MEAN FIELD GAMES AND MASTER EQUATION

FRANÇOIS DELARUE

Abstract. The purpose of the lecture is to introduce the master equation for mean field games. It is a partial differential equation defined on the space of probability measures, which reads as the analogue of the so-called Nash system in game theory but for a continuum of players. This equation was originally set by J.M. Lasry and P.L. Lions to address mean field games when players in the game are all subjected to the same noise, which is then said to be "a common noise". In short, the equation provides an Eulerian point of view on the game.

The notes below should be regarded as a preliminary draft of the lecture that will be given in Denver. This draft will serve as a basis for the final notes that will be published after the courses.

Contents

1. A typical framework

Throughout, we use a quite simple framework. The purpose of this section is precisely to make this framework clear and to introduce step by step the objects that are needed to arrive to the master equation in this context.

1.1. **Finite dimensional version.** In order to capture the philosophy of the master equation, the best is to focus on the following case. To make it clear, we start with the finite player version (N players). The dynamics of particle number $i \in \{1, \dots, N\}$ are assumed to be (in $\mathbb{R}^d$)

$$dX_t^i = \alpha_t^i dt + dW_t^i, \quad t \in [0,T]$$

where $(W_t^1, \dots, W_t^N)_{0 \le t \le T}$ are independent Brownian motions with values in $\mathbb{R}^d$, $(\alpha_t^i)_{0 \le t \le T}$ is a progressively-measurable processes with values in $\mathbb{R}^d$. The initial conditions are IID.

Then, we assume that the cost functional to player $i \in \{1, \dots, N\}$ is

$$J^i\big(\boldsymbol{\alpha}^1, \boldsymbol{\alpha}^2, \dots, \boldsymbol{\alpha}^N\big) = \mathbb{E}\bigg[g\big(X_T^i, \bar{\mu}_T^N\big) + \int_0^T \Big[f\big(X_t^i, \bar{\mu}_t^N\big) + \frac{1}{2}|\alpha_t^i|^2\Big] dt\bigg],$$

where we use bold letters to denote stochastic processes. Here, we encode the population at time t through the so-called empirical measure:

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

Moreover, (f,g) are the same for all i but, importantly, J^i depends on the strategies chosen by the other players, through $\bar{\mu}^N$: This is precisely this sort of interaction that justifies the notion of *mean field* games. Also, it must be clear that

$$f, g : \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R},$$

where $\mathcal{P}_2(\mathbb{R}^d) \rightsquigarrow$ set of probabilities on $\mathbb{R}^d$ with finite 2nd moments. To make it simple, we may assume f and g to be bounded. Later on in the lecture, we require them to be smooth in both directions, which requires (among others) to clarify the notion of derivative to a function on the space $\mathcal{P}_2(\mathbb{R}^d)$.

A typical instance for the cost functional is

$$J^i = \mathbb{E}\Big[\frac{1}{N}\sum_{j=1}^N g\big(X_T^i - X_T^j\big) + \int_0^T \Big(\frac{1}{N}\sum_{j=1}^N f\big(X_t^i - X_t^j\big) + |\alpha_t^i|^2\Big) dt\Big],$$

which reads as a potential energy plus a (sort of) kinetic energy. Interestingly enough, the terminal and running costs may be here rewritten in the form:

$$\frac{1}{N}\sum_{j=1}^N f\big(X_T^i - X_T^j\big) = \big(f \star \bar{\mu}_t^N\big)(X_t^i),$$

$$\frac{1}{N}\sum_{j=1}^N g\big(X_T^i - X_T^j\big) = \big(g \star \bar{\mu}_t^N\big)(X_t^i),$$

where $\star$ stands for the convolution product.

1.2. **MFG problem.** The guess in mean field game theory is that, in the limit N tends to ∞, the (at least some) Nash equilibria are described in terms of a fixed point problem set on the time-indexed trajectories with values in $\mathcal{P}_2(\mathbb{R}^d)$. Whilst it may seem a purely mathematical question, the interest for taking the limit $N \to \infty$ is in fact motivated by practical purposes (as the other lectures should clarify): In short, the limit $N \to \infty$ is expected to allow for a decrease of complexity. The latter is due to the fact that the players (in equilibrium) are expected to become independent in the limit; in the asymptotic regime, they no longer interact with one another, but they just feel the theoretical distribution of the population. To wit, we postulate that the (asymptotic) equilibria should be defined as the solutions of the following fixed point problem:

(i) fix a flow of probability measures $\boldsymbol{\mu} = (\mu_t)_{0\le t\le T}$ (with values in $\mathcal{P}_2(\mathbb{R}^d)$) ;

(ii) solve the stochastic optimal control problem in the environment $(\mu_t)_{0\le t\le T}$

$$dX_t = \alpha_t dt + dW_t$$

with X_0 being fixed on some set-up $(\Omega, \mathbb{F}, \mathbb{P})$ with a d-dimensional Brownian motion, and, importantly, with the cost functional

$$J(\boldsymbol{\alpha}) = \mathbb{E}\Big[g(X_T, \mu_T) + \int_0^T \big(f(X_t, \mu_t) + \frac{1}{2}|\alpha_t|^2\big) dt\Big] ;$$

(*iii*) let $(X_t^{\star,\boldsymbol{\mu}})_{0\le t\le T}$ be the unique optimizer (under nice assumptions). The goal is to find $\boldsymbol{\mu} = (\mu_t)_{0\le t\le T}$ such that

$$\mu_t = \mathcal{L}\big(X_t^{\star,\boldsymbol{\mu}}\big), \quad t \in [0,T].$$

Pay attention that this is not a proof of convergence from the finite game. This is just a guess about the limiting form. And this is part of the other lectures to justify this convergence.

1.3. **PDE formulation.** Part of the MFG theory relies on the PDE characterization of the optimal control problem. Indeed, we may first define the value function in environment $(\mu_t)_{0\le t\le T}$, the latter being defined as in item (i) in §1.2,

$$u(t,x) = \inf_{\alpha \text{ processes}} \mathbb{E}\Big[g(X_T,\mu_T) + \int_t^T \Big(f(X_s,\mu_s) + \tfrac{1}{2}|\alpha_s|^2\Big)ds \,|\, X_t = x\Big].$$

Then, a standard fact in stochastic control theory (a vey good reference on the subject is [6]) is that u solves a backward HJB equation

$$\Big(\partial_t u + \frac{\Delta_x u}{2}\Big)(t,x) - \frac{1}{2}|\partial_x u(t,x)|^2 + f(x,\mu_t) = 0 \tag{1}$$

with the terminal boundary condition: $u(T,\cdot) = g(\cdot,\mu_T)$. Of course, the difficulty is that u depends on $(\mu_t)_{0\le t\le T}$! So, we need for a PDE characterization of $(\mathcal{L}(X_t^{\star,\boldsymbol{\mu}}))_{0\le t\le T}$ (the star is here used to denote the optimizer) in (2) in §1.2. In fact, we know that the dynamics of $X^{\star,\boldsymbol{\mu}}$ are given by

$$dX_t^{\star,\boldsymbol{\mu}} = -\partial_x u(t, X_t^{\star,\boldsymbol{\mu}})dt + dW_t.$$

Recalling that $(X_t^{\star,\boldsymbol{\mu}})_{0\le t\le T}$ satisfies a Fokker-Planck (FP) equation, condition (3) gives

$$\partial_t \mu_t = \operatorname{div}_x\big(\partial_x u(t,x)\mu_t\big) + \frac{1}{2}\Delta_x \mu_t \tag{2}$$

We end up with the fact that MFG equilibria are described by a $\boxed{\text{forward-backward}}$ system in infinite dimension:

Fokker-Planck (forward) / HJB (backward)

This is an infinite dimensional analogue of a two point boundary problem ($x_\circ$ being prescribed)

$$\dot{x}_t = b(x_t,y_t)dt, \quad x_0 = x_\circ\,; \quad \dot{y}_t = -f(x_t,y_t)dt, \quad y_T = g(x_T),$$

for which Cauchy-Lipschitz theory holds in small time only.

1.4. **Existence and uniqueness.** For sure, solvability is a key subject in mean field games. And in fact, the very first papers in the field were precisely devoted to this question, recalling from the previous paragraph that Picard fixed point theorem only applies in very specific situations (somehow, assuming T small is the same as assuming the coupling between the two forward and backward equation is weak enough). Generally speaking, existence may be proved by a fixed point theorem without uniqueness, like Schauder theorem: We refer to [2] (although it is certainly not the first reference on this question). Uniqueness is a more delicate issue, but is known to hold true under the following conditions (due to Lasry and Lions, see for instance [1, 2, 3]):

$$\int_{\mathbb{R}^d} \big(f(x,\mu) - f(x,\mu')\big) d\big(\mu - \mu'\big)(x) \geq 0$$
$$\int_{\mathbb{R}^d} \big(g(x,\mu) - g(x,\mu')\big) d\big(\mu - \mu'\big)(x) \geq 0$$

Examples are provided in [3].

1.5. **MFG system as a system of characteristics.** Assume now that existence and uniqueness to the MFG system holds true. Consider an initial condition μ° of the population at time t°. Since uniqueness holds true, we may denote by $(\mu_t)_{t^\circ \leq t \leq T}$ the equilibrium starting from μ°. Accordingly, we may condider the solution of the optimal control problem starting from x° under $\boldsymbol{\mu} = (\mu_t)_{t^\circ \leq t \leq T}$

$$dX_t = -\partial_x u^{\boldsymbol{\mu}}(t, X_t)dt + dW_t \quad t \in [t^\circ, T], \tag{3}$$

with $X_{t^\circ} = x^\circ$ and

$$\partial_t u^{\boldsymbol{\mu}}(t,x) = -\tfrac{1}{2}\Delta u^{\boldsymbol{\mu}}(t,x) + \tfrac{1}{2}|\partial_x u^{\boldsymbol{\mu}}(t,x)|^2 - f(x,\mu_t)$$
$$u^{\boldsymbol{\mu}}(T,x) = g(x,\mu_T)$$

Then, we may introduce a generalized value function :

$$\mathcal{U}(t^\circ, x^\circ, \mu^\circ) = u^{\boldsymbol{\mu}}(t^\circ, x^\circ). \tag{4}$$

In short, we take the value, at time (t°, x°) of the optimal control problem in environment $\boldsymbol{\mu}$! The good point is that we expect for a dynamic programming, meaning that, for ϵ small,

$$\mathcal{U}(t^\circ, x^\circ, \mu^\circ) = \mathbb{E}\bigg[\int_{t^\circ}^{t^\circ+\epsilon} \Big[f\big(X_s, \mu_s\big) + \tfrac{1}{2}|\partial_x u^{\boldsymbol{\mu}}(s, X_s)|^2\Big] ds + \mathcal{U}\big(t^\circ + \epsilon, X_{t^\circ+\epsilon}, \mu_{t^\circ+\epsilon}\big)\bigg]. \tag{5}$$

This relationship is the key step toward the master equation.

2. Differentiation on the space of probability measures

The trick now is to expand (with respect to ϵ) the right-hand side of (5) at order 1. Obviously, this requires a suitable chain rule (or Itô's formula).

2.1. **Wasserstein distance.** Throughout, we equip the space of probability measures $\mathcal{P}_2(\mathbb{R}^d)$ with a metric, called Wasserstein distance (or more precisely 2-Wasserstein distance). The definition is as follows

$$\forall \mu, \nu \in \mathcal{P}_2(\mathbb{R}^d), \quad W_2(\mu,\nu) = \bigg(\inf_\pi \int_{\mathbb{R}^d \times \mathbb{R}^d} |x-y|^2 d\pi(x,y) \bigg)^{1/2},$$

where π has μ and ν as marginals on $\mathbb{R}^d \times \mathbb{R}^d$.

It must be understood that, in the rest of the notes, the probability measures we have in hand are (most of the time) associated in a somewhat canonical fashion with random variables (meaning that there are natural candidates for representing those probability measures as laws of random variables). So, for two probability measures μ and μ', we may think of random variables X and X' having μ and μ' as respective laws. Then, the joint law of X

and X' is a candidate for being a π in the above infimum. This leads us to the following upper bound:

$$W_2(\mathcal{L}(X), \mathcal{L}(X')) \leq \mathbb{E}[|X - X'|^2]^{1/2}.$$

Exercise: Prove that for any two N-tuples $(x_1, \cdots, x_N)$ and $(x'_1, \cdots, x'_N)$ in $(\mathbb{R}^d)^N$:

$$W_2\Big(\frac{1}{N}\sum_{i=1}^N \delta_{x_i}, \frac{1}{N}\sum_{i=1}^N \delta_{x'_i}\Big) \leq \Big(\frac{1}{N}\sum_{i=1}^N |x_i - x'_i|^2\Big)^{1/2}.$$

Hint: We may regarde $\frac{1}{N}\sum_{i=1}^N \delta_{x_i}$ as the law of x_Θ for Θ a random variable (on some probability space) with uniform law on $\{1, \cdots, N\}$. Doing the same with $\frac{1}{N}\sum_{i=1}^N \delta_{x_i}$, the proof follows.

2.2. **Differential calculus on Wasserstein space (à la Lions).** The difficulty is that $\mathcal{P}_2(\mathbb{R}^d)$ is not a flat space. For sure, one may embed it into the space of finite measures, but we may also look for a more intrinsic way of differentiating on the space of probability measures. There are several approaches to do so (they are in the end more or less equivalent). The one we present here is due to Lions.

To make it clear, we are given a function

$$\mathcal{U} : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}.$$

Lions' idea is to lift $\mathcal{U}$ onto an arbitrary L^2 space lying above the space of probability measures. In words, we let

$$\hat{\mathcal{U}} : L^2(\Omega, \mathbb{P}) \ni X \mapsto \mathcal{U}(\mathcal{L}(X) = \mathrm{Law}(X))$$

We then say that $\mathcal{U}$ differentiable if $\hat{\mathcal{U}}$ is Fréchet differentiable. In fact, the main question lies in the form of the Fréchet derivative: How does it keep track of the structure of $\mathcal{U}$? The answer is given by the following identity (see [1, 3]):

$$D\hat{\mathcal{U}}(X) = \partial_\mu \mathcal{U}(\mu)(X),$$

where $\partial_\mu \mathcal{U}(\mu) : \mathbb{R}^d \ni x \mapsto \partial_\mu \mathcal{U}(\mu)(x) \quad \mu = \mathcal{L}(X) \in L^2(\mathbb{R}^d, \mu; \mathbb{R}^d)$. The reader must understand the subtlety: In the left-hand side, the map DU takes the whole random variable (seen as an object of infinite dimension) as input; in the right-hand side, the map $\partial_\mu \mathcal{U}(\mu)(\cdot)$ just takes as input $X(\omega)$ as input for ω being the current random outcome. The very nice thing is that the latter does not depend on the choice of the lift (nor on the choice of X such that $\mathcal{L}(X) = \mu$) and hence may be regarded as a gradient of $\mathcal{U}$ at μ.

Exercise: (The following is checked in [3, Chapter 5].) We may associate with $\mathcal{U}$ the finite-dimensional projection:

$$(\mathbb{R}^d)^N \ni (x_1, \cdots, x_N) \mapsto \mathcal{U}\Big(\frac{1}{N}\sum_{j=1}^N \delta_{x_j}\Big).$$

Then,

$$\partial_{x_i}\Big[\mathcal{U}\Big(\frac{1}{N}\sum_{j=1}^N \delta_{x_j}\Big)\Big] = \frac{1}{N}\partial_\mu \mathcal{U}\Big(\frac{1}{N}\sum_{j=1}^N \delta_{x_j}\Big)(x_i), \quad x_1, \ldots, x_N \in \mathbb{R}^d.$$

Exercise: This is a very standard example:

$$\mathcal{U}(\mu) = \int_{\mathbb{R}^d} h(y) d\mu(y).$$

Compute the derivative when h is continuously differentiable and ∇h at most of linear growth.

Hint: For two random variables X, Y, we have

$$\begin{aligned}\hat{\mathcal{U}}(X+Y) = \mathbb{E}[h(X+Y)] &= \mathbb{E}[h(X)] + \mathbb{E}[\nabla h(X) \cdot Y] + o(\|Y\|_2) \\ &\Rightarrow D\hat{\mathcal{U}}(X) = \nabla h(X).\end{aligned}$$

Deduce that $\partial_\mu \mathcal{U}(\mu)(v) = \nabla h(v)$.

Example: Take a smooth vector field $b : \mathbb{R}^d \to \mathbb{R}^d$ and solve the SDE

$$dX_t = b(X_t)dt, \quad X_0 \sim \mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$$

Then,

$$\begin{aligned}\frac{d}{dt}_{|t=0} \mathcal{U}(\mu_t) = \frac{d}{dt}_{|t=0} \mathbb{E}\big[\hat{\mathcal{U}}(X_t)\big] &= \mathbb{E}\big[\partial_\mu \mathcal{U}(\mu)(X_0) \cdot b(X_0)\big] \\ &= \int_{\mathbb{R}^d} \partial_\mu \mathcal{U}(\mu)(v) \cdot b(v) d\mu_0(v)\end{aligned}$$

2.3. **Second-order differentiability.** As we are dealing with stochastic processes, we need second-order derivatives of $\mathcal{U}$. However, asking the lift to be twice Fréchet is too strong: There would be very few examples of functions on the space of probability measures with a lift that is twice Fréchet (by the way, this is an interesting exercise...). Hence, we only discuss the existence of second-order partial derivatives. Basically, we require $\partial_\mu \mathcal{U}(\mu)(v)$ to be differentiable in v and μ and then

$$\partial_v \partial_\mu \mathcal{U}(\mu)(v) \qquad \partial^2_\mu \mathcal{U}(\mu)(v, v')$$

to be continuous in (μ, v, v') (for W_2 in μ) with suitable growth

Similarly as before, we have

$$\begin{aligned}\partial^2_{x_i\ x_j}\left[\mathcal{U}\left(\frac{1}{N}\sum_{k=1}^N \delta_{x_k}\right)\right] &= \frac{1}{N}\partial_v \partial_\mu \mathcal{U}\left(\frac{1}{N}\sum_{k=1}^N \delta_{x_k}\right)(x_i)\ \delta_{i,j} \\ &+ \frac{1}{N^2}\partial^2_\mu \mathcal{U}\left(\frac{1}{N}\sum_{k=1}^N \delta_{x_k}\right)(x_i, x_j)\end{aligned}$$

2.4. **Itô's formula on $\mathcal{P}_2(\mathbb{R}^d)$.** One key question to expand the right-hand side of (5) is to get an Itô formula (or a chain rule) on the space of probability measures. To make it clear, assume that we have a process satisfying the equation

$$dX_t = b_t dt + dW_t$$

with

$$\mathbb{E}\int_0^T |b_t|^2 dt < \infty.$$

Denoting by $\mu_t = \mathcal{L}(X_t)$ the marginal of X_t for every $t \in [0,T]$ and taking a map $\mathcal{U}$, differentiable in the sense of Lions as we explained right above, such that $\mathbb{R}^d \ni v \mapsto \partial_\mu \mathcal{U}(\mu)(v)$ is differentiable in v and μ, is it possible to expand

$$(\mathcal{U}(\mu_t))_{t\geq 0}?$$

One way to do so is to perform a space discretization, namely to approximate μ_t by a a finite cloud, or by a finite particle system. Intuitively, it becomes clear that the right candidate is the right-hand side below:

$$\mu_t \sim \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}$$

where $((X_t^i)_{t\geq 0})_{1\leq i\leq N}$ are independent copies of $(X_t)_{t\geq 0}$. One then may expand

$$\mathcal{U}\Big(\frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}\Big),$$

by using standard Itô's formula together with the dictionary given in last two paragraphs to write the derivative of the finite-dimensional projection of $\mathcal{U}$ in terms of the Wasserstein derivative. Passing to the limit in the formula, we get

$$d\mathcal{U}(\mu_t) = \mathbb{E}\big[b_t \cdot \partial_\mu \mathcal{U}(\mu_t)(X_t)\big] + \tfrac{1}{2}\mathbb{E}\big[\mathrm{Trace}\big(\partial_v \partial_\mu \mathcal{U}(\mu_t)(X_t)\big)\big]dt.$$

In case when $b_t = b(t, X_t)$, it is then an easy exercice to check that the right-hand side can be written in the sole term of μ_t (and so independently of the choice of the variable X_t used to lift μ_t). We get

$$d\mathcal{U}(\mu_t) = \Big(\int_{\mathbb{R}^d} b(t,x)\cdot \partial_\mu \mathcal{U}(\mu_t)(x) d\mu_t(x) + \tfrac{1}{2}\int_{\mathbb{R}^d}\big[\mathrm{Trace}\big(\partial_v\partial_\mu \mathcal{U}(\mu_t)(x)\big)\big]d\mu_t(x)\Big)dt.$$

2.5. **Connection with flat derivative.** What we have understood above is that the Wasserstein derivative arises when performing linear perturbations of the random variables that live above the probability measures. As we accounted for, another strategy is to define a notion of derivative based upon convex (or linear) perturbations on the space of measures itself. For instance, we may say that $\mathcal{V} : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ is $\mathcal{C}^1$ if

$$\frac{d}{d\varepsilon}_{|\varepsilon=0+} \mathcal{V}\big((1-\varepsilon)\mu + \varepsilon\mu'\big) = \underbrace{\int_{\mathbb{R}^d} \frac{\delta \mathcal{V}}{\delta m}(\mu)(v) d(\mu' - \mu)(v)}_{\frac{\delta \mathcal{V}}{\delta m}(\mu)(\cdot)\cdot(\mu'-\mu)} \tag{6}$$

for a continuous map

$$\frac{\delta \mathcal{V}}{\delta m} : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}$$

that is at most of quadratic growth in the variable v, say uniformly in μ (in order to guarantee the well-posedness of the integral). Obviously some care is needed: the derivative is uniquely defined up to an additive constant (for instance, we may choose the derivative that has zero mean with respect to μ).

A natural question is to address the connection with the previously defined Wasserstein derivative. The typical relationship is

$$\partial_\mu \mathcal{V}(\mu)(v) = \partial_v \frac{\delta \mathcal{V}}{\delta m}(\mu)(v).$$

And we know conditions under which equality holds true.

Exercise: Assume that $\mathcal{V}$ is flat differentiable and that the function

$$\partial_v \frac{\delta \mathcal{V}}{\delta m}(\mu)(v)$$

is jointly Lipschitz continuous (for the Wasserstein distance in the direction μ) and, for any two random variables X and Y with values in $\mathbb{R}^d$ (on a given probability space), expand

$$\begin{aligned}
&\mathcal{V}(\mathcal{L}(Y)) - \mathcal{V}(\mathcal{L}(X)) \\
&= \int_0^1 \mathbb{E}\Big[\frac{\delta \mathcal{V}}{\delta m}\big(\lambda \mathcal{L}(Y) + (1-\lambda)\mathcal{L}(X), Y\big) - \frac{\delta \mathcal{V}}{\delta m}\big(\lambda \mathcal{L}(Y) + (1-\lambda)\mathcal{L}(X), X\big)\Big] d\lambda \\
&= \int_0^1 \mathbb{E}\Big[\partial_v \frac{\delta \mathcal{V}}{\delta m}\big(\lambda \mathcal{L}(Y) + (1-\lambda)\mathcal{L}(X), \lambda' Y + (1-\lambda')X\big) \cdot \big(Y - X\big)\Big] d\lambda d\lambda'.
\end{aligned}$$

Deduce that (2.5) holds true.

3. Master equation

3.1. **Formal derivation of the master equation.** Back to (5), we may expand the right-hand side with respect to ϵ by using a variant of the chain rule derived in §2.4. The need for a variant is clear: Whilst $\mathcal{U}$ in the right-hand side of (5) depends on (t, x, μ), the chain rule proved in §2.4 just holds for functional a sole measure argument. The extension is easily taken: It consists in combining the standard version of Itô's formula (for functions of (t, x)) with the chain rule obtained in §2.4.

By expanding the right-hand side of (5), by subtracting $\mathcal{U}(t^\circ, x^\circ, \mu^\circ)$ in both sides, by dividing by ϵ and then by letting ϵ tend to 0, we finally obtain that $\mathcal{U}$ should satisfy the following equation

$$\begin{aligned}
&\partial_t \mathcal{U}(t,x,\mu) - \tfrac{1}{2}|\partial_x \mathcal{U}(t,x,\mu)|^2 + f(x,\mu) + \tfrac{1}{2}\mathrm{Trace}\Big(\partial_x^2 \mathcal{U}(t,x,\mu)\Big) \\
&\quad - \int_{\mathbb{R}^d} \partial_x \mathcal{U}(t,v,\mu) \cdot \partial_\mu \mathcal{U}(t,x,\mu,v) d\mu(v) + \tfrac{1}{2} \int_{\mathbb{R}^d} \mathrm{Trace}\Big(\partial_v \partial_\mu \mathcal{U}(t,x,\mu)(v)\Big) d\mu(v) = 0.
\end{aligned} \tag{7}$$

The boundary condition reads

$$\mathcal{U}(T, x, \mu) = g(x, \mu).$$

The first line in equation (7) accounts for the backward HJB equation (1). The second line accounts for the Fokker Planck equation (2). Somehow, the master equation encapsulates all the information underpinning the MFG system addressed in the first section. The reader who is aware of forward-backward stochastic differential equations should call $\mathcal{U}$ a decoupling field. The reader who is more familiar with PDEs should understand that we derived the equation for the system of characteristics (1)–(2): The master equation should be seen as a nonlinear PDE in infinite dimension.

Remark. Although the first line in (7) inherits the HJB structure of (1), the master equation is not a HJB equation on the space of probability measures. This is by the way an interesting question because, in some cases, it can be proved that the master equation derives in fact from an HJB equation on the space of probability measures.

3.2. **Potential case.** To see how the master equation may derive from a HJB equation on the space of probability, we may consider the so-called potential case. We refer to [7] for an earlier reference. More generally, this example makes the connection between mean field games and so-called mean field control. The idea is as follows and we consider a *social* optimization problem that consists in minimizing the cost functional

$$J(\boldsymbol{\alpha}) = G\big(\mathcal{L}(X_T)\big) + \int_0^T F\big(\mathcal{L}(X_t)\big)dt + \tfrac{1}{2}\mathbb{E}\int_0^T |\alpha_t|^2 dt$$

over dynamics of the form

$$dX_t = \alpha_t dt + \sigma dW_t.$$

Above, G and F are functions defined on $\mathcal{P}_2(\mathbb{R}^d)$. The reader should be aware of the difference with the mean field game: When α varies here, the law in J varies! This is absolutely different from the paradigm of mean field games.

With this optimization problem, we may associate a value function:

$$U(t,\mu) = \inf_{\alpha \text{ processes}} G\big(\mathcal{L}(X_T)\big) + \int_t^T F\big(\mathcal{L}(X_s)\big) + \tfrac{1}{2}\mathbb{E}\int_t^T |\alpha_s|^2 ds$$

over dynamics of the form

$$dX_s = \alpha_s dt + \sigma dW_t,$$

with $X_t \sim \mu$ as initial condition. Since U is the value function, it is expected to be a solution (in some sense) of a HJB equation, which here writes:

$$\partial_t U(t,\mu) - \tfrac{1}{2}\int_{\mathbb{R}^d} |\partial_\mu U(\mu)(v)|^2 d\mu(v) + \tfrac{1}{2}\int_{\mathbb{R}^d} \mathrm{Trace}\big(\partial_v \partial_\mu U(\mu)(v)\big) d\mu(v) + F(\mu) = 0$$

with G as terminal boundary condition.

This is then an interesting exercise to recover the master equation (7) by differentiating (in the flat sense!) the above equation with respect to μ, provided that f and g in the mean field game are chosen as

$$f(x,\mu) = \frac{\delta F}{\delta m}(\mu)(x), \quad g(x,\mu) = \frac{\delta G}{\delta m}(\mu)(x).$$

3.3. **Solving the master equation.** Solving the master equation is a difficult question. We refer to [1, 2, 4, 5, 7] for the main existing results of solvability.

The first question we may wonder is the notion of solution. Of course, one may ask for a classical solution, but the reader who knows PDE theory may also think of viscosity solutions. In short, the known results could be summarized as follows: Provided that the Lasry Lions monotonicity condition described in the first section is in force, $\mathcal{U}$ should solve the master equation in a viscosity sense provided that f and g satisfy mild regularity conditions. If f and g are still monotonous and satisfy stronger regularity conditions (including differentiability with respect to the measure argument), then $\mathcal{U}$ should should solve the master equation in the classical sense. In fact, the main restriction here is the monotonicity condition: It is very convenient because it ensures uniqueness of the solution to the mean field game in hand, but it has a limited scope in practice. Also, this remains a quite general open question so far to say something about the master equation when uniqueness does not hold.

Here is a typical form of result, see for instance [4]. Assume that f and g are continuously differentiable in (x,μ). Assume also that the derivatives of f (and similarly for g) satisfies:

- $\partial_x f(x,\mu)$ is bounded and Lipschitz in (x,μ);
- $\partial_\mu f(x,\mu)(v)$ is bounded and Lipschitz in (x,μ,v).

Require also the existence and the smoothness of the second-order derivatives of f and g and assume the derivatives of f (similarly for g) satisfy:

- $\partial_x^2 f(x,\mu)$ bounded and Lipschitz in (x,μ);
- $\partial_\mu f(x,\mu)(v)$ is differentiable in x, v and μ;
- $\partial_x\partial_\mu f(x,\mu)(v)$, $\partial_v\partial_\mu f(x,\mu)(v)$ are bounded and Lipschitz;
- $\partial_\mu^2 f(x,\mu)(v,v')$ is bounded and Lipschitz

Then, $\mathcal{U}$ is a classical solution to the mean field game, with the same smoothness as f and g.

3.4. **Road map to regularity of $\mathcal{U}$.** We now explain the main lines of the proof, as given in [2]. Therein, the analysis is reduced to the torus $\mathbb{T}^d$ to avoid any problem at infinity. The main point is to understand the smoothness of $\mathcal{U}$ with respect to μ: In the finite dimensional direction x, things are simpler because of the Laplace term. To make it clear, we are thus interested in

$$\frac{d}{d\varepsilon}_{|\varepsilon=0+}\mathcal{U}\big(t_0,\cdot,(1-\varepsilon)\mu+\varepsilon\mu'\big) \quad \mu,\ \mu'\in\mathcal{P}(\mathbb{T}^d),$$

where $\mathcal{P}(\mathbb{T}^d)$ is the set of probability measures on the torus. Here, t_0 is an initial time in $[0,T]$. Back to (4), we thus focus on

$$\frac{d}{d\varepsilon}_{|\varepsilon=0+}u^{(1-\varepsilon)\mu+\varepsilon\mu'}(t_0,\cdot),$$

where the superscript in u denotes the initial state of the population at time t_0. The whole point to address the term right above is to linearize the MFG system (1)–(2). This prompts us to let

$$\begin{aligned} z_t &= \underbrace{\frac{d}{d\varepsilon}_{|\varepsilon=0+}u^{(1-\varepsilon)\mu+\varepsilon\mu'}(t,\cdot)}_{\text{function}}, \\ m_t &= \underbrace{\frac{d}{d\varepsilon}_{|\varepsilon=0+}\mu_t^{(1-\varepsilon)\mu+\varepsilon\mu'}}_{\text{distribution}}. \end{aligned} \tag{8}$$

The first term is seen as a function, and the second one as a distribution. The guess is that the pair (z,m) should satisfy the linearized system

$$\begin{aligned} &\partial_t m_t - \tfrac{1}{2}\Delta m_t - \operatorname{div}\Big(m_t\partial_x u(t,x) + \mu_t\partial_x z(t,x)\Big) = 0 \\ &\partial_t z(t,x) + \tfrac{1}{2}\Delta z(t,x) - \partial_x u(t,x)\cdot\partial_x z(t,x) + \underbrace{\frac{\delta f}{\delta m}(x,\mu_t)(\cdot)\cdot m_t(\cdot)}_{\text{balance reg in } v\ /\ \text{singularity } m} = 0 \\ &z_T(x) = \frac{\delta g}{\delta m}(x,\mu_T)(\cdot)\cdot m_T(\cdot). \end{aligned} \tag{9}$$

Since m_t is a distribution, $\delta f/\delta m$ needs to be smooth enough in the argument $\cdot$, as mentioned right below the corresponding term.

The key point in the proof is to show that (9) is well-posed and that its solution is indeed the derivative of the original MFG system (1)–(2), see (8). Well-posedness is here derived by using monotonicity of f and g again and, in fact, monotonicity provides more: It gives a form of stability on the linearized system with permits to check that the solution indeed satisfies (8).

We then observe that, in (9), we can play for free with the initial condition of the forward equation. For instance, we may choose $m_{t_0} = \delta_v$, in which case we may call $\mathcal{V}(t_0, x, \mu_0)(v)$ the value of $z(t_0, x)$. By linearity, we get that, if the initial condition is a finite signed measure m_{t_0}, then the corresponding $z(t_0, \cdot)$ in (9) takes the form:

$$z(t_0, x) = \int_{\mathbb{T}^d} \mathcal{V}(t_0, x, \mu_0)(v) dm_{t_0}(v).$$

Thinking of m_{t_0} as $\mu' - \mu$, the above is the analogue of (6). Equivalently, $\mathcal{V}(t_0, x, \mu_0)(v)$ should be thought as $[\delta\mathcal{U}/\delta m](t_0, x, \mu_0)(v)$.

References

[1] P. Cardaliaguet. Notes on mean field games. https://www.ceremade.dauphine.fr/∼cardalia/MFG20130420.pdf/, 2013.

[2] P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions. *The master equation and the convergence problem in mean field games*, volume 201 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2019.

[3] R. Carmona and F. Delarue. *Probabilistic theory of mean field games with applications. I*, volume 83 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2018. Mean field FBSDEs, control, and games.

[4] R. Carmona and F. Delarue. *Probabilistic theory of mean field games with applications. II*, volume 84 of *Probability Theory and Stochastic Modelling*. Springer, Cham, 2018. Mean field games with common noise and master equations.

[5] J.F. Chassagneux, D. Crisan, and F. Delarue. A Probabilistic approach to classical solutions of the master equation for large population equilibria. To appear in Memoirs of the AMS.

[6] W. Fleming and M. Soner. *Controlled Markov Processes and Viscosity Solutions*. Springer Verlag, 2010.

[7] W. Gangbo and A. Swiech. Existence of a solution to an equation arising from the theory of mean field games. J. Differential Equations 259, no. 11, 6573–6643, 2015.