Robust mean-field games under entropy-based uncertainty

François Delarue * Pierre Lavigne †

dedicated to Alain Bensoussan

Abstract

In this article, we introduce a new class of entropy-penalized robust mean-field game problems in which the representative agent is opposed to Nature. The agent's objective is formulated as a min–max stochastic control problem, in which Nature distorts the reference probability measure at an entropic cost. As a consequence, the distribution of the continuum of agents represented by the player is given by the effective measure induced by Nature. Existence of a mean-field game equilibrium is established via a Schauder fixed point argument. To ensure uniqueness, we introduce a joint flat anti-monotonicity and displacement monotonicity condition, extending the classical Lasry–Lions monotonicity framework. Finally, we present two classes of N-player games for which the mean-field game limit yields ε-Nash equilibria.

Keywords: Risk-averse mean-field games, Quadratic backward stochastic differential equations, Monotonicity on the space of probability measures, ε-Nash equilibria, Entropic penalties

MSC2020. Primary: 49N80, 91A16; Secondary: 60H10

1 Introduction

In this article, we introduce a class of robust (or risk-averse) mean-field game problems in which, for a given mean-field configuration, a representative player optimizes against an adversarial agent, referred to as Nature, who acts on the probability distribution by emphasizing worst-case scenarios. For a fixed control of the representative player, the mean-field configuration is defined as the marginal law of the controlled state under the probability measure induced by Nature.

Formulation of the problem. Let $[0,T]$ be a finite time horizon and let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting a d-dimensional Brownian motion $W = (W_t)_{0\leq t\leq T}$ and an independent $\mathbb{R}^n$-valued random variable η, representing the initial state of the representative player. Here, $n \in \mathbb{N}^\star$ denotes the state dimension and $d \in \mathbb{N}^\star$ the dimension of the driving noise. The $\mathbb{P}$-complete filtration generated by (η, W) is denoted by $\mathbb{F} = (\mathcal{F}_t)_{0\leq t\leq T}$.

*Université Côte d'Azur, CNRS, Laboratoire J.A. Dieudonné, 06108 Nice, France. Email: francois.delarue@univ-cotedazur.fr

†Université Côte d'Azur, CNRS, Laboratoire J.A. Dieudonné, 06108 Nice, France. Email: pierre.lavigne@univ-cotedazur.fr

Given a mean field coupling μ, viewed as an element of the space $\mathcal{M}(\mathbb{R}^d)$ of non-negative (possibly non-normalized) measures on $\mathbb{R}^d$, the representative agent seeks to minimize a robust (risk-averse) objective functional of min–max type:

$$\begin{aligned} &\inf_{\psi\in\mathcal{A}}\sup_{q\in\mathcal{Q}}\mathcal{J}[\mu](\psi,q), \\ &\text{where}\quad \mathcal{J}[\mu](\psi,q) := \mathbb{E}\bigg[q_T g(X_T^{\psi},\mu)+\int_0^T q_s\ell(s,\psi_s)\,\mathrm{d}s\bigg]-\mathcal{S}(q). \end{aligned}\tag{MinMax[μ]}$$

Nature optimizes over a flow $q=(q_t)_{t\in[0,T]}$ of random non-normalized density process, while the representative player optimizes over a control ψ, with associated state process $X^{\psi}=(X_t^{\psi})_{t\in[0,T]}$. The functional $\mathcal{S}(q)$, defined precisely below, is referred to as a generalized entropy, as it extends the classical relative entropy with respect to $\mathbb{P}$. The functions $\ell:\Omega\times[0,T]\times\mathbb{R}^n\to\mathbb{R}$ and $g:\mathbb{R}^n\times\mathcal{M}(\mathbb{R}^n)\to\mathbb{R}$ denote respectively the running cost and the mean-field terminal cost of the representative player. In this setting, min–max equilibria are sought within the class of open-loop controls. The mean field consistency condition requires finding a measure μ such that, if (q,ψ) is a saddle point of (MinMax[μ]) (whose existence and uniqueness are ensured under the assumptions stated below), then μ coincides with the law of X_T^{ψ} under the effective measure $q_T\mathbb{P}$ selected by Nature.

Admissible processes q are assumed to admit the representation

$$q_t=e^{\int_0^t Y_s^{\star}\mathrm{d}s}\mathcal{E}_t\bigg(\int_0^{\cdot} Z_s^{\star}\cdot\mathrm{d}W_s\bigg),\quad t\in[0,T],\tag{1}$$

where $(\mathcal{E}_t(M)=\exp(M_t-\frac{1}{2}\langle M\rangle_t))_{t\in[0,T]}$ denotes the Doléans-Dade exponential of a local martingale $(M_t)_{t\in[0,T]}$ with $(\langle M\rangle_t)_{t\in[0,T]}$ as bracket, $Y^{\star}=(Y_t^{\star})_{t\in[0,T]}$ and $Z^{\star}=(Z_t^{\star})_{t\in[0,T]}$ are two $\mathbb{F}$-progressively measurable processes with values in $\mathbb{R}$ and $\mathbb{R}^d$, respectively, such that, $\mathbb{P}$-almost surely, $\int_0^T|Y_s^{\star}|\mathrm{d}s$ and $\int_0^T|Z_s^{\star}|^2\mathrm{d}s$ are finite. In particular, $(q_t)_{t\in[0,T]}$ satisfies the equation

$$q_t=1+\int_0^t q_sY_s^{\star}\mathrm{d}s+\int_0^t q_sZ_s^{\star}\cdot\mathrm{d}W_s,\quad t\in[0,T].\tag{2}$$

The generalized entropy of q is defined as

$$\mathcal{S}(q):=\mathbb{E}\left[\int_0^T q_sf^{\star}(s,Y_s^{\star},Z_s^{\star})\mathrm{d}s\right]\tag{3}$$

and is required to be finite. The function $f^{\star}\colon\Omega\times[0,T]\times\mathbb{R}\times\mathbb{R}^d\to\mathbb{R}$ is the 'convex dual' driver (as it it the Legendre-Fenchel transform of some f introduced in the assumptions below). When $f^{\star}(s,y^{\star},z^{\star})=\frac{1}{2}|z^{\star}|^2$ and $Y^{\star}\equiv 0$ in (1), $\mathcal{S}(q)$ coincides with the standard relative entropy $\mathbb{E}[q_T\ln(q_T)]$. The set of admissible process q is denoted $\mathcal{Q}$.

The admissible control set $\mathcal{A}$ consists of all $\mathbb{F}$-progressively measurable, $\mathbb{R}^n$-valued processes $\psi=(\psi_t)_{t\in[0,T]}$ satisfying

$$\mathcal{S}^{\star}(\psi)<+\infty,\qquad \mathcal{S}^{\star}(\psi):=\sup_{q\in\mathcal{Q}}\left\{\mathbb{E}\bigg[\int_0^T q_s|\psi_s|^2\,\mathrm{d}s\bigg]-\gamma\,\mathcal{S}(q)\right\},\tag{4}$$

where the constant $\gamma>0$ is specified below in accordance with the assumptions on the model coefficients. Although this definition may appear technical at first glance, it in fact captures the duality between the state of Nature and that of the representative player. In particular, when $\mathcal{S}$ coincides with the relative entropy with respect to $\mathbb{P}$, condition (4) echoes the Donsker–Varadhan duality formula and effectively enforces the existence of an exponential moment for $\int_0^T |\psi_s|^2\,\mathrm{d}s$. The role played by $\mathcal{S}^\star$ in the analysis of problem (MinMax[μ]) has been highlighted in our parallel work [16], where the present framework is adapted to robust mean field control. For a given control $\psi \in \mathcal{A}$, the state $X^\psi = (X_t^\psi)_{t\in[0,T]}$ of the representative player is the solution to

$$\mathrm{d}X_t = b(t, X_t, \psi_t)\mathrm{d}t + \sigma(t, \psi_t)\mathrm{d}W_t, \quad X_0 = \eta, \tag{5}$$

where the drift $b\colon \Omega \times [0,T] \times \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ and the volatility $\sigma\colon \Omega \times [0,T] \times \mathbb{R}^n \to \mathbb{R}^{n\times d}$ are possibly random. Implicitly, the coefficients b and σ are assumed to be $\mathbb{F}$-progressively measurable. Their precise regularity and structural assumptions are specified later in the paper. In particular, although the state equation (5) will ultimately be taken to be linear, we keep its general form here for expositional purposes.

We now provide a more detailed description of the equilibrium condition. Let q and ψ denote the optimal state of Nature and the optimal control of the representative agent respectively (see Theorem 2 for the existence and uniqueness of such a saddle point). In the mean-field framework, the representative agent is assumed to be typical of a continuum of statistically identical and independent agents, all playing the same game. In particular, the mean-field equilibrium condition requires that the coupling measure μ coincides with the law of the terminal state X_T^ψ but under the *effective* measure induced by Nature. Formally, the latter writes $q\mathbb{P} : \mathcal{F} \ni A \mapsto \mathbb{E}[q\mathbb{1}_A]$, and the law of X_T^ψ under $q\mathbb{P}$ is $q\mathbb{P} \circ (X_T^\psi)^{-1}$ (which we also write $(q\mathbb{P})_{X_T^\psi}$). This leads to the fixed point condition:

$$\mu = (q\mathbb{P}) \circ (X_T^\psi)^{-1}. \tag{MFG-eq}$$

We refer to such a mean-field equilibrium as *consistent*, meaning that the mean-field coupling observed by each agent at equilibrium is determined by the probability measure induced by Nature. In this approach, the agent's risk sensitivity is fully reflected in the mean-field coupling.

Alternatively, one may consider the classical fixed point condition

$$\mu = \mathbb{P} \circ (X_T^\psi)^{-1}.$$

However, under this prescription, the equilibrium is *inconsistent*: the representative agent remains risk-averse with respect to their own idiosyncratic noise, yet anticipates that the aggregate mean-field will materialize under the reference measure $\mathbb{P}$. This inconsistency arises from a mismatch between the agent's risk perception and the formation of the mean-field interaction.

In the rest of the article, the mean-field game problem thus consists in finding a triple $(q, \psi, \mu) \in \mathcal{Q} \times \mathcal{A} \times \mathcal{M}(\mathbb{R}^n)$ solving (MinMax[μ])-(MFG-eq), which may be summarized as

$$\mathcal{J}[\mu](\psi, q) = \inf_{\psi'\in\mathcal{A}} \sup_{q'\in\mathcal{Q}} \mathcal{J}[\mu](\psi', q'), \quad \mu = (q\mathbb{P}) \circ (X_T^\psi)^{-1}. \tag{MFG}$$

FBSDE formulation In the existing literature, mean-field game equilibria are typically characterized by a system of forward-backward partial differential equations (PDEs)—specifically two Fokker-Planck and Hamilton-Jacobi-Bellman equations, see

[23, 24]—or, from a probabilistic standpoint, a system of forward-backward stochastic differential equations (FBSDEs), see [13]. In this work, we adopt the probabilistic perspective. Under this formulation, the mean-field game FBSDE is viewed as the first-order system describing the optimal control problem of a representative agent interacting with a prescribed population distribution. The equilibrium is established via a fixed-point condition, which requires that this distribution coincides with the law induced by the agent's optimal strategy under Nature's density process. Within our framework, the FBSDE associated with the representative agent is coupled with another FBSDE, describing Nature's optimal control. The resulting equilibrium must therefore account for the simultaneous optimization of the agent and Nature, coupled with the aggregate consistency stemming from the mean-field interaction.

Given an arbitrary measure $\mu \in \mathcal{M}(\mathbb{R}^n)$, candidate for being a mean-field equilibrium, the two FBSDEs are driven by the following two pre-Hamiltonians, associated with the representative player and Nature respectively:

$$\begin{aligned} H(t,x,\psi,p,k,q) &:= q\ell(t,\psi) + p\cdot b(t,x,\psi) + \mathrm{Tr}(k\sigma^\top(t,\psi)), \\ F(t,q,y^\star,z^\star,y,z,\psi) &:= q\left(yy^\star + z\cdot z^\star - f^\star(t,y^\star,z^\star) + \ell(t,\psi)\right). \end{aligned} \tag{6}$$

Given $q \in \mathcal{Q}$, we say that a tuple (ψ,p,k,X) satisfies the first order condition (RP[μ]) for the representative player problem if (ψ,p,k,X) is a solution to

$$\begin{cases} -\mathrm{d}p_t &= \nabla_x H(t,X_t,\psi_t,p_t,k_t,q_t)\mathrm{d}t - k_t\mathrm{d}W_t, \quad p_T = q_T\nabla g(X_T^\psi,\mu), \\ \mathrm{d}X_t &= b(t,X_t,\psi_t)\mathrm{d}t + \sigma(t,\psi_t)\mathrm{d}W_t, \quad X_0 = \eta, \\ \psi_t &\in \arg\min_\alpha H(t,X_t,\alpha,p_t,k_t,q_t), \quad \mathrm{d}\mathbb{P}\otimes\mathrm{d}t\text{-a.e.} \end{cases}$$

The first equation is interpreted as the adjoint equation for the representative player, the second equation as the state equation, and the last equation as the optimality condition. Because the last equation creates a coupling between the first two equations, the system above is an FBSDE.

Given $\psi \in \mathcal{A}$, we say that a tuple (Y,Z,q) satisfies the first order condition (N[μ]) for Nature problem if (Y,Z,q) is a solution to

$$\begin{cases} -\mathrm{d}Y_t &= \partial_q F(t,q_t,Y_t^\star,Z_t^\star,Y_t,Z_t,\psi_t)\mathrm{d}t - Z_t\cdot\mathrm{d}W_t, \quad Y_T = g(X_T^\psi,\mu), \\ \mathrm{d}q_t &= q_tY_t^\star\mathrm{d}t + q_tZ_t^\star\cdot\mathrm{d}W_t, \quad q_0 = 1, \\ (Y_t^\star,Z_t^\star) &\in \arg\max_{(Y^{\star\prime},Z^{\star\prime})} F(t,q_t,Y^{\star\prime},Z^{\star\prime},Y_t,Z_t,\psi_t), \quad \mathrm{d}\mathbb{P}\otimes\mathrm{d}t\text{-a.e.} \end{cases}$$

The first equation is interpreted as the adjoint equation to Nature' state. The process Y describes the time instantaneous value of the representative player. Indeed, when ψ is optimal for the representative player, Y can be seen as the solution to a (risk averse) dynamic programming principle for the representative player. When $(Y^\star,Z^\star)\equiv 0$ (and thus $q\equiv 1$), and (say) $f^\star(s,0,0)=0$, we clearly recover the standard dynamic programming principle. The second equation describes the dynamics of the control variable, and the last equation is the optimality condition. This system of equations is also an FBSDE for the same reason as the previous system. The two FBSDEs are obviously coupled.

The FBSDE characterizing the Nash equilibria is given by (RP[μ])-(N[μ]) complemented with the equilibrium condition (MFG-eq).

In the benchmark case where $f^\star(s,y^\star,z^\star) = \frac{1}{2}|z^\star|^2$ and $f(s,y,z) = \frac{1}{2}|z|^2$, the FBSDE for Y becomes quadratic. This structure significantly complicates the solvability of the coupled systems, particularly when the terminal reward g is an unbounded function of the state. Under these conditions, the solvability of the quadratic BSDE satisfied by Y appears to lie beyond the scope of standard results in the literature, such as those found in [8, 9]. This difficulty prompted a dedicated investigation into the properties of this BSDE in our companion work [16], where we provide a tailored analysis leveraging the specific min-max structure of problem (MinMax[μ]). Broadly speaking, for a fixed measure μ, we establish existence and uniqueness by exploiting the underlying concavity in the variable q and convexity in the variable ψ. These structural properties are further utilized in our approach to resolving the fixed-point condition (MFG-eq).

A typical situation where these conditions arise is that of a financial investor seeking to maximize utility while being subject to trading costs. Assuming the market consists of n assets, each evolving according to the dynamics

$$\frac{\mathrm{d}S_t^i}{S_t^i} = c_t^i \mathrm{d}t + \sigma_t^i \mathrm{d}W_t,$$

where W is a d-dimensional noise process and the (potentially random) coefficients c^i and σ^i are of appropriate dimensions, the investor's self-financing portfolio X^ψ evolves according to the equation

$$\mathrm{d}X_t^\psi = \sum_{i=1}^n \psi_t^i \frac{\mathrm{d}S_t^i}{S_t^i}, \quad X_0 = 1,$$

where the initial condition is arbitrarily chosen to be unitary. The problem of the risk-averse investor, expressed in a min-max form, is given by

$$\sup_{q \in \mathcal{Q}} \inf_{\psi \in \mathcal{A}} \mathcal{J}(q, \psi), \tag{7}$$

where

$$\mathcal{J}(q,\psi) = \mathbb{E}^{\mathbb{Q}}\left[g(X_T^\psi) + \frac{1}{2}\int_0^T |\psi_s|^2 ds\right] - \gamma \mathcal{H}(\mathbb{Q}|\mathbb{P}), \quad \mathbb{Q} = q_T \mathbb{P},$$

with $q_T = \mathcal{E}_T(\int_0^\cdot Z_s^\star dW_s)$, and g denotes a payoff function and $\mathcal{H}(\mathbb{Q}|\mathbb{P})$ the relative entropy of $\mathbb{Q}$ with respect to $\mathbb{P}$.

Under this formulation, given a probability measure $\mathbb{Q}$ equivalent to $\mathbb{P}$, the investor optimizes the portfolio return while incurring a quadratic trading cost. Conversely, for a given investor strategy, Nature selects the worst-case probability measure $\mathbb{Q}$, subject to an entropic penalty. The parameter $\gamma > 0$ models the investor's level of risk aversion. In a mean-field regime, the payoff may take the form $g(X_T^\psi - \lambda \bar{\mu})$, where λ represents an interaction parameter and $\bar{\mu}$ denotes the mean of μ, representing the average wealth under the effective probability measure.

Literature. Mean-field games (MFGs) are competitive problems involving a continuum of agents whose interactions occur through a mean-field functional. They naturally arise as the limit of large, symmetric, and anonymous finite-player games, much like stochastic mean-field control problems. The theory was independently introduced in [21] and [23, 24] and has since been extensively developed [6, 7, 10, 11, 13, 14]. MFGs have found numerous applications, including economics and finance [1, 12, 19, 27], environmental studies [22, 25], and electricity markets [2]. For a comprehensive exposition of the theory, see the monograph [13].

The classical theory of mean field games (MFGs) primarily considers risk-neutral agents minimizing expected costs. Extending this framework to account for risk aversion has been the focus of several lines of research, each introducing different ways to capture the agents' attitudes toward uncertainty. In risk-sensitive MFGs [26, 28, 29], agents optimize exponential or variance-sensitive criteria, which penalize high variability in costs (see also [5] for the mean-field control analogue problem). More generally, risk-averse MFGs [15, 18, 20] incorporate abstract risk measures into the cost functional, allowing a wide class of preferences beyond variance-based penalties. Finally, robust or worst-case MFGs, often inspired by H^∞ control, introduce an explicit adversarial player (Nature) that acts to worsen the representative agent's outcome [4, 30], capturing ambiguity and model uncertainty. These different approaches reflect complementary ways to model agent sensitivity to risk and uncertainty in large populations.

The robust mean-field game studied here is closely related to the risk-sensitive framework: by the Gibbs-variational (Donsker–Varadhan) principle, minimizing an exponential cost is equivalent to a min–max game in which Nature selects a worst-case measure penalized by relative entropy. That said, unlike the standard risk-sensitive setting, we model Nature explicitly as acting on the weighting of events, which allows the equilibrium condition (MFG-eq) to be defined under the effective measure. To our knowledge, this explicit incorporation of the effective measure into the equilibrium definition is novel and provides a new perspective on robust mean-field equilibria.

Contributions. Beyond the model itself, which we find interesting, we contribute the following results. First, based on Schauder's theorem, we establish a general existence result for equilibria (see Theorem 14). Compared to standard MFGs, the proof requires careful treatment of Nature's state. Continuity with respect to the state of Nature is obtained by combining entropy-type inequalities, established under ad-hoc convexity assumptions in [16], with Pinsker's inequality, which ultimately controls the total variation of Nature's state. Next, we identify a general uniqueness criterion (see Proposition 18), which can be seen as an analogue of the Lasry–Lions monotonicity conditions (or displacement monotonicity in certain cases) in the risk-neutral setting. When the game is derived from a potential, this criterion reduces to a joint condition of flat concavity and displacement convexity for the potential. We also provide examples of non-potential games where the condition holds. Finally, we discuss the connection with finite-player models. This question is subtle, since the law of large numbers underlying the derivation of the mean-field model is perturbed by Nature's behavior. We present two approaches to show how the mean-field regime can emerge asymptotically and quantify to what extent the asymptotic equilibrium induces approximate equilibria (Lemmas 25, 29, and 30). As in the rich literature on convergence in MFG theory, the passage from finite-player games to the continuum remains a challenging problem and certainly calls for further study.

Organization of the article. The article is organized as follows. In Section 2, we useful introduce notations and definitions. In Section 3, we present the stochastic maximum principle recently obtained in [16], which allows us to handle (MinMax[μ]) when μ is fixed. Section 4 addresses the solvability of the mean-field game, providing both existence and uniqueness results. Finally, Section 5 is dedicated to the passage from two forms of finite-player game to the mean-field limit.

Several results in the text are directly taken from [16]; nevertheless, we have written the exposition to maintain a smooth and coherent flow.

2 Notations and definitions

This section introduces the notation used throughout the paper.

Spaces of random variables and processes. We work on the same filtered, complete probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ as in the introduction. When another probability measure is used, this will be indicated explicitly in the notation of the corresponding spaces of random variables or processes; for instance, we write $L^p(\cdot, \mathbb{Q})$, which is the first space defined in the list below.

For each $t \in [0,T]$, we denote by $L^0(\mathcal{F}_t, \mathbb{R}^d)$ the set of $\mathbb{R}^d$ valued and $\mathcal{F}_t$-measurable random variables (r.v.'s in short). And then, for each $p > 0$, we define the sets

- $L^p(\mathcal{F}_t, \mathbb{R}^d)$ of r.v.'s $X \in L^0(\mathcal{F}_t, \mathbb{R}^d)$ s.t. $\|X\|_{L^p(\mathcal{F}_t,\mathbb{R}^d)} := \mathbb{E}[|X|^p] < +\infty$;
- $L^\infty(\mathcal{F}_t, \mathbb{R}^d)$ of r.v.'s $X \in L^0(\mathcal{F}_t, \mathbb{R}^d)$ s.t. $\|X\|_{L^\infty(\mathcal{F}_t,\mathbb{R}^d)} := \operatorname*{ess\,sup}_{\omega\in\Omega} \sup_{i\in\{1,\dots,d\}} |X^i(\omega)| < +\infty$.

We denote by $L^0(\mathbb{F}, \mathbb{R}^d)$ the space of $\mathbb{F}$-progressively measurable random processes (r.p.'s in short) with values in $\mathbb{R}^d$, and by $S^0(\mathbb{F}, \mathbb{R}^d)$ the subset of $L^0(\mathbb{F}, \mathbb{R}^d)$ comprising processes with continuous trajectories. Given $p > 0$, we define the sets

- $L^p(\mathbb{F}, \mathbb{R}^d)$ of r.p.'s $X \in L^0(\mathbb{F}, \mathbb{R}^d)$ s.t. $\|X\|_{L^p(\mathbb{F},\mathbb{R}^d)} := \mathbb{E}\left[\left(\int_0^T |X_t|^p \mathrm{d}t\right)^{1/p}\right] < +\infty$,
- $M^p(\mathbb{F}, \mathbb{R}^d)$ of r.p.'s $X \in L^0(\mathbb{F}, \mathbb{R}^d)$ s.t. $\|X\|_{M^p(\mathbb{F},\mathbb{R}^d)} := \mathbb{E}\left[\left(\int_0^T |X_t|^2 \mathrm{d}t\right)^{p}\right] < +\infty$,
- $L^\infty(\mathbb{F}, \mathbb{R}^d)$ of r.p.'s $X \in L^0(\mathbb{F}, \mathbb{R}^d)$ s.t. $\|X\|_{L^\infty(\mathbb{F},\mathbb{R}^d)} := \sup_{t\in[0,T]} \|X_t\|_{L^\infty(\mathcal{F}_t,\mathbb{R}^d)} < +\infty$,
- $S^p(\mathbb{F}, \mathbb{R}^d)$ of r.p.'s $X \in S^0(\mathbb{F}, \mathbb{R}^d)$ s.t. $\|X\|_{S^p(\mathbb{F},\mathbb{R}^d)} := \mathbb{E}\left[\sup_{t\in[0,T]} |X_t|^p\right] < +\infty$.
- $D(\mathbb{F}, \mathbb{R}^d)$ of r.p.'s $X \in S^0(\mathbb{F}, \mathbb{R}^d)$ such that the family $(|X_\tau|)_\tau$, with τ running over the set of $[0,T]$-valued $\mathbb{F}$-stopping times, is uniformly integrable.

For more details on the class $D(\mathbb{F}, \mathbb{R}^d)$, we refer to [17, Definition 20]. For each space defined above, we omit the notation $\mathbb{R}^d$ when $d = 1$.

Spaces of measures. We call $\mathcal{P}(\mathbb{R}^n)$ the set of probability measures on $\mathbb{R}^n$, and $\mathcal{M}(\mathbb{R}^n)$ the set of finite non-negative measures on $\mathbb{R}^n$. For any $p \geq 1$, we define the sets

- $\mathcal{P}_p(\mathbb{R}^n)$ of $\mu \in \mathcal{P}(\mathbb{R}^n)$ s.t. $\int_{\mathbb{R}^n} |x|^p \mathrm{d}\mu(x) < +\infty$,
- $\mathcal{M}_p(\mathbb{R}^n)$ of $\mu \in \mathcal{M}(\mathbb{R}^n)$ s.t. $M_p(\mu) := \int_{\mathbb{R}^d} |x|^p \mathrm{d}\mu(x) < +\infty$.

For any finite measure $\mathbb{Q}$ on Ω and any measurable mapping $X \colon \Omega \to \mathbb{R}^n$, we denote by $\mathbb{Q}_X = \mathbb{Q} \circ X^{-1}$ the image measure of $\mathbb{Q}$ under X. And, for any non-negative measurable function f on Ω, not necessarily normalized, we denote by $f\mathbb{P}$ the associated (possibly non-normalized) measure $\mathbb{Q}$, defined by

$$\mathbb{Q}(A) = \int_A f \,\mathrm{d}\mathbb{P}, \qquad A \in \mathcal{F}.$$

Duality. By definition of $\mathcal{S}$ and $\mathcal{S}^\star$ in (3) and (4), we have for any $\mathbb{F}$-progressively measurable processes q and ζ, valued in $\mathbb{R}$,

$$\mathcal{S}(q) + \mathcal{S}^\star(\zeta) \geq \frac{1}{\gamma}\mathbb{E}\left[\int_0^T q_s|\zeta_s|^2 \mathrm{d}s\right], \tag{8}$$

where $\mathcal{S}(q)$ and $\mathcal{S}^\star(\zeta)$ might take infinite values. Equality holds whenever

$$q \in \operatorname*{arg\,max}_{q' \in \mathcal{Q}} \left\{\mathbb{E}\left[\int_0^T q_s|\zeta_s|^2 \mathrm{d}s\right] - \gamma\mathcal{S}(q)\right\}.$$

Miscellaneous. For finite-dimensional vectors x and y (in the same space), $x \cdot y$ denotes their scalar product. We also define the entropy function $\mathrm{Ent} \colon \mathbb{R}_+ \to \mathbb{R}$:

$$\mathrm{Ent}(x) := x(\ln(x) - 1). \tag{9}$$

3 Robust control within a fixed environment

In this section, we address the problem (MinMax[μ]) when $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$ is fixed. The set-up is clarified in Subsection 3.1. In Subsection 3.2, we present an existence and uniqueness result to (MinMax[μ]), which directly follows from our companion work [16]. In Subsection 3.3, we derive stability estimates, which are key in the analysis of the mean-field game carried out in the next section.

3.1 Set-up

The assumptions are mostly derived from the analysis introduced in [16].

We use repeatedly the notion of *progressive-measurable* field. For a metric space $(\mathcal{X}, d)$ and an integer $k \in \mathbb{N}^*$, a random field $\mathcal{F} : \Omega \times [0,T] \times \mathcal{X} \to \mathbb{R}^k$ is progressively-measurable if, for any $t \in [0,T]$, its restriction to $\Omega \times [0,t] \times \mathcal{X}$ is $\mathcal{F}_t \otimes \mathcal{B}([0,t]) \otimes \mathcal{B}(\mathcal{X})/\mathcal{B}(\mathbb{R}^k)$ measurable.

Throughout, L and r are two constants, with $L > 0$ and $r \in \{0,1\}$. The assumptions hold true for any fixed $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$, and the constants L and r are assumed to be independent of μ. In fact, the only assumption in which μ appears is (A7). Therein, we pay special attention to introduce a tailored notation for the constants that genuinely depend on μ.

A1 *Initial condition and drift.* The initial condition η belongs to $L^\infty(\mathcal{F}_0, \mathbb{R}^n)$, i.e. $\|\eta\|_{L^\infty(\mathcal{F}_0,\mathbb{R}^n)} < +\infty$, and the drift $b\colon \Omega \times [0,T] \times \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}^n$ is linear, i.e.,

$$b(t, x, \psi) = a_t + b_t x + c_t \psi,$$

with a, b and c in $L^\infty(\mathbb{F}, \mathbb{R}^n)$, $L^\infty(\mathbb{F}, \mathbb{R}^{n\times n})$ and $L^\infty(\mathbb{F}, \mathbb{R}^{n\times n})$, and $\|a\|_{L^\infty(\mathbb{F},\mathbb{R}^n)} + \|b\|_{L^\infty(\mathbb{F},\mathbb{R}^{n\times n})} + \|c\|_{L^\infty(\mathbb{F},\mathbb{R}^{n\times n})} \le L$.

A2 *Volatility.* The volatility $\sigma\colon \Omega \times [0,T] \times \mathbb{R}^n \to \mathbb{R}^{n\times d}$ is linear, the $n \times d$ entries of the matrix σ being of the form

$$(\sigma(t,\psi))_{i,j} = (\nu_t)_{i,j} + r(\sigma_t)_{i,j,k}\psi_k,$$

for $(i,j,k) \in \{1,\dots,n\} \times \{1,\dots,d\} \times \{1,\dots,n\}$. Above, ν and σ are in $L^\infty(\mathbb{F}, \mathbb{R}^{n\times d})$ and $L^\infty(\mathbb{F}, \mathbb{R}^{n\times d})$, and $\|\nu\|_{L^\infty(\mathbb{F},\mathbb{R}^{n\times d})} + \|\sigma\|_{L^\infty(\mathbb{F},\mathbb{R}^{n\times d\times n})} \le L$.

A3 *Driver.* The driver $f\colon \Omega \times [0,T] \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ is progressively-measurable, and twice continuously differentiable and convex in its last two last arguments. There exist two positive constants α, β such that, almost surely in ω and almost everywhere in t,

$$f(t,y,z) \le |f_t^0| + \alpha|y| + \frac{\beta}{2}|z|^2, \quad (y,z) \in \mathbb{R} \times \mathbb{R}^d,$$

where $f^0 := f(0,0) \in L^\infty(\mathbb{F})$. The second order derivatives in y and z are bounded by L.

A4 *Dual driver.* We call $f^\star\colon \Omega \times [0,T] \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$ the Fenchel transform of the driver f with respect to its variables (y,z),

$$f^\star(t, y^\star, z^\star) := \sup_{(y,z)\in\mathbb{R}\times\mathbb{R}^d} \left\{ \langle (y^\star, z^\star), (y,z) \rangle - f(t,y,z) \right\}.$$

It is shown in [16] that $f^\star$ is progressively-measurable.

A5 *Running cost.* The cost $\ell\colon \Omega \times [0,T] \times \mathbb{R}^n \to \mathbb{R}$ is progressively-measurable and strongly convex, twice differentiable, and has a quadratic growth with respect to the control variable:

$$\big(\nabla_\psi \ell(t,\psi) - \nabla_\psi \ell(t,\psi')\big) \cdot (\psi - \psi') \ge \frac{1}{L}|\psi - \psi'|^2, \quad |\nabla_\psi^2 \ell(t,\psi)| \le L,$$

and $|\ell(t,0)| \le L$ for any $t \in [0,T]$ and $\psi, \psi' \in \mathbb{R}^n$.

A6 *Coefficients.* We fix the coefficient γ in (4) to be given by

$$\gamma = 2Le^{\alpha T}\|\Gamma\|_{L^\infty}\left(\|\nu\|_{L^\infty} + 6e^{\alpha T}\|\sigma\|_{L^\infty}\right).$$

When $r = 0$ we assume that the coefficients satisfy the condition

$$2\beta e^{\alpha T} L \|\Gamma\|_{L^\infty(\mathbb{F},\mathbb{R}^n)} \|\nu\|_{L^\infty(\mathbb{F},\mathbb{R}^{n\times d})} < 1,$$

where Γ is the solution to

$$\frac{\mathrm{d}}{\mathrm{d}t}\Gamma_t = b_t\Gamma_t, \quad t \in [0,T], \quad \Gamma_0 = I_n,$$

with I_n standing for the $n \times n$ identity matrix.

A7 *Terminal cost.* We assume that $g : \mathbb{R}^n \times \mathcal{M}_{2-r}(\mathbb{R}^n) \to \mathbb{R}$ is convex and twice differentiable in x and, for any real $C \geq 0$, there exists a constant $L_C \geq 0$ such that, for any $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$ with $M_{2-r}(\mu) \leq C$,

$$\begin{aligned} -L_C\,(1+|x|) \leq g(x,\mu) &\leq L_C\left(1+|x|^{2-r}\right), \\ |\nabla_x g(x,\mu)| &\leq L_C\left(1+|x|^{1-r}\right), \\ |\nabla_x^2 g(x,\mu)| &\leq L_C. \end{aligned} \tag{10}$$

Remark 1. The growth condition (A3) implies that $f^\star(t, y^\star, z^\star) = +\infty$ if $|y^\star| > \alpha$. In particular, $Y^\star$ in (2) is necessarily bounded by α if $\mathcal{S}(q)$ is finite (as it is required). Moreover, it is proven in [16, Proposition 23] that there exists a constant C, only depending on the parameters in the standing assumption such that, for $q \in \mathbb{Q}$,

$$\mathbb{E}[\mathrm{Ent}(q)] \leq C(1+\mathcal{S}(q)), \tag{11}$$

where Ent is given by (9).

In A6, the choice of γ is dictated by the analysis carried out in [16].

In A7, g is assumed to be deterministic (contrary to the other coefficients). In fact, g could be allowed to be random in some of the statements, but it is typically deterministic in the whole discussion on uniqueness and on the N-player approximation.

3.2 Solvability of the robust control problem

Following [16], we study the optimization problem (MinMax[μ]) (for a fixed $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$) via the associated Pontryagin system. Under the standing assumptions, the cost functionals are convex with respect to the state variable X and concave with respect to the control variable q. As a consequence, the Pontryagin principle yields a full characterization of the saddle point. This constitutes one of the main results of [16].

As already explained in Introduction, the Pontryagin system takes the form of two FBSDEs, each backward equation being driven by the derivative (with respect to the corresponding coordinate) of the corresponding pre-Hamiltonian introduced in (6).

Given $q \in \mathcal{Q}$, the first order condition (RP[μ]) for the representative player problem writes in the form of a foward-backward system, with (ψ, p, k, X) as unknown:

$$\begin{cases} -\mathrm{d}p_t &= \nabla_x H(t, X_t, \psi_t, p_t, k_t, q_t)\mathrm{d}t - k_t \mathrm{d}W_t \\ &= b_t^\top p_t \mathrm{d}t - k_t \mathrm{d}W_t, \qquad\qquad p_T = q_T \nabla_x g(X_T^\psi, \mu), \\ \mathrm{d}X_t &= b(t, X_t, \psi_t)\mathrm{d}t + \sigma(t,\psi_t)\mathrm{d}W_t, \qquad X_0 = \eta, \\ \psi_t &\in \arg\min_\alpha H(t, X_t, \alpha, p_t, k_t, q_t), \\ \text{i.e. } 0 &= q_t \nabla_\psi \ell(t,\psi_t) + p_t \cdot c_t + r\mathrm{Tr}(\sigma_t^\top k_t), \quad \mathrm{d}\mathbb{P}\otimes \mathrm{d}t\text{-a.s.}, \end{cases} \tag{RP[μ]}$$

where we denote by convention

$$\mathrm{Tr}\left(\sigma_t^\top k_t\right) = \left(\sum_{i=1}^n \sum_{j=1}^d (\sigma_t)_{i,j,\ell}(k_t)_{i,j}\right)_{\ell=1,\dots,d}. \tag{12}$$

Solutions (ψ, p, k, X) to (RP[μ]) are sought within the space

$$\mathscr{A} := \mathcal{A} \times D(\mathbb{F}) \times (\underset{\beta\in(0,1)}{\cap} M^\beta(\mathbb{F}, \mathbb{R}^d)) \times S^{2-r}(\mathbb{F}, \mathbb{R}^n, \mathbb{Q}), \tag{13}$$

where $\mathbb{Q}$ in the first line is the equivalent measure associated to q, i.e., $\mathbb{Q} = q\mathbb{P}$.

Given $\psi \in \mathcal{A}$, the first order condition (N[μ]) for the nature problem writes in the form of another forward-backward system, with (Y, Z, q) as unknown:

$$\begin{cases} -\mathrm{d}Y_t & = \partial_q F(t, q_t, Y_t^\star, Z_t^\star, Y_t, Z_t, \psi_t)\mathrm{d}t - Z_t \cdot \mathrm{d}W_t \\ & = (f(t, Y_t, Z_t) + \ell(t, \psi_t))\mathrm{d}t - Z_t \cdot \mathrm{d}W_t, \qquad Y_T = g(X_T^\psi, \mu), \\ \mathrm{d}q_t & = q_t Y_t^\star \mathrm{d}t + q_t Z_t^\star \cdot \mathrm{d}W_t, \qquad q_0 = 1, \\ (Y_t^\star, Z_t^\star) & \in \underset{(Y^{\star\prime}, Z^{\star\prime})}{\arg\max} F(t, q_t, Y^{\star\prime}, Z^{\star\prime}, Y_t, Z_t, \psi_t) \\ \Leftrightarrow (Y_t^\star, Z_t^\star) & = (\partial_y f(t, Y_t, Z_t), \nabla_z f(t, Y_t, Z_t)), \qquad \mathrm{d}\mathbb{P} \otimes \mathrm{d}t\text{-a.s.} \end{cases} \tag{N[μ]}$$

Solutions (q, Y, Z) to (N[μ]) are sought in the space

$$\mathscr{Q} := \mathcal{Q} \times D(\mathbb{F}, \mathbb{Q}) \times (\underset{\beta\in(0,1)}{\cap} M^\beta(\mathbb{F}, \mathbb{R}^d, \mathbb{Q})). \tag{14}$$

Here is now the main statement of [16] regarding the inf-sup mean-field stochastic control problem (MinMax[μ]).

Theorem 2. *Let $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$. Then, there exists a unique saddle point $(\psi, q) \in \mathcal{A} \times \mathcal{Q}$ to Problem (MinMax[μ]), i.e.*

$$\min_{\psi'\in\mathcal{A}} \max_{q'\in\mathcal{Q}} \mathcal{J}(q', \psi') = \max_{q'\in\mathcal{Q}} \min_{\psi'\in\mathcal{A}} \mathcal{J}(q', \psi') = \mathcal{J}(q, \psi).$$

Moreover, if a pair $(\psi, q) \in \mathcal{A} \times \mathcal{Q}$ is a solution to the problem (MinMax[μ]), then the tuples (ψ, p, k, X), obtained by solving in $\mathscr{A}$ the two decoupled equations in (RP[μ]), and (q, Y, Z), obtained by solving in $\mathscr{Q}$ the two decoupled equations in (N[μ]), satisfy the optimality conditions in (RP[μ]) and (N[μ]) respectively. Conversely, if $(\psi, p, k, X, q, Y, Z) \in \mathscr{A} \times \mathscr{Q}$ is the solution to (RP[μ])-(N[μ]), then the pair $(q, \psi) \in \mathcal{Q} \times \mathcal{A}$ is a solution to the problem (MinMax[μ]).

In the rest of the article, we denote the unique saddle point by (q^μ, ψ^μ). Accordingly, the solution to (RP[μ]) is denoted by $(\psi^\mu, p^\mu, k^\mu, X^\mu)$ and the solution to (N[μ]) is denoted by (q^μ, Y^μ, Z^μ). The representatives $(Y^\star, Z^\star)$ of q^μ in (2) are denoted by $(Y^{\star,\mu}, Z^{\star,\mu})$.

The fact that the cost $\mathbb{E}[q_T^\mu g(X_T^{\tilde{\mu}}, \tilde{\mu})]$ is well-defined is the consequence of A7 and of the fact that $X^\mu \in S^{2-r}(\mathbb{F}, \mathbb{R}^n, \mathbb{Q})$. Generally speaking, the latter is a consequence of the following lemma, which corresponds to [16, Lemma 43]:

Lemma 3. *Let $(q, \psi) \in \mathcal{Q} \times \mathcal{A}$. Then, X^ψ belongs to $S^{2-r}(\mathbb{F}, \mathbb{Q}, \mathbb{R}^n)$, where $r \in \{0, 1\}$ is as in A2, and there exists a constant C, independent of q and ψ, such that*

$$\mathbb{E}\left[q_T |X_T^*|^{2-r}\right] \le C\left(1 + \mathcal{S}(q) + \mathcal{S}^\star(\psi)\right).$$

3.3 Stability Estimates

The next result is taken from [16]. It is in fact part of the proof on which the derivation of the stochastic minimum principle for the problem (RP[μ]) relies.

Lemma 4. *Let $\mu, \tilde{\mu} \in \mathcal{M}_{2-r}(\mathbb{R}^n)$. Then,*

$$\mathbb{E}\left[p_T^{\mu} \cdot \left(X_T^{\tilde{\mu}} - X_T^{\mu}\right)\right] = -\mathbb{E}\left[\int_0^T q_s^{\mu} \nabla_{\psi} \ell\left(s, \psi_s^{\mu}\right) \cdot \left(\psi_s^{\tilde{\mu}} - \psi_s^{\mu}\right) \mathrm{d}s\right], \tag{15}$$

which implicitly implies that the expectations right above are well-defined.

Proof. This is the penultimate display in the proof of [16, Lemma 40]. □

We now recall the following result from convex analysis.

Lemma 5. *There exists a constant $c > 0$, only depending on the parameters in A1-A7, such that, for any $t \in [0, T]$, $y \in \mathbb{R}, z \in \mathbb{R}^d$, and $h_y \in \mathbb{R}, h_z \in \mathbb{R}^d$,*

$$\begin{aligned}
&f^{\star}\left(t, \partial_y f(t, y, z) + h_y, \partial_z f(t, y, z) + h_z\right) \\
&\geq f^{\star}(\partial_y f(t, y, z), \partial_z f(t, y, z)) + y \cdot h_y + z \cdot h_z + c\left(|h_y|^2 + |h_z|^2\right).
\end{aligned}$$

Proof. By Fenchel's duality and A3-A4, we have

$$f^{\star}\left(t, \partial_y f(t, y, z), \partial_z f(t, y, z)\right) + f(t, y, z) = y \partial_y f(t, y, z) + z \cdot \partial_z f(t, y, z).$$

And, for any $k_y \in \mathbb{R}$ and $k_z \in \mathbb{R}^d$,

$$\begin{aligned}
&f^{\star}\left(t, \partial_y f(t, y, z) + h_y, \partial_z f(t, y, z) + h_z\right) \\
&\geq (y + k_y)\left(\partial_y f(t, y, z) + h_y\right) + (z + k_z) \cdot \left(\partial_z f(t, y, z) + h_z\right) - f(t, y + k_y, z + k_z) \\
&\geq f^{\star}\left(t, \partial_y f(t, y, z), \partial_z f(t, y, z)\right) + y h_y + z \cdot h_z \\
&\quad - \left[f(t, y + k_y, z + k_z) - f(t, y, z) - k_y \partial_y f(t, y, z) - k_z \cdot \partial_z f(t, y, z)\right] \\
&\quad + k_y h_y + k_z \cdot h_z.
\end{aligned}$$

By A3, there exists a constant $C > 0$ such that the penultimate term on the above display is bounded by $C(|k_y|^2 + |k_z|^2)$. We deduce that

$$\begin{aligned}
&f^{\star}\left(t, \partial_y f(t, y, z) + h_y, \partial_z f(t, y, z) + h_z\right) \\
&\geq f^{\star}\left(t, \partial_y f(t, y, z), \partial_z f(t, y, z)\right) + y h_y + z \cdot h_z \\
&\quad + k_y h_y + k_z \cdot h_z - C(|k_y|^2 + |k_z|^2).
\end{aligned}$$

Choosing $k_y = h_y/(2C)$ and $k_z = h_z/(2C)$, we get

$$\begin{aligned}
&f^{\star}\left(t, \partial_y f(t, y, z) + h_y, \partial_z f(t, y, z) + h_z\right) \\
&\geq f^{\star}\left(t, \partial_y f(t, y, z), \partial_z f(t, y, z)\right) + y h_y + z \cdot h_z + \frac{1}{4C^2}\left(|h_y|^2 + |h_z|^2\right).
\end{aligned}$$

This completes the proof. □

Here is now the main result of this subsection:

Proposition 6. *There exists a constant $c > 0$, only depending on the parameters in A1-A7, such that, the following two inequalities hold true, for any $\mu, \tilde{\mu} \in \mathcal{M}_{2-r}(\mathbb{R}^n)$:*

$$\begin{aligned} &\mathbb{E}\left[q_T^{\tilde{\mu}}\left(g(X_T^{\mu},\tilde{\mu}) - g(X_T^{\mu},\mu)\right) + q_T^{\mu}\left(g(X_T^{\tilde{\mu}},\mu) - g(X_T^{\tilde{\mu}},\tilde{\mu})\right)\right] \\ &\geq c\mathbb{E}\left[\int_0^T \left(q_t^{\mu} + q_t^{\tilde{\mu}}\right)\left(|\psi_t^{\tilde{\mu}} - \psi_t^{\mu}|^2 + |Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right], \end{aligned} \tag{16}$$

and

$$\begin{aligned} &\mathbb{E}\left[\left(q_T^{\mu} - q_T^{\tilde{\mu}}\right)\left(g(X_T^{\mu},\mu) - g(X_T^{\tilde{\mu}},\tilde{\mu})\right)\right] \\ &\geq \mathbb{E}\left[\left(q_T^{\mu}\nabla_x g(X_T^{\mu},\mu) - q_T^{\tilde{\mu}}\nabla_x g(X_T^{\tilde{\mu}},\tilde{\mu})\right)\cdot(X_T^{\mu} - X_T^{\tilde{\mu}})\right] \\ &\quad + c\mathbb{E}\left[\int_0^T \left(q_t^{\mu} + q_t^{\tilde{\mu}}\right)\left(|\psi_t^{\tilde{\mu}} - \psi_t^{\mu}|^2 + |Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right]. \end{aligned} \tag{17}$$

The fact that the expectations on the first line of (16) and on the first and second lines of (17) are well-defined is a consequence of Lemma 3.

Proof. The proof is divided in three steps.

Step 1. From Step 3 in the proof of [16, Lemma 32], we have

$$\begin{aligned} &\mathbb{E}\left[q_T^{\mu} g(X_T^{\mu},\mu) + \int_0^T q_t^{\mu}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\mu}) \\ &\geq \mathbb{E}\left[q_T^{\tilde{\mu}} g(X_T^{\mu},\mu) + \int_0^T q_t^{\tilde{\mu}}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\tilde{\mu}}) + \lim_{A\to\infty}\mathbb{E}\left[\int_0^{T\wedge\tau_A} q_t^{\tilde{\mu}}\Delta f_t^{\star}\mathrm{d}t\right], \end{aligned}$$

where $(\tau_A)_{A>0}$ is a collection of stopping time that converges almost surely to T (as A tends to $+\infty$), and

$$\begin{aligned} \Delta f_t^{\star} = f^{\star}&\left(t, Y_t^{\star,\tilde{\mu}}, Z_t^{\star,\tilde{\mu}}\right) - f^{\star}\left(t, Y_t^{\star,\mu}, Z_t^{\star,\mu}\right) \\ &- \left(Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}\right)Y_t^{\mu} - \left(Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}\right)\cdot Z_t^{\mu}. \end{aligned}$$

By Lemma 5, we obtain

$$\begin{aligned} &\mathbb{E}\left[q_T^{\mu} g(X_T^{\mu},\mu) + \int_0^T q_t^{\mu}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\mu}) \\ &\geq \mathbb{E}\left[q_T^{\tilde{\mu}} g(X_T^{\mu},\mu) + \int_0^T q_t^{\tilde{\mu}}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\tilde{\mu}}) \\ &\quad + c\mathbb{E}\left[\int_0^T q_t^{\tilde{\mu}}\left(|Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right]. \end{aligned}$$

And then, by strong convexity of ℓ (see A5), we obtain (for a new value of c),

$$\begin{aligned} &\mathbb{E}\left[q_T^{\mu} g(X_T^{\mu},\mu) + \int_0^T q_t^{\mu}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\mu}) \\ &\geq \mathbb{E}\left[q_T^{\tilde{\mu}} g(X_T^{\mu},\mu) + \int_0^T q_t^{\tilde{\mu}}\ell(t,\psi_t^{\tilde{\mu}})\mathrm{d}t + \int_0^t q_t^{\tilde{\mu}}\nabla_{\psi}\ell(t,\psi_t^{\tilde{\mu}})\cdot\left(\psi_t^{\mu} - \psi_t^{\tilde{\mu}}\right)\mathrm{d}t\right] - \mathcal{S}(q^{\tilde{\mu}}) \\ &\quad + c\mathbb{E}\left[\int_0^T q_t^{\tilde{\mu}}\left(|\psi_t^{\tilde{\mu}} - \psi_t^{\mu}|^2 + |Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right]. \end{aligned}$$

By Lemma 4,

$$\begin{aligned}
&\mathbb{E}\left[q_T^{\mu} g(X_T^{\mu},\mu) + \int_0^T q_t^{\mu}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\mu}) \\
&\geq \mathbb{E}\left[q_T^{\tilde\mu} g(X_T^{\mu},\mu) - q_T^{\tilde\mu}\nabla_x g(X_T^{\tilde\mu},\tilde\mu)\cdot(X_T^{\mu}-X_T^{\tilde\mu}) + \int_0^T q_t^{\tilde\mu}\ell(t,\psi_t^{\tilde\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\tilde\mu}) \qquad (18)\\
&\quad + c\mathbb{E}\left[\int_0^T q_t^{\tilde\mu}\left(|\psi_t^{\tilde\mu}-\psi_t^{\mu}|^2 + |Y_t^{\star,\tilde\mu}-Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde\mu}-Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right].
\end{aligned}$$

Step 2. We derive the first claim. By convexity of g in the first variable, we get the following bound for the first term on the second line of (18)

$$g(X_T^{\mu},\tilde\mu) \geq g(X_T^{\tilde\mu},\tilde\mu) + \nabla_x g(X_T^{\tilde\mu},\tilde\mu)\cdot(X_T^{\mu}-X_T^{\tilde\mu}),$$

from which we deduce that

$$\begin{aligned}
&\mathbb{E}\left[q_T^{\mu} g(X_T^{\mu},\mu) + \int_0^T q_t^{\mu}\ell(t,\psi_t^{\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\mu}) \\
&\geq \mathbb{E}\left[q_T^{\tilde\mu}\left(g(X_T^{\mu},\mu) - g(X_T^{\mu},\tilde\mu) + g(X_T^{\tilde\mu},\tilde\mu)\right) + \int_0^T q_t^{\tilde\mu}\ell(t,\psi_t^{\tilde\mu})\mathrm{d}t\right] - \mathcal{S}(q^{\tilde\mu}) \\
&\quad + c\mathbb{E}\left[\int_0^T q_t^{\tilde\mu}\left(|\psi_t^{\tilde\mu}-\psi_t^{\mu}|^2 + |Y_t^{\star,\tilde\mu}-Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde\mu}-Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right].
\end{aligned}$$

By exchanging the roles of μ and μ' and by adding the resulting two inequalities, we obtain

$$\begin{aligned}
0 &\geq \mathbb{E}\left[q_T^{\tilde\mu}\left(g(X_T^{\mu},\mu) - g(X_T^{\mu},\tilde\mu)\right) + q_T^{\mu}\left(g(X_T^{\tilde\mu},\tilde\mu) - g(X_T^{\tilde\mu},\mu)\right)\right] \\
&\quad + c\mathbb{E}\left[\int_0^T \left(q_t^{\mu}+q_t^{\tilde\mu}\right)\left(|\psi_t^{\tilde\mu}-\psi_t^{\mu}|^2 + |Y_t^{\star,\tilde\mu}-Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde\mu}-Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right],
\end{aligned}$$

which completes the proof of (16).

Step 3. We now derive the second claim. We come back to (18). Exchanging the roles of μ and $\tilde\mu$ therein, and then summing the resulting two inequalities, we get

$$\begin{aligned}
&\mathbb{E}\left[\left(q_T^{\mu}-q_T^{\tilde\mu}\right)\left(g(X_T^{\mu},\mu) - g(X_T^{\tilde\mu},\tilde\mu)\right)\right] \\
&\geq \mathbb{E}\left[\left(q_T^{\mu}\nabla_x g(X_T^{\mu},\mu) - q_T^{\tilde\mu}\nabla_x g(X_T^{\tilde\mu},\tilde\mu)\right)\cdot(X_T^{\mu}-X_T^{\tilde\mu})\right] \\
&\quad + c\mathbb{E}\left[\int_0^T \left(q_t^{\mu}+q_t^{\tilde\mu}\right)\left(|\psi_t^{\tilde\mu}-\psi_t^{\mu}|^2 + |Y_t^{\star,\tilde\mu}-Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde\mu}-Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right].
\end{aligned}$$

This completes the proof. $\square$

4 Mean-field games

This section is devoted to the study of the mean-field game problem (MFG). In Subsection 4.1, we define the notion of a mean-field game equilibrium and introduce the topology underlying the existence result, together with additional assumptions on the interaction mapping g. In Subsection 4.2, we derive uniform estimates on

(q^μ, ψ^μ) with respect to μ, which are required to apply Schauder's fixed point theorem. Subsection 4.3 is devoted to our main existence result, stated in Theorem 14. Finally, in Subsection 4.4, we establish a uniqueness result under a joint flat non-increasing and displacement non-decreasing condition on the mapping g, as defined in Definition 15; see Proposition 18.

4.1 Definition of an equilibrium

Definition 7. *For r as in (A2), we say that $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$ is an equilibrium to the robust mean-field game set over (MinMax[μ]) if the unique saddle point (q^μ, ψ^μ) of (MinMax[μ]) satisfies*

$$\mu = (q_T^\mu \mathbb{P})_{X_T^\mu}.$$

Topology. Below, we study existence and uniqueness separately. For this, we equip the space of non-negative measures with the narrow topology, a sequence $(\mu_k)_{k\geq 1}$ in $\mathcal{M}(\mathbb{R}^k)$ converging narrowly to some μ in $\mathcal{M}(\mathbb{R}^n)$ if, for any bounded and continuous function f on $\mathbb{R}^n$, it holds

$$\lim_{k\to+\infty} \int_{\mathbb{R}^n} f(x)\mathrm{d}\mu_k(x) = \int_{\mathbb{R}^n} f(x)\mathrm{d}\mu(x).$$

In fact, we are only interested in elements $\mu \in \mathcal{M}(\mathbb{R}^n)$ whose mass $\mu(\mathbb{R}^n)$ is less than $\exp(\alpha T)$. The reason is that, for any $\mu \in \mathcal{M}(\mathbb{R}^n)$, $\mathbb{E}[q_T^\mu] \leq \exp(\alpha T)$. In this regard, it is important to remember that Prokhorov's theorem extends easily to non-negative measures with a mass less than a fixed constant:

Lemma 8. *Let $\mathcal{C}$ be a subset of $\mathcal{M}(\mathbb{R}^n)$ such that*

$$\sup_{\mu\in\mathcal{C}} \mu(\mathbb{R}^n) < +\infty.$$

Then, $\mathcal{C}$ is relatively compact for the narrow topology if it is tight, i.e., for any $\varepsilon > 0$, there exists a compact subset $K \subset \mathbb{R}^n$ such that

$$\sup_{\mu\in\mathcal{C}} \mu(\mathbb{R}^n \setminus K) \leq \varepsilon.$$

In what follows (see the forthcoming condition A9), we require the function g to be continuous in μ with respect to the narrow topology, but only on bounded subsets of $\mathcal{M}_{2-r}(\mathbb{R}^n)$, i.e., on subsets of the form

$$\mathcal{B}_{\mathcal{M}_{2-r}}(C) := \left\{ \mu \in \mathcal{M}_{2-r}(\mathbb{R}^n) : \int_{\mathbb{R}^n} \left(1 + |x|^{2-r}\right) \mathrm{d}\mu(x) \leq C \right\}.$$

This notion is motivated by the following standard lemma:

Lemma 9. *Let $h : \mathbb{R}^n \to \mathbb{R}$ be a continuous function such that, for some $c > 0$ and $\eta \in (0, 2-r)$, $|h(x)| \leq c(1 + |x|^{2-r-\eta})$. Then, for any $C > 0$, the function*

$$\mathcal{B}_{\mathcal{M}_{2-r}(C)} \ni \mu \mapsto \int_{\mathbb{R}^n} h(x)\mathrm{d}\mu(x)$$

is continuous for the narrow topology.

As it is well-known, the result becomes false when $\eta = 0$. In this case, continuity just holds but on subsets of $\mathcal{M}_{2-r}(\mathbb{R}^n)$ that are uniformly integrable. In our framework, we are not able to prove that, in for generality, the collection $((q_T^\mu \mathbb{P})_{X_T^\mu})_{\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)}$ is uniformly integrable, which explains why continuity with respect to the narrow topology is required on larger subsets (and thus leads to less general examples, as clearly illustrated by the above lemma).

We thus require further regularity properties on the cost function g with respect to the measure argument:

Assumptions *(continued)*

A8 For the same L as in A1-A7, condition (10) holds true with $L = L_C$, for any $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$ such that $\mu(\mathbb{R}^n) \leq \exp(\alpha T)$.

A9 For any $C > 0$ and for any sequence $(\mu_\ell)_{\ell \geq 1}$ in $\mathcal{B}_{\mathcal{M}_{2-r}}(C)$ that converges narrowly to some μ, it holds

$$\lim_{n \to \infty} \sup_{x \in \mathbb{R}^n} \left[\frac{1}{1 + |x|^{2-r}} |g(x, \mu_n) - g(x, \mu)| \right] = 0. \tag{19}$$

Remark 10. The following comments are in order.

- Thanks to Lemma 8, it is plain to see that, for a given $C > 0$, the *ball*

$$\mathcal{B}_{\mathcal{M}_{2-r}}(C) := \left\{ \mu \in \mathcal{M}_{2-r}(\mathbb{R}^n) : \int_{\mathbb{R}^n} (1 + |x|) \, \mathrm{d}\mu(x) \leq C \right\}$$

  is relatively compact for the narrow topology. In fact, it is also closed and hence compact. In particular, the measure μ in A9 is necessarily in $\mathcal{B}_{\mathcal{M}_{2-r}(C)}$.

- Following the above item, we notice that any real-valued function on $\mathcal{M}(\mathbb{R}^n)$ that is continuous on $\mathcal{B}_{\mathcal{M}_{2-r}}(C)$ with respect to the narrow topology, for some $C > 0$, is in fact uniformly continuous. In particular, for each $x \in \mathbb{R}^n$, the function $\mu \mapsto g(x, \mu)$ is, under condition A9, uniformly continuous on $\mathcal{B}_{\mathcal{M}_{2-r}}(C)$. Somehow, condition A9 imposes an additional constraint on the modulus of continuity, but uniformly in x.

- Following Lemma 9, a standard example of a function g that satisfies all the requirements A7-A9 is

$$g(x, \mu) := \int_{\mathbb{R}^n} \Gamma(x, y) \mathrm{d}\mu(y),$$

  where Γ is convex in the variable x and satisfies (all the derivatives below being implicitly assumed to exist),

$$\begin{aligned}
-L \left(1 + |x|^{1-r}\right) \leq \Gamma(x, y) \quad &\leq L \left(1 + |x|^{2-r}\right), \\
|\nabla_x \Gamma(x, y)| \quad &\leq L \left(1 + |x|^{1-r}\right), \\
|\nabla_y \Gamma(x, y)| \quad &\leq L \left(1 + |x|^{2-r}\right), \\
|\nabla^2_{x,x} \Gamma(x, y)| \quad &\leq L.
\end{aligned}$$

The proof is as follows. Let $C > 0$ and $\varepsilon > 0$. By the first line above, we can find a compact subset $K \subset \mathbb{R}^n$ such that, for any $\mu \in \mathcal{B}_{\mathcal{M}_{2-r}}(C)$,

$$\sup_{x\in\mathbb{R}^n}\left[\frac{1}{1+|x|^{2-r}}\left|\int_{\mathbb{R}^n}\Gamma(x,y)\mathrm{d}\mu(y)-\int_K\Gamma(x,y)\mathrm{d}\mu(y)\right|\right]\leq\varepsilon.$$

By the penultimate point, the functions $(y \mapsto \Gamma(x,y)/(1+|x|^{2-r}))_{x\in\mathbb{R}^n}$ are equicontinuous on K. Therefore, we can approximate any of them, to any fixed accuracy for the sup norm on K, by a continuous function in a finite collection. The proof is then easily completed.

- Similar to [16], the presentation is restricted to games in which only the terminal cost has a mean-field structure. That said, we could also consider mean-field running cost of the separated form

$$\ell'(t,\psi_t,\mu_t)=\ell(t,\psi_t)+c(X_t,\mu_t),$$

where μ_t is the marginal law of X_t under the probability measure $q_T\mathbb{P}$.

In the rest of this subsection, Assumptions A7-A9 are in force.

4.2 Entropy and moment estimates

In this subsection, we provide a series of bounds that are satisfied for any $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$.

We start with the following lemma:

Lemma 11. *There exists a constant C_1, only depending on the parameters in the standing assumptions, such that*

$$\sup_{\mu\in\mathcal{M}_{2-r}(\mathbb{R}^n)}\mathcal{S}(q^\mu)\leq C_1.$$

In particular, (up to a possibly new value of C_1)

$$\sup_{\mu\in\mathcal{M}_{2-r}(\mathbb{R}^n)}\sup_{t\in[0,T]}\mathbb{E}[h(q_t^\mu)]\leq C_1.$$

Proof. The result is a direct consequence of [16, Lemma 28]. The main point is to observe that the quantity $\mathcal{G}(q^0, X_T^\psi)$ appearing in the first step of the proof is here equal to $\mathbb{E}[q_T^0 g(X_T^\psi,\mu)]$. By convexity of g in the variable x and then by condition A8, it is greater than

$$\begin{aligned}\mathbb{E}[q_T^0 g(X_T^\psi,\mu)] &\geq \mathbb{E}[q_T^0 g(0,\mu)]+\mathbb{E}[q_T^0\nabla_x g(0,\mu)\cdot X_T^\psi]\\ &\geq -L\left(1+\mathbb{E}[q_T^0|X_T^\psi|]\right).\end{aligned}$$

The key fact is that the constant L is here independent of μ. This is the reason why the constant C_1 in the statement can be chosen independently of μ. □

Lemma 12. *There exists a constant C_2, only depending on the parameters in the standing assumptions, such that*

$$\sup_{\mu\in\mathcal{M}_{2-r}(\mathbb{R}^n)}\mathcal{S}^\star(\psi^\mu)\leq C_2.$$

Proof. The proof is an adaptation of [16, Lemma 34]. The bound established therein depends on g through g and $\nabla_x g$ at $x = 0$, but the latter two are bounded independently of μ, see A8.

One also needs a bound for the cost driven by the null control. Thanks again to A8, it is independent of μ. $\square$

Lemma 13. *There exists a constant C_3, only depending on the parameters in the standing assumptions, such that*

$$\sup_{\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)} \mathbb{E}\left[q_T^\mu |X_T^\mu|^{2-r}\right] \leq C_3.$$

Proof. This is a consequence of Lemmas 11 and 12 and 3. $\square$

4.3 Existence

Here is the first main result of the article.

Theorem 14. *Let Assumptions A1-A9 be in force. Then, there exists at least one equilibrium to the mean-field game set over (MinMax[μ]), in the sense of Definition 7.*

Proof. The proof is an application of Schauder's theorem, see [3, Corollary 17.56]. Throughout, we metricize the narrow topology introduced in Subsection 4.1 by means of the Fortet-Mourier distance. In fact, the latter extends to a norm on the whole space $\mathcal{M}_{\text{sign}}(\mathbb{R}^n)$ of signed measures on $\mathbb{R}^n$, given by

$$\|\mu\|_{\text{FM}} := \sup_{\varphi} \left[\int_{\mathbb{R}^n} \varphi(x) \mathrm{d}\mu(x)\right],$$

where the supremum is taken over functions φ that are bounded by 1 and 1-Lipschitz continuous.

Given the constant C_3 from Lemma 13, we consider the collection $\mathcal{C}$ of measures $\mu \in \mathcal{M}_{2-r}(\mathbb{R}^n)$ such that $\mu(\mathbb{R}^n) \leq \exp(\alpha T)$ and $M_{2-r}(\mu) \leq C_3$. By Lemma 8, we easily deduce that $\mathcal{C}$ is compact for $\|\cdot\|_{\text{FM}}$. Obviously, it is convex.

We then consider the mappping

$$\Phi : \mathcal{C} \ni \mu \mapsto (q_T^\mu \mathbb{P})_{X_T^\mu}.$$

By Lemma 13, $\mathcal{C}$ is stable by Φ.

It remains to check that Φ is continuous. We thus consider a sequence $(\mu_\ell)_{\ell \geq 1}$ in $\mathcal{C}$ that converges to μ for the narrow topology. By closedness of $\mathcal{C}$, $\mu \in \mathcal{C}$. For simplicity, we write $(q, X, \psi, Y^\star, Z^\star)$ for $(q^\mu, X^\mu, \psi^\mu, Y^{\star,\mu}, Z^{\star,\mu})$ and $(q^\ell, X^\ell, \psi^\ell, Y^{\star,\ell}, Z^{\star,\ell})$ for $(q^{\mu_\ell}, X^{\mu_\ell}, \psi^{\mu_\ell}, Y^{\star,\mu_\ell}, Z^{\star,\mu_\ell})$. By (16) in Proposition 6, there exists a constant $c > 0$ such that, for any $\ell \geq 1$,

$$\begin{aligned}
&\mathbb{E}\left[q_T^\ell \left(g(X_T, \mu^\ell) - g(X_T, \mu)\right) + q_T \left(g(X_T^\ell, \mu) - g(X_T^\ell, \mu^\ell)\right)\right] \\
&\geq c\mathbb{E}\left[\int_0^T \left(q_t + q_t^\ell\right) \left(|\psi_t^\ell - \psi_t|^2 + |Y_t^{\star,\ell} - Y_t^\star|^2 + |Z_t^{\star,\ell} - Z_t^\star|^2\right) \mathrm{d}t\right].
\end{aligned}$$

By A9, there exists a sequence ε_ℓ that tends to 0 such that

$$\mathbb{E}\left[q_T^\ell\left(g(X_T,\mu^\ell)-g(X_T,\mu)\right)+q_T\left(g(X_T^\ell,\mu)-g(X_T^\ell,\mu^\ell)\right)\right]$$
$$\leq \varepsilon_\ell \mathbb{E}\left[q_T^\ell\left(1+|X_T|^{2-r}\right)+q_T\left(1+|X_T^\ell|^{2-r}\right)\right].$$

By Lemmas 11, 12 and 3,

$$\sup_{\ell\geq 1}\mathbb{E}\left[q_T^\ell\left(1+|X_T|^{2-r}\right)+q_T\left(1+|X_T^\ell|^{2-r}\right)\right]<+\infty.$$

By the last three displays we deduce that

$$\lim_{\ell\to+\infty}\mathbb{E}\left[\int_0^T\left(q_t+q_t^\ell\right)\left(|\psi_t^\ell-\psi_t|^2+|Y_t^{\star,\ell}-Y_t^\star|^2+|Z_t^{\star,\ell}-Z_t^\star|^2\right)\mathrm{d}t\right]=0. \tag{20}$$

Using the linearity of the dynamics of X, we observe that

$$X_T-X_T^\ell=X_T^{\psi-\psi^\ell}.$$

And then, following the proof [16, Lemma 43] (which corresponds to Lemma 3), we deduce that

$$\lim_{\ell\to\infty}\mathbb{E}\left[\left(q_T+q_T^\ell\right)|X_T-X_T^\ell|^{2-r}\right]=0. \tag{21}$$

It then remains to prove that

$$\lim_{\ell\to\infty}\mathbb{E}\left[|q_T-q_T^\ell|\right]=0. \tag{22}$$

Assume indeed that the above holds true. Then, by combining (21) and (22), we obtain, for any test function $\varphi:\mathbb{R}^n\to\mathbb{R}$ that is bounded by 1 and 1-Lipschitz,

$$\mathbb{E}\left[|q_T\varphi(X_T)-q_T^\ell\varphi(X_T^\ell)|\right]\leq\mathbb{E}\left[|q_T-q_T^\ell|\right]+\mathbb{E}\left[q_T|X_T-X_T^\ell|\right].$$

Since the right-hand side tends to 0 (as ℓ tends to ∞) and is independent of φ, this gives $\|(q_T\mathbb{P})_{X_T}-(q_T^\ell\mathbb{P})_{X_T^\ell}\|_{\mathrm{FM}}\to 0$ as ℓ tends to ∞, which yields the required continuity property.

We now prove (22). We let $\mathcal{E}_T:=\mathcal{E}_T(\int_0^\cdot Z_s^\star\cdot\mathrm{d}W_s)$ and $\mathcal{E}_T^\ell:=\mathcal{E}_T(\int_0^\cdot Z_s^{\star,\ell}\cdot\mathrm{d}W_s)$. By definition of q_T and q_T^ℓ, we have

$$q_T=\exp\left(\int_0^T Y_t^\star\mathrm{d}t\right)\mathcal{E}_T,\quad q_T^\ell=\exp\left(\int_0^T Y_t^{\star,\ell}\mathrm{d}t\right)\mathcal{E}_T^\ell.$$

Therefore,

$$\begin{aligned}\mathbb{E}\left[|q_T-q_T^\ell|\right]&\leq\mathbb{E}\left[\mathcal{E}_T\left|\exp\left(\int_0^T Y_t^\star\mathrm{d}t\right)-\exp\left(\int_0^T Y_t^{\star,\ell}\mathrm{d}t\right)\right|\right]\\&\quad+\mathbb{E}\left[\exp\left(\int_0^T Y_t^{\star,\ell}\mathrm{d}t\right)\left|\mathcal{E}_T-\mathcal{E}_T^\ell\right|\right]\\&\leq\mathbb{E}^{\mathbb{Q}}\left[\left|\exp\left(\int_0^T Y_t^\star\mathrm{d}t\right)-\exp\left(\int_0^T Y_t^{\star,\ell}\mathrm{d}t\right)\right|\right]\\&\quad+\exp(\alpha T)\mathbb{E}\left[\left|\mathcal{E}_T-\mathcal{E}_T^\ell\right|\right],\end{aligned} \tag{23}$$

where we used the fact that $Y^{\star,\ell}$ is bounded by α in the last inequality.

We first consider the the first term on the last inequality (23). Since $Y^\star$ and $Y^{\star,\ell}$ are bounded by α, we have that

$$
\begin{aligned}
\mathbb{E}^{\mathbb{Q}}&\left[\left|\exp\left(\int_0^T Y_t^\star \mathrm{d}t\right) - \exp\left(\int_0^T Y_t^{\star,\ell}\mathrm{d}t\right)\right|\right] \\
&\leq \mathbb{E}^{\mathbb{Q}}\left[\int_0^T |Y_t^\star - Y_t^{\star,\ell}|\mathrm{d}t\right] \\
&\leq \exp(\alpha T)\mathbb{E}\left[q_T\int_0^T |Y_t^\star - Y_t^{\star,\ell}|^2\mathrm{d}t\right]^{1/2}, \qquad (24)
\end{aligned}
$$

where the last line follows by Jensen's inequality, the definition of q and the boundedness of $Y^\star$ again.

We now turn to the second term in (23). By Pinsker's inequality, we know that there exists a (universal) constant c_0 such that

$$
\mathbb{E}\left[|\mathcal{E}_T - \mathcal{E}_T^\ell|\right] \leq c_0\sqrt{\mathcal{H}\left(\mathcal{E}_T\mathbb{P}|\mathcal{E}_T^\ell\mathbb{P}\right)}, \quad \text{where} \quad \mathcal{H}\left(\mathcal{E}_T\mathbb{P}|\mathcal{E}_T^\ell\mathbb{P}\right) := \mathbb{E}\left[\ln\left(\frac{\mathcal{E}_T}{\mathcal{E}_T^\ell}\right)\mathcal{E}_T\right].
$$

It is standard to prove that

$$
\begin{aligned}
\mathbb{E}\left[\ln\left(\frac{\mathcal{E}_T}{\mathcal{E}_T^\ell}\right)\mathcal{E}_T\right] &= \mathbb{E}\left[\mathcal{E}_T\int_0^T |Z_t^\star - Z_t^{\star,\ell}|^2\mathrm{d}t\right] \\
&\leq \exp(\alpha T)\mathbb{E}\left[q_T\int_0^T |Z_t^\star - Z_t^{\star,\ell}|^2\mathrm{d}t\right].
\end{aligned}
$$

And then, there exists a constant C_0, independent of ℓ, such that

$$
\mathbb{E}\left[|\mathcal{E}_T - \mathcal{E}_T^\ell|\right] \leq C_0\mathbb{E}\left[q_T\int_0^T |Z_t^\star - Z_t^{\star,\ell}|^2\mathrm{d}t\right]^{1/2}. \qquad (25)
$$

Finally combining (23)–(24)–(25), we obtain that

$$
\mathbb{E}\left[|q_T - q_T^\ell|\right] \leq C\left(\mathbb{E}\left[q_T\int_0^T |Y_t^\star - Y_t^{\star,\ell}|^2\mathrm{d}t\right]^{1/2} + \mathbb{E}\left[q_T\int_0^T |Z_t^\star - Z_t^{\star,\ell}|^2\mathrm{d}t\right]^{1/2}\right),
$$

for some $C > 0$ independent on ℓ, and the conclusion follows by (20). □

4.4 Uniqueness criterion

Uniqueness is a more subtle issue than in standard mean-field games, due the presence of Nature. To understand this, we may just focus on the situation where T is small. Of course, we want to use the stability inequality (17), which we recall here for convenience:

$$
\begin{aligned}
&\mathbb{E}\left[\left(q_T^\mu - q_T^{\tilde{\mu}}\right)\left(g(X_T^\mu,\mu) - g(X_T^{\tilde{\mu}},\tilde{\mu})\right)\right] \\
&\geq \mathbb{E}\left[\left(q_T^\mu\nabla_x g(X_T^\mu,\mu) - q_T^{\tilde{\mu}}\nabla_x g(X_T^{\tilde{\mu}},\tilde{\mu})\right)\cdot(X_T^\mu - X_T^{\tilde{\mu}})\right] \qquad (26) \\
&\quad + c\mathbb{E}\left[\int_0^T \left(q_t^\mu + q_t^{\tilde{\mu}}\right)\left(|\psi_t^{\tilde{\mu}} - \psi_t^\mu|^2 + |Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}|^2\right)\mathrm{d}t\right].
\end{aligned}
$$

In absence of Nature, this inequality becomes very much simpler and just writes

$$\mathbb{E}\left[\left(\nabla_x g(X_T^{\mu},\mu)-\nabla_x g(X_T^{\tilde\mu},\tilde\mu)\right)\cdot(X_T^{\mu}-X_T^{\tilde\mu})\right]\geq c\mathbb{E}\left[\int_0^T|\psi_t^{\tilde\mu}-\psi_t^{\mu}|^2\mathrm{d}t\right].$$

Although we do not pretend that the derivation of the above display is rigorous, it turns out that this is indeed what can be obtained by using the so-called 'probabilistic approach to mean-field games', see for instance [13, Chapter 4]. When μ is understood as the law of X_T and μ' as the law of X'_T (under the common probability measure $\mathbb{P}$), the left-hand side can be upper bounded, under Lipschitz assumptions on the derivatives of g (the Lipschitz constant being denoted by the generic letter L), by

$$\mathbb{E}\left[\left(\nabla_x g(X_T^{\mu},\mu)-\nabla_x g(X_T^{\tilde\mu},\tilde\mu)\right)\cdot(X_T^{\mu}-X_T^{\tilde\mu})\right]\leq LT\mathbb{E}\left[\int_0^T|\psi_t^{\tilde\mu}-\psi_t^{\mu}|^2\mathrm{d}t\right].$$

The extra factor T on the right-hand side makes it possible to guarantee uniqueness in small time.

Here, we want to argue, at least informally, that the same computation can not be reproduced in the robust setting. When μ and $\tilde\mu$ are understood as $(q_T\mathbb{P})_{X_T}$ and $(\tilde q_T\mathbb{P})_{\tilde X_T}$ respectively (with (q_T,X_T) standing for (q_T^{μ},X_T^{μ}), and $(\tilde q_T,\tilde X_T)$ for $(q_T^{\tilde\mu},X_T^{\tilde\mu})$), the first term on (26) can be estimated as follows, at least in the simpler situation where X_T and $\tilde X_T$ are equal (which is of course not true in general, but which cannot make the difficulty worse). In the latter situation, we are led to estimate $g(X_T,(q_T\mathbb{P}_{X_T}))-g(X_T,(\tilde q_T\mathbb{P})_{X_T})$. At best, we can expect to upper bound it by $\mathbb{E}[|q_T-\tilde q_T|]$. Therefore, the first term on (26) can be bounded by $\mathbb{E}[|q_T-\tilde q_T|]^2$, up to a multiplicative constant. Then, we know from the proof of Theorem 14 that this term can be bounded by means of Pinsker inequality by $\mathbb{E}[q_T\int_0^T|Z_t^{\star}-\tilde Z_t^{\star}|^2\mathrm{d}t]$ (with an obvious meaning for $Z^{\star}$ and $\tilde Z^{\star}$). In particular, there is no extra factor T that could render its contribution smaller than the contribution of the corresponding term on the right-hand side of (26).

Of course, a more direct way to obtain uniqueness is to multiply g by a small parameter and then obtain the desired 'contraction' (in the sense that the right hand side on (26) dominates the left hand side when μ and $\tilde\mu$ are equilibria). Although this could be one result towards uniqueness, we feel better to follow another route. The main point is to focus on the difference

$$\begin{aligned}&\mathbb{E}\left[\left(q_T^{\mu}-q_T^{\tilde\mu}\right)\left(g(X_T^{\mu},\mu)-g(X_T^{\tilde\mu},\tilde\mu)\right)\right]\\&\quad-\mathbb{E}\left[\left(q_T^{\mu}\nabla_x g(X_T^{\mu},\mu)-q_T^{\tilde\mu}\nabla_x g(X_T^{\tilde\mu},\tilde\mu)\right)\cdot(X_T^{\mu}-X_T^{\tilde\mu})\right],\end{aligned}$$

when μ and $\tilde\mu$ satisfy the fixed point conditions

$$\mu=(q_T^{\mu}\mathbb{P})_{X_T^{\mu}},\quad \tilde\mu=(q_T^{\tilde\mu}\mathbb{P})_{X_T^{\tilde\mu}}.$$

This prompts us to introduce the following definition:

Definition 15. *The function g is said to be (jointly) flat non-increasing/displacement non-decreasing if, for any non-negative-valued random variables q and q' with $\mathbb{E}[q],\mathbb{E}[q']\leq\exp(\alpha T)$, and any $\mathbb{R}^n$-valued random variables X and X' satisfying $\mathbb{E}[(q+q')(|X|^{2-r}+|X'|^{2-r})]<+\infty$, it holds*

$$\begin{aligned}&\mathbb{E}\left[(q-q')\left(g(X,(q\mathbb{P})_X)-g(X',(q'\mathbb{P})_{X'})\right)\right]\\&\quad-\mathbb{E}\left[\left(q\nabla_x g(X,(q\mathbb{P})_X)-q'\nabla_x g(X',(q'\mathbb{P})_{X'})\right)\cdot(X-X')\right]\leq 0.\end{aligned}\tag{27}$$

Remark 16. The following remarks are in order.

- It is easy to see that the property (27) only depends on the joint law of (q, q', X, X') under $\mathbb{P}$. In particular, the property (27) can be transferred from one probability space to another.

  In fact, since the space $(\Omega, \mathcal{F}, \mathbb{P})$ is here equipped with a Brownian motion, we can construct, for any given law on $(0, +\infty) \times (0, +\infty) \times \mathbb{R}^n \times \mathbb{R}^n$, a 4-tuple (q, q', X, X') having this law under $\mathbb{P}$ (as we can 'reconstruct' any random variable from a random variable with uniform distribution on $[0, 1]$). This guarantees that, on any probability space, the above inequality is guaranteed for any random variables q, q', X, X' (taking values in the required spaces, and satisfying the required integrability properties).

- Choose $q = q' = 1$ in (27), and deduce that, for any $\mathbb{R}^n$-valued random variables X, X' satisfying $\mathbb{E}[|X|^{2-r}], \mathbb{E}[|X'|^{2-r}] < +\infty$,

$$\mathbb{E}\left[\left(\nabla_x g(X, \mathbb{P}_X) - \nabla_x g(X', \mathbb{P}_{X'})\right) \cdot (X - X')\right] \geq 0,$$

  which is the standard displacement monotonicity property.

- Choose now $X = X'$ in (27), and deduce that, for any q, q' with positive values,

$$\mathbb{E}\left[\left(q - q'\right)\left(g(X, (q\mathbb{P})_X) - g(X, (q'\mathbb{P})_X)\right)\right] \leq 0.$$

  Choose now X as being uniformly distributed on a given domain Borel subset $A \subset \mathbb{R}^n$ with finite Lebesgue measure (denoted $\mathrm{Leb}_n(A)$) and then (q, q', X) such that $q = \mathrm{Leb}_n(A) f(X)$ and $q' = \mathrm{Leb}_n(A) f'(X)$ for two non-negative functions f and f' with support included in A, and satisfying $\int_{\mathbb{R}^n} f(x)\mathrm{d}x, \int_{\mathbb{R}^n} f'(x)\mathrm{d}x \leq \exp(\alpha T)$ and $\int_{\mathbb{R}^n} |x|^{2-r} f(x)\mathrm{d}x, \int_{\mathbb{R}^n} |x|^{2-r} f'(x)\mathrm{d}x < +\infty$. The above inequality can be rewritten

$$\int_A \left(f(x) - f'(x)\right)\left(g(x, f\mathrm{Leb}_n) - g(x, f'\mathrm{Leb}_n)\right)\mathrm{d}x \leq 0.$$

  Obviously $\mathbb{R}^n$ can be substituted for A in the above display. And, then, by a standard approximation argument (using the regularity of g in the measure argument), we deduce that the inequality holds true for f, f' with $\int_{\mathbb{R}^n} f(x)\mathrm{d}x, \int_{\mathbb{R}^n} f'(x)\mathrm{d}x \leq \exp(\alpha T)$. And then, approximating any (finite non-negative) measure on $\mathbb{R}^n$ by measures with densities, we deduce that, for any measures $m, m' \in \mathcal{M}_{2-r}(\mathbb{R}^n)$,

$$\int_A \left(g(x, m) - g(x, m')\right)\mathrm{d}\left(m - m'\right)(x) \leq 0,$$

  which is an anti-Lasry-Lions monotonicity condition.

- It is not clear to us whether a function that is non-increasing in the flat sense (as in the previous item) and, separately displacement non-decreasing (as in the penultimate item), is (jointly) flat non-increasing/displacement non-decreasing as in Definition 15.

We provide below a canonical example of a function g satisfying Definition 15.

Lemma 17. *Let $\mathcal{G}$ be a function from $\mathcal{M}_{2-r}(\mathbb{R}^n)$ that is flat concave and displacement convex, in the sense that, for all $\mu, \mu' \in \mathcal{M}_{2-r}(\mathbb{R}^n)$,*

$$
\begin{aligned}
\mathcal{G}(\mu') &\leq \mathcal{G}(\mu) + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}(\mu, x) \mathrm{d}\left(\mu' - \mu\right)(x), \\
\mathcal{G}(\mu') &\geq \mathcal{G}(\mu) + \int_{\mathbb{R}^n \times \mathcal{R}^n} \partial_\mu \mathcal{G}(\mu, x) \cdot (y - x) \mathrm{d}\pi(x, y),
\end{aligned}
$$

where π on the last term is a coupling between μ and μ', i.e. π has μ as first marginal on $\mathbb{R}^n$ and μ' as second marginal.

Then, the function

$$
(x, \mu) \mapsto g(x, \mu) := \frac{\delta \mathcal{G}}{\delta \mu}(\mu, x)
$$

is flat non-increasing/displacement non-decreasing.

The notions of derivatives used in the statement are standard. In brief, the flat derivative $\delta \mathcal{G} / \delta \mu$ is defined as

$$
\frac{\delta \mathcal{G}}{\delta \mu}(\mu, x) = \frac{\mathrm{d}}{\mathrm{d}\varepsilon}_{|\varepsilon = 0+} \mathcal{G}\left(\mu + \varepsilon \delta_x\right), \quad \mu \in \mathcal{M}_{2-r}(\mathbb{R}^n), \quad x \in \mathbb{R}^n,
$$

and the intrinsic derivative $\partial_\mu \mathcal{G}$ as

$$
\partial_\mu \mathcal{G}(\mu, x) = \nabla_x \frac{\delta \mathcal{G}}{\delta \mu}(\mu, x), \quad \mu \in \mathcal{M}_{2-r}(\mathbb{R}^n), \quad x \in \mathbb{R}^n,
$$

Precise definitions and examples of flat-concave/displacement-convex functions are provided in [16, Subsection 4].

Proof. By [16, Corollary 17], the following two inequalities hold true for any q, q' and any X, X' as in Definition 15:

$$
\begin{aligned}
&\mathcal{G}\left(q\mathbb{P}\right)_{X'} \geq \mathcal{G}((q\mathbb{P})_X) + \mathbb{E}\left[q \partial_\mu \mathcal{G}\left((q\mathbb{P})_X, X\right) \cdot (X' - X)\right], \\
&\mathcal{G}\left((q'\mathbb{P})_X\right) \leq \mathcal{G}((q\mathbb{P})_X) + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}\left((q\mathbb{P})_X, x\right) \mathrm{d}\left[(q'\mathbb{P})_X - (q\mathbb{P})_X\right](x).
\end{aligned} \tag{28}
$$

We rewrite the second line as

$$
\mathcal{G}\left((q\mathbb{P})_X\right) \geq \mathcal{G}((q'\mathbb{P})_X) + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}\left((q\mathbb{P})_X, x\right) \mathrm{d}\left[(q\mathbb{P})_X - (q'\mathbb{P})_X\right](x).
$$

And then, exchanging the roles of (q, X) and of (q', X') in the above inequality and then using the first line of (28),

$$
\begin{aligned}
\mathcal{G}\left((q'\mathbb{P})_{X'}\right) &\geq \mathcal{G}((q\mathbb{P})_{X'}) + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}\left((q'\mathbb{P})_{X'}, x\right) \mathrm{d}\left[(q'\mathbb{P})_{X'} - (q\mathbb{P})_{X'}\right](x) \\
&\geq \mathcal{G}((q\mathbb{P})_X) + \mathbb{E}\left[q \partial_\mu \mathcal{G}\left((q\mathbb{P})_X, X\right) \cdot (X' - X)\right] \\
&\quad + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}\left((q'\mathbb{P})_{X'}, x\right) \mathrm{d}\left[(q'\mathbb{P})_{X'} - (q\mathbb{P})_{X'}\right](x).
\end{aligned}
$$

Next, we exchange once again the roles of (q, X) and (q', X') and then sum the two resulting inqualities. We get

$$
\begin{aligned}
&- \mathbb{E}\left[\left(q'\partial_\mu \mathcal{G}\left((q'\mathbb{P})_{X'}, X'\right) - q\partial_\mu \mathcal{G}\left((q\mathbb{P})_X, X\right)\right) \cdot (X' - X)\right] \\
&\quad + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}\left((q'\mathbb{P})_{X'}, x\right) \mathrm{d}\left[(q'\mathbb{P})_{X'} - (q\mathbb{P})_{X'}\right](x) \\
&\quad + \int_{\mathbb{R}^n} \frac{\delta \mathcal{G}}{\delta \mu}\left((q\mathbb{P})_X, x\right) \mathrm{d}\left[(q\mathbb{P})_X - (q'\mathbb{P})_X\right](x) \leq 0.
\end{aligned}
$$

Letting $g(x, \mu) = [\delta \mathcal{G}/\delta \mu](\mu, x)$ as done in the statement, and recalling that $\nabla_x g(x, \mu) = \partial_\mu \mathcal{G}(\mu, x)$, the above display can be rewritten as

$$
\begin{aligned}
&- \mathbb{E}\left[\left(q\nabla_x g\left(X, (q\mathbb{P})_X\right) - q'\nabla_x g\left(X', (q'\mathbb{P})_{X'}\right)\right) \cdot (X - X')\right] \\
&\quad + \int_{\mathbb{R}^n} g\left(x, (q'\mathbb{P})_{X'}\right) \mathrm{d}\left[(q'\mathbb{P})_{X'} - (q\mathbb{P})_{X'}\right](x) \\
&\quad + \int_{\mathbb{R}^n} g\left(x, (q\mathbb{P})_X\right) \mathrm{d}\left[(q\mathbb{P})_X - (q'\mathbb{P})_X\right](x) \leq 0,
\end{aligned}
$$

and then

$$
\begin{aligned}
&- \mathbb{E}\left[\left(q\nabla_x g\left(X, (q\mathbb{P})_X\right) - q'\nabla_x g\left(X', (q'\mathbb{P})_{X'}\right)\right) \cdot (X - X')\right] \\
&\quad + \mathbb{E}\left[(q' - q)\left(g\left(X', (q'\mathbb{P})_{X'}\right) - g\left(X, (q\mathbb{P})_X\right)\right)\right] \leq 0,
\end{aligned}
$$

which completes the proof. $\square$

Here is now the main result of this section:

Proposition 18. *Let Assumptions A1-A9 be in force. If further, the function g is flat non-increasing/displacement non-decreasing, then there exists a unique equilibrium in the sense of Definition 7.*

Proof. The proof is a straightforward consequence of (17).

Consider indeed two equilibria, denoted by $\mu, \tilde{\mu}$. Using the fact that

$$
\mu = (q_T^\mu \mathbb{P})_{X_T^\mu}, \quad \tilde{\mu} = (q_T^{\tilde{\mu}} \mathbb{P})_{X_T^{\tilde{\mu}}},
$$

we get

$$
\begin{aligned}
&\mathbb{E}\left[\left(q_T^\mu - q_T^{\tilde{\mu}}\right)\left(g\left(X_T^\mu, (q_T^\mu \mathbb{P})_{X_T^\mu}\right) - g\left(X_T^{\tilde{\mu}}, (q_T^{\tilde{\mu}} \mathbb{P})_{X_T^{\tilde{\mu}}}\right)\right)\right] \\
&\quad - \mathbb{E}\left[\left(q_T^\mu \nabla_x g\left(X_T^\mu, (q_T^\mu \mathbb{P})_{X_T^\mu}\right) - q_T^{\tilde{\mu}} \nabla_x g\left(X_T^{\tilde{\mu}}, (q_T^{\tilde{\mu}} \mathbb{P})_{X_T^{\tilde{\mu}}}\right)\right) \cdot (X_T^\mu - X_T^{\tilde{\mu}})\right] \\
&\geq c\mathbb{E}\left[\int_0^T \left(q_t^\mu + q_t^{\tilde{\mu}}\right)\left(|\psi_t^{\tilde{\mu}} - \psi_t^\mu|^2 + |Y_t^{\star,\tilde{\mu}} - Y_t^{\star,\mu}|^2 + |Z_t^{\star,\tilde{\mu}} - Z_t^{\star,\mu}|^2\right) \mathrm{d}t\right].
\end{aligned}
$$

By (27), the left-hand side is less than (or equal to) 0, from which we deduce that the right-hand side is equal to 0. $\square$

Constructing flat non-decreasing/displacement non-increasing functions The purpose of this paragraph is to provide a tractable condition ensuring that g satisfies Definition 15, beyond the potential regime introduced in the statement of Lemma 17.

Typically, we require the function g to be jointly convex in the flat sense, i.e.

$$g(x',\mu') \geq g(x,\mu) + \nabla_x g(x,\mu)\cdot(x'-x) + \int_{\mathbb{R}^n} \frac{\delta g}{\delta \mu}(x,\mu,y)\mathrm{d}\left(\mu'-\mu\right)(y), \qquad (29)$$

for any $x, x' \in \mathbb{R}^n$ and $\mu, \mu' \in \mathcal{M}_{2-r}(\mathbb{R}^n)$. Implicitly, the function g is assumed to be differentiable (in the flat sense) with respect to the measure argument, and the integral on the right-hand side is assumed to make sense.

Back to Definition 15, the purpose is to upper bound the left-hand side on (27). Thanks to (29), we have

$$\begin{aligned}
&\mathbb{E}\left[q\left(g(X,(q\mathbb{P})_X) - g(X',(q'\mathbb{P})_{X'}) - \nabla_x g(X,(q\mathbb{P})_X)\cdot(X-X')\right)\right]\\
&= \mathbb{E}\left[q\left(g(X,(q\mathbb{P})_X) - g(X',(q'\mathbb{P})_{X'}) + \nabla_x g(X,(q\mathbb{P})_X)\cdot(X'-X)\right)\right]\\
&\leq -\mathbb{E}\left[q\int_{\mathbb{R}^n}\frac{\delta g}{\delta\mu}\left(X,(q\mathbb{P})_X,y\right)\mathrm{d}\left((q'\mathbb{P})_{X'}-(q\mathbb{P})_X\right)(y)\right].
\end{aligned}$$

Exchanging the roles of (q,X) and (q',X') and summing the two resulting inequalities, we get

$$\begin{aligned}
&\mathbb{E}\left[\left(q-q'\right)\left(g(X,(q\mathbb{P})_X) - g(X',(q'\mathbb{P})_{X'})\right)\right]\\
&\quad - \mathbb{E}\left[\left(q\nabla_x g(X,(q\mathbb{P})_X) - q'\nabla_x g(X',(q'\mathbb{P})_{X'})\right)\cdot(X-X')\right]\\
&\leq \mathbb{E}\left[\int_{\mathbb{R}^n}\left(q'\frac{\delta g}{\delta\mu}\left(X',(q'\mathbb{P})_{X'},y\right) - q\frac{\delta g}{\delta\mu}\left(X,(q\mathbb{P})_X,y\right)\right)\mathrm{d}\left((q'\mathbb{P})_{X'}-(q\mathbb{P})_X\right)(y)\right]\\
&= \int_{\mathbb{R}^n}\left[\int_{\mathbb{R}^n}\frac{\delta g}{\delta\mu}\left(z,(q'\mathbb{P})_{X'},y\right)\mathrm{d}(q'\mathbb{P})_{X'}(z)\right]\mathrm{d}\left((q'\mathbb{P})_{X'}-(q\mathbb{P})_X\right)(y)\\
&\quad - \int_{\mathbb{R}^n}\left[\int_{\mathbb{R}^n}\frac{\delta g}{\delta\mu}\left(z,(q\mathbb{P})_X,y\right)\mathrm{d}(q\mathbb{P})_X(z)\right]\mathrm{d}\left((q'\mathbb{P})_{X'}-(q\mathbb{P})_X\right)(y).
\end{aligned}$$

And then, in order to guarantee (27), it suffices to have

$$\int_{\mathbb{R}^n}\left[\int_{\mathbb{R}^n}\frac{\delta g}{\delta\mu}\left(z,\mu',y\right)\mathrm{d}\mu'(z) - \int_{\mathbb{R}^n}\frac{\delta g}{\delta\mu}\left(z,\mu,y\right)\mathrm{d}\mu(z)\right]\mathrm{d}\left(\mu'-\mu\right)(y) \leq 0, \qquad (30)$$

for any $\mu, \mu' \in \mathcal{M}_{2-r}(\mathbb{R}^n)$.

Here is a typical example:

Lemma 19. *Let $K : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ be a smooth function, bounded with bounded derivatives of any order, of negative type, i.e., satisfying for any smooth function $h : \mathbb{R}^n \to \mathbb{R}$ with a compact support,*

$$\int_{\mathbb{R}^n\times\mathbb{R}^n} K(x,y)h(x)h(y)\mathrm{d}x\mathrm{d}y \leq 0. \qquad (31)$$

Then, the function g defined by

$$g(x,\mu) = \int_{\mathbb{R}^n\times\mathbb{R}^n} K(x,y)\mathrm{d}\mu(y), \quad x\in\mathbb{R}^n,\ \mu\in\mathcal{M}_{2-r}(\mathbb{R}^n),$$

satisfies (30).

Proof. The proof is quite obvious as the left-hand side on (30) rewrites

$$\int_{\mathbb{R}^n} \left[\int_{\mathbb{R}^n} \frac{\delta g}{\delta \mu} \left(z, \mu', y\right) \mathrm{d}\mu'(z) - \int_{\mathbb{R}^n} \frac{\delta g}{\delta \mu} \left(z, \mu, y\right) \mathrm{d}\mu(z) \right] \mathrm{d} \left(\mu' - \mu\right)(y)$$
$$= \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} K(z, y) \mathrm{d} \left(\mu' - \mu\right)(z) \mathrm{d} \left(\mu' - \mu\right)(y).$$

By (31), it is easy to see that the left-hand side is negative. □

Example 20.

- A first example for K satisfying (31) is

  $$K(x, y) = -\phi(x)\phi(y),$$

  for a smooth function $\phi : \mathbb{R}^n \to \mathbb{R}$.

- Another example is

  $$K(x, y) = -\int_{\mathbb{R}^n} \phi(x, r)\phi(y, r)\mathrm{d}\lambda(r),$$

  where λ is a compactly supported positive finite measure on $\mathbb{R}^k$, and ϕ is a smooth function from $\mathbb{R}^n \times \mathbb{R}^k$ to $\mathbb{R}$.

- The first two examples are symmetric in (x, y), as a result of which the function g, as defined in Lemma 19, derives from a potential.

  That said, any (smooth) function K that is anti-symmetric, i.e. $K(x, y) = -K(y, x)$, satisfies Lemma 19.

Of course, the function g defined in the statement of Lemma 19 does not satisfy the joint convexity condition 29. To make it jointly convex, we may add a function that is convex in the variable x. Following the examples constructed in [16, Subsection 4], we claim

Lemma 21. *Let K be as in the statement of Lemma 19.*

1. *If $r = 0$, we can find λ large enough such that the function*

   $$g(x, \mu) = \frac{\lambda}{2}|x|^2 + \int_{\mathbb{R}^n} K(x, y)\mathrm{d}\mu(y)$$

   satisfies (29) and, therefore, is jointly flat non-decreasing/displacement non-increasing functions.

2. *If $r = 1$ and K is compactly supported, we can find λ large enough such that the function*

   $$g(x, \mu) = \lambda \left(1 + |x|^2\right)^{1/2} + \int_{\mathbb{R}^n} K(x, y)\mathrm{d}\mu(y)$$

   satisfies (29) and, therefore, is jointly flat non-decreasing/displacement non-increasing functions.

Here, the choice of the convex perturbation is adapted to the value of r, so that g satisfies the required growth properties in A7.

5 Limiting theory

In this section, we investigate the connection between the mean-field game problem (MFG) and a finite-player game in which N players interact with Nature. The model is presented in Subsection 5.1. In Subsection 5.2, we establish an ε_N-Nash equilibrium result for the finite-player game.

5.1 Game with N competitive players vs. Nature

A primer on the law of large numbers. The construction of the N-player game relies on the following variant of the law of large numbers:

Lemma 22. *Let (q, X) be a random variable with values in $(0, +\infty) \times \mathbb{R}^n$. Assume that $\mathbb{E}[q] = 1$. Let $(q^i, X^i)_{i\geq 1}$ be an I.I.D. sequence with the law of (q, X) as common distribution (the sequence being constructed on $(\Omega, \mathcal{F}, \mathbb{P})$). Then,*

$$\forall \varepsilon > 0, \quad \lim_{N\to+\infty} \mathbb{E}\left[\left(\prod_{i=1}^{N} q^i\right) \mathbb{1}_{\{d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{i=1}^{N}\delta_{X^i}, (q\mathbb{P})_X\right) > \varepsilon\}}\right] = 0,$$

where d_{FM} is the Fortet-Mourier distance $d_{\mathrm{FM}}(\mu, \nu) = \|\mu - \nu\|_{\mathrm{FM}}$, for $\mu, \nu \in \mathcal{P}(\mathbb{R}^n)$, and $\|\cdot\|_{\mathrm{FM}}$ is defined in the proof of Theorem 14.

This result says that the standard empirical measure converges, in probability under $q^1 \dots q^N \mathbb{P}$, to $(q\mathbb{P})_X$.

Proof. We consider an I.I.D. sequence $(\tilde{X}^i)_{i\geq 1}$ with common distribution $(q\mathbb{P})_X$ under $\mathbb{P}$. It is easy to see that, for each $N \geq 1$, the law of $(X^1, \dots, X^N)$ under $q^1 \dots q^N \mathbb{P}$ is equal to $[(q\mathbb{P})_X]^{\times N}$, which is also the law of $(\tilde{X}^1, \dots, \tilde{X}^N)$ (but under $\mathbb{P}$). In particular,

$$\begin{aligned}\forall \varepsilon > 0, \quad & \mathbb{E}\left[\prod_{i=1}^{N} q_i \mathbb{1}_{\{d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{i=1}^{N}\delta_{X^i}, (q\mathbb{P})_X\right) > \varepsilon\}}\right] \\ & = \mathbb{P}\left(\left\{d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{i=1}^{N}\delta_{\tilde{X}^i}, (q\mathbb{P})_X\right) > \varepsilon\right\}\right),\end{aligned}$$

but the right-hand side tends to 0, as a consequence of the law of large numbers. $\square$

Presentation of the game. Based on Lemma 22, we now construct a game with N competitive players playing against Nature, whose asymptotic version corresponds to the game studied in Section 4. Due to the restriction on the mass of q imposed in Lemma 22, Nature' state in (1) is assumed to be a Doléans-Dade exponential, i.e. $Y^\star \equiv 0$ in the mean-field game.

We consider the product space $(\Omega^{\times N}, \mathcal{F}^{\times N}, \mathbb{P}^{\times N})$, and we equip its i-th factor with an $\mathbb{R}^d$-valued Brownian motion $W^i = (W^i_t)_{t\in[0,T]}$ and an initial condition η^i, η^i and W^i being independent. We assume that all the random variables $\eta^1, \dots, \eta^N$ are identically distributed, the support of their common statistical law being bounded. We denote by $\mathbb{F}^N = (\mathcal{F}^N_t)_{0\leq t\leq T}$ the completion of the filtration generated by $(\eta^1, \dots, \eta^N, W^1, \dots, W^n)$.

Below the function b^i is a copy of b on the i-th factor of $\Omega^{\times N}$, i.e., for any $\omega^{(N)} = (\omega^1, \dots, \omega^N) \in \Omega^{\times N}$, the quantity $b^i(\omega^{(N)}, t, x, \psi)$ depends on $\omega^{(N)}$ only through ω^i and is thus equal to $b^i(\omega^i, t, x, \psi)$. The functions σ^i, ℓ^i and $f^{\star,i}$ are constructed from b, σ, ℓ and $f^\star$ in the same way.

The admissible set of Nature, denoted by $\mathcal{Q}^{(N)}$, is the class of $\mathbb{F}^N$-progressively measurable, positive valued processes $(Q_t^N)_{t\in[0,T]}$ such that (compare with (1))

$$\mathcal{S}^{(N)}(Q^N) < +\infty, \qquad Q_t^N := \prod_{i=1}^N \mathcal{E}_t\left(\int_0^\cdot Z_s^{\star,i}\cdot \mathrm{d}W_s^i\right), \quad t\in[0,T],$$

with $(Z^{\star,1},\dots,Z^{\star,N})$ acting as Nature's control. The mapping $\mathcal{S}^N$ denotes the N-player generalized entropy counterpart, defined as follows

$$\mathcal{S}^{(N)}(Q^N) := \mathbb{E}^{\times N}\left[\int_0^T Q_s^N\left(\sum_{i=1}^N f^{\star,i}(s, Z_s^{\star,i})\right)\mathrm{d}s\right].$$

Here, the function

$$\mathbb{R}^{d\times N} \ni \left(z^{\star,1},\dots,z^{\star,N}\right) \mapsto \sum_{i=1}^N f^{\star,i}(t, z^{\star,i})$$

is understood as the Fenchel-Legendre transform of the function

$$\mathbb{R}^{d\times N} \ni \left(z^1\dots,z^N\right) \mapsto \sum_{i=1}^N f^i(t, z^i),$$

where, as before, $f^i(\omega^{(N)}, t, z)$ is equal to $f(t,\omega^i,z)$ (with f being now independent of y as Nature's mass remains equal to 1).

The control and state processes to player $i\in\{1,\dots,N\}$ are denoted by $\psi^i = (\psi_t^i)_{t\in[0,T]}$ and $X^i = (X_t^i)_{t\in[0,T]}$ respectively, both processes taking values in $\mathbb{R}^n$. When needed, we write X^{i,ψ^i} to emphasize the fact that X^i is controlled by ψ^i. Following (5), the dynamics of X^i write

$$\mathrm{d}X_t^i = b^i(t, X_t^i, \psi_t^i)\mathrm{d}t + \sigma^i(t,\psi_t^i)\mathrm{d}W_t^i, \quad t\in[0,T], \quad X_0^i = \eta^i.$$

The admissible set of each player $i\in\{1,\dots,N\}$ is denoted by $\mathcal{A}^{(N)}$ (it does not depend on i) and consists in a class of $\mathbb{F}^N$-progressively-measurable, $\mathbb{R}^n$-valued processes $\psi = (\psi_t)_{0\le t\le T}$ such that

$$\mathcal{S}^{\star,(N)}(\psi) < +\infty, \quad \mathcal{S}^{\star,(N)}(\psi) := \sup_{Q^N\in\mathcal{Q}^{(N)}}\left\{\mathbb{E}\left[\int_0^T Q_s^N|\psi_s|^2\mathrm{d}s\right] - \gamma\mathcal{S}^{(N)}(Q^N)\right\}.$$

Importantly, the parameter γ remains unchanged and is thus independent of N.

Given a control $q\in\mathcal{Q}^N$ of Nature, the cost to player $i\in\{1,\dots,N\}$ is defined as

$$\begin{aligned}\mathcal{R}^i(Q^N,\psi^1,\dots,\psi^N) := \mathbb{E}^{\times N}&\left[Q_T^N\left(g^i\left(X_T^{\psi^i}, \mu_T^{\psi^1,\dots,\psi^N}\right)\right)\right] \\ &+ \mathbb{E}^{\times N}\left[\int_0^T Q_s^N \ell^i(s,\psi_s^i)\mathrm{d}s\right],\end{aligned}$$

where

$$\mu_T^{\psi^1,\dots,\psi^N} := \frac{1}{N}\sum_{i=1}^N \delta_{X_T^{i,\psi^i}}. \tag{32}$$

The N adversarial players are also in competition against Nature, whose reward is given by

$$\begin{aligned}\mathcal{J}^{(N)}\left(Q^N,\psi^1,\ldots,\psi^N\right):=\mathbb{E}^{\times N}&\left[Q_T^N\sum_{i=1}^N h\left(X_T^{i,\psi^i},\mu_T^{\psi^1,\ldots,\psi^N}\right)\right]\\&+\mathbb{E}^{\times N}\left[\int_0^T Q_s^N\sum_{i=1}^N \ell^i(s,\psi_s^i)\mathrm{d}s\right]-\mathcal{S}^{(N)}\left(Q^N\right),\end{aligned}\tag{33}$$

where $h:\mathbb{R}^n\times\mathcal{P}_{2-r}(\mathbb{R}^n)\to\mathbb{R}$ is a function satisfying the same properties as g (It may be equal to g, but not necessarily).

The intuition is as follows: when Nature is frozen, players act as a in 'standard' N-game under the measure $q\mathbb{P}$; but Nature penalizes them by choosing the worst (standing from players' viewpoint) q according to the reward $\mathcal{J}^{(N)}$.

In this framework, we have

Definition 23. *A tuple $(Q^N,(\psi^1,\ldots,\psi^N))\in\mathcal{Q}^{(N)}\times[\mathcal{A}^{(N)}]^N$ is said to be a Nash equilibrium (over open loop controls) if, for any other tuple $(\tilde{Q}^N,(\tilde{\psi}^1,\ldots,\tilde{\psi}^N))$ in the same class, the following $N+1$ inequalities hold true:*

$$\mathcal{R}^i\left(Q^N,(\psi^1,\ldots,\psi^N)\right)\le\mathcal{R}^i\left(Q^N,(\psi^1,\ldots,\psi^{i-1},\tilde{\psi}^i,\psi^{i+1},\ldots,\psi^N)\right),$$

for $i=1,\ldots,N$, and

$$\mathcal{J}^{(N)}\left(Q^N,(\psi^1,\ldots,\psi^N)\right)\ge\mathcal{J}^{(N)}\left(\tilde{Q}^N,(\psi^1,\ldots,\psi^N)\right).$$

5.2 Approximate Nash equilibria

Strategy induced by a mean-field equilibrium. Thanks to Theorem 14, we can consider one equilibrium to the mean-field game set over (MinMax[μ]). We denote q^* the Nature equilibrium state, and ψ^* the player equilibrium control. Both q^* and ψ^* are defined on the space $(\Omega,\mathcal{F},\mathbb{P})$. On the extended product space $(\Omega^{\times N},\mathcal{F}^{\times N},\mathbb{P}^{\times N})$, we let, for any $i=1,\ldots,N$,

$$q^{*,i}(\omega_1,\ldots,\omega_N)=q^*(\omega^i),\quad \psi^{*,i}(\omega_1,\ldots,\omega_N)=\psi^*(\omega^i),\quad (\omega_1,\ldots,\omega_N)\in\Omega^{\times N},$$

which makes it possible to define

$$Q^{*,N}:=\prod_{i=1}^N q^i.$$

Below, we write $X^{*,i}$ for $X^{i,\psi^{*,i}}$, and we represent $q^{*,i}$ in the form $q^{*,i}=\mathcal{E}.(\int_0^\cdot Z_s^{\star,*,i}\cdot\mathrm{d}W_s^i)$.

The strategy constructed in this way is called a *mean-field* strategy. It could also be referred to as a *distributed* strategy in the following sense :

Definition 24. *A strategy $(Q=q^1\ldots q^N,(\psi^1,\ldots,\psi^N))$ is said to be distributed if, for each $i\in\{1,\ldots,N\}$, (q^i,ψ^i) is $\sigma(\eta^i,W^i)$-measurable.*

Players deviating from the mean-field equilibrium. The purpose of the paragraphs below is to show that, in a certain sense, the mean-field strategy is an *approximate equilibrium* of the N-player game. The whole analysis is carried out under the following assumption, which is rather stronger than A7-A9 but which suffices to illustrate our approach:

Assumptions *(continued)*

A10 The function g is the sum of two functions

$$g(x, \mu) = g_0(x) + g_1(x, \mu),$$

where g_0 satisfies A7, and g_1 also satisfies A7, and is bounded and Lipschitz continuous in (x, μ), when $\mathcal{P}_{2-r}(\mathbb{R}^n)$ is equipped with the Fortet-Mourier distance.

It is easy to see that A8 and A9 are necessarily satisfied under A10.

In the statement below, we show that a player who unilaterally deviates from the mean-field strategy can only expect a modest reduction in their loss. This corresponds to the classical result in mean-field game theory.

Lemma 25. *Let A1–A10 be in force. Then, there exists a sequence $(\varepsilon_N)_{N\geq 1}$ converging to 0 such that, for any $i = 1, \dots, N$, and any $\tilde{\psi}^i \in \mathcal{A}^{(N)}$, it holds*

$$\mathcal{R}^i \left(Q^{*,N}, \psi^{*,1}, \dots, \psi^{*,i-1}, \tilde{\psi}^i, \psi^{*,i+1}, \dots, \psi^{*,N} \right) \geq \mathcal{R}^i \left(Q^{*,N}, \psi^{*,1}, \dots, \psi^{*,N} \right) - \varepsilon_N.$$

Proof. Throughout the proof, we use the convenient notation $\tilde{X}^i := X^{i,\tilde{\psi}^i}$.

By Assumption A10, we can find a constant C, independent of N, such that, for any $i \in \{1, \dots, N\}$.

$$\begin{aligned}
&\left| \mathbb{E}^{\times N} \left[Q_T^{*,N} \left(g \left(\tilde{X}_T^i, \mu_T^{\psi^{*,1},\dots,\tilde{\psi}^i,\dots,\psi^{*,N}} \right) - g \left(\tilde{X}_T^i, (q^*\mathbb{P})_{X^*} \right) \right) \right] \right| \\
&\leq C \mathbb{E}^{\times N} \left[Q_T^{*,N} \min \left(1, d_{\mathrm{FM}} \left(\frac{1}{N} \sum_{j \neq i} \delta_{X_T^{*,j}} + \frac{1}{N} \delta_{\tilde{X}_T^i}, (q^*\mathbb{P})_{X^*} \right) \right) \right].
\end{aligned}$$

By Lemma 22, the term on the second line tends to 0 as N tends to $+\infty$. Therefore, we can find a sequence $(\varepsilon_N)_{N\geq 1}$, independent of $\tilde{\psi}^i$, such that

$$\begin{aligned}
&\mathcal{R}^i \left(Q^{*N}, \psi^{*,1}, \dots, \psi^{*,i-1}, \tilde{\psi}^i, \psi^{*,i+1}, \dots, \psi^{*,N} \right) \\
&\geq \mathbb{E}^{\times N} \left[Q_T^{*,N} \left(g^i \left(\tilde{X}_T^i, (q^*\mathbb{P})_{X^*} \right) \right) \right] + \mathbb{E}^{\times N} \left[\int_0^T Q_s^{*,N} \ell^i(s, \tilde{\psi}_s^i) \mathrm{d}s \right] - \varepsilon_N.
\end{aligned}$$

By convexity properties of the function g^i in the space variable and of the function ℓ^i in the variable ψ, we get

$$\begin{aligned}
&\mathcal{R}^i \left(Q^{*,N}, \psi^{*,1}, \dots, \psi^{*,i-1}, \tilde{\psi}^i, \psi^{*,i+1}, \dots, \psi^{*,N} \right) \\
&\geq \mathbb{E}^{\times N} \left[Q_T^{*,N} \left(g \left(X_T^{*,i}, (q^*\mathbb{P})_{X^*} \right) \right) \right] + \mathbb{E}^{\times N} \left[\int_0^T Q_s^{*,N} \ell^i(s, \psi_s^{*,i}) \mathrm{d}s \right] \\
&\quad + \mathbb{E}^{\times N} \left[Q_T^{*,N} \nabla_x g \left(X_T^{*,i}, (q^*\mathbb{P})_{X^*} \right) \cdot (\tilde{X}_T^{*,i} - X_T^{*,i}) \right] \\
&\quad + \mathbb{E}^{\times N} \left[\int_0^T Q_s^N \nabla_\psi \ell^i(s, \psi_s^{*,i}) \cdot (\tilde{\psi}_s^i - \psi_s^{*,i}) \mathrm{d}s \right] - \varepsilon_N.
\end{aligned}$$

The rest of the proof is quite standard and just consists in verifying that the strategy $\psi^{*,i}$ is optimal. The only difficulty is that the process $\tilde{\psi}^i$ is defined on the product space $\Omega^{\times N}$. A careful inspection shows that the proof of [16, Lemma 40], which is based on Itô calculus arguments and from which we already derived Lemma 4, remains the same. This shows that the sum of the third and fourth terms on the right-hand side is equal to 0.

Reverting the computations, we deduce that

$$\mathcal{R}^i\left(Q^{*,N},\psi^{*,1},\dots,\psi^{*,i-1},\tilde{\psi}^i,\psi^{*,i+1},\dots,\psi^N\right) \geq \mathcal{R}^i\left(Q^{*,N},\psi^{*,1},\dots,\psi^{*,N}\right) - \varepsilon_N.$$

This completes the proof. □

Mimicking the empirical distribution under $\tilde{Q}^N \in \mathcal{Q}^{(N)}$ Deviations by Nature are more difficult to understand, due to the multiple correlations that may arise when modifying $Q^{*,N}$. To overcome this difficulty, we rely on a rewriting of the cost function, whose principle is as follows and applies only to distributed strategies.

Lemma 26. *Let h in (33) satisfy A7 and A10, and let $c>0$. Then, there exists a sequence $(\varepsilon_N)_{N\geq 1}$ converging to 0 such that, for any distributed strategy $(\tilde{Q}^N,(\tilde{\psi}^1,\dots,\tilde{\psi}^N)) \in \mathcal{Q}^{(N)} \times [\mathcal{A}^{(N)}]^N$ with $\sup_{i=1,\dots,N}\mathcal{S}(\tilde{q}^i) \leq c$ and $\sup_{i=1,\dots,N}\mathcal{S}^\star(\tilde{\psi}^i) \leq c$, it holds*

$$\begin{aligned}\Bigg|\mathbb{E}^{\times N}&\left[\tilde{Q}^N_T\frac{1}{N}\sum_{i=1}^N h\left(\tilde{X}^i_T,\mu_T^{\tilde{\psi}^1,\dots,\tilde{\psi}^N}\right)\right]\\ &-\frac{1}{N}\sum_{i=1}^N\mathbb{E}^{\times N}\left[\tilde{q}^i_T h\left(\tilde{X}^i_T,\frac{1}{N}\sum_{j=1}^N \tilde{q}^j_T\delta_{\tilde{X}^j_T}\right)\right]\Bigg| \leq \varepsilon_N.\end{aligned}$$

Proof. Throughout the proof, ϵ denotes a fixed positive real. Moreover, we let φ be a compactly supported function from $\mathbb{R}^n$ to $\mathbb{R}$ that is equal to the identity on the ball $B_n(0,A)$ of center 0 and of radius A, for a certain $A>0$, and satisfies $|\varphi(x)|\leq|x|$ for all $x\in\mathbb{R}^n$. With the shorthand notation

$$\hat{X}^i_T := \varphi(\tilde{X}^i_T), \quad i\in\{1,\dots,N\},$$

we deduce from condition A10 (for h) that, for a constant C independent of N and of $(\tilde{Q}^N,(\tilde{\psi}^1,\dots,\tilde{\psi}^N))$,

$$\begin{aligned}&\left|\mathbb{E}^{\times N}\left[\tilde{Q}^N_T h\left(\tilde{X}^1_T,\frac{1}{N}\sum_{i=1}^N\delta_{\tilde{X}^i_T}\right)\right]-\mathbb{E}^{\times N}\left[\tilde{Q}^N_T h\left(\hat{X}^1_T,\frac{1}{N}\sum_{i=1}^N\delta_{\hat{X}^i_T}\right)\right]\right|\\ &\leq C\mathbb{E}^{\times N}\left[\tilde{Q}^N_T\left(1\wedge|\tilde{X}^1_T-\hat{X}^1_T|+\frac{1}{N}\sum_{i=1}^N 1\wedge|\tilde{X}^i_T-\hat{X}^i_T|\right)\right]\\ &\leq C\mathbb{E}^{\times N}\left[\tilde{Q}^N_T\mathbb{1}_{\{|\tilde{X}^1_T|\geq A\}}+\frac{1}{N}\sum_{i=1}^N\tilde{Q}^N_T\mathbb{1}_{\{|\tilde{X}^i_T|\geq A\}}\right].\end{aligned} \tag{34}$$

Because the strategy $\tilde{Q}^N$ is distributed, we have

$$\mathbb{E}^{\times N}\left[\tilde{Q}^N_T\mathbb{1}_{\{|\tilde{X}^i_T|\geq A\}}\right]=\mathbb{E}^{\times N}\left[\tilde{q}^i_T\mathbb{1}_{\{|\tilde{X}^i_T|\geq A\}}\right].$$

Since $\mathcal{S}(\tilde{q}^i) \leq c$ and $\mathcal{S}^\star(\tilde{\psi}^i) \leq c$, we can choose A large enough, only depending on ϵ and c, such that the right-hand side is less than $\epsilon/(2C)$; see Lemma 3.

Next, we follow the proof of Lemma 22 and consider an N-tuple $(Y_1, \ldots, Y_N)$ of independent random variables, constructed on $(\Omega, \mathcal{F}, \mathbb{P})$, such that $Y_i \sim (\tilde{q}^i_T \mathbb{P}^{\times N})_{\hat{X}^i_T}$ for each $i \in \{1, \ldots, N\}$. We have

$$\mathbb{E}^{\times N}\left[\tilde{Q}^N_T h\left(\hat{X}^1, \frac{1}{N}\sum_{i=1}^N \delta_{\hat{X}^i_T}\right)\right] = \mathbb{E}\left[h\left(Y^1, \frac{1}{N}\sum_{i=1}^N \delta_{Y^i}\right)\right].$$

Following the standard L^4-proof of the law of large numbers, we can find a universal constant C such that, for any real-valued function φ that is bounded by 1 and 1-Lipschitz continuous on the ball $B_n(0, A)$,

$$\mathbb{E}\left[\left|\frac{1}{N}\sum_{i=1}^N \varphi(Y^i) - \int_{\mathbb{R}^n} \varphi(x)\mathrm{d}\mu^N(x)\right|^4\right] \leq \frac{C}{N^2}, \tag{35}$$

where

$$\mu^N = \frac{1}{N}\sum_{i=1}^N (\tilde{q}^i_T \mathbb{P}^{\times N})_{\hat{X}^i_T}.$$

Call now $(\varphi_k)_{k\geq 1}$ a sequence that is dense (for the sup norm topology on the ball $B_n(0, A)$) in the set of real-valued functions that are bounded by 1 and 1-Lipschitz continuous on the ball $B_n(0, A)$. We deduce from the above bound (together with Markov inequality) that, $\mathbb{P}$-a.s.,

$$\lim_{N\to\infty}\left|\frac{1}{N}\sum_{i=1}^N \varphi_k(Y^i) - \int_{\mathbb{R}^n} \varphi_k(x)\mathrm{d}\mu^N(x)\right| = 0.$$

And then, using the compactness of the collection of real-valued functions on $B_n(0, A)$ that are bounded by 1 and 1-Lipschitz continuous, we deduce that, $\mathbb{P}$-a.s.,

$$\lim_{N\to\infty} d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{i=1}^N \delta_{Y^i}, \mu^N\right) = 0.$$

Since the constant C in (35) is universal, it is easy to see that the rate is independent of $\tilde{Q}^N$, in the sense that the rate at which the sequence

$$\mathbb{P}\left(\left\{d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{i=1}^N \delta_{Y^i}, \mu^N\right) > \varepsilon\right\}\right)$$

tends to 0, for any given $\varepsilon > 0$, is independent of $\tilde{Q}^N$.

Combining with (34), we deduce that there exists a sequence $(\varepsilon_N)_{N\geq 1}$, as in the statement, but depending on A, such that

$$\left|\mathbb{E}^{\times N}\left[\tilde{Q}^N_T h\left(\tilde{X}^1_T, \frac{1}{N}\sum_{i=1}^N \delta_{\tilde{X}^i_T}\right)\right] - \mathbb{E}\left[h\left(Y^1, \mu^N\right)\right]\right| \leq \epsilon + \varepsilon_N,$$

which we rewrite as

$$\left|\mathbb{E}^{\times N}\left[\tilde{Q}^N_T h\left(\tilde{X}^1_T, \frac{1}{N}\sum_{i=1}^N \delta_{\tilde{X}^i_T}\right)\right] - \mathbb{E}^{\times N}\left[\tilde{q}^1_T h\left(\hat{X}^1_T, \mu^N\right)\right]\right| \leq \epsilon + \varepsilon_N.$$

For a given $a > 0$, we now let

$$\hat{q}_T^i := q_T^i \mathbb{1}_{\{\tilde{q}_T^i \le a\}}, \quad i = 1, \dots, N$$

and then,

$$\hat{\mu}^N := \frac{1}{N} \sum_{i=1}^N (\hat{q}_T^i \mathbb{P}^{\times N})_{\hat{X}_T^i}.$$

We have

$$d_{\mathrm{FM}}\left(\hat{\mu}^N, \mu^N\right) \le \frac{1}{N} \sum_{i=1}^N \mathbb{E}^{\times N}\left[\tilde{q}_T^i \mathbb{1}_{\{\tilde{q}_T^i \ge a\}}\right].$$

And, by (11), we can choose a large enough, only depending on c, such that the right-hand side is less than ϵ/C, where C is the Lipschitz constant of h with respect to the Fortet-Mourier distance. This shows that

$$\left| \mathbb{E}^{\times N}\left[\tilde{Q}_T^N h\left(\tilde{X}_T^1, \frac{1}{N}\sum_{i=1}^N \delta_{\tilde{X}_T^i}\right)\right] - \mathbb{E}^{\times N}\left[\tilde{q}_T^1 h\left(\hat{X}_T^1, \hat{\mu}^N\right)\right] \right| \le 2\epsilon + \varepsilon_N.$$

Following (35), we can find another constant, still denoted by C, only depending on a, such that, for any real-valued function φ on the ball $B_n(0, A)$, bounded by 1 and 1-Lipschitz continuous,

$$\mathbb{E}\left[\left| \frac{1}{N}\sum_{i=1}^N \hat{q}_T^i \varphi(\hat{X}_T^i) - \int_{\mathbb{R}^n} \varphi(x) \mathrm{d}\hat{\mu}^N(x) \right|^4\right] \le \frac{C}{N^2}. \tag{36}$$

Proceeding as before, this shows that, for any $\varepsilon > 0$, the sequence

$$\mathbb{P}\left(\left\{ d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{i=1}^N \hat{q}_T^i \delta_{\hat{X}_T^i}, \hat{\mu}^N\right) > \varepsilon \right\}\right)$$

tends to 0, with a rate that is independent of $\tilde{Q}^N$. And then, for a possibly new choice of the sequence $(\varepsilon_N)_{N\ge1}$,

$$\left| \mathbb{E}^{\times N}\left[\tilde{Q}_T^N h\left(\tilde{X}_T^1, \frac{1}{N}\sum_{i=1}^N \delta_{\tilde{X}_T^i}\right)\right] - \mathbb{E}^{\times N}\left[\tilde{q}_T^1 h\left(\hat{X}_T^1, \frac{1}{N}\sum_{i=1}^N \hat{q}_T^i \delta_{\hat{X}_T^i}\right)\right] \right| \le 2\epsilon + \varepsilon_N.$$

It remains to see from condition A10 that, for a new constant C,

$$\begin{aligned}
&\left| \mathbb{E}^{\times N}\left[\tilde{q}_T^1 h\left(\hat{X}_T^1, \frac{1}{N}\sum_{i=1}^N \tilde{q}_T^i \delta_{\hat{X}_T^i}\right)\right] - \mathbb{E}^{\times N}\left[\tilde{q}_T^1 h\left(\hat{X}_T^1, \frac{1}{N}\sum_{i=1}^N \hat{q}_T^i \delta_{\hat{X}_T^i}\right)\right] \right| \\
&\le C\mathbb{E}^{\times N}\left[\tilde{q}_T^1 \min\left(1, \frac{1}{N}\sum_{i=1}^N |\tilde{q}_T^i - \hat{q}_T^i|\right)\right] \\
&\le C\mathbb{E}^{\times N}\left[\tilde{q}_T^1 \min\left(1, \frac{1}{N}\sum_{i=1}^N \tilde{q}_T^i \mathbb{1}_{\{\tilde{q}_T^i \ge a\}}\right)\right].
\end{aligned}$$

And, thanks to (11), we can increase the value of a, only in function of c, so that the right-hand side is less than ϵ. This gives

$$\left| \mathbb{E}^{\times N}\left[\tilde{Q}_T^N h\left(\tilde{X}_T^1, \frac{1}{N}\sum_{i=1}^N \delta_{\tilde{X}_T^i}\right)\right] - \mathbb{E}^{\times N}\left[\tilde{q}_T^1 h\left(\hat{X}_T^1, \frac{1}{N}\sum_{i=1}^N \tilde{q}_T^i \delta_{\hat{X}_T^i}\right)\right] \right| \le 3\epsilon + \varepsilon_N.$$

Arguing in the same way, we can substitute $\tilde{X}^i_T$ for $\hat{X}^i_T$ in the above display, assuming that A is large enough and adding a new ϵ in the right-hand side. Substituting $(\tilde{X}^j_T, \hat{X}^j_T)$ for $(\tilde{X}^1_T, \hat{X}^1_T)$, for any $j = 2, \dots, N$, and averaging over the indices $j \in \{1, \dots, N\}$, we complete the proof. $\square$

Lemma 26 leads us to introduce a *surrogate reward* for nature:

Definition 27. *Given $(Q^N, (\psi^1, \dots, \psi^N)) \in \mathcal{Q}^{(N)} \times [\mathcal{A}^{(N)}]^N$, we define the surrogate reward for nature as*

$$\mathcal{J}^{(N)}_{\text{surrog}}\left(Q^N, \psi^1, \dots, \psi^N\right) := \mathbb{E}^{\times N}\left[\sum_{i=1}^N q^i_T h^i\left(X^{i,\psi^i}_T, \frac{1}{N}\sum_{j=1}^N q^j_T \delta_{X^{j,\psi^j}_T}\right)\right] + \mathbb{E}^{\times N}\left[\sum_{i=1}^N \int_0^T q^i_s \left(\ell^i(s, \psi^i_s) - f^{\star,i}(s, Z^{\star,i}_s)\right) \mathrm{d}s\right].$$

Lemma 26 ensures that, for distributed strategies $(Q^N, (\psi^1, \dots, \psi^N))$, the costs $\mathcal{J}^{(N)}_{\text{surrog}}\left(Q^N, \psi^1, \dots, \psi^N\right)$ and $\mathcal{J}^{(N)}\left(Q^N, \psi^1, \dots, \psi^N\right)$ are asymptotically close as $N \to \infty$. Although this result is restricted to distributed strategies, we focus below on the surrogate reward, even for non-distributed strategies. Implicitly, this leads to the construction of approximate Nash equilibria, but for the surrogate game. When the game is restricted to distributed strategies, approximate equilibria of the surrogate game are also approximate equilibria of the original game.

When the strategy derives from a mean-field equilibrium, the empirical measure in the surrogate reward is governed by the following form of large of large numbers, which can be established as in the second part of the proof of Lemma 26:

Lemma 28. *Let $(q^i, X^i)_{i \geq 1}$ be an I.I.D sequence with the law of (q, X) as common distribution on $(0, +\infty) \times \mathbb{R}^n$ (the sequence being constructed on $(\Omega, \mathcal{F}, \mathbb{P})$), where $\mathbb{E}[q] = 1$. Then, $\mathbb{P}$-almost surely,*

$$\lim_{N \to +\infty} \frac{1}{N}\sum_{i=1}^N q^i \delta_{X^i} = (q\mathbb{P})_X,$$

the limiting being understood for the narrow convergence.

Nature locally deviating from the mean equilibrium. In this paragraph, we choose $h = g$ in (33).

Our goal is to show that, in this case, the surrogate reward cannot increase as a result of a local deviation, that is, when only the weight q^i corresponding to the noise W^i to which player i is subjected is modified. Below, we denote by $\mathcal{Q}[i]$ the collection of Doléans-Dade exponentials $\tilde{q}^i$ of the form $(\tilde{q}^i_t := \mathcal{E}_t(\int_0^\cdot \tilde{Z}^{\star,i}_s \cdot \mathrm{d}W^i_s))_{t \in [0,T]}$ with $\tilde{Z}^{\star,i}$ being $\mathbb{F}^N$-progressively measurable and satisfying

$$\mathcal{S}(\tilde{q}^i) := \mathbb{E}\int_0^T \tilde{q}^i_s f^{\star,i}(s, \tilde{Z}^{\star,i}_s)\mathrm{d}s < +\infty.$$

Lemma 29. *Let $c > 0$. There exists a sequence $(\varepsilon_N)_{N>0}$ converging to 0 such that, for any $i \in \{1, \dots, N\}$ and any $\tilde{q}^i \in \mathcal{Q}[i]$ such that $\mathcal{S}(\tilde{q}^i) \leq c$,*

$$\mathcal{J}^{(N)}_{\text{surrog}}\left(\tilde{Q}^{*,N}, \psi^{*,1}, \dots, \psi^{*,N}\right) \leq \mathcal{J}^{(N)}_{\text{surrog}}\left(Q^{*,N}, \psi^{*,1}, \dots, \psi^{*,N}\right) + N\varepsilon_N,$$

with

$$\tilde{Q}_T^N = q_T^{*,1} \dots q_T^{*,i-1} \tilde{q}_T^i q_T^{*,i+1} \dots q_T^{*,N}.$$

Proof. By Lipschitz property A10 of g,

$$\begin{aligned}
&\mathbb{E}^{\times N}\left[\tilde{q}_T^i \left(g\left(\tilde{X}_T^i, \frac{1}{N}\sum_{j\neq i} q_T^{*,j}\delta_{X_T^{*,j}} + \frac{1}{N}\tilde{q}_T^i \delta_{X_T^{*,i}} \right) - g\left(\tilde{X}_T^i, (q^*\mathbb{P})_{X_T^*}\right)\right)\right] \\
&\leq C\mathbb{E}^{\times N}\left[\tilde{q}_T^i \min\left(1, d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{j=1}^N q_T^{*,j}\delta_{X_T^{*,j}}, (q^*\mathbb{P})_{X_T^*}\right)\right)\right] \\
&\quad + C\mathbb{E}^{\times N}\left[\min\left(1, \frac{1}{N}|q_T^{*,i} - \tilde{q}_T^i|\right)\right].
\end{aligned}$$

By Lemma 28, we know that $\mathbb{P}$-a.s.,

$$\lim_{N\to\infty} d_{\mathrm{FM}}\left(\frac{1}{N}\sum_{j=1}^N q_T^{*,j}\delta_{X_T^{*,j}}, (q^*\mathbb{P})_{X_T^*}\right) = 0.$$

Since $\mathcal{S}(\tilde{q}^i) \leq c$, we can use (11) to find a sequence $(\varepsilon_N)_{N>0}$ converging to 0, only depending on $\tilde{q}^i$ via c, such that

$$\mathbb{E}^{\times N}\left[\tilde{q}_T^i \left(g\left(\tilde{X}_T^i, \frac{1}{N}\sum_{j=1}^N q_T^{*,j}\delta_{X_T^{*,j}} \right) - g\left(\tilde{X}_T^i, (q^*\mathbb{P})_{X_T^*}\right)\right)\right] \leq \varepsilon_N.$$

Of course, we can proceed similarly with the coordinates $k \neq i$.

Hence, denoting by $\tilde{Z}^{\star,i}$ the representative of $\tilde{q}^i$, i.e. $\tilde{q}^i = \mathcal{E}_\cdot(\int_0^\cdot \tilde{Z}_s^{\star,i}\cdot \mathrm{d}W_s^i)$, we obtain

$$\begin{aligned}
\mathcal{J}_{\mathrm{surrog}}^{(N)}\left(\tilde{Q}^N, \psi^{*,1}, \dots, \psi^{*,N}\right) \leq\ & \mathbb{E}^{\times N}\left[\tilde{q}_T^i g\left(X_T^{*,i}, (q_T^*\mathbb{P})_{X_T^*}\right)\right] \\
& + \mathbb{E}^{\times N}\left[\int_0^T \tilde{q}_s^i \left(\ell^i(s, \psi_s^{*,i}) - f^{\star,i}(s, \tilde{Z}_s^{\star,i})\right)\mathrm{d}s\right] \\
& + \sum_{j\neq i}\Big\{ \mathbb{E}^{\times N}\left[q_T^{*,j} g\left(X_T^{*,j}, (q_T^*\mathbb{P})_{X_T^*}\right)\right] \\
& + \mathbb{E}^{\times N}\left[\int_0^T q_s^{*,j}\left(\ell^j(s, \psi_s^{*,j}) - f^{\star,i}(s, Z_s^{\star,*,j})\right)\mathrm{d}s\right]\Big\} \\
& + N\varepsilon_N.
\end{aligned}$$

To handle the first two terms on the right-hand side, we use the optimality of $q^{*,i}$. As in the proof of Lemma 25, the main subtlety comes from the fact that the probability space is not supported by Ω but by $\Omega^{\times N}$. That said, we can apply the same Itô expansion as in the proof of [16, Lemma 32] to show that

$$\begin{aligned}
&\mathbb{E}^{\times N}\left[\tilde{q}_T^i g\left(X_T^{*,i}, (q_T^*\mathbb{P})_{X_T^*}\right)\right] + \mathbb{E}^{\times N}\left[\int_0^T \tilde{q}_s^i\left(\ell^i(s, \psi_s^{*,i}) - f^{\star,i}(s, \tilde{Z}_s^{\star,i})\right)\mathrm{d}s\right] \\
&\leq \mathbb{E}^{\times N}\left[q_T^{*,i} g\left(X_T^{*,i}, (q_T^*\mathbb{P})_{X_T^*}\right)\right] + \mathbb{E}^{\times N}\left[\int_0^T q_s^{*,i}\left(\ell^i(s, \psi_s^{*,i}) - f^{\star,i}(s, Z_s^{\star,*,i})\right)\mathrm{d}s\right],
\end{aligned}$$

from which we deduce that

$$
\begin{aligned}
\mathcal{J}^{(N)}_{\text{surrog}}\left(\tilde{Q}^N, \psi^{*,1}, \dots, \psi^{*,N}\right) \le \sum_{j=1}^N \mathbb{E}^{\times N}\left[q_T^{*,j} g\left(X_T^{*,j}, (q_T^*\mathbb{P})_{X_T^*}\right)\right] \\
+ \mathbb{E}^{\times N}\left[\int_0^T q_s^{*,j}\left(\ell^j(s,\psi_s^{*,j}) - f^{\star,i}(s, Z_s^{\star,*,j})\right) \mathrm{d}s\right] \\
+ N\varepsilon_N.
\end{aligned}
$$

Repeating the computations, but with $q^{*,i}$ substituted for $\tilde{q}^i$, we get

$$
\mathcal{J}^{(N)}_{\text{surrog}}\left(\tilde{Q}^N, \psi^{*,1}, \dots, \psi^{*,N}\right) \le \mathcal{J}^{(N)}_{\text{surrog}}\left(Q^{*,N}, \psi^{*,1}, \dots, \psi^{*,N}\right) + N\varepsilon_N.
$$

This completes the proof. □

Nature globally deviating when the game is potential In this paragraph, we assume that there exists a smooth function $G : \mathcal{M}(\mathbb{R}^n) \to \mathbb{R}$ such that

$$
g_1(x,\mu) = \frac{\delta G}{\delta \mu}(\mu, x) := \frac{d}{d\varepsilon}|_{\varepsilon=0+} G\left(\mu + \varepsilon\delta_x\right), \quad x \in \mathbb{R}^n.
$$

We still assume A7 and A10.

We choose, as surrogate cost,

$$
\begin{aligned}
\mathcal{J}^{(N)}_{\text{surrog}}\left(Q^N, \psi^1, \dots, \psi^N\right) := \mathbb{E}^{\times N}\left[\sum_{i=1}^N q_T^i g_0(X_T^i) + NG\left(\frac{1}{N}\sum_{j=1}^N q_T^j \delta_{X_T^{j,\psi^j}}\right)\right] \\
+ \mathbb{E}^{\times N}\left[\sum_{i=1}^N \int_0^T q_s^i\left(\ell^i(s,\psi_s^i) - f^{\star,i}(s, Z_s^{\star,i})\right) \mathrm{d}s\right].
\end{aligned}
$$

Lemma 30. *Assume that the function G is flat concave on the cone of non-negative measures. Then, for any constant $c > 0$, there exists a sequence $(\varepsilon_N)_{N\ge 1}$, converging to 0 such that, for any $\tilde{Q}^N \in \mathcal{Q}^{(N)}$, with $\sup_{i=1,\dots,N} \mathcal{S}(\tilde{q}^i) \le c$,*

$$
\mathcal{J}^{(N)}_{\text{surrog}}\left(\tilde{Q}^N, \psi^{*,1}, \dots, \psi^{*,N}\right) \le \mathcal{J}^{(N)}_{\text{surrog}}\left(\tilde{Q}^N, \psi^{*,1}, \dots, \psi^{*,N}\right) + N\varepsilon_N,
$$

Proof. By concavity of G,

$$
\begin{aligned}
\mathbb{E}^{\times N}\left[NG\left(\frac{1}{N}\sum_{j=1}^N \tilde{q}_T^j \delta_{X_T^{*,j}}\right)\right] \le \mathbb{E}^{\times N}\left[NG\left(\frac{1}{N}\sum_{j=1}^N q_T^{*,j} \delta_{X_T^{*,j}}\right)\right] \\
+ \sum_{i=1}^N \mathbb{E}^{\times N}\left[\left(\tilde{q}_T^i - q_T^{*,i}\right)\frac{\delta G}{\delta\mu}\left(\frac{1}{N}\sum_{j=1}^N q_T^{*,j}\delta_{X_T^{*,j}}, X_T^{*,i}\right)\right].
\end{aligned}
$$

By proceeding as in the proof of Lemma 29, we deduce that there exists a sequence $(\varepsilon_N)_{N\ge1}$ converging to 0 and only depending on $\tilde{Q}^N$ via c such that

$$
\begin{aligned}
\mathbb{E}^{\times N}\left[NG\left(\frac{1}{N}\sum_{j=1}^N \tilde{q}_T^j \delta_{X_T^{*,j}}\right)\right] \le \mathbb{E}^{\times N}\left[NG\left(\frac{1}{N}\sum_{j=1}^N q_T^{*,j} \delta_{X_T^{*,j}}\right)\right] \\
+ \sum_{i=1}^N \mathbb{E}^{\times N}\left[\left(\tilde{q}_T^i - q_T^{*,i}\right)\frac{\delta G}{\delta\mu}\left((q_T^*\mathbb{P})_{X_T^*}, X_T^{*,i}\right)\right] + N\varepsilon_N.
\end{aligned}
$$

Once again by [16, Lemma 32],

$$\sum_{i=1}^{N} \mathbb{E}^{\times N} \left[\left(\tilde{q}_T^i - q_T^{*,i} \right) \frac{\delta G}{\delta \mu} \left((q_T^* \mathbb{P})_{X_T^*}, X_T^{*,i} \right) \right]$$
$$\leq - \sum_{i=1}^{N} \mathbb{E}^{\times N} \left[\int_0^T \left(\tilde{q}_s^i - q_s^{*,i} \right) \ell^i (s, \psi_s^{*,i}) \mathrm{d}s \right]$$
$$- \sum_{i=1}^{N} \mathbb{E}^{\times N} \left[\int_0^T \left(\tilde{q}_s^i f^{i,\star}(s, \tilde{Z}_s^{\star,i}) - q_s^{*,i} f^{i,\star}(s, Z_s^{\star,*,i}) \right) \mathrm{d}s \right].$$

Combining the last two displays, we easily complete the proof. □

Acknowledgment

F. Delarue and P. Lavigne acknowledge the financial support of the European Research Council (ERC) under the European Union's Horizon Europe research and innovation program (ELISA project, Grant agreement No. 101054746). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

[1] Yves Achdou, Guillaume Carlier, Quentin Petit, and Daniela Tonon. A simple city equilibrium model with an application to teleworking. Applied Mathematics & Optimization, 88(2):60, 2023.

[2] Clémence Alasseur, Imen Ben Tahar, and Anis Matoussi. An extended mean field game for storage in smart grids. Journal of Optimization Theory and Applications, 184(2):644–670, 2020.

[3] Charalambos D. Aliprantis and Kim C. Border. Infinite dimensional analysis. Springer, Berlin, third edition, 2006. A hitchhiker's guide.

[4] Dario Bauso, Hamidou Tembine, and Tamer Başar. Robust mean field games. Dynamic games and applications, 6(3):277–303, 2016.

[5] Alain Bensoussan, Boualem Djehiche, Hamidou Tembine, and Phillip Yam. Risk-sensitive mean-field-type control. In 2017 IEEE 56th Annual Conference on Decision and Control (CDC), pages 33–38. IEEE, 2017.

[6] Charles Bertucci, Jean-Michel Lasry, and Pierre-Louis Lions. Some remarks on mean field games. Communications in Partial Differential Equations, 44(3):205–227, 2019.

[7] J. Frédéric Bonnans, Saeed Hadikhanloo, and Laurent Pfeiffer. Schauder estimates for a class of potential mean field games of controls. Applied Mathematics & Optimization, 83:1431–1464, 2021.

[8] Philippe Briand and Ying Hu. BSDE with quadratic growth and unbounded terminal value. Probability Theory and Related Fields, 136(4):604–618, 2006.

[9] Philippe Briand and Ying Hu. Quadratic BSDEs with convex generators and unbounded terminal conditions. Probability Theory and Related Fields, 141:543–567, 2008.

[10] Pierre Cardaliaguet, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. The master equation and the convergence problem in mean field games, volume 201 of Ann. Math. Stud. Princeton, NJ: Princeton University Press, 2019.

[11] Pierre Cardaliaguet, P. Jameson Graber, Alessio Porretta, and Daniela Tonon. Second order mean field games with degenerate diffusion and local coupling. Nonlinear Differential Equations and Applications NoDEA, 22(5):1287–1317, 2015.

[12] Pierre Cardaliaguet and Charles-Albert Lehalle. Mean field game of controls and an application to trade crowding. Mathematics and Financial Economics, 12(3):335–363, 2018.

[13] René Carmona and François Delarue. Probabilistic theory of mean field games with applications I. Mean field FBSDEs, control, and games, volume 83 of Probab. Theory Stoch. Model. Cham: Springer, 2018.

[14] René Carmona, François Delarue, and Daniel Lacker. Mean field games with common noise. The Annals of Probability, 44(6):3740 – 3803, 2016.

[15] Ziteng Cheng and Sebastian Jaimungal. Risk-averse mean field games: exploitability and non-asymptotic analysis. arXiv, 2301.06930, 2023.

[16] François Delarue and Pierre Lavigne. Robust mean field control: stochastic maximum principle and variational mean field games. Arxiv, 2026.

[17] Claude Dellacherie and Paul-André Meyer. Probabilities and potential. B, volume 72 of North-Holland Mathematics Studies. North-Holland Publishing Co., Amsterdam, 1982.

[18] Célia Escribe, Josselin Garnier, and Emmanuel Gobet. A mean field game model for renewable investment under long-term uncertainty and risk aversion. Dynamic Games and Applications, 14(5):1093–1130, 2024.

[19] Olivier Féron, Peter Tankov, and Laura Tinsi. Price formation and optimal trading in intraday electricity markets. In Network Games, Control and Optimization: 10th International Conference, NetGCooP 2020, France, September 22–24, 2021, Proceedings 10, pages 294–305. Springer, 2021.

[20] J. Frédéric Bonnans, Pierre Lavigne, and Laurent Pfeiffer. Discrete-time mean field games with risk-averse agents. ESAIM: Control, Optimisation and Calculus of Variations, 27:44, 2021.

[21] Minyi Huang, Peter E. Caines, and Roland P. Malhamé. Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior

and decentralized ε-nash equilibria. IEEE transactions on automatic control, 52(9):1560–1571, 2007.

[22] Ziad Kobeissi, Idriss Mazari-Fouquer, and Domènec Ruiz-Balet. The tragedy of the commons: A Mean-Field Game approach to the reversal of travelling waves. Nonlinearity, 37(11):115010, 2024.

[23] Jean-Michel Lasry and Pierre-Louis Lions. Jeux à champ moyen. i–le cas stationnaire. Comptes Rendus Mathématique, 343(9):619–625, 2006.

[24] Jean-Michel Lasry and Pierre-Louis Lions. Jeux à champ moyen. ii–horizon fini et contrôle optimal. Comptes Rendus Mathématique, 343(10):679–684, 2006.

[25] Pierre Lavigne and Peter Tankov. Decarbonization of financial markets: a mean-field game approach. arXiv, 2301.09163, 2023.

[26] Jun Moon and Tamer Başar. Linear quadratic risk-sensitive and robust mean field games. IEEE Transactions on Automatic Control, 62(3):1062–1077, 2016.

[27] Charafeddine Mouzouni. Topic in mean field games theory & applications in economics and quantitative finance. PhD thesis, Université de Lyon, 2019.

[28] Naci Saldi, Tamer Başar, and Maxim Raginsky. Approximate markov-nash equilibria for discrete-time risk-sensitive mean-field games. Mathematics of Operations Research, 45(4):1596–1620, 2020.

[29] Hamidou Tembine, Quanyan Zhu, and Tamer Başar. Risk-sensitive mean-field games. IEEE Transactions on Automatic Control, 59(4):835–850, 2013.

[30] Muhammad Aneeq Uz Zaman, Mathieu Lauriere, Alec Koppel, and Tamer Başar. Robust cooperative multi-agent reinforcement learning: A mean-field type game perspective. In 6th Annual Learning for Dynamics & Control Conference, pages 770–783. PMLR, 2024.