SpezialForschungsBereich F 32

Karl–Franzens Universität Graz
Technische Universität Graz
Medizinische Universität Graz

Optimality conditions for mean-field type optimal control problems

L. Pfeiffer

SFB-Report No. 2015-015 December 2015

A–8010 GRAZ, HEINRICHSTRASSE 36, AUSTRIA

Supported by the
Austrian Science Fund (FWF)

Der Wissenschaftsfonds.

SFB sponsors:

- **Austrian Science Fund** (FWF)
- **University of Graz**
- **Graz University of Technology**
- **Medical University of Graz**
- **Government of Styria**
- **City of Graz**

Optimality conditions for mean-field type optimal control problems

Laurent Pfeiffer*

December 1, 2015

Abstract

In this article, a maximum principle and a dynamic programming principle for mean-field type control problems are proved. These problems involve a cost functional depending on the probability distribution of the state variable, which is itself the solution to a controlled stochastic differential equation. The control processes are adapted processes with respect to the Brownian motion. Our approach uses in an essential manner a convexity property of the set of reachable probability measures.

Keywords Mean-field type control, stochastic optimization, maximum principle, dynamic programming, relaxation.

AMS subject classification 90C15, 93E20.

1 Introduction

1.1 Contributions of the article

Context This article is dedicated to the derivation of optimality conditions for mean-field type optimal control problems. More precisely, we provide a maximum principle and a dynamic programming principle. The state variable is modelled by a controlled stochastic differential equation (SDE). The specificity of the problems under study lies in the fact that the cost function is formulated as a function of the probability distribution of the state variable. These problems include what we call in this article "standard" problems: stochastic optimal control problems for which the cost function is the expectation of a function of the state variable. As it is well-known, standard problems can be tackled with a dynamic programming approach, consisting in solving the Hamilton-Jacobi-Bellman (HJB) equation.

In this paper, we give optimality conditions for control processes which are adapted processes with respect to the Brownian motion. Note that if the control processes are sufficiently regular – typically if they take the form a feedback control which is Lipschitz-continuous with respect to the state variable –, the evolution of the probability distribution of the state variable (with respect to time) can be described with the Fokker-Planck equation, a linear parabolic partial differential equation (PDE). Thus, the whole problem can be reformulated as a PDE-constrained problem. We refer to [20, 21] for general references on this equation and its relation with SDEs. For the considered problems, the first-order optimality conditions consist in a system of two coupled PDEs: a forward PDE, the Fokker-Planck equation and a

*Institute for Mathematics and Scientific Computing, Karl-Franzens-Universität, Heinrichstraße 36, 8010 Graz, Austria. Mail: laurent.pfeiffer@uni-graz.at.

backward PDE, the HJB equation. The HJB equation can be seen in this setting as an adjoint equation to the Fokker-Planck equation. This forward-backward coupling has received much attention in the last years, with the development of the mean-field game theory with the seminal papers [14] and [11]. For an introduction to this topic, we refer to [6] and [3].

Mean-field type control problems arise typically in risk-averse optimization: when one tries to minimize a random cost, it may be unsatisfactory to minimize the expectation of the cost, and one may prefer to take into account the "risk" induced by the dispersion of the cost. Different models for the risk can be considered, and many of them can be formulated as functions of the probability distribution of the state variable, see for example [23, Chapter 6]. Mean-field type control problems also arise in potential mean-field games, an example is given in [13] to modelize the emergence of a new technology. Let us also mention that model predictive approaches have been recently developed for the optimal control of the Fokker-Planck equation, motivated by problems in quantum mechanics, see [1, 2]. Finally, let us also mention that stochastic problems with a constraint on the probability distribution of the state variable have been recently studied: in [5], the final probability distribution is prescribed, in [18], the constraint is simply an expectation constraint.

Description of the main results Recently, first-order optimality conditions taking the form of a maximum principle were derived for mean-field type control problems in [3, Chapter 4] and [15, Proposition 3], in the usual form of a backward-forward coupling, as explained above. The control processes considered in these references are regular feedback controls, for which the Fokker-Planck equation is well-posed. Thus, it is possible to derive a maximum principle with the usual "needle perturbation" method of the calculus of variations. Let us mention that a "stochastic maximum principle" is also derived in [3, Chapter 4], consisting in a coupling of a forward SDE and a backward SDE.

Our approach is different in so far as the set of admissible control processes is the set of adapted processes (with respect to the Brownian motion) taking values in a given set. In this framework, the probability distribution of the state variable is a well-defined object, but cannot be anymore expressed as the solution to the Fokker-Planck equation. Thus, the approach mentioned above cannot be used. As a consequence, the maximum principle cannot be formulated anymore as a forward-backward system, but takes the following form: any optimal control process is also the solution to a standard problem, obtained by linearizing the cost function. Our analysis is mainly based on the following property: the closure of the set of reachable probability distributions is, as we show, convex. The convexity is proved by constructing controls imitating the behavior of relaxed controls. Let us mention that the convexity is used in a closely related paper to build a numerical method [17]. Note that our result is rather general, but requires that the cost function is defined (and differentiable in some sense) on a large set of probability measures. In particular, cost functions described with the density of the probability measure are discarded. Some examples are discussed in [17].

In [15, Proposition 2], a dynamic programming principle is derived. The underlying idea is the usual one: the problem is embedded into a family of problems, parameterized by the initial time and the initial probability distribution of the state variable. The derivation of a dynamic programming principle is rather easy when only feedback controls are allowed. In the framework of this article, a new difficulty arises. For two different initial times $t < s$, there is less available information when starting at time s than when starting at time t, since the Brownian motion cannot be observed between times t and s. We are able to overcome this difficulty by building once again control processes imitating relaxed control processes.

Structure of the article We provide the main notations of the paper and describe the problem under study below. Section 2 gathers technical lemmas related to controlled SDEs. In Lemma 6, we prove the convexity of the closure of the set of reachable probability measures. In Lemma 8, we show how to build control processes imitating the behavior of relaxed processes. In section 3, we prove the maximum principle (Theorem 9) and we prove a dynamic programming principle (Theorem 14), parameterizing the problem by the initial time and the initial probability distribution. We also prove that the corresponding value function is Lipschitz continuous (Proposition 15). Elements on the optimal transportation theory are given in the appendix.

1.2 Formulation of the problem and assumptions

General notations The set of probability measures on $\mathbb{R}^n$ is denoted by $\mathcal{P}(\mathbb{R}^n)$. For a function $\phi : \mathbb{R}^n \to \mathbb{R}$, its integral (if well-defined) with respect to the measure $m \in \mathcal{P}(\mathbb{R}^n)$ is denoted by

$$\int_{\mathbb{R}^n} \phi(x) \, \mathrm{d}m(x) \quad \text{or} \quad \int_{\mathbb{R}^n} \phi \, \mathrm{d}m.$$

Given two measures m_1 and $m_2 \in \mathcal{P}(\mathbb{R}^n)$, we denote:

$$\int_{\mathbb{R}^n} \phi(x) \, \mathrm{d}(m_2(x) - m_1(x)) := \int_{\mathbb{R}^n} \phi(x) \, \mathrm{d}m_2(x) - \int_{\mathbb{R}^n} \phi(x) \, \mathrm{d}m_1(x).$$

For a given random variable X with values in $\mathbb{R}^n$, its probability distribution is denoted by $\mathcal{L}(X) \in \mathcal{P}(\mathbb{R}^n)$. If $m = \mathcal{L}(X) \in \mathcal{P}(\mathbb{R}^n)$, then for any continuous and bounded function $\phi : \mathbb{R}^n \to \mathbb{R}$,

$$\mathbb{E}\big[\phi(X)\big] = \int_{\mathbb{R}^n} \phi \, \mathrm{d}m.$$

We also denote by $\sigma(X)$ the σ-algebra generated by X.

For $p \geq 1$, we denote by $\mathcal{P}_p(\mathbb{R}^n)$ the set of probability measures having a finite p-th moment:

$$\mathcal{P}_p(\mathbb{R}^n) := \Big\{ m \in \mathcal{P}(\mathbb{R}^n) \,|\, \int_{\mathbb{R}^n} |x|^p \, \mathrm{d}m(x) < +\infty \Big\}.$$

We equip $\mathcal{P}_p(\mathbb{R}^n)$ with the Wasserstein distance d_p, see the definition in the appendix as well as the dual representation of d_p for $p = 1$.

For all $R \geq 0$, we define:

$$\bar{B}_p(R) := \Big\{ m \in \mathcal{P}_p(\mathbb{R}^n) \,|\, \int_{\mathbb{R}^n} |x|^p \, \mathrm{d}m(x) \leq R \Big\}. \tag{1}$$

The open (resp. closed) ball of radius $r \geq 0$ and center 0 is denoted by B_r (resp. $\bar{B}_r$), its complement B_r^{c} (resp. $\bar{B}_r^{\mathrm{c}}$). The set of real-valued Lipschitz continuous functions with modulus 1 defined on $\mathbb{R}^n$ is denoted by $1 - \mathrm{Lip}(\mathbb{R}^n)$.

For a given $p \geq 1$, a function $\phi : \mathbb{R}^n \to \mathbb{R}^n$ is said to be dominated by $|x|^p$ if for all $\varepsilon > 0$, there exists $r > 0$ such that for all $x \in B_r^{\mathrm{c}}$,

$$|\phi(x)| \leq \varepsilon |x|^p. \tag{2}$$

Controlled SDEs We fix $T > 0$ and a Brownian motion $(W_t)_{t \in [0,T]}$ of dimension d. For all $0 \leq t \leq T$, $(W_s - W_t)_{s \in [t,T]}$ is a standard Brownian motion. For all $s \in [t, T]$, we denote by $\mathcal{F}_{t,s}$ the σ-algebra generated by $(W_\theta - W_t)_{\theta \in [t,s]}$.

Let U be a compact subset of $\mathbb{R}^n$. For a given random variable Y_t independent of $\mathcal{F}_{t,T}$ with values in $\mathbb{R}^n$, we denote by $\mathcal{U}_t(Y_t)$ the set of control processes $(u_s)_{s \in [t,T]}$ with values in U which are such that for all s, u_s is $(\sigma(Y_t) \times \mathcal{F}_{t,s})$-measurable.

The drift $b : \mathbb{R}^n \times U \to \mathbb{R}^n$ and the volatility $\sigma : \mathbb{R}^n \times U \to \mathbb{R}^{n\times d}$ are given. For all $u \in \mathcal{U}_t(Y_t)$, we denote by $\big(X_s^{t,Y_t,u}\big)_{s\in[t,T]}$ the solution to the following SDE:

$$\mathrm{d}X_s^{t,Y_t,u} = b(X_s^{t,Y_t,u}, u_s)\,\mathrm{d}s + \sigma(X_s^{t,Y_t,u}, u_s)\,\mathrm{d}W_s, \quad \forall s \in [t,T], \quad X_t^{t,Y_t,u} = Y_t. \tag{3}$$

The well-posedness of this SDE is ensured by Assumption 1 [16, Section 5] below. We also denote by $m_s^{t,Y_t,u}$ the probability distribution of $X_s^{t,Y_t,u}$:

$$m_s^{t,Y_t,u} = \mathcal{L}(X_s^{t,Y_t,u}).$$

All along the article, we assume that the following assumption holds true.

Assumption 1. *There exists $L > 0$ such that for all $x, y \in \mathbb{R}^n$, for all $u, v \in U$,*

$$\begin{aligned}
&|b(x,u)| + |\sigma(x,u)| \leq L(1 + |x| + |u|),\\
&|b(x,u) - b(y,v)| + |\sigma(x,u) - \sigma(y,v)| \leq L(|y-x| + |v-u|).
\end{aligned}$$

Formulation of the problem We fix an initial condition Y_0 (independent of $\mathcal{F}_{0,T}$) and $p \geq 2$ such that $\mathcal{L}(Y_0) \in \mathcal{P}_p(\mathbb{R}^n)$. Let χ be a real-valued function defined on $\mathcal{P}_p(\mathbb{R}^n)$. We aim at studying the following problem:

$$\inf_{u\in\mathcal{U}_0(Y_0)} \chi(m_T^{0,Y_0,u}). \tag{P}$$

Throughout the article, we assume that the next two assumptions (dealing with the continuity and the differentiability of χ) are satisfied.

Assumption 2. *The mapping χ is continuous for the d_1-distance.*

In order to state optimality conditions, we will need a notion of derivative for the mapping χ. There are different ways to define the derivative of χ and we refer to [6, Section 6] for a discussion on this topic.

Assumption 3. *The mapping χ is directionally differentiable in the following sense: for all m_1 in $\mathcal{P}_p(\mathbb{R}^n)$, there exists a linear form $D\chi(m_1)$ on $\mathcal{P}_p(\mathbb{R}^n)$ such that for all m_2 in $\mathcal{P}(\mathbb{R}^n)$, for all $\theta \in [0,1]$,*

$$\chi\big((1-\theta)m_1 + \theta m_2\big) = \chi(m_1) + \theta D\chi(m_1)(m_2 - m_1) + o(\theta). \tag{4}$$

Moreover, we assume that the linear form can be identified with a continuous function denoted by $x \in \mathbb{R}^n \mapsto D\chi(m_1, x)$ which is dominated by $|x|^p$.

Under Assumption 3, equation (4) reads:

$$\chi\big((1-\theta)m_1 + \theta m_2\big) = \chi(m_1) + \theta\Big[\int_{\mathbb{R}^n} D\chi(m_1,x)\,\mathrm{d}(m_2(x) - m_1(x))\Big] + o(\theta). \tag{5}$$

Observe that the directional derivative of χ is not uniquely defined: equation (5) remains true if a constant is added to $D\chi(m_1,\cdot)$.

2 Controlled stochastic differential equations

2.1 Fokker-Planck equation and Feynmann-Kack formula

We recall in this subsection in an informal way the connections between the Fokker-Planck equation, the Feynmann-Kack formula, and SDEs. As we mentioned in the introduction, our analysis for deriving optimality conditions is not based on the Fokker-Planck equation. However, the connections between these equations offer an interesting insight into the problem. In this subsection only, we consider the solution $(X_s)_{s\in[0,T]}$ to the uncontrolled SDE:

$$\mathrm{d}X_s = b(X_s)\,\mathrm{d}t + \sigma(X_s)\,\mathrm{d}W_s, \quad \forall s \in [0,T], \quad X_0 = Y_0, \tag{6}$$

where Y_0 is a random variable independent of the Brownian motion. We set $a(x) = \sigma(x)\sigma(x)^{\mathrm{t}}$. For all s, we denote by $m_s \in \mathcal{P}(\mathbb{R}^n)$ the probability distribution of X_s. It is well-known that it is the weak solution to the following parabolic equation, called Fokker-Planck equation:

$$\partial_s m_s - \frac{1}{2}\Delta(ma) + \nabla\cdot(mb) = 0, \quad m_0 = \mathcal{L}(Y_0). \tag{7}$$

The divergence operator $\nabla\cdot$ and the Laplace operator Δ are only considered with respect to the variable x. If $f : \mathbb{R}^n \to \mathbb{R}^n$ (resp. $g : \mathbb{R}^n \to \mathbb{R}$) is differentiable (resp. twice differentiable), the divergence operator (resp. the Laplace operator) is defined by

$$\nabla\cdot f(x) = \sum_{i=1}^n \partial_{x_i} f_i(x) \quad \Big(\text{resp. } \Delta g(x) = \sum_{i,j=1}^n \partial^2_{x_i x_j} g_{ij}(x)\Big).$$

Observe that the Fokker-Planck equation (7) is linear with respect to m. A justification can be found in [6, Lemma 3.3], for example.

For row vectors $p \in \mathbb{R}^n$ and symmetric matrices $Q \in \mathbb{R}^{n\times n}$, we define the Hamiltonian H:

$$H(x,p,Q) = pb(x) + \frac{1}{2}\mathrm{tr}(a(x)Q). \tag{8}$$

If for all s, m_s has a density and if $(s,x) \mapsto m_s(x)$ is sufficiently smooth, then for all $\phi : \mathbb{R}^n \to \mathbb{R}$ sufficiently smooth, one can check that:

$$\int_{\mathbb{R}^n} H(x, \partial_x\phi(x), \partial_{xx}\phi(x)) m_s(x)\,\mathrm{d}x = \int_{\mathbb{R}^n} \frac{1}{2}\Delta(m_s(x)a(x)) - \nabla\cdot(m_s(x)b(x))\phi(x)\,\mathrm{d}x. \tag{9}$$

For $t \in [0,T]$ and $x \in \mathbb{R}^n$, we denote now by $(X_s^{t,x})$ the solution to (6) with the initial condition $X_t = x$. Given $\psi : \mathbb{R}^n \to \mathbb{R}$, the value function $V(t,x) = \mathbb{E}\big[\psi(X_T^{t,x})\big]$ is a viscosity solution to the following PDE: $\forall x \in \mathbb{R}^n$,

$$-\partial_t V(t,x) = H(x, \partial_x V(t,x), \partial_{xx} V(t,x)), \ \forall t \in [0,T], \quad V(T,x) = \psi(x).$$

This the Feynmann-Kack equation [19, page 25], that can be seen (because of (9)) as an adjoint equation to the Fokker-Planck equation.

2.2 Controlled stochastic differential equations

We return to controlled SDEs and first introduce different sets of control processes, differing by their measurability properties. For all t, given a random variable Z_t independent of $\mathcal{F}_{t,T}$ and given a $\sigma(Z_t)$-measurable random variable Y_t with values in $\mathbb{R}^n$, we define the sets $\mathcal{U}_t^0$, $\mathcal{U}_t(Y_t)$, and $\mathcal{U}_t(Z_t)$ as the sets of control processes $(u_s)_{s\in[t,T]}$ taking values in U such that for all $s \in [t,T]$, u_s is respectively $\mathcal{F}_{t,s}$-measurable, $(\sigma(Y_t) \times \mathcal{F}_{t,s})$-measurable and $(\sigma(Z_t) \times \mathcal{F}_{t,s})$-measurable.

Note that the space $\mathcal{U}_t(Z_t)$ is here introduced for a technical purpose. The role of the random variable Z_t will be clear in subsection 2.3. In the definition of the optimal control problem P in the introduction, we only consider control processes in $\mathcal{U}_t(Y_t)$, that is to say, depending on the Brownian motion and the initial condition.

For a given $u \in \mathcal{U}_t(Z_t)$, we still denote by $X_s^{t,Y_t,u}$ the solution to (3), and for a given $u \in \mathcal{U}_t^0$ and $x \in \mathbb{R}^n$, we denote by $X_s^{t,x,u}$ the solution to (3) with the deterministic initial condition $X_t^{t,x,u} = x$. We use the same conventions for the associated probability distributions.

The following lemma is classical, see for example [12, Section 2.5].

Lemma 4. *For all $p \geq 2$, there exists a constant $C > 0$ depending on L, T, and p such that for all $\sigma(Z_t)$-measurable random variables Y_t and $\tilde{Y}_t$ with values in $\mathbb{R}^n$, for all $0 \leq h \leq T-t$, for all $0 \leq s \leq T-h$, the following estimates hold:*

1. $\mathbb{E}\Big[\sup_{t \leq \theta \leq T} |X_\theta^{t,Y_t,u}|^p\Big] \leq C\big(1 + \mathbb{E}\big[|Y_t|^p\big]\big)$
2. $\mathbb{E}\Big[\sup_{s \leq \theta \leq s+h} |X_\theta^{t,Y_t,u} - X_s^{t,Y_t,u}|^p\Big] \leq Ch^{p/2}\big(1 + \mathbb{E}\big[|Y_t|^p\big]\big)$
3. $\mathbb{E}\Big[\sup_{t \leq \theta \leq T} |X_\theta^{t,\tilde{Y}_t,u} - X_\theta^{t,Y_t,u}|^p\Big] \leq C\mathbb{E}\big[|\tilde{Y}_t - Y_t|^p\big]$.

In the following corollary, we obtain a Lipschitz-continuity property of $Y_t \mapsto m_T^{t,Y_t,u}$, uniformly with respect to u.

Corollary 5. *For all $\sigma(Z_t)$-measurable random variables Y_t and $\tilde{Y}_t$ with values in $\mathbb{R}^n$, for all $u \in \mathcal{U}_t(Z_t)$, the following estimate holds:*

$$d_1(m_T^{t,\tilde{Y}_t,u}, m_T^{t,Y_t,u}) \leq \mathbb{E}\big[|X_T^{t,\tilde{Y}_t,u} - X_T^{t,Y_t,u}|\big] \leq \sqrt{C}\big(\mathbb{E}\big[|\tilde{Y}_t - Y_t|^2\big]\big)^{1/2}.$$

Proof. The estimate is obtained by applying successively the dual representation (33), Cauchy-Schwarz inequality and finally the third estimate given by Lemma 4:

$$\begin{aligned} d_1(m_T^{t,\tilde{Y}_t,u}, m_T^{t,Y_t,u}) &\leq \sup_{\phi \in 1-\mathrm{Lip}(\mathbb{R}^n)} \Big\{\mathbb{E}\big[\phi(X_T^{t,\tilde{Y}_t,u}) - \phi(X_T^{t,Y_t,u})\big]\Big\} \\ &\leq \mathbb{E}\big[|X_T^{t,\tilde{Y}_t,u} - X_T^{t,Y_t,u}|\big] \leq \big(\mathbb{E}\big[|X_T^{t,\tilde{Y}_t,u} - X_T^{t,Y_t,u}|^2\big]\big)^{1/2} \leq \sqrt{C}\big(\mathbb{E}\big[|\tilde{Y}_t - Y_t|^2\big]\big)^{1/2}. \end{aligned}$$

□

2.3 Two technical lemmas on relaxation

This subsection contains two technical lemmas which will be used to prove the main results of section 3. Let us explain the underlying purpose with a simple example. Consider two processes u_1 and $u_2 \in \mathcal{U}_t(Y_t)$ and the corresponding final probability distributions m_T^{t,Y_t,u_1} and m_T^{t,Y_t,u_2}. We aim at building a control process u such that:

$$m_T^{t,Y_t,u} = \frac{1}{2}\big(m_T^{t,Y_t,u_1} + m_T^{t,Y_t,u_2}\big). \tag{10}$$

A very simple way of building such a control process is to define a random variable S independent of Y_t and $F_{t,T}$ taking two different values with probability 1/2. A control u realizing (10) can then be constructed in $\mathcal{U}_t\big((S, Y_t)\big)$: it suffices that $u = u_1$ for one value of S and that $u = u_2$ for the other. The obtained controled process can be seen as a relaxed control, since it is now measurable with respect to a larger filtration.

A property like (10) is crucial to derive easily optimality conditions, therefore we need to justify the relaxation process, this is what we do in Lemma 6 and Lemma 8. Relaxed controls (in $\mathcal{U}_t(Z_t)$) are also useful to prove other properties, see subsection 3.4.

Throughout this subsection, the initial time t is fixed, as well as a random variable Y_t independent of $\mathcal{F}_{t,T}$ with a finite second-order moment. We denote by $\mathcal{R}(t,Y_t)$ the set of reachable probability measures at time T:

$$\mathcal{R}(t,Y_t) = \{m_T^{t,Y_t,u} \mid u \in \mathcal{U}_t(Y_t)\}. \tag{11}$$

We denote by $\mathrm{cl}\big(\mathcal{R}(t,Y_t)\big)$ its closure for the d_1-distance.

In Lemma 6, we prove that $\mathrm{cl}\big(\mathcal{R}(t,Y_t)\big)$ convex, for control processes in $\mathcal{U}_t(Y_t)$. This lemma uses a technique introduced in [18, Theorem 22]. The second lemma, Lemma 8, states that $\mathrm{cl}\big(\mathcal{R}(t,Y_t)\big)$ is invariant if we consider processes in $\mathcal{U}_t(Z_t)$ (instead of taking control processes in $\mathcal{U}_t(Y_t)$).

Lemma 6. *For all random variables Y_t with values in $\mathbb{R}^n$, independent of $\mathcal{F}_{t,T}$, with a probability distribution in $\mathcal{P}_2(\mathbb{R}^n)$, the closure of the set of reachable probability measures $cl\big(\mathcal{R}(t,Y_t)\big)$ is convex.*

Proof. Let $K \in \mathbb{N}^*$, let $u^1,...,u^K$ in $\mathcal{U}_t(Y_t)$, let $\theta_1,...,\theta_K$ in $\mathbb{R}_+\backslash\{0\}$ with $\sum_{k=1}^K \theta_k = 1$. To prove the lemma, it suffices to prove that there exists a sequence $(u^\varepsilon)_{\varepsilon\geq 0}$ in $\mathcal{U}_t(Y_t)$ such that

$$d_1\left(m_T^{t,Y_t,u^\varepsilon}, \textstyle\sum_{k=1}^K \theta_k m_T^{t,Y_t,u^k}\right) \underset{\varepsilon\to 0}{\longrightarrow} 0.$$

Let $0<\varepsilon<T$, let $\tilde{u}^1,...,\tilde{u}^K$ be K processes in $\mathcal{U}_{t+\varepsilon}(Y_t)$ such that for all k, the processes $(u^k)_{s\in[t,T-\varepsilon]}$ and $(\tilde{u}^k)_{s\in[t+\varepsilon,T]}$ can be seen as the same measurable function of respectively $(Y_t,(W_s-W_t)_{s\in[t,T-\varepsilon]})$ and $(Y_t,(W_s-W_{t+\varepsilon})_{s\in[t+\varepsilon,T]})$. In other words, we simply delay the observation of the variation of the Brownian motion of a time ε. We view each control process $\tilde{u}^k$ as an element of $\mathcal{U}_{t+\varepsilon}\big((Y_t,(W_s-W_t)_{s\in[t,t+\varepsilon]})\big)$. Let $-\infty = r_0 < ... < r_K = +\infty$ be such that for all $k=1,...,K$,

$$\int_{r_{k-1}}^{r_k} e^{-z^2}\,\mathrm{d}z = \sqrt{2\pi}\theta_k$$

and let us denote by A_k the following event:

$$(W^1_{t+\varepsilon}-W^1_t)/\sqrt{\varepsilon} \in (r_{k-1},r_k),$$

where W^1 is the first Brownian motion. For all k, we have $\mathbb{P}\big[A_k\big]=\theta_k$. Fixing $u^0\in U$, we define $u^\varepsilon \in \mathcal{U}_t(Y_t)$ as follows:

$$u^\varepsilon_t = u^0, \quad \text{for a.a. } t\in[0,\varepsilon], \quad u^\varepsilon_t = \tilde{u}^k_t, \quad \text{for a.a. } t\in[\varepsilon,T], \text{ when } A_k \text{ is realised.}$$

Let $\phi \in$ 1-Lip$(\mathbb{R}^n)$. It holds:

$$\mathbb{E}\big[\phi\big(X_T^{t,Y_t,u^\varepsilon}\big)\big] = \sum_{k=1}^K \theta_k\,\mathbb{E}\Big[\phi\Big(X_T^{t+\varepsilon,X_{t+\varepsilon}^{t,Y_t,u^0},\tilde{u}^k}\Big) \mid A_k\Big]. \tag{12}$$

Since ϕ is 1-Lipschitz, we obtain for all $k=1,...,K$,

$$\Big|\mathbb{E}\Big[\phi\Big(X_T^{t+\varepsilon,X_{t+\varepsilon}^{t,Y_t,u^0},\tilde{u}^k}\Big) - \phi\big(X_T^{t+\varepsilon,Y_t,\tilde{u}^k}\big)|A_k\Big]\Big| \leq \frac{1}{\theta_k}\mathbb{E}\big[|X_T^{t+\varepsilon,X_{t+\varepsilon}^{t,Y_t,u^0},\tilde{u}^k} - X_T^{t+\varepsilon,Y_t,\tilde{u}^k}|\big]. \tag{13}$$

Let us estimate the right-hand-side of (13). By Corollary 5 and by the second estimates of Lemma 4, we obtain:

$$\mathbb{E}\big[|X_T^{t+\varepsilon,X_{t+\varepsilon}^{t,Y_t,u^0},\tilde{u}^k} - X_T^{t+\varepsilon,Y_t,\tilde{u}^k}|\big] = O\Big(\big(\mathbb{E}\big[|X_{t+\varepsilon}^{t,Y_t,u^0}-Y_t|^2\big]\big)^{1/2}\Big) = O(\sqrt{\varepsilon}). \tag{14}$$

Let us estimate the left-hand-side of (13). We notice first that:

$$\mathbb{E}\big[\phi\big(X_T^{t+\varepsilon,Y_t,\tilde{u}^k}\big)\big|A_k\big] = \mathbb{E}\big[\phi\big(X_{T-\varepsilon}^{t,Y_t,u^k}\big)\big]. \tag{15}$$

Moreover, by Corollary 5 and the second estimate of Lemma 4,

$$\big|\mathbb{E}\big[\phi\big(X_T^{t,Y_t,u^k}\big)\big] - \mathbb{E}\big[\phi\big(X_{T-\varepsilon}^{t,Y_t,u^k}\big)\big]\big| \le \big(\mathbb{E}\big[\big|X_T^{t,Y_t,u^k} - X_{T-\varepsilon}^{t,Y_t,u^k}\big|^2\big]\big)^{1/2} = O(\sqrt{\varepsilon}). \tag{16}$$

Combining (13)-(16), we obtain:

$$\left|\mathbb{E}\Big[\phi\Big(X_T^{t+\varepsilon,X_{t+\varepsilon}^{t,Y_t,u^0},\tilde{u}^k}\Big)\big|A_k\Big] - \mathbb{E}\big[\phi\big(X_T^{t,Y_t,u^k}\big)\big]\right| = O(\sqrt{\varepsilon}). \tag{17}$$

Using the dual representation of d_1 (given by (33)) and estimates (12) and (17), we finally obtain:

$$d_1\Big(m_T^{t,Y_t,u^\varepsilon},\sum_{k=1}^K \theta_k m_T^{t,Y_t,u_k}\Big) \le \sup_{\phi\in 1-\mathrm{Lip}} \Big\{\mathbb{E}\big[\phi\big(X_T^{t,Y_t,u^\varepsilon}\big)\big] - \sum_{k=1}^K \theta_k \mathbb{E}\big[\phi\big(X_T^{t,Y_t,u^k}\big)\big]\Big\} = O(\sqrt{\varepsilon}).$$

This concludes the proof. □

Corollary 7. *Given a sequence of control processes $(u^k)_{k\in\mathbb{N}}$ in $\mathcal{U}_t(Y_t)$ and a sequence of nonnegative coefficients $(\theta_k)_{k\in\mathbb{N}}$ such that $\sum_{k\in\mathbb{N}}\theta_k = 1$, we can build a sequence $(u^\varepsilon)_\varepsilon$ in $\mathcal{U}_t(Y_t)$ such that:*

$$d_1\Big(m_T^{t,Y_t,u^\varepsilon},\sum_{k\in\mathbb{N}} \theta_k m_T^{t,Y_t,u^k}\Big) \underset{\varepsilon\to 0}{\longrightarrow} 0.$$

The second technical lemma states that any probability measure obtained with a control process in $\mathcal{U}_t(Z_t)$ can be approximated (as precisely as desired) with a control process in $\mathcal{U}_t(Y_t)$.

Lemma 8. *Let Z_t be a random variable independent of $\mathcal{F}_{t,T}$ and assume that Y_t is $\sigma(Z_t)$-measurable and that $\mathcal{L}(Y_t)\in\mathcal{P}_2(\mathbb{R}^n)$. Then, for all $u\in\mathcal{U}_t(Z_t)$, there exists a sequence $(u^\varepsilon)_{\varepsilon>0}$ in $\mathcal{U}_t(Y_t)$ such that*

$$d_1\big(m_T^{t,Y_t,u}, m_T^{t,Y_t,u^\varepsilon}\big) \underset{\varepsilon\to 0}{\longrightarrow} 0. \tag{18}$$

Proof. Notations. Let us fix $\varepsilon>0$. Let $(A_i)_{i\in\mathbb{N}}$ be a countable measurable partition of $\mathbb{R}^n$, such that there exists a sequence $(y_i)_{i\in\mathbb{N}}$ in $\mathbb{R}^n$ satisfying the following:

$$\forall i\in\mathbb{N},\ \forall x\in A_i,\ |x-y_i|\le\varepsilon. \tag{19}$$

Let $I:\mathbb{R}^n\to\mathbb{N}$ be the unique mapping satisfying: $\forall x\in\mathbb{R}^n,\ x\in A_{I(x)}$. Since the set $\mathcal{P}_1(\mathbb{R}^n)$ is separable for the d_1-distance, the set $\{m_T^{t,y_i,u}\,|\,u\in\mathcal{U}_t^0\}$ is itself separable for all i. We can the define a sequence $(u^{ij})_{j\in\mathbb{N}}$ in $\mathcal{U}_t^0$ satisfying: for all $u\in\mathcal{U}_t^0$, there exists $j\in\mathbb{N}$ such that:

$$d_1\big(m_T^{t,y_i,u}, m_T^{t,y_i,u^{ij}}\big)\le\varepsilon. \tag{20}$$

We denote now by $\mathcal{S}_i$ the following set of sequences: $\big\{\big(\mathbb{E}\big[\varphi^k(X_T^{t,y_i,u})\big]\big)_{k\in\mathbb{N}}\,|\,u\in\mathcal{U}_t^0\big\}$, where $(\varphi^k)_{k\in\mathbb{N}}$ is given by Lemma 18. For all i, j, k, we set $\psi_{ij}^k = \mathbb{E}\big[\varphi^k(X_T^{t,y_i,u^{ij}})\big]$. By (20), for all sequences $(\zeta^k)_{k\in\mathbb{N}}$ in $\mathcal{S}_i$, there exists $j\in\mathbb{N}$ such that

$$\sup_{k\in\mathbb{N}}\{\zeta^k - \psi_{ij}^k\}\le\varepsilon. \tag{21}$$

Construction of u^ε. We fix now a control process u in $\mathcal{U}_t(Z_t)$. For all i and k, we define the following Z_t-measurable random variable:

$$\Psi_i^k = \mathbb{E}\big[\varphi^k(X_T^{t,y_i,u}) \,|\, Z_t\big].$$

The random sequence $(\Psi_i^k)_{k\in\mathbb{N}}$ takes its values in $\mathcal{S}_i$ almost surely. We then define the event B_{ij}:

$$B_{ij} = \Big\{ Y_t \in A_i,\, \sup_{k\in\mathbb{N}} \big(\Psi_i^k - \psi_{ij'}^k \big) > \varepsilon,\, \forall j' = 0, ..., j-1,\, \sup_{k\in\mathbb{N}} \big(\Psi_i^k - \psi_{ij}^k \big) \leq \varepsilon \Big\}$$

and denote by β_{ij} its probability. Observe that the events B_{ij} are disjoint. By (21), $\cup_{j\in\mathbb{N}} B_{ij} = \{Y_i \in A_i\}$ and finally, $\sum_{i,j\in\mathbb{N}} \beta_{ij} = 1$. We set: $\beta_i := \sum_{j\in\mathbb{N}} \beta_{ij} = \mathbb{P}\big[Y_i \in A_i\big]$. By Corollary 7, there exists for all i a control $u^i \in \mathcal{U}_t^0$ such that

$$d_1\Big(\frac{1}{\beta_i} \sum_{j\in\mathbb{N}} \beta_{ij} m_T^{t,y_i,u^{ij}}, m_T^{t,y_i,u^i} \Big) \leq \varepsilon. \tag{22}$$

We finally define $u^\varepsilon \in U_t(Y_t)$ by: $u^\varepsilon = u_{I(Y_t)}$.

Error estimate. We decompose the error (18) into 4 error terms.

$$r_1 = d_1\Big(m_T^{t,Y_t,u^\varepsilon}, m_T^{t,y_{I(Y_t)},u^\varepsilon} \Big), \quad r_2 = d_1\Big(m_T^{t,y_{I(Y_t)},u^\varepsilon}, \sum_{i\in\mathbb{N}} \sum_{j\in\mathbb{N}} \beta_{ij} m_T^{t,y_i,u^{ij}} \Big),$$

$$r_3 = d_1\Big(\sum_{i\in\mathbb{N}} \sum_{j\in\mathbb{N}} \beta_{ij} m_T^{t,y_i,u^{ij}}, m_T^{t,y_{I(Y_t)},u} \Big), \quad r_4 = d_1\big(m_T^{t,y_{I(Y_t)},u}, m_T^{t,Y_t,u} \big).$$

We easily check with Corollary 5 and (19) that: $r_1 = O(\varepsilon)$, $r_4 = O(\varepsilon)$. Then, observe that: $m_T^{t,y_{I(Y_t)},u^\varepsilon} = \sum_{i\in\mathbb{N}} \beta_i m_T^{t,y_i,u^i}$ and thus by (22),

$$r_2 \leq \sum_{i\in\mathbb{N}} \beta_i d_1\Big(\frac{1}{\beta_i} \sum_{j\in\mathbb{N}} \beta_{ij} m_T^{t,y_i,u^{ij}}, m_T^{t,y_i,u^i} \Big) \leq \varepsilon.$$

Finally by the definition of B_{ij} and by Lemma 18,

$$\begin{aligned} r_3 = \;& \sup_{k\in\mathbb{N}} \Big\{ \mathbb{E}\big[\varphi^k\big(X_T^{t,y_{I(Y_t)},u}\big)\big] - \sum_{i,j\in\mathbb{N}} \beta_{ij} \mathbb{E}\big[\varphi^k(X_T^{t,y_i,u^{ij}})\big] \Big\} \\ = \;& \sup_{k\in\mathbb{N}} \Big\{ \sum_{i,j\in\mathbb{N}} \beta_{ij} \big(\mathbb{E}\big[\varphi^k(X_T^{t,y_i,u}) \,|\, B_{ij}\big] - \psi_{ij}^k \big) \Big\} \\ \leq \;& \sum_{i,j\in\mathbb{N}} \beta_{ij} \sup_{k\in\mathbb{N}} \big\{ \mathbb{E}\big[\varphi^k(X_T^{t,y_i,u}) \,|\, B_{ij}\big] - \psi_{ij}^k \big\} \leq \varepsilon. \end{aligned}$$

The lemma follows now with the triangle inequality. □

3 Analysis of the problem

In this section, we give a maximum principle for problem P (defined in the introduction, page 4) and state a dynamic programming principle.

3.1 On the existence of a solution

Recall the definition of $\mathcal{R}(0, Y_0)$ given in (11) of the set of reachable probability measures. Observe that problem P can be reformulated as follows:

$$\inf_{m\in\mathcal{R}(0,Y_0)} \chi(m).$$

By continuity of χ,

$$\inf_{m\in\mathcal{R}(0,Y_0)} \chi(m) = \inf_{m\in \mathrm{cl}(\mathcal{R}(0,Y_0))} \chi(m), \tag{23}$$

where $\mathrm{cl}(\mathcal{R}(0,Y_0))$ is the closure for the d_1-distance.

It is proved in Lemma 16 that for all $R > 0$, the set $\bar{B}_p(R)$ (defined by (1), standing for the set of probability measures having a p-th moment bounded by R) is compact for the d_1-distance. By Lemma 4, $\mathrm{cl}(\mathcal{R}(0,Y_0))$ is a subset of $\bar{B}_p(R)$, for a sufficiently large value of R. Thus $\mathrm{cl}(\mathcal{R}(0,Y_0))$ is compact, as a closed subset of a compact set. Since χ is continuous for the d_1-distance, the right-hand-side of (23) has a minimizer.

However, we are not able to prove the existence of an optimal solution to the original problem P (or, in other words, to prove the existence of an optimal solution in $\mathcal{R}(0,Y_0)$). Note that the existence of an optimal solution is proved in [25, Theorem 5.3] for problems that we call in the next subsection standard problems, under a convexity assumption and considering a weak formulation of the SDE.

3.2 Optimality conditions

We investigate optimality conditions for the stochastic problem P. We start by recalling the approach by dynamic programming for a *linear* cost function of the form:

$$\chi(m) = \int_{\mathbb{R}^n} \phi(x)\,\mathrm{d}m(x),$$

where $\phi : \mathbb{R}^n \to \mathbb{R}$ is continuous and dominated by $|x|^p$. The term "linear" refers here to the following property: for all m_1, $m_2 \in \mathcal{P}_p(\mathbb{R}^n)$, for all $\theta \in [0,1]$,

$$\chi\big(\theta m_1 + (1-\theta)m_2\big) = \theta\chi(m_1) + (1-\theta)\chi(m_2).$$

In this case, the derivative introduced in Assumption 3 is equal to ϕ for all m.

We denote by $P(\phi)$ the following problem:

$$\inf_{u\in\mathcal{U}_0(Y_0)} \mathbb{E}\big[\phi(X_T^{0,Y_0,u})\big]. \tag{$P(\phi)$}$$

We call such a problem *standard* problem. As in section 2, we set $a(x,u) = \sigma(x,u)\sigma(x,u)^{\mathrm{t}}$ and define the unminimized Hamiltonian $h(u,x,p,Q)$ and the Hamiltonian $H(x,p,Q)$ by

$$h(u,x,p,Q) = pb(u,x) + \frac{1}{2}\mathrm{tr}(a(x,u)Q) \quad \text{and} \quad H(x,p,Q) = \inf_{u\in U} h(u,x,p,Q), \tag{24}$$

where p is a row vector of size n and Q a symmetric matrix of size n. As is well-known, the standard problem $P(\phi)$ can be solved by dynamic programming. We introduce the value function $V : [0,T]\times\mathbb{R}^n \to \mathbb{R}$, defined by:

$$V(t,x) = \inf_{u\in\mathcal{U}_t^0} \mathbb{E}\big[\phi(X_T^{t,x,u})\big].$$

It is the viscosity solution to the following Hamilton-Jacobi-Bellman (HJB) equation:

$$-\partial_t V(t,x) = H(x, \partial_x V(t,x), \partial_{xx} V(t,x)), \quad V(T,x) = \phi(x). \tag{25}$$

An optimal solution $\bar{u}$ to $P(\phi)$ is then such that for a.a. t, $\bar{u}_t$ minimizes almost surely $h\big(\cdot, X_t^{0,Y_0,\bar{u}}, \partial_x V(t,X_t^{0,Y_0,\bar{u}}), \partial_{xx}V(t,X_t^{0,Y_0,\bar{u}})\big)$, if V is sufficiently regular. The dynamic programming principle associated with V states also that

$$\mathbb{E}\big[V(t,X_t^{0,Y_0,\bar{u}})\big] \ \Big(= \int_{\mathbb{R}^n} V(t,x)\,\mathrm{d}m_t^{0,Y_0,\bar{u}}(x)\Big)$$

is independent of $t \in [0,T]$ and equal to the value of $P(\phi)$.

General case The following theorem is a maximum principle for problem P.

Theorem 9. *Let $\bar{u} \in \mathcal{U}_0(Y_0)$ be an optimal solution to problem P. Then, $\bar{u}$ is a solution to the standard problem $P(\phi)$, where $\phi(\cdot) = D\chi(\bar{m}, \cdot)$, with $\bar{m} = m_T^{0,Y_0,\bar{u}}$.*

In the sequel, we call problem $P(\phi)$ *linearized problem* when $\phi(\cdot) = D\chi(\bar{m}, \cdot)$, in order to emphasize its connection with problem P.

Proof. Using the continuity of χ of the d_1-distance, we obtain that $\bar{m}$ is optimal on $\mathrm{cl}(\mathcal{R}(0, Y_0))$, which is by Lemma 6 a convex set. Therefore, for all $u \in \mathcal{U}_0(Y_0)$, for all $\theta \in [0,1]$,

$$\begin{aligned} 0 \leq \chi\big((1-\theta)\bar{m} + \theta m_T^{0,Y_0,u}\big) - \chi(\bar{m}) &= \theta D\chi(\bar{m})(m_T^{0,Y_0,u} - \bar{m}) + o(\theta) \\ &= \theta \mathbb{E}\big[D\chi\big(\bar{m}, X_T^{0,Y_0,u}\big) - D\chi\big(\bar{m}, X_T^{0,Y_0,\bar{u}}\big)\big] + o(\theta). \end{aligned} \tag{26}$$

Therefore, $\mathbb{E}\big[D\chi\big(\bar{m}, X_T^{0,Y_0,u}\big)\big] \geq \mathbb{E}\big[D\chi\big(\bar{m}, X_T^{0,Y_0,\bar{u}}\big)\big]$. The theorem is proved. □

Here, it is not possible in general to compute directly the value function in order to obtain a characterization of the optimal solution (as we would do to solve problem $P(\phi)$), since the terminal condition $D\chi(\bar{m}, \cdot)$ itself depends on the optimal control.

The following lemma explains the role of the value function associated with the linearized problem when χ is convex. Note that in this case, the necessary condition of Theorem 9 is also a sufficient condition.

Lemma 10. *Assume that χ is convex on $\mathcal{P}_p(\mathbb{R}^n)$, that is to say, for all $\theta \in [0,1]$, for all m_1 and $m_2 \in \mathcal{P}_p(\mathbb{R}^n)$,*

$$\chi(\theta m_1 + (1-\theta) m_2) \leq \theta \chi(m_1) + (1-\theta)\chi(m_2). \tag{27}$$

Then, for all $\bar{m} \in cl(\mathcal{R}(0, Y_0))$, the following upper estimate holds:

$$\chi(\bar{m}) - \Big(\inf_{m \in cl(\mathcal{R}(0,Y_0))} \chi(m) \Big) \leq D\chi(\bar{m})\bar{m} - Val(P(\phi)), \tag{28}$$

where $\phi = D\chi(\bar{m}, \cdot)$. In particular, if a control process $\bar{u}$ is a solution to $P(\phi)$, with $\phi = D\chi(m_T^{0,Y_0,\bar{u}})$, then $\bar{u}$ is an optimal solution to P.

Proof. Since χ is convex, the following inequality holds true for all $m \in \mathcal{P}_p(\mathbb{R}^n)$:

$$\chi(m) - \chi(\bar{m}) \geq D\chi(\bar{m})(m - \bar{m}).$$

The lemma follows directly, minimizing both sides of the last inequality. □

We finally give an extension of Theorem 9.

Proposition 11. *Let $(\alpha_k)_{k\in\mathbb{N}}$ be a sequence in $\mathbb{R}_+ \backslash \{0\}$ such that $\sum_{k=0}^{+\infty} \alpha_k = 1$, let $(u^k)_{k\in\mathbb{N}}$ be a sequence in $\mathcal{U}_0(Y_0)$ and let $\bar{m} = \sum_{k=0}^{\infty} \alpha_k m_T^{0,Y_0,u^k}$. By Corollary 7, $\bar{m} \in cl(\mathcal{R}(0, Y_0))$. If $\bar{m}$ minimizes χ over $cl(\mathcal{R}(0, Y_0))$, then the control processes u^0,... are all solutions to the same linearized problem $P(\phi)$ with $\phi = D\chi(\bar{m}, \cdot)$.*

Proof. The proof is the same as the proof of Theorem 9. Let $k \in \mathbb{N}$, let $u \in \mathcal{U}_0(Y_0)$, we define then:

$$m = \Big(\sum_{i \in \mathbb{N}\backslash\{k\}} \alpha_i m_T^{0,Y_0,u^i} \Big) + \alpha_k m_T^{0,Y_0,u}.$$

By Corollary 7, $m \in \mathrm{cl}(\mathcal{R}(0, Y_0))$, as well as $\theta \bar{m} + (1-\theta)m$, for all $\theta \in [0,1]$. Thus,

$$\begin{aligned} 0 \leq \chi\big((1-\theta)\bar{m} + \theta m\big) - \chi(\bar{m}) &= \theta D\chi(\bar{m})(m - \bar{m}) + o(\theta) \\ &= \alpha_k \theta D\chi(\bar{m})(m_T^{0,Y_0,u} - m_T^{0,Y_0,u^k}) + o(\theta), \end{aligned} \tag{29}$$

and the conclusion follows as in the proof Theorem 9. □

As a consequence of Proposition 11, if the linearized problem has a unique solution, then $\bar{m}$ cannot be written as a convex combination of m_T^{0,Y_0,u^k} with different values of u^k.

Remark 12. *Theorem 9 and Proposition 11 can be easily extended to the following cost function χ', in which an integral term is involved:*

$$\chi'(m,u) = \chi(m) + \mathbb{E}\Big[\int_0^T \ell\big(X_t^{0,Y_0,u}, u_t\big)\, dt \Big],$$

where $\ell : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is given and satisfies the same assumptions as b and σ. It suffices to introduce a supplementary state variable with dynamics ℓ, as it is customary. We believe that the optimality conditions can be extended to an integral cost ℓ depending also on the current probability distribution $m_t^{0,Y_0,u}$, but such results cannot be obtained by a simple reformulation of the problem.

3.3 Discussion

Mean-Field Game interpretation The optimality conditions state that an optimal control $\bar{u}$ is a Nash equilibrium for the following Mean-Field Game: there are infinitely many players, having an initial distribution of Y_0. A given player with initial position x_0 minimizes the cost: $\mathbb{E}\big[D\chi(m, X_T^{0,x_0,u})\big]$, with $u \in \mathcal{U}_0(x_0)$, where m is the probability distribution of the players at time T. Our problem is a particular Mean-Field Game due to the particular structure of $D\chi(m,x)$, which can be interpreted as a derivative (in the sense of Assumption 3).

Formulation of P as an optimal control problem of the Fokker-Planck equation As we already mentioned in the introduction, optimality conditions have been derived so far for mean-field type problems formulated as optimal control of the Fokker-Planck equation as follows:

$$\inf_{u \in \bar{\mathcal{U}}} \chi(m(T,\cdot)), \quad \text{subject to:} \tag{30}$$
$$\begin{cases} \partial_t m(t,x) - \frac{1}{2}\Delta\big(m(t,x)a(x,u(t,x))\big) + \operatorname{div}\big(m(t,x)b(x,u(t,x))\big) = 0, \\ \quad m(0,x) = m_0(x), \end{cases}$$

where $\bar{\mathcal{U}} \subset L^\infty([0,T]\times\mathbb{R}^n, U)$ is an appropriate set of feedback controls ensuring the well-posedness of the Fokker-Planck equation. We refer to [9] for recent results dealing with the well-posedness of the Fokker-Planck equation for rough coefficients. In the references [3, Chapter 4] and [15, Proposition 3], a maximum principle is derived for an optimal feedback control which is Lipschitz with respect to the state variable. Observe that in this setting, the value function introduced in (25) plays the role of an adjoint variable for the Fokker-Planck equation, as a consequence of equation (9).

In our framework, a much richer structure is allowed, since the controls are adapted processes (possibly depending on the past of the Brownian motion). Our approach is similar to the one based on relaxation with Young measures for optimal control of deterministic nonlinear optimal control problems, which was already known in the 60s. This approach is explained in [4, Section 3] for example, Pontryagin's principle is directly deduced from the convexity of the set of reachable linearized states.

3.4 Properties of the general value function

We introduce in this subsection a general value function $\mathcal{V}(t, m_t)$, taking for state the whole probability distribution at time t. We use the term *general* to emphasize the difference with the value function associated with the linearized problem $P(\phi)$ (with $\phi = D\chi(\bar{m})$). The general value function is defined as follows: for all t, for all $p \geq 2$, for all $m_t \in \mathcal{P}_p(\mathbb{R}^n)$, for all random variable Y_t independent of $\mathcal{F}_{t,T}$ with probability distribution m_t,

$$\mathcal{V}(t, m_t) = \inf_{u \in \mathcal{U}_t(Y_t)} \chi\big(m_T^{t,Y_t,u}\big).$$

The following remark justifies that $\mathcal{V}(t, m_t)$ is independent of the choice of Y_t.

Remark 13. *Let Y_t and $\tilde{Y}_t$ be two random variables independent of $\mathcal{F}_{t,T}$ having the same probability distribution, let u and $\tilde{u}$ be respectively in $\mathcal{U}_t(Y_t)$ and $\mathcal{U}_t(\tilde{Y}_t)$ having the same probability distribution. Then, for all $s \in [t, T]$,*

$$m_s^{t,Y_t,u} = m_s^{t,\tilde{Y}_t,\tilde{u}}.$$

In this subsection, we prove that $\mathcal{V}$ satisfies a dynamic programming principle and that if χ is Lipschitz continuous for the d_1-distance, then $\mathcal{V}$ is Lipschitz continuous with respect to m_t, for the d_1-distance.

Dynamic programming principle The next theorem provides a dynamic programming principle. As mentioned in the introduction, the main difficulty lies in the fact that for two different initial times, the available information is not the same.

Theorem 14. *The following dynamic programming principle holds: for all $p \geq 2$, for all $t \leq s \leq T$, for all $m_t \in \mathcal{P}_p(\mathbb{R}^n)$, for all random variable Y_t, independent of $\mathcal{F}_{t,T}$ and with probability distribution m_t, it holds:*

$$\mathcal{V}(t, m_t) = \inf_{u \in \mathcal{U}_t(Y_t)} \mathcal{V}\big(s, m_s^{t,Y_t,u}\big). \tag{31}$$

Proof. Let us fix s, t, m_t, and Y_t. We denote by $\mathcal{V}'(t, m_t)$ the r.h.s. of (31).

- *First inequality: $\mathcal{V}(t, m_t) \leq \mathcal{V}'(t, m_t)$.*

Let $\varepsilon > 0$, let $u^1 \in \mathcal{U}_t(Y_t)$ be ε-optimal for $\mathcal{V}'(t, m_t)$. Let also $u^2 \in \mathcal{U}_s(X_s^{t,Y_t,u_1})$ be ε-optimal for $\mathcal{V}(s, m_s^{t,Y_t,u_1})$. We define then the following control process $\bar{u} \in \mathcal{U}_t(Y_t)$:

$$\bar{u}_\theta = u_\theta^1, \text{ if } \theta \in [t, s], \quad \bar{u}_\theta = u_\theta^2, \text{ if } \theta \in (s, T].$$

By construction, $m_T^{t,Y_t,\bar{u}} = m_T^{s,X_s^{t,Y_t,u^1},u^2}$ and therefore

$$\mathcal{V}(t, m_t) \leq \chi\Big(m_T^{s,X_s^{t,Y_t,u^1},u^2}\Big) \leq \mathcal{V}\big(s, m_s^{t,Y_t,u^1}\big) + \varepsilon \leq \mathcal{V}'(t, m_t) + 2\varepsilon.$$

The first inequality follows.

- *Second inequality: $\mathcal{V}(t, m_t) \geq \mathcal{V}'(t, m_t)$.*

Let $u \in \mathcal{U}_t(Y_t)$ be ε-optimal for $\mathcal{V}(t, m_t)$. Let us denote by Z_t the following random variable: $(Y_t, (W_\theta)_{\theta \in [t,s]})$. Then, the control process $(u_\theta)_{\theta \in [s,T]}$ belongs to $\mathcal{U}_s(Z_t)$. We set $Y_s = X_s^{t,Y_t,u}$. By continuity of χ, there exists $\alpha > 0$ such that for all $m \in \mathcal{P}_p(\mathbb{R}^n)$, if $d_1(m_T^{s,Y_s,u}, m) \leq \alpha$, then, $\chi(m_T^{s,Y_s,u}) \geq \chi(m) - \varepsilon$. By Lemma 8, there exists $u' \in \mathcal{U}_s(Y_s)$ such that $d_1(m_T^{s,Z_t,u}, m_T^{s,Y_s,u'}) \leq \alpha$. We obtain:

$$\mathcal{V}(t, m_t) \geq \chi\big(m_T^{s,Y_s,u}\big) - \varepsilon = \chi\big(m_T^{s,Y_s,u_\alpha}\big) - 2\varepsilon \geq \mathcal{V}(s, m_s^{t,Y_t,u}) - 2\varepsilon \geq \mathcal{V}'(t, m_t) - 2\varepsilon.$$

The second inequality follows then when $\varepsilon \to 0$. □

An approach by dynamic programming for problem P is tempting, since it would allow to compute a global minimizer. This is however numerically intractable, since the whole probability distribution should be taken as a state variable.

The dynamic programming principle ensures that an optimal strategy is time-consistent (see the papers [7, 22] for a discussion of this concept). It would also enable us to derive a Hamilton-Jacobi-Bellman equation for V and to derive a notion of viscosity solutions, as in [8, 10].

Lipschitz continuity

Proposition 15. *If χ is Lipschitz for the d_1-distance, then, for all $t \in [0,T]$, for all $p \geq 2$, the general value function $\mathcal{V}(t,\cdot)$ is Lipschitz continuous for the d_p-distance.*

Proof. Let m_t^1 and $m_t^2 \in \mathcal{P}_p(\mathbb{R}^n)$, let π be an optimal transportation plan between m_t^1 and m_t^2 for the d_p-distance. Let Z_t be a random variable in $\mathbb{R}^{2n}$, independent of $\mathcal{F}_{t,T}$, with probability distribution π. We write $Z_t = (Y_t^1, Y_t^2)$ with Y_t^1 and Y_t^2 taking values in $\mathbb{R}^n$. By construction, the probability distribution of Y_t^1 and Y_t^2 is equal to respectively m_t^1 and m_t^2. Let $u \in \mathcal{U}_t(Y_t^1)$. This control process can of course be seen as an element of $\mathcal{U}_t(Z_t)$. By Hölder's inequality and Lemma 4,

$$d_1(m_T^{t,Y_t^1,u}, m_T^{t,Y_t^2,u}) \leq d_p(m_T^{t,Y_t^1,u}, m_T^{t,Y_t^2,u}) \leq C d_p(m_t^1, m_t^2).$$

Moreover, by Lemma 8, for all ε, there exists $u^\varepsilon \in \mathcal{U}_t(Y_t^2)$ such that:

$$d_1\big(m_T^{t,Y_t^2,u}, m_T^{t,Y_t^2,u^\varepsilon}\big) \leq \varepsilon.$$

Thus, denoting by K the Lipschitz constant of χ,

$$\begin{aligned} \chi(m_T^{t,Y_t^2,u^\varepsilon}) - \chi(m_T^{t,Y_t^1,u}) &\leq K\big(d_1(m_T^{t,Y_1,u}, m_T^{t,Y_2,u}) + d_1(m_T^{t,Y_2,u}, m_T^{t,Y_2,u^\varepsilon})\big) \\ &\leq K(C d_p(m_t^1, m_t^2) + \varepsilon). \end{aligned}$$

Minimizing the left-hand side, then the right-hand side, then making ε converge to 0, we obtain:

$$\mathcal{V}(t, m_t^2) \leq \mathcal{V}(t, m_t^1) + KC d_p(m_t^1, m_t^2).$$

Exchanging m_t^1 and m_t^2, we obtain: $\mathcal{V}(t, m_t^1) \leq \mathcal{V}(t, m_t^2) + KC d_p(m_t^1, m_t^2)$, and the Lipschitz-continuity follows. ☐

Acknowledgements The author gratefully acknowledges the Austrian Science Fund (FWF) for financial support under SFB F32, "Mathematical Optimization and Applications in Biomedical Sciences."

A Elements on optimal transportation

Wasserstein distance For $p \geq 1$, let us recall the definition of the Wasserstein distance, denoted by d_p in the article. For all m_1 and m_2 in $\mathcal{P}_p(\mathbb{R}^n)$,

$$d_p(m_1, m_2) = \Big[\inf_{\pi \in \Pi(m_1,m_2)} \int_{\mathbb{R}^n \times \mathbb{R}^n} |y - x|^p \, \mathrm{d}\pi(x,y) \Big]^{1/p}, \tag{32}$$

$\Pi(m_1, m_2)$ being the set of transportation mappings from m_1 to m_2 defined as:

$$\left\{ \pi \in \mathcal{P}(\mathbb{R}^{2n}) \,|\, \left\{ \begin{array}{l} \pi(A \times \mathbb{R}^n) = m_1(A), \\ \pi(\mathbb{R}^n \times A) = m_2(A), \end{array} \right. \text{for all measurable } A \subset \mathbb{R}^n \right\}.$$

By Hölder's inequality, for all $1 \leq p \leq p'$, $\mathcal{P}_{p'}(\mathbb{R}^n) \subset \mathcal{P}_p(\mathbb{R}^n)$. Moreover, for all $1 \leq p \leq p'$ and for all m and m' in $\mathcal{P}_{p'}(\mathbb{R}^n)$,

$$d_p(m, m') \leq d_{p'}(m, m').$$

Note that $\mathcal{P}_p(\mathbb{R}^n)$ equipped with d_p is complete and separable [24, Theorem 6.18]. Note also the dual representation of d_1 [24, Remark 6.5]: for all m_1, $m_2 \in \mathcal{P}_1(\mathbb{R}^n)$,

$$d_1(m_1, m_2) = \sup_{\phi \in 1-\mathrm{Lip}(\mathbb{R}^n)} \int_{\mathbb{R}^n} \phi \, \mathrm{d}(m_2 - m_1), \tag{33}$$

where $1 - \mathrm{Lip}(\mathbb{R}^n)$ is the set of Lipschitz function with modulus 1.

A compactness property

Lemma 16. *For all $p > 1$ and $R \geq 0$, the subset $\bar{B}_p(R)$ of $\mathcal{P}_1(\mathbb{R}^n)$ (defined in (1)) is compact for the d_1-distance.*

Proof. We first prove that $\bar{B}_p(R)$ is compact for the weak topology of $\mathcal{P}(\mathbb{R}^n)$. For all $r \geq 0$, for all $m \in \bar{B}_p(R)$,

$$R \geq \int_{\bar{B}_r^c} |x|^p \, \mathrm{d}m(x) \geq r^p \int_{\bar{B}_r^c} 1 \, \mathrm{d}m(x), \tag{34}$$

and thus, $m(\bar{B}_r^c) \leq R/r^p \to 0$, meaning that $\bar{B}_p(R)$ is tight. By Prokhorov's theorem [24, Page 43], $\bar{B}_p(R)$ is therefore precompact for the weak-topology. Now, let $(m_k)_{k \in \mathbb{N}}$ be a sequence in $\bar{B}_p(R)$ weakly converging to $\bar{m} \in \mathcal{P}(\mathbb{R}^n)$. For all $r \geq 0$, the function $\min(|x|^p, r)$ is continuous and bounded, thus:

$$\int_{\mathbb{R}^n} \min(|x|^p, r) \, \mathrm{d}\bar{m}(x) = \lim_{k \to \infty} \int_{\mathbb{R}^n} \min(|x|^p, r) \, \mathrm{d}m_k(x) \leq R.$$

We obtain, using the monotone convergence theorem:

$$\int_{\mathbb{R}^n} |x|^p \, \mathrm{d}\bar{m}(x) = \lim_{r \to \infty} \int_{\mathbb{R}^n} \min(|x|^p, r) \, \mathrm{d}\bar{m}(x) \leq R,$$

thus $\bar{m} \in \bar{B}_p(R)$. Therefore, $\bar{B}_p(R)$ is weakly closed, and thus weakly compact.

Finally, we need to prove that any weakly converging sequence $(m_k)_{k \in \mathbb{N}}$ in $\bar{B}_p(R)$ to some $\bar{m} \in \bar{B}_p(R)$ also converges for the d_1-distance. By [24, Definition 6.8/Theorem 6.9], it suffices to prove that:

$$\int_{\mathbb{R}^n} |x| \, \mathrm{d}m_k(x) \underset{k \to \infty}{\longrightarrow} \int_{\mathbb{R}^n} |x| \, \mathrm{d}\bar{m}(x). \tag{35}$$

Observe that for all $r \geq 0$, for all $m \in \bar{B}_p(R)$, similarly to (34),

$$\Big| \int_{\mathbb{R}^n} |x| - \min(|x|, r) \, \mathrm{d}m(x) \Big| \leq \int_{\bar{B}_r^c} |x| \, \mathrm{d}m(x) \leq \frac{R}{r^{p-1}}.$$

Therefore, for all $r \geq 0$,

$$\liminf_{k \to \infty} \Big| \int_{\mathbb{R}^n} |x| \, \mathrm{d}(m_k(x) - \bar{m}(x)) \Big| \leq \underbrace{\Big(\lim_{k \to \infty} \int_{\mathbb{R}^n} \min(|x|, r) \, \mathrm{d}(m_k(x) - \bar{m}(x)) \Big)}_{=0} + \frac{2R}{r^{p-1}}.$$

We obtain (35), making r tend to $+\infty$. □

On the dual representation of the d_1-distance The following technical lemma extends a 1-Lipschitz function defined on $\bar{B}_r$ to a 1-Lipschitz continuous function on $\mathbb{R}^n$.

Lemma 17. *Let $r > 0$, let $f : \bar{B}_r \to \mathbb{R}$ be 1-Lipschitz continuous. For all $\varepsilon > 0$, there exists $\hat{f} \in$ 1-Lip$(\mathbb{R}^n)$ satisfying:*

1. $\hat{f}(x) \underset{|x|\to+\infty}{\longrightarrow} 0$
2. $|f(x) - \hat{f}(x)| \leq \varepsilon$, *for all* $x \in \bar{B}_r$
3. $|\hat{f}(x)| \leq \sup_{|y|=r} |f(y)|$, *for all* $x \in B_r^c$.

Proof. Let $\varepsilon > 0$, let us set $h(0) = 1 - \min\left(\frac{\varepsilon}{\|f\|_\infty}, \frac{1}{2}\right) \in (0,1)$ and define

$$h : a \in \mathbb{R}_+ \mapsto \begin{cases} h(0) & \text{if } a \in [0,r] \\ h(0) - \frac{1-h(0)}{\|f\|_\infty}(a-r) & \text{if } a \in \left[r, r + \frac{h(0)\|f\|_\infty}{1-h(0)}\right] \\ 0 & \text{otherwise.} \end{cases}$$

Observe that h is decreasing, takes values in $(0,1)$, converges to 0 at infinity and is $\frac{1-h(0)}{\|f\|_\infty}$-Lipschitz continuous. We define then $\hat{f} : x \in \mathbb{R}^n \mapsto \hat{f}(x) = h(|x|)f(P(x))$, where $P : x \in \mathbb{R}^n \mapsto P(x) = \min(|x|, R)\frac{x}{|x|}$ is the projection on $\bar{B}_r$. Firstly, we have:

$$|\hat{f}(x)| \leq h(|x|)\|f\|_\infty \underset{|x|\to+\infty}{\longrightarrow} 0.$$

Secondly, for all $x \in \bar{B}_r$, $h(|x|) = h(0)$ and $P(x) = x$, thus:

$$|\hat{f}(x) - f(x)| = (1 - h(0))|f(x)| = \min\Big(\frac{\varepsilon}{\|f\|_\infty}, \frac{1}{2}\Big)|f(x)| \leq \varepsilon.$$

Thirdly, for all $x \in B_r^c$, $P(x) = \frac{rx}{|x|}$, thus, $|\hat{f}(x)| = h(|x|)|f\big(\frac{rx}{|x|}\big)| \leq \sup_{|y|=r} |f(y)|$. The three conditions are therefore satisfied. Finally, for all x and $y \in \mathbb{R}^n$,

$$\begin{aligned} |\hat{f}(y) - \hat{f}(x)| &\leq |h(|y|) - h(|x|)f(P(y))| + h(|x|)|f(P(y)) - f(P(x))| \\ &\leq \Big(\frac{1-h(0)}{\|f\|_\infty}\big||y| - |x|\big| \cdot \|f\|_\infty\Big) + \Big(h(0)|y-x|\Big) \leq |y-x|, \end{aligned}$$

thus $\hat{f}$ is 1-Lipschitz. □

The following lemma states that the set of functions to consider in the dual representation of d_1 can be reduced to a (countable) sequence of 1-Lipschitz functions. It is used in the proof of Lemma 8.

Lemma 18. *There exists a sequence of 1-Lipschitz continuous functions $(\varphi^k)_{k\in\mathbb{N}}$ satisfying the following property: for all m_1 and $m_2 \in \mathcal{P}_1(\mathbb{R}^n)$,*

$$d_1(m_1, m_2) = \sup_{k\in\mathbb{N}} \Big\{ \int_{\mathbb{R}^n} \varphi^k \, d(m_2 - m_1) \Big\}. \tag{36}$$

Proof. It is well-known that the set of continuous real-valued functions on a compact set of $\mathbb{R}^n$ is separable for the L^∞-norm. Therefore, for all $r > 0$, the set of 1-Lipschitz continuous functions on B_r is itself separable for the L^∞-norm, as a subset of $C(B_r)$. For all $r \in \mathbb{N}\backslash\{0\}$, we denote by Φ_r a dense and countable subset of 1-Lip(B_r). For all $N \in \mathbb{N}\backslash\{0\}$, we denote by $\Phi_{r,N}$ a subset of 1-Lip($\mathbb{R}^n$) obtained by applying Lemma 17 to all the functions of Φ_r with the value $\varepsilon = 1/N$.

We now prove that Lemma 18 holds with a sequence $(\varphi^k)_{k\in\mathbb{N}}$ obtained by listing all the elements of $\cup_{r,N\in\mathbb{N}\setminus\{0\}}\Phi_{r,N}$. Let m_1 and $m_2 \in \mathcal{P}_1(\mathbb{R}^n)$, let $\phi \in$ 1-Lip($\mathbb{R}^n$), let $\varepsilon > 0$. We have to find an index k so that φ^k suitably approximates ϕ. For this purpose, we first define two intermediate functions, $\hat{\phi}$ and $\bar{\phi}$. Without loss of generality, we assume that $\phi(0) = 0$. By the dominated convergence theorem, there exists $r > 0$ such that:

$$\int_{B^c(0,r)} |x| \, \mathrm{d}m_1(x) \leq \varepsilon \quad \text{and} \quad \int_{B^c(0,r)} |x| \, \mathrm{d}m_2(x) \leq \varepsilon.$$

Let $\hat{\phi}$ be a function in 1-Lip($\mathbb{R}^n$) obtained by applying Lemma 17 to the restriction of ϕ to B_r with the parameter ε. Then, let $r' \in \mathbb{N}^*$ be such that for all $x \in B_{r'}^c$, $|\hat{\phi}(x)| \leq \varepsilon$. Let $\tilde{\phi} \in \Phi_{r'}$ be such that $|\bar{\phi}(x) - \hat{\phi}(x)| \leq \varepsilon$ for all $x \in B_{r'}$. Let $N \in \mathbb{N}\setminus\{0\}$ be such that $1/N \leq \varepsilon$ and let k the index for which φ^k is the function obtained by the application of Lemma 17 to $\bar{\phi}$ with the value $1/N$.

We derive now two estimates. First, for all $x \in \mathbb{R}^n$, $|\phi(x)| \leq |x|$ and thus if $|x| \geq r$, $|\hat{\phi}(x)| \leq \sup_{|y|=r} |\phi(y)| \leq r \leq |x|$. Therefore,

$$\Big| \int_{\mathbb{R}^n} \phi - \hat{\phi} \, \mathrm{d}m_1 \Big| = \int_{B_R^c} |\phi| + |\hat{\phi}| \, \mathrm{d}m_1 + \int_{B_r} |\phi - \hat{\phi}| \, \mathrm{d}m_1 \leq 3\varepsilon. \tag{37}$$

Secondly, for all $|x| \geq R$,

$$|\varphi^k(x)| \leq \sup_{|y|=r'} |\bar{\phi}(y)| \leq \Big(\sup_{|y|=r'} |\hat{\phi}(y)| \Big) + \Big(\sup_{|y|=r'} |\bar{\phi}(y) - \hat{\phi}(y)| \Big) \leq 2\varepsilon.$$

Therefore, $\Big| \int_{\mathbb{R}^n} \hat{\phi} - \varphi^k \, \mathrm{d}m_1 \Big|$ is upper bounded by:

$$\int_{B_{r'}} |\hat{\phi} - \bar{\phi}| \, \mathrm{d}m_1 + \int_{B_{r'}} |\bar{\phi} - \varphi^k| \, \mathrm{d}m_1 + \int_{B_{r'}^c} |\hat{\phi}| \, \mathrm{d}m_1 + \int_{B_{r'}^c} |\varphi^k| \, \mathrm{d}m_1 \leq 5\varepsilon. \tag{38}$$

Combining inequalities (37) and (38) (which also hold for m_2), we obtain:

$$\int_{\mathbb{R}^n} \phi \, \mathrm{d}(m_2 - m_1) \leq \Big(\int_{\mathbb{R}^n} \varphi^k \, \mathrm{d}(m_2 - m_1) \Big) + 16\varepsilon.$$

Maximizing the right-hand-side of the previous inequality with respect to k, then the left-hand-side with respect to $\phi \in$ 1-Lip($\mathbb{R}^n$), and making ε tend to 0, we obtain that:

$$d_1(m_1, m_2) \leq \sup_{k\in\mathbb{N}} \Big\{ \int_{\mathbb{R}^n} \varphi^k \, \mathrm{d}(m_2 - m_1) \Big\}.$$

The converse inequality is trivial, since the functions φ^k are 1-Lipschitz. □

References

[1] M. Annunziato and A. Borzì. Optimal control of probability density functions of stochastic processes. *Mathematical Modelling and Analysis*, 15(4):393–407, 2010.

[2] M. Annunziato and A. Borzì. Fokker-Planck-based control of a two-level open quantum system. *Math. Models Methods Appl. Sci.*, 23(11):2039–2064, 2013.

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

[4] J. F. Bonnans, L. Pfeiffer, and O. S. Serea. Sensitivity analysis for relaxed optimal control problems with final-state constraints. *Nonlinear Anal.*, 89:55–80, 2013.

[5] J. F. Bonnans and X. Tan. A model-free no-arbitrage price bound for variance options. *Applied Mathematics & Optimization*, 68(1):43–73, 2013.

[6] P. Cardaliaguet. Notes on Mean Field Games. 2012.

[7] P. Carpentier, J.-P. Chancelier, G. Cohen, M. De Lara, and P. Girardeau. Dynamic consistency for stochastic optimal control problems. *Annals of Operations Research*, 200(1):247–263, 2012.

[8] J. Feng and M. Katsoulakis. A comparison principle for Hamilton-Jacobi equations related to controlled gradient flows in infinite dimensions. *Archive for Rational Mechanics and Analysis*, 192(2):275–310, 2009.

[9] A. Figalli. Existence and uniqueness of martingale solutions for SDEs with rough or degenerate coefficients. *J. Funct. Anal.*, 254(1):109–153, 2008.

[10] W. Gangbo, T. Nguyen, and A. Tudorascu. Hamilton-Jacobi equations in the wasserstein space. *Methods and Applications of Analysis*, 15(2):155–184, 06 2008.

[11] M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.*, 6(3):221–252, 2006.

[12] N. V. Krylov. *Controlled diffusion processes*, volume 14 of *Applications of Mathematics*. Springer-Verlag, New York-Berlin, 1980. Translated from the Russian by A. B. Aries.

[13] A. Lachapelle, J. Salomon, and G. Turinici. Computation of mean field equilibria in economics. *Math. Models Methods Appl. Sci.*, 20(4):567–588, 2010.

[14] J.-M. Lasry and P.-L. Lions. Mean field games. *Jpn. J. Math.*, 2(1):229–260, 2007.

[15] M. Laurière and O. Pironneau. Dynamic programming for mean-field type control. *C. R. Math. Acad. Sci. Paris*, 352(9):707–713, 2014.

[16] B. Øksendal. *Stochastic Differential Equations: An Introduction with Applications.* Hochschultext / Universitext. Springer, 2003.

[17] L. Pfeiffer. Numerical methods for mean-field type optimal control problems. *SFB Report 2015-016*, 2015.

[18] L. Pfeiffer. Two approaches to constrained stochastic optimal control problems. *SFB Report 2015-007*, 2015.

[19] H. Pham. *Continuous-time Stochastic Control and Optimization with Financial Applications*, volume 61 of *Stochastic Modelling and Applied Probability.* Springer-Verlag, Berlin, 2009.

[20] S. Primak, V. Kontorovitch, and V. Lyandres. *Stochastic Methods and their Applications to Communications: Stochastic Differential Equations Approach.* Wiley, 2005.

[21] H. Risken. *The Fokker-Planck Equation: Methods of Solution and Applications*. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 1996.

[22] A. Shapiro. On a time consistency concept in risk averse multistage stochastic programming. *Oper. Res. Lett.*, 37(3):143–147, May 2009.

[23] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming*. MOS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA; Mathematical Optimization Society, Philadelphia, PA, second edition, 2014. Modeling and theory.

[24] C. Villani. *Optimal transport*, volume 338 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2009. Old and new.

[25] J. Yong and X. Y. Zhou. *Stochastic controls*, volume 43 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1999. Hamiltonian systems and HJB equations.