SELECTED TOPICS IN MEAN FIELD GAMES

PIERRE CARDALIAGUET AND FRANÇOIS DELARUE

ABSTRACT

Mean field game theory was initiated a little more than 15 years ago with the aim of simplifying the search for Nash equilibria in games with a large number of weakly interacting players. Since then, a lot has been done. Numerous equilibrium existence results have been obtained, using different characterizations and in various contexts. The analysis of the master equation, which describes the evolution of the value of the game, has also seen significant progress, which has, for example, allowed establishing in certain cases the convergence of games with a finite number of players. However, mean field games remain of a complex nature. For instance, the typical lack of uniqueness of solutions raises selection issues that are still poorly understood. The objective of the note is to present some of the latest advances, as well as some avenues to address further challenging questions.

MATHEMATICS SUBJECT CLASSIFICATION 2020

Primary 91A13; Secondary 91A15, 60H30

KEYWORDS

Mean field games, master equation, convergence problem, common noise, long time behaviour

INTERNATIONAL CONGRESS OF MATHEMATICIANS 2022 JULY 6–14

Proc. Int. Cong. Math. 2022, Vol. 5, pp. 3660–3703
DOI 10.4171/ICM2022/17

Published by EMS Press

The theory of mean field games (MFG) aims at providing an asymptotic description of differential games with a large number of interacting players. The number of applications of the theory is huge, ranging from macroeconomics to crowd motions, and from finance to power grid models. In all these models, each player controls his/her own dynamical state which evolves in time according to a deterministic or stochastic differential (or difference if the state space is at most countable) equation. The individual goal is to minimize some cost depending on his/her own control but also on the behavior of the whole population of agents, which is described through the empirical distribution of their states. In this setting, the central concept is the notion of Nash equilibria, which explains how agents play in an optimal way by taking into account the others' strategies. The MFG theory is precisely intended to simplify the search for these Nash equilibria. In this respect, the key idea is to postulate that, asymptotically, the single theoretical (and not empirical) statistical distribution of the states is sufficient to compute the individual goal of each player.

The MFG theory was introduced and largely developed by Lasry and Lions through a series of papers around 2005 and during the famous lectures of Lions at the Collège de France [85–88]. At about the same time, Caines, Malhamé, and Huang discussed similar models under the terminology of "Nash certainty equivalence principle" [73,74]. The MFG theory is also reminiscent of the so-called heterogenous agent models developed in economics at the end of the 1990s by Aiyagari [7] and by Krusell and Smith [79] or, more recently, by Lucas and Moll [91]. One of the main achievements of the MFG theory—though not discussed here—is a better formulation and understanding of these models (see, for instance, Achdou et al. [1]). After a decade and a half of research, the theory has answered—at least partially—several important questions and has developed a number of mathematical techniques and tools for this purpose. A large part of the material can be found in the monographs or in the surveys [6,15,25,35,36,71].

From a mathematical perspective, the MFG theory lies at the intersection of probability and partial differential equations (PDEs). The connection between games with finitely many players and MFGs is addressed by means of statistical averaging arguments, which are made possible by the symmetric structure of the interactions. This approach is, of course, reminiscent of the very typical issues and techniques underpinning the standard mean field theory and the related propagation of chaos properties for large weakly interacting particle systems (see [76,92] for the earliest papers in the field and [101] for a review). However, unlike the standard mean field theory, in which the interacting particles obey a given dynamics, the dynamics of the agents is not given a priori in the MFG theory but rather is obtained after an optimization procedure. This seemingly innocuous difference dramatically increases the level of complexity of the problem, as it introduces several nonlinearities in the equations describing the mean field models. These nonlinearities manifest themselves in several ways, depending on the formulation used to characterize the equilibria and, implicitly, on the approach chosen to manage the optimization step in the definition of these equilibria. In this respect, let us say that both probabilistic and PDE arguments have been successfully developed. In short, the probabilistic approach aims at following the dynamics of a reference player in the population, while the PDE one aims at following the dynamics of the statistical state of the whole population. The key feature is that both approaches lead to the study of a form of forward–backward system that couples either two stochastic differential equations or two PDEs: the probabilistic system is usually referred to as a forward–backward McKean–Vlasov system and the PDE system is usually known as the "MFG system." Regardless of the system, the strong coupling between the forward and backward components therein raises many issues. Obviously, one knows in general how to pass from one approach to the other and, generally speaking, PDE tools are useful to obtain better regularity of the solutions. Very importantly, these two systems can be regarded as the characteristics of a common infinite-dimensional PDE of hyperbolic type set on the space of probability measures. It is called the "master equation." This master equation has become a challenging object in the field and has attracted much attention in analysis, probability, and calculus of variation. At present, it is only well-understood in certain cases where the solutions are known to be regular. A theory allowing less regular solutions and thus covering a wider scope is totally lacking. Needless to say, this is a very exciting area of research.

In addition to the analysis of mean field games themselves, the study of the convergence problem, namely the convergence of games with a finite number of players to a mean field game, is another challenge, which has also required the development of appropriate arguments. As already mentioned, this asks for a nontrivial adaptation of the existing results on the convergence of weakly interacting particle systems. Among others, a key idea is to test classical solutions of the master equation onto the equilibria of the games with finitely many players. The main contributions in this direction are presented in the notes, but many questions remain open. To wit, solutions to mean field games are typically nonunique and identifying those that are selected by taking the limit in large games is a fascinating, but really difficult question.

Before presenting the rest of the contents of these notes in a more exhaustive way, we insist on the fact that the MFG theory provides a concept that has proven to be effective in the analysis of some typical examples of game theory. However, the same concept can be applied to many other cases. We give an overview of some of them at the very end of the notes. For example, mean field games with common noise is an extension of the original concept that has stimulated many recent works. In short, this corresponds to the case where the state of the population itself is random. Understanding the precise impact of noise on equilibria is another challenge in the field. To emphasize the importance of this research direction, we have therefore decided to write these notes by systematically including common noise in the models we present. We hope that this will help the reader to grasp the essence of it.

Contents. After a short presentation of the PDE formulation of MFGs in Section 1, we concentrate ourselves on the following three fundamental aspects of the theory:

1) The analysis of the convergence problem, which, as we have said, investigates how Nash equilibria in differential games with finitely many players converge to MFG equilibria. This point is essential to justify the MFG models and is one of the main mathematical achievement of the MFG theory. We provide an overview in Section 2, which includes a presentation of the master equation.

2) The long time behavior of the MFG equilibria. Since time-dependent models are difficult to handle and to approximate numerically, the analysis of "stationary models" and their robustness is essential in both theory and application. For instance, economists often concentrate on these stationary solutions. We present the main results in Section 3.

3) The regularizing aspect of the common noise in MFG. Since the MFG equilibria are in general not unique, it is crucial to understand the extent to which a common noise can force uniqueness. This question is addressed in Section 4.

We complete the notes by providing in the final Section 5 a general overview of other topics from MFG theory that are not discussed in the first four sections. We give some references that may be useful for the reader and we provide some open problems.

Notation. We denote by $\mathcal{P}_2(\mathbb{R}^d)$ the set of Borel probability measures on $\mathbb{R}^d$ with a finite second order moment, endowed with the Wasserstein distance (see, for instance, [9]). If $x \in \mathbb{R}^d$, we denote by δ_x the Dirac mass at x. For X a random variable, we denote by $\mathcal{L}(X)$ the law of X.

1. THE MFG EQUILIBRIA

In this section we introduce the main problems of the MFG theory. The simplest for this is to start with a game with a large number of players and then to pass (at least formally) to the limit as the number of players tends to infinity.

1.1. The N-player problem

The N-player game. Let $N \in \mathbb{N}$, with $N \geq 1$ being the (large) number of players. Player i (where $i \in \{1, \dots, N\}$) controls her own state X_t^i, which is an element of $\mathbb{R}^d$ and evolves in time according to the stochastic differential equation (SDE)

$$dX_t^i = \alpha_t^i + \sqrt{2} dB_t^i + \sqrt{2\epsilon} dW_t,$$

for prescribed initial conditions $(X_0^i)_{i=1,\dots,N}$. Here the processes $((B_t^i)_{t\geq 0})_{i=1,\dots,N}$ and $(W_t)_{t\geq 0}$ are independent d-dimensional Brownian motions. The noise $(B_t^i)_t$, which affects only the dynamics of player i, is called the idiosyncratic (or the individual) noise. The Brownian motion $(W_t)_t$, on the contrary, impacts all the dynamics and is called the common noise; the nonnegative real ϵ denotes (up to the square root) the intensity of the effective common noise that is felt by all the players. The initial conditions $(X_0^i)_{i\geq 1}$ are independent and identically distributed (i.i.d.) random variables with common distribution $\tilde{m}_0 \in \mathcal{P}_2(\mathbb{R}^d)$. We assume that the random variables $(X_0^i)_{i=1,\dots,N}$ and the Brownian motions $((B_t^i)_t)_{i=1,\dots,N}$ and W are independent. Player i chooses a bounded control $(\alpha_t^i)_t$ that takes values in $\mathbb{R}^d$ and that is adapted to the filtration $(\mathbb{F}_t = \sigma\{X_0^j, B_s^j, W_s, s \leq t, j = 1, \dots, N\})$.

The cost of player i is given by

$$\mathcal{J}_i^N\big(\alpha^i,(\alpha^j)_{j\neq i}\big) = \mathbb{E}\bigg[\int_0^T \Big(\frac{1}{2}|\alpha_t^i|^2 + F(X_t^i, m_{\boldsymbol{X}_t}^{N,i})\Big)dt + G(X_T^i, m_{\boldsymbol{X}_T}^{N,i})\bigg],$$

where $\boldsymbol{X}_t = (X_t^1,\dots,X_t^N)$ and $m_{\boldsymbol{X}_t}^{N,i} = \frac{1}{N-1}\sum_{j=1,\dots,N, j\neq i}\delta_{X_t^j}$. To fix the ideas, we work here with a finite-horizon problem (where $T>0$ is the horizon) and we assume the maps $F:\mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d)\to\mathbb{R}$ and $G:\mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d)\to\mathbb{R}$ to be continuous and bounded. Here we make the assumption that the running cost of player i depends only on her own control, her own position, and on the distribution of the other players' positions, while the terminal cost depends only on her position and on the distribution of the other players' positions at terminal time. The important point is the symmetry of the problem: for a player, the other players play exactly the same role. The specific form of the cost and dynamics is made here for simplicity.

Nash equilibria. In that setting, a natural notion of equilibrium is the the notion of *Nash equilibrium*. We say that a family $(\bar{\alpha}^1,\dots,\bar{\alpha}^N)$ of (time-dependent stochastic) controls is a Nash equilibrium of the N-player game if, for any $i\in\{1,\dots,N\}$ and any control α^i,

$$\mathcal{J}_i^N\big(\bar{\alpha}^i,(\bar{\alpha}^j)_{j\neq i}\big) \le \mathcal{J}_i^N\big(\alpha^i,(\bar{\alpha}^j)_{j\neq i}\big).$$

We are intentionally fuzzy in the definition of what a control is. There are actually many possibilities and we feel better to restrict ourselves to two of them. The controls can be either (i) open-loop, which means that they are regarded as adapted functions of the initial conditions $(X_0^i)_i$ and of the noises $((B_t^i)_t)_i$ and W, or (ii) closed-loop controls, in which case they are considered as adapted functions of the trajectories $((X^i)_t)_i$ (when the closed-loop structure is Markov, the dependence just occurs through the current states of the players). The main difference between the two notions is as follows: when one player deviates, the function underpinning the definition is kept fixed. As such, the controls played by the other players remain the same in the open-loop case while they change in the closed-loop case. In the rest of the note, we always mean Markov closed-loop control when speaking about a closed-loop control.

The Nash system. A key fact with games involving closed-loop controls is that they have a PDE interpretation, in the form of a system of equations for the equilibrium value of the game. In our setting, one can show that if $v^N:[0,T]\times(\mathbb{R}^d)^N\to\mathbb{R}^N$ is the classical solution to the following backward parabolic system (called here the Nash system)

$$\begin{cases} -\partial_t v_t^{N,i} - \displaystyle\sum_{j=1}^N \Delta_{x_j} v_t^{N,i} - \epsilon \sum_{j,k=1}^N \mathrm{Tr}(D^2_{x_j x_k} v_t^{N,i}) + \frac{1}{2}\big|D_{x_i} v_t^{N,i}\big|^2 \\ \qquad + \displaystyle\sum_{j\neq i} D_{x_i} v_t^{N,i}\cdot D_{x_j} v_t^{N,j} = F(x_i, m_{\boldsymbol{x}}^{N,i}) \quad \text{in } (0,T)\times(\mathbb{R}^d)^N, \quad i\in\{1,\dots,N\}, \\ v_T^{N,i}(\boldsymbol{x}) = G(x_i, m_{\boldsymbol{x}}^{N,i}) \qquad\qquad\qquad\quad \text{in } (\mathbb{R}^d)^N, \quad i\in\{1,\dots,N\}, \end{cases} \tag{1.1}$$

then $(\bar{\alpha}^i(t, \boldsymbol{x}) := -D_{x_i} v_t^{N,i}(\boldsymbol{x}))_{i=1,\ldots,N}$ is a Nash equilibrium of the N-player game in closed-loop form. Here, the notation $\boldsymbol{x}$ stands for an N-tuple $(x_1, \ldots, x_N)$, in which case x_i is the entry number i in $\boldsymbol{x}$. The existence and uniqueness of the solution to the above system, called the Nash system, is classical under suitable assumptions on F and G and discussed, for instance, in [84]. Under similar assumptions on F and G, this equilibrium can be shown to be unique (within the class of bounded Markov closed-loop controls); see, for instance, [36, CHAPTER 6].

The main question raised by MFG theory is the characterization of the limit, as N tends to infinity, of the Nash equilibria of the game (or of the Nash system) and the analysis of the resulting limit.

1.2. The MFG equilibria

In this part we derive from the N-player problem several (equivalent) formulations of an MFG equilibrium. The derivation is formal at this stage, but will be justified more rigorously in Section 2.

MFG equilibria without common noise ($\epsilon = 0$). The are several ways to guess and write the limit of the Nash equilibrium or of the Nash system as N tends to infinity. We start with the problem without common noise, which is easier to grasp. As players are symmetric, one can expect, using classical ideas of mean field theory [101], that the in-equilibrium trajectories $((\overline{X}_t^{N,i})_t)_i$ associated with the Nash equilibrium identified right above become more and more decorrelated as N increases and eventually become asymptotically independent. In this case the empirical measure $m^{N,i}_{\overline{\boldsymbol{X}}_t^N}$ should become asymptotically deterministic and, as N gets larger and larger, the impact of the deviation of a player over $m^{N,i}_{\overline{\boldsymbol{X}}_t^N}$ should be negligible. Therefore players can solve their own optimization problem as if $m^{N,i}_{\overline{\boldsymbol{X}}_t^N}$ were given and independent of i. Implementing this idea, one finds the notion of MFG equilibrium in its probabilistic formulation:

Probabilistic formulation of the MFG equilibrium ($\epsilon = 0$). One searches for a pair (m, α), where $m = (m_t)_t \in C^0([0, T], \mathcal{P}_2(\mathbb{R}^d))$, and $\alpha = (\alpha_t)_t$ is a control such that

(i) α is optimal for the control problem

$$\inf_{\beta} \mathbb{E}\left[\int_0^T \left(\frac{1}{2}|\beta_t|^2 + F(X_t^\beta, m_t)\right) dt + G(X_T^\beta, m_T)\right], \tag{1.2}$$

where the infimum is taken over the controls $\beta = (\beta_t)_t$ (that are $(X_0^1, (B_t^1)_t)$-progressively measurable) and where X^β is the solution to

$$dX_t^\beta = \beta_t dt + \sqrt{2} dB_t^1, \quad X_0^\beta = X_0^1. \tag{1.3}$$

(ii) For any $t \in [0, T]$, the law of X_t^α is m_t.

Other probabilistic formulations of MFG equilibria are possible: Carmona and Delarue discuss in [34] a formulation involving the stochastic maximum principle. Mainly, the optimizers in item (i) are described by means of a forward–backward stochastic differential equation depending on the input $(m_t)_t$. Under the fixed point condition (ii), this input is identified with the marginal law of the solution of the forward equation, which gives rise to a so-called forward–backward system of McKean–Vlasov type. While Pontryagin's principle provides the dynamics of an equilibrium feedback along a corresponding equilibrium trajectory, an alternative approach is to provide a representation of the equilibrium value. This approach is usually known as the weak formulation as it may rely on a convenient change of noise in the dynamics. In short, it provides another form of the forward–backward system of McKean–Vlasov type, see Carmona and Lacker [39] and [35, CHAPTER 3]. The latter is useful for proving existence results. In comparison, the stochastic Pontryagin principle provides, in general, only sufficient conditions satisfied by an arbitrary equilibrium.

PDE formulation of the MFG equilibria: the MFG system ($\epsilon = 0$). Another characterization of the MFG equilibria goes through a forward–backward system of PDEs known as the MFG system: the unknown are (u, m) where u corresponds to the value function associated with the optimal control problem described in the probabilistic formulation while m solves the Kolmogorov equation satisfied by the marginal law of the equilibrium. It reads therefore

$$\begin{cases} -\partial_t u_t(x) - \Delta u_t(x) + \frac{1}{2}\left|Du_t(x)\right|^2 = F(x, m_t) & \text{in } (0,T) \times \mathbb{R}^d, \\ \partial_t m_t(x) - \Delta m_t(x) - \operatorname{div}(m_t(x) Du_t(x)) = 0 & \text{in } (0,T) \times \mathbb{R}^d, \\ m_0(x) = \tilde{m}_0, \quad u_T(x) = G(x, m_T) & \text{in } \mathbb{R}^d. \end{cases} \tag{1.4}$$

This system is unusual: the first equation (a Hamilton–Jacobi equation) is backward in time, while the Kolmogorov equation is forward in time. The main issue is that both equations are strongly coupled, in the sense that each of the two unknowns shows up in the other equation. Since the two equations are set in opposite time directions, this creates a conflict which makes the spice of the analysis. The existence of a solution has been proved by Lasry and Lions [85–87] under suitable assumptions on the coupling functions F and G (regularity and growth conditions). In general, there is no uniqueness: this is a typical feature of equilibria in game theory (in contrast, uniqueness holds in the finite game because of the smoothing effect of the Laplacians in the related Nash system (1.1); we will come back to this observation in Section 4). However, the solution of (1.4) is unique if the following monotonicity condition, introduced in [85–87], is satisfied:

$$\begin{aligned} &\int_{\mathbb{R}^d} \big(F(x,m) - F(x,m')\big)(m - m')(dx) \geq 0, \\ &\int_{\mathbb{R}^d} \big(G(x,m) - G(x,m')\big)(m - m')(dx) \geq 0. \end{aligned} \tag{1.5}$$

There is by now a huge literature on the MFG system, including different types of coupling functions, different types of boundary conditions, etc. We briefly present some aspects of this literature in Section 5.1.

The MFG equilibria with common noise ($\epsilon > 0$). In the presence of common noise, the heuristic analysis of the limit problem is more subtle. Indeed, even if the players do not take into account the idiosyncratic noises of the other players, their dynamics are perturbed by the common noise $(W_t)_t$. Therefore, the limit $(m_t)_t$ (if it exists) of the marginal empirical measure $(m^{N,i}_{\boldsymbol{X}^N_t})_t$ associated with the equilibrium trajectories $((\overline{X}^{N,j}_t)_t)_{j \neq i}$ becomes random and is typically expected to be adapted to the Brownian motion $(W_t)_t$ (very much as before, this limit is expected to be independent of i).

Probabilistic formulation of the MFG equilibrium with common noise ($\epsilon > 0$). One searches for a pair (m, α), where the stochastic process $(m_t)_t$ is adapted to $(W_t)_t$ and takes values in $C^0([0,T], \mathscr{P}_2(\mathbb{R}^d))$ and $\alpha = (\alpha_t)_t$ is a control such that

(i) α is optimal for the control problem

$$\inf_{\beta} \mathbb{E}\left[\int_0^T \left(\frac{1}{2}|\beta_t|^2 + F(X_t^\beta, m_t)\right)dt + G(X_T^\beta, m_T)\right], \tag{1.6}$$

where the infimum is taken over the controls $\beta = (\beta_t)_t$ (that are $(X_0^1, (B_t^1, W_t)_t)$-progressively measurable) and where X^β is the solution to

$$dX_t^\beta = \beta_t dt + \sqrt{2} dB_t^1 + \sqrt{2\epsilon} W_t, \quad X_0^\beta = X_0^1. \tag{1.7}$$

(ii) For any $t \in [0,T]$, the (conditional) law of X_t^α given $(W_s)_s$ is m_t.

In general, it is difficult to prove the existence of MFG equilibria because the fixed-point condition (ii) is defined, in the presence of common noise, on a very wide space. To overcome this issue, a possible path is to discretize the common noise into a noise with finitely many outcomes (see [38]). In that case, it is much easier to adapt the arguments used when $\epsilon = 0$. However, much may be lost when passing to the limit over the discretization of the common noise. Very similar to weak solutions to stochastic differential equations, equilibria that are obtained in this way may no longer be adapted with respect to the original common noise $(W_t)_t$. This requires a relevant notion of weak MFG equilibria, in which the flow of measures $(m_t)_t$ is adapted to a larger filtration than that generated by $(W_t)_t$. When the monotonicity property (1.5) is in force, it can be proved that these weak solutions are in fact strong, i.e., they are adapted with respect to $(W_t)_t$.

PDE formulation of the MFG equilibria with common noise: the stochastic MFG system ($\epsilon > 0$). In the probabilistic formulation of the MFG equilibria with a common noise, the optimal control problem (1.6)–(1.7) (which is solved by a reference player in the population) is driven by random coefficients (because $(m_t)_t$ is random). The associated value function is no longer deterministic. Following Peng [96], it should be regarded as the solution of a backward stochastic Hamilton–Jacobi equation. Moreover, the Kolmogorov equation satisfied by the random flow $(m_t)_t$ is stochastic. The resulting MFG system therefore reads:

$$
\begin{cases}
du_t = \Big[-(1+\epsilon)\Delta u_t + \frac{1}{2}|Du_t|^2 - F(x,m_t) - 2\epsilon \operatorname{div}(v_t)\Big]dt \\
\qquad\qquad + v_t \cdot \sqrt{2\epsilon}\, dW_t & \text{in } (0,T)\times\mathbb{R}^d, \\
dm_t = \big[(1+\epsilon)\Delta m_t + \operatorname{div}(m_t Du_t)\big]dt - \operatorname{div}(m_t\sqrt{2\epsilon}\, dW_t) & \text{in } (0,T)\times\mathbb{R}^d, \\
m_0(x) = \tilde{m}_0, \quad u_T(x) = G(x, m_T) & \text{in } \mathbb{R}^d.
\end{cases}
\tag{1.8}
$$

Note that now the unknown is the triplet (u, v, m). As explained in Peng [96], the role of the random field v is to ensure the solution u to the backward Hamilton–Jacobi equation to be adapted to the common noise $(W_t)_t$. The existence of a solution for (1.8) is subtle and has been achieved, under suitable conditions on F and G including monotonicity, in [25] (see also [36]).

2. THE MASTER EQUATION AND THE CONVERGENCE OF THE NASH SYSTEM

In this part we address the rigorous derivation of the MFG equilibria and the convergence of the Nash system. This analysis requires the introduction of a new equation, the master equation, which is a nonlinear equation stated on the infinite-dimensional space $\mathcal{P}_2(\mathbb{R}^d)$. In order to restrict the technicality of the exposition, we will often be fuzzy in the assumption and in the statement of the results and refer to [25,36], that we follow closely, for details.

2.1. Derivatives of maps defined on the space of probability measures

There are several notions of derivatives for a map $U : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$: we refer, for instance, to [8,9,25,35] and the references therein for several possible notions together with an overview of the connections between all of them. Here we mostly discuss an idea of Lions which consists in lifting the map U to a suitable space of random variables.

Let us consider the space $L^2 := L^2((\Omega, \mathbb{F}, \mathbb{P}), \mathbb{R}^d)$ of square-integrable random variables on $\mathbb{R}^d$, with Ω being a Polish space, $\mathbb{F}$ its Borel σ-algebra, and $\mathbb{P}$ an atomless probability measure. The space L^2 is endowed with the usual Hilbert scalar product. It is known that, for any $m \in \mathcal{P}_2(\mathbb{R}^d)$, there exists a random variable X with law m.

Given a map $U : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$, we lift U to L^2 by setting

$$
\tilde{U}(X) = U\big(\mathcal{L}(X)\big) \quad \forall X \in L^2.
$$

Definition 2.1 (The L-derivative). We say that U is L-differentiable at $m \in \mathcal{P}_2(\mathbb{R}^d)$ if there exists a random variable $X \in L^2$ with law m such that $\tilde{U}$ is Fréchet differentiable at X (we denote by $\nabla\tilde{U}(X)$ its gradient).

Theorem 2.1 (Structure of the L-derivative). *Assume that U is L-differentiable at $m \in \mathcal{P}_2(\mathbb{R}^d)$. Then there exists a map $D_m U(m,\cdot) : \mathbb{R}^d \to \mathbb{R}^d$ which is Borel measurable* *and such that*

$$\nabla \tilde{U}(X) = D_m U(m, X)$$

for any random variable X with $\mathcal{L}(X) = m$. We call the map $D_m U(m, \cdot)$ the L-derivative of U at m.

The first version of this result goes back to Lions [88]. The version given here is due to Gangbo and Tudorascu [65], who also explain the connection with the notion of subdifferential introduced in [9].

Finite dimensional projection. A key principle to establish the link between the Nash system and the master equation is to associate with any function defined on $\mathcal{P}_2(\mathbb{R}^d)$ a finite dimensional projection, whose definition is as follows.

Given a continuous map $U : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ and a nonzero integer N, we define the projection U^N of U as the map $U^N : (\mathbb{R}^d)^N \to \mathbb{R}$ defined by

$$U^N(x_1, \dots, x_N) = U(m_{\boldsymbol{x}}^N), \quad \text{where } m_{\boldsymbol{x}}^N := \frac{1}{N} \sum_{i=1}^N \delta_{x_i}, \ \boldsymbol{x} = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N.$$

The following statement clarifies the meaning of the derivative D_m:

Proposition 2.2. *Assume that U is L-differentiable with a Lipschitz-continuous derivative. Then U^N is of class C^1 and*

$$D_{x_i} U^N(x_1, \dots, x_N) = \frac{1}{N} D_m U(m_{\boldsymbol{x}}^N, x_i),$$

for $(x_1, \dots, x_N) \in (\mathbb{R}^d)^N$.

One can, of course, introduce higher-order derivatives of a map $U : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ in a similar way and extend Proposition 2.2 to higher-order derivatives, see [35, CHAPTER 5].

Itô's formula along a flow of conditional measures. The following Itô's formula, needed in the proofs below and of independent interest, is a generalization of Itô rule for flows of measures and functions defined on the space of measures. Let $(X_t)_{t \geq 0}$ be an Itô process of the form

$$dX_t = b_t dt + \sigma_t dB_t + \sigma_t^0 dW_t, \quad t \geq 0, \tag{2.1}$$

with a given (possibly random) initial condition X_0, where $(B_t)_t$ and $(W_t)_t$ are two d-dimensional Brownian motions, X_0, $(B_t)_t$, and $(W_t)_t$ being independent. Above, $(b_t)_t$, $(\sigma_t)_t$, and $(\sigma_t^0)_t$ are progressively-measurable processes with respect to the filtration generated by X_0, $(B_t)_t$, and $(W_t)_t$, with values in (respectively) $\mathbb{R}^d$, $\mathbb{R}^{d \times d}$, and $\mathbb{R}^{d \times d}$. For simplicity, we assume that that the probability space is given in a product form $(\Omega_0 \times \Omega_1, \mathbb{F}_0 \otimes \mathbb{F}_1, \mathbb{P}_0 \otimes \mathbb{P}_1)$, where $(\Omega_0, \mathbb{F}_0, \mathbb{P}_0)$ supports W, while $(\Omega_1, \mathbb{F}_1, \mathbb{P}_1)$ supports (X_0, B). We denote by $\mathbb{E}^0$ the expectation with respect to $\mathbb{P}^0$ and by $\mathbb{E}^1$ the expectation with respect to $\mathbb{P}^1$. We assume that

$$\mathbb{E}\left[|X_0|^2 + \int_0^T \left(|b_t|^2 + |\sigma_t|^4 + \left|\sigma_t^0\right|^4 \right) dt \right] < +\infty,$$

where $\mathbb{E} = \mathbb{E}^0 \mathbb{E}^1$.

The following result is taken from [36] (see also [20, 45]).

Theorem 2.3. *Let $(X_t)_t$ be as in (2.1) and, for any $t \geq 0$, let m_t be the conditional law of X_t given $(W_s)_s$. Then, for $U : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ a sufficiently smooth mapping on $\mathcal{P}_2(\mathbb{R}^d)$,*

$$\begin{aligned} U(m_t) = U(m_0) &+ \int_0^t \mathbb{E}^1\big[D_m U(m_s, X_s) \cdot b_s ds\big] + \int_0^t \mathbb{E}^1\big[(\sigma^0)^* D_m U(m_s, X_s)\big] \cdot dW_s \\ &+ \frac{1}{2}\int_0^t \mathbb{E}^1\big[\mathrm{Tr}\big(D_y D_m U(m_s, X_s)\big(\sigma_s \sigma_s^* + \sigma_s^0 (\sigma_s^0)^*\big)\big]ds \\ &+ \frac{1}{2}\int_0^t \mathbb{E}^1 \tilde{\mathbb{E}}^1\big[\mathrm{Tr}\big(D_m^2 U(m_s, X_s, \tilde{X}_s)\sigma_s^0 (\tilde{\sigma}_s^0)^*\big)\big]ds, \end{aligned}$$

where $\tilde{X}_s$ and $\tilde{\sigma}_s^0$ are independent copies of X_s and σ_s^0 defined on $\Omega_0 \times \tilde{\Omega}_1$, for a copy $\tilde{\Omega}_1$ of Ω_1 equipped with the expectation $\tilde{\mathbb{E}}^1$.

Briefly, $D_y D_m U$ is the y-derivative of the function $y \mapsto D_m U(m, y)$ for a fixed m. Similarly, D_m^2 is the m-derivative of the function $m \mapsto D_m U(m, y)$ for a fixed y, which implies that $D_m^2 U$ can be written in the form $(m, y, y') \mapsto D_m^2 U(m, y, y')$. Under the regularity assumptions mentioned in the statement, all these derivatives exist implicitly and are jointly continuous. They also satisfy appropriate growth conditions that permit giving a meaning to the various expectations appearing in the expansion. The symbol Tr is for the trace.

Potential games. We feel it useful to provide another application of the derivative D_m. There is indeed one special class of mean field games, for which the corresponding MFG system coincides with the first-order condition (or equivalently, with the Pontryagin system) of a control problem. Such games are called potential games, and the control problem lying above a potential game is usually called a mean field control problem. The connection between both can be thus formulated in this way: The minimizers to the mean field control problem are equilibria of the corresponding potential game. This was noted in the earlier articles by Lasry and Lions [85–87], see also [88].

The potential structure turns out to be very useful in practice for the simple reason that it might be easier to work with minimizers than with Nash equilibria. We provide a longer discussion in Section 4 about possible applications to the selection of equilibria when there is no uniqueness.

In the simple framework of (1.2)–(1.4), the potential game typically requires that the cost coefficients F and G derive from a potential, namely

$$\partial_x F(x, m) = D_m \mathcal{F}(m, x), \quad \partial_x G(x, m) = D_m \mathcal{G}(m, x), \tag{2.2}$$

for two smooth functionals $\mathcal{F}$ and $\mathcal{G}$ on $\mathcal{P}_2(\mathbb{R}^d)$. With the trajectory $(X_t^\beta)_t$ as in (1.3), we can associate the cost

$$\mathcal{J}\big((\beta_t)_{0 \leq t \leq T}\big) = \int_0^T \Big(\mathcal{F}\big(\mathcal{L}(X_t)\big) + \frac{1}{2}\mathbb{E}\big[|\beta_t|^2\big]\Big)dt + \mathcal{G}\big(\mathcal{L}(X_T)\big).$$

The following statement may be found under more precise assumptions in [35, CHAPTER 6]:

Proposition 2.4. *Under suitable regularity properties on $\mathcal{F}$ and $\mathcal{G}$, and for given initial distribution $\tilde{m}_0 \in \mathcal{P}_2(\mathbb{R}^d)$ for X_0 in (1.3), the optimal trajectories of $\mathcal{J}$ with respect to $(X_0, (B_t)_t)$-progressively measurable controls $(\beta_t)_t$ are solutions of the mean field game* (1.2)–(1.3).

When $(\beta_t)_t$ is identified with a feedback function $\beta : [0, T] \times \mathbb{R}^d \to \mathbb{R}^d$, equation (1.3) becomes a stochastic differential equation whose marginal law solves the Kolmogorov equation

$$\partial_t m_t - \Delta m_t + \operatorname{div}(\beta_t m_t) = 0,$$

with $m_0 = \tilde{m}_0$. It is then possible to write $\mathcal{J}$ as

$$\mathcal{J}\big((\beta_t)_{0\le t\le T}\big) = \int_0^T \left(\mathcal{F}(m_t) + \frac{1}{2}\int_{\mathbb{R}^d} \big|\beta_t(x)\big|^2 m_t(dx)\right) dt + \mathcal{G}(m_T),$$

which is more in line with the formulation (1.4) of the mean field game. Although the resulting class of controls is obviously smaller when restricted to feedback controls, the infimum of $\mathcal{J}$ is the same, see Lacker [82].

2.2. The master equation

The master equation was first derived by Lions in [88] as the formal limit of the Nash system (1.1). It is a PDE with unknown $U : [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ (with U writing $(t, x, m) \mapsto U(t, x, m)$) and reads

$$\begin{cases} \text{(i)} \quad -\partial_t U_t - (1+\epsilon)\Delta_x U_t \\ \qquad + \frac{1}{2}|D_x U_t|^2 + \int_{\mathbb{R}^d} D_m U_t(t,x,m,y) \cdot D_x U_t(t,y,m) m(dy) \\ \qquad - (1+\epsilon)\int_{\mathbb{R}^d} \operatorname{div}_y(D_m U_t(t,x,m,y)) m(dy) \\ \qquad - 2\epsilon \int_{\mathbb{R}^d} \operatorname{div}_x(D_m U_t(t,x,m,y)) m(dy) \\ \qquad - \epsilon \int_{\mathbb{R}^d\times\mathbb{R}^d} \operatorname{Tr}_y(D^2_{mm} U_t(t,x,m,y,y')) m(dy) m(dy') = F(x,m) \\ \qquad\qquad \text{in } (0,T)\times\mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d), \\ \text{(ii)} \quad U_T(x,m) = G(x,m) \quad \text{in } \mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d). \end{cases} \tag{2.3}$$

This is a kind of hyperbolic equation stated on the infinite-dimensional space $\mathcal{P}_2(\mathbb{R}^d)$. Indeed, when F and G are monotone (recall (1.5)), the solution can be (at least formally) built by the method of characteristics. To ease the presentation, let us explain this when there is no common noise ($\epsilon = 0$). Let $(t_0, m_0) \in [0, T) \times \mathcal{P}_2(\mathbb{R}^d)$ and $(u_t, m_t)_t$ be the unique solution of the MFG system (1.4) stated on $(t_0, T) \times \mathbb{R}^d$ with initial condition $m_{t_0} = \tilde{m}_0$. Let us set $U_{t_0}(x, m_0) = u_{t_0}(x)$. Assuming that U is sufficiently smooth, one can easily check that U solves (2.3) by expanding it along the path $(m_t)_t$ (see [25]). The main issue is to prove that U is indeed smooth. When there is a common noise ($\epsilon > 0$), similar ideas can be implemented, using the stochastic MFG system (1.8) instead of the deterministic one.

Let us loosely summarize the main result concerning (2.3) (see [25, 36]).

Theorem 2.5. *Assume that F and G are smooth enough and are monotone in the sense of* (1.5)*. Then there exists a unique classical solution to* (2.3).

By "a classical solution," we mean a map U for which all the derivatives in (2.3) exist, are bounded and globally Lipschitz continuous. This strong notion of solution is needed below for the convergence results.

It is known that, if one removes the monotonicity condition, then the solution of (2.3) exists on a short time interval but may develop a discontinuity after a while. Most formal properties of the master equation have been introduced and discussed by Lions in [88], who also introduced the so-called Hilbertian approach (lifting the equation to the space of random variables). The actual proof of the existence of a solution of the master equation is a tedious verification that the map U, as defined above from the MFG system, actually gives rise to a classical solution. This required several steps in the literature before the proof was completed: The first paper in this direction is [20], where the classical solutions to the linear Kolmogorov equation associated with a standard Fokker–Planck equation are studied; Gangbo and Swiech [64] address the master equation in short time and without any diffusion term; Chassagneux et al. [45] obtain the existence and uniqueness for the master equation without common noise; Cardaliaguet et al. [25] establish the existence and uniqueness of solutions for the master equation with common noise under the monotonicity condition (see also [36]). Since then, there have been many works on the subject. We provide some references in Section 5.1.

2.3. Convergence of the Nash system

One key feature of the master equation is that it allows building approximate solutions of the Nash system (1.1) whose regularity is independent of the number of players.

Proposition 2.6. *Assume that U is a classical solution of* (2.3) *and let $(u^{N,i})_{i\in\{1,\dots,N\}}$ be its finite-dimensional projections:*

$$u_t^{N,i}(x_1,\dots,x_N) = U_t(x_i, m_{\boldsymbol{x}}^{N,i}), \quad \textit{where } m_{\boldsymbol{x}}^{N,i} := \frac{1}{N-1}\sum_{j=1,\dots,N:\, j\neq i} \delta_{x_j}.$$

Then u^N almost solves the Nash system (1.1)*:*

$$\begin{cases} -\partial_t u_t^{N,i} - \displaystyle\sum_{j=1}^N \Delta_{x_j} u_t^{N,i} - \epsilon \sum_{j,k=1}^N \mathrm{Tr}(D^2_{x_j x_k} u_t^{N,i}) + \frac{1}{2}|D_{x_i} u_t^{N,i}|^2 \\ \quad + \displaystyle\sum_{j\neq i} D_{x_i} u_t^{N,i} \cdot D_{x_j} u_t^{N,j} = F(x_i, m_{\boldsymbol{x}}^{N,i}) + r_t^{N,i}(\boldsymbol{x}) \\ \qquad\qquad \textit{in } (0,T)\times(\mathbb{R}^d)^N, \quad i \in \{1,\dots,N\}, \\ u_T^{N,i}(\boldsymbol{x}) = G(x_i, m_{\boldsymbol{x}}^{N,i}) \quad \textit{in } (\mathbb{R}^d)^N, \quad i\in\{1,\dots,N\}, \end{cases}$$

where

$$\left|r_t^{N,i}(\boldsymbol{x})\right| \le \frac{1}{N}\left(1 + \frac{1}{N}\sum_{i=1}^N |x_i - x_j|\right).$$

The proof relies on Proposition 2.2 and on its extension to higher-order derivatives. Note, however, that the result does not show directly that u^N is close to the solution v^N of (1.1) because each $u^{N,i}$ solves (1.1) only up to an error term of size $1/N$ while the system counts exactly N equations.

The main convergence result of [25] and [36] is the following:

Theorem 2.7. *Let v^N be the solution of the Nash system and assume that U is a classical solution of* (2.3) *with bounded derivatives. Then there exists a constant $C > 0$ such that, for all $i \in \{1, \ldots, N\}$ and all $(t, \boldsymbol{x}) \in [0, T] \times (\mathbb{R}^d)^N$,*

$$\left|U_t(x_i, m_{\boldsymbol{x}}^{N,i}) - v_t^{N,i}(\boldsymbol{x})\right| \le \frac{C}{N}\left(1 + |x_i|^2 + \frac{1}{N}\sum_{j=1}^N |x_j|^2\right)^{1/2}.$$

Theorem 2.7 provides an obvious comparison between the equilibrium values of the finite game and of the mean field game. Even though it is not obvious at first sight, the statement is in fact reminiscent of earlier results on the convergence of classical mean field particle systems to Fokker–Planck equations. An alternative strategy to the standard coupling argument for proving propagation of chaos (see [101]) consists indeed in studying the action of the semigroup generated by the McKean–Vlasov equation onto the marginal empirical measure of the particle system (see Kolokoltsov [78] and the works of Mouhot, Mischler, and Wennberg [94]). In comparison, the game setting involves an additional optimization step, which makes the analysis really difficult. In order to account for this optimization step, we work instead with forward–backward McKean–Vlasov equations, following the approach developed in [34,36]. We describe the main lines below.

Sketch of the proof of Theorem 2.7. The first step is to provide a probabilistic representation of the solution v^N of the Nash system. This goes through the representation of the equilibrium paths. To this end, we recall that $(\bar\alpha^i(t,\boldsymbol{x}) := -D_{x_i} v^{N,i}(t,\boldsymbol{x}))_{i=1,\ldots,N}$ is the Nash equilibrium of the N-player game in closed-loop form. For a given starting point $\boldsymbol{x} = (x_1, \ldots, x_N) \in (\mathbb{R}^d)^N$, the equilibrium trajectories $\boldsymbol{X}^{*N} = (X^{*N,i})_{i\in\{1,\ldots,N\}}$ associated to the Nash equilibrium are the solutions to the system

$$\begin{cases} dX_t^{*N,i} = -D_{x_i} v_t^{N,i}(\boldsymbol{X}_t^{*N})dt + \sqrt{2}dB_t^i + \sqrt{2\epsilon}dW_t, \\ X_0^{*N,i} = x_i. \end{cases} \tag{2.4}$$

Adopting a Lagrangian point of view, we may then follow the evolution of the cost and of the control along the system, which prompts us to let

$$Y_t^{*N,i} = v_t^{N,i}(\boldsymbol{X}_t^{*N}), \quad Z_t^{*N,i,j} = D_{x_j} v_t^{N,i}(\boldsymbol{X}_t^{*N}).$$

Classical Itô's formula, combined with the form of the Nash system, leads to the following expansion:

$$dY_t^{*N,i} = -\Big(\frac{1}{2}|Z_t^{*N,i,i}|^2 + F(X_t^{*N,i}, m^{N,i}_{\boldsymbol{X}_t^{*N}})\Big)dt + \sqrt{2}\sum_{j=1}^N Z_t^{*N,i,j}\cdot(dB_t^j + \sqrt{\epsilon}dW_t). \tag{2.5}$$

In order to test, as suggested before, the action of the solution U to the master equation (which is somehow the analogue of the semigroup generated by a McKean–Vlasov equation but in the nonlinear setting induced by the game structure) on the Nash equilibrium of the N-player game, we need to perform a similar computation, but for the processes

$$\mathcal{Y}_t^{*N,i} = u^{N,i}(t, X_t^{*N}), \quad \mathcal{Z}_t^{*N,i,j} = D_{x_j}u_t^{N,i}(\boldsymbol{X}_t^{*N}), \quad t\in[0,T],$$

with $(u^{N,i})_i$ being as in Proposition 2.6. In fact, Proposition 2.6 now permits expanding $(\mathcal{Y}_t^{*N,i})_t$. We get

$$\begin{aligned} d\mathcal{Y}_t^{*N,i} &= -\Big(\frac{1}{2}|\mathcal{Z}_t^{*N,i,i}|^2 + F(X_t^{*N,i}, m^{N,i}_{\boldsymbol{X}_t^{*N}}) + r_t^{N,i}(\boldsymbol{X}_t^{*N,i})\Big)dt \\ &\quad + \sum_{j=1}^N \mathcal{Z}_t^{*N,i,j}\cdot(\mathcal{Z}_t^{*N,j,j} - Z_t^{*N,j,j})dt + \sqrt{2}\sum_{j=1}^N \mathcal{Z}_t^{*N,i,j}\cdot(dB_t^j + \sqrt{\epsilon}dW_t). \end{aligned}$$

Importantly, the two processes $(Y_t^{*N,i})_t$ and $(\mathcal{Y}_t^{*N,i})_t$ satisfy the same boundary conditions at time T, namely $Y_T^{*N,i} = \mathcal{Y}_T^{*N,i} = g(X_T^{*N,i}, m^{N,i}_{\boldsymbol{X}_T^{*N}})$, which prompts us to address the difference process $(Y_t^{*N,i} - \mathcal{Y}_t^{*N,i})_{0\le t\le T}$. We get

$$\begin{aligned} d(\mathcal{Y}_t^{*N,i} - Y_t^{*N,i}) = &-\Big(\frac{1}{2}|\mathcal{Z}_t^{*N,i,i}|^2 - \frac{1}{2}|Z_t^{*N,i,i}|^2 + r_t^{N,i}(\boldsymbol{X}_t^{*N,i})\Big)dt \\ &+ \sum_{j=1}^N \mathcal{Z}_t^{*N,i,j}\cdot(\mathcal{Z}_t^{*N,j,j} - Z_t^{*N,j,j})dt \\ &+ \sqrt{2}\sum_{j=1}^N(\mathcal{Z}_t^{*N,i,j} - Z_t^{*N,i,j})\cdot(dB_t^j + \sqrt{\epsilon}dW_t). \end{aligned} \tag{2.6}$$

The last term yields a stochastic integral. If there were no dt-term in the right-hand side, then the simple fact that the terminal condition is equal to 0 would say that the stochastic integral is also null. In turn, this would say that $\mathcal{Z}_t^{*N,i,j} - Z_t^{*N,i,j} = 0$ for any t. In other words, the noise provides a strong form of stability in the above equation. This is consistent with the fact that, in the Nash system, the Laplace operator dissipates the energy when time runs backwards. The sum on the second line is also challenging, at least at first sight. However, Proposition 2.2 says that, except when $j = i$, all the terms are of order $1/N$, which guarantees that the whole sum is of order 1. On the first line of the right-hand side, the remainder $r_t^{N,i}$ is also known to be of order $1/N$ on compact sets. In the end, we are thus left with a backward stochastic differential inequation of the form

$$
\begin{aligned}
d(\mathcal{Y}_t^{*N,i} - Y_t^{*N,i}) = & - \left[\frac{1}{2} \left| \mathcal{Z}_t^{*N,i,i} \right|^2 - \frac{1}{2} \left| Z_t^{*N,i,i} \right|^2 \right. \\
& \left. + O\left(\frac{1}{N} + \frac{1}{N^2} \sum_{i,j=1}^{N} \left| X_t^{*i,N} - X_t^{*j,N} \right| \right) \right] dt \\
& + O\left(\frac{1}{N} \sum_{j=1}^{N} \left| \mathcal{Z}_t^{*N,j,j} - Z_t^{*N,j,j} \right| + \left| \mathcal{Z}_t^{*N,i,i} - Z_t^{*N,i,i} \right| \right) dt \\
& + \sqrt{2} \sum_{j=1}^{N} (\mathcal{Z}_t^{*N,i,j} - Z_t^{*N,i,j}) \cdot (dB_t^j + \sqrt{\epsilon} dW_t).
\end{aligned}
$$

Above, the symbol $O(\cdot)$ is used for the Landau notation, the underlying constant being (in our setting) deterministic and independent of N and t. Obviously, the goal is to provide a stability analysis of this equation. Needless to say, the main difficulty in this regard is the difference of the two quadratic terms on the first line of the right-hand side. Invoking Proposition 2.2 once again and using the fact that the solution to the master equation is assumed to have bounded derivatives, it is pretty easy to get L^∞-bounds on the process $(\mathcal{Z}_t^{*N,i,i,})_t$, independently of i and N. However, there are no similar inequalities for the process $(Z_t^{*N,i,i,})_t$. This is in fact the main challenge in this proof: Known estimates on the regularity of $v^{N,i}$, and in particular on its gradient, depend on N. Accordingly, most of the analysis relies on the sole properties of the solution U to the master equation. In words, there is no easy way here to linearize the difference of the two quadratic terms in the backward equation. The idea is then to adapt some of the tricks that have been developed in the literature on backward stochastic differential equations with a quadratic dependence on the martingale representation term (here denoted by $(\mathcal{Z}_t^{*N,i,j} - Z_t^{*N,i,j})_t$). In the analysis of the well-posedness of a backward stochastic differential equation, quadratic growth (with respect to the same martingale representation term) is indeed known to be a threshold. This is consistent with the results on nonlinear parabolic PDEs: quadratic growth in the gradient of the solution is also known to be a threshold. Noticeably, the unknown in the backward equation should be in fact regarded as being multidimensional since it comprises all the coordinates $(\mathcal{Y}_t^{*N,i} - Y_t^{*N,i})_{i=1,\dots,N}$. In general, this is known to render the analysis in the quadratic case even more challenging. Anyway, the symmetric structure of the equation here is very helpful and somehow permits thinking as if the equation were set in dimension 1. In the end, a suitable form of exponential transform (very much inspired from the Cole–Hopf transform in the analysis of Hamilton–Jacobi–Bellman equations, see [36, CHAPTER 6] for the details) allows transforming the quadratic equation into a linear one, and then concluding by using standard stability arguments from the theory of backward stochastic differential equations. Essentially, the size of the difference terms $((\mathcal{Y}_t^{*N,i} - Y_t^{*N,i})_t)_{i=1,\dots,N}$ is dictated by the remainder in the equation and is thus of order $1/N$. It then remains to observe that that, at time $t = 0$,

$$
\mathcal{Y}_0^{*N,i} - Y_0^{*N,i} = u_0^{N,i}(\boldsymbol{x}) - v_0^{N,i}(\boldsymbol{x}) = U_0(x_i, m_{\boldsymbol{x}}^{N,i}) - v_0^{N,i}(\boldsymbol{x}).
$$

Our sketch of proof hence shows that the left-hand side is of order $1/N$. In fact, a careful inspection would permit tracking the dependence on the initial conditions and recovering the same rate as in the statement. ■

2.4. Propagation of chaos for the N-player game

In fact, the proof of Theorem 2.7 kills two birds with one stone. Indeed, it also permits addressing the large-N behavior of the equilibrium trajectories of the N-player game. Recall indeed from (2.4) that these equilibrium trajectories solve the system of stochastic differential equations

$$dX_t^{*N,i} = -D_{x_i} v_t^{N,i}(X_t^{*N})dt + \sqrt{2}dB_t^i + \sqrt{2\epsilon}dW_t, \tag{2.7}$$

for a given choice of initial conditions. In order to state propagation of chaos in a proper manner, we assume, as in our preliminary description of a mean field game in Section 2, that these initial conditions are given as independent samples $X_0^1, \dots, X_0^N$ from a common distribution $\tilde{m}_0 \in \mathcal{P}_2(\mathbb{R}^d)$.

Noticeably, the drift in (2.7) may be rewritten in terms of the notations introduced in the proof of Theorem 2.7. Indeed, this drift is nothing but $(-Z_t^{*N,i,i})_{0\le t\le T}$, which is a key quantity in the proof of Theorem 2.7. It is then worth emphasizing that stability arguments for backward stochastic differential equations like those we used in this proof provide more than what is eventually contained in the result. They also provide a similar bound on the quadratic variation (or, equivalently, on the energy) of the martingale representation term in (2.6). Using the fact that $\tilde{m}_0$ is square-integrable, we end up with the fact that

$$\mathbb{E}\int_0^T \left|\mathcal{Z}_t^{*N,i,i} - Z_t^{*N,i,i}\right|^2 dt \le \frac{C}{N^2},$$

for a constant C that is independent of N. Implicitly, the constant C depends on $\tilde{m}_0$ through its second-order moment. Moreover, it is worth recalling that, on the left-hand side, $\mathcal{Z}_t^{*N,i,i} = -D_x U_t(X_t^{*N,i}, m^{N,i}_{\boldsymbol{X}^{*N}})$. In turn, this says that, up to an error of order $1/N$, we can replace the drift in (2.7) by $-D_x U_t(X_t^{*N,i}, m^{N,i}_{\boldsymbol{X}^{*N}})$. Equivalently, by using the regularity properties of $D_x U$, we have

$$\sup_{i=1,\dots,N} \mathbb{E}\Big[\sup_{0\le t\le T} \left|\mathcal{X}_t^{*N,i} - X_t^{*N,i}\right|^2\Big] \le \frac{C}{N^2}, \tag{2.8}$$

where

$$\begin{cases} d\mathcal{X}_t^{*N,i} = -D_x U_t(\mathcal{X}_t^{*N,i}, m^{N,i}_{\boldsymbol{\mathcal{X}}_t^{*N}})dt + \sqrt{2}dB_t^i + \sqrt{2\epsilon}dW_t, \\ \mathcal{X}_0^i = X_0^i. \end{cases} \tag{2.9}$$

Very differently from (2.7), whose structure is made intricate by the presence of v^N, (2.9) is a standard weakly interacting particle system. As such, it is known to converge to the solution of the conditional McKean–Vlasov equation

$$\begin{cases} d\mathcal{X}_t^i = -D_x U_t\big(\mathcal{X}_t^i, \mathcal{L}(\mathcal{X}_t^i|W)\big)dt + \sqrt{2}dB_t^i + \sqrt{2\epsilon}dW_t, \\ \mathcal{X}_0^i = X_0^i. \end{cases}$$

The analysis of the above equation is standard. Under our standing assumptions on U, it follows from a classical contraction argument. In particular, uniqueness for the above equation implies that the conditional law $\mathcal{L}(\mathcal{X}_t^i|W)$ that appears in the dynamics is in fact independent of i. Sznitman's coupling argument **[101]** then allows estimating the distance between the solution of (2.9) and the solution of the above conditional McKean–Vlasov equation. We get

Theorem 2.8. *For any $\eta > 0$, there exists a constant $C_\eta > 0$ such that, for all $N \geq 1$ and for all $i \in \{1, \ldots, N\}$,*

$$\mathbb{E}\Big[\sup_{t\in[0,T]} \big|X^{*N,i} - X_t^i\big|\Big] \leq C_\eta N^{-1/\max\{d,2+\eta\}}.$$

When $d \geq 3$, we can choose $\eta = 0$.

Noticeably, the rate in Theorem 2.8 is much weaker than the rate in (2.8). In fact, the bound in Theorem 2.8 is the same as the bound for the mean 1-Wasserstein distance between a probability distribution in $\mathcal{P}_2(\mathbb{R}^d)$ and the empirical law of an independent sample of it. We refer to Fournier and Guillin's idea **[62]** for a complete review of the subject.

Some bibliographical comments on the propagation of chaos in Nash equilibria are now in order. The first results concerning this question are due to Fischer **[61]** and Lacker **[81]** for open-loop controls (in which players observe only the initial states and the Brownian motions) in problems without common noise: Lacker **[81]**, in particular, identified completely the possible limits, which are always MFG equilibria (in a weak form, with a notion of weak solution similar to **[38]**). The question of convergence of closed-loop equilibria is more subtle. As shown in a counterexample in **[36, I.7.2.5]** (inspired from **[56]**), this convergence does not hold in full generality. At present, the minimal conditions to obtain it are still not clear. Theorem 2.8, proved first in the periodic setting in **[25]** and then extended to the Euclidean framework in **[35]**, shows that the convergence holds if there exists a classical solution (with bounded derivatives) to the master equation (which implies that equilibria are unique) and if the idiosyncratic noise is nondegenerate (which implies that it is not null). In the same framework (and with $\mathbb{R}^d$ as state space), Delarue et al. **[51]** and **[50]** established a central limit theorem and a large deviation principle, using the same idea as in the proof of Theorem 2.8: the main point is to show that the fluctuations and the deviations in the convergence of the N-player game equilibria are mainly due to the fluctuations and the deviations in the convergence of the standard particle system (2.9). In a beautiful work, Lacker **[83]** extended the result by establishing convergence without assuming the existence of the master equation or any monotonicity property (but keeping the assumption that the idiosyncratic noise is nondegenerate): the limit points are weak MFG equilibria. The main difference with Theorem 2.8 is that **[83]** is based on a compactness argument (obtained by using the theory of relaxed controls, in which controls are regarded as being measure-valued) and provides no convergence rate. The result relies on the fact that, in some average sense, the deviation of a player barely affects the distribution of the players when N is large. Heuristically, this is due to the presence of the noise, which prevents the players to guess if another has deviated or not. However, in Lacker's approach, there might be a lot of (weak) MFG equilibria, apart from the monotone case where they are unique. This raises subtle questions of selection since only some of these equilibria may be selected when passing to limit: this is what happens in the examples discussed in Bayraktar and Zhang [14], Cecchin, Dai Pra, Fischer, and Pelino, [44], and Delarue and Foguen [52]. We provide more details in Section 4. Let us underline another limitation: The result presented above, as well as Lacker's approach, rely in a crucial way on the presence of a nondegenerate idiosyncratic noise and, to date, nothing is known outside this framework.

Finally, it is important to note that, historically, another approach was first implemented to relate the N-player and mean field games. In short, any solution to the mean field game gives rise to an approximate Nash equilibrium to the N-player game, with an accuracy that gets better and better as N increases. This idea dates back to the earliest papers in the field [73,75]. We refer to [36, CHAPTER 6] for a complete review.

3. THE LONG-TIME BEHAVIOR

In this section we discuss the behavior of MFG equilibria (without common noise) as the time horizon T tends to infinity. This is an interesting question both in terms of theory and applications: for instance, in economics, it is related to the existence of stationary equilibria or business cycles. On the other hand, the answer is not obvious because the MFG system has two boundary conditions, one at the initial time and one at the terminal time. One can therefore expect that convergence holds only far from the initial and terminal times. In order to perform this analysis, it is necessary to require that the solution of the stochastic control problem remains confined in an appropriate sense: the simplest setting in which this is possible is the spatially periodic one. We make this assumption here: we set $\mathbb{T}^d := \mathbb{R}^d/\mathbb{Z}^d$ and denote by $\mathcal{P}(\mathbb{T}^d)$ the set of Borel probability measures on $\mathbb{T}^d$ endowed with the corresponding 2-Wasserstein distance. We consider the solution $(u^T, m^T) = (u_t^T, m_t^T)_{0\le t\le T}$ of the MFG system (1.4), now stated on $(0,T)\times\mathbb{T}^d$, in which $F, G : \mathbb{T}^d \times \mathcal{P}(\mathbb{T}^d) \to \mathbb{T}$ are "smooth."

3.1. The ergodic MFG system

As explained by Lions in [88], the limit of the MFG system (1.4), as the time horizon T tends to infinity, is expected to be given by the ergodic MFG system

$$\begin{cases} \bar\lambda - \Delta \bar u + \frac{1}{2}|D\bar u|^2 = F(x,\bar m) & \text{in } \mathbb{T}^d, \\ -\Delta \bar m - \operatorname{div}(\bar m\, D\bar u) = 0 & \text{in } \mathbb{T}^d, \\ \int_{\mathbb{T}^d} \bar m = 1, \quad \int_{\mathbb{T}^d} \bar u = 0. \end{cases} \tag{3.1}$$

Here the unknowns are $(\bar\lambda, \bar u, \bar m)$, where $\bar\lambda \in \mathbb{R}$ is the so-called ergodic constant. The interpretation of the system is the following: each player wants to minimize her ergodic cost

$$J(x,\alpha) := \inf_{\alpha} \limsup_{T\to+\infty} \mathbb{E}\left[\frac{1}{T}\int_0^T \left\{\frac{1}{2}|\alpha_t|^2 + F(X_t,\bar m)\right\} dt\right]$$

where $(X_t)_{t\geq 0}$ is the solution to

$$\begin{cases} dX_t = \alpha_t dt + \sqrt{2} dB_t, \\ X_0 = x. \end{cases}$$

The measure $\bar{m}$ in (3.1) is then understood as the invariant ergodic measure associated to the optimal trajectory (the existence of which is much easier to prove in the periodic setting). The solution to (3.1) is known to exist under fairly general assumptions on F and to be unique when the coupling function F is monotone (i.e., satisfies (1.5)); see [85, 87].

3.2. The convergence in the monotone setting

In this part we assume that F is smooth and monotone. Under this monotonicity assumption, one can show that the "long-time stability" takes the form of a *turnpike pattern*; namely, the solution (u^T, m^T) of (1.4) becomes nearly stationary for most of the time. The strongest way to formulate this type of behavior is the following exponential estimate:

Theorem 3.1. *There exist $K, \omega > 0$ such that for (u^T, m^T) and $(\bar{u}, \bar{m})$ solving respectively (1.4) and (3.1),*

$$\left\| m^T(t) - \bar{m} \right\|_\infty + \left\| Du^T(t) - D\bar{u} \right\|_\infty \leq K\big(e^{-\omega t} + e^{-\omega(T-t)}\big), \quad \forall t \in (0, T). \tag{3.2}$$

The reader may notice that the initial condition $\tilde{m}_0$ for m^T and the terminal condition G for u^T are lost at the limit (as $(\bar{\lambda}, \bar{u}, \bar{m})$ does not depend on $\tilde{m}_0$ or G). This result was first stated in Cardaliaguet, Lasry, Lions, and Porretta [29] when the coupling F is monotone and local and in [30] when this coupling is monotone and regularizing. The proof is based in a crucial way on uniform (in t and T) semiconcavity estimates for u^T and on the energy identity established by Lasry and Lions [87]:

$$\begin{aligned} &\int_0^T \int_{\mathbb{T}^d} \frac{1}{2}(m_t^T + \bar{m})(x) \left| Du_t^T(x) - D\bar{u}(x) \right|^2 dt dx \\ &\quad = -\int_0^T \int_{\mathbb{T}^d} \big(F(x, m_t^T) - F(x, \bar{m})\big)(m_t^T - \bar{m})(x) dt dx \\ &\qquad - \int_{\mathbb{T}^d} \big(G(x, m_T^T) - \bar{u}(x)\big)(m_T^T - \bar{m})(x) dx + \int_{\mathbb{T}^d} (u_0^T - \bar{u})(x)(m_0^T - \bar{m})(x) dx. \end{aligned}$$

This energy identity shows the role of the monotonicity property (1.5) in the analysis.

A consequence of the exponential estimate (3.2) is the existence of a constant C such that

$$\left| u^T(t, x) - \bar{u}(x) - \bar{\lambda}(T - t) \right| \leq C.$$

Following ideas of weak KAM theory (see, for instance, the ICM proceeding by Fathi [60] in the calculus of variation framework), one could expect the existence of a limit for $u^T(t, x) - \bar{\lambda}(T - t)$ as T tends to ∞; moreover, this limit should be given (up to an additive constant) by $\bar{u}$. However, this heuristic is not completely correct and the description of the asymptotic behavior of u^T (eventually established in the paper by Cardaliaguet and Porretta [33]) happens to be more subtle.

To overcome the difficulty that the MFG system is forward–backward, a possible path (towards a long-time expansion of u^T) is to use the master equation (2.3), which is just backward in time. One of the main results of **[33]** states that the solution of the master equation converges to the solution of the following ergodic master equation:

$$\lambda - \Delta_x \chi(x,m) + \frac{1}{2}\left|D_x \chi(x,m)\right|^2 - \int_{\mathbb{T}^d} \mathrm{div}_y (D_m \chi(x,m,y)) dm(y)$$
$$+ \int_{\mathbb{T}^d} D_m \chi(x,m,y) \cdot D_x \chi(y,m) dm(y) = F(x,m) \quad \text{in } \mathbb{T}^d \times \mathcal{P}(\mathbb{T}^d). \tag{3.3}$$

Concerning the existence of (3.3), the following result holds:

Theorem 3.2. *There is a unique constant $\lambda \in \mathbb{R}$ for which the master cell problem* (3.3) *has a (weak) solution. The constant λ coincides with the unique constant $\bar{\lambda}$ for which the ergodic MFG problem* (3.1) *has a solution. Besides, if χ is a solution to* (3.3)*, then $\chi(\cdot, m)$ is of class C^2 (in space) for any $m \in \mathcal{P}(\mathbb{T}^d)$ and*

$$D_x \chi(x, \bar{m}) = D\bar{u}(x) \quad \forall x \in \mathbb{T}^d,$$

where $(\bar{u}, \bar{m})$ is the solution to (3.1).

As in many constructions of a solution to an ergodic problem, the first step consists in building solutions to approximating compact problems and then in proving uniform estimates on these solutions. Here, the compact problems are discounted master equations which can be solved by a method of (infinite-dimensional) characteristics (as for (2.3)). The main issue is to prove estimates on these solutions, independently of the discount rate. In contrast with standard constructions in this area (see Lions–Papanicolau–Varadhan **[90]** or **[60]**, which analyze the ergodic behavior of (pure) Hamilton–Jacobi equations with a coercive Hamiltonian), the proof of these estimates cannot rely on the coercivity properties of the equation, but must use in a very strong way the bound (3.2), which describes the long-time behavior of the characteristics.

We are now ready to discuss the convergence, as $t \to -\infty$, of the solution U of the master equation (2.3) (now defined in the time interval $(-\infty, 0]$ with terminal condition $U(0,x,m) = G(x,m)$).

Theorem 3.3. *Let χ be a weak solution to the master cell problem* (3.3)*. Then, there exists a constant $c \in \mathbb{R}$ such that*

$$\lim_{t \to -\infty} U(t,x,m) + \bar{\lambda} t = \chi(x,m) + c,$$

uniformly with respect to $(x,m) \in \mathbb{T}^d \times \mathcal{P}(\mathbb{T}^d)$.

Moreover, we also have that $D_x U(t,x,m) \to D_x \chi(x,m)$ as $t \to -\infty$, uniformly with respect to (x,m).

This result looks like an extension of the famous Fathi's result on the convergence of the Lax–Oleinik semigroup in weak-KAM theory **[60]**. This parallel is not completely correct since the master equation is not a Hamilton–Jacobi equation in an infinite-dimensional setting: the comparison principle does not hold, for instance. One has to rely instead on the energy identity described above.

From Theorem 3.3 one can derive the full convergence of the solution (u^T, m^T) of the MFG system:

Corollary 3.4. *Let c be the constant given in Theorem 3.3. For $T > 0$ and $\tilde{m}_0 \in \mathcal{P}(\mathbb{T}^d)$, let (u^T, m^T) be the solution to (1.4). Then, for any $t \geq 0$,*

$$\lim_{T\to+\infty} \big(u^T(t,x) - \bar{\lambda}(T-t)\big) = \chi\big(x, m(t)\big) + c,$$

where the convergence is uniform in x and m solves

$$\partial_t m - \Delta m - \operatorname{div}\big(m D_x \chi(x,m)\big) = 0, \quad m(0) = \tilde{m}_0.$$

In the recent paper [47], Cirant and Porretta managed to show the above corollary without relying on the master equation.

Among many open problems in this area, let us point out the following ones: we have explained in Section 2 that the master equation can be obtained as the limit of the Nash system (1.1). Now that we understand the behavior of the master equation on long time intervals, it would be interesting to see if this convergence holds uniformly in time. Similar results have been obtained, for instance, by Mischler and Mouhot [93] in the framework of kinetic theory. Another very intriguing issue is the long-time behavior of the MFG system in the presence of a common noise: the existence of stationary measures is a completely open problem.

3.3. The long-time behavior without monotonicity

The long-time behavior of the MFG equilibria when the coupling is not monotone is poorly understood and only partial results are known.

The potential case. When the MFG is potential (see (2.2)), then one can extend weak-KAM theory to the infinite-dimensional setup and describe the possible ω-limit sets of the solution of the time-dependent MFG system minimizing a natural energy in terms of a "Mather set." The main point is that this set may not contain an ergodic MFG equilibrium (i.e., an $\bar{m} \in \mathcal{P}(\mathbb{T}^d)$ for which there exists $(\bar{\lambda}, \bar{u})$ such that $(\bar{\lambda}, \bar{u}, \bar{m})$ solves (3.1)): this shows that the ω-limit set of the solutions of the time-dependent MFG system (1.4) that additionally minimize the natural energy may not contain an MFG ergodic equilibrium. In other words, the ergodic MFG system (3.1) may not describe the long-time behavior of these trajectories.

Periodic solutions. The existence of a periodic solution to the MFG system is a fascinating topic on which little is known. The main result in that direction is the analysis by Cirant [46] of a class of examples. It relies on local and global bifurcation methods based on the analysis of eigenfunction expansions of solutions to a suitable linearized problem. Note, however, that the stability of these solutions is not known.

Traveling waves. Intimately related to the notion of equilibria and to periodic solutions, the question of traveling waves has been discussed in the framework of an MFG problem of knowledge growth, first introduced in economics by Lucas and Moll [91]. In this setting the construction of a traveling wave solution is crucial (it is called a balanced-growth path solution in economics) and has been documented by Papanicolaou, Ryzhik and Velcheva [95] and by Porretta and Rossi [98]. The convergence of the solution of the time-dependent problem to this solution remains an open problem.

4. SMOOTHING EFFECT OF THE COMMON NOISE

A natural question is to address the impact of the common noise on the well-posedness of a mean field game. It is indeed useful to observe that, most often, standard mean field games (without common noise) have multiple solutions. In this respect, condition (1.5) is rather restrictive. Just as an additive Brownian motion can restore uniqueness of differential equations driven by nonsmooth vector fields, we can then wonder whether a form of common noise could force equilibria to be unique in a fairly large class of mean field games.

4.1. The linear–quadratic case as a warm-up

It is pretty clear that the form of common noise that is inserted into (1.8) is certainly not sufficient to reach such an aim in full generality. Indeed, the noise is just finite-dimensional whereas the model is infinite-dimensional because of the mean field structure. For sure, we could think of some hypoelliptic structure that could allow the finite-dimensional noise to be transmitted to all the components of the space of probability measures, but this looks a very challenging question. A much easier (but much less ambitious) alternative is to restrict oneself to mean field games whose equilibria are *a priori* known to live in a finite-dimensional subset or, using a standard concept from statistics, to belong to a parametric model of statistical distributions. The typical example in this direction is the class of linear–quadratic mean field games, which has been studied with a lot of attention (see Bardi [10], Bensoussan, Sung, Yam and Yung [16], Carmona, Delarue, and Lachapelle [37], and the works [73, 75] by Caines, Huang, and Malhamé for a tiny example). In short, it corresponds to the case when F and G in (1.8) have the form

$$F(x,m) = \frac{1}{2}\big|Qx + f(\overline{m})\big|^2, \quad G(x,m) = \frac{1}{2}\big|Rx + g(\overline{m})\big|^2, \tag{4.1}$$

where Q and R are matrices of size $d \times e$ (with e being another integer), f, g are Borel functions from $\mathbb{R}^d$ to $\mathbb{R}^e$ and $\overline{m}$ is the mean of m, i.e., $\overline{m} = \int_{\mathbb{R}^d} x dm(x)$ (which implicitly requires m to have a finite first moment). Referring back to Section 1.2, we see that the control problem (1.2)–(1.3) ((1.6)–(1.7) in the presence of common noise) becomes a stochastic control problem with linear–quadratic coefficients depending on the (possibly random) path $(m_t)_t \in C^0([0,T], \mathcal{P}_2(\mathbb{R}^d))$ injected into the coefficients. The key point is that this stochastic control problem has a unique solution (depending on $(m_t)_t$), with the optimal feedback being affine (regardless of the value of the intensity of the noises). In turn, this implies that the equilibrium trajectories must be Gaussian processes (conditional on the initial condition whenever the latter is random). Therefore, for the above choice of F and G, the equilibria are necessarily Gaussian (once again, conditional on the initial condition). Even more, since the volatility coefficient is prescribed in the state dynamics, the variances of the marginal conditional laws of the equilibria given the initial condition are also fixed. In the end, only the means count for determining the equilibria: As expected, the model is parametric. It is then an interesting question to address the impact of the common noise in this specific framework and to see whether the existing well-posedness results can be improved under the action of $(W_t)_t$. A very convenient approach is to use the Pontryagin principle, which provides, under the standing x-convex structure of F and G, a characterization of the equilibria in the form of a forward–backward system of the McKean–Vlasov type. Standard computations (see the aforementioned references together with [35, CHAPTER 3]) then show that, for a given $(W_t)_t$-adapted path $(m_t)_t$ with values in $C^0([0,T], \mathcal{P}_2(\mathbb{R}^d))$, the optimal control in the stochastic control problem described in (1.6)–(1.7) has the following feedback form:

$$\alpha_t^* = -(\eta_t X_t^* + h_t), \tag{4.2}$$

where $(\eta_t)_t$ is the solution of an autonomous deterministic Riccati equation (the form of which is completely independent of the input $(m_t)_t$) and $(h_t)_t$ solves the finite-dimensional backward stochastic differential equation

$$h_t = R^\dagger g(\overline{m}_T) + \int_t^T \big[Q^\dagger f(\overline{m}_s) - \eta_s h_s\big] ds - \int_t^T k_s dW_s, \quad t \in [0,T]. \tag{4.3}$$

Obviously, this equation should be regarded as a finite-dimensional version of the backward equation in (1.8) when the value function therein is sought in a quadratic form.

Forcing uniqueness. Inserting the relationship (4.2) for the optimal feedback into the dynamics (1.7), taking the conditional mean of $(X_t^*)_t$ (with the exponent $*$ being used to denote the optimal trajectory) given the common noise $(W_t)_t$, and then identifying $\mathbb{E}[X_t^* | (W_s)_s]$ with $\overline{m}_t$ (in full consistency with the probabilistic fixed-point formulation of a mean field game), we end up with the following forward–backward system (which is now the finite-dimensional analogue of the whole system (1.8)):

$$\begin{cases} d\overline{m}_t = -(\eta_t \overline{m}_t + h_t)dt + \sqrt{2\epsilon}\, dW_t, & m_0 = \mathbb{E}(X_0), \\ dh_t = -\big(Q^\dagger f(\overline{m}_t) - \eta_t h_t\big)dt + k_t dW_t, & h_T = R^\dagger g(\overline{m}_T). \end{cases} \tag{4.4}$$

Similar to $(v_t)_t$ in (1.8), the $(W_t)_t$-adapted process $(k_t)_t$ is here designed to render the solution $(h_t)_t$ $(W_t)_t$-adapted. Remarkably, system (4.4) just involves the conditional expectation $(\overline{m}_t)_t$. This is in line with the fact that equilibria are known to belong to a parametric model. It then remains to interpret the forward–backward system (4.4) as the system of characteristics of a parabolic PDE. We obtain

$$h_t = \theta_t(\overline{m}_t),$$

where θ solves

$$\partial_t \theta_t(x) + \epsilon \Delta_x^2 \theta_t(x) - \big(\eta_t x + \theta_t(x)\big) \cdot \nabla_x \theta_t(x) + Q^\dagger f(x) - \eta_t \theta_t(x) = 0, \tag{4.5}$$

for $(t,x) \in (0,T) \times \mathbb{R}^d$, with the terminal condition $\theta_T(x) = g(x)$. This PDE is a finite-dimensional version of the master equation (2.3). Obviously, it is much easier to solve.

In particular, when $\epsilon > 0$, the sole presence of the Laplacian forces the existence of a classical solution when f and g are bounded and regular coefficients. In turn, this forces the well-posedness of the system of characteristics (4.4) (see Foguen [102] or [36, CHAPTER 3]):

Proposition 4.1. *Let the cost coefficients F and G be of the same form as in* (4.1)*, with f and g therein being bounded and sufficiently regular coefficients. Then, for any $\epsilon > 0$, the mean field game has a unique solution.*

It must be stressed that the statement becomes false when $\epsilon = 0$ (see the same references for explicit examples). One must then assume more about the coefficients f and g to force uniqueness. For instance, it is easy to reformulate the monotonicity condition (1.5) in terms of f and g: The point is then to require $Q^\dagger f$ and $R^\dagger g$ to satisfy $(Q^\dagger f(x') - Q^\dagger f(x)) \cdot (x' - x) \geq 0$ for any $x, x' \in \mathbb{R}^d$, and similarly for $R^\dagger g$. Regarding the explicit conditions of regularity that f and g must satisfy in Proposition 4.1, a typical instance is to assume that f is bounded and Hölder continuous on the whole space and g, together with its first and second-order derivatives, are bounded and Hölder continuous on the whole space.

4.2. Finite-state mean field games

Another obvious manner to get a parametric model is to force the state space to be finite, in which case the space of probability measures itself becomes finite-dimensional. This requires, however, a modicum of care since the state dynamics can no longer be formulated as in (1.3)–(1.7). In particular, the common noise cannot be chosen in a mere additive fashion.

Games without common noise. When the state space is finite (and is thus chosen as a finite set E), the dynamics of the reference player are usually postulated in the form of a Markov controlled process taking values in E. Typically, the transition rates are explicitly prescribed as functions of the control (see Gomes, Mohr, and Souza [66,67] and Guéant [70]). A simple, but convenient, choice is then to identify the control with the entire transition matrix. In that case, using the same notation $(X_t)_t$ as in (1.3) to denote the trajectory of the reference player, the transition probabilities read (with $\mathbb{P}$ being implicitly identified with $\mathbb{P}^1$ since there is no common noise at this stage of the discussion)

$$\begin{aligned} &\mathbb{P}(X_{t+dt} = j | X_t = i) = \beta_t^{i,j} dt + o(dt), \quad i \neq j, \\ &\mathbb{P}(X_{t+dt} = i | X_t = i) = 1 + \beta_t^{i,i} dt + o(dt), \end{aligned} \tag{4.6}$$

with $((\beta_t^{i,j})_{i,j \in E})_t$ standing for a deterministic path with values in the set of E-indexed matrices satisfying the following two standard prescriptions:

$$\begin{cases} \beta_t^{i,j} \geq 0, \quad i \neq j, \\ \beta_t^{i,i} = -\sum_{j \neq i} \beta_t^{i,j}. \end{cases} \tag{4.7}$$

This formulation is reminiscent of (1.3) in the sense that the transitions do not depend on the choice of the environment $(m_t)_t$ that underpins the cost functional (1.2). In particular, the

Fokker–Planck equation for the marginal law of the $(\beta_t)_t$-controlled process $(X_t)_t$ can be written as

$$\frac{d}{dt} p_t^i = \sum_{j \in E} p_t^j \beta_t^{j,i}, \quad t \in [0,T], \quad i \in E, \tag{4.8}$$

with p_t^i being understood as $\mathbb{P}(X_t = i)$. As for the cost functional, we may choose it as in (1.2) provided that the functions F and G are now defined on $E \times \mathcal{P}(E)$, with $\mathcal{P}(E)$ denoting the space of probability measures (which can be obviously identified with the simplex of dimension $|E| - 1$). To give a clear account that the state space is finite, we will write (in this subsection) $F^x(m)$ and $G^x(m)$ instead of $F(x,m)$ and $G(x,m)$. Of course, there is another slight difference with (1.2), which lies in the interpretation of $(\beta_t)_t$. In (1.2), β_t is implicitly chosen as a control in feedback form: Loosely speaking, we write $\tilde{\beta}_t(X_t)$ for a d-dimensional vector field $\tilde{\beta}_t$; In other words, the quadratic cost in (1.2) is calculated from the pointwise value of the feedback function at X_t. Differently, β_t in (4.7) encodes the entire feedback function: Somehow, it coincides with the entire function $\tilde{\beta}_t$. In this framework, the cost functional (1.2) should read

$$\begin{aligned}
&\mathbb{E}\bigg[\int_0^T \bigg(\frac{1}{2}\sum_{j \neq X_t} |\beta_t^{X_t^\beta, j}|^2 + F^{X_t^\beta}(m_t)\bigg) dt + G^{X_T^\beta}(m_T)\bigg] \\
&\quad = \sum_{i \in E} \bigg[\int_0^T p_t^i \bigg(\frac{1}{2}\sum_{j \neq i} |\beta_t^{i,j}|^2 + F^i(m_t)\bigg) dt + p_T^i G^i(m_T)\bigg] \\
&\quad := J\big((\beta_t)_t; (p_t)_t; (m_t)_t\big),
\end{aligned} \tag{4.9}$$

for a given continuous (and here deterministic) path $(m_t)_t$ with values in $\mathcal{P}(E)$.

It it then quite standard to compute the corresponding HJB equation. Since E is finite, it becomes a mere ordinary differential equation. Accordingly, the MFG system (1.4) becomes

$$\begin{cases}
-\partial_t u_t^i + \dfrac{1}{2}\displaystyle\sum_{j \in E} (u_t^i - u_t^j)_+^2 = F^i(m_t), \\
\partial_t m_t^i - \displaystyle\sum_{j \in E} \big[m_t^j (u_t^j - u_t^i)_+ - m_t^i (u_t^i - u_t^j)_+\big] = 0, \quad i \in E, \ t \in [0,T].
\end{cases} \tag{4.10}$$

Once the system (4.10) is solved, the optimal feedback is given by $\alpha_t^{i,j} = (u_t^i - u_t^j)_+, i \neq j$. Consistently with the notation introduced in (4.8), the probability measure m_t is identified with the collection of nonnegative weights $(m_t^i)_{i \in E}$, with the latter satisfying $\sum_{j \in E} m_t^j = 1$.

Adding a common noise. Differently from (1.4), (4.10) is a finite-dimensional forward–backward system. The question is then how to find a suitable form of finite-dimensional common noise that forces existence and uniqueness. Although it is very similar to the question addressed in Section 4.1 for linear–quadratic quadratic mean field games, the problem is in fact formulated in a different way. Indeed, the analysis carried out in Section 4.1 mostly relies on the probabilistic formulation of the mean field game or, equivalently, on the equation for the dynamics of the reference player. Instead, we want to use here the equation for the dynamics of the population, as it is more adapted to the model in hand. This raises some subtle issues on the structure of the common noise as we want the resulting Fokker–Planck equation to preserve the simplex. In other words, we want to find a form of simplex-valued diffusion process. A very famous instance is the so-called Wright–Fisher process, originally introduced in stochastic models for population genetics (see Kimura [77]). Recast in our framework (the analysis of which is taken from Bayraktar, Cecchin, Cohen, and Delarue [13]), it leads to the following stochastic version of the MFG system (4.10):

$$\begin{cases} d_t u_t^i = \Big(\frac{1}{2} \sum_{j\in E} (u_t^i - u_t^j)_+^2 - F^i(m_t) - \sqrt{\epsilon} \sum_{j \in E} \sqrt{m_t^i m_t^j} (v_t^{i,i,j} - v_t^{i,j,i}) \Big) dt \\ \qquad\qquad + \sum_{j,k\in E} v_t^{i,j,k} dW_t^{j,k}, \\ d_t m_t^i = \sum_{j\in E} [m_t^j (u_t^j - u_t^i)_+ - m_t^i (u_t^i - u_t^j)_+] dt + \sqrt{\epsilon} \sum_{j\in E} \sqrt{m_t^i m_t^j} d[W_t^{i,j} - W_t^{j,i}], \end{cases} \tag{4.11}$$

for $i \in E$ and $t \in [0,T]$. In the above, $(W_t)_t = ((W_t^{i,j})_{0\le t\le T})_{(i,j)\in E^2}$ is a collection of independent Brownian motions. Following the notations introduced in the statement of Theorem 2.3, it is very useful to distinguish the space carrying $(W_t)_t$ from the space carrying the idiosyncratic noise underpinning the transition rates (4.6): The former will be denoted by $(\Omega_0, \mathbb{F}_0, \mathbb{P}_0)$ and the latter by $(\Omega_1, \mathbb{F}_1, \mathbb{P}_1)$. Accordingly, the expectations are respectively denoted by $\mathbb{E}_0$ and $\mathbb{E}_1$. The product measure on the product space is denoted by $\mathbb{P}$ and the corresponding expectation by $\mathbb{E}$. Intuitively, the process $(v_t)_t$ in the above backward equation plays the same role as the process $(v_t)_t$ in (1.8). In particular, it is worth observing that, in both cases, the process $(v_t)_t$ appears in the dt term of the backward equation.

Before we provide the interpretation of the above system in terms of a mean field game, we write down the resulting form of the master equation (see again [13]):

$$\begin{aligned} &\partial_t U_t^i(m) + \epsilon \sum_{j,k\in E} (m_j \delta_{j,k} - m_j m_k) \partial^2_{m_j m_k} U_t^i(m) \\ &\quad + \sum_{j,k\in E} p_k \big(U_t^k(m) - U_t^j(m)\big)_+ \big(\partial_{m_j} U_t^i(m) - \partial_{m_k} U_t^i(m)\big) \\ &\quad + 2\epsilon \sum_{j\in E} p_j \big(\partial_{m_i} U_t^i(m) - \partial_{m_j} U_t^i(m)\big) - \frac{1}{2} \sum_{j\in E} \big(U_t^i(m) - U_t^j(m)\big)_+^2 \\ &\quad + F^i(m) = 0, \end{aligned} \tag{4.12}$$

with the boundary condition $U_T^i(m) = g^i(m)$. The terms induced by the common noise are those featuring the prefactor ϵ. In particular, the master equation without common noise is obtained by letting $\epsilon = 0$. The main impact of the common noise is to generate the second-order differential operator

$$\epsilon \sum_{j,k\in E} (m_j \delta_{j,k} - m_j m_k) \partial^2_{m_j m_k}, \tag{4.13}$$

which is called a (purely second-order) Kimura operator on the simplex of dimension $|E| - 1$. In both (4.12) and (4.13), the derivatives should be formally regarded as intrinsic derivatives on the simplex, with gradients being of dimension $|E| - 1$. However, it is also possible to assume that U has a smooth extension to an $|E|$-dimensional subset of the simplex and then to consider the derivatives as standard $|E|$-dimensional derivatives. It is worth noticing that the resulting derivative used in (4.12) and (4.13) does not coincide with the derivative D_m introduced in Theorem 2.1. The infinite-dimensional analogue of the derivative used in (4.12) and (4.13) is the so-called flat, or linear, functional derivative. In short, it is the restriction, to the space of probability measures, of the derivative on the space of signed measures. It is a potential of the derivative D_m.

Forcing uniqueness. A key feature of the Kimura operator (4.13) lies in the structure of the diffusion matrix: it degenerates near the boundary of the simplex. This is somehow the price to pay to construct a diffusion process that does not leave the simplex. As an issue, it makes much more difficult any attempt to prove smoothing properties (which are precisely what we need in order to force uniqueness to the system (4.11), in full analogy with the result stated in Proposition 4.1). However, a relevant form of Schauder's theory was established in the monograph by Epstein and Mazzeo [59]. In short, it says that linear equations driven by Kimura operators have classical solutions (with a suitable behavior at the boundary) if the first-order and source terms are just Hölder continuous in time and space. This is, however, not sufficient to get similar results for the nonlinear equation (4.12), as the first-order term therein is driven by the solution itself. As U is easily shown to be bounded from a straightforward application of the maximum principle, the next step to fill the gap is thus to prove the following form of *a priori* estimate: For some Hölder exponent, the Hölder norm of a classical solution to a homogeneous parabolic equation driven by a Kimura operator is bounded in terms of the L^∞-norm of the solution and the Hölder norm of the initial condition (if the equation is set forward) or of the terminal condition (if the equation is set backward as (4.12) is). Nevertheless, it is not possible to prove this in full generality. In short, the best results that are known require the presence in (4.13) of a first-order term with strictly positive components along inward normal directions to the boundary. When applying this principle to (4.12), we are led to consider the following modified version of the master equation:

$$\begin{aligned}
&\partial_t U_t^i(m) + \epsilon \sum_{j,k\in E} (m_j \delta_{j,k} - m_j m_k)\partial^2_{m_j m_k} U_t^i(m) \\
&\quad + \sum_{j,k\in E} p_k\big[\varphi(m_j) + \big(U_t^k(m) - U_t^j(m)\big)_+\big]\big(\partial_{m_j} U_t^i(m) - \partial_{m_k} U_t^i(m)\big) \\
&\quad + 2\epsilon \sum_{j\in E} m_j\big(\partial_{m_i} U_t^i(m) - \partial_{m_j} U_t^i(m)\big) - \frac{1}{2}\sum_{j\in E}\big(U_t^i(m) - U_t^j(m)\big)_+^2 \\
&\quad + F^i(m) + \sum_{j\in E} \varphi(m_j)\big(U_t^j(m) - U_t^i(m)\big) = 0, \qquad (4.14)
\end{aligned}$$

with the terminal condition $U_T^i(m) = g^i(m)$, for a smooth function φ from $[0,\infty)$ into itself that is nonzero in the neighborhood of 0. This function φ should be regarded as a penalty: when inserted in the transition rates (4.6), it forces the corresponding solution to the Fokker–Planck equation (4.8) to leave the boundary of the simplex (here and below, the notions of boundary and interior of the simplex are understood when $\mathcal{P}(E)$ is regarded as a subset of

$\mathbb{R}^{|E|-1}$). Notice that this additional penalty φ appears in the first-order term on the second line, which is consistent with our preliminary discussion, but also in the zeroth-order term on the last line, which is necessary to have a relevant interpretation of (4.14) as the master equation of a mean field game (see Definition 4.1 below).

The next statement is also taken from [13]:

Theorem 4.2. *We can find a threshold $\kappa_0 > 0$, only depending on ϵ ($\epsilon > 0$), $\|F\|_\infty$, $\|G\|_\infty$, and T, such that, if $\varphi(0) > \kappa_0$, and if F and G are smooth enough, then equation (4.14) has a classical solution, with first-order derivatives in space that are bounded on the whole domain and second-order derivatives in space that are bounded on $[0,T] \times \mathcal{K}$, for any compact subset $\mathcal{K}$ included in the interior of $\mathcal{P}(E)$.*

The existence of a classical solution is then shown to force uniqueness to the corresponding system of characteristics. Due to the presence of the penalty φ, this system does not exactly fit (4.11). The right-version is

$$\begin{cases} d_t u_t^i = \Big(\frac{1}{2} \sum_{j \in E} (u_t^i - u_t^j)_+^2 - F^i(m_t) - \sqrt{\epsilon} \sum_{j \in E} \sqrt{m_t^i m_t^j} (v_t^{i,i,j} - v_t^{i,j,i}) \Big) dt \\ \qquad - \sum_{j \in E} \varphi(m_t^j)(u_t^j - u_t^i) + \sum_{j,k \in E} v_t^{i,j,k} dW_t^{j,k}, \\ d_t m_t^i = \sum_{j \in E} \big[m_t^j \big(\varphi(m_t^i) + (u_t^j - u_t^i)_+ \big) - m_t^i \big(\varphi(m_t^j) + (u_t^i - u_t^j)_+ \big) \big] dt \\ \qquad + \sqrt{\epsilon} \sum_{j \in E} \sqrt{m_t^i m_t^j} \, d\big[W_t^{i,j} - W_t^{j,i} \big], \end{cases} \tag{4.15}$$

for $i \in E$ and $t \in [0,T]$. In line with Theorem 4.2, we have (see again [13]):

Theorem 4.3. *We can find a threshold $\kappa_0 > 0$, only depending on ϵ ($\epsilon > 0$), $\|F\|_\infty$, $\|G\|_\infty$, and T, such that, if $\varphi(0) > \kappa_0$, and if F and G are smooth enough, then the forward–backward system (4.15) has a unique solution when the initial condition $m_0 = (m_0^i)_{i \in E}$ is prescribed in the interior of the simplex.*

To be fair, we should mention that uniqueness holds within a class of solutions with suitable integrability properties. We refer to [13] for the complete version of the statement. As for the constraint on the initial condition, it says that $m_0^i > 0$ for any $i \in E$. The resulting solution $(m_t)_t$ is then shown to stay away from the boundary (which is helpful since the diffusion coefficient in the dynamics of $(m_t)_t$ becomes singular on the boundary). Implicitly, all the statements below are also limited to initial conditions in the interior of the simplex.

It now remains to provide an interpretation of the two systems (4.11) and (4.15) in terms of a mean field game. This goes through the following definition:

Definition 4.1. We say that a $(W_t)_t$-adapted continuous stochastic process $(m_t)_{0 \le t \le T}$ with values in the interior of $\mathcal{P}(E)$ is a solution to the mean field game with common noise of intensity $\sqrt{\epsilon}$ (and without the penalization φ) if $(m_t)_{0 \le t \le T}$ satisfies an equation of the form

$$dm_t^i = \sum_{j \in E} m_t^j \alpha_t^{j,i} dt + \sqrt{\epsilon} \sum_{j \in E} \sqrt{m_t^i m_t^j} \, d(W_t^{i,j} - W_t^{j,i}), \quad t \in [0,T], \tag{4.16}$$

for a bounded $(W_t)_t$-progressively-measurable process $((\alpha_t^{i,j})_{i,j \in E})_t$ satisfying (4.7) and, for any other bounded $(W_t)_t$-progressively-measurable process $((\beta_t^{i,j})_{i,j \in E})_t$ satisfying (4.7), the solution of the equation

$$dp_t^i = \sum_{j \in E} p_t^j \beta_t^{j,i} dt + \sqrt{\epsilon} p_t^i \sum_{j \in E} \sqrt{\frac{m_t^j}{m_t^i}} \, d(W_t^{i,j} - W_t^{j,i}), \quad t \in [0,T], \tag{4.17}$$

satisfies the inequality

$$\mathbb{E}^0\big[J\big((\beta_t)_t; (p_t)_t; (m_t)_t\big)\big] \ge \mathbb{E}^0\big[J\big((\alpha_t)_t; (m_t)_t; (m_t)_t\big)\big],$$

with J being defined as in (4.9).

A similar definition holds for the mean field game with common noise of intensity $\sqrt{\epsilon}$ in the presence of the penalization φ. It suffices to replace $(\alpha_t^{j,i})_t$ by $(\varphi(m_t^i) + \alpha_t^{j,i})_t$ in (4.16) and $(\beta_t^{j,i})_t$ by $(\varphi(p_t^i) + \beta_t^{j,i})_t$ in (4.17).

In fact, Definition 4.1 is rather subtle. Differently from the formulation (1.6)–(1.7) used for continuous state spaces, the current one does not provide an explicit formulation of the (private) dynamics of the reference player within the population. In short, Definition 4.1 is missing an equation similar to (1.7). Instead, equation (4.17) should be regarded as a form of Fokker–Planck equation for some marginal statistics of the reference player given the common noise. Actually, it can be proven that there exists a stochastic process $(X_t, Y_t)_{0 \le t \le T}$ with values in the space $E \times \mathbb{R}_+$ such that

$$p_t^i = \mathbb{E}^1[Y_t \mathbf{1}_{\{X_t = i\}}], \quad t \in [0,T], \quad i \in E,$$

with $(Y_t)_t$ satisfying $\mathbb{E}^0\mathbb{E}^1[Y_t] = 1$. In this formulation, X_t should be regarded as the physical state, at time t, of the reference player, with the latter being also assigned a mass Y_t. The mass of the tagged particle is in fact a density on the entire probability space carrying both types of noise. It is a density accounting for the way the reference player perceives the world. In this respect, it is important to note that the process $(p_t)_t$ does not take values in the simplex, but only in the orthant $(\mathbb{R}_+)^{|E|}$. This follows from the linear structure of equation (4.17) (with $(p_t)_t$ as unknown). The linear structure, here with stochastic coefficients, is consistent with the linear structure of the Fokker–Planck equation (4.8). In order to obtain solutions in a relevant space, integrability conditions on these stochastic coefficients are thus necessary, whence the assumption that $((\alpha_t^{i,j})_{i,j \in E})_t$ and $((\beta_t^{i,j})_{i,j \in E})_t$ are bounded.

4.3. Vanishing viscosity

Following the latter two subsections, a natural question is to address the vanishing viscosity limits of the solutions to the mean field game with common noise and to the corresponding parabolic master equation. Both for linear quadratic mean field games and finite state mean field games, uniqueness of the equilibria may be lost in the framework of Proposition 4.1 and Theorem 4.3 when the common noise is removed. This is the same for the corresponding master equation: Classical solutions may cease to exist and, accordingly, weaker notions of solutions are needed; Uniqueness is then a challenging question.

For sure, we could think of other methods for selecting equilibria. For instance, we could think of returning back to the game with N players and then identifying which equilibria coincide with a limit point of the N-equilibrium as N tends to infinity. This is, however, a very difficult road. As easily seen from the uniformly parabolic structure of the system (1.1), the N-player game (at least in the form studied there) satisfies a form of non-degeneracy that is asymptotically lost when N tends to infinity. The study of the large-N limit thus combines two difficulties at the same time: The whole system becomes more and more degenerate (this is a vanishing viscosity limit) and, meanwhile, some propagation of chaos is expected to occur (this is the mean field limit). In contrast, taking the small noise limit in a mean field game with common noise just raises one of these two issues since the mean field limit has already been taken.

Earlier selection results can be found in Bayraktar and Zhang [14], Cecchin, Dai Pra, Fischer, and Pelino, [44] and Delarue and Foguen [52]. Generally speaking, they are stated for mean field games whose equilibria are known to belong to a one-dimensional parametric model. This covers the following two examples: Linear–quadratic mean field games of the same type as in Section 4.1, but with d therein being equal to 1 (which implies in particular that, conditional on the initial state, the equilibria follow Gaussian distributions with a known variance but an unknown mean); Finite state mean field games on a set E containing two elements only (in which case the simplex is one-dimensional). In all these aforementioned works, selection is directly proved by taking the large-N limit in the finite game. Basically, this is possible thanks to the totally ordered structure of $\mathbb{R}$. Moreover, the master equation then reduces to a scalar conservation law and the selected solution is the entropy solution.

When the effective dimension of the model is greater than or equal to 2, things become much more challenging. A way to make the problem simpler is to address the so-called potential case. As explained in Proposition 2.4, potential games are a special kind of mean field games that coincide with the first-order condition of a mean field control problem. When the state space E is finite, this corresponds to the case where F and G satisfy

$$F^i(x,m) = \partial_{m_i}\mathcal{F}(m), \quad G^i(x,\mu) = \partial_{m_i}\mathcal{G}(m), \quad i \in E, \tag{4.18}$$

for two real-valued functions $\mathcal{F}$ and $\mathcal{G}$ defined on $\mathcal{P}(E)$. In words, F and G are identified with (respectively) the gradient of $\mathcal{F}$ and the gradient of $\mathcal{G}$. The identification is, however, a bit subtle since, formally, these two gradients should be identified with vectors of dimension $|E|-1$. In turn, this says that the above condition could be slightly relaxed: In short, it would suffice to identify the projections of F and G onto the orthogonal complement of $(1,\dots,1)$ (which should be regarded as the tangent space to the simplex) with the corresponding intrinsic gradient. Anyway, given $\mathcal{F}$ and $\mathcal{G}$, we can consider the deterministic optimal control problem

$$\inf_{(\beta_t)_t} \mathcal{J}\big((\beta_t)_{0\le t\le T}\big),$$

associated with the cost functional

$$\mathcal{J}\big((\beta_t)_{0\le t\le T}\big) = \int_0^T \left(\frac{1}{2}\sum_{i,j\in E: i\neq j} p_t^i |\beta_t^{i,j}|^2 + \mathcal{F}(p_t)\right) dt + \mathcal{G}(p_T), \tag{4.19}$$

and with the dynamics (4.8) (for a given initial condition $(p_0^i)_{i\in E}$), the function $(\beta_t)_t$ satisfying the constraint (4.7) at any time. Then, very similar to Proposition 2.4, we have

Proposition 4.4. *Let $m_0 = (m_0^i)_{i\in E}$ be an initial condition in the interior of the simplex. Under condition (4.18), any bounded minimizer $((\beta_t^{i,j})_{i,j\in E})_{0\le t\le T}$ of the cost functional (4.19), with $(p_0^i)_{i\in E} = (m_0^i)_{i\in E}$ as initial condition in (4.8), yields a solution to the mean field game (4.9).*

The proof follows from a standard application of the Pontryagin principle. The adjoint variable then identifies with $(u_t^i)_{i\in E}$ in the system (4.10). In the statement, the two constraints on $(p_0^i)_{i\in E}$ (which is required to have strictly positive coordinates) and on $((\beta_t^{i,j})_{i,j\in E})_{0\le t\le T}$ (which is required to be bounded) force the corresponding trajectory (4.8) to stay away from the boundary of the simplex (when the latter is viewed as an open subset of dimension of $|E|-1$). This guarantees that, along the trajectory (4.8), the extended Hamiltonian has a unique minimizer, as required in the application of the Pontryagin principle.

Selection of equilibria. Obviously, there is no converse to Proposition 4.4: The set of equilibria of a potential mean field game may be strictly larger than the set of minimizers to the corresponding mean field control. In this respect, a natural selection principle would consist in ruling out the equilibria that are not minimizers of the corresponding mean field control. Very interestingly, this principle is consistent with the results mentioned above when the state space E is of cardinality 2. Indeed, any mean field game on a finite state space with two elements is potential. As such, it derives from a mean field control problem. In particular, a natural question is to ask whether the solutions to the mean field game that are selected by taking the large-N limit in the finite game associated with (4.8)–(4.9) are also minimizers of the corresponding mean field control problem. The answer is yes. The same result remains open when $|E| \ge 3$. However, a simpler (but still interesting) question is to ask whether, under the same property (4.18) as before, the vanishing viscosity limits of the mean field game with common noise, as defined in Section 4.2, are minimizers of the corresponding mean field control problem. Formulated in this way, this question is also open. The main issue is that, in the presence of the common noise (and of the additional penalization φ that is necessary to guarantee the conclusion of Theorem 4.3), the mean field game is no longer potential. In order to get a potential form in (4.15), an additional penalization is necessary. Once the game with common noise is potential, it is pretty easy to take the vanishing viscosity limit in the mean field control problem that lies above. The following result is taken from Cecchin and Delarue [43]:

Theorem 4.5. *Let $m_0 = (m_0^i)_{i\in E}$ be an initial condition in the interior of the simplex. For any $\epsilon > 0$, we can find two functions $\varphi^\epsilon : [0,+\infty) \to [0,+\infty)$ and $(F^{\epsilon,i} : \mathcal{P}(E) \to \mathbb{R})_{i\in E}$, with φ^ϵ converging to 0 uniformly on any compact subset of $(0,+\infty)$, such that:*

(1) *The system (4.11) obtained by replacing (F,φ) by $(F^\epsilon,\varphi^\epsilon)$ is uniquely solvable, the solution of the forward equation being denoted by $((m_t^{\epsilon,i})_{0\le t\le T})_{i\in E}$;*

(2) *Any weak limits of the sequence of the laws of the processes $((m_t^{\epsilon,i})_t)_{i\in E}$ has a support included in the set of minimizers of $\mathcal{J}$ in (4.19) (with the same initial condition);*

(3) *There exists a family of positive reals $(\delta_\epsilon)_\epsilon$, satisfying $\lim_{\epsilon\to 0}\delta_\epsilon = 0$, such that the trajectories $((m_t^{\epsilon,i})_t)_{i\in E}$ form $(\delta_\epsilon)_\epsilon$-approximate solutions of the mean field game with common noise of intensity $\sqrt{\epsilon}$ and with penalty φ^ϵ, as defined in Definition 4.1. In clear, if $((p_t^i)_{0\le t\le T})_{i\in E}$ solves (4.17) (with $(\beta_t^{j,i}, m_t)_t$ being replaced by $(\varphi^\epsilon(p_t^i) + \beta_t^{j,i}, m_t^\epsilon)_t$ and with the prescription that $(\beta_t)_t$ is bounded by a fixed constant), then*

$$\left|\mathbb{E}^0 \sum_{i\in E}\left[\int_0^T p_t^i F^i(m_t^\epsilon)dt\right] - \mathbb{E}^0 \sum_{i\in E}\left[\int_0^T p_t^i F^{\epsilon,i}(m_t^\epsilon)dt\right]\right| \le \delta_\epsilon.$$

Obviously, item (3) says that the additional penalization in the definition of F^ϵ has a limited impact: The solution to the mean field game associated with the cost functional driven by F^ϵ is almost a solution of the same mean field game but associated with the cost functional driven by F. For sure, the notion of approximated solution is here consistent with the standard notion of approximated Nash equilibria: when the reference player in the population chooses a feedback function different from that chosen by the others, the best possible improvement (in the cost functional) tends to 0 with ϵ.

Interestingly, uniqueness of the minimizers (and thus of the limit points) in the second item of Theorem 4.5 is in fact the typical situation. Indeed, standard control theory says that the control problem (4.19)–(4.8) has in fact a unique minimizer at any point in time and space where the corresponding value function, which we denote by $\mathcal{V}$, is differentiable (see **[23]**). However, it is a standard exercise to prove that $\mathcal{V}$ is Lipschitz continuous, hence the fact that uniqueness holds for almost every starting point (in time and space) when the simplex is equipped with the $(|E|-1)$-dimensional Lebesgue measure. Obviously, in the formulation (4.19)–(4.8), the initial time is 0, but there is no difficulty in adapting the definition to any other time $t \in [0,T]$.

Selection of solutions to the master equation. In fact, $\mathcal{V}$ plays an even more important role in the analysis of the vanishing viscosity limit as it also permits characterizing the limit of the solutions to the second-order master equation (4.14) associated with the common noise of intensity $\sqrt{\epsilon}$, with the penalty φ^ϵ and with the penalization F^ϵ (for the same choices φ^ϵ and F^ϵ as in the statement of Theorem 4.5). The next statement result is also taken from **[43]**:

Theorem 4.6. *With the same notation as in the statement of Theorem 4.5 and with U^ϵ denoting the solution to equation (4.14) when $\varphi \equiv \varphi^\epsilon$ and $F \equiv F^\epsilon$ therein, the limit*

$$\lim_{\epsilon \to 0}\big[U^{\epsilon,i}(t,q) - U^{\epsilon,j}(t,q)\big] = \partial_{m_i}\mathcal{V}(t,q) - \partial_{m_j}\mathcal{V}(t,q)$$

holds for almost-every $(t,q) \in [0,T] \times \mathcal{P}(E)$ and for any $i, j \in E$, where $\mathcal{V}$ is the value function of the control problem (4.19)–(4.8).

As we have already explained, the gradient of the value function exists almost-everywhere in time and space. It also important to note that the argument in the limit is not the solution of the master equation itself but the finite differences of it. In short, the limit of the master equation is just identified in dimension $|E| - 1$, which is fully consistent with the fact that the gradient of $\mathcal{V}$ is a vector of dimension $|E| - 1$. Alternatively, the above statement provides the limiting form of the feedback function used in the mean field game with a common noise of intensity $\sqrt{\epsilon}$. The $|E|$-dimensional limit of the function U^ϵ itself can be found by computing the minimal in cost (4.9) when the environment $(m_t)_t$ therein is the solution of the control problem (4.19)–(4.8).

In accordance with the program outlined above, it is a natural question to ask whether the limit established in Theorem 4.6 can be characterized in terms of the original master equation itself (i.e., the master equation (4.12) but with $\epsilon = 0$ therein). The answer is positive. As shown in [43], the master equation can be written in a conservative form. Following earlier results of Kružkov [80] and Lions [89], this conservative form has a unique solution that is bounded and satisfies a weak one-sided Lipschitz condition in space. It coincides the gradient of the value function $\mathcal{V}$. This recovers the existing results when $|E| = 2$.

4.4. Complements and open problems

Even when the state space is finite, the extension of the above results to the nonpotential case is a highly difficult problem.

Another interesting problem is to extend the same results to mean field games on continuous state spaces. The main issue is to define a suitable form of common noise. In short, this requires addressing stochastic processes with values in the infinite-dimensional space $\mathcal{P}_2(\mathbb{R}^d)$ and with sufficiently strong smoothing properties, which is known to be a challenging problem in the literature. There are earlier results in this direction, but they are not sufficient to handle the nonlinearities that make the spice of mean field games: We refer, for instance, to Stannat [100] for smoothing estimates of the Fleming–Viot process, which is an infinite-dimensional version of the Wright–Fisher noise underpinning the forward–backward system. In short, the Dirichlet form of the Fleming–Viot process is driven by the aforementioned linear-functional derivative (which provides a potential of the derivative D_m). In the meantime, the construction of a process with a Dirichlet form associated with the derivative D_m has been addressed in a series of works initiated in von Renesse and Sturm [103], but no canonical definition has yet been given. Another strategy in order to force uniqueness consists in embedding the problem in some L^2 space: following the idea underpinning

Definition 2.1, we can indeed see the unknown in a mean field game as a flow of random variables and not as a flow of probability measures. This makes it possible to use noises in Hilbert spaces. However, this destroys the mean field structure of the problem. We refer to Delarue [49] for results in this direction.

From another perspective, it is important to note that common noises in finite state mean field games can be defined a manner different from (4.11). We refer in particular to Bertucci, Lasry, and Lions [19], the key idea of which is to force the finite-player system to have many simultaneous jumps at some random times prescribed by the common noise. The reader may also have a look at [6], which provides a discrete point of view on the system (1.8). As far as the formulation (4.11) is concerned, a study of the convergence problem, very much in the spirit of Theorem 2.7, is available in [12].

5. FURTHER PROSPECTIVES AND RELATED OPEN PROBLEMS

We will now briefly review some aspects of the theory that we have not covered so far. This is only a summary presentation which demonstrates (if needed) that the field has diversified into many active branches.

5.1. Analysis of the MFG system and of the master equation

The MFG system. In the last two decades there has been a large amount of research on MFG systems of the type (which generalize (1.4)):

$$\begin{cases} \text{(i)} \quad -\partial_t u_t(x) - \Delta u_t(x) + H\big(t, x, Du_t(x), m_t\big) = 0 & \text{in } (0,T) \times \mathbb{R}^d, \\ \text{(ii)} \quad \partial_t m_t(x) - \Delta m_t(x) - \operatorname{div}\Big(m_t(x) D_p H\big(t, x, Du_t(x), m_t\big)\Big) = 0 & \text{in } (0,T) \times \mathbb{R}^d, \\ \text{(iii)} \quad m_0(x) = \tilde{m}_0(x), \quad u(T, x) = g(x, m_T) & \text{in } \mathbb{R}^d, \end{cases}$$

and of more general (fully nonlinear) MFG systems (where $D_p H$ is the derivative of the Hamiltonian $H(t, x, p, m)$ with respect to p). It is impossible to give a complete overview of this literature: we refer to the survey [6] and to the references therein for a general presentation of this literature. The question of the existence and regularity of the solutions has been investigated in several frameworks: When the dependence of the Hamiltonian is local (depending on the pointwise value of the density), existence of classical solutions is discussed, for instance, by Cardaliaguet, Lasry, Lions, and Porretta in [29] and by Gomes, Pimentel, and Voskanyan in [68]; Porretta introduced in [97] a notion of a weak solution for these problems and proved uniqueness in this framework. The MFG system can also be set with other boundary conditions: for instance, Neumann boundary condition (Bardi and Cirant [11]), optimal stopping (Bertucci [17]), state constraints (Cannarsa, Capuani, and Cardaliaguet [22]). Mean field games can be also stated in networks (Camilli and Marchi [21] or Achdou, Dao, Ley, and Tchou [3]). Problems with congestion or with density constraints are discussed by Lions [88], Achdou and Porretta [5] and Cardaliaguet, Mészáros, and Santambrogio [32].

Variational aspects. In general, the analysis of the MFG system relies on fixed point techniques. In some frameworks (the local coupling case, for instance) it is possible to use variational methods. This turns out to be very useful for problems in which the diffusion is degenerate and for which this approach allows building weak solutions; see, for instance, the papers by Cardaliaguet and Graber [26] (first-order problems with a local coupling) and by Cardaliaguet, Graber, Porretta, and Tonon [27] (for degenerate second-order problems with a local coupling). Refining earlier result by Lions [88], Santambrogio [99] combined variational techniques with ideas from optimal transport to obtain nice regularity results of first-order MFG systems (system without diffusion). Many other references and results can be found in the survey by Santambrogio in [6].

The master equation. The analysis of the master equation has attracted some attention in the recent years, refining the earlier results [25,36,45,64,88]. Without trying to be exhaustive, one can quote the recent papers: Bertucci [18] for a notion of a weak solution under monotonicity conditions; Cardaliaguet, Porretta, and Cirant [24] for the construction of solutions to general master equations (with common noise or for a major player, see the paragraph below) on short time intervals using a kind of Trotter–Kato scheme; Gangbo, Mészáros, Mou, and Zhang [63] for the existence of a classical solution to the master equation outside the classical monotone framework, obtained by using instead conditions related with displacement convexity. Let us underline that a suitable notion of weak (discontinuous) notion of solution for the master equation is still missing (see, however, Section 4 and [43] by Cecchin and Delarue in the finite-state framework and for potential problems).

MFG problem with a major player. In general, mean field games address problems with a single homogeneous population. It is, of course, not the only interesting configuration. Among the many possible generalizations, one can mention the MFG problems with a major player, in which a controller (the major player) interacts with a population. This problem, first introduced by Huang [72], has been studied (among many other references) by Carmona and Zhu [42] by a probabilistic approach, and in [24] using the master equation. It is related with the principal–agent problems with one principal and infinitely many agents, as explained by Elie, Mastrolia, and Possamaï [57].

5.2. Mean field games of control

Mean field games of controls (sometimes also called extended mean field games) are mean field games in which players interact through the joint distribution of their positions and their controls. Many models in economics are of this type (for instance, agents interact through the price of a good that depends directly on their collective decisions to buy or sell). This kind of problem was first discussed by Gomes and Voskanyan [69]. Weak solutions have been built through a probabilistic approach by Carmona and Lacker [39]. In [35, CHAPTER 4], Carmona and Delarue pointed out the specific structure of the corresponding MFG system, which involves two fixed point problems (the classical one and a static one used to build the distribution of positions and controls from the distribution of positions and the control feedback). This MFG system was also studied in Cardaliaguet and Lehalle [31] (existence of weak solutions for problems with degenerate diffusions) and in Achdou and Kobeissi [4] (classical solution in the diffusive case and with very general interactions). Very recently Djete [54] proved the convergence of open-loop Nash equilibria for the N-player game as N tends to infinity.

5.3. Numerical methods and learning

The fixed-point nature of MFG equilibria makes them difficult to approximate and implement in practice. In the work by Achdou and Capuzzo Dolcetta [2], the authors explain how to reproduce numerically the forward–backward nature of the MFG system in order to obtain convergent numerical schemes, thus starting a series of works of the subject. An up-to-date literature on the numerical methods for mean field games, including effective methods for decoupling the two equations, can be found in Achdou's survey on this topic [6]. Recently, other works have also demonstrated the possible efficiency of tools from machine learning within this complex framework: standard equations for characterizing the equilibria may be approximately solved by means of a neural network; see, for instance, Carmona and Laurière [40,41].

Intimately related to the numerical approximation, the intriguing question of learning ("how do the MFG equilibria actually appear?") has attracted some attention. One of the first results in this direction is the transposition to MFG games of the classical fictitious play by Cardaliaguet and Hadikhanloo [28]: assuming that players know the model and that the MFG problem is potential, the method explains how players could converge to an MFG equilibrium after playing the game many times. Elie, Pérolat, Laurière, Geist, and Pietquin [58] study the effects of diverse reinforcement learning algorithms for agents with no prior information on an MFG equilibrium and learn their policy through repeated experiments. The very recent paper Delarue and Vasileiadis [53] shows that common noise may also serve as an exploration noise for learning the solution of a mean field game.

5.4. Mean field control

Mean field control (MFC) is a distinct theory from mean field games, but both theories are connected in many ways. For instance, potential games are a typical instance of mean field games that are solved by the minimizers of an MFC problem, see Proposition 2.4. The very aim of MFC theory is to address minimization problems set over Kolmogorov equations (when formulated by means of PDEs) or over McKean–Valsov equations (when formulated in a probabilistic fashion). From a particle point of view, MFC problems provide an asymptotic description of large systems of weakly interacting controlled agents who cooperate in order to minimize some common cost. Therefore, in contrast with MFGs, the agents no longer compete, and the solutions of the two problems are different. As such, this asks for a new proof of the corresponding convergence problem. We refer, for instance, to Lacker [82] for a proof based on compactness arguments, and to Djete [55] for a similar result but for models including the law of the control in the mean field interaction. As for the analysis of MFC themselves, the related value function satisfies a form of Hamilton–Jacobi equation. Similar to the master equation, the Hamilton–Jacobi equation is set on the space of probability measures, but Lions' lifting procedure allows lifting it onto an L^2 space (see [88]). This observation can be used in order to adapt the notion of viscosity solutions and thus to handle less regular solutions. We refer to [65] for a recent contribution in this direction in the first-order case, namely when the dynamics of the players are deterministic. In the presence of an idiosyncratic noise in the dynamics (so-called second-order case), the theory is still in progress and a complete theory of existence and uniqueness of viscosity solutions has not yet been achieved. We refer to Cosso and Pham [48] for an overview of the stakes.

FUNDING

F. Delarue acknowledges the financial support of French ANR project ANR-19-P3IA-0002—3IA Côte d'Azur—Nice—Interdisciplinary Institute for Artificial Intelligence.

REFERENCES

[1] Y. Achdou, F. J. Buera, J.-M. Lasry, P.-L. Lions, and B. Moll, Partial differential equation models in macroeconomics. *Philos. Trans. R. Soc. A* **372** (2014), 20130397.

[2] Y. Achdou and I. Capuzzo-Dolcetta, Mean field games: numerical methods. *SIAM J. Numer. Anal.* **48** (2010), 1136–1162.

[3] Y. Achdou, M. K. Dao, O. Ley, and N. Tchou, A class of infinite horizon mean field games on networks. *Netw. Heterog. Media* **14** (2019), 537–566.

[4] Y. Achdou and Z. Kobeissi, Mean field games of controls: finite difference approximations. 2020, arXiv:2003.03968.

[5] Y. Achdou and A. Porretta, Mean field games with congestion. *Ann. Inst. H. Poincaré Anal. Non Linéaire* **35** (2018), 443–480.

[6] Y. Achdou P, F. Cardaliaguet, A. Delarue, Porretta, and F. Santambrogio, *Mean field games*. Lecture Notes in Math., CIME 2281, Springer, 2020.

[7] S. R. Aiyagari, Uninsured idiosyncratic risk and aggregate saving. *Q. J. Econ.* **109** (1994), 659–684.

[8] S. Albeverio, Y. G. Kondratiev, and M. Röckner, Analysis and geometry on configuration spaces. *J. Funct. Anal.* **154** (1998), 444–500.

[9] L. Ambrosio, N. Gigli, and G. Savaré, *Gradient flows: in metric spaces and in the space of probability measures*. Springer, 2006.

[10] M. Bardi, Explicit solutions of some linear quadratic mean field games. *Netw. Heterog. Media* **7** (2012), 243–261.

[11] M. Bardi and M. Cirant, Uniqueness of solutions in mean field games with several populations and Neumann conditions. In *PDE models for multi-agent phenomena*, pp. 1–20, Springer, 2018.

[12] E. Bayraktar, A. Cecchin, A. Cohen, and F. Delarue, Finite state mean field games with Wright–Fisher common noise as limits of N-player weighted games. 2020, arXiv:2012.04845.

[13] E. Bayraktar, A. Cecchin, A. Cohen, and F. Delarue, Finite state mean field games with Wright–Fisher common noise. *J. Math. Pures Appl.* **147** (2021), 98–162.

[14] E. Bayraktar and X. Zhang, On non-uniqueness in mean field games. *Proc. Amer. Math. Soc.* **148** (2020), 4091–4106.

[15] A. Bensoussan, J. Frehse, and P. Yam, *Mean field games and mean field type control theory*. Springer Briefs in Mathematics 101, Springer, 2013.

[16] A. Bensoussan, K. C. J. Sung, S. C. P. Yam, and S. P. Yung, Linear quadratic mean field games. *J. Optim. Theory Appl.* **169** (2016), 469–529.

[17] C. Bertucci, Optimal stopping in mean field games, an obstacle problem approach. *J. Math. Pures Appl.* **120** (2018), 165–194.

[18] C. Bertucci, Monotone solutions for mean field games master equations: finite state space and optimal stopping. 2020, arXiv:2007.11854.

[19] C. Bertucci, J.-M. Lasry, and P.-L. Lions, Some remarks on mean field games. *Comm. Partial Differential Equations* **44** (2019), 205–227.

[20] R. Buckdahn, J. Li, S. Peng, and C. Rainer, Mean-field stochastic differential equations and associated PDEs. *Ann. Probab.* **45** (2017), 824–878.

[21] F. Camilli and C. Marchi, Stationary mean field games systems defined on networks. *SIAM J. Control Optim.* **54** (2016), no. 2, 1085–1103.

[22] P. Cannarsa, R. Capuani, and P. Cardaliaguet, Mean field games with state constraints: from mild to pointwise solutions of the pde system. 2018, arXiv:1812.11374.

[23] P. Cannarsa and C. Sinestrari, *Semiconcave functions, Hamilton–Jacobi equations, and optimal control*. Progr. Nonlinear Differential Equations Appl. 58, Birkhäuser Boston, Inc., Boston, MA, 2004.

[24] P. Cardaliaguet, M. Cirant, and A. Porretta, Splitting methods and short time existence for the master equations in mean field games. 2020, arXiv:2001.10406.

[25] P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions, *The master equation and the convergence problem in mean field games*. Ann. of Math. Stud. 201, Princeton University Press, 2019.

[26] P. Cardaliaguet and P. J. Graber, Mean field games systems of first order. *ESAIM Control Optim. Calc. Var.* **21** (2015), 690–722.

[27] P. Cardaliaguet, P. J. Graber, A. Porretta, and D. Tonon, Second order mean field games with degenerate diffusion and local coupling. *NoDEA Nonlinear Differential Equations Appl.* **22** (2015), 1287–1317.

[28] P. Cardaliaguet and S. Hadikhanloo, Learning in mean field games: the fictitious play. *ESAIM Control Optim. Calc. Var.* **23** (2017), no. 2, 569–591.

[29] P. Cardaliaguet, J.-M. Lasry, P.-L. Lions, and A. Porretta, Long time average of mean field games. *Netw. Heterog. Media* **7** (2012).

[30] P. Cardaliaguet, J.-M. Lasry, P.-L. Lions, and A. Porretta, Long time average of mean field games with a nonlocal coupling. *SIAM J. Control Optim.* **51** (2013), 3558–3591.

[31] P. Cardaliaguet and C.-A. Lehalle, Mean field game of controls and an application to trade crowding. *Math. Financ. Econ.* **12** (2018), 335–363.
[32] P. Cardaliaguet, A. R. Mészáros, and F. Santambrogio, First order mean field games with density constraints: pressure equals price. *SIAM J. Control Optim.* **54** (2016), 2672–2709.
[33] P. Cardaliaguet and A. Porretta, Long time behavior of the master equation in mean field game theory. *Anal. PDE* **12** (2019), 1397–1453.
[34] R. Carmona and F. Delarue, Probabilistic analysis of mean-field games. *SIAM J. Control Optim.* **51** (2013), 2705–2734.
[35] R. Carmona and F. Delarue, *Probabilistic theory of mean field games with applications. I. Mean field FBSDEs, control, and games*. Probab. Theory Stoch. Model. 83, Springer, Cham, 2018.
[36] R. Carmona and F. Delarue, *Probabilistic theory of mean field games with applications. II. Mean field games with common noise and master equations*. Probab. Theory Stoch. Model. 84, Springer, Cham, 2018.
[37] R. Carmona, F. Delarue, and A. Lachapelle, Control of McKean–Vlasov versus mean field games. *Math. Financ. Econ.* **7** (2013), 131–166.
[38] R. Carmona, F. Delarue, and D. Lacker, Mean field games with common noise. *Ann. Probab.* **44** (2016), 3740–3803.
[39] R. Carmona and D. Lacker, A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.* **25** (2015), 1189–1231.
[40] R. Carmona and M. Laurière, Convergence analysis of machine learning algorithms for the numerical solution of mean-field control and games: I – The ergodic case. 2019, arXiv:1907.05980.
[41] R. Carmona and M. Laurière, Convergence analysis of machine learning algorithms for the numerical solution of mean-field control and games: II – The finite horizon case. 2019, arXiv:1908.01613.
[42] R. Carmona and X. Zhu, A probabilistic approach to mean field games with major and minor players. *Ann. Appl. Probab.* **26** (2016), 1535–1580.
[43] A. Cecchin and F. Delarue, Selection by vanishing common noise for potential finite state mean field games. *Comm. Partial Differential Equations* (2021). DOI 10.1080/03605302.2021.1955256.
[44] A. Cecchin, P. D. Pra, M. Fischer, and G. Pelino, On the convergence problem in mean field games: a two state model without uniqueness. *SIAM J. Control Optim.* **57** (2019), 2443–2466.
[45] J.-F. Chassagneux, D. Crisan, and F. Delarue, Classical solutions to the master equation for large population equilibria. 2015, arXiv:1411.3009.
[46] M. Cirant, On the existence of oscillating solutions in non-monotone mean-field games. *J. Differential Equations* **266** (2019), 8067–8093.
[47] M. Cirant and A. Porretta, Long time behavior and turnpike solutions in mildly non-monotone mean field games. *ESAIM Control Optim. Calc. Var.* **27** (2021), 86, 40 pp.

[48] A. Cosso and H. Pham, Zero-sum stochastic differential games of generalized McKean–Vlasov type. *J. Math. Pures Appl.* **129** (2019), 180–212.

[49] F. Delarue, Restoring uniqueness to mean-field games by randomizing the equilibria. *Stoch. Partial Differ. Equ. Anal. Comput.* **7** (2019), 598–678.

[50] F. Delarue, D. Lacker, and K. Ramanan, From the master equation to mean field game limit theory: a central limit theorem. *Electron. J. Probab.* **24** (2019), 1–54.

[51] F. Delarue, D. Lacker, and K. Ramanan, From the master equation to mean field game limit theory: large deviations and concentration of measure. *Ann. Probab.* **48** (2020), 211–263.

[52] F. Delarue and R. F. Tchuendom, Selection of equilibria in a linear quadratic mean-field game. *Stochastic Process. Appl.* **130** (2020), 1000–1040.

[53] F. Delarue and A. Vasileiadis, Exploration noise for learning linear–quadratic mean field games. 2021, arXiv:2107.00839.

[54] M. F. Djete, Mean field games of controls: on the convergence of Nash equilibria. 2020, arXiv:2006.12993.

[55] M. F. Djete, Extended mean field control problem: a propagation of chaos result. 2020, arXiv:2006.12996.

[56] J. Doncel, N. Gast, and B. Gaujal, Are mean-field games the limits of finite stochastic games? *ACM SIGMETRICS Perform. Eval. Rev.* **44** (2016), no. 2, 18–20. Special Issue on the Workshop on MAthematical performance Modeling and Analysis (MAMA 2016).

[57] R. Elie, T. Mastrolia, and D. Possamaï, A tale of a principal and many, many agents. *Math. Oper. Res.* **44** (2019), 440–467.

[58] R. Elie, J. Pérolat, M. Laurière, M. Geist, and O. Pietquin, Approximate fictitious play for mean field games. 2019, arXiv:1907.02633.

[59] C. L. Epstein and R. Mazzeo, *Degenerate diffusion operators arising in population biology*. Ann. of Math. Stud. 185, Princeton University Press, Princeton, NJ, 2013.

[60] A. Fathi, Weak KAM theory: the connection between Aubry–Mather theory and viscosity solutions of the Hamilton–Jacobi equation. In *Proceedings of the International Congress of Mathematicians* 3, pp. 597–621, Kyung Moon Sa, Seoul, 2014.

[61] M. Fischer, On the connection between symmetric N-player games and mean field games. *Ann. Appl. Probab.* **27** (2017), 757–810.

[62] N. Fournier and A. Guillin, On the rate of convergence in Wasserstein distance of the empirical measure. *Probab. Theory Related Fields* **162** (2015), 707–738.

[63] W. Gangbo, A. R. Mészáros, C. Mou, and J. Zhang, Mean field games master equations with non-separable Hamiltonians and displacement monotonicity. 2021, arXiv:2101.12362.

[64] W. Gangbo and A. Swiech, Existence of a solution to an equation arising from the theory of mean field games. *J. Differential Equations* **259** (2015), 6573–6643.

[65] W. Gangbo and A. Tudorascu, On differentiability in the Wasserstein space and well-posedness for Hamilton–Jacobi equations. *J. Math. Pures Appl.* **125** (2019), 119–174.

[66] D. A. Gomes, J. Mohr, and R. R. Souza, Discrete time, finite state space mean field games. *J. Math. Pures Appl.* **93** (2010), 308–328.

[67] D. A. Gomes, J. Mohr, and R. R. Souza, Continuous time finite state mean field games. *Appl. Math. Optim.* **68** (2013), 99–143.

[68] D. A. Gomes, E. A. Pimentel, and V. Voskanyan, *Regularity theory for mean-field game systems*. Springer, Berlin, 2016.

[69] D. A. Gomes and V. K. Voskanyan, Extended deterministic mean-field games. *SIAM J. Control Optim.* **54** (2016), 1030–1055.

[70] O. Guéant, Existence and uniqueness result for mean field games with congestion effect on graphs. *Appl. Math. Optim.* **72** (2015), 291–303.

[71] O. Gueant, J.-M. Lasry, and P.-L. Lions, Mean field games and applications. In *Paris–Princeton lectures on mathematical finance*, pp. 205–266, Springer, Berlin, Heidelberg, 2010.

[72] M. Huang, Large-population LQG games involving a major player: the Nash certainty equivalence principle. *SIAM J. Control Optim.* **48** (2010), 3318–3353.

[73] M. Huang, P. E. Caines, and R. P. Malhamé, Individual and mass behaviour in large population stochastic wireless power control problems: centralized and Nash equilibrium solutions. In *Proceedings of the 42nd IEEE Conference on Decision and Control* 1, pp. 98–103, IEEE, 2003.

[74] M. Huang, R. P. Malhamé, and P. E. Caines, Large population stochastic dynamic games: closed-loop Mckean–Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.* **6** (2006), 221–252.

[75] M. Huang, R. P. Malhamé, and P. E. Caines, Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior and decentralized ε-Nash equilibria. *IEEE Trans. Automat. Control* **52** (2007), 1560–1571.

[76] M. Kac, Foundations of kinetic theory. In *Proceedings of the third Berkeley symposium on mathematical statistics and probability, 1954–1955* III, pp. 171–197, University of California Press, Berkeley and Los Angeles, CA, 1956.

[77] M. Kimura, Diffusion models in population genetics. *J. Appl. Probab.* **1** (1964), 177–232.

[78] V. N. Kolokoltsov, *Nonlinear Markov processes and kinetic equations* 182, Cambridge University Press, 2010.

[79] P. Krusell and A. A. Smith Jr., Income and wealth heterogeneity in the macroeconomy. *J. Polit. Econ.* **106** (1998), 867–896.

[80] S. N. Kružkov, Generalized solutions of nonlinear equations of the first order with several independent variables. II. *Mat. Sb.* **72** (1967), 108–134.

[81] D. Lacker, A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields* **165** (2016), 581–648.

[82] D. Lacker, Limit theory for controlled McKean–Vlasov dynamics. *SIAM J. Control Optim.* **55** (2017), 1641–1672.
[83] D. Lacker, On the convergence of closed-loop Nash equilibria to the mean field game limit. *Ann. Appl. Probab.* **30** (2020), 1693–1761.
[84] O. A. Ladyzhenskaia, V. A. Solonnikov, and N. Uraltseva, *Linear and quasi-linear equations of parabolic type* 23, Amer. Math. Soc., Providence, RI, 1998.
[85] J.-M. Lasry and P.-L. Lions, Jeux à champ moyen. I – le cas stationnaire. *C. R. Math.* **343** (2006), 619–625.
[86] J.-M. Lasry and P.-L. Lions, Jeux à champ moyen. II – horizon fini et contrôle optimal. *C. R. Math.* **343** (2006), 679–684.
[87] J.-M. Lasry and P.-L. Lions, Mean field games. *Jpn. J. Math.* **2** (2007), 229–260.
[88] P.-L. Lions, Cours au College de France, 2007–2012.
[89] P.-L. Lions, *Generalized solutions of Hamilton-Jacobi equations*. Res. Notes Math. 69, Pitman (Advanced Publishing Program), Boston, Mass.-London, 1982.
[90] P.-L. Lions, G. Papanicolaou, and S. S. Varadhan, *Homogenization of Hamilton–Jacobi equations*. 1986. Unpublished work.
[91] R. E. Lucas Jr. and B. Moll, Knowledge growth and the allocation of time. *J. Polit. Econ.* **122** (2014), no. 1, 1–51.
[92] H. P. McKean, A class of Markov processes associated with nonlinear parabolic equations. *Proc. Natl. Acad. Sci. USA* **56** (1966), 1907–1911.
[93] S. Mischler and C. Mouhot, Kac's program in kinetic theory. *Invent. Math.* **193** (2013), 1–147.
[94] S. Mischler, C. Mouhot, and B. Wennberg, A new approach to quantitative propagation of chaos for drift, diffusion and jump processes. *Probab. Theory Related Fields* **161** (2015), 1–59.
[95] G. Papanicolaou, L. Ryzhik, and K. Velcheva, Travelling waves in a mean field learning model. 2020, arXiv:2002.06287.
[96] S. Peng, Stochastic Hamilton–Jacobi–Bellman equations. *SIAM J. Control Optim.* **30** (1992), 284–304.
[97] A. Porretta, Weak solutions to Fokker–Planck equations and mean field games. *Arch. Ration. Mech. Anal.* **216** (2015), 1–62.
[98] A. Porretta and L. Rossi, Travelling waves in a mean field game model of knowledge diffusion. 2020, arXiv:2010.10828.
[99] F. Santambrogio, Regularity via duality in calculus of variations and degenerate elliptic pdes. *J. Math. Anal. Appl.* **457** (2018), 1649–1674.
[100] W. Stannat, Long-time behaviour and regularity properties of transition semigroups of Fleming–Viot processes. *Probab. Theory Related Fields* **122** (2002), 431–469.
[101] A.-S. Sznitman, Topics in propagation of chaos. In *Ecole d'été de probabilités de Saint-Flour XIX—1989*, pp. 165–251, Springer, 1991.
[102] R. F. Tchuendom, Uniqueness for linear–quadratic mean field games with common noise. *Dyn. Games Appl.* **8** (2018), 199–210.

[103] M.-K. von Renesse and K. T. Sturm, Entropic measure and Wasserstein diffusion. *Ann. Probab.* **37** (2009), 1114–1191.

PIERRE CARDALIAGUET

Université Paris Dauphine-PSL, Place du Mal de Lattre de Tassigny, 75775 Paris Cedex 16, France, cardaliaguet@ceremade.dauphine.fr

FRANÇOIS DELARUE

Université Côte d'Azur, Laboratoire J. A. Dieudonné, CNRS, Parc Valrose, 06108 Nice Cedex 02, France, delarue@univ-cotedazur.fr