Stochastic maximum principle for optimal control of non exchangeable mean field systems.

Samy Mekkaoui

CMAP, École Polytechnique and LPSM

Advances in Financial Mathematics, Paris

Joint works with

Idris Kharroubi (LPSM, Sorbonne Université)
Huyên Pham (CMAP, École Polytechnique)
Xavier Warin (EDF)

27 January 2026

Motivations for non exchangeable mean field systems

- MFC theory : Interactions through symmetric particles and homogeneous interactions, through empirical measure $\frac{1}{N}\sum_{j=1}^{N}\delta_{X_t^{j,N}}$.

- NE-MFC : Particle $i \in [\![1, N]\!]$ interacts through $\frac{\sum_{j=1}^{N}\xi^{i,j}\delta_{X_t^{j,N}}}{\sum_{j=1}^{N}\xi^{i,j}}$ where $(\xi^{i,j})_{1\leqslant j\leqslant N}$ refers to interactions weights between i and j assuming no isolated particle, i.e. $\sum_{j=1}^{N}\xi^{i,j} > 0$.
  $\rightarrow$ Graphon case : $\xi^{i,j} = G(\frac{i}{N}, \frac{j}{N})$.

- Taking heuristically the limit as $N \nearrow \infty$, agent labeled by $u \in I := [0,1]$ interacts through weighted probability measure

$$I \ni u \mapsto \frac{\int_I G(u,v)\mathbb{P}_{X_t^v}(\mathrm{d}x)\mathrm{d}v}{\int_I G(u,v)\mathrm{d}v}) \in \mathcal{P}(\mathbb{R}^d), \quad 0 \leqslant t \leqslant T, \quad u \in I,$$

and dynamics of the controlled state system

$$\begin{cases} \mathrm{d}X_t^u &= b\big(u, X_t^u, \alpha_t^u, \frac{\int_I G(u,v)\mathbb{P}_{X_t^v}(\mathrm{d}x)\mathrm{d}v}{\int_I G(u,v)\mathrm{d}v}\big)\mathrm{d}t + \sigma\big(u, X_t^u, \alpha_t^u, \frac{\int_I G(u,v)\mathbb{P}_{X_t^v}(\mathrm{d}x)\mathrm{d}v}{\int_I G(u,v)\mathrm{d}v}\big)\mathrm{d}W_t^u \\ X_0^u &= \xi^u. \end{cases}$$

Motivations for non exchangeable mean field systems

- MFC theory : Interactions through symmetric particles and homogeneous interactions, through empirical measure $\frac{1}{N}\sum_{j=1}^{N}\delta_{X_t^{j,N}}$.
- NE-MFC : Particle $i \in [\![1,N]\!]$ interacts through $\frac{\sum_{j=1}^{N}\xi^{i,j}\delta_{X_t^{j,N}}}{\sum_{j=1}^{N}\xi^{i,j}}$ where $(\xi^{i,j})_{1\leqslant j\leqslant N}$ refers to interactions weights between i and j assuming no isolated particle, i.e. $\sum_{j=1}^{N}\xi^{i,j} > 0$.
  $\rightarrow$ Graphon case : $\xi^{i,j} = G(\frac{i}{N},\frac{j}{N})$.
- Taking heuristically the limit as $N \nearrow \infty$, agent labeled by $u \in I := [0,1]$ interacts through weighted probability measure

$$I \ni u \mapsto \frac{\int_I G(u,v)\mathbb{P}_{X_t^v}(\mathrm{d}x)\mathrm{d}v}{\int_I G(u,v)\mathrm{d}v}) \in \mathcal{P}(\mathbb{R}^d), \quad 0 \leqslant t \leqslant T, \quad u \in I,$$

and dynamics of the controlled state system

$$\begin{cases} \mathrm{d}X_t^u &= b\big(u, X_t^u, \alpha_t^u, \frac{\int_I G(u,v)\mathbb{P}_{X_t^v}(\mathrm{d}x)\mathrm{d}v}{\int_I G(u,v)\mathrm{d}v}\big)\mathrm{d}t + \sigma\big(u, X_t^u, \alpha_t^u, \frac{\int_I G(u,v)\mathbb{P}_{X_t^v}(\mathrm{d}x)\mathrm{d}v}{\int_I G(u,v)\mathrm{d}v}\big)\mathrm{d}W_t^u \\ X_0^u &= \xi^u. \end{cases}$$

A strong formulation

Dynamics of the controlled state processes:

$$\begin{cases} dX_t^u &= b(u, X_t^u, \alpha_t^u, (\mathbb{P}_{X_t^v})_{v\in I})\mathrm{d}t + \sigma(u, X_t^u, \alpha_t^u, (\mathbb{P}_{X_t^v})_{v\in I})\mathrm{d}W_t^u, \quad 0 \leqslant t \leqslant T, u \in I, \\ X_0^u &= \xi^u. \end{cases} \tag{1}$$

Cost Functional : Aim to minimize over a collection of processes $\boldsymbol{\alpha} = (\alpha^u)_{u\in I}$ in a suitable class $\mathcal{A}$

$$J^S(\boldsymbol{\alpha}) = \int_I \mathbb{E}\Big[\int_0^T f(u, X_t^u, \alpha_t^u, (\mathbb{P}_{X_t^v})_{v\in I})\mathrm{d}t + g(u, X_T^u, (\mathbb{P}_{X_T^v})_{v\in I})\Big]\mathrm{d}u \tag{2}$$

→ Compute $V_0^S = J^S(\boldsymbol{\alpha}^\star)$ where $\boldsymbol{\alpha}^\star$ is a minimizer of J^S.

• Maps (b, σ, f, g) are defined over the space

$$L^2(I; \mathcal{P}_2(\mathbb{R}^d)) = \big\{u \to \mu^u \text{ is measurable and} \int_I \mathcal{W}_2(\mu^u, \delta_0)^2 \mathrm{d}u < +\infty\big\}.$$

• Lack of measurability of the map $(u, \omega) \mapsto X^u(\omega)$ on the space product $(I \times \Omega, \mathcal{B}(I) \otimes \mathcal{A})$. Control problem is defined at the level of the law of the processes $\mathbb{P}_{(X^u, \alpha^u, W^u)}$.

Since we are working at the level of the laws of the processes and for numerical purposes, we can relax the control problem formulation (1)-(2)

A label state formulation

Dynamics of the controlled state processes:

$$\begin{cases} dX_t &= b(U, X_t, \alpha_t, \mathbb{P}_{(U,X_t)})\mathrm{d}t + \sigma(U, X_t, \alpha_t, \mathbb{P}_{(U,X_t)})\mathrm{d}W_t, \quad 0 \leqslant t \leqslant T \\ X_0 &= \xi. \end{cases} \tag{3}$$

Cost Functional : Aim to minimize over α over a suitable class $\mathcal{A}$

$$J^W(\alpha) = \mathbb{E}\Big[\int_0^T f(U, X_t, \alpha_t, \mathbb{P}_{(U,X_t)})\mathrm{d}t + g(U, X_T, \mathbb{P}_{(U,X_T)})\Big] \tag{4}$$

$\rightarrow$ Compute $V_0^W = J^W(\alpha^\star)$ where $\alpha^\star$ is a minimizer of J^W.

Maps (b, σ, f, g) are defined over the space

$$\mathcal{P}_2^\lambda(I \times \mathbb{R}^d) = \big\{\mu \in \mathcal{P}_2(I \times \mathbb{R}^d) : \mathrm{pr}_{1\sharp}\mu = \lambda\big\}.$$

Connection between strong and label-state formulation

- We prove $V_0^S = V_0^W$. It relies essentially on

$$\mathbb{P}_{(X_t^u, \alpha_t^u)} = \mathbb{P}_{(X_t, \alpha_t)|U=u}, \quad \mathrm{d}u \text{ a.e}, \tag{5}$$

  when given the same policy map $\hat{a}$.
- Label-state formulation is more suitable for numerical methods.
- Strong formulation is more suitable for path-wise interpretation.

Objectives

- Adapt the Pontryagin Maximum Principle to mean field control for non exchangeable mean field systems (NE-MFC) to find necessary and sufficient conditions for an admissible optimal control α.
- Propose an illustration in the Linear Quadratic (LQ) case with numerical illustrations.

Gateaux derivative on $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$

Let $f : \mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \to \mathbb{R}$. For $U, X, Y \in L^2(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{R}^d)$ such that $\mathbb{P}_{(U,X)}, \mathbb{P}_{(U,Y)} \in \mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$, we have

$$\lim_{\epsilon \to 0} \frac{1}{\epsilon}\big(f(\mathbb{P}_{(U,X+\epsilon Y)}) - f(\mathbb{P}_{(U,X)})\big) = \tilde{\mathbb{E}}\Big[\partial_{\tilde{x}} \frac{\delta}{\delta m} f(\mathbb{P}_{(U,X)})(\tilde{U}, \tilde{X}) \cdot \tilde{Y}\Big] \tag{6}$$

where $(\tilde{U}, \tilde{X}, \tilde{Y})$ is an independent copy of (U, X, Y) on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

→ Such function $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \times I \times \mathbb{R}^d \ni (\mu, \tilde{u}, \tilde{x}) \mapsto \frac{\delta}{\delta m} f(\mu)(\tilde{u}, \tilde{x}) \in \mathbb{R}$ is called linear functional derivative of f.

Gateaux derivative on $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$

Let $f : \mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \to \mathbb{R}$. For $U, X, Y \in L^2\big(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{R}^d\big)$ such that $\mathbb{P}_{(U,X)}, \mathbb{P}_{(U,Y)} \in \mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$, we have

$$\lim_{\epsilon \to 0} \frac{1}{\epsilon}\big(f(\mathbb{P}_{(U,X+\epsilon Y)}) - f(\mathbb{P}_{(U,X)})\big) = \tilde{\mathbb{E}}\Big[\partial_{\tilde{x}} \frac{\delta}{\delta m} f(\mathbb{P}_{(U,X)})(\tilde{U}, \tilde{X}) \cdot \tilde{Y})\Big] \tag{6}$$

where $(\tilde{U}, \tilde{X}, \tilde{Y})$ is an independent copy of (U, X, Y) on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

→ Such function $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \times I \times \mathbb{R}^d \ni (\mu, \tilde{u}, \tilde{x}) \mapsto \frac{\delta}{\delta m} f(\mu)(\tilde{u}, \tilde{x}) \in \mathbb{R}$ is called linear functional derivative of f.

Convexity on $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$

Let $f : I \times \mathbb{R}^d \times \mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \to \mathbb{R}$. f is said to be convex if for every $u \in I$, $x, x' \in \mathbb{R}^d$, $\mu, \mu' \in L^2\big(\mathcal{P}_2(\mathbb{R}^d)\big)$, we have :

$$\begin{aligned} f(u, x', \mu') - f(u, x, \mu) \geqslant\ & \partial_x f(u, x, \mu).(x' - x) \\ & + \mathbb{E}\Big[\partial_{\tilde{x}} \frac{\delta}{\delta m} f(u, x, \mu)(U, X).(X' - X)\Big]. \end{aligned} \tag{7}$$

where $(U, X') \sim \mu'$ and $(U, X) \sim \mu$.

Definition of the Hamiltonian H

The Hamiltonian $\mathbb{R}$-valued function H of the stochastic optimization problem is defined as :

$$H(u, x, \mu, y, z, a) = b(u, x, \mu, a) \cdot y + \sigma(u, x, \mu, a) : z + f(u, x, \mu, a) \tag{8}$$

where $(u, x, \mu, y, z, a) \in I \times \mathbb{R}^d \times \mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^{d \times n} \times A$.

$\longrightarrow$ Compute an optimality criterion involving the Hamiltonian H assuming differentiability and convexity as defined previously over the space $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$.
$\longrightarrow$ In the following, A will denote a convex subset of $\mathbb{R}^m$ for $m \in \mathbb{N}^*$.

Adjoint Equations to X

We call adjoint processes of X any pair (Y, Z) in $\mathbb{S}^2([0,T];\mathbb{R}^d) \times \mathbb{H}^2([0,T];\mathbb{R}^{d\times n})$ such that (Y, Z) is solution to the adjoint equation

$$\begin{cases} dY_t = -\partial_x H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, \alpha_t)\mathrm{d}t + Z_t \mathrm{d}W_t \\ \qquad -\tilde{\mathbb{E}}\left[\partial_x \frac{\delta}{\delta m} H(\tilde{U}, \tilde{X}_t, \mathbb{P}_{(U,X_t)}, \tilde{Y}_t, \tilde{Z}_t, \tilde{\alpha}_t)(U, X_t)\right]\mathrm{d}t\,, \quad t \in [0,T]\,, \\ Y_T = \partial_x g(U, X_T, \mathbb{P}_{(U,X_T)}) + \tilde{\mathbb{E}}\left[\partial_x \frac{\delta}{\delta m} g(\tilde{U}, \tilde{X}_T, \mathbb{P}_{(U,X_T)})(U, X_T)\right]\,, \end{cases} \tag{9}$$

where $(\tilde{X}, \tilde{Y}, \tilde{Z}, \tilde{\alpha})$ is an independent copy of (X, Y, Z, α) defined on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

Adjoint Equations to X

We call adjoint processes of X any pair (Y, Z) in $\mathbb{S}^2([0,T];\mathbb{R}^d) \times \mathbb{H}^2([0,T];\mathbb{R}^{d\times n})$ such that (Y, Z) is solution to the adjoint equation

$$\begin{cases} dY_t = -\partial_x H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, \alpha_t)\mathrm{d}t + Z_t \mathrm{d}W_t \\ \qquad -\tilde{\mathbb{E}}\left[\partial_x \frac{\delta}{\delta m} H(\tilde{U}, \tilde{X}_t, \mathbb{P}_{(U,X_t)}, \tilde{Y}_t, \tilde{Z}_t, \tilde{\alpha}_t)(U, X_t)\right]\mathrm{d}t\,, \quad t \in [0,T]\,, \\ Y_T = \partial_x g(U, X_T, \mathbb{P}_{(U,X_T)}) + \tilde{\mathbb{E}}\left[\partial_x \frac{\delta}{\delta m} g(\tilde{U}, \tilde{X}_T, \mathbb{P}_{(U,X_T)})(U, X_T)\right]\,, \end{cases} \tag{9}$$

where $(\tilde{X}, \tilde{Y}, \tilde{Z}, \tilde{\alpha})$ is an independent copy of (X, Y, Z, α) defined on $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$.

$\rightarrow$ We retrieve the adjoint equations of the standard Pontryagin formulation but here with the addition of the label randomization U, i.e via $\mathbb{P}_{(U,X_t)}$ and extension of Lions's derivative over $\mathcal{P}_2^{\lambda}(I \times \mathbb{R}^d)$.

We now state the main results which are obtained under some regularity assumptions on b, σ, f and g.

Gâteaux derivative of J

For $\beta \in \mathcal{A}$ such that $\alpha + \epsilon\beta \in \mathcal{A}$ for $\epsilon > 0$ small enough, we have :

$$\lim_{\epsilon \to 0} \frac{1}{\epsilon}\left(J^W(\alpha + \epsilon\beta) - J^W(\alpha)\right) = \mathbb{E}\left[\int_0^T \left(\partial_\alpha H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, \alpha_t) \cdot \beta_t\right)\mathrm{d}t\right]$$

where X is given by (3), (Y, Z) are given by (9) and the Hamiltonian function H is given by (8).

Necessary condition for optimality of α

Moreover, if we assume that H is convex in $a \in A$, that $\alpha = (\alpha_t)_{,0 \leqslant t \leqslant T}$ is optimal, that $X = (X_t)_{0 \leqslant t \leqslant T}$ is the associated optimal control state given by (3) and that $(Y, Z) = (Y_t, Z_t)_{0 \leqslant t \leqslant T}$ are the associated adjoint processes solving (9), then we have :

$$\forall a \in A, \quad H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, \alpha_t) \leqslant H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, a) \quad dt \otimes d\mathbb{P} \text{ a.e} \quad (10)$$

Sufficient condition for optimality of α

Let $\alpha \in \mathcal{A}$, $X = (X_t)_{0 \leqslant t \leqslant T}$ the corresponding controlled state process and $(Y, Z) = (Y_t, Z_t)_{0 \leqslant t \leqslant T}$ the corresponding adjoint processes.

(1) $\mathbb{R}^d \times \mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \ni (x, \mu) \to g(U, x, \mu)$ is convex $\mathrm{d}\mathbb{P}$ a .e

(2) $\mathbb{R}^d \times \mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \times A \ni (x, \mu, a) \to H(U, x, \mu, Y_t, Z_t, a)$ is convex $dt \otimes d\mathbb{P}$ a.e

If we assume also following the necessary condition for optimality :

$$H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, \alpha_t) = \inf_{\beta \in A} H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, Z_t, \beta), \qquad dt \otimes d\mathbb{P} \text{ a.e}$$

Then, α is an optimal control in the sense that $J(\alpha) = \inf_{\alpha' \in \mathcal{A}} J(\alpha')$

Linear quadratic optimal control problem

We consider the following class of models (assuming for sake of simplicity σ constant and $A = \mathbb{R}^m$).

$$\begin{cases} dX_t &= \Big[\beta(U) + A(U)X_t + \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}}\big[G_A(U,\tilde{U})\tilde{X}_t\big] + B(U)\alpha_t\Big]\mathrm{d}t + \gamma(U)\mathrm{d}W_t, t \in [0,T] \\ &= \Big[\beta(U) + A(U)X_t + \int_{I\times\mathbb{R}^d}\big[G_A(U,v)x\big]\mathbb{P}_{(U,X_t)}(\mathrm{d}v,\mathrm{d}x) + B(U)\alpha_t\Big] + \gamma(U)\mathrm{d}W_t, t \in [0,T] \\ X_0 &= \xi, \end{cases}$$

where $\beta \in L^\infty(I;\mathbb{R}^d)$, $\gamma \in L^\infty(I;\mathbb{R}^d)$, $A \in L^\infty(I,\mathbb{R}^{d\times d})$, $B \in L^\infty(I;\mathbb{R}^{d\times m})$ and $G_A \in L^2(I \times I;\mathbb{R}^{d\times d})$.

Linear quadratic optimal control problem

We consider the following class of models (assuming for sake of simplicity σ constant and $A = \mathbb{R}^m$).

$$\begin{cases} dX_t &= \Big[\beta(U) + A(U)X_t + \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}}\big[G_A(U,\tilde{U})\tilde{X}_t\big] + B(U)\alpha_t\Big]\mathrm{d}t + \gamma(U)\mathrm{d}W_t, t \in [0,T] \\ &= \Big[\beta(U) + A(U)X_t + \int_{I\times\mathbb{R}^d}\big[G_A(U,v)x\big]\mathbb{P}_{(U,X_t)}(\mathrm{d}v,\mathrm{d}x) + B(U)\alpha_t\Big] + \gamma(U)\mathrm{d}W_t, t \in [0,T] \\ X_0 &= \xi, \end{cases}$$

where $\beta \in L^\infty(I;\mathbb{R}^d)$, $\gamma \in L^\infty(I;\mathbb{R}^d)$, $A \in L^\infty(I,\mathbb{R}^{d\times d})$, $B \in L^\infty(I;\mathbb{R}^{d\times m})$ and $G_A \in L^2(I \times I;\mathbb{R}^{d\times d})$.

Quadratic cost functional

$$\begin{aligned} J(\boldsymbol{\alpha}) = \mathbb{E}\Big[\int_0^T \Big[& Q(U)(X_t - \tilde{\mathbb{E}}_{[\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}]}[\tilde{G}_Q(U,\tilde{U})\tilde{X}_t]) \cdot (X_t - -\tilde{\mathbb{E}}_{[\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}}[\tilde{G}_Q(U,\tilde{U})\tilde{X}_t] \\ & + \alpha_t \cdot R(U)\alpha_t + 2\alpha_t \cdot \Gamma(U)X_t + 2\alpha_t \cdot \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}}\big[G_I(U,\tilde{U})\tilde{X}_t\big]\Big]dt \\ & + H(U)\big(X_T - \tilde{\mathbb{E}}_{[\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}}[\tilde{G}_H(U,\tilde{U})\tilde{X}_t]\big) \cdot \big(X_T - \tilde{\mathbb{E}}_{[\tilde{U},\tilde{X}_t)\sim\mathbb{P}_{(U,X_t)}}[\tilde{G}_H(U,\tilde{U})\tilde{X}_t]\big)\Big] \end{aligned} \quad (11)$$

Proposition : Optimal control in the LQ case

In the Linear quadratic model, the unique optimal control $\hat{\alpha} = (\hat{\alpha}_t)_{0 \leqslant t \leqslant T}$ is given by

$$\hat{\alpha}_t = S_t(U)X_t + \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim \mathbb{P}_{(U,X_t)}}\Big[\bar{S}_t(U,\tilde{U})\tilde{X}_t\Big] + \Gamma_t(U), \quad 0 \leqslant t \leqslant T, \tag{12}$$

where $X = (X_t)_{0 \leqslant t \leqslant T}$ is the unique solution to the SDE obtained after replacing $\hat{\alpha}_t$ by (12) and where we denoted

$$\begin{cases} S_t(U) = -(R(U))^{-1}\Big((B(U))^\top K_t(U) + \Gamma(U)\Big), \\ \bar{S}_t(U,\tilde{U}) = -(R(U))^{-1}\Big((B(U)^\top \bar{K}_t(U,\tilde{U}) + G_I(U,\tilde{U})\Big) \\ \Gamma_t(U) = -(R(U))^{-1}(B(U))^\top \Lambda_t(U) \end{cases}$$

where $K \in C^1\big([0,T]; L^\infty(I;\mathbb{S}^d_+)\big)$, $\bar{K} \in C^1\big([0,T], L^2(I\times I;\mathbb{R}^{d\times d})\big)$ and $\Lambda \in C^1\big([0,T]; L^2(I;\mathbb{R}^d)\big)$ are to be determined through infinite dimensional Riccati equations.

Systemic risk model with heterogeneous banks

Dynamics of the controlled state processes:

$$\begin{cases} dX_t &= \Big[\kappa(X_t - \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim \mathbb{P}_{(U,X_t)}}[G_\kappa(U,\tilde{U})\tilde{X}_t] + \alpha_t\Big]dt + \sigma\, dW_t, \quad 0 \leqslant t \leqslant T, \\ X_0 &= \xi. \end{cases}$$

Cost Functional :

$$J(\alpha) = \mathbb{E}\Bigg[\int_0^T \Bigg\{\eta(U)\Big(X_t - \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim \mathbb{P}_{(U,X_t)}}[G_\eta(U,\tilde{U})\,\tilde{X}_t]\Big)^2 + q(U)\,\alpha_t\Big(X_t - \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim \mathbb{P}_{(U,X_t)}}[G_q(U,\tilde{U})\,\tilde{X}_t]\Big) + \alpha_t^2\Bigg\}\mathrm{d}t + r(U)\Big(X_T - \tilde{\mathbb{E}}_{(\tilde{U},\tilde{X}_t)\sim \mathbb{P}_{(U,X_t)}}[G_r(U,\tilde{U})\,\tilde{X}_T]\Big)^2\Bigg]. \tag{13}$$

Numerical methods for learning the optimal feedback control map $\hat{\mathrm{a}}$.

Deep Graphon :

- Direct parametrization of the control via a Neural Network in feedback form in view of (12) solved by standard gradient descent algorithm.
- To learn functions defined over $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \to$ conditional moment neural network where we approximate $\mu \in \mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$ by its conditional moments.

Numerical methods for learning the optimal feedback control map $\hat{\mathrm{a}}$.

Deep Graphon :

- Direct parametrization of the control via a Neural Network in feedback form in view of (12) solved by standard gradient descent algorithm.
- To learn functions defined over $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$ $\rightarrow$ conditional moment neural network where we approximate $\mu \in \mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$ by its conditional moments.

Deep BSDE Graphon :

- Optimal control learnt in view of (10) since

$$\hat{\alpha}_t = \mathrm{a}(U, t, X_t, \mathbb{P}_{(U,X_t)}, Y_t) = \arg\min_{a \in A} H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, a).$$

$\rightarrow$ Learn (X, Y) by exploiting the FBSDE equation.

- We use 2 neural networks $\mathcal{U}_\theta(\mu)(u, x)$ and $\mathcal{Z}_\theta(t, \mu)(u, x)$ to approximate initial value of Y and the Z component and we minimiser over θ the cost functional

$$\theta \mapsto L(\theta) = \mathbb{E}\Big[\big|Y_T^\theta - G(X_T^\theta, \mathbb{P}_{(U,X_T^\theta)})\big|^2\Big],$$

Starting from $\mathcal{U}_\theta(\mathbb{P}_{(U,X_0)})(U, X_0)$, we diffuse

$$\begin{cases} X_{t_{i+1}}^\theta &= X_{t_i}^\theta + B(U, X_{t_i}^\theta, Y_{t_i}^\theta, \mathbb{P}_{(U,X_{t_i}^\theta)})\Delta t + \sigma \Delta W_{t_{i+1}} \\ Y_{t_{i+1}}^\theta &= Y_{t_i}^\theta + H(U, X_{t_i}^\theta, Y_{t_i}^\theta, \mathcal{Z}_\theta(t_i, \mathbb{P}_{(U,X_{t_i}^\theta)})(U, X_{t_i}^\theta), \mathbb{P}_{(U,X_{t_i}^\theta)})\Delta t + \mathcal{Z}_\theta(t_i, \mathbb{P}_{(U,X_{t_i})})(U, X_{t_i})\Delta W_{t_{i+1}} \end{cases}$$

for certains maps B, H depending on model coefficients.

Numerical methods for learning the optimal feedback control map $\hat{\mathrm{a}}$.

Deep Graphon :

- Direct parametrization of the control via a Neural Network in feedback form in view of (12) solved by standard gradient descent algorithm.
- To learn functions defined over $\mathcal{P}_2^\lambda(I \times \mathbb{R}^d) \to$ conditional moment neural network where we approximate $\mu \in \mathcal{P}_2^\lambda(I \times \mathbb{R}^d)$ by its conditional moments.

Deep BSDE Graphon :

- Optimal control learnt in view of (10) since

$$\hat{\alpha}_t = \mathrm{a}(U, t, X_t, \mathbb{P}_{(U,X_t)}, Y_t) = \arg\min_{a \in A} H(U, X_t, \mathbb{P}_{(U,X_t)}, Y_t, a).$$

$\to$ Learn (X, Y) by exploiting the FBSDE equation.

- We use 2 neural networks $\mathcal{U}_\theta(\mu)(u, x)$ and $\mathcal{Z}_\theta(t, \mu)(u, x)$ to approximate initial value of Y and the Z component and we minimiser over θ the cost functional

$$\theta \mapsto L(\theta) = \mathbb{E}\Big[\big|Y_T^\theta - G(X_T^\theta, \mathbb{P}_{(U,X_T^\theta)})\big|^2\Big],$$

Starting from $\mathcal{U}_\theta(\mathbb{P}_{(U,X_0)})(U, X_0)$, we diffuse

$$\begin{cases} X_{t_{i+1}}^\theta &= X_{t_i}^\theta + B(U, X_{t_i}^\theta, Y_{t_i}^\theta, \mathbb{P}_{(U,X_{t_i}^\theta)})\Delta t + \sigma \Delta W_{t_{i+1}} \\ Y_{t_{i+1}}^\theta &= Y_{t_i}^\theta + H(U, X_{t_i}^\theta, Y_{t_i}^\theta, \mathcal{Z}_\theta(t_i, \mathbb{P}_{(U,X_{t_i}^\theta)})(U, X_{t_i}^\theta), \mathbb{P}_{(U,X_{t_i}^\theta)})\Delta t + \mathcal{Z}_\theta(t_i, \mathbb{P}_{(U,X_{t_i})})(U, X_{t_i})\Delta W_{t_{i+1}} \end{cases}$$

for certains maps B, H depending on model coefficients.

Application with $\sigma = 1$, $\eta = 0.73$, $q = 0.8$, $r = 0.22$ and $\kappa = 0.62$:

Method	Riccati	Deep Graphon	Deep BSDE Graphon
Value	0.58830	0.58826	0.58820

Table: Expected cost function using $M = 10000$ in simulation with $G(u, v) = e^{-uv}$.

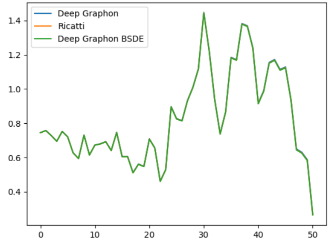

Figure: Optimal trajectory of X with $u = 0.708$

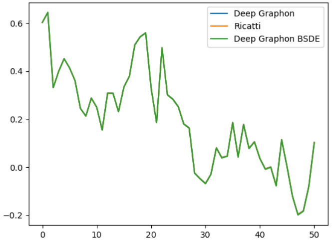

Figure: Optimal trajectory of X with $u = 0.599$

Figure: Comparison between NN solvers and the Riccati one with $G(u, v) = e^{-uv}$

Summary of the talk

- We introduced a general class of non exchangeable mean field systems.
- We present an application of our framework to the case of LQ optimal control problem where we exhibit a new infinite-dimensional system of Riccati equations and we show numerically how to solve the optimal control problem through Deep learning algorithms.

Summary of the talk

- We introduced a general class of non exchangeable mean field systems.
- We present an application of our framework to the case of LQ optimal control problem where we exhibit a new infinite-dimensional system of Riccati equations and we show numerically how to solve the optimal control problem through Deep learning algorithms.

Future Works

- In the present setting, agents interact through a specified graph/graphon structure but it could be interesting to add a control perspective on the agent's interactions.
- Adding some randomness in the graph structure would lead to the study of dynamical systems with random interactions $\Longrightarrow$ Bridge with Random Matrix Theory and Operator-Theory.

References

I. Kharroubi, S.Mekkaoui and H.Pham. *Stochastic maximum principle for optimal control problem of non exchangeable mean field systems, arXiv preprint arXiv:2506.05595*

F.de Feo and S.Mekkaoui. *Optimal control of heterogeneous mean-field stochastic differential equations with common noise and applications to financial models, arXiv preprint arXiv:2511.18636*

S. Mekkaoui, H.Pham and X.Warin *Learning mappings on labeled conditional distributions, Work in Progress*

S.Mekkaoui, H.Pham *Analysis of Non-Exchangeable Mean Field Markov Decision Processes with common noise : From Bellman equation to quantitative propagation of chaos. Work in Progress*

THANK YOU FOR YOUR ATTENTION