‘

AMS short course on mean field games: The convergence problem

Daniel Lacker

Columbia University

January 14, 2020

Outline

- The convergence problem: basic questions
- An example (linear-quadratic)
- Warmup: Interacting diffusion processes
- Compactness approach to the convergence problem
- Master equation approach to the convergence problem

n-player games

Agents $i = 1, \ldots, n$ have state process dynamics

$$dX_t^i = \alpha_t^i dt + dW_t^i,$$

with $W^1, \ldots, W^n$ independent Brownian, $(X_0^1, \ldots, X_0^n)$ i.i.d.
Agent i chooses α^i to maximize

$$J_i^n(\alpha^1, \ldots, \alpha^n) = \mathbb{E}\left[\int_0^T f(X_t^i, \mu_t^n, \alpha_t^i)dt + g(X_T^i, \mu_T^n)\right],$$

$$\mu_t^n = \frac{1}{n}\sum_{k=1}^n \delta_{X_t^k},$$

Say $(\alpha^1, \ldots, \alpha^n)$ form an Nash equilibrium if

$$J_i^n(\alpha^1, \ldots, \alpha^n) \geq \sup_\beta J_i^n(\ldots, \alpha^{i-1}, \beta, \alpha^{i+1}, \ldots), \quad \forall i.$$

n-player games

Agents $i = 1, \ldots, n$ have state process dynamics

$$dX_t^i = \alpha_t^i dt + dW_t^i,$$

with $W^1, \ldots, W^n$ independent Brownian, $(X_0^1, \ldots, X_0^n)$ i.i.d.
Agent i chooses α^i to maximize

$$J_i^n(\alpha^1, \ldots, \alpha^n) = \mathbb{E}\left[\int_0^T f(X_t^i, \mu_t^n, \alpha_t^i)dt + g(X_T^i, \mu_T^n)\right],$$

$$\mu_t^n = \frac{1}{n}\sum_{k=1}^n \delta_{X_t^k},$$

Say $(\alpha^1, \ldots, \alpha^n)$ form an ϵ-Nash equilibrium if

$$J_i^n(\alpha^1, \ldots, \alpha^n) \geq \sup_\beta J_i^n(\ldots, \alpha^{i-1}, \beta, \alpha^{i+1}, \ldots) - \epsilon, \quad \forall i.$$

Standing assumptions

Assume f and g bounded & continuous.

Controls take values in a compact convex set $A \subset \mathbb{R}^d$.

Equilibrium concepts

Two common choices of admissible control classes:

- **Open loop**: Agent i chooses a process α^i, adapted to the filtration generated by $(W^1, \ldots, W^n)$
- **Closed loop (Markovian)**: Agent i chooses a function

$$\alpha^i = \alpha^i(t, X_t^1, \ldots, X_t^n).$$

(such that $dX_t^i = \alpha^i(t, X_t^1, \ldots, X_t^n)dt + dW_t^i$ is well-posed)

These choices lead to distinct equilibria!

Even though a Markovian *control* induces an open loop *control*,

$$\hat{\alpha}_t^i(\omega) := \alpha^i(t, X_t^1(\omega), \ldots, X_t^n(\omega)).$$

Open-loop equilibrium: players do not "react" to others' controls.
cf. examples in Carmona-Fouque-Sun '13

Connecting n-player games and MFGs

Question:

How do we rigorously connect the n-player games to their proposed $n \to \infty$ limit, the mean field game (MFG)?

Idea #1: Construction of approximate equlibria

Suppose we can directly solve the candidate MFG, with associated equilibrium control $\widehat{\alpha}(t,x)$. Let player i in n-player game use control $\alpha^i_t = \widehat{\alpha}(t, X^i_t)$. Can often show this yields an ϵ_n-Nash equilibrium, where $\lim_n \epsilon_n = 0$. (see R. Malhamé's talk)

Idea #2: Limits of true equilibria (the convergence problem)

Suppose for each n we are given a Nash equilibrium for the n-player game. Can we identify a limit as $n \to \infty$, say of the empirical measure μ^n?

An example

Players $i = 1, \dots, n$ control

$$dX_t^i = \alpha_t^i dt + dW_t^i, \quad X_0^i = 0.$$

Goal of player i: minimize

$$\mathbb{E}\left[\int_0^T |\alpha_t^i|^2 dt + \lambda |X_T^i - \overline{X}_T|^2\right], \quad \overline{X}_T = \frac{1}{n}\sum_{i=1}^n X_T^i, \quad \lambda > 0.$$

Unique (closed-loop) Nash equilibrium:

$$\alpha_t^i = c_t^n(\overline{X}_t - X_t^i), \quad c_t^n := \frac{\lambda(1 - \frac{1}{n})}{1 + \lambda(1 - \frac{1}{n^2})(T - t)}.$$

An example

Optimal state processes:

$$dX_t^i = c_t^n(\overline{X}_t - X_t^i)dt + dW_t^i, \quad X_0^i = 0.$$

Average over $i = 1, \dots, n$ to get

$$\overline{X}_t = \frac{1}{n}\sum_{i=1}^{n} W_t^i \to 0 \quad \text{as } n \to \infty. \qquad (*)$$

$$c_t^n = \frac{\lambda(1 - \frac{1}{n})}{1 + \lambda(1 - \frac{1}{n^2})(T - t)} \to \frac{\lambda}{1 + \lambda(T - t)} =: c_t^\infty. \qquad (**)$$

$$dX_t^i \approx -c_t^\infty X_t^i dt + dW_t^i.$$

Convergence problem: Direct using (*) and (**).

An example

Optimal state processes:

$$dX_t^i = c_t^n(\overline{X}_t - X_t^i)dt + dW_t^i, \quad X_0^i = 0.$$

Average over $i = 1, \dots, n$ to get

$$\overline{X}_t = \frac{1}{n}\sum_{i=1}^{n} W_t^i \to 0 \quad \text{as } n \to \infty. \qquad (*)$$

$$c_t^n = \frac{\lambda(1-\frac{1}{n})}{1+\lambda(1-\frac{1}{n^2})(T-t)} \to \frac{\lambda}{1+\lambda(T-t)} =: c_t^\infty. \qquad (**)$$

$$dX_t^i \approx -c_t^\infty X_t^i dt + dW_t^i.$$

Constructing approximate equilibria: Use $\alpha_t^i = -c_t^\infty X_t^i$ for n-player game. ($\rightsquigarrow X^i$ i.i.d.)

An example

Optimal state processes:

$$dX_t^i = c_t^n(\overline{X}_t - X_t^i)dt + dW_t^i, \quad X_0^i = 0.$$

Average over $i = 1, \ldots, n$ to get

$$\overline{X}_t = \frac{1}{n}\sum_{i=1}^n W_t^i \to 0 \quad \text{as } n \to \infty. \qquad (*)$$

$$c_t^n = \frac{\lambda(1 - \frac{1}{n})}{1 + \lambda(1 - \frac{1}{n^2})(T - t)} \to \frac{\lambda}{1 + \lambda(T - t)} =: c_t^\infty. \qquad (**)$$

$$dX_t^i \approx -c_t^\infty X_t^i dt + dW_t^i.$$

Better construction of approximate equilibria:
Use $\alpha_t^i = c_t^\infty(\overline{X}_t - X_t^i)$ "$= \partial_x U(t, X_t^i, \mu_t^n)$" (master equation)

Warmup for general theory: Interacting diffusions

Suppose particles $i = 1, \dots, n$ interact through their empirical measure according to

$$dX_t^i = b(X_t^i, \mu_t^n)dt + dW_t^i, \qquad \mu_t^n = \frac{1}{n}\sum_{k=1}^n \delta_{X_t^k},$$

where $W^1, \dots, W^n$ are independent Brownian, $(X_0^1, \dots, X_0^n)$ i.i.d.

Under "nice" assumptions on b, we have $\mu_t^n \to \mu_t$, where μ_t solves the **McKean-Vlasov** equation,

$$dX_t = b(X_t, \mu_t)dt + dW_t, \qquad \mu_t = \mathrm{Law}(X_t),$$

or (nonlinear) Fokker-Planck equation

$$\frac{d}{dt}\langle \mu_t, \varphi\rangle = \left\langle \mu_t,\ b(\cdot, \mu_t) \cdot \nabla\varphi + \frac{1}{2}\Delta\varphi \right\rangle, \quad \forall \varphi \in C_c^\infty.$$

Convergence proof method #1

Step 1: Pre-compactness/tightness of (μ^n) in $C([0,T];\mathcal{P}(\mathbb{R}^d))$.

Step 2: Characterize limit points. Let $\varphi \in C_c^\infty$. Apply Itô's formula to $\varphi(X_t^i)$, then average over $i = 1,\ldots,n$ to get

$$d\langle \mu_t^n, \varphi\rangle = \left\langle \mu_t^n, b(\cdot,\mu_t^n)\cdot\nabla\varphi + \frac{1}{2}\Delta\varphi\right\rangle dt + \frac{1}{n}\sum_{i=1}^n \nabla\varphi(X_t^i)\cdot dW_t^i.$$

The red term is a martingale with quadratic variation $O(1/n)$. Hence, any limit μ of (μ^n) satisfies

$$\frac{d}{dt}\langle \mu_t, \varphi\rangle = \left\langle \mu_t, b(\cdot,\mu_t)\cdot\nabla\varphi + \frac{1}{2}\Delta\varphi\right\rangle.$$

Step 3: Prove uniqueness for this Fokker-Planck equation.

Even without uniqueness: limits of (μ^n) concentrate on the set of McKean-Vlasov solutions.

Convergence proof method #2

Step 1: Show existence/uniqueness of McKean-Vlasov SDE:

$$dX_t = b(X_t, \mu_t)dt + dW_t, \qquad \mu_t = \mathrm{Law}(X_t),$$

Construct i.i.d. driven by same Brownian motions as n-particle system:

$$dY_t^i = b(Y_t^i, \mu_t)dt + dW_t^i, \quad Y_0^i = X_0^i, \quad \rightsquigarrow \mathrm{Law}(Y_t^i) = \mu_t.$$

Step 2: Use this coupling to estimate $\frac{1}{n}\sum_{i=1}^n |X_t^i - Y_t^i|$.
If b Lipschitz w.r.t. suitable (e.g., Wasserstein) metric, can show

$$\lim_{n\to\infty} \mathbb{E}\left[\sup_{t\in[0,T]} |X_t^k - Y_t^k|^2\right] = 0, \qquad \forall k.$$

Note $(X^1, \ldots, X^k) \to (Y^1, \ldots, Y^k)$ for fixed k as $n \to \infty$, so particles are asymptotically i.i.d. $\rightsquigarrow$ "propagation of chaos"

MFG recap

(Take $f \equiv 0$ for simplicity.)

Representative player's state process:

$$dX_t^\alpha = \alpha_t dt + dW_t, \qquad X_0 = \xi.$$

- For fixed $\mu = (\mu_t)_{t\in[0,T]} \in C([0,T];\mathcal{P}(\mathbb{R}^d))$, solve
$$\alpha^* \in \operatorname*{argmax}_\alpha \mathbb{E}\left[g(X_T^\alpha, \mu_T)\right].$$
- Let $\mu_t^* = \mathrm{Law}(X_t^{\alpha^*})$.
- Say μ is a mean field equilibrium (MFE) if $\mu^* = \mu$.

The equilibrium control is α^*.

Constructing approximate equilibria

Let μ be a MFE with equilibrium control $\alpha^* = \alpha^*(t, X_t)$.

Tell each player i in n-player game to use control $\alpha^*(t, X_t^i)$, so state processes are

$$dX_t^i = \alpha^*(t, X_t^i)dt + dW_t^i.$$

These are i.i.d. with $X_t^i \sim \mu_t$, so empirical measure $\mu_t^n = \frac{1}{n}\sum_{i=1}^n \delta_{X_t^i}$ converges to μ_t by LLN. If player i alone switches to some β, then μ_t^n does not change much, so

$$\begin{aligned} J_i^n(\ldots,\alpha^*,\beta,\alpha^*,\ldots) &= \mathbb{E}\left[g(X_T^{i,\beta},\tilde{\mu}_T^n)\right] \approx \mathbb{E}\left[g(X_T^{\beta},\mu_T)\right] \\ &\le \mathbb{E}\left[g(X_T^{\alpha^*},\mu_T)\right] \approx \mathbb{E}\left[g(X_T^{i,\alpha^*},\mu_T^n)\right] \\ &= J_i^n(\ldots,\alpha^*,\alpha^*,\alpha^*,\ldots). \end{aligned}$$

$\rightsquigarrow$ $(\alpha^*,\ldots,\alpha^*)$ is ϵ_n-Nash with $\epsilon_n \to 0$.

The convergence problem

Suppose for each n that $(\alpha^{n,1}, \dots, \alpha^{n,n})$ is a Nash equilibrium.

$$dX_t^{n,i} = \alpha_t^{n,i} dt + dW_t^i.$$

The problem: Does $\mu_t^n = \frac{1}{n}\sum_{i=1}^n \delta_{X_t^{n,i}}$ converge to a MFE?

Two main approaches (so far):

1. Compactness + characterize limit points
   - Works even when MFE is non-unique.
   - Focuses on limit of μ^n.
2. Show n-player PDE system "converges" to master equation
   - Quantiative, yields rates. See K. Ramanan's lecture.
   - Focuses on limit of value functions.

The convergence problem via compactness

Suppose for each n that $(\alpha^{n,1}, \ldots, \alpha^{n,n})$ is a Nash equilibrium.

$$dX_t^{n,i} = \alpha_t^{n,i} dt + dW_t^i.$$

The problem: Does $\mu_t^n = \frac{1}{n} \sum_{i=1}^n \delta_{X_t^{n,i}}$ converge to a MFE?

Three main steps:

1. Show tightness (precompactness) of sequence $(\mathrm{Law}(\mu^n))$, e.g. in $\mathcal{P}(C([0,T];\mathcal{P}(\mathbb{R}^d))$.
2. Characterize dynamics of limit points, i.e., show limit measure flow "corresponds to" $dX_t^\alpha = \alpha_t dt + dW_t$, for some α.
3. Check the MFG optimality property, i.e., show $\mathbb{E}[g(X_T^\alpha, \mu_T)] \geq \mathbb{E}[g(X_T^\beta, \mu_T)]$ for all β.

The convergence problem via compactness

Step 1: Tightness/precompactness requires some machinery but is straightforward, e.g., if:

- controls α take values in a compact space A.
- Initial positions $X_0^{n,i} = \xi^i$ are i.i.d.

The convergence problem via compactness

Step 2: Limiting dynamics of $(\mu_t^n)_{t\in[0,T]}$.

A key issue is the lack of regularity of controls: $(\alpha_t^{n,i})_{t\in[0,T]}$ need not be continuous or even right-continuous!

A solution:
Bounded control set $A \Rightarrow (\alpha_\cdot^{n,i})_{n,i}$ uniformly bounded in $L^2[0,T]$
$\Rightarrow (\alpha_\cdot^{n,i})_{n,i}$ tight as random elements of $(L^2[0,T], \mathrm{weak})$.

A more versatile approach uses relaxed controls, a.k.a. Young measures

The convergence problem via compactness

Step 3: Optimality of limiting control α.

Recall: $dX_t^\beta = \beta_t dt + dW_t$

Goal: Show

$$\mathbb{E}[g(X_T^\alpha, \mu_T)] \geq \mathbb{E}[g(X_T^\beta, \mu_T)], \quad \forall \beta.$$

Idea: Let β be any alternative control, and "give it" to each player in n-player game, one at a time. Since $(\alpha^{n,1}, \ldots, \alpha^{n,n})$ is Nash,

$$\mathbb{E}[g(X_T^{n,i}, \mu_T^n)] \geq \mathbb{E}[g(\widetilde{X}_T^{n,i}, \mu_T^{n,-i})], \quad \forall i = 1, \ldots, n,$$

where:

$$\widetilde{X}^{n,i} = \text{player } i\text{'s state after switching to } \beta,$$
$$\mu^{n,-i} = \text{empirical measure after } i \text{ switches to } \beta.$$

The convergence problem via compactness

Step 3 continued: Optimality of limiting control α.

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[g(X_T^{n,i},\mu_T^n)] \geq \frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[g(\widetilde{X}_T^{n,i},\mu_T^{n,-i})]$$

Step 2 shows $\mu^n \to \mu$, so

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[g(X_T^{n,i},\mu_T^n)] = \mathbb{E}[\langle \mu_T^n, g(\cdot,\mu_T^n)\rangle]$$
$$\to \mathbb{E}[\langle \mu_T, g(\cdot,\mu_T)\rangle] = \mathbb{E}[g(X_T^{\alpha},\mu_T)].$$

Remains to show

$$\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[g(\widetilde{X}_T^{n,i},\mu_T^{n,-i})] \to \mathbb{E}[g(X_T^{\beta},\mu_T)].$$

The convergence problem via compactness

Step 3 continued: Optimality of limiting control α.

Remains to show

$$\frac{1}{n}\sum_{i=1}^{n} \underbrace{\mathbb{E}[g(\widetilde{X}^{n,i}_T, \mu^{n,-i}_T)]}_{\text{player } i \text{ switches to } \beta} \to \mathbb{E}[g(X^{\beta}_T, \mu_T)], \quad \forall \beta. \qquad (*)$$

This is technical & challenging!

The issue: Need to show rigorously $\mu^{n,-i}_T \to \mu_T$, i.e., empirical measure limit is insensitive to one-player deviations.

- **Open-loop case**: Not hard, since player's deviation effects only $X^{n,i}$, not $(X^{n,k})_{k \neq i}$. E.g. $\|\mu^n_t - \mu^{n,-i}_t\|_{\mathrm{TV}} \leq 2/n$.
- **Closed-loop case**: Players react to each other, so one player's deviation can effect entire system!

The convergence problem via PDEs

Consider n-player game where player i controls

$$dX_t^i = \alpha_t^i dt + dW_t^i,$$

by choosing $\alpha_t^i \in \mathbb{R}^d$ to try to maximize

$$\mathbb{E}\left[g(X_T^i, \mu_T^n) - \frac{1}{2}\int_0^T |\alpha_t^i|^2 dt\right].$$

The value functions $(v_i^n)_{i=1}^n$, $v_i^n : [0, T] \times (\mathbb{R}^d)^n \to \mathbb{R}$, satisfy

$$0 = \partial_t v_i^n(t, \boldsymbol{x}) - \frac{1}{2}|\nabla_i v_i^n(t, \boldsymbol{x})|^2 + \frac{1}{2}\Delta v_i^n(t, \boldsymbol{x})$$
$$+ \sum_{k=1}^n \nabla_k v_i^n(t, \boldsymbol{x}) \cdot \nabla_k v_k^n(t, \boldsymbol{x}), \qquad v_i^n(t, \boldsymbol{x}) = g(x_i, m_{\boldsymbol{x}}^n),$$

$$\boldsymbol{x} = (x_1, \ldots, x_n) \in (\mathbb{R}^d)^n, \qquad m_{\boldsymbol{x}}^n = \frac{1}{n}\sum_{i=1}^n \delta_{x_i}.$$

Control of player k is $\alpha_t^k = \nabla_k v_k^n(t, \boldsymbol{X_t})$

The convergence problem via PDEs: heuristics

Symmetry (+ uniqueness) $\Rightarrow v_i^n(t, \boldsymbol{x}) = V_n(t, x_i, m_{\boldsymbol{x}}^n)$ for some $V_n : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$, assumed smooth.

Differentiation rules $\Rightarrow$

$$\nabla_k v_i^n(t, \boldsymbol{x}) = \frac{1}{n} D_m V_n(t, x_i, m_{\boldsymbol{x}}^n, x_k) + \delta_{ik} \nabla_x V_n(t, x_i, m_{\boldsymbol{x}}^n)$$

$$\Delta_k v_i^n(t, \boldsymbol{x}) = \frac{1}{n} \mathrm{div}_v D_m V_n(t, x_i, \boldsymbol{x}, x_k) + \delta_{ik} \Delta_x V_n(t, x_i, m_{\boldsymbol{x}}^n) + O(1/n)$$

PDE system becomes "almost" the master equation:

$$O(1/n) = \partial_t V_n(t, x_i, m_{\boldsymbol{x}}^n) - \frac{1}{2} |\nabla_x V_n(t, x_i, m_{\boldsymbol{x}}^n)|^2 + \frac{1}{2} \Delta_x V_n(t, x_i, m_{\boldsymbol{x}}^n)$$
$$+ \frac{1}{n} \sum_{k=1}^n \mathrm{div}_v D_m V_n(t, x_i, m_{\boldsymbol{x}}^n, x_k)$$
$$+ \frac{1}{n} \sum_{k=1}^n D_m V_n(t, x_i, m_{\boldsymbol{x}}^n, x_k) \cdot \nabla_x V_n(t, x_k, m_{\boldsymbol{x}}^n)$$

The convergence problem via PDEs: heuristics

Averages rewrite as integrals w.r.t. empirical measure:

$$
\begin{aligned}
O(1/n) &= \partial_t V_n(t, x_i, m^n_{\boldsymbol{x}}) - \frac{1}{2}|\nabla_x V_n(t, x_i, m^n_{\boldsymbol{x}})|^2 + \frac{1}{2}\Delta_x V_n(t, x_i, m^n_{\boldsymbol{x}}) \\
&\quad + \int_{\mathbb{R}^d} \mathrm{div}_v D_m V_n(t, x_i, m^n_{\boldsymbol{x}}, v)\, m^n_{\boldsymbol{x}}(dv) \\
&\quad + \int_{\mathbb{R}^d} D_m V_n(t, x_i, m^n_{\boldsymbol{x}}, v) \cdot \nabla_x V_n(t, v, m^n_{\boldsymbol{x}})\, m^n_{\boldsymbol{x}}(dv)
\end{aligned}
$$

Replace $x_i \to x$ and $m^n_{\boldsymbol{x}} \to m$ to get master equation:

$$
\begin{aligned}
O(1/n) &= \partial_t V_n(t, x, m) - \frac{1}{2}|\nabla_x V_n(t, x, m)|^2 + \frac{1}{2}\Delta_x V_n(t, x, m) \\
&\quad + \int_{\mathbb{R}^d} \mathrm{div}_v D_m V_n(t, x, m, v)\, m(dv) \\
&\quad + \int_{\mathbb{R}^d} D_m V_n(t, x, m, v) \cdot \nabla_x V_n(t, v, m)\, m(dv).
\end{aligned}
$$

The convergence problem via PDEs

The master equation:

$$0 = \partial_t U(t,x,m) - \frac{1}{2}|\nabla_x U(t,x,m)|^2 + \frac{1}{2}\Delta_x U(t,x,m)$$
$$+ \int_{\mathbb{R}^d} \mathrm{div}_v D_m U(t,x,m,v)\, m(dv)$$
$$+ \int_{\mathbb{R}^d} D_m U(t,x,m,v) \cdot \nabla_x U(t,v,m)\, m(dv),$$
$$U(T,x,m) = g(x,m).$$

Proper argument: Show existence/uniqueness of classical solution U of master equation. Define $u_i^n(t,\boldsymbol{x}) := U(t,x_i,m_{\boldsymbol{x}}^n)$.

Show: $(u_i^n)_{i=1}^n$ nearly solves n-player Nash PDE system, via same idea as before.

The convergence problem via PDEs

Proper argument: Show existence/uniqueness of classical solution U of master equation. Define $u_i^n(t, \boldsymbol{x}) := U(t, x_i, m_{\boldsymbol{x}}^n)$.

Show: $(u_i^n)_{i=1}^n$ nearly solves n-player Nash PDE system. Derive estimates on $|u_i^n - v_i^n|$ and $|\nabla_i u_i^n - \nabla_i v_i^n|$.

Equilibrium states, n-player game:

$$dX_t^i = \nabla_i v_i^n(t, \boldsymbol{X}_t)dt + dW_t^i, \quad X_0^i \text{ i.i.d.}$$

Auxiliary n-particle system (McKean-Vlasov):

$$d\overline{X}_t^i = \nabla_x U(t, \overline{X}_t^i, m_{\overline{\boldsymbol{X}}_t}^n)dt + dW_t^i, \quad X_0^i = \overline{X}_0^i.$$

(1) Estimates on $|\nabla_i u_i^n - \nabla_i v_i^n| \Rightarrow X^i \approx \overline{X}^i$

(2) Propagation of chaos $\Rightarrow m_{\overline{\boldsymbol{X}}_t}^n \to \mu_t$, unique solution of

$$dY_t = \nabla_x U(t, Y_t, \mu_t)dt + dW_t, \quad \mu_t = \mathrm{Law}(Y_t).$$

(3) (1) and (2) $\Rightarrow m_{\boldsymbol{X}_t}^n \to \mu_t$

Outlook & open problems

- **Master equation** approach requires smoothness. What to do without classical solutions?
  - Is there a good weak solution theory?
  - Does convergence argument still work?
  - Same rate? Are fluctuations & large deviations at the same scale? (see K. Ramanan's lecture)
- **Non-unique regime** is important & challenging.
  - We know all (sub-sequential) limit points of n-player equilibria are (weak) MFE. Which MFE arise as a limit point?
  - When is a single MFE "selected" in the limit?