

A PROBABILISTIC WEAK FORMULATION OF MEAN FIELD GAMES AND APPLICATIONS

RENÉ CARMONA AND DANIEL LACKER

ABSTRACT. Mean field games are studied by means of the weak formulation of stochastic optimal control. This approach allows the mean field interactions to enter through both state and control processes and take a form which is general enough to include rank and nearest-neighbor effects. Moreover, the data may depend discontinuously on the state variable, and more generally its entire history. Existence and uniqueness results are proven, along with a procedure for identifying and constructing distributed strategies which provide approximate Nash equilibria for finite-player games. Our results are applied to a new class of multi-agent price impact models and a class of flocking models for which we prove existence of equilibria.

1. INTRODUCTION

The methodology of mean field games initiated by Lasry and Lions [32] has provided an elegant and tractable way to study approximate Nash equilibria for large-population stochastic differential games with a so-called mean field interaction. In such games, the players' private state processes are coupled only through their empirical distribution. Borrowing intuition from statistical physics, Lasry and Lions study the system which should arise in the limit as the number of players tends to infinity. A set of strategies for the finite-player game is then derived from the solution of this limiting problem. These strategies form an approximate Nash equilibrium for the n -player game if n is large, in the sense that no player can improve his expected reward by more than ϵ_n by unilaterally changing his strategy, where $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$ (see [25]). An attractive feature of these strategies is that they are distributed, in the sense that the strategy of a single player depends only on his own private state.

Mean field games have seen a wide variety of applications, including models of oil production, volatility formation, population dynamics, and economic growth (see [32, 22, 33, 31] for some examples). Independently, Huang, Malhamé, and Caines developed a similar research program under the name of Nash Certainty Equivalent. The interested reader is referred to [25] and [26] and the papers cited therein. They have since generalized the framework, allowing for several different types of players and one major player.

The finite-player games studied in this paper are summarized as follows. For $i = 1, \dots, n$, the dynamics of player i 's private state process are given by a stochastic differential equation (SDE):

$$dX_t^i = b(t, X^i, \mu^n, \alpha_t^i)dt + \sigma(t, X^i)dW_t^i, \quad X_0^i = \xi^i, \quad (1)$$

where μ^n is the empirical distribution of the states:

$$\mu^n = \frac{1}{n} \sum_{j=1}^n \delta_{X^j}. \quad (2)$$

Key words and phrases. Mean Field Games, Weak Formulation, Price Impact, Flocking Models.
Partially supported by NSF: DMS-0806591.
Partially supported by NSF: DMS-0806591.

The drift b may depend on time, player i 's private state (possibly its history), the distribution of the private states (possibly their histories), and player i 's own choice of control α_t^i . Here, W^i are independent Wiener processes and ξ^i are independent identically distributed random variables independent of the Wiener processes, and each player has the same drift and volatility coefficients. Moreover, each player i has the same objective, which is to maximize

$$\mathbb{E} \left[\int_0^T f(t, X^i, \mu^n, q_t^n, \alpha_t^i) dt + g(X^i, \mu^n) \right], \text{ where } q_t^n = \frac{1}{n} \sum_{j=1}^n \delta_{\alpha_t^j}$$

over all admissible choices of α^i , subject to the constraint (1). Note that the running reward function f may depend upon the empirical distribution of the controls at time t , in addition to the same arguments as b . This is part of the thrust of the paper. Of course, each player's objective depends on the actions of the other players, and so we look for Nash equilibria.

Intuitively, if n is large, because of the symmetry of the model, player i 's contribution to μ^n is negligible, and he may as well treat μ^n as fixed. This line of argument leads to the derivation of the mean field game problem, which has the following structure:

- (1) Fix a probability measure μ on path space and a flow $\nu : t \mapsto \nu_t$ of measures on the control space;
- (2) With μ and ν frozen, solve the standard optimal control problem:

$$\begin{cases} \sup_{\alpha} \mathbb{E} \left[\int_0^T f(t, X, \mu, \nu_t, \alpha_t) dt + g(X, \mu) \right], \text{ s.t.} \\ dX_t = b(t, X, \mu, \alpha_t) dt + \sigma(t, X) dW_t, \quad X_0 = \xi; \end{cases} \quad (3)$$

- (3) Find an optimal control α , inject it into the dynamics of (3), and find the law $\Phi_x(\mu, \nu)$ of the optimally controlled state process, and the flow $\Phi_\alpha(\mu, \nu)$ of marginal laws of the optimal control process;
- (4) Find a fixed point $\mu = \Phi_x(\mu, \nu)$, $\nu = \Phi_\alpha(\mu, \nu)$.

This is to be interpreted as the optimization problem faced by a single representative player in a game consisting of infinitely many independent and identically distributed (i.i.d.) players. In the first three steps, the representative player determines his best response to the other players' states and controls which he treats as given. The final step is an equilibrium condition; if each player takes this approach, and there is to be any consistency, then there should be a fixed point. Once existence and perhaps uniqueness of a fixed point are established, the second problem is to use this fixed point to construct approximate Nash equilibrium strategies for the original finite-player game. These strategies will be constructed from the optimal control for the problem of step (2), corresponding to the choosing (μ, ν) to be the fixed point in step (1).

The literature on mean field games comprises two streams of papers: one based on analytic methods and one on a probabilistic approach.

Lasry and Lions (e.g. [32], [22], etc.) study these problems via a system of partial differential equations (PDEs). The control problem gives rise to a Hamilton-Jacobi-Bellman equation for the value function, which evolves backward in time. The law of the state process is described by a Kolmogorov equation, which evolves forward in time. These equations are coupled through the dependence on the law of the state process, in light of the consistency requirement (4). This approach applies in the Markovian case, when the data b , σ , f , and g are smooth or at least continuous functions of the states and not of their pasts. Results in this direction include two broad classes of mean field interactions: some have considered local dependence of the data on the measure argument, such as functions $(x, \mu) \mapsto G(d\mu(x)/dx)$ of the density, while others have studied nonlocal functionals, which are continuous with respect to a weak or Wasserstein topology.

More recently, several authors have taken a probabilistic approach to this problem by using the Pontryagin maximum principle to solve the optimal control problem. See, for example, [10, 6, 11]. Typically in a stochastic optimal control problem, the backward stochastic differential equations (BSDEs) satisfied by the adjoint processes are coupled with the forward SDE for the state process through the optimal control, which is generally a function of both the forward and backward parts. When the maximum principle is applied to mean field games, the forward and backward equations are coupled additionally through the law of the forward part. Carmona and Delarue investigate this new type of forward-backward stochastic differential equations (FBSDEs) in [9]. It should be noted that there is a similar but distinct way to analyze the infinite-player limit of large-population games, leading to the optimal control of stochastic dynamics of McKean-Vlasov type. Early forms of a stochastic maximum principle for this new type of control problem were given in [4, 7, 34]. A general form of this principle was given in [10] where it was applied to the solution of the control problem. A comparison of these two asymptotic regimes is given in [12].

The aim of this paper is to present a new probabilistic approach to the analysis of mean field games with uncontrolled diffusion coefficients. Assuming $\sigma = \sigma(t, x)$ contains neither a mean field term nor a control, we obtain a general existence result. Under stronger assumptions, we prove a modest extension of the uniqueness result of Lasry and Lions [32]. Finally, we provide a construction of approximate Nash equilibria for finite-player games in the spirit of [11], in the case that b has no mean field term.

Our analysis is based on the *weak formulation* of stochastic optimal control problems, sometimes known as the *martingale approach*; see for example, [15, 30, 37]. This approach depends heavily on the non-degeneracy of σ and its independence of the control, and in our case, it is also important that σ has no mean field term. The *strong* formulation of the problem, as in [11], would require that the state SDEs have strong solutions when controls are applied. The two formulations are compared in Remark 7.12. One of the main conveniences of the weak formulation is that weak existence and uniqueness of the state SDE require much less regularity in the coefficients, which are allowed to be path-dependent and merely measurable in the state variable. Also, the value function solves a backward stochastic differential equation (BSDE), and necessary and sufficient conditions for the optimality of a control follow easily from the comparison principle for BSDEs. This method is discussed by El Karoui and Quenez in [30], Peng in [37], and perhaps most thoroughly by Hamadene and Lepeltier in [23].

Our results allow for the mean field interaction (at least in the running reward function f) to occur through the control processes in addition to the state processes. This seems quite important for many practical applications and has received very little attention thusfar in the literature of mean field games. A very recent paper of Gomes and Voskanyan [21] uses PDE methods to study these types of interactions in the deterministic case, $\sigma \equiv 0$, under the name *extended mean field games*. Under strong continuity and convexity assumptions, they obtain existence as well as some regularity of the solutions, and interestingly they are able to allow for general dependence of the running objective f on the *joint* law of the state and control processes. Our setting is very different: notably $\sigma > 0$, and our convexity and continuity assumptions are much weaker.

We also allow for very general nonlocal mean field interactions, including but not limited to weakly or Wasserstein continuous functionals. Among the natural interactions that have not yet been addressed in the mean field games literature which we are able to treat, we mention the case of coefficients which depend on the rank (Example 5.9 in Section 5), or on the mean field of the individual's nearest neighbors (Section 2.2). Our framework also includes models with different types of agents, similar to [25]. Moreover, f does not need to be strictly convex, and may in fact be identically zero. A final novelty of our results worth emphasizing is that they apply in non-Markovian settings and require no continuity in the state variable.

For the sake of illustration, we present two applications which had been touted as models for mean field games, without being solved in full generality. First we study price impact models in which asset price dynamics depend naturally on the *rates of change* of investors' positions, inspired by the model of Carlin et al. [8]. As a second application of our theoretical results, we discuss a model of flocking proposed by Nourian et al. in [35] to provide a mechanism by which flocking behavior emerges as an equilibrium, as a game counterpart of the well-known Cucker-Smale model, [14]. In [35], the authors identify the mean field limit and, *under the assumption that there exists a unique solution to the limiting mean field game*, construct approximate Nash equilibria for the finite-player games. While flocking is often defined mathematically as a large time phenomenon (case in point, the stationary form of the mean field game strategy is considered in [35]), we treat the finite horizon case to be consistent with the set-up of the paper, even though this case is most often technically more challenging. We provide existence and approximation results for both their model and two related nearest-neighbor models.

This paper is organized as follows. We introduce the two practical applications in Section 2. The price impact models of Section 2.1 motivate the analysis of mean field games in which players interact through their controls, while Section 2.2 describes the flocking model of [35] as well as two related nearest-neighbor models. Then, Section 3 provides precise statements of the assumptions used throughout the paper and the main existence and uniqueness results. Section 4 explains the construction of approximate Nash equilibria for the finite-player game. The assumptions of the main theorems are discussed in more detail in Section 5, along with important examples. In Section 6 the general theory is specialized to the applications of Section 2. The proofs of the main theorems of Sections 3 and 4 are given in Sections 7 and 8, respectively.

2. APPLICATIONS

2.1. Price impact models. To motivate our generalization of the class of mean field games worthy of investigation, we present a simple multi-agent model of price impact which leads to mean field interaction through the control processes. The model is along the lines of Almgren and Chriss's model [3] for price impact, or rather its natural extension to an n -player competitive game given by Carlin, Lobo, and Viswanathan in [8]. The latter model is highly tractable, modeling a flat order book from which each agent must execute a fixed order. We instead model a nonlinear order book and use fairly general reward functions. See [1, 20] for a discussion of order book mechanics as well as a discussion of resilience, a concept we do not address. In our model, after each trade, the order book reconstructs itself instantly around a new mid-price S_t , and with the same shape. At each time t , each agent faces a cost structure given by the same transaction cost curve $c : \mathbb{R} \rightarrow [0, \infty]$, which is convex and satisfies $c(0) = 0$. We consider only order books with finite volume; an infinite value for $c(\alpha)$ simply means that the volume α is not available. Flat order books are common in the literature, though not realistic: they correspond to quadratic transaction costs c .

We work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P)$ supporting $n + 1$ independent Wiener processes, W^1, \dots, W^n and B . Let S denote the asset price, K^i the cash of agent i , and X^i his position. Each agent controls his trading rate α_t^i and his position evolves according to

$$dX_t^i = \alpha_t^i dt + \sigma dW_t^i.$$

The noise term σdW_t^i models a random stream of demand that a broker may receive from his clients. If a single agent i places a market order of α_t^i when the mid-price is S_t , the transaction costs him $\alpha_t^i S_t + c(\alpha_t^i)$. Hence, the changes in cash of agent i are naturally given by

$$dK_t^i = -(\alpha_t^i S_t + c(\alpha_t^i)) dt.$$

Assuming c is differentiable on its domain, the marginal price per share of this trade is $S_t + c'(\alpha_t^i)$, meaning that the agent receives all of the volume on the order book between the prices S_t and $S_t + c'(\alpha_t^i)$. The order book should recenter somewhere in this price range, say at $S_t + \gamma c'(\alpha_t^i)/n$,

where $\gamma > 0$. The factor of $1/n$ is irrelevant when n is fixed, but it is the right scaling factor for obtaining a mean field approximation.

In a continuous-time, continuous-trading model with multiple agents, it is not clear how simultaneous trades should be handled. Somewhat more realistic are continuous-time, discrete-trade models, which many continuous-trade models are designed to approximate. In a continuous-time, discrete-trade model, it is reasonable to assume that agents never trade simultaneously, given that there is a continuum of trade times to choose from. We choose to model this in our continuous-trade setting in the following manner: When the n agents trade at rates $\alpha_t^1, \dots, \alpha_t^n$ at time t , agent i still pays $\alpha_t^i S_t + c(\alpha_t^i)$, but the total change in price is

$$\frac{\gamma}{n} \sum_{i=1}^n c'(\alpha_t^i).$$

Finally, the mid-price is modeled as an underlying martingale plus a drift representing a form of permanent price impact:

$$dS_t = \frac{\gamma}{n} \sum_{i=1}^n c'(\alpha_t^i) dt + \sigma_0 dB_t.$$

Note that the particular case $c(\alpha) = \alpha^2$ corresponds to the influential Almgren-Chriss model [3]. The wealth V_t^i of agent i at time t , as pegged to the mid-price, is given by $V_0^i + X_t^i S_t + K_t^i$, which leads to the following dynamics:

$$dV_t^i = \left(\frac{\gamma}{n} \sum_{j=1}^n c'(\alpha_t^j) X_t^i - c(\alpha_t^i) \right) dt + \sigma_0 X_t^i dB_t + \sigma S_t dW_t^i. \quad (4)$$

We assume that the agents are risk-neutral and seek to maximize their expected terminal wealths at the end of the trading period, including some agency costs given by functions f and g , so that the objective of agent i is to maximize:

$$J^i(\alpha^1, \dots, \alpha^n) = \mathbb{E} \left[V_T^i - \int_0^T f(t, X_t^i) dt - g(X_T^i) \right].$$

Price impact models are most often used in optimal execution problems for high frequency trading. Because of their short time scale, the fact that S_t as defined above can become negative is not an issue in practice. In these problems, one often chooses $g(x) = mx^2$ for some $m > 0$ in order to penalize left over inventory. The function f is usually designed to provide an incentive for tracking a benchmark, say the frequently used market volume weighted average price (VWAP) and a penalty slippage.

If the control processes are square integrable and the cost function c has at most quadratic growth, the volumes X_t^i and the transaction price S_t are also square integrable and the quadratic variation terms in (4) are true martingales. So after using Itô's formula we find

$$J^i(\alpha^1, \dots, \alpha^n) = \mathbb{E} \left[\int_0^T \left(\frac{\gamma}{n} \sum_{j=1}^n c'(\alpha_t^j) X_t^i - c(\alpha_t^i) - f(t, X_t^i) \right) dt - g(X_T^i) \right].$$

Treating X^i as the state processes, this problem is of the form described in the introduction. The general theory presented in the sequel will apply to this model under modest assumptions on the functions c , f , and g , ensuring existence of approximate Nash equilibria. Intuitively, when n is large, a single agent may ignore his price impact without losing much in the way of optimality. This model could be made more realistic in many ways, but we believe any improvement will preserve the basic structure of the price impact, which naturally depends on the mean field of the control processes. It should be mentioned that the risk-neutrality assumption is crucial and hides a much more difficult problem. Without risk-neutrality, we would have to keep track of V and S as state processes. More

importantly, the Brownian motion B would not disappear after taking expectations, and this would substantially complicate the mean field limit.

2.2. Flocking models. The position X_t^i and velocity V_t^i of individual i change according to

$$\begin{aligned} dX_t^i &= V_t^i dt, \\ dV_t^i &= \alpha_t^i dt + \sigma dW_t^i, \end{aligned}$$

where α_t^i is the individual's acceleration vector, W^i are independent d -dimensional Wiener processes, and $\sigma > 0$ is a $d \times d$ matrix (usually $d = 2$ or $d = 3$). The objective of individual i is to choose α^i to minimize

$$\mathbb{E} \left[\int_0^T |\alpha_t^i|_R^2 + \left| \frac{1}{n} \sum_{j=1}^n (V_t^j - V_t^i) \phi(|X_t^j - X_t^i|) \right|_Q^2 dt \right]. \quad (5)$$

Here, $\phi : [0, \infty) \rightarrow [0, \infty)$ is a nonincreasing function, and $|x|_Q := x^\top Q x$ and $|x|_R := x^\top R x$ for $x \in \mathbb{R}^d$, where Q and R are positive semidefinite $d \times d$ matrices. The $|\alpha_t^i|_R^2$ term penalizes too rapid an acceleration, while the second term provides an incentive for an individual to align his velocity vector with the average velocity of the flock. The weights $\phi(|X_t^j - X_t^i|)$ emphasize the velocities of nearby (in position) individuals more than distant ones. In [35], drawing inspiration from [14], ϕ is of the form

$$\phi(x) = c(1 + x^2)^{-\beta}, \quad \beta \geq 0, \quad c > 0. \quad (6)$$

Our existence and approximation results apply to the model above as well as a related model in which the weights in (5) take a different form. Namely, individual i may give non-zero weight only to those individuals it considers to be neighbors, where the set of neighbors may be determined in two different ways. *Nearest neighbor rules* pre-specify a radius $r > 0$, and an individual i 's neighbors at time t are those individuals j with $|X_t^j - X_t^i| \leq r$. Letting N_t^i denote the set of such j and $|N_t^i|$ its cardinality, the objective function is

$$\mathbb{E} \left[\int_0^T |\alpha_t^i|_R^2 + \left| \frac{c}{|N_t^i|} \sum_{j \in N_t^i} (V_t^j - V_t^i) \right|_Q^2 dt \right]. \quad (7)$$

This is inspired by what is now known as Vicsek's model, proposed in [40] and studied mathematically in [29]. On the other hand, recent studies such as [5] provide evidence that birds in flocks follow so-called *k-nearest neighbor rules*, which track only a fixed number $k \leq n$ of neighbors at each time. The corresponding objective function is the same, if we instead define N_t^i to be the set of indices j of the k closest individuals to i (so of course $|N_t^i| = k$). Note that there are no "ties"; that is, for each distinct $i, j, l \leq n$ and $t > 0$, we have $P(|X_t^i - X_t^j| = |X_t^i - X_t^l|) = 0$.

3. MEAN FIELD GAMES

We turn now to a general discussion of the mean field game models which we consider in this paper. We collect the necessary notation and assumptions in order to state the main existence, uniqueness, and approximation theorems.

3.1. Construction of the mean field game. Let $\mathcal{B}(E, \tau)$ denote the Borel σ -field of a topological space (E, τ) . When the choice of topology is clear, we use the abbreviated form $\mathcal{B}(E)$. For a measurable space (Ω, \mathcal{F}) , let $\mathcal{P}(\Omega)$ denote the set of probability measures on (Ω, \mathcal{F}) . We write

$\mu \ll \mu'$ when μ is absolutely continuous with respect to μ' , and $\mu \sim \mu'$ when the measures are equivalent. Given a measurable function $\psi : \Omega \rightarrow [1, \infty)$, we set:

$$\begin{aligned} \mathcal{P}_\psi(\Omega) &= \left\{ \mu \in \mathcal{P}(\Omega) : \int \psi d\mu < \infty \right\}, \\ \mathcal{B}_\psi(\Omega) &= \left\{ f : \Omega \rightarrow \mathbb{R} \text{ measurable, } \sup_\omega |f(\omega)|/\psi(\omega) < \infty \right\}. \end{aligned}$$

We define $\tau_\psi(\Omega)$ to be the weakest topology on $\mathcal{P}_\psi(\Omega)$ making the map $\mu \mapsto \int f d\mu$ continuous for each $f \in \mathcal{B}_\psi(\Omega)$. The space $(\mathcal{P}_\psi(\Omega), \tau_\psi(\Omega))$ is generally neither metrizable nor separable, which will pose some problems. We define the empirical measure map $e_n : \Omega^n \rightarrow \mathcal{P}(\Omega)$ by

$$e_n(\omega_1, \dots, \omega_n) = \frac{1}{n} \sum_{j=1}^n \delta_{\omega_j}.$$

Notice that e_n need not be $\mathcal{B}(\mathcal{P}_\psi(\Omega), \tau_\psi(\Omega))$ -measurable, but this will not be an issue.

Definition 3.1. Given measurable spaces E and F , we say that a function $f : \mathcal{P}(\Omega) \times E \rightarrow F$ is *empirically measurable* if

$$\Omega^n \times E \ni (x, y) \mapsto f(e_n(x), y) \in F$$

is jointly measurable for all $n \geq 1$.

Let $\mathcal{C} := C([0, T]; \mathbb{R}^d)$ be the space of \mathbb{R}^d -valued continuous functions on $[0, T]$ endowed with the sup-norm $\|x\| := \sup_{s \in [0, T]} |x(s)|$ and fix a Borel measurable function $\psi : \mathcal{C} \rightarrow [1, \infty)$ throughout. It will play a role similar to the ‘‘Lyapunov-like’’ function of Gärtner [19], controlling a tradeoff between integrability and continuity requirements. Some comments on the choice of ψ follow in Remark 3.7. For any $\mu \in \mathcal{P}(\mathcal{C})$ and $t \in [0, T]$, the marginal μ_t denotes the image of μ under the coordinate map $\mathcal{C} \ni x \mapsto x_t \in \mathbb{R}^d$.

We use the notation $\lambda_0 \in \mathcal{P}(\mathbb{R}^d)$ for the initial distribution of the infinitely many players’ state processes. Let $\Omega := \mathbb{R}^d \times \mathcal{C}$, define $\xi(x, \omega) := x$ and $W(x, \omega) := \omega$, and let P denote the product of λ_0 and the Wiener measure, defined on $\mathcal{B}(\Omega)$. Define \mathcal{F}_t to be the completion of $\sigma((\xi, W_s) : 0 \leq s \leq t)$ by P -null sets of $\mathcal{B}(\Omega)$, and set $\mathbb{F} := (\mathcal{F}_t)_{0 \leq t \leq T}$. We work with the filtered probability space $(\Omega, \mathcal{F}_T, \mathbb{F}, P)$ for the remainder of the section. For $k \in \mathbb{N}$ and $q \geq 1$ define the spaces

$$\mathbb{H}^{q,k} := \left\{ h : [0, T] \times \Omega \rightarrow \mathbb{R}^k \text{ progressively measurable with } \mathbb{E} \left[\left(\int_0^T |h_t|^2 dt \right)^{q/2} \right] < \infty \right\}.$$

For a martingale M , we denote by $\mathcal{E}(M)$ its Doleans stochastic exponential. We now state assumptions on the data which will stand throughout the paper. Unless otherwise stated, $\mathcal{P}_\psi(\mathcal{C})$ is equipped with the topology $\tau_\psi(\mathcal{C})$.

The following assumptions **(S)** are implicitly assumed throughout the paper.

Assumption (S) (Standing assumptions).

(S.1) The *control space* A is a compact convex subset of a normed vector space, and the set \mathbb{A} of *admissible controls* consists of all progressively measurable A -valued processes. The volatility $\sigma : [0, T] \times \mathcal{C} \rightarrow \mathbb{R}^{d \times d}$ is progressively measurable. The drift $b : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times A \rightarrow \mathbb{R}^d$ is such that $(t, x) \mapsto b(t, x, \mu, a)$ is progressively measurable for each (μ, a) , and $a \mapsto b(t, x, \mu, a)$ is continuous for each (t, x, μ) .

(S.2) There exists a unique strong solution X of the driftless state equation

$$dX_t = \sigma(t, X) dW_t, \quad X_0 = \xi, \tag{8}$$

such that $\mathbb{E}[\psi^2(X)] < \infty$, $\sigma(t, X) > 0$ for all $t \in [0, T]$ almost surely, and $\sigma^{-1}(t, X)b(t, X, \mu, a)$ is uniformly bounded.

We will elaborate on these and the subsequent assumptions in Section 6 below, but for now let us make a few remarks. If σ has linear growth, $\psi(x) = 1 + \|x\|^p$, and $\int_{\mathbb{R}^d} |x|^{2p} \lambda_0(dx) < \infty$, then indeed $\mathbb{E}[\psi^2(X)] < \infty$. Compactness of A is a strong assumption which will be used in several places, in particular to ensure that $\mathcal{P}(A)$ is compact. Boundedness of $\sigma^{-1}b$ is also restrictive, but it will be crucial to ensure that the Hamiltonian is a uniformly Lipschitz function of the adjoint variable. See Remark 5.8 for more details and some comments about relaxing these assumption.

From now on, X denotes the unique solution of (8). For each $\mu \in \mathcal{P}_\psi(\mathcal{C})$ and $\alpha \in \mathbb{A}$, define a measure $P^{\mu, \alpha}$ on (Ω, \mathcal{F}_T) by

$$\frac{dP^{\mu, \alpha}}{dP} = \mathcal{E} \left(\int_0^\cdot \sigma^{-1}b(t, X, \mu, \alpha_t) dW_t \right)_T.$$

By Girsanov's theorem and boundedness of $\sigma^{-1}b$, the process $W^{\mu, \alpha}$ defined by

$$W_t^{\mu, \alpha} := W_t - \int_0^t \sigma^{-1}b(s, X, \mu, \alpha_s) ds$$

is a Wiener process under $P^{\mu, \alpha}$, and

$$dX_t = b(t, X, \mu, \alpha_t) dt + \sigma(t, X) dW_t^{\mu, \alpha}.$$

That is, under $P^{\mu, \alpha}$, X is a weak solution of the state equation. Note that $P^{\mu, \alpha}$ and P agree on \mathcal{F}_0 ; in particular, the law of $X_0 = \xi$ is still λ_0 . Moreover, ξ and W remain independent under $P^{\mu, \alpha}$.

Remark 3.2. It is well-known that the nonsingularity assumption (S.2) of σ guarantees that \mathbb{F} coincides with the completion of the filtration generated by X . It is thus implicit in the definition of \mathbb{A} that our admissible controls can be written in *closed-loop* form, that is as deterministic functions of (t, X) .

We now state the assumptions on the *reward* functions entering the objectives to be maximized by the players. Throughout, $\mathcal{P}(A)$ is endowed with the weak topology and its corresponding Borel σ -field.

(S.3) The running reward $f : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$ is such that $(t, x) \mapsto f(t, x, \mu, q, a)$ is progressively measurable for each (μ, q, a) and $a \mapsto f(t, x, \mu, q, a)$ is continuous for each (t, x, μ, q) . The terminal reward function $g : \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \rightarrow \mathbb{R}$ is such that $x \mapsto g(x, \mu)$ is Borel measurable for each μ .

(S.4) There exist $c > 0$ and an increasing function $\rho : [0, \infty) \rightarrow [0, \infty)$ such that

$$|g(x, \mu)| + |f(t, x, \mu, q, a)| \leq c \left(\psi(x) + \rho \left(\int \psi d\mu \right) \right), \quad \forall (t, x, \mu, q, a).$$

Since $\psi \geq 1$, this is equivalent to the same assumption but with ψ replaced by $1 + \psi$.

(S.5) The function f is of the form

$$f(t, x, \mu, q, a) = f_1(t, x, \mu, a) + f_2(t, x, \mu, q).$$

Remark 3.3. The only restrictive assumption among (S.3-5) is (S.5). Combined with the assumption that b does not depend on q , assumption (S.5) renders the maximizer(s) of the Hamiltonian independent of the $\mathcal{P}(A)$ argument. Separation assumptions of this sort are common in mean field games literature, largely for this reason (c.f. [32]).

Given a measure $\mu \in \mathcal{P}_\psi(\mathcal{C})$, a control $\alpha \in \mathbb{A}$, and a measurable map $[0, T] \ni t \mapsto q_t \in \mathcal{P}(A)$, we define the associated expected reward by

$$J^{\mu, q}(\alpha) := \mathbb{E}^{\mu, \alpha} \left[\int_0^T f(t, X, \mu, q_t, \alpha_t) dt + g(X, \mu) \right]$$

where $\mathbb{E}^{\mu,\alpha}$ denotes expectation with respect to the measure $P^{\mu,\alpha}$. Considering μ and q as fixed, we are faced with a standard stochastic optimal control problem, the value of which is given by

$$V^{\mu,q} = \sup_{\alpha \in \mathbb{A}} J^{\mu,q}(\alpha).$$

Definition 3.4. We say a measure $\mu \in \mathcal{P}_\psi(\mathcal{C})$ and a measurable function $q : [0, T] \rightarrow \mathcal{P}(A)$ form a *solution of the MFG* if there exists $\alpha \in \mathbb{A}$ such that $V^{\mu,q} = J^{\mu,q}(\alpha)$, $P^{\mu,\alpha} \circ X^{-1} = \mu$, and $P^{\mu,\alpha} \circ \alpha_t^{-1} = q_t$ for almost every t .

3.2. Existence and uniqueness. Some additional assumptions are needed for the existence and uniqueness results. Define the Hamiltonian $h : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times \mathbb{R}^d \times A \rightarrow \mathbb{R}$, the maximized Hamiltonian $H : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times \mathbb{R}^d \rightarrow \mathbb{R}$, and the set on which the supremum is attained by

$$\begin{aligned} h(t, x, \mu, q, z, a) &:= f(t, x, \mu, q, a) + z \cdot \sigma^{-1} b(t, x, \mu, a), \\ H(t, x, \mu, q, z) &:= \sup_{a \in A} h(t, x, \mu, q, z, a), \\ A(t, x, \mu, q, z) &:= \{a \in A : h(t, x, \mu, q, z, a) = H(t, x, \mu, q, z)\}, \end{aligned} \tag{9}$$

respectively. Note that $A(t, x, \mu, q, z)$ does not depend on q , in light of assumption (S.5), so we shall often drop q from the list of arguments of A and use the notation $A(t, x, \mu, z)$. Note also that $A(t, x, \mu, z)$ is always nonempty, since A is compact and h is continuous in a by assumptions (S.1) and (S.3).

Assumption (C). For each (t, x, μ, z) , the set $A(t, x, \mu, z)$ is convex.

It will be useful to have notation for the driftless law and the set of equivalent laws:

$$\begin{aligned} \mathcal{X} &:= P \circ X^{-1} \in \mathcal{P}_\psi(\mathcal{C}) \\ \mathcal{P}_X &:= \{\mu \in \mathcal{P}_\psi(\mathcal{C}) : \mu \sim \mathcal{X}\}. \end{aligned}$$

Assumption (E) (Existence assumptions). For each $(t, x) \in [0, T] \times \mathcal{C}$ the following maps are sequentially continuous, using $\tau_\psi(\mathcal{C})$ on \mathcal{P}_X and the weak topology on $\mathcal{P}(A)$:

$$\begin{aligned} \mathcal{P}_X \times A \ni (\mu, a) &\mapsto b(t, x, \mu, a), \\ \mathcal{P}_X \times \mathcal{P}(A) \times A \ni (\mu, q, a) &\mapsto f(t, x, \mu, q, a), \\ \mathcal{P}_X \ni \mu &\mapsto g(x, \mu). \end{aligned}$$

Theorem 3.5. *Suppose (E) and (C) hold. Then there exists a solution of the MFG.*

Remark 3.6. It is worth emphasizing that sequential continuity is often easier to check for $\tau_\psi(\mathcal{C})$, owing in part to the failure of the dominated convergence theorem for nets. For example, functions like $\mu \mapsto \int \int \phi(x, y) \mu(dx) \mu(dy)$ for bounded measurable ϕ are always sequentially continuous but may fail to be continuous.

Remark 3.7. The function ψ enters the assumptions in two essential ways. On the one hand, the functions b , f , and g should be $\tau_\psi(\mathcal{C})$ -continuous in their measure arguments as in (E). On the other hand, the solution of the SDE $dX_t = \sigma(t, X) dW_t$ should possess ψ^2 -moments as in (S.2), and the growth of f and g should be controlled by ψ , as in (S.4). There is a tradeoff in the choice of ψ : larger ψ makes the latter point more constraining and the former less constraining.

The following uniqueness theorem is inspired by Lasry and Lions [32]. They provide counterexamples to show that one should not expect uniqueness in much generality, unless one assumes that the time horizon is small and the coefficients are suitably Lipschitz (e.g. [25]).

Assumption (U).

(U.1) For each (t, x, μ, z) , the set $A(t, x, \mu, z)$ is a singleton;

- (U.2) $b = b(t, x, a)$ has no mean field term;
 (U.3) $f(t, x, \mu, a) = f_1(t, x, \mu) + f_2(t, \mu, q) + f_3(t, x, a)$ for some f_1, f_2 , and f_3 ;
 (U.4) For all $\mu, \mu' \in \mathcal{P}_\psi(\mathcal{C})$,

$$\int_{\mathcal{C}} \left[g(x, \mu) - g(x, \mu') + \int_0^T (f_1(t, x, \mu) - f_1(t, x, \mu')) dt \right] (\mu - \mu')(dx) \leq 0. \quad (10)$$

Theorem 3.8. *Suppose (U) holds. Then there is at most one solution of the MFG.*

Corollary 3.9. *Suppose (E) and (U) hold. Then there exists a unique solution of the MFG.*

Remark 3.10. The following simple extension of the above formulation allows more heterogeneity among agents. Work instead on a probability space $\Omega = \Omega' \times \mathbb{R}^d \times \mathcal{C}$, where Ω' is some measurable space which will model additional time-zero randomness. We may then fix an initial law $\lambda_0 \in \mathcal{P}(\Omega' \times \mathbb{R}^d)$, and let P be the product of λ_0 and Wiener measure. Letting (θ, ξ, W) denote the coordinate maps, we work with the filtration generated by the process $(\theta, \xi, W_s)_{0 \leq s \leq T}$. The data b, σ, f , and g may all depend on θ . In the finite-player game, the agents have i.i.d. initial data (θ^i, ξ^i) , known at time zero, where ξ^i is the initial state and θ^i can encode other differences between the agents. For example, in a price impact model, perhaps a fraction $\rho \in [0, 1]$ of the agents need to liquidate but the rest do not; this can be modeled using such a θ which equals $c > 0$ with probability ρ and 0 otherwise, and setting $g(X, \theta) = \theta |X_T|^2$ for some $c > 0$. This generalization complicates the notation but changes essentially none of the analysis.

4. APPROXIMATE NASH EQUILIBRIA FOR FINITE-PLAYER GAMES

Before proving these theorems, we discuss how a solution of the MFG may be used to construct an approximate Nash equilibrium for the finite-player game, using only distributed controls. Additional assumptions are needed for the approximation results:

Assumption (F).

- (F.1) $b = b(t, x, a)$ has no mean field term;
 (F.2) For all (t, x, μ, q, a) , $f(t, x, \mu, q, a) = f(t, x, \mu^t, q, a)$, where μ^t denotes the image of μ under the map $\mathcal{C} \ni x \mapsto x_{\cdot \wedge t} \in \mathcal{C}$;
 (F.3) The functions b, f , and g are empirically measurable, as in Definition 3.1, using the progressive σ -field on $[0, T] \times \mathcal{C}$, and Borel σ -fields elsewhere;
 (F.4) For each (t, x) , the following functions are continuous at each point satisfying $\mu \sim \mathcal{X}$:

$$\begin{aligned} \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times A &\ni (\mu, q, a) \mapsto f(t, x, \mu, q, a), \\ \mathcal{P}_\psi(\mathcal{C}) \ni \mu &\mapsto g(x, \mu); \end{aligned}$$

- (F.5) There exists $c > 0$ such that, for all (t, x, μ, q, a) ,

$$|g(x, \mu)| + |f(t, x, \mu, q, a)| \leq c \left(\psi(x) + \int \psi d\mu \right).$$

Remark 4.1. The continuity assumption (F.4) is stronger than assumption (E). Indeed, in (E) we required only *sequential* continuity on a subset of the space $\mathcal{P}_\psi(\mathcal{C})$. Assumption (F.2) is simply progressive measurability of f with respect to the measure argument, which in fact was not needed for the results of Section 3. Analogs of the result of this section are possible when (F.1) fails, under stronger continuity requirements. Namely, $\sigma^{-1}b, f$, and g should be continuous in μ uniformly in the other arguments, and $\sigma^{-1}b$ should be uniformly Lipschitz in μ with respect to total variation. However, we refrain from elaborating on this result, as it seems suboptimal and the proof is quite long.

Adhering to the philosophy of the weak formulation, we choose a single convenient probability space on which we define the n -player games, simultaneously for all n . Assumptions (C) and (F) stand throughout this section (as does (S), as always). We fix a solution of the MFG $(\hat{\mu}, \hat{q})$ throughout, whose existence is guaranteed by Theorem 3.5, with corresponding closed-loop control $\hat{\alpha}(t, x)$ (see Remark 3.2). Consider a probability space (Ω, \mathcal{F}, P) supporting a sequence (W^1, W^2, \dots) of independent d -dimensional Wiener processes, independent \mathbb{R}^d -valued random variables (ξ^1, ξ^2, \dots) with common law λ_0 , and processes (X^1, X^2, \dots) satisfying

$$dX_t^i = b(t, X^i, \hat{\alpha}(t, X^i))dt + \sigma(t, X^i)dW_t^i, \quad X_0^i = \xi^i.$$

For each n , let $\mathbb{F}^n = (\mathcal{F}_t^n)_{t \in [0, T]}$ denote the completion of the filtration generated by (X^1, \dots, X^n) by null sets of \mathcal{F} . Let \mathbb{X}^i denote the completion of the filtration generated by X^i . Note that X^i are independent and identically distributed and that the process $(\xi^i, W_t^i)_{0 \leq t \leq T}$ generates the same filtration \mathbb{X}^i , as in Remark 3.2. Abbreviate $\alpha_t^i = \hat{\alpha}(t, X^i)$. These controls are known as distributed controls.

We now describe the n -player game for fixed n . The control space \mathbb{A}_n is the set of all \mathbb{F}^n -progressively measurable A -valued processes; the players have complete information of the other players' state processes. On the other hand, \mathbb{A}_n^n is the n -fold Cartesian product of \mathbb{A}_n , or the set of \mathbb{F}^n -progressively measurable A^n -valued processes. Let μ^n denote the empirical measure of the first n state processes as defined in the introduction by (2). For $\beta = (\beta^1, \dots, \beta^n) \in \mathbb{A}_n^n$, define a measure $P_n(\beta)$ on $(\Omega, \mathcal{F}_T^n)$ by the density

$$\frac{dP_n(\beta)}{dP} := \mathcal{E} \left(\int_0^\cdot \sum_{i=1}^n (\sigma^{-1}b(t, X^i, \beta_t^i) - \sigma^{-1}b(t, X^i, \alpha_t^i)) dW_t^i \right)_T.$$

Under $P_n(\beta)$, for each $i = 1, \dots, n$, X^i is a weak solution of the SDE

$$dX_t^i = b(t, X^i, \beta_t^i)dt + \sigma(t, X^i)dW_t^{\beta^i, i},$$

where

$$W_t^{\beta^i, i} := W_t - \int_0^t [\sigma^{-1}b(t, X^i, \beta_t^i) - \sigma^{-1}b(t, X^i, \alpha_t^i)] dt$$

is a d -dimensional $P_n(\beta)$ -Wiener process. Note that X_0^i are i.i.d. with common law λ_0 under any of the measures $P_n(\beta)$ with $\beta \in \mathbb{A}_n^n$. For $\beta = (\beta^1, \dots, \beta^n) \in \mathbb{A}_n^n$, the value to player i of the strategies β is defined by

$$J_{n,i}(\beta) := \mathbb{E}^{P_n(\beta)} \left[\int_0^T f(t, X^i, \mu^n, q^n(\beta_t), \beta_t^i) dt + g(X^i, \mu^n) \right],$$

where, for $a = (a^1, \dots, a^n) \in A^n$, we define

$$q^n(a) := \frac{1}{n} \sum_{i=1}^n \delta_{a^i}.$$

Note that the joint measurability assumption (F.3) guarantees that $g(X^i, \mu^n)$ is \mathcal{F}_T^n -measurable, while (F.2) and (F.3) ensure that $(f(t, X^i, \mu^n, q^n(\beta_t), \beta_t^i))_{t \in [0, T]}$ and $(b(t, X^i, \beta_t^i))_{t \in [0, T]}$ are progressively measurable with respect to \mathbb{F}^n .

Theorem 4.2. *Assume (C) and (F) hold, and let $(\hat{\mu}, \hat{q})$ denote a solution of the MFG, with corresponding closed-loop control $\hat{\alpha} = \hat{\alpha}(t, x)$ (see Remark 3.2). Then the strategies $\alpha_t^i := \hat{\alpha}(t, X^i)$ form an approximate Nash equilibrium for the finite-player game in the sense that there exists a sequence $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$ such that, for $1 \leq i \leq n$ and $\beta \in \mathbb{A}_n^n$,*

$$J_{n,i}(\alpha^1, \dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots, \alpha^n) \leq J_{n,i}(\alpha^1, \dots, \alpha^n) + \epsilon_n.$$

Remark 4.3. The punchline is that α^i is \mathbb{X}^i -adapted for each i . That is, player i determines his strategy based only on his own state process. As explained earlier, such strategies are said to be *distributed*. The theorem tells us that even with full information, there is an approximate Nash equilibrium consisting of distributed controls, and we know precisely how to construct one using a solution of the MFG. Note that the strategies $(\alpha^i)_{i \in \mathbb{N}}$ also form an approximate Nash equilibrium for any partial-information version of the game, as long as player i has access to (at least) the filtration \mathbb{X}^i generated by his own state process.

5. DISCUSSION OF THE ASSUMPTIONS AND EXAMPLES

This section discusses some important special cases of the assumptions of Sections 3 and 4. Assumptions (C) and (U) are examined first, before we turn to assumptions (S), (E), and (F).

5.1. Assumptions (C) and (U). Condition (C) (resp. (U.1)) is crucial for the fixed point (resp. uniqueness) argument and holds when the Hamiltonian $h(t, x, \mu, q, z, a)$ is concave (resp. strictly concave) in a , for each (t, x, μ, q, z) , which is a common assumption in control theory. For example, condition (C) (resp. (U.1)) holds if b is affine in a and f is concave (resp. strictly concave) in a . More generally, we can get away with quasiconcavity in the previous statements. Note that if $f \equiv 0$ then $A(t, x, \mu, 0) = A$, and thus condition (U.1) fails except in trivial cases. However, condition (C) frequently holds even in the absence of a running reward function $f \equiv 0$; the optimal control in such a case is typically a bang-bang control.

Example 5.1 (Monotone functionals of measures). Here we provide some examples of the monotonicity assumption (U.4) of Theorem 3.8. For any of the following g , we have

$$\int_{\mathcal{C}} [g(x, \mu) - g(x, \mu')] (\mu - \mu')(dx) \leq 0, \quad \forall \mu, \mu' \in \mathcal{P}_{\psi}(\mathcal{C}).$$

- $g(x, \mu) = \phi_1(x) + \phi_2(\mu)$ for some $\phi_1 : \mathcal{C} \rightarrow \mathbb{R}$ and $\phi_2 : \mathcal{P}_{\psi}(\mathcal{C}) \rightarrow \mathbb{R}$. In this case, there is equality for all μ, μ' .
- $g(x, \mu) = |\phi(x) - \int_{\mathcal{C}} \phi(y) \mu(dy)|^2$ for some $\phi : \mathcal{C} \rightarrow \mathbb{R}$. If, for example, $\phi(x) = x$, then this payoff function rewards a player if his state process deviates from the average.
- $g(x, \mu) = -\int_{\mathbb{R}^d} \phi(|x - y|) \mu_T(dy)$, where $\phi : [0, \infty) \rightarrow [0, \infty)$ is bounded, continuous, and positive definite. A special case is when ϕ is bounded, nonincreasing, and convex; see Proposition 2.6 of [20].

5.2. Assumptions (S), (E), and (F). Standard arguments give:

Lemma 5.2. *Assume that $\psi_0 : \mathbb{R}^d \rightarrow [1, \infty)$ is either $\psi_0(x) = 1 + |x|^p$ for some $p \geq 1$ or $\psi_0(x) = e^{p|x|}$ for some $p > 0$, and let $\psi(x) = \sup_{t \in [0, T]} \psi_0(x_t)$. If $\int_{\mathbb{R}^d} \psi_0(x)^2 \lambda_0(dx) < \infty$, $\sigma > 0$, $|\sigma(\cdot, 0)| \in L^2[0, T]$, and $|\sigma(t, x) - \sigma(t, y)| \leq c\|x - y\|$ for some $c > 0$, then (S.2) holds as long as $\sigma^{-1}b$ is bounded.*

The measurability requirement (F.3) is unusual, but not terribly restrictive. The more difficult assumption to verify is that of continuity, (F.4). Common assumptions in the literature involve continuity with respect to the topology of weak convergence or more generally a Wasserstein metric. For a separable Banach space $(E, \|\cdot\|_E)$ and $p \geq 1$, let

$$\mathcal{W}_{E,p}^p(\mu, \mu') := \inf \left\{ \int_E \|x - y\|_E^p \pi(dx, dy) : \pi \in \mathcal{P}(E \times E) \text{ with marginals } \mu, \mu' \right\},$$

When $\psi_{E,p}(x) = 1 + \|x\|_E^p$, it is known that $\mathcal{W}_{E,p}$ metrizes the weakest topology making the map $\mathcal{P}_{\psi_{E,p}}(E) \ni \mu \mapsto \int \phi d\mu$ continuous for each *continuous* function $\phi \in B_{\psi_{E,p}}(E)$ (see Theorem 7.12 of [41]). Thus $\mathcal{W}_{E,p}$ is weaker than $\tau_{\psi_{E,p}}(\mathcal{C})$, which proves the following result.

Lemma 5.3. *Let $\psi = \psi_{\mathcal{C},p}$, $p \geq 1$. Suppose f and g are (sequentially) continuous in (μ, q, a) at points with $\mu \sim \mathcal{X}$, for each (t, x) , using the metric $\mathcal{W}_{\mathcal{C},p}$ on $\mathcal{P}_\psi(\mathcal{C})$. Then (F.4) holds.*

In most applications the coefficients are Markovian; that is, $f(t, x, \mu, q, a) = \hat{f}(t, x_t, \mu_t, q, a)$ for some $\hat{f} : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$. Note that for any $\mu, \mu' \in \mathcal{P}(\mathcal{C})$, $p \geq 1$, and $t \in [0, T]$,

$$\mathcal{W}_{\mathbb{R}^d,p}(\mu_t, \mu'_t) \leq \mathcal{W}_{\mathcal{C},p}(\mu, \mu'),$$

and thus the previous proposition includes Markovian data. Note also that assumption (F.4) demands continuity in the measure argument only at the points which are equivalent to \mathcal{X} . Of course, if σ does not depend on X or is uniformly bounded from below, then $\mathcal{X}_t \sim \mathcal{L}$ for all $t > 0$, and thus in the Markovian case we need only to check that \hat{f} is continuous at points which are equivalent to Lebesgue measure. At no point was a Markov property of any use, and this is why we chose to allow path-dependence in each of the coefficients. Moreover, continuity in the spatial variable was never necessary either. Indeed, we require only that $dX_t = \sigma(t, X)dW_t$ admits a strong solution, as in assumption (S.2), which of course covers the usual Lipschitz assumption. The most common type of mean field interaction is scalar and Markovian, so we investigate such cases carefully.

Proposition 5.4 (Scalar dependence on the measure). *Consider a function of the form*

$$f(t, x, \mu, q, a) = \int_{\mathcal{C}} F(t, x_t, y_t, q, a) \mu(dy) = \int_{\mathbb{R}^d} F(t, x_t, y, q, a) \mu_t(dy)$$

where $F : [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$ is jointly measurable and jointly continuous in its last two arguments whenever the first three are fixed. Let $\psi_0 : \mathbb{R}^d \rightarrow [1, \infty)$ be lower semicontinuous, and suppose there exists $c > 0$ such that

$$\sup_{(t,a) \in [0,T] \times A} |F(t, x, y, q, a)| \leq c(\psi_0(x) + \psi_0(y))$$

for all $(x, y) \in \mathbb{R}^d \times \mathbb{R}^d$. Let $\psi(x) = \sup_{t \in [0, T]} \psi_0(x_t)$ for $x \in \mathcal{C}$. Then f satisfies the relevant parts of assumptions (S.3), (S.4), (E), (F).

Proof. Note that $\psi : \mathcal{C} \rightarrow [1, \infty)$ is lower-semicontinuous and thus measurable. Note also that the function $\mathcal{C} \ni y \mapsto F(t, x, y_t, q, a) \in \mathbb{R}$ is in $B_\psi(\mathcal{C})$ for each $(t, x, q, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}(A) \times A$, and thus f is indeed well defined for $\mu \in \mathcal{P}_\psi(\mathcal{C})$. Property (F.2) is obvious, and property (F.5) follows from the inequality

$$|f(t, x, \mu, q, a)| \leq c \left(\psi_0(x_t) + \int_{\mathcal{C}} \psi_0(y_t) \mu(dy) \right).$$

The measurability assumption (F.3) is easy to verify. Condition (E) will follow from (F.4), which we prove now.

Fix $(t, x) \in [0, T] \times \mathcal{C}$, and let $E = \mathcal{P}(A) \times A$. Let $F_0(y, \eta) := F(t, x_t, y, \eta)$ for $(y, \eta) \in \mathbb{R}^d \times E$. Fix $(\mu, \eta) \in \mathcal{P}_\psi(\mathcal{C}) \times E$ and a net $(\mu^\alpha, \eta^\alpha)$ converging to (μ, η) . We also have $\mu_t^\alpha \rightarrow \mu_t$ in $\tau_{\psi_0}(\mathbb{R}^d)$. Note that

$$f(t, x, \mu^\alpha, \eta^\alpha) - f(t, x, \mu, \eta) = \int_{\mathbb{R}^d} (F_0(y, \eta^\alpha) - F_0(y, \eta)) \mu_t^\alpha(dy) + \int_{\mathbb{R}^d} F_0(y, \eta) (\mu_t^\alpha - \mu_t)(dy)$$

The second term clearly tends to zero. For the first term, fix $\epsilon > 0$. Since E is compact metric, the function $\mathbb{R}^d \ni y \mapsto F_0(y, \cdot) \in C(E)$ is measurable, using the Borel σ -field generated by the supremum norm on the space $C(E)$ of continuous real-valued functions of E ; see Theorem 4.55 of [2]. Thus, by Lusin's theorem (12.8 of [2]), there exists a compact set $K \subset \mathbb{R}^d$ such that $\int_{K^c} \psi_0 d\mu_t < \epsilon$ and $K \ni y \mapsto F_0(y, \cdot) \in C(E)$ is continuous. Since $|F_0(y, \eta')| \leq c(\psi_0(x_t) + \psi_0(y))$ for all $(y, \eta') \in \mathbb{R}^d \times E$,

$$\left| \int_{\mathbb{R}^d} (F_0(y, \eta^\alpha) - F_0(y, \eta)) \mu_t^\alpha(dy) \right| \leq \sup_{y \in K} |F_0(y, \eta^\alpha) - F_0(y, \eta)| + 2c \int_{K^c} (\psi_0(x_t) + \psi_0(y)) \mu_t^\alpha(dy).$$

It follows from the compactness of E and Lemma 5.6 below that the restriction of F_0 to $K \times E$ is uniformly continuous. Since K is compact, we use Lemma 5.6 again in the other direction to get $\sup_{y \in K} |F_0(y, \eta^\alpha) - F_0(y, \eta)| \rightarrow 0$. Since also

$$\lim \int_{K^c} (\psi_0(x_t) + \psi_0(y)) \mu_t^\alpha(dy) = \int_{K^c} (\psi_0(x_t) + \psi_0(y)) \mu_t(dy) \leq (1 + \psi_0(x_t))\epsilon,$$

we have

$$\limsup \left| \int_{\mathbb{R}^d} (F_0(y, \eta^\alpha) - F_0(y, \eta)) \mu_t^\alpha(dy) \right| \leq 2c(1 + \psi_0(x_t))\epsilon.$$

□

Corollary 5.5. *Let F and ψ_0 be as in Proposition 5.4, and suppose*

$$f(t, x, \mu, q, a) = G \left(t, x_t, \int_{\mathbb{R}^d} F(t, x_t, y, q, a) \mu_t(dy), q, a \right),$$

where $G : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$ is jointly measurable and continuous in its last three arguments. If also

$$|G(t, x, y, q, a)| \leq c(\psi_0(x) + |y|)$$

for some $c > 0$, then f satisfies the relevant parts of assumptions (S.3), (S.4), (E), (F).

We will occasionally need the following simple lemma, which was used in the proof of Proposition 5.4. Its proof is straightforward and thus omitted.

Lemma 5.6. *Let E and K be topological spaces with K compact, let $G : E \times K \rightarrow \mathbb{R}$, and let $x_0 \in E$ be fixed. Then G is jointly continuous at points of $\{x_0\} \times K$ if and only if $G(x_0, \cdot)$ is continuous and $x \mapsto \sup_{y \in K} |G(x, y) - G(x_0, y)|$ is continuous at x_0 .*

Example 5.7 (Geometric Brownian motion). Requiring $\sigma^{-1}b$ to be bounded rather than σ^{-1} and b each to be bounded notably allows for state processes of a geometric Brownian motion type. For example, if $d = 1$, our assumptions allow for coefficients of the form

$$\begin{aligned} b(t, x, \mu, a) &= \hat{b}(t, \mu, a)x_t, \\ \sigma(t, x) &= \hat{\sigma}(t)x_t, \end{aligned}$$

where $\hat{\sigma}(t) > 0$ for all t and $\hat{\sigma}^{-1}\hat{b}$ is bounded.

Remark 5.8. We close the subsection with a remark on the assumption of boundedness of $\sigma^{-1}b$, which could certainly be relaxed. The reason for this assumption lies in the BSDE (11) for the value function; boundedness of $\sigma^{-1}b$ equates to a standard Lipschitz driver, as covered in [36]. The results of Hamadene and Lepeltier in [23] may be applied if b and σ have linear growth in x and σ is bounded below, but this increases the technicalities and rules out a direct application of the results of [24]. However, we only really need [24] in order to treat mean field interactions in the control, and thus our analysis should still work under appropriate linear growth assumptions. Our assumptions of boundedness of $\sigma^{-1}b$ and compactness of A unfortunately rule out common linear-quadratic models, but, nonetheless, the same general techniques could be used to study a large class of linear-quadratic problems (still, of course, with uncontrolled volatility) in which both these assumptions fail. More care is required in the choice of admissible controls, and the BSDE for the value function becomes quadratic in z ; this program was carried out for stochastic optimal control problems in [18], and could presumably be adapted to mean field games.

5.3. Additional Examples. Corollary 5.5 allows us to treat many mean field interactions which are not weakly continuous, as they may involve integrals of discontinuous functions. This is useful in the following examples.

Example 5.9 (Rank effects). Suppose an agent's reward depends on the rank of his state process among the population. That is, suppose $d = 1$ and $f(t, x, \mu, q, a)$ involves a term of the form $G(\mu_t(-\infty, x_t])$, where $G : [0, 1] \rightarrow \mathbb{R}$ is continuous. Such terms with G monotone are particularly interesting for applications, as suggested for a model of oil production in [22]. The intuition is that an oil producer prefers to produce before his competitors, in light of the uncertainty about the longevity of the oil supply. The state process X represents oil reserves, and G should be decreasing in their model. Proposition 5.4 shows that the inclusion of such terms as $\mu_t(-\infty, x_t]$ in f or g is compatible with all of our assumptions. If b contains such rank effects, no problem is posed for assumptions (S) and (E), but of course (F.1) is violated.

Example 5.10 (Types). In [25], Huang, Caines, and Malhamé consider multiple types of agents, and a dependence on the mean field within each type. The number of types is fixed, and an agent cannot change type during the course of the game. Using the construction of Remark 3.10, we may model this by giving each agent a random but i.i.d. type at time zero. Alternatively, in some models an agent's type may change with his state (or with time, or with his strategy); for example, a person's income bracket depends on his wealth. Suppose, for example, that $A_1, A_2, \dots, A_m \subset \mathbb{R}^d$ are Borel sets of positive Lebesgue measure, and define $F_i : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathcal{P}(\mathbb{R}^d)$ by $F_i(\mu)(B) := \mu(B \cap A_i) / \mu(A_i)$ when $\mu(A_i) > 0$ and $F_i(\mu) = 0$ otherwise. As long as σ is bounded away from zero, then $\mathcal{X}_t \sim \mathcal{L}$ where \mathcal{L} is again Lebesgue measure on \mathbb{R}^d , and indeed F_i are $\tau_1(\mathbb{R}^d)$ -continuous at points $\mu \sim \mathcal{X}_t$. So we can treat functionals of the form

$$f(t, x, \mu, q, a) = G(t, x_t, F(\mu_t), q, a),$$

where $F = (F_1, \dots, F_m)$, and $G : [0, T] \times \mathbb{R}^d \times (\mathcal{P}(\mathbb{R}^d))^m \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$.

6. APPLICATIONS REVISITED

Before proving the main results, we return briefly to the models presented in Section 2, for which we demonstrate the applicability of the existence and approximation theorems (3.5 and 4.2).

6.1. Price impact models. We restrict our attention to finite-volume order books. We suppose that $A \subset \mathbb{R}$ is a compact interval containing the origin, $c' : A \rightarrow \mathbb{R}$ is continuous and nondecreasing, $\sigma > 0$, $f : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ and $g : \mathbb{R} \rightarrow \mathbb{R}$ are measurable, and finally that there exists $c_1 > 0$ such that

$$|f(t, x)| + |g(x)| \leq c_1 e^{c_1 |x|}, \text{ for all } (t, x) \in [0, T] \times \mathbb{R}.$$

Let $c(x) = \int_0^x c'(a) da$. Assume that X_0^i are i.i.d. and that their common distribution $\lambda_0 \in \mathcal{P}(\mathbb{R})$ satisfies $\int_{\mathbb{R}} e^{p|x|} \lambda_0(dx) < \infty$ for all $p > 0$. In the notation of the paper, we have $b(t, x, \mu, a) = a$, $\sigma(t, x) = \sigma$, $f(t, x, \mu, q, a) = \gamma x_t \int_A c' dq - c(a) - f(t, x_t)$, $g(x, \mu) = g(x_T)$, and $\psi(x) = e^{c_1 \|x\|}$.

It is quite easy to check the assumptions of the previous sections, at least with the help of Lemma 5.2 below, yielding the following theorem. Moreover, in this simple case we can estimate the rate of convergence, as proven at the end of Section 8.

Proposition 6.1. *Under the above assumptions, the existence and approximation theorems 3.5 and 4.2 apply to the price impact model. Moreover, in the approximation theorem, there exists a constant $C > 0$ such that*

$$\epsilon_n \leq C / \sqrt{n}.$$

6.2. Flocking models. To work around the degeneracy of the diffusion (X^i, V^i) , we consider only V^i as the state variable, and recover X^i by making the coefficients path-dependent. Let $b(t, v, \mu, a) = a$, $\sigma > 0$ constant, $g \equiv 0$, and $A \subset \mathbb{R}^d$ compact convex. Define $\iota : [0, T] \times \mathcal{C} \rightarrow \mathbb{R}^d$ and $I : [0, T] \times \mathcal{P}(\mathcal{C}) \rightarrow \mathcal{P}(\mathbb{R}^d)$ by

$$\iota(t, v) := \int_0^t v_s ds, \quad I(t, \mu) := \mu \circ \iota(t, \cdot)^{-1}.$$

Note that $\iota(t, V^i)$ represents the position of the individual at time t ; we are assuming each individual starts at the origin to keep the notation simple and consistent, although any initial distribution of positions could be accounted for by using the construction of Remark 3.10. For flocking models, (5) is captured by choosing a running reward function of the form:

$$f^{(1)}(t, v, \mu, a) = -|\alpha|_R^2 - \left| \int_{\mathcal{C}} \mu(dv')(v'_t - v_t) \phi(|\iota(t, v' - v)|) \right|_Q^2.$$

The minus signs are only to turn the problem into a maximization, to be consistent with the notation of the rest of the paper. Recall that $\phi : [0, \infty) \rightarrow [0, \infty)$ is nonincreasing and thus Borel measurable. Assume the initial data V^i are i.i.d. and square-integrable, with law $\lambda_0 \in \mathcal{P}_2(\mathbb{R}^d)$. Take $\psi(x) = 1 + \|x\|^2$ for $x \in \mathcal{C}$. For the nearest neighbor model, we use

$$f^{(2)}(t, v, \mu, a) = -|\alpha|_R^2 - \left| \frac{c}{I(t, \mu)(B(\iota(t, v), r))} \int_{\mathcal{C}} \mu(dv')(v'_t - v_t) \mathbf{1}_{B(\iota(t, v), r)}(\iota(t, v')) \right|_Q^2,$$

where $r > 0$ was given, and $B(x, r')$ denotes the closed ball of radius r' centered at x . Consider the second term above to be zero whenever $I(t, \mu)(B(\iota(t, v), r)) = 0$. Finally, for the k -nearest-neighbor model, we choose $\eta \in (0, 1)$ to represent a fixed percentage of neighbors, which amounts to keeping k/n fixed in the finite-player game as we send $n \rightarrow \infty$. We define $r : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \rightarrow [0, \infty)$ by

$$r(\mu, x, y) := \inf \{ r' > 0 : \mu(B(x, r')) \geq y \},$$

and

$$f^{(3)}(t, v, \mu, a) = -|\alpha|_R^2 - \left| \frac{c}{\eta} \int_{\mathcal{C}} \mu(dv')(v'_t - v_t) \mathbf{1}_{B(\iota(t, v), r(I(t, \mu), \iota(t, v), \eta))}(\iota(t, v')) \right|_Q^2.$$

It is straightforward to check that the objective (7) for the nearest neighbor models is equivalent to maximizing

$$\mathbb{E} \int_0^T f^{(1)}(t, V^i, \mu^n, \alpha_t^i) dt, \quad \text{where } \mu^n = \frac{1}{n} \sum_{j=1}^n \delta_{V^j},$$

replacing $f^{(1)}$ by $f^{(2)}$ in the case of the k -nearest neighbor model.

Proposition 6.2. *Under the above assumptions, the existence and approximation theorems 3.5 and 4.2 apply to each of the flocking models.*

Proof. Assumptions (S.1), (S.4), (S.5), (C), (F.1), (F.2), and (F.5) are easy to check. Lemma 5.2 below takes care of (S.2). Also, (S.3) and (F.3) are clear for $f^{(1)}$ and $f^{(2)}$, and follow from Lemma 6.3 below for $f^{(3)}$. It remains to check the continuity assumption (F.4). For $f^{(1)}$, this follows from Proposition 5.4 below. Apply Itô's formula to tW_t to get

$$\iota(t, X) = \int_0^t X_s ds = \int_0^t (\xi + \sigma W_s) ds = t\xi + \sigma t W_t - \sigma \int_0^t s dW_s.$$

Since ξ and W are independent, we see that $I(t, \mathcal{X}) \sim \mathcal{L}$ for $t \in (0, T]$, where \mathcal{L} denotes Lebesgue measure on \mathbb{R}^d . Hence $I(t, \mu) \sim \mathcal{L}$ for $\mu \sim \mathcal{X}$, and so $\mu \mapsto 1/I(t, \mu)(B(x, r))$ is $\tau_\psi(\mathcal{C})$ -continuous at points $\mu \sim \mathcal{X}$, for each $(x, r) \in \mathbb{R}^d \times (0, \infty)$. This along with Proposition 5.4 below establish (F.4) for $f^{(2)}$. Finally, we prove (F.4) for $f^{(3)}$. Fix $(t, v) \in (0, T] \times \mathcal{C}$, and define

$$B_\mu := B(\iota(t, v), r(I(t, \mu), \iota(t, v), \eta)),$$

$$F(\mu) := \int_{\mathcal{C}} (v'_t - v_t) \mathbf{1}_{B_\mu}(\iota(t, v')) \mu(dv'),$$

for $\mu \in \mathcal{P}_\psi(\mathcal{C})$. In light of Lemma 5.3 and the discussion preceding it, it suffices to show F is $\mathcal{W}_{\mathcal{C}, 2}$ -continuous at points $\mu \sim \mathcal{X}$. Let $\mu^n \rightarrow \mu$ in $\mathcal{W}_{\mathcal{C}, 2}$ with $\mu \sim \mathcal{X}$, and note that $I(t, \mu) \sim I(t, \mathcal{X}) \sim \mathcal{L}$.

Then

$$\begin{aligned} F(\mu^n) - F(\mu) &= \int_{\mathcal{C}} (v'_t - v_t) (1_{B_{\mu^n}} - 1_{B_{\mu}}) (\iota(t, v')) \mu(dv') \\ &\quad + \int_{\mathcal{C}} (v'_t - v_t) 1_{B_{\mu^n}} (\iota(t, v')) [\mu^n - \mu](dv') \\ &=: I_n + II_n. \end{aligned}$$

Note that $I(t, \mu^n) \rightarrow I(t, \mu)$ weakly, and thus $r(I(t, \mu^n), \iota(t, x), \eta) \rightarrow r(I(t, \mu), \iota(t, x), \eta)$ by Lemma 6.3. Since $1_{B_{\mu^n}} \rightarrow 1_{B_{\mu}}$ holds \mathcal{L} -a.e. (and thus $I(t, \mu)$ -a.e.) and $\int_{\mathcal{C}} (v'_t - v_t) [\mu^n - \mu](dv') \rightarrow 0$, the dominated convergence theorem yields $I_n \rightarrow 0$. To show $II_n \rightarrow 0$, note that

$$I(t, (v'_t - v_t) \mu^n(dv')) \rightarrow I(t, (v'_t - v_t) \mu(dv')), \text{ weakly.}$$

Since the latter measure is absolutely continuous with respect to Lebesgue measure, Theorem 4.2 of [38] implies

$$II_n = [I(t, (v'_t - v_t) \mu^n(dv')) - I(t, (v'_t - v_t) \mu(dv'))](B_{\mu^n}) \rightarrow 0.$$

In fact, we should consider separately the positive and negative parts of each of the d components of the signed vector measures $(v'_t - v_t) \mu^n(dv')$, since Theorem 4.2 of [38] is stated only for nonnegative real-valued measures. \square

Lemma 6.3. *The function r is empirically measurable, and $r(\cdot, x, y)$ is weakly continuous at points $\mu \sim \mathcal{L}$.*

Proof. To prove measurability, note that for any $c > 0$

$$\{(z, x, y) : r(e_n(z), x, y) > c\} = \left\{ (z, x, y) : \frac{1}{n} \sum_{i=1}^n 1_{B(x, c)}(z_i) < y \right\}$$

is clearly a Borel set in $(\mathbb{R}^d)^n \times \mathbb{R}^d \times (0, 1)$ for each n . To prove continuity, let $\mu_n \rightarrow \mu$ weakly in $\mathcal{P}(\mathbb{R}^d)$ with $\mu \sim \mathcal{L}$. Let $\epsilon > 0$. Since $\mu \sim \mathcal{L}$, the map $r \mapsto \mu(B(x, r))$ is continuous and strictly increasing. Thus the inverse function $r(\mu, x, \cdot)$ is also continuous, and we may find $\delta > 0$ such that $|r(\mu, x, y) - r(\mu, x, z)| < \epsilon$ whenever $|z - y| \leq \delta$. Theorem 4.2 of [38] tells us that $\mu_n(B) \rightarrow \mu(B)$ uniformly over measurable convex sets B , since $\mu \ll \mathcal{L}$. Hence, for n sufficiently large,

$$\sup_{(x, r) \in \mathbb{R}^d \times (0, \infty)} |\mu(B(x, r)) - \mu_n(B(x, r))| < \delta.$$

Thus, for sufficiently large n ,

$$\begin{aligned} r(\mu_n, x, y) &= \inf \{r' > 0 : \mu(B(x, r')) \geq y + (\mu - \mu_n)(B(x, r'))\} \\ &\geq \inf \{r' > 0 : \mu(B(x, r')) \geq y - \delta\} \\ &= r(\mu, x, y - \delta) \geq r(\mu, x, y) - \epsilon, \end{aligned}$$

and similarly

$$r(\mu_n, x, y) \leq \inf \{r' > 0 : \mu(B(x, r')) \geq y + \delta\} = r(\mu, x, y + \delta) \leq r(\mu, x, y) + \epsilon.$$

\square

7. PROOFS OF EXISTENCE AND UNIQUENESS THEOREMS

This section is devoted to the proofs of the existence and uniqueness results of Theorems 3.5 and 3.8. Throughout the section, we work with the canonical probability space described in the second paragraph of Section 3. Since BSDEs will be used repeatedly, it is important to note that the classical existence, uniqueness, and comparison results for BSDEs do indeed hold in our setting, despite the fact that \mathbb{F} is not the Brownian filtration. The purpose of working with the Brownian filtration is of course for martingale representation, which we still have with our slightly larger filtration: It

follows from Theorem 4.33 of [28], for example, that every square integrable \mathbb{F} -martingale $(M_t)_{0 \leq t \leq T}$ admits the representation $M_t = M_0 + \int_0^t \phi_s dW_s$ for some $\phi \in \mathbb{H}^{2,d}$. However, note that in our case the initial value of the solution of a BSDE is random, since \mathcal{F}_0 is not trivial.

To find a fixed point for the law of the control, we will make use of the space \mathcal{M} of positive Borel measures ν on $[0, T] \times \mathcal{P}(A)$ (using the weak topology on $\mathcal{P}(A)$) whose first projection is Lebesgue measure; that is, $\nu([s, t] \times \mathcal{P}(A)) = t - s$ for $0 \leq s \leq t \leq T$. Endow \mathcal{M} with the weakest topology making the map $\nu \mapsto \int \phi d\nu$ continuous for each bounded measurable function $\phi : [0, T] \times \mathcal{P}(A) \rightarrow \mathbb{R}$ for which $\phi(t, \cdot)$ is continuous for each t . This is known as the stable topology, which was studied thoroughly by Jacod and Mémin in [27]. In particular, since A is a compact metrizable space, so is $\mathcal{P}(A)$, and thus so is \mathcal{M} . Note that a measure $\nu \in \mathcal{M}$ disintegrates into $\nu(dt, dq) = \nu_t(dq)dt$, where the measurable map $[0, T] \ni t \mapsto \nu_t \in \mathcal{P}(\mathcal{P}(A))$ is uniquely determined up to almost everywhere equality. For any bounded measurable function $F : \mathcal{P}(A) \rightarrow \mathbb{R}^k$, we extend F to $\mathcal{P}(\mathcal{P}(A))$ in the natural way by defining

$$F(\nu) := \int_{\mathcal{P}(A)} \nu(dq)F(q).$$

In this way, $F(\delta_q) = F(q)$ for $q \in \mathcal{P}(A)$.

Remark 7.1. Because of condition (S.5), the aforementioned convention will not lead to any confusion regarding the meaning of $H(t, x, \mu, \nu, z)$, for $\nu \in \mathcal{P}(\mathcal{P}(A))$. In particular, it is consistent with the relationship $H(t, x, \mu, \nu, a) := \sup_{a \in A} h(t, x, \mu, \nu, z, a)$, since the only dependence of h on ν is outside of the supremum.

For each $(\mu, \nu) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{M}$, we now construct the corresponding control problem. The standing assumptions (S) are in force throughout, and the following construction is valid without any of the other assumptions. Recall the definitions of h and H from (9) in Section 3. That $(t, x, z) \mapsto H(t, x, \mu, \nu_t, z)$ is jointly measurable for each (μ, ν) follows, for example, from the measurable maximum Theorem 18.19 of [2]. Boundedness of $\sigma^{-1}b$ guarantees that H is uniformly Lipschitz in z . Since $\mu \in \mathcal{P}_\psi(\mathcal{C})$, it follows from assumptions (S.2) and (S.4) that $g(X, \mu) \in L^2(P)$ and that $(H(t, X, \mu, \nu_t, 0))_{0 \leq t \leq T} = (\sup_a f(t, X, \mu, \nu_t, a))_{0 \leq t \leq T} \in \mathbb{H}^{2,1}$. Hence the classical result of Pardoux and Peng [36] (or rather a slight extension thereof, as remarked above) applies, and there exists a unique solution $(Y^{\mu, \nu}, Z^{\mu, \nu}) \in \mathbb{H}^{2,1} \times \mathbb{H}^{2,d}$ of the BSDE

$$Y_t^{\mu, \nu} = g(X, \mu) + \int_t^T H(s, X, \mu, \nu_s, Z_s^{\mu, \nu}) ds - \int_t^T Z_s^{\mu, \nu} dW_s. \quad (11)$$

For each $\alpha \in \mathbb{A}$, we may similarly solve the BSDE

$$\begin{aligned} Y_t^{\mu, \nu, \alpha} &= g(X, \mu) + \int_t^T h(s, X, \mu, \nu_s, Z_s^{\mu, \nu, \alpha}, \alpha_s) ds - \int_t^T Z_s^{\mu, \nu, \alpha} dW_s \\ &= g(X, \mu) + \int_t^T f(s, X, \mu, \nu_s, \alpha_s) ds - \int_t^T Z_s^{\mu, \nu, \alpha} dW_s^{\mu, \alpha}. \end{aligned}$$

Since $W^{\mu, \alpha}$ is a Wiener process under $P^{\mu, \alpha}$ and $Y^{\mu, \alpha}$ is adapted, we get

$$Y_t^{\mu, \nu, \alpha} = \mathbb{E}^{\mu, \alpha} \left[g(X, \mu) + \int_t^T f(s, X, \mu, \nu_s, \alpha_s) ds \middle| \mathcal{F}_t \right].$$

In particular, $\mathbb{E}[Y_0^{\mu, \nu, \alpha}] = J^{\mu, \nu}(\alpha)$.

It is immediate from the comparison principle for BSDEs (e.g. Theorem 2.2 of [30]) that $\mathbb{E}[Y_0^{\mu, \nu}] \geq \mathbb{E}[Y_0^{\mu, \nu, \alpha}] = J^{\mu, \nu}(\alpha)$ for each $\alpha \in \mathbb{A}$, and thus $\mathbb{E}[Y_0^{\mu, \nu}] \geq V^{\mu, \nu}$. By a well-known measurable selection theorem (e.g. Theorem 18.19 of [2]), there exists a function $\hat{\alpha} : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathbb{R}^d \rightarrow A$ such that

$$\hat{\alpha}(t, x, \mu, z) \in A(t, x, \mu, z), \quad \text{for all } (t, x, \mu, z), \quad (12)$$

and such that for each μ the map $(t, x, z) \mapsto \hat{\alpha}(t, x, \mu, z)$ is jointly measurable with respect to the progressive σ -field on $[0, T] \times \mathcal{C}$ and $\mathcal{B}(\mathbb{R}^d)$. Letting

$$\alpha_t^{\mu, \nu} := \hat{\alpha}(t, X, \mu, Z_t^{\mu, \nu}), \quad (13)$$

the uniqueness of solutions of BSDEs implies $Y_t^{\mu, \nu} = Y_t^{\mu, \nu, \alpha^{\mu, \nu}}$, which in turn implies $V^{\mu, \nu} = J^{\mu, \nu}(\alpha^{\mu, \nu})$ since $J^{\mu, \nu}(\alpha^{\mu, \nu}) \leq V^{\mu, \nu}$.

The process $\alpha^{\mu, \nu}$ is an optimal control, but so is any process in the set

$$\mathbb{A}(\mu, \nu) := \{\alpha \in \mathbb{A} : \alpha_t \in A(t, X, \mu, Z_t^{\mu, \nu}) \text{ dt} \times dP - a.e.\}. \quad (14)$$

Define $\Phi : \mathcal{P}_\psi(\mathcal{C}) \times \mathbb{A} \rightarrow \mathcal{P}(\mathcal{C}) \times \mathcal{M}$ by

$$\Phi(\mu, \alpha) := (P^{\mu, \alpha} \circ X^{-1}, \delta_{P^{\mu, \alpha} \circ \alpha_t^{-1}}(dq)dt)$$

The goal now is to find a point $(\mu, \nu) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{M}$ for which there exists $\alpha \in \mathbb{A}(\mu, \nu)$ such that $(\mu, \nu) = \Phi(\mu, \alpha)$. In other words, we seek a fixed point of the set-valued map $(\mu, \nu) \mapsto \Phi(\mu, \mathbb{A}(\mu, \nu)) := \{\Phi(\mu, \alpha) : \alpha \in \mathbb{A}(\mu, \nu)\}$. Note that under condition (U.1), $\alpha^{\mu, \nu}$ is the unique element of $\mathbb{A}(\mu, \nu)$ (up to almost everywhere equality), and this reduces to a fixed point problem for a single-valued function.

Remark 7.2. It is worth emphasizing that the preceding argument demonstrates that the set $\mathbb{A}(\mu, \nu)$ is always nonempty, under only the standing assumptions (S).

Remark 7.3. The main difficulty in the analysis is the adjoint process $Z^{\mu, \nu}$. Note that for each (μ, ν) there exists a progressively measurable function $\zeta_{\mu, \nu} : [0, T] \times \mathcal{C} \rightarrow \mathbb{R}^d$ such that $Z_t^{\mu, \nu} = \zeta_{\mu, \nu}(t, X)$. If we choose a measurable selection $\hat{\alpha}$ as in (12), any weak solution of the following McKean-Vlasov SDE provides a solution of the MFG:

$$\begin{cases} dX_t = b(t, X, \mu, \hat{\alpha}(t, X, \mu, \zeta_{\mu, \nu}(t, X)))dt + \sigma(t, X)dW_t, \\ X \sim \mu, \mu \circ (\hat{\alpha}(t, \cdot, \mu, \zeta_{\mu, \nu}(t, \cdot)))^{-1} = \nu_t \text{ a.e.} \end{cases}$$

The notation $X \sim \mu$ means that μ should equal the law of X . This map $\zeta_{\mu, \nu}$ is typically quite inaccessible, which is why we do not appeal to any existing results on McKean-Vlasov equations, even when ν is not present. All such results require some kind of continuity of the map $(x, \mu) \mapsto b(t, x, \mu, \hat{\alpha}(t, x, \mu, \zeta_{\mu, \nu}(t, x)))$, as far as the authors know. It is possible to make assumptions on the data which would guarantee, for example, that $\zeta_{\mu, \nu}(t, \cdot)$ is continuous, but continuous dependence on μ would be a much trickier matter.

7.1. Some results of set-valued analysis. We precede the main proofs with some useful lemmas. Without assumption (U), the optimal controls need not be unique, and thus we will need a fixed point theorem for set-valued maps. We first summarize some terminology from set-valued analysis.

For a point y in a metric space (E, d) and $\delta > 0$, let $B(y, \delta)$ denote the open ball of radius δ centered at y . Similarly, for $F \subset E$, let $B(F, \delta) = \{x \in E : \inf_{y \in F} d(x, y) < \delta\}$. For two subsets F, G of E , we (abusively) define

$$d(F, G) := \sup_{y \in G} d(F, y) = \sup_{y \in G} \inf_{x \in F} d(x, y).$$

Note that d is not symmetric. If K is another metric space, a set-valued function $\Gamma : K \rightarrow 2^E$ is said to be *upper hemicontinuous* at $x \in K$ if for all $\epsilon > 0$ there exists $\delta > 0$ such that $\Gamma(B(x, \delta)) \subset B(\Gamma(x), \epsilon)$. It is straightforward to prove that Γ is upper hemicontinuous at $x \in K$ if and only if $d(\Gamma(x), \Gamma(x_n)) \rightarrow 0$ for every sequence x_n converging to x .

In order to relax somewhat the convexity assumption of Kakutani's fixed point theorem, we adapt results of Cellina in [13] to derive a slight generalization of Kakutani's theorem, which will assist in the proof of Theorem 3.5.

Proposition 7.4. *Let K be a compact convex metrizable subset of a locally convex topological vector space, and let E be a normed vector space. Suppose $\Gamma : K \rightarrow 2^E$ is upper hemicontinuous and has closed and convex values, and suppose $\phi : K \times E \rightarrow K$ is continuous. Then there exists $x \in K$ such that $x \in \phi(x, \Gamma(x)) := \{\phi(x, y) : y \in \Gamma(x)\}$.*

Proof. Let $Gr(\Gamma) := \{(x, y) \in K \times E : y \in \Gamma(x)\}$ denote the graph of Γ . By Cellina's result (Theorem 1 of [13]), for each positive integer n we may find a continuous (single-valued) function $\gamma_n : K \rightarrow E$ such that the graph of γ_n is contained in the $1/n$ neighborhood of $Gr(\Gamma)$. That is, for all $x \in K$,

$$d((x, \gamma_n(x)), Gr(\Gamma)) := \inf \{d((x, \gamma_n(x)), (y, z)) : y \in K, z \in \Gamma(y)\} < 1/n,$$

where d denotes some metric on $K \times E$. Since $K \ni x \mapsto \phi(x, \gamma_n(x)) \in K$ is continuous, Schauder's fixed point theorem implies that there exists $x_n \in K$ such that $x_n = \phi(x_n, \gamma_n(x_n))$. By Lemma 17.8 and Theorem 17.10 of [2], $\Gamma(K) := \bigcup_{x \in K} \Gamma(x) \subset X$ is compact and $Gr(\Gamma)$ is closed. Thus $Gr(\Gamma) \subset K \times \Gamma(K)$ is compact. Since $d((x_n, \gamma_n(x_n)), Gr(\Gamma)) \rightarrow 0$ and $Gr(\Gamma)$ is compact, there exist a subsequence x_{n_k} and a point $(x, y) \in Gr(\Gamma)$ such that $(x_{n_k}, \gamma_{n_k}(x_{n_k})) \rightarrow (x, y)$. This completes the proof, since $y \in \Gamma(x)$ and since continuity of ϕ yields

$$x = \lim x_{n_k} = \lim \phi(x_{n_k}, \gamma_{n_k}(x_{n_k})) = \phi(x, y).$$

□

A special case of Berge's maximum theorem (17.31 of [2]) will be useful:

Theorem 7.5 (Berge's Theorem). *Let E be a metric space, K a compact metric space, and $\phi : E \times K \rightarrow \mathbb{R}$ a continuous function. Then $\gamma(x) := \max_{y \in K} \phi(x, y)$ is continuous, and the following set-valued function is upper hemicontinuous and compact-valued:*

$$E \ni x \mapsto \arg \max_{y \in K} \phi(x, y) := \{y \in K : \gamma(x) = \phi(x, y)\} \in 2^K$$

7.2. Proof of Theorem 3.5 (existence). We now turn toward the proof of Theorem 3.5. In what follows, we always use the topology $\tau_\psi(\mathcal{C})$ on $\mathcal{P}_\psi(\mathcal{C})$, except when stated otherwise. Despite its simplicity, we state the following result as a lemma for later references.

Lemma 7.6. *Let (E, \mathcal{E}) and (F, \mathcal{F}) be measurable spaces, and let $\mu, \nu \in \mathcal{P}(E)$ with $\nu \ll \mu$. If $X : E \rightarrow F$ is measurable, then*

$$\frac{d\nu \circ X^{-1}}{d\mu \circ X^{-1}} \circ X = \mathbb{E}^\mu \left[\frac{d\nu}{d\mu} \middle| X \right] \quad \mu - a.s.$$

Lemma 7.7. *For any $q \in \mathbb{R}$ with $|q| \geq 1$, we have (recall that $\mathcal{X} := P \circ X^{-1}$)*

$$M_q := \sup_{(\mu, \alpha) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathbb{A}} \int (d\Phi(\mu, \alpha)/d\mathcal{X})^q d\mathcal{X} < \infty. \quad (15)$$

Proof. Recall that $\sigma^{-1}b$ is bounded, say by $c > 0$. Fix $(\mu, \alpha) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathbb{A}$. Letting $N_t := \int_0^t \sigma^{-1}b(t, X, \mu, \alpha_t) dW_t$, we see that $[N, N]_T \leq Tc^2$, and thus, since $q(q-1) \geq 0$,

$$\mathcal{E}(N)_T^q = \mathcal{E}(qN)_T \exp(q(q-1)[N, N]_T/2) \leq \mathcal{E}(qN)_T \exp(q(q-1)Tc^2/2).$$

Hence, Lemma 7.6 and Jensen's inequality yield

$$\int (d\Phi(\mu, \alpha)/d\mathcal{X})^q d\mathcal{X} = \mathbb{E}[\mathbb{E}[dP^{\mu, \alpha}/dP | X]^q] \leq \mathbb{E}[(dP^{\mu, \alpha}/dP)^q] \leq \exp(q(q-1)Tc^2/2).$$

Since this bound is independent of (μ, α) , we indeed have $M_q < \infty$. □

In terms of the notation from Lemma 7.7, let $M := \max(M_2, M_{-1})$. Let

$$\mathcal{Q} := \left\{ \mu \in \mathcal{P}_\psi(\mathcal{C}) : \mu \sim \mathcal{X}, \int (d\mu/d\mathcal{X})^2 d\mathcal{X} \leq M, \int (d\mathcal{X}/d\mu) d\mathcal{X} \leq M \right\}. \quad (16)$$

By construction, the range of Φ is contained in $\mathcal{Q} \times \mathcal{M}$. Critical to our fixed point theorem is the following compactness result, which probably exists in various forms elsewhere in the literature. Part of the result may be found, for example, in Lemma 6.2.16 of [16]. But, for lack of a concise reference, and to keep the paper fairly self-contained, we include a complete proof of the following:

Proposition 7.8. *The space $(\mathcal{Q}, \tau_\psi(\mathcal{C}))$ is convex, compact, and metrizable. Moreover, $\tau_1(\mathcal{C})$ and $\tau_\psi(\mathcal{C})$ induce the same topology on \mathcal{Q} .*

Proof. Of course, by $\tau_1(\mathcal{C})$ we mean $\tau_\phi(\mathcal{C})$ with $\phi \equiv 1$. Define

$$\begin{aligned} \mathcal{Q}_1 &= \left\{ \mu \in \mathcal{P}(\mathcal{C}) : \mu \ll \mathcal{X}, \int (d\mu/d\mathcal{X})^2 d\mathcal{X} \leq M \right\}, \\ \mathcal{Q}_2 &= \left\{ \mu \in \mathcal{P}(\mathcal{C}) : \mu \sim \mathcal{X}, \int (d\mathcal{X}/d\mu) d\mathcal{X} \leq M \right\}. \end{aligned}$$

Clearly each set is convex. We will show that \mathcal{Q}_1 is compact and metrizable under $\tau_1(\mathcal{C})$, that \mathcal{Q}_2 is $\tau_1(\mathcal{C})$ -closed, and that $\tau_1(\mathcal{C})$ and $\tau_\psi(\mathcal{C})$ induce the same topology on \mathcal{Q}_1 .

Let $q \in \mathbb{R}$ with $|q| \geq 1$. The set $K_q := \{Z \in L^1(\mathcal{X}) : Z \geq 0 \text{ } \mathcal{X} - \text{a.s.}, \int |Z|^q d\mathcal{X} \leq M\}$ is clearly convex. It is also norm-closed: if $Z_n \rightarrow Z$ in $L^1(\mathcal{X})$ with $Z_n \in K_q$, then $Z_n \rightarrow Z$ \mathcal{X} -a.s. along a subsequence, and thus Fatou's lemma yields $\int |Z|^q d\mathcal{X} \leq \liminf \int |Z_n|^q d\mathcal{X} \leq M$. Hence, K_q is weakly closed (see Theorem 5.98 of [2]). For $q > 1$, the set K_q is uniformly integrable and thus weakly compact, by the Dunford-Pettis theorem; moreover, K_q is metrizable, since it is a weakly compact subset of separable Banach space (Theorem V.6.3 of [17]). Now, for $\mu \ll \mathcal{X}$, define $F(\mu) := d\mu/d\mathcal{X}$. Then F is a homeomorphism from $(\mathcal{Q}_2, \tau_1(\mathcal{C}))$ to K_{-1} equipped with the weak topology of $L^1(\mathcal{X})$, and so \mathcal{Q}_2 is $\tau_1(\mathcal{C})$ -closed. Similarly, F is a homeomorphism from $(\mathcal{Q}_1, \tau_1(\mathcal{C}))$ to K_2 with the weak topology, and so $(\mathcal{Q}_1, \tau_1(\mathcal{C}))$ is compact and metrizable.

It remains to prove that $\tau_1(\mathcal{C})$ and $\tau_\psi(\mathcal{C})$ coincide on \mathcal{Q}_1 . Let $\phi \in B_\psi(\mathcal{C})$ with $|\phi| \leq \psi$, $\mu \in \mathcal{P}_\psi(\mathcal{C})$, and $\epsilon > 0$, and define $U = \{\nu \in \mathcal{P}_\psi(\mathcal{C}) : |\int \phi d(\nu - \mu)| < \epsilon\}$. Since $\tau_\psi(\mathcal{C})$ is stronger than $\tau_1(\mathcal{C})$, it suffices to find a $\tau_1(\mathcal{C})$ -neighborhood V of μ with $V \cap \mathcal{Q}_1 \subset U \cap \mathcal{Q}_1$. First, note that for any $c > 0$ and $\nu \in \mathcal{Q}_1$, the Cauchy-Schwarz inequality yields

$$\left(\int_{\{\psi \geq c\}} \psi d\nu \right)^2 \leq \int \left(\frac{d\nu}{d\mathcal{X}} \right)^2 d\mathcal{X} \int_{\{\psi \geq c\}} \psi^2 d\mathcal{X} \leq M \int_{\{\psi \geq c\}} \psi^2 d\mathcal{X}.$$

Since $\int \psi^2 d\mathcal{X} < \infty$ by (S.2), we may find $c > 0$ such that $\int_{\{\psi \geq c\}} \psi d\nu \leq \epsilon/3$ for all $\nu \in \mathcal{Q}_1$. Then, for any $\nu \in \mathcal{Q}_1$,

$$\left| \int \phi d(\nu - \mu) \right| \leq \left| \int_{\{\psi < c\}} \phi d(\nu - \mu) \right| + \left| \int_{\{\psi \geq c\}} \phi d\nu \right| + \left| \int_{\{\psi \geq c\}} \phi d\mu \right| \leq \frac{2\epsilon}{3} + \left| \int_{\{\psi < c\}} \phi d(\nu - \mu) \right|.$$

Set $V = \{\nu \in \mathcal{P}(\mathcal{C}) : |\int_{\{\psi < c\}} \phi d(\nu - \mu)| < \epsilon/3\}$, so that $V \cap \mathcal{Q}_1 \subset U \cap \mathcal{Q}_1$. Since $|\phi| \leq \psi$, we have $\phi 1_{\{\psi < c\}} \in B_1(\mathcal{C})$, and thus $V \in \tau_1(\mathcal{C})$. \square

The next two lemmas pertain to the $Z^{\mu, \nu}$ terms that arise in the BSDE representations above; in particular, a kind of continuity of the map $(\mu, \nu) \mapsto Z^{\mu, \nu}$ is needed.

Lemma 7.9. *Suppose assumption (E) holds. Then for each $(t, x) \in [0, T] \times \mathcal{C}$, the function $\mathcal{Q} \times \mathcal{P}(A) \times \mathbb{R}^d \ni (\mu, q, z) \mapsto H(t, x, \mu, q, z)$ is continuous, and the set-valued function $\mathcal{Q} \times \mathbb{R}^d \ni (\mu, z) \mapsto A(t, x, \mu, z)$ is upper hemicontinuous.*

Proof. Since \mathcal{Q} is metrizable by Lemma 7.8, this is simply a combination of assumption (E) with Theorem 7.5, using $E = \mathcal{Q} \times \mathcal{P}(A) \times \mathbb{R}^d$ and $K = A$. Recall from (S.1) that A is compact. \square

Lemma 7.10. *Suppose assumption (E) holds. Suppose $(\mu^n, \nu^n) \rightarrow (\mu, \nu)$ in $\mathcal{Q} \times \mathcal{M}$, using $\tau_\psi(\mathcal{C})$ on \mathcal{Q} . Then*

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\int_0^T |Z_t^{\mu^n, \nu^n} - Z_t^{\mu, \nu}|^2 dt \right] = 0.$$

Proof. Note that the functions $H(s, x, \mu', \nu', \cdot)$ have the same Lipschitz constant for each (t, x, μ', ν') , coinciding with the uniform bound for $\sigma^{-1}b$. Assumption (S.4) implies

$$\begin{aligned} \mathbb{E} \left[\int_0^T |H(t, X, \mu^n, \nu_t^n, 0)|^2 dt \right] &= \mathbb{E} \left[\int_0^T \sup_{a \in A} |f(t, X, \mu', \nu_t', a)|^2 dt \right] \\ &\leq 2c^2 T \mathbb{E}[\psi^2(X)] + 2c^2 T \rho^2 \left(\int \psi d\mu^n \right) \end{aligned}$$

for all $1 \leq n \leq \infty$, where $(\mu^\infty, \nu^\infty) := (\mu, \nu)$. Since $\mu^n \in \mathcal{P}_\psi(\mathcal{C})$ and $\mu^n \rightarrow \mu$ in $\tau_\psi(\mathcal{C})$ it follows that $\sup_n \int \psi d\mu^n < \infty$. Since ρ is increasing and nonnegative,

$$\sup_n \rho^2 \left(\int \psi d\mu^n \right) = \rho^2 \left(\sup_n \int \psi d\mu^n \right) < \infty. \quad (17)$$

Assumption (S.2) yields $\mathbb{E}[\psi^2(X)] < \infty$. Hence, we will be able to conclude via a convergence result for BSDEs proven by Hu and Peng in [24], as soon as we show that

$$I_n := \mathbb{E} \left[|g(X, \mu^n) - g(X, \mu)|^2 \right] \rightarrow 0,$$

and

$$II_n := \mathbb{E} \left[\left(\int_t^T (H(s, X, \mu, \nu_s, Z_s^{\mu, \nu}) - H(s, X, \mu^n, \nu_s^n, Z_s^{\mu, \nu})) ds \right)^2 \right] \rightarrow 0,$$

for all $t \in [0, T]$.

We first check that the integrands of I_n and II_n are uniformly integrable. Assumption (S.4) gives

$$|g(X, \mu^n) - g(X, \mu)| \leq c \left(2\psi(X) + \rho \left(\int \psi d\mu \right) + \rho \left(\int \psi d\mu_n \right) \right),$$

which is indeed square integrable in light of (S.2) and (17). Note that

$$\begin{aligned} |H(t, X, \mu, \nu_t, Z_t^{\mu, \nu}) - H(t, X, \mu^n, \nu_t^n, Z_t^{\mu, \nu})| &\leq \sup_{a \in A} |f(t, X, \mu, \nu_t, a) + Z_t^{\mu, \nu} \cdot \sigma^{-1}b(t, X, \mu, a) \\ &\quad - f(t, X, \mu^n, \nu_t^n, a) - Z_t^{\mu, \nu} \cdot \sigma^{-1}b(t, X, \mu^n, a)| \\ &\leq |\Delta_t^{f, n}| + |Z_t^{\mu, \nu}| |\Delta_t^{b, n}|. \end{aligned} \quad (18)$$

where

$$\begin{aligned} \Delta_t^{f, n} &:= \sup_{a \in A} |f(t, X, \mu, \nu_t, a) - f(t, X, \mu^n, \nu_t^n, a)|, \text{ and} \\ \Delta_t^{b, n} &:= \sup_{a \in A} |\sigma^{-1}b(t, X, \mu, a) - \sigma^{-1}b(t, X, \mu^n, a)|. \end{aligned}$$

Again, (S.4) lets us bound $|\Delta_t^{f, n}|$ by the same term with which we bounded $|g(X, \mu^n) - g(X, \mu)|$. Since $Z^{\mu, \nu} \in \mathbb{H}^{2,1}$ and $|\Delta_t^{b, n}|$ is bounded, the integrands are indeed uniformly integrable.

It is clear now that $I_n \rightarrow 0$, because of assumption (E) and the dominated convergence theorem. Rewrite II_n as

$$II_n = \mathbb{E} \left[\left| \int_t^T ds \left(\int_{\mathcal{P}(A)} \nu_s(dq) H(s, X, \mu, q, Z_s^{\mu, \nu}) - \int_{\mathcal{P}(A)} \nu_s^n(dq) H(s, X, \mu^n, q, Z_s^{\mu, \nu}) \right) \right|^2 \right].$$

For fixed s and ω , the function $\mathcal{Q} \times \mathcal{P}(A) \ni (\mu', q) \mapsto H(s, X, \mu', q, Z_s^{\mu, \nu})$ is continuous, by Lemma 7.9. Compactness of $\mathcal{P}(A)$ implies that the function $\mathcal{Q} \ni \mu' \mapsto H(s, X, \mu', q, Z_s^{\mu, \nu})$ is continuous, uniformly in q (see Lemma 5.6). Thus

$$\int_{\mathcal{P}(A)} \nu_s^n(dq) H(s, X, \mu^n, q, Z_s^{\mu, \nu}) - \int_{\mathcal{P}(A)} \nu_s^n(dq) H(s, X, \mu, q, Z_s^{\mu, \nu}) \rightarrow 0.$$

By definition of the stable topology of \mathcal{M} , we also have

$$\int_t^T ds \int_{\mathcal{P}(A)} \nu_s^n(dq) H(s, X, \mu, q, Z_s^{\mu, \nu}) \rightarrow \int_t^T ds \int_{\mathcal{P}(A)} \nu_s(dq) H(s, X, \mu, q, Z_s^{\mu, \nu}).$$

It is now clear that $II_n \rightarrow 0$, and the proof is complete. \square

The last ingredient of the proof is to establish the applicability of Proposition 7.4. Note that A is a compact subset of a normed space, say $(A', \|\cdot\|_A)$, and thus \mathbb{A} may also be viewed as a subset of the normed space of (equivalence classes of $dt \times dP$ -a.e. equal) progressively measurable A' -valued processes, with the norm

$$\|\alpha\|_{\mathbb{A}} := \mathbb{E} \int_0^T \|\alpha_t\|_A dt.$$

Lemma 7.11. *Under assumptions (E) and (C), the function $\mathbb{A} : \mathcal{Q} \times \mathcal{M} \rightarrow 2^{\mathbb{A}}$ defined by (14) is upper hemicontinuous and has closed and convex values.*

Proof. Convexity follows immediately from assumption (C). We first show $\mathbb{A}(\cdot)$ has closed values. Let \mathcal{L} denote Lebesgue measure on $[0, T]$. Note that $\|\cdot\|_A$ is bounded on A , and thus $\|\cdot\|_{\mathbb{A}}$ metrizes convergence in $\mathcal{L} \times P$ -measure. To prove closedness, fix a sequence $\alpha^n \in \mathbb{A}(\mu, \nu)$ such that $\|\alpha^n - \alpha\|_{\mathbb{A}} \rightarrow 0$ for some $\alpha \in \mathbb{A}$. By passing to a subsequence, we may assume $\alpha_t^n(\omega) \rightarrow \alpha_t(\omega)$ for all $(t, \omega) \in N$, for some $N \subset [0, T] \times \Omega$ with $\mathcal{L} \times P(N) = 1$. We may assume also that $\alpha_t^n(\omega) \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))$ for all n and $(t, \omega) \in N$. By Theorem 7.5, for each (t, ω) the set $A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega)) \subset A$ is compact, and thus $\alpha_t(\omega) \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))$ for all $(t, \omega) \in N$.

To prove upper hemicontinuity, let $(\mu^n, \nu^n) \rightarrow (\mu, \nu)$ in $\mathcal{Q} \times \mathcal{M}$. We must show that

$$d(\mathbb{A}(\mu, \nu), \mathbb{A}(\mu^n, \nu^n)) = \sup_{\alpha^n \in \mathbb{A}(\mu^n, \nu^n)} \inf_{\alpha \in \mathbb{A}(\mu, \nu)} \mathbb{E} \int_0^T \|\alpha_t^n - \alpha_t\|_A dt \rightarrow 0.$$

Define

$$\begin{aligned} c_t^n(\omega) &:= d\left(A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega)), A(t, X(\omega), \mu^n, Z_t^{\mu^n, \nu^n}(\omega))\right) \\ &= \sup \left\{ \inf \{ \|a - b\|_A : b \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega)) \} : a \in A(t, X(\omega), \mu^n, Z_t^{\mu^n, \nu^n}(\omega)) \right\}. \end{aligned}$$

Lemma 7.10 implies that $Z^{\mu^n, \nu^n} \rightarrow Z^{\mu, \nu}$ in $\mathcal{L} \times P$ -measure; it follows then from upper hemicontinuity of $A(t, x, \cdot, \cdot)$ (Lemma 7.9) that $c^n \rightarrow 0$ in $\mathcal{L} \times P$ -measure as well. Since of course c^n is bounded, the proof will be complete once we establish

$$\sup_{\alpha^n \in \mathbb{A}(\mu^n, \nu^n)} \inf_{\alpha \in \mathbb{A}(\mu, \nu)} \mathbb{E} \int_0^T \|\alpha_t^n - \alpha_t\|_A dt = \mathbb{E} \int_0^T c_t^n dt.$$

To prove that we can pass the infimum and supremum inside of the integrals, we first use Theorem 18.19 of [2] to draw a number of conclusions. First, the map $(t, \omega) \mapsto A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))$

is measurable, in the sense of Definition 18.1 of [2], and thus also weakly measurable since it is compact-valued (see Lemma 18.2 of [2]). Second, there exists a measurable function $\hat{\beta} : [0, T] \times \Omega \times A \rightarrow A$ such that

$$\begin{aligned} \|a - \hat{\beta}(t, \omega, a)\|_A &= \inf \{\|a - b\|_A : b \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))\}, \\ \hat{\beta}(t, \omega, a) &\in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega)). \end{aligned}$$

Note that for any $\alpha^n \in \mathbb{A}$, the process $\hat{\beta}(t, \omega, \alpha_t^n(\omega))$ is in $\mathbb{A}(\mu, \nu)$. Hence, we may exchange the infimum and the expectation to get

$$\inf_{\alpha \in \mathbb{A}(\mu, \nu)} \mathbb{E} \int_0^T \|\alpha_t^n - \alpha_t\|_A dt = \mathbb{E} \int_0^T \inf \{\|\alpha_t^n - b\|_A : b \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))\} dt,$$

It follows from Theorem 7.5 that $a \mapsto \inf \{\|a - b\|_A : b \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))\}$ is continuous for each (t, ω) . Hence, Theorem 18.19 of [2] also tells us that there exists a measurable selection $\hat{\beta}^n : [0, T] \times \Omega \rightarrow A$ such that

$$\begin{aligned} c_t^n(\omega) &= \inf \left\{ \|\hat{\beta}^n(t, \omega) - b\|_A : b \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega)) \right\}, \\ \hat{\beta}^n(t, \omega) &\in A(t, X(\omega), \mu^n, Z_t^{\mu^n, \nu^n}(\omega)). \end{aligned}$$

The process $\hat{\beta}^n(t, \omega)$ is in $\mathbb{A}(\mu^n, \nu^n)$, and so we exchange the supremum and the expectation to get

$$\sup_{\alpha^n \in \mathbb{A}(\mu^n, \nu^n)} \mathbb{E} \int_0^T \inf \{\|\alpha_t^n - b\|_A : b \in A(t, X(\omega), \mu, Z_t^{\mu, \nu}(\omega))\} dt = \mathbb{E} \int_0^T c_t^n dt.$$

□

Proof of Theorem 3.5. The proof of Theorem 3.5 is an application of Proposition 7.4, with $K = \mathcal{Q} \times \mathcal{M}$ and $E = \mathbb{A}$. Let \mathcal{S} denote the vector space of bounded measurable functions $\phi : [0, T] \times \mathcal{P}(A) \rightarrow \mathbb{R}$ such that $\phi(t, \cdot)$ is continuous for each t . Endow \mathcal{S} with the supremum norm, and let \mathcal{S}^* denote its continuous dual space. Note that $\mathcal{M} \subset \mathcal{S}^*$. Let $\mathcal{Y} := B_\psi(\mathcal{C}) \oplus \mathcal{S}$, endowed with the norm

$$\|(\phi, \eta)\|_{\mathcal{Y}} := \sup_{x \in \mathcal{C}} \frac{|\phi(x)|}{\psi(x)} + \sup_{(t, q) \in [0, T] \times \mathcal{P}(A)} |\eta(t, q)|.$$

The dual of \mathcal{Y} is $\mathcal{Y}^* = B_\psi^*(\mathcal{C}) \oplus \mathcal{S}^*$, which contains $\mathcal{Q} \times \mathcal{M}$ as a subset. Using $\tau_\psi(\mathcal{C})$ on \mathcal{Q} , the product topology of $\mathcal{Q} \times \mathcal{M}$ coincides with the topology induced by the weak*-topology of \mathcal{Y}^* . By Lemma 7.7, the function Φ takes values in $\mathcal{Q} \times \mathcal{M}$, noting that $\mathcal{Q} \times \mathcal{M}$ is convex and compact by Lemma 7.8. Let $\tau_{\mathcal{M}}$ denote the topology of \mathcal{M} . To prove that $\Phi : (\mathcal{Q}, \tau_\psi(\mathcal{C})) \times (\mathbb{A}, \|\cdot\|_{\mathbb{A}}) \rightarrow (\mathcal{Q}, \tau_\psi(\mathcal{C})) \times (\mathcal{M}, \tau_{\mathcal{M}})$ is continuous, Lemma 7.8 tells us that it suffices to show that $\Phi : (\mathcal{Q}, \tau_\psi(\mathcal{C})) \times (\mathbb{A}, \|\cdot\|_{\mathbb{A}}) \rightarrow (\mathcal{Q}, \tau_1(\mathcal{C})) \times (\mathcal{M}, \tau_{\mathcal{M}})$ is sequentially continuous. We will instead prove the stronger statement that $\Phi : (\mathcal{Q}, \tau_\psi(\mathcal{C})) \times (\mathbb{A}, \|\cdot\|_{\mathbb{A}}) \rightarrow (\mathcal{Q}, \mathcal{V}_1) \times (\mathcal{M}, \tau_{\mathcal{M}})$ is sequentially continuous, where \mathcal{V}_1 denotes the total variation metric,

$$\mathcal{V}_1(\mu, \nu) := \sup \int \phi d(\mu - \nu),$$

where the supremum is over measurable real-valued functions ϕ with $|\phi| \leq 1$. Denote by $\mathcal{H}(\nu|\mu)$ the relative entropy,

$$\mathcal{H}(\nu|\mu) = \begin{cases} \int \log \frac{d\nu}{d\mu} d\nu & \text{if } \nu \ll \mu \\ +\infty & \text{otherwise.} \end{cases}$$

Now let $(\mu^n, \alpha^n) \rightarrow (\mu, \alpha)$ in $(\mathcal{Q}, \tau_\psi(\mathcal{C})) \times (\mathbb{A}, \|\cdot\|_{\mathbb{A}})$. We first show that $P^{\mu^n, \alpha^n} \rightarrow P^{\mu, \alpha}$. By Pinsker's inequality, it suffices to show $\mathcal{H}(P^{\mu, \alpha} | P^{\mu^n, \alpha^n}) \rightarrow 0$. Since

$$\frac{dP^{\mu^n, \alpha^n}}{dP^{\mu, \alpha}} = \mathcal{E} \left(\int_0^\cdot (\sigma^{-1}b(t, X, \mu^n, \alpha_t^n) - \sigma^{-1}b(t, X, \mu, \alpha_t)) dW_t^{\mu, \alpha} \right)_T,$$

and since $\sigma^{-1}b$ is bounded, we compute

$$\begin{aligned} \mathcal{H}(P^{\mu,\alpha}|P^{\mu^n,\alpha^n}) &= -\mathbb{E}^{\mu,\alpha} \left[\log \frac{dP^{\mu^n,\alpha^n}}{dP^{\mu,\alpha}} \right] \\ &= \frac{1}{2} \mathbb{E}^{\mu,\alpha} \left[\int_0^T |\sigma^{-1}b(t, X, \mu^n, \alpha_t^n) - \sigma^{-1}b(t, X, \mu, \alpha_t)|^2 dt \right]. \end{aligned}$$

Since $P^{\mu,\alpha} \sim P$ and $\alpha^n \rightarrow \alpha$ in $\mathcal{L} \times P$ -measure, it follows from Lemma 7.10 that $Z^{\mu^n,\nu^n} \rightarrow Z^{\mu,\nu}$ in $\mathcal{L} \times P^{\mu,\alpha}$ -measure, where \mathcal{L} denotes Lebesgue measure on $[0, T]$. By assumption (E), the map $\sigma^{-1}b(t, x, \cdot, \cdot)$ is continuous for each (t, x) . Conclude from the bounded convergence theorem that $P^{\mu^n,\alpha^n} \rightarrow P^{\mu,\alpha}$ in total variation. It follows immediately that $P^{\mu^n,\alpha^n} \circ X^{-1} \rightarrow P^{\mu,\alpha} \circ X^{-1}$ in total variation, and that

$$\mathcal{V}_1 \left(P^{\mu^n,\alpha^n} \circ (\alpha_t^n)^{-1}, P^{\mu,\alpha} \circ (\alpha_t^n)^{-1} \right) \leq \mathcal{V}_1 \left(P^{\mu^n,\alpha^n}, P^{\mu,\alpha} \right) \rightarrow 0.$$

Moreover, $P^{\mu,\alpha} \circ (\alpha_t^n)^{-1} \rightarrow P^{\mu,\alpha} \circ \alpha_t^{-1}$ in \mathcal{L} -measure, since $\alpha^n \rightarrow \alpha$ in $\mathcal{L} \times P$ -measure. Thus $P^{\mu^n,\alpha^n} \circ (\alpha_t^n)^{-1} \rightarrow P^{\mu,\alpha} \circ \alpha_t^{-1}$ in \mathcal{L} -measure, which finally implies

$$\delta_{P^{\mu^n,\alpha^n} \circ (\alpha_t^n)^{-1}}(dq)dt \rightarrow \delta_{P^{\mu,\alpha} \circ \alpha_t^{-1}}(dq)dt, \text{ in } \mathcal{M}.$$

With continuity of Φ established, Φ and $\mathbb{A}(\cdot)$ verify the assumptions of Proposition 7.4, and thus there exists a fixed point $(\mu, \nu) \in \Phi(\mu, \mathbb{A}(\mu, \nu)) = \{\Phi(\mu, \alpha) : \alpha \in \mathbb{A}(\mu, \nu)\}$. It remains to notice that the function Φ takes values in $\mathcal{Q} \times \mathcal{M}^0$, where

$$\mathcal{M}^0 := \{ \nu \in \mathcal{M} : \nu(dt, dq) = \delta_{\hat{q}(t)}(dq)dt \text{ for some measurable map } \hat{q} : [0, T] \rightarrow \mathcal{P}(A) \}.$$

For an element in \mathcal{M}^0 , the corresponding map \hat{q} is uniquely determined, up to almost everywhere equality. Hence, for our fixed point (μ, ν) , we know that there exist $\alpha \in \mathbb{A}(\mu, \nu)$ and a measurable function $\hat{q} : [0, T] \rightarrow \mathcal{P}(A)$ such that $\nu_t = \delta_{\hat{q}(t)}$ and $\hat{q}(t) = P^{\mu,\alpha} \circ \alpha_t^{-1}$ for almost every t . \square

Remark 7.12. Assume for the moment that there is no mean field interaction in the control. Following the notation of Remark 7.3, we may ask if the SDE

$$dX_t = b(t, X, \mu, \hat{\alpha}(t, X, \mu, \zeta_\mu(t, X)))dt + \sigma(t, X)dW_t,$$

admits a strong solution, with μ equal to the law of X . This would allow us to solve the mean field game in a strong sense, on a given probability space, as is required in [11] and [6]. Since $\zeta_\mu(t, X) = Z_t^\mu$, this forward SDE is coupled with the backward SDE:

$$\begin{cases} dX_t &= b(t, X, \mu, \hat{\alpha}(t, X, \mu, Z_t))dt + \sigma(t, X)dW_t, \\ dY_t &= -H(t, X, \mu, Z_t)dt + Z_t dW_t, \\ \mu_0 &= \lambda_0, \quad X \sim \mu, \quad Y_T = g(X, \mu). \end{cases}$$

To solve the mean field game in a strong sense, one must therefore resolve this ‘‘mean field FBSDE’’, studied in some generality in [9]. The solution must consist of (X, Y, Z, μ) , such that (X, Y, Z) are processes adapted to the filtration generated by $(W_t, X_0)_{t \in [0, T]}$ and satisfying the above SDEs, and such that the law of X is μ . Our formulation is a relaxation of the more common formulation (e.g. [11] and [6]) in that the forward SDEs no longer need to be solved in a strong sense. Note, however, that the FBSDE written here is of a different nature from those of [11, 6], which were obtained from the maximum principle. Our FBSDE is more like a stochastic form of the PDE systems of Lasry and Lions; indeed, in the Markovian case, the Feynman-Kac formula for the backward part is nothing but the HJB equation.

7.3. Proof of Theorem 3.8 (uniqueness).

Proof of Theorem 3.8. Recall that $\mathbb{A}(\mu, \nu)$ is always nonempty, as in Remark 7.2. By condition (U.1), we know $\mathbb{A}(\mu, \nu)$ is a singleton for each $(\mu, \nu) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{M}$. Its unique element $\alpha^{\mu, \nu}$ is defined given by

$$\alpha_t^{\mu, \nu} = \hat{\alpha}(t, X, Z_t^{\mu, \nu}),$$

where the function $\hat{\alpha}$ is defined as in (12); note that assumptions (U.2) and (U.3) imply that $\hat{\alpha} = \hat{\alpha}(t, x, z)$ does not depend on μ or ν . Suppose now that $(\mu^1, \nu^1), (\mu^2, \nu^2) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{M}$ are two solutions of the MFG; that is, they are fixed points of the (single-valued) function $\Phi(\cdot, \mathbb{A}(\cdot))$. Abbreviate $Y^i = Y^{\mu^i, \nu^i}$, $Z^i = Z^{\mu^i, \nu^i}$, $\alpha^i = \alpha^{\mu^i, \nu^i}$, $f_t^i := f(t, X, \mu^i, \nu_t^i, \alpha_t^i)$ and $b_t^i := \sigma^{-1}b(t, X, \alpha_t^i)$. We begin by rewriting the BSDEs (11) in two ways:

$$\begin{aligned} d(Y_t^1 - Y_t^2) &= - [f_t^1 - f_t^2 + Z_t^1 \cdot b_t^1 - Z_t^2 \cdot b_t^2] dt + (Z_t^1 - Z_t^2) dW_t \\ &= - [f_t^1 - f_t^2 + Z_t^2 \cdot (b_t^1 - b_t^2)] dt + (Z_t^1 - Z_t^2) dW_t^{\mu^1, \alpha^1} \\ &= - [f_t^1 - f_t^2 + Z_t^1 \cdot (b_t^1 - b_t^2)] dt + (Z_t^1 - Z_t^2) dW_t^{\mu^2, \alpha^2}, \end{aligned}$$

with $Y_T^1 - Y_T^2 = g(X, \mu^1) - g(X, \mu^2)$. Recall that $P^{\mu, \alpha}$ agrees with P on \mathcal{F}_0 for each $\mu \in \mathcal{P}_\psi(\mathcal{C})$ and $\alpha \in \mathbb{A}$. In particular,

$$\mathbb{E}^{\mu^1, \alpha^1} [Y_0^1 - Y_0^2] = \mathbb{E} [Y_0^1 - Y_0^2] = \mathbb{E}^{\mu^2, \alpha^2} [Y_0^1 - Y_0^2].$$

Thus

$$\mathbb{E} [Y_0^1 - Y_0^2] = \mathbb{E}^{\mu^1, \alpha^1} \left[g(X, \mu^1) - g(X, \mu^2) + \int_0^T (f_t^1 - f_t^2 + Z_t^2 \cdot (b_t^1 - b_t^2)) dt \right] \quad (19)$$

$$= \mathbb{E}^{\mu^2, \alpha^2} \left[g(X, \mu^1) - g(X, \mu^2) + \int_0^T (f_t^1 - f_t^2 + Z_t^1 \cdot (b_t^1 - b_t^2)) dt \right]. \quad (20)$$

Since the optimal control maximizes the Hamiltonian,

$$\begin{aligned} f_t^1 + Z_t^2 \cdot b_t^1 &= h(t, X, \mu^1, \nu_t^1, Z_t^2, \alpha_t^1) \leq H(t, X, \mu^1, \nu_t^1, Z_t^2) \\ &= f_1(t, X, \mu^1) + f_2(t, \mu^1, \nu_t^1) + f_3(t, X, \alpha_t^2) + Z_t^2 \cdot b_t^2, \end{aligned}$$

and thus, since $f_t^2 = f_1(t, X, \mu^2) + f_2(t, \mu^2, \nu_t^2) + f_3(t, X, \alpha_t^2)$,

$$f_t^1 - f_t^2 + Z_t^2 \cdot (b_t^1 - b_t^2) \leq f_1(t, X, \mu^1) - f_1(t, X, \mu^2) + f_2(t, \mu^1, \nu_t^1) - f_2(t, \mu^2, \nu_t^2). \quad (21)$$

By switching the place of the indices, the same argument yields

$$f_t^1 - f_t^2 + Z_t^1 \cdot (b_t^1 - b_t^2) \geq f_1(t, X, \mu^1) - f_1(t, X, \mu^2) + f_2(t, \mu^1, \nu_t^1) - f_2(t, \mu^2, \nu_t^2). \quad (22)$$

Since $f_2(t, \mu^i, \nu_t^i)$ are deterministic, applying inequality (21) to (19) and (22) to (20) yields

$$0 \leq \left[\mathbb{E}^{\mu^1, \alpha^1} - \mathbb{E}^{\mu^2, \alpha^2} \right] \left[g(X, \mu^1) - g(X, \mu^2) + \int_0^T (f_1(t, X, \mu^1) - f_1(t, X, \mu^2)) dt \right].$$

Hypothesis (U.4) implies that the right side is at most zero, so in fact

$$0 = \left[\mathbb{E}^{\mu^1, \alpha^1} - \mathbb{E}^{\mu^2, \alpha^2} \right] \left[g(X, \mu^1) - g(X, \mu^2) + \int_0^T (f_1(t, X, \mu^1) - f_1(t, X, \mu^2)) dt \right]. \quad (23)$$

Suppose $\alpha^1 \neq \alpha^2$ holds on a (t, ω) -set of strictly positive $\mathcal{L} \times P$ -measure, where \mathcal{L} is again Lebesgue measure. Then assumption (U.1) implies that the inequalities (21) and (22) are strict on a set of positive $\mathcal{L} \times P$ -measure. Since $P \sim P^{\mu^1, \alpha^1} \sim P^{\mu^2, \alpha^2}$, this implies

$$0 < \left[\mathbb{E}^{\mu^1, \alpha^1} - \mathbb{E}^{\mu^2, \alpha^2} \right] \left[g(X, \mu^1) - g(X, \mu^2) + \int_0^T (f_1(t, X, \mu^1) - f_1(t, X, \mu^2)) dt \right],$$

which contradicts (23). Thus $\alpha^1 \neq \alpha^2$ must hold $\mathcal{L} \times P$ -a.e., which yields

$$\begin{aligned} \frac{dP^{\mu^1, \alpha^1}}{dP} &= \mathcal{E} \left(\int_0^\cdot \sigma^{-1} b(t, X, \alpha_t^1) dW_t \right)_T \\ &= \mathcal{E} \left(\int_0^\cdot \sigma^{-1} b(t, X, \alpha_t^2) dW_t \right)_T = \frac{dP^{\mu^2, \alpha^2}}{dP}, \quad a.s. \end{aligned}$$

Thus $\mu^1 = P^{\mu^1, \alpha^1} \circ X^{-1} = P^{\mu^2, \alpha^2} \circ X^{-1} = \mu^2$, and $\nu_t^1 = \delta_{P^{\mu^1, \alpha^1} \circ (\alpha_t^1)^{-1}} = \delta_{P^{\mu^2, \alpha^2} \circ (\alpha_t^2)^{-1}} = \nu_t^2$ a.e. \square

8. PROOF OF FINITE-PLAYER APPROXIMATION THEOREMS

This section justifies the mean field approximation by proving Theorem 4.2, the general approximation result, as well as Proposition 6.1, the rate of convergence for the price impact model.

8.1. Proof of Theorem 4.2. We work on the probability space of Section 4. Recall that under P , X^1, X^2, \dots are i.i.d. with common law $\hat{\mu}$ and $\alpha_t^1, \alpha_t^2, \dots$ are i.i.d. with common law \hat{q}_t , for almost every t . By symmetry, we may prove the result for player 1 only. For $\beta \in \mathbb{A}_n$, define $\beta^\alpha := (\beta, \alpha^2, \dots, \alpha^n) \in \mathbb{A}_n^n$. We abuse notation somewhat by writing α in place of $(\alpha^1, \dots, \alpha^n) \in \mathbb{A}_n^n$. Note that $(\alpha^1)^\alpha = \alpha$ and $P_n(\alpha) = P$, in our notation. For $\beta \in \mathbb{A}_n$, let

$$J'_n(\beta) := \mathbb{E}^{P_n(\beta^\alpha)} \left[\int_0^T f(t, X^1, \hat{\mu}, \hat{q}_t, \beta_t) dt + g(X^1, \hat{\mu}) \right].$$

Note that $J'_n(\alpha^1)$ does not depend on n . We divide the proof into three lemmas.

Lemma 8.1. *Let $F : \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \rightarrow \mathbb{R}$ be empirically measurable, and suppose $F(x, \cdot)$ is $\tau_\psi(\mathcal{C})$ continuous at $\hat{\mu}$ for each $x \in \mathcal{C}$. Assume also that there exists $c > 0$ such that*

$$|F(x, \mu)| \leq c \left(\psi(x) + \int \psi d\mu \right), \quad \text{for all } (x, \mu) \in \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}).$$

Then $\lim_{n \rightarrow \infty} \mathbb{E}[|F(X^i, \mu^n) - F(X^i, \hat{\mu})|^p] = 0$ for each i and $p \in [1, 2)$.

Proof. By symmetry, it suffices to prove this for $i = 1$. By replacing $F(x, \mu)$ with $|F(x, \mu) - F(x, \hat{\mu})|$, assume without loss of generality that $F \geq 0$ and $F(x, \hat{\mu}) = 0$ for all x . Define

$$\nu^n := \frac{1}{n-1} \sum_{i=2}^n \delta_{X^i}.$$

By independence of X^1 and ν^n , we have

$$\mathbb{E}[F(X^1, \mu^n)] = \mathbb{E} \left[\mathbb{E} \left[F \left(x, \frac{1}{n} \delta_x + \frac{n-1}{n} \nu^n \right) \right]_{x=X^1} \right].$$

Now let $\epsilon > 0$. By continuity of $F(x, \cdot)$, there exist $\delta > 0$ and $\phi_1, \dots, \phi_k \in B_\psi(\mathcal{C})$ such that $F(x, \nu) < \epsilon$ whenever $|\int \phi_i d(\nu - \hat{\mu})| < \delta$ for all $i = 1, \dots, k$. By the law of large numbers,

$$\lim_{n \rightarrow \infty} \left| \int \phi_i d \left(\frac{1}{n} \delta_x + \frac{n-1}{n} \nu^n - \hat{\mu} \right) \right| = 0, \quad a.s.$$

Thus $\limsup_{n \rightarrow \infty} F \left(x, \frac{1}{n} \delta_x + \frac{n-1}{n} \nu^n \right) \leq \epsilon$ a.s., for each $\epsilon > 0$, and so $F \left(x, \frac{1}{n} \delta_x + \frac{n-1}{n} \nu^n \right) \rightarrow 0$ a.s. for each x . The growth assumption along with (S.2) yield

$$\mathbb{E}[F^2(X^1, \mu^n)] \leq 2c^2 \mathbb{E} \left[\psi^2(X^1) + \left(\int \psi d\mu^n \right)^2 \right] \leq 4c^2 \mathbb{E}[\psi^2(X^1)] < \infty,$$

and we conclude by the dominated convergence theorem. \square

Lemma 8.2. *We have $\lim_{n \rightarrow \infty} \sup_{\beta \in \mathbb{A}_n} |J_{n,1}(\beta^\alpha) - J'_n(\beta)| = 0$.*

Proof. Note that, for any $\beta \in \mathbb{A}_n$,

$$\begin{aligned} |J_{n,1}(\beta^\alpha) - J'_n(\beta)| &\leq \int_0^T \mathbb{E}^{P_n(\beta^\alpha)} [F_t(X^1, \mu^n) + G_t(X^1, q^n(\beta_t^\alpha))] dt \\ &\quad + \mathbb{E}^{P_n(\beta^\alpha)} [|g(X^1, \mu^n) - g(X^1, \hat{\mu})|], \end{aligned} \quad (24)$$

where $F : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \rightarrow \mathbb{R}$ and $G : [0, T] \times \mathcal{C} \times \mathcal{P}(A) \rightarrow \mathbb{R}$ are defined by

$$\begin{aligned} F_t(x, \mu) &:= \sup_{(a,q) \in A \times \mathcal{P}(A)} |f(t, x, \mu, q, a) - f(t, x, \hat{\mu}, q, a)|, \\ G_t(x, q) &:= \sup_{a \in A} |f(t, x, \hat{\mu}, q, a) - f(t, x, \hat{\mu}, \hat{q}_t, a)|. \end{aligned}$$

Theorem 18.19 of [2] ensures that both functions are (empirically) measurable. Since A and $\mathcal{P}(A)$ are compact, Lemma 5.6 assures us that $F_t(x, \cdot)$ is $\tau_\psi(\mathcal{C})$ -continuous at $\hat{\mu}$ and that $G_t(x, \cdot)$ is weakly continuous, for each (t, x) . Similar to the proof of Lemma 7.7, $\{dP_n(\beta^\alpha)/dP : \beta \in \mathbb{A}_n, n \geq 1\}$ are bounded in $L^p(P)$, for any $p \geq 1$. Since assumption (F.5) is uniform in t for f , we deduce from Lemma 8.1 and the dominated convergence theorem that

$$\lim_{n \rightarrow \infty} \sup_{\beta \in \mathbb{A}_n} \left[\int_0^T \mathbb{E}^{P_n(\beta^\alpha)} [F_t(X^1, \mu^n)] dt + \mathbb{E}^{P_n(\beta^\alpha)} [|g(X^1, \mu^n) - g(X^1, \hat{\mu})|] \right] = 0.$$

It remains to check that the G_t term converges. Note that $G_t(x, \cdot)$ is uniformly continuous, as $\mathcal{P}(A)$ is compact. Also $\mathcal{V}_1(q^n(\beta_t^\alpha), q^n(\alpha_t)) \leq 2/n$, since these are empirical measures of n points which differ in only one point (recall that \mathcal{V}_1 denotes total variation). Hence

$$\lim_{n \rightarrow \infty} \sup_{\beta \in \mathbb{A}_n} |G_t(X^1, q^n(\alpha_t)) - G_t(X^1, q^n(\beta_t^\alpha))| = 0, \quad a.s.$$

Since $\alpha_t^1, \alpha_t^2, \dots$ are i.i.d. with common law \hat{q}_t , we have $q^n(\alpha_t) \rightarrow \hat{q}_t$ weakly a.s. (see [39]), and thus $G_t(X^1, q^n(\alpha_t)) \rightarrow 0$ a.s. Note that $dP_n(\beta^\alpha)/dP$ are bounded in $L^p(P)$ for any $p \geq 1$ and that the integrands above are bounded in $L^p(P)$ for any $p \in [1, 2)$, by (F.5) and the same argument as in the proof of Lemma 8.2. The dominated convergence theorem completes the proof. \square

Lemma 8.3. *For any $\beta \in \mathbb{A}_n$, $J'_n(\alpha^1) \geq J'_n(\beta)$.*

Proof. We use the comparison principle for BSDEs. Fix n and $\beta \in \mathbb{A}_n$. Define $\phi, \tilde{\phi} : [0, T] \times \Omega \times \mathbb{R}^d \rightarrow \mathbb{R}$ by

$$\begin{aligned} \phi(t, z) &:= \sup_{a \in A} [f(t, X^1, \hat{\mu}, \hat{q}_t, a) + z \cdot (\sigma^{-1}b(t, X^1, a) - \sigma^{-1}b(t, X^1, \alpha_t^1))] \\ \tilde{\phi}(t, z) &:= f(t, X^1, \hat{\mu}, \hat{q}_t, \beta_t) + z \cdot (\sigma^{-1}b(t, X^1, \beta_t) - \sigma^{-1}b(t, X^1, \alpha_t^1)) \end{aligned}$$

By Pardoux and Peng [36], there exist unique solutions (Y, Z^1, \dots, Z^n) and $(\tilde{Y}, \tilde{Z}^1, \dots, \tilde{Z}^n)$ of the BSDEs

$$\begin{cases} dY_t = -\phi_n(t, Z_t^1) dt + \sum_{j=1}^n Z_t^j dW_t^j \\ Y_T = g(X^1, \hat{\mu}), \\ d\tilde{Y}_t = -\tilde{\phi}_n(t, \tilde{Z}_t^1) dt + \sum_{j=1}^n \tilde{Z}_t^j dW_t^j \\ \tilde{Y}_T = g(X^1, \hat{\mu}). \end{cases}$$

The unique solution of the first BSDE is in fact given by $Z^2 \equiv \dots \equiv Z^n \equiv 0$, where (Y, Z^1) are \mathbb{X}^1 -progressively measurable and solve the BSDE

$$\begin{cases} dY_t = -[H(t, X^1, \hat{\mu}, \hat{q}_t, Z_t^1) - Z_t^1 \cdot \sigma^{-1}b(t, X^1, \alpha_t^1)] dt + Z_t^1 dW_t^1 \\ Y_T = g(X^1, \hat{\mu}). \end{cases}$$

This is due to the \mathbb{X}^1 -measurability of the driver and terminal condition of this BSDE. Recall that α^1 is optimal for the mean field problem, and thus it must maximize the Hamiltonian; that is,

$$H(t, X^1, \hat{\mu}, \hat{q}_t, Z_t^1) = h(t, X^1, \hat{\mu}, \hat{q}_t, Z_t^1, \alpha_t^1) = f(t, X^1, \hat{\mu}, \hat{q}_t, \alpha_t^1) + Z_t^1 \cdot \sigma^{-1} b(t, X^1, \alpha_t^1).$$

Thus $dY_t = -f(t, X^1, \hat{\mu}, \hat{q}_t, \alpha_t^1)dt + Z_t^1 dW_t^1$. Since W^1 is a Wiener process under P , taking expectations yields $\mathbb{E}[Y_0] = J'_n(\alpha^1)$, which we note does not depend on n .

Similarly, note that W^j , $j \geq 2$ are Wiener processes under $P_n(\beta^\alpha)$, as is $W^{\beta,1}$. Hence, we rewrite \tilde{Y} as follows:

$$\begin{cases} d\tilde{Y}_t &= -f(t, X^1, \hat{\mu}, \hat{q}_t, \beta_t)dt + \tilde{Z}_t^1 dW_t^{\beta,1} + \sum_{j=2}^n \tilde{Z}_t^j dW_t^j \\ \tilde{Y}_T &= g(X^1, \hat{\mu}). \end{cases}$$

Take expectations, noting that $P = P_n(\beta^\alpha)$ on \mathcal{F}_0^n , to see $\mathbb{E}[\tilde{Y}_0] = \mathbb{E}^{P_n(\beta^\alpha)}[\tilde{Y}_0] = J'_n(\beta)$. Finally, since $\phi \geq \tilde{\phi}$, the comparison principle for BSDEs yields $Y_0 \geq \tilde{Y}_0$, and thus $J'_n(\beta) \leq J'_n(\alpha^1)$. \square

Proof of Theorem 4.2. Simply let $\epsilon_n = 2 \sup_{\beta \in \mathbb{A}_n} |J_{n,1}(\beta^\alpha) - J'_n(\beta)|$. Then $\epsilon_n \rightarrow 0$ by Lemma 8.2, and Lemma 8.3 yields, for all $\beta \in \mathbb{A}_n$,

$$J_{n,1}(\beta^\alpha) \leq \frac{1}{2}\epsilon_n + J'_n(\beta) \leq \frac{1}{2}\epsilon_n + J'_n(\alpha^1) \leq \epsilon_n + J_{n,1}(\alpha).$$

\square

8.2. Proof of Proposition 6.1.

Proof of Proposition 6.1. We simply modify the proof of Theorem 4.2, in light of the special structure of the price impact model. Namely, the inequality (24) becomes

$$\epsilon_n = 2 \sup_{\beta \in \mathbb{A}_n} |J_{n,1}(\beta^\alpha) - J'_n(\beta)| \leq 2 \sup_{\beta \in \mathbb{A}_n} \mathbb{E}^{P_n(\beta)} \int_0^T \left| \gamma X_t^1 \int_A c' d(q_t^n(\beta^\alpha) - \hat{q}_t) \right| dt.$$

Use Hölder's inequality to get

$$\epsilon_n \leq 2\gamma \mathbb{E} [\|X^1\|^4]^{1/4} \sup_{\beta \in \mathbb{A}_n} \mathbb{E} \left[\left(\frac{dP_n(\beta)}{dP} \right)^4 \right]^{1/4} \int_0^T dt \mathbb{E} \left[\left(\int_A c' d(q_t^n(\beta^\alpha) - \hat{q}_t) \right)^2 \right]^{1/2}.$$

Assumption (S.2) with $\psi(x) = e^{c_1 \|x\|}$ implies that $\|X^1\|$ has finite moments of all orders. Again, $\{dP_n(\beta^\alpha)/dP : \beta \in \mathbb{A}_n, n \geq 1\}$ are bounded in $L^p(P)$ for any $p \geq 1$. So it suffices to show

$$\sup_{\beta \in \mathbb{A}_n} \mathbb{E} \left[\left(\int_A c' d(q_t^n(\beta^\alpha) - \hat{q}_t) \right)^2 \right]^{1/2} \leq C/\sqrt{n},$$

for some $C > 0$. This will follow from two inequalities: An easy calculation gives

$$\left| \int_A c' d(q_t^n(\beta^\alpha) - \hat{q}_t) \right| \leq 2C_1/n + \left| \int_A c' d(q_t^n(\alpha) - \hat{q}_t) \right|,$$

where $C_1 = \sup_{a \in A} |c'(a)|$. Since $\alpha_t^1, \alpha_t^2, \dots$ are i.i.d. with common law \hat{q}_t ,

$$\mathbb{E} \left[\left(\int_A c' d(q_t^n(\alpha) - \hat{q}_t) \right)^2 \right] = \text{Var}(c'(\alpha_t^1))/n \leq 4C_1^2/n.$$

\square

9. CONCLUSIONS

This paper provide a theoretical framework for fairly general mean field games with uncontrolled volatility, allowing to prove new existence and uniqueness results for several types of mean field interactions which arise naturally in applications, and which were not studied before. Such results include models with rank effects, nearest-neighbor (e.g. quantile) interactions, and mean field interactions through the controls. The strength of our approach is its generality; the existence, uniqueness, and approximation results apply easily to many concrete models. More refined analysis, for example of regularity of solutions or numerics, could in theory be based on the McKean-Vlasov FBSDE discussed in Remark 7.12, which essentially provides a probabilistic representation of the PDE approach of [32]. This is left for further investigation.

REFERENCES

1. A. Alfonsi, A. Fruth, and A. Schied, *Optimal execution strategies in limit order books with general shape functions*, Quantitative Finance **10** (2010), no. 2, 143–157.
2. C. Aliprantis and K. Border, *Infinite dimensional analysis: A hitchhiker's guide*, 3 ed., Springer, 2007.
3. R. Almgren and N. Chriss, *Optimal execution of portfolio transactions*, Journal of Risk **3** (2000), no. 2, 5–39.
4. D. Andersson and B. Djehiche, *A maximum principle for SDEs of mean-field type*, Applied Mathematics & Optimization **63** (2011), 341–356.
5. M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardinà, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, et al., *Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study*, Proceedings of the National Academy of Sciences **105** (2008), no. 4, 1232–1237.
6. A. Bensoussan, K.C.J. Sung, S.C.P. Yam, and S.P. Yung, *Linear-quadratic mean field games*, Preprint (2011).
7. R. Buckdahn, B. Djehiche, and J. Li, *A general stochastic maximum principle for sdes of mean-field type*, Applied Mathematics & Optimization **64** (2011), no. 2, 197–216.
8. B.I. Carlin, M.S. Lobo, and S. Viswanathan, *Episodic liquidity crises: Cooperative and predatory trading*, The Journal of Finance **62** (2007), no. 5, 2235–2274.
9. R. Carmona and F. Delarue, *Mean field forward-backward stochastic differential equations*, arXiv preprint arXiv:1211.4186 (2012).
10. ———, *Forward-backward stochastic differential equations and controlled McKean Vlasov dynamics*, arXiv preprint arXiv:1303.5835 (2013).
11. ———, *Probabilistic analysis of mean field games*, SIAM Journal of Control and Optimization (2013).
12. R. Carmona, F. Delarue, and A. Lachapelle, *Control of McKean-Vlasov dynamics versus mean field games*, Mathematics and Financial Economics **7** (2013), no. 2, 131–166.
13. A. Cellina, *Approximation of set valued functions and fixed point theorems*, Annali di matematica pura ed applicata **82** (1969), no. 1, 17–24.
14. F. Cucker and S. Smale, *Emergent behavior in flocks*, IEEE Transactions on Automatic Control **52** (2007), no. 5, 852–862.
15. M. Davis, *Martingale methods in stochastic control*, Stochastic Control Theory and Stochastic Differential Systems (M. Kohlmann and W. Vogel, eds.), Lecture Notes in Control and Information Sciences, vol. 16, Springer Berlin / Heidelberg, 1979, pp. 85–117.
16. A. Dembo and O. Zeitouni, *Large deviations techniques and applications*, Stochastic Modelling and Applied Probability, Springer, 2009.
17. N. Dunford and J.T. Schwartz, *Linear operators, part I: General theory*, Interscience, 1957.
18. M. Fuhrman, Y. Hu, and G. Tessitore, *On a class of stochastic optimal control problems related to bsdes with quadratic growth*, SIAM journal on control and optimization **45** (2006), no. 4, 1279–1296.
19. J. Gärtner, *On the McKean-Vlasov limit for interacting diffusions*, Mathematische Nachrichten **137** (1988), no. 1, 197–248.
20. J. Gatheral, A. Schied, and A. Slynko, *Transient linear price impact and Fredholm integral equations*, Mathematical Finance **22** (2012), no. 3, 445–474.
21. D.A. Gomes and V.K. Voskanyan, *Extended mean field games-formulation, existence, uniqueness and examples*, arXiv preprint arXiv:1305.2600 (2013).
22. O. Guéant, J.M. Lasry, and P.L. Lions, *Mean field games and applications*, Paris-Princeton Lectures on Mathematical Finance 2010, Lecture Notes in Mathematics, vol. 2003, Springer Berlin / Heidelberg, 2011, pp. 205–266.
23. S. Hamadene and J.P. Lepeltier, *Backward equations, stochastic control and zero-sum stochastic differential games*, Stochastics: An International Journal of Probability and Stochastic Processes **54** (1995), no. 3-4, 221–231.

24. Y. Hu and S. Peng, *A stability theorem of backward stochastic differential equations and its application*, Comptes Rendus de l'Académie des Sciences-Series I-Mathematics **324** (1997), no. 9, 1059–1064.
25. M. Huang, R. Malhamé, and P. Caines, *Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle*, Communications in Information and Systems **6** (2006), no. 3, 221–252.
26. ———, *The Nash certainty equivalence principle and McKean-Vlasov systems: An invariance principle and entry adaptation*, Proceedings of the 46th IEEE Conference on Decision and Control, 2007, pp. 121–126.
27. J. Jacod and J. Mémin, *Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité*, Séminaire de probabilités de Strasbourg **15** (1981), 529–546.
28. J. Jacod and A.N. Shiryaev, *Limit theorems for stochastic processes*, Grundlehren der mathematischen Wissenschaften, Springer-Verlag, 1987.
29. A. Jadbabaie, J. Lin, and A.S. Morse, *Coordination of groups of mobile autonomous agents using nearest neighbor rules*, IEEE Transactions on Automatic Control **48** (2003), no. 6, 988–1001.
30. N. El Karoui, S. Peng, and M.C. Quenez, *Backward stochastic differential equations in finance*, Mathematical Finance **7** (1997), no. 1, 1–71.
31. A. Lachapelle and M.T. Wolfram, *On a mean field game approach modeling congestion and aversion in pedestrian crowds*, Transportation Research Part B: Methodological **45** (2011), no. 10, 1572 – 1589.
32. J.M. Lasry and P.L. Lions, *Mean field games*, Japanese Journal of Mathematics **2** (2007), 229–260.
33. J.M. Lasry, P.L. Lions, and O. Guéant, *Application of mean field games to growth theory*, (2008).
34. T. Meyer-Brandis, B. Øksendal, and X.Y. Zhou, *A mean-field stochastic maximum principle via Malliavin calculus*, Stochastics (2012), 1–24.
35. M. Nourian, P. Caines, and R. Malhamé, *Mean field analysis of controlled cucker-smale type flocking: Linear analysis and perturbation equations*, Proceedings of the 18th IFAC World Congress, Milan, August 2011, 2011, pp. 4471–4476.
36. E. Pardoux and S. Peng, *Adapted solution of a backward stochastic differential equation*, Systems & Control Letters **14** (1990), 55–61.
37. S. Peng, *Filtration consistent nonlinear expectations and evaluations of contingent claims*, Acta Mathematica Applicata Sinica, English Series **20** (2004), no. 2, 191–214.
38. R. Rao, *Relations between weak and uniform convergence of measures with applications*, The Annals of Mathematical Statistics (1962), 659–680.
39. V.S. Varadarajan, *On the convergence of sample probability distributions*, Sankhyā: The Indian Journal of Statistics (1933-1960) **19** (1958), no. 1/2, 23–26.
40. T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, *Novel type of phase transition in a system of self-driven particles*, Physical Review Letters **75** (1995), no. 6, 1226–1229.
41. C. Villani, *Topics in optimal transportation*, Graduate Studies in Mathematics, American Mathematical Society, 2003.

ORFE, BENDHEIM CENTER FOR FINANCE, PRINCETON UNIVERSITY, PRINCETON, NJ 08544, USA.

E-mail address: rcarmona@princeton.edu

ORFE, PRINCETON UNIVERSITY, PRINCETON, NJ 08544, USA.

E-mail address: dlacker@princeton.edu