $\tilde{\eta}_t$ are defined as

$$\gamma_t = \partial_x b(t, X_t, \mathbb{P}_{X_t}, \alpha_t), \quad \text{and} \quad \tilde{\gamma}_t = \partial_x \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)$$
$$\eta_t = \partial \alpha b(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\beta_t, \quad \text{and} \quad \tilde{\eta}_t = \partial_\alpha \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\beta_t$$
$$\gamma_t = \partial_x b(t, X_t, \mathbb{P}_{X_t}, \alpha_t), \quad \text{and} \quad \tilde{\gamma}_t = \partial_x \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)$$

and

$$\delta_t = \tilde{\mathbb{E}}\partial_\mu b(t, x, \mathbb{P}_{X_t}, \alpha)(\tilde{X}_t) \cdot \tilde{V}_t\Big|_{\substack{x=X_t \\ \alpha=\alpha_t}}, \quad \text{and} \quad \tilde{\delta}_t = \tilde{\mathbb{E}}\partial_\mu \sigma(t, x, \mathbb{P}_{X_t}, \alpha)(\tilde{X}_t) \cdot \tilde{V}_t\Big|_{\substack{x=X_t \\ \alpha=\alpha_t}}$$

where $(\tilde{X}_t, \tilde{V}_t)$ is an independent copy of (X_t, V_t).

Pontryagin Minimum Principle (Sufficiency)

Assume

1. Coefficients continuously differentiable with bounded derivatives;
2. Terminal cost function g is convex;
3. α admissible control, X corresponding dynamics, (Y, Z) 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.

$$J(\alpha) = \inf_{\overline{\alpha} \in \mathcal{A}} J(\overline{\alpha})$$

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.

There exists a solution $(X_t, Y_t, Z_t)_0$ of the McKean-Vlasov FBSDE

$$\begin{cases} dX_t = b_0(t)\mathbb{E}(X_t)dt + b_1(t)X_t dt + b_2(t)\hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t)dt + \sigma dW_t, \\ dY_t = -\partial_x H\big(t, X_t, \mathbb{P}_{X_t}, Y_t, \hat{\alpha}_t\big)dt \\ \qquad\qquad - \mathbb{E}\big[\partial_\mu \underline{H}(t, X_t', X_t, Y_t', \hat{\alpha}_t')\big]dt + Z_t dW_t. \end{cases}$$

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 $(W^1, \dots, W^N)$ and for a square integrable exchangeable process $(\beta^1, \dots, \beta^N)$, consider the system of particles

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

and define the common cost

$$J^N(\beta) = \mathbb{E}\Big[+ \int_0^T f(s, X_s^i, \bar\mu_s^N, \beta_s^i) ds + g(X_T^1, \bar\mu_T^N) \Big], \quad \text{with } \bar\mu_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}.$$

Then, there exists a sequence of positive reals $(\epsilon_N)_{N \geq 1}$, independent of β and converging toward 0, such that

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

where, if X is the solution to the **controlled McKean Vlasov problem**, $(\tilde X^1, \dots, \tilde X^N)$ solves

$$d\tilde X_t^i = \frac{1}{N} \sum_{j=1}^N b_0(t) \tilde X_t^j + b_1(t) \tilde X_t^i + b_2(t) \hat\alpha(s, \tilde X_s^i, u(s, \tilde X_s^i), \mathbb{P}_{X_s}) + \sigma dW_t^i, \quad \tilde X_0^i = \xi_0^i.$$

Merci, et Bon Anniversaire Freddy!