

On the connection between symmetric N -player games and mean field games

Markus Fischer*

May 14, 2014; revised September 7, 2015; modified April 15, 2016

Abstract

Mean field games are limit models for symmetric N -player games with interaction of mean field type as $N \rightarrow \infty$. The limit relation is often understood in the sense that a solution of a mean field game allows to construct approximate Nash equilibria for the corresponding N -player games. The opposite direction is of interest, too: When do sequences of Nash equilibria converge to solutions of an associated mean field game? In this direction, rigorous results are mostly available for stationary problems with ergodic costs. Here, we identify limit points of sequences of certain approximate Nash equilibria as solutions to mean field games for problems with Itô-type dynamics and costs over a finite time horizon. Limits are studied through weak convergence of associated normalized occupation measures and identified using a probabilistic notion of solution for mean field games.

2000 AMS subject classifications: 60B10, 60K35, 91A06, 93E20

Key words and phrases: Nash equilibrium; mean field game; McKean-Vlasov limit; weak convergence; martingale problem; optimal control

1 Introduction

Mean field games, as introduced by J.M. Lasry and P.-L. Lions [Lasry and Lions, 2006a,b, 2007] and, independently, by M. Huang, R.P. Malhamé, and

*Department of Mathematics, University of Padua, via Trieste 63, 35121 Padua, Italy. The starting point for the present work was a course on mean field games given by Pierre Cardaliaguet at the University of Padua in 2013. The author is grateful to Martino Bardi and Pierre Cardaliaguet for stimulating discussions. The author thanks two anonymous Referees for their critique, comments, and helpful suggestions. Financial support was provided by the University of Padua through the Project “Stochastic Processes and Applications to Complex Systems” (CPDA123182).

P.E. Caines [Huang et al., 2006, and subsequent works], are limit models for symmetric non-zero-sum non-cooperative N -player games with interaction of mean field type as the number of players tends to infinity. The limit relation is often understood in the sense that a solution of the mean field game allows to construct approximate Nash equilibria for the corresponding N -player games if N is sufficiently large; see, for instance, Huang et al. [2006], Kolokoltsov et al. [2011], Carmona and Delarue [2013], and Carmona and Lacker [2015]. This direction is useful from a practical point of view since the model of interest is commonly the N -player game with N big so that a direct computation of Nash equilibria is not feasible.

The opposite direction in the limit relation is of interest, too: When and in which sense do sequences of Nash equilibria for the N -player games converge to solutions of a corresponding mean field game? An answer to this question is useful as it provides information on what kind of Nash equilibria can be captured by the mean field game approach. In view of the theory of McKean-Vlasov limits and propagation of chaos for uncontrolled weakly interacting systems [cf. McKean, 1966, Sznitman, 1991], one may expect to obtain convergence results for broad classes of systems, at least under some symmetry conditions on the Nash equilibria. This heuristic was the original motivation in the introduction of mean field games by Lasry and Lions. Rigorous results supporting it are nonetheless few, and they mostly apply to stationary problems with ergodic costs and special structure (in particular, affine-linear dynamics and convex costs); see Lasry and Lions [2007], Feleqi [2013], Bardi and Priuli [2013, 2014]. For non-stationary problems, the passage to the limit has been established rigorously in Gomes et al. [2013] for a class of continuous-time finite horizon problems with finite state space, but only if the time horizon is sufficiently small. Moreover, in the situation studied there, Nash equilibria for the N -player games are unique in a class of symmetric Markovian feedback strategies. The above cited works on the passage to the limit all employ methods from the theory of ordinary or partial differential equations, in particular, equations of Hamilton-Jacobi-Bellman-type.

In Lacker [2015b], which appeared as preprint three months after submission of the present paper, a general characterization of the limit points of N -player Nash equilibria is obtained through probabilistic methods. We come back to that work, which also covers mean field games with common noise, in the second but last paragraph of this section. Lastly, we mention the even more recent work by Cardaliaguet et al. [2015] on the passage to the mean field game limit; see Remark 3.2 below.

Here, we study the limit relation between symmetric N -player games and

mean field games in the direction of the Lasry-Lions heuristic for continuous time finite horizon problems with fairly general cost structure and Itô-type dynamics. The aim is to identify limit points of sequences of symmetric Nash equilibria for the N -player games as solutions of a mean field game. For a general introduction to mean field games, see Cardaliaguet [2013] or Carmona et al. [2013]. The latter work also explains the difference in the limit relation that distinguishes mean field games from optimal control problems of McKean-Vlasov-type.

To describe the prelimit systems, let $X_i^N(t)$ denote the state of player i at time t in the N -player game, and denote by $u_i(t)$ the control action that he or she chooses at time t . Individual states will be elements of \mathbb{R}^d , while control actions will be elements of some closed set $\Gamma \subset \mathbb{R}^{d_2}$. The evolution of the individual states is then described by the Itô stochastic differential equations

$$(1.1) \quad dX_i^N(t) = b(t, X_i^N(t), \mu^N(t), u_i(t))dt + \sigma(t, X_i^N(t), \mu^N(t))dW_i^N(t),$$

$i \in \{1, \dots, N\}$, where W_1^N, \dots, W_N^N are independent standard Wiener processes, and $\mu^N(t)$ is the empirical measure of the system at time t :

$$\mu^N(t) \doteq \frac{1}{N} \sum_{i=1}^N \delta_{X_i^N(t)}.$$

Notice that the coefficients b, σ in Eq. (1.1) are the same for all players. We will assume b, σ to be continuous in the time variable, Lipschitz continuous in the state and measure variable, where we use the square Wasserstein metric as a distance on probability measures, and of sub-linear growth. The dispersion coefficient σ does not depend on the control variable, but it may depend on the measure-variable. Moreover, σ is allowed to be degenerate. Deterministic systems are thus covered as a special case.

The individual dynamics are explicitly coupled only through the empirical measure process μ^N . There is also an implicit coupling, namely through the strategies u_1, \dots, u_N , which may depend on non-local information; in particular, a strategy u_i might depend, in a non-anticipative way, on X_j^N or W_j^N for $j \neq i$. In this paper, strategies will always be stochastic open-loop. In particular, strategies will be processes adapted to a filtration that represents the information available to the players. We consider two types of information: full information, which is the same for all players and is represented by a filtration at least as big as the one generated by the initial states and the Wiener processes, and local information, which is player-dependent and, for player i , is represented by the filtration generated by his/her own initial state and the Wiener process W_i^N .

Let $\mathbf{u} = (u_1, \dots, u_N)$ be a strategy vector, i.e., an N -vector of Γ -valued processes such that u_i is a strategy for player i , $i \in \{1, \dots, N\}$. Player i evaluates the effect of the strategy vector \mathbf{u} according to the cost functional

$$J_i^N(\mathbf{u}) \doteq \mathbf{E} \left[\int_0^T f(s, X_i^N(s), \mu^N(s), u_i(s)) ds + F(X_i^N(T), \mu^N(T)) \right],$$

where $T > 0$ is the finite time horizon, (X_1^N, \dots, X_N^N) the solution of the system (1.1) under \mathbf{u} , and μ^N the corresponding empirical measure process. The cost coefficients f, F , which quantify running and terminal costs, respectively, are assumed to be continuous in the time and control variable, locally Lipschitz continuous in the state and measure variable, and of sub-quadratic growth. The action space Γ is assumed to be closed, but not necessarily compact; in the non-compact case, f will be quadratically coercive in the control. The assumptions on the coefficients are chosen so that they cover some linear-quadratic problems, in addition to many genuinely non-linear problems.

If there were no control in Eq. (1.1) (i.e., b independent of the control variable) and if the initial states for the N -player games were independent and identically distributed with common distribution not depending on N , then (X_1^N, \dots, X_N^N) would be exchangeable for every $N \in \mathbb{N}$ and, under our assumptions on b, σ , the sequence (μ^N) of empirical measure processes would converge to some deterministic flow of probability measures:

$$\mu^N(t) \xrightarrow{N \rightarrow \infty} \mathbf{m}(t), \quad \text{in distribution / probability.}$$

This convergence would also hold for the sequence of path-space empirical measures, which, by symmetry and the Tanaka-Sznitman theorem, is equivalent to the propagation of chaos property for the triangular array $(X_i^N)_{i \in \{1, \dots, N\}, N \in \mathbb{N}}$. In particular, $\text{Law}(X_i^N(t)) \rightarrow \mathbf{m}(t)$ as $N \rightarrow \infty$ for each fixed index i , and \mathbf{m} would be the flow of laws for the uncontrolled McKean-Vlasov equation

$$dX(t) = b(t, X(t), \mathbf{m}(t))dt + \sigma(t, X(t), \mathbf{m}(t))dW(t), \quad \mathbf{m}(t) = \text{Law}(X(t)).$$

The above equation would determine the flow of measures \mathbf{m} .

Now, for $N \in \mathbb{N}$, let \mathbf{u}^N be a strategy vector for the N -player game. For the sake of argument, let us suppose that $\mathbf{u}^N = (u_1^N, \dots, u_N^N)$ is a symmetric Nash equilibrium for each N (symmetric in the sense that the finite sequence $((X_1^N(0), u_1^N, W_1^N), \dots, (X_N^N(0), u_N^N, W_N^N))$ is exchangeable). If the mean field heuristic applies, then the associated sequence of empirical measure processes $(\mu^N)_{N \in \mathbb{N}}$ converges in distribution to some deterministic flow

of probability measures \mathbf{m} . In this case, \mathbf{m} should be the flow of measures induced by the solution of the controlled equation

$$(1.2) \quad dX(t) = b(t, X(t), \mathbf{m}(t), u(t))dt + \sigma(t, X(t), \mathbf{m}(t))dW(t),$$

where the control process u should, by the Nash equilibrium property of the N -player strategies, be optimal for the control problem

$$(1.3) \quad \text{minimize } J_{\mathbf{m}}(v) \doteq \mathbf{E} \left[\int_0^T f(s, X(s), \mathbf{m}(s), v(s)) ds + F(X(T), \mathbf{m}(T)) \right]$$

over all admissible v subject to: X solves Eq. (1.2) under v .

The mean field game, which is the limit system for the N -player games, can now be described as follows: For each flow of measures \mathbf{m} , solve the optimal control problem (1.3) to find an optimal control $u^{\mathbf{m}}$ with corresponding state process $X^{\mathbf{m}}$. Then choose a flow of measures \mathbf{m} according to the mean field condition $\mathbf{m}(\cdot) = \text{Law}(X^{\mathbf{m}}(\cdot))$. This yields a solution of the mean field game, which can be identified with the pair $(\text{Law}(X^{\mathbf{m}}, u^{\mathbf{m}}, W), \mathbf{m})$; see Definition 4.3 below. We include the driving noise process W in the definition of the solution, as it is the joint distribution of initial condition, control process and noise process that determines the law of a solution to Eq. (1.2). If $(\text{Law}(X^{\mathbf{m}}, u^{\mathbf{m}}, W), \mathbf{m})$ is a solution of the mean field game, then, thanks to the mean field condition, $X^{\mathbf{m}}$ is a McKean-Vlasov solution of the controlled equation (1.2); moreover, $X^{\mathbf{m}}$ is an optimally controlled process for the standard optimal control problem (1.3) with cost functional $J_{\mathbf{m}}$. Clearly, neither existence nor uniqueness of solutions of the mean field game are a priori guaranteed.

In order to connect sequences of Nash equilibria with solutions of the mean field game in a rigorous way, we associate strategy vectors for the N -player games with normalized occupation measures or path-space empirical measures; see Eq. (5.1) in Section 5 below. Those occupation measures are random variables with values in the space of probability measures on an extended canonical space $\mathcal{Z} \doteq \mathcal{X} \times \mathcal{R}_2 \times \mathcal{W}$, where \mathcal{X} , \mathcal{W} are path spaces for the individual state processes and the driving Wiener processes, respectively, and \mathcal{R}_2 is a space of Γ -valued relaxed controls. Observe that \mathcal{Z} contains a component for the trajectories of the driving noise process. Let (\mathbf{u}^N) be a sequence such that, for each $N \in \mathbb{N}$, \mathbf{u}^N is a strategy vector for the N -player game (not necessarily a Nash equilibrium). Let (Q_N) be the associated normalized occupation measures; thus, Q_N is the empirical measure of $((X_1^N, u_1^N, W_1^N), \dots, (X_N^N, u_N^N, W_N^N))$ seen as a random element of $\mathcal{P}(\mathcal{Z})$. We then show the following:

1. The family $(Q_N)_{N \in \mathbb{N}}$ is pre-compact under a mild uniform integrability condition on strategies and initial states; see Lemma 5.1.
2. Any limit random variable Q of (Q_N) takes values in the set of McKean-Vlasov solutions of Eq. (1.2) with probability one; see Lemma 5.3.
3. Suppose that (\mathbf{u}^N) is a sequence of local approximate Nash equilibria (cf. Definition 3.1). If Q is a limit point of (Q_N) such that the flow of measures induced by Q is deterministic with probability one, then Q takes values in the set of solutions of the mean field game with probability one; see Theorem 5.1.

The hypothesis in Point 3 above that the flow of measures induced by Q is deterministic with probability one means that the corresponding subsequence of (μ^N) , the empirical measure processes, converges in distribution to a deterministic flow of probability measures \mathbf{m} . This is a strong hypothesis, essentially part of the mean field heuristic; nonetheless, it is satisfied if \mathbf{u}^N is a vector of independent and identically distributed individual strategies for each N , where the common distribution is allowed to vary with N ; see Corollary 5.2. While Nash equilibria for the N -player games with independent and identically distributed individual strategies do not exist in general, local approximate Nash equilibria with i.i.d. components do exist, at least under the additional assumption of compact action space Γ and bounded coefficients; see Proposition 3.1. In this situation, the passage to the mean field game limit is justified.

For the passage to the limit required by Point 2 above, we have to identify solutions of Eq. (1.2), which describes the controlled dynamics of the limit system. To this end, we employ a local martingale problem in the spirit of Stroock and Varadhan [1979]. The use of martingale problems, together with weak convergence methods, has a long tradition in the analysis of McKean-Vlasov limits for uncontrolled weakly interacting systems [for instance, Funaki, 1984, Oelschläger, 1984] as well as in the study of stochastic optimal control problems. Controlled martingale problems are especially powerful in combination with relaxed controls; see El Karoui et al. [1987], Kushner [1990], and the references therein. In the context of mean field games, a martingale problem formulation has been used by Carmona and Lacker [2015] to establish existence and uniqueness results for non-degenerate systems and, more recently, by Lacker [2015a], where existence of solutions is established for mean field games of the type studied here; the assumptions on the coefficients are rather mild, allowing for degenerate as well as control-dependent diffusion coefficient. The notion of solution we give in Definition 4.3 below

corresponds to the notion of “relaxed mean field game solution” introduced in Lacker [2015a].

The martingale problem formulation for the controlled limit dynamics we use here is actually adapted from the joint work Budhiraja et al. [2012], where we studied large deviations for weakly interacting Itô processes through weak convergence methods. While the passage to the limit needed there for obtaining convergence of certain Laplace functionals is analogous to the convergence result of Point 2 above, the limit problems in Budhiraja et al. [2012] are not mean field games; they are, in fact, optimal control problems of McKean-Vlasov type, albeit with a particular structure. As a consequence, optimality has to be verified in a different way: In order to establish Point 3 above, we construct an asymptotically approximately optimal competitor strategy in noise feedback form (i.e., as a function of time, initial condition, and the trajectory of the player’s noise process up to current time), which is then applied to exactly one of the N players for each N ; this yields optimality of the limit points thanks to the Nash equilibrium property of the prelimit strategies. If the limit problem were of McKean-Vlasov type, one would use a strategy selected according to a different optimality criterion and apply it to all components (or players) of the prelimit systems.

In the work by Lacker [2015b] mentioned in the second paragraph, limit points of normalized occupation measures associated with a sequence of N -player Nash equilibria are shown to be concentrated on solutions of the corresponding mean field game even if the induced limit flow of measures is stochastic (in contrast to Point 3 above). This characterization is established for mean field systems over a finite time horizon as here, but possibly with a common noise (represented as an additional independent Wiener process common to all players). There as here, Nash equilibria are considered in stochastic open-loop strategies, and the methods of proof are similar to ours. The characterization of limit points in Lacker [2015b] relies, even in the situation without common noise studied here, on a new notion of solution of the mean field game (“weak MFG solution”) that applies to probability measures on an extended canonical space (extended with respect to our \mathcal{Z} to keep track of the possibly stochastic flow of measures). In terms of that notion of solution a complete characterization of limit points is achieved. In particular, the assumption in Point 3 that the flow of measures induced by Q is deterministic can be removed. However, if that assumption is dropped, then the claim that Q takes values in the set of solutions of the mean field game with probability one will in general be false. A counterexample illustrating this fact can be deduced from the discussion of subsection 3.3 in Lacker [2015b]. The notion of “weak MFG solution” is indeed strictly weaker than what one

obtains by randomization of the usual notion of solution (“strong” solution with probability one), and this is what makes the complete characterization of Nash limit points possible.

The rest of this work is organized as follows. Notation, basic objects as well as the standing assumptions on the coefficients b , σ , f , F are introduced in Section 2. Section 3 contains a precise description of the N -player games. Nash equilibria are defined and an existence result for certain local approximate Nash equilibria is given; see Proposition 3.1. In Section 4, the limit dynamics for the N -player games are introduced. The corresponding notions of McKean-Vlasov solution and solution of the mean field game are defined and discussed. An approximation result in terms of noise feedback strategies, needed in the construction of competitor strategies, is given in Lemma 4.3. In Section 5, the convergence analysis is carried out, leading to Theorem 5.1 and its Corollary 5.2, which are our main results. Existence of solutions of the mean field game falls out as a by-product of the analysis.

2 Preliminaries and assumptions

Let $d, d_1, d_2 \in \mathbb{N}$, which will be the dimensions of the space of private states, noise values, and control actions, respectively. Choose $T > 0$, the finite time horizon. Set

$$\mathcal{X} \doteq \mathbf{C}([0, T], \mathbb{R}^d), \quad \mathcal{W} \doteq \mathbf{C}([0, T], \mathbb{R}^{d_1}),$$

and, as usual, equip \mathcal{X} , \mathcal{W} with the topology of uniform convergence, which turns them into Polish spaces. Let $\|\cdot\|_{\mathcal{X}}$, $\|\cdot\|_{\mathcal{W}}$ denote the supremum norm on \mathcal{X} and \mathcal{W} , respectively. The spaces \mathbb{R}^n with $n \in \mathbb{N}$ are equipped with the standard Euclidean norm, always indicated by $|\cdot|$.

For \mathcal{S} a Polish space, let $\mathcal{P}(\mathcal{S})$ denote the space of probability measures on $\mathcal{B}(\mathcal{S})$, the Borel sets of \mathcal{S} . For $s \in \mathcal{S}$, let δ_s indicate the Dirac measure concentrated in s . Equip $\mathcal{P}(\mathcal{S})$ with the topology of weak convergence of probability measures. Then $\mathcal{P}(\mathcal{S})$ is again a Polish space. Let $d_{\mathcal{S}}$ be a metric compatible with the topology of \mathcal{S} such that $(\mathcal{S}, d_{\mathcal{S}})$ is a complete and separable metric space. A metric that turns $\mathcal{P}(\mathcal{S})$ into a complete and separable metric space is then given by the bounded Lipschitz metric

$$d_{\mathcal{P}(\mathcal{S})}(\nu, \tilde{\nu}) \doteq \sup \left\{ \int_{\mathcal{S}} g d\nu - \int_{\mathcal{S}} g d\tilde{\nu} : g : \mathcal{S} \rightarrow \mathbb{R} \text{ such that } \|g\|_{\text{bLip}} \leq 1 \right\},$$

where

$$\|g\|_{\text{bLip}} \doteq \sup_{s \in \mathcal{S}} |g(s)| + \sup_{s, \tilde{s} \in \mathcal{S} : s \neq \tilde{s}} \frac{|g(s) - g(\tilde{s})|}{d_{\mathcal{S}}(s, \tilde{s})}.$$

Given a complete compatible metric $d_{\mathcal{S}}$ on \mathcal{S} , we also consider the space of probability measures on $\mathcal{B}(\mathcal{S})$ with finite second moments:

$$\mathcal{P}_2(\mathcal{S}) \doteq \left\{ \nu \in \mathcal{P}(\mathcal{S}) : \exists s_0 \in \mathcal{S} : \int_{\mathcal{S}} d_{\mathcal{S}}(s, s_0)^2 \nu(ds) < \infty \right\}.$$

Notice that $\int d_{\mathcal{S}}(s, s_0)^2 \nu(ds) < \infty$ for some $s_0 \in \mathcal{S}$ implies that the integral is finite for every $s_0 \in \mathcal{S}$. The topology of weak convergence of measures plus convergence of second moments turns $\mathcal{P}_2(\mathcal{S})$ into a Polish space. A compatible complete metric is given by

$$d_{\mathcal{P}_2(\mathcal{S})}(\nu, \tilde{\nu}) \doteq \left(\inf_{\alpha \in \mathcal{P}(\mathcal{S} \times \mathcal{S}) : [\alpha]_1 = \nu \text{ and } [\alpha]_2 = \tilde{\nu}} \int_{\mathcal{S} \times \mathcal{S}} d_{\mathcal{S}}(s, \tilde{s})^2 \alpha(ds, d\tilde{s}) \right)^{1/2},$$

where $[\alpha]_1$ ($[\alpha]_2$) denotes the first (second) marginal of α ; $d_{\mathcal{P}_2(\mathcal{S})}$ is often referred to as the square Wasserstein (or Vasershtein) metric. An immediate consequence of the definition of $d_{\mathcal{P}_2(\mathcal{S})}$ is the following observation: for all $N \in \mathbb{N}$, $s_1, \dots, s_N, \tilde{s}_1, \dots, \tilde{s}_N \in \mathcal{S}$,

$$(2.1) \quad d_{\mathcal{P}_2(\mathcal{S})} \left(\frac{1}{N} \sum_{i=1}^N \delta_{s_i}, \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{s}_i} \right) \leq \sqrt{\frac{1}{N} \sum_{i=1}^N d_{\mathcal{S}}(s_i, \tilde{s}_i)^2}.$$

The bounded Lipschitz metric and the square Wasserstein metric on $\mathcal{P}(\mathcal{S})$ and $\mathcal{P}_2(\mathcal{S})$, respectively, depend on the choice of the metric $d_{\mathcal{S}}$ on the underlying space \mathcal{S} . This dependence will be clear from context. If $\mathcal{S} = \mathbb{R}^d$ with the metric induced by Euclidean norm, we may write d_2 to indicate the square Wasserstein metric $d_{\mathcal{P}_2(\mathbb{R}^d)}$.

Let \mathcal{M} , \mathcal{M}_2 denote the spaces of continuous functions on $[0, T]$ with values in $\mathcal{P}(\mathbb{R}^d)$ and $\mathcal{P}_2(\mathbb{R}^d)$, respectively:

$$\mathcal{M} \doteq \mathbf{C}([0, T], \mathcal{P}(\mathbb{R}^d)), \quad \mathcal{M}_2 \doteq \mathbf{C}([0, T], \mathcal{P}_2(\mathbb{R}^d)).$$

Let Γ be a closed subset of \mathbb{R}^{d_2} , the set of control actions, or action space. Given a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ and a filtration (\mathcal{F}_t) in \mathcal{F} , let $\mathcal{H}_2((\mathcal{F}_t), \mathbf{P}; \Gamma)$ denote the space of all Γ -valued (\mathcal{F}_t) -progressively measurable processes u such that $\mathbf{E} \left[\int_0^T |u(t)|^2 dt \right] < \infty$. The elements of $\mathcal{H}_2((\mathcal{F}_t), \mathbf{P}; \Gamma)$ might be referred to as (individual) strategies.

Denote by \mathcal{R} the space of all deterministic relaxed controls on $\Gamma \times [0, T]$, that is,

$$\mathcal{R} \doteq \{r : r \text{ positive measure on } \mathcal{B}(\Gamma \times [0, T]) : r(\Gamma \times [0, t]) = t \forall t \in [0, T]\}.$$

If $r \in \mathcal{R}$ and $B \in \mathcal{B}(\Gamma)$, then the mapping $[0, T] \ni t \mapsto r(B \times [0, t])$ is absolutely continuous, hence differentiable almost everywhere. Since $\mathcal{B}(\Gamma)$ is

countably generated, the time derivative of r exists almost everywhere and is a measurable mapping $\dot{r}_t: [0, T] \rightarrow \mathcal{P}(\Gamma)$ such that $r(dy, dt) = \dot{r}_t(dy)dt$. Denote by \mathcal{R}_2 the space of deterministic relaxed controls with finite second moments:

$$\mathcal{R}_2 \doteq \left\{ r \in \mathcal{R} : \int_{\Gamma \times [0, T]} |y|^2 r(dy, dt) < \infty \right\}.$$

By definition, $\mathcal{R}_2 \subset \mathcal{R}$. The topology of weak convergence of measures turns \mathcal{R} into a Polish space (not compact unless Γ is bounded). Equip \mathcal{R}_2 with the topology of weak convergence of measures plus convergence of second moments, which makes \mathcal{R}_2 a Polish space, too.

Any Γ -valued process v defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$ induces an \mathcal{R} -valued random variable ρ according to

$$(2.2) \quad \rho_\omega(B \times I) \doteq \int_I \delta_{v(t, \omega)}(B) dt, \quad B \in \mathcal{B}(\Gamma), \quad I \in \mathcal{B}([0, T]), \quad \omega \in \Omega.$$

If v is such that $\int_0^T |v(t)|^2 dt < \infty$ \mathbf{P} -almost surely, then the induced random variable ρ takes values in \mathcal{R}_2 \mathbf{P} -almost surely. If v is progressively measurable with respect to a filtration (\mathcal{F}_t) in \mathcal{F} , then ρ is adapted in the sense that the mapping $t \mapsto \rho(B \times [0, t])$ is (\mathcal{F}_t) -adapted for every $B \in \mathcal{B}(\Gamma)$ [cf. Kushner, 1990, Section 3.3]. More generally, an \mathcal{R} -valued random variable ρ defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$ is called adapted to a filtration (\mathcal{F}_t) in \mathcal{F} if the process $t \mapsto \rho(B \times [0, t])$ is (\mathcal{F}_t) -adapted for every $B \in \mathcal{B}(\Gamma)$.

Below, we will make use of the following canonical space. Set

$$\mathcal{Z} \doteq \mathcal{X} \times \mathcal{R}_2 \times \mathcal{W},$$

and endow \mathcal{Z} with the product topology, which makes it a Polish space. Let $d_{\mathcal{R}_2}$ be a complete metric compatible with the topology of \mathcal{R}_2 . Set

$$d_{\mathcal{Z}}((\varphi, r, w), (\tilde{\varphi}, \tilde{r}, \tilde{w})) \doteq \|\varphi - \tilde{\varphi}\|_{\mathcal{X}} + \frac{d_{\mathcal{R}_2}(r, \tilde{r})}{1 + d_{\mathcal{R}_2}(r, \tilde{r})} + \frac{\|w - \tilde{w}\|_{\mathcal{W}}}{1 + \|w - \tilde{w}\|_{\mathcal{W}}},$$

where $(\varphi, r, w), (\tilde{\varphi}, \tilde{r}, \tilde{w})$ are elements of \mathcal{Z} written component-wise. This defines a complete metric compatible with the topology of \mathcal{Z} . Let $d_{\mathcal{P}_2(\mathcal{Z})}$ be the square Wasserstein metric on $\mathcal{P}_2(\mathcal{Z})$ induced by $d_{\mathcal{Z}}$. Since $d_{\mathcal{Z}}$ is bounded with respect to the second and third component of \mathcal{Z} , the condition of finite second moment is a restriction only on the first marginal of the probability measures on $\mathcal{B}(\mathcal{Z})$. Let us indicate by $d_{\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))}$ the bounded Lipschitz metric on $\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))$ induced by $d_{\mathcal{P}_2(\mathcal{Z})}$. Denote by $(\hat{X}, \hat{\rho}, \hat{W})$ the coordinate process on \mathcal{Z} :

$$\hat{X}(t, (\varphi, r, w)) \doteq \varphi(t), \quad \hat{\rho}(t, (\varphi, r, w)) \doteq r|_{\mathcal{B}(\Gamma \times [0, t])}, \quad \hat{W}(t, (\varphi, r, w)) \doteq w(t).$$

Let (\mathcal{G}_t) be the canonical filtration in $\mathcal{B}(\mathcal{Z})$, that is,

$$\mathcal{G}_t \doteq \sigma\left((\hat{X}, \hat{\rho}, \hat{W})(s) : 0 \leq s \leq t\right), \quad t \in [0, T].$$

Let b denote the drift coefficient and σ the dispersion coefficient of the dynamics, and let f, F quantify the running costs and terminal costs, respectively; we take

$$\begin{aligned} b: [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \Gamma &\rightarrow \mathbb{R}^d, & \sigma: [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) &\rightarrow \mathbb{R}^{d \times d_1}, \\ f: [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \times \Gamma &\rightarrow [0, \infty), & F: \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) &\rightarrow [0, \infty). \end{aligned}$$

Notice that the dispersion coefficient σ does not depend on the control variable and that the cost coefficients f, F are non-negative functions. We make the following assumptions, where K, L are some finite positive constants:

(A1) Measurability and continuity in time and control: b, σ, f, F are Borel measurable and such that, for all $(x, \nu) \in \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, $b(\cdot, x, \nu, \cdot)$, $\sigma(\cdot, x, \nu)$, $f(\cdot, x, \nu, \cdot)$ are continuous, uniformly over compact subsets of $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$.

(A2) Lipschitz continuity of b, σ : for all $x, \tilde{x} \in \mathbb{R}^d$, $\nu, \tilde{\nu} \in \mathcal{P}_2(\mathbb{R}^d)$,

$$\begin{aligned} \sup_{t \in [0, T]} \sup_{\gamma \in \Gamma} \{ &|b(t, x, \nu, \gamma) - b(t, \tilde{x}, \tilde{\nu}, \gamma)| \vee |\sigma(t, x, \nu) - \sigma(t, \tilde{x}, \tilde{\nu})| \} \\ &\leq L(|x - \tilde{x}| + d_2(\nu, \tilde{\nu})). \end{aligned}$$

(A3) Sublinear growth of b, σ : for all $x \in \mathbb{R}^d$, $\nu \in \mathcal{P}_2(\mathbb{R}^d)$, $\gamma \in \Gamma$,

$$\begin{aligned} \sup_{t \in [0, T]} |b(t, x, \nu, \gamma)| &\leq K \left(1 + |x| + |\gamma| + \sqrt{\int |y|^2 \nu(dy)} \right), \\ \sup_{t \in [0, T]} |\sigma(t, x, \nu)| &\leq K \left(1 + |x| + \sqrt{\int |y|^2 \nu(dy)} \right). \end{aligned}$$

(A4) Local Lipschitz continuity of f, F : for all $x, \tilde{x} \in \mathbb{R}^d$, $\nu, \tilde{\nu} \in \mathcal{P}_2(\mathbb{R}^d)$

$$\begin{aligned} \sup_{t \in [0, T]} \sup_{\gamma \in \Gamma} \{ &|f(t, x, \nu, \gamma) - f(t, \tilde{x}, \tilde{\nu}, \gamma)| + |F(x, \nu) - F(\tilde{x}, \tilde{\nu})| \} \\ &\leq L(|x - \tilde{x}| + d_2(\nu, \tilde{\nu})) \left(1 + |x| + |\tilde{x}| + \sqrt{\int |y|^2 \nu(dy)} + \sqrt{\int |y|^2 \tilde{\nu}(dy)} \right). \end{aligned}$$

(A5) Subquadratic growth of f, F : for all $x \in \mathbb{R}^d$, $\nu \in \mathcal{P}_2(\mathbb{R}^d)$, $\gamma \in \Gamma$,

$$\sup_{t \in [0, T]} \{ |f(t, x, \nu, \gamma)| \vee |F(x, \nu)| \} \leq K \left(1 + |x|^2 + |\gamma|^2 + \int |y|^2 \nu(dy) \right).$$

(A6) Action space and coercivity: $\Gamma \subset \mathbb{R}^{d_2}$ is closed, and there exist $c_0 > 0$ and $\Gamma_0 \subset \Gamma$ such that Γ_0 is compact and for every $\gamma \in \Gamma \setminus \Gamma_0$

$$\inf_{(t,x,\nu) \in [0,T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)} f(t, x, \nu, \gamma) \geq c_0 |\gamma|^2.$$

3 N -player games

Let $N \in \mathbb{N}$. Let $(\Omega_N, \mathcal{F}^N, \mathbf{P}_N)$ be a complete probability space equipped with a filtration (\mathcal{F}_t^N) in \mathcal{F}^N that satisfies the usual hypotheses and carrying N independent d_1 -dimensional (\mathcal{F}_t^N) -Wiener processes W_1^N, \dots, W_N^N . For each $i \in \{1, \dots, N\}$, choose a random variable $\xi_i^N \in L^2(\Omega_N, \mathcal{F}_0^N, \mathbf{P}_N; \mathbb{R}^d)$, the initial state of player i in the prelimit game with N players. In addition, we assume that the stochastic basis is rich enough to carry a sequence $(\vartheta_i^N)_{i \in \{1, \dots, N\}}$ of independent random variables with values in the interval $[0, 1]$ such that each ϑ_i^N is \mathcal{F}_0^N -measurable and uniformly distributed on $[0, 1]$, and $(\vartheta_i^N)_{i \in \{1, \dots, N\}}$ is independent of the σ -algebra generated by ξ_1^N, \dots, ξ_N^N and the Wiener processes W_1^N, \dots, W_N^N . The random variables $\vartheta_i^N, i \in \{1, \dots, N\}$, are a technical device which we may use without loss of generality; see Remark 3.3 below.

A vector of individual strategies, that is, a vector $\mathbf{u} = (u_1, \dots, u_N)$ such that $u_1, \dots, u_N \in \mathcal{H}_2((\mathcal{F}_t^N), \mathbf{P}_N; \Gamma)$, is called a strategy vector. Given a strategy vector $\mathbf{u} = (u_1, \dots, u_N)$, consider the system of Itô stochastic integral equations

$$(3.1) \quad \begin{aligned} X_i^N(t) &= \xi_i^N + \int_0^t b(s, X_i^N(s), \mu^N(s), u_i(s)) ds \\ &+ \int_0^t \sigma(s, X_i^N(s), \mu^N(s)) dW_i^N(s), \quad t \in [0, T], \end{aligned}$$

$i \in \{1, \dots, N\}$, where $\mu^N(s)$ is the empirical measure of the processes X_1^N, \dots, X_N^N at time $s \in [0, T]$, that is,

$$\mu_\omega^N(s) \doteq \frac{1}{N} \sum_{j=1}^N \delta_{X_j^N(s, \omega)}, \quad \omega \in \Omega_N.$$

The process X_i^N describes the evolution of the private state of player i if he/she uses strategy u_i while the other players use strategies $u_j, j \neq i$. Thanks to assumptions (A2) and (A3), the system of equations (3.1) possesses a unique solution in the following sense: given any strategy vector $\mathbf{u} = (u_1, \dots, u_N)$, there exists a vector (X_1^N, \dots, X_N^N) of continuous \mathbb{R}^d -valued (\mathcal{F}_t^N) -adapted processes such that (3.1) holds \mathbf{P}_N -almost surely, and

(X_1^N, \dots, X_N^N) is unique (up to \mathbf{P}_N -indistinguishability) among all continuous (\mathcal{F}_t^N) -adapted solutions.

The following estimates on the controlled state process and the associated empirical measure process will be useful in Section 5.

Lemma 3.1. *There exists a finite constant $C_{T,K}$ depending on T, K , but not on N , such that if $\mathbf{u}^N = (u_1^N, \dots, u_N^N)$ is a strategy vector for the N -player game and (X_1^N, \dots, X_N^N) the solution of the system (3.1) under \mathbf{u}^N , then*

$$\begin{aligned} & \sup_{t \in [0, T]} \mathbf{E}_N [|X_i^N(t)|^2] \\ & \leq C_{T,K} \left(1 + \mathbf{E}_N [|\xi_i^N|^2] + \mathbf{E}_N \left[\int_0^T \left(d_2(\mu^N(t), \delta_0)^2 + |u_i^N(t)|^2 \right) dt \right] \right) \end{aligned}$$

for every $i \in \{1, \dots, N\}$, and

$$\begin{aligned} \sup_{t \in [0, T]} \mathbf{E}_N [d_2(\mu^N(t), \delta_0)^2] & \leq \sup_{t \in [0, T]} \mathbf{E}_N \left[\frac{1}{N} \sum_{j=1}^N |X_j^N(t)|^2 \right] \\ & \leq C_{T,K} \left(1 + \frac{1}{N} \sum_{j=1}^N \mathbf{E}_N \left[|\xi_j^N|^2 + \int_0^T |u_j^N(t)|^2 dt \right] \right). \end{aligned}$$

Proof. By Jensen's inequality, Hölder's inequality, Itô's isometry, assumption (A3), and the Fubini-Tonelli theorem, we have for every $t \in [0, T]$,

$$\begin{aligned} \mathbf{E}_N [|X_i^N(t)|^2] & \leq 3 \mathbf{E}_N [|\xi_i^N|^2] + 12(T+1)K^2 \int_0^t \mathbf{E}_N [|X_i^N(s)|^2] ds \\ & \quad + 12(T+1)K^2 \mathbf{E}_N \left[\int_0^T \left(1 + d_2(\mu^N(s), \delta_0)^2 + |u_i^N(s)|^2 \right) ds \right], \end{aligned}$$

and the first estimate follows by Gronwall's lemma.

By definition of the square Wasserstein metric d_2 , we have for every $t \in [0, T]$, every $\omega \in \Omega_N$,

$$d_2(\mu_\omega^N(t), \delta_0)^2 = \frac{1}{N} \sum_{j=1}^N |X_j^N(t, \omega)|^2.$$

Thus, using again assumption (A3) and the same inequalities as above, we

have for every $t \in [0, T]$,

$$\begin{aligned} & \mathbf{E}_N \left[\frac{1}{N} \sum_{j=1}^N |X_j^N(t)|^2 \right] \\ & \leq 3 \mathbf{E}_N \left[\frac{1}{N} \sum_{j=1}^N |\xi_j^N|^2 \right] + 12(T+1)K^2 \int_0^T \mathbf{E}_N \left[1 + \frac{1}{N} \sum_{j=1}^N |u_j^N(s)|^2 \right] ds \\ & \quad + 24(T+1)K^2 \int_0^t \mathbf{E}_N \left[\frac{1}{N} \sum_{j=1}^N |X_j^N(s)|^2 \right] ds, \end{aligned}$$

and we conclude again by Gronwall's lemma. The constant $C_{T,K}$ for both estimates need not be greater than $12(T \vee 1)(T+1)(K \vee 1)^2 \exp(24(T+1)K^2T)$. \square

Lemma 3.2. *Let $p \geq 2$. Then there exists a finite constant $\tilde{C}_{p,T,K,d}$ depending on p, T, K, d , but not on N such that if $\mathbf{u}^N = (u_1^N, \dots, u_N^N)$ is a strategy vector for the N -player game and (X_1^N, \dots, X_N^N) the solution of the system (3.1) under \mathbf{u}^N , then*

$$\begin{aligned} \mathbf{E}_N \left[\sup_{t \in [0, T]} d_2(\mu^N(t), \delta_0)^p \right] & \leq \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N [\|X_i^N\|_{\mathcal{X}}^p] \\ & \leq \tilde{C}_{p,T,K,d} \left(1 + \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[|\xi_i^N|^p + \int_0^T |u_i^N(t)|^p dt \right] \right). \end{aligned}$$

Proof. The inequality

$$\mathbf{E}_N \left[\sup_{t \in [0, T]} d_2(\mu^N(t), \delta_0)^p \right] \leq \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N [\|X_i^N\|_{\mathcal{X}}^p]$$

follows by (2.1) and Jensen's inequality. In verifying the second part of the assertion, we may assume that

$$\frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[|\xi_i^N|^p + \int_0^T |u_i^N(t)|^p dt \right] < \infty.$$

By Jensen's inequality, Hölder's inequality, (A3), the Fubini-Tonelli theorem,

and the Burkholder-Davis-Gundy inequalities, we have for every $t \in [0, T]$,

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{s \in [0, t]} |X_i^N(s)|^p \right] \\ & \leq \hat{C}_{p, T, K, d} \left(1 + \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[|\xi_i^N|^p + \int_0^T |u_i^N(s)|^p ds \right] \right) \\ & \quad + 2\hat{C}_{p, T, K, d} \int_0^t \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{s \in [0, \tilde{s}]} |X_i^N(s)|^p \right] d\tilde{s}, \end{aligned}$$

where $\hat{C}_{p, T, K, d} \doteq 12^{p-1}(T \vee 1)^p(K \vee 1)^p(1 + \hat{C}_{p, d})$ and $\hat{C}_{p, d}$, which depends only on p and d , is the finite ‘‘universal’’ constant from the Burkholder-Davis-Gundy inequalities [for instance, Theorem 3.3.28 and Remark 3.3.30 in Karatzas and Shreve, 1991, pp. 166-167]. The assertion now follows thanks to Gronwall’s lemma. \square

Player i evaluates a strategy vector $\mathbf{u} = (u_1, \dots, u_N)$ according to the cost functional

$$J_i^N(\mathbf{u}) \doteq \mathbf{E}_N \left[\int_0^T f(s, X_i^N(s), \mu^N(s), u_i(s)) ds + F(X_i^N(T), \mu^N(T)) \right],$$

where (X_1^N, \dots, X_N^N) is the solution of the system (3.1) under \mathbf{u} and μ^N is the empirical measure process induced by (X_1^N, \dots, X_N^N) .

Given a strategy vector $\mathbf{u} = (u_1, \dots, u_N)$ and an individual strategy $v \in \mathcal{H}_2((\mathcal{F}_t^N), \mathbf{P}_N; \Gamma)$, let $[\mathbf{u}^{-i}, v] \doteq (u_1, \dots, u_{i-1}, v, u_{i+1}, \dots, u_N)$ indicate the strategy vector that is obtained from \mathbf{u} by replacing u_i , the strategy of player i , with v . Let $(\mathcal{F}_t^{N, i})$ denote the filtration generated by ϑ_i^N , ξ_i^N , and the Wiener process W_i^N , that is,

$$\mathcal{F}_t^{N, i} \doteq \sigma(\vartheta_i^N, \xi_i^N, W_i^N(s) : s \in [0, t]), \quad t \in [0, T].$$

The filtration $(\mathcal{F}_t^{N, i})$ represents the local information available to player i . Clearly, $(\mathcal{F}_t^{N, i}) \subset \mathcal{F}_t^N$ and $\mathcal{H}_2((\mathcal{F}_t^{N, i}), \mathbf{P}_N; \Gamma) \subset \mathcal{H}_2((\mathcal{F}_t^N), \mathbf{P}_N; \Gamma)$. We may refer to the elements of $\mathcal{H}_2((\mathcal{F}_t^{N, i}), \mathbf{P}_N; \Gamma)$ as decentralized strategies for player i .

Definition 3.1. Let $\varepsilon \geq 0$, $u_1, \dots, u_N \in \mathcal{H}_2((\mathcal{F}_t^N), \mathbf{P}_N; \Gamma)$. The strategy vector $\mathbf{u} \doteq (u_1, \dots, u_N)$ is called a *local ε -Nash equilibrium* for the N -player game if for every $i \in \{1, \dots, N\}$, every $v \in \mathcal{H}_2((\mathcal{F}_t^{N, i}), \mathbf{P}_N; \Gamma)$,

$$(3.2) \quad J_i^N(\mathbf{u}) \leq J_i^N([\mathbf{u}^{-i}, v]) + \varepsilon.$$

If inequality (3.2) holds for all $v \in \mathcal{H}_2((\mathcal{F}_t^N), \mathbf{P}_N; \Gamma)$, then \mathbf{u} is called an ε -Nash equilibrium.

If \mathbf{u} is a (local) ε -Nash equilibrium with $\varepsilon = 0$, then \mathbf{u} is called a (local) Nash equilibrium.

Remark 3.1. The attribute “local” in the expression “local Nash equilibrium” or “local ε -Nash equilibrium” refers to the information that is available to the deviating player for choosing competitor strategies. Based on local information, those strategies have to be decentralized. In this sense, decentralized strategies are “local,” too. Notice that the decentralized strategies here are not necessarily representable as functionals of time and the corresponding individual state process alone.

Remark 3.2. In Definition 3.1, Nash equilibria are defined with respect to stochastic open-loop strategies. This is the same notion as the one used in the probabilistic approach to mean field games; see Carmona and Delarue [2013]. A Nash equilibrium in stochastic open-loop strategies may be induced by a Markov feedback strategy (or a more general closed-loop strategy); still, it need not correspond to a Nash equilibrium in feedback strategies. Given a vector of feedback strategies, varying the strategy of exactly one player means that the feedback functions defining the strategies of the other players are kept frozen. Since in general the state processes of the other players depend on the state process of the deviating player (namely, through the empirical measure of the system), the strategies of the other players seen as control processes may change when one player deviates. This is in contrast with the stochastic open-loop formulation where the control processes of the other players are frozen when one player varies her/his strategy. Now, suppose we had a Nash equilibrium in Markov feedback strategies for the N -player game. If the feedback functions defining that Nash equilibrium depend only on time, the current individual state, and the current empirical measure, and if they are regular in the sense of being Lipschitz continuous, then they will induce an ε_N -Nash equilibrium in stochastic open-loop strategies with ε_N also depending on the Lipschitz constants of the feedback functions. Here, we do not address the question of when Nash equilibria in regular feedback strategies exist nor of how their Lipschitz constants would depend on the number of players N . Neither do we address the more general question of convergence of N -player Nash equilibria in feedback strategies, regular or not. That difficult problem was posed in Lasry and Lions [2006b, 2007] and is beyond the scope of the present work. It was solved very recently by Cardaliaguet et al. [2015] in the situation where the noise is additive (allowing for an extra common noise), the cost structure satisfies the Lasry-

Lions monotonicity conditions, and the N -player game possesses, for each N , a unique Nash equilibrium in feedback strategies. The authors introduce an infinite-dimensional partial differential equation (the “master equation”) that characterizes solutions of the mean field game and allows to capture the dependence on the measure variable. They establish existence of a unique regular solution to the master equation. That solution is then used in proving convergence “on average” of the (symmetric) equilibrium strategies to the mean field game limit.

Remark 3.3. The random variables ϑ_i^N appearing in the definition of the local information filtrations $(\mathcal{F}_t^{N,i})$ are a technical device for randomization. They will be used in the sequel only in two places, namely in the proof of Proposition 3.1 on existence of local ε -Nash equilibria, where they allow to pass from optimal relaxed controls to nearly optimal ordinary controls, and in the proof of Lemma 5.2, where they serve to generate a coupling of initial conditions. The presence of the random variables ϑ_i^N causes no loss of generality in the following sense. Suppose that $\mathbf{u} \doteq (u_1, \dots, u_N)$ is a strategy vector adapted to the filtration generated by ξ_1^N, \dots, ξ_N^N and the Wiener processes W_1^N, \dots, W_N^N such that, for some $\varepsilon \geq 0$, every $i \in \{1, \dots, N\}$, inequality (3.2) holds for all individual strategies v that are adapted to the filtration generated by ξ_i^N and the Wiener process W_i^N . Then inequality (3.2) holds for all $v \in \mathcal{H}_2((\mathcal{F}_t^{N,i}), \mathbf{P}_N; \Gamma)$; hence \mathbf{u} is a local ε -Nash equilibrium. To check this, take conditional expectation with respect to ϑ_i^N inside the expectation defining the cost functional J_i^N and use the independence of ϑ_i^N from the σ -algebra generated by ξ_1^N, \dots, ξ_N^N and W_1^N, \dots, W_N^N . An analogous reasoning applies to the situation of non-local (approximate) Nash equilibria provided the strategy vector \mathbf{u} is independent of the family $(\vartheta_i^N)_{i \in \{1, \dots, N\}}$.

By Definition 3.1, an ε -Nash equilibrium is also a local ε -Nash equilibrium. Observe that the individual strategies of a local ε -Nash equilibrium are adapted to the full filtration (\mathcal{F}_t^N) ; only the competitor strategies in the verification of the local equilibrium property have to be decentralized (or “local”), that is, strategies adapted to one of the smaller filtrations $(\mathcal{F}_t^{N,1}), \dots, (\mathcal{F}_t^{N,N})$.

If ξ_1^N, \dots, ξ_N^N are independent and $\mathbf{u} = (u_1, \dots, u_N)$ is a vector of decentralized strategies, that is, $u_i \in \mathcal{H}_2((\mathcal{F}_t^{N,i}), \mathbf{P}_N; \Gamma)$ for every $i \in \{1, \dots, N\}$, then $(\xi_1^N, u_1, W_1^N), \dots, (\xi_N^N, u_N, W_N^N)$, interpreted as $\mathbb{R}^d \times \mathcal{R}_2 \times \mathcal{W}$ -valued random variables, are independent. This allows to deduce existence of local approximate Nash equilibria through Fan’s fixed point theorem in a way similar to that for one-shot games [cf. Appendix 8.1 in Cardaliaguet, 2013]. For

simplicity, we give the result for a compact action space, bounded coefficients and in the fully symmetric situation. In the sequel, Proposition 3.1 will be used only to provide an example of a situation in which all the hypotheses of our main result can be easily verified.

Proposition 3.1. *In addition to (A1) – (A6), assume that Γ is compact and that b, σ, f, F are bounded. Suppose that ξ_1^N, \dots, ξ_N^N are independent and identically distributed. Given any $\varepsilon > 0$, there exist decentralized strategies $u_i^\varepsilon \in \mathcal{H}_2((\mathcal{F}_t^{N,i}), \mathbf{P}_N; \Gamma)$, $i \in \{1, \dots, N\}$, such that $\mathbf{u}^\varepsilon \doteq (u_1^\varepsilon, \dots, u_N^\varepsilon)$ is a local ε -Nash equilibrium for the N -player game and the random variables $(\xi_1^N, u_1^\varepsilon, W_1^N), \dots, (\xi_N^N, u_N^\varepsilon, W_N^N)$ are independent and identically distributed.*

Proof. Since Γ is compact by hypothesis, we have $\mathcal{R} = \mathcal{R}_2$ as topological spaces, and $\mathcal{P}(\mathcal{R})$ is compact.

Let \mathbf{m}_0 denote the common distribution of the initial states ξ_1^N, \dots, ξ_N^N ; thus $\mathbf{m}_0 \in \mathcal{P}_2(\mathbb{R}^d)$. With a slight abuse of notation, let $(\hat{X}(0), \rho, \hat{W})$ denote the restriction to $\mathbb{R}^d \times \mathcal{R} \times \mathcal{W}$ of the canonical process on \mathcal{Z} . Let $(\tilde{\mathcal{G}}_t)$ indicate the corresponding canonical filtration, that is, $\tilde{\mathcal{G}}_t \doteq \sigma(\hat{X}(0), \rho(s), \hat{W}(s) : s \leq t)$, $t \in [0, T]$. Let \mathcal{Y} be the space of all $\nu \in \mathcal{P}(\mathbb{R}^d \times \mathcal{R} \times \mathcal{W})$ such that $[\nu]_1 = \mathbf{m}_0$ and \hat{W} is a $(\tilde{\mathcal{G}}_t)$ -Wiener process under ν (in particular, $\hat{W}(0) = 0$ ν -almost surely). Then \mathcal{Y} is a non-empty compact convex subset of $\mathcal{P}(\mathbb{R}^d \times \mathcal{R} \times \mathcal{W})$, which in turn is contained in a locally convex topological linear space (under the topology of weak convergence of measures).

The proof proceeds in two steps. First, we show that there exists $\nu_* \in \mathcal{Y}$ such that $\otimes^N \nu_*$ corresponds to a local Nash equilibrium in relaxed controls on the canonical space \mathcal{Z}^N . In the second step, given any $\varepsilon > 0$, we use ν_* to construct a local ε -Nash equilibrium for the N -player game.

First step. Let $\nu, \bar{\nu} \in \mathcal{Y}$. Then there exists a unique $\Psi(\nu; \bar{\nu}) \in \mathcal{P}_2(\mathcal{Z}^N)$ such that

$$\Psi(\nu; \bar{\nu}) = \mathbf{P} \circ (\mathbf{X}, \boldsymbol{\rho}, \mathbf{W})^{-1},$$

where $\mathbf{W} = (W_1, \dots, W_N)$ is a vector of independent d_1 -dimensional (\mathcal{F}_t) -adapted Wiener processes defined on some stochastic basis $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t))$ satisfying the usual hypotheses and carrying a vector $\boldsymbol{\rho} = (\rho^1, \dots, \rho^N)$ of (\mathcal{F}_t) -adapted \mathcal{R} -valued random variables such that

$$\mathbf{P} \circ (\mathbf{X}(0), \boldsymbol{\rho}, \mathbf{W})^{-1} = \nu \otimes^{N-1} \bar{\nu},$$

and $\mathbf{X} = (X_1, \dots, X_N)$ is the vector of continuous \mathbb{R}^d -valued (\mathcal{F}_t) -adapted

processes determined through the system of equations

$$(3.3) \quad \begin{aligned} X_i(t) = & X_i(0) + \int_{\Gamma \times [0,t]} b \left(s, X_i(s), \frac{1}{N} \sum_{j=1}^N \delta_{X_j^N(s)}, \gamma \right) \rho^i(d\gamma, ds) \\ & + \int_0^t \sigma \left(s, X_i(s), \frac{1}{N} \sum_{j=1}^N \delta_{X_j^N(s)} \right) dW_i(s), \quad t \in [0, T], \end{aligned}$$

$i \in \{1, \dots, N\}$, which is the relaxed version of (3.1). The mapping

$$(\nu, \bar{\nu}) \mapsto \Psi(\nu; \bar{\nu})$$

defines a continuous function $\mathcal{Y} \times \mathcal{Y} \rightarrow \mathcal{P}_2(\mathcal{Z}^N)$. The continuity of Ψ can be checked by using a martingale problem characterization of solutions to (3.3); cf. El Karoui et al. [1987], Kushner [1990], and also Section 4 below. Define a function $J: \mathcal{Y} \times \mathcal{Y} \rightarrow [0, \infty)$ by

$$\begin{aligned} & J(\nu; \bar{\nu}) \\ & \doteq \mathbf{E}_{\Psi(\nu; \bar{\nu})} \left[\int_{\Gamma \times [0, T]} f \left(s, \hat{X}_1(s), \hat{\mu}(s), \gamma \right) d\hat{\rho}^1(d\gamma, ds) + F \left(\hat{X}_1(T), \hat{\mu}(T) \right) \right], \end{aligned}$$

where $\hat{\mu}(s) \doteq \frac{1}{N} \sum_{j=1}^N \delta_{\hat{X}_j(s)}$ and $(\hat{X}_1, \dots, \hat{X}_N), (\hat{\rho}^1, \dots, \hat{\rho}^N)$ are components of the canonical process on \mathcal{Z}^N with the obvious interpretation. Thanks to the continuity of Ψ and the boundedness and continuity of f, F , we have that J is a continuous mapping on $\mathcal{Y} \times \mathcal{Y}$. On the other hand, for any fixed $\bar{\nu} \in \mathcal{Y}$, all $\nu, \tilde{\nu} \in \mathcal{Y}$, all $\lambda \in [0, 1]$,

$$\begin{aligned} \Psi(\lambda\nu + (1-\lambda)\tilde{\nu}; \bar{\nu}) &= \lambda\Psi(\nu; \bar{\nu}) + (1-\lambda)\Psi(\tilde{\nu}; \bar{\nu}), \\ J(\lambda\nu + (1-\lambda)\tilde{\nu}; \bar{\nu}) &= \lambda J(\nu; \bar{\nu}) + (1-\lambda)J(\tilde{\nu}; \bar{\nu}). \end{aligned}$$

Define a function $\chi: \mathcal{Y} \rightarrow \mathcal{B}(\mathcal{Y})$ by

$$\chi(\bar{\nu}) \doteq \left\{ \nu \in \mathcal{Y} : J(\nu; \bar{\nu}) = \min_{\tilde{\nu} \in \mathcal{Y}} J(\tilde{\nu}; \bar{\nu}) \right\}.$$

Observe that $\chi(\bar{\nu})$ is non-empty, compact and convex for every $\bar{\nu} \in \mathcal{Y}$. Thus, χ is well-defined as a mapping from \mathcal{Y} to $\mathcal{K}(\mathcal{Y})$, the set of all non-empty compact convex subsets of \mathcal{Y} . Moreover, χ is upper semicontinuous in the sense that $\nu \in \chi(\bar{\nu})$ whenever $(\nu_n) \subset \mathcal{Y}$, $(\bar{\nu}_n) \subset \mathcal{Y}$ are sequences such that $\lim_{n \rightarrow \infty} \bar{\nu}_n = \bar{\nu}$, $\lim_{n \rightarrow \infty} \nu_n = \nu$, and $\nu_n \in \chi(\bar{\nu}_n)$ for each $n \in \mathbb{N}$ (recall that \mathcal{Y} is metrizable). We are therefore in the situation of Theorem 1 in Fan [1952], which guarantees the existence of a fixed point for χ , that is, there exists $\nu_* \in \mathcal{Y}$ such that $\nu_* \in \chi(\nu_*)$.

Second step. Let $\varepsilon > 0$, and let $\nu_* \in \mathcal{Y}$ be such that $\nu_* \in \chi(\nu_*)$. Let $d_{\mathcal{Y}}$ be a compatible metric on the compact Polish space \mathcal{Y} , and define a corresponding metric on $\mathcal{Y} \times \mathcal{Y}$ by $d_{\mathcal{Y} \times \mathcal{Y}}((\nu, \bar{\nu}), (\mu, \bar{\mu})) \doteq d_{\mathcal{Y}}(\nu, \mu) + d_{\mathcal{Y}}(\bar{\nu}, \bar{\mu})$. Choose a stochastic basis $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t))$ satisfying the usual hypotheses and carrying a vector $\mathbf{W} = (W_1, \dots, W_N)$ of independent d_1 -dimensional (\mathcal{F}_t) -adapted Wiener processes, a vector $\boldsymbol{\rho} = (\rho^1, \dots, \rho^N)$ of (\mathcal{F}_t) -adapted \mathcal{R} -valued random variables as well as a vector $\boldsymbol{\xi} = (\xi_1, \dots, \xi_N)$ of \mathbb{R}^d -valued \mathcal{F}_0 -measurable random variables such that

$$\mathbf{P} \circ (\boldsymbol{\xi}, \boldsymbol{\rho}, \mathbf{W})^{-1} = \otimes^N \nu_*.$$

For $i \in \{1, \dots, N\}$, let $(\mathcal{F}_t^{\circ, i})$ be the filtration generated by ξ_i, ρ^i, W_i , that is, $\mathcal{F}_t^{\circ, i} \doteq \sigma(\xi_i, \rho^i(s), W_i(s) : s \leq t), t \in [0, T]$. By independence and a version of the chattering lemma [for instance, Theorem 3.5.2 in Kushner, 1990, p. 59], for every $\delta > 0$, there exists a vector $\boldsymbol{\rho}^\delta = (\rho^{\delta, 1}, \dots, \rho^{\delta, N})$ of \mathcal{R} -valued random variables such that:

- (i) for every $i \in \{1, \dots, N\}$, $\rho^{\delta, i}$ is the relaxed control induced by a piecewise constant $(\mathcal{F}_t^{\circ, i})$ -progressively measurable Γ -valued process;
- (ii) the random variables $(\xi_1, \rho^{\delta, 1}, W_1), \dots, (\xi_N, \rho^{\delta, N}, W_N)$ are independent and identically distributed;
- (iii) setting $\nu_\delta \doteq \mathbf{P} \circ (\xi_1, \rho^{\delta, 1}, W_1)^{-1}$, we have $d_{\mathcal{Y}}(\nu_\delta, \nu_*) \leq \delta$.

Since J is continuous on the compact space $\mathcal{Y} \times \mathcal{Y}$, it is uniformly continuous. We can therefore find $\delta = \delta(\varepsilon) > 0$ such that

$$(3.4) \quad |J(\nu_\delta; \nu_\delta) - J(\nu_*; \nu_*)| + \max_{\nu \in \mathcal{Y}} |J(\nu; \nu_\delta) - J(\nu; \nu_*)| \leq \varepsilon.$$

The law ν_δ (with $\delta = \delta(\varepsilon)$) and the corresponding product measure can be reproduced on the stochastic basis of the N -player game. More precisely, there exists a measurable function $\psi: [0, T] \times [0, 1] \times \mathbb{R}^d \times \mathcal{W} \rightarrow \Gamma$ such that, upon setting

$$u_i(t, \omega) \doteq \psi(t, \vartheta_i^N(\omega), \xi_i^N(\omega), W_i^N(\cdot, \omega)), \quad (t, \omega) \in [0, T] \times \Omega_N,$$

the following hold:

- (i) $u_i \in \mathcal{H}_2((\mathcal{F}_t^{N, i}), \mathbf{P}_N; \Gamma)$ for every $i \in \{1, \dots, N\}$;
- (ii) $(\xi_1^N, u_1, W_1^N), \dots, (\xi_N^N, u_N, W_N)$, interpreted as $\mathbb{R}^d \times \mathcal{R} \times \mathcal{W}$ -valued random variables, are independent and identically distributed;
- (iii) $\mathbf{P}_N \circ (\xi_1^N, u_1, W_1^N)^{-1} = \nu_\delta$.

The relaxed controls $\rho^{\delta,1}, \dots, \rho^{\delta,1}$ are, in fact, induced by Γ -valued processes that may be taken to be piece-wise constant in time with respect to a common equidistant grid in $[0, T]$. Existence of a function ψ with the desired properties can therefore be established by iteration along the grid points, repeatedly invoking Theorem 6.10 in Kallenberg [2001, p. 112] on measurable transfers; this procedure also yields progressive measurability of ψ .

Set $\mathbf{u} \doteq (u_1, \dots, u_N)$ with $u_i \in \mathcal{H}_2((\mathcal{F}_t^{N,i}), \mathbf{P}_N; \Gamma)$ as above. Then

$$J_1^N(\mathbf{u}) = J(\nu_\delta; \nu_\delta).$$

Let $v \in \mathcal{H}_2((\mathcal{F}_t^{N,1}), \mathbf{P}_N; \Gamma)$, and set $\nu \doteq \mathbf{P}_N \circ (\xi_1^N, v, W_1^N)^{-1}$, where v is identified with its relaxed control. By independence and construction,

$$J_1^N([\mathbf{u}^{-1}, v]) = J(\nu; \nu_\delta).$$

Now, thanks to (3.4) and the equilibrium property of ν_* ,

$$\begin{aligned} & J(\nu; \nu_\delta) - J(\nu_\delta; \nu_\delta) \\ &= J(\nu; \nu_\delta) - J(\nu; \nu_*) + J(\nu_*; \nu_*) - J(\nu_\delta; \nu_\delta) + J(\nu; \nu_*) - J(\nu_*; \nu_*) \\ &\geq -\varepsilon. \end{aligned}$$

It follows that

$$J_1^N(\mathbf{u}) \leq J_1^N([\mathbf{u}^{-1}, v]) + \varepsilon \quad \text{for all } v \in \mathcal{H}_2((\mathcal{F}_t^{N,1}), \mathbf{P}_N; \Gamma).$$

This establishes the local approximate equilibrium property of the strategy vector \mathbf{u} with respect to deviations in decentralized strategies of player one. By symmetry, the property also holds with respect to deviations of the other players. We conclude that \mathbf{u} is a local ε -Nash equilibrium. \square

4 Mean field games

In order to describe the limit system for the N -player games introduced above, consider the stochastic integral equation

$$(4.1) \quad \begin{aligned} X(t) &= X(0) + \int_0^t b(s, X(s), \mathbf{m}(s), u(s)) ds \\ &\quad + \int_0^t \sigma(s, X(s), \mathbf{m}(s)) dW(s), \quad t \in [0, T], \end{aligned}$$

where $\mathbf{m} \in \mathcal{M}_2$ is a flow of probability measures, W a d_1 -dimensional Wiener process defined on some stochastic basis, and u a Γ -valued square-integrable adapted process.

The notion of solution of the mean field game we introduce here makes use of a version of Eq. (4.1) involving relaxed controls and varying stochastic bases. Given a flow of measures $\mathbf{m} \in \mathcal{M}_2$, consider the stochastic integral equation

$$(4.2) \quad \begin{aligned} X(t) = X(0) &+ \int_{\Gamma \times [0,t]} b(s, X(s), \mathbf{m}(s), \gamma) \rho(d\gamma, ds) \\ &+ \int_0^t \sigma(s, X(s), \mathbf{m}(s)) dW(s), \quad t \in [0, T]. \end{aligned}$$

A solution of Eq. (4.2) with flow of measures $\mathbf{m} \in \mathcal{M}_2$ is a quintuple $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), X, \rho, W)$ such that $(\Omega, \mathcal{F}, \mathbf{P})$ is a complete probability space, (\mathcal{F}_t) a filtration in \mathcal{F} satisfying the usual hypotheses, W a d_1 -dimensional (\mathcal{F}_t) -Wiener process, ρ an \mathcal{R}_2 -valued random variable adapted to (\mathcal{F}_t) , and X an \mathbb{R}^d -valued (\mathcal{F}_t) -adapted continuous process satisfying Eq. (4.2) with flow of measures \mathbf{m} \mathbf{P} -almost surely. Under our assumptions on b and σ , existence and uniqueness of solutions hold for Eq. (4.2) given any flow of measures $\mathbf{m} \in \mathcal{M}_2$. Moreover, if $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), X, \rho, W)$ is a solution, then the joint distribution of (X, ρ, W) with respect to \mathbf{P} can be identified with a probability measure on $\mathcal{B}(\mathcal{Z})$. Conversely, the set of probability measures $\Theta \in \mathcal{P}(\mathcal{Z})$ that correspond to a solution of Eq. (4.2) with respect to some stochastic basis carrying a d_1 -dimensional Wiener process can be characterized through a local martingale problem. To this end, for $f \in \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$, $\mathbf{m} \in \mathcal{M}_2$, define the process $M_f^{\mathbf{m}}$ on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ by

$$(4.3) \quad \begin{aligned} M_f^{\mathbf{m}}(t, (\varphi, r, w)) &\doteq f(\varphi(t), w(t)) - f(\varphi(0), 0) \\ &- \int_{\Gamma \times [0,t]} \mathcal{A}_{\gamma,s}^{\mathbf{m}}(f)(\varphi(s), w(s)) r(d\gamma, ds), \quad t \in [0, T], \end{aligned}$$

where

$$(4.4) \quad \begin{aligned} \mathcal{A}_{\gamma,s}^{\mathbf{m}}(f)(x, y) &\doteq \sum_{j=1}^d b_j(s, x, \mathbf{m}(s), \gamma) \frac{\partial f}{\partial x_j}(x, y) \\ &+ \frac{1}{2} \sum_{j=1}^d \sum_{k=1}^d (\sigma \sigma^\top)_{jk}(s, x, \mathbf{m}(s)) \frac{\partial^2 f}{\partial x_j \partial x_k}(x, y) \\ &+ \frac{1}{2} \sum_{l=1}^{d_1} \frac{\partial^2 f}{\partial y_l^2}(x, y) + \sum_{k=1}^d \sum_{l=1}^{d_1} \sigma_{kl}(s, x, \mathbf{m}(s)) \frac{\partial^2 f}{\partial x_k \partial y_l}(x, y). \end{aligned}$$

Recall that (\mathcal{G}_t) denotes the canonical filtration in $\mathcal{B}(\mathcal{Z})$ and $(\hat{X}, \hat{\rho}, \hat{W})$ the

coordinate process on \mathcal{Z} . By construction,

$$M_f^{\mathbf{m}}(t) = f(\hat{X}(t), \hat{W}(t)) - f(\hat{X}(0), 0) - \int_{\Gamma \times [0, t]} \mathcal{A}_{\gamma, s}^{\mathbf{m}}(f)(\hat{X}(s), \hat{W}(s)) \hat{\rho}(d\gamma, ds),$$

and $M_f^{\mathbf{m}}$ is (\mathcal{G}_t) -adapted.

Definition 4.1. A probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ is called a *solution of Eq. (4.2) with flow of measures \mathbf{m}* if the following hold:

- (i) $\mathbf{m} \in \mathcal{M}_2$;
- (ii) $\hat{W}(0) = 0$ Θ -almost surely;
- (iii) $M_f^{\mathbf{m}}$ is a local martingale with respect to the filtration (\mathcal{G}_t) and the probability measure Θ for every f monomial of first or second order.

Remark 4.1. The test functions f in (iii) of Definition 4.1 are the functions $\mathbb{R}^d \times \mathbb{R}^{d_1} \rightarrow \mathbb{R}$ given by $(x, y) \mapsto x_j$, $(x, y) \mapsto y_l$, $(x, y) \mapsto x_j \cdot x_k$, $(x, y) \mapsto y_l \cdot y_{\tilde{l}}$, and $(x, y) \mapsto x_j \cdot y_l$, where $j, k \in \{1, \dots, d\}$, $l, \tilde{l} \in \{1, \dots, d_1\}$.

The following lemma justifies the terminology of Definition 4.1.

Lemma 4.1. *Let $\mathbf{m} \in \mathcal{M}_2$. If $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), X, \rho, W)$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} , then $\Theta \doteq \mathbf{P} \circ (X, \rho, W)^{-1} \in \mathcal{P}(\mathcal{Z})$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} in the sense of Definition 4.1.*

Conversely, if $\Theta \in \mathcal{P}(\mathcal{Z})$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} in the sense of Definition 4.1, then the quintuple $((\mathcal{Z}, \mathcal{G}^\Theta, \Theta), (\mathcal{G}_{t+}^\Theta), \hat{X}, \hat{\rho}, \hat{W})$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} , where \mathcal{G}^Θ is the Θ -completion of $\mathcal{G} \doteq \mathcal{B}(\mathcal{Z})$ and $(\mathcal{G}_{t+}^\Theta)$ the right-continuous version of the Θ -augmentation of the canonical filtration (\mathcal{G}_t) .

Proof. The first part of the assertion is a consequence of Itô's formula and the local martingale property of the stochastic integral. The local martingale property of $M_f^{\mathbf{m}}$ clearly holds for any $f \in \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$.

The proof of the second part is similar to the proof of Proposition 5.4.6 in Karatzas and Shreve [1991, pp. 315-316], though here we do not need to extend the probability space; see Appendix A below. \square

A particular class of solutions of Eq. (4.2) in the sense of Definition 4.1 are those where the flow of measures $\mathbf{m} \in \mathcal{M}_2$ is induced by the probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ in the sense that $\mathbf{m}(t)$ coincides with the law of $\hat{X}(t)$ under Θ . We call those solutions McKean-Vlasov solutions:

Definition 4.2. A probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ is called a *McKean-Vlasov solution of Eq. (4.2)* if there exists $\mathbf{m} \in \mathcal{M}_2$ such that

(i) Θ is a solution of Eq. (4.2) with flow of measures \mathbf{m} ;

(ii) $\Theta \circ (\hat{X}(t))^{-1} = \mathbf{m}(t)$ for every $t \in [0, T]$.

Remark 4.2. If $\Theta \in \mathcal{P}_2(\mathcal{Z})$, then the induced flow of measures is in \mathcal{M}_2 . More precisely, let $\Theta \in \mathcal{P}_2(\mathcal{Z})$ and set $\mathbf{m}(t) \doteq \Theta \circ (\hat{X}(t))^{-1}$, $t \in [0, T]$. By definition of $\mathcal{P}_2(\mathcal{Z})$ and the metric $d_{\mathcal{Z}}$,

$$\mathbf{E}_{\Theta} \left[\|\hat{X}\|_{\mathcal{X}}^2 \right] = \int_{\mathcal{Z}} \|\varphi\|_{\mathcal{X}}^2 \Theta(d\varphi, dr, dw) < \infty.$$

This implies, in particular, that $\mathbf{m}(t) \in \mathcal{P}_2(\mathbb{R}^d)$ for every $t \in [0, T]$. By construction and definition of the square Wasserstein metric, for all $s, t \in [0, T]$,

$$d_2(\mathbf{m}(t), \mathbf{m}(s))^2 \leq \mathbf{E}_{\Theta} \left[|\hat{X}(t) - \hat{X}(s)|^2 \right].$$

Continuity of the trajectories of \hat{X} and the dominated convergence theorem with $2\|\hat{X}\|_{\mathcal{X}}^2$ as dominating Θ -integrable random variable imply that $d_2(\mathbf{m}(t), \mathbf{m}(s)) \rightarrow 0$ whenever $|t - s| \rightarrow 0$. It follows that $\mathbf{m} \in \mathcal{M}_2$.

Uniqueness holds not only for solutions of Eq. (4.2) with fixed flow of measures $\mathbf{m} \in \mathcal{M}_2$, but also for McKean-Vlasov solutions of Eq. (4.2).

Lemma 4.2. *Let $\Theta, \tilde{\Theta} \in \mathcal{P}_2(\mathcal{Z})$. If $\Theta, \tilde{\Theta}$ are McKean-Vlasov solutions of Eq. (4.2) such that $\Theta \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1} = \tilde{\Theta} \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1}$, then $\Theta = \tilde{\Theta}$.*

Proof. Let $\Theta, \tilde{\Theta} \in \mathcal{P}_2(\mathcal{Z})$ be McKean-Vlasov solutions of Eq. (4.2) such that $\Theta \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1} = \tilde{\Theta} \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1}$. Set

$$\mathbf{m}(t) \doteq \Theta \circ \hat{X}(t)^{-1}, \quad \tilde{\mathbf{m}}(t) \doteq \tilde{\Theta} \circ \hat{X}(t)^{-1}, \quad t \in [0, T].$$

In view of Remark 4.2, we have $\mathbf{m}, \tilde{\mathbf{m}} \in \mathcal{M}_2$. Define an extended canonical space $\bar{\mathcal{Z}}$ by

$$\bar{\mathcal{Z}} \doteq \mathcal{X} \times \mathcal{X} \times \mathcal{R}_2 \times \mathcal{W}.$$

Let $(\bar{\mathcal{G}})_{t \geq 0}$ denote the canonical filtration in $\bar{\mathcal{G}} \doteq \mathcal{B}(\bar{\mathcal{Z}})$, and let $(X, \tilde{X}, \hat{\rho}, \hat{W})$ be the canonical process. A construction analogous to the one used in the proof of Proposition 1 in Yamada and Watanabe [1971] (also see Section 5.3.D in Karatzas and Shreve [1991]) yields a measure $Q \in \mathcal{P}(\bar{\mathcal{Z}})$ such that

$$Q \circ (X, \hat{\rho}, \hat{W})^{-1} = \Theta, \quad Q \circ (\tilde{X}, \hat{\rho}, \hat{W})^{-1} = \tilde{\Theta}, \quad Q \left\{ X(0) = \tilde{X}(0) \right\} = 1.$$

By Lemma 4.1, $((\bar{\mathcal{Z}}, \bar{\mathcal{G}}^Q, Q), (\bar{\mathcal{G}}_{t+}^Q), X, \hat{\rho}, \hat{W}), ((\bar{\mathcal{Z}}, \bar{\mathcal{G}}^Q, Q), (\bar{\mathcal{G}}_{t+}^Q), \tilde{X}, \hat{\rho}, \hat{W})$ are solutions of Eq. (4.2) with flow of measures \mathbf{m} and $\tilde{\mathbf{m}}$, respectively, where

$\bar{\mathcal{G}}^Q$ is the Q -completion of $\bar{\mathcal{G}}$ and $(\bar{\mathcal{G}}_{t+}^Q)$ the right-continuous version of the Q -augmentation of $(\bar{\mathcal{G}}_t)$.

By construction and definition of the square Wasserstein distance,

$$d_2(\mathbf{m}(t), \tilde{\mathbf{m}}(t))^2 \leq \mathbf{E}_Q \left[\left| X(t) - \tilde{X}(t) \right|^2 \right] \quad \text{for all } t \in [0, T].$$

Using (A2), Hölder's inequality, Itô's isometry, Fubini's theorem and the fact that $X(0) = \tilde{X}(0)$ Q -almost surely, we find that for every $t \in [0, T]$,

$$\begin{aligned} & \mathbf{E}_Q \left[\left| X(t) - \tilde{X}(t) \right|^2 \right] \\ & \leq 4(T+1)L^2 \int_0^t \mathbf{E}_Q \left[\left| X(s) - \tilde{X}(s) \right|^2 + d_2(\mathbf{m}(s), \tilde{\mathbf{m}}(s))^2 \right] ds \\ & \leq 8(T+1)L^2 \int_0^t \mathbf{E}_Q \left[\left| X(s) - \tilde{X}(s) \right|^2 \right] ds. \end{aligned}$$

Gronwall's lemma and the continuity of trajectories imply that $X = \tilde{X}$ Q -almost surely and that $\mathbf{m} = \tilde{\mathbf{m}}$. It follows that $\Theta = \tilde{\Theta}$. \square

Define the costs associated with a flow of measures $\mathbf{m} \in \mathcal{M}_2$, an initial distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$ and a probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ by

$$\hat{J}(\nu, \Theta; \mathbf{m}) \doteq \begin{cases} \mathbf{E}_\Theta \left[\int_{\Gamma \times [0, T]} f(s, \hat{X}(s), \mathbf{m}(s), \gamma) \hat{\rho}(d\gamma, ds) + F(\hat{X}(T), \mathbf{m}(T)) \right] \\ \quad \text{if } \Theta \text{ is a solution of Eq. (4.2) with flow of measures } \mathbf{m} \\ \quad \text{and } \Theta \circ \hat{X}(0)^{-1} = \nu, \\ \infty \quad \text{otherwise.} \end{cases}$$

This defines a measurable mapping $\hat{J}: \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(\mathcal{Z}) \times \mathcal{M}_2 \rightarrow [0, \infty]$. The corresponding value function $\hat{V}: \mathcal{P}(\mathbb{R}^d) \times \mathcal{M}_2 \rightarrow [0, \infty]$ is given by

$$\hat{V}(\nu; \mathbf{m}) \doteq \inf_{\Theta \in \mathcal{P}(\mathcal{Z})} \hat{J}(\nu, \Theta; \mathbf{m}).$$

Definition 4.3. A pair (Θ, \mathbf{m}) is called a *solution of the mean field game* if the following hold:

- (i) $\mathbf{m} \in \mathcal{M}_2$, $\Theta \in \mathcal{P}(\mathcal{Z})$, and Θ is a solution of Eq. (4.2) with flow of measures \mathbf{m} ;
- (ii) Mean field condition: $\Theta \circ \hat{X}(t)^{-1} = \mathbf{m}(t)$ for every $t \in [0, T]$;
- (iii) Optimality condition: $\hat{J}(\mathbf{m}(0), \Theta; \mathbf{m}) \leq \hat{J}(\mathbf{m}(0), \tilde{\Theta}; \mathbf{m})$ for every $\tilde{\Theta} \in \mathcal{P}(\mathcal{Z})$.

In Definition 4.3, there is some redundancy in the choice of the pair (Θ, \mathbf{m}) as solution of the mean field game in that, thanks to the mean field condition, the flow of measures \mathbf{m} is completely determined by the probability measure Θ . Consequently, we may call a probability measure $\Theta \in \mathcal{P}(\mathcal{Z})$ a *solution of the mean field game* if the pair (Θ, \mathbf{m}) is a solution of the mean field game in the sense of Definition 4.3 where \mathbf{m} is the flow of measures induced by Θ , that is, $\mathbf{m}(t) \doteq \Theta \circ \hat{X}(t)^{-1}$, $t \in [0, T]$.

If Θ is a solution of the mean field game, then, again thanks to the mean field condition, it is also a McKean-Vlasov solution of Eq. (4.2). In general, however, Θ is not optimal as a controlled McKean-Vlasov solution. In the optimality condition of Definition 4.3, in fact, the flow of measures is frozen at the flow of measures induced by Θ , while in an optimization problem of McKean-Vlasov type the flow of measures would have to vary with the controlled solution.

Remark 4.3. The use of relaxed controls in Definition 4.3 has a twofold motivation. The first is pragmatic and well known [for instance, El Karoui et al., 1987, Kushner, 1990], namely the fact that relaxed controls allow one to embed the space of control processes into a nice topological space (if Γ is compact, then $\mathcal{R} = \mathcal{R}_2$ is compact; for unbounded Γ , \mathcal{R}_2 is still Polish) without changing the minimal costs. In particular, existence of optimal controls is guaranteed in the space of relaxed controls. The second motivation is related to this fact, but more conceptual. The mean field condition in the mean field game is required to hold for the law of the state process under an optimal control only. Thus, existence of optimal controls (for a given flow of measures) is crucial for the existence of solutions to the mean field game. For ordinary optimal control problems, on the other hand, it suffices that the minimal costs be well defined. Still, it is natural to ask for conditions ensuring that a solution of the mean field game can be obtained in ordinary control processes, not just in relaxed controls. Sufficient conditions of this kind have been established in Lacker [2015a]. One simple sufficient condition is that the dynamics be linear and the costs convex in the control.

The next lemma will be an essential ingredient in the construction of competitor strategies in the proof of Theorem 5.1 below.

Lemma 4.3. *Let $\mathbf{m} \in \mathcal{M}_2$. Given any $\varepsilon > 0$, there exists a measurable function $\psi_\varepsilon^{\mathbf{m}} : [0, T] \times \mathbb{R}^d \times \mathcal{W} \rightarrow \Gamma$ such that the following hold:*

- (i) $\psi_\varepsilon^{\mathbf{m}}$ is progressively measurable in the sense that, for every $t \in [0, T]$, every $x \in \mathbb{R}^d$, we have $\psi_\varepsilon^{\mathbf{m}}(t, x, w) = \psi_\varepsilon^{\mathbf{m}}(t, x, \tilde{w})$ whenever $w(s) = \tilde{w}(s)$ for all $s \in [0, t]$;

(ii) $\psi_\varepsilon^{\mathbf{m}}$ takes values in a finite subset of Γ ;

(iii) $\hat{J}(\mathbf{m}(0), \Theta_\varepsilon^{\mathbf{m}}; \mathbf{m}) \leq \hat{V}(\mathbf{m}(0); \mathbf{m}) + \varepsilon$, where $\Theta_\varepsilon^{\mathbf{m}}$ is the unique probability measure in $\mathcal{P}_2(\mathcal{Z})$ such that $\Theta_\varepsilon^{\mathbf{m}}$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} , $\Theta_\varepsilon^{\mathbf{m}} \circ (\hat{X}(0))^{-1} = \mathbf{m}(0)$, and

$$\hat{\rho}(d\gamma, dt) = \delta_{\psi_\varepsilon^{\mathbf{m}}(t, \hat{X}(0), \hat{W})}(d\gamma) dt \quad \Theta_\varepsilon^{\mathbf{m}}\text{-almost surely.}$$

Proof. The proof is based on time discretization and dynamic programming; see Appendix B. \square

Remark 4.4. The conditions of Lemma 4.3 do not determine $\psi_\varepsilon^{\mathbf{m}}$ in a unique way. On the other hand, once $\psi_\varepsilon^{\mathbf{m}}$ has been constructed, the probability measure $\Theta_\varepsilon^{\mathbf{m}}$ is uniquely determined as the law of the solution of Eq. (4.2) with flow of measures \mathbf{m} , initial distribution $\mathbf{m}(0)$ and control process u given by $u(t) \doteq \psi_\varepsilon^{\mathbf{m}}(t, X(0), W)$, $t \in [0, T]$, where W is the driving Wiener process and u is identified with its relaxed control random variable. Notice that u is square-integrable since $\psi_\varepsilon^{\mathbf{m}}$ takes values in a finite subset of Γ .

5 Convergence of Nash equilibria

For $N \in \mathbb{N}$, let $u_1^N, \dots, u_N^N \in \mathcal{H}_2((\mathcal{F}_t^N), \mathbf{P}_N; \Gamma)$ be individual strategies for the N -player game, and let $\mathbf{u}^N \doteq (u_1^N, \dots, u_N^N)$ be the corresponding strategy vector. Let Q^N be the normalized occupation measure associated with \mathbf{u}^N . More precisely, Q^N is the $\mathcal{P}_2(\mathcal{Z})$ -valued random variable determined by setting, for $B \in \mathcal{B}(\mathcal{X})$, $R \in \mathcal{B}(\mathcal{R}_2)$, $D \in \mathcal{B}(\mathcal{W})$,

$$(5.1) \quad Q_\omega^N(B \times R \times D) \doteq \frac{1}{N} \sum_{i=1}^N \delta_{X_i^N(\cdot, \omega)}(B) \cdot \delta_{\rho_\omega^{N,i}}(R) \cdot \delta_{W_i^N(\cdot, \omega)}(D), \quad \omega \in \Omega_N,$$

where (X_1^N, \dots, X_N^N) is the solution of the system of equations (3.1) under strategy vector \mathbf{u}^N , and $\rho^{N,i}$ is the relaxed control associated with individual strategy u_i^N , $i \in \{1, \dots, N\}$.

Convergence results will be obtained under the hypothesis that

$$(T) \quad \exists \delta_0 > 0 : \sup_{N \in \mathbb{N}} \mathbf{E}_N \left[\frac{1}{N} \sum_{i=1}^N \left(|\xi_i^N|^{2+\delta_0} + \int_0^T |u_i^N(t)|^{2+\delta_0} dt \right) \right] < \infty.$$

Whenever (T) holds, we will—as we may—suppose that $\delta_0 \in (0, 1 \wedge T]$.

Remark 5.1. Condition (T) is automatically satisfied if the action space Γ is compact and the initial states, i.e., the random variables ξ_i^N , $N \in \mathbb{N}$, $i \in \{1, \dots, N\}$, are uniformly bounded.

Lemma 5.1. *If condition (T) holds, then the family $(\mathbf{P}_N \circ (Q^N)^{-1})_{N \in \mathbb{N}}$ is pre-compact in $\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))$.*

Proof. We verify that condition (T) implies the pre-compactness of the family $(\mathbf{P}_N \circ (Q^N)^{-1})_{N \in \mathbb{N}}$ by using a suitable tightness function on $\mathcal{P}_2(\mathcal{Z})$. For a function ψ on $[0, T]$ with values in \mathbb{R}^d or \mathbb{R}^{d_1} , let $\mathbf{w}_\psi(\cdot, T)$ denote the modulus of continuity of ψ on $[0, T]$, that is, the function

$$[0, \infty) \ni h \mapsto \mathbf{w}_\psi(h, T) \doteq \sup_{t, s \in [0, T]: |t-s| \leq h} |\psi(t) - \psi(s)| \in [0, \infty].$$

If ψ is continuous, then the modulus of continuity of ψ takes values in $[0, \infty)$. Clearly, $\mathbf{w}_\psi(h, T) = \mathbf{w}_\psi(T, T)$ whenever $h > T$. Choose $\delta_0 > 0$ according to condition (T), and set $\alpha \doteq \frac{\delta_0}{2(8+\delta_0)}$. Define the function $g: \mathcal{P}_2(\mathcal{Z}) \rightarrow [0, \infty]$ by

$$(5.2) \quad g(\Theta) \doteq \int_{\mathcal{Z}} \left(\|\varphi\|_{\mathcal{X}}^{2+\delta_0} + |w(0)| + \int_{\Gamma \times [0, T]} |\gamma|^{2+\delta_0} r(d\gamma, dt) \right. \\ \left. + \sup_{h \in (0, 1]} \{h^{-\alpha} (\mathbf{w}_\varphi(h, T) + \mathbf{w}_w(h, T))\} \right) \Theta(dw, dr, d\varphi).$$

Then g is a tightness function on $\mathcal{P}_2(\mathcal{Z})$; see Appendix C.2. It is therefore enough to check that condition (T) entails $\sup_{N \in \mathbb{N}} \mathbf{E}_N [g(Q^N)] < \infty$. By definition of Q^N and g ,

$$\mathbf{E}_N [g(Q^N)] = \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\|X_i^N\|_{\mathcal{X}}^{2+\delta_0} + \int_0^T |u_i^N(t)|^{2+\delta_0} dt \right] \\ + \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{h \in (0, 1]} \left\{ h^{-\alpha} \left(\mathbf{w}_{X_i^N}(h, T) + \mathbf{w}_{W_i^N}(h, T) \right) \right\} \right].$$

By Lemma 3.2 and condition (T),

$$\sup_{N \in \mathbb{N}} \left\{ \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\|X_i^N\|_{\mathcal{X}}^{2+\delta_0} + \int_0^T |u_i^N(t)|^{2+\delta_0} dt \right] \right\} < \infty.$$

As to the terms involving the moduli of continuity, set $p \doteq 2 + \delta_0/2$; then, by monotonicity of $h \mapsto h^{-\alpha}$ and Markov's inequality (as well as Jensen's

inequality),

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{h \in (0,1]} \left\{ h^{-\alpha} \left(\mathbf{w}_{X_i^N}(h, T) + \mathbf{w}_{W_i^N}(h, T) \right) \right\} \right] \\
& \leq \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{k \in \mathbb{N}: k \geq 1/T} \left\{ (k+1)^\alpha \left(\mathbf{w}_{X_i^N} \left(\frac{1}{k}, T \right) + \mathbf{w}_{W_i^N} \left(\frac{1}{k}, T \right) \right) \right\} \right] \\
& \leq 1 + \frac{1}{N} \sum_{i=1}^N \int_1^\infty \sum_{k=1}^\infty \mathbf{P}_N \left(\mathbf{w}_{X_i^N} \left(\frac{1}{k}, T \right) + \mathbf{w}_{W_i^N} \left(\frac{1}{k}, T \right) \geq \frac{M}{(k+1)^\alpha} \right) dM \\
& \leq 1 + \sum_{k=1}^\infty (k+1)^{\alpha p} \left(\frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\mathbf{w}_{X_i^N} \left(\frac{1}{k}, T \right)^p + \mathbf{w}_{W_i^N} \left(\frac{1}{k}, T \right)^p \right] \right)^{\frac{2^{p-1}}{p-1}},
\end{aligned}$$

where we have used that $\int_1^\infty M^{-p} dM = 1/(p-1)$ since $p > 1$. To find an upper bound for the above sums that does not depend on N , we employ estimates on the moments of the modulus of continuity of Itô processes; cf. Fischer and Nappo [2010] and the references therein. Since W_1^N, \dots, W_N^N are standard d_1 -dimensional Wiener processes, we have by Lemma 3 of that paper and Hölder's inequality that there exists a finite constant \bar{C}_{p,d_1} depending only on p and d_1 such that, for every $i \in \{1, \dots, N\}$, every $k \in \mathbb{N}$ with $k \geq 1/T$,

$$\mathbf{E}_N \left[\mathbf{w}_{W_i^N} \left(\frac{1}{k}, T \right)^p \right] \leq \bar{C}_{p,d_1} \left(\frac{\log(2Tk)}{k} \right)^{p/2}.$$

Recall that $p = 2 + \delta_0/2$. By Theorem 1 in Fischer and Nappo [2010], there exists a finite constant $\bar{C}_{\delta_0, d, d_1}$ depending only on δ_0 (through $p = 2 + \delta_0/2$ and $\delta_0/2 = 2 + \delta_0 - p$), d , and d_1 such that, for every $k \in \mathbb{N}$ with $k \geq 1/T$,

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\mathbf{w}_{X_i^N} \left(\frac{1}{k}, T \right)^{2+\delta_0/2} \right] \\
& \leq \bar{C}_{\delta_0, d, d_1} \left(\frac{\log(2Tk)}{k} \right)^{1+\delta_0/4} \\
& \quad \cdot \left(\frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{s, t \in [0, T]: s < t} \left(\frac{\int_s^t |b(\tilde{s}, X_i^N(\tilde{s}), \mu^N(\tilde{s}), u_i^N(\tilde{s}))| d\tilde{s}}{\sqrt{|t-s|}} \right)^{2+\delta_0/2} \right] \right. \\
& \quad \left. + \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{s \in [0, T]} |\sigma(s, X_i^N(s), \mu^N(s))|^{2+\delta_0} \right] + 1 \right).
\end{aligned}$$

Thanks to assumption (A3), Lemma 3.2 and condition (T), we have

$$\sup_{N \in \mathbb{N}} \left\{ \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{s \in [0, T]} |\sigma(s, X_i^N(s), \mu^N(s))|^{2+\delta_0} \right] \right\} < \infty.$$

On the other hand, by Hölder's inequality,

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{s,t \in [0,T]: s < t} \left(\frac{\int_s^t |b(\tilde{s}, X_i^N(\tilde{s}), \mu^N(\tilde{s}), u_i^N(\tilde{s}))| d\tilde{s}}{\sqrt{|t-s|}} \right)^{2+\delta_0/2} \right] \\ & \leq T^{\delta_0/4} \cdot \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\int_0^T |b(\tilde{s}, X_i^N(\tilde{s}), \mu^N(\tilde{s}), u_i^N(\tilde{s}))|^{2+\delta_0/2} d\tilde{s} \right], \end{aligned}$$

and, thanks to assumption (A3), Lemma 3.1 and condition (T),

$$\sup_{N \in \mathbb{N}} \left\{ \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\int_0^T |b(\tilde{s}, X_i^N(\tilde{s}), \mu^N(\tilde{s}), u_i^N(\tilde{s}))|^{2+\delta_0/2} d\tilde{s} \right] \right\} < \infty.$$

Recall that $\alpha = \frac{\delta_0}{2(8+\delta_0)}$ and $p = 2 + \delta_0/2$. It follows that, for some finite constant $\bar{C}_{K,T,\delta_0,d,d_1}$ not depending on N ,

$$\begin{aligned} & \sup_{N \in \mathbb{N}} \left\{ \frac{1}{N} \sum_{i=1}^N \mathbf{E}_N \left[\sup_{h \in (0,1]} \left\{ h^{-\alpha} \left(\mathbf{w}_{X_i^N}(h, T) + \mathbf{w}_{W_i^N}(h, T) \right) \right\} \right] \right\} \\ & \leq \bar{C}_{K,T,\delta_0,d,d_1} \left(1 + \sum_{k=1}^{\infty} (k+1)^{\alpha \cdot p} \left(\frac{\log(2Tk)}{k} \right)^{p/2} \right), \end{aligned}$$

where the infinite sum on the right-hand side above has a finite limit since $p/2 - \alpha \cdot p = (8 + 2\delta_0)/(8 + \delta_0) > 1$. \square

Below, we will use the symbol \mathbb{I} to indicate the index set of a (convergent) subsequence; thus \mathbb{I} is a subset of \mathbb{N} with the natural ordering and $\#\mathbb{I} = \infty$.

Lemma 5.2. *Suppose that $(\mathbf{P}_n \circ \xi_{i_*^n}^n)_{n \in \mathbb{I}}$ converges in $\mathcal{P}_2(\mathbb{R}^d)$ to some $\bar{\nu} \in \mathcal{P}_2(\mathbb{R}^d)$, where, for each $n \in \mathbb{I}$, $i_*^n \in \{1, \dots, n\}$. Then there exists a sequence $(\bar{\xi}^n)_{n \in \mathbb{I}}$ of \mathbb{R}^d -valued random variables such that the following hold:*

(i) *for every $n \in \mathbb{I}$, $\bar{\xi}^n$ is defined on $(\Omega_n, \mathcal{F}^n)$, measurable with respect to $\sigma(\xi_{i_*^n}^n, \vartheta_{i_*^n}^n) \subset \mathcal{F}_0^n$, and such that $\mathbf{P}_n \circ (\bar{\xi}^n)^{-1} = \bar{\nu}$;*

(ii) $\mathbf{E}_n \left[|\xi_{i_*^n}^n - \bar{\xi}^n|^2 \right] \rightarrow 0$ as $n \rightarrow \infty$.

Proof. Set $\nu_n \doteq \mathbf{P}_n \circ (\xi_{i_*^n}^n)^{-1}$. By hypothesis,

$$d_2(\nu_n, \bar{\nu}) \xrightarrow{n \rightarrow \infty} 0.$$

Let $n \in \mathbb{I}$. By definition of the square Wasserstein metric,

$$d_2(\nu_n, \bar{\nu})^2 = \inf_{\alpha \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d): [\alpha]_1 = \nu_n \text{ and } [\alpha]_2 = \bar{\nu}} \int_{\mathbb{R}^d \times \mathbb{R}^d} |x - \tilde{x}|^2 \alpha(dx, d\tilde{x}).$$

The infimum in the above equation is attained; see, for instance, Theorem 1.3 (Kantorovich's theorem) in Villani [2003, pp. 19-20]. Thus, there exists $\alpha_*^n \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ such that $[\alpha_*^n]_1 = \nu_n$, $[\alpha_*^n]_2 = \bar{\nu}$ and

$$d_2(\nu_n, \bar{\nu})^2 = \int_{\mathbb{R}^d \times \mathbb{R}^d} |x - \tilde{x}|^2 \alpha_*^n(dx, d\tilde{x}).$$

Recall that $\vartheta_1^n, \dots, \vartheta_n^n$ are independent \mathcal{F}_0^n -measurable random variables which are uniformly distributed on $[0, 1]$ and independent of the σ -algebra generated by $\xi_1^n, \dots, \xi_n^n, W_1^n, \dots, W_n^n$. By Theorem 6.10 in Kallenberg [2001, p. 112] on measurable transfers, there exists a measurable function $\varphi_n : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ such that

$$\mathbf{P}_n \circ \left(\xi_{i_*}^n, \varphi_n(\xi_{i_*}^n, \vartheta_{i_*}^n) \right)^{-1} = \alpha_*^n.$$

Set $\bar{\xi}^n \doteq \varphi_n(\xi_{i_*}^n, \vartheta_{i_*}^n)$. Then $\bar{\xi}^n$ is $\sigma(\xi_{i_*}^n, \vartheta_{i_*}^n)$ -measurable, $\mathbf{P}_n \circ (\bar{\xi}^n)^{-1} = \bar{\nu}$, and

$$\mathbf{E}_n \left[|\xi_{i_*}^n - \bar{\xi}^n|^2 \right] = d_2(\nu_n, \bar{\nu})^2,$$

which tends to zero as $n \rightarrow \infty$. \square

Lemma 5.3. *Grant condition (T). Let $(Q^n)_{n \in \mathbb{I}}$ be a subsequence that converges in distribution to some $\mathcal{P}_2(\mathcal{Z})$ -valued random variable Q defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Set*

$$\mu_\omega(t) \doteq Q_\omega \circ \hat{X}(t)^{-1}, \quad t \in [0, T], \quad \omega \in \Omega.$$

Then for \mathbf{P} -almost every $\omega \in \Omega$, $\mu_\omega \in \mathcal{M}_2$ and Q_ω is a solution of Eq. (4.2) with flow of measures μ_ω . Moreover,

$$\liminf_{\mathbb{I} \ni n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n J_i^n(\mathbf{u}^n) \geq \int_{\Omega} \hat{J}(\mu_\omega(0), Q_\omega, \mu_\omega) \mathbf{P}(d\omega).$$

Proof. By Lemma 5.1, $(\mathbf{P}_N \circ (Q^N)^{-1})_{N \in \mathbb{N}}$ is pre-compact in $\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))$. Let $(Q^n)_{n \in \mathbb{I}}$ be a subsequence that converges in distribution to some $\mathcal{P}_2(\mathcal{Z})$ -valued random variable Q , defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Set $\mu_\omega(t) \doteq Q_\omega \circ \hat{X}(t)^{-1}$, $t \in [0, T]$, $\omega \in \Omega$. Since $Q_\omega \in \mathcal{P}_2(\mathcal{Z})$ for every $\omega \in \Omega$, we have $\mu_\omega \in \mathcal{M}_2$ for every $\omega \in \Omega$; cf. Remark 4.2 above. By construction, $\hat{W}(0) = 0$ Q_ω^n -almost surely for \mathbf{P}_n -almost every $\omega \in \Omega_n$. Convergence in distribution implies $\hat{W}(0) = 0$ Q_ω -almost surely for \mathbf{P} -almost every $\omega \in \Omega$.

In order to verify that Q_ω is a solution of Eq. (4.2) with flow of measures μ_ω for \mathbf{P} -almost every $\omega \in \Omega$, it suffices to check that condition (iii) of Definition 4.1 holds. The proof of this fact is analogous to the proof of

Lemma 5.2 in Budhiraja et al. [2012]. Since the situation here is somewhat different, we give details in Appendix D below.

The asymptotic lower bound for the average costs is a consequence of a version of Fatou's lemma [cf. Theorem A.3.12 Dupuis and Ellis, 1997, p. 307] since, for every $n \in \mathbb{I}$,

$$\frac{1}{n} \sum_{i=1}^n J_i^n(\mathbf{u}^n) = \int_{\Omega_n} \int_{\mathcal{Z}} \left(\int_{\Gamma \times [0, T]} f(t, \varphi(t), Q_\omega^n \circ \hat{X}(t)^{-1}, \gamma) r(d\gamma, dt) \right. \\ \left. F(T, \varphi(T), Q_\omega^n \circ \hat{X}(T)^{-1}) \right) Q_\omega^n(d\varphi, dr, dw) \mathbf{P}_n(d\omega)$$

and $Q_\omega^n \circ \hat{X}(t)^{-1} \rightarrow \mu(t)$ in distribution as $n \rightarrow \infty$. \square

Remark 5.2. Lemma 5.3 shows that, under condition (T), all limit points of the normalized occupation measures $(Q^N)_{N \in \mathbb{N}}$ are concentrated on those random variables that, with probability one, take values in the set of McKean-Vlasov solutions of Eq. (4.2). The mean field condition of Definition 4.3 is therefore always satisfied.

In addition to (T), we will need the following weak symmetry condition on the costs:

\exists a sequence of indices $(i_*^N)_{N \in \mathbb{N}}$ with $i_*^N \in \{1, \dots, N\}$ such that

$$(S) \quad \sup_{N \in \mathbb{N}} J_{i_*^N}^N(\mathbf{u}^N) < \infty \text{ and } \limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N J_i^N(\mathbf{u}^N) \leq \limsup_{N \rightarrow \infty} J_{i_*^N}^N(\mathbf{u}^N).$$

Remark 5.3. Condition (S) is automatically satisfied if the cost coefficients f, F are bounded functions. If f, F are unbounded and the costs associated with \mathbf{u}^N are symmetric in the sense that, for every N , every $i \in \{2, \dots, N\}$, $J_1^N(\mathbf{u}^N) = J_i^N(\mathbf{u}^N)$, then, thanks to assumption (A5) and Lemma 3.1, condition (S) follows from condition (T).

Theorem 5.1. *Let $(\varepsilon_N)_{N \in \mathbb{N}} \subset [0, \infty)$ be a sequence converging to zero. Suppose that $(\xi^N)_{N \in \mathbb{N}}$ and $(\mathbf{u}^N)_{N \in \mathbb{N}}$ are such that (T) and (S) hold and, for each $N \in \mathbb{N}$, $\xi^N = (\xi_1^N, \dots, \xi_N^N)$ is exchangeable and \mathbf{u}^N is a local ε_N -Nash equilibrium for the N -player game. Let $(Q^n)_{n \in \mathbb{I}}$ be a subsequence that converges in distribution to some $\mathcal{P}_2(\mathcal{Z})$ -valued random variable Q defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. If there is $\mathbf{m} \in \mathcal{M}_2$ such that, for \mathbf{P} -almost every $\omega \in \Omega$,*

$$Q_\omega \circ \hat{X}(t)^{-1} = \mathbf{m}(t), \quad t \in [0, T],$$

then (Q_ω, \mathbf{m}) is a solution of the mean field game for \mathbf{P} -almost every $\omega \in \Omega$.

We postpone the proof of Theorem 5.1 to the end of this section. The crucial hypothesis in Theorem 5.1 is the almost sure non-randomness of the flow of measures induced by a limit random variable Q . Thus, under the rather general conditions (T) and (S), we prove convergence to solutions of a mean field game for subsequences with limit random variable Q such that $\mathbf{P} \circ (Q \circ (\hat{X}(t))_{t \in [0, T]}^{-1})^{-1} = \delta_{\mathbf{m}}$ for some $\mathbf{m} \in \mathcal{M}_2$. This condition is reminiscent of the characterization of propagation of chaos in the Tanaka-Sznitman theorem. The non-randomness of the induced flow of measures is implied by the non-randomness of the joint law of initial condition, relaxed control and noise process, that is, by the condition $\mathbf{P} \circ (Q \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1})^{-1} = \delta_{\nu}$ for some $\nu \in \mathcal{P}(\mathbb{R}^d \times \mathcal{R}_2 \times \mathcal{W})$. This condition, in turn, is satisfied if the initial states and individual strategies of each N -player game are independent and identically distributed, where the marginal distributions are allowed to vary with N .

Corollary 5.2. *Let $(\varepsilon_N)_{N \in \mathbb{N}} \subset [0, \infty)$ be a sequence converging to zero. Suppose that $(\xi^N)_{N \in \mathbb{N}}$ and $(\mathbf{u}^N)_{N \in \mathbb{N}}$ are such that (T) holds and, for each $N \in \mathbb{N}$, \mathbf{u}^N is a local ε_N -Nash equilibrium for the N -player game and the random variables $(\xi_1^N, u_1^N, W_1^N), \dots, (\xi_N^N, u_N^N, W_N^N)$ are independent and identically distributed. Let $(Q^n)_{n \in \mathbb{I}}$ be a subsequence that converges in distribution to some $\mathcal{P}_2(\mathcal{Z})$ -valued random variable Q defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$. Then Q_ω is a solution of the mean field game for \mathbf{P} -almost every $\omega \in \Omega$.*

Proof. By distributional symmetry of the vectors of initial states and individual strategies, the costs are symmetric and condition (T) entails condition (S); cf. Remark 5.3 above.

Let $\mathcal{T} \subset \mathbf{C}_b(\mathbb{R}^d \times \mathcal{R}_2 \times \mathcal{W})$ be a countable and measure determining set of functions. Let $(Q^n)_{n \in \mathbb{I}}$ be a convergent subsequence with limit random variable Q on $(\Omega, \mathcal{F}, \mathbf{P})$. Let $\Psi \in \mathcal{T}$, and set

$$\begin{aligned} m_\Psi &\doteq \mathbf{E}_{\mathbf{P}} \left[\mathbf{E}_Q \left[\Psi(\hat{X}(0), \hat{\rho}, \hat{W}) \right] \right], \\ v_\Psi &\doteq \mathbf{E}_{\mathbf{P}} \left[\left(\mathbf{E}_Q \left[\Psi(\hat{X}(0), \hat{\rho}, \hat{W}) \right] - m_\Psi \right)^2 \right], \\ m_\Psi^n &\doteq \mathbf{E}_n \left[\mathbf{E}_{Q^n} \left[\Psi(\hat{X}(0), \hat{\rho}, \hat{W}) \right] \right], \quad n \in \mathbb{I}. \end{aligned}$$

The mapping $\Theta \mapsto \int \Psi d\Theta$ is continuous on $\mathcal{P}_2(\mathcal{Z})$. By convergence of (Q^n)

to Q and the continuous mapping theorem,

$$\begin{aligned} v_\Psi &= \lim_{n \rightarrow \infty} \mathbf{E}_n \left[\left(\mathbf{E}_{Q^n} \left[\Psi(\hat{X}(0), \hat{\rho}, \hat{W}) \right] - m_\Psi^n \right)^2 \right] \\ &= \lim_{n \rightarrow \infty} \mathbf{E}_n \left[\left(\frac{1}{n} \sum_{i=1}^n \Psi(\xi_i^n, \rho^{n,i}, W_i^n) - m_\Psi^n \right)^2 \right], \end{aligned}$$

where $\rho^{n,i}$ is the relaxed control random variable induced by u_i^n . As a consequence of the i.i.d. hypothesis, the random variables $\Psi(\xi_i^n, \rho^{n,i}, W_i^n)$, $i \in \{1, \dots, n\}$, are independent and identically distributed with common mean equal to m_Ψ^n . Since Ψ is bounded, it follows that $v_\Psi = 0$. This implies

$$\mathbf{E}_Q \left[\Psi(\hat{X}(0), \hat{\rho}, \hat{W}) \right] = m_\Psi \quad \mathbf{P}\text{-almost surely.}$$

Since \mathcal{T} is countable, we have with \mathbf{P} -probability one

$$\mathbf{E}_Q \left[\Psi(\hat{X}(0), \hat{\rho}, \hat{W}) \right] = m_\Psi \quad \text{for all } \Psi \in \mathcal{T}.$$

Since \mathcal{T} is also measure determining, it follows that there exists a measure $\nu \in \mathcal{P}(\mathbb{R}^d \times \mathcal{R}_2 \times \mathcal{W})$ such that, for \mathbf{P} -almost every $\omega \in \Omega$,

$$Q_\omega \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1} = \nu.$$

On the other hand, we know by Lemma 5.3 that $Q_\omega \in \mathcal{P}_2(\mathcal{Z})$ is a McKean-Vlasov solution of Eq. (4.2) for \mathbf{P} -almost every $\omega \in \Omega$. Uniqueness of such solutions according to Lemma 4.2 yields the existence of a measure $\Theta \in \mathcal{P}_2(\mathcal{Z})$ such that $Q_\omega = \Theta$ for \mathbf{P} -almost every $\omega \in \Omega$. Let $\mathbf{m} \in \mathcal{M}_2$ be the flow of measures induced by Θ . Then, for \mathbf{P} -almost every $\omega \in \Omega$,

$$Q_\omega \circ \hat{X}(t)^{-1} = \mathbf{m}(t), \quad t \in [0, T].$$

The assertion is now a consequence of Theorem 5.1. \square

Existence of local approximate Nash equilibria as required in Corollary 5.2 is guaranteed, in particular, under the hypotheses of Proposition 3.1 above (compact action space, bounded coefficients). Suppose that (ξ^N) is such that, for each $N \in \mathbb{N}$, ξ^N is a vector of independent and identically distributed random variables with common marginal $\mathbf{m}_0^N \in \mathcal{P}_2(\mathbb{R}^d)$ and that, for some $\delta_0 > 0$, $\sup_{N \in \mathbb{N}} \int |x|^{2+\delta_0} \mathbf{m}_0^N(dx) < \infty$. Then, by Proposition 3.1, there exists a corresponding sequence (\mathbf{u}^N) of local approximate Nash equilibria such that the hypotheses of Corollary 5.2 are satisfied. In addition to the desired limit relation, we thus obtain a proof of existence of solutions for the mean field game. Note that existence of solutions is just a by-product of

our analysis; analogous existence results can in fact be obtained by directly working with the mean field game; see Lacker [2015a]. The proof there is based, as in Proposition 3.1 here, on relaxed controls and a version of Fan's fixed point theorem.

Proof of Theorem 5.1. By hypothesis, $Q \circ \hat{X}(\cdot)^{-1} = \mathbf{m}(\cdot)$ \mathbf{P} -almost surely for some deterministic $\mathbf{m} \in \mathcal{M}_2$. In view of Lemma 5.3, it is enough to show that the pair (Q_ω, \mathbf{m}) satisfies the optimality condition of Definition 4.3 with \mathbf{P} -probability one. This is equivalent to showing that $\hat{J}(\mathbf{m}(0), Q_\omega; \mathbf{m}) = \hat{V}(\mathbf{m}(0); \mathbf{m})$ for \mathbf{P} -almost all $\omega \in \Omega$.

Let $\varepsilon > 0$. Choose a function $\psi_\varepsilon^{\mathbf{m}} : [0, T] \times \mathbb{R}^d \times \mathcal{W} \rightarrow \Gamma$ and a probability measure $\Theta_\varepsilon^{\mathbf{m}} \in \mathcal{P}_2(\mathcal{Z})$ according to Lemma 4.3. Choose a sequence of indices $(i_*^n)_{n \in \mathbb{I}}$ according to condition (S). We will, as we may, assume that $i_*^n = 1$ for every $n \in \mathbb{I}$; otherwise renumber the components of the n -player games.

The proof proceeds in five steps. First, we construct a coupling for the initial conditions. In the second step, based on that coupling and the feedback function $\psi_\varepsilon^{\mathbf{m}}$, we define a competitor strategy $\tilde{\mathbf{u}}^n$ that differs from \mathbf{u}^n only in component one ($= i_*^n$). As verified in step three, the associated normalized occupation measures have the same limit Q as the sequence (Q^n) . This is used in the fourth step to show that $\limsup_{n \rightarrow \infty} J_1^n(\tilde{\mathbf{u}}^n) \leq \hat{V}(\mathbf{m}(0); \mathbf{m}) + \varepsilon$. Thanks to this upper limit, the local approximate Nash equilibrium property of \mathbf{u}^n together with condition (S), and the asymptotic lower bound on the average costs from Lemma 5.3, we establish optimality in the fifth and last step.

First step. By hypothesis, $(\mathbf{P}_n \circ (Q^n)^{-1})_{n \in \mathbb{I}}$ converges to $\mathbf{P} \circ Q^{-1}$ in $\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))$. By the choice of the metric on \mathcal{Z} , the continuity of the map $\mathcal{Z} \ni (\varphi, r, w) \mapsto \varphi(0) \in \mathbb{R}^d$, and the mapping theorem [for instance, Theorem 5.1 in Billingsley, 1968, p. 30], we have that

$$\mathcal{P}_2(\mathcal{Z}) \ni \Theta \mapsto \Theta \circ (\hat{X}(0))^{-1} \in \mathcal{P}_2(\mathbb{R}^d)$$

is continuous. This implies, again by the continuous mapping theorem, that

$$\mathbf{P}_n \circ \left(Q^n \circ (\hat{X}(0))^{-1} \right)^{-1} \xrightarrow{n \rightarrow \infty} \mathbf{P} \circ \left(Q \circ (\hat{X}(0))^{-1} \right)^{-1} \text{ in } \mathcal{P}(\mathcal{P}_2(\mathbb{R}^d)).$$

By construction and hypothesis, respectively,

$$Q^n \circ (\hat{X}(0))^{-1} = \frac{1}{n} \sum_{i=1}^n \delta_{\xi_i^n}, \quad \text{while} \quad \mathbf{P} \circ \left(Q \circ (\hat{X}(0))^{-1} \right)^{-1} = \delta_{\mathbf{m}(0)}.$$

It follows that $(\frac{1}{n} \sum_{i=1}^n \delta_{\xi_i^n})_{n \in \mathbb{I}}$ converges to $\mathbf{m}(0)$ in distribution as $\mathcal{P}_2(\mathbb{R}^d)$ -valued random variables, where $\mathbf{m}(0)$ is deterministic. This convergence

implies, in particular, that

$$\mathbf{E}_n \left[\frac{1}{n} \sum_{i=1}^n |\xi_i^n|^2 \right] \xrightarrow{n \rightarrow \infty} \int_{\mathbb{R}^d} |x|^2 \mathbf{m}(0)(dx).$$

By hypothesis, $\boldsymbol{\xi}^n = (\xi_1^n, \dots, \xi_n^n)$ is exchangeable for every $n \in \mathbb{I}$. Convergence of the associated empirical measures, by the Tanaka-Sznitman theorem [for instance, Theorem 3.2 in Gottlieb, 1998, p. 27], implies that

$$\mathbf{P}_n \circ (\xi_1^n)^{-1} \xrightarrow{n \rightarrow \infty} \mathbf{m}(0) \text{ in } \mathcal{P}(\mathbb{R}^d).$$

Actually, we have convergence in $\mathcal{P}_2(\mathbb{R}^d)$ since, by exchangeability,

$$\mathbf{E}_n [|\xi_1^n|^2] = \mathbf{E}_n \left[\frac{1}{n} \sum_{i=1}^n |\xi_i^n|^2 \right] \text{ for every } n \in \mathbb{I},$$

and the expectations on the right-hand side above converge to the second moment of $\mathbf{m}(0)$. We are therefore in the situation of Lemma 5.2, and we apply that result with the choice $i_*^n = 1$ to obtain a sequence $(\bar{\xi}^n)_{n \in \mathbb{I}}$ of \mathbb{R}^d -valued random variables such that $\bar{\xi}^n$ is $\sigma(\xi_{i_*^n}^n, \vartheta_{i_*^n}^n)$ -measurable, $\mathbf{P}_n \circ (\bar{\xi}^n)^{-1} = \mathbf{m}(0)$ and $\mathbf{E}_n [|\xi_1^n - \bar{\xi}^n|^2] \rightarrow 0$ as $n \rightarrow \infty$.

Second step. Define a strategy vector $\tilde{\mathbf{u}}^n = (\tilde{u}_1^n, \dots, \tilde{u}_n^n)$ by setting, for $(t, \omega) \in [0, T] \times \Omega_n$,

$$\tilde{u}_i^n(t, \omega) \doteq \begin{cases} \psi_\varepsilon^{\mathbf{m}}(t, \bar{\xi}^n(\omega), W_1^n(\cdot, \omega)) & \text{if } i = 1, \\ u_i^n(t, \omega) & \text{if } i \in \{2, \dots, n\}. \end{cases}$$

Notice that $\tilde{\mathbf{u}}^n$ is indeed a strategy vector for the game with n players. Moreover, $\tilde{u}_i^n = u_i^n$ for $i \in \{2, \dots, n\}$, while $\tilde{u}_1^n \in \mathcal{H}_2((\mathcal{F}_t^{n,1}), \mathbf{P}_n; \Gamma)$. Let $\tilde{\rho}^{n,i}$ be the relaxed control induced by \tilde{u}_i^n , $i \in \{1, \dots, n\}$. Clearly, $\tilde{\rho}^{n,i} = \rho^{n,i}$ for $i \geq 2$. On the other hand, by construction and since $\bar{\xi}^n$ and W_1^n are independent,

$$\mathbf{P}_n \circ (\bar{\xi}^n, \tilde{\rho}^{n,1}, W_1^n)^{-1} = \Theta_\varepsilon^{\mathbf{m}} \circ (\hat{X}(0), \hat{\rho}, \hat{W})^{-1} \text{ for every } n \in \mathbb{I}.$$

The law of \tilde{u}_1^n , in particular, does not change with n . It follows that

$$\sup_{n \in \mathbb{I}} \mathbf{E}_n \left[\int_0^T |\tilde{u}_1^n(t)|^2 dt \right] < \infty.$$

The coercivity assumption (A6) implies that there exists $C > 0$ such that for every $n \in \mathbb{I}$,

$$\mathbf{E}_n \left[\int_0^T |u_1^n(t)|^2 dt \right] \leq C (1 + J_1^n(\mathbf{u}^n)).$$

By choice of the index $i_*^n = 1$ according to (S), we have $\sup_{n \in \mathbb{N}} J_1^n(\mathbf{u}^n) < \infty$. Since $\mathbf{E}_n [|\xi_1^n|^2] = \frac{1}{n} \sum_{i=1}^n \mathbf{E}_n [|\xi_i^n|^2]$ by exchangeability, it follows that

$$(5.3) \quad \sup_{n \in \mathbb{I}} \mathbf{E}_n \left[|\xi_1^n|^2 + \int_0^T (|u_1^n(t)|^2 + |\tilde{u}_1^n(t)|^2) dt \right] < \infty.$$

Third step. Let $(\tilde{X}_1^n, \dots, \tilde{X}_n^n)$ be the solution of the system of equations (3.1) under strategy vector $\tilde{\mathbf{u}}^n$, and let $\tilde{\mu}^n$ denote the empirical measure process associated with $(\tilde{X}_1^n, \dots, \tilde{X}_n^n)$. Let \tilde{Q}^n be the normalized occupation measure associated with $\tilde{\mathbf{u}}^n$, i.e., the $\mathcal{P}_2(\mathcal{Z})$ -valued random variable determined by

$$\tilde{Q}_\omega^n(B \times R \times D) \doteq \frac{1}{n} \sum_{i=1}^n \delta_{\tilde{X}_i^n(\cdot, \omega)}(B) \cdot \delta_{\tilde{\rho}_\omega^{n,i}}(R) \cdot \delta_{W_i^n(\cdot, \omega)}(D), \quad \omega \in \Omega_n,$$

$B \in \mathcal{B}(\mathcal{X})$, $R \in \mathcal{B}(\mathcal{R}_2)$, $D \in \mathcal{B}(\mathcal{W})$. We are going to show that

$$(5.4) \quad \tilde{Q}^n \xrightarrow{n \rightarrow \infty} Q \text{ in distribution as } \mathcal{P}_2(\mathcal{Z})\text{-valued random variables.}$$

Since $Q^n \rightarrow Q$ in distribution, it suffices to show that

$$d_{\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))} \left(\mathbf{P}_n \circ (\tilde{Q}^n)^{-1}, \mathbf{P}_n \circ (Q^n)^{-1} \right) \xrightarrow{n \rightarrow \infty} 0.$$

Let $n \in \mathbb{I}$. By construction, definition of the bounded Lipschitz metric, inequality (2.1), and Hölder's inequality,

$$\begin{aligned} & d_{\mathcal{P}(\mathcal{P}_2(\mathcal{Z}))} \left(\mathbf{P}_n \circ (\tilde{Q}^n)^{-1}, \mathbf{P}_n \circ (Q^n)^{-1} \right) \\ &= \sup_{G \in \mathcal{C}(\mathcal{P}_2(\mathcal{Z})) : \|G\|_{\text{bLip}} \leq 1} \mathbf{E}_n \left[G(Q^n) - G(\tilde{Q}^n) \right] \\ &\leq \mathbf{E}_n \left[d_{\mathcal{P}_2(\mathcal{Z})}(Q^n, \tilde{Q}^n) \right] \\ &\leq \sqrt{\mathbf{E}_n \left[\frac{1}{n} \sum_{i=1}^n d_{\mathcal{Z}} \left((X_i^n, \rho^{n,i}, W_i^n), (\tilde{X}_i^n, \tilde{\rho}^{n,i}, W_i^n) \right)^2 \right]} \\ &\leq \frac{1}{\sqrt{n}} + \sqrt{\mathbf{E}_n \left[\frac{1}{n} \sum_{i=1}^n \sup_{t \in [0, T]} |X_i^n(t) - \tilde{X}_i^n(t)|^2 \right]}, \end{aligned}$$

where the last inequality follows by definition of $d_{\mathcal{Z}}$ and from the fact that $\rho^{n,i} = \tilde{\rho}^{n,i}$ for $i \in \{2, \dots, n\}$. Using assumption (A2), Hölder's inequality, Doob's maximal inequality, Itô's isometry, inequality (2.1), and Fubini's

theorem, we find that for $i \in \{2, \dots, n\}$, every $t \in [0, T]$,

$$\begin{aligned} & \mathbf{E}_n \left[\sup_{s \in [0, t]} |X_i^n(s) - \tilde{X}_i^n(s)|^2 \right] \\ & \leq 4(T+4)L^2 \mathbf{E}_n \left[\int_0^t |X_i^n(s) - \tilde{X}_i^n(s)|^2 ds + \int_0^t d_2(\mu^N(s), \tilde{\mu}^N(s))^2 ds \right] \\ & \leq 4(T+4)L^2 \int_0^t \mathbf{E}_n \left[|X_i^n(s) - \tilde{X}_i^n(s)|^2 + \frac{1}{n} \sum_{k=1}^n |X_k^n(s) - \tilde{X}_k^n(s)|^2 \right] ds. \end{aligned}$$

Similarly, but also using assumption (A3),

$$\begin{aligned} & \mathbf{E}_n \left[\sup_{s \in [0, t]} |X_1^n(s) - \tilde{X}_1^n(s)|^2 \right] \leq C_n \\ & \quad + 8(T+2)L^2 \int_0^t \mathbf{E}_n \left[|X_1^n(s) - \tilde{X}_1^n(s)|^2 + \frac{1}{n} \sum_{k=1}^n |X_k^n(s) - \tilde{X}_k^n(s)|^2 \right] ds, \end{aligned}$$

where C_n is equal to

$$80TK^2 \int_0^T \mathbf{E}_n \left[1 + |X_1^n(s)|^2 + |u_1^n(s)|^2 + |\tilde{u}_1^n(s)|^2 + \frac{1}{n} \sum_{k=1}^n |X_k^n(s)|^2 \right] ds.$$

It follows that, for every $t \in [0, T]$,

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^n \mathbf{E}_n \left[\sup_{s \in [0, t]} |X_i^n(s) - \tilde{X}_i^n(s)|^2 \right] \\ & \leq \frac{C_n}{n} + 8(T+4)L^2 \int_0^t \mathbf{E}_n \left[\frac{1}{n} \sum_{i=1}^n \sup_{\tilde{s} \in [0, s]} |X_i^n(\tilde{s}) - \tilde{X}_i^n(\tilde{s})|^2 \right] ds. \end{aligned}$$

Therefore, by Gronwall's lemma,

$$\mathbf{E}_n \left[\frac{1}{n} \sum_{i=1}^n \sup_{t \in [0, T]} |X_i^n(t) - \tilde{X}_i^n(t)|^2 \right] \leq \frac{C_n}{n} \exp(8T(T+4)L^2).$$

To complete the proof of (5.4), one checks that $\sup_{n \in \mathbb{I}} C_n < \infty$. But this is a consequence of (5.3), condition (T), and Lemma 3.1.

Fourth step. We are going to show that

$$(5.5) \quad \limsup_{n \rightarrow \infty} J_1^n(\tilde{\mathbf{u}}^n) \leq \hat{J}(\mathbf{m}(0), \Theta_\varepsilon^{\mathbf{m}}; \mathbf{m}).$$

Let $n \in \mathbb{I}$. Recall that \tilde{X}_1^n solves the equation

$$\begin{aligned} \tilde{X}_1^n(t) &= \xi_1^n + \int_0^t b\left(s, \tilde{X}_1^n(s), \tilde{\mu}^n(s), \tilde{u}_1^n(s)\right) ds \\ & \quad + \int_0^t \sigma\left(s, \tilde{X}_1^n(s), \tilde{\mu}^n(s)\right) dW_1^n(s), \quad t \in [0, T]. \end{aligned}$$

Let \bar{X}_1^n be the unique solution to

$$\begin{aligned}\bar{X}_1^n(t) &= \bar{\xi}^n + \int_0^t b(s, \bar{X}_1^n(s), \mathbf{m}(s), \tilde{u}_1^n(s)) ds \\ &\quad + \int_0^t \sigma(s, \bar{X}_1^n(s), \mathbf{m}(s)) dW_1^n(s), \quad t \in [0, T].\end{aligned}$$

Then, by uniqueness in law and construction, for every $n \in \mathbb{I}$,

$$\begin{aligned}\hat{J}(\mathbf{m}(0), \Theta_\varepsilon^{\mathbf{m}}; \mathbf{m}) \\ = \mathbf{E}_n \left[\int_0^T f(t, \bar{X}_1^n(t), \mathbf{m}(t), \tilde{u}_1^n(t)) dt + F(\bar{X}_1^n(T), \mathbf{m}(T)) \right].\end{aligned}$$

Using assumption (A2), Hölder's inequality, Itô's isometry, and Fubini's theorem, we find that for every $t \in [0, T]$,

$$\begin{aligned}\mathbf{E}_n \left[\left| \tilde{X}_1^n(t) - \bar{X}_1^n(t) \right|^2 \right] \\ \leq 3 \mathbf{E}_n \left[\left| \xi_1^n - \bar{\xi}^n \right|^2 \right] + 6(T+1)L^2 \mathbf{E}_n \left[\int_0^T d_2(\tilde{\mu}^n(s), \mathbf{m}(s))^2 ds \right] \\ + 6(T+1)L^2 \int_0^t \mathbf{E}_n \left[\left| \tilde{X}_1^n(s) - \bar{X}_1^n(s) \right|^2 \right] ds.\end{aligned}$$

The limit relation (5.4) implies that $(\tilde{\mu}^n(0))_{n \in \mathbb{I}}$ converges to $\mathbf{m}(0)$ in distribution as $\mathcal{P}_2(\mathbb{R}^d)$ -valued random variables and that

$$\sup_{t \in [0, T]} \mathbf{E}_n \left[d_2(\tilde{\mu}^n(t), \mathbf{m}(t))^2 \right] \xrightarrow{n \rightarrow \infty} 0.$$

By choice of the random variables $\bar{\xi}^n$ according to Lemma 5.2,

$$\mathbf{E}_n \left[\left| \xi_1^n - \bar{\xi}^n \right|^2 \right] \xrightarrow{n \rightarrow \infty} 0.$$

Therefore, by Gronwall's lemma,

$$\sup_{t \in [0, T]} \mathbf{E}_n \left[\left| \tilde{X}_1^n(t) - \bar{X}_1^n(t) \right|^2 \right] \xrightarrow{n \rightarrow \infty} 0.$$

Thanks to assumption (A4) and Hölder's inequality,

$$\begin{aligned}\left| J_1^n(\tilde{\mathbf{u}}^n) - \hat{J}(\mathbf{m}(0), \Theta_\varepsilon^{\mathbf{m}}; \mathbf{m}) \right| \\ \leq \mathbf{E}_n \left[\int_0^T \left| f(t, \tilde{X}_1^n(t), \tilde{\mu}^n(t), \tilde{u}_1^n(t)) - f(t, \bar{X}_1^n(t), \mathbf{m}(t), \tilde{u}_1^n(t)) \right| dt \right] \\ + \mathbf{E}_n \left[\left| F(\tilde{X}_1^n(T), \tilde{\mu}^n(T)) - F(\bar{X}_1^n(T), \mathbf{m}(T)) \right| \right], \\ \leq \sqrt{10}L(1 + \sqrt{T}) \sup_{t \in [0, T]} \mathbf{E}_n \left[\left| \tilde{X}_1^n(t) - \bar{X}_1^n(t) \right|^2 + d_2(\tilde{\mu}^n(t), \mathbf{m}(t))^2 \right]^{1/2} \\ \cdot \sup_{t \in [0, T]} \mathbf{E}_n \left[1 + \left| \tilde{X}_1^n(t) \right|^2 + \left| \bar{X}_1^n(t) \right|^2 + d_2(\tilde{\mu}^n(t), \delta_0)^2 + d_2(\mathbf{m}(t), \delta_0)^2 \right]^{1/2}.\end{aligned}$$

By (5.3) together with Lemma 3.1 and an analogous estimate applied to \bar{X}_1^n , and since $\sup_{t \in [0, T]} d_2(\mathbf{m}(t), \delta_0)^2 < \infty$ by continuity, we have

$$\sup_{n \in \mathbb{I}} \sup_{t \in [0, T]} \mathbf{E}_n \left[|\tilde{X}_1^n(t)|^2 + |\bar{X}_1^n(t)|^2 + d_2(\tilde{\mu}^n(t), \delta_0)^2 + d_2(\mathbf{m}(t), \delta_0)^2 \right] < \infty.$$

It follows that $J_1^n(\tilde{\mathbf{u}}^n) \rightarrow \hat{J}(\mathbf{m}(0), \Theta_\varepsilon^{\mathbf{m}}; \mathbf{m})$ as $n \rightarrow \infty$, which establishes (5.5).

Fifth step. The limit relation (5.5) and the choice of $\Theta_\varepsilon^{\mathbf{m}}$ imply that

$$\limsup_{j \rightarrow \infty} J_1^{N_j}(\tilde{\mathbf{u}}^{N_j}) \leq \hat{V}(\mathbf{m}(0); \mathbf{m}) + \varepsilon.$$

By hypothesis, \mathbf{u}^n is a local ε_n -Nash equilibrium. By construction, $\tilde{\mathbf{u}}^n$ differs from \mathbf{u}^n only in component number one ($= i_*^n$), and \tilde{u}_1^n is $(\mathcal{F}_t^{n,1})$ -adapted. Therefore,

$$J_1^n(\mathbf{u}^n) \leq J_1^n(\tilde{\mathbf{u}}^n) + \varepsilon_n.$$

By choice of the index $1 = i_*^n$ according to (S) and since $\varepsilon_n \rightarrow 0$ by hypothesis,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n J_i^n(\mathbf{u}^n) \leq \limsup_{n \rightarrow \infty} J_1^n(\mathbf{u}^n) \leq \limsup_{n \rightarrow \infty} J_1^n(\tilde{\mathbf{u}}^n).$$

It follows that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n J_i^n(\mathbf{u}^n) \leq \hat{V}(\mathbf{m}(0); \mathbf{m}) + \varepsilon.$$

On the other hand, thanks to the second part of Lemma 5.3,

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n J_i^n(\mathbf{u}^n) \geq \int_{\Omega} \hat{J}(\mathbf{m}(0), Q_\omega, \mathbf{m}) \mathbf{P}(d\omega).$$

It follows that

$$\int_{\Omega} \hat{J}(\mathbf{m}(0), Q_\omega, \mathbf{m}) \mathbf{P}(d\omega) \leq \hat{V}(\mathbf{m}(0); \mathbf{m}) + \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary and $\hat{J}(\mathbf{m}(0), Q_\omega, \mathbf{m}) \geq \hat{V}(\mathbf{m}(0); \mathbf{m})$ for every $\omega \in \Omega$ by definition of \hat{V} , we conclude that

$$\hat{J}(\mathbf{m}(0), Q_\omega, \mathbf{m}) = \hat{V}(\mathbf{m}(0); \mathbf{m}) \quad \text{for } \mathbf{P}\text{-almost all } \omega \in \Omega.$$

□

Remark 5.4. The proof of Theorem 5.1 gives some insight into why the assumption that the limit flow of measures \mathbf{m} is deterministic cannot simply be dropped. In the second step of the proof, we define a competitor

strategy \tilde{u}_1^n for the deviating player (player one after relabeling) in terms of the noise feedback function $\psi_\varepsilon^{\mathbf{m}}$. In general, for any $t \in [0, T]$, $\psi_\varepsilon^{\mathbf{m}}(t, \cdot, \cdot)$ depends on \mathbf{m} through its values for all times, not only through its values up to time t . Therefore, if \mathbf{m} were random, even taking for granted the measurable dependence of $\psi_\varepsilon^{\mathbf{m}}$ on \mathbf{m} , we might end up with a non-adapted competitor strategy. Indeed, the natural choice for \tilde{u}_1^n , namely $\tilde{u}_1^n(t, \omega) \doteq \psi_\varepsilon^{\mu_\omega^n(\cdot)}(t, \bar{\xi}^n(\omega), W_1^n(\cdot, \omega))$, would in general yield a Γ -valued process that would not be an admissible strategy for player one in the n -player game.

Appendix

A Proof of Lemma 4.1, second part

Let $\Theta \in \mathcal{P}(\mathcal{Z})$ be a solution of Eq. (4.2) with flow of measures \mathbf{m} in the sense of Definition 4.1. Using the local martingale property of $M_f^{\mathbf{m}}$ for f a monomial of first or second order as in the proof of Proposition 5.4.6 in Karatzas and Shreve [1991, pp. 315-316], we find that, under Θ and with respect to the filtration (\mathcal{G}_t) :

- \hat{W} is a d_1 -dimensional vector of continuous local martingales with $\hat{W}(0) = 0$ and quadratic covariations

$$\langle \hat{W}_l, \hat{W}_{\tilde{l}} \rangle(t) = t \cdot \delta_{l, \tilde{l}}, \quad l, \tilde{l} \in \{1, \dots, d_1\};$$

- $\bar{X} \doteq \hat{X} - \hat{X}(0) - \int_{\Gamma \times [0, \cdot]} b(s, \hat{X}(s), \mathbf{m}(s), \gamma) \hat{\rho}(d\gamma, ds)$ is a d -dimensional vector of continuous local martingales with quadratic covariations

$$\langle \bar{X}_j, \bar{X}_k \rangle(t) = \int_0^t (\sigma \sigma^\top)_{jk}(s, \hat{X}(s), \mathbf{m}(s)) ds, \quad j, k \in \{1, \dots, d\};$$

- \hat{W} , \bar{X} have quadratic covariations

$$\langle \bar{X}_k, \hat{W}_l \rangle(t) = \int_0^t \sigma_{kl}(s, \hat{X}(s), \mathbf{m}(s)) ds,$$

where $k \in \{1, \dots, d\}$, $l \in \{1, \dots, d_1\}$.

The local martingale property also holds with respect to the filtration $(\mathcal{G}_{t+}^\Theta)$; see the solution to Problem 5.4.13 in Karatzas and Shreve [1991, pp. 318-319, 392] and Remark 4.2 in Budhiraja et al. [2012]. By Lévy's characterization of Brownian motion [for instance, Theorem 3.3.16 in Karatzas and

Shreve, 1991, p. 157], we see that \hat{W} is a standard Wiener process with respect to $(\mathcal{G}_{t+}^\Theta)$. As a consequence, the process

$$Y(t) \doteq \int_0^t \sigma(s, \hat{X}(s), \mathbf{m}(s)) d\hat{W}(s), \quad t \in [0, T],$$

is well defined and a d -dimensional vector of continuous local martingales (under Θ with respect to $(\mathcal{G}_{t+}^\Theta)$) with quadratic covariations

$$\begin{aligned} \langle Y_j, Y_k \rangle(t) &= \int_0^t (\sigma \sigma^\top)_{jk}(s, \hat{X}(s), \mathbf{m}(s)) ds, & j, k \in \{1, \dots, d\}, \\ \langle Y_j, \hat{W}_l \rangle(t) &= \int_0^t \sigma_{jl}(s, \hat{X}(s), \mathbf{m}(s)) ds, & j \in \{1, \dots, d\}, l \in \{1, \dots, d_1\}. \end{aligned}$$

The quadratic covariations between the components of the vectors of continuous local martingales \bar{X}, Y are given by [cf. Proposition 3.2.24 in Karatzas and Shreve, 1991, p. 147]

$$\begin{aligned} \langle Y_j, \bar{X}_k \rangle(t) &= \sum_{l=1}^{d_1} \int_0^t \sigma_{jl}(s, \hat{X}(s), \mathbf{m}(s)) d\langle \bar{X}_k, \hat{W}_l \rangle(s) \\ &= \int_0^t (\sigma \sigma^\top)_{jk}(s, \hat{X}(s), \mathbf{m}(s)) ds, \quad j, k \in \{1, \dots, d\}. \end{aligned}$$

It follows that $\bar{X} - Y$ is a d -dimensional vector of continuous local martingales with $\bar{X}(0) = 0 = Y(0)$ and quadratic covariations

$$\langle \bar{X}_j - Y_j, \bar{X}_k - Y_k \rangle = \langle \bar{X}_j, \bar{X}_k \rangle - \langle Y_j, \bar{X}_k \rangle - \langle \bar{X}_j, Y_k \rangle + \langle Y_j, Y_k \rangle \equiv 0.$$

This implies [cf. Problem 1.5.12 in Karatzas and Shreve, 1991, p. 35] that $\bar{X} = Y$ Θ -almost surely, which establishes the solution property.

B Proof of Lemma 4.3

Fix $\mathbf{m} \in \mathcal{M}_2$, and set, for $(t, x, \gamma) \in [0, T] \times \mathbb{R}^d \times \Gamma$,

$$\begin{aligned} b_{\mathbf{m}}(t, x, \gamma) &\doteq b(t, x, \mathbf{m}(t), \gamma), & \sigma_{\mathbf{m}}(t, x) &\doteq \sigma(t, x, \mathbf{m}(t)), \\ f_{\mathbf{m}}(t, x, \gamma) &\doteq f(t, x, \mathbf{m}(t), \gamma), & F_{\mathbf{m}}(x) &\doteq F(x, \mathbf{m}(T)). \end{aligned}$$

Thanks to assumptions (A1), (A2), (A4), and the continuity of \mathbf{m} , we have that $b_{\mathbf{m}}, \sigma_{\mathbf{m}}, f_{\mathbf{m}}$ are continuous in the time and control variable, uniformly over compact subsets of \mathbb{R}^d , $b_{\mathbf{m}}, \sigma_{\mathbf{m}}$ are globally Lipschitz continuous in the state variable, uniformly in the other variables, and $f_{\mathbf{m}}, F_{\mathbf{m}}$ are locally

Lipschitz continuous in the state variable, uniformly in the other variables, with local Lipschitz constants that grow sublinearly in the state variable.

The function $\psi_\varepsilon^{\mathbf{m}}$ will be constructed based on the principle of dynamic programming applied in discrete time. To this end, we first introduce an original control problem corresponding to the minimal costs $\hat{V}(\cdot, \mathbf{m})$, then we build a sequence of approximating optimal control problems by successively restricting the set of admissible strategies. The proof proceeds in six steps.

First step. Let \mathcal{U} be the set of all quadruples $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), \rho, W)$ such that the pair $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t))$ forms a stochastic basis satisfying the usual hypotheses, W is a d_1 -dimensional (\mathcal{F}_t) -Wiener process, and ρ is an (\mathcal{F}_t) -adapted \mathcal{R}_2 -valued random variable such that $\mathbf{E} \left[\int_{\Gamma \times [0, T]} |\gamma|^2 \rho(d\gamma, ds) \right] < \infty$. For simplicity, we may write $\rho \in \mathcal{U}$ instead of $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), \rho, W) \in \mathcal{U}$. Given any $\rho \in \mathcal{U}$, $(t_0, x) \in [0, T] \times \mathbb{R}^d$, the stochastic integral equation

$$(B.1) \quad \begin{aligned} X(t) = x + \int_{\Gamma \times [0, t]} b_{\mathbf{m}}(t_0 + s, X(s), \gamma) \rho(d\gamma, ds) \\ + \int_0^t \sigma_{\mathbf{m}}(t_0 + s, X(s)) dW(s), \quad t \in [0, T - t_0], \end{aligned}$$

has a unique solution $X = X^{t_0, x, \rho}$, that is, X is the unique (up to indistinguishability with respect to \mathbf{P}) \mathbb{R}^d -valued (\mathcal{F}_t) -adapted continuous process that satisfies (B.1) with \mathbf{P} -probability one. Although the solution X of Eq. (B.1) starts in x at time zero, it corresponds to the solution of Eq. (4.2) starting in x at time t_0 . Define the costs associated with strategy ρ and initial condition $(t_0, x) \in [0, T] \times \mathbb{R}^d$ by

$$J_{\mathbf{m}}(t_0, x, \rho) \doteq \mathbf{E} \left[\int_{\Gamma \times [0, T - t_0]} f_{\mathbf{m}}(t_0 + s, X(s), \gamma) \rho(d\gamma, ds) + F_{\mathbf{m}}(X(T - t_0)) \right],$$

where $X = X^{t_0, x, \rho}$. The corresponding value function $V_{\mathbf{m}}$ is given by

$$V_{\mathbf{m}}(t, x) \doteq \inf_{\rho \in \mathcal{U}} J_{\mathbf{m}}(t, x, \rho),$$

which is well-defined as a measurable function $[0, T] \times \mathbb{R}^d \rightarrow [0, \infty)$. Actually, $V_{\mathbf{m}}$ is continuous. For $x \in \mathbb{R}^d$, $\rho \in \mathcal{U}$, set

$$\Theta^{x, \rho} \doteq \mathbf{P} \circ (X^{0, x, \rho}, \rho, W)^{-1}.$$

Then $\Theta^{x, \rho}$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} and

$$J_{\mathbf{m}}(0, x, \rho) = \hat{J}(\delta_x, \Theta^{x, \rho}; \mathbf{m}).$$

Conversely, in view of Lemma 4.1 and thanks to Assumption (A6), any $\Theta \in \mathcal{P}(\mathcal{Z})$ with $\hat{J}(\delta_x, \Theta; \mathbf{m}) < \infty$ induces a strategy $\rho \in \mathcal{U}$ such that $\Theta^{x, \rho} = \Theta$.

It follows that $V_{\mathbf{m}}(0, x) = \hat{V}(\delta_x; \mathbf{m})$ for every $x \in \mathbb{R}^d$ and, by conditioning on the initial state at time zero,

$$\int_{\mathbb{R}^d} V_{\mathbf{m}}(0, x) \mathbf{m}(0)(dx) = \hat{V}(\mathbf{m}(0); \mathbf{m}).$$

Second step. The function $V_{\mathbf{m}}(0, \cdot)$ is locally Lipschitz continuous. To be more precise, choose $c_0 > 0$, $\Gamma_0 \subset \Gamma$ according to (A6), and let $r_0 > 0$ be such that $\Gamma_0 \subset \{\gamma \in \mathbb{R}^{d_2} : |\gamma| \leq r_0\}$. We are going to show that there exists a constant $C_1 \in (0, \infty)$ depending only on K, L, T, \mathbf{m}, r_0 , and c_0 such that

$$(B.2) \quad |V_{\mathbf{m}}(0, x) - V_{\mathbf{m}}(0, \tilde{x})| \leq C_1 (1 + R) |x - \tilde{x}| \text{ whenever } |x| \vee |\tilde{x}| \leq R.$$

To establish (B.2), set, for $\varepsilon > 0$, $R > 0$,

$$\mathcal{U}_{\varepsilon, R} \doteq \{\rho \in \mathcal{U} : J_{\mathbf{m}}(0, x; \rho) \leq V_{\mathbf{m}}(0, x) + \varepsilon \text{ for some } x \text{ with } |x| \leq R\}.$$

Then for all $x, \tilde{x} \in \mathbb{R}^d$ with $|x| \vee |\tilde{x}| \leq R$,

$$|V_{\mathbf{m}}(0, x) - V_{\mathbf{m}}(0, \tilde{x})| \leq \inf_{\varepsilon > 0} \sup_{\rho \in \mathcal{U}_{\varepsilon, R}} |J_{\mathbf{m}}(0, x; \rho) - J_{\mathbf{m}}(0, \tilde{x}; \rho)|.$$

Let $x, \tilde{x} \in \mathbb{R}^d$, $\rho \in \mathcal{U}$, and let X, \tilde{X} be the solutions of (B.1) under ρ with initial state x and \tilde{x} , respectively. Using Hölder's inequality, Jensen's inequality, Itô's isometry, Fubini's theorem, assumption (A2), and Gronwall's lemma, we find that there exists a constant $C_{L, T}$ depending only on L, T such that

$$\sup_{t \in [0, T]} \mathbf{E} \left[\left| X(t) - \tilde{X}(t) \right|^2 \right] \leq C_{L, T} |x - \tilde{x}|.$$

Reusing the same tools but with assumption (A3) in place of (A2) (also cf. Lemma 3.1), we find that there exists a constant $C_{K, T, \mathbf{m}}$ depending only on K, T , and on \mathbf{m} (through $\sup_{t \in [0, T]} \int |y|^2 \mathbf{m}(t)(dy)$, which is finite since \mathbf{m} is continuous in time) such that

$$\sup_{t \in [0, T]} \mathbf{E} \left[|X(t)|^2 \right] \leq C_{K, T, \mathbf{m}} \left(1 + |x|^2 + \mathbf{E} \left[\int_{\Gamma \times [0, T]} |\gamma|^2 \rho(d\gamma, dt) \right] \right).$$

Thanks to the above estimates and assumption (A4), we have that there exist a constant $C_{L, T, \mathbf{m}}$ depending only on L, T , and \mathbf{m} , and a constant $C_{K, L, T, \mathbf{m}}$ depending only on K, L, T , and \mathbf{m} such that

$$\begin{aligned} & |J_{\mathbf{m}}(0, x; \rho) - J_{\mathbf{m}}(0, \tilde{x}; \rho)| \\ & \leq C_{L, T, \mathbf{m}} \left(1 + \sup_{t \in [0, T]} \sqrt{\mathbf{E} \left[|X(t)|^2 \right]} + \sup_{t \in [0, T]} \sqrt{\mathbf{E} \left[|\tilde{X}(t)|^2 \right]} \right) \cdot |x - \tilde{x}| \\ & \leq C_{K, L, T, \mathbf{m}} \left(1 + |x| \vee |\tilde{x}| + \sqrt{\mathbf{E} \left[\int_{\Gamma \times [0, T]} |\gamma|^2 \rho(d\gamma, dt) \right]} \right) \cdot |x - \tilde{x}|. \end{aligned}$$

It follows that for all $x, \tilde{x} \in \mathbb{R}^d$ with $|x| \vee |\tilde{x}| \leq R$,

$$\begin{aligned} & |V_{\mathbf{m}}(0, x) - V_{\mathbf{m}}(0, \tilde{x})| \\ & \leq C_{K,L,T,\mathbf{m}} \cdot \inf_{\varepsilon > 0} \left(1 + R + \sup_{\rho \in \mathcal{U}_{\varepsilon,R}} \sqrt{\mathbf{E} \left[\int_{\Gamma \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right]} \right) \cdot |x - \tilde{x}|. \end{aligned}$$

By the same estimates as above, but using (A5) instead of (A4), we find that there exists a constant $\tilde{C}_{K,T,\mathbf{m}}$ depending only on K, T, \mathbf{m} such that, for all $x \in \mathbb{R}^d$, all $\rho \in \mathcal{U}$,

$$J_{\mathbf{m}}(0, x; \rho) \leq \tilde{C}_{K,T,\mathbf{m}} \left(1 + |x|^2 + \mathbf{E} \left[\int_{\Gamma \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right] \right).$$

This implies that there exists a constant $C_{K,T,\mathbf{m},\Gamma}$ depending only on K, T, \mathbf{m} , and on Γ (through $\min_{\gamma \in \Gamma} |\gamma|^2$) such that, for all $x \in \mathbb{R}^d$,

$$V_{\mathbf{m}}(0, x) \leq C_{K,T,\mathbf{m},\Gamma} (1 + |x|^2).$$

Let $\rho \in \mathcal{U}_{\varepsilon,R}$ for some $\varepsilon > 0$. Choose $x \in \mathbb{R}^d$ with $|x| \leq R$ such that $J_{\mathbf{m}}(0, x; \rho) \leq V_{\mathbf{m}}(0, x) + \varepsilon$ (possible by definition of $\mathcal{U}_{\varepsilon,R}$). By the coercivity assumption (A6),

$$J_{\mathbf{m}}(0, x; \rho) \geq c_0 \mathbf{E} \left[\int_{(\Gamma \setminus \Gamma_0) \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right],$$

hence

$$c_0 \mathbf{E} \left[\int_{(\Gamma \setminus \Gamma_0) \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right] \leq C_{K,T,\mathbf{m},\Gamma} (1 + R^2) + \varepsilon.$$

By construction,

$$\mathbf{E} \left[\int_{\Gamma \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right] \leq T \cdot r_0^2 + \mathbf{E} \left[\int_{(\Gamma \setminus \Gamma_0) \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right].$$

It follows that there exists a constant $C_{K,T,\mathbf{m},c_0,r_0}$ depending only on K, T, \mathbf{m}, c_0 , and on r_0 (clearly, $\min_{\gamma \in \Gamma} |\gamma|^2 \leq r_0^2$) such that

$$\sup_{\rho \in \mathcal{U}_{\varepsilon,R}} \sqrt{\mathbf{E} \left[\int_{\Gamma \times [0,T]} |\gamma|^2 \rho(d\gamma, dt) \right]} \leq C_{K,T,\mathbf{m},c_0,r_0} (1 + R + \sqrt{\varepsilon}).$$

This establishes (B.2).

Third Step. For $M \in \mathbb{N}$, set $\Gamma_M \doteq \{\gamma \in \Gamma : |\gamma| \leq M\}$. For M big enough, say $M \geq M_0$, Γ_M is non-empty. Choose $\gamma_0 \in \Gamma_{M_0}$, and set $\Gamma_M \doteq \{\gamma_0\}$ if $M < M_0$. Then, for every $M \in \mathbb{N}$, Γ_M is compact (and non-empty) and $\Gamma_M \subset \Gamma_{M+1}$. Set

$$\mathcal{U}_M \doteq \{\rho \in \mathcal{U} : \rho(\Gamma_M \times [0, T]) = T \text{ } \mathbf{P}\text{-almost surely}\},$$

and let $V_{\mathfrak{m}, M}$ be the value function defined with respect to \mathcal{U}_M instead of \mathcal{U} . We claim that

$$(B.3) \quad V_{\mathfrak{m}, M}(0, \cdot) \xrightarrow{M \rightarrow \infty} V_{\mathfrak{m}}(0, \cdot) \text{ uniformly over compact subsets of } \mathbb{R}^d.$$

Notice that, by construction, $V_{\mathfrak{m}, M}(0, \cdot) \geq V_{\mathfrak{m}, M+1}(0, \cdot) \geq V_{\mathfrak{m}}(0, \cdot)$ for every $M \in \mathbb{N}$. By Step 2, we know that $V_{\mathfrak{m}}(0, \cdot)$ is locally Lipschitz. Repeating the arguments of Step 2 (notice that $\mathcal{U}_M \subset \mathcal{U}$ by definition), we find that inequality (B.2) also holds for $V_{\mathfrak{m}, M}(0, \cdot)$ in place of $V_{\mathfrak{m}}(0, \cdot)$ and that the constant C_1 can be chosen independently of $M \in \mathbb{N}$. To establish (B.3), it is therefore enough to check that point-wise convergence holds. Fix $x \in \mathbb{R}^d$. It suffices to show that given $\rho \in \mathcal{U}$ there exists a sequence $(\rho^{(M)}) \subset \mathcal{U}$ such that $\rho^{(M)} \in \mathcal{U}_M$ for every M and $J_{\mathfrak{m}}(0, x; \rho^{(M)}) \rightarrow J_{\mathfrak{m}}(0, x; \rho)$ as $M \rightarrow \infty$.

Let $\rho \in \mathcal{U}$. For $M \in \mathbb{N}$, let $\rho^{(M)} \in \mathcal{U}_M$ be such that for every $B \in \mathcal{B}(\Gamma)$, every $I \in \mathcal{B}([0, T])$,

$$\rho^{(M)}(B \times I) = \rho((B \cap \Gamma_M) \times I) + \rho((\Gamma \setminus \Gamma_M) \times I) \cdot \delta_{\gamma_0}(B).$$

This determines a unique strategy $\rho^{(M)} \in \mathcal{U}_M$. Clearly, $\rho^{(M)}$ comes with the same stochastic basis as ρ . If $(\dot{\rho}_t)$ is a version of the time derivative process associated with ρ (thus, $\rho(d\gamma, dt) = \dot{\rho}_t(d\gamma)dt$), then a version of the time derivative process of $\rho^{(M)}$ is given by

$$\dot{\rho}_t^{(M)}(d\gamma) = \mathbf{1}_{\Gamma_M}(\gamma) \cdot \rho_t(d\gamma) + \rho_t(\Gamma \setminus \Gamma_M) \cdot \delta_{\gamma_0}(d\gamma).$$

Let $X, X^{(M)}$ be the solutions of (B.1) under ρ and $\rho^{(M)}$, respectively. Thanks to Hölder's inequality, Jensen's inequality, Itô's isometry, Fubini's theorem, and assumption (A2), there exists a constant $C_{L,T}$ depending only on L, T such that, for every $t \in [0, T]$,

$$\begin{aligned} \mathbf{E} \left[\left| X(t) - X^{(M)}(t) \right|^2 \right] &\leq C_{L,T} \int_0^t \mathbf{E} \left[\left| X(s) - X^{(M)}(s) \right|^2 \right] ds \\ &\quad + C_{L,T} \mathbf{E} \left[\left| \int_{\Gamma \times [0, t]} b_{\mathfrak{m}}(s, X(s), \gamma) \left(\rho^{(M)} - \rho \right) (d\gamma, ds) \right|^2 \right]. \end{aligned}$$

Using the definition of $\rho^{(M)}$, Hölder's inequality and assumption (A3), we find that, for some constant $C_{K,T,m}$ depending only on K, T and \mathbf{m} ,

$$\begin{aligned}
& \mathbf{E} \left[\left| \int_{\Gamma \times [0,t]} b_{\mathbf{m}}(s, X(s), \gamma) (\rho^{(M)} - \rho) (d\gamma, ds) \right|^2 \right] \\
& \leq 2T \mathbf{E} \left[\int_0^T \int_{\Gamma \setminus \Gamma_M} |b_{\mathbf{m}}(s, X(s), \gamma)|^2 \dot{\rho}_s(d\gamma) ds \right] \\
& \quad + 2 \mathbf{E} \left[\rho((\Gamma \setminus \Gamma_M) \times [0, T]) \cdot \int_0^T |b_{\mathbf{m}}(s, X(s), \gamma_0)|^2 ds \right] \\
& \leq C_{K,T,m} \mathbf{E} \left[\rho((\Gamma \setminus \Gamma_M) \times [0, T]) \cdot \left(1 + \sup_{r \in [0, T]} |X(r)|^2 \right) \right] \\
& \quad + C_{K,T,m} \mathbf{E} \left[\int_{\Gamma \times [0, T]} \mathbf{1}_{\Gamma \setminus \Gamma_M}(\gamma) \cdot |\gamma|^2 \rho(d\gamma, ds) \right].
\end{aligned}$$

By (A3) and the usual estimates, including Gronwall's lemma, we have $\mathbf{E} \left[\sup_{r \in [0, T]} |X(r)|^2 \right] < \infty$. Since ρ_ω is a measure with total mass T for every $\omega \in \Omega$, we have $\rho((\Gamma \setminus \Gamma_M) \times [0, T]) \rightarrow 0$ as $M \rightarrow \infty$ \mathbf{P} -almost surely. This implies, by dominated convergence,

$$\mathbf{E} \left[\rho((\Gamma \setminus \Gamma_M) \times [0, T]) \cdot \left(1 + \sup_{r \in [0, T]} |X(r)|^2 \right) \right] \xrightarrow{M \rightarrow \infty} 0.$$

On the other hand, $\mathbf{E} \left[\int_{\Gamma \times [0, T]} |\gamma|^2 \rho(d\gamma, ds) \right] < \infty$ by definition of \mathcal{U} . This means that

$$\mathbf{E} \left[\int_{\Gamma \times [0, T]} \mathbf{1}_{\Gamma \setminus \Gamma_M}(\gamma) \cdot |\gamma|^2 \rho(d\gamma, ds) \right] \xrightarrow{M \rightarrow \infty} 0.$$

An application of Gronwall's lemma now yields

$$\mathbf{E} \left[|X(t) - X^{(M)}(t)|^2 \right] \xrightarrow{M \rightarrow \infty} 0.$$

This convergence together with assumption (A5) (and an estimate completely analogous to the one above) implies that $J_{\mathbf{m}}(0, x; \rho^{(M)}) \rightarrow J_{\mathbf{m}}(0, x; \rho)$ as $M \rightarrow \infty$.

Fourth Step. Choose a family $(\Gamma_{M,k})_{M,k \in \mathbb{N}}$ of finite subsets of Γ such that $\Gamma_{M,k} \subset \Gamma_{M,k+1} \subset \Gamma_M$, $\Gamma_{M,k} \subset \Gamma_{M+1,k}$, and $\min_{\tilde{\gamma} \in \Gamma_{M,k}} |\gamma - \tilde{\gamma}| \leq 1/k$ for any $\gamma \in \Gamma_M$. Let $\mathcal{U}_{M,k}$ be the set of all $\rho \in \mathcal{U}$ such that ρ is the \mathcal{R}_2 -valued random variable induced by a $\Gamma_{M,k}$ -valued adapted process that is piece-wise

constant in time with respect to the equidistant grid of step size $T \cdot 2^{-k}$. Thus, $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), \rho, W) \in \mathcal{U}_{M,k}$ if and only if $\rho_\omega(d\gamma, dt) = \delta_{u(t,\omega)}(d\gamma)dt$ for \mathbf{P} -almost every $\omega \in \Omega$, where u is a $\Gamma_{M,k}$ -valued (\mathcal{F}_t) -progressively measurable process with càdlàg trajectories that are piece-wise constant over the grid $\{0, T \cdot 2^{-k}, 2T \cdot 2^{-k}, 3T \cdot 2^{-k}, \dots, T\}$. Let $V_{m,M,k}$ be the value function defined with respect to $\mathcal{U}_{M,k}$. Then, thanks to the continuity in time and control of the coefficients according to (A1), a version of the chattering lemma [for instance, Theorem 3.5.2 in Kushner, 1990, p. 59], and the local Lipschitz continuity of $V_{m,M,k}(0, \cdot)$, which holds uniformly in k and M (one repeats the arguments of Step 2), we find that

$$V_{m,M,k}(0, \cdot) \xrightarrow{k \rightarrow \infty} V_{m,M}(0, \cdot) \text{ uniformly over compact subsets of } \mathbb{R}^d.$$

By (B.3) and since $\mathcal{U}_{M,k} \subset \mathcal{U}_{M,k+1} \subset \mathcal{U}_M$ and $\mathcal{U}_{M,k} \subset \mathcal{U}_{M+1,k}$, it follows that

$$(B.4) \quad V_{m,M,M}(0, \cdot) \xrightarrow{M \rightarrow \infty} V_m(0, \cdot) \text{ uniformly over compact subsets of } \mathbb{R}^d.$$

Fifth step. The value function $V_{m,M,k}$ coincides with the value function of a discrete-time optimal control problem defined as follows. Set $h \doteq T \cdot 2^{-k}$. Thanks to Theorem 1 in Kallenberg [1996] and because $\Gamma_{M,k}$ is finite, we find a measurable and universally predictable function

$$\Phi_{m,M,k}: \mathbb{N}_0 \times \mathbb{R}^d \times \Gamma_{M,k} \times \mathbf{C}([0, h], \mathbb{R}^{d_1}) \rightarrow \mathbb{R}^d$$

such that $\Phi_{m,M,k}(j, x, \gamma, W) = X((j+1)h)$ \mathbf{P} -almost surely whenever X is the unique strong solution to

$$\begin{aligned} X(t) &= x + \int_0^t b_m(j \cdot h + s, X(s), \gamma) ds \\ &\quad + \int_0^t \sigma_m(j \cdot h + s, X(s)) dW(s), \quad t \in [0, h], \end{aligned}$$

where W is a d_1 -dimensional standard Wiener process defined on some stochastic basis $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t))$. The function $\Phi_{m,M,k}$ is the system function of the control problem in the sense of Bertsekas and Shreve [1996]. Let $\bar{\mathcal{U}}_{M,k}$ denote the set of discrete-time Markov feedback strategies with values in $\Gamma_{M,k}$, that is, the set of all Borel measurable functions $v: \mathbb{N}_0 \times \mathbb{R}^d \rightarrow \Gamma_{M,k}$. To describe the path-wise evolution of the system, choose a complete probability space $(\Omega_\circ, \mathcal{F}^\circ, \mathbf{P}_\circ)$ rich enough to carry a d_1 -dimensional standard Wiener process W_\circ . For $j \in \mathbb{N}_0$, set $\zeta_j \doteq (W(jh+s) - W(jh))_{s \in [0, h]}$, which defines a $\mathbf{C}([0, h], \mathbb{R}^{d_1})$ -valued random variable. Given any Markov feedback

strategy $v \in \bar{\mathcal{U}}_{M,k}$ and initial condition $(j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d$, the corresponding state sequence is defined recursively, for each $\omega \in \Omega_\circ$, by

$$(B.5) \quad \begin{aligned} X_0(\omega) &\doteq x, \\ X_{l+1}(\omega) &\doteq \Phi_{\mathbf{m},M,k}(j+l, X_l(\omega), v(j+l, X_l(\omega)), \zeta_l(\omega)), \end{aligned}$$

$l \in \{0, \dots, 2^k - j - 1\}$. The associated costs are given by

$$\bar{J}_{\mathbf{m},M,k}(j, x, v) \doteq \mathbf{E}_\circ \left[\sum_{l=0}^{2^k-j-1} f_{\mathbf{m}}((j+l)h, X_l, v(j+l, X_l)) \cdot h + F_{\mathbf{m}}(X_{k-j}) \right],$$

where (X_l) is the state sequence generated according to (B.5) with feedback strategy v and initial condition (j, x) . Let $\bar{V}_{\mathbf{m},M,k}$ be the value function of the control problem just defined:

$$\bar{V}_{\mathbf{m},M,k}(j, x) \doteq \inf_{v \in \bar{\mathcal{U}}_{M,k}} \bar{J}_{\mathbf{m},M,k}(j, x, v), \quad (j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d.$$

By Proposition 8.6 in Bertsekas and Shreve [1996, pp. 209-210], the principle of dynamic programming applies to $\bar{V}_{\mathbf{m},M,k}$. This has two consequences. First, notice that any feedback strategy $v \in \bar{\mathcal{U}}_{M,k}$ induces, for any given initial condition $(j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d$, a relaxed control variable $\rho \in \mathcal{U}_{M,k}$ such that

$$\bar{J}_{\mathbf{m},M,k}(j, x, v) = J_{\mathbf{m}}(jh, x, \rho).$$

This implies $\bar{V}_{\mathbf{m},M,k}(j, x) \geq V_{\mathbf{m},M,k}(jh, x)$ for all $(j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d$. Since $\bar{V}_{\mathbf{m},M,k}(2^k, \cdot) = F_{\mathbf{m}}(\cdot) = V_{\mathbf{m},M,k}(2^k h, \cdot)$, it follows by dynamic programming for $\bar{V}_{\mathbf{m},M,k}$ and backward induction that

$$\bar{V}_{\mathbf{m},M,k}(j, x) = V_{\mathbf{m},M,k}(jh, x) \quad \text{for all } (j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d.$$

As a second consequence of the principle of dynamic programming, there exists an optimal Markov feedback strategy. More precisely, we can choose $v_* \in \bar{\mathcal{U}}_{M,k}$ such that, for every $(j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d$,

$$\begin{aligned} v_*(j, x) \in \operatorname{argmin}_{\gamma \in \Gamma_{M,k}} & \left\{ f_{\mathbf{m}}(jh, x, \gamma) \cdot h \right. \\ & \left. + \int_{\mathbf{C}([0,h], \mathbb{R}^{d_1})} \bar{V}_{\mathbf{m},M,k}(j+1, \Phi_{\mathbf{m},M,k}(j, x, \gamma, y)) \eta_h(dy) \right\}, \end{aligned}$$

where η_h is standard Wiener measure on $\mathcal{B}(\mathbf{C}([0, h], \mathbb{R}^{d_1}))$. Then

$$\bar{J}_{\mathbf{m},M,k}(j, x, v_*) = \bar{V}_{\mathbf{m},M,k}(j, x) \quad \text{for all } (j, x) \in \{0, \dots, 2^k\} \times \mathbb{R}^d.$$

Sixth step. Define a function $\psi_{M,k}^{\mathbf{m}} : [0, T] \times \mathbb{R}^d \times \mathcal{W} \rightarrow \Gamma_{M,k}$ as follows. Let $x \in \mathbb{R}^d$, $w \in \mathcal{W}$. In analogy with (B.5), recursively define a sequence $(x_j)_{j \in \{0, \dots, 2^k\}}$ by

$$x_0 \doteq x, \quad x_{j+1} \doteq \Phi_{\mathbf{m}, M, k}(j, x_j, v_*(j, x_j), (w(jh + s) - w(jh))_{s \in [0, h]}).$$

For $j \in \{0, \dots, 2^k - 1\}$, $s \in [0, h]$, set

$$\psi_{M,k}^{\mathbf{m}}(jh + s, x, w) \doteq v_*(j, x_j),$$

and set $\psi_{M,k}^{\mathbf{m}}(T, x, w) \doteq v_*(2^k, x_k)$. By construction, $\psi_{M,k}^{\mathbf{m}}$ is progressively measurable with values in a finite set. Let $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t))$ be a stochastic basis rich enough to carry a d_1 -dimensional (\mathcal{F}_t) -Wiener process W and an \mathbb{R}^d -valued \mathcal{F}_0 -measurable random variable ξ such that $\mathbf{P} \circ \xi^{-1} = \mathbf{m}(0)$. For every $x \in \mathbb{R}^d$, the process $\psi_{M,k}^{\mathbf{m}}(t, x, W)$ induces a relaxed control random variable ρ such that $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), \rho, W) \in \mathcal{U}_{M,k}$ and $J_{\mathbf{m}}(0, x, \rho) = V_{\mathbf{m}, M, k}(0, x)$. Let $\rho^{M,k}$ be the relaxed control random variable in $\mathcal{U}_{M,k}$ induced by the process $\psi_{M,k}^{\mathbf{m}}(t, \xi, W)$. Let $X_{M,k}$ be the unique continuous (\mathcal{F}_t) -adapted process such that $X_{M,k}(0) = \xi$ and $((\Omega, \mathcal{F}, \mathbf{P}), (\mathcal{F}_t), X_{M,k}, \rho^{M,k}, W)$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} . Set

$$\Theta_{M,k}^{\mathbf{m}} \doteq \mathbf{P} \circ (X_{M,k}, \rho^{M,k}, W)^{-1}.$$

Then $\Theta_{M,k}^{\mathbf{m}} \in \mathcal{P}_2(\mathcal{Z})$ and $\Theta_{M,k}^{\mathbf{m}}$ is a solution of Eq. (4.2) with flow of measures \mathbf{m} such that $\Theta_{M,k}^{\mathbf{m}} \circ (\hat{X}(0))^{-1} = \mathbf{m}(0)$, $\hat{\rho}(d\gamma, dt) = \delta_{\psi_{M,k}^{\mathbf{m}}(t, \hat{X}(0), \hat{W})}(d\gamma)dt$ with probability one under $\Theta_{M,k}^{\mathbf{m}}$, and

$$\hat{J}(\mathbf{m}(0), \Theta_{M,k}^{\mathbf{m}}; \mathbf{m}) = \int_{\mathbb{R}^d} V_{\mathbf{m}, M, k}(0, x) \mathbf{m}(0)(dx) < \infty.$$

By (B.4) and dominated convergence, it follows that

$$\hat{J}(\mathbf{m}(0), \Theta_{M,M}^{\mathbf{m}}; \mathbf{m}) \xrightarrow{M \rightarrow \infty} \hat{V}(\mathbf{m}(0); \mathbf{m}).$$

Hence, given any $\varepsilon > 0$, there exists $M(\varepsilon) \in \mathbb{N}$ such that, for all $M \geq M(\varepsilon)$, $\hat{J}(\mathbf{m}(0), \Theta_{M,M}^{\mathbf{m}}; \mathbf{m}) \leq \hat{V}(\mathbf{m}(0); \mathbf{m}) + \varepsilon$. This ends the proof.

C Tightness functions

Let \mathcal{S} be a Polish space. A function $g : \mathcal{S} \rightarrow [0, \infty]$ is called a *tightness function* on \mathcal{S} if it is measurable and its sublevel sets $\{s \in \mathcal{S} : g(s) \leq c\}$ are pre-compact in \mathcal{S} for all $c \in [0, \infty)$. If g is a tightness function on \mathcal{S} , then the function $\mathcal{P}(\mathcal{S}) \ni \Theta \mapsto \int_{\mathcal{S}} g(s) \Theta(ds) \in [0, \infty]$ is a tightness function on $\mathcal{P}(\mathcal{S})$; see, for instance, Theorem A.3.17 in Dupuis and Ellis [1997, p. 309].

C.1 A tightness function on \mathcal{R}_2

Let $\delta_0 > 0$. Define a function $\tilde{g}: \mathcal{R}_2 \rightarrow [0, \infty]$ by

$$\tilde{g}(r) \doteq \int_{\Gamma \times [0, T]} |\gamma|^{2+\delta_0} r(d\gamma, dt).$$

We check that \tilde{g} is a tightness function on \mathcal{R}_2 . By construction, \tilde{g} is measurable. For $c \in [0, \infty)$, set

$$A_c \doteq \{r \in \mathcal{R}_2 : \tilde{g}(r) \leq c\}.$$

Fix $c \in [0, \infty)$. Then we have to show that A_c is pre-compact in \mathcal{R}_2 . This is equivalent to showing that

- a) A_c is pre-compact in \mathcal{R} ,
- b) if $(r_n)_{n \in \mathbb{N}} \subset A_c$ is such that $r_n \rightarrow r$ in \mathcal{R} for some $r \in \mathcal{R}$, then $r \in \mathcal{R}_2$ and $\int_{\Gamma \times [0, T]} |\gamma|^2 r_n(d\gamma, dt) \rightarrow \int_{\Gamma \times [0, T]} |\gamma|^2 r(d\gamma, dt)$ as $n \rightarrow \infty$.

Pre-compactness of A_c in \mathcal{R} is equivalent to tightness of A_c . This holds since, for every $M > 0$, the set $\{\gamma \in \Gamma : |\gamma| \leq M\}$ is compact (by assumption (A6), Γ is closed) and, by Markov's inequality,

$$\sup_{r \in A_c} r\{(\gamma, t) \in \Gamma \times [0, T] : |\gamma| > M\} \leq \frac{1}{M^{2+\delta_0}} \cdot \sup_{r \in A_c} \tilde{g}(r) \leq \frac{c}{M^{2+\delta_0}},$$

which tends to zero as $M \rightarrow \infty$.

As to the convergence of moments, let $(r_n)_{n \in \mathbb{N}} \subset A_c$ be such that $r_n \rightarrow r$ in \mathcal{R} for some $r \in \mathcal{R}$. Then, by Fatou's lemma and Hölder's inequality,

$$\liminf_{n \rightarrow \infty} \int_{\Gamma \times [0, T]} |\gamma|^2 r_n(d\gamma, dt) \geq \int_{\Gamma \times [0, T]} |\gamma|^2 r(d\gamma, dt),$$

hence $r \in A_c \subset \mathcal{R}_2$. By convergence in \mathcal{R} , we have, for every $M > 0$,

$$\lim_{n \rightarrow \infty} \int_{\Gamma \times [0, T]} |\gamma|^2 \wedge M r_n(d\gamma, dt) = \int_{\Gamma \times [0, T]} |\gamma|^2 \wedge M r(d\gamma, dt).$$

On the other hand, again by Hölder's and Markov's inequality, for every $n \in \mathbb{N}$, every $M > 0$,

$$\begin{aligned} & \int_{\Gamma \times [0, T]} |\gamma|^2 \cdot \mathbf{1}_{[M, \infty)}(|\gamma|^2) r_n(d\gamma, dt) \\ & \leq \left(\int_{\Gamma \times [0, T]} |\gamma|^{2+\delta_0} r_n(d\gamma, dt) \right)^{\frac{1}{2+\delta_0}} \cdot r_n\{(\gamma, t) \in \Gamma \times [0, T] : |\gamma|^2 > M\}^{\frac{1+\delta_0}{2+\delta_0}} \\ & \leq c^{\frac{1}{2+\delta_0}} \cdot c^{\frac{1+\delta_0}{2+\delta_0}} \cdot M^{-(1+\delta_0/2)}. \end{aligned}$$

It follows that

$$\sup_{n \in \mathbb{N}} \int_{\Gamma \times [0, T]} |\gamma|^2 \cdot \mathbf{1}_{[M, \infty)} (|\gamma|^2) r_n(d\gamma, dt) \xrightarrow{M \rightarrow \infty} 0,$$

hence $\lim_{n \rightarrow \infty} \int_{\Gamma \times [0, T]} |\gamma|^2 r_n(d\gamma, dt) = \int_{\Gamma \times [0, T]} |\gamma|^2 r(d\gamma, dt)$.

C.2 A tightness function on $\mathcal{P}_2(\mathcal{Z})$

We check that the function g defined by (5.2) is a tightness function on $\mathcal{P}_2(\mathcal{Z})$. By construction, g is measurable (by continuity, the suprema appearing inside the second integral and in the definition of the modulus of continuity can be restricted to countable index sets). Thus we have to show that, given any $c \in [0, \infty)$, the set

$$A(c) \doteq \{\Theta \in \mathcal{P}_2(\mathcal{Z}) : g(\Theta) \leq c\}$$

is pre-compact in $\mathcal{P}_2(\mathcal{Z})$. Fix $c \in [0, \infty)$. The pre-compactness of $A(c)$ in $\mathcal{P}_2(\mathcal{Z})$ is equivalent to the following two conditions:

- a) $A(c)$ is tight in $\mathcal{P}(\mathcal{Z})$;
- b) if $(\Theta^n)_{n \in \mathbb{N}} \subset A(c)$ is such that Θ^n converges to $\bar{\Theta}$ in $\mathcal{P}(\mathcal{Z})$ for some $\bar{\Theta} \in \mathcal{P}(\mathcal{Z})$, then $\bar{\Theta} \in \mathcal{P}_2(\mathcal{Z})$ and $\int_{\mathcal{Z}} d_{\mathcal{Z}}(s, s_0)^2 \Theta^n(ds) \rightarrow \int_{\mathcal{Z}} d_{\mathcal{Z}}(s, s_0)^2 \bar{\Theta}(ds)$, where s_0 is some arbitrarily fixed element of \mathcal{Z} .

To verify a), it is enough to check tightness of marginals, that is, to verify that $A_{\mathcal{X}}(c) \doteq \{[\Theta]_{\mathcal{X}} : \Theta \in A(c)\}$ is tight in $\mathcal{P}(\mathcal{X})$, $A_{\mathcal{R}_2}(c) \doteq \{[\Theta]_{\mathcal{R}_2} : \Theta \in A(c)\}$ is tight in $\mathcal{P}(\mathcal{R}_2)$, and $A_{\mathcal{W}}(c) \doteq \{[\Theta]_{\mathcal{W}} : \Theta \in A(c)\}$ is tight in $\mathcal{P}(\mathcal{W})$, where $[\Theta]_{\mathcal{X}}$, $[\Theta]_{\mathcal{R}_2}$, $[\Theta]_{\mathcal{W}}$ denote the marginal distributions of Θ on \mathcal{X} , \mathcal{R}_2 , and \mathcal{W} , respectively. Thanks to Markov's inequality and the Ascoli-Arzelà criterion [for instance, Theorem 8.2 in Billingsley, 1968, p. 55], $A_{\mathcal{X}}(c)$, $A_{\mathcal{W}}(c)$ are tight in $\mathcal{P}(\mathcal{X})$ and $\mathcal{P}(\mathcal{W})$, respectively. The tightness of $A_{\mathcal{R}_2}(c)$ in $\mathcal{P}(\mathcal{R}_2)$ follows from the fact that the mapping

$$\mathcal{R}_2 \ni r \mapsto \int_{\Gamma \times [0, T]} |\gamma|^{2+\delta_0} r(d\gamma, dt) \in [0, \infty]$$

is a tightness function on \mathcal{R}_2 ; see Appendix C.1.

In order to check b), let $(\Theta^n)_{n \in \mathbb{N}} \subset A(c)$ be such that Θ^n converges to $\bar{\Theta}$ in $\mathcal{P}(\mathcal{Z})$ for some $\bar{\Theta} \in \mathcal{P}(\mathcal{Z})$. By a version of Fatou's lemma [cf. Theorem A.3.12 Dupuis and Ellis, 1997, p. 307],

$$\liminf_{n \rightarrow \infty} \int_{\mathcal{Z}} \|\varphi\|_{\mathcal{X}}^{2+\delta_0} \Theta^n(d\varphi, dr, dw) \geq \int_{\mathcal{Z}} \|\varphi\|_{\mathcal{X}}^{2+\delta_0} \bar{\Theta}(d\varphi, dr, dw).$$

By definition of $d_{\mathcal{Z}}$ and of g , and thanks to Hölder's inequality, it follows that $\Theta \in \mathcal{P}_2(\mathcal{Z})$. By convergence of $(\Theta^n)_{n \in \mathbb{N}}$ to $\bar{\Theta}$ in $\mathcal{P}(\mathcal{Z})$, we have for every $M > 0$,

$$\lim_{n \rightarrow \infty} \int_{\mathcal{Z}} M \wedge \|\varphi\|_{\mathcal{X}}^2 \Theta^n(d\varphi, dr, dw) = \int_{\mathcal{Z}} M \wedge \|\varphi\|_{\mathcal{X}}^2 \bar{\Theta}(d\varphi, dr, dw).$$

It suffices to show that (recall the notation for the marginal distributions)

$$\limsup_{M \rightarrow \infty} \sup_{n \in \mathbb{N}} \int_{\mathcal{X}} \mathbf{1}_{\{\|\varphi\|_{\mathcal{X}}^2 \geq M\}} \cdot \|\varphi\|_{\mathcal{X}}^2 [\Theta^n]_{\mathcal{X}}(d\varphi) = 0.$$

But this is true by Hölder's inequality, the Markov inequality and the fact that $\sup_{n \in \mathbb{N}} g(\Theta^n) \leq c < \infty$ by hypothesis since

$$\begin{aligned} & \sup_{n \in \mathbb{N}} \int_{\mathcal{X}} \mathbf{1}_{\{\|\varphi\|_{\mathcal{X}}^2 \geq M\}} \cdot \|\varphi\|_{\mathcal{X}}^2 [\Theta^n]_{\mathcal{X}}(d\varphi) \\ & \leq \sup_{n \in \mathbb{N}} \left\{ [\Theta^n]_{\mathcal{X}}(\{\|\varphi\|_{\mathcal{X}}^2 \geq M\})^{\frac{\delta_0}{2+\delta_0}} \cdot \left(\int_{\mathcal{X}} \|\varphi\|_{\mathcal{X}}^{2+\delta_0} [\Theta^n]_{\mathcal{X}}(d\varphi) \right)^{\frac{2}{2+\delta_0}} \right\} \\ & \leq M^{-\frac{\delta_0}{2+\delta_0}} \cdot c^{\frac{2\delta_0}{(2+\delta_0)^2}} \cdot c^{\frac{2}{2+\delta_0}}, \end{aligned}$$

which tends to zero as $M \rightarrow \infty$.

D Proof of Lemma 5.3: local martingale property

We have to show that, for \mathbf{P} -almost every $\omega \in \Omega$, any $f : \mathbb{R}^d \times \mathbb{R}^{d_1} \rightarrow \mathbb{R}$ monomial of first or second order, $M_f^{\mu_\omega}$ is a (\mathcal{G}_t) -local martingale under Q_ω ; cf. (iii) in Definition 4.1. Recall that μ_ω is the flow of measures in \mathcal{M}_2 induced by Q_ω , that is, $\mu_\omega(t) = Q_\omega \circ (\hat{X}(t))^{-1}$, $t \in [0, T]$. If $\Theta \in \mathcal{P}_2(\mathcal{Z})$, then the flow of measures induced by Θ is in \mathcal{M}_2 ; cf. Remark 4.2 above. Thus, we may write M_f^Θ meaning the process $M_f^{\mathfrak{m}}$ with \mathfrak{m} the flow of measures in \mathcal{M}_2 given by $\mathfrak{m}(t) \doteq \Theta \circ (\hat{X}(t))^{-1}$, $t \in [0, T]$.

We closely follow the proof of Lemma 5.2 in Budhiraja et al. [2012]. The canonical space \mathcal{Z} there is slightly bigger than our \mathcal{Z} here (relaxed controls in \mathcal{R}_1 instead of \mathcal{R}_2), but this causes no problems since the smaller space gives L^2 -integrability of controls (instead of L^1) and we have the corresponding distributional convergence of Q^n to Q as $\mathcal{P}_2(\mathcal{Z})$ -valued random variables; cf. Lemma 5.1 above.

In verifying the local martingale property of $M_f^{\mu_\omega}$, we will work with randomized stopping times. This will ensure almost sure continuity of certain mappings even if the diffusion coefficient $\sigma\sigma^\top$ is degenerate. The randomized stopping times live on an extension $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$ of the measurable space

$(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ and are adapted to a canonical filtration $(\hat{\mathcal{G}}_t)$ in $\mathcal{B}(\hat{\mathcal{Z}})$ given by

$$\hat{\mathcal{Z}} \doteq \mathcal{Z} \times [0, 1], \quad \hat{\mathcal{G}}_t \doteq \mathcal{G}_t \times \mathcal{B}([0, 1]), \quad t \in [0, T].$$

Any random object defined on $(\mathcal{Z}, \mathcal{B}(\mathcal{Z}))$ also lives on $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$, and no notational distinction will be made. Let λ denote the uniform distribution on $\mathcal{B}([0, 1])$. Any probability measure Θ on $\mathcal{B}(\mathcal{Z})$ induces a probability measure on $\mathcal{B}(\hat{\mathcal{Z}})$ given by $\Theta \otimes \lambda$. For $k \in \mathbb{N}$, define a stopping time τ_k on $(\hat{\mathcal{Z}}, \mathcal{B}(\hat{\mathcal{Z}}))$ with respect to the filtration $(\hat{\mathcal{G}}_t)$ by setting, for $((\varphi, r, w), a) \in \mathcal{Z} \times [0, 1]$,

$$\tau_k((\varphi, r, w), a) \doteq \inf \{t \in [0, T] : v((\varphi, r, w), t) \geq k + a\},$$

where

$$v((\varphi, r, w), t) \doteq \int_{\Gamma \times [0, t]} |y|^2 r(dy, ds) + \sup_{s \in [0, t]} |\varphi(s)| + \sup_{s \in [0, t]} |w(s)|.$$

Then, given any $\Theta \in \mathcal{P}(\mathcal{Z})$, $\tau_k \nearrow T$ as $k \rightarrow \infty$ and the mapping

$$\mathcal{Z} \times [0, 1] \ni ((\varphi, r, w), a) \mapsto \tau_k((\varphi, r, w), a) \in [0, T]$$

is continuous with probability one under $\Theta \otimes \lambda$.

Notice that if M_f^Θ is a local martingale with respect to $(\hat{\mathcal{G}}_t)$ under $\Theta \otimes \lambda$ with localizing sequence of stopping times $(\tau_k)_{k \in \mathbb{N}}$, then M_f^Θ is also a local martingale with respect to (\mathcal{G}_t) under Θ with localizing sequence of stopping times $(\tau_k(\cdot, 0))_{k \in \mathbb{N}}$; see the appendix in Budhiraja et al. [2012]. Thus, it suffices to prove the martingale property of M_f^Θ up till time τ_k with respect to the filtration $(\hat{\mathcal{G}}_t)$ and the probability measure $\Theta \otimes \lambda$.

Clearly, the process $M_f^\Theta(\cdot \wedge \tau_k)$ is a $(\hat{\mathcal{G}}_t)$ -martingale under $\Theta \otimes \lambda$ if and only if

$$(D.1) \quad \mathbf{E}_{\Theta \otimes \lambda} [\Psi \cdot (M_f^\Theta(t_1 \wedge \tau_k) - M_f^\Theta(t_0 \wedge \tau_k))] = 0$$

for all $t_0, t_1 \in [0, T]$ with $t_0 \leq t_1$, and $\hat{\mathcal{G}}_{t_0}$ -measurable $\Psi \in \mathbf{C}_b(\hat{\mathcal{Z}})$. To verify the martingale property of $M_f^\Theta(\cdot \wedge \tau_k)$, it is enough to check that (D.1) holds for any countable collection of times t_0, t_1 which is dense in $[0, T]$ and any countable collection of functions $\Psi \in \mathbf{C}_b(\hat{\mathcal{Z}})$ that generates the (countably many) σ -algebras $\hat{\mathcal{G}}_{t_0}$. Recall that the collection of test functions f for which a martingale property must be verified consists of just monomials of degree one or two, and hence is finite. Thus, we can choose a countable collection $\mathcal{T} \subset \mathbb{N} \times [0, T]^2 \times \mathbf{C}_b(\hat{\mathcal{Z}}) \times \mathbf{C}^2(\mathbb{R}^d \times \mathbb{R}^{d_1})$ of test parameters such that whenever $\Theta \in \mathcal{P}_2(\mathcal{Z})$ satisfies (D.1) for all $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$, then M_f^Θ is a (\mathcal{G}_t) -local martingale under Θ .

Let $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$. Define a mapping $\Phi = \Phi_{(k, t_0, t_1, \Psi, f)} : \mathcal{P}_2(\mathcal{Z}) \rightarrow \mathbb{R}$ by

$$\Phi(\Theta) \doteq \mathbf{E}_{\Theta \otimes \lambda} \left[\Psi \cdot (M_f^\Theta(t_1 \wedge \tau_k) - M_f^\Theta(t_0 \wedge \tau_k)) \right].$$

We claim that Φ is continuous on $\mathcal{P}_2(\mathcal{Z})$. To check this, take $\Theta \in \mathcal{P}_2(\mathcal{Z})$ and any sequence $(\Theta_l)_{l \in \mathbb{N}} \subset \mathcal{P}_2(\mathcal{Z})$ that converges to Θ in $\mathcal{P}_2(\mathcal{Z})$. Let $\mathbf{m}_l, l \in \mathbb{N}$, \mathbf{m} be the induced flows of measures in \mathcal{M}_2 , that is, $\mathbf{m}_l(t) \doteq \Theta_l \circ (\hat{X}(t))^{-1}$, $\mathbf{m}(t) \doteq \Theta \circ (\hat{X}(t))^{-1}$, $t \in [0, T]$. Recall the definition of $M_f^\Theta = M_f^{\mathbf{m}}$ in (4.3) and (4.4) above. By Assumption (A3) and definition of the stopping time τ_k , the integrand in (4.3) is bounded. By continuity of b, σ according to Assumption (A2), the almost sure continuity of τ_k under $\Theta \otimes \lambda$, the extended mapping theorem [Theorem 5.5 in Billingsley, 1968, p. 34] applied to the relaxed controls in (4.3) (plus convergence of first moments by choice of the topology on \mathcal{R}_2), and the fact that $\Psi \in \mathbf{C}_b(\hat{\mathcal{Z}})$, it follows that the mapping

$$\hat{\mathcal{Z}} \ni \hat{z} \mapsto \Psi(\hat{z}) \cdot (M_f^{\mathbf{m}}(t_1 \wedge \tau_k(\hat{z}), \hat{z}) - M_f^{\mathbf{m}}(t_0 \wedge \tau_k(\hat{z}), \hat{z})) \in \mathbb{R}$$

is bounded and $\Theta \otimes \lambda$ -almost surely continuous. By weak convergence and the mapping theorem [Theorem 5.1 in Billingsley, 1968, p. 30], it follows that

$$(D.2) \quad \begin{aligned} & \mathbf{E}_{\Theta_l \otimes \lambda} \left[\Psi \cdot (M_f^{\mathbf{m}_l}(t_1 \wedge \tau_k) - M_f^{\mathbf{m}_l}(t_0 \wedge \tau_k)) \right] \\ & \xrightarrow{l \rightarrow \infty} \mathbf{E}_{\Theta \otimes \lambda} \left[\Psi \cdot (M_f^{\mathbf{m}}(t_1 \wedge \tau_k) - M_f^{\mathbf{m}}(t_0 \wedge \tau_k)) \right]. \end{aligned}$$

Since $(\Theta_l)_{l \in \mathbb{N}}$ converges to Θ in $\mathcal{P}_2(\mathcal{Z})$, we have that $\{\Theta_l : l \in \mathbb{N}\} \cup \{\Theta\}$ is compact in $\mathcal{P}_2(\mathcal{Z})$. By continuity of projections, dominated convergence and the definition of $d_{\mathcal{Z}}$, we have $\lim_{l \rightarrow \infty} d_2(\mathbf{m}_l(t), \mathbf{m}(t)) = 0$ uniformly in $t \in [0, T]$. This together with Assumption (A2) and the construction of τ_k implies that

$$\sup_{t \in [0, T], \hat{z} \in \hat{\mathcal{Z}}} \left| M_f^{\mathbf{m}_l}(t \wedge \tau_k(\hat{z}), \hat{z}) - M_f^{\mathbf{m}}(t \wedge \tau_k(\hat{z}), \hat{z}) \right| \xrightarrow{l \rightarrow \infty} 0.$$

Since Ψ is bounded, it follows by dominated convergence that

$$\begin{aligned} & \left| \mathbf{E}_{\Theta_l \otimes \lambda} \left[\Psi \cdot (M_f^{\mathbf{m}_l}(t_1 \wedge \tau_k) - M_f^{\mathbf{m}_l}(t_0 \wedge \tau_k)) \right] \right. \\ & \quad \left. - \mathbf{E}_{\Theta \otimes \lambda} \left[\Psi \cdot (M_f^{\mathbf{m}_l}(t_1 \wedge \tau_k) - M_f^{\mathbf{m}_l}(t_0 \wedge \tau_k)) \right] \right| \xrightarrow{l \rightarrow \infty} 0. \end{aligned}$$

In combination with (D.2), this implies $\Phi(\Theta_l) \rightarrow \Phi(\Theta)$ as $l \rightarrow \infty$.

By hypothesis, the sequence $(Q^n)_{n \in \mathbb{I}}$ of $\mathcal{P}_2(\mathcal{Z})$ -valued random variables converges to Q in distribution. Hence the mapping theorem and the continuity of Φ imply $\Phi(Q^n) \rightarrow \Phi(Q)$ in distribution as $n \rightarrow \infty$. Let $n \in \mathbb{I}$. By

construction of Q^n and Fubini's theorem, for every $\omega \in \Omega_n$,

$$\begin{aligned}
\Phi(Q_\omega^n) &= \mathbf{E}_{Q_\omega^n \otimes \lambda} \left[\Psi \cdot (M_f^{\mu_\omega^n}(t_1 \wedge \tau_k) - M_f^{\mu_\omega^n}(t_0 \wedge \tau_k)) \right] \\
&= \frac{1}{n} \sum_{i=1}^n \int_0^1 \Psi((X_i^n(\cdot, \omega), \rho_\omega^{n,i}, W_i^n(\cdot, \omega)), a) \\
&\quad \cdot \left(f(X_i^n(t_1 \wedge \tau_k^{n,i}(\omega, a), \omega), W_i^n(t_1 \wedge \tau_k^{n,i}(\omega, a), \omega)) \right. \\
&\quad - f(X_i^n(t_0 \wedge \tau_k^{n,i}(\omega, a), \omega), W_i^n(t_0 \wedge \tau_k^{n,i}(\omega, a), \omega)) \\
&\quad \left. - \int_{t_0 \wedge \tau_k^{n,i}(\omega, a)}^{t_1 \wedge \tau_k^{n,i}(\omega, a)} \mathcal{A}_{u_i^n(s, \omega), s}^{\mu_\omega^n}(f)(X_i^n(s, \omega), W_i^n(s, \omega)) ds \right) da,
\end{aligned}$$

where \mathcal{A} is defined by (4.4) and $\tau_k^{n,i}(\omega, a)$ is defined like $\tau_k((\varphi, r, w), a)$ with φ replaced by $X_i^n(\cdot, \omega)$, r replaced by $\rho_\omega^{n,i}$, the relaxed control corresponding to $u_i^n(\cdot, \omega)$, and w replaced by $W_i^n(\cdot, \omega)$.

Let $a \in [0, 1]$. By Itô's formula, it holds \mathbf{P}_n -almost surely that

$$\begin{aligned}
&f(X_i^n(t_1 \wedge \tau_k^{n,i}), W_i^n(t_1 \wedge \tau_k^{n,i})) - f(X_i^n(t_0 \wedge \tau_k^{n,i}), W_i^n(t_0 \wedge \tau_k^{n,i})) \\
&\quad - \int_{t_0 \wedge \tau_k^{n,i}}^{t_1 \wedge \tau_k^{n,i}} \mathcal{A}_{u_i^n(s), s}^{\mu^n}(f)(X_i^n(s), W_i^n(s)) ds \\
&= \int_{t_0 \wedge \tau_k^{n,i}}^{t_1 \wedge \tau_k^{n,i}} \nabla_x f(X_i^n(s), W_i^n(s))^\top \sigma(s, X_i^n(s), \mu^n(s)) dW_i^n(s) \\
&\quad + \int_{t_0 \wedge \tau_k^{n,i}}^{t_1 \wedge \tau_k^{n,i}} \nabla_y f(X_i^n(s), W_i^n(s))^\top dW_i^n(s),
\end{aligned}$$

where $\tau_k^{n,i} = \tau_k^{n,i}(\cdot, a)$ and $\tau_k^{n,i}, \mu^n, X_i^n, u_i^n$ all live on $(\Omega_n, \mathcal{F}^n)$. By Fubini's theorem and Jensen's inequality, it follows that

$$\begin{aligned}
&\mathbf{E}_n [\Phi(Q^n)^2] \\
&\leq \int_0^1 \mathbf{E}_n \left[\mathbf{E}_{Q_\omega^n} \left[\Psi(\cdot, a) \cdot \left(M_f^{Q_\omega^n}(t_1 \wedge \tau_k(\cdot, a)) - M_f^{Q_\omega^n}(t_0 \wedge \tau_k(\cdot, a)) \right) \right]^2 \right] da.
\end{aligned}$$

Let again $a \in [0, 1]$. By the Itô isometry, the independence of the Wiener processes W_1^n, \dots, W_n^n , and because $\Psi(\cdot, a)$ is \mathcal{G}_{t_0} -measurable and $\tau_k(\cdot, a)$ is a stopping time with respect to (\mathcal{G}_t) , it holds that

$$\mathbf{E}_n \left[\mathbf{E}_{Q_\omega^n} \left[\Psi(\cdot, a) \cdot \left(M_f^{Q_\omega^n}(t_1 \wedge \tau_k(\cdot, a)) - M_f^{Q_\omega^n}(t_0 \wedge \tau_k(\cdot, a)) \right) \right]^2 \right]$$

$$\begin{aligned}
&= \mathbf{E}_n \left[\left(\frac{1}{n} \sum_{i=1}^n \int_{t_0 \wedge \tau_k^{n,i}(\cdot, a)}^{t_1 \wedge \tau_k^{n,i}(\cdot, a)} \Psi(\cdot, a) \cdot \mathbf{1}_{\{\tau_k^{n,i}(\cdot, a) \geq t_0\}} \cdot \left(\nabla_y f(X_i^n(s), W_i^n(s))^\top \right. \right. \right. \\
&\quad \left. \left. \left. + \nabla_x f(X_i^n(s), W_i^n(s))^\top \sigma(s, X_i^n(s), \mu^n(s)) \right) dW_i^n(s) \right)^2 \right] \\
&= \frac{1}{n^2} \sum_{i=1}^n \mathbf{E}_n \left[\int_{t_0 \wedge \tau_k^{n,i}(\cdot, a)}^{t_1 \wedge \tau_k^{n,i}(\cdot, a)} \left| \Psi(\cdot, a) \cdot \mathbf{1}_{\{\tau_k^{n,i}(\cdot, a) \geq t_0\}} \cdot \left(\nabla_y f(X_i^n(s), W_i^n(s))^\top \right. \right. \right. \\
&\quad \left. \left. \left. + \nabla_x f(X_i^n(s), W_i^n(s))^\top \sigma(s, X_i^n(s), \mu^n(s)) \right) \right|^2 ds \right] \\
&\xrightarrow{n \rightarrow \infty} 0.
\end{aligned}$$

Since $(\Phi(Q^n))_{n \in \mathbb{I}}$ converges to $\Phi(Q)$ in distribution, it follows that for each $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$ we can choose a set $Z_{(k, t_0, t_1, \Psi, f)} \in \mathcal{F}$ such that $\mathbf{P}(Z_{(k, t_0, t_1, \Psi, f)}) = 0$ and

$$\Phi(Q_\omega) = \Phi_{(k, t_0, t_1, \Psi, f)}(Q_\omega) = 0 \text{ for all } \omega \in \Omega \setminus Z_{(k, t_0, t_1, \Psi, f)}.$$

Let Z be the union of all sets $Z_{(k, t_0, t_1, \Psi, f)}$, $(k, t_0, t_1, \Psi, f) \in \mathcal{T}$. Since \mathcal{T} is countable, we have $Z \in \mathcal{F}$, $\mathbf{P}(Z) = 0$ and

$$\Phi_{(k, t_0, t_1, \Psi, f)}(Q_\omega) = 0 \text{ for all } \omega \in \Omega \setminus Z, \text{ all } (k, t_0, t_1, \Psi, f) \in \mathcal{T}.$$

By definition of Φ , this implies that, for every test function f , $M_f^{\mu_\omega}$ is a (\mathcal{G}_t) -local martingale under Q_ω for \mathbf{P} -almost every $\omega \in \Omega$.

References

- M. Bardi and F. S. Priuli. LQG mean-field games with ergodic costs. In *Proc. 52nd IEEE-CDC*, 2013.
- M. Bardi and F. S. Priuli. Linear-quadratic N -person and mean-field games with ergodic cost. *SIAM J. Control Optim.*, 52(5):3022–3052, 2014.
- D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, Belmont, Massachusetts, reprint of the 1978 edition, 1996.
- P. Billingsley. *Convergence of Probability Measures*. Wiley series in Probability and Statistics. John Wiley & Sons, New York, 1968.
- A. Budhiraja, P. Dupuis, and M. Fischer. Large deviation properties of weakly interacting processes via weak convergence methods. *Ann. Probab.*, 40(1):74–102, 2012.

- P. Cardaliaguet. Notes on mean field games. Technical report, Université de Paris - Dauphine, September 2013.
- P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions. The master equation and the convergence problem in mean field games. arXiv:1509.02505 [math.AP], September 2015.
- R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SIAM J. Control Optim.*, 51(4):2705–2734, 2013.
- R. Carmona and D. Lacker. A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.*, 25(3):1189–1231, 2015.
- R. Carmona, F. Delarue, and A. Lachapelle. Control of McKean-Vlasov dynamics versus mean field games. *Math. Financ. Econ.*, 7(2):131–166, 2013.
- P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. Wiley Series in Probability and Statistics. John Wiley & Sons, New York, 1997.
- N. El Karoui, D. H. Nguyen, and M. Jeanblanc-Picqué. Compactification methods in the control of degenerate diffusions: existence of an optimal control. *Stochastics*, 20(3):169–219, 1987.
- K. Fan. Fixed-point and minimax theorems in locally convex topological linear spaces. *Proc. Nat. Acad. Sci. U.S.A.*, 38:121–126, 1952.
- E. Feleqi. The derivation of ergodic mean field game equations for several populations of players. *Dyn. Games Appl.*, 3(4):523–536, 2013.
- M. Fischer and G. Nappo. On the moments of the modulus of continuity of Itô processes. *Stochastic Anal. Appl.*, 28(1):103–122, 2010.
- T. Funaki. A certain class of diffusion processes associated with nonlinear parabolic equations. *Z. Wahrscheinlichkeitstheor. verw. Gebiete*, 67(3):331–348, 1984.
- D. A. Gomes, J. Mohr, and R. R. Souza. Continuous time finite state mean field games. *Appl. Math. Optim.*, 68(1):99–143, 2013.
- A. D. Gottlieb. *Markov transitions and the propagation of chaos*. PhD thesis, Lawrence Berkeley National Laboratory, 1998.

- M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.*, 6(3):221–252, 2006.
- O. Kallenberg. On the existence of universal functional solutions to classical SDEs. *Ann. Probab.*, 24(1):196–205, 1996.
- O. Kallenberg. *Foundations of Modern Probability*. Probability and Its Applications. Springer, New York, 2nd edition, 2001.
- I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer, New York, 2nd edition, 1991.
- V. N. Kolokoltsov, J. Li, and W. Yang. Mean field games and nonlinear Markov processes. arXiv:1112.3744 [math.PR], 2011.
- H. J. Kushner. *Weak Convergence Methods and Singularly Perturbed Stochastic Control and Filtering Problems*, volume 3 of *Systems & Control: Foundations & Applications*. Birkhäuser, Boston, 1990.
- D. Lacker. Mean field games via controlled martingale problems: existence of Markovian equilibria. *Stochastic Processes Appl.*, 125(7):2856–2894, 2015a.
- D. Lacker. A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields*, to appear, 2015b. arXiv:1408.2708 [math.PR].
- J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. I. Le cas stationnaire. *C. R. Math. Acad. Sci. Paris*, 343(9):619–625, 2006a.
- J.-M. Lasry and P.-L. Lions. Jeux à champ moyen. II. Horizon fini et contrôle optimal. *C. R. Math. Acad. Sci. Paris*, 343(10):679–684, 2006b.
- J.-M. Lasry and P.-L. Lions. Mean field games. *Japan. J. Math.*, 2(1):229–260, 2007.
- H. P. McKean. A class of Markov processes associated with nonlinear parabolic equations. *Proc. Nat. Acad. Sci. U.S.A.*, 56(6):1907–1911, 1966.
- K. Oelschläger. A martingale approach to the law of large numbers for weakly interacting stochastic processes. *Ann. Probab.*, 12(2):458–479, 1984.

- D. W. Stroock and S. R. S. Varadhan. *Multidimensional Diffusion Processes*, volume 233 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1979.
- A.-S. Sznitman. Topics in propagation of chaos. In P.-L. Hennequin, editor, *Ecole d'Eté de Probabilités de Saint-Flour XIX - 1989*, volume 1464 of *Lecture Notes in Math.*, pages 165–251, Berlin, 1991. Springer-Verlag.
- C. Villani. *Topics in Optimal Transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, Rhode Island, 2003.
- T. Yamada and S. Watanabe. On the uniqueness of solutions of stochastic differential equations. *J. Math. Kyoto Univ.*, 11(1):155–167, 1971.