N-player games and mean field games with absorption*

Luciano Campi† Markus Fischer‡

December 12, 2016; revised August 7, 2017

Abstract

We introduce a simple class of mean field games with absorbing boundary over a finite time horizon. In the corresponding N-player games, the evolution of players' states is described by a system of weakly interacting Itô equations with absorption on first exit from a bounded open set. Once a player exits, her/his contribution is removed from the empirical measure of the system. Players thus interact through a renormalized empirical measure. In the definition of solution to the mean field game, the renormalization appears in form of a conditional law. We justify our definition of solution in the usual way, that is, by showing that a solution of the mean field game induces approximate Nash equilibria for the N-player games with approximation error tending to zero as N tends to infinity. This convergence is established provided the diffusion coefficient is non-degenerate. The degenerate case is more delicate and gives rise to counter-examples.

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

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

*The second author is partially supported by the research projects "Mean Field Games and Nonlinear PDEs" (CPDA157835) of the University of Padua and "Nonlinear Partial Differential Equations: Asymptotic Problems and Mean-Field Games" of the Fondazione CaRiPaRo. Both authors thank two anonymous Referees for their careful reading of the manuscript and helpful comments and suggestions.

†Department of Statistics, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, United Kingdom.

‡Department of Mathematics, University of Padua, via Trieste 63, 35121 Padova, Italy.

1 Introduction

Mean field games (MFGs, henceforth) were introduced by Lasry and Lions [2006a,b, 2007] and, independently, by Huang et al. [2006], as limit models for symmetric nonzero-sum non-cooperative N-player games with interaction of mean field type as the number of players tends to infinity. The limit relation is commonly understood in the sense that a solution of the MFG 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 approximation result is useful from a practical point of view since the model of interest is commonly the N-player game with N very large so that a direct computation of Nash equilibria is not feasible.

The purpose of this paper is to study N-player games and related MFGs in the presence of an absorbing set. Thus, a player is eliminated from the N-player game as soon as her/his private state process leaves a given open set $O \subset \mathbb{R}^d$, the set of non-absorbing states, which is the same for all players. We carry out our study for a simple class of continuous-time models with Itô-type dynamics with mean field interaction over a bounded time horizon. More specifically, the vector of private states $\mathbf{X}^N = (X_1^N, \dots, X_N^N)$ in the N-player game is assumed to evolve according to

(1.1)
$$X_i^N(t) = X_i^N(0) + \int_0^t \left(u_i(s, \boldsymbol{X}^N) + \bar{b}\left(t, X_i^N(t), \int_{\mathbb{R}^d} w(y)\pi^N(s, dy)\right)\right) ds + \sigma W_i^N(t), \quad t \in [0,T],\ i \in \{1, \dots, N\},$$

where $\boldsymbol{u} = (u_1, \dots, u_N)$ is a vector of feedback strategies with full state information (to be specified below), $W_1^N, \dots, W_N^N$ are independent d-dimensional Wiener processes defined on some filtered probability space, σ is some dispersion matrix, which we assume to be constant for simplicity, and $\bar{b}$, w are given deterministic functions. Moreover, $\pi^N(t, \cdot)$ is the (random) empirical measure of the players' states at time t that have not left O, that is,

$$\pi^N(t, \cdot) \doteq \begin{cases} \frac{1}{\bar{N}^N(t)} \sum_{j=1}^N \mathbf{1}_{[0, \tau^{X_j^N})}(t) \cdot \delta_{X_j^N(t)}(\cdot) & \text{if } \bar{N}^N(t) > 0, \\ \delta_0(\cdot) & \text{if } \bar{N}^N(t) = 0, \end{cases}$$

where $\bar{N}^N(t) \doteq \sum_{j=1}^N \mathbf{1}_{[0, \tau^{X_j^N})}(t)$ is the number of players still in the game at time t and

$$\tau^{X_j^N} \doteq \inf\left\{t \in [0,T] : X_j^N(t) \notin O\right\},$$

denotes the time of first exit of X_j^N from O, with the convention $\inf \emptyset = \infty$. By definition, $\pi^N(t, \cdot)$ equals the Dirac measure in $0 \in \mathbb{R}^d$ if all players have left O by time t. The choice of δ_0 in this case is arbitrary and has no impact on the game; this will be clear from the definition of the cost functionals J_i^N below. Heuristically, when some player exits the set O, he/she does not contribute anymore to the empirical measure $\pi^N(t, \cdot)$, which is computed with respect to the "survivors" only. Notice that the controls appear linearly in (1.1) only for the sake of simplicity. Even though more general dependencies could be considered, we do not aim at giving the minimal set of assumptions under which our results hold true.

Each player wants to minimize expected costs according to a given cost functional over the random time horizon determined as the minimum between the player's time of first exit from O and the overall time horizon T. More precisely, player i evaluates a strategy vector $\boldsymbol{u} = (u_1, \dots, u_N)$ according to

$$J_i^N(\boldsymbol{u}) \doteq \mathbf{E}\Bigg[\int_0^{\tau_i^N} f\left(s, X_i^N(s), \int_{\mathbb{R}^d} w(y)\pi^N(s, dy), u_i\left(s, \boldsymbol{X}^N\right)\right) ds \\ + F\left(\tau_i^N, X_i^N(\tau_i^N)\right)\Bigg],$$

where $\boldsymbol{X}^N$ is the solution of Eq. (1.1) under $\boldsymbol{u}$ and $\tau_i^N \doteq \tau^{X_i^N} \wedge T$ is the (random) duration of the game for player i. Notice that the cost coefficients f, F are the same for all players. As in the dynamics, we have mean field interaction in the costs through the renormalized empirical measures $\pi^N(t, \cdot)$. For simplicity, we only consider finite-dimensional dependencies on the measure variable, namely through integrals of the vector-valued function w. More details on the setting with all the technical assumptions will be given in the next sections.

The presence of an absorbing set can be motivated by economic models for credit risk and corporate finance. There, the players can be interpreted as interacting firms striving to maximize some objective functional, for instance the expected value of discounted future dividends, or as banks controlling their rate of borrowing/lending to a central bank as in the systemic risk model proposed by Carmona et al. [2015]. Within both contexts, the absorbing boundary can be naturally seen as a default barrier as in structural models for credit risk. In this paper, we concentrate on the mathematical properties of this family of games, while we postpone their possible applications to future research.

For our class of games, we focus on the construction of approximate Nash

equilibria for the N-player games through solutions of the corresponding MFG. Our main contributions are as follows:

- We introduce the limit model corresponding to the above N-player games as $N \to \infty$, namely the MFG with absorption. For a solution of the MFG, the renormalized empirical measures $\pi^N(t, \cdot)$ are replaced by a flow of conditional laws; see Definition 4.1.

- Under a non-degeneracy condition on the dispersion matrix σ, we prove that any regular feedback solution of the MFG induces a sequence of approximate Nash equilibria for the N-player games with approximation error tending to zero as $N \to \infty$; see Theorem 5.1. Here, "regular" means that the optimal feedback strategy is (almost surely) continuous with respect to the state variable.

- Under the same non-degeneracy condition on σ, we prove existence of a solution in feedback form to the MFG with absorption; see Theorem 6.2. Moreover, under some additional conditions, we show that the optimal feedback strategies are Markovian and continuous in the state variable; see Proposition 6.1 and Corollary 6.3. In that situation, we briefly sketch what would be the PDE approach to mean field games with absorption.

- In the degenerate case, i.e. when σ may vanish, we provide a counterexample where the solution of the limit MFG is not even nearly optimal (in the sense of inducing approximate Nash equilibria) for the N-player games. This is in contrast with what happens in the absence of an absorbing set.

The proof of Theorem 5.1 on approximate Nash equilibria is based on weak convergence arguments, controlled martingale problems and a reformulation of the original dynamics and costs using path-dependent coefficients; see, in particular, (3.2) and (5.4). This allows to work with solutions defined over the entire time interval $[0, T]$. The resulting description of the systems should be compared to the set-up used in Carmona and Lacker [2015]. There, questions of existence and uniqueness of solutions finite horizon MFGs with non-degenerate noise and functional coefficients are studied through probabilistic methods, and approximate Nash equilibria are constructed from the MFG. Nonetheless, the results of Carmona and Lacker [2015] cannot be applied directly to our situation, due in part to different continuity assumptions. What is more, approximate Nash equilibria are constructed there only for dynamics without mean field interaction [Carmona

and Lacker, 2015, Theorem 4.2]; this assumption implies an independence property not warranted in more general situations. The use of martingale problems in proving convergence to the McKean-Vlasov limit and propagation of chaos for weakly interacting systems has a long tradition; see, for instance, Funaki [1984], Oelschläger [1984], or Méléard [1996]. In those works, the N-particle systems are usually assumed to be fully symmetric, and the dynamics of all particles are determined by the same coefficients. Here, we have to study the passage to the many player (particle) limit also in the presence of a deviating player, which destroys the symmetry of the prelimit systems. The absorbing boundary introduces a discontinuity into the dynamics so that a single deviating player might have a non-negligible effect on the limit behavior of the N-player empirical measures. This is the reason why we give a detailed proof of convergence in Appendix A.

The proof of Theorem 6.2 on existence of a feedback solution for the limit MFG is based on results on BSDEs with random terminal horizon as in Darling and Pardoux [1997], Briand and Hu [1998] and the use of the Brouwer-Schauder-Tychonov fixed point theorem applied to a suitable map in the same spirit as Carmona and Lacker [2015]. For the continuity of the optimal feedback strategy, we rely on the classical PDE approach to optimal control and adapt to our setting the proof of a regularity result due to Fleming and Rishel [1975].

For the counter-example (given in Section 7), we consider systems with dispersion coefficient σ equal to zero; the only source of randomness comes from the initial conditions. We construct, for a specific choice of the initial distribution, a feedback solution of the MFG that is Lipschitz continuous in the state variable and such that the induced strategy vectors are not asymptotically optimal for the corresponding N-player games. The reason for this non-convergence is a change in the controllability of the individual player dynamics between limit and prelimit systems in conjunction with the discontinuity of the costs introduced by the absorbing boundary. Also notice that the initial distribution we choose is singular with respect to Lebesgue measure. The counter-example thus holds little surprise from the point of view of optimal control theory. In the context of MFGs without absorption, on the other hand, the connection between solutions of the limit MFG and approximate Nash equilibria for the N-player games is known to be robust and persists even for systems with fully degenerate noise; see Theorem 2.11 in Lacker [2016] for a general result in this direction. From this point of view, the counter-example seems to be interesting.

Related literature. Mean-field models similar to ours have been studied before in different contexts. In a first group of papers, such as Giesecke et al. [2013], Spiliopoulos et al. [2014], Giesecke et al. [2015], a point process model of correlated defaults timing in a portfolio of firms has been introduced. More specifically, a firm defaults with some intensity which follows a mean-reverting jump-diffusion process that is driven by several terms, one of them being the default rate in the pool. This naturally induces a mean-field among the default intensities: every time one firm defaults a jump occurs in the default rate of the pool and hence in the intensities of the survivors. The effects of defaults fade away with time. A law of large number (LLN) for the default rate as the number N of firms goes to infinity is proved in Giesecke et al. [2013], while other results on the loss from default are analyzed in the companion papers Spiliopoulos et al. [2014], Giesecke et al. [2015]. Similar results been obtained in the interacting particle models proposed by Cvitanić et al. [2012] and by Dai Pra et al. [2009], where the effect of defaults on the survivors is permanent.

Apart from the fact that their setting is not controlled, the main difference between our model and theirs is that whenever, in our model, some diffusion is absorbed and the conditional empirical measure is updated accordingly, the diffusions still in O are affected by it in a continuous way since the empirical measure appears only in the drift coefficient. Moreover, while in the setting of Giesecke, Spiliopoulos and coauthors the intensities can have a common noise, this is not included in our model for the sake of simplicity.

More recently, Hambly and Ledger [2016], motivated by various applications from large credit portfolios to neuroscience, have proposed a system of N uncontrolled diffusions that are killed as soon as they go negative. Furthermore, their coefficients are functions of the current proportional loss, which is defined as the proportion out of N of killed diffusion. They prove a LLN for the empirical measure of the population using some energy estimates in combination with weak convergence techniques.

Two more papers, which are related to ours, are those by Delarue et al. [2015a,b]. Motivated by applications in neuroscience, these authors study the well-posedness of an integrate-and-fire model and its approximation via particle systems. Mathematically speaking, they look at a nonlinear scalar SDE with a jump component responsible for resetting trajectories as soon as they hit a given threshold (occurrence of a "spike") and a singular drift term of mean field type (the "nonlinear" interaction) depending on the average number of spikes before current time. When the nonlinear interaction is too strong, any solution can have a blow-up in finite time. In our model, we

could have a similar cascade effect, which would correspond to a situation when all players get absorbed before the finite horizon T. Investigating such a phenomenon in our setting is interesting on its own but goes beyond the scope of this paper.

In the applied literature, MFGs with absorption at zero have been recently considered in Chan and Sircar [2015a,b] within the context of oligopolistic models with exhaustible resources. A common feature that their model shares with ours is that they also keep track of the fraction of active players remaining at time t, which appears in the objective functions (through the control) but not in the state variable. Moreover, they look at some particular cases, which are relevant from an economic perspective, and perform some asymptotic expansions corresponding to the case of "small competition". A rigorous study addressing existence and uniqueness issues in Chan and Sircar's model has been subsequently done by Graber and Bensoussan [2016].

The last work we mention here is by Bensoussan et al. [2014] and bears a different relation to ours. There, the authors construct Nash equilibria using PDE methods for a class of stochastic differential games with a varying number of players (the dynamics are not of mean field type). While the maximum number of players is prescribed, players may leave the game (be pushed out or "die"), and new players may join. In our case, we only allow players to leave (through absorption). A natural extension of our model would include a mechanism by which players enter the game.

Structure of the paper. The rest of this paper is organized as follows. Section 2 introduces some terminology and notation and sets the main assumptions on the dynamics as well as on the cost functionals. Section 3 describes the setting of N-player games with absorption, while Section 4 introduces the corresponding MFG. In Section 5, one of the main results, namely the construction of approximate Nash equilibria for the N-player game from a solution of the limit problem, is stated and proved. Section 6 contains the results on the existence of (regular) feedback solutions for the MFG, in particular those with Markov feedback controls, as well as a sketch of the PDE approach to mean field games with absorption. In Section 7, we provide the aforementioned counter-example in the case of degenerate noise. The technical results used in the paper are all gathered in the Appendix, including the aforementioned propagation-of-chaos-type results in Appendix A and a uniqueness result for McKean-Vlasov equations in Appendix B.

2 Preliminaries and assumptions

Let $d \in \mathbb{N}$, which will be the dimension of the space of private states, noise values, as well as control actions. The spaces $\mathbb{R}^n$ with $n \in \mathbb{N}$ are equipped with the standard Euclidean norm, always indicated by $|.|$. Choose $T > 0$, the finite time horizon.

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.

Set $\mathcal{X} \doteq \mathbf{C}([0,T], \mathbb{R}^d)$, which can be seen as the space of individual state trajectories. As usual, equip $\mathcal{X}$ with the topology of uniform convergence, which turns it into a Polish space. Let $\|\cdot\|_{\mathcal{X}}$ denote the supremum norm on $\mathcal{X}$. Denote by $\hat{X}$ the coordinate process on $\mathcal{X}$:

$$\hat{X}(t, \varphi) \doteq \varphi(t), \qquad t \in [0,T], \qquad \varphi \in \mathcal{X}.$$

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

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

Given $N \in \mathbb{N}$, we will use the usual identification of $\mathcal{X}^N = \times^N \mathcal{X}$ with the space $\mathbf{C}([0,T], \mathbb{R}^{N \cdot d})$; $\mathcal{X}^N$, too, will be equipped with the topology of uniform convergence. We call a function u defined on $[0,T] \times \mathcal{X}^N$ with values in some measurable space *progressively measurable* if u is measurable and, for every $t \in [0,T]$, all $\boldsymbol{\varphi}, \tilde{\boldsymbol{\varphi}} \in \mathcal{X}^N$,

$$u(t, \boldsymbol{\varphi}) = u(t, \tilde{\boldsymbol{\varphi}}) \text{ whenever } \boldsymbol{\varphi}(s) = \tilde{\boldsymbol{\varphi}}(s) \text{ for all } s \in [0,t].$$

Let Γ be a closed subset of $\mathbb{R}^d$, the set of control actions, or action space.

Let $O \subset \mathbb{R}^d$ be an open set, the set of non-absorbing states. For Y an $\mathbb{R}^d$-valued process defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$ over the time interval $[0,T]$, let

$$\tau^Y(\omega) \doteq \inf\{t \in [0,T] : Y(t,\omega) \notin O\}, \quad \omega \in \Omega,$$

denote the random time of first exit of Y from O, with the convention $\inf \emptyset = \infty$. Clearly, if Y has continuous trajectories and is adapted to some filtration $(\mathcal{F}_t)$ in $\mathcal{F}$, then τ^Y is an $(\mathcal{F}_t)$-stopping time. In this case, $\tau^Y = \inf\{t \in [0,T] : Y(t) \in \partial O\}$, where ∂O denotes the boundary of O.

Let $d_0 \in \mathbb{N}$, and let $w \colon \mathbb{R}^d \to \mathbb{R}^{d_0}$,

$$\bar{b} \colon [0,T] \times \mathbb{R}^d \times \mathbb{R}^{d_0} \to \mathbb{R}^d, \qquad \sigma \in \mathbb{R}^{d \times d},$$
$$f \colon [0,T] \times \mathbb{R}^d \times \mathbb{R}^{d_0} \times \Gamma \to [0,\infty), \qquad F \colon [0,T] \times \mathbb{R}^d \to [0,\infty).$$

The function w will denote an integrand for the measure variable in the drift and the running costs, respectively, $\bar{b}$ a function of the drift integral, σ the dispersion coefficient of the dynamics, while f, F quantify the running and terminal costs, respectively. Notice that the dispersion coefficient σ is a constant matrix and that the cost coefficients f, F are non-negative functions. Let us make the following assumptions:

(H1) Boundedness and measurability: w, $\bar{b}$, f, F are Borel measurable functions uniformly bounded by some constant $K > 0$.

(H2) Continuity: w, f, F are continuous, and $\bar{b}(t, \cdot, \cdot)$ is continuous, uniformly in $t \in [0, T]$.

(H3) Lipschitz continuity of $\bar{b}$: there exists $\bar{L} > 0$ such that for all $x, \tilde{x} \in \mathbb{R}^d$, $y, \tilde{y} \in \mathbb{R}^{d_0}$,

$$\sup_{t\in[0,T]} \left|\bar{b}(t, x, y) - \bar{b}(t, \tilde{x}, \tilde{y})\right| \leq \bar{L} \left(|x - \tilde{x}| + |y - \tilde{y}|\right).$$

(H4) Action space: $\Gamma \subset \mathbb{R}^d$ is compact and convex (and non-empty).

(H5) State space: $O \subset \mathbb{R}^d$ is non-empty, open, and bounded such that ∂O is a C^2-manifold.

The results of Sections 5 and 6 will be established under the following additional assumption:

(ND) Non-degeneracy: σ is a matrix of full rank.

3 N-player games

Let $N \in \mathbb{N}$ be the number of players. Denote by $X_i^N(t)$ the private state of player i at time $t \in [0, T]$. The evolution of the players' states depends on the strategies they choose as well as the initial distribution of states, which we indicate by ν_N (thus, $\nu_N \in \mathcal{P}(\mathbb{R}^{N\times d})$). We assume that $\mathrm{supp}(\nu_N) \subset \times^N O$ and that ν_N is symmetric in the sense that

$$\nu^N \circ \mathfrak{s}^{-1} = \nu^N$$

for all maps $\mathfrak{s}\colon (\mathbb{R}^d)^N \to (\mathbb{R}^d)^N$ of the form $(x_1, \ldots, x_N) \mapsto (x_{p(1)}, \ldots, x_{p(N)})$ for some permutation p of $(1, \ldots, N)$.

Here, we consider players using feedback strategies with full state information (up to the current time). Thus, let $\mathcal{U}_N$ denote the set of all progressively measurable functions $u : [0, T] \times \mathcal{X}^N \to \Gamma$. Elements of $\mathcal{U}_N$

represent individual strategies. A vector $(u_1, \dots, u_N)$ of individual strategies is called a strategy vector or strategy profile. Given a strategy vector $\boldsymbol{u} = (u_1, \dots, u_N) \in \times^N \mathcal{U}_N$, consider the system of equations

(3.1)
$$X_i^N(t) = X_i^N(0) + \int_0^t \left(u_i(s, \boldsymbol{X}^N) + \bar{b}\left(s, X_i^N(s), \int_{\mathbb{R}^d} w(y)\pi^N(s, dy) \right) \right) ds + \sigma W_i^N(t), \quad t \in [0,T],\ i \in \{1, \dots, N\},$$

where $\boldsymbol{X}^N = (X_1^N, \dots, X_N^N)$, $W_1^N, \dots, W_N^N$ are independent d-dimensional Wiener processes defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ and $\pi^N(t, \cdot)$ is the empirical measure of the players' states at time t that have not left O, that is,

$$\pi_\omega^N(t, \cdot) \doteq \begin{cases} \frac{1}{\bar{N}_\omega^N(t)} \sum_{j=1}^N \mathbf{1}_{[0, \tau^{X_j^N}(\omega))}(t) \cdot \delta_{X_j^N(t,\omega)}(\cdot) & \text{if } \bar{N}_\omega^N(t) > 0, \\ \delta_0(\cdot) & \text{if } \bar{N}_\omega^N(t) = 0, \end{cases}$$

where

$$\bar{N}_\omega^N(t) \doteq \sum_{j=1}^N \mathbf{1}_{[0, \tau^{X_j^N}(\omega))}(t), \quad \omega \in \Omega.$$

It will be convenient to rewrite (3.1) as a system of particles interacting through their unconditional empirical measure on the path space. To this end, set $\tau \doteq \tau^{\hat{X}}$, that is,

$$\tau(\varphi) \doteq \inf\{t \in [0,T] : \varphi(t) \notin O\}, \quad \varphi \in \mathcal{X},$$

and define $b\colon [0,T] \times \mathcal{X} \times \mathcal{P}(\mathcal{X}) \times \Gamma \to \mathbb{R}^d$ by

(3.2)
$$b(t, \varphi, \theta, \gamma) \doteq \begin{cases} \gamma + \bar{b}\left(t, \varphi(t), \frac{\int w(\tilde{\varphi}(t)) \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t) \theta(d\tilde{\varphi})}{\int \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t) \theta(d\tilde{\varphi})} \right) & \text{if } \theta(\tau > t) > 0, \\ \gamma + \bar{b}\left(t, \varphi(t), w(0) \right) & \text{if } \theta(\tau > t) = 0. \end{cases}$$

Then b is measurable and progressive in the sense that, for all $t \in [0,T]$, all $\gamma \in \Gamma$,

$$b(t, \varphi, \theta, \gamma) = b(t, \tilde{\varphi}, \tilde{\theta}, \gamma) \text{ whenever } \varphi_{|[0,t]} = \tilde{\varphi}_{|[0,t]} \text{ and } \theta_{|\mathcal{G}_t} = \tilde{\theta}_{|\mathcal{G}_t}.$$

The solutions of (3.1) are thus equivalently described by:

$$X_i^N(t) = X_i^N(0) + \int_0^t b\left(s, X_i^N, \mu^N, u_i(s, \boldsymbol{X}^N) \right) ds + \sigma W_i^N(t), \quad t \in [0,T],\ i \in \{1, \dots, N\}, \tag{3.3}$$

where $W_1^N, \dots, W_N^N$ are independent d-dimensional Wiener processes defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$, and μ^N is the empirical measure of the players' state trajectories, that is,

$$\mu_\omega^N(B) \doteq \frac{1}{N} \sum_{j=1}^{N} \delta_{X_j^N(\cdot,\omega)}(B), \quad B \in \mathcal{B}(\mathcal{X}),\ \omega \in \Omega.$$

A solution of Eq. (3.1) (equivalently, of Eq. (3.3)) under $\boldsymbol{u} \in \times^N \mathcal{U}_N$ with initial distribution ν_N is therefore a triple $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \boldsymbol{W}^N, \boldsymbol{X}^N)$ where $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ is a filtered probability space satisfying the usual hypotheses, $\boldsymbol{W}^N = (W_1^N, \dots, W_N^N)$ a vector of independent d-dimensional $(\mathcal{F}_t)$-Wiener processes, and $\boldsymbol{X}^N = (X_1^N, \dots, X_N^N)$ a vector of continuous $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-adapted processes such that Eq. (3.1) (resp. Eq. (3.3)) holds $\mathbf{P}$-almost surely with strategy vector $\boldsymbol{u}$ and $\mathbf{P} \circ (\boldsymbol{X}^N(0))^{-1} = \nu_N$.

Let $\mathcal{U}_{fb}^N$ be the set of all strategy vectors $\boldsymbol{u} \in \times^N \mathcal{U}_N$ such that Eq. (3.1) under $\boldsymbol{u}$ with initial distribution ν_N possesses a solution that is unique in law. If the non-degeneracy assumption (ND) holds in addition to (H1) – (H5), then Eq. (3.1) is well posed given any strategy vector:

Proposition 3.1. *Grant (ND) in addition to (H1) – (H5). Then* $\mathcal{U}_{fb}^N = \times^N \mathcal{U}_N$.

Proof. Let $\boldsymbol{u} \in \times^N \mathcal{U}_N$. The system of equations (3.3) can be rewritten as one stochastic differential equation with state space $\mathbb{R}^{N \times d}$ driven by an $N \times d$-dimensional standard Wiener process with drift coefficient $\boldsymbol{b_u} \colon [0,T] \times \mathcal{X}^N \to \mathbb{R}^{N \times d}$ given by

$$\begin{aligned} &\boldsymbol{b_u}(t, \boldsymbol{\varphi}) \\ &\doteq \left(b\left(t, \varphi_1, \frac{1}{N}\sum_{j=1}^{N} \delta_{\varphi_j}, u_1(s, \boldsymbol{\varphi})\right), \dots, b\left(t, \varphi_N, \frac{1}{N}\sum_{j=1}^{N} \delta_{\varphi_j}, u_N(s, \boldsymbol{\varphi})\right)\right)^{\mathsf{T}} \end{aligned}$$

and non-degenerate constant diffusion coefficient. Notice that $\boldsymbol{b_u}$ is bounded and progressive with respect to the natural filtration in $\mathcal{B}(\mathcal{X}^N)$. Existence of a weak solution and uniqueness in law are now a consequence of Girsanov's theorem and the Stroock-Varadhan martingale problem; cf. V.27 in Rogers and Williams [2000, pp. 177-178]. □

The i-th player evaluates a strategy vector $\boldsymbol{u} = (u_1, \dots, u_N) \in \mathcal{U}_{fb}^N$ ac-

cording to the cost functional

$$J_i^N(\boldsymbol{u}) \doteq \mathbf{E}\Bigg[\int_0^{\tau_i^N} f\left(s, X_i^N(s), \int_{\mathbb{R}^d} w(y)\pi^N(s,dy), u_i\left(s, \boldsymbol{X}^N\right)\right) ds \\ + F\left(\tau_i^N, X_i^N(\tau_i^N)\right)\Bigg],$$

where $\boldsymbol{X}^N = (X_1^N, \ldots, X_N^N)$ and $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \boldsymbol{W}^N, \boldsymbol{X}^N)$ is a solution of Eq. (3.1) under $\boldsymbol{u}$ with initial distribution ν_N,

$$\tau_i^N(\omega) \doteq \tau^{X_i^N}(\omega) \wedge T, \quad \omega \in \Omega,$$

the random time horizon for player $i \in \{1, \ldots, N\}$, and $\pi^N(\cdot)$ the conditional empirical measure process induced by $(X_1^N, \ldots, X_N^N)$. The cost functional is well defined, and it is finite thanks to assumption (H1).

Given a strategy vector $\boldsymbol{u} = (u_1, \ldots, u_N)$ and an individual strategy $v \in \mathcal{U}_N$, let $[\boldsymbol{u}^{-i}, v] \doteq (u_1, \ldots, u_{i-1}, v, u_{i+1}, \ldots, u_N)$ indicate the strategy vector that is obtained from $\boldsymbol{u}$ by replacing u_i, the strategy of player i, with v.

Definition 3.1. Let $\varepsilon \geq 0$. A strategy vector $\boldsymbol{u} = (u_1, \ldots, u_N) \in \mathcal{U}_{fb}^N$ is called an *ε-Nash equilibrium* for the N-player game if for every $i \in \{1, \ldots, N\}$, every $v \in \mathcal{U}_N$ such that $[\boldsymbol{u}^{-i}, v] \in \mathcal{U}_{fb}^N$,

$$J_i^N(\boldsymbol{u}) \leq J_i^N\left([\boldsymbol{u}^{-i}, v]\right) + \varepsilon. \tag{3.4}$$

If $\boldsymbol{u}$ is an ε-Nash equilibrium with $\varepsilon = 0$, then $\boldsymbol{u}$ is called a *Nash equilibrium.*

According to Definition 3.1, we consider the Nash equilibrium property with respect to feedback strategies with full state information (i.e., the states of the vector of individual processes up to current time).

4 Mean field games

Let $\mathcal{M}$ denote the space of measurable flows of measures, that is,

$$\mathcal{M} \doteq \left\{\mathfrak{p} \colon [0,T] \to \mathcal{P}(\mathbb{R}^d) : \mathfrak{p} \text{ is Borel measurable}\right\}.$$

Given a flow of measures $\mathfrak{p} \in \mathcal{M}$ and a feedback strategy $u \in \mathcal{U}_1$, consider the equation

$$X(t) = X(0) + \int_0^t \left(u(s,X) + \bar{b}\left(s, X(s), \int_{\mathbb{R}^d} w(y)\mathfrak{p}(s,dy)\right)\right) ds \\ + \sigma W(t), \quad t \in [0,T], \tag{4.1}$$

where W is a d-dimensional Wiener process defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$.

Let $\mathcal{U}_{fb}$ denote the set of all feedback strategies $u \in \mathcal{U}_1$ such that Eq. (4.1) possesses a solution that is unique in law given any initial distribution with support contained in O.

Proposition 4.1. *Grant (ND) in addition to (H1) – (H5). Then* $\mathcal{U}_{fb} = \mathcal{U}_1$.

Proof. Existence and uniqueness in law are again a consequence of Girsanov's theorem; cf. Proposition 3.1. □

The costs associated with a strategy $u \in \mathcal{U}_{fb}$, a flow of measures $\mathfrak{p}$, and an initial distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O are given by

$$J(\nu, u; \mathfrak{p}) \doteq \mathbf{E}\left[\int_0^\tau f\left(s, X(s), \int_{\mathbb{R}^d} w(y)\mathfrak{p}(s, dy), u(s, X)\right) ds \right. \\ \left. + F\left(\tau, X(\tau)\right)\right],$$

where $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), W, X)$ is a solution of (4.1) under u with initial distribution ν, and $\tau \doteq \tau^X \wedge T$ the random time horizon.

We will measure the minimal costs with respect to a class of stochastic open-loop strategies. To this end, let $\mathcal{A}$ denote the set of all quadruples $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \alpha, W)$ such that $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ is a filtered probability space satisfying the usual hypotheses, ξ an $\mathcal{F}_0$-measurable random variable with values in O, α a Γ-valued $(\mathcal{F}_t)$-progressively measurable process, and W a d-dimensional $(\mathcal{F}_t)$-Wiener process. Any strategy $u \in \mathcal{U}_{fb}$, together with an initial distribution, induces an element of $\mathcal{A}$.

Given $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \alpha, W) \in \mathcal{A}$ and a flow of probability measures $\mathfrak{p} \in \mathcal{M}$, consider the stochastic integral equation

$$X(t) = \xi + \int_0^t \left(\alpha(s) + \bar{b}\left(s, X(s), \int_{\mathbb{R}^d} w(y)\mathfrak{p}(s, dy)\right)\right) ds \\ + \sigma W(t), \quad t \in [0, T]. \tag{4.2}$$

Thanks to (H3), X is determined through Eq. (4.2) with $\mathbf{P}$-probability one; in particular, X is defined on the given stochastic basis. The minimal costs associated with a flow of measures $\mathfrak{p}$ and an initial distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O are now given by

$$V(\nu; \mathfrak{p}) \doteq \inf_{((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \alpha, W) \in \mathcal{A}: \mathbf{P} \circ \xi^{-1} = \nu} \\ \mathbf{E}\left[\int_0^\tau f\left(s, X(s), \int_{\mathbb{R}^d} w(y)\mathfrak{p}(s, dy), \alpha(s)\right) ds + F\left(\tau, X(\tau)\right)\right],$$

where X is the process determined by $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \alpha, W)$ via Eq. (4.2), and $\tau \doteq \tau^X \wedge T$ the random time horizon.

Remark 4.1. Since any admissible feedback strategy induces a stochastic open-loop strategy, we always have

$$\inf_{u \in \mathcal{U}_{fb}} J(\nu, u; \mathfrak{p}) \geq V(\nu; \mathfrak{p}).$$

If the non-degeneracy assumption (ND) holds in addition to (H1) – (H5), if $\bar{b}$ is continuous also in the time variable, and if the flow of measures $\mathfrak{p}$ is such that the mapping $t \mapsto \int w(y)\mathfrak{p}(t, dy)$ is continuous, then

$$\inf_{u \in \mathcal{U}_{fb}} J(\nu, u; \mathfrak{p}) = V(\nu; \mathfrak{p}).$$

This follows, for instance, from the results of El Karoui et al. [1987]; see, in particular, Proposition 2.6 and Remark 2.6.b) as well as Theorem 6.7 and Section 7 therein. Alternatively, one may use time discretization and discrete-time dynamic programming in analogy to Lemma 4.3 in Fischer [2017].

The notion of solution we consider for the mean field game is the following:

Definition 4.1. A *feedback solution of the mean field game* is a triple $(\nu, u, \mathfrak{p})$ such that

(i) $\nu \in \mathcal{P}(\mathbb{R}^d)$ with $\mathrm{supp}(\nu) \subset O$, $u \in \mathcal{U}_{fb}$, and $\mathfrak{p} \in \mathcal{M}$;

(ii) optimality property: strategy u is optimal for $\mathfrak{p}$ and initial distribution ν, that is,
$$J(\nu, u; \mathfrak{p}) = V(\nu; \mathfrak{p});$$

(iii) conditional mean field property: if $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), W, X)$ is a solution of Eq. (4.1) with flow of measures $\mathfrak{p}$, strategy u, and initial distribution ν, then $\mathfrak{p}(t) = \mathbf{P}\left(X(t) \in \cdot \mid \tau^X > t\right)$ for every $t \in [0, T]$ such that $\mathbf{P}\left(\tau^X > t\right) > 0$.

5 Approximate Nash equilibria from the mean field game

Throughout this section, we assume that the non-degeneracy condition (ND) holds. If we have a feedback solution of the mean field game that satisfies a mild regularity condition, then we can construct a sequence of approximate

Nash equilibria for the corresponding N-player game. This approximation result is the content of Theorem 5.1 below.

In order to state the regularity condition, we set, for $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O,

$$\Theta_\nu \doteq \mathrm{Law}\left((\xi + \sigma W(t))_{t\in[0,T]}\right), \tag{5.1}$$

where W is a d-dimensional Wiener process and ξ an independent $\mathbb{R}^d$-valued random variable with distribution ν. Clearly, Θ_ν is well-defined as an element of $\mathcal{P}(\mathcal{X})$.

The proof of Theorem 5.1 below relies on the convergence, uniqueness and regularity results given in the Appendix. The following subsets of probability measures will play an important role as they characterize the possible distributions of the limit processes. For $c \geq 0$, let $\mathcal{Q}_{\nu,c}$ be the set of all laws $\theta \in \mathcal{P}(\mathcal{X})$ such that $\theta = \mathbf{P} \circ X^{-1}$ where

$$X(t) \doteq \xi + \int_0^t v(s)ds + \sigma W(t), \quad t \in [0,T], \tag{5.2}$$

W is an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-Wiener process defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$, ξ is an $\mathbb{R}^d$-valued $\mathcal{F}_0$-measurable random variable with distribution $\mathbf{P} \circ \xi^{-1} = \nu$, and v is an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-progressively measurable bounded process with $\|v\|_\infty \leq c$. Clearly, $\mathcal{Q}_{\nu,0} = \{\Theta_\nu\}$. Also note that $\mathcal{Q}_{\nu,c}$ is compact and that any measure $\theta \in \mathcal{Q}_{\nu,c}$ is equivalent to Θ_ν; see Lemmata C.1 and C.2, respectively, in the Appendix.

We recall that a sequence of symmetric probability measures $(\nu_N)_{N\in\mathbb{N}}$, with $\nu_N \in \mathcal{P}(\mathbb{R}^{N\cdot d})$ for all N, is called ν-chaotic for some $\nu \in \mathcal{P}(\mathbb{R}^d)$ if, for each $k \in \mathbb{N}$ and for any choice of bounded and continuous functions $\psi_i : \mathbb{R}^d \to \mathbb{R}$, $i \in \{1, \ldots, k\}$, we have

$$\lim_{N\to\infty} \int_{\mathbb{R}^{N\cdot d}} \prod_{i=1}^k \psi_i(x_i)\, d\nu_N(x_1, \ldots, x_N) = \prod_{i=1}^k \int_{\mathbb{R}^d} \psi_i(x_i)\nu(dx_i).$$

Theorem 5.1. *Grant (ND) in addition to (H1) – (H5). Suppose the sequence of initial distributions $(\nu_N)_{N\in\mathbb{N}}$ is ν-chaotic for some $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O. If $(\nu, u, \mathfrak{p})$ is a feedback solution of the mean field game such that, for Lebesgue-almost every $t \in [0,T]$,*

$$\Theta_\nu\left(\{\varphi \in \mathcal{X} \,:\, u(t,\cdot) \text{ is discontinuous at } \varphi\}\right) = 0,$$

then

$$u_i^N(t, \boldsymbol{\varphi}) \doteq u(t, \varphi_i), \quad t \in [0,T],\ \boldsymbol{\varphi} = (\varphi_1, \ldots, \varphi_N) \in \mathcal{X}^N,\ i \in \{1, \ldots, N\},$$

defines a strategy vector $\boldsymbol{u}^N = (u_1^N, \dots, u_N^N) \in \mathcal{U}_{fb}^N$. *Moreover, for every* $\varepsilon > 0$, *there exists* $N_0 = N_0(\varepsilon) \in \mathbb{N}$ *such that* $\boldsymbol{u}^N$ *is an* ε*-Nash equilibrium for the* N*-player game whenever* $N \geq N_0$.

Proof. Let $(\nu, u, \mathfrak{p})$ be a feedback solution of the mean field game according to Definition 4.1; such a solution exists by hypothesis. For $N \in \mathbb{N}$, consider the strategy $\boldsymbol{u}^N = (u_1^N, \dots, u_N^N)$ as in the statement. By Proposition 3.1, $\mathcal{U}_{fb}^N = \times^N \mathcal{U}_N$; in particular, $\boldsymbol{u}^N$ is a feedback strategy vector in $\mathcal{U}_{fb}^N$.

Let $\varepsilon > 0$. By symmetry, it is enough to verify the ε-Nash property for player one only. Thus, we have to show that there exists $N_0 = N_0(\varepsilon) \in \mathbb{N}$ such that for all $N \geq N_0$,

$$J_1^N(\boldsymbol{u}^N) \leq \inf_{v \in \mathcal{U}_N} J_1^N\left([\boldsymbol{u}^{N,-1}, v]\right) + \varepsilon. \tag{5.3}$$

Step 1. We rewrite the system dynamics and costs in such way that we can apply the convergence, uniqueness and regularity results of the Appendix. In accordance with (3.2), the definition of b from Section 3, set, for $(t, \varphi, \theta) \in [0, T] \times \mathcal{X} \times \mathcal{P}(\mathcal{X})$,

(5.4)

$$\begin{aligned}\hat{b}(t, \varphi, \theta) &\doteq b\left(t, \varphi, \theta, u(t, \varphi)\right) \\ &= \begin{cases} u(t,\varphi) + \bar{b}\left(t, \varphi(t), \frac{\int w(\tilde{\varphi}(t)) \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t) \theta(d\tilde{\varphi})}{\int \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t) \theta(d\tilde{\varphi})}\right) & \text{if } \theta(\tau > t) > 0, \\ u(t,\varphi) + \bar{b}\left(t, \varphi(t), w(0)\right) & \text{if } \theta(\tau > t) = 0, \end{cases}\end{aligned}$$

where $\tau(\varphi) \doteq \inf\{t \geq 0 : \varphi(t) \notin O\}$. Then, thanks to hypotheses (H1) and (H4), $\hat{b}$ is Borel measurable, progressive, and bounded. Thus, the measurability and boundedness assumptions (M) and (B) of the Appendix hold for $\hat{b}$. Moreover, assumption (C) and (L) of the Appendix, the conditions of almost continuity and of partial Lipschitz continuity, respectively, are satisfied for this choice of $\hat{b}$; see Section D in the Appendix.

Let $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), W, X)$ be a solution of Eq. (4.1) with flow of measures $\mathfrak{p}$, feedback strategy u, and initial distribution ν; such a solution exists and determines a unique measure

$$\theta_* \doteq \mathbf{P} \circ X^{-1}.$$

Notice that $\theta_* \in \mathcal{Q}_{\nu,K}$ whenever $K \geq \|\hat{b}\|_\infty$. In terms of θ_* and the canonical process $\hat{X}$, we can rewrite the costs associated with strategy u, flow of

measures $\mathfrak{p}$, and initial distribution ν as

$$J(\nu,u;\mathfrak{p}) = \mathbf{E}_{\theta_*}\Bigg[\int_0^\tau f\left(s,\hat{X}(s),\int_{\mathbb{R}^d} w(y)\mathfrak{p}(s,dy),u(s,\hat{X})\right)ds \\ + F\left(\tau,\hat{X}(\tau)\right)\Bigg],$$

where $\tau \doteq \tau^{\hat{X}} \wedge T$. Define $\hat{f}\colon [0,T]\times\mathcal{X}\times\mathcal{P}(\mathcal{X})\to\mathbb{R}$ by

$$\hat{f}(t,\varphi,\theta) \doteq \begin{cases} f\left(t,\varphi(t),\frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}{\int \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})},u(t,\varphi)\right) & \text{if } \theta(\tau>t)>0,\\ f\left(t,\varphi(t),w(0),u(t,\varphi)\right) & \text{if } \theta(\tau>t)=0,\end{cases}$$

and a function $\hat{G}\colon \mathcal{X}\times\mathcal{P}(\mathcal{X})\to\mathbb{R}$ by

$$\hat{G}(\varphi,\theta) \doteq \int_0^{\tau(\varphi)\wedge T} \hat{f}\left(s,\varphi,\theta\right)ds + F\big(\tau(\varphi)\wedge T,\varphi(\tau(\varphi)\wedge T)\big). \tag{5.5}$$

Then

$$J(\nu,u;\mathfrak{p}) = \mathbf{E}_{\theta_*}\left[\hat{G}(\hat{X},\theta_*)\right].$$

The function $\hat{G}$ is bounded thanks to hypothesis (H1). Moreover, since $\theta_*\in\mathcal{Q}_{\nu,K}$ for $K\geq\|\hat{b}\|_\infty$ and thanks to (H2),

$$\Theta_\nu\big(\varphi\in\mathcal{X}:\exists\,(\varphi_n,\theta_n)\subset\mathcal{X}\times\mathcal{P}(\mathcal{X}) \text{ such that } \hat{G}(\varphi_n,\theta_n)\not\to\hat{G}(\varphi,\theta_*) \\ \text{while } (\varphi_n,\theta_n)\to(\varphi,\theta_*)\big)=0, \tag{5.6}$$

To be more precise, one checks that $\hat{f}$ satisfies condition (C) as does $\hat{b}$; see Section D of the Appendix. Eq. (5.6) is then a consequence of the dominated convergence theorem, the continuity of F and Part (b) of Lemma C.3. Eq. (5.6) says that $\hat{G}$ is $\Theta_\nu\otimes\delta_{\theta_*}$-almost surely continuous. By Lemma C.2, any $\theta\in\mathcal{Q}_{\nu,c}$, $c\geq 0$, is equivalent to Θ_ν. It follows that $\hat{G}$ is $\theta\otimes\delta_{\theta_*}$-almost surely continuous given any $\theta\in\mathcal{Q}_{\nu,c}$, any $c\geq 0$.

Recall that $(\nu,u,\mathfrak{p})$ is a feedback solution of the mean field game. Thanks to the conditional mean field property in Definition 4.1 and the construction of $\hat{b}$ according to (5.4), we have an alternative characterization of θ_*, namely as a McKean-Vlasov solution of Eq. (A.3); see Definition A.1 in Section A of the Appendix. By Proposition B.1 in Section B of the Appendix and since $\hat{b}$ satisfies assumption (L) there, we actually have that θ_* is the unique McKean-Vlasov solution of Eq. (A.3) with initial distribution ν.

Step 2. For $N\in\mathbb{N}$, let $((\Omega_N,\mathcal{F}_N,(\mathcal{F}_t^N),\mathbf{P}_N),\boldsymbol{W}^N,\boldsymbol{X}^N)$ be a solution of Eq. (3.3) under strategy vector $\boldsymbol{u}^N$ with initial distribution ν_N. Let μ^N

denote the associated empirical measure on the path space $\mathcal{X}$. We are going to show that

$$\lim_{N\to\infty} J_1^N(\boldsymbol{u}^N) = J(\nu, u; \mathfrak{p}).$$

To this end, for $N \in \mathbb{N}$, set

$$\tilde{b}_N(t, \boldsymbol{\varphi}, \theta) \doteq \hat{b}(t, \varphi_1, \theta), \quad (t, \boldsymbol{\varphi}, \theta) \in [0, T] \times \mathcal{X}^N \times \mathcal{P}(\mathcal{X}),$$

where $\hat{b}$ is given by (5.4). From Step 1 we know that $\hat{b}$ satisfies the continuity assumptions (C) and (L) of the Appendix as well as the measurability and boundedness assumptions (M) and (B), respectively. Assumptions (M) and (B) also apply to the coefficients $\tilde{b}_N$. Assumptions (ND) and (I) of the Appendix (non-degeneracy of the diffusion matrix and chaoticity of the initial distributions) hold by hypothesis. Moreover, we have that $((\Omega_N, \mathcal{F}_N, (\mathcal{F}_t^N), \mathbf{P}_N), \boldsymbol{W}^N, \boldsymbol{X}^N)$ is a solution of Eq. (A.1) with initial distribution ν_N. We can therefore apply Lemmata A.1 and A.2 to the sequence of empirical measures $(\mu_N)_{N\in\mathbb{N}}$; in combination with Proposition B.1, this yields

$$\mathbf{P}_N \circ \left(\mu^N\right)^{-1} \overset{N\to\infty}{\longrightarrow} \delta_{\theta_*} \text{ in } \mathcal{P}(\mathcal{P}(\mathcal{X})),$$

where θ_* is the measure identified in Step 1 as the unique McKean-Vlasov solution of Eq. (A.3) with initial distribution ν.

Symmetry of the coefficients, chaoticity of the initial distributions, and the Tanaka-Sznitman theorem [Proposition 2.2 in Sznitman, 1991] now imply that

$$\mathbf{P}_N \circ \left(X_1^N, \mu^N\right)^{-1} \overset{N\to\infty}{\longrightarrow} \theta_* \otimes \delta_{\theta_*} \text{ in } \mathcal{P}(\mathcal{X} \times \mathcal{P}(\mathcal{X})), \tag{5.7}$$

Recalling (5.5) and since $u_1^N(t, \boldsymbol{\varphi}) = u(t, \varphi_1)$, we can rewrite the costs associated with $\boldsymbol{u}^N$ as

$$J_1^N(\boldsymbol{u}^N) = \mathbf{E}_{\mathbf{P}_N}\left[\hat{G}\left(X_1^N, \mu^N\right)\right].$$

By Step 1, we know that $\hat{G}$ is bounded and, since $\theta_* \in \mathcal{Q}_{\nu,K}$ for $K \geq \|\hat{b}\|_\infty$, also that $\hat{G}$ is $\theta_* \otimes \delta_{\theta_*}$-almost surely continuous. By the mapping theorem [Theorem 5.1 in Billingsley, 1968, p. 30], convergence according to (5.7) therefore implies that

$$J_1^N(\boldsymbol{u}^N) = \mathbf{E}_{\mathbf{P}_N}\left[\hat{G}\left(X_1^N, \mu^N\right)\right] \overset{N\to\infty}{\longrightarrow} \mathbf{E}_{\theta_*}\left[\hat{G}(\hat{X}, \theta_*)\right] = J(\nu, u; \mathfrak{p}).$$

Step 3. For $N \in \mathbb{N} \setminus \{1\}$, choose $v_1^N \in \mathcal{U}_N$ such that

$$J_1^N\left([\boldsymbol{u}^{N,-1}, v_1^N]\right) \leq \inf_{v\in\mathcal{U}_N} J_1^N\left([\boldsymbol{u}^{N,-1}, v]\right) + \varepsilon/2.$$

Let $((\tilde{\Omega}_N, \tilde{\mathcal{F}}_N, (\tilde{\mathcal{F}}_t^N), \tilde{\mathbf{P}}_N), \tilde{\boldsymbol{W}}^N, \tilde{\boldsymbol{X}}^N)$ be a solution of Eq. (3.3) under the strategy vector $[\boldsymbol{u}^{N,-1}, v_1^N]$ with initial distribution ν_N. Let $\tilde{\mu}^N$ denote the associated empirical measure on path space. We are going to show that

$$\liminf_{N\to\infty} J_1^N([\boldsymbol{u}^{N,-1}, v_1^N]) \geq V(\nu; \mathfrak{p}).$$

To this end, redefine $\tilde{b}_N$, for $N \in \mathbb{N} \setminus \{1\}$, by

$$\tilde{b}_N(t, \boldsymbol{\varphi}, \theta) \doteq b\left(t, \varphi_1, \theta, v_1^N(t, \boldsymbol{\varphi})\right), \quad (t, \boldsymbol{\varphi}, \theta) \in [0, T] \times \mathcal{X}^N \times \mathcal{P}(\mathcal{X}).$$

Assumptions (M), (B), (ND), and (I) of the Appendix continue to hold. With the new definition of the coefficient $\tilde{b}_N$, $((\tilde{\Omega}_N, \tilde{\mathcal{F}}_N, (\tilde{\mathcal{F}}_t^N), \tilde{\mathbf{P}}_N), \tilde{\boldsymbol{W}}^N, \tilde{\boldsymbol{X}}^N)$ is a solution of Eq. (A.1) with initial distribution ν_N. As in Step 2, we can apply Lemmata A.1 and A.2 in combination with Proposition B.1, but now to the sequence of empirical measures $(\tilde{\mu}_N)_{N\in\mathbb{N}}$. This yields

$$\tilde{\mathbf{P}}_N \circ \left(\tilde{\mu}^N\right)^{-1} \overset{N\to\infty}{\longrightarrow} \delta_{\theta_*} \text{ in } \mathcal{P}(\mathcal{P}(\mathcal{X})), \tag{5.8}$$

where θ_* is still the unique McKean-Vlasov solution of Eq. (A.3) with initial distribution ν found in Steps 1 and 2.

Now, interpret the feedback controls of the first player as stochastic open-loop relaxed controls. To this end, let $\mathcal{R} \doteq \mathcal{R}_\Gamma$ be the set of Γ-valued deterministic relaxed controls over $[0, T]$; see Appendix E. For $N \in \mathbb{N}$, let $\tilde{\rho}_1^N$ be the stochastic relaxed control induced by the feedback strategy v_1^N, that is, $\tilde{\rho}_1^N$ is the $\mathcal{R}$-valued $(\tilde{\mathcal{F}}_t^N)$-adapted random variable determined by

$$\tilde{\rho}_{1,\omega}^N(B \times I) \doteq \int_I \delta_{v_1^N(t, \tilde{X}_1^N(\cdot,\omega))}(B) dt, \quad B \in \mathcal{B}(\Gamma),\ I \in \mathcal{B}([0, T]),\ \omega \in \tilde{\Omega}_N.$$

Let $\mathcal{A}_{rel}$ denote the set of all quadruples $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \rho, W)$ such that $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ is a filtered probability space satisfying the usual hypotheses, ξ an $\mathcal{F}_0$-measurable random variable with values in O, ρ an $\mathcal{R}$-valued $(\mathcal{F}_t)$-adapted random variable, and W a d-dimensional $(\mathcal{F}_t)$-Wiener process. Using the convergence of $(\tilde{\mu}^N)$ according to (5.8) and weak convergence arguments analogous to those of Section A in the Appendix, one verifies that the sequence

$$\left(\tilde{\mathbf{P}}_N \circ \left(\tilde{X}_1^N, \tilde{\rho}_1^N, \tilde{W}_1^N, \tilde{\mu}^N\right)^{-1}\right)_{N\in\mathbb{N}} \text{ is tight in } \mathcal{P}(\mathcal{X} \times \mathcal{R} \times \mathcal{X} \times \mathcal{P}(\mathcal{X})) \tag{5.9}$$

and that its limit points are concentrated on measures $\mathbf{P} \circ (X, \rho, W)^{-1} \otimes \delta_{\theta_*}$ where $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \rho, W) \in \mathcal{A}_{rel}$ for some $\mathcal{F}_0$-measurable $\mathbb{R}^d$-valued

random variable ξ with $\mathbf{P}\circ\xi^{-1}=\nu$, and X is a continuous $(\mathcal{F}_t)$-adapted process satisfying

$$X(t)=\xi+\int_{\Gamma\times[0,t]}\gamma\,\rho(d\gamma,ds)+\int_0^t\bar{b}\left(s,X(s),\int_{\mathbb{R}^d}w(y)\mathfrak{p}(s,dy)\right)ds \\ +\sigma W(t),\quad t\in[0,T], \tag{5.10}$$

with flow of measures $\mathfrak{p}$ given by

$$\mathfrak{p}(t)\doteq\theta_*\left(\hat{X}(t)\in\cdot\,|\,\tau^{\hat{X}}>t\right),\quad t\in[0,T].$$

Notice that X is uniquely determined (with $\mathbf{P}$-probability one) by Eq. (5.10). The stochastic relaxed control in (5.10) actually corresponds to an ordinary stochastic open-loop control, namely to the Γ-valued process

$$\alpha(t,\omega)\doteq\int_\Gamma\gamma\dot{\rho}_{s,\omega}(d\gamma),\quad(t,\omega)\in[0,T]\times\Omega. \tag{5.11}$$

The process X is therefore the unique solution of Eq. (4.2) corresponding to $((\Omega,\mathcal{F},(\mathcal{F}_t),\mathbf{P}),\xi,\alpha,W)\in\mathcal{A}$ and the flow of measures $\mathfrak{p}$ induced by θ_*.

In analogy with Step 1, define $\bar{f}\colon[0,T]\times\mathcal{X}\times\mathcal{P}(\mathcal{X})\times\Gamma\to\mathbb{R}$ by

$$\bar{f}(t,\varphi,\theta,\gamma)\doteq\begin{cases}f\left(t,\varphi(t),\frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}{\int\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})},\gamma\right) & \text{if }\theta(\tau>t)>0,\\ f\left(t,\varphi(t),w(0),u(t,\varphi)\right) & \text{if }\theta(\tau>t)=0,\end{cases}$$

and a function $\bar{G}\colon\mathcal{X}\times\mathcal{P}(\mathcal{X})\times\mathcal{R}\to\mathbb{R}$ by

$$\bar{G}(\varphi,\theta,r)\doteq\int_{\Gamma\times[0,T]}\mathbf{1}_{[0,\tau(\varphi)\wedge T)}(s)\cdot\bar{f}\left(s,\varphi,\theta,\gamma\right)r(d\gamma,ds) \\ +F\big(\tau(\varphi)\wedge T,\varphi(\tau(\varphi)\wedge T)\big). \tag{5.12}$$

Then

$$J_1^N([\boldsymbol{u}^{N,-1},v_1^N])=\mathbf{E}_{\tilde{\mathbf{P}}_N}\left[\bar{G}\left(\tilde{X}_1^N,\tilde{\mu}^N,\tilde{\rho}_1^N\right)\right].$$

The function $\bar{G}$ is bounded thanks to hypothesis (H1). By (H2) and arguments analogous to those of Step 1, one checks that

$$\Theta_\nu\big(\varphi\in\mathcal{X}:\exists\,r\in\mathcal{R},\ (\varphi_n,\theta_n,r_n)\subset\mathcal{X}\times\mathcal{P}(\mathcal{X})\times\mathcal{R}\text{ such that} \\ \hat{G}(\varphi_n,\theta_n,r_n)\not\to\hat{G}(\varphi,\theta_*,r)\text{ while }(\varphi_n,\theta_n,r_n)\to(\varphi,\theta_*,r)\big)=0, \tag{5.13}$$

By Lemma C.2, any $\theta\in\mathcal{Q}_{\nu,c}$, $c\geq0$, is equivalent to Θ_ν. In view of (5.13), it follows that $\bar{G}$ is Q-almost surely continuous given any $Q\in\mathcal{P}(\mathcal{X}\times\mathcal{P}(\mathcal{X})\times\mathcal{R})$ such that $Q\circ\pi^{-1}_{\mathcal{X}\times\mathcal{P}(\mathcal{X})}=\theta\otimes\delta_{\theta_*}$ for some $\theta\in\mathcal{Q}_{\nu,c}$, some $c\geq0$, where $\pi_{\mathcal{X}\times\mathcal{P}(\mathcal{X})}$

denotes the projection of $\mathcal{X} \times \mathcal{P}(\mathcal{X}) \times \mathcal{R}$ onto its first two components. Using the mapping theorem of weak convergence, we conclude that

$$\liminf_{N\to\infty} J_1^N([\boldsymbol{u}^{N,-1}, v_1^N]) \geq \inf_{((\Omega,\mathcal{F},(\mathcal{F}_t),\mathbf{P}),\xi,\rho,W)\in\mathcal{A}_{rel}:\mathbf{P}\circ\xi^{-1}=\nu} \mathbf{E}\left[\bar{G}\left(X,\theta_*,\rho\right)\right].$$

It remains to check that

$$\inf_{((\Omega,\mathcal{F},(\mathcal{F}_t),\mathbf{P}),\xi,\rho,W)\in\mathcal{A}_{rel}:\mathbf{P}\circ\xi^{-1}=\nu} \mathbf{E}\left[\bar{G}\left(X,\theta_*,\rho\right)\right] = V(\nu;\mathfrak{p}).$$

Inequality "$\geq$" holds by definition of $\bar{G}$ and because $\mathcal{A} \subseteq \mathcal{A}_{rel}$ if one identifies ordinary stochastic open-loop controls with the induced stochastic relaxed controls. The opposite inequality is a consequence of what is called "chattering lemma", i.e., the fact that relaxed controls can be approximated by ordinary controls; see, for instance, Theorem 3.5.2 in Kushner [1990, p. 59] or Section 4 in El Karoui et al. [1987]. In our situation, the stochastic relaxed control in Eq. (5.10) corresponds to an ordinary stochastic open-loop control. If the running costs f are additive convex in the control, then inequality "$\leq$" can be verified more directly by using (5.11) and Jensen's inequality.

Step 4. For every $N \in \mathbb{N} \setminus \{1\}$,

$$\begin{aligned} &J_1^N(\boldsymbol{u}^N) - \inf_{v\in\mathcal{U}_N} J_1^N([\boldsymbol{u}^{N,-1}, v]) \\ &\leq J_1^N(\boldsymbol{u}^N) - J(\nu,u;\mathfrak{p}) + J(\nu,u;\mathfrak{p}) - J_1^N([\boldsymbol{u}^{N,-1}, v_1^N]) + \varepsilon/2. \end{aligned}$$

By Steps 2 and 3, there exists $N_0 = N_0(\varepsilon)$ such that for all $N \geq N_0$,

$$J_1^N(\boldsymbol{u}^N) - J(\nu,u;\mathfrak{p}) + V(\nu;\mathfrak{p}) - J_1^N([\boldsymbol{u}^{N,-1}, v_1^N]) \leq \varepsilon/2.$$

Since $(\nu, u; \mathfrak{p})$ is a solution of the mean field game, $J(\nu,u;\mathfrak{p}) = V(\nu;\mathfrak{p})$. It follows that for all $N \geq N_0$,

$$J_1^N(\boldsymbol{u}^N) - \inf_{v\in\mathcal{U}_N} J_1^N([\boldsymbol{u}^{N,-1}, v]) \leq \varepsilon,$$

which establishes (5.3). □

6 Existence of solutions

Throughout this section, hypotheses (H1) – (H5) as well as (ND) are in force. Under these and some additional assumptions, we show that, given any initial distribution, a feedback solution of the mean field game exists in the sense of Definition 4.1; see Theorem 6.2 in Subsection 6.2. To prove Theorem 6.2, we use a BSDE to approach to the control problem for fixed flow of measures

(Subsection 6.1) and a fixed point argument, both in the spirit of Carmona and Lacker [2015].

In Subsection 6.3, we provide conditions that guarantee the existence of a feedback solution to the mean field game with continuous Markovian feedback strategy; see Corollary 6.3. Such a solution satisfies, in particular, the requirements of Theorem 5.1.

As above, let $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O denote the initial distribution. Since ν will be fixed, the dependence on ν will often be omitted in the notation.

6.1 Optimal controls for a given flow of measures

Let $\mathfrak{p} \in \mathcal{M}$ be a given flow of measures, and define the corresponding flow $m_{\mathfrak{p}}$ of means by

$$m_{\mathfrak{p}}(t) \doteq \int_{\mathbb{R}^d} w(y)\mathfrak{p}(t, dy), \quad t \in [0, T]. \tag{6.1}$$

Recall that $\hat{X}$ denotes the canonical process on the path space $\mathcal{X} \doteq \mathbf{C}([0, T]; \mathbb{R}^d)$ (equipped with the sup-norm topology). Also recall (5.1), the definition of Θ_ν. Set $\Omega \doteq \mathcal{X}$, $X \doteq \hat{X}$, let $\mathcal{F}$ be the Θ_ν-completion of the Borel σ-algebra $\mathcal{B}(\mathcal{X})$, and set $\mathbf{P} \doteq \Theta_\nu$ (extended to all null sets). Let $(\mathcal{F}_t)$ be the $\mathbf{P}$-augmentation of the filtration generated by X. Lastly, set

$$\xi(\omega) \doteq X(0, \omega), \quad W(t, \omega) \doteq \sigma^{-1}\left(X(t, \omega) - \xi(\omega)\right), \quad (t, \omega) \in [0, T] \times \Omega,$$

which is well-defined thanks to (ND). Then $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ is a filtered probability space satisfying the usual assumptions, W a d-dimensional $(\mathcal{F}_t)$-Wiener process, and ξ an $\mathcal{F}_0$-measurable random variable with law ν independent of W. Moreover, by construction,

$$X(t) = \xi + \sigma W(t), \quad t \in [0, T], \tag{6.2}$$

and $\mathbf{P} \circ X^{-1} = \Theta_\nu$.

Let $\mathcal{U}$ denote the set of all Γ-valued $(\mathcal{F}_t)$-progressively measurable processes. The filtration $(\mathcal{F}_t)$ is the $\mathbf{P}$-augmentation of the filtration generated by X. Any feedback strategy $u \in \mathcal{U}_{fb}$ therefore induces an element of $\mathcal{U}$ through $\mathrm{u}_t \doteq u(t, X)$, $t \in [0, T]$. On the other hand, given $\mathrm{u} \in \mathcal{U}$, we can find a progressive functional $u \in \mathcal{U}_1$ such that $u(\cdot, X) = \mathrm{u}$. $\mathrm{Leb}_T \otimes \mathbf{P}$-almost surely, where Leb_T denotes Lebesgue measure on $[0, T]$. Lastly, by Proposition 4.1 and (ND), $\mathcal{U}_{fb} = \mathcal{U}_1$. We can therefore identify elements of $\mathcal{U}$ with those of $\mathcal{U}_{fb}$, and vice versa.

Let $\mathrm{u} \in \mathcal{U}$. By Girsanov's theorem and the boundedness of $\bar{b}$, there exists a probability measure $\mathbf{P}^{\mathfrak{p},\mathrm{u}} \sim \mathbf{P}$ and a $\mathbf{P}^{\mathfrak{p},\mathrm{u}}$-Wiener process $W^{\mathfrak{p},\mathrm{u}}$ such that

$$dX(t) = \left(\mathrm{u}_t + \bar{b}\left(t, X(t), m_{\mathfrak{p}}(t)\right)\right) dt + \sigma dW^{\mathfrak{p},\mathrm{u}}(t).$$

The costs associated with u are then given by

$$J^{\mathfrak{p}}(\mathrm{u}) \doteq \mathbf{E}^{\mathfrak{p},\mathrm{u}}\left[\int_0^\tau f\left(t, X(t), m_{\mathfrak{p}}(t), \mathrm{u}_t\right) dt + F(\tau, X_\tau)\right].$$

If $u \in \mathcal{U}_{fb} = \mathcal{U}_1$ is such that $u(\cdot, X) = \mathrm{u}$. $\mathrm{Leb}_T \otimes \mathbf{P}$-almost surely, then (cf. Section 4)

$$J^{\mathfrak{p}}(\mathrm{u}) = J(\nu, u; \mathfrak{p}).$$

Since elements of $\mathcal{U}$ can be identified with those of $\mathcal{U}_{fb}$, and vice versa, we have

$$V^{\mathfrak{p}} \doteq \inf_{\mathrm{u} \in \mathcal{U}} J^{\mathfrak{p}}(\mathrm{u}) = \inf_{u \in \mathcal{U}_{fb}} J(\nu, u; \mathfrak{p}).$$

In view of Remark 4.1, if $\mathfrak{p}$ is such that $t \mapsto m_{\mathfrak{p}}(t)$ is continuous, then

$$V^{\mathfrak{p}} = V(\nu; \mathfrak{p}).$$

Define the Hamiltonian h and the minimized Hamiltonian H as follows

$$\begin{aligned} h(t, x, m, z, \gamma) &\doteq f(t, x, m, \gamma) + z \cdot \sigma^{-1} b(t, x, m, \gamma), \\ H(t, x, m, z) &\doteq \inf_{\gamma \in \Gamma} h(t, x, m, z, \gamma), \end{aligned}$$

where $b(t, x, m, \gamma) \doteq \gamma + \bar{b}(t, x, m)$. The set in which the infimum is attained is denoted by

$$A(t, x, m, z) \doteq \left\{\gamma \in \Gamma : h(t, x, m, z, \gamma) = H(t, x, m, z)\right\}.$$

This set is nonempty since Γ is compact and the function h is continuous in γ thanks to (H4) and (H2), respectively.

Now, consider the following BSDE with random terminal date $\tau \doteq \tau^X \wedge T$:

$$\begin{aligned} Y^{\mathfrak{p}}(t) = F(\tau, X(\tau)) &+ \int_t^T H\left(s, X(s), m_{\mathfrak{p}}(s), Z^{\mathfrak{p}}(s)\right) \mathbf{1}_{(s<\tau)} ds \\ &- \int_t^T Z^{\mathfrak{p}}(s) dW(s), \end{aligned} \tag{6.3}$$

where X follows the forward dynamics (6.2) and W is a Wiener process under $\mathbf{P}$ as before. As in Darling and Pardoux [1997], we adopt the convention that for any solution to the BSDE above:

$$\mathbf{1}_{(s>\tau)} Y(s) = F(\tau, X(\tau)), \quad \mathbf{1}_{(s>\tau)} Z(s) = 0, \quad \mathbf{1}_{(s>\tau)} H(s, x, m, z) = 0.$$

Observe that since the driver is Lipschitz in z and does not depend on y, assumptions (4), (5), (7) and (21) in Darling and Pardoux [1997] are satisfied. Moreover, their assumption (25) is also fulfilled due to the fact that the terminal time τ, the terminal value F, the drift $\bar{b}$ and the running cost f are all bounded; cf. (H1). Hence Theorem 3.4 of Darling and Pardoux [1997] applies, yielding that the BSDE (6.3) has a unique solution $(Y^{\mathfrak{p}}, Z^{\mathfrak{p}})$ in the space of all progressively measurable processes $K \doteq (Y, Z)$ with values in $\mathbb{R}^d \times \mathbb{R}^{d\times d}$ such that

$$\mathbf{E}\left[\int_0^{\tau\wedge T} |K(t)|^2 dt\right] < \infty.$$

For each $\mathrm{u} \in \mathcal{U}$ we can proceed in a similar way to get existence and uniqueness (in the same space of processes as above) of the solution $(Y^{\mathfrak{p},\mathrm{u}}, Z^{\mathfrak{p},\mathrm{u}})$ to the following BSDE

$$\begin{aligned} Y^{\mathfrak{p},\mathrm{u}}(t) = F(\tau, X(\tau)) &+ \int_t^T h\left(s, X(s), m_{\mathfrak{p}}(s), Z^{\mathfrak{p},\mathrm{u}}(s), \mathrm{u}_s\right) \mathbf{1}_{(s<\tau)} ds \\ &- \int_t^T Z^{\mathfrak{p},\mathrm{u}}(s) dW(s). \end{aligned} \tag{6.4}$$

Changing measure from $\mathbf{P}$ to $\mathbf{P}^{\mathfrak{p},\mathrm{u}}$ we obtain

$$\begin{aligned} Y^{\mathfrak{p},\mathrm{u}}(t) = F(\tau, X(\tau)) &+ \int_t^T f\left(s, X(s), m_{\mathfrak{p}}(s), \mathrm{u}_s\right) \mathbf{1}_{(s<\tau)} ds \\ &- \int_t^T Z^{\mathfrak{p},\mathrm{u}}(s) dW^{\mathfrak{p},\mathrm{u}}(s). \end{aligned}$$

Hence, taking conditional expectation with respect to $\mathbf{P}^{\mathfrak{p},\mathrm{u}}$ and $\mathcal{F}_t$ on both sides, we have

$$Y^{\mathfrak{p},\mathrm{u}}(t) = \mathbf{E}^{\mathfrak{p},\mathrm{u}}\left[F(\tau, X(\tau)) + \int_t^T f\left(s, X(s), m_{\mathfrak{p}}(s), \mathrm{u}_s\right) \mathbf{1}_{(s<\tau)} ds \mid \mathcal{F}_t\right],$$

which implies $\mathbf{E}[Y^{\mathfrak{p},\mathrm{u}}(0)] = \mathbf{E}^{\mathfrak{p},\mathrm{u}}[Y^{\mathfrak{p},\mathrm{u}}(0)] = J^{\mathfrak{p}}(\mathrm{u})$. The first equality is due to the fact that the law of ξ is the same under both probability measures.

Since $H \leq h$ and the forward state variable X is the same for both BSDEs (6.3) and (6.4) (so that they both have the same random terminal time), we can apply Darling and Pardoux [1997, Corollary 4.4.2] yielding $Y_t^{\mathfrak{p}} \leq Y_t^{\mathfrak{p},\mathrm{u}}$ a.s. for all $t \in [0, T]$. This implies that $\mathbf{E}[Y^{\mathfrak{p}}(0)] \leq \mathbf{E}[Y^{\mathfrak{p},\mathrm{u}}(0)] = J^{\mathfrak{p}}(\mathrm{u})$ for all $\mathrm{u} \in \mathcal{U}$, and therefore $\mathbf{E}[Y^{\mathfrak{p}}(0)] \leq V^{\mathfrak{p}}$.

By a standard measurable selection argument, there exists a measurable function $\hat{u}\colon [0, T] \times \mathbb{R}^d \times \mathbb{R}^{d_0} \times \mathbb{R}^{d\times d} \to \Gamma$ such that

$$\hat{\mathrm{u}}(t, x, m, z) \in A(t, x, m, z) \text{ for all } (t, x, m, z).$$

Hence, letting

$$\mathrm{u}_t^{\mathfrak{p}} \doteq \hat{\mathrm{u}}\left(t, X(t), m_{\mathfrak{p}}(t), Z^{\mathfrak{p}}(t)\right), \tag{6.5}$$

the uniqueness of the solution of BSDEs with random terminal time as in Darling and Pardoux [1997, Theorem 3.4] gives $Y^{\mathfrak{p}}(t) = Y^{\mathfrak{p},\mathrm{u}^{\mathfrak{p}}}(t)$, which in turn implies $V^{\mathfrak{p}} = J^{\mathfrak{p}}(\mathrm{u}^{\mathfrak{p}})$.

6.2 Existence of solutions of the mean field game

We are going to apply the Brouwer-Schauder-Tychonoff fixed point theorem [for instance, Aliprantis and Border, 2006, 17.56 Corollary], which we recall for the reader's convenience:

Theorem 6.1 (Brouwer-Schauder-Tychonoff)**.** *Let $\mathcal{K}$ be a nonempty compact convex subset of a locally convex Hausdorff space, and let $\Psi : \mathcal{K} \to \mathcal{K}$ be a continuous function. Then the set of fixed points of Ψ is compact and nonempty.*

We need to identify a suitable space $\mathcal{K}$ and a good function Ψ in order to apply the fixed theorem above to our setting and obtain the existence of a solution to the mean field game.

In addition to (H1) – (H5) and (ND), let us make the following assumptions:

Assumptions 6.1.

(i) The coefficients $\bar{b}$ and f are continuous (jointly in all their variables).

(ii) The set $A(t, x, m, z)$ is a singleton for all $(t, x, m, z) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^{d_0} \times \mathbb{R}^{d \times d}$.

By Assumption 6.1(ii), the function $\hat{u}$ appearing in (6.5) is uniquely determined. Moreover, $\hat{u}$ is continuous:

Lemma 6.1. *The function $\hat{u}$ is continuous (jointly in all its variables).*

Proof. Assumption 6.1(i) implies that the (pre-)Hamiltonian h is (jointly) continuous. The continuity of $\hat{u}$ now follows from Berge's Maximum Theorem [for instance, Aliprantis and Border, 2006, Theorem 17.31], Assumption 6.1(ii) and the compactness of Γ according to (H4). □

We will apply the fixed point theorem to a restriction of the mapping

$$\Psi \colon \mathcal{P}(\mathcal{X}) \ni \mu \mapsto \mathbf{P}^{\mu} \circ X^{-1} \in \mathcal{P}(\mathcal{X}) \tag{6.6}$$

where $\mathbf{P}^\mu \doteq \mathbf{P}^{\mathfrak{p}^\mu, \hat{\mathrm{u}}^\mu}$ according to Subsection 6.1 with $\mathfrak{p}^\mu$ and $\hat{u}^\mu$ defined as follows:

$$\mathfrak{p}^\mu(t,\cdot) \doteq \mu\left(X(t) \in \cdot \mid \tau^X > t\right), \tag{6.7a}$$

$$\hat{\mathrm{u}}^\mu_t \doteq \mathrm{u}^{\mathfrak{p}^\mu}_t = \hat{\mathrm{u}}\left(t, X(t), \int w(y)\mathfrak{p}^\mu(t,dy), Z^{\mathfrak{p}^\mu}(t)\right), \quad t \in [0,T]. \tag{6.7b}$$

Thanks to Assumption 6.1(ii), the optimal control process $\hat{\mathrm{u}}^\mu$ is unique $\mathrm{Leb}_T \otimes \mathbf{P}$-almost surely. Let $\mathbf{E}_\mu$ denote expectation with respect to μ (not $\mathbf{P}^\mu$!).

Now, we need to specify the subspace of $\mathcal{P}(\mathcal{X})$ with the right properties for applying the fixed point theorem. Recall the definition of $\mathcal{Q}_{\nu,c}$ with $c \geq 0$ from Section 5. Choose $K > 0$ such that $K \geq \|b\|_\infty$, and set

$$\mathcal{K} \doteq \mathcal{Q}_{\nu,K}.$$

Then $\mathcal{K}$ is non-empty and convex, and it is compact with respect to the topology inherited from $\mathcal{P}(\mathcal{X})$ (i.e., the topology of weak convergence of measures), see Lemma C.1 in the Appendix. Moreover, $\mathcal{K}$ can be seen as a subset of the dual space $\mathbf{C}_b(\mathcal{X})^*$, where $\mathbf{C}_b(\mathcal{X})$ is the space of all continuous bounded functions on $\mathcal{X}$. The dual space $\mathbf{C}_b(\mathcal{X})^*$, equipped with the weak* topology $\sigma(\mathbf{C}_b(\mathcal{X})^*, \mathbf{C}_b(\mathcal{X}))$, is a locally convex topological vector space, inducing the weak convergence on $\mathcal{P}(\mathcal{X}) \supset \mathcal{K}$.

If $\mu \in \mathcal{K}$, then the flow of conditional means induced by μ is continuous:

Lemma 6.2. *Let $\mu \in \mathcal{K}$. Then the mapping $t \mapsto m^\mu(t) \doteq \int w(y)\mathfrak{p}^\mu(t,dy)$ is continuous, where $\mathfrak{p}^\mu$ is the flow of conditional measures induced by μ according to* (6.7a).

Proof. By Lemma C.2, $\inf_{t\in[0,T]} \mu(\tau^X > t) > 0$. Therefore and by construction, for every $t \in [0,T]$,

$$m^\mu(t) = \frac{1}{\mu(\tau^X > t)} \mathbf{E}_\mu\left[w(X(t))\mathbf{1}_{[0,\tau^X)}(t)\right].$$

Let $t \in [0,T]$, and let $(t_n) \subset [0,T]$ be such that $t_n \to t$ as $n \to \infty$. For $\omega \in \Omega$, the mapping $s \mapsto \mathbf{1}_{[0,\tau^X(\omega))}(s)$ is discontinuous only at $s = \tau^X(\omega)$. By part (c) of Lemma C.3, $\mu\left(\tau^X = t\right) = 0$. Since w is bounded and continuous and X has continuous trajectories, we have with μ-probability one,

$$\mathbf{1}_{[0,\tau^X)}(t_n) \overset{n\to\infty}{\longrightarrow} \mathbf{1}_{[0,\tau^X)}(t), \quad w\left(X(t_n)\right)\mathbf{1}_{[0,\tau^X)}(t_n) \overset{n\to\infty}{\longrightarrow} w\left(X(t)\right)\mathbf{1}_{[0,\tau^X)}(t).$$

The dominated convergence theorem now implies $\lim_{n\to\infty} m(t_n) = m(t)$. □

In order to apply Theorem 6.1 above, we have to show that $\Psi\colon \mathcal{K} \to \mathcal{K}$ is (sequentially) continuous. Let (μ^n) be a sequence in $\mathcal{K}$ converging to some measure $\mu \equiv \mu^\infty \in \mathcal{K}$. We want to show that $\Psi(\mu^n) \to \Psi(\mu)$ in $\mathcal{K} \subset \mathcal{P}(\mathcal{X})$. To ease the notation, we set

$$\mathfrak{p}^n \doteq \mathfrak{p}^{\mu^n}, \qquad Z^n \doteq Z^{\mathfrak{p}^n}, \qquad \mathrm{u}_t^n \doteq \mathrm{u}_t^{\mathfrak{p}^n}, \qquad 1 \le n \le \infty.$$

We proceed as in Carmona and Lacker [2015, Proof of Theorem 3.5] and show that

$$\mathcal{H}(\mathbf{P}^\mu \mid \mathbf{P}^{\mu^n}) \overset{n\to\infty}{\longrightarrow} 0, \tag{6.8}$$

where $\mathcal{H}(\cdot \mid \cdot)$ denotes the relative entropy:

$$\mathcal{H}(\tilde{\mu} \mid \bar{\mu}) \doteq \begin{cases} \int \log\left(\frac{d\tilde{\mu}}{d\bar{\mu}}\right) d\tilde{\mu} & \text{if } \tilde{\mu} \text{ absolutely continuous w.r.t. } \bar{\mu} \\ \infty & \text{otherwise.} \end{cases}$$

Observe that (6.8) implies $\Psi(\mu^n) \to \Psi(\mu)$ in the topology of weak convergence as well as in the topology of convergence in total variation.

In our situation, denoting by $\mathcal{E}(\cdot)$ the stochastic exponential of a local martingale, we have

$$\frac{d\mathbf{P}^{\mu^n}}{d\mathbf{P}^\mu} = \mathcal{E}\left(\int_0^\cdot \sigma^{-1}\left(b(t, X(t), m^n(t), \mathrm{u}_t^n) - b(t, X(t), m^\infty(t), \mathrm{u}_t^\infty)\right) \mathbf{1}_{[0,\tau^X)}(t)\right.$$
$$\left. dW^{\mathfrak{p}^\infty}(t)\right)(T),$$

where m^n denotes the flow of means induced by $\mathfrak{p}^n$ according to (6.1):

$$m^n(t) \doteq \int w(y)\mathfrak{p}^n(t, dy), \quad 1 \le n \le \infty.$$

Since $\sigma^{-1}b$ is bounded, we obtain

$$\mathcal{H}(\mathbf{P}^\mu \mid \mathbf{P}^{\mu^n}) = -\mathbf{E}^{\mathbf{P}^\mu}\left[\log \frac{d\mathbf{P}^{\mu^n}}{d\mathbf{P}^\mu}\right]$$
$$= \frac{1}{2}\mathbf{E}^{\mathbf{P}^\mu}\left[\int_0^\tau \left|\sigma^{-1}b\left(t, X(t), m^n(t), \mathrm{u}_t^n\right) - \sigma^{-1}b\left(t, X(t), m^\infty(t), \mathrm{u}_t^\infty\right)\right|^2 dt\right].$$

By (H2), the map $b(t, x, \cdot, \cdot)$ is continuous for all (t, x). Hence, we can conclude by an application of the bounded convergence theorem provided we show that

$$(m^n, \mathrm{u}^n) \overset{n\to\infty}{\longrightarrow} (m^\infty, \mathrm{u}^\infty) \text{ in } \mathrm{Leb}_T \otimes \mathbf{P}^\mu\text{-measure.} \tag{6.9}$$

First, we show that

$$m^n \overset{n\to\infty}{\longrightarrow} m^\infty \text{ in } \mathrm{Leb}_T\text{-measure.} \tag{6.10}$$

By Lemma C.2, $\inf_{n\in\mathbb{N}\cup\{\infty\}} \inf_{t\in[0,T]} \mu^n\left(\tau^X > t\right) > 0$. Hence, by construction, for every $t \in [0,T]$,

$$m^n(t) = \frac{1}{\mu^n(\tau^X > t)} \mathbf{E}_{\mu^n}\left[w(X(t))\mathbf{1}_{[0,\tau^X)}(t)\right], \quad n \in \mathbb{N}.$$

The mapping $\mathbf{1}_{[0,\tau^X)}(t)$ is μ-almost surely continuous for each fixed t by part (d) of Lemma C.3. By hypothesis, w is bounded and continuous. The mapping theorem [Theorem 5.1 in Billingsley, 1968, p. 30] therefore implies that

$$\mathbf{E}_{\mu^n}\left[w(X(t))\mathbf{1}_{[0,\tau^X)}(t)\right] \to \mathbf{E}_{\mu}\left[w(X(t))\mathbf{1}_{[0,\tau^X)}(t)\right],$$

as well as $\mu^n(\tau^X > t) \to \mu(\tau^X > t)$. This gives that $m^n(t) \to m^\infty(t)$ for all t, hence $m^n \to m^\infty$ pointwise and in Leb_T-measure.

Next, let us check that

$$\lim_{n\to\infty} \mathbf{E}\left[\int_0^T |Z^n(t) - Z^\infty(t)|^2\, dt\right] = 0. \tag{6.11}$$

To this end, we use a stability result for BSDEs with random terminal time as in Briand and Hu [1998, Theorem 2.4]. According to that result, in order to get the convergence (6.11), it suffices to show that

$$\mathbf{E}\left[\int_0^\tau |H(s, X(s), m^n(s), Z^\infty(s)) - H(s, X(s), m^\infty(s), Z^\infty(s))|^2 ds\right] \overset{n\to\infty}{\longrightarrow} 0.$$

Since the functions f and b in the definition of the minimized Hamiltonian H are bounded and the controls take values in a compact set, we easily have

$$|H(s, X(s), m^n(s), Z^\infty(s)) - H(s, X(s), m^\infty(s), Z^\infty(s))|^2 \le c_1 + c_2|Z^\infty(s)|^2,$$

where c_1, c_2 are some positive constants. Moreover, $\mathbf{E}[\int_0^\tau |Z^\infty(t)|^2 dt] < \infty$ (cf. Darling and Pardoux [1997, Theorem 3.4]). We can therefore apply the dominated convergence theorem to obtain (6.11). The convergence in $\mathrm{Leb}_T \times \mathbf{P}^\mu$-measure follows from the equivalence $\mathbf{P}^\mu \sim \mathbf{P}$.

In view of (6.10), in order to establish (6.9), it remains to show that $\mathrm{u}_t^n \to \mathrm{u}_t^\infty$ $\mathbf{P}^\mu$-almost surely for all $t \in [0,T]$. Notice that for all $n \in \mathbb{N}\cup\{\infty\}$

$$\hat{\mathrm{u}}_t^n = \hat{\mathrm{u}}\left(t, X(t), m^n(t), Z^n(t)\right).$$

By Lemma 6.1, $\hat{\mathrm{u}}$ is jointly continuous in all its variables. Equations (6.10) and (6.11) together imply that $(m^n, Z^n) \to (m^\infty, Z^\infty)$ in $\mathrm{Leb}_T \times \mathbf{P}^\mu$-measure. It follows that $\hat{\mathrm{u}}^n \to \hat{\mathrm{u}}^\infty$ in $\mathrm{Leb}_T \times \mathbf{P}^\mu$-measure as well.

Now, we can state and prove the main result of this section.

Theorem 6.2. *Under (H1) – (H5), (ND), and Assumptions 6.1, there exists a feedback solution of the mean field game.*

Proof. In view of the discussion above, we can apply Theorem 6.1 (the Brouwer-Schauder-Tychonoff fixed point theorem) to the map $\Psi : \mathcal{K} \to \mathcal{K}$. Thus, there exists $\mu \in \mathcal{K}$ such that $\Psi(\mu) = \mu$. Let $\mathfrak{p} \doteq \mathfrak{p}^\mu$ be the flow of conditional measures induced by the fixed point μ according to (6.7a). Let $m_\mathfrak{p}$ be the flow of (conditional) means induced by $\mathfrak{p}$ according to (6.1), and let $(Y^\mathfrak{p}, Z^\mathfrak{p})$ be the unique solution of the BSDE (6.3). Notice that $Z^\mathfrak{p}$ is progressively measurable with respect to $(\mathcal{F}_t)$, the $\mathbf{P}$-augmentation of the filtration generated by the canonical process $X = \hat{X}$. Hence we can find $u \in \mathcal{U}_1$ such that

$$u(\cdot, X) = \hat{u}\left(\cdot, X(\cdot), m_\mathfrak{p}(\cdot), Z^\mathfrak{p}(\cdot)\right) \quad \text{Leb}_T \otimes \mathbf{P}\text{-almost surely.}$$

Recalling the discussion of Subsection 6.1 and the fact that in our situation $\mathcal{U}_1 = \mathcal{U}_{fb}$, we see that

$$J(\nu, u; \mathfrak{p}) = V^\mathfrak{p}.$$

The flow of means $m_\mathfrak{p}$ is continuous thanks to Lemma 6.2. This implies, by Remark 4.1, that $V^\mathfrak{p} = V(\nu; \mathfrak{p})$. The triple $(\nu, u, \mathfrak{p})$ therefore satisfies the optimality property of Definition 4.1. By construction of Ψ and the fixed point property of μ, $(\nu, u, \mathfrak{p})$ also satisfies the conditional mean field property of Definition 4.1. It follows that $(\nu, u, \mathfrak{p})$ is a feedback solution of the mean field game. □

6.3 Continuity of optimal controls

Here, we work under all the hypotheses of the previous subsections plus a couple of additional assumptions to be introduced below. Let $\mathfrak{p} \in \mathcal{M}$ be a flow of measures such that the induced flow of means $m_\mathfrak{p}$ is continuous, that is, we assume that the mapping

$$t \mapsto m_\mathfrak{p}(t) \doteq \int_{\mathbb{R}^d} w(y)\mathfrak{p}(t, dy)$$

is continuous. From Lemma 6.2 we know that this is the case if $\mathfrak{p}$ is the flow of measures induced by some $\mu \in \mathcal{K}$.

The aim here is to show that $\hat{\mathrm{u}}(\cdot, X(\cdot), m_\mathfrak{p}(\cdot), Z^\mathfrak{p}(\cdot))$, the optimal control process appearing in (6.5), corresponds to a Markovian feedback strategy that is continuous both in its time and its state variable. In view of Lemma 6.1 and the continuity of $m_\mathfrak{p}$, it suffices to show that $Z^\mathfrak{p}(t)$ can be expressed as a continuous function of time t and the forward state variable $X(t)$.

Set, for $(t,x,z) \in [0,T] \times \mathbb{R}^d \times \mathbb{R}^{d\times d}$,

$$\tilde{h}(t,x,z) \doteq h(t,x,m(t),z,\hat{\mathrm{u}}(t,x,m_{\mathfrak{p}}(t),z)).$$

Notice that $\tilde{h}(t,x,z) = H(t,x,m_{\mathfrak{p}}(t),z)$. Moreover, by Assumptions 6.1, Lemma 6.1, and the continuity of $m_{\mathfrak{p}}$, we have that $\tilde{h}$ is continuous (jointly in all its variables). By (H1) and (H4), $\tilde{h}(\cdot,\cdot,z)$ is bounded for each $z \in \mathbb{R}^{d\times d}$.

Consider the following BSDE under $\mathbf{P}$:

$$Y(t) = F(\tau, X(\tau)) + \int_t^T \tilde{h}\,(s, X(s), Z(s))\,\mathbf{1}_{(s<\tau)} ds - \int_t^T Z(s) dW(s), \tag{6.12}$$

with the same convention regarding times exceeding τ as in Subsection 6.1. The BSDE above possesses a unique L^2-solution (Y,Z), which coincides with the unique solution $(Y^{\mathfrak{p}}, Z^{\mathfrak{p}})$ of the BSDE (6.3). To prove the continuity property, we adopt the following procedure: first, we consider the PDE corresponding to the BSDE (6.12) and prove some regularity implying, in particular, the continuity of the gradient of the solution; second, we provide a verification argument essentially identifying the solution of the PDE and its gradient with the unique solution (Y,Z) of (6.12). This is performed in the proof of the next proposition. Before that, we have to introduce some notation on Sobolev and Hölder-type norms.

Set $D \doteq [0,T) \times O$, and let $\partial_p D \doteq ([0,T) \times \partial O) \cup (\{T\} \times \bar{O})$ be the parabolic boundary of D. Moreover, let $\mathcal{A}$ denote the (parabolic) Dynkin operator associated with X under $\mathbf{P}$, that is,

$$\mathcal{A} \doteq \partial_t + \frac{1}{2}\mathrm{Tr}(a\partial^2_{xx}),$$

where $a \doteq \sigma\sigma^{\mathsf{T}}$. Let $L^\lambda(D)$ denote the space of λ-th power integrable functions on D, with $\|\cdot\|_{\lambda,D}$ norm in $L^\lambda(D)$. For $1 < \lambda < \infty$, let $\mathcal{H}^\lambda(D)$ denote the space of functions ψ that together with their generalized partial derivatives $\partial_t\psi$, $\partial_{x_i}\psi$, $\partial_{x_ix_j}\psi$, $i,j = 1,\dots,n$, are in $L^\lambda(D)$. In $\mathcal{H}^\lambda(D)$ let us introduce the norm (of Sobolev type)

$$\|\psi\|^{(2)}_{\lambda,D} \doteq \|\psi\|_{\lambda,D} + \|\partial_t\psi\|_{\lambda,D} + \sum_{i=1}^d \|\partial_{x_i}\psi\|_{\lambda,D} + \sum_{i,j=1}^d \|\partial_{x_ix_j}\psi\|_{\lambda,D}. \tag{6.13}$$

Let us also consider Hölder-type norms as follows. For $0 < \alpha \le 1$, let

$$\|\psi\|_D \doteq \sup_{(t,x)\in D} |\psi(t,x)|,$$

$$|\psi|^\alpha_D \doteq \|\psi\|_D + \sup_{x,y\in\bar{O},t\in[0,T]} \frac{|\psi(t,x)-\psi(t,y)|}{|x-y|^\alpha} + \sup_{s,t\in[0,T],x\in\bar{O}} \frac{|\psi(s,x)-\psi(t,x)|}{|s-t|^{\alpha/2}}.$$

Finally, we let

$$|\psi|_D^{1+\alpha} \doteq |\psi|_D^{\alpha} + \sum_{i=1}^{d} |\partial_{x_i}\psi|_D^{\alpha}.$$

Assumptions 6.2.

(i) $\bar{b}$, f belong to $\mathbf{C}^{0,1}([0,T]\times(\bar{O}\times\mathbb{R}^{d_0}))$ and $\mathbf{C}^{0,1}([0,T]\times(\bar{O}\times\mathbb{R}^{d_0}\times\Gamma))$, respectively. Thus, $b \in \mathbf{C}^{0,1}([0,T]\times(\bar{O}\times\mathbb{R}^{d_0}\times\Gamma))$.

(ii) $F(T,\cdot) \in \mathbf{C}^2(\bar{O})$ and $F(t,x) = \tilde{F}(t,x)$ for all $t \in [0,T]$ and $x \in \partial O$, where $\tilde{F} \in \mathbf{C}^{1,2}(\bar{D})$. Moreover, for some $\lambda > \frac{d+2}{2}$,

$$\|F\|_{\lambda,\partial_p D}^{(2)} < \infty.$$

Assumption 6.2(ii) on the terminal costs F corresponds to assumption (E7) in Fleming and Rishel [1975, Appendix E].

Proposition 6.1. *Grant (H1) – (H5), (ND) as well as Assumptions 6.1 and 6.2. The optimal control is then given by*

$$\mathrm{u}_t^{\mathfrak{p}} = \hat{\mathrm{u}}(t, X(t), m_{\mathfrak{p}}(t), \partial_x\varphi(t,X(t))\sigma), \quad t \in [0,\tau],$$

where the function φ is the unique solution in $\mathcal{H}^\lambda(D)$ of the Cauchy problem

$$-\mathcal{A}\varphi - \tilde{h}(\cdot,\varphi,\partial_x\varphi\sigma) = 0 \text{ on } D, \qquad \varphi = F \text{ on } \partial_p O, \tag{6.14}$$

and it satisfies

$$\|\varphi\|_{\lambda,D}^{(2)} \le M\left(\|f\|_{\lambda,D} + \|F\|_{\lambda,\partial_p D}^{(2)}\right), \tag{6.15}$$

with constant M depending only on bounds on b and σ on D. Moreover,

$$|\varphi|_D^{1+\alpha} \le M'\|\varphi\|_{\lambda,D}^{(2)}, \quad \alpha = 1 - \frac{d+2}{\lambda}, \tag{6.16}$$

for some constant M' depending on D and λ.

As a consequence, the function $(t,x) \mapsto \hat{\mathrm{u}}(t,x,m(t),\partial_x\varphi(t,x)\sigma)$ is continuous over $[0,T]\times\bar{O}$.

Proof. First, consider the Cauchy problem (6.14). The existence and uniqueness of a solution in $\mathcal{H}^\lambda(D)$ satisfying the estimate (6.15) can be found in Fleming and Rishel [1975, Appendix E, Proof of Theorem VI.6.1, pp. 208-209]. Observe that our assumptions together with the continuity of $\tilde{h}$ imply the hypotheses (6.1)–(6.3) in Fleming and Rishel [1975, p. 167], with the exception of (6.3)(b). In fact, in our setting the drift b and the running cost

f are only continuous in t and not necessarily $\mathbf{C}^1$; cf. Assumption 6.2(i). However, a careful inspection of the proof by Fleming and Rishel shows that we can apply their approximation argument excluding the last sentence. Hence, there exists a non-increasing sequence (φ_n) converging pointwise to the unique solution $\varphi \in \mathcal{H}^\lambda(D)$ of the Cauchy problem (6.14). Moreover, each φ_n satisfies the inequality (6.15). Using Ladyženskaja et al. [1968, p. 80, p. 342], the estimate (6.15) gives (6.16) if $\lambda > d+2$, and in this case $\alpha = 1 - \frac{d+2}{\lambda}$ where the constant M' depends on D and λ. Notice that the upper bounds in (6.15) and (6.16) are both uniform in n. Therefore, letting $n \to \infty$, we have that the solution φ satisfies the same bounds. In particular, (6.16) yields that the gradient $\partial_x \varphi$ is well-defined in the classical sense and it is continuous on D.

To complete the proof, it suffices to show that the pair of processes $(\varphi(t, X(t)), \partial_x \varphi(t, X(t))\sigma)_{t\in[0,\tau]}$ coincides with the unique L^2 solution of the BSDE (6.12). We use a verification argument based on a generalization of Itô's formula as in Krylov [1980, Theorem 1, Sect. 2.10]. Applying that formula between t and τ, on the random set $\{t \le \tau\}$, we have with $\mathbf{P}$-probability one

$$\begin{aligned}\varphi(\tau, X(\tau)) = \varphi(t, X(t)) &+ \int_t^\tau \mathcal{A}\varphi(s, X(s))ds \\ &+ \int_t^\tau \partial_x \varphi(s, X(s))\sigma dW(s).\end{aligned}$$

Being φ a solution in $\mathcal{H}^\lambda(D)$ of the Cauchy problem (6.14), we have (by re-arranging the terms)

$$\begin{aligned}\varphi(t, X(t)) = F(\tau, X(\tau)) &- \int_t^\tau \tilde{h}(s, X(s), \partial_x \varphi(s, X(s))\sigma)ds \\ &- \int_t^\tau \partial_x \varphi(s, X(s))\sigma dW(s),\end{aligned}$$

which gives, by uniqueness of the solution (Y, Z) of the BSDE (6.12), that

$$Y(t) = \varphi(t, X(t)), \quad Z(t) = \partial_x \varphi(t, X(t))\sigma,$$

on the event $\{t \le \tau\}$ for all $t \in [0, T]$. □

Proposition 6.1 and Theorem 6.2 yield (together with Lemma 6.2) the following corollary. Observe that the control actions can be chosen in an arbitrary way when the state process is outside $\mathrm{cl}(O)$.

Corollary 6.3. *Grant (H1) – (H5), (ND) as well as Assumptions 6.1 and 6.2. Then there exist a feedback solution of the mean field game $(\nu, u, \mathfrak{p})$ and a continuous function $\tilde{u}\colon [0,T] \times \mathbb{R}^d \to \Gamma$ such that*

$$u(t,\varphi) = \tilde{u}(t,\varphi(t)) \text{ for all } (t,\varphi) \in [0,T] \times \mathcal{X}.$$

In particular, u is continuous on $[0,T] \times \mathcal{X}$.

We end this section by sketching the PDE approach to mean field games with absorption. For simplicity, we only consider a setting analogous to that of Lasry and Lions [2006b], where the dynamics are of calculus-of-variation-type with additive noise and the running costs split into two parts, one depending on the control, the other on the measure variable. Thus, we assume that

$$\sigma \equiv \sigma \operatorname{Id}_d, \qquad \bar{b}(t,x,y) \equiv 0, \qquad f(t,x,y,\gamma) = f_0(t,x,\gamma) + f_1(t,x,y)$$

for some scalar constant $\sigma > 0$ and suitable functions $f_0\colon [0,T] \times \mathbb{R}^d \times \Gamma \to \mathbb{R}$ and $f_1\colon [0,T] \times \mathbb{R}^d \times \mathbb{R}^{d_0} \to \mathbb{R}$. Let us also assume that the initial distribution ν is absolutely continuous with respect to Lebesgue measure with density m_0.

Let $(\nu, u, \mathfrak{p})$ be a solution according to Corollary 6.3, and let $\tilde{u}$ be the associated Markov feedback strategy.

Set $H(t,x,z) \doteq \max_{\gamma \in \Gamma} \{-\gamma \cdot z - f_0(t,x,\gamma)\}$. Let V be the unique solution of the Hamilton-Jacobi Bellman equation

(6.17)
$$-\partial_t V - \frac{\sigma^2}{2}\Delta V + H(t,x,\nabla V) = f_1\left(t,x,\int w(y)\mathfrak{p}(t,dy)\right) \text{ in } [0,T) \times O$$

with boundary condition $V(t,x) = F(t,x)$ in $\{T\} \times \operatorname{cl}(O) \cup [0,T) \times \partial O$, and let m be the unique solution of the Kolmogorov forward equation

$$\partial_t m - \frac{\sigma^2}{2}\Delta m + \operatorname{div}\left(m(t,x) \cdot \tilde{u}(t,x)\right) = 0 \text{ in } (0,T] \times O \tag{6.18}$$

with initial condition $m(0,x) = m_0(x)$ and boundary condition $m(t,x) = 0$ in $(0,T] \times \partial O$. Then

$$\tilde{u}(t,x) = -D_z H\left(t,x,\nabla V(t,x)\right), \quad \mathfrak{p}(t,dx) = \frac{m(t,x)}{\int_O m(t,y)dy}dx. \tag{6.19}$$

A non-standard feature in (6.19) is the renormalization by the factor $1/\int_O m(t,y)dy$. This is related to the fact that $m(t,\cdot)$ will be the density of a sub-probability measure, not necessarily a probability measure if $t > 0$. Relationship (6.19) clearly allows to eliminate the measure flow $\mathfrak{p}$ from

Eq. (6.17) and the feedback control $\tilde{u}$ from Eq. (6.18), yielding a coupled system of (backward) Hamilton-Jacobi Bellman equation and Kolmogorov forward equation. This gives the PDE system for the mean field game with absorption.

7 A counterexample

If the diffusion coefficient is degenerate and no additional controllability assumptions are made, then a regular feedback solution of the mean field game need not induce a sequence of approximate Nash equilibria with vanishing error.

The following (counter-)example is constructed in such a way that, for a certain initial condition, the individual state processes of the N-player games cannot leave the set O of non-absorbing states before terminal time, while this is possible in the mean field game limit. Exit before terminal time can actually occur only at one fixed instant, if at all. With the choice of the costs we make, leaving before terminal time is preferable. The individual strategies that the solution of the mean field game induces in the N-player games will therefore push the state processes towards the boundary if they start from certain initial points, but will fail to make them exit (for N odd, the probability of exit equals zero, for N even it will tend to zero as $N \to \infty$). The failure to leave the set O in the N-player games will produce costs that are higher than those of an alternative individual response strategy which does not try to make its state process exit.

To achieve the described effect, we choose deterministic state dynamics; the only source of randomness comes from the initial conditions. Individual states will have three components, where one component (the third) corresponds to time, while the second component simply keeps its own initial condition. Only the first component is controlled; its evolution is driven, apart from the control, by an average over the second component of the states of all players (a mean with respect to the measure variable in the limit). That interaction term is responsible for the different controllability of the N-player systems with respect to the limit system. Notice that the construction also relies on the particular choice of the initial conditions, which are singular with respect to Lebesgue measure. To be specific, consider the following data for our systems:

- dimensions: $d = 3$, $d_0 = 1$;
- time horizon: $T = 2$;

- set of control actions: $\Gamma \doteq \{\gamma \in \mathbb{R}^3 : \gamma_1 \in [-1,1], \gamma_2 = 0 = \gamma_3\}$;
- set of non-absorbing states

$$O \doteq \Big\{x \in \mathbb{R}^3 : -4 < x_1\,,\ -2 < x_2 < 2\,,\ -1 < x_3 < \frac{11}{5}\,,\\ x_1 < 1 + e^{x_3 - 1}\Big\}$$

- measure integrand: w bounded Lipschitz and such that $w(x) = x_2$ for all $x \in \mathrm{cl}(O)$;
- drift function:

$$\bar{b}(t,x,y) \doteq \begin{pmatrix} -|y| \wedge \frac{1}{4} \\ 0 \\ 1 \end{pmatrix}, \quad (t,x,y) \in [0,2] \times \mathbb{R}^3 \times \mathbb{R};$$

- dispersion coefficient: $\sigma \equiv 0$;
- running costs: $f \equiv 1$;
- terminal costs: F non-negative bounded Lipschitz and such that

$$F(t,x) = 1 + \frac{x_3}{12} \cdot x_1 \text{ for all } (t,x) \in [0,2] \times \mathrm{cl}(O).$$

Remark 7.1. The set of non-absorbing states O defined above has a boundary that is only piecewise smooth. Let $B \subset \mathbb{R}^3$ be the line segment given by

$$B \doteq \big\{x \in \mathbb{R}^3 : x_1 = 2\,,\ -1 \le x_2 \le 1\,,\ x_3 = 1\big\}\,.$$

As will become clear below, the (counter-)example works for any bounded open set O as long as O contains

$$\Big\{x \in \mathbb{R}^3 : -1 \le x_2 \le 1\,,\ 0 \le x_3 \le 2\,,\ -1 - \frac{5}{4}x_3 \le x_1 \le 1 + x_3\Big\} \setminus B,$$

while $B \cap O = \emptyset$ (hence $B \subset \partial O$). There are bounded open sets with this property and smooth ($\mathbf{C}^2$ or even $\mathbf{C}^\infty$) boundary.

Let ρ denote the Rademacher distribution on $\mathcal{B}(\mathbb{R})$, that is, $\mathrm{supp}(\rho) = \{-1,1\}$ and $\rho(\{-1\}) = 1/2 = \rho(\{1\})$. Define a probability measure ν on $\mathcal{B}(\mathbb{R}^3)$ by

$$\nu \doteq \rho \otimes \rho \otimes \delta_0,$$

and choose, for $N \in \mathbb{N}$, the initial distribution for the N-player game according to $\nu_N \doteq \otimes^N \nu$.

The dynamics of the N-player game are thus given by

$$\begin{pmatrix} X^N_{i,1}(t) \\ X^N_{i,2}(t) \\ X^N_{i,3}(t) \end{pmatrix} = \begin{pmatrix} X^N_{i,1}(0) \\ X^N_{i,2}(0) \\ 0 \end{pmatrix} + \int_0^t \begin{pmatrix} u_{i,1}(s, \boldsymbol{X}^N) - \left| \frac{1}{\bar{N}^N(s)} \sum_{j=1}^N \mathbf{1}_{[0,\tau_j^N)}(s) \cdot X^N_{j,2}(s) \right| \wedge \frac{1}{4} \\ 0 \\ 1 \end{pmatrix} ds,$$

where $X^N_{i,k}(0)$, $k \in \{1,2\}$, $i \in \{1,\dots,N\}$, are i.i.d. Rademacher and $\boldsymbol{u} = (u_1,\dots,u_N) \in \mathcal{U}^N_{fb}$ is an admissible strategy vector. Clearly, for all $s \in [0,2]$,

$$\frac{1}{\bar{N}^N(s)} \sum_{j=1}^N \mathbf{1}_{[0,\tau_j^N)}(s) \cdot X^N_{j,2}(s) = \frac{1}{\bar{N}^N(s)} \sum_{j=1}^N \mathbf{1}_{[0,\tau_j^N)}(s) \cdot X^N_{j,2}(0).$$

Randomness thus enters the system only through the initial condition. We may therefore fix the stochastic basis. To this end, let $\xi^N_{i,k}$, $k \in \{1,2\}$, $i \in \{1,\dots,N\}$, be i.i.d. Rademacher random variables defined on some probability space $(\Omega_N, \mathcal{F}^N, \mathbf{P}_N)$. As filtration, we may take any filtration that makes the $\xi^N_{i,k}$ measurable at time zero. The dynamics of the N-player system can be rewritten as

$$\begin{pmatrix} X^N_{i,1}(t) \\ X^N_{i,2}(t) \\ X^N_{i,3}(t) \end{pmatrix} = \begin{pmatrix} \xi^N_{i,1} + \int_0^t \left(u_{i,1}(s, \boldsymbol{X}^N) - \left| \frac{1}{\bar{N}^N(s)} \sum_{j=1}^N \mathbf{1}_{[0,\tau_j^N)}(s) \cdot \xi^N_{j,2} \right| \wedge \frac{1}{4} \right) ds \\ \xi^N_{i,2} \\ t \end{pmatrix}$$

Since $u_{i,1}$ takes values in $[-1,1]$ and $\xi^N_{i,1}$ values in $\{-1,1\}$, we have, for $\mathbf{P}_N$-almost all $\omega \in \Omega_N$,

$$-1 - \frac{5}{4}t \le X^N_{i,1}(t,\omega) \le 1 + t \text{ for all } t \in [0,2]. \tag{7.1}$$

By construction of O, $X^N_i(\cdot,\omega)$ can leave O before the terminal time only if $X^N_{i,1}(1,\omega) = 2$; this is possible only if $\sum_{j=1}^N \xi^N_{j,2}(\omega) = 0$. But

$$\mathbf{P}_N \left(\sum_{j=1}^N \xi^N_{j,2}(\omega) = 0 \right) = \begin{cases} 0 & \text{if } N \text{ is odd,} \\ \binom{N}{N/2} \frac{1}{2^N} & \text{if } N \text{ is even.} \end{cases}$$

Since $\binom{N}{N/2} \frac{1}{2^N} \to 0$ as $N \to \infty$, we may assume for simplicity that N is odd. The dynamics of the N-player game then reduce to

$$\begin{pmatrix} X^N_{i,1}(t) \\ X^N_{i,2}(t) \\ X^N_{i,3}(t) \end{pmatrix} = \begin{pmatrix} \xi^N_{i,1} + \int_0^t u_{i,1}(s, \boldsymbol{X}^N) ds - t \cdot \left(\left| \frac{1}{N} \sum_{j=1}^N \xi^N_{j,2} \right| \wedge \frac{1}{4} \right) \\ \xi^N_{i,2} \\ t \end{pmatrix} \tag{7.2}$$

$t \in [0,2]$, for any admissible strategy vector $\boldsymbol{u}$. The associated costs for player i are, in view of (7.1),

$$J_i^N(\boldsymbol{u}) = 2 + \mathbf{E}_N \left[1 + \frac{1}{6} \int_0^2 u_{i,1}(s, \boldsymbol{X}^N) ds - \frac{1}{3} \left(\left| \frac{1}{N} \sum_{j=1}^N \xi_{j,2}^N \right| \wedge \frac{1}{4} \right) \right].^1$$

Let us turn to the limit model. Given a flow of measures $\mathfrak{p} \in \mathcal{M}$ and a stochastic-open loop control $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \alpha, W) \in \mathcal{A}$ such that $\mathbf{P} \circ \xi^{-1} = \nu$, the dynamics are given by

(7.3)
$$\begin{pmatrix} X_1(t) \\ X_2(t) \\ X_3(t) \end{pmatrix} = \begin{pmatrix} \xi_1 \\ \xi_2 \\ 0 \end{pmatrix} + \int_0^t \begin{pmatrix} \alpha_1(s) - \left| \int_{\mathbb{R}^3} w(y) \mathfrak{p}(s, dy) \right| \wedge \frac{1}{4} \\ 0 \\ 1 \end{pmatrix} ds, \quad t \in [0,2].$$

Notice that ξ_1, ξ_2 are independent Rademacher variables.

Suppose that $\mathfrak{p}$ is such that, for all $t \in [0,2]$, $\mathrm{supp}(\mathfrak{p}(t)) \subseteq \mathrm{cl}(O)$ and $\int_{\mathbb{R}^3} w(y) \mathfrak{p}(t, dy) = 0$. The dynamics then reduce to

$$\begin{pmatrix} X_1(t) \\ X_2(t) \\ X_3(t) \end{pmatrix} = \begin{pmatrix} \xi_1 + \int_0^t \alpha_1(s) ds \\ \xi_2 \\ t \end{pmatrix}, \tag{7.4}$$

while the associated expected costs are equal to $\mathbf{E}[\zeta_\alpha]$ where

$$\zeta_\alpha(\omega) \doteq \tau^X(\omega) \wedge 2 + 1 + \frac{\tau^X(\omega) \wedge 2}{12} \cdot X_1(\tau^X(\omega) \wedge 2, \omega), \quad \omega \in \Omega.$$

Since α_1 takes values in $[-1,1]$ and ξ_1 values in $\{-1,1\}$, we have, $\mathbf{P}$-almost surely,

$$-1 - t \leq X_1(t) \leq 1 + t \text{ for all } t \in [0,2]. \tag{7.5}$$

By construction of O, it follows that, for $\mathbf{P}$-almost every $\omega \in \Omega$, $X(\cdot, \omega)$ leaves O before time $T = 2$ if and only if $\xi_1(\omega) = 1$ and $\alpha_1(t, \omega) = 1$ for Lesbegue almost every $t \in [0,1]$. In this case, $\tau^X(\omega) = 1$, $X_1(1, \omega) = 2$, and $\zeta_\alpha(\omega) = 2 + 1/6$. If $X(\cdot, \omega)$ does not leave O before the terminal time, then, by (7.5), $\zeta_\alpha(\omega) \geq 2 + 1/2$, and the optimal control is to choose $\alpha_1(t, \omega) = -1$ for almost every $t \in [0,2]$. Therefore, if for $\mathbf{P}$-almost every $\omega \in \Omega$,

$$\alpha(t, \omega) = \begin{cases} (1,0,0)^\mathsf{T} & \text{if } \xi_1(\omega) = 1 \text{ and } t \in [0,1], \\ (-1,0,0)^\mathsf{T} & \text{if } \xi_1(\omega) = -1 \text{ or } t > 1, \end{cases} \tag{7.6}$$

[1] It is easy to see that if $\boldsymbol{u}$ is such that $u_{i,1} \equiv -1$ for all $i \in \{1, \ldots, N\}$, then $\boldsymbol{u}$ is a Nash equilibrium for the N-player game.

then

$$\mathbf{E}[\zeta_\alpha] = \frac{1}{2}\left(2 + \frac{1}{6}\right) + \frac{1}{2}\left(2 + \frac{1}{2}\right) = 2 + \frac{1}{3}$$

and α is optimal in the sense that

$$\mathbf{E}[\zeta_\alpha] = V(\nu; \mathfrak{p}).$$

Now, choose $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), \xi, \alpha, W) \in \mathcal{A}$ such that $\mathbf{P} \circ \xi^{-1} = \nu$ and (ξ, α) satisfies (7.6) $\mathbf{P}$-almost surely. Let X be the unique strong solution of Eq. (7.4), and define a flow of measures $\mathfrak{p}_*$ according to

$$\mathfrak{p}_*(t, \cdot) \doteq \mathbf{P}(X \in \cdot \mid \tau^X > t), \quad t \in [0, 2].$$

Notice that $\mathbf{P}(\tau^X > t) \geq 1/2$ for all $t \in [0, 2]$; thus, $\mathfrak{p}_*$ is well defined. By construction, $\mathrm{supp}(\mathfrak{p}_*) \subseteq \mathrm{cl}(O)$, which implies

$$\int_{\mathbb{R}^3} w(y)\mathfrak{p}_*(t, dy) = \int_{\mathbb{R}^3} y_2 \mathfrak{p}_*(t, dy).$$

We are going to show that $\int_{\mathbb{R}^3} w(y)\mathfrak{p}_*(t, dy) = 0$ for all $t \in [0, 2]$. By definition of ν, $X_1(0)$, $X_2(0)$, $X_3(0)$ are independent. Moreover, X_3 is (almost surely) deterministic, while α_1 is measurable with respect to $\sigma(\xi_1) = \sigma(X_1(0))$. This implies that the real-valued processes X_1, X_2, X_3 are independent. The time of first exit τ^X can be rewritten in terms of X_1 and X_3 only. It follows that τ^X and X_2 are independent, hence

$$\begin{aligned}\int_{\mathbb{R}^3} w(y)\mathfrak{p}_*(t, dy) &= \int_{\mathbb{R}^3} y_2 \mathfrak{p}_*(t, dy) \\ &= \mathbf{E}_{\mathbf{P}}\left[X_2(t)\right] = \mathbf{E}_{\mathbf{P}}\left[X_2(0)\right] = \int_{\mathbb{R}} z\rho(dz) = 0.\end{aligned}$$

As a consequence, X solves Eq. (7.3) with flow of measures $\mathfrak{p} = \mathfrak{p}_*$ and the associated costs are optimal in the sense that

$$\mathbf{E}[\zeta_\alpha] = V(\nu; \mathfrak{p}_*).$$

Recall that $\mathfrak{p}_*$ was defined as the conditional flow of measures for the law of X. Since X solves Eq. (7.3) with flow of measures $\mathfrak{p} = \mathfrak{p}_*$, the conditional mean field property of Definition 4.1 holds.

Relation (7.6) between the open-loop control α and the initial condition ξ induces a feedback strategy in $\mathcal{U}_1$, namely

$$u^*(t, \varphi) = \begin{cases} (1, 0, 0)^\mathsf{T} & \text{if } \varphi_1(0) = 1 \text{ and } t \in [0, 1], \\ (-1, 0, 0)^\mathsf{T} & \text{if } \varphi_1(0) = -1 \text{ or } t > 1, \\ \text{arbitrary} & \text{otherwise.} \end{cases}$$

For pairs (t, φ) where the control is unspecified, we may choose the control actions in such a way that u^* becomes Lipschitz continuous in the state variable at every point in time; to be specific, set

$$u^*(t,\varphi) \doteq \begin{cases} (-1 \vee \varphi_1(0) \wedge 1, 0, 0)^{\mathsf{T}} & \text{if } t \in [0,1], \\ (-1,0,0)^{\mathsf{T}} & \text{if } t > 1. \end{cases}$$

Eq. (4.1) is well posed under u^* given the flow of measures $\mathfrak{p}_*$ and any initial distribution with support in O. Let $((\Omega_*, \mathcal{F}_*, \mathbf{P}_*), (\mathcal{F}_t^*), \bar{W}, X^*)$ be a solution of Eq. (4.1) under u^* with flow of measures $\mathfrak{p}_*$ and initial distribution $\nu = \rho \otimes \rho \otimes \delta_0$. Then, by construction,

$$\mathbf{P}_* \circ (X^*)^{-1} = \mathbf{P} \circ X^{-1}.$$

This implies the optimality property

$$J(\nu, u^*; \mathfrak{p}_*) = V(\nu; \mathfrak{p}_*),$$

but also the conditional mean field property of Definition 4.1, namely

$$\mathfrak{p}_*(t) = \mathbf{P}(X \in \cdot \mid \tau^X > t) = \mathbf{P}_*(X^* \in \cdot \mid \tau^{X^*} > t) \text{ for all } t \in [0,2].$$

It follows that $(\nu, u^*, \mathfrak{p}_*)$ is a feedback solution of the mean field game.

Let us check whether or not the feedback strategy u^* induces a sequence of approximate Nash equilibria in analogy with Theorem 5.1. Thus, for $N \in \mathbb{N}$, define $\boldsymbol{u}^N = (u_1^N, \dots, u_N^N)$ by

$$u_i^N(t, \boldsymbol{\varphi}) \doteq u^*(t, \varphi_i), \quad t \in [0,T],\ \boldsymbol{\varphi} = (\varphi_1, \dots, \varphi_N) \in \mathcal{X}^N,\ i \in \{1, \dots, N\}.$$

Equation (3.1) has a unique solution under $\boldsymbol{u}^N$ with initial distribution ν_N because $\bar{b}$, w are Lipschitz continuous in the state variable and $\boldsymbol{u}^N$ is measurable with respect to $\sigma(\boldsymbol{X}^N(0))$. Therefore, $\boldsymbol{u}^N \in \mathcal{U}_{fb}^N$. For simplicity, assume again that N is odd. Let $\boldsymbol{X}^N$ be the unique strong solution of Equation (7.2) under $\boldsymbol{u}^N$. Then, for every $i \in \{1, \dots, N\}$,

$$\begin{aligned} J_i^N(\boldsymbol{u}^N) &= 2 + \mathbf{E}_N \left[1 + \frac{1}{6} \int_0^2 u^*(s, X_i^N) ds - \frac{1}{3} \left(\left| \frac{1}{N} \sum_{j=1}^N \xi_{j,2}^N \right| \wedge \frac{1}{4} \right) \right] \\ &= 3 - \frac{1}{2} \cdot \frac{2}{6} - \frac{1}{3} \mathbf{E}_N \left[\left| \frac{1}{N} \sum_{j=1}^N \xi_{j,2}^N \right| \wedge \frac{1}{4} \right] \\ &\geq \frac{33}{12}. \end{aligned}$$

Suppose player one deviates from $\boldsymbol{u}^N$ by choosing the strategy that is constant and equal to -1. Denote that strategy by v. Notice that $[\boldsymbol{u}^{N,-1}, v] \in \mathcal{U}_{fb}^N$. For the associated costs, we have

$$\begin{aligned} J_1^N([\boldsymbol{u}^{N,-1}, v]) &= 2 + \mathbf{E}_N \left[1 - \frac{2}{6} - \frac{1}{3} \left(\Big| \frac{1}{N} \sum_{j=1}^N \xi_{j,2}^N \Big| \wedge \frac{1}{4} \right) \right] \\ &= 3 - \frac{1}{3} - \frac{1}{3} \mathbf{E}_N \left[\Big| \frac{1}{N} \sum_{j=1}^N \xi_{j,2}^N \Big| \wedge \frac{1}{4} \right] \\ &\leq \frac{32}{12}. \end{aligned}$$

Player one can thus save costs of at least $1/12$ by deviating from $\boldsymbol{u}^N$ for every N odd (asymptotically, also for N even). It follows that the strategy vectors induced by the feedback solution $(\nu, u^*, \mathfrak{p}_*)$ of the mean field game do not yield approximate Nash equilibria with vanishing error.

Appendix

Let $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O. Recall the definition of $\Theta_\nu \in \mathcal{P}(\mathcal{X})$ in (5.1) at the beginning of Section 5, and also the definition of the sets $\mathcal{Q}_{\nu,c} \subset \mathcal{P}(\mathcal{X})$, $c \geq 0$, given there.

Let σ be a $d \times d$-matrix, and let

$$\hat{b} \colon [0,T] \times \mathcal{X} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}^d, \quad \tilde{b}_N \colon [0,T] \times \mathcal{X}^N \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}^d, \quad N \in \mathbb{N},$$

be functions such that the following hold:

(ND) Non-degeneracy: σ is invertible.

(M) Measurability: $\hat{b}$, $\tilde{b}_N$, $N \in \mathbb{N}$, are Borel measurable and progressive in the sense that, for all $t \in [0,T]$,

$$\hat{b}(t, \varphi, \theta) = \hat{b}(t, \tilde{\varphi}, \tilde{\theta}) \text{ whenever } \varphi_{|[0,t]} = \tilde{\varphi}_{|[0,t]} \text{ and } \theta_{|\mathcal{G}_t} = \tilde{\theta}_{|\mathcal{G}_t},$$

and analogously for $\tilde{b}_N$.

(B) Boundedness: there exists a finite constant $K > 0$ such that

$$\|\hat{b}\|_\infty \vee \sup_{N \in \mathbb{N}} \|\tilde{b}_N\|_\infty \leq K.$$

For each $N \in \mathbb{N}$, let $\nu_N \in \mathcal{P}(\mathbb{R}^{N \times d})$ be symmetric with $\mathrm{supp}(\nu_N) \subset O^N$, as above. We assume, in addition to (ND), (M), and (B):

(I) Initial distributions: the sequence $(\nu_N)_{N\in\mathbb{N}}$ is ν-chaotic, where ν has support in O.

(C) Almost continuity: for Lebesgue a.e. $t \in [0,T]$, every $\theta \in \mathcal{Q}_{\nu,K}$,

$$\Theta_\nu\Big(\varphi \in \mathcal{X} : \exists\, (\varphi_n, \theta_n) \subset \mathcal{X} \times \mathcal{P}(\mathcal{X}) \text{ such that } \hat{b}(t,\varphi_n,\theta_n) \not\to \hat{b}(t,\varphi,\theta) \text{ while } (\varphi_n,\theta_n) \to (\varphi,\theta)\Big) = 0.$$

In Section B of this Appendix, we will make the following assumption of partial Lipschitz continuity:

(L) Lipschitz continuity in the measure variable: There exists a finite constant $L > 0$ such that for all $t \in [0,T]$, all $\varphi \in \mathcal{X}$,

$$\left|\hat{b}(t,\varphi,\theta) - \hat{b}(t,\varphi,\tilde{\theta})\right| \le L \cdot \mathrm{d}_t(\theta,\tilde{\theta}) \text{ whenever } \theta,\tilde{\theta} \in \mathcal{Q}_{\nu,K},$$

where the distances d_t are pseudo-metrics derived from the total variation distance; see below.

A Convergence of empirical measures

For $N \in \mathbb{N}$, consider the system of equations

$$\begin{aligned} X_1^N(t) &= X_1^N(0) + \int_0^t \tilde{b}_N\left(s, \boldsymbol{X}^N, \mu^N\right) ds + \sigma W_1^N(t), \\ X_i^N(t) &= X_i^N(0) + \int_0^t \hat{b}\left(s, X_i^N, \mu^N\right) ds + \sigma W_i^N(t), \\ & \qquad i \in \{2,\dots,N\},\ t \in [0,T], \end{aligned} \tag{A.1}$$

where $W_1^N, \dots, W_N^N$ are independent d-dimensional Wiener processes defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$, and μ^N is the empirical measure of the players' state trajectories, that is,

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

A solution of Eq. (A.1) with initial distribution ν_N is given by a triple $((\Omega,\mathcal{F},(\mathcal{F}_t),\mathbf{P}), \boldsymbol{W}^N, \boldsymbol{X}^N)$ such that $(\Omega,\mathcal{F},(\mathcal{F}_t),\mathbf{P})$ is a filtered probability space satisfying the usual hypotheses, $\boldsymbol{W}^N = (W_1^N,\dots,W_N^N)$ a vector of independent d-dimensional $(\mathcal{F}_t)$-Wiener processes, and $\boldsymbol{X}^N = (X_1^N,\dots,X_N^N)$ a vector of continuous $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-adapted processes such that Eq. (A.1) holds $\mathbf{P}$-almost surely with $\mathbf{P} \circ (\boldsymbol{X}^N(0))^{-1} = \nu_N$.

As in Section 3, existence and uniqueness in law of solutions to Eq. (A.1) hold thanks to Girsanov's theorem and assumptions (ND), (M), and (B). Now, for each $N \in \mathbb{N}$, take a solution $((\Omega_N, \mathcal{F}_N, (\mathcal{F}_t^N), \mathbf{P}_N), \boldsymbol{W}^N, \boldsymbol{X}^N)$ of Eq. (A.1) with initial distribution ν_N, and let μ^N be the associated empirical measure on the path space $\mathcal{X}$.

Lemma A.1. *Grant (ND), (M), (B), (I), and (C). Then $(\mathbf{P}_N \circ (\mu^N)^{-1})_{N\in\mathbb{N}}$ is tight in $\mathcal{P}(\mathcal{P}(\mathcal{X}))$, and its limit points have support in $\mathcal{Q}_{\nu,K}$.*

Proof. For $N \in \mathbb{N}$, let ι_N be the intensity measure of $\mathbf{P}_N \circ (\mu^N)^{-1}$, that is,

$$\iota_N(A) \doteq \mathbf{E}_N \left[\mu^N(A)\right], \quad A \in \mathcal{B}(\mathcal{X}).$$

Tightness of $(\mathbf{P}_N \circ (\mu^N)^{-1})$ in $\mathcal{P}(\mathcal{P}(\mathcal{X}))$ is then equivalent to the tightness of (ι_N) in $\mathcal{P}(\mathcal{X})$ [cf. (2.5) in Sznitman, 1991, p. 178]. By construction,

$$\iota_N(A) = \frac{1}{N} \sum_{i=1}^{N} \mathbf{P}_N \left(X_i^N \in A\right), \quad A \in \mathcal{B}(\mathcal{X}).$$

It is therefore enough to check that the family $(\mathbf{P} \circ (X_i^N)^{-1})_{N\in\mathbb{N}, i\in\{1,\dots,N\}}$ is tight. Now, for $N \in \mathbb{N}$, $i \in \{1, \dots, N\}$,

$$\mathbf{P}_N \left(X_i^N(0) \in \mathrm{cl}(O)\right) = 1,$$

and $\mathrm{cl}(O)$ is compact since O is open and bounded. Moreover, thanks to assumption (B), for all $s, t \in [0, T]$,

$$\left|X_i^N(t) - X_i^N(s)\right| \leq K \cdot |t - s| + |\sigma| \cdot \left|W_i^N(t) - W_i^N(s)\right|,$$

where we recall that W_i^N is a standard Wiener process under $\mathbf{P}_N$. Tightness of $(\mathbf{P}_N \circ (X_i^N)^{-1})$ is now a consequence of the Kolmogorov-Chentsov tightness criterion [for instance, Corollary 16.9 in Kallenberg, 2001, p. 313].

As to the support of the limit points of $(\mathbf{P}_N \circ (\mu^N)^{-1})$, we interpret the drift terms appearing in (A.1) as stochastic relaxed controls. To this end, set $\mathcal{R} \doteq \mathcal{R}_{B_K(0)}$, where $B_K(0) \subset \mathbb{R}^d$ is the closed ball of radius K around the origin. Then $\mathcal{R}$ is compact (cf. Appendix E). For $N \in \mathbb{N}$, let ρ_i^N be the $\mathcal{R}$-valued $(\mathcal{F}_t^N)$-adapted random measure determined by

$$\rho_{i,\omega}^N(B \times I) \doteq \begin{cases} \int_I \delta_{\tilde{b}_N(t, \boldsymbol{X}^N(\cdot,\omega), \mu_\omega^N)}(B)dt & \text{if } i = 1, \\ \int_I \delta_{\hat{b}(t, X_i^N(\cdot,\omega), \mu_\omega^N)}(B)dt & \text{if } i > 1, \end{cases}$$
$$B \in \mathcal{B}(\Gamma),\ I \in \mathcal{B}([0,T]),\ \omega \in \Omega_N.$$

We can rewrite Eq. (A.1) in terms of $\rho_1^N, \dots, \rho_N^N$:

$$X_i^N(t) = X_i^N(0) + \int_{B_K(0)\times[0,t]} y\, \rho_i^N(dy, ds) + \sigma W_i^N(t), \quad i \in \{1, \dots, N\},\ t \in [0, T]. \tag{A.2}$$

Now, form the extended empirical measure

$$Q_\omega^N \doteq \frac{1}{N} \sum_{i=1}^N \delta_{(X_i^N(\cdot,\omega), \rho_{i,\omega}^N)}, \quad \omega \in \Omega_N.$$

Thus, Q^N is a $\mathcal{P}(\mathcal{X} \times \mathcal{R})$-valued random variable, and the projection on its first component coincides with μ^N. The family $(\mathbf{P}_N \circ (Q^N)^{-1})_{N\in\mathbb{N}}$ is tight in $\mathcal{P}(\mathcal{P}(\mathcal{X} \times \mathcal{R}))$ thanks to the first part of the proof and the fact that $\mathcal{R}$ is compact.

With a slight abuse of notation, let $(\hat{X}, \hat{\rho})$ denote the canonical process on $\mathcal{X}\times\mathcal{R}$. Let $(\mathbf{P}_n \circ (Q^n)^{-1})_{n\in I}$ be a convergent subsequence of $(\mathbf{P}_N \circ (Q^N)^{-1})_{N\in\mathbb{N}}$, and let Q be a $\mathcal{P}(\mathcal{X}\times\mathcal{R})$-valued random variable defined on some probability space $(\Omega, \mathcal{F}, \mathbf{P})$ such that

$$Q^n \to Q \text{ in distribution as } I \ni n \to \infty.$$

We have to show that $Q_\omega \circ \hat{X}^{-1} \in \mathcal{Q}_{\nu,K}$ for $\mathbf{P}$-almost every $\omega \in \Omega$. To this end, define a process $\hat{W}$ on $\mathcal{X} \times \mathcal{R}$ by

$$\hat{W}(t) \doteq \sigma^{-1} \left(\hat{X}(t) - \hat{X}(0) - \int_{\bar{B}_K(0)\times[0,t]} y\, \hat{\rho}(dy, ds) \right), \quad t \in [0, T].$$

By a martingale argument similar to that in the proof of Lemma A.2, but using Eq. (A.2), one checks that $\hat{W}$ is a standard Wiener process under Q_ω for $\mathbf{P}$-almost every $\omega \in \Omega$. This entails that $\hat{X}$ solves Eq. (5.2) under Q_ω with $B_K(0)$-valued control process

$$v(t, (\varphi, r)) \doteq \int_{B_K(0)} y\, \dot{r}_t(dy), \quad t \in [0, T],\ (\varphi, r) \in \mathcal{X} \times \mathcal{R}.$$

It follows that $Q_\omega \circ \hat{X}^{-1} \in \mathcal{Q}_{\nu,K}$ for $\mathbf{P}$-almost every $\omega \in \Omega$. □

In order to further characterize the limit points of $(\mathbf{P}_N \circ (\mu^N)^{-1})_{N\in\mathbb{N}}$, consider, for $\theta \in \mathcal{P}(\mathcal{X})$, the equation

$$X(t) = X(0) + \int_0^t \hat{b}\,(s, X, \theta)\, ds + \sigma W(t), \quad t \in [0, T], \tag{A.3}$$

where W is a d-dimensional Wiener process defined on some filtered probability space. A solution of Eq. (A.3) with measure $\theta \in \mathcal{P}(\mathcal{X})$ and initial distribution ν is a triple $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), W, X)$ such that $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ is a filtered probability space satisfying the usual hypotheses, W a d-dimensional $(\mathcal{F}_t)$-Wiener process, and X an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-adapted process such that (A.3) holds $\mathbf{P}$-almost surely with $\mathbf{P} \circ (X(0))^{-1} = \nu$ and drift coefficient $\hat{b}(\cdot, \cdot, \theta)$. Again thanks to Girsanov's theorem and assumptions (ND), (M), and (B), existence and uniqueness in law of solutions to Eq. (A.3) hold for each fixed $\theta \in \mathcal{P}(\mathcal{X})$.

Recall that $\hat{X}$ denotes the canonical process on $\mathcal{X}$.

Definition A.1. A measure $\theta \in \mathcal{P}(\mathcal{X})$ is called a *McKean-Vlasov solution* of Eq. (A.3) if there exists a solution $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), W, X)$ of Eq. (A.3) with initial distribution $\theta \circ (\hat{X}(0))^{-1}$ and measure θ such that $\mathbf{P} \circ X^{-1} = \theta$.

By uniqueness in law (with fixed measure θ), if $\mathbf{P} \circ X^{-1} = \theta$ holds for one solution $((\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P}), W, X)$ of Eq. (A.3) with measure θ and initial distribution $\theta \circ (\hat{X}(0))^{-1}$, then it holds for any such solution of Eq. (A.3). According to the next lemma, limit points of the sequence of empirical measures $(\mu^N)_{N \in \mathbb{N}}$ are almost surely concentrated on McKean-Vlasov solutions of Eq. (A.3). This yields, in particular, existence of McKean-Vlasov solutions. Those solutions do not necessarily have the same law.

Lemma A.2. *Grant (ND), (M), (B), (I), and (C). Let* $(\mathbf{P}_n \circ (\mu^n)^{-1})_{n \in I}$ *be a convergent subsequence of the family* $(\mathbf{P}_N \circ (\mu^N)^{-1})_{N \in \mathbb{N}}$*, and let* μ *be a* $\mathcal{P}(\mathcal{X})$*-valued random variable defined on some probability space* $(\Omega, \mathcal{F}, \mathbf{P})$ *such that*

$$\mu^n \to \mu \text{ in distribution as } I \ni n \to \infty.$$

Then μ_ω *is a McKean-Vlasov solution of Eq.* (A.3) *for* $\mathbf{P}$*-almost every* $\omega \in \Omega$*.*

Proof. Thanks to hypothesis (I), we have $\mu_\omega \circ (\hat{X}(0))^{-1} = \nu$ for $\mathbf{P}$-almost every $\omega \in \Omega$.

In the proof, we use the characterization of solutions to Eq. (A.3) with fixed measure variable through a martingale problem in the sense of Stroock and Varadhan [1979]; also see Karatzas and Shreve [1991, Section 5.4]. Since the coefficients in (A.3) are bounded, we can employ a "true" martingale problem instead of a local martingale problem and work with test functions that have compact support. For a test function $g \in \mathbf{C}_c^2(\mathbb{R}^d)$ and a measure

$\theta \in \mathcal{P}(\mathcal{X})$, define the process M_g^θ on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$ by

$$M_g^\theta(t,\varphi) \doteq g\big(\varphi(t)\big) - g\big(\varphi(0)\big) - \int_0^t \left(\hat{b}(s,\varphi,\theta)\cdot\nabla g + \frac{1}{2}\sum_{i,j=1}^{d} (\sigma\sigma^{\mathsf{T}})_{ij}\frac{\partial^2 g}{\partial x_i \partial x_j} \right)\big(\varphi(s)\big)ds.$$

Recall that $(\mathcal{G}_t)$ is the canonical filtration in $\mathcal{B}(\mathcal{X})$ and that $\hat{X}$ is the coordinate process on $\mathcal{X}$. We have to show that, for $\mathbf{P}$-almost every $\omega \in \Omega$, μ_ω solves the martingale problem associated with $\hat{b}(\cdot,\cdot,\mu_\omega)$ and $\sigma\sigma^{\mathsf{T}}$, that is, for every test function $g \in \mathbf{C}_c^2(\mathbb{R}^d)$, the process $M_g^{\mu_\omega}$ is a $(\mathcal{G}_t)$-martingale under μ_ω. Although $(\mathcal{G}_t)$ is a "raw" filtration (i.e., not necessarily right-continuous or μ_ω-complete), checking the martingale property with respect to $(\mathcal{G}_t)$ is sufficient; see, for instance, Problem 5.4.13 in Karatzas and Shreve [1991, pp. 318-319, 392]. Indeed, if $M_g^{\mu_\omega}$ is a $(\mathcal{G}_t)$-martingale under μ_ω for every $g \in \mathbf{C}_c^2(\mathbb{R}^d)$, then $((\mathcal{X}, \mathcal{B}(\mathcal{X}), (\mathcal{G}_{t+}^{\mu_\omega}), \mu_\omega), \hat{W}, \hat{X})$ is a solution of Eq. (A.3) with initial distribution ν, where $(\mathcal{G}_{t+}^{\mu_\omega})$ indicates the right-continuous μ_ω-augmentation of $(\mathcal{G}_t)$ and $\hat{W}$ is defined by

$$\hat{W}(t) \doteq \sigma^{-1}\left(\hat{X}(t) - \hat{X}(0) - \int_0^t \hat{b}\left(s, \hat{X}, \mu_\omega\right) ds \right), \quad t \in [0,T].$$

Since $\mu_\omega \circ \hat{X}^{-1} = \mu_\omega$, it then follows that μ_ω is a McKean-Vlasov solution of Eq. (A.3) in the sense of Definition A.1. For this implication to hold, it suffices to take a countable collection of test functions $g \in \mathbf{C}_c^2(\mathbb{R}^d)$ that approximate the d-variate monomials of first and second order.

Let $\theta \in \mathcal{P}(\mathcal{X})$. The processes M_g^θ are, by construction and thanks to assumptions (B) and (M), bounded, measurable, and $(\mathcal{G}_t)$-adapted. The martingale property of M_g^θ is equivalent to having

$$\mathbf{E}_\theta\left[\psi \cdot \left(M_g^\theta(t_1) - M_g^\theta(t_0)\right)\right] = 0 \tag{A.4}$$

for every choice of $(t_0, t_1, \psi) \in [0,T]^2 \times \mathbf{C}_b(\mathcal{X})$ such that $t_0 \le t_1$ and ψ is $\mathcal{G}_{t_0}$-measurable. Since the σ-algebras $\mathcal{G}_t$ are countably generated, the processes M_g^θ have continuous trajectories, and since the test functions can be taken from a countable family, we can choose a countable collection of test parameters $\mathcal{T} \subset [0,T]^2 \times \mathbf{C}_b(\mathcal{X}) \times \mathbf{C}_c^2(\mathbb{R}^d)$ with the following two properties: First, for every $(t_0, t_1, \psi, g) \in \mathcal{T}$ we have $t_0 \le t_1$ and ψ is $\mathcal{G}_{t_0}$-measurable; second, if $\theta \in \mathcal{P}(\mathcal{X})$ is such that (A.4) holds for every $(t_0, t_1, \psi, g) \in \mathcal{T}$, then θ is a McKean-Vlasov solution of Eq. (A.3). In the following three steps, we will show that there exists $\bar{\Omega} \in \mathcal{F}$ such that $\mathbf{P}(\bar{\Omega}) = 1$ and, for every $\omega \in \bar{\Omega}$, (A.4) with $\theta = \mu_\omega$ holds for all $(t_0, t_1, \psi, g) \in \mathcal{T}$.

Step 1. Let $(t_0, t_1, \psi, g) \in \mathcal{T}$. Define a function $\Psi = \Psi_{(t_0,t_1,\psi,g)} \colon \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ by

$$\Psi(\theta) = \Psi_{(t_0,t_1,\psi,g)}(\theta) \doteq \mathbf{E}_\theta \left[\psi \cdot \left(M_g^\theta(t_1) - M_g^\theta(t_0) \right) \right].$$

Notice that Ψ is well defined (the expectation on the right-hand side is finite), Borel measurable, and bounded. We claim that Ψ is continuous at every $\theta \in \mathcal{Q}_{\nu,K}$. To see this, let $\theta \in \mathcal{Q}_{\nu,K}$ and $(\theta_n)_{n\in\mathbb{N}} \subset \mathcal{P}(\mathcal{X})$ be such that $\theta_n \to \theta$ as $n \to \infty$. Define bounded measurable functions $h_n, h : [0,T] \times \mathcal{X} \to \mathbb{R}$ according to

$$h_n(s,\varphi) \doteq \psi(\varphi) \cdot \hat{b}(s,\varphi,\theta_n) \cdot \left(\sum_{i=1}^d \frac{\partial g}{\partial x_i}\big(\varphi(s)\big) \right), \qquad n \in \mathbb{N},$$

$$h(s,\varphi) \doteq \psi(\varphi) \cdot \hat{b}(s,\varphi,\theta) \cdot \left(\sum_{i=1}^d \frac{\partial g}{\partial x_i}\big(\varphi(s)\big) \right), \qquad (s,\varphi) \in [0,T] \times \mathcal{X}.$$

By hypothesis, $\theta_n \to \theta$ in the sense of weak convergence of probability measures. The functions ψ, σ (constant) as well as g together with its first and second order partial derivatives are all bounded and continuous. In order to establish the convergence of $\Psi(\theta_n)$ to $\Psi(\theta)$ it is therefore enough to verify that

$$\int_{\mathcal{X}} \int_{t_0}^{t_1} h_n(s,\varphi) ds\, \theta_n(d\varphi) \overset{n\to\infty}{\longrightarrow} \int_{\mathcal{X}} \int_{t_0}^{t_1} h(s,\varphi) ds\, \theta(d\varphi).$$

The functions h, h_n, $n \in \mathbb{N}$, are uniformly bounded. By dominated convergence and the Fubini-Tonelli theorem, it thus suffices to check that

$$\int_{\mathcal{X}} h_n(s,\varphi)\, \theta_n(d\varphi) \overset{n\to\infty}{\longrightarrow} \int_{\mathcal{X}} h(s,\varphi)\, \theta(d\varphi) \text{ for almost every } s \in [0,T]. \tag{A.5}$$

For $s \in [0,T]$, set

$$E_s \doteq \{\varphi \in \mathcal{X} : \exists (\varphi_n) \subset \mathcal{X} : h_n(s,\varphi_n) \not\to h(s,\varphi) \text{ while } \varphi_n \to \varphi\}.$$

The functions ψ, $\frac{\partial g}{\partial x_i}$, $i \in \{1,\dots,d\}$, are continuous and bounded. The choice of $\theta \in \mathcal{Q}_{\nu,K}$ entails that θ is absolutely continuous with respect to Θ_ν. By (C), the assumption of almost continuity, it follows that

$$\theta(E_s) = 0 \text{ for Lebesgue almost every } s \in [0,T].$$

The extended mapping theorem [Theorem 5.5 in Billingsley, 1968, p. 34] now implies that (A.5) holds. It follows that Ψ is continuous at $\theta \in \mathcal{Q}_{\nu,K}$.

Step 2. Let again $(t_0, t_1, \psi, g) \in \mathcal{T}$. We are going to show that

$$\mathbf{E}_\mathbf{P} \left[\big(\Psi_{(t_0,t_1,\psi,g)}(\mu) \big)^2 \right] = 0. \tag{A.6}$$

Recall that $\Psi = \Psi_{(t_0,t_1,\psi,g)}$ is bounded and, by the previous step, continuous at every $\theta \in \mathcal{Q}_{\nu,K}$. By Lemma A.1, we have $\mathbf{P}\left(\mu \in \mathcal{Q}_{\nu,K}\right) = 1$. By hypothesis, $\mu^n \to \mu$ in distribution as $I \ni n \to \infty$. The mapping theorem [Theorem 5.1 in Billingsley, 1968, p. 30] thus implies that

$$\mathbf{E}_{\mathbf{P}}\left[\left(\Psi(\mu)\right)^2\right] = \lim_{I \ni n \to \infty} \mathbf{E}_{\mathbf{P}_n}\left[\left(\Psi(\mu^n)\right)^2\right].$$

Now, for every $n \in I$, using Itô's formula and Eq. (A.1),

$$\begin{aligned}
&\mathbf{E}_{\mathbf{P}_n}\left[\left(\Psi(\mu^n)\right)^2\right] \\
&= \mathbf{E}_{\mathbf{P}_n}\left[\left(\frac{1}{n}\sum_{i=1}^n \psi(X_i^n)\cdot\left(M_g^{\mu^n}(t_1,X_i^n) - M_g^{\mu^n}(t_0,X_i^n)\right)\right)^2\right] \\
&= \frac{1}{n^2}\,\mathbf{E}_{\mathbf{P}_n}\left[\left(\psi(X_1^n)\cdot\int_{t_0}^{t_1}\left(\tilde{b}_n(s,\boldsymbol{X}^n,\mu^n) - \hat{b}(s,X_1^n,\mu_\omega)\right)\cdot\nabla g\big(X_1^n(s)\big)ds\right.\right. \\
&\qquad \left.\left. + \sum_{i=1}^n \psi(X_i^n)\cdot\int_{t_0}^{t_1}\nabla g\big(X_i^n(s)\big)^{\mathsf{T}}\sigma dW_i^n(s)\right)^2\right].
\end{aligned}$$

The functions ψ, σ (constant), ∇g, b, and $\hat{b}_n$ are bounded, uniformly in $n \in I$; cf. assumption (B). Since ψ is $\mathcal{G}_{t_0}$-measurable, the random variables $\psi(X_1^n),\ldots,\psi(X_n^n)$ are $\mathcal{F}_{t_0}^n$-measurable. The Wiener processes $W_1^n,\ldots,W_n^n$ are independent. Setting

$$\tilde{B}_1^n \doteq \int_{t_0}^{t_1}\left(\tilde{b}_n(s,\boldsymbol{X}^n,\mu^n) - \hat{b}(s,X_1^n,\mu_\omega)\right)\cdot\nabla g\big(X_1^n(s)\big)ds,$$

we obtain

$$\begin{aligned}
&\mathbf{E}_{\mathbf{P}_n}\left[\left(\Psi(\mu^n)\right)^2\right] \\
&= \frac{1}{n^2}\,\mathbf{E}_{\mathbf{P}_n}\left[\left(\psi(X_1^n)\cdot\tilde{B}_1^n\right)^2\right] \\
&\quad + \frac{2}{n^2}\,\mathbf{E}_{\mathbf{P}_n}\left[\sum_{i=1}^n \psi(X_i^n)\cdot\psi(X_1^n)\cdot\tilde{B}_1^n\cdot\int_{t_0}^{t_1}\nabla g\big(X_i^n(s)\big)^{\mathsf{T}}\sigma dW_i^n(s)\right] \\
&\quad + \frac{1}{n^2}\,\mathbf{E}_{\mathbf{P}_n}\left[\sum_{i=1}^n \left(\psi(X_i^n)\right)^2\cdot\int_{t_0}^{t_1}\nabla g\big(X_i^n(s)\big)^{\mathsf{T}}\sigma\sigma^{\mathsf{T}}\nabla g\big(X_i^n(s)\big)ds\right] \\
&\overset{I\ni n\to\infty}{\longrightarrow} 0,
\end{aligned}$$

where convergence to zero follows from the uniform boundedness of the integrands (and the Cauchy-Schwarz inequality).

Step 3. By Step 2, we have that (A.6) holds for every $\mathfrak{t} = (t_0, t_1, \psi, g) \in \mathcal{T}$. For every $\mathfrak{t} \in \mathcal{T}$, we can therefore select $A_\mathfrak{t} \in \mathcal{F}$ such that $\mathbf{P}(A_\mathfrak{t}) = 1$ and $\Psi(\mu_\omega) = 0$ for every $\omega \in A_\mathfrak{t}$. Set $\bar{\Omega} \doteq \bigcap_{\mathfrak{t}\in\mathcal{T}} A_\mathfrak{t}$. Then $\mathbf{P}(\bar{\Omega}) = 1$ since $\mathcal{T}$ is countable, and $\Psi(\mu_\omega) = 0$ for every $\omega \in \bar{\Omega}$. Hence, for $\mathbf{P}$-almost every $\omega \in \Omega$, μ_ω is such that (A.4) with $\theta = \mu_\omega$ holds for every $(t_0, t_1, \psi, g) \in \mathcal{T}$. □

B Uniqueness of McKean-Vlasov solutions

Here, we consider McKean-Vlasov solutions of Eq. (A.3) in the sense of Definition A.1 and obtain a uniqueness result.

On $\mathcal{P}(\mathcal{X})$ define the following distances derived from the total variation distance by setting, for $t \in [0, T]$,

$$\mathrm{d}_t(\theta, \tilde{\theta}) \doteq \sup_{A\in\mathcal{G}_t} \left|\theta(A) - \tilde{\theta}(A)\right|, \quad \theta, \tilde{\theta} \in \mathcal{P}(\mathcal{X}).$$

Notice that d_t are pseudo-metrics for $t \in [0, T)$, while d_T is a true metric since $\mathcal{G}_T = \mathcal{B}(\mathcal{X})$.

Let K be the constant from assumption (B), $\nu \in \mathcal{P}(\mathbb{R}^d)$ the initial distribution. We will assume the following partial Lipschitz property with respect to the measure argument of $\hat{b}$:

(L) There exists a finite constant $L > 0$ such that for all $t \in [0, T]$, all $\varphi \in \mathcal{X}$,

$$\left|\hat{b}(t, \varphi, \theta) - \hat{b}(t, \varphi, \tilde{\theta})\right| \leq L \cdot \mathrm{d}_t(\theta, \tilde{\theta}) \text{ whenever } \theta, \tilde{\theta} \in \mathcal{Q}_{\nu,K}.$$

Notice that condition (L) does not require $\hat{b}$ to be Lipschitz continuous or simply continuous with respect to the trajectory variable $\varphi \in \mathcal{X}$. In fact, for the following uniqueness result to be valid, $\hat{b}$ need not satisfy (C), the assumption of almost continuity, either.

Proposition B.1. *Grant (ND), (M), (B), and (L). Then there exists at most one McKean-Vlasov solution of Eq.* (A.3) *with initial distribution* ν.

Proof. Suppose $\theta, \tilde{\theta} \in \mathcal{P}(\mathcal{X})$ are McKean-Vlasov solutions of Eq. (A.3) with initial distribution $\nu = \theta \circ (\hat{X}(0))^{-1} = \tilde{\theta} \circ (\hat{X}(0))^{-1}$. We have to show that $\theta = \tilde{\theta}$.

Observe that θ and $\tilde{\theta}$, being McKean-Vlasov solutions of Eq. (A.3) with initial distribution ν, are elements of $\mathcal{Q}_{\nu,K}$. In particular, the Lipschitz property expressed in (L) applies to $\theta, \tilde{\theta}$.

Step 1. Set $\gamma \doteq \Theta_\nu$ and, for $t \in [0,T]$,

$$\begin{aligned} Y(t) \doteq \exp\Bigg(& \int_0^t \left(\sigma\sigma^{\mathsf{T}}\right)^{-1} \hat{b}(s,\hat{X},\theta)\cdot d\hat{X}(s) \\ & -\frac{1}{2}\int_0^t \hat{b}(s,\hat{X},\theta)^{\mathsf{T}} \left(\sigma\sigma^{\mathsf{T}}\right)^{-1} \hat{b}(s,\hat{X},\theta)ds\Bigg), \\ \tilde{Y}(t) \doteq \exp\Bigg(& \int_0^t \left(\sigma\sigma^{\mathsf{T}}\right)^{-1} \hat{b}(s,\hat{X},\tilde{\theta})\cdot d\hat{X}(s) \\ & -\frac{1}{2}\int_0^t \hat{b}(s,\hat{X},\tilde{\theta})^{\mathsf{T}} \left(\sigma\sigma^{\mathsf{T}}\right)^{-1} \hat{b}(s,\hat{X},\tilde{\theta})ds\Bigg). \end{aligned}$$

Then Y, $\tilde{Y}$ are well defined under γ and, for every $t \in [0,T]$,

$$\frac{d\theta}{d\gamma}\Big|_{\mathcal{G}_t} = Y(t), \qquad \frac{d\tilde{\theta}}{d\gamma}\Big|_{\mathcal{G}_t} = \tilde{Y}(t). \tag{B.1}$$

The equalities in (B.1) are a consequence of the Cameron-Martin-Girsanov formula; see Theorem 6.4.2 in Stroock and Varadhan [1979, p. 154]. For the sake of completeness, we derive (B.1) from a standard version of Girsanov's formula. First observe that Y, $\tilde{Y}$ are well defined if the stochastic integral that appears in the exponential is well defined. This is the case if we take γ as our reference probability measure since $\hat{X}$ is a vector of continuous martingales under γ, while the integrands $\hat{b}(s,\hat{X},\theta)$ and $\hat{b}(s,\hat{X},\tilde{\theta})$, respectively, are bounded progressively measurable processes with respect to $(\mathcal{G}_t)$ thanks to assumptions (M) and (B).

As to the densities given in (B.1), it is enough to prove the assertion for θ, the proof for $\tilde{\theta}$ being completely analogous. Let $((\Omega,\mathcal{F},(\mathcal{F}_t),\mathbf{P}),W,X)$ be a solution of Eq. (A.3) with initial distribution ν and measure θ such that $\mathbf{P}\circ X^{-1} = \theta$; such a solution exists by hypothesis.

Define a process Z on $(\Omega,\mathcal{F})$ by setting

$$Z(t) \doteq e^{-\int_0^t \sigma^{-1}\hat{b}(s,X,\theta)\cdot dW(s) - \frac{1}{2}\int_0^t \left|\sigma^{-1}\hat{b}(s,X,\theta)\right|^2 ds}, \quad t \in [0,T].$$

Then Z is a martingale with respect to $\mathbf{P}$ and $(\mathcal{F}_t)$, and

$$Q(A) \doteq \mathbf{E}_{\mathbf{P}}\left[Z(T)\mathbf{1}_A\right], \quad A \in \mathcal{F},$$

defines a probability measure such that

$$\frac{dQ}{d\,\mathbf{P}}\Big|_{\mathcal{F}_t} = Z(t) \text{ for every } t \in [0,T].$$

By Girsanov's Theorem [Theorem 3.5.1 Karatzas and Shreve, 1991, p. 191] and Eq. (A.3), we have that

$$\sigma^{-1}\left(X(\cdot) - X(0)\right) = W(\cdot) + \int_0^{\cdot} \sigma^{-1}\hat{b}(s,X,\theta)ds$$

is an $(\mathcal{F}_t)$-Wiener process under Q. Thus, since $X(0)$ is $\mathcal{F}_0$-measurable, $X(0)$ and $X(\cdot) - X(0)$ are independent under Q. It follows that

$$Q \circ X^{-1} = \mathbf{P} \circ (X(0) + \sigma W)^{-1} = \gamma.$$

Using again Eq. (A.3), we obtain with probability one under $\mathbf{P}$ as well as under Q that for all $t \in [0,T]$,

$$\begin{aligned} Z(t) &= e^{-\int_0^t \left(\sigma\sigma^\mathsf{T}\right)^{-1}\hat{b}(s,X,\theta)\cdot dX(s) + \frac{1}{2}\int_0^t \left|\sigma^{-1}\hat{b}(s,X,\theta)\right|^2 ds}, \\ \frac{1}{Z(t)} &= e^{\int_0^t \left(\sigma\sigma^\mathsf{T}\right)^{-1}\hat{b}(s,X,\theta)\cdot dX(s) - \frac{1}{2}\int_0^t \hat{b}(s,X,\theta)^\mathsf{T}\left(\sigma\sigma^\mathsf{T}\right)^{-1}\hat{b}(s,X,\theta) ds}. \end{aligned}$$

The process X is a vector of continuous semimartingales with respect to $\mathbf{P}$ as well as Q, with quadratic covariation processes

$$\langle X_i, X_j \rangle (t) = t \cdot \left(\sigma\sigma^\mathsf{T}\right)_{ij}, \quad t \in [0,T], \ \mathbf{P}\text{-}/Q\text{-almost surely.}$$

It follows that $1/Z$ is a stochastic exponential, hence a local martingale, under Q. Since $\hat{b}$ is bounded, $1/Z$ is a true martingale under Q. As a consequence,

$$\frac{d\mathbf{P}}{dQ}\Big|_{\mathcal{F}_t} = \frac{1}{Z(t)} \text{ for every } t \in [0,T].$$

Recall that $Q \circ X^{-1} = \gamma$. Comparing the expressions for $1/Z$ and Y, we find that

$$Q \circ \left(\frac{1}{Z(t)}, X\right)^{-1} = \gamma \circ \left(Y(t), \hat{X}\right)^{-1} \text{ for all } t \in [0,T].$$

Since $\theta = \mathbf{P} \circ X^{-1}$, it follows that, for every $B \in \mathcal{G}_t$,

$$\theta(B) = \mathbf{E}_\mathbf{P}\left[\mathbf{1}_B(X)\right] = \mathbf{E}_Q\left[\frac{1}{Z(t)}\mathbf{1}_B(X)\right] = \mathbf{E}_\gamma\left[Y(t)\mathbf{1}_B(\hat{X})\right],$$

hence $\frac{d\theta}{d\gamma}\big|_{\mathcal{G}_t} = Y(t)$ for all $t \in [0,T]$.

Step 2. We are going to show that there exists a constant $C \in (0,\infty)$ depending only on T, K, and σ, such that for every bounded and progressively measurable functional $\psi\colon [0,T] \times \mathcal{X} \to \mathbb{R}$, every $t \in [0,T]$,

$$\left|\mathbf{E}_\theta\left[\psi(t,\hat{X})\right] - \mathbf{E}_{\tilde{\theta}}\left[\psi(t,\hat{X})\right]\right|^2 \le C \cdot \|\psi\|_\infty^2 \int_0^t \mathrm{d}_s(\theta,\tilde{\theta})^2 ds.$$

Indeed, by (B.1), for every $t \in [0,T]$,

$$\mathbf{E}_\theta\left[\psi(t,\hat{X})\right] - \mathbf{E}_{\tilde{\theta}}\left[\psi(t,\hat{X})\right] = \mathbf{E}_\gamma\left[\left(Y(t) - \tilde{Y}(t)\right)\psi(t,\hat{X})\right]$$

Under γ, $\hat{X}$ is a martingale with quadratic covariation processes

$$\langle X_i, X_j \rangle (t) = t \cdot \left(\sigma\sigma^{\mathsf{T}} \right)_{ij}, \quad t \in [0,T],$$

while Y, $\tilde{Y}$ are stochastic exponentials. In fact, setting

$$M(t) \doteq \int_0^t \left(\sigma\sigma^{\mathsf{T}} \right)^{-1} \hat{b}(s, \hat{X}, \theta) \cdot d\hat{X}(s),$$
$$\tilde{M}(t) \doteq \int_0^t \left(\sigma\sigma^{\mathsf{T}} \right)^{-1} \hat{b}(s, \hat{X}, \tilde{\theta}) \cdot d\hat{X}(s),$$

we have, with probability one under γ,

$$Y(t) = \exp\left(M(t) - \frac{1}{2}\langle M \rangle(t) \right), \quad \tilde{Y}(t) = \exp\left(\tilde{M}(t) - \frac{1}{2}\langle \tilde{M} \rangle(t) \right).$$

Therefore (by Itô's formula), with probability one under γ,

$$Y(t) = 1 + \int_0^t Y(s) dM(s), \quad \tilde{Y}(t) = 1 + \int_0^t \tilde{Y}(s) d\tilde{M}(s), \quad t \in [0,T].$$

The invertibility of σ and the boundedness assumption (B) imply that

$$\sup_{t\in[0,T]} \left\{ \mathbf{E}_\gamma \left[|Y(t)|^2 \right] \vee \mathbf{E}_\gamma \left[\left| \tilde{Y}(t) \right|^2 \right] \right\} \leq e^{T \cdot K^2 \|\sigma^{-1}\|^2} \doteq e^{TC_0}, \tag{B.2}$$

where $\|\cdot\|$ denotes the Hilbert-Schmidt matrix norm. The bound (B.2) holds since

$$Y(t)^2 = \underbrace{\exp\left(2M(t) - \frac{1}{2}\langle 2M \rangle(t) \right)}_{\text{stochastic exponential under } \gamma} \exp\left(\langle M \rangle(t) \right),$$

while $\sup_{t\in[0,T]} |\langle M \rangle(t)| \leq T \cdot C_0$ γ-almost surely by the Cauchy-Schwarz inequality; analogously for the tilde part.

Using Itô's isometry, (B.2), and assumption (L), we obtain, for every $t \in [0,T]$,

$$\begin{aligned}
&\mathbf{E}_\gamma \left[\left| Y(t) - \tilde{Y}(t) \right|^2 \right] \\
&= \mathbf{E}_\gamma \left[\left| \int_0^t \left(\sigma\sigma^{\mathsf{T}} \right)^{-1} \left(Y(s)\hat{b}(s, \hat{X}, \theta) - \tilde{Y}(s)\hat{b}(s, \hat{X}, \tilde{\theta}) \right) \cdot d\hat{X}(s) \right|^2 \right] \\
&\leq \int_0^t \mathbf{E}_\gamma \left[\left\| \sigma^{-1} \right\|^2 \left| Y(s)\hat{b}(s, \hat{X}, \theta) - \tilde{Y}(s)\hat{b}(s, \hat{X}, \tilde{\theta}) \right|^2 \right] ds \\
&\leq 2 \left\| \sigma^{-1} \right\|^2 L^2 e^{TC_0} \int_0^t \mathrm{d}_s(\theta, \tilde{\theta})^2 ds + 2C_0 \int_0^t \mathbf{E}_\gamma \left[\left| Y(s) - \tilde{Y}(s) \right|^2 \right] ds,
\end{aligned}$$

hence, thanks to Gronwall's lemma,

$$\mathbf{E}_\gamma\left[\left|Y(t)-\tilde{Y}(t)\right|^2\right] \le 2\left\|\sigma^{-1}\right\|^2 L^2 e^{3TC_0}\int_0^t \mathrm{d}_s(\theta,\tilde{\theta})^2 ds, \quad t\in[0,T].$$

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

$$\begin{aligned}\left|\mathbf{E}_\theta\left[\psi(t,\hat{X})\right]-\mathbf{E}_{\tilde{\theta}}\left[\psi(t,\hat{X})\right]\right|^2 &\\ \le \mathbf{E}_\gamma\left[\left|Y(t)-\tilde{Y}(t)\right|^2|\psi(t,\hat{X})|^2\right] &\\ \le \|\psi\|_\infty^2\cdot 2\left\|\sigma^{-1}\right\|^2 L^2 e^{3TC_0}\int_0^t \mathrm{d}_s(\theta,\tilde{\theta})^2 ds.&\end{aligned}$$

Step 3. For $t\in[0,T]$, $A\in\mathcal{G}_t$, define $\psi_{(A,t)}\colon[0,T]\times\mathcal{X}\to\mathbb{R}$ through

$$\psi_{(A,t)}(s,\varphi)\doteq\begin{cases}\mathbf{1}_A(\varphi) & \text{if } s\ge t,\\ 0 & \text{otherwise.}\end{cases}$$

Then $\psi_{(A,t)}$ is bounded and progressively measurable with $\|\psi_{(A,t)}\|_\infty=1$. By Step 2, there exists a finite constant C, depending only on T, K, σ, such that for every $t\in[0,T]$,

$$\begin{aligned}\sup_{A\in\mathcal{G}_t}\left|\theta(A)-\tilde{\theta}(A)\right|^2 &\\ = \sup_{A\in\mathcal{G}_t}\left|\mathbf{E}_\theta\left[\psi_{(A,t)}(t,\hat{X})\right]-\mathbf{E}_{\tilde{\theta}}\left[\psi_{(A,t)}(t,\hat{X})\right]\right|^2 &\\ \le C\cdot\int_0^t \mathrm{d}_s(\theta,\tilde{\theta})^2 ds.&\end{aligned}$$

By the definition of the total variation semi-distances, it follows that

$$\mathrm{d}_t(\theta,\tilde{\theta})^2\le C\cdot\int_0^t \mathrm{d}_s(\theta,\tilde{\theta})^2 ds \text{ for all } t\in[0,T],$$

hence, thanks to Gronwall's lemma, $\mathrm{d}_T(\theta,\tilde{\theta})=0$. Since d_T is a true metric, we obtain $\theta=\tilde{\theta}$. □

C Regularity results

Here we collect some (well known) regularity results on the exit time $\tau^{\hat{X}}$ with respect to measures in $\mathcal{Q}_{\nu,K}$. Recall that $\tau^{\hat{X}}$ is the time of first exit from O on path space:

$$\tau^{\hat{X}}(\varphi)=\tau(\varphi)\doteq\inf\{t\ge 0:\varphi(t)\notin O\},\quad \varphi\in\mathcal{X},$$

where O is a bounded open set satisfying (H5).

Lemma C.1. *Let $K \geq 0$. Then $\mathcal{Q}_{\nu,K}$ is compact in $\mathcal{P}(\mathcal{X})$.*

Proof. Recall that $\mathcal{Q}_{\nu,K}$ is the set of all laws of processes of the form

$$X(t) = \xi + \int_0^t v(s)ds + \sigma W(t), \quad t \in [0,T],$$

where W is an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-Wiener process defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$, ξ is an $\mathbb{R}^d$-valued $\mathcal{F}_0$-measurable random variable with distribution $\mathbf{P} \circ \xi^{-1} = \nu$, and v is an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-progressively measurable bounded process with $\|v\|_\infty \leq K$. By the boundedness of the control processes v and the Kolmogorov-Chentsov tightness criterion [for instance, Corollary 16.9 in Kallenberg, 2001, p. 313], we have that $\mathcal{Q}_{\nu,K}$ is tight, hence precompact in $\mathcal{P}(\mathcal{X})$. By interpreting the control processes v as relaxed controls, using arguments analogous to those of the second part of the proof of Lemma A.1, one checks that $\mathcal{Q}_{\nu,K}$ coincides with its own closure. It follows that $\mathcal{Q}_{\nu,K}$ is compact. □

Lemma C.2. *Let $K > 0$. Any measure $\theta \in \mathcal{Q}_{\nu,K}$ is equivalent to Θ_ν. Moreover, there exists a strictly positive constant c_0 depending only on d, ν, T, and K, such that*

$$\inf_{\theta \in \mathcal{Q}_{\nu,K}} \mathbf{E}_\theta \left[\mathbf{1}_{[0,\tau^{\hat{X}})}(T) \right] \geq c_0 > 0.$$

Proof. Let $\theta \in \mathcal{Q}_{\nu,K}$. Let $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbf{P})$ be a filtered probability space carrying an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-Wiener process W, an $\mathbb{R}^d$-valued $\mathcal{F}_0$-measurable random variable ξ with $\mathbf{P} \circ \xi^{-1} = \nu$, and an $\mathbb{R}^d$-valued $(\mathcal{F}_t)$-progressively measurable bounded process v with $\|v\|_\infty \leq K$ such that the process

$$X(t) \doteq \xi + \int_0^t v(s)ds + \sigma W(t), \quad t \in [0,T],$$

has law $\mathbf{P} \circ X^{-1} = \theta$. Set

$$B(t) \doteq \xi + \sigma W(t), \quad t \in [0,T].$$

Define a process Z on $(\Omega, \mathcal{F})$ by setting

$$Z(t) \doteq e^{-\int_0^t \sigma^{-1} v(s) \cdot dW(s) - \frac{1}{2}\int_0^t |\sigma^{-1} v(s)|^2 ds}, \quad t \in [0,T].$$

Then Z is a martingale with respect to $\mathbf{P}$ and $(\mathcal{F}_t)$, and

$$Q(A) \doteq \mathbf{E}_\mathbf{P} \left[Z(T) \mathbf{1}_A \right], \quad A \in \mathcal{F},$$

defines a probability measure such that

$$\frac{dQ}{d\,\mathbf{P}}_{|\mathcal{F}_t} = Z(t) \text{ for every } t \in [0,T].$$

By using Girsanov's theorem as in the first step of the proof of Proposition B.1 and thanks to the boundedness of v, we find that

$$\tilde{W} \doteq \sigma^{-1}\left(X(\cdot) - \xi\right) = W(\cdot) + \int_0^{\cdot} \sigma^{-1} v(s) ds$$

is an $(\mathcal{F}_t)$-Wiener process under Q,

$$Q \circ X^{-1} = \mathbf{P} \circ (\xi + \sigma W)^{-1} = \mathbf{P} \circ B^{-1},$$

and

$$\frac{d\,\mathbf{P}}{dQ}_{|\mathcal{F}_t} = \frac{1}{Z(t)} \text{ for every } t \in [0,T].$$

Since a d-dimensional Brownian motion stays in an open ball for any fixed finite time with strictly positive probability, there exists a constant $c_0 > 0$ depending only on d, ν, and T, such that

$$\inf_{t\in[0,T]} \mathbf{E}_{\mathbf{P}}\left[\mathbf{1}_{[0,\tau^B)}(t)\right] = \mathbf{E}_{\mathbf{P}}\left[\mathbf{1}_{[0,\tau^B)}(T)\right] \geq c_0 > 0.$$

Using Hölder's inequality, we find that

$$\begin{aligned}
\mathbf{E}_\theta\left[\mathbf{1}_{[0,\tau^{\hat{X}})}(T)\right] = \mathbf{E}_{\mathbf{P}}\left[\mathbf{1}_{[0,\tau^X)}(T)\right] &= \mathbf{E}_Q\left[\frac{1}{Z(T)} \cdot \mathbf{1}_{[0,\tau^X)}(T)\right] \\
&\geq \mathbf{E}_Q\left[\sqrt{\mathbf{1}_{[0,\tau^X)}(T)}\right]^2 \mathbf{E}_Q\left[Z(T)\right]^{-1} \\
&= \mathbf{E}_P\left[\mathbf{1}_{[0,\tau^B)}(T)\right]^2 \mathbf{E}_Q\left[Z(T)\right]^{-1} \geq \frac{c_0^2}{\mathbf{E}_Q\left[Z(T)\right]}.
\end{aligned}$$

Now,

$$Z(T) = \underbrace{e^{-\int_0^T \sigma^{-1}v(s)\cdot d\tilde{W}(s) - \frac{1}{2}\int_0^T \left|\sigma^{-1}v(s)\right|^2 ds}}_{\text{corresponds to a } Q\text{-stochastic exponential}} \cdot e^{\int_0^T \left|\sigma^{-1}v(s)\right|^2 ds},$$

hence

$$\frac{1}{\mathbf{E}_Q\left[Z(T)\right]} \geq e^{-T\cdot\|\sigma^{-1}\|^2\cdot\|v\|_\infty^2}.$$

□

Lemma C.3. *Let $K > 0$, and let $\theta \in \mathcal{Q}_{\nu,K}$. Then the following hold:*

(a) $\tau^{\hat{X}} < \infty$ θ-almost surely;

(b) the mapping $\mathcal{X} \ni \varphi \mapsto \tau^{\hat{X}}(\varphi) \in [0,\infty]$ *is continuous* θ*-almost surely;*

(c) $\theta\left(\tau^{\hat{X}} = t\right) = 0$ *for every* $t \geq 0$*;*

(d) the mapping $\mathcal{X} \ni \varphi \mapsto \mathbf{1}_{[0,\tau^{\hat{X}}(\varphi))}(t) \in \mathbb{R}$ *is continuous* θ*-almost surely for every* $t \geq 0$*;*

Proof. Since θ is equivalent to Θ_ν by Lemma C.2, it is enough to check the above properties for $\theta = \Theta_\nu$.

Property (a) is a consequence of the law of the iterated logarithm (as time tends to infinity), the non-degeneracy of σ, and the fact that O is bounded. Property (b) follows again from the law of the iterated logarithm, now as time goes to zero, the non-degeneracy of σ and the fact that O is open with smooth boundary; cf., for instance, Kushner and Dupuis [2001, pp. 260-261].

Property (d) is a consequence of properties (b) and (c). As to (c), by continuity of trajectories,

$$\Theta_\nu\left(\tau^{\hat{X}} = t\right) \leq \Theta_\nu\left(\hat{X}(t) \in \partial O\right) \text{ for every } t \geq 0.$$

The boundary ∂O has Lebesgue measure zero as it is bounded and of class $\boldsymbol{C}^2$ by hypothesis. By the non-degeneracy of σ, $\hat{X}(t)$ is absolutely continuous w.r.t. Lebesgue measure, hence $\Theta_\nu(\hat{X}(t) \in \partial O) = 0$. It follows that $\Theta_\nu(\tau^{\hat{X}} = t) = 0$.

□

D Assumptions (C) and (L)

Throughout this section, we assume that $\hat{b}$ is defined by (5.4), that is, for $(t,\varphi,\theta) \in [0,T] \times \mathcal{X} \times \mathcal{P}(\mathcal{X})$,

$$\hat{b}(t,\varphi,\theta) = \begin{cases} u(t,\varphi) + \bar{b}\left(t, \varphi(t), \frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}{\int \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}\right) & \text{if } \theta(\tau > t) > 0, \\ u(t,\varphi) + \bar{b}\left(t, \varphi(t), w(0)\right) & \text{if } \theta(\tau > t) = 0, \end{cases}$$

where w is bounded continuous according to hypotheses (H1) and (H2), $\bar{b}$ is bounded measurable and Lipschitz in its third variable according to (H1) and (H3), $\bar{b}(t,\cdot,\cdot)$ is continuous uniformly in $t \in [0,T]$ thanks to (H2), and u is a feedback strategy such that, for Lebesgue-almost every $t \in [0,T]$,

$$\Theta_\nu\left(\{\varphi \in \mathcal{X} \,:\, u(t,\cdot) \text{ is discontinuous at } \varphi\}\right) = 0.$$

We are going to show that $\hat{b}$ thus defined satisfies conditions (C) and (L) above.

To establish condition (C), choose $K \geq \|\hat{b}\|_\infty$. We have to show that for Lebesgue almost every $t \in [0,T]$, every $\theta \in \mathcal{Q}_{\nu,K}$,

$$\Theta_\nu\Big(\varphi \in \mathcal{X} : \exists\, (\varphi_n, \theta_n) \subset \mathcal{X} \times \mathcal{P}(\mathcal{X}) \text{ s.th. } \hat{b}(t,\varphi_n,\theta_n) \not\to \hat{b}(t,\varphi,\theta) \text{ while } (\varphi_n,\theta_n) \to (\varphi,\theta)\Big) = 0. \tag{D.1}$$

Let $t \in [0,T]$ be such that $u(t,\cdot)$ is Θ_ν-almost surely continuous; this is true for Lebesgue almost every $t \in [0,T]$ by the continuity assumption on u. Let $\theta \in \mathcal{Q}_{\nu,K}$. Using Part (d) of Lemma C.3, we find $A_{t,\theta} \in \mathcal{B}(\mathcal{X})$ such that $\Theta_\nu(A_{t,\theta}) = 1$ and, for every $\varphi \in A_{t,\theta}$, the mappings $u(t,\cdot)$ and $\mathbf{1}_{[0,\tau^{\hat{X}}(\cdot))}(t)$ are continuous at φ. In view of Part (b) of Lemma C.3, one can choose $A_{t,\theta}$ such that also the mapping $\tau^{\hat{X}}(\cdot)$ is continuous on $A_{t,\theta}$. Since Θ_ν and θ are equivalent, we have $\theta(A_{t,\theta}) = 1$. Now, let $\varphi \in A_{t,\theta}$, and let $(\varphi_n,\theta_n) \subset \mathcal{X} \times \mathcal{P}(\mathcal{X})$ be such that $(\varphi_n,\theta_n) \to (\varphi,\theta)$ as $n \to \infty$. Then, by the mapping theorem and since $\mathbf{1}_{[0,\tau^{\hat{X}}(\cdot))}(t)$ is θ-almost surely continuous,

$$\int_{\mathcal{X}} \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta_n(d\tilde{\varphi}) \overset{n\to\infty}{\longrightarrow} \int_{\mathcal{X}} \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi}).$$

Since w is bounded and continuous, we also have

$$\int_{\mathcal{X}} w(\tilde{\varphi}(t)) \cdot \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta_n(d\tilde{\varphi}) \overset{n\to\infty}{\longrightarrow} \int_{\mathcal{X}} w(\tilde{\varphi}(t)) \cdot \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi}).$$

In view of Lemma C.2,

$$\int_{\mathcal{X}} \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi}) = \mathbf{E}_\theta\left[\mathbf{1}_{[0,\tau^{\hat{X}})}(t)\right] = \theta\left(\tau > t\right) > 0.$$

Since $\bar{b}(t,\cdot,\cdot)$ is bounded and continuous, it follows that

$$\begin{aligned}\hat{b}(t,\varphi_n,\theta_n) &= u(t,\varphi_n) + \bar{b}\left(t,\varphi(t),\frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta_n(d\tilde{\varphi})}{\int \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta_n(d\tilde{\varphi})}\right)\\ &\overset{n\to\infty}{\longrightarrow} u(t,\varphi) + \bar{b}\left(t,\varphi(t),\frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}{\int \mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}\right) = \hat{b}(t,\varphi,\theta),\end{aligned}$$

and we conclude that (D.1) holds.

As to condition (L) from Section B, we have to show that, given $K \geq \|\hat{b}\|_\infty$, there exists $L > 0$ (possibly depending on d, ν, T, and K) such that for all $t \in [0,T]$, all $\varphi \in \mathcal{X}$,

$$\left|\hat{b}(t,\varphi,\theta) - \hat{b}(t,\varphi,\tilde{\theta})\right| \leq L \cdot \mathrm{d}_t(\theta,\tilde{\theta}) \text{ whenever } \theta,\tilde{\theta} \in \mathcal{Q}_{\nu,K}. \tag{D.2}$$

Thanks to Lemma C.2,

$$\inf_{t\in[0,T]}\inf_{\theta\in\mathcal{Q}_{\nu,K}}\theta\left(\tau>t\right)=\inf_{\theta\in\mathcal{Q}_{\nu,K}}\theta\left(\tau>T\right)\geq c_0>0$$

for some strictly positive constant c_0 depending only on d, ν, T, and K. Denote by $\bar{L}$ the Lipschitz constant of $\bar{b}$ with respect to its third variable according to (H3). Let $\theta,\tilde{\theta}\in\mathcal{Q}_{\nu,K}$. Then for all $\varphi\in\mathcal{X}$,

$$\begin{aligned}
&\left|\hat{b}(t,\varphi,\theta)-\hat{b}(t,\varphi,\tilde{\theta})\right|\\
&=\left|\bar{b}\left(t,\varphi(t),\frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})}{\theta(\tau>t)}\right)\right.\\
&\quad\left.-\bar{b}\left(t,\varphi(t),\frac{\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\tilde{\theta}(d\tilde{\varphi})}{\tilde{\theta}(\tau>t)}\right)\right|\\
&\leq\frac{\bar{L}}{c_0}\left|\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})-\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\tilde{\theta}(d\tilde{\varphi})\right|\\
&\quad+\frac{\bar{L}}{c_0^2}\cdot\|w\|_\infty\cdot\left|\theta(\tau>t)-\tilde{\theta}(\tau>t)\right|.
\end{aligned}$$

By definition and since $\{\tau>t\}\in\mathcal{G}_t$,

$$\left|\theta(\tau>t)-\tilde{\theta}(\tau>t)\right|\leq\mathrm{d}_t(\theta,\tilde{\theta}).$$

The measures θ, $\tilde{\theta}$ are both equivalent to Θ_ν, and their restrictions to $\mathcal{G}_t$ are equivalent to the restriction of Θ_ν to $\mathcal{G}_t$. Denoting by Z_t and $\tilde{Z}_t$, respectively, the densities with respect to the restriction of Θ_ν, we have

$$\mathrm{d}_t(\theta,\tilde{\theta})=\frac{1}{2}\int_{\mathcal{X}}\left|Z_t(\tilde{\varphi})-\tilde{Z}_t(\tilde{\varphi})\right|\Theta_\nu(d\tilde{\varphi}).$$

It follows that

$$\begin{aligned}
&\left|\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\theta(d\tilde{\varphi})-\int w(\tilde{\varphi}(t))\mathbf{1}_{[0,\tau(\tilde{\varphi}))}(t)\tilde{\theta}(d\tilde{\varphi})\right|\\
&\qquad\qquad\leq 2\|w\|_\infty\cdot\mathrm{d}_t(\theta,\tilde{\theta}).
\end{aligned}$$

Consequently, for all $t\in[0,T]$, all $\varphi\in\mathcal{X}$, all $\theta,\tilde{\theta}\in\mathcal{Q}_{\nu,K}$,

$$\left|\hat{b}(t,\varphi,\theta)-\hat{b}(t,\varphi,\tilde{\theta})\right|\leq\frac{\bar{L}}{c_0^2}\cdot\|w\|_\infty\left(2c_0+1\right)\cdot\mathrm{d}_t(\theta,\tilde{\theta}),$$

which implies (D.2).

E Relaxed controls

In the proofs of Section 5, Appendix A and Appendix C, we need the concept of relaxed controls. For a Polish space $\mathcal{S}$, let $\mathcal{R}_\mathcal{S}$ denote the space of all deterministic $\mathcal{S}$-valued relaxed controls over the time interval $[0,T]$, that is,

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

If $r \in \mathcal{R}_\mathcal{S}$, then the time derivative of r exists almost everywhere as a measurable mapping $\dot{r}_t\colon [0,T]\to\mathcal{P}(\mathcal{S})$ such that $r(dy,dt) = \dot{r}_t(dy)dt$. The topology of weak convergence of measures turns $\mathcal{R}_\mathcal{S}$ into a Polish space. Notice that $\mathcal{R}_\mathcal{S}$ is compact if $\mathcal{S}$ is compact. Any $\mathcal{S}$-valued $(\mathcal{F}_t)$-adapted process α defined on some filtered probability space $(\Omega,\mathcal{F},\mathbf{P})$ induces an $\mathcal{R}_\mathcal{S}$-valued random variable ρ, the corresponding stochastic relaxed control, according to

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

The random measure ρ is $(\mathcal{F}_t)$-adapted in the sense that its restriction to $\mathcal{S}\times[0,t]$ is $\mathcal{F}_t$-measurable for every $t\in[0,T]$ or, equivalently, that (a version of) the time derivative process $\dot{\rho}_\cdot$ is $(\mathcal{F}_t)$-adapted. For details on relaxed controls, see, for instance, El Karoui et al. [1987] or Kushner [1990].

References

C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, Berlin, 3rd edition, 2006.

P. Billingsley. *Convergence of Probability Measures*. Wiley series in Probability and Statistics. John Wiley & Sons, New York, 1968.

A. Bensoussan, J. Frehse, and C. Grün. Stochastic differential games with a varying number of players. *Commun. Pure Appl. Anal.*, 13(5):1719–1736, 2014.

P. Briand and Y. Hu. Stability of BSDEs with random terminal time and homogenization of semilinear elliptic PDEs. *J. Funct. Anal.*, 155(2):455–494, 1998.

R. Carmona and F. Delarue. Probabilistic analysis of mean-field games. *SIAM J. Control Optim.*, 51(4):2705–2734, 2013.

R. Carmona, J.-P. Fouque, and L.-H. Sun. Mean field games and systemic risk. *Commun. Math. Sci.*, 13 (4):911–933, 2015.

R. Carmona and D. Lacker. A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.*, 25(3):1189–1231, 2015.

P. Chan and R. Sircar. Bertrand and Cournot mean field games. *Appl. Math. Optim.*, 71(3):533–569, 2015a.

P. Chan and R. Sircar. Fracking, renewables and mean field games. Available at SSRN: https://ssrn.com/abstract=2632504, 2015b.

J. Cvitanić, J. Ma, and J. Zhang. The law of large numbers for self-exciting correlated defaults. *Stochastic Processes Appl.*, 122(8):2781–2810, 2012.

P. Dai Pra, W. J. Runggaldier, E. Sartori, and M. Tolotti. Large portfolio losses: A dynamic contagion model. *Ann. Appl. Probab.*, 19(1):347–394, 2009.

R. W. R. Darling, and E. Pardoux. Backwards SDE with random terminal time and applications to semilinear elliptic PDE. *Ann. Probab.*, 25(3): 1135–1159, 1997.

F. Delarue, J. Inglis, S. Rubenthaler, and E. Tanré. Global solvability of a networked integrate-and-fire model of McKean–Vlasov type. *Ann. Appl. Probab.*, 25(4):2096–2133, 2015a.

F. Delarue, J. Inglis, S. Rubenthaler, and E. Tanré. Particle systems with a singular mean-field self-excitation. Application to neuronal networks. *Stochastic Processes Appl.*, 125(6):2451–2492, 2015b.

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.

M. Fischer. On the connection between symmetric N-player games and mean field games. *Ann. Appl. Probab.*, 27(2):757–810, 2017.

W. H. Fleming and R. W. Rishel. *Deterministic and Stochastic Optimal Control*, volume 1 of *Applications of Mathematics*. Springer, New York, 1975.

T. Funaki. A certain class of diffusion processes associated with nonlinear parabolic equations. *Z. Wahrscheinlichkeitstheor. verw. Gebiete*, 67(3): 331–348, 1984.

K. Giesecke, K. Spiliopoulos, and R. B. Sowers. Default clustering in large portfolios: Typical events. *Ann. Appl. Probab.*, 23(1):348–385, 2013.

K. Giesecke, K. Spiliopoulos, R. B. Sowers, and J. A. Sirignano. Large portfolio asymptotics for loss from default. *Math. Finance*, 25(1):77–114, 2015.

P. J. Graber and A. Bensoussan. Existence and uniqueness of solutions for Bertrand and Cournot mean field games. *Appl. Math. Optim.*, 2016, doi:10.1007/s00245-016-9366-0.

B. Hambly and S. Ledger. A stochastic McKean-Vlasov equation for absorbing diffusions on the half-line. arXiv:1605.00669 [math.PR], May 2016.

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. *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.

N. V. Krylov. *Controlled Diffusion Processes*, volume 14 of *Applications of Mathematics*. Springer, New York, 1980.

H. J. Kushner. *Weak Convergence Methods and Singularly Perturbed Stochastic Control and Filtering Problems*, volume 3 of *Systems & Control: Foundations & Applications*. Birkhäuser, Boston, 1990.

H. J. Kushner and P. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*, volume 24 of *Applications of Mathematics*. Springer, New York, 2nd edition, 2001.

O. A. Ladyženskaja, V. A. Solonnikov, and N. N. Ural'ceva. *Linear and Quasi-linear Equations of Parabolic Type*. American Mathematical Society, Providence, Rhode Island, 1968.

D. Lacker. A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields*, 165(3):581–648, 2016.

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.

S. Méléard. Asymptotic behaviour of some interacting particle systems; McKean-Vlasov and Boltzmann models. In D. Talay and L. Tubaro, editors, *Probabilistic Models for Nonlinear Partial Differential Equations. 1st C.I.M.E. Summer School, Montecatini Terme, 1995*, volume 1627 of *Lecture Notes in Math.*, pages 42–95. Springer-Verlag, Berlin, 1996.

K. Oelschläger. A martingale approach to the law of large numbers for weakly interacting stochastic processes. *Ann. Probab.*, 12(2):458–479, 1984.

L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes and Martingales. Volume 2: Itô Calculus*. Cambridge University Press, Cambridge, 2000.

K. Spiliopoulos, J. A. Sirignano, and K. Giesecke. Fluctuation analysis for the loss from default. *Stochastic Processes Appl.*, 124(7):2322–2362, 2014.

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, Springer-Verlag, Berlin, 1991.