The Annals of Probability
2015, Vol. 43, No. 5, 2647–2700
DOI: 10.1214/14-AOP946

FORWARD–BACKWARD STOCHASTIC DIFFERENTIAL EQUATIONS AND CONTROLLED MCKEAN–VLASOV DYNAMICS

BY RENÉ CARMONA[1] AND FRANÇOIS DELARUE

Princeton University and Université de Nice

The purpose of this paper is to provide a detailed probabilistic analysis of the optimal control of nonlinear stochastic dynamical systems of McKean–Vlasov type. Motivated by the recent interest in mean-field games, we highlight the connection and the differences between the two sets of problems. We prove a new version of the stochastic maximum principle and give sufficient conditions for existence of an optimal control. We also provide examples for which our sufficient conditions for existence of an optimal solution are satisfied. Finally we show that our solution to the control problem provides approximate equilibria for large stochastic controlled systems with mean-field interactions when subject to a common policy.

1. Introduction. The purpose of this paper is to provide a detailed probabilistic analysis of the optimal control of nonlinear stochastic dynamical systems of McKean–Vlasov type. The present study is motivated in part by a recent surge of interest in mean-field games.

Stochastic differential equations (SDEs) of McKean–Vlasov type are usually referred to as *nonlinear* SDEs, the term *nonlinear* emphasizing the possible dependence of the coefficients upon the marginal distributions of the solutions, and having no bearing on a possible nonlinear dependence upon the state variable. This special feature of the coefficients, even when the latter are nonrandom, creates nonlocal feedback effects which rule out the standard Markov property. Including the marginal distribution in the state of the system could restore the Markov property at the cost of a leap in complexity of the state of the process. The latter would have to include a probability measure and subsequently become infinite dimensional. While the analysis of the infinitesimal generator could be done with tools developed for infinite dimensional differential operators, the standard differential calculus, even in infinite dimension, would have a hard time capturing the fact that the second component of the state process would have to match the statistical distribution of the first component. Yet, the complexity of a pathwise analysis is still reasonable and can be efficiently handled by means of standard tools in stochastic calculus. This suggests the design of a completely probabilistic approach for tackling the optimal control of these *nonlinear* systems.

Received March 2013; revised March 2014.

[1]Supported in part by NSF Grant DMS-08-06591.

MSC2010 subject classifications. Primary 93E20; secondary 60H10, 60K35.

Key words and phrases. Stochastic control, McKean–Vlasov diffusion, stochastic Pontryagin principle, mean-field interaction, mean-field forward–backward stochastic differential equation.

The analysis of SDEs of McKean–Vlasov type has a long history. These equations were first introduced by McKean [18] to provide a rigorous treatment of special nonlinear partial differential equations (PDEs). Later on, they were studied for their own sake and in a more general setting. See, for example, [13, 19, 23] for existence and uniqueness results. Properties of the solutions have been studied in the framework of the propagation of chaos, as McKean–Vlasov equations appear as effective equations describing dynamical systems of large populations subject to a mean-field interaction.

However, the optimal control of dynamics driven by McKean–Vlasov SDEs seems to be a brand new problem, incredulously ignored in the stochastic control literature. Solving a McKean–Vlasov SDE is commonly done by a fixed point argument: First one fixes a set of candidates for the marginal distributions, and then one solves the resulting standard SDE, the fixed point argument being to demand that the marginal distributions of the solution be equal to the marginal distributions one started from. A stochastic control problem adds an extra optimization layer to the fixed point. This formulation bears a lot of resemblance to the *mean-field game* (MFG) problem as originally formulated by Lasry and Lions [14–16] and, simultaneously, by Caines, Huang and Malhamé [12]. The similarities and the differences between the two problems were identified and discussed in [10], where it is clearly emphasized that optimizing first and searching for a fixed point afterwards leads to the solution of a mean-field game problem, while finding the fixed point first and then optimizing afterwards leads to the solution of the optimal control of McKean–Vlasov SDEs. The solutions to both problems describe equilibrium states of large populations of individuals whose interactions and objective functions are of mean-field type. The differences between these notions of equilibrium are subtle and depend upon the formulation of the optimization component of the equilibrium model. Simple examples of linear quadratic models are provided in [10] to illustrate these differences. The connection between the optimal control of McKean–Vlasov SDEs and large population equilibriums is addressed in Section 6 below.

Since the dynamics described by McKean–Vlasov SDEs are genuinely non-Markovian, it is natural to approach the optimization problem using a suitable version of the Pontryagin stochastic maximum principle, rather than a contrived adaptation of the Hamilton–Jacobi–Bellman paradigm. The stochastic Pontryagin approach is based on the Introduction (and the manipulation) of adjoint processes defined as solutions of adjoint backward stochastic differential equations (BSDEs). Writing these equations involves partial derivatives of the Hamiltonian function with respect to the state variable, but, in the case of McKean–Vlasov SDEs, the marginal distributions of the solutions are full-fledged variables of the Hamiltonian function and need to be differentiated in search for criticality. We believe that this is the main reason for the stalemate in the existing literature: Only dynamics depending upon moments of the marginal distributions have been considered so far, and in those cases, differentiability with respect to the measure can be done by standard calculus chain rules; see, for example, [2, 4]. The right notion of differentiability was identified by Lions in his Lectures at the Collège de France [17], and thoroughly explained in Cardaliaguet's notes [6]. Our first contribution is to review this notion of differentiability, generalize the definition of adjoint processes to the case of measure dependent Hamiltonians and enhance the standard proofs of the necessary condition for optimality to cover the case of McKean–Vlasov dynamics. We also generalize the sufficient condition for convex Hamiltonians to the case of McKean–Vlasov dynamics.

The stochastic Pontryagin principle is a very powerful tool. However, the insights it provides come at the price of restrictive assumptions on the models. Case in point, our results rely on a set of technical assumptions which limit the class of models to dynamics given by coefficients which are essentially linear in the state, control and measure variables and costs which are convex in the state and the control variables. While seemingly restrictive, these assumptions are typical in the applications of the stochastic Pontryagin principle to control problems. Note that the convexity of the space of controls is an assumption which is only made for the sake of simplicity. More general spaces could be handled at the cost of using spike variation techniques and adding one extra adjoint equation. See, for example, [26], Chapter 3, for a discussion of the classical (i.e., non-McKean–Vlasov) case. Without the motivation from specific applications, we chose to refrain from providing this level of generality and avoid an excessive overhead in notation and technicalities which, in our mind, can only obscure the thrust of the paper.

The necessary part of the stochastic Pontryagin principle suggests to search the control set for a candidate minimizing the Hamiltonian function (this goes often under the name of satisfying the Isaacs condition), while the sufficient part proposes to insert the formula for the minimizer into the forward equation governing the dynamics, and the adjoint backward equation providing the adjoint processes. The presence of the minimizer in both of these equations creates a strong coupling between the forward and backward equations, and the solution of the control problem reduces de facto to the solution of a forward backward stochastic differential equation (FBSDE). Implementing this strategy in the present situation leads to the analysis of FBSDEs where the marginal distributions of the solutions appear in the coefficients of the equations. We call these equations mean-field FBSDEs, or FBSDEs of McKean–Vlasov type. To the best of our knowledge, these equations have not been studied before. A rather general existence result was recently proposed in [8], but one of the assumptions (boundedness of the coefficients with respect to the state variable) precludes the application of this result to the linear quadratic (LQ) models often used as benchmarks in stochastic control. Here, we take advantage of the convexity of the Hamiltonian to apply the *continuation method* introduced in [21], and prove existence and uniqueness of the solution of the FBSDE at hand, extending and refining the results of [8] to the models considered in this paper. Restoring the Markov property by extending the state space, as alluded to earlier, we identify the backward component of the solution of this FBSDE to a function of the forward component and its marginal distribution. This function is known as the *decoupling field* of the FBSDE. In the classical cases, it can be found by solving a PDE. In the present set-up, such a PDE would be infinite dimensional as it would involve differentiation with respect to the state of the forward dynamics as well as its distribution. Precisely, it would read as an infinite dimensional Hamilton–Jacobi–Bellman equation. This PDE is related to what Lions calls the *master equation* (or *master PDE*) in his lectures on mean-field games [17]. However, we emphasize the fact that, in the framework of mean-field games, the master equation is not a Hamilton–Jacobi–Bellman equation, as the MFG equilibriums do not solve an optimization problem. As explained in [17], formulating the master PDE is already very delicate and technical. Solving it is even more elusive. For that reason, we choose to leave the discussion of the connections between McKean–Vlasov FBSDEs and master equations to a forthcoming paper; see [7].

As already mentioned, McKean–Vlasov SDEs describe the asymptotic behavior of mean-field interacting particle systems as the number of particles tends to infinity. Asymptotically, particles become independent of each others, and each single one of them satisfies the same McKean–Vlasov SDE; see [13, 23]. Such a phenomenon is usually referred to as *propagation of chaos*. When particle evolutions are controlled, each particle attempting to minimize an energy functional, it is natural to investigate equilibriums inside the population, especially when the number of particles tends to infinity. In such a framework, the theory of mean-field games focuses on asymptotic Nash equilibriums, which correspond to particles (or agents) choosing their strategies in an *individual* way. Alternatively, we show that our solution of the optimal control of McKean–Vlasov SDEs provides strategies leading to approximate equilibriums when agents use a *common* policy. Not surprisingly, the identification of approximate equilibriums in *feedback form* requires strong regularity properties of the decoupling field of the FBSDE. We prove these properties by tailor-made arguments.

We end this Introduction with a quick summary of the contents of the paper. The notation and definitions specific to McKean–Vlasov SDEs are given in Section 2. The following Section 3 provides the definitions and first properties of the differentiation and convexity of functions of measures. It also introduces the Hamiltonian function of the stochastic control problem, and the definition of the adjoint processes as solutions of McKean–Vlasov BSDEs. Section 4 is devoted to the generalization of the Pontryagin stochastic maximum principle. Proofs of the sufficient and necessary conditions for optimality are given in full detail. Section 5 contains the technical details of the solution of the mean-field FBSDE derived from the stochastic maximum principle. This FBSDE is central to the analysis of the control problem. In Section 6, we make the connection with controlled interacting particle systems. We show how our strategy permits to construct approximate equilibriums for large populations using a common control policy function. The paper closes with an Appendix devoted to the proof of a crucial technical result about the Lipschitz property of functions of state variables and measures.

2. Probabilistic set-up of McKean–Vlasov equations. In what follows, we assume that $W = (W_t)_{0\le t\le T}$ is an m-dimensional standard Wiener process defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and $\mathbb{F} = (\mathcal{F}_t)_{0\le t\le T}$ is its natural filtration possibly augmented with an independent σ-algebra $\mathcal{F}_0$. For each random variable/vector or stochastic process X, we denote by $\mathbb{P}_X$ the law (alternatively called the distribution) of X.

The stochastic dynamics of interest in this paper are given by a stochastic process $X = (X_t)_{0\le t\le T}$, satisfying a nonlinear SDE of the form

$$dX_t = b(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dt + \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dW_t, \qquad 0 \le t \le T, \tag{1}$$

where the drift and diffusion coefficients of the state X_t of the system are given by the pair of deterministic functions $(b, \sigma) : [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times A \to \mathbb{R}^d \times \mathbb{R}^{d\times m}$, and $\alpha = (\alpha_t)_{0\le t\le T}$ is a progressively measurable process with values in a measurable space $(A, \mathcal{A})$. Typically, A will be an open subset of an Euclidean space $\mathbb{R}^k$ and $\mathcal{A}$ the σ-field induced by the Borel σ-field of this Euclidean space. Also, for each measurable space $(E, \mathcal{E})$, we use the notation $\mathcal{P}(E)$ for the space of probability measures on $(E, \mathcal{E})$, assuming that the σ-field $\mathcal{E}$ on which the measures are defined is understood. When E is a metric or a normed space (most often $\mathbb{R}^d$), we denote by $\mathcal{P}_p(E)$ the subspace of $\mathcal{P}(E)$ of the probability measures of order p, namely those probability measures which integrate the pth power of the distance to a fixed point whose choice is irrelevant in the definition of $\mathcal{P}_p(E)$. The term *nonlinear*, used for describing (1), does not refer to the fact that the coefficients b and σ could be nonlinear functions of x, but instead to the fact that they depend not only on the value of the unknown process X_t at time t, but also on its marginal distribution $\mathbb{P}_{X_t}$. In that framework, we shall assume that the drift coefficient b and the volatility σ satisfy the following assumptions:

(A1) the function $[0, T] \ni t \mapsto (b, \sigma)(t, 0, \delta_0, 0) \in \mathbb{R}^d \times \mathbb{R}^{d\times m}$ is square integrable;

(A2) $\exists c > 0, \forall t \in [0, T], \forall \alpha, \alpha' \in A, \forall x, x' \in \mathbb{R}^d, \forall \mu, \mu' \in \mathcal{P}_2(\mathbb{R}^d)$,

$$\begin{aligned}
&\bigl|b(t, x, \mu, \alpha) - b\bigl(t, x', \mu', \alpha'\bigr)\bigr| + \bigl|\sigma(t, x, \mu, \alpha) - \sigma\bigl(t, x', \mu', \alpha'\bigr)\bigr| \\
&\qquad \le c\bigl[\bigl|x - x'\bigr| + \bigl|\alpha - \alpha'\bigr| + W_2\bigl(\mu, \mu'\bigr)\bigr],
\end{aligned}$$

where $W_2(\mu, \mu')$ denotes the 2-Wasserstein distance. For a general $p > 1$, the p-Wasserstein distance $W_p(\mu, \mu')$ is defined on $\mathcal{P}_p(E)$ by

$$\begin{aligned}
W_p\bigl(\mu, \mu'\bigr) = \inf\biggl\{\biggl[\int_{E\times E} |x - y|^p \pi(dx, dy)\biggr]^{1/p}; \\
\pi \in \mathcal{P}_2(E \times E) \text{ with marginals } \mu \text{ and } \mu'\biggr\}.
\end{aligned}$$

Notice that if X and X' are random variables of order p, then $W_p(\mathbb{P}_X, \mathbb{P}_{X'}) \le [\mathbb{E}|X - X'|^p]^{1/p}$.

The set $\mathbb{A}$ of so-called *admissible* control processes α is defined as the set of A-valued progressively measurable processes $\alpha \in \mathbb{H}^{2,k}$, where $\mathbb{H}^{2,n}$ denotes the Hilbert space

$$\mathbb{H}^{2,n} := \left\{ Z \in \mathbb{H}^{0,n};\ \mathbb{E} \int_0^T |Z_s|^2 \, ds < +\infty \right\}$$

with $\mathbb{H}^{0,n}$ standing for the collection of all $\mathbb{R}^n$-valued progressively measurable processes on $[0, T]$. By (A.1) and (A.2), any $\alpha \in \mathbb{A}$ satisfies

$$\mathbb{E} \int_0^T \big[\big| b(t, 0, \delta_0, \alpha_t) \big|^2 + \big| \sigma(t, 0, \delta_0, \alpha_t) \big|^2 \big] \, dt < +\infty.$$

Together with the Lipschitz assumption (A.2), this guarantees that, for any $\alpha \in \mathbb{A}$, there exists a unique solution $X = X^{\alpha}$ of (1), and that moreover this solution satisfies

$$\mathbb{E} \sup_{0 \le t \le T} |X_t|^p < +\infty \tag{2}$$

for every $p \in [1, 2]$. See, for example, [13, 23] for a proof. The stochastic optimization problem which we consider is to minimize the objective function

$$J(\alpha) = \mathbb{E} \left\{ \int_0^T f(t, X_t, \mathbb{P}_{X_t}, \alpha_t) \, dt + g(X_T, \mathbb{P}_{X_T}) \right\}, \tag{3}$$

over the set $\mathbb{A}$ of admissible control processes $\alpha = (\alpha_t)_{0 \le t \le T}$. The *running cost* function f is a real valued deterministic function on $[0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times A$, and the *terminal cost* function g is a real valued deterministic function on $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$. In particular, it is worth mentioning that all the coefficients involved in the definition of the stochastic optimization problem are assumed to be deterministic. Part of the results obtained in the paper can be extended to random coefficients. We refer the reader to Section 5.7. Assumptions on the cost functions f and g will be spelled out later.

The McKean–Vlasov dynamics posited in (1) are sometimes called of *mean-field type*. This is justified by the fact that the uncontrolled stochastic differential equations of McKean–Vlasov type first appeared in the infinite particle limit of large systems of particles with mean-field interactions; see, for example, [13, 19, 23]. Typically, the dynamics of such a system of N particles are given by a system of N SDEs of the form

$$dX_t^i = b^i\big(t, X_t^1, \dots, X_t^N\big) \, dt + \sigma^i\big(t, X_t^1, \dots, X_t^N\big) \, dW_t^i,$$

where the W^i's are N independent standard Wiener processes in $\mathbb{R}^m$, the σ^i's are N deterministic functions from $[0, T] \times \mathbb{R}^{N \times d}$ into the space of $d \times m$ real matrices and the b^i's are N deterministic functions from $[0, T] \times \mathbb{R}^{N \times d}$ into $\mathbb{R}^d$.

The interaction between the particles is said to be of mean-field type when the functions b^i and σ^i are of the form

$$b^i(t, x_1, \dots, x_N) = b(t, x_i, \bar{\mu}^N_{\underline{x}}),$$
$$\sigma^i(t, x_1, \dots, x_N) = \sigma(t, x_i, \bar{\mu}^N_{\underline{x}}), \qquad i = 1, \dots, N,$$

for some deterministic function b from $[0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$ into $\mathbb{R}^d$, and σ from $[0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$ into the space of $d \times m$ real matrices. Here, for each N-tuple $\underline{x} = (x_1, \dots, x_N)$, we denote by $\bar{\mu}^N_{\underline{x}}$, or $\bar{\mu}^N$ when no confusion is possible, the empirical probability measure defined by

$$\bar{\mu}^N_{\underline{x}}(dx') = \frac{1}{N} \sum_{j=1}^{N} \delta_{x_j}(dx'), \tag{4}$$

and for each x, by δ_x the unit point mass (Dirac) measure at x. We shall come back to this formulation of the problem in the last section of the paper when we use results from the propagation of chaos to construct approximate equilibriums. Indeed, as highlighted in Section 6, the optimization problem for controlled McKean–Vlasov dynamics considered here, also reads as the limit as N tends to infinity, of the optimal states of N interacting players using a common policy.

We emphasize that the optimization problem (3) differs from the optimization problem encountered in the theory of mean-field games. Differences between these optimization problems are discussed in [10]. When solving a mean-field game problem, the optimization of the cost functional (3) is performed for a fixed flow of probability measures. In other words, the argument $(\mathbb{P}_{X_t})_{0 \le t \le T}$ in (1) and (3) is kept fixed as α varies, and the controlled processes are driven by the same flow of measures, which is not necessarily the flow of marginal distributions of the process $(X_t)_{0 \le t \le T}$, but merely an input. Solving the corresponding mean-field game then consists of identifying a flow of probability measures, that is, an input, such that the optimal states have precisely the input as flow of statistical distributions.

Useful notation. Given a function $h : \mathbb{R}^d \to \mathbb{R}$ and a vector $p \in \mathbb{R}^d$, we will denote by $\partial h(x) \cdot p$ the action of the gradient of h onto p. When $h : \mathbb{R}^d \to \mathbb{R}^\ell$, we will also denote by $\partial h(x) \cdot p$ the action of the gradient of h onto p, the resulting quantity being an element of $\mathbb{R}^\ell$. When $h : \mathbb{R}^d \to \mathbb{R}^\ell$ and $p \in \mathbb{R}^\ell$, we will denote by $\partial h(x) \odot p$ the element of $\mathbb{R}^d$ defined by $\partial_x[h(x) \cdot p]$ where $\cdot$ is understood as the inner product in $\mathbb{R}^\ell$.

3. Preliminaries. We now introduce the notation and concepts needed for the analysis of the stochastic optimization problem associated with the control of McKean–Vlasov dynamics.

3.1. *Differentiability and convexity of functions of measures*. There are many notions of differentiability for functions defined on spaces of measures, and recent progress in the theory of optimal transportation have put some of them in the limelight. See, for example, [1, 24] for exposés of these geometric approaches in textbook form. However, the notion of differentiability, which we find convenient for the type of stochastic control problem studied in this paper, is slightly different. It is more of a functional analytic nature. We believe that it was introduced by Lions in his lectures at the *Collège de France*. See [6] for a readable account. This notion of differentiability is based on the *lifting* of functions $\mathcal{P}_2(\mathbb{R}^d) \ni \mu \mapsto H(\mu)$ into functions $\tilde{H}$ defined on the Hilbert space $L^2(\tilde{\Omega}; \mathbb{R}^d)$ over some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ by setting $\tilde{H}(\tilde{X}) = H(\tilde{\mathbb{P}}_{\tilde{X}})$ for $\tilde{X} \in L^2(\tilde{\Omega}; \mathbb{R}^d)$, $\tilde{\Omega}$ being a Polish space and $\tilde{\mathbb{P}}$ an atomless measure. Then a function H is said to be differentiable at $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$ if there exists a random variable $\tilde{X}_0$ with law μ_0, in other words satisfying $\tilde{\mathbb{P}}_{\tilde{X}_0} = \mu_0$, such that the lifted function $\tilde{H}$ is Fréchet differentiable at $\tilde{X}_0$. Whenever this is the case, the Fréchet derivative of $\tilde{H}$ at $\tilde{X}_0$ can be viewed as an element of $L^2(\tilde{\Omega}; \mathbb{R}^d)$ by identifying $L^2(\tilde{\Omega}; \mathbb{R}^d)$ and its dual. It then turns out that its distribution depends only upon the law μ_0 and not upon the particular random variable $\tilde{X}_0$ having distribution μ_0. See Section 6 in [6] for details. This Fréchet derivative $[D\tilde{H}](\tilde{X}_0)$ is called the representation of the derivative of H at μ_0 along the variable $\tilde{X}_0$. Since it is viewed as an element of $L^2(\tilde{\Omega}; \mathbb{R}^d)$, by definition,

$$H(\mu) = H(\mu_0) + [D\tilde{H}](\tilde{X}_0) \cdot (\tilde{X} - \tilde{X}_0) + o(\|\tilde{X} - \tilde{X}_0\|_2) \tag{5}$$

whenever $\tilde{X}$ and $\tilde{X}_0$ are random variables with distributions μ and μ_0, respectively, the dot product being here the L^2-inner product over $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ and $\|\cdot\|_2$ the associated norm. It is shown in [6] that, as a random variable, this Fréchet derivative is of the form $\tilde{h}(\tilde{X}_0)$ for some deterministic measurable function $\tilde{h} : \mathbb{R}^d \to \mathbb{R}^d$, which is uniquely defined μ_0-almost everywhere on $\mathbb{R}^d$. The equivalence class of $\tilde{h}$ in $L^2(\mathbb{R}^d, \mu_0)$ being uniquely defined, we can denote it by $\partial_\mu H(\mu_0)$ [or $\partial H(\mu_0)$ when no confusion is possible]. We shall call $\partial_\mu H(\mu_0)$ the derivative of H at μ_0 and most often identify it with a function $\partial_\mu H(\mu_0)(\cdot) : \mathbb{R}^d \ni x \mapsto \partial_\mu H(\mu_0)(x) \in \mathbb{R}^d$ [or by $\partial H(\mu_0)(\cdot)$ when no confusion is possible]. Notice that $\partial_\mu H(\mu_0)$ allows us to express $[D\tilde{H}](\tilde{X}_0)$ as a function of any random variable $\tilde{X}_0$ with distribution μ_0, irrespective of where this random variable is defined. In particular, the differentiation formula (5) is invariant by modification of the space $\tilde{\Omega}$ and of the variables $\tilde{X}_0$ and $\tilde{X}$ used for the representation of H, in the sense that $[D\tilde{H}](\tilde{X}_0)$ always reads as $\partial_\mu H(\mu_0)(\tilde{X}_0)$, whatever the choices of $\tilde{\Omega}$, $\tilde{X}_0$ and $\tilde{X}$ are. It is plain to see how this works when the function H is of the form

$$H(\mu) = \int_{\mathbb{R}^d} h(x)\mu(dx) = \langle h, \mu \rangle \tag{6}$$

for some scalar differentiable function h defined on $\mathbb{R}^d$. Indeed, in this case, $\tilde{H}(\tilde{X}) = \tilde{\mathbb{E}}[h(\tilde{X})]$ and $D\tilde{H}(\tilde{X}) \cdot \tilde{Y} = \tilde{\mathbb{E}}[\partial h(\tilde{X}) \cdot \tilde{Y}]$, and we can think of $\partial_\mu H(\mu)$

as the deterministic function ∂h. We shall use this particular example to recover the Pontryagin principle proved in [2] for scalar interactions, as a particular case of the general Pontryagin principle which we prove below. Example (6) highlights the fact that this notion of differentiability is very different from the usual one. Indeed, given the fact that the function H defined by (6) is linear in the measure μ, when viewed as an element of the dual of a function space, one should expect the derivative to be h and NOT h'!

A nice, though quite technical, feature of this notion of differentiation is given by the following result.

LEMMA 3.1. *Given a differentiable function $H : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$, one can redefine, for each $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\partial_\mu H(\mu)(\cdot) : \mathbb{R}^d \ni x \mapsto \partial_\mu H(\mu)(x)$ on a μ-negligible set in such a way that the mapping $\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (\mu, x) \mapsto \partial_\mu H(\mu)(x)$ is jointly measurable when $\mathcal{P}_2(\mathbb{R}^d)$ is equipped with the Borel σ-field generated by the weak convergence topology or, indistinguishably, by the 2-Wasserstein topology. Whenever $\partial_\mu H(\mu)(\cdot)$ has a continuous version, the version constructed above for measurability reasons coincides with it.*

The proof is deferred to the Appendix. The notion of differentiability used in this paper is best understood as differentiation of functions of limits of empirical measures (or linear combinations of Dirac point masses) in the directions of the atoms of the measures. We illustrate this fact in the next two propositions.

PROPOSITION 3.2. *Given a function $u : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$ and an integer $N \geq 1$, we define the empirical projection of u onto $\mathbb{R}^d$ by*

$$\bar{u}^N : (\mathbb{R}^d)^N \ni \underline{x} = (x_1, \dots, x_N) \mapsto u\left(\frac{1}{N}\sum_{i=1}^N \delta_{x_i}\right).$$

If u is differentiable on $\mathcal{P}_2(\mathbb{R}^d)$, then $\bar{u}^N$ is differentiable on $(\mathbb{R}^d)^N$ and, for all $i \in \{1, \dots, N\}$,

$$\partial_{x_i} \bar{u}^N(\underline{x}) = \partial_{x_i} \bar{u}^N(x_1, \dots, x_N) = \frac{1}{N} \partial u\left(\frac{1}{N}\sum_{j=1}^N \delta_{x_j}\right)(x_i).$$

PROOF. On $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, consider a uniformly distributed random variable ϑ over the set $\{1, \dots, N\}$. Then, for any fixed $\underline{x} = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N$, x_ϑ is a random variable having the distribution $\bar{\mu}^N = N^{-1}\sum_{i=1}^N \delta_{x_i}$. In particular, with the same notation as above for $\tilde{u}$,

$$\bar{u}^N(\underline{x}) = \bar{u}^N(x_1, \dots, x_N) = \tilde{u}(x_\vartheta).$$

Therefore, for $\underline{h} = (h_1, \dots, h_N) \in (\mathbb{R}^d)^N$,

$$\bar{u}^N(\underline{x} + \underline{h}) = \tilde{u}(x_\vartheta + h_\vartheta) = \tilde{u}(x_\vartheta) + D\tilde{u}(x_\vartheta) \cdot h_\vartheta + o(|h|),$$

the dot product being here the L^2-inner product over $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$, from which we deduce

$$\bar{u}^N(\underline{x} + \underline{h}) = \bar{u}^N(\underline{x}) + \frac{1}{N}\sum_{i=1}^{N} \partial u(\bar{\mu}^N)(x_i)h_i + o(|h|),$$

which is the desired result. $\square$

The mapping $D\tilde{u} : L^2(\tilde{\Omega}; \mathbb{R}^d) \to L^2(\tilde{\Omega}; \mathbb{R}^d)$ is said to be Lipschitz continuous if there exists a constant $C > 0$ such that, for any square integrable random variables $\tilde{X}$ and $\tilde{Y}$ in $L^2(\tilde{\Omega}; \mathbb{R}^d)$, it holds $\|D\tilde{u}(\tilde{X}) - D\tilde{u}(\tilde{Y})\|_2 \le C\|\tilde{X} - \tilde{Y}\|_2$. In such a case, the Lipschitz property can be transferred onto $L^2(\Omega)$ and then rewritten as

$$\mathbb{E}\big[\big|\partial u(\mathbb{P}_X)(X) - \partial u(\mathbb{P}_Y)(Y)\big|^2\big] \le C^2\mathbb{E}\big[|X - Y|^2\big], \tag{7}$$

for any square integrable random variables X and Y in $L^2(\Omega; \mathbb{R}^d)$. From our discussion of the construction of ∂u, notice that, for each μ, $\partial u(\mu)(\cdot)$ is only uniquely defined μ-almost everywhere. The following lemma (the proof of which is deferred to the Appendix) then says that, in the current framework, there is a Lipschitz continuous version of $\partial u(\mu)(\cdot)$:

LEMMA 3.3. *Given a family of Borel-measurable mappings $(v(\mu)(\cdot) : \mathbb{R}^d \to \mathbb{R}^d)_{\mu \in \mathcal{P}_2(\mathbb{R}^d)}$ indexed by the probability measures of order 2 on $\mathbb{R}^d$, assume that there exists a constant C such that, for any square integrable random variables ξ and ξ' in $L^2(\Omega; \mathbb{R}^d)$, it holds*

$$\mathbb{E}\big[\big|v(\mathbb{P}_\xi)(\xi) - v(\mathbb{P}_{\xi'})(\xi')\big|^2\big] \le C^2\mathbb{E}\big[|\xi - \xi'|^2\big]. \tag{8}$$

Then, for each $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, one can redefine $v(\mu)(\cdot)$ on a μ-negligeable set in such a way that

$$\forall x, x' \in \mathbb{R}^d \qquad \big|v(\mu)(x) - v(\mu)(x')\big| \le C|x - x'|,$$

for the same C as in (8).

By (7), we can use Lemma 3.3 in order to define $\partial u(\mu)(x)$ for every μ and every x while preserving the Lipschitz property in the variable x. From now on, we shall use this version of ∂u. So, if $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ and X and Y are random variables such that $\mathbb{P}_X = \mu$ and $\mathbb{P}_Y = \nu$, we have

$$\begin{aligned}
&\mathbb{E}\big[\big|\partial u(\mu)(X) - \partial u(\nu)(X)\big|^2\big] \\
&\qquad \le 2\big(\mathbb{E}\big[\big|\partial u(\mu)(X) - \partial u(\nu)(Y)\big|^2\big] + \mathbb{E}\big[\big|\partial u(\nu)(Y) - \partial u(\nu)(X)\big|^2\big]\big) \\
&\qquad \le 4C^2\mathbb{E}\big[|Y - X|^2\big],
\end{aligned}$$

where we used the Lipschitz property (7) of the derivative together with the result of Lemma 3.3 applied to the function $\partial u(\nu)$. Now, taking the infimum over all the couplings (X, Y) with marginals μ and ν, we obtain

$$\inf_{X, \mathbb{P}_X=\mu} \mathbb{E}\big[\big|\partial u(\mu)(X) - \partial u(\nu)(X)\big|^2\big] \le 4C^2 W_2(\mathbb{P}_X, \mathbb{P}_Y)^2,$$

and since the left-hand side depends only upon μ and not on X as long as $\mathbb{P}_X = \mu$, we get

$$\mathbb{E}\big[\big|\partial u(\mu)(X) - \partial u(\nu)(X)\big|^2\big] \le 4C^2 W_2(\mu, \nu)^2. \tag{9}$$

We will use the following consequence of this estimate:

PROPOSITION 3.4. *Let u be a differentiable function on $\mathcal{P}_2(\mathbb{R}^d)$ with a Lipschitz derivative, and let $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $\underline{x} = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N$ and $\underline{y} = (y_1, \dots, y_N) \in (\mathbb{R}^d)^N$. Then, with the same notation as in the statement of Proposition 3.2, we have*

$$\begin{aligned}
&\partial \bar{u}^N(\underline{x}) \cdot (\underline{y} - \underline{x}) \\
&\quad = \frac{1}{N} \sum_{i=1}^N \partial u(\mu)(x_i)(y_i - x_i) + \mathcal{O}\left[W_2(\bar{\mu}_N, \mu) \left(N^{-1} \sum_{i=1}^N |x_i - y_i|^2 \right)^{1/2} \right],
\end{aligned}$$

the dot product being here the usual Euclidean inner product and $\mathcal{O}$ standing for the Landau notation.

PROOF. Using Proposition 3.2, we get

$$\begin{aligned}
&\partial \bar{u}^N(\underline{x}) \cdot (\underline{y} - \underline{x}) \\
&\quad = \sum_{i=1}^N \partial_{x_i} \bar{u}^N(\underline{x})(y_i - x_i) = \frac{1}{N} \sum_{i=1}^N \partial u(\bar{\mu}^N)(x_i)(y_i - x_i) \\
&\quad = \frac{1}{N} \sum_{i=1}^N \partial u(\mu)(x_i)(y_i - x_i) \\
&\qquad + \frac{1}{N} \sum_{i=1}^N \big[\partial u(\bar{\mu}^N)(x_i) - \partial u(\mu)(x_i)\big](y_i - x_i).
\end{aligned}$$

Now, by the Cauchy–Schwarz inequality,

$$\begin{aligned}
&\left| \frac{1}{N} \sum_{i=1}^N \big[\partial u(\bar{\mu}^N)(x_i) - \partial u(\mu)(x_i)\big](y_i - x_i) \right| \\
&\quad \le \left(\frac{1}{N} \sum_{i=1}^N \big|\partial u(\bar{\mu}^N)(x_i) - \partial u(\mu)(x_i)\big|^2 \right)^{1/2} \left(\frac{1}{N} \sum_{i=1}^N |y_i - x_i|^2 \right)^{1/2}
\end{aligned}$$

$$= \left(\tilde{\mathbb{E}}\left[\left|\partial u(\bar{\mu}^N)(x_\vartheta) - \partial u(\mu)(x_\vartheta)\right|^2\right]\right)^{1/2}\left(\frac{1}{N}\sum_{i=1}^N |y_i - x_i|^2\right)^{1/2}$$

$$\leq 2C W_2(\bar{\mu}^N, \mu)\left(\frac{1}{N}\sum_{i=1}^N |y_i - x_i|^2\right)^{1/2},$$

if we use the same notation for ϑ as in the proof of Proposition 3.2, and apply the estimate (9) with $X = x_\vartheta$, $\mu = \bar{\mu}^N$ and $\nu = \mu$. □

REMARK 3.5. We shall use the estimate of Proposition 3.4 when $x_i = X_i$ when the X_i's are independent $\mathbb{R}^d$-valued random variables with common distribution μ. Whenever $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, the law of large numbers ensures that the Wasserstein distance between μ and the empirical measure $\bar{\mu}^N$ tends to 0 a.s., that is,

$$\mathbb{P}\left[\lim_{n\to+\infty} W_2(\bar{\mu}^N, \mu) = 0\right] = 1;$$

see, for example, Section 10 in [22]. Since we can find a constant $C > 0$, independent of N, such that

$$W_2^2(\bar{\mu}^N, \mu) \leq C\left(1 + \frac{1}{N}\sum_{i=1}^N |X_i|^2\right),$$

we deduce from the law of large numbers again that $(W_2^2(\bar{\mu}^N, \mu))_{N\geq 1}$ is uniformly integrable, so that the convergence to 0 also holds in the L^2 sense,

$$\lim_{n\to+\infty} \mathbb{E}\left[W_2^2(\bar{\mu}^N, \mu)\right] = 0. \tag{10}$$

Whenever $\int_{\mathbb{R}^d} |x|^{d+5}\mu(dx) < \infty$, the rate of convergence can be specified. We indeed have the following standard estimate on the Wasserstein distance between μ and the empirical measure $\bar{\mu}^N$:

$$\mathbb{E}\left[W_2^2(\bar{\mu}^N, \mu)\right] \leq C N^{-2/(d+4)}, \tag{11}$$

for some constant $C > 0$. See, for example, Section 10 in [22]. Proposition 3.4 then says that, when N is large, the gradient of $\bar{u}^N$ at the empirical sample $(X_i)_{1\leq i\leq N}$ is close to the sample $(\partial u(\mu)(X_i))_{1\leq i\leq N}$, the accuracy of the approximation being specified in the $L^2(\Omega)$ norm by (11) when μ is sufficiently integrable.

3.2. *Joint differentiability and convexity.* *Joint differentiability.* We often consider functions $h : \mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^d) \ni (x, \mu) \to h(x, \mu) \in \mathbb{R}$ depending on both an n-dimensional x and a probability measure μ. Joint differentiability is then defined according to the same procedure: h is said to be jointly differentiable if the *lifting* $\tilde{h} : \mathbb{R}^n \times L^2(\tilde{\Omega}; \mathbb{R}^d) \ni (x, \tilde{X}) \mapsto h(x, \tilde{\mathbb{P}}_{\tilde{X}})$ is jointly differentiable. In such a case, we can define the partial derivatives in x and μ: they read $\mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^d) \ni (x,\mu) \mapsto \partial_x h(x,\mu)$ and $\mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^d) \ni (x,\mu) \mapsto \partial_\mu h(x,\mu)(\cdot) \in L^2(\mathbb{R}^d,\mu)$, respectively. The partial Fréchet derivative of $\tilde{h}$ in the direction $\tilde{X}$ thus reads $L^2(\tilde{\Omega};\mathbb{R}^d) \ni (x,\tilde{X}) \mapsto D_{\tilde{X}}\tilde{h}(x,\tilde{X}) = \partial_\mu h(x,\tilde{\mathbb{P}}_{\tilde{X}})(\tilde{X}) \in L^2(\tilde{\Omega};\mathbb{R}^d)$. The statement and the proof of Lemma 3.1 can be easily adapted to the joint measurability of $\mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (x,\mu,x') \mapsto \partial_\mu h(x,\mu)(x')$.

We often use the fact that joint continuous differentiability in the two arguments is equivalent with partial differentiability in each of the two arguments and joint continuity of the partial derivatives. Here, the joint continuity of $\partial_x h$ is understood as the joint continuity with respect to the Euclidean distance on $\mathbb{R}^n$ and the Wasserstein distance on $\mathcal{P}_2(\mathbb{R}^d)$. The joint continuity of $\partial_\mu h$ is understood as the joint continuity of the mapping $(x,\tilde{X}) \mapsto \partial_\mu h(x,\tilde{\mathbb{P}}_{\tilde{X}})(\tilde{X})$ from $\mathbb{R}^n \times L^2(\tilde{\Omega};\mathbb{R}^d)$ into $L^2(\tilde{\Omega};\mathbb{R}^d)$.

When the partial derivatives of h are assumed to be Lipschitz-continuous, we can benefit from Lemma 3.3. It says that, for any (x,μ), the *representation* $\mathbb{R}^d \ni x' \mapsto \partial_\mu h(x,\mu)(x')$ makes sense as a Lipschitz function in x' and that an appropriate version of (9) holds true.

Convex functions of measures. We define a notion of convexity associated with this notion of differentiability. A function h on $\mathcal{P}_2(\mathbb{R}^d)$ which is differentiable in the above sense is said to be convex if, for all μ and μ' in $\mathcal{P}_2(\mathbb{R}^d)$, we have

$$h(\mu') - h(\mu) - \tilde{\mathbb{E}}\big[\partial_\mu h(\mu)(\tilde{X}) \cdot (\tilde{X}' - \tilde{X})\big] \geq 0 \tag{12}$$

whenever $\tilde{X}$ and $\tilde{X}'$ are square integrable random variables with distributions μ and μ', respectively. Examples are given in Section 4.3.

More generally, a function h on $\mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^d)$, which is jointly differentiable in the above sense, is said to be convex if for every (x,μ) and (x',μ') in $\mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^d)$, we have

$$\begin{aligned} &h(x',\mu') - h(x,\mu) - \partial_x h(x,\mu) \cdot (x'-x) \\ &\qquad - \tilde{\mathbb{E}}\big[\partial_\mu h(x,\mu)(\tilde{X}) \cdot (\tilde{X}' - \tilde{X})\big] \geq 0 \end{aligned} \tag{13}$$

whenever $\tilde{X}$ and $\tilde{X}'$ are square integrable random variables with distributions μ and μ', respectively.

3.3. *The Hamiltonian and the dual equations.* The Hamiltonian of the stochastic optimization problem is defined as the function H given by

$$H(t,x,\mu,y,z,\alpha) = b(t,x,\mu,\alpha) \cdot y + \sigma(t,x,\mu,\alpha) \cdot z + f(t,x,\mu,\alpha), \tag{14}$$

where the dot notation stands for the inner product in an Euclidean space. Because we need to compute derivatives of H with respect to its variable μ, we consider the lifting $\tilde{H}$ defined by

$$\tilde{H}(t,x,\tilde{X},y,z,\alpha) = H(t,x,\mu,y,z,\alpha) \tag{15}$$

for any random variable $\tilde{X}$ with distribution μ, and we shall denote by $\partial_\mu H(t,x,\mu_0,y,z,\alpha)$ the derivative with respect to μ computed at μ_0 (as defined above) whenever all the other variables t, x, y, z and α are held fixed. We recall that $\partial_\mu H(t,x,\mu_0,y,z,\alpha)$ is an element of $L^2(\mathbb{R}^d,\mu_0)$ and that we identify it with a function $\partial_\mu H(t,x,\mu_0,y,z,\alpha)(\cdot):\mathbb{R}^d \ni \tilde{x} \mapsto \partial_\mu H(t,x,\mu_0,y,z,\alpha)(\tilde{x})$. It satisfies

$$D\tilde{H}(t,x,\tilde{X},y,z,\alpha) = \partial_\mu H(t,x,\mu_0,y,z,\alpha)(\tilde{X})$$

almost-surely under $\tilde{\mathbb{P}}$.

DEFINITION 3.6. In addition to (A1)–(A2) for b and σ, assume that the coefficients b,σ, f and g are (jointly) differentiable with respect to x and μ. Then, given an admissible control $\alpha = (\alpha_t)_{0\le t\le T} \in \mathbb{A}$, we denote by $X = X^\alpha$ the corresponding controlled state process. Whenever

$$\mathbb{E}\int_0^T \{|\partial_x f(t,X_t,\mathbb{P}_{X_t},\alpha_t)|^2 + \tilde{\mathbb{E}}[|\partial_\mu f(t,X_t,\mathbb{P}_{X_t},\alpha_t)(\tilde{X}_t)|^2]\}\,dt < +\infty \tag{16}$$

and

$$\mathbb{E}\{|\partial_x g(X_T,\mathbb{P}_{X_T})|^2 + \tilde{\mathbb{E}}[|\partial_\mu g(X_T,\mathbb{P}_{X_T})(\tilde{X}_T)|^2]\} < +\infty, \tag{17}$$

we call adjoint processes of X any couple $((Y_t)_{0\le t\le T},(Z_t)_{0\le t\le T})$ of progressively measurable stochastic processes in $\mathbb{H}^{2,d}\times\mathbb{H}^{2,d\times m}$ satisfying the equation (which we call the adjoint equation)

$$\begin{cases} dY_t = -\partial_x H(t,X_t,\mathbb{P}_{X_t},Y_t,Z_t,\alpha_t)\,dt + Z_t\,dW_t \\ \qquad\qquad - \tilde{\mathbb{E}}[\partial_\mu H(t,\tilde{X}_t,\mathbb{P}_{X_t},\tilde{Y}_t,\tilde{Z}_t,\tilde{\alpha}_t)(X_t)]\,dt, \\ Y_T = \partial_x g(X_T,\mathbb{P}_{X_T}) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T,\mathbb{P}_{X_T})(X_T)], \end{cases} \tag{18}$$

where $(\tilde{X},\tilde{Y},\tilde{Z},\tilde{\alpha})$ is an independent copy of (X,Y,Z,α) defined on the space $L^2(\tilde{\Omega},\tilde{\mathcal{F}},\tilde{\mathbb{P}})$, and $\tilde{\mathbb{E}}$ denotes the expectation on $(\tilde{\Omega},\tilde{\mathcal{F}},\tilde{\mathbb{P}})$.

Notice that $\tilde{\mathbb{E}}[\partial_\mu H(t,\tilde{X}_t,\mathbb{P}_{X_t},\tilde{Y}_t,\tilde{Z}_t,\tilde{\alpha}_t)(X_t)]$ is a (measurable) function of the random variable X_t as it stands for $\tilde{\mathbb{E}}[\partial_\mu H(t,\tilde{X}_t,\mathbb{P}_{X_t},\tilde{Y}_t,\tilde{Z}_t,\tilde{\alpha}_t)(x)]|_{x=X_t}$ (and similarly for $\tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T,\mathbb{P}_{X_T})(X_T)]$). Notice that, when b, σ, f and g do not depend upon the marginal distributions of the controlled state process, the extra terms appearing in the adjoint equation and its terminal condition disappear and this equation coincides with the classical adjoint equation of stochastic control.

Using the interpretation of the symbol $\odot$ explained in Section 2, and extending this notation to derivatives of the form $\partial_\mu h(\mu)(x)\odot p = (\partial_\mu[h(\mu)\cdot p])(x)$, the adjoint equation rewrites

$$\begin{aligned} dY_t = {} & -[\partial_x b(t,X_t,\mathbb{P}_{X_t},\alpha_t)\odot Y_t + \partial_x\sigma(t,X_t,\mathbb{P}_{X_t},\alpha_t)\odot Z_t \\ & \qquad + \partial_x f(t,X_t,\mathbb{P}_{X_t},\alpha_t)]\,dt + Z_t\,dW_t \\ & - \tilde{\mathbb{E}}[\partial_\mu b(t,\tilde{X}_t,\mathbb{P}_{X_t},\tilde{\alpha}_t)(X_t)\odot\tilde{Y}_t + \partial_\mu\sigma(t,\tilde{X}_t,\mathbb{P}_{X_t},\tilde{\alpha}_t)(X_t)\odot\tilde{Z}_t \\ & \qquad + \partial_\mu f(t,\tilde{X}_t,\mathbb{P}_{X_t},\tilde{\alpha}_t)(X_t)]\,dt, \end{aligned} \tag{19}$$

with the terminal condition $Y_T = \partial_x g(X_T, \mathbb{P}_{X_T}) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T, \mathbb{P}_{X_T})(X_T)]$. Notice that $\partial_x b$ and $\partial_x \sigma$ are bounded since b and σ are assumed to be c-Lipschitz continuous in the variable x; see (A2). Notice also that the terms $\mathbb{E}[|\partial_\mu b(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t)|^2]^{1/2}$ and $\mathbb{E}[|\partial_\mu \sigma(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t)|^2]^{1/2}$ are also bounded by c since b and σ are assumed to be c-Lipschitz continuous in the variable μ with respect to the 2-Wasserstein distance. It is indeed plain to check that, given a differentiable function $h : \mathcal{P}_2(\mathbb{R}^d) \to \mathbb{R}$, the notion of differentiability being defined as above, it holds $\mathbb{E}[|\partial_\mu h(X)|^2]^{1/2} \le c$, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and any random variable X having μ as distribution, when h is c-Lipschitz continuous in μ with respect to the 2-Wasserstein distance.

Notice finally that, given an admissible control $\alpha \in \mathbb{A}$ and the corresponding controlled state process $X = X^\alpha$, despite conditions (16)–(17) and despite the fact that the first part of the equation appears to be linear in the unknown processes Y_t and Z_t, existence and uniqueness of a solution (Y, Z) of the adjoint equation is not provided by standard results on backward stochastic differential equations (BSDEs) as the distributions of the solution processes (more precisely their joint distributions with the control and state processes α and X) appear in the coefficients of the equation. However, a slight modification of the original existence and uniqueness result of Pardoux and Peng [20] shows that existence and uniqueness still hold in our more general setting. The main lines of the proof are given in [5], Proposition 3.1 and Lemma 3.1. However, Lemma 3.1 in [5] does not apply directly since the coefficients $(\partial_\mu b(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t) \odot \tilde{Y}_t)_{0\le t\le T}$ and $(\partial_\mu \sigma(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t) \odot \tilde{Z}_t)_{0\le t\le T}$ are not Lipschitz continuous in $\tilde{Y}$ and $\tilde{Z}$ uniformly in the randomness; see Condition (C1) in [5]. Actually, a careful inspection of the proof shows that the bounds

$$\mathbb{E}\tilde{\mathbb{E}}\big[\big|\partial_\mu b(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t) \odot \tilde{Y}_t\big|^2\big] \le c' \mathbb{E}\big[|Y_t|^2\big],$$
$$\mathbb{E}\tilde{\mathbb{E}}\big[\big|\partial_\mu \sigma(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t) \odot \tilde{Z}_t\big|^2\big] \le c' \mathbb{E}\big[|Z_t|^2\big],$$

are sufficient to make the whole argument work and thus to prove existence and uniqueness of a solution (Y, Z) satisfying

$$\mathbb{E}\left[\sup_{0\le t\le T} |Y_t|^2 + \int_0^T |Z_t|^2 \, dt\right] < +\infty.$$

4. Pontryagin principle for optimality. In this section, we discuss sufficient and necessary conditions for optimality when the Hamiltonian satisfies appropriate assumptions of convexity. Generally speaking, the role of convexity is twofold. On one hand, convexity permits one to give a quite elegant proof of the necessary part of the Pontryagin principle. The proof of the necessary part is indeed based on a perturbation argument of the optimal controls. In the case when the set A is convex (recalling that controls are A-valued), and the Hamiltonian H is convex in α, the choice of the perturbation is quite simple. In the standard non-McKean–Vlasov case, other kinds of perturbations have been considered, allowing for extensions of the necessary condition to nonconvex cases; see, for instance, the monograph [26]. To keep the paper at a reasonable level of technicality, we refrain from discussing such a relaxation here; see, however, Proposition 4.6 for a slight relaxation. On the other hand, convexity also plays a major role in the proof of the sufficient part. It is indeed mandatory to require the Hamiltonian H to be convex in x, μ and α for establishing the converse of the Pontryagin principle. As explained in Section 5, the typical example is to require b and σ to be linear in x, μ and α, which imposes the same kind of limitation as in the non-McKean–Vlasov case. All these convexity conditions will be specified next, depending upon the framework.

For the time being, we state the regularity assumptions which will be used throughout the section. Referring to Section 3.2 for definitions of joint differentiability, we assume:

(A3) The functions b, σ and f are differentiable with respect to (x,α), the maps $(x,\mu,\alpha) \mapsto \partial_x(b,\sigma,f)(t,x,\mu,\alpha)$ and $(x,\mu,\alpha) \mapsto \partial_\alpha(b,\sigma,f)(t,x,\mu,\alpha)$ being continuous for any $t \in [0,T]$. The functions b, σ and f are also differentiable with respect to the variable μ in the sense given above, the mapping $\mathbb{R}^d \times L^2(\Omega;\mathbb{R}^d) \times A \ni (x,X,\alpha) \mapsto \partial_\mu(b,\sigma,f)(t,x,\mathbb{P}_X,\alpha)(X) \in L^2(\Omega;\mathbb{R}^{d\times d} \times \mathbb{R}^{(d\times m)\times d} \times \mathbb{R}^d)$ being continuous for any $t \in [0,T]$. Similarly, the function g is differentiable with respect to x, the mapping $(x,\mu) \mapsto \partial_x g(x,\mu)$ being continuous. The function g is also differentiable with respect to the variable μ, the mapping $\mathbb{R}^d \times L^2(\Omega;\mathbb{R}^d) \ni (x,X) \mapsto \partial_\mu g(x,\mathbb{P}_X)(X) \in L^2(\Omega;\mathbb{R}^d)$ being continuous.

(A4) The coefficients $((b,\sigma,f)(t,0,\delta_0,0))_{0\le t\le T}$ are uniformly bounded. The derivatives $\partial_x(b,\sigma)$ and $\partial_\alpha(b,\sigma)$ are uniformly bounded and the norm of the mapping $x' \mapsto \partial_\mu(b,\sigma)(t,x,\mu,\alpha)(x')$ in $L^2(\mathbb{R}^d,\mu)$ is also uniformly bounded [i.e., uniformly in (t,x,μ,α)]. There exists a constant L such that, for any $R \ge 0$ and (t,x,μ,α) such that $|x|, \|\mu\|_2, |\alpha| \le R$, $|\partial_x(f,g)(t,x,\mu,\alpha)|$ and $|\partial_\alpha f(t,x,\mu,\alpha)|$ are bounded by $L(1+R)$, and the $L^2(\mathbb{R}^d,\mu)$-norm of $x' \mapsto \partial_\mu(f,g)(t,x,\mu,\alpha)(x')$ is bounded by $L(1+R)$. Here, we have used the notation

$$\|\mu\|_2^2 = \int_{\mathbb{R}^d} |x|^2\, d\mu(x), \qquad \mu \in \mathcal{P}_2(\mathbb{R}^d).$$

Notice that (A3)–(A4) covers (A1)–(A2).

4.1. *A necessary condition.* We assume that the sets A and $\mathbb{A}$ of admissible controls are convex, we fix $\alpha \in \mathbb{A}$, and as before, we denote by $X = X^\alpha$ the corresponding controlled state process, namely the solution of (1) with given initial condition $X_0 = x_0$. Our first task is to compute the Gâteaux derivative of the cost functional J at α in all directions. In order to do so, we choose $\beta \in \mathbb{H}^{2,k}$ such that $\alpha + \varepsilon\beta \in \mathbb{A}$ for $\varepsilon > 0$ small enough. We then compute the variation of J at α in the direction of β (think of β as the difference between another element of $\mathbb{A}$ and α).

Letting $(\theta_t = (X_t, \mathbb{P}_{X_t}, \alpha_t))_{0\le t\le T}$, we define the variation process $V = (V_t)_{0\le t\le T}$ to be the solution of the equation

$$dV_t = [\gamma_t \cdot V_t + \rho_t(\mathbb{P}_{(X_t,V_t)}) + \eta_t]\,dt + [\tilde{\gamma}_t \cdot V_t + \tilde{\rho}_t(\mathbb{P}_{(X_t,V_t)}) + \tilde{\eta}_t]\,dW_t, \tag{20}$$

with $V_0 = 0$, where the coefficients γ_t, δ_t, ρ_t, $\tilde{\gamma}_t$, $\tilde{\rho}_t$ and $\tilde{\eta}_t$ are defined as

$$\gamma_t = \partial_x b(t, \theta_t), \qquad \tilde{\gamma}_t = \partial_x \sigma(t, \theta_t),$$
$$\eta_t = \partial_\alpha b(t, \theta_t) \cdot \beta_t, \qquad \tilde{\eta}_t = \partial_\alpha \sigma(t, \theta_t) \cdot \beta_t,$$

which are progressively measurable bounded processes with values in $\mathbb{R}^{d\times d}$, $\mathbb{R}^{(d\times m)\times d}$, $\mathbb{R}^d$ and $\mathbb{R}^{d\times m}$, respectively (the parentheses around $d \times m$ indicating that $\tilde{\gamma}_t \cdot u$ is seen as an element of $\mathbb{R}^{d\times m}$ whenever $u \in \mathbb{R}^d$), and

$$\begin{aligned}\rho_t &= \tilde{\mathbb{E}}[\partial_\mu b(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t] = \tilde{\mathbb{E}}[\partial_\mu b(t, x, \mathbb{P}_{X_t}, \alpha)(\tilde{X}_t) \cdot \tilde{V}_t]\big|_{\substack{x=X_t\\ \alpha=\alpha_t}},\\ \tilde{\rho}_t &= \tilde{\mathbb{E}}[\partial_\mu \sigma(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t] = \tilde{\mathbb{E}}[\partial_\mu \sigma(t, x, \mathbb{P}_{X_t}, \alpha)(\tilde{X}_t) \cdot \tilde{V}_t]\big|_{\substack{x=X_t\\ \alpha=\alpha_t}},\end{aligned} \tag{21}$$

which are progressively measurable bounded processes with values in $\mathbb{R}^d$ and $\mathbb{R}^{d\times m}$, respectively, and where $(\tilde{X}_t, \tilde{V}_t)$ is an independent copy of (X_t, V_t). As expectations of functions of $(\tilde{X}_t, \tilde{V}_t)$, ρ_t and $\tilde{\rho}_t$ depend upon the joint distribution of X_t and V_t. In (20) we wrote $\rho_t(\mathbb{P}_{(X_t,V_t)})$ and $\tilde{\rho}_t(\mathbb{P}_{(X_t,V_t)})$ in order to stress the dependence upon the joint distribution of X_t and V_t. Even though we are dealing with possibly random coefficients, the existence and uniqueness of the variation process is guaranteed by Proposition 2.1 of [13] applied to the couple (X, V) and the system formed by (1) and (20). Because of our assumption on the boundedness of the partial derivatives of the coefficients, V satisfies $\mathbb{E}\sup_{0\le t\le T} |V_t|^p < \infty$ for every finite $p \ge 1$.

LEMMA 4.1. *For each $\varepsilon > 0$ small enough, we denote by α^ε the admissible control defined by $\alpha^\varepsilon_t = \alpha_t + \varepsilon\beta_t$, and by $X^\varepsilon = X^{\alpha^\varepsilon}$ the corresponding controlled state. We have*

$$\lim_{\varepsilon\searrow 0} \mathbb{E}\Big[\sup_{0\le t\le T}\Big|\frac{X^\varepsilon_t - X_t}{\varepsilon} - V_t\Big|^2\Big] = 0. \tag{22}$$

PROOF. For the purpose of this proof we set $\theta^\varepsilon_t = (X^\varepsilon_t, \mathbb{P}_{X^\varepsilon_t}, \alpha^\varepsilon_t)$ and $V^\varepsilon_t = \varepsilon^{-1}(X^\varepsilon_t - X_t) - V_t$. Notice that $V^\varepsilon_0 = 0$ and that

$$\begin{aligned}dV^\varepsilon_t = \Big[&\frac{1}{\varepsilon}[b(t, \theta^\varepsilon_t) - b(t, \theta_t)] - \partial_x b(t, \theta_t) \cdot V_t - \partial_\alpha b(t, \theta_t) \cdot \beta_t\\ &\qquad - \tilde{\mathbb{E}}[\partial_\mu b(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t]\Big]\,dt\\ &+ \Big[\frac{1}{\varepsilon}[\sigma(t, \theta^\varepsilon_t) - \sigma(t, \theta_t)] - \partial_x \sigma(t, \theta_t) \cdot V_t - \partial_\alpha \sigma(t, \theta_t) \cdot \beta_t\end{aligned} \tag{23}$$

$$
- \tilde{\mathbb{E}}\big[\partial_\mu \sigma(t,\theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big]\Big] dW_t
$$

$$
= V_t^{\varepsilon,1}\,dt + V_t^{\varepsilon,2}\,dW_t.
$$

Now for each $t \in [0,T]$ and each $\varepsilon > 0$, we have

$$
\begin{aligned}
\frac{1}{\varepsilon}\big[b(t,\theta_t^\varepsilon) - b(t,\theta_t)\big] &= \int_0^1 \partial_x b\big(t,\theta_t^{\lambda,\varepsilon}\big) \cdot \big(V_t^\varepsilon + V_t\big)\,d\lambda + \int_0^1 \partial_\alpha b\big(t,\theta_t^{\lambda,\varepsilon}\big) \cdot \beta_t\,d\lambda \\
&\quad + \int_0^1 \tilde{\mathbb{E}}\big[\partial_\mu b\big(t,\theta_t^{\lambda,\varepsilon}\big)\big(\tilde{X}_t^{\lambda,\varepsilon}\big) \cdot \big(\tilde{V}_t^\varepsilon + \tilde{V}_t\big)\big]\,d\lambda,
\end{aligned}
$$

where, in order to simplify a little bit the notation, we have set $X_t^{\lambda,\varepsilon} = X_t + \lambda\varepsilon(V_t^\varepsilon + V_t)$, $\alpha_t^{\lambda,\varepsilon} = \alpha_t + \lambda\varepsilon\beta_t$ and $\theta_t^{\lambda,\varepsilon} = (X_t^{\lambda,\varepsilon}, \mathbb{P}_{X_t^{\lambda,\varepsilon}}, \alpha_t^{\lambda,\varepsilon})$. Computing the "$dt$"-term, we get

$$
\begin{aligned}
V_t^{\varepsilon,1} &= \int_0^1 \partial_x b\big(t,\theta_t^{\lambda,\varepsilon}\big) \cdot V_t^\varepsilon\,d\lambda + \int_0^1 \tilde{\mathbb{E}}\big[\partial_\mu b\big(t,\theta_t^{\lambda,\varepsilon}\big)\big(\tilde{X}_t^{\lambda,\varepsilon}\big) \cdot \tilde{V}_t^\varepsilon\big]\,d\lambda \\
&\quad + \int_0^1 \big[\partial_x b\big(t,\theta_t^{\lambda,\varepsilon}\big) - \partial_x b(t,\theta_t)\big] \cdot V_t\,d\lambda \\
&\quad + \int_0^1 \big[\partial_\alpha b\big(t,\theta_t^{\lambda,\varepsilon}\big) - \partial_\alpha b(t,\theta_t)\big] \cdot \beta_t\,d\lambda \\
&\quad + \int_0^1 \tilde{\mathbb{E}}\big[\big(\partial_\mu b\big(t,\theta_t^{\lambda,\varepsilon}\big)\big(\tilde{X}_t^{\lambda,\varepsilon}\big) - \partial_\mu b(t,\theta_t)(\tilde{X}_t)\big) \cdot \tilde{V}_t\big]\,d\lambda \\
&= \int_0^1 \partial_x b\big(t,\theta^{\lambda,\varepsilon}\big) \cdot V_t^\varepsilon\,d\lambda + \int_0^1 \tilde{\mathbb{E}}\big[\partial_\mu b\big(t,\theta_t^{\lambda,\varepsilon}\big)\big(\tilde{X}_t^{\lambda,\varepsilon}\big) \cdot \tilde{V}_t^\varepsilon\big]\,d\lambda \\
&\quad + I_t^{\varepsilon,1} + I_t^{\varepsilon,2} + I_t^{\varepsilon,3}.
\end{aligned}
$$

By (A4), the three last terms of the above right-hand side are bounded in $L^2([0,T] \times \Omega)$, uniformly in ε. Next, we treat the diffusion part $V_t^{\varepsilon,2}$ in the same way using Jensen's inequality and the Burkholder–Davis–Gundy inequality to control the quadratic variation of the stochastic integrals. Consequently, going back to (23), we see that, for any $S \in [0,T]$,

$$
\mathbb{E}\Big[\sup_{0 \le t \le S} |V_t^\varepsilon|^2\Big] \le c' + c' \int_0^S \mathbb{E}\Big[\sup_{0 \le s \le t} |V_s^\varepsilon|^2\Big]\,dt,
$$

where as usual $c' > 0$ is a generic constant which can change from line to line. Applying Gronwall's inequality, we deduce that $\mathbb{E}[\sup_{0 \le t \le T} |V_t^\varepsilon|^2] \le c'$. Therefore, we have

$$
\lim_{\varepsilon \searrow 0} \mathbb{E}\Big[\sup_{0 \le \lambda \le 1} \sup_{0 \le t \le T} \big|X_t^{\lambda,\varepsilon} - X_t\big|^2\Big] = 0.
$$

We then prove that $I^{\varepsilon,1}$, $I^{\varepsilon,2}$ and $I^{\varepsilon,3}$ converge to 0 in $L^2([0,T]\times\Omega)$ as $\varepsilon \searrow 0$. Indeed,

$$\begin{aligned}\mathbb{E}\int_0^T |I_t^{\varepsilon,1}|^2\,dt &= \mathbb{E}\int_0^T \left|\int_0^1 \big(\big[\partial_x b\big(t,\theta_t^{\lambda,\varepsilon}\big) - \partial_x b(t,\theta_t)\big]\cdot V_t\big)\,d\lambda\right|^2 dt \\ &\le \mathbb{E}\int_0^T \int_0^1 \big|\partial_x b\big(t,\theta_t^{\lambda,\varepsilon}\big) - \partial_x b(t,\theta_t)\big|^2 |V_t|^2\,d\lambda\,dt.\end{aligned}$$

Since the function $\partial_x b$ is bounded and continuous in x, μ and α, the above right-hand side converges to 0 as $\varepsilon \searrow 0$. A similar argument applies to $I_t^{\varepsilon,2}$ and $I_t^{\varepsilon,3}$. Again, we treat the diffusion part $V_t^{\varepsilon,2}$ in the same way using Jensen's inequality and the Burkholder–Davis–Gundy inequality. Consequently, going back to (23), we finally see that, for any $S \in [0,T]$,

$$\mathbb{E}\Big[\sup_{0\le t\le S}|V_t^{\varepsilon}|^2\Big] \le c' \int_0^S \mathbb{E}\Big[\sup_{0\le s\le t}|V_s^{\varepsilon}|^2\Big]\,dt + c_\varepsilon,$$

where $\lim_{\varepsilon\searrow 0} c_\varepsilon = 0$. Finally, we get the desired result applying Gronwall's inequality. $\square$

We now compute the Gâteaux derivative of the objective function.

LEMMA 4.2. *The function $\alpha \mapsto J(\alpha)$ is Gâteaux differentiable in the direction β, and its derivative is given by*

$$\begin{aligned}&\frac{d}{d\varepsilon} J(\alpha+\varepsilon\beta)\Big|_{\varepsilon=0} \\ &\quad = \mathbb{E}\int_0^T \big[\partial_x f(t,\theta_t)\cdot V_t + \tilde{\mathbb{E}}\big[\partial_\mu f(t,\theta_t)(\tilde X_t)\cdot \tilde V_t\big] + \partial_\alpha f(t,\theta_t)\cdot\beta_t\big]\,dt \\ &\qquad + \mathbb{E}\big[\partial_x g(X_T,\mathbb{P}_{X_T})\cdot V_T + \tilde{\mathbb{E}}\big[\partial_\mu g(X_T,\mathbb{P}_{X_T})(\tilde X_T)\cdot \tilde V_T\big]\big].\end{aligned} \tag{24}$$

PROOF. We use freely the notation introduced in the proof of the previous lemma:

$$\begin{aligned}&\frac{d}{d\varepsilon} J(\alpha+\varepsilon\beta)\Big|_{\varepsilon=0} \\ &\quad = \lim_{\varepsilon\searrow 0}\frac{1}{\varepsilon}\mathbb{E}\Big[\int_0^T \big[f\big(t,\theta_t^{\varepsilon}\big) - f(t,\theta_t)\big]\,dt + \big[g\big(X_T^{\varepsilon},\mathbb{P}_{X_T^{\varepsilon}}\big) - g(X_T,\mathbb{P}_{X_T})\big]\Big].\end{aligned} \tag{25}$$

Computing the two limits separately, we get

$$\begin{aligned}&\lim_{\varepsilon\searrow 0}\frac{1}{\varepsilon}\mathbb{E}\int_0^T \big[f\big(t,\theta_t^{\varepsilon}\big) - f(t,\theta_t)\big]\,dt \\ &\quad = \lim_{\varepsilon\searrow 0}\frac{1}{\varepsilon}\mathbb{E}\int_0^T \int_0^1 \frac{d}{d\lambda} f\big(t,\theta_t^{\lambda,\varepsilon}\big)\,d\lambda\,dt\end{aligned}$$

$$
\begin{aligned}
&= \lim_{\varepsilon \searrow 0} \mathbb{E} \int_0^T \int_0^1 \big[\partial_x f\big(t, \theta_t^{\lambda,\varepsilon}\big) \cdot \big(V_t^\varepsilon + V_t\big) \\
&\qquad\qquad + \tilde{\mathbb{E}}\big[\partial_\mu f\big(t, \theta_t^{\lambda,\varepsilon}\big)\big(\tilde{X}_t^{\lambda,\varepsilon}\big) \cdot \big(\tilde{V}_t^\varepsilon + \tilde{V}_t\big)\big] \\
&\qquad\qquad\qquad + \partial_\alpha f\big(t, \theta_t^{\lambda,\varepsilon}\big) \cdot \beta_t\big]\, d\lambda\, dt \\
&= \mathbb{E} \int_0^T \big[\partial_x f(t, \theta_t) \cdot V_t + \tilde{\mathbb{E}}\big[\partial_\mu f(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big] + \partial_\alpha f(t, \theta_t) \cdot \beta_t\big]\, dt,
\end{aligned}
$$

using the hypothesis on the continuity and growth of the derivatives of f, the uniform convergence proven in the previous lemma and standard uniform integrability arguments. The second term in (25) is tackled in a similar way. □

Observing that conditions (16)–(17) are satisfied under (A3)–(A4), the duality relationship is given by:

LEMMA 4.3. *Given* $(Y_t, Z_t)_{0 \le t \le T}$ *as in Definition* 3.6, *it holds*

$$
\begin{aligned}
\mathbb{E}[Y_T \cdot V_T] = \mathbb{E} \int_0^T & \big[Y_t \cdot \big(\partial_\alpha b(t, \theta_t) \cdot \beta_t\big) + Z_t \cdot \big(\partial_\alpha \sigma(t, \theta_t) \cdot \beta_t\big) \\
& - \partial_x f(t, \theta_t) \cdot V_t - \tilde{\mathbb{E}}\big[\partial_\mu f(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big]\big]\, dt.
\end{aligned}
\tag{26}
$$

PROOF. Letting $\Theta_t = (X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t)$ and using the definitions (20) of the variation process V, and (18) or (19) of the adjoint process Y, integration by parts gives

$$
\begin{aligned}
Y_T \cdot V_T &= Y_0 \cdot V_0 + \int_0^T Y_t \cdot dV_t + \int_0^T dY_t \cdot V_t + \int_0^T d[Y, V]_t \\
&= M_T + \int_0^T \big[Y_t \cdot \big(\partial_x b(t, \theta_t) \cdot V_t\big) + Y_t \cdot \tilde{\mathbb{E}}\big[\partial_\mu b(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big] \\
&\qquad + Y_t \cdot \big(\partial_\alpha b(t, \theta_t) \cdot \beta_t\big) \\
&\qquad - \partial_x H(t, \Theta_t) \cdot V_t - \tilde{\mathbb{E}}\big[\partial_\mu H(t, \tilde{\Theta}_t)(X_t) \cdot V_t\big] \\
&\qquad + Z_t \cdot \big(\partial_x \sigma(t, \theta_t) \cdot V_t\big) + Z_t \cdot \tilde{\mathbb{E}}\big[\partial_\mu \sigma(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big] \\
&\qquad\qquad + Z_t \cdot \big(\partial_\alpha \sigma(t, \theta_t) \cdot \beta_t\big)\big]\, dt,
\end{aligned}
$$

where $(M_t)_{0 \le t \le T}$ is a mean zero integrable martingale. By taking expectations on both sides and applying Fubini's theorem,

$$
\begin{aligned}
&\mathbb{E}\tilde{\mathbb{E}}\big[\partial_\mu H(t, \tilde{\Theta}_t)(X_t) \cdot V_t\big] \\
&\quad = \mathbb{E}\tilde{\mathbb{E}}\big[\partial_\mu H(t, \Theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big] \\
&\quad = \mathbb{E}\tilde{\mathbb{E}}\big[\big(\partial_\mu b(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big) \cdot Y_t + \big(\partial_\mu \sigma(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big) \cdot Z_t \\
&\qquad\qquad + \partial_\mu f(t, \theta_t)(\tilde{X}_t) \cdot \tilde{V}_t\big].
\end{aligned}
$$

By commutativity of the inner product, cancellations occur, and we get the desired equality (26). □

By putting together the duality relation (26) and (24) we get:

COROLLARY 4.4. *The Gâteaux derivative of J at α in the direction of β can be written as*

$$\frac{d}{d\varepsilon} J(\alpha + \varepsilon\beta)\Big|_{\varepsilon=0} = \mathbb{E}\int_0^T \big[\partial_\alpha H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) \cdot \beta_t\big]\, dt. \tag{27}$$

PROOF. Using Fubini's theorem, the second expectation appearing in the expression (24) of the Gâteaux derivative of J given in Lemma 4.2 can be rewritten as

$$\begin{aligned}
&\mathbb{E}\big[\partial_x g(X_T, \mathbb{P}_{X_T}) \cdot V_T + \tilde{\mathbb{E}}\big(\partial_\mu g(X_T, \mathbb{P}_{X_T})(\tilde{X}_T) \cdot \tilde{V}_T\big)\big] \\
&\quad = \mathbb{E}\big[\partial_x g(X_T, \mathbb{P}_{X_T}) \cdot V_T\big] + \mathbb{E}\tilde{\mathbb{E}}\big[\partial_\mu g(\tilde{X}_T, \mathbb{P}_{X_T})(X_T) \cdot V_T\big] \\
&\quad = \mathbb{E}[Y_T \cdot V_T],
\end{aligned}$$

and using the expression derived in Lemma 4.3 for $\mathbb{E}[Y_T \cdot V_T]$ in (24) gives the desired result. □

The main result of this subsection is the following theorem.

THEOREM 4.5. *Under the above assumptions, if we assume further that the Hamiltonian H is convex in α, the admissible control $(\alpha_t)_{0\le t\le T} \in \mathbb{A}$ is optimal, $(X_t)_{0\le t\le T}$ is the associated (optimally) controlled state, and $(Y_t, Z_t)_{0\le t\le T}$ are the associated adjoint processes solving the adjoint equation (18), then we have*

$$\forall \alpha \in A \qquad H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) \le H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha), \qquad dt \otimes d\mathbb{P}\text{-a.e.} \tag{28}$$

PROOF. Since A is convex, given $\beta \in \mathbb{A}$, we can choose the perturbation $\alpha_t^\varepsilon = \alpha_t + \varepsilon(\beta_t - \alpha_t)$ which is still in $\mathbb{A}$ for $0 \le \varepsilon \le 1$. Since α is optimal, we have the inequality

$$\frac{d}{d\varepsilon} J\big(\alpha + \varepsilon(\beta - \alpha)\big)\Big|_{\varepsilon=0} = \mathbb{E}\int_0^T \big[\partial_\alpha H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) \cdot (\beta_t - \alpha_t)\big]\, dt \ge 0.$$

By convexity of the Hamiltonian with respect to the control variable $\alpha \in A$, we conclude that

$$\mathbb{E}\int_0^T \big[H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \beta_t) - H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t)\big]\, dt \ge 0,$$

for all β. Now, if for a given (deterministic) $\alpha \in A$ we choose β in the following way:

$$\beta_t(\omega) = \begin{cases} \alpha, & \text{if } (t,\omega) \in C, \\ \alpha_t(\omega), & \text{otherwise,} \end{cases}$$

for an arbitrary progressively-measurable set $C \subset [0,T] \times \Omega$ (i.e., $C \cap [0,t] \in \mathcal{B}([0,t]) \otimes \mathcal{F}_t$ for any $t \in [0,T]$), we see that

$$\mathbb{E}\int_0^T \mathbf{1}_C \big[H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha) - H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t)\big]\,dt \geq 0,$$

from which we conclude that

$$H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha) - H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) \geq 0, \qquad dt \otimes d\mathbb{P}\text{-a.e.,}$$

which is the desired conclusion. $\square$

When convexity of the set A fails, the following weaker version of the stochastic Pontryagin principle holds:

PROPOSITION 4.6. *Keep the same assumptions as above, but do not require A to be convex and H to be convex in α. Assume that the admissible control $(\alpha_t)_{0\leq t\leq T} \in \mathbb{A}$ is optimal, $(X_t)_{0\leq t\leq T}$ is the associated (optimally) controlled state and $(Y_t, Z_t)_{0\leq t\leq T}$ are the associated adjoint processes solving the adjoint equation (18). Then we have*

$$\partial_\alpha H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) = 0, \qquad dt \otimes d\mathbb{P}\text{-}a.e.$$

PROOF. Given $\varepsilon_0 > 0$, $\beta \in \mathbb{R}^k$ with $|\beta| = 1$, and a progressively-measurable set $C \subset [0,T] \times \Omega$, we let

$$\beta_t = \beta \mathbf{1}_{C \cap \{\mathrm{dist}(\alpha_t, A^{\complement}) > \varepsilon_0\}},$$

for $t \in [0,T]$. By construction, $\alpha_t + \varepsilon\beta_t \in A$ for all $t \in [0,T]$ and $\varepsilon \in (0, \varepsilon_0)$. Following the proof of Theorem 4.5, we claim

$$\mathbb{E}\int_0^T \big[\partial_\alpha H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) \cdot \beta_t\big]\,dt \geq 0,$$

from which we deduce that

$$\mathbf{1}_{\{\mathrm{dist}(\alpha_t, A^{\complement}) > \varepsilon_0\}} \partial_\alpha H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) \cdot \beta \geq 0, \qquad dt \otimes d\mathbb{P}\text{-a.e.}$$

As β and ε_0 are arbitrary, we finally get

$$\mathbf{1}_{\{\mathrm{dist}(\alpha_t, A^{\complement}) > 0\}} \partial_\alpha H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) = 0, \qquad dt \otimes d\mathbb{P}\text{-a.e.}$$

Recalling that A is open, the result follows. $\square$

4.2. *A sufficient condition.* The necessary condition for optimality identified in the previous subsection can be turned into a sufficient condition for optimality under some technical assumptions.

THEOREM 4.7. *Under the same assumptions of regularity on the coefficients as before, let $\alpha \in \mathbb{A}$ be an admissible control, $(X_t = X_t^\alpha)_{0\le t\le T}$ the corresponding controlled state process and $(Y_t, Z_t)_{0\le t\le T}$ the corresponding adjoint processes. Let us also assume that*:

(1) $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \ni (x,\mu) \mapsto g(x,\mu)$ *is convex*;

(2) $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times A \ni (x,\mu,\alpha) \mapsto H(t,x,\mu,Y_t,Z_t,\alpha)$ *is convex* $dt \otimes d\mathbb{P}$ *almost everywhere.*

If

$$H(t,X_t,\mathbb{P}_{X_t},Y_t,Z_t,\alpha_t) = \inf_{\alpha\in A} H(t,X_t,\mathbb{P}_{X_t},Y_t,Z_t,\alpha), \qquad dt\otimes d\mathbb{P}\text{-a.e.} \tag{29}$$

then α is an optimal control, that is, $J(\alpha) = \inf_{\alpha'\in\mathbb{A}} J(\alpha')$.

PROOF. Let $\alpha' \in \mathbb{A}$ be a generic admissible control, and $X' = X^{\alpha'}$ the corresponding controlled state. By definition of the objective function of the control problem, we have

$$\begin{aligned}
&J(\alpha) - J(\alpha') \\
&\quad = \mathbb{E}\big[g(X_T,\mathbb{P}_{X_T}) - g\big(X'_T,\mathbb{P}_{X'_T}\big)\big] + \mathbb{E}\int_0^T \big[f(t,\theta_t) - f\big(t,\theta'_t\big)\big]\,dt \\
&\quad = \mathbb{E}\big[g(X_T,\mathbb{P}_{X_T}) - g\big(X'_T,\mathbb{P}_{X'_T}\big)\big] + \mathbb{E}\int_0^T \big[H(t,\Theta_t) - H\big(t,\Theta'_t\big)\big]\,dt \\
&\qquad - \mathbb{E}\int_0^T \big\{\big[b(t,\theta_t) - b\big(t,\theta'_t\big)\big]\cdot Y_t + \big[\sigma(t,\theta_t) - \sigma\big(t,\theta'_t\big)\big]\cdot Z_t\big\}\,dt
\end{aligned} \tag{30}$$

by definition of the Hamiltonian, with the notation $\theta_t = (X_t,\mathbb{P}_{X_t},\alpha_t)$ and $\Theta_t = (X_t,\mathbb{P}_{X_t},Y_t,Z_t,\alpha_t)$ (and similarly for θ'_t and Θ'_t). The function g being convex, we have

$$g(x,\mu) - g\big(x',\mu'\big) \le \partial_x g(x,\mu)\cdot\big(x-x'\big) + \tilde{\mathbb{E}}\big[\partial_\mu g(x,\mu)(\tilde X)\cdot\big(\tilde X - \tilde X'\big)\big],$$

so that

$$\begin{aligned}
&\mathbb{E}\big[g(X_T,\mathbb{P}_{X_T}) - g\big(X'_T,\mathbb{P}_{X'_T}\big)\big] \\
&\quad \le \mathbb{E}\big[\partial_x g(X_T,\mathbb{P}_{X_T})\cdot\big(X_T - X'_T\big) + \tilde{\mathbb{E}}\big[\partial_\mu g(X_T,\mathbb{P}_{X_T})(\tilde X_T)\cdot\big(\tilde X_T - \tilde{X}'_T\big)\big]\big] \\
&\quad = \mathbb{E}\big[\big(\partial_x g(X_T,\mathbb{P}_{X_T}) + \tilde{\mathbb{E}}\big[\partial_\mu g(\tilde X_T,\mathbb{P}_{X_T})(X_T)\big]\big)\cdot\big(X_T - X'_T\big)\big] \\
&\quad = \mathbb{E}\big[Y_T\cdot\big(X_T - X'_T\big)\big] = \mathbb{E}\big[\big(X_T - X'_T\big)\cdot Y_T\big],
\end{aligned} \tag{31}$$

where we used Fubini's theorem and the fact that the "tilde random variables" are independent copies of the "nontilde variables." Using the adjoint equation and taking the expectation, we get

$$\begin{aligned}
&\mathbb{E}\big[(X_T - X'_T)\cdot Y_T\big] \\
&\quad = \mathbb{E}\biggl[\int_0^T (X_t - X'_t)\cdot dY_t + \int_0^T Y_t \cdot d[X_t - X'_t] \\
&\qquad\qquad + \int_0^T \big[\sigma(t,\theta_t) - \sigma(t,\theta'_t)\big]\cdot Z_t\,dt\biggr] \\
&\quad = -\mathbb{E}\int_0^T \big[\partial_x H(t,\Theta_t)\cdot(X_t - X'_t) + \tilde{\mathbb{E}}\big[\partial_\mu H(t,\tilde{\Theta}_t)(X_t)\big]\cdot(X_t - X'_t)\big]\,dt \\
&\qquad + \mathbb{E}\int_0^T \big[\big[b(t,\theta_t) - b(t,\theta'_t)\big]\cdot Y_t + \big[\sigma(t,\theta_t) - \sigma(t,\theta'_t)\big]\cdot Z_t\big]\,dt,
\end{aligned}$$

where we used integration by parts and the fact that Y_t solves the adjoint equation. Using Fubini's theorem and the fact that $\tilde{\Theta}_t$ is an independent copy of Θ_t, the expectation of the second term in the second line can be rewritten as

$$\begin{aligned}
&\mathbb{E}\int_0^T \big\{\tilde{\mathbb{E}}\big[\partial_\mu H(t,\tilde{\Theta}_t)(X_t)\big]\cdot(X_t - X'_t)\big\}\,dt \\
&\quad = \mathbb{E}\tilde{\mathbb{E}}\int_0^T \big\{\big[\partial_\mu H(t,\Theta_t)(\tilde{X}_t)\big]\cdot(\tilde{X}_t - \tilde{X}'_t)\big\}\,dt \qquad (32)\\
&\quad = \mathbb{E}\int_0^T \tilde{\mathbb{E}}\big[\partial_\mu H(t,\Theta_t)(\tilde{X}_t)\cdot(\tilde{X}_t - \tilde{X}'_t)\big]\,dt.
\end{aligned}$$

Consequently, by (30), (31) and (32), we obtain

$$\begin{aligned}
&J(\alpha) - J(\alpha') \\
&\quad \le \mathbb{E}\int_0^T \big[H(t,\Theta_t) - H(t,\Theta'_t)\big]\,dt \\
&\qquad - \mathbb{E}\int_0^T \big\{\partial_x H(t,\Theta_t)\cdot(X_t - X'_t) + \tilde{\mathbb{E}}\big[\partial_\mu H(t,\tilde{\Theta}_t)(X_t)\cdot(\tilde{X}_t - \tilde{X}'_t)\big]\big\}\,dt \qquad (33)\\
&\quad \le 0,
\end{aligned}$$

because of the convexity assumption on H [see, in particular, (13)], and because of the criticality of the admissible control $(\alpha_t)_{0\le t\le T}$, see (29), which says the first order derivative in α vanishes. $\square$

4.3. *Special cases*. We consider a set of particular cases which already appeared in the literature, and we provide the special forms of the stochastic Pontryagin principle which apply in these cases. We discuss only sufficient conditions for optimality for the sake of definiteness. The corresponding necessary conditions can easily be derived from the results of Section 4.1.

Scalar interactions. In this subsection we show how the model handled in [2] appears as a specific example of our more general formulation. We consider scalar interactions for which the dependence upon the probability measure of the coefficients is through functions of scalar moments of the measure. More specifically, we assume that

$$b(t,x,\mu,\alpha) = \hat{b}(t,x,\langle\psi,\mu\rangle,\alpha), \qquad \sigma(t,x,\mu,\alpha) = \hat{\sigma}(t,x,\langle\phi,\mu\rangle,\alpha),$$
$$f(t,x,\mu,\alpha) = \hat{f}(t,x,\langle\gamma,\mu\rangle,\alpha), \qquad g(x,\mu) = \hat{g}(x,\langle\zeta,\mu\rangle)$$

for some scalar functions ψ, ϕ, γ and ζ with at most quadratic growth at ∞, and functions $\hat{b}$, $\hat{\sigma}$ and $\hat{f}$ defined on $[0,T]\times\mathbb{R}^d\times\mathbb{R}\times A$ with values in $\mathbb{R}^d$, $\mathbb{R}^{d\times m}$ and $\mathbb{R}$, respectively, and a real valued function $\hat{g}$ defined on $\mathbb{R}^d\times\mathbb{R}$. We use the bracket notation $\langle h,\mu\rangle$ to denote the integral of the function h with respect to the measure μ. The functions $\hat{b}$, $\hat{\sigma}$, $\hat{f}$ and $\hat{g}$ are similar to the functions b, σ, f and g with the variable μ, which was a measure, replaced by a numeric variable, say r. Reserving the notation H for the Hamiltonian we defined above, we have

$$\begin{aligned}&H(t,x,\mu,y,z,\alpha)\\&\quad= \hat{b}(t,x,\langle\psi,\mu\rangle,\alpha)\cdot y + \hat{\sigma}(t,x,\langle\phi,\mu\rangle,\alpha)\cdot z + \hat{f}(t,x,\langle\gamma,\mu\rangle,\alpha).\end{aligned}$$

We then proceed to derive the particular form taken by the adjoint equation in the present situation. We start with the terminal condition as it is easier to identify. According to (18), it reads

$$Y_T = \partial_x g(X_T,\mathbb{P}_{X_T}) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T,\mathbb{P}_{\tilde{X}_T})(X_T)].$$

Since the terminal cost is of the form $g(x,\mu) = \hat{g}(x,\langle\zeta,\mu\rangle)$, given our definition of differentiability with respect to the variable μ, we know, as a generalization of (6), that $\partial_\mu g(x,\mu)(\cdot)$ reads

$$\partial_\mu g(x,\mu)(x') = \partial_r \hat{g}(x,\langle\zeta,\mu\rangle)\partial\zeta(x'), \qquad x'\in\mathbb{R}^d.$$

Therefore, the terminal condition Y_T can be rewritten as

$$Y_T = \partial_x \hat{g}(X_T,\mathbb{E}[\zeta(X_T)]) + \tilde{\mathbb{E}}[\partial_r \hat{g}(\tilde{X}_T,\mathbb{E}[\zeta(X_T)])]\partial\zeta(X_T),$$

which is exactly the terminal condition used in [2] once we remark that the "tildes" can be removed since $\tilde{X}_T$ has the same distribution as X_T. Within this framework, convexity in μ is quite easy to check. Here is a typical example borrowed from [2]: if g and $\hat{g}$ do not depend on x, then the function $\mathcal{P}_2(\mathbb{R}^d)\ni\mu\mapsto g(\mu) = \hat{g}(\langle\zeta,\mu\rangle)$ is convex if ζ is convex and $\hat{g}$ is nondecreasing and convex.

Similarly, $\partial_\mu H(t,x,\mu,y,z,\alpha)$ can be identified to the $\mathbb{R}^d$-valued function defined by

$$\begin{aligned}&\partial_\mu H(t,x,\mu,y,z,\alpha)(x')\\&\quad= [\partial_r\hat{b}(t,x,\langle\psi,\mu\rangle,\alpha)\odot y]\partial\psi(x') + [\partial_r\hat{\sigma}(t,x,\langle\phi,\mu\rangle,\alpha)\odot z]\partial\phi(x')\\&\qquad + \partial_r\hat{f}(t,x,\langle\gamma,\mu\rangle,\alpha)\partial\gamma(x'),\end{aligned}$$

and the dynamic part of the adjoint equation (18) rewrites

$$
\begin{aligned}
dY_t = &-\big\{\partial_x \hat{b}\big(t, X_t, \mathbb{E}[\psi(X_t)], \alpha_t\big) \odot Y_t + \partial_x \hat{\sigma}\big(t, X_t, \mathbb{E}[\phi(X_t)], \alpha_t\big) \odot Z_t \\
&\qquad + \partial_x \hat{f}\big(t, X_t, \mathbb{E}[\gamma(X_t)], \alpha_t\big)\big\}\,dt \\
&+ Z_t\,dW_t \\
&- \big\{\tilde{\mathbb{E}}\big[\partial_r \hat{b}\big(t, \tilde{X}_t, \mathbb{E}[\psi(X_t)], \tilde{\alpha}_t\big) \odot \tilde{Y}_t\big]\partial\psi(X_t) \\
&\quad + \tilde{\mathbb{E}}\big[\partial_r \hat{\sigma}\big(t, \tilde{X}_t, \mathbb{E}[\phi(X_t)], \tilde{\alpha}_t\big) \odot \tilde{Z}_t\big]\partial\phi(X_t) \\
&\quad + \tilde{\mathbb{E}}\big[\partial_r \hat{f}\big(t, \tilde{X}_t, \mathbb{E}[\gamma(X_t)], \tilde{\alpha}_t\big)\big]\partial\gamma(X_t)\big\}\,dt,
\end{aligned}
$$

which, again, is exactly the adjoint equation used in [2] once we *remove the "tildes."*

REMARK 4.8. The mean–variance portfolio optimization example discussed in [2] and the solution proposed in [3] and [10] of the optimal control of linear-quadratic (LQ) McKean–Vlasov dynamics are based on the general form of the Pontryagin principle proven in this section as applied to the scalar interactions considered in this subsection.

First order interactions. In the case of first order interactions, the dependence upon the probability measure is linear in the sense that the coefficients b, σ, f and g are given in the form

$$
\begin{aligned}
b(t,x,\mu,\alpha) &= \langle \hat{b}(t,x,\cdot,\alpha), \mu\rangle, &\qquad \sigma(t,x,\mu,\alpha) &= \langle \hat{\sigma}(t,x,\cdot,\alpha), \mu\rangle, \\
f(t,x,\mu,\alpha) &= \langle \hat{f}(t,x,\cdot,\alpha), \mu\rangle, &\qquad g(x,\mu) &= \langle \hat{g}(x,\cdot), \mu\rangle
\end{aligned}
$$

for some functions $\hat{b}$, $\hat{\sigma}$ and $\hat{f}$ defined on $[0,T] \times \mathbb{R}^d \times \mathbb{R}^d \times A$ with values in $\mathbb{R}^d$, $\mathbb{R}^{d\times m}$ and $\mathbb{R}$, respectively, and a real valued function $\hat{g}$ defined on $\mathbb{R}^d \times \mathbb{R}^d$. The form of this dependence comes from the original derivation of the McKean–Vlasov equation as limit of the dynamics of a large system of particles evolving according to a system of stochastic differential equations with *mean-field* interactions of the form

$$
dX_t^i = \frac{1}{N}\sum_{j=1}^N \hat{b}\big(t, X_t^i, X_t^j\big)\,dt + \frac{1}{N}\sum_{j=1}^N \hat{\sigma}\big(t, X_t^i, X_t^j\big)\,dW_t^j, \qquad i = 1,\dots,N, \tag{34}
$$

for $t \in [0,T]$, where W^i's are N independent standard Wiener processes in $\mathbb{R}^d$. In the present situation the linearity in μ implies that $\partial_\mu g(x,\mu)(x') = \partial_{x'}\hat{g}(x,x')$, and similarly,

$$
\begin{aligned}
\partial_\mu H(t,x,\mu,y,z,\alpha)(x') = &\ \partial_{x'}\hat{b}\big(t,x,x',\alpha\big) \odot y + \partial_{x'}\hat{\sigma}\big(t,x,x',\alpha\big) \odot z \\
&+ \partial_{x'}\hat{f}\big(t,x,x',\alpha\big),
\end{aligned}
$$

and the dynamic part of the adjoint equation (18) rewrites

$$dY_t = -\tilde{\mathbb{E}}\big[\partial_x \hat{H}(t, X_t, \tilde{X}_t, Y_t, Z_t, \alpha_t) + \partial_{x'} \hat{H}(t, \tilde{X}_t, X_t, \tilde{Y}_t, \tilde{Z}_t, \tilde{\alpha}_t)\big]\,dt + Z_t\,dW_t,$$

if we use the obvious notation

$$\hat{H}\big(t, x, x', y, z, \alpha\big) = \hat{b}\big(t, x, x', \alpha\big) \cdot y + \hat{\sigma}\big(t, x, x', \alpha\big) \cdot z + \hat{f}\big(t, x, x', \alpha\big),$$

and the terminal condition is given by

$$Y_T = \tilde{\mathbb{E}}\big[\partial_x \hat{g}(X_T, \tilde{X}_T) + \partial_{x'} \hat{g}(\tilde{X}_T, X_T)\big].$$

5. Solvability of forward–backward systems. We now turn to the application of the Pontryagin stochastic maximum principle to the solution of the optimal control of McKean–Vlasov dynamics. The strategy is to identify a minimizer of the Hamiltonian, and to use it in the forward dynamics and the adjoint equation. This creates a coupling between these equations, leading to the study of an FBSDE of mean-field type. As explained in the Introduction, the existence results proven in [9] and [8] do not cover some of the solvable models (such as the LQ models). Here we establish existence and uniqueness by taking advantage of the specific structure of the equation, inherited from the underlying optimization problem. Assuming that the terminal cost and the Hamiltonian satisfy the same convexity assumptions as in the statement of Theorem 4.7, we indeed prove that unique solvability holds by applying the *continuation method*, originally exposed within the framework of FBSDEs in [21]. Some of the results of this section were announced in the note [10].

5.1. *Technical assumptions.* We state the conditions we shall use from now on. These assumptions subsume assumptions (A1)–(A4) introduced in Sections 2 and 4. As it is most often the case in applications of the stochastic maximum principle, we choose $A = \mathbb{R}^k$, and we consider a *linear* model for the forward dynamics of the state.

(B1) The drift b and the volatility σ are linear in μ, x and α. They read

$$b(t, x, \mu, \alpha) = b_0(t) + b_1(t)\bar{\mu} + b_2(t)x + b_3(t)\alpha,$$

$$\sigma(t, x, \mu, \alpha) = \sigma_0(t) + \sigma_1(t)\bar{\mu} + \sigma_2(t)x + \sigma_3(t)\alpha,$$

for some bounded measurable deterministic functions b_0, b_1, b_2 and b_3 with values in $\mathbb{R}^d$, $\mathbb{R}^{d\times d}$, $\mathbb{R}^{d\times d}$ and $\mathbb{R}^{d\times k}$, and σ_0, σ_1, σ_2 and σ_3 with values in $\mathbb{R}^{d\times m}$, $\mathbb{R}^{(d\times m)\times d}$, $\mathbb{R}^{(d\times m)\times d}$ and $\mathbb{R}^{(d\times m)\times k}$ [the parentheses around $d \times m$ indicating that $\sigma_i(t)u_i$ is seen as an element of $\mathbb{R}^{d\times m}$ whenever $u_i \in \mathbb{R}^d$, with $i = 1, 2$, or $u_i \in \mathbb{R}^k$, with $i = 3$], and where we use the notation $\bar{\mu} = \int x\,d\mu(x)$ for the mean of a measure μ.

(B2) The functions f and g satisfy the same assumptions as in (A.3)–(A4) in Section 4 (with respect to some constant L). In particular, there exists a constant $\hat{L}$ such that

$$\begin{aligned}
&\bigl|f(t,x',\mu',\alpha') - f(t,x,\mu,\alpha)\bigr| + \bigl|g(x',\mu') - g(x,\mu)\bigr| \\
&\quad \le \hat{L}\bigl[1 + |x'| + |x| + |\alpha'| + |\alpha| + \|\mu\|_2 + \|\mu'\|_2\bigr] \\
&\qquad \times \bigl[\bigl|(x',\alpha') - (x,\alpha)\bigr| + W_2(\mu',\mu)\bigr].
\end{aligned}$$

(B3) There exists a constant $\hat{c} > 0$ such that the derivatives of f and g with respect to (x,α) and x, respectively, are $\hat{c}$-Lipschitz continuous with respect to (x,α,μ) and (x,μ), respectively, the Lipschitz property in the variable μ being understood in the sense of the 2-Wasserstein distance. Moreover, for any $t \in [0,T]$, any $x, x' \in \mathbb{R}^d$, any $\alpha, \alpha' \in \mathbb{R}^k$, any $\mu, \mu' \in \mathcal{P}_2(\mathbb{R}^d)$ and any $\mathbb{R}^d$-valued random variables X and X' having μ and μ' as distributions,

$$\begin{aligned}
&\mathbb{E}\bigl[\bigl|\partial_\mu f(t,x',\mu',\alpha')(X') - \partial_\mu f(t,x,\mu,\alpha)(X)\bigr|^2\bigr] \\
&\quad \le \hat{c}\bigl(\bigl|(x',\alpha') - (x,\alpha)\bigr|^2 + \mathbb{E}\bigl[|X' - X|^2\bigr]\bigr), \\
&\mathbb{E}\bigl[\bigl|\partial_\mu g(x',\mu')(X') - \partial_\mu g(x,\mu)(X)\bigr|^2\bigr] \\
&\quad \le \hat{c}\bigl(|x' - x|^2 + \mathbb{E}\bigl[|X' - X|^2\bigr]\bigr).
\end{aligned}$$

(B4) The function f is convex with respect to (x,μ,α) for t fixed, in such a way that, for some $\lambda > 0$,

$$\begin{aligned}
&f(t,x',\mu',\alpha') - f(t,x,\mu,\alpha) \\
&\quad - \partial_{(x,\alpha)} f(t,x,\mu,\alpha) \cdot (x'-x, \alpha'-\alpha) - \tilde{\mathbb{E}}\bigl[\partial_\mu f(t,x,\mu,\alpha)(\tilde{X}) \cdot (\tilde{X}' - \tilde{X})\bigr] \\
&\qquad \ge \lambda |\alpha' - \alpha|^2,
\end{aligned}$$

whenever $\tilde{X}, \tilde{X}' \in L^2(\tilde{\Omega}, \tilde{\mathcal{A}}, \tilde{\mathbb{P}}; \mathbb{R}^d)$ with distributions μ and μ', respectively. The function g is also assumed to be convex in (x,μ) (on the same model, but with $\lambda = 0$).

We refer to Section 3.2 for a discussion of the conditions (B3) and (B4). By comparing (7) with (B3), we notice that the liftings $L^2(\tilde{\Omega}; \mathbb{R}^d) \ni \tilde{X} \mapsto f(t,x,\tilde{\mathbb{P}}_{\tilde{X}},\alpha)$ and $L^2(\tilde{\Omega}; \mathbb{R}^d) \ni \tilde{X} \mapsto g(x,\tilde{\mathbb{P}}_{\tilde{X}})$ have Lipschitz continuous derivatives. As a consequence, Lemma 3.3 applies, and for any $t \in [0,T]$, $x \in \mathbb{R}^d$, $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\alpha \in \mathbb{R}^k$, there exist versions of $\mathbb{R}^d \ni x' \mapsto \partial_\mu f(t,x,\mu,\alpha)(x')$ and $\mathbb{R}^d \ni x' \mapsto \partial_\mu g(x,\mu)(x')$ which are $\hat{c}$-Lipschitz continuous.

Following example (6), we also emphasize that b and σ obviously satisfy (B3).

5.2. *The Hamiltonian and the adjoint equations*. The drift and the volatility being linear, the Hamiltonian takes the particular form

$$\begin{aligned}
H(t,x,\mu,y,z,\alpha) = {}& \bigl[b_0(t) + b_1(t)\bar{\mu} + b_2(t)x + b_3(t)\alpha\bigr] \cdot y \\
&+ \bigl[\sigma_0(t) + \sigma_1(t)\bar{\mu} + \sigma_2(t)x + \sigma_3(t)\alpha\bigr] \cdot z + f(t,x,\mu,\alpha),
\end{aligned}$$

for $t \in [0, T]$, $x, y \in \mathbb{R}^d$, $z \in \mathbb{R}^{d\times m}$, $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $\alpha \in \mathbb{R}^k$. Given $(t, x, \mu, y, z) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^{d\times m}$, the function $\mathbb{R}^k \ni \alpha \mapsto H(t, x, \mu, y, z, \alpha)$ is strictly convex so that there exists a unique minimizer $\hat{\alpha}(t, x, \mu, y, z)$

$$\hat{\alpha}(t, x, \mu, y, z) = \operatorname*{argmin}_{\alpha\in\mathbb{R}^k} H(t, x, \mu, y, z, \alpha). \tag{35}$$

Assumptions (B1)–(B4) above being slightly stronger than the assumptions used in [9], we can follow the arguments given in the proof of Lemma 2.1 of [9] in order to prove that, for all $(t, x, \mu, y, z) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^{d\times m}$, the function $[0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^{d\times m} \ni (t, x, \mu, y, z) \mapsto \hat{\alpha}(t, x, \mu, y, z)$ is measurable, locally bounded and Lipschitz-continuous with respect to (x, μ, y, z), uniformly in $t \in [0, T]$, the Lipschitz constant depending only upon λ, the supremum norms of b_3 and σ_3 and the Lipschitz constant of $\partial_\alpha f$ in (x, μ). Except maybe for the Lipschitz property with respect to the measure argument, these facts were explicitly proved in [9]. The regularity of $\hat{\alpha}$ with respect to μ follows from the following remark. If $(t, x, y, z) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^{d\times m}$ is fixed and μ, μ' are generic elements in $\mathcal{P}_2(\mathbb{R}^d)$, $\hat{\alpha}$ and $\hat{\alpha}'$ denoting the associated minimizers, we deduce from the convexity assumption (B4),

$$\begin{aligned}
2\lambda|\hat{\alpha}' - \hat{\alpha}|^2 &\le (\hat{\alpha}' - \hat{\alpha}) \cdot \big[\partial_\alpha f\big(t, x, \mu, \hat{\alpha}'\big) - \partial_\alpha f(t, x, \mu, \hat{\alpha})\big] \\
&= (\hat{\alpha}' - \hat{\alpha}) \cdot \big[\partial_\alpha H\big(t, x, \mu, y, z, \hat{\alpha}'\big) - \partial_\alpha H(t, x, \mu, y, z, \hat{\alpha})\big] \\
&= (\hat{\alpha}' - \hat{\alpha}) \cdot \big[\partial_\alpha H\big(t, x, \mu, y, z, \hat{\alpha}'\big) - \partial_\alpha H\big(t, x, \mu', y, z, \hat{\alpha}'\big)\big] \\
&= (\hat{\alpha}' - \hat{\alpha}) \cdot \big[\partial_\alpha f\big(t, x, \mu, \hat{\alpha}'\big) - \partial_\alpha f\big(t, x, \mu', \hat{\alpha}'\big)\big] \\
&\le C|\hat{\alpha}' - \hat{\alpha}|W_2(\mu', \mu),
\end{aligned} \tag{36}$$

the passage from the second to the third line following from the identity

$$\partial_\alpha H(t, x, \mu, y, z, \hat{\alpha}) = \partial_\alpha H\big(t, x, \mu', y, z, \hat{\alpha}'\big) = 0.$$

For each admissible control $\alpha = (\alpha_t)_{0\le t\le T}$, if we denote the corresponding solution of the state equation by $X = (X_t^\alpha)_{0\le t\le T}$, then the adjoint BSDE (18) introduced in Definition 3.6 reads

$$\begin{aligned}
dY_t = &-\partial_x f(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dt - b_2^\dagger(t)Y_t\,dt - \sigma_2^\dagger(t)Z_t\,dt + Z_t\,dW_t \\
&- \tilde{\mathbb{E}}\big[\partial_\mu f(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{\alpha}_t)(X_t)\big]\,dt - b_1^\dagger(t)\mathbb{E}[Y_t]\,dt - \sigma_1^\dagger(t)\mathbb{E}[Z_t]\,dt.
\end{aligned} \tag{37}$$

Given the necessary and sufficient conditions proven in the previous section, our goal is to use the control $(\hat{\alpha}_t)_{0\le t\le T}$ defined by $\hat{\alpha}_t = \hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t)$ where $\hat{\alpha}$ is the minimizer function constructed above, and $(X_t, Y_t, Z_t)_{0\le t\le T}$ is a solution of the FBSDE

$$\begin{aligned}
dX_t = &\big[b_0(t) + b_1(t)\mathbb{E}[X_t] + b_2(t)X_t + b_3(t)\hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t)\big]\,dt \\
&+ \big[\sigma_0(t) + \sigma_1(t)\mathbb{E}[X_t] + \sigma_2(t)X_t + \sigma_3(t)\hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t)\big]\,dW_t,
\end{aligned} \tag{38}$$

$$
\begin{aligned}
dY_t = &-\big[\partial_x f\big(t, X_t, \mathbb{P}_{X_t}, \hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t)\big) + b_2^\dagger(t)Y_t + \sigma_2^\dagger(t)Z_t\big]\,dt \\
&+ Z_t\,dW_t \\
&- \big\{\tilde{\mathbb{E}}\big[\partial_\mu f\big(t, \tilde{X}_t, \mathbb{P}_{X_t}, \hat{\alpha}(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{Y}_t, \tilde{Z}_t)\big)(X_t)\big] \\
&\qquad\qquad + b_1^\dagger(t)\mathbb{E}[Y_t] + \sigma_1^\dagger(t)\mathbb{E}[Z_t]\big\}\,dt,
\end{aligned}
$$

with the initial condition $X_0 = x_0$, for a given deterministic point $x_0 \in \mathbb{R}^d$, and the terminal condition $Y_T = \partial_x g(X_T, \mathbb{P}_{X_T}) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T, \mathbb{P}_{X_T})(X_T)]$.

5.3. *Main result.* Here is the main existence and uniqueness result:

THEOREM 5.1. *Under* (B1)–(B4), *the forward–backward system* (38) *is uniquely solvable.*

PROOF. The proof is an adaptation of the *continuation method* used in [21] to handle standard FBSDEs satisfying appropriate monotonicity conditions. Generally speaking, it consists in proving that existence and uniqueness are kept preserved when the coefficients in (38) are slightly perturbed. Starting from an initial case for which existence and uniqueness are known to hold, we then establish Theorem 5.1 by modifying iteratively the coefficients so that (38) is eventually shown to belong to the class of uniquely solvable systems.

A natural and simple strategy then consists in modifying the coefficients in a linear way. Unfortunately, this might generate heavy notation. For this reason, we use the following conventions.

First, as in Section 4.1, the notation $(\Theta_t)_{0\le t\le T}$ stands for the generic notation for denoting a process of the form $(X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t)_{0\le t\le T}$ with values in $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^{d\times m} \times \mathbb{R}^k$. We will then denote by $\mathbb{S}$ the space of processes $(\Theta_t)_{0\le t\le T}$ such that $(X_t, Y_t, Z_t, \alpha_t)_{0\le t\le T}$ is $(\mathcal{F}_t)_{0\le t\le T}$ progressively-measurable, $(X_t)_{0\le t\le T}$ and $(Y_t)_{0\le t\le T}$ have continuous trajectories and

$$
\|\Theta\|_{\mathbb{S}} = \mathbb{E}\Big[\sup_{0\le t\le T}\big[|X_t|^2 + |Y_t|^2\big] + \int_0^T \big[|Z_t|^2 + |\alpha_t|^2\big]\,dt\Big]^{1/2} < +\infty. \tag{39}
$$

Similarly, the notation $(\theta_t)_{0\le t\le T}$ is the generic notation for denoting a process $(X_t, \mathbb{P}_{X_t}, \alpha_t)_{0\le t\le T}$ with values in $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^k$. All the processes $(\theta_t)_{0\le t\le T}$ that are considered below appear as the restrictions of an *extended* process $(\Theta_t)_{0\le t\le T} \in \mathbb{S}$.

Moreover, we call an initial condition for (38) a square-integrable $\mathcal{F}_0$-measurable random variable ξ with values in $\mathbb{R}^d$, that is, an element of $L^2(\Omega, \mathcal{F}_0, \mathbb{P}; \mathbb{R}^d)$. Recall indeed that $\mathcal{F}_0$ can be chosen as a σ-algebra independent of $(W_t)_{0\le t\le T}$. In comparison with the statement of Theorem 5.1, this permits us to generalize the case when ξ is deterministic.

Finally, we call an input for (38) a four-tuple $\mathcal{I} = ((\mathcal{I}_t^b, \mathcal{I}_t^\sigma, \mathcal{I}_t^f)_{0\le t\le T}, \mathcal{I}_T^g)$, $(\mathcal{I}_t^b)_{0\le t\le T}$, $(\mathcal{I}_t^\sigma)_{0\le t\le T}$ and $(\mathcal{I}_t^f)_{0\le t\le T}$ being three square-integrable progressively-measurable processes with values in $\mathbb{R}^d$, $\mathbb{R}^{d\times m}$ and $\mathbb{R}^d$, respectively, and $\mathcal{I}_T^g$ denoting a square-integrable $\mathcal{F}_T$-measurable random variable with values in $\mathbb{R}^d$. Such an input is specifically designed to be injected into the dynamics of (38), $\mathcal{I}^b$ being plugged into the drift of the forward equation, $\mathcal{I}^\sigma$ into the volatility of the forward equation, $\mathcal{I}^f$ into the bounded variation term of the backward equation and $\mathcal{I}^g$ into the terminal condition of the backward equation. The space of inputs is denoted by $\mathbb{I}$. It is endowed with the norm

$$\|\mathcal{I}\|_{\mathbb{I}} = \mathbb{E}\biggl[|\mathcal{I}_T^g|^2 + \int_0^T \bigl[|\mathcal{I}_t^b|^2 + |\mathcal{I}_t^\sigma|^2 + |\mathcal{I}_t^f|^2\bigr]\,dt\biggr]^{1/2}. \tag{40}$$

We then put:

DEFINITION 5.2. For any $\gamma \in [0, 1]$, any $\xi \in L^2(\Omega, \mathcal{F}_0, \mathbb{P}; \mathbb{R}^d)$ and any input $\mathcal{I} \in \mathbb{I}$, the FBSDE

$$\begin{aligned} dX_t &= \bigl(\gamma b(t, \theta_t) + \mathcal{I}_t^b\bigr)\,dt + \bigl(\gamma\sigma(t, \theta_t) + \mathcal{I}_t^\sigma\bigr)\,dW_t, \\ dY_t &= -\bigl(\gamma\bigl\{\partial_x H(t, \Theta_t) + \tilde{\mathbb{E}}\bigl[\partial_\mu H(t, \tilde{\Theta}_t)(X_t)\bigr]\bigr\} + \mathcal{I}_t^f\bigr)\,dt + Z_t\,dW_t, \end{aligned} \tag{41}$$

for $t \in [0, T]$, with the optimality condition

$$\alpha_t = \hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t), \qquad t \in [0, T], \tag{42}$$

and with $X_0 = \xi$ as initial condition and

$$Y_T = \gamma\bigl\{\partial_x g(X_T, \mathbb{P}_{X_T}) + \tilde{\mathbb{E}}\bigl[\partial_\mu g(\tilde{X}_T, \mathbb{P}_{X_T})(X_T)\bigr]\bigr\} + \mathcal{I}_T^g$$

as terminal condition, is referred to as $\mathcal{E}(\gamma, \xi, \mathcal{I})$.

Whenever $(X_t, Y_t, Z_t)_{0\le t\le T}$ is a solution, $(X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t)_{0\le t\le T}$ is referred to as the *associated extended solution*.

REMARK 5.3. The way the coupling is summarized between the forward and backward equations in (41) is a bit different from the way equation (38) is written. In the formulation used in the statement of Lemma 5.2, the coupling between the forward and the backward equations follows from the optimality condition (42). Because of that optimality condition, the two formulations are equivalent: When $\gamma = 1$ and $\mathcal{I} \equiv 0$, the pair (41)–(42) coincides with (38).

The following lemma is proved in the next subsection:

LEMMA 5.4. *Given $\gamma \in [0, 1]$, we say that property $(\mathcal{S}_\gamma)$ holds true if, for any $\xi \in L^2(\Omega, \mathcal{F}_0, \mathbb{P}; \mathbb{R}^d)$ and any $\mathcal{I} \in \mathbb{I}$, the FBSDE $\mathcal{E}(\gamma, \xi, \mathcal{I})$ has a unique extended solution in $\mathbb{S}$. With this definition, there exists $\delta_0 > 0$ such that, if $(\mathcal{S}_\gamma)$ holds true for some $\gamma \in [0, 1)$, then $(\mathcal{S}_{\gamma+\eta})$ holds true for any $\eta \in (0, \delta_0]$ satisfying $\gamma + \eta \le 1$.*

Given Lemma 5.4, Theorem 5.1 follows from a straightforward induction as $(\mathcal{S}_0)$ obviously holds true. $\square$

5.4. *Proof of Lemma* 5.4. The proof follows from Picard's contraction theorem. As in the statement, consider indeed γ such that $(\mathcal{S}_\gamma)$ holds true. For $\eta > 0$, $\xi \in L^2(\Omega, \mathcal{F}_0, \mathbb{P}; \mathbb{R}^d)$ and $\mathcal{I} \in \mathbb{I}$, we then define a mapping Φ from $\mathbb{S}$ into itself whose fixed points coincide with the solutions of $\mathcal{E}(\gamma + \eta, \xi, \mathcal{I})$.

The definition of Φ is as follows. Given a process $\Theta \in \mathbb{S}$, we denote by Θ' the extended solution of the FBSDE $\mathcal{E}(\gamma, \xi, \mathcal{I}')$ with

$$\begin{aligned}
&\mathcal{I}_t^{b,\prime} = \eta b(t, \theta_t) + \mathcal{I}_t^b, \\
&\mathcal{I}_t^{\sigma,\prime} = \eta \sigma(t, \theta_t) + \mathcal{I}_t^\sigma, \\
&\mathcal{I}_t^{f,\prime} = \eta \partial_x H(t, \Theta_t) + \eta \tilde{\mathbb{E}}\big[\partial_\mu H(t, \tilde{\Theta}_t)(X_t)\big] + \mathcal{I}_t^f, \\
&\mathcal{I}_T^{g,\prime} = \eta \partial_x g(X_T, \mathbb{P}_{X_T}) + \eta \tilde{\mathbb{E}}\big[\partial_\mu g(\tilde{X}_T, \mathbb{P}_{X_T})(X_T)\big] + \mathcal{I}_T^g.
\end{aligned}$$

By assumption, it is uniquely defined, and it belongs to $\mathbb{S}$, so that the mapping $\Phi : \Theta \mapsto \Theta'$ maps $\mathbb{S}$ into itself. It is then clear that a process $\Theta \in \mathbb{S}$ is a fixed point of Φ if and only if Θ is an extended solution of $\mathcal{E}(\gamma + \eta, \xi, \mathcal{I})$. So we only need to prove that Φ is a contraction when η is small enough. This is a consequence of the following lemma:

LEMMA 5.5. *Let $\gamma \in [0, 1]$ such that $(\mathcal{S}_\gamma)$ holds true. Then there exists a constant C, independent of γ, such that, for any $\xi, \xi' \in L^2(\Omega, \mathcal{F}_0, \mathbb{P}; \mathbb{R}^d)$ and $\mathcal{I}, \mathcal{I}' \in \mathbb{I}$, the respective extended solutions Θ and Θ' of $\mathcal{E}(\gamma, \xi, \mathcal{I})$ and $\mathcal{E}(\gamma, \xi', \mathcal{I}')$ satisfy*

$$\|\Theta - \Theta'\|_{\mathbb{S}} \le C\big(\mathbb{E}\big[|\xi - \xi'|^2\big]^{1/2} + \|\mathcal{I} - \mathcal{I}'\|_{\mathbb{I}}\big).$$

Given Lemma 5.5, we indeed check that Φ is a contraction when η is small enough. Given Θ^1 and Θ^2 two processes in $\mathbb{S}$ and denoting by $\Theta'^{,1}$ and $\Theta'^{,2}$ their respective images by Φ, we deduce from Lemma 5.5 that

$$\|\Theta'^{,1} - \Theta'^{,2}\|_{\mathbb{S}} \le C\eta \|\Theta^1 - \Theta^2\|_{\mathbb{S}},$$

which is enough to conclude.

5.5. *Proof of Lemma* 5.5. The strategy follows from a mere variation on the proof of the classical stochastic maximum principle. With the same notation as in the statement, and with the convention of expanding Θ as $(X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t)_{0 \le t \le T}$ and letting $(\theta_t = (X_t, \mathbb{P}_{X_t}, \alpha_t))_{0 \le t \le T}$, we compute

$$\begin{aligned}
&\mathbb{E}\big[(X_T' - X_T) \cdot Y_T\big] \\
&\quad = \mathbb{E}\big[(\xi' - \xi) \cdot Y_0\big] \\
&\qquad - \gamma \Big\{ \mathbb{E} \int_0^T \big[\partial_x H(t, \Theta_t) \cdot (X_t' - X_t) + \tilde{\mathbb{E}}\big[\partial_\mu H(t, \tilde{\Theta}_t)(X_t)\big] \cdot (X_t' - X_t)\big]\, dt
\end{aligned}$$

$$
\begin{aligned}
&\quad - \mathbb{E}\int_0^T \big[\big[b(t,\theta_t') - b(t,\theta_t)\big]\cdot Y_t + \big[\sigma(t,\theta_t') - \sigma(t,\theta_t)\big]\cdot Z_t\big]\,dt\Big\} \\
&\quad - \Big\{\mathbb{E}\int_0^T \big[(X_t' - X_t)\cdot \mathcal{I}_t^f + (\mathcal{I}_t^b - \mathcal{I}_t^{b,\prime})\cdot Y_t + (\mathcal{I}_t^\sigma - \mathcal{I}_t^{\sigma,\prime})\cdot Z_t\big]\,dt\Big\} \\
&= T_0 - \gamma T_1 - T_2.
\end{aligned}
$$

Following (31),

$$
\begin{aligned}
&\mathbb{E}\big[(X_T' - X_T)\cdot Y_T\big] \\
&\quad = \gamma\mathbb{E}\big[\big(\partial_x g(X_T,\mathbb{P}_{X_T}) + \tilde{\mathbb{E}}\big[\partial_\mu g(\tilde{X}_T,\mathbb{P}_{X_T})(X_T)\big]\big)\cdot (X_T' - X_T)\big] \\
&\qquad + \mathbb{E}\big[(\mathcal{I}_T^{g,\prime} - \mathcal{I}_T^g)\cdot Y_T\big] \\
&\quad \le \gamma\mathbb{E}\big[g(X_T',\mathbb{P}_{X_T'}) - g(X_T,\mathbb{P}_{X_T})\big] + \mathbb{E}\big[(\mathcal{I}_T^{g,\prime} - \mathcal{I}_T^g)\cdot Y_T\big].
\end{aligned}
$$

Identifying the two expressions above and repeating the proof of Theorem 4.7, we obtain

$$
\begin{aligned}
&\gamma J(\alpha') - \gamma J(\alpha) \\
&\quad \ge \gamma\lambda\mathbb{E}\int_0^T |\alpha_t - \alpha_t'|^2\,dt + T_0 - T_2 + \mathbb{E}\big[(\mathcal{I}_T^g - \mathcal{I}_T^{g,\prime})\cdot Y_T\big].
\end{aligned}
\tag{43}
$$

Now we can reverse the roles of α and α' in (43). Denoting by T_0' and T_2' the corresponding terms in the inequality and summing both inequalities, we deduce that

$$
2\gamma\lambda\mathbb{E}\int_0^T |\alpha_t - \alpha_t'|^2\,dt + T_0 + T_0' - (T_2 + T_2') + \mathbb{E}\big[(\mathcal{I}_T^g - \mathcal{I}_T^{g,\prime})\cdot (Y_T - Y_T')\big] \le 0.
$$

The sum $T_2 + T_2'$ reads

$$
\begin{aligned}
T_2 + T_2' = \mathbb{E}\int_0^T \big[&-(\mathcal{I}_t^f - \mathcal{I}_t^{f,\prime})\cdot (X_t - X_t') + (\mathcal{I}_t^b - \mathcal{I}_t^{b,\prime})\cdot (Y_t - Y_t') \\
&+ (\mathcal{I}_t^\sigma - \mathcal{I}_t^{\sigma,\prime})\cdot (Z_t - Z_t')\big]\,dt.
\end{aligned}
$$

Similarly,

$$
T_0 + T_0' = -\mathbb{E}\big[(\xi - \xi')\cdot (Y_0 - Y_0')\big].
$$

Therefore, using Young's inequality, there exists a constant C (the value of which may change from line to line), C being independent of γ, such that, for any $\varepsilon > 0$,

$$
\gamma\mathbb{E}\int_0^T |\alpha_t - \alpha_t'|^2\,dt \le \varepsilon\|\Theta - \Theta'\|_{\mathbb{S}}^2 + \frac{C}{\varepsilon}\big(\mathbb{E}\big[|\xi - \xi'|^2\big] + \|\mathcal{I} - \mathcal{I}'\|_{\mathbb{I}}^2\big).
\tag{44}
$$

Now, we observe that, by standard estimates for BSDEs, there exists a constant C, independent of γ, such that

$$
\begin{aligned}
&\mathbb{E}\Big[\sup_{0\le t\le T}|Y_t - Y_t'|^2 + \int_0^T |Z_t - Z_t'|^2\,dt\Big] \\
&\quad \le C\gamma\mathbb{E}\Big[\sup_{0\le t\le T}|X_t - X_t'|^2 + \int_0^T |\alpha_t - \alpha_t'|^2\,dt\Big] + C\|\mathcal{I} - \mathcal{I}'\|_{\mathbb{I}}^2.
\end{aligned}
\tag{45}
$$

Similarly,

$$
\begin{aligned}
&\mathbb{E}\Big[\sup_{0\le t\le T}|X_t - X_t'|^2\Big] \\
&\quad \le \mathbb{E}\big[|\xi - \xi'|^2\big] + C\gamma\mathbb{E}\int_0^T |\alpha_t - \alpha_t'|^2\,dt + C\|\mathcal{I} - \mathcal{I}'\|_{\mathbb{I}}^2.
\end{aligned}
\tag{46}
$$

From (45), (46) and (44), we deduce that

$$
\begin{aligned}
&\mathbb{E}\Big[\sup_{0\le t\le T}|X_t - X_t'|^2 + \sup_{0\le t\le T}|Y_t - Y_t'|^2 + \int_0^T |Z_t - Z_t'|^2\,dt\Big] \\
&\quad \le C\gamma\mathbb{E}\int_0^T |\alpha_t - \alpha_t'|^2\,dt + C\big(\mathbb{E}\big[|\xi - \xi'|^2\big] + \|\mathcal{I} - \mathcal{I}'\|_{\mathbb{I}}^2\big) \\
&\quad \le C\varepsilon\|\Theta - \Theta'\|_{\mathbb{S}}^2 + \frac{C}{\varepsilon}\big(\mathbb{E}\big[|\xi - \xi'|^2\big] + \|\mathcal{I} - \mathcal{I}'\|_{\mathbb{I}}^2\big).
\end{aligned}
\tag{47}
$$

Using the Lispchitz property of $\hat{\alpha}(t,\cdot,\cdot,\cdot,\cdot)$ and choosing ε small enough, we complete the proof.

5.6. *Decoupling field.* The notion of decoupling field, also referred to as an *FBSDE value function*, plays a main role in the machinery of forward–backward equations. Indeed, it provides a representation of the value Y_t of the backward component at time t, as a function of the value X_t of the forward component. When the coefficients of the forward–backward equation are random, the decoupling field is a random field. When the coefficients are deterministic, the decoupling field is a deterministic function, which solves a specific partial differential equation. Here is the structure of the decoupling field in the McKean–Vlasov framework:

LEMMA 5.6. *For any $t \in [0,T]$ and any $\xi \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$, there exists a unique solution, denoted by $(X_s^{t,\xi}, Y_s^{t,\xi}, Z_s^{t,\xi})_{t\le s\le T}$, of (38) on $[t,T]$ with $X_t^{t,\xi} = \xi$ as initial condition.*

For any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, there exists a measurable mapping $u(t,\cdot,\mu): \mathbb{R}^d \ni x \mapsto u(t,x,\mu)$ such that

$$
\mathbb{P}\big(Y_t^{t,\xi} = u(t,\xi,\mathbb{P}_\xi)\big) = 1. \tag{48}
$$

Moreover, there exists a constant C, depending only on the parameters in (B1)–(B4), *such that, for any $t \in [0, T]$ and any $\xi^1, \xi^2 \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$,*

$$\mathbb{E}\big[\big|u\big(t, \xi^1, \mathbb{P}_{\xi^1}\big) - u\big(t, \xi^2, \mathbb{P}_{\xi^2}\big)\big|^2\big] \le C\mathbb{E}\big[|\xi^1 - \xi^2|^2\big]. \tag{49}$$

The proof is given below. For the time being, we notice that the additional variable $\mathbb{P}_\xi$ is for free in the above writing since we could set $v(t, \cdot) = u(t, \cdot, \mathbb{P}_\xi)$ and then have $Y_t^{t,\xi} = v(t, \xi)$. The additional variable $\mathbb{P}_\xi$ is specified to emphasize the non-Markovian nature of the equation over the state space $\mathbb{R}^d$: starting from two different initial conditions, the decoupling fields might not be the same, since the law of the initial conditions might be different. It is important to keep indeed in mind that, in the Markovian framework, the decoupling field is the same for all possible initial conditions, thus yielding the connection with partial differential equations. Here the Markov property holds, but over the enlarged space $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, thus justifying the use of the extra variable $\mathbb{P}_\xi$. Nevertheless, we often remove the dependence upon $\mathbb{P}_\xi$ in the remaining of the paper.

An important fact is that the representation formula (48) can be extended to the whole path:

PROPOSITION 5.7. *Under* (B1)–(B4), *for any $\xi \in L^2(\Omega, \mathcal{F}_0, \mathbb{P}; \mathbb{R}^d)$, there exists a measurable mapping $v : [0, T] \times \mathbb{R}^d \to \mathbb{R}^d$ such that*

$$\mathbb{P}\big(\forall t \in [0, T], Y_t^{0,\xi} = v\big(t, X_t^{0,\xi}\big)\big) = 1.$$

It satisfies $\sup_{0 \le t \le T} |v(t, 0)| < +\infty$. Moreover, there exists a constant C such that $v(t, \cdot)$ is C-Lipschitz continuous for any $t \in [0, T]$.

We start with the following:

PROOF OF LEMMA 5.6. Given $t \in [0, T)$ and $\xi \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$, existence and uniqueness of a solution of (38) on $[t, T]$ with ξ as initial condition is a direct consequence of Theorem 5.1, more precisely of its proof since we are starting from a random initial condition. Using as underlying filtration the augmented filtration $\mathbb{F}^t$ generated by ξ and by $(W_s - W_t)_{t \le s \le T}$, we deduce that $Y_t^{t,\xi}$ coincides a.s. with a $\sigma(\xi)$-measurable $\mathbb{R}^d$-valued random variable. In particular, there exists a measurable function $u_\xi(t, \cdot) : \mathbb{R}^d \to \mathbb{R}^d$ such that $\mathbb{P}[Y_t^{t,\xi} = u_\xi(t, \xi)] = 1$.

We now claim that the law of $(\xi, Y_t^{t,\xi})$ only depends upon the law of ξ. This directly follows from the version of the Yamada–Watanabe theorem for FBSDEs [11]. Since uniqueness holds pathwise, it also holds in law, so that given two initial conditions with the same law, the solutions also have the same laws. Therefore, given another $\mathbb{R}^d$-valued random vector ξ' with the same law as ξ, it holds $(\xi, u_\xi(t, \xi)) \sim (\xi', u_{\xi'}(t, \xi'))$. In particular, for any measurable function $v : \mathbb{R}^d \to \mathbb{R}^d$, the random variables $u_\xi(t, \xi) - v(\xi)$ and $u_{\xi'}(t, \xi') - v(\xi')$ have the same law. Choosing $v = u_\xi(t,\cdot)$, we deduce that $u_{\xi'}(t,\cdot)$ and $u_\xi(t,\cdot)$ are a.e. equal under the probability measure $\mathbb{P}_\xi$. To put it differently, denoting by μ the law of ξ, there exists an element $u(t,\cdot,\mu) \in L^2(\mathbb{R}^d,\mu)$ such that $u_\xi(t,\cdot)$ and $u_{\xi'}(t,\cdot)$ coincide μ a.e. with $u(t,\cdot,\mu)$. Identifying $u(t,\cdot,\mu)$ with one of its version, this proves that

$$\mathbb{P}\big(Y_t^{t,\xi} = u(t,\xi,\mu)\big) = 1.$$

When $t > 0$, we notice that, for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, there exists an $\mathcal{F}_t$-measurable random variable ξ such that $\mu = \mathbb{P}_\xi$. In such a case, the procedure we just described permits us to define $u(t,\cdot,\mu)$ for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. The situation may be different when $t = 0$ as $\mathcal{F}_0$ may reduce to events of measure zero or one. In such a case, $\mathcal{F}_0$ can be enlarged without any loss of generality in order to support $\mathbb{R}^d$-valued random variables with arbitrary distributions.

The Lipschitz property (49) of $u(0,\cdot,\cdot)$ is a direct consequence of Lemma 5.5 with $\gamma = 1$. By a time shift, the same argument applies to $u(t,\cdot,\cdot)$. □

We now turn to the following:

PROOF OF PROPOSITION 5.7. For the sake of simplicity, we denote the process $(X_t^{0,\xi}, Y_t^{0,\xi}, Z_t^{0,\xi})_{0\le t\le T}$ by $(X_t, Y_t, Z_t)_{0\le t\le T}$. The proof is then a combination of Lemmas 3.3 and 5.6. Indeed, given $t \in (0,T]$, Lemma 5.6 says that the family $(u(t,\cdot,\mu))_{\mu\in\mathcal{P}_2(\mathbb{R}^d)}$ satisfies (8) since any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ can be seen as the law of some $\mathcal{F}_t$-measurable random vector ζ. Therefore, for $\mu = \mathbb{P}_{X_t}$, we can find a mapping $w(t,\cdot)$ that is C-Lipschitz continuous [for the same C as in (49)] and that coincides with $u(t,\cdot,\mathbb{P}_{X_t})$ a.e. under the probability measure $\mathbb{P}_{X_t}$. It satisfies

$$\forall t \in [0,T] \qquad \mathbb{P}\big[Y_t = w(t,X_t)\big] = 1, \tag{50}$$

since $Y_t = Y_t^{t,X_t}$. In particular,

$$\begin{aligned}\sup_{0\le t\le T}\big|w(t,0)\big| &\le \sup_{0\le t\le T}\mathbb{E}\big[|Y_t|\big] + \sup_{0\le t\le T}\mathbb{E}\big[\big|w(t,X_t) - w(t,0)\big|\big]\\ &\le \sup_{0\le t\le T}\mathbb{E}\big[|Y_t|\big] + C\sup_{0\le t\le T}\mathbb{E}\big[|X_t|\big] < +\infty.\end{aligned} \tag{51}$$

For any integer $n \ge 1$, we then let

$$v^{(n)}(t,x) = \mathbf{1}_{[0,T/2^n]}(t)\, w\left(\frac{T}{2^n}, x\right) + \sum_{k=2}^{2^n} \mathbf{1}_{((k-1)T/2^n, kT/2^n]}(t)\, w\left(\frac{kT}{2^n}, x\right),$$

for $t \in [0,T]$ and $x \in \mathbb{R}^d$. Denoting by $v^{(n),i}$ the ith coordinate of $v^{(n)}$ for any $i \in \{1,\dots,d\}$, we also let

$$v^i(t,x) = \limsup_{n\to+\infty} v^{(n),i}(t,x), \qquad t \in [0,T], x \in \mathbb{R}^d,$$

and then $v(t,x) = (v^1(t,x), \dots, v^d(t,x))$. As each of the $v^{(n)}$ is a Borel measurable function on $[0,T] \times \mathbb{R}^d$, so is v. Similarly, v satisfies (51) and, for any $t \in [0,T]$, $v(t,\cdot)$ is C-Lipschitz continuous.

Finally, we notice that, for any $t \in \mathbb{D}_n = \{kT/2^n, k \in \{1,\dots,2^n\}\}$, with $n \in \mathbb{N} \setminus \{0\}$, and for $\ell \geq n$,

$$v^{(\ell)}(t,\cdot) = w(t,\cdot) = v(t,\cdot),$$

so that, $w(t,\cdot) = v(t,\cdot)$ for any $t \in \mathbb{D} = \bigcup_{n \geq 1} \mathbb{D}_n$. Therefore, $\mathbb{D}$ being countable, we deduce from (50) that the event

$$A = \{\omega \in \Omega : \forall t \in \mathbb{D}, Y_t(\omega) = w(t, X_t(\omega)) = v(t, X_t(\omega))\}$$

has full measure, that is, $\mathbb{P}(A) = 1$. On the event A, we notice that, for any $t \in (0,T]$,

$$Y_t = \lim_{n \to +\infty} Y_{t_n} = \lim_{n \to +\infty} v(t_n, X_{t_n}),$$

where $(t_n)_{n \geq 1}$ is the sequence of points in $\mathbb{D}$ such that, for any $n \geq 1$, $t_n \in \mathbb{D}_n$ and $t_n - T/2^n < t \leq t_n$. Since $v(t_n,\cdot)$ is C-Lipschitz continuous, we deduce that, on A,

$$Y_t = \lim_{n \to +\infty} v(t_n, X_t),$$

which shows that the sequence $(v(t_n, X_t))_{n \geq 1}$ is convergent. Now we observe that $v(t_n, X_t)$ is also equal to $v^{(n)}(t, X_t)$. Therefore, the limit must coincide with $v(t, X_t)$. This proves that, on the event A, $Y_t = v(t, X_t)$ for any $t \in (0,T]$. By the same argument, the same equality holds true for $t = 0$. The case $t = 0$ is handled separately for notational reasons since the definition of v_n at time 0 is different. □

5.7. *Comments*, *conjectures and future prospects*. A first question concerns the possible extension of the above results to the case of random coefficients. It is indeed well known that the classical Pontryagin stochastic maximum principle also applies to systems with random coefficients. The same should hold true in the McKean–Vlasov case with a modicum of care.

Basically, Theorems 4.5 (necessary condition in the Pontryagin principle), 4.7 (sufficient condition in the Pontryagin principle) and 5.1 are still valid for random coefficients. Allowing random coefficients means that b, σ and f may depend upon the realization $\omega \in \Omega$ in a progressively-measurable way with respect to the filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$, and that g may also depend upon the randomness in a measurable way with respect to the σ-field $\mathcal{F}_T$. The various assumptions, which may differ from one theorem to another, are then supposed to hold true pathwise under the probability measure $\mathbb{P}$. The shape of the underlying adjoint backward equation is then the same as that one in (18), provided that the *independent copy* made from the space $(\Omega, \mathcal{F}, \mathbb{P})$ to the space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ also takes into account the dependence upon the randomness. This means that, in (18), the value of $\partial_\mu H$ in

$\partial_\mu H(t, \tilde{X}_t, \mathbb{P}_{X_t}, \tilde{Y}_t, \tilde{Z}_t, \tilde{\alpha}_t)$ must be computed along the corresponding realization $\tilde{\omega} \in \tilde{\Omega}$ (as $\tilde{X}_t$, $\tilde{Y}_t$, $\tilde{Z}_t$ and $\tilde{\alpha}_t$ actually stand for $\tilde{X}_t(\tilde{\omega})$, $\tilde{Y}_t(\tilde{\omega})$, $\tilde{Z}_t(\tilde{\omega})$ and $\tilde{\alpha}_t(\tilde{\omega})$). This principle holds true for all the "copies" considered in the computations.

It seems much less clear how the results about the decoupling field can be extended to the random setting. Indeed $Y_t^{t,\xi}$ is already random, so that the decoupling field u_ξ introduced in the proof of Lemma 5.6 is also random. Following the proof of Lemma 5.6, we can write it as $u_\xi(\omega, t, x)$, but the shape of u_ξ cannot be entirely described by the law of ξ as it also relies on the joint law of ξ and (b, f, σ, g). Unfortunately, the proofs of Lemma 3.3 and Proposition 5.7 rely on a coupling argument which fails when the decoupling field depends on the joint distribution of ξ and (b, f, σ, g).

Going back to the original case of deterministic coefficients, an interesting challenge is to identify an analytical counterpart of the probabilistic approach to the optimization problem. In the standard (non-McKean–Vlasov) setting, the optimization problem can be tackled by solving the Hamilton–Jacobi–Bellman equation satisfied by the value function. In the McKean–Vlasov setting, the problem is much more challenging. As highlighted by Proposition 5.7, the problem could be given a Markov structure on the enlarged space $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$. As a result, the corresponding Hamilton–Jacobi–Bellman equation, if it exists, would hold as a partial differential equation on $[0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$. We do believe that the differential calculus presented in Section 3 should allow to write such a Hamilton–Jacobi–Bellman equation rigorously. It would be of the same general structure as the so-called *master equation* of mean-field games suggested by Lions in his lectures at the Collège de France, and described in the notes by Cardaliaguet [6] in a case of some first order differential games. We believe that the differences between those two equations should be borne by the following fundamental distinction. The *master equation* for controlled McKean–Vlasov processes is expected to be a Hamilton–Jacobi–Bellman equation as it derives from an optimization problem, whereas the *master equation* for mean-field games is not a Hamilton–Jacobi–Bellman equation since the fixed point condition describing equilibriums in a mean-field game does not derive from an optimization criterion. We investigate this question in the forthcoming work [7].

6. Propagation of chaos and approximate equilibrium.

In this section, we show how the solution of the optimal control of McKean–Vlasov dynamics can provide equilibriums for N-player games when N tends to $+\infty$.

Throughout this section, assumptions (B1)–(B4) are in force. For each integer $N \geq 1$, we consider a stochastic system whose time evolution is given by a system of N coupled stochastic differential equations of the form

$$
\begin{aligned}
&dU_t^i = b\big(t, U_t^i, \bar{\nu}_t^N, \beta_t^i\big)\,dt + \sigma\big(t, U_t^i, \bar{\nu}_t^N, \beta_t^i\big)\,dW_t^i, \\
&\qquad\qquad 1 \leq i \leq N, \text{ with } \bar{\nu}_t^N = \frac{1}{N}\sum_{j=1}^N \delta_{U_t^j},
\end{aligned}
\tag{52}
$$

with $t \in [0, T]$ and $U_0^i = x_0$, $1 \le i \le N$. Here $((\beta_t^i)_{0\le t\le T})_{1\le i\le N}$ are assumed to be progressively measurable with respect to the filtration generated by $(W^1, \dots, W^N)$. They take values in $\mathbb{R}^k$, and have finite L^2 norms over $[0, T] \times \Omega$

$$\forall i \in \{1, \dots, N\} \qquad \mathbb{E}\int_0^T |\beta_t^i|^2\, dt < +\infty,$$

where, for convenience, we have fixed an infinite sequence $((W_t^i)_{0\le t\le T})_{i\ge 1}$ of independent m-dimensional Brownian motions. One should think of U_t^i as the (private) state at time t of agent or player $i \in \{1, \dots, N\}$, β_t^i being the action taken at time t by player i. For each $1 \le i \le N$, we denote by

$$J^{N,i}(\beta^1, \dots, \beta^N) = \mathbb{E}\left[g\big(U_T^i, \bar{\nu}_T^N\big) + \int_0^T f\big(t, U_t^i, \bar{\nu}_t^N, \beta_t^i\big)\, dt\right] \tag{53}$$

the cost to the ith player. We frame the problem in the same set-up as in the case of the mean-field game models studied in [9], but the rule we apply for minimizing the cost is different. Indeed, we now minimize the cost over exchangeable strategies: when the family $(\beta^i, W^i)_{1\le i\le N}$ is exchangeable (with a slight abuse of terminology, we shall say that the strategy is exchangeable), the costs to all the players are the same. We use the notation $J^{N,i}(\underline{\beta}) = J^N(\underline{\beta})$ for this common cost. From a practical point of view, restricting the minimization to exchangeable strategies means that the players agree to use a common policy, which is not the case in the standard mean-field game approach.

Our first goal is to compute the limit

$$\lim_{N\to+\infty} \inf_{\underline{\beta}} J^N(\underline{\beta}),$$

the infimum being taken over exchangeable strategies. Another one is to identify, for each integer N, a specific set of ε-optimal strategies and the corresponding state evolutions.

6.1. *Limit of the costs and non-Markovian approximate equilibriums.* Recall that we denote by J the optimal cost,

$$J = \mathbb{E}\left[g(X_T, \mu_T) + \int_0^T f\big(t, X_t, \mu_t, \hat{\alpha}(t, X_t, \mu_t, Y_t, Z_t)\big)\, dt\right], \tag{54}$$

where $(X_t, Y_t, Z_t)_{0\le t\le T}$ is the solution to (38) with $X_0 = x_0$ as initial condition, $(\mu_t)_{0\le t\le T}$ denoting the flow of marginal probability measures $\mu_t = \mathbb{P}_{X_t}$, for $0 \le t \le T$.

For the purpose of comparison, we introduce $(\bar{X}^1, \dots, \bar{X}^N)$, each $\bar{X}^i$ standing for the solution of the forward equation in (38) when driven by the Brownian motion W^i. To put it differently, $(\bar{X}^1, \dots, \bar{X}^N)$ solves the system (52) when the empirical distribution $\bar{\nu}_t^N$ is replaced by μ_t, and β_t^i is given by $\beta_t^i = \bar{\alpha}_t^i$ with

$$\bar{\alpha}_t^i = \hat{\alpha}\big(t, \bar{X}_t^i, \mu_t, \bar{Y}_t^i, \bar{Z}_t^i\big),$$

the pair $(\bar{Y}^i, \bar{Z}^i)$ solving the backward equation in (38) when driven by W^i. Notice that the processes $((\bar{\Theta}_t^i = (\bar{X}_t^i, \mu_t, \bar{Y}_t^i, \bar{Z}_t^i, \bar{\alpha}_t^i))_{0\le t\le T})_{1\le i\le N}$ are independent. Our first result in this direction is the following.

THEOREM 6.1. *Under assumptions* (B1)–(B4),

$$\lim_{N\to+\infty} \inf_{\underline{\beta}} J^N(\underline{\beta}) = J,$$

the infimum being taken over (*square integrable*) *strategies* $\underline{\beta} = (\beta^1, \dots, \beta^N)$ *such that the family* $(\beta^i, W^i)_{1\le i\le N}$ *is exchangeable. Moreover, the non-Markovian control* $\underline{\bar{\alpha}} = (\bar{\alpha}^1, \dots, \bar{\alpha}^N)$ *is an approximate optimal control in the sense that*

$$\lim_{N\to+\infty} J^N(\underline{\bar{\alpha}}) = J.$$

PROOF. The proof consists in comparing $J^N(\underline{\beta})$ to J for a given exchangeable strategy $\underline{\beta}$. Once again, it relies on a variant of the Pontryagin stochastic maximum principle proven in Section 4. With the above notation, we have

$$\begin{aligned}
J^N(\underline{\beta}) - J &= \mathbb{E}\big[g\big(U_T^i, \bar{\nu}_T^N\big) - g\big(\bar{X}_T^i, \mu_T\big)\big] \\
&\quad + \mathbb{E}\biggl[\int_0^T \big(f\big(s, U_s^i, \bar{\nu}_s^N, \beta_s^i\big) - f\big(s, \bar{X}_s^i, \mu_s, \bar{\alpha}_s^i\big)\big)\,ds\biggr],
\end{aligned}$$

the identity holding for any $1 \le i \le N$. Therefore, we can write

$$J^N(\underline{\beta}) - J = T_1^i + T_2^i, \tag{55}$$

with

$$\begin{aligned}
T_1^i &= \mathbb{E}\big[\big(U_T^i - \bar{X}_T^i\big)\cdot \bar{Y}_T^i\big] + \mathbb{E}\biggl[\int_0^T \big(f\big(s, U_s^i, \bar{\nu}_s^N, \beta_s^i\big) - f\big(s, \bar{X}_s^i, \mu_s, \bar{\alpha}_s^i\big)\big)\,ds\biggr],\\
T_2^i &= \mathbb{E}\big[g\big(U_T^i, \bar{\nu}_T^N\big) - g\big(\bar{X}_T^i, \mu_T\big)\big] - \mathbb{E}\big[\big(U_T^i - \bar{X}_T^i\big)\cdot \partial_x g\big(\bar{X}_T^i, \mu_T\big)\big]\\
&\quad - \mathbb{E}\tilde{\mathbb{E}}\big[\big(\tilde{U}_T^i - \tilde{\bar{X}}_T^i\big)\cdot \partial_\mu g\big(\bar{X}_T^i, \mu_T\big)\big(\tilde{\bar{X}}_T^i\big)\big]\\
&= T_{2,1}^i - T_{2,2}^i - T_{2,3}^i,
\end{aligned}$$

where we used Fubini's theorem with the independent copies denoted with a tilde "$\tilde{\cdot}$".

Analysis of T_2^i. Using the diffusive effect of independence, we claim

$$\begin{aligned}
T_{2,3}^i &= \mathbb{E}\tilde{\mathbb{E}}\big[\big(\tilde{U}_T^i - \tilde{\bar{X}}_T^i\big)\cdot \partial_\mu g\big(\bar{X}_T^i, \mu_T\big)\big(\tilde{\bar{X}}_T^i\big)\big]\\
&= \frac{1}{N}\sum_{j=1}^N \tilde{\mathbb{E}}\big[\big(\tilde{U}_T^i - \tilde{\bar{X}}_T^i\big)\cdot \partial_\mu g\big(\tilde{\bar{X}}_T^j, \mu_T\big)\big(\tilde{\bar{X}}_T^i\big)\big]
\end{aligned}$$

$$+\mathcal{O}\biggl(\tilde{\mathbb{E}}\bigl[|\tilde{U}_T^i-\tilde{\bar{X}}_T^i|^2\bigr]^{1/2}\tilde{\mathbb{E}}\biggl[\biggl|\frac{1}{N}\sum_{j=1}^N\partial_\mu g\bigl(\tilde{\bar{X}}_T^j,\mu_T\bigr)\bigl(\tilde{\bar{X}}_T^i\bigr)$$

$$-\mathbb{E}\bigl[\partial_\mu g\bigl(\bar{X}_T^i,\mu_T\bigr)\bigl(\tilde{\bar{X}}_T^i\bigr)\bigr]\biggr|^2\biggr]^{1/2}\biggr)$$

$$=\frac{1}{N}\sum_{j=1}^N\mathbb{E}\bigl[\bigl(U_T^i-\bar{X}_T^i\bigr)\cdot\partial_\mu g\bigl(\bar{X}_T^j,\mu_T\bigr)\bigl(\bar{X}_T^i\bigr)\bigr]$$

$$+\mathbb{E}\bigl[|U_T^i-\bar{X}_T^i|^2\bigr]^{1/2}\mathcal{O}\bigl(N^{-1/2}\bigr),$$

where $\mathcal{O}(\cdot)$ stands for the Landau notation. Therefore, taking advantage of the exchangeability in order to handle the remainder, we obtain

$$\frac{1}{N}\sum_{i=1}^N T_{2,3}^i=\frac{1}{N^2}\sum_{j=1}^N\sum_{i=1}^N\mathbb{E}\bigl[\bigl(U_T^i-\bar{X}_T^i\bigr)\cdot\partial_\mu g\bigl(\bar{X}_T^j,\mu_T\bigr)\bigl(\bar{X}_T^i\bigr)\bigr]$$

$$+\mathbb{E}\bigl[|U_T^1-\bar{X}_T^1|^2\bigr]^{1/2}\mathcal{O}\bigl(N^{-1/2}\bigr).$$

Introducing a random variable ϑ from $(\tilde{\Omega},\tilde{\mathcal{F}},\tilde{\mathbb{P}})$ into $\mathbb{R}$ with uniform distribution on $\{1,\dots,N\}$ as in the proof of Proposition 3.2, we can write

$$\frac{1}{N}\sum_{i=1}^N T_{2,3}^i=\frac{1}{N}\sum_{j=1}^N\mathbb{E}\tilde{\mathbb{E}}\bigl[\bigl(U_T^\vartheta-\bar{X}_T^\vartheta\bigr)\cdot\partial_\mu g\bigl(\bar{X}_T^j,\mu_T\bigr)\bigl(\bar{X}_T^\vartheta\bigr)\bigr]$$

$$+\mathbb{E}\bigl[|U_T^1-\bar{X}_T^1|^2\bigr]^{1/2}\mathcal{O}\bigl(N^{-1/2}\bigr).$$

Finally, defining the flow of empirical measures

$$\bar{\mu}_t^N=\frac{1}{N}\sum_{j=1}^N\delta_{\bar{X}_t^j},\qquad t\in[0,T],$$

and using (B3), Propositions 3.2 and 3.4, and Remark 3.5 to estimate the distance $W_2(\bar{\mu}_T^N,\mu_T)$, the above estimate gives

$$\frac{1}{N}\sum_{i=1}^N T_{2,3}^i=\frac{1}{N}\sum_{j=1}^N\mathbb{E}\tilde{\mathbb{E}}\bigl[\bigl(U_T^\vartheta-\bar{X}_T^\vartheta\bigr)\cdot\partial_\mu g\bigl(\bar{X}_T^j,\bar{\mu}_T^N\bigr)\bigl(\bar{X}_T^\vartheta\bigr)\bigr]$$

$$+\mathbb{E}\bigl[|U_T^1-\bar{X}_T^1|^2\bigr]^{1/2}\mathcal{O}\bigl(\ell_N(d)\bigr),$$

where we used the notation $\ell_N(d)$ for any function of N which could be used as an upper bound for

$$\mathbb{E}\bigl[W_2^2\bigl(\bar{\mu}_T^N,\mu_T\bigr)\bigr]^{1/2}+\biggl(\int_0^T\mathbb{E}\bigl[W_2^2\bigl(\bar{\mu}_t^N,\mu_t\bigr)\bigr]\,dt\biggr)^{1/2}=\mathcal{O}\bigl(\ell_N(d)\bigr).\tag{56}$$

By Remark 3.5, the left-hand side tends to 0 as N tends to $+\infty$, since the function $[0,T] \ni t \mapsto \mathbb{E}[W_2^2(\bar{\mu}_t^N, \mu_t)]$ can be bounded independently of N. Therefore, $(\ell_N(d))_{N\geq 1}$ is always chosen as a sequence that converges to 0 as N tends to $+\infty$. When $\sup_{0\leq t\leq T} |\bar{X}_t^1|$ has finite moment of order $d+5$, Remark 3.5 says that $\ell_N(d)$ can be chosen as $N^{-1/(d+4)}$. In any case, we will assume that $\ell_N(d) \geq N^{-1/2}$. Going back to (55),

$$
\begin{aligned}
&\frac{1}{N}\sum_{i=1}^N T_2^i \\
&\quad = \frac{1}{N}\sum_{i=1}^N \{\mathbb{E}[g(U_T^i, \bar{\nu}_T^N) - g(\bar{X}_T^i, \bar{\mu}_T^N)] - \mathbb{E}[(U_T^i - \bar{X}_T^i)\cdot \partial_x g(\bar{X}_T^i, \bar{\mu}_T^N)] \\
&\qquad\qquad - \mathbb{E}\tilde{\mathbb{E}}[(U_T^\vartheta - \bar{X}_T^\vartheta)\cdot \partial_\mu g(\bar{X}_T^i, \bar{\mu}_T^N)(\bar{X}_T^\vartheta)]\} \\
&\qquad + (1 + \mathbb{E}[|U_T^1 - \bar{X}_T^1|^2]^{1/2})\mathcal{O}(\ell_N(d)),
\end{aligned}
$$

where we used the local Lipschitz property of g and Remark 3.5 to replace μ_T by $\bar{\mu}_T^N$.

Noticing that a.s. under $\mathbb{P}$, the law of U_T^ϑ (resp., $\bar{X}_T^\vartheta$) under $\tilde{\mathbb{P}}$ is the empirical distribution $\bar{\nu}_T^N$ (resp., $\bar{\mu}_T^N$), we can apply the convexity property of g [see (13)] to get

$$
\frac{1}{N}\sum_{i=1}^N T_2^i \geq (1 + \mathbb{E}[|U_T^1 - \bar{X}_T^1|^2]^{1/2})\mathcal{O}(\ell_N(d)). \tag{57}
$$

Analysis of T_1^i. Using Itô's formula and Fubini's theorem, we obtain

$$
\begin{aligned}
T_1^i &= \mathbb{E}\bigg[\int_0^T (H(s, U_s^i, \bar{\nu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \beta_s^i) - H(s, \bar{X}_s^i, \mu_s, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i))\,ds\bigg] \\
&\quad - \mathbb{E}\bigg[\int_0^T (U_s^i - \bar{X}_s^i)\cdot \partial_x H(s, \bar{X}_s^i, \mu_s, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)\,ds\bigg] \\
&\quad - \mathbb{E}\tilde{\mathbb{E}}\bigg[\int_0^T (\tilde{U}_s^i - \tilde{\bar{X}}_s^i)\cdot \partial_\mu H(s, \bar{X}_s^i, \mu_s, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)(\tilde{\bar{X}}_s^i)\,ds\bigg] \\
&= T_{1,1}^i - T_{1,2}^i - T_{1,3}^i.
\end{aligned} \tag{58}
$$

Using the regularity properties (B2) and (B3) of the Hamiltonian, (56), and recalling that the limit process $(\bar{X}_t^i, \mu_t, \bar{Y}_t^i, \bar{Z}_t^i, \bar{\alpha}_t^i)_{0\leq t\leq T}$ has finite $\mathbb{S}$-norm [see (39)], we get

$$
\begin{aligned}
T_{1,1}^i &= \mathbb{E}\bigg[\int_0^T (H(s, U_s^i, \bar{\nu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \beta_s^i) - H(s, \bar{X}_s^i, \bar{\mu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i))\,ds\bigg] \\
&\quad + \mathcal{O}(\ell_N(d)),
\end{aligned} \tag{59}
$$

$$T_{1,2}^i = \mathbb{E}\Big[\int_0^T (U_s^i - \bar{X}_s^i) \cdot \partial_x H(s, \bar{X}_s^i, \bar{\mu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)\, ds\Big]$$
$$+ \Big(\mathbb{E}\int_0^T |U_s^1 - \bar{X}_s^1|^2\, ds\Big)^{1/2} \mathcal{O}(\ell_N(d)).$$

Finally, using the diffusive effect of independence, we have

$$\begin{aligned}
&\frac{1}{N}\sum_{i=1}^N T_{1,3}^i \\
&\quad = \frac{1}{N}\sum_{i=1}^N \mathbb{E}\tilde{\mathbb{E}}\Big[\int_0^T (U_s^i - \bar{X}_s^i) \cdot \partial_\mu H(s, \tilde{\bar{X}}_s^i, \mu_s, \tilde{\bar{Y}}_s^i, \tilde{\bar{Z}}_s^i, \tilde{\bar{\alpha}}_s^i)(\bar{X}_s^i)\, ds\Big] \\
&\quad = \frac{1}{N^2}\sum_{j=1}^N\sum_{i=1}^N \mathbb{E}\Big[\int_0^T (U_s^i - \bar{X}_s^i) \cdot \partial_\mu H(s, \bar{X}_s^j, \mu_s, \bar{Y}_s^j, \bar{Z}_s^j, \bar{\alpha}_s^j)(\bar{X}_s^i)\, ds\Big] \\
&\qquad + \Big(\mathbb{E}\int_0^T |U_s^1 - \bar{X}_s^1|^2\, ds\Big)^{1/2} \mathcal{O}(N^{-1/2}).
\end{aligned}$$

By (B3), Propositions 3.2 and 3.4, we have

$$\begin{aligned}
&\frac{1}{N}\sum_{i=1}^N T_{1,3}^i \\
&\quad = \frac{1}{N}\sum_{i=1}^N \mathbb{E}\tilde{\mathbb{E}}\Big[\int_0^T (U_s^\vartheta - \bar{X}_s^\vartheta) \cdot \partial_\mu H(s, \bar{X}_s^i, \mu_s, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)(\bar{X}_s^\vartheta)\, ds\Big] \\
&\qquad + \Big(\mathbb{E}\int_0^T |U_s^1 - \bar{X}_s^1|^2\, ds\Big)^{1/2} \mathcal{O}(N^{-1/2}) \\
&\quad = \frac{1}{N}\sum_{i=1}^N \mathbb{E}\tilde{\mathbb{E}}\Big[\int_0^T (U_s^\vartheta - \bar{X}_s^\vartheta) \cdot \partial_\mu H(s, \bar{X}_s^i, \bar{\mu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)(\bar{X}_s^\vartheta)\, ds\Big] \\
&\qquad + \Big(\mathbb{E}\int_0^T |U_s^1 - \bar{X}_s^1|^2\, ds\Big)^{1/2} \mathcal{O}(\ell_N(d)).
\end{aligned} \tag{60}$$

In order to complete the proof, we evaluate the missing term in the Taylor expansion of T_1^i in (58), namely

$$\frac{1}{N}\sum_{i=1}^N \mathbb{E}\Big[\int_0^T (\beta_s^i - \bar{\alpha}_s^i) \cdot \partial_\alpha H(s, \bar{X}_s^i, \bar{\mu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)\, ds\Big],$$

in order to benefit from the convexity of H. We use Remark 3.5 once more,

$$
\begin{aligned}
\mathbb{E}\Big[\int_0^T (\beta_s^i - \bar{\alpha}_s^i) \cdot \partial_\alpha H(s, \bar{X}_s^i, \bar{\mu}_s^N, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)\, ds\Big] \\
&= \mathbb{E}\Big[\int_0^T (\beta_s^i - \bar{\alpha}_s^i) \cdot \partial_\alpha H(s, \bar{X}_s^i, \mu_s, \bar{Y}_s^i, \bar{Z}_s^i, \bar{\alpha}_s^i)\, ds\Big] \\
&\quad + \Big(\mathbb{E}\int_0^T |\beta_s^i - \bar{\alpha}_s^i|^2\, ds\Big)^{1/2} \mathcal{O}(\ell_N(d)) \\
&= \Big(\mathbb{E}\int_0^T |\beta_s^i - \bar{\alpha}_s^i|^2\, ds\Big)^{1/2} \mathcal{O}(\ell_N(d)),
\end{aligned}
\tag{61}
$$

since $\bar{\alpha}$ is an optimizer for H. Using the convexity of H and taking advantage of the exchangeability, we finally deduce from (58), (59), (60) and (61) that there exists a constant $c > 0$ such that

$$
\begin{aligned}
\frac{1}{N}\sum_{i=1}^N T_1^i &\geq c\mathbb{E}\int_0^T |\beta_s^1 - \bar{\alpha}_s^1|^2\, ds \\
&\quad + \mathcal{O}(\ell_N(d))\Big(1 + \sup_{0\leq t\leq T} \mathbb{E}\big[|U_t^1 - \bar{X}_t^1|^2\big] + \mathbb{E}\int_0^T |\beta_s^1 - \bar{\alpha}_s^1|^2\, ds\Big)^{1/2}.
\end{aligned}
$$

By (57) and (55), we deduce that

$$
\begin{aligned}
J^N(\underline{\beta}) &\geq J + c\mathbb{E}\int_0^T |\beta_s^1 - \bar{\alpha}_s^1|^2\, ds \\
&\quad + \mathcal{O}(\ell_N(d))\Big(1 + \sup_{0\leq t\leq T} \mathbb{E}\big[|U_t^1 - \bar{X}_t^1|^2\big] + \mathbb{E}\int_0^T |\beta_s^1 - \bar{\alpha}_s^1|^2\, ds\Big)^{1/2}.
\end{aligned}
$$

From the inequality

$$
\sup_{0\leq t\leq T} \mathbb{E}\big[|U_t^1 - \bar{X}_t^1|^2\big] \leq C\mathbb{E}\int_0^T |\beta_s^1 - \bar{\alpha}_s^1|^2\, ds,
$$

which holds for some constant C independent of N, we deduce that

$$
J^N(\underline{\beta}) \geq J - C\ell_N(d), \tag{62}
$$

for a possibly new value of C. This proves that

$$
\liminf_{N\to+\infty} \inf_{\underline{\beta}} J^N(\underline{\beta}) \geq J.
$$

In order to prove Theorem 6.1, it only remains to find a sequence of controls $(\underline{\beta}^N)_{N\geq 1}$ such that

$$
\limsup_{N\to+\infty} J^N(\underline{\beta}^N) \leq J.
$$

More precisely, we are about to show that

$$\limsup_{N\to+\infty} J^N(\underline{\bar{\alpha}}) \le J,$$

thus proving that $\underline{\bar{\alpha}} = (\bar{\alpha}^1, \dots, \bar{\alpha}^N)$ is an approximate equilibrium, though non-Markovian. Denoting by $(X^1, \dots, X^N)$ the solution of (52) with $\beta_t^i = \bar{\alpha}_t^i$, classical estimates from the theory of propagation of chaos (see, e.g., [23] or [13]) imply that

$$\sup_{0\le t\le T} \mathbb{E}\big[\big|X_t^i - \bar{X}_t^i\big|^2\big] = \sup_{0\le t\le T} \mathbb{E}\big[\big|X_t^1 - \bar{X}_t^1\big|^2\big] = \mathcal{O}(N^{-1}).$$

It is then plain to deduce that

$$\limsup_{N\to+\infty} J^N(\underline{\bar{\alpha}}) \le J.$$

This completes the proof. $\square$

6.2. *Approximate equilibriums with distributed closed loop controls.* When σ does not depend upon α, we are able to provide an approximate equilibrium using only distributed controls in closed loop form. This is of real interest from the practical point of view. Indeed, in a such case, the optimizer $\hat{\alpha}$ of the Hamiltonian, as defined in (35), does not depend on z. It thus reads as $\hat{\alpha}(t, x, \mu, y)$. By Proposition 5.7, this says that the optimal control $(\alpha_t)_{0\le t\le T}$ in Theorem 5.1 has the *feedback* form

$$\alpha_t = \hat{\alpha}\big(t, X_t, \mu_t, v(t, X_t)\big), \qquad t \in [0, T]. \tag{63}$$

The reader may wonder why we make this assumption on σ. Indeed, when σ depends upon α, the process Z_t at time t is also expected to read as a function of t and X_t, since such a representation is known to hold in the classical *decoupled* forward–backward setting. Even if we feel that it is indeed possible to prove such a representation in our more general setting, we refrain from doing it here for the following reasons: (i) from a practical point of view, for equation (63) to be meaningful, one would want the feedback function to be Lipschitz-continuous, as the Lipschitz property ensures that the stochastic differential equation obtained by plugging (63) into the forward equation in (38) is solvable; (ii) in the current framework, the function v is known to be Lipschitz continuous by Proposition 5.7, but proving the same result for the representation of Z_t in terms of X_t seems to be really challenging—notice that it is already challenging in the standard case, that is, without any McKean–Vlasov interaction; (iii) we finally mention that, in any case, the relationship between Z_t and X_t, if it exists, must be rather intricate as Z_t is expected to solve the equation $Z_t = \partial_x v(t, X_t)\sigma(t, X_t, \mathbb{P}_{X_t}, \hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t))$, which can be formally derived by identifying martingale integrands when expanding $Y_t = v(t, X_t)$ by a formal application of Itô's formula. This equation has been investigated in [25] in the standard case, but we are convinced that extending this analysis to the current setting, though possible, would be very technical and lengthy, and distract us from the main thrust of this section.

Now, for each integer N, we can consider the solution $(X_t^1, \dots, X_t^N)_{0 \le t \le T}$ of the system of N stochastic differential equations

$$dX_t^i = b(t, X_t^i, \mu_t^N, \hat{\alpha}(t, X_t^i, \mu_t, v(t, X_t^i)))\, dt + \sigma(t, X_t^i, \mu_t^N)\, dW_t^i \quad \text{with } \mu_t^N = \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}, \tag{64}$$

with $t \in [0, T]$ and $X_0^i = x_0$. The system (64) is well posed since v satisfies Proposition 5.7, and the minimizer $\hat{\alpha}(t, x, \mu_t, y)$ is Lipschitz continuous and at most of linear growth in the variables x, μ and y, uniformly in $t \in [0, T]$. The processes $(X^i)_{1 \le i \le N}$ give the dynamics of the private states of the N players in the stochastic differential game of interest when the players use the strategies

$$\alpha_t^{N,i} = \hat{\alpha}(t, X_t^i, \mu_t, v(t, X_t^i)), \qquad 0 \le t \le T, i \in \{1, \dots, N\}. \tag{65}$$

These strategies are in closed loop form. They are even *distributed* since, at each time $t \in [0, T]$, a player only needs to know his own private state in order to compute the value of the action to take at that time. By the linear growth of v and of the minimizer $\hat{\alpha}$, it holds, for any $p \ge 2$,

$$\sup_{N \ge 1} \max_{1 \le i \le N} \mathbb{E}\Big[\sup_{0 \le t \le T} |X_t^i|^p\Big] < +\infty, \tag{66}$$

the expectation being actually independent of i since the strategy is obviously exchangeable. We then have the following approximate equilibrium property:

THEOREM 6.2. *In addition to assumptions* (B1)–(B4), *assume that σ does not depend upon α. Then*

$$J^N(\underline{\beta}) \ge J^N(\underline{\alpha}^N) - \mathcal{O}(N^{-1/(d+4)}),$$

for any $\underline{\beta} = (\beta^1, \dots, \beta^N)$ such that $(\beta^i, W^i)_{1 \le i \le N}$ is exchangeable, *where $\underline{\alpha}^N$ is defined in* (65).

PROOF. We use the same notation as in the proof of Theorem 6.1.

Since $\bar{\alpha}_t^1$ now reads as $\hat{\alpha}(t, \bar{X}_t^1, \mu_t, v(t, \bar{X}_t^1))$ for $0 \le t \le T$, we first notice, by the growth property of v, that $\mathbb{E}[\sup_{0 \le t \le T} |\bar{X}_t^1|^p] < +\infty$ for any $p \ge 1$. As mentioned in (11) in Remark 3.5, this says that $\ell_N(d)$ in the lower bound

$$J^N(\underline{\beta}) \ge J - C\ell_N(d),$$

[see (62)] can be chosen as $N^{-1/(d+4)}$.

Moreover, since $v(t,\cdot)$ is Lipschitz continuous, using once again classical estimates from the theory of propagation of chaos (see, e.g., [23] or [13]), we also have

$$\sup_{0\le t\le T} \mathbb{E}\big[|X_t^i - \bar{X}_t^i|^2\big] = \sup_{0\le t\le T} \mathbb{E}\big[|X_t^1 - \bar{X}_t^1|^2\big] = \mathcal{O}(N^{-1}),$$

so that

$$\sup_{0\le t\le T} \mathbb{E}\big[|\alpha_t^{N,i} - \bar{\alpha}_t^i|^2\big] = \sup_{0\le t\le T} \mathbb{E}\big[|\alpha_t^{N,1} - \bar{\alpha}_t^1|^2\big] = \mathcal{O}(N^{-1}),$$

for any $1 \le i \le N$. It is then plain to deduce that

$$J^N(\underline{\alpha}^N) \le J + C\ell_N(d).$$

This completes the proof. $\square$

APPENDIX

A.1. Proof of Lemma 3.1. *First step.* The proof is based on the fact that, for a bounded continuous function $\ell : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}^d$, the mapping $L^2(\tilde{\Omega};\mathbb{R}^d) \times \mathbb{R}^d \ni (\tilde{X},x) \mapsto \ell(\tilde{X},x) \in L^2(\tilde{\Omega};\mathbb{R}^d)$ is continuous and thus measurable, where $L^2(\tilde{\Omega};\mathbb{R}^d)$ is the quotient of the space of square-integrable random variables by the $\tilde{\mathbb{P}}$ almost sure equality and $\ell(\tilde{X},x)$ is an abuse of notation for denoting the class (in $L^2(\tilde{\Omega};\mathbb{R}^d)$) of $\ell(\tilde{\chi},x)$, with $\tilde{\chi}$ a random variable matching $\tilde{X}$ almost surely. Measurability is preserved by replacing ℓ by $\mathbf{1}_I \circ \ell$ for an interval I, as $\mathbf{1}_I$ can be written as the pointwise limit of continuous functions.

Here is the way we apply this simple remark. Denoting by "$\cdot$" the inner product in $L^2(\tilde{\Omega};\mathbb{R}^d)$, the mapping $[L^2(\tilde{\Omega};\mathbb{R}^d)]^2 \ni (\tilde{X},\tilde{Y}) \mapsto [D\tilde{H}](\tilde{X})\cdot\tilde{Y}$ is measurable as the pointwise limit of measurable mappings. Therefore, for any vector $e \in \mathbb{R}^d$ and any $\varepsilon > 0$, the mapping $L^2(\tilde{\Omega};\mathbb{R}^d) \times \mathbb{R}^d \ni (\tilde{X},x) \mapsto [D\tilde{H}](\tilde{X}) \cdot (e\mathbf{1}_{\{|\tilde{X}-x|\le\varepsilon\}})$ is jointly measurable. Then the mapping $\psi = (\psi_1,\dots,\psi_d) : L^2(\tilde{\Omega};\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}^d$ given by

$$\psi_i(\tilde{X},x) = \liminf_{\varepsilon \searrow 0} \left[\frac{[D\tilde{H}](\tilde{X}) \cdot (e_i \mathbf{1}_{\{|\tilde{X}-x|\le\varepsilon\}})}{\tilde{\mathbb{P}}(|\tilde{X}-x|\le\varepsilon)} \mathbf{1}_{\{\tilde{\mathbb{P}}(|\tilde{X}-x|\le\varepsilon)>0\}} \right]$$

is also jointly measurable, where $(e_1,\dots,e_d)$ is the canonical basis. By Lebesgue–Besicovitch differentiation theorem, $\psi(\tilde{X},\cdot)$ is a version of $\partial_\mu H(\tilde{\mathbb{P}}_{\tilde{X}})(\cdot)$ in $L^2(\mathbb{R}^d, \mathbb{P}_{\tilde{X}})$. If $\partial_\mu H(\tilde{\mathbb{P}}_{\tilde{X}})(\cdot)$ admits a continuous version, then it coincides with it.

Second step. In order to complete the proof, we prove that we can find a measurable mapping $\mathcal{P}_2(\mathbb{R}^d) \ni \mu \mapsto \tilde{X}^\mu \in L^2(\tilde{\Omega};\mathbb{R}^d)$ with $\tilde{X}^\mu \sim \mu$. By chain rule, this will show that the mapping $\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (\mu,x) \mapsto \psi(\tilde{X}^\mu,x)$ is measurable. We first consider the case $\mu([0,1)^d) = 1$. Given some $n \ge 0$, we split the hypercube $[0,1)^d$ into $(2^n)^d$ hypercubes of the form $Q^n(k_1,\dots,k_d) =$

$\prod_{i=1}^{d}[k_i/2^n,(k_i+1)/2^n)$, with $(k_1,\dots,k_d)\in(\mathbb{Z}\cap[0,2^n-1])^d$. For any d-tuple $(k_1,\dots,k_d)$, we let $M^{n,\mu}(k_1,\dots,k_d)=\mu(Q^n(k_1,\dots,k_d))$.

The strategy is to arrange the cylinders $Q^n(k_1,\dots,k_d)$ increasingly according to some order. To this end, we observe that, for any $1\le i\le d$, $k_i/2^n$ may be uniquely written as

$$\frac{k_i}{2^n}=\sum_{j=1}^{n}\frac{\varepsilon_j(k_i)}{2^j}, \tag{67}$$

with $\varepsilon_j(k_i)\in\{0,1\}$. Given $(k_1,\dots,k_d)$ and $(k_1',\dots,k_d')$ in $(\mathbb{Z}\cap[0,2^n-1])^d$, with $(k_1,\dots,k_d)\neq(k_1',\dots,k_d')$, we say that $(k_1,\dots,k_d)\prec(k_1',\dots,k_d')$ if, letting

$$p=\inf\{j\in\{1,\dots,n\}:(\varepsilon_j(k_1),\dots,\varepsilon_j(k_d))\neq(\varepsilon_j(k_1'),\dots,\varepsilon_j(k_d'))\},$$
$$q=\inf\{i\in\{1,\dots,d\}:\varepsilon_p(k_i)\neq\varepsilon_p(k_i')\},$$

it holds $0=\varepsilon_p(k_q)<\varepsilon_p(k_q')=1$. In other words, the order is defined by taking into account first the index j in (67) and then the coordinate i.

Then we can divide the interval $[0,1)$ into a family $(I^{n,\mu}(k_1,\dots,k_d))$, with $(k_1,\dots,k_d)\in(\mathbb{Z}\cap[0,2^n-1])^d$, of $(2^n)^d$ disjoint (possibly empty) intervals, closed at the left end and open at the right end, of length $M^{n,\mu}(k_1,\dots,k_d)$ each, and ordered increasingly according to $\prec$. This means that, for any $x\in I^{n,\mu}(k_1,\dots,k_d)$ and $x'\in I^{n,\mu}(k_1',\dots,k_d')$, $x<x'$ if $(k_1,\dots,k_d)\prec(k_1',\dots,k_d')$. Then we let

$$\tilde{X}^{n,\mu}=\sum_{(k_1,\dots,k_d)\in(\mathbb{Z}\cap[0,2^n-1])^d}(2^{-n}k_1,\dots,2^{-n}k_d)\mathbf{1}_{\{\tilde{\eta}\in I^{n,\mu}(k_1,\dots,k_d)\}},$$

where $\tilde{\eta}:\tilde{\Omega}\to(0,1)$ is uniformly distributed. It is then plain to check that the mapping $\mu\mapsto\tilde{X}^{n,\mu}$ is measurable from the Borel subset of $\mathcal{P}_2(\mathbb{R}^d)$ made of probability measures μ satisfying $\mu([0,1)^d)=1$ to $L^2(\tilde{\Omega};\mathbb{R}^d)$. Indeed, writing $I^{n,\mu}(k_1,\dots,k_d)$ as $[a^{n,\mu}(k_1,\dots,k_d),b^{n,\mu}(k_1,\dots,k_d))$, the mapping $\mu\mapsto(a^{n,\mu}(k_1,\dots,k_d),b^{n,\mu}(k_1,\dots,k_d))$ is measurable [the mapping $\mu\mapsto M^{n,\mu}(k_1,\dots,k_d)$ is measurable, and the intervals are constructed in a canonical way from their lengths]. Then, as explained in the beginning of the proof, the functions $\mu\mapsto\mathbf{1}_{\{\tilde{\eta}\ge a^{n,\mu}(k_1,\dots,k_d)\}}$ and $\mu\mapsto\mathbf{1}_{\{\tilde{\eta}<b^{n,\mu}(k_1,\dots,k_d)\}}$ can be proved to be measurable by approximating the indicator functions by continuous functions.

We then notice that, for any bounded and continuous function ℓ,

$$\begin{aligned}&\tilde{\mathbb{E}}[\ell(\tilde{X}^{n,\mu})]\\&\quad=\sum_{(k_1,\dots,k_d)\in(\mathbb{Z}\cap[0,2^n-1])^d}\ell\Big(\frac{k_1}{2^n},\dots,\frac{k_d}{2^n}\Big)\mu(Q^n(k_1,\dots,k_d))\\&\quad\to\int_{\mathbb{R}^d}\ell(x)\,d\mu(x),\end{aligned}$$

proving that $\tilde{X}^{n,\mu}$ converges in law to μ as n tends to $+\infty$. Moreover, because of our choice of ordering, we have

$$I^{n,\mu}(k_1,\dots,k_d) = \bigcup_{(\varepsilon_i)_{i=1,\dots,d}\in\{0,1\}^d} I^{n+1,\mu}(2k_1+\varepsilon_1,\dots,2k_d+\varepsilon_d).$$

The reason is that, for any other $(k'_1,\dots,k'_d) \prec (k_1,\dots,k_d)$ in $\{0,2^n-1\}^d$, it must hold $I^{n+1,\mu}(2k'_1+\varepsilon'_1,\dots,2k'_d+\varepsilon'_d) \prec I^{n+1,\mu}(2k_1+\varepsilon_1,\dots,2k_d+\varepsilon_d)$ for any $\varepsilon_1,\varepsilon'_1,\dots,\varepsilon_d,\varepsilon'_d \in \{0,1\}$ (and the same with $\prec$ replaced by $\succ$). As a by-product, there exists a constant C, independent of n and μ, such that

$$\tilde{\mathbb{E}}\big[\big|\tilde{X}^{n,\mu} - \tilde{X}^{n+1,\mu}\big|^2\big]^{1/2} \le \frac{C}{2^n}.$$

This proves $(\tilde{X}^{n,\mu})_{n\ge1}$ is Cauchy in $L^2(\tilde{\Omega};\mathbb{R}^d)$. The limit is denoted by $\tilde{X}^{\infty,\mu}$. Therefore, the mapping $\mu \mapsto \tilde{X}^{\infty,\mu}$ is measurable on the set of probability measures μ satisfying $\mu([0,1)^d)=1$.

When the support of μ is general, we define $\phi\sharp\mu$ as the push-forward (or image) of μ by the mapping

$$\begin{aligned}&\phi(x_1,\dots,x_d)\\ &\quad= \left(\frac{1}{\pi}\arctan(x_1)+\frac{1}{2},\dots,\frac{1}{\pi}\arctan(x_d)+\frac{1}{2}\right), \qquad x_1,\dots,x_d\in\mathbb{R}^d.\end{aligned}$$

We then let $\tilde{X}^\mu = (\tilde{X}_i^\mu)_{1\le i\le d}$, with $\tilde{X}_i^\mu = \tan(\pi\tilde{X}_i^{\infty,\phi\sharp\mu} - \pi/2)$, for $i\in\{1,\dots,d\}$. Clearly, $\tilde{X}^\mu$ has μ as distribution. Considering a sequence of bounded continuous functions $(\zeta_n)_{n\ge1}$ converging to the identity, $\tilde{X}^\mu$ is the limit of $(\zeta_n(\tan(\pi\tilde{X}_i^{\infty,\phi\sharp\mu} - \pi/2)))_{1\le i\le d}$, which is measurable with respect to μ. This proves that $\mu\mapsto\tilde{X}^\mu$ is measurable.

A.2. Proof of Lemma 3.3. *First step.* We first consider the case v bounded and assume that μ has a strictly positive continuous density p on the whole $\mathbb{R}^d$, p and its derivatives being of exponential decay at infinity. We claim that there exists a continuously differentiable one-to-one function from $(0,1)^d$ onto $\mathbb{R}^d$ such that, whenever $\eta_1,\dots,\eta_d$ are d independent random variables uniformly distributed on $(0,1)$, $U(\eta_1,\dots,\eta_d)$ has distribution μ. It satisfies for any $(z_1,\dots,z_d)\in(0,1)^d$

$$\frac{\partial U_i}{\partial z_i}(z_1,\dots,z_d)\neq 0, \qquad \frac{\partial U_j}{\partial z_i}(z_1,\dots,z_d)=0, \qquad 1\le i<j\le d.$$

The result is well known when $d=1$. In such a case, U is the inverse of the cumulative distribution function of μ. In higher dimension, U can be constructed by an induction argument on the dimension. Assume indeed that some $\hat{U}$ has been constructed for the first marginal distribution $\hat{\mu}$ of μ on $\mathbb{R}^{d-1}$, that is, for the push-forward of μ by the projection mapping $\mathbb{R}^d \ni (x_1,\dots,x_d)\mapsto(x_1,\dots,x_{d-1})$.

Given $(x_1, \dots, x_{d-1}) \in \mathbb{R}^{d-1}$, we then denote by $p(\cdot|x_1, \dots, x_{d-1})$ the conditional density of μ given the $d-1$ first coordinates

$$p(x_d|x_1, \dots, x_{d-1}) = \frac{p(x_1, \dots, x_d)}{\hat{p}(x_1, \dots, x_{d-1})}, \qquad x_1, \dots, x_{d-1} \in \mathbb{R}^{d-1},$$

where $\hat{p}$ denotes the density of $\hat{\mu}$ (which is continuously differentiable and positive). We then denote by $(0,1) \ni z_d \mapsto U^{(d)}(z_d|x_1, \dots, x_{d-1})$ the inverse of the cumulative distribution function of the law of density $p(\cdot|x_1, \dots, x_{d-1})$. It satisfies

$$F_d\big(U^{(d)}(z_d|x_1, \dots, x_{d-1})|x_1, \dots, x_{d-1}\big) = z_d,$$

with

$$F_d(x_d|x_1, \dots, x_{d-1}) = \int_{-\infty}^{x_d} p(y|x_1, \dots, x_{d-1})\, dy,$$

which is continuously differentiable in $(x_1, \dots, x_d)$ (using the exponential decay of the density at infinity). By the implicit function theorem, the mapping $\mathbb{R}^{d-1} \times (0,1) \ni (x_1, \dots, x_{d-1}, z_d) \mapsto U^{(d)}(z_d|x_1, \dots, x_{d-1})$ is continuously differentiable. The partial derivative with respect to z_d is given by

$$\frac{\partial U^{(d)}}{\partial z_d}(z_d|x_1, \dots, x_{d-1}) = \frac{1}{p(U^{(d)}(z_d|x_1, \dots, x_{d-1})|x_1, \dots, x_{d-1})},$$

which is nonzero. We now let, for $(z_1, \dots, z_d) \in (0,1)^d$,

$$U(z_1, \dots, z_d) = \big(\hat{U}(z_1, \dots, z_{d-1}), U^{(d)}\big(z_d|\hat{U}(z_1, \dots, z_{d-1})\big)\big).$$

By construction, $U(\eta_1, \dots, \eta_d)$ has distribution μ: $\hat{U}(\eta_1, \dots, \eta_{d-1})$ has distribution $\hat{\mu}$, and the conditional law of $U_d(\eta_1, \dots, \eta_d)$ given $\eta_1, \dots, \eta_{d-1}$ is the conditional law of μ given the $d-1$ first coordinates, since $U_d(\eta_1, \dots, \eta_d) = U^{(d)}(\eta_d|\hat{U}(\eta_1, \dots, \eta_{d-1}))$. It satisfies $[\partial U_d/\partial z_d](z_1, \dots, z_d) > 0$ and, for $i < d$, $[\partial U_i/\partial z_d](z_1, \dots, z_d) = 0$. In particular, since $\hat{U}$ is assumed (by induction) to be one-to-one and $[\partial U_d/\partial z_d](z_1, \dots, z_d) > 0$, U must be one-to-one as well. As the Jacobian matrix of U is triangular with nonzero elements on the diagonal, it is invertible. By the global inversion theorem, U is a diffeomorphism: the range of U is the support of μ, that is, $\mathbb{R}^d$. This proves that U is one-to-one from $(0,1)^d$ onto $\mathbb{R}^d$.

Second step. We still consider the case v bounded and assume that μ has a strictly positive continuous density p on the whole $\mathbb{R}^d$, p and its derivatives being of exponential decay at infinity. We will use the mapping U constructed in the first step. For three random variables ξ, ξ' and G in $L^2(\Omega; \mathbb{R}^d)$, the pair (ξ, ξ') being independent of G, the random variables ξ and ξ' having the same distribution and G being normally distributed with mean 0 and covariance matrix given by the identity I_d in dimension d, in notation $G \sim \mathcal{N}_d(0, I_d)$, then (8) implies that, for any integer $n \geq 1$,

$$\mathbb{E}\big[\big|v(\xi + n^{-1}G, \mathbb{P}_{\xi+n^{-1}G}) - v(\xi' + n^{-1}G, \mathbb{P}_{\xi+n^{-1}G})\big|^2\big] \leq C^2 \mathbb{E}\big[|\xi - \xi'|^2\big].$$

In particular, setting

$$
\begin{aligned}
v_n(x) &= \mathbb{E}\big[v\big(x+n^{-1}G, \mathbb{P}_{\xi+n^{-1}G}\big)\big] \\
&= \frac{n^d}{(2\pi)^{d/2}} \int_{\mathbb{R}^d} v(y, \mathbb{P}_{\xi+n^{-1}G}) \exp\Big(-n^2 \frac{|x-y|^2}{2}\Big)\, dy,
\end{aligned}
$$

we have

$$
\mathbb{E}\big[\big|v_n(\xi) - v_n(\xi')\big|^2\big] \le C^2 \mathbb{E}\big[|\xi - \xi'|^2\big]. \tag{68}
$$

Notice that v_n is infinitely differentiable with bounded derivatives.

We now choose a specific coupling for ξ and ξ'. Indeed, we know that for any $\eta = (\eta_1, \dots, \eta_d)$ and $\eta' = (\eta'_1, \dots, \eta'_d)$, with uniform distributions on $(0,1)^d$, $U(\eta)$ and $U(\eta')$ have the same distribution as ξ. Without any loss of generality, we then assume that the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is given by $(0,1)^d \times \mathbb{R}^d$ endowed with its Borel σ-algebra and the product of the Lebesgue measure on $(0,1)^d$ and of the Gaussian measure $\mathcal{N}_d(0, I_d)$. The random variables η and G are then chosen as the canonical mappings $\eta : (0,1)^d \times \mathbb{R}^d \ni (z,y) \mapsto z$ and $G : (0,1)^d \times \mathbb{R}^d \ni (z,y) \mapsto y$.

We then define η' as a function of the variable $z \in (0,1)^d$ only. For a given $z^0 = (z_1^0, \dots, z_d^0) \in (0,1)^d$ and for h small enough so that the open ball $B(z^0, h)$ of center z^0 and radius h is included in $(0,1)^d$, we let

$$
\eta'(z) = \begin{cases} z - 2(z_d - z_d^0) e_d, & \text{if } z \in B(z^0, h), \\ z, & \text{outside,} \end{cases}
$$

where e_d is the dth vector of the canonical basis, that is, η' matches locally the symmetry with respect to the hyperplane containing z^0 and orthogonal to e_d. Clearly, η' preserves the Lebesgue measure. We rewrite (68) as

$$
\begin{aligned}
&\int_{(0,1)^d} \big|v_n\big(U\big(\eta(z)\big)\big) - v_n\big(U\big(\eta'(z)\big)\big)\big|^2\, dz \\
&\qquad \le C^2 \int_{(0,1)^d} \big|U\big(\eta(z)\big) - U\big(\eta'(z)\big)\big|^2\, dz,
\end{aligned}
$$

or equivalently

$$
\begin{aligned}
&\int_{|r|<h} \big|v_n\big[U\big(z^0 + r - 2r_d e_d\big)\big] - v_n\big(U\big(z^0 + r\big)\big)\big|^2\, dr \\
&\qquad \le C^2 \int_{|r|<h} \big|U\big(z^0 + r - 2r_d e_d\big) - U\big(z^0 + r\big)\big|^2\, dr.
\end{aligned} \tag{69}
$$

Since U is continuously differentiable, we have

$$
v_n\big(U\big(z^0 + r\big)\big) = v_n\big(U\big(z^0\big)\big) + \partial v_n\big(U\big(z^0\big)\big) \cdot \big[\partial U\big(z^0\big) \cdot r\big] + o(r),
$$

where $\partial U(z^0)$ is a $d \times d$ matrix. We deduce that

$$\begin{aligned} v_n\big[U\big(z^0 + r - 2r_d e_d\big)\big] - v_n\big(U\big(z^0 + r\big)\big) \\ &= -2\sum_{i=1}^d \frac{\partial v_n}{\partial x_i}\big(U\big(z^0\big)\big)\frac{\partial U_i}{\partial z_d}\big(z^0\big) r_d + o(r) \\ &= -2\frac{\partial v_n}{\partial x_d}\big(U\big(z^0\big)\big)\frac{\partial U_d}{\partial z_d}\big(z^0\big) r_d + o(r), \end{aligned}$$

since $\partial U_i/\partial z_d = 0$ for $i \neq d$, and

$$\begin{aligned} &\int_{|r|<h} \big|v_n\big[U\big(z^0 + r - 2r_d e_d\big)\big] - v_n\big(U\big(z^0 - r\big)\big)\big|^2\, dr \\ &\quad = 4\left|\frac{\partial v_n}{\partial x_d}\big(U\big(z^0\big)\big)\frac{\partial U_d}{\partial z_d}\big(z^0\big)\right|^2 \int_{|r|<h} r_d^2\, dr + o\big(h^{d+2}\big). \end{aligned} \tag{70}$$

Similarly,

$$\begin{aligned} &\int_{|r|<h} \big|U\big(z^0 + r - 2r_d e_d\big) - U\big(z^0 + r\big)\big|^2\, dr \\ &\quad = 4\left|\frac{\partial U_d}{\partial z_d}\big(z^0\big)\right|^2 \int_{|r|<h} r_d^2\, dr + o\big(h^{d+2}\big), \end{aligned} \tag{71}$$

and putting together (69), (70) and (71), we obtain

$$\left|\frac{\partial v_n}{\partial x_d}\big(U\big(z^0\big)\big)\frac{\partial U_d}{\partial z_d}\big(z^0\big)\right|^2 \le C^2 \left|\frac{\partial U_d}{\partial z_d}\big(z^0\big)\right|^2.$$

Since $[\partial U_d/\partial z_d](z^0)$ is different from zero, we deduce that

$$\left|\frac{\partial v_n}{\partial x_d}\big(U\big(z^0\big)\big)\right|^2 \le C^2,$$

and since U is a one-to-one mapping from $(0, 1)^d$ onto $\mathbb{R}^d$, and $z^0 \in (0, 1)^d$ is arbitrary, we conclude that $|[\partial v_n/\partial x_d](x)| \le C$, for any $x \in \mathbb{R}^d$. By changing the basis used for the construction of U (we used the canonical basis but we could use any orthonormal basis), we have $|\nabla v_n(x)e| \le C$ for any $x, e \in \mathbb{R}^d$ with $|e| = 1$. This proves that the functions $(v_n)_{n\ge1}$ are uniformly bounded and C-Lipschitz continuous. We then denote by $\hat{v}$ the limit of a subsequence converging for the topology of uniform convergence on compact subsets. For simplicity, we keep the index n to denote the subsequence. Assumption (8) implies

$$\mathbb{E}\big[\big|v_n(\xi) - v(\xi, \mathbb{P}_\xi)\big|^2\big] \le \mathbb{E}\big[\big|v\big(\xi + n^{-1}G, \mathbb{P}_{\xi+n^{-1}G}\big) - v(\xi, \mathbb{P}_\xi)\big|^2\big] \le C^2 n^{-2},$$

and taking the limit $n \to +\infty$, we deduce that $\hat{v}$ and $v(\cdot, \mathbb{P}_\xi)$ coincide $\mathbb{P}_\xi$ almost everywhere. This completes the proof when v is bounded, and ξ has a continuous positive density p, p and its derivatives being of exponential decay at infinity.

Third step. When v is bounded, and ξ is bounded and has a general distribution, we approximate ξ by $\xi + n^{-1}G$ again. Then $\xi + n^{-1}G$ has a positive continuous density, the density and its derivatives being of Gaussian decay at infinity, so that, by the second step, the function $\mathbb{R}^d \ni x \mapsto v(x, \mathbb{P}_{\xi+n^{-1}G})$ can be assumed to be C-Lipschitz continuous for each $n \geq 1$. Extracting a convergent subsequence and passing to the limit as above, we deduce that $v(\cdot, \mathbb{P}_\xi)$ admits a C-Lipschitz continuous version.

When v is bounded but ξ is not bounded, we approximate ξ by its orthogonal projection on the ball of center 0 and radius n. We then complete the proof in a similar way.

Finally when v is not bounded, we approximate v by $(\psi_n(v))_{n\geq 1}$ where, for each $n \geq 1$, ψ_n is a bounded smooth function from $\mathbb{R}$ into itself such that $\psi_n(r) = r$ for $r \in [-n, n]$ and $|[d\psi_n/dr](r)| \leq 1$ for all $r \in \mathbb{R}$. Then, for each $n \geq 1$, there exists a C-Lipschitz continuous version of $\psi_n(v(\cdot, \mathbb{P}_\xi))$. Choosing some $x_0 \in \mathbb{R}^d$ such that $|v(x_0, \mathbb{P}_\xi)| < +\infty$, the sequence $\psi_n(v(x_0, \mathbb{P}_\xi))$ is bounded so that the sequence of functions $(\psi_n(v(\cdot, \mathbb{P}_\xi)))_{n\geq 1}$ is uniformly bounded and continuous on compact subsets. Extracting a converging subsequence, we complete the proof in the same way as before.

REFERENCES

[1] AMBROSIO, L., GIGLI, N. and SAVARÉ, G. (2008). *Gradient Flows in Metric Spaces and in the Space of Probability Measures*, 2nd ed. Birkhäuser, Basel. MR2401600

[2] ANDERSSON, D. and DJEHICHE, B. (2011). A maximum principle for SDEs of mean-field type. *Appl. Math. Optim.* **63** 341–356. MR2784835

[3] BENSOUSSAN, A., SUNG, K. C. J., YAM, S. C. P. and YUNG, S. P. (2011). Linear quadratic mean field games. Technical report.

[4] BUCKDAHN, R., DJEHICHE, B. and LI, J. (2011). A general stochastic maximum principle for SDEs of mean-field type. *Appl. Math. Optim.* **64** 197–216. MR2822408

[5] BUCKDAHN, R., DJEHICHE, B., LI, J. and PENG, S. (2009). Mean-field backward stochastic differential equations: A limit approach. *Ann. Probab.* **37** 1524–1565. MR2546754

[6] CARDALIAGUET, P. (2012). Notes on mean field games. Notes from P. L. Lions' lectures at the Collège de France. Available at https://www.ceremade.dauphine.fr/~cardalia/MFG100629.pdf.

[7] CARMONA, R. and DELARUE, F. (2014). The master equation for large population equilibriums. In *Stochastic Analysis and Applications 2014* (B. Hambly, D. Crisan, T. Zariphopoulou and M. Reizakis, eds.) 77–128. Springer, Cham. MR3332710.

[8] CARMONA, R. and DELARUE, F. (2013). Mean field forward–backward stochastic differential equations. *Electron. Commun. Probab.* **18** 1–15. MR3091726

[9] CARMONA, R. and DELARUE, F. (2013). Probabilistic analysis of mean-field games. *SIAM J. Control Optim.* **51** 2705–2734. MR3072222

[10] CARMONA, R., DELARUE, F. and LACHAPELLE, A. (2013). Control of McKean–Vlasov dynamics versus mean field games. *Math. Financ. Econ.* **7** 131–166. MR3045029

[11] DELARUE, F. (2002). On the existence and uniqueness of solutions to FBSDEs in a non-degenerate case. *Stochastic Process. Appl.* **99** 209–286. MR1901154

[12] HUANG, M., MALHAMÉ, R. P. and CAINES, P. E. (2006). Large population stochastic dynamic games: Closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.* **6** 221–251. MR2346927

[13] JOURDAIN, B., MÉLÉARD, S. and WOYCZYNSKI, W. A. (2008). Nonlinear SDEs driven by Lévy processes and related PDEs. *ALEA Lat. Am. J. Probab. Math. Stat.* **4** 1–29. MR2383731

[14] LASRY, J.-M. and LIONS, P.-L. (2006). Jeux à champ moyen. I. Le cas stationnaire. *C. R. Math. Acad. Sci. Paris* **343** 619–625. MR2269875

[15] LASRY, J.-M. and LIONS, P.-L. (2006). Jeux à champ moyen. II. Horizon fini et contrôle optimal. *C. R. Math. Acad. Sci. Paris* **343** 679–684. MR2271747

[16] LASRY, J.-M. and LIONS, P.-L. (2007). Mean field games. *Jpn. J. Math.* **2** 229–260. MR2295621

[17] LIONS, P. L. (2007/2008). Théorie des jeux à champs moyen et applications. Technical report.

[18] MCKEAN, H. P. JR. (1966). A class of Markov processes associated with nonlinear parabolic equations. *Proc. Natl. Acad. Sci. USA* **56** 1907–1911. MR0221595

[19] MCKEAN, H. P. JR. (1967). Propagation of chaos for a class of non-linear parabolic equations. In *Stochastic Differential Equations* (*Lecture Series in Differential Equations*, *Session* 7, *Catholic Univ.*, 1967) 41–57. Air Force Office Sci. Res., Arlington, VA. MR0233437

[20] PARDOUX, É. and PENG, S. G. (1990). Adapted solution of a backward stochastic differential equation. *Systems Control Lett.* **14** 55–61. MR1037747

[21] PENG, S. and WU, Z. (1999). Fully coupled forward–backward stochastic differential equations and applications to optimal control. *SIAM J. Control Optim.* **37** 825–843. MR1675098

[22] RACHEV, S. T. and RÜSCHENDORF, L. (1998). *Mass Transportation Problems*: *Applications*. *Vol. II*. Springer, New York. MR1619171

[23] SZNITMAN, A.-S. (1991). Topics in propagation of chaos. In *École D'Été de Probabilités de Saint-Flour XIX*—1989 (D. L. Burkholder et al., eds.). *Lecture Notes in Math.* **1464** 165–251. Springer, Berlin. MR1108185

[24] VILLANI, C. (2009). *Optimal Transport*: *Old and New*. *Grundlehren der Mathematischen Wissenschaften* **338**. Springer, Berlin. MR2459454

[25] WU, Z. and YU, Z. (2014). Probabilistic interpretation for a system of quasilinear parabolic partial differential equation combined with algebra equations. *Stochastic Process. Appl.* **124** 3921–3947. MR3264433

[26] YONG, J. and ZHOU, X. Y. (1999). *Stochastic Controls*: *Hamiltonian Systems and HJB Equations*. *Applications of Mathematics* (*New York*) **43**. Springer, New York. MR1696772

ORFE, BENDHEIM CENTER FOR FINANCE
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: rcarmona@princeton.edu

LABORATOIRE JEAN-ALEXANDRE DIEUDONNÉ
UNIVERSITÉ NICE-SOPHIA ANTIPOLIS
PARC VALROSE
06108 CEDEX 02, NICE
FRANCE
E-MAIL: delarue@unice.fr