From the master equation to mean field game asymptotics

Daniel Lacker

Division of Applied Mathematics, *Brown University*

8th Western Conference in Mathematical Finance, March 24, 2017

Joint work with Francois Delarue and Kavita Ramanan

Overview

A **mean field game** (MFG) is a game with a continuum of players.

In various contexts, we know rigorously that the MFG arises as the limit of n-player games as $n \to \infty$.

But how close of an approximation is an MFG for the n-player game?

This talk: Refined MFG asymptotics in the form of a central limit theorem and large deviation principle, as well as non-asymptotic concentration bounds.

Interacting diffusions

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

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

where $W^1, \ldots, W^n$ are independent Brownian motions.

Under "nice" assumptions on b, we have $\bar{\nu}_t^n \to \nu_t$, where ν_t solves the **McKean-Vlasov** equation,

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

Empirical measure limit theory

There is a rich literature on asymptotics of $\bar{\nu}^n_t$:

1. LLN: $\bar{\nu}^n \to \nu$, where ν solves a McKean-Vlasov equation.
   (Oelschläger '84, Gärtner '88, Sznitman '91, etc.)
2. Fluctuations: $\sqrt{n}(\bar{\nu}^n_t - \nu_t)$ converges to a distribution-valued process driven by space-time Brownian motion.
   (Tanaka '84, Sznitman '85, Kurtz-Xiong '04, etc.)
3. Large deviations: $\bar{\nu}^n$ has an explicit LDP.
   (Dawson-Gärtner '87, Budhiraja-Dupius-Fischer '12)
4. Concentration: Finite-n bounds are available for $\mathbb{P}(d(\bar{\nu}^n, \nu) > \epsilon)$, for various metrics d.
   (Bolley-Guillin-Villani '07, etc.)

The idea: The McKean-Vlasov system is often more amenable to analysis than the more physical n-particle system.

From particle systems to mean field games

Interacting diffusion systems are zero-intelligence models.

Mean field games are often more suitable in financial/economic applications, replacing particles with decision-makers. The dynamics of X^i become controlled, and the n-particle system becomes a game.

The idea: Approximate the realistic n-player game equilibrium using the more tractable MFG limit ($n \to \infty$).

This talk: Quantitatively relate the n-player equilibrium to an interacting diffusion system, then bootstrap existing results for the latter.

A class of mean field 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 minimize

$$J_i^n(\alpha^1, \ldots, \alpha^n) = \mathbb{E}\left[\int_0^T \left(f(X_t^i, \bar{\mu}_t^n) + \frac{1}{2}|\alpha_t^i|^2\right) dt + g(X_T^i, \bar{\mu}_T^n)\right],$$

$$\bar{\mu}_t^n = \frac{1}{n}\sum_{k=1}^n \delta_{X_t^k}.$$

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

$$J_i^n(\alpha^1, \ldots, \alpha^n) \leq \epsilon + \inf_\beta J_i^n(\ldots, \alpha^{i-1}, \beta, \alpha^{i+1}, \ldots), \forall i = 1, \ldots, n$$

The n-player HJB system

The value function $v_i^n(t,\boldsymbol{x})$, for $\boldsymbol{x} = (x_1, \dots, x_n)$, for agent i in the n-player game solves

$$\partial_t v_i^n(t,\boldsymbol{x}) + \frac{1}{2}\sum_{k=1}^n \Delta_{x_k} v_i^n(t,\boldsymbol{x}) + \frac{1}{2}|\nabla_{x_i} v_i^n(t,\boldsymbol{x})|^2 \\ + \sum_{k\neq i} \nabla_{x_k} v_k^n(t,\boldsymbol{x}) \cdot \nabla_{x_k} v_i^n(t,\boldsymbol{x}) = f\left(x_i, \frac{1}{n}\sum_{k=1}^n \delta_{x_k}\right).$$

A Nash equilibrium is given by

$$\alpha_t^i = \nabla_{x_i} v_i^n(t, X_t^1, \dots, X_t^n).$$

But v_i^n is generally hard to find, especially for large n.

Mean field limit $n \to \infty$?

The problem

Given a Nash equilibrium $(\alpha^{n,1}, \ldots, \alpha^{n,n})$ for each n, can we describe the limit(s) of $\bar{\mu}^n_t$?

Previous results

Lasry/ Lions '06, Feleqi '13, Fischer '14, Lacker '15, Cardaliaguet-Delarue-Lasry-Lions '15, Cardaliaguet '16...

A related, better-understood problem

Find a mean field game solution directly, and use it to construct an ϵ_n-Nash equilibrium for the n-player game, where $\epsilon_n \to 0$.
See Huang/Malhamé/Caines '06 & many others.

Proposed mean field game limit

A deterministic measure flow $(\mu_t)_{t \in [0,T]} \in C([0,T];\mathcal{P}(\mathbb{R}^d))$ is a mean field equilibrium (MFE) if:

$$\begin{cases} \alpha^* & \in \arg\min_\alpha \mathbb{E}\left[\int_0^T \left(f(X_t^\alpha, \mu_t) + \frac{1}{2}|\alpha_t|^2\right) dt + g(X_T^\alpha, \mu_T)\right], \\ dX_t^\alpha & = \alpha_t dt + dW_t, \\ \mu_t & = \mathrm{Law}(X_t^{\alpha^*}). \end{cases}$$

Theorem (Law of large numbers)

Under very strong assumptions, there exists a unique MFE μ, and $\bar{\mu}^n \to \mu$ in probability in $C([0,T];\mathcal{P}(\mathbb{R}^d))$.

MFG value function

The MFE is completely described by the master equation, when it is solvable.

1. Fix $t \in [0, T)$ and $m \in \mathcal{P}(\mathbb{R}^d)$.
2. Solve the MFG starting from (t, m), i.e., find (α^*, μ) s.t.

$$\begin{cases} \alpha^* & \in \arg\min_\alpha \mathbb{E}\left[\int_t^T \left(f(X_s^\alpha, \mu_s) + \frac{1}{2}|\alpha_s|^2\right) ds + g(X_T^\alpha, \mu_T)\right], \\ dX_s^\alpha & = \alpha_s ds + dW_s, \quad s \in (t, T) \\ \mu_s & = \mathrm{Law}(X_s^{\alpha^*}), \quad \mu_t = m \end{cases}$$

3. Define the value function, for $x \in \mathbb{R}^d$, by

$$U(t, x, m) = \mathbb{E}\left[\int_t^T \left(f(X_s^{\alpha^*}, \mu_s) + \frac{1}{2}|\alpha_s^*|^2\right) ds + g(X_T^{\alpha^*}, \mu_T) \middle| X_t^{\alpha^*} = x\right]$$

Derivatives

There is a dynamic programming principle for U if the MFE is unique. To derive a PDE, we need to differentiate in m:

Definition

Say $u : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ is C^1 if $\exists \frac{\delta u}{\delta m} : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}$ continuous such that, for $m, \widetilde{m} \in \mathcal{P}(\mathbb{R}^d)$,

$$\lim_{h \downarrow 0} \frac{u(m + t(\widetilde{m} - m)) - u(m)}{t} = \int_{\mathbb{R}^d} \frac{\delta u}{\delta m}(m, y)\, d(\widetilde{m} - m)(y).$$

Define also (when it exists)

$$D_m u(m, y) = \nabla_y \left(\frac{\delta u}{\delta m}(m, y) \right).$$

Key tool: The master equation

Heuristically, using the DPP along with an Itô formula for functions of measures, one derives the master equation for the value function:

$$\begin{aligned}
&\partial_t U(t,x,m) - \int_{\mathbb{R}^d} \nabla_x U(t,y,m) \cdot D_m U(t,x,m,y) m(dy) \\
&\quad + f(x,m) - \frac{1}{2}|\nabla_x U(t,x,m)|^2 + \frac{1}{2}\Delta_x U(t,x,m) \\
&\quad + \frac{1}{2}\int_{\mathbb{R}^d} \mathrm{div}_y D_m U(t,x,m,y) m(dy) = 0,
\end{aligned}$$

Refer to Cardaliaguet-Delarue-Lasry-Lions '15, Chassagneux-Crisan-Delarue '14, Carmona-Delarue '14, Bensoussan-Frehse-Yam '15

Key tool: The master equation

Heuristically, using the DPP along with an Itô formula for functions of measures, one derives the master equation for the value function:

$$
\begin{aligned}
&\partial_t U(t,x,m) - \int_{\mathbb{R}^d} \nabla_x U(t,y,m) \cdot D_m U(t,x,m,y) m(dy) \\
&\quad + f(x,m) - \frac{1}{2}\left|\nabla_x U(t,x,m)\right|^2 + \frac{1}{2}\Delta_x U(t,x,m) \\
&\quad + \frac{1}{2}\int_{\mathbb{R}^d} \mathrm{div}_y D_m U(t,x,m,y) m(dy) = 0,
\end{aligned}
$$

Assume henceforth that there is a smooth classical solution!

A first n-particle approximation

The MFE μ is the unique solution of the McKean-Vlasov equation

$$dX_t = \underbrace{\nabla_x U(t, X_t, \mu_t)}_{\alpha_t^*} dt + dW_t, \qquad \mu_t = \mathrm{Law}(X_t).$$

Old idea: Consider the system of n independent processes,

$$dX_t^i = \underbrace{\nabla_x U(t, X_t^i, \mu_t)}_{\alpha_t^i} dt + dW_t^i.$$

These controls α_t^i can be proven to form an ϵ_n-equilibrium for the n-player game, where $\epsilon_n \to 0$.

Note X_t^i are i.i.d. $\sim \mu_t$, so their empirical measure tends to μ_t.

A better n-particle approximation

Key idea of Cardaliaguet et al.: Consider the McKean-Vlasov system

$$dY_t^i = \underbrace{\nabla_x U(t, Y_t^i, \bar{\nu}_t^n)}_{\alpha_t^i} dt + dW_t^i, \qquad \bar{\nu}_t^n = \frac{1}{n}\sum_{k=1}^n \delta_{Y_t^k}.$$

Classical theory says that $\bar{\nu}^n \to \nu$, where ν solves the McKean-Vlasov equation,

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

We had the same equation for the MFE μ, so uniqueness implies

$$\mu \equiv \nu.$$

So to prove $\bar{\mu}^n \to \mu$, it suffices to show $\bar{\mu}^n$ and $\bar{\nu}^n$ are **close**.

A better n-particle approximation

Theorem (Cardaliaguet et al. '15)

Recalling that $\bar{\mu}^n_t$ denotes the n-player Nash equilibrium empirical measure, $\bar{\mu}^n$ and $\bar{\nu}^n$ are very close.

Proof idea: Show that

$$u_i^n(t, x_1, \ldots, x_n) = U\left(t, x_i, \frac{1}{n-1}\sum_{k\neq i}\delta_{x_k}\right)$$

nearly solves the n-player HJB system.

Note: This requires smoothness assumptions on the master equation U, but not on the n-player HJB system!

The n-player HJB system revisited

Define

$$u_i^n(t, x_1, \ldots, x_n) = U\left(t, x_i, \frac{1}{n-1}\sum_{k \neq i} \delta_{x_k}\right).$$

Assuming $\nabla_x U$ is Lipschitz and $D_m U$ is bounded, we have

$$\partial_t u_i^n(t, \boldsymbol{x}) + \frac{1}{2}\sum_{k=1}^n \Delta_{x_k} u_i^n(t, \boldsymbol{x}) + \frac{1}{2}|\nabla_{x_i} u_i^n(t, \boldsymbol{x})|^2$$

$$+ \sum_{k \neq i} \nabla_{x_k} u_k^n(t, \boldsymbol{x}) \cdot \nabla_{x_k} u_i^n(t, \boldsymbol{x}) = f\left(x_i, \frac{1}{n}\sum_{k=1}^n \delta_{x_k}\right) + r_i^n(t, \boldsymbol{x}),$$

where r_i^n is continuous, with $\|r_i^n\|_\infty \leq C/n$.

Nash system vs. McKean-Vlasov system

The n-player Nash equilibrium state processes solve

$$dX_t^i = \nabla_{x_i} v_i^n(t, X_t^1, \ldots, X_t^n)dt + dW_t^i.$$

Compare this to the McKean-Vlasov system,

$$\begin{aligned} dY_t^i &= \nabla_x U(t, Y_t^i, \bar{\nu}_t^n)dt + dW_t^i, \quad \text{where } \bar{\nu}_t^n = \frac{1}{n}\sum_{k=1} \delta_{Y_t^k}, \\ &\approx \nabla_x U(t, Y_t^i, \bar{\nu}_t^{n,i})dt + dW_t^i, \quad \text{where } \bar{\nu}_t^{n,i} = \frac{1}{n-1}\sum_{k \neq i} \delta_{Y_t^k}, \\ &= \nabla_{x_i} u_i^n(t, Y_t^1, \ldots, Y_t^n)dt + dW_t^i. \end{aligned}$$

Apply Itô to $|v_i^n(t, Y^1, \ldots, Y_t^n) - u_i^n(t, X^1, \ldots, X_t^n)|^2$ and use the PDEs, along with Lipschitz estimates on $\nabla_x U$.

Toward refined mean field game asymptotics

Main idea: Compare the Nash EQ empirical measure $\bar{\mu}^n$ to the McKean-Vlasov empirical measure $\bar{\nu}^n$, and then apply...

Known results on McKean-Vlasov limits:

1. LLN: $\bar{\nu}^n \to \mu$, where μ is the unique MFE.
2. Fluctuations: $\sqrt{n}(\bar{\nu}^n_t - \mu_t)$ converges.
3. Large deviations: $\mathbb{P}(\bar{\nu}^n \in A) \approx \exp(-c_A n)$ asymptotically.
4. Concentration: $\mathbb{P}(d(\bar{\nu}^n, \mu) \geq \epsilon) \leq C \exp(-Cn\epsilon^2)$.

Note: In linear-quadratic systems, we can instead describe the asymptotics of the mean $\int_{\mathbb{R}^d} x \, d\bar{\mu}^n_t(x)$ in a self-contained manner.

Fluctiuations

Assuming the master equation has a sufficiently smooth solution,

Theorem

The sequences $\sqrt{n}(\bar{\mu}_t^n - \mu_t)$ *and* $\sqrt{n}(\bar{\nu}_t^n - \mu_t)$ *both converge to the unique solution of the SPDE:*

$$\partial_t S_t(x) = \mathcal{A}^*_{t,\mu_t} S_t(x) - \mathrm{div}_x(\sqrt{\mu_t(x)}\dot{B}(t,x)),$$

where B *is a space-time Brownian motion and*

$$\mathcal{A}_{t,m}\varphi(x) := \mathcal{L}_{t,m}\varphi(x) + \int_{\mathbb{R}^d} \frac{\delta}{\delta m}\left(\nabla_x U(t,y,m)\right)(x) \cdot \nabla\varphi(y) m(dy),$$

$$\mathcal{L}_{t,m}\varphi(x) := \nabla_x U(t,x,m) \cdot \nabla\varphi(x) + \frac{1}{2}\Delta\varphi(x),$$

Proof idea: Show $\sqrt{n}(\bar{\mu}_t^n - \bar{\nu}_t^n) \to 0$ using master equation estimates. Kurtz-Xiong '04 identifies limit of $\sqrt{n}(\bar{\nu}_t^n - \mu_t)$.

Large deviations

Assuming the master equation has a sufficiently smooth solution,

Theorem

The sequences $\bar{\mu}^n$ *and* $\bar{\nu}^n$ *both satisfy a large deviation principle on* $C([0,T];\mathcal{P}(\mathbb{R}^d))$, *with the same rate function.*

$$I(m.) = \begin{cases} \frac{1}{2}\int_0^T \|\partial_t m_t - \mathcal{L}^*_{t,m_t} m_t\|_S^2 dt & \text{if } m \text{ abs. cont.} \\ \infty & \text{otherwise,} \end{cases}$$

where $\|\cdot\|_S$ *acts on Schwartz distributions by*

$$\|\gamma\|_S^2 = \sup_{\varphi \in C_c^\infty} \langle \gamma, \varphi \rangle^2 / \langle \gamma, |\nabla \varphi|^2 \rangle.$$

Large deviations

Proof idea: Show exponential equivalence

$$\lim_{n\to\infty}\frac{1}{n}\log\mathbb{P}\left(\sup_{t\in[0,T]}W_1(\bar{\mu}^n_t,\bar{\nu}^n_t)>\epsilon\right)=-\infty,\ \forall\epsilon>0,$$

using master equation estimates, namely $\|\nabla_x U\|_\infty<\infty$.

Identify the LDP for $\bar{\nu}^n$ using Dawson-Gärtner '87 or Budhiraja-Dupuis-Fischer '12.

Concentration

Assuming the master equation has a sufficiently smooth solution,

Theorem

There exist $C, \delta > 0$ such that for all $a > 0$ and $n \geq C/\sqrt{a}$ we have

$$\mathbb{P}\left(\sup_{t\in[0,T]} W_1(\bar{\mu}_t^n, \mu_t) > a\right) \leq C\left(ne^{-\delta a n^2} + e^{-\delta a^2 n}\right),$$

where W_1 is the Wasserstein distance.

Proof.

Use McKean-Vlasov results after showing

$$\mathbb{P}\left(\sup_{t\in[0,T]} W_1(\bar{\mu}_t^n, \bar{\nu}_t^n) > a\right) \leq 2n\exp(-\delta a n^2).$$

□

The moral of the story

Sufficiently smooth solution of master equation
$\Longrightarrow$ refined asymptotics for mean field game equilibria,
by comparing the n-player equilibrium to an n-particle system and then applying existing results on McKean-Vlasov systems.

Major challenges

- Requires a lot of smoothness for the master equation.
- Uniqueness at the limit (i.e., of the MFG) is a restrictive assumption, not needed for McKean-Vlasov large deviations! (c.f. Dawson-Gärtner '88)