

MEAN FIELD GAMES WITH COMMON NOISE

BY RENÉ CARMONA^{*,1}, FRANÇOIS DELARUE[†] AND DANIEL LACKER^{*,2}

*Princeton University** and *Université de Nice*[†]

A theory of existence and uniqueness is developed for general stochastic differential mean field games with common noise. The concepts of strong and weak solutions are introduced in analogy with the theory of stochastic differential equations, and existence of weak solutions for mean field games is shown to hold under very general assumptions. Examples and counterexamples are provided to enlighten the underpinnings of the existence theory. Finally, an analog of the famous result of Yamada and Watanabe is derived, and it is used to prove existence and uniqueness of a strong solution under additional assumptions.

CONTENTS

1. Introduction	3741
2. Strong MFG solutions and discretization	3747
2.1. General set-up and standing assumption	3747
2.2. General objective	3748
2.3. Relaxed controls	3751
2.4. Mean field games with a discretized conditioning	3753
2.4.1. Set-up	3753
2.4.2. Discretization procedure	3754
2.4.3. Measures parameterized by discretized trajectories	3755
2.4.4. Control problems	3756
2.4.5. Strong MFG solutions	3757
3. Weak limits of discretized MFG	3761
3.1. Weak MFG solution with weak control	3761
3.2. Canonical space	3763
3.3. Relative compactness and MFG pre-solution	3765
3.4. Existence of a MFG solution under Assumption B	3771
3.4.1. Strategy	3772
3.4.2. Proof of Theorem 3.8	3773
3.4.3. Proof of Lemma 3.12	3774
3.4.4. Proof of Lemma 3.13	3774
3.4.5. Proof of Lemma 3.14	3775
3.5. Unbounded coefficients	3775
3.5.1. Relative compactness	3776

Received July 2014; revised July 2015.

¹Supported in part by NSF Grant DMS-08-06591.

²Supported in part by NSF Grant DMS-07-39195.

MSC2010 subject classifications. Primary 93E20; secondary 60H10, 91A13.

Key words and phrases. Mean field games, stochastic optimal control, McKean–Vlasov equations, weak solutions, relaxed controls.

3.5.2. Limit points 3777
 3.5.3. Optimality 3779
 3.5.4. Proof of Lemma 3.17 3780
 3.5.5. Proof of Lemma 3.18 3781
 4. Strict and strong controls 3781
 4.1. Strict controls 3781
 4.2. Strong controls 3784
 5. Counter-examples 3786
 5.1. Nonexistence of strong solutions 3787
 5.2. Nonexistence of weak solutions 3787
 6. Uniqueness 3788
 6.1. Pathwise uniqueness and uniqueness in law 3789
 6.2. Lasry–Lions monotonicity condition 3790
 Appendix A: Topology of Wasserstein spaces 3794
 Appendix B: A compactness result for Itô processes 3796
 Appendix C: Density of adapted controls 3798
 References 3801

1. Introduction. While mean field games have been around for quite some time in one form or another, especially in economics, the theoretical framework underlying the present work goes back to the pioneering works of Lasry and Lions [39], and Huang, Malhamé and Caines [30]. The basic idea is to describe asymptotic consensus among a large population of optimizing individuals interacting with each other in a mean-field way, and subject to constraints of energetic or economical type. The strategy is to take advantage of the mean-field interaction to reduce the analysis of the consensus to a control problem for one single representative individual evolving in, and interacting with, the environment created by the aggregation of the other individuals. Intuitively, when consensus occurs, symmetries in the system are expected to force the individuals to obey a form of law of large numbers and satisfy a propagation of chaos phenomenon as the size of the population grows. In most of the existing works following [39], and [30], the sources of randomness in the dynamics of the population are assumed to be independent from one individual to another. The purpose of this paper is to analyze the case of correlated randomness in a general setting.

We concentrate on stochastic differential games for which the epitome of the models can be described as follows. Given a finite time horizon $T > 0$, we start with an N -player stochastic differential game, in which the private state process X^i of player i is given by the solution of the stochastic differential equation

$$dX_t^i = b(t, X_t^i, \bar{\mu}_t^N, \alpha_t^i) dt + \sigma(t, X_t^i, \bar{\mu}_t^N) dW_t^i + \sigma_0(t, X_t^i, \bar{\mu}_t^N) dB_t$$

for $t \in [0, T]$, with $\bar{\mu}_t^N = \frac{1}{N} \sum_{j=1}^n \delta_{X_t^j}$.

Here, B is a Wiener process called the *common noise*, and W^1, \dots, W^N are independent Wiener processes, independent of B . The processes W^1, \dots, W^N are

called the *independent* or *idiosyncratic noises*. The objective of player i is to choose a control α^i in order to maximize the quantity

$$J^i(\alpha^1, \dots, \alpha^N) := \mathbb{E} \left[\int_0^T f(t, X_t^i, \bar{\mu}_t^N, \alpha_t^i) dt + g(X_T^i, \bar{\mu}_T^N) \right],$$

the difficulty coming from the fact that these N optimizations are conducted simultaneously. Besides the correlations coming through the common noise B , the optimization problems are coupled through the marginal empirical distributions $(\bar{\mu}_t^N)_{t \in [0, T]}$ of the state processes. Additionally, the individuals share the same coefficients and objective functions, and thus the game is symmetric as long as the initial conditions X_0^1, \dots, X_0^N are exchangeable.

The symmetry is a very important feature of mean field games. However, since the controls are allowed to differ from one player to another, the expected reward functionals J^1, \dots, J^N may not be the same. In particular, except for some very specific cases, there is no hope to find controls $\alpha^1, \dots, \alpha^N$ that maximize simultaneously all the reward functionals J^1, \dots, J^N . Instead of a global maximizer, the idea of consensus is formalized by the concept of *Nash equilibrium*. In short, an N -tuple $(\alpha^{1,*}, \dots, \alpha^{N,*})$ is a Nash equilibrium if the reward J^i of the particle i attains a maximum at $\alpha^{i,*}$ when all the other particles $j \neq i$ use the controls $\alpha^{j,*}$. Because of the symmetric structure of the game, it then makes sense to investigate the asymptotic behavior of exchangeable equilibria. Drawing intuition from the theory of propagation of chaos, one may anticipate that effective equations may hold in the limit as the number of players N tends to infinity, and hope that their solutions may be more manageable than the search for Nash equilibria for large stochastic differential games of the type described above. This is the rationale for the formulation of the *mean field game* (MFG) problem introduced in [39], and [30]. See also [11] for recent developments. We stress that the goal of the present paper is not to justify the passage to the limit, but to study the resulting asymptotic optimization problem.

This informal discussion suggests that the MFG is essentially an asymptotic formulation of the game, in which the influence of each player on the empirical measure is small, hinting at the fact that the asymptotic optimization problems could be decoupled and identical in nature. Put differently, the limiting equilibrium problem consists of a standard optimization problem for one representative player only (instead of N) interacting (competing) with the environment provided by the asymptotic behavior (as N tends to ∞) of the marginal empirical measures $(\bar{\mu}_t^N)_{t \in [0, T]}$ corresponding to an exchangeable equilibrium $(\alpha^{1,*}, \dots, \alpha^{N,*})$. Without common noise, the classical law of large numbers says that the limit environment should be a deterministic flow of probability measures $(\mu_t)_{t \in [0, T]}$ describing the statistical distribution of the population in equilibrium. When σ_0 is nonzero, the impact of the common noise does not average out, and since it does not disappear, the limiting environment must be given by a stochastic flow $(\mu_t)_{t \in [0, T]}$ of probability measures describing the conditional distribution of the population in

equilibrium given the common noise. Therefore, we introduce the following generalization to the MFG problem proposed in [10, 30, 39] in the absence of common noise:

(1) For a fixed adapted process $(\mu_t)_{t \in [0, T]}$ with values in the space $\mathcal{P}(\mathbb{R}^d)$ of probability measures on \mathbb{R}^d , solve the optimal control problem given by

$$(1.1) \quad \sup_{\alpha} \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \alpha_t) dt + g(X_T, \mu_T) \right] \quad \text{such that}$$

$$(1.2) \quad dX_t = b(t, X_t, \mu_t, \alpha_t) dt + \sigma(t, X_t, \mu_t) dW_t + \sigma_0(t, X_t, \mu_t) dB_t.$$

(2) Given an optimal control, find the corresponding conditional laws $(\mu_t^*)_{t \in [0, T]}$ of the optimally controlled state process $(X_t^*)_{t \in [0, T]}$ given B .

(3) Find a fixed point $(\mu_t)_{t \in [0, T]}$, such that the resulting μ_t^* equals μ_t for all $t \in [0, T]$.

The *fixed-point problem* or *consistency condition* (3) characterizes the fact that, under the conditional equilibrium measure, the optimal state (conditional on B) must be typical of the population. This is exactly the usual MFG problem except for the fact that the solution $(\mu_t)_{t \in [0, T]}$ is now a *random* measure flow. Again, the conditioning on B appears because the effect of the independent noises W^i on the empirical measure flow averages out as N tends to infinity, but the effect of the common noise B does not.

The goal of this paper is to discuss the existence and, possibly, the uniqueness of an equilibrium in the presence of a common noise. Except perhaps when the time horizon is small and contractions are available as in [30], it may not be possible to invoke Picard’s theorem, even in the simpler case without common noise. While more diverse techniques continue to develop in the literature in order to accommodate rather singular variants of the problem (regarding the form of the mean-field interaction or the structure of the differential operator underpinning (1.1) and (1.2); see, e.g., the PDE contributions [9, 24, 25]), one basic strategy for proving the existence of an equilibrium without common noise relies on Schauder’s fixed-point theorem. We refer to [8] for the implementation of this strategy under the PDE formulation when the differential operator driving (1.2) has sufficiently strong smoothing properties and to [10] for a probabilistic account when the optimal stochastic control problem (2) can be solved by means of the stochastic Pontryagin principle. From a probabilistic point of view, Schauder’s theorem is typically applied in the space $\mathcal{C}([0, T], \mathcal{P}(\mathbb{R}^d))$ of continuous functions from $[0, T]$ into the space of probability measures on \mathbb{R}^d , as it is straightforward via tightness arguments to identify compact subsets under very natural topologies. In the presence of a common noise, the problem is much more complicated, as the natural space in which one searches for the fixed point is $[\mathcal{C}([0, T], \mathcal{P}(\mathbb{R}^d))]^\Omega$, where Ω denotes the underlying probability space carrying the common noise. Except when Ω is finite, this space is far too large and it is too difficult to find compact

subsets left invariant by the transformations of interest. For that reason, the existence proof is done first on the level of a discretized version of the mean field game, in which the conditioning on the common noise B in the step (2) of the MFG procedure is replaced by a conditioning on a finitely-supported approximation of B . The introduction of such a discretization procedure seems to be original in the context of MFG problems. It must be stressed that it leaves completely unchanged the space of values for the time and state variables, which remain continuous, as it preserves the Brownian nature of the noises B and W . In a word, only the support of the conditioning in the step (2) of the MFG procedure is impacted by the discretization. In this regard, our approach does not consist in an approximation of the full fledge MFG by finite-time finite-space MFG, although this latter idea appears to be powerful on its own. For the reader interested in this subject, we refer to the papers [23] and [22], which are devoted the analysis of MFGs with finite state space, in continuous and discrete time, respectively, but not with the objective of taking a limit between the two. Aside from a similar application of Kakutani's fixed-point theorem, the methods of [22, 23] differ substantially from ours and do not apply directly in our setting.

Most importantly, this discretization procedure crucially bypasses a key technical difficulty: in general, the operation of *conditioning* fails to be continuous in any useful sense, and this puts a wrench in any effort to directly apply fixed-point theorems. However, when the conditioning σ -field is finite, enough continuity is recovered; for example, if X, Y, Y_n are random variables, X is nonatomic, and \mathcal{G} is a finite sub- σ -field of $\sigma(X)$, then $(X, Y_n) \rightarrow (X, Y)$ in distribution implies $\text{Law}(Y_n|\mathcal{G}) \rightarrow \text{Law}(Y|\mathcal{G})$ (weakly) in distribution. Exploiting this remark, the existence proof for the discretized MFG becomes a simple application of Kakutani's fixed-point theorem; Kakutani's fixed-point theorem being preferred to Schauder's because of the possible existence of multiple optimal controls.

The existence for the true mean field game is then obtained by refining the discretization, proving tightness of the sequence of solutions and taking limits. In this way, solutions are constructed as *weak limits* and read as *weak MFG solutions*. The word *weak* refers to the fact that in the limit, the fixed point $(\mu_t)_{t \in [0, T]}$ may not be adapted to the filtration of the common noise B any longer. Such a phenomenon is well known in stochastic calculus: when solving a stochastic differential equation, solutions need not be adapted with respect to the noise driving the equation, in which case they are called *weak*. We use here the same terminology. Because of that lack of adaptedness, we weaken the fixed-point condition by requiring instead that $\mu = \text{Law}(X|B, \mu)$ and that (μ, B) and W are independent. Intuitively, the first of these two new conditions is natural in the sense that the same relationship holds in the n -player game, that is, $\bar{\mu}^N = \text{Law}(X^1|B, \bar{\mu}^N)$ when the agents choose symmetric strategies; the independence with W arises only in the limit $N \rightarrow \infty$. The follow-up paper [37] elaborates on this point and further justifies the notion of weak solution by proving that it characterizes the limits of N -player equilibria.

We refer to a solution of the fixed-point problem (1)–(3) with the more desirable fixed-point condition $\mu = \text{Law}(X|B)$ as a *strong MFG solution*. A strong solution is then a weak solution for which the measure flow μ happens to be measurable with respect to the common noise B . Again, the terminology *strong* is used in analogy with the theory of stochastic differential equations. This brings us to the famous result by Yamada and Watanabe [45] in stochastic analysis: whenever a stochastic differential equation has the pathwise uniqueness property, any weak solution is in fact a strong solution. In this paper, we develop a similar notion of pathwise uniqueness for mean field games and provide an analog of the theorem of Yamada and Watanabe in this context. From this result, we conclude that, whenever pathwise uniqueness holds for a MFG with common noise, the unique weak solution is in fact a strong solution, which then completes our program.

Our analysis relies on one important additional ingredient. In order to guarantee compactness (or at least closedness) of the sets of controls in a sufficiently weak sense, it is also useful for existence proofs to enlarge the family of admissible controls. Precisely, we allow for relaxed (i.e., measure-valued) controls which may be randomized externally to the inputs of the control problems. With this extension, we first treat the case when controls take values in a compact set and the state coefficients b , σ and σ_0 are bounded. Another approximation procedure is then needed to derive the general case. Existence and the limiting arguments are all derived at the level of the joint law of (B, W, μ, α, X) in a suitable function space. In the search for a weak MFG solution, the filtration of the control problem is generated by the two Wiener processes B and W but also by the measure flow μ , which we do not require to be adapted to B or W . Allowing the controls to be randomized externally to the inputs (B, W, μ) requires specifying an admissible family of enlargements of the probability space supporting these inputs. Because the original filtration is not necessarily Brownian, special care is needed in choosing the correct type of allowable extensions. This leads to the important, though rather technical, notion of *compatibility*. The delicate measure theoretic arguments required for the proof are described in detail in Section 3.2.

The main contributions of the paper are as follows. We prove first that there exists a weak MFG solution under general assumptions. Under additional convexity assumptions, we derive existence results without relaxed or externally randomized controls. Under a monotonicity assumption due to Lasry and Lions [39], we prove that pathwise uniqueness holds and, as a consequence, that existence and uniqueness hold in the strong sense. Our results appear to be the first general existence and uniqueness results for mean field games with common noise, introduced in [41] and touted in various forms in [1, 5, 11, 13, 26, 27, 40]. The latter papers [5, 11, 26] discuss the formulation of the problem in terms of the *master equation*, which is a single partial differential equation (PDE) in infinite dimension which summarizes the entire system. Ahuja [1] finds (in our terminology) strong solutions of a class of essentially linear-quadratic mean field games with common noise, but with nonquadratic terminal objective g . The papers [27, 40] of Guéant

et al. solve explicitly some specific common noise models of income distribution. On the other hand, Carmona et al. [13] compute explicit solutions for both the finite-player game and the mean field game in a certain linear-quadratic common noise model, verifying directly the convergence as the number of agents tends to infinity. Although we will not discuss finite-player games here, we draw attention to the follow-up paper [37] that provides rigorous convergence results. Finally, the reader may be also interested in the different, but related, MFG problem with a major player; see [15, 29, 42] for a presentation together with some solvability results. We believe that our strategy based on the discretization of the support of the conditioning could be successfully implemented in this slightly different framework (as it could permit to relax the assumptions used to guarantee the existence of an equilibrium).

The analysis of this paper allows for degenerate volatilities, and thus includes mean field games without common noise (where $\sigma_0 \equiv 0$) and deterministic mean field games (where $\sigma_0 \equiv \sigma \equiv 0$). However, since it is especially intended to accommodate the common noise, our construction is not so well fitted to these two cases. The reason is that the solutions we obtain by applying our results still involve *random* measure flows and are thus somehow weaker than the MFG solutions typically considered in the literature. That said, the follow-up paper [37] explains why this notion of randomized solution is relevant from the point of view of finite-player games, even when $\sigma_0 \equiv 0$. The reader interested in existence results when $\sigma_0 \equiv 0$ is encouraged to look at [14] and [38], both works being related to ours in that they employ weak formulations of the optimal control problems. The paper [38] especially mirrors ours in several ways, in particular in its use of relaxed controls in conjunction with Kakutani's theorem as well as measurable selection arguments for constructing strict (nonrelaxed) controls (however the presence of common noise necessitates here a much more careful formulation and analysis of the problem). See [38] also for an explanation of the strengths and weaknesses of this relaxed control approach to existence theory for mean field games without common noise, compared to other analytic and probabilistic techniques. To name a few of the most influential of these references, we refer to [8, 39] for a PDE-based analysis and [4, 10] for a more probabilistic point of view.

The paper is organized as follows. First, Section 2 discusses the main Assumptions A, definitions of strong MFG solutions and existence of discretized MFG solutions. Section 3 defines weak MFG solutions in detail, discusses some of their properties and proves existence by refining the discretizations of the previous section and taking limits. Section 4 discusses how to strengthen the notion of control, providing general existence results without relaxed controls under additional convexity hypotheses. The brief Section 5 discusses two counter-examples, which explain why we must work with weak solutions and why we cannot relax the growth assumptions placed on the coefficients. Uniqueness is studied in Section 6, discussing our analog of the Yamada–Watanabe theorem and its application to an existence and uniqueness result for strong MFG solutions.

2. Strong MFG solutions and discretization.

2.1. *General set-up and standing assumption.* Fix a time horizon $T > 0$. For a measurable space (Ω, \mathcal{F}) , let $\mathcal{P}(\Omega, \mathcal{F})$ denote the set of probability measures on (Ω, \mathcal{F}) . When the σ -field is understood, we write simply $\mathcal{P}(\Omega)$. When Ω is a metric space, let $\mathcal{B}(\Omega)$ denote its Borel σ -field, and endow $\mathcal{P}(\Omega)$ with the topology of weak convergence. Let $\mathcal{C}^k = C([0, T]; \mathbb{R}^k)$ denote the set of continuous functions from $[0, T]$ to \mathbb{R}^k . Define the evaluation mappings π_t on \mathcal{C}^k by $\pi_t(x) = x_t$ and the truncated supremum norms $\|\cdot\|_t$ on \mathcal{C}^k by

$$\|x\|_t := \sup_{s \in [0, t]} |x_s|, \quad t \in [0, T].$$

Unless otherwise stated, \mathcal{C}^k is endowed with the norm $\|\cdot\|_T$. Let \mathcal{W}^k denote Wiener measure on \mathcal{C}^k . For $\mu \in \mathcal{P}(\mathcal{C}^k)$, let $\mu_t \in \mathcal{P}(\mathbb{R}^k)$ denote the image of μ under π_t . For $p \geq 0$ and a separable metric space (E, ℓ) , let $\mathcal{P}^p(E)$ denote the set of $\mu \in \mathcal{P}(E)$ with $\int_E \ell^p(x, x^0) \mu(dx) < \infty$ for some (and thus for any) $x^0 \in E$. For $p \geq 1$ and $\mu, \nu \in \mathcal{P}^p(E)$, let $\ell_{E,p}$ denote the p -Wasserstein distance, given by

$$(2.1) \quad \ell_{E,p}(\mu, \nu) := \inf \left\{ \left(\int_{E \times E} \gamma(dx, dy) \ell^p(x, y) \right)^{1/p} : \gamma \in \mathcal{P}(E \times E) \text{ has marginals } \mu, \nu \right\}.$$

Unless otherwise stated, the space $\mathcal{P}^p(E)$ is equipped with the metric $\ell_{E,p}$, and $\mathcal{P}(E)$ has the topology of weak convergence. Both are equipped with the corresponding Borel σ -fields, which coincide with the σ -field generated by the mappings $\mathcal{P}^p(E)$ (resp. $\mathcal{P}(E)$) $\ni \mu \mapsto \mu(F)$, F being any Borel subset of E . Appendix A discusses the topological properties of Wasserstein distances relevant to this paper.

We are given two exponents $p', p \geq 1$, a control space A , and the following functions:

$$\begin{aligned} (b, f) &: [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A \rightarrow \mathbb{R}^d \times \mathbb{R}, \\ (\sigma, \sigma_0) &: [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \rightarrow \mathbb{R}^{d \times m} \times \mathbb{R}^{d \times m_0}, \\ g &: \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \rightarrow \mathbb{R}. \end{aligned}$$

The standing assumptions for our existence and convergence theorems are as follows. Continuity and measurability statements involving $\mathcal{P}^p(\mathbb{R}^d)$ are with respect to the Wasserstein distance $\ell_{\mathbb{R}^d,p}$ and its corresponding Borel σ -field.

ASSUMPTION A. The main results of the paper will be proved under the following assumptions, which we assume to hold throughout the paper:

(A.1) A is a closed subset of a Euclidean space. (More generally, as in [28], a closed σ -compact subset of a Banach space would suffice.)

(A.2) $p' > p \geq 1 \vee p_\sigma$, $p_\sigma \in [0, 2]$, and $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$. [Here, $a \vee b := \max(a, b)$.]

(A.3) The functions b, σ, σ_0, f , and g of (t, x, μ, a) are jointly measurable and are continuous in (x, μ, a) for each t .

(A.4) There exists $c_1 > 0$ such that, for all $(t, x, y, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

$$|b(t, x, \mu, a) - b(t, y, \mu, a)| + |(\sigma, \sigma_0)(t, x, \mu) - (\sigma, \sigma_0)(t, y, \mu)| \leq c_1|x - y|$$

and

$$|b(t, 0, \mu, a)| \leq c_1 \left[1 + \left(\int_{\mathbb{R}^d} |z|^p \mu(dz) \right)^{1/p} + |a| \right],$$

$$|\sigma(t, x, \mu)|^2 + |\sigma_0(t, x, \mu)|^2 \leq c_1 \left[1 + |x|^{p_\sigma} + \left(\int_{\mathbb{R}^d} |z|^p \mu(dz) \right)^{p_\sigma/p} \right].$$

(A.5) There exist $c_2, c_3 > 0$ such that, for each $(t, x, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

$$\begin{aligned} & -c_2 \left(1 + |x|^p + \int_{\mathbb{R}^d} |z|^p \mu(dz) \right) \\ & \leq g(x, \mu) \leq c_2 \left(1 + |x|^p + \int_{\mathbb{R}^d} |z|^p \mu(dz) \right), \\ & -c_2 \left(1 + |x|^p + \int_{\mathbb{R}^d} |z|^p \mu(dz) + |a|^{p'} \right) \\ & \leq f(t, x, \mu, a) \leq c_2 \left(1 + |x|^p + \int_{\mathbb{R}^d} |z|^p \mu(dz) \right) - c_3|a|^{p'}. \end{aligned}$$

Examples under which Assumption A holds will be discussed in Section 5.

2.2. *General objective.* Ideally, we are interested in the following notion of strong MFG solution.

DEFINITION 2.1 (Strong MFG solution with strong control). A *strong MFG solution with strong control* and with initial condition λ is a tuple $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \alpha, X)$, where $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P)$ is a filtered probability space supporting (B, W, μ, α, X) satisfying:

(1) $(\mathcal{F}_t)_{t \in [0, T]}$ is the P -complete filtration generated by the process $(X_0, B_t, W_t)_{t \in [0, T]}$.

(2) The processes $(B_t)_{t \in [0, T]}$ and $(W_t)_{t \in [0, T]}$ are independent $(\mathcal{F}_t)_{t \in [0, T]}$ Wiener processes of respective dimension m_0 and m , the processes $(\mu_t = \mu \circ \pi_t^{-1})_{t \in [0, T]}$ and $(X_t)_{t \in [0, T]}$ are $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted processes [with values in $\mathcal{P}^p(\mathbb{R}^d)$ and \mathbb{R}^d respectively], and $P \circ X_0^{-1} = \lambda$.

(3) $(\alpha_t)_{t \in [0, T]}$ is $(\mathcal{F}_t)_{t \in [0, T]}$ -progressively measurable with values in A and $\mathbb{E} \int_0^T |\alpha_t|^p dt < \infty$.

(4) The state equation holds

$$(2.2) \quad \begin{aligned} dX_t &= b(t, X_t, \mu_t, \alpha_t) dt + \sigma(t, X_t, \mu_t) dW_t \\ &\quad + \sigma_0(t, X_t, \mu_t) dB_t, \quad t \in [0, T]. \end{aligned}$$

(5) If $(\alpha'_t)_{t \in [0, T]}$ is another $(\mathcal{F}_t)_{t \in [0, T]}$ -progressively measurable A -valued process satisfying $\mathbb{E} \int_0^T |\alpha'_t|^p dt < \infty$, and X' is the unique strong solution of

$$dX'_t = b(t, X'_t, \mu_t, \alpha'_t) dt + \sigma(t, X'_t, \mu_t) dW_t + \sigma_0(t, X'_t, \mu_t) dB_t, \quad X'_0 = X_0,$$

then

$$\begin{aligned} &\mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \alpha_t) dt + g(X_T, \mu_T) \right] \\ &\quad \geq \mathbb{E} \left[\int_0^T f(t, X'_t, \mu_t, \alpha'_t) dt + g(X'_T, \mu_T) \right]. \end{aligned}$$

(6) P -almost surely, $\mu(\cdot) = P(X \in \cdot | B)$. That is, μ is a version of the conditional law of X given B .

Pay attention that $E \int_0^T |\alpha_t|^{p'} dt$ is not required to be finite. Thanks to (A.5), there is no need. When $E \int_0^T |\alpha_t|^{p'} dt = \infty$, the reward functional is well defined and is equal to $-\infty$.

Definition 2.1 may be understood as follows. Points (1), (2) and (3) are somewhat technical requirements that fix the probabilistic set-up under which the MFG solution is defined. Given μ as in the definition, (4) and (5) postulate that $(X_t)_{t \in [0, T]}$ is a solution of the stochastic optimal control problem driven by the reward functionals f and g in the random environment μ . Condition (6) is a fixed-point condition. It is an adaptation of the condition $\mu = P(X \in \cdot)$ used in the MFG literature to describe asymptotic Nash equilibria between interacting particles X^1, \dots, X^N submitted to independent noises:

$$dX_t^i = b(t, X_t^i, \bar{\mu}_t^N, \alpha_t^i) dt + \sigma(t, X_t^i, \bar{\mu}_t^N) dW_t^i, \quad i = 1, \dots, N,$$

where W^1, \dots, W^N are independent Wiener processes, $\bar{\mu}_t^N$ is the empirical distribution of the N -tuple (X_t^1, \dots, X_t^N) and $\alpha^1, \dots, \alpha^N$ are control processes. In (6), the conditioning by B reflects correlations between the particles when their dynamics are governed by a common noise:

$$\begin{aligned} dX_t^i &= b(t, X_t^i, \bar{\mu}_t^N, \alpha_t^i) dt + \sigma(t, X_t^i, \bar{\mu}_t^N) dW_t^i + \sigma_0(t, X_t^i, \bar{\mu}_t^N) dB_t, \\ &\quad i = 1, \dots, N, \end{aligned}$$

where B, W^1, \dots, W^N are independent Wiener processes. Intuitively, conditioning in (6) follows from a conditional application of the law of large numbers; see [11] for an overview. In our definition, the equilibrium is called *strong* as it is entirely described by the common noise B .

REMARK 2.2. The fact that $(X_t)_{t \in [0, T]}$ is $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted and $(B_t)_{t \in [0, T]}$ is an $(\mathcal{F}_t)_{t \in [0, T]}$ -Wiener process in the above definition implies, with (6), that $\mu_t = P(X_t \in \cdot | \mathcal{B}) = P(X_t \in \cdot | \sigma(B_s : s \leq t))$ P a.s. The filtration being complete, $(\mu_t)_{t \in [0, T]}$ is automatically $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted [without requiring it in (2)]. Note also that $(\mu_t)_{t \in [0, T]}$ has continuous trajectories [in $\mathcal{P}^p(\mathbb{R}^d)$] as μ is $\mathcal{P}^p(\mathcal{C}^d)$ -valued.

We will not be able to prove existence of such a solution under the general Assumptions A. It is not until Section 6 that we find additional assumptions which do ensure the existence and uniqueness of a strong MFG solution (either in the sense of Definition 2.1 or the following weaker Definitions 2.3). Assuming only Assumption A, a general existence theorem will hold if we relax the notion of solution. As the first of two relaxations, the class of admissible controls will be enlarged to include what we call *weak controls*. Weak controls are essentially $\mathcal{P}(A)$ -valued processes rather than A -valued processes, which may be interpreted as a randomization of the control; moreover, weak controls are also allowed to be randomized externally to the given sources of randomness (X_0, B, W) . Weak controls are useful for both convexity and compactness purposes; since we have not required the control space A to be convex, it is well known even in deterministic control theory that measure-valued controls are essential. See [19] for an excellent introduction to this approach to stochastic optimal control.

DEFINITION 2.3 (Strong MFG solution with weak control). A *strong MFG solution with weak control* is a tuple $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$, where $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P)$ is a probability space with a complete filtration supporting (B, W, μ, Λ, X) satisfying:

- (1) The processes $(B_t)_{t \in [0, T]}$ and $(W_t)_{t \in [0, T]}$ are independent $(\mathcal{F}_t)_{t \in [0, T]}$ Wiener processes of respective dimension m_0 and m , the processes $(\mu_t = \mu \circ \pi_t^{-1})_{t \in [0, T]}$ and $(X_t)_{t \in [0, T]}$ are $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted processes [with values in $\mathcal{P}^p(\mathbb{R}^d)$ and \mathbb{R}^d , respectively] and $P \circ X_0^{-1} = \lambda$.
- (2) $(\Lambda_t)_{t \in [0, T]}$ is $(\mathcal{F}_t)_{t \in [0, T]}$ -progressively measurable with values in $\mathcal{P}(A)$ and

$$\mathbb{E} \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty.$$

(3) The state equation holds:³

$$dX_t = \left\{ \int_A b(t, X_t, \mu_t, a) \Lambda_t(da) \right\} dt + \sigma(t, X_t, \mu_t) dW_t + \sigma_0(t, X_t, \mu_t) dB_t.$$

(4) If $(\Omega', \mathcal{F}'_t, P')$ is another filtered probability space supporting processes $(B', W', \mu', \Lambda', X')$ satisfying (1)–(3) and $P \circ (B, \mu)^{-1} = P' \circ (B', \mu')^{-1}$, then

$$\begin{aligned} &\mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda_t(da) dt + g(X_T, \mu_T) \right] \\ &\quad \geq \mathbb{E} \left[\int_0^T \int_A f(t, X'_t, \mu'_t, a) \Lambda'_t(da) dt + g(X'_T, \mu'_T) \right]. \end{aligned}$$

(5) μ is a version of the conditional law of X given B .

Note that (1) and (5) are the same as (2) and (6) of Definition 2.1. Given a MFG solution in either of the above senses, we call the corresponding measure flow $(\mu_t)_{t \in [0, T]}$ an *equilibrium*.

2.3. *Relaxed controls.* We now specify the notion of relaxed controls. Recall that Assumption A is in force at all times. Define \mathcal{V} to be the set of measures q on $[0, T] \times A$ satisfying both $q(\cdot \times A) = \text{Lebesgue}$ (that is the image of q by the projection on $[0, T]$ is the Lebesgue measure on $[0, T]$) and

$$\int_{[0, T] \times A} q(dt, da) |a|^p < \infty.$$

An element of \mathcal{V} is called a *relaxed control*. Any element $q \in \mathcal{V}$ may be rescaled into $q/T \in \mathcal{P}^p([0, T] \times A)$. This permits to endow \mathcal{V} with the p -Wasserstein metric, denoted by $\ell_{\mathcal{V}}$. It follows from results of [31] that \mathcal{V} is a Polish space (since A is), and in fact if A is compact then so is \mathcal{V} , and in this case $\ell_{\mathcal{V}}$ metrizes the topology of weak convergence. See Appendix A for some details about this space.

Each $q \in \mathcal{V}$ may be identified with a measurable function $[0, T] \ni t \mapsto q_t \in \mathcal{P}^p(A)$, determined uniquely (up to a.e. equality) by $dtq_t(da) = q(dt, da)$. As in [38], Lemma 3.8, we can find a predictable version of $(q_t)_{t \in [0, T]}$ in the following sense. Let Λ denote the identity map on \mathcal{V} , and let

$$(2.3) \quad \mathcal{F}_t^\Lambda := \sigma(\Lambda(C) : C \in \mathcal{B}([0, t] \times A)), \quad t \in [0, T].$$

Then there exists an $(\mathcal{F}_t^\Lambda)_{t \in [0, T]}$ -predictable process $\bar{\Lambda} : [0, T] \times \mathcal{V} \rightarrow \mathcal{P}(A)$ such that, for each $q \in \mathcal{V}$, $\bar{\Lambda}(t, q) = q_t$ for almost every $t \in [0, T]$. In particular, $q = dt[\bar{\Lambda}(t, q)](da)$ for each $q \in \mathcal{V}$, and it is immediate that $\mathcal{F}_t^\Lambda = \sigma(\bar{\Lambda}(s, \cdot) : s \leq t)$.

³Throughout the paper, we avoid augmenting filtrations to be right-continuous, mostly because it could cause real problem in point (3) of Definition 3.1. The concerned reader is referred to [43], Lemma 4.3.3, for a carefully discussion of stochastic integration without completeness or right-continuity of the filtration.

We will abuse notation somewhat by writing $\Lambda_t := \overline{\Lambda}(t, \cdot)$. Before we proceed, we first state a frequently useful version of a standard moment estimate for the state equation (4) in Definition 2.3.

LEMMA 2.4. *On some filtered probability space $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P)$, suppose B and W are independent $(\mathcal{F}_t)_{t \in [0, T]}$ -Wiener processes, suppose μ is a $\mathcal{P}^p(\mathbb{C}^d)$ -valued random variable such that $(\mu_t = \mu \circ \pi_t^{-1})$ is $(\mathcal{F}_t)_{t \in [0, T]}$ -progressive, suppose $(\Lambda_t)_{t \in [0, T]}$ is an $(\mathcal{F}_t)_{t \in [0, T]}$ -progressive $\mathcal{P}^p(A)$ -valued process, and suppose ξ is a \mathcal{F}_0 -measurable random vector with law λ . Assume Assumption A holds. Then there exists a unique solution X of the state equation (4) in Definition 2.3 with $X_0 = \xi$ as initial condition.*

For each $\gamma \in [p, p']$, there exists a constant $c_4 > 0$, depending only on γ, λ, T , and the constant c_1 of (A.4) such that

$$\mathbb{E} \|X\|_T^\gamma \leq c_4 \left(1 + \int_{\mathbb{C}^d} \|z\|_T^\gamma \mu(dz) + \mathbb{E} \int_0^T \int_A |a|^\gamma \Lambda_t(da) dt \right).$$

Moreover, if $P(X \in \cdot | B) = \mu$, then we have

$$\mathbb{E} \int_{\mathbb{C}^d} \|z\|_T^\gamma \mu(dz) = \mathbb{E} \|X\|_T^\gamma \leq c_4 \left(1 + \mathbb{E} \int_0^T \int_A |a|^\gamma \Lambda_t(da) dt \right).$$

PROOF. Existence and uniqueness are standard. The Burkholder–Davis–Gundy inequality and Jensen’s inequality yield a constant C (depending only on γ, λ, c_1 and T , and which may then change from line to line) such that, if $\Sigma := \sigma \sigma^\top + \sigma_0 \sigma_0^\top$, then

$$\begin{aligned} \mathbb{E} \|X\|_t^\gamma &\leq C \mathbb{E} \left[|X_0|^\gamma + \int_0^t ds \int_A \Lambda_s(da) |b(s, X_s, \mu_s, a)|^\gamma \right. \\ &\quad \left. + \left(\int_0^t ds |\Sigma(s, X_s, \mu_s)| \right)^{\gamma/2} \right] \\ &\leq C \mathbb{E} \left\{ |X_0|^\gamma + c_1^\gamma \int_0^t ds \left[1 + \|X\|_s^\gamma + \left(\int_{\mathbb{C}^d} \|z\|_s^p \mu(dz) \right)^{\gamma/p} \right. \right. \\ &\quad \left. \left. + \int_A |a|^\gamma \Lambda_s(da) \right] \right. \\ &\quad \left. + \left[c_1 \int_0^t ds \left(1 + \|X\|_s^{p\sigma} + \left(\int_{\mathbb{C}^d} \|z\|_s^p \mu(dz) \right)^{p\sigma/p} \right) \right]^{\gamma/2} \right\} \\ &\leq C \mathbb{E} \left[1 + |X_0|^\gamma + \int_0^t ds \left(1 + \|X\|_s^\gamma + \int_{\mathbb{C}^d} \|z\|_s^\gamma \mu(dz) \right. \right. \\ &\quad \left. \left. + \int_A |a|^\gamma \Lambda_s(da) \right) \right]. \end{aligned}$$

To pass from the second to the last line, we used the bound $(\int \|z\|_s^p \mu(dz))^{\gamma/p} \leq \int \|z\|_s^\gamma \mu(dz)$, which holds true since $\gamma \geq p$. To bound $(\int \|z\|_s^p \mu(dz))^{p\sigma/p}$ in the third line, we used the following argument. If $\gamma \geq 2$, we can pass the power $\gamma/2$ inside the integral in time by means of Jensen’s inequality and then use the inequality $|x|^{p\sigma\gamma/2} \leq 1 + |x|^\gamma$, which holds since $p\sigma \leq 2$. If $\gamma \leq 2$, we can use the inequality $|x|^{\gamma/2} \leq 1 + |x|$ followed by $|x|^{p\sigma} \leq 1 + |x|^\gamma$, which holds since $\gamma \geq p\sigma$. The first claim follows now from Gronwall’s inequality. If $P(X \in \cdot | B) = \mu$, then the above becomes

$$\begin{aligned} \mathbb{E} \int_{\mathcal{C}^d} \|z\|_t^\gamma \mu(dz) &= \mathbb{E} \|X\|_t^\gamma \\ &\leq C \mathbb{E} \left[|X|_0^\gamma + \int_0^t \left(1 + 2 \int_{\mathcal{C}^d} \|z\|_s^\gamma \mu(dz) + \int_A |a|^\gamma \Lambda_s(da) \right) ds \right]. \end{aligned}$$

The second claim now also follows from Gronwall’s inequality. \square

2.4. Mean field games with a discretized conditioning. Quite often, existence of a solution to a mean-field game without common noise is proved by means of Schauder’s fixed-point theorem. See, for instance [8, 10], where Schauder’s theorem is applied on $\mathcal{P}^p(\mathcal{C}^d)$ (with $p = 2$ in usual cases), for which compact subsets may be easily described. In the current setting, the presence of the common noise makes things much more complicated. Indeed, an equilibrium, denoted by μ in Definitions 2.1 and 2.3, is an element of the much larger space $[\mathcal{P}^p(\mathcal{C}^d)]^{\mathcal{C}^{m_0}}$, and the difficulty is to identify compact sets which could be stable under the transformations we consider.

2.4.1. Set-up. In this subsection, we thus define a discretization of the mean field game for which equilibria only depend on a finite number of random outcomes. Roughly speaking, equilibria can then be viewed as elements of the space $[\mathcal{P}^p(\mathcal{C}^d)]^k$ for some integer $k \geq 1$, the compact sets of which may be described quite simply. Compactness will be much easier to come by when the state coefficients are bounded and the control space compact, and so we will begin the search for MFG solutions by working under the following assumptions.

ASSUMPTION B. Assume that the following (B.1)–(B.5) hold for the rest of the subsection:

- (B.1) A is a compact metric space, and (b, σ, σ_0) is uniformly bounded.
- (B.2) $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$, and $p' > p \geq 1$.
- (B.3) The functions b, σ, σ_0, f , and g of (t, x, μ, a) are jointly measurable and continuous in (x, μ, a) for each t .
- (B.4) There exists $c_1 > 0$ such that, for all $(t, x, y, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

$$|b(t, x, \mu, a) - b(t, y, \mu, a)| + |(\sigma, \sigma_0)(t, x, \mu) - (\sigma, \sigma_0)(t, y, \mu)| \leq c_1|x - y|.$$

(B.5) There exists $c_2 > 0$ such that, for each $(t, x, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

$$|f(t, x, \mu, a)| + |g(x, \mu)| \leq c_2 \left(1 + |x|^p + \int_{\mathbb{R}^d} |z|^p \mu(dz) \right).$$

Note in particular that \mathcal{V} is compact since A is. Define then the canonical spaces

$$\Omega_0 := \mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m, \quad \Omega_f := \mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{V} \times \mathcal{C}^d.$$

Let ξ, B, W, Λ and X denote the identity maps on $\mathbb{R}^d, \mathcal{C}^{m_0}, \mathcal{C}^m, \mathcal{V}$ and \mathcal{C}^d , respectively. With a slight abuse of notation, we will also denote by ξ, B and W the projections from Ω_0 onto $\mathbb{R}^d, \mathcal{C}^{m_0}$ and \mathcal{C}^m , respectively, and by ξ, B, W, Λ and X the projections from Ω_f onto $\mathbb{R}^d, \mathcal{C}^{m_0}, \mathcal{C}^m, \mathcal{V}$ and \mathcal{C}^d , respectively.

The canonical processes B, W , and X generate obvious natural filtrations on Ω_f : $(\mathcal{F}_t^B)_{t \in [0, T]}$, $(\mathcal{F}_t^W)_{t \in [0, T]}$, and $(\mathcal{F}_t^X)_{t \in [0, T]}$. Recall the definition of $(\mathcal{F}_t^\Lambda)_{t \in [0, T]}$ on \mathcal{V} from (2.3). We will frequently work with filtrations generated by several canonical processes, such as $\mathcal{F}_t^{\xi, B, W} := \sigma(\xi, B_s, W_s : s \leq t)$ defined on Ω_0 , and $\mathcal{F}_t^{\xi, B, W, \Lambda} = \mathcal{F}_t^{\xi, B, W} \otimes \mathcal{F}_t^\Lambda$ defined on $\Omega_0 \times \mathcal{V}$. When needed, we will use the same symbol $(\mathcal{F}_t)_{t \in [0, T]}$ to denote the natural extension of a filtration $(\mathcal{F}_t)_{t \in [0, T]}$ on a space Ω to any product $\Omega \times \Omega'$, given by $(\mathcal{F}_t \otimes \{\emptyset, \Omega'\})_{t \in [0, T]}$. This permits to use $(\mathcal{F}_t^{\xi, B, W})_{t \in [0, T]}$ for the filtration on $\Omega_0 \times \mathcal{V}$ generated by (ξ, B, W) , and it should be clear from the context on which space the filtration is defined.

REMARK 2.5. In much of the analysis, it is useful to transfer our MFG solution concepts to a canonical space, as is often done in connecting weak solutions of SDEs with martingale problems. We find Definitions 2.1 and 2.3 (and later 3.1) most intuitive to state on an abstract probability space, but we favor a canonical space for carrying out the analysis. The rest of this section, for example, works exclusively on the canonical space outlined above. Similarly, in Section 3, after the initial Definition 3.1, we work mostly with the canonical space introduced in Section 3.2.

2.4.2. *Discretization procedure.* To define the discretized MFG problem, we discretize both time and the space of the common noise B . For each $n \geq 1$, let $t_i^n = i2^{-n}T$ for $i = 0, \dots, 2^n$. For each positive integer n , we choose a partition $\pi^n := \{C_1^n, \dots, C_n^n\}$ of \mathbb{R}^{m_0} into n measurable sets of strictly positive Lebesgue measure, such that π^{n+1} is a refinement of π^n for each n , and $\mathcal{B}(\mathbb{R}^{m_0}) = \sigma(\bigcup_{n=1}^\infty \pi^n)$. For a given n , the time mesh $(t_i^n)_{i=0, \dots, 2^n}$ and the spatial partition π^n yield a time–space grid along which we can discretize the trajectories in \mathcal{C}^{m_0} (which is the space carrying the common noise B). Intuitively, the idea is to project the increments of the trajectories between two consecutive times of the mesh $(t_i^n)_{i=0, \dots, 2^n}$ onto the spatial partition π^n . For $1 \leq k \leq 2^n$ and $\underline{i} = (i_1, \dots, i_k) \in \{1, \dots, n\}^k$, we thus

define $S_{\underline{i}}^{n,k}$ as the set of trajectories with increments up until time t_k in $C_{i_1}^n, \dots, C_{i_k}^n$, that is,

$$S_{\underline{i}}^{n,k} = \{\beta \in C^{m_0} : \beta_{t_j^n} - \beta_{t_{j-1}^n} \in C_{i_j}^n, \forall j = 1, \dots, k\}.$$

Obviously, the $S_{\underline{i}}^{n,k}$ s, $\underline{i} \in \{1, \dots, n\}^k$, form a finite partition (of cardinal n^k) of C^{m_0} , each $S_{\underline{i}}^{n,k}$ writing as a set of trajectories having the same discretization up until t_k and having a strictly positive \mathcal{W}^{m_0} -measure. The collection of all the possible discretization classes up until t_k thus reads

$$\Pi_k^n := \{S_{\underline{i}}^{n,k} : \underline{i} \in \{1, \dots, n\}^k\}.$$

When $k = 0$, we let $\Pi_0^n := \{C^{m_0}\}$, since all the trajectories are in the same discretization class.

For any $n \geq 0$, the filtration $(\sigma(\Pi_k^n))_{k=0, \dots, 2^n}$ is the filtration generated by the discretization of the canonical process. Clearly, $\sigma(\Pi_k^n) \subset \mathcal{F}_{t_k^n}^B$ and $\sigma(\Pi_k^n) \subset \sigma(\Pi_k^{n+1})$. For each $t \in [0, T]$, define

$$\lfloor t \rfloor_n := \max\{t_k^n : 0 \leq k \leq 2^n, t_k^n \leq t\}.$$

Let $\Pi^n(t)$ equal Π_k^n , where k is the largest integer such that $t_k^n \leq t$, and let $\mathcal{G}_t^n := \sigma(\Pi^n(t)) = \mathcal{G}_{\lfloor t \rfloor_n}^n$. It is straightforward to verify that $(\mathcal{G}_t^n)_{t \in [0, T]}$ is a filtration (i.e., $\mathcal{G}_s^n \subset \mathcal{G}_t^n$ when $s < t$) for each n and that

$$\mathcal{F}_t^B = \sigma\left(\bigcup_{n=1}^{\infty} \mathcal{G}_t^n\right).$$

2.4.3. Measures parameterized by discretized trajectories. The purpose of the discretization procedure described right below is to reduce the complexity of the scenarios upon which an equilibrium μ depends in Definitions 2.1 and 2.3. Roughly speaking, the strategy is to force μ to depend only on the discretization of the canonical process B on C^{m_0} . A natural way to do so is to restrict (in some way) the analysis to functions $\mu : \Pi_{2^n}^n \rightarrow \mathcal{P}^p(C^d)$ [instead of $\mu : C^{m_0} \rightarrow \mathcal{P}^p(C^d)$] or equivalently to functions $\mu : C^{m_0} \rightarrow \mathcal{P}^p(C^d)$ that are \mathcal{G}_T^n -measurable. In addition, some adaptedness is needed. We thus let \mathcal{M}_n denote the set of functions $\mu : C^{m_0} \rightarrow \mathcal{P}^p(C^d)$ that are \mathcal{G}_T^n -measurable such that for each $t \in [0, T]$ and $C \in \mathcal{F}_t^X$ the map $\beta \mapsto [\mu(\beta)](C)$ is \mathcal{G}_t^n -measurable. In particular, for $\mu \in \mathcal{M}_n$, the process $(\mu_t := \mu \circ \pi_t^{-1})_{t \in [0, T]}$ is $(\mathcal{G}_t^n)_{t \in [0, T]}$ -adapted and càdlàg [with values in $\mathcal{P}^p(\mathbb{R}^d)$].

Note that any $\mu \in \mathcal{M}_n$ is constant on S for each $S \in \Pi_{2^n}^n$ in the sense that $\beta \mapsto [\mu(\beta)](F)$ (which depends on the discretized trajectory) is constant on S for each Borel subset F of $\mathcal{P}^p(C^d)$. Endow \mathcal{M}_n with the topology of pointwise convergence, which of course is the same as the topology of uniform convergence since the common domain of each $\mu \in \mathcal{M}_n$ is effectively $\Pi_{2^n}^n$, which is finite.

Since $\mathcal{G}_T^n = \sigma(\Pi_{2^n}^n)$ is finite, the space \mathcal{M}_n is homeomorphic to a closed subset of $\mathcal{P}^p(\mathcal{C}^d)^{|\Pi_{2^n}^n|}$. Hence, \mathcal{M}_n is a metrizable closed convex subset of a locally convex topological vector space.

2.4.4. *Control problems.* Control problems will be described in terms of measures on $\Omega_0 \times \mathcal{V}$. Let

$$(2.4) \quad \mathcal{W}_\lambda := \lambda \times \mathcal{W}^{m_0} \times \mathcal{W}^m \in \mathcal{P}^{p'}(\Omega_0)$$

denote the distribution of the given sources of randomness on Ω_0 ; note that p' -integrability follows from the assumption $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$. The set of *admissible control rules* \mathcal{A}_f is defined to be the set of $Q \in \mathcal{P}(\Omega_0 \times \mathcal{V})$ such that B and W are independent $(\mathcal{F}_t^{\xi, B, W, \Lambda})_{t \in [0, T]}$ -Wiener processes under Q and $Q \circ (\xi, B, W)^{-1} = \mathcal{W}_\lambda$. Equivalently, $Q \in \mathcal{P}(\Omega_0 \times \mathcal{V})$ is in \mathcal{A}_f if $Q \circ (\xi, B, W)^{-1} = \mathcal{W}_\lambda$ and $(B_t - B_s, W_t - W_s)$ is Q -independent of $\mathcal{F}_s^{\xi, B, W, \Lambda}$ for each $0 \leq s < t \leq T$. Intuitively, this is just the set of “reasonable” joint laws of the control process with the given randomness. It is easy to check that \mathcal{A}_f is closed in the topology of weak convergence.

Given $\mu \in \mathcal{M}_n$ and $Q \in \mathcal{A}_f$, on the completion of the filtered probability space $(\Omega_0 \times \mathcal{V}, (\mathcal{F}_t^{\xi, B, W, \Lambda})_{t \in [0, T]}, Q)$ we may find a process Y such that (ξ, B, W, Λ, Y) satisfy the SDE

$$(2.5) \quad \begin{aligned} Y_t &= \xi + \int_0^t ds \int_A \Lambda_s(da) b(s, Y_s, \mu_s(B), a) \\ &\quad + \int_0^t \sigma(s, Y_s, \mu_s(B)) dW_s + \int_0^t \sigma_0(s, Y_s, \mu_s(B)) dB_s. \end{aligned}$$

Define the law of the solution and the interpolated solution by

$$\mathcal{R}_f(\mu, Q) := Q \circ (\xi, B, W, \Lambda, Y)^{-1}, \quad \mathcal{R}_f^n(\mu, Q) := Q \circ (\xi, B, W, \Lambda, \hat{Y}^n)^{-1},$$

where, for an element $x \in \mathcal{C}^d$, \hat{x}^n is the (delayed) linear interpolation of x along the mesh $(t_i^n)_{i=0, \dots, 2^n}$:

$$(2.6) \quad \hat{x}_t^n = \frac{2^n}{T}(t - t_i^n)x_{t_i^n} + \frac{2^n}{T}(t_{i+1}^n - t)x_{t_{i+1}^n} \quad \text{for } t \in [t_i^n, t_{i+1}^n], i = 0, \dots, 2^n - 1.$$

The delay ensures that \hat{X}^n is $(\mathcal{F}_t^X)_{t \in [0, T]}$ -adapted. By Lemma 2.4 and compactness of A , $\mathcal{R}_f(\mu, Q)$ and $\mathcal{R}_f^n(\mu, Q)$ are in $\mathcal{P}^p(\Omega_f)$. Note that \mathcal{R}_f and \mathcal{R}_f^n are well defined; by the uniqueness part in Lemma 2.4, $\mathcal{R}_f(\mu, Q)$ is the unique element P of $\mathcal{P}(\Omega_f)$ such that $P \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q$ and such that the canonical processes verify the SDE (2.5) under P . Again, as in footnote 3 on page 3751, it is no cause for concern that the Q -completion of the canonical filtration $(\mathcal{F}_t^{\xi, B, W, \Lambda})_{t \in [0, T]}$ may fail to be right-continuous.

The objective of the discretized control problem is as follows. Define the reward functional $\Gamma : \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$ by

$$(2.7) \quad \Gamma(\mu, q, x) := \int_0^T dt \int_A q_t(da) f(t, x_t, \mu_t, a) + g(x_T, \mu_T),$$

and the expected reward functional $J_f : \mathcal{M}_n \times \mathcal{P}^p(\Omega_f) \rightarrow \mathbb{R}$ by

$$J_f(\mu, P) := \mathbb{E}^P[\Gamma(\mu(B), \Lambda, X)].$$

For a given $\mu \in \mathcal{M}_n$, we are then dealing with the optimal control problem (with random coefficients) consisting in maximizing $J_f(\mu, P)$ over $P \in \mathcal{R}_f^n(\mu, \mathcal{A}_f)$. The set of maximizers is given by

$$\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f) := \arg \max_{P \in \mathcal{R}_f^n(\mu, \mathcal{A}_f)} J_f(\mu, P).$$

The set $\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ represents the optimal controls for the n th discretization corresponding to μ . A priori, it may be empty.

2.4.5. *Strong MFG solutions.* The main result of this section is the following theorem, which proves the existence of a strong MFG solution with weak control for our discretized mean field game.

THEOREM 2.6. *For each n , there exist $\mu \in \mathcal{M}_n$ and $P \in \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ such that $\mu = P(X \in \cdot | \mathcal{G}_T^n)$ [$P(X \in \cdot | \mathcal{G}_T^n)$ being seen as a map from \mathcal{C}^{m_0} to $\mathcal{P}^p(\mathcal{C}^d)$, constant on each $S \in \Pi_{2^n}^n$.]*

PROOF. A MFG equilibrium may be viewed as a fixed point of a set-valued function. Defining the set-valued map $F : \mathcal{M}_n \rightarrow 2^{\mathcal{M}_n}$ (where $2^{\mathcal{M}_n}$ is seen as the collection of subsets of \mathcal{M}_n) by

$$F(\mu) := \{P(X \in \cdot | \mathcal{G}_T^n) : P \in \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)\},$$

the point is indeed to prove that F admits a fixed point, that is, a point $\mu \in F(\mu)$. Since the unique event in \mathcal{G}_T^n of null probability under P is the empty set, we notice that $G(P) := P(X \in \cdot | \mathcal{G}_T^n)$ is uniquely defined for each $P \in \mathcal{P}^p(\Omega_f)$. Let \mathcal{P}_f^p denote those elements P of $\mathcal{P}^p(\Omega_f)$ for which $P \circ (\xi, B, W, \Lambda)^{-1}$ is admissible, that is $\mathcal{P}_f^p := \{P \in \mathcal{P}^p(\Omega_f) : P \circ (\xi, B, W, \Lambda)^{-1} \in \mathcal{A}_f\}$. For $P \in \mathcal{P}_f^p$, $G(P)$ is given by

$$(2.8) \quad \begin{aligned} G(P) : \mathcal{C}^{m_0} \ni \beta &\mapsto \sum_{S \in \Pi_{2^n}^n} P(X \in \cdot | B \in S) 1_S(\beta) \\ &= \sum_{S \in \Pi_{2^n}^n} \frac{P(\{X \in \cdot\} \cap \{B \in S\})}{\mathcal{W}^{m_0}(S)} 1_S(\beta). \end{aligned}$$

The very first step is then to check that $F(\mu) \subset \mathcal{M}_n$ for each $\mu \in \mathcal{M}_n$. The above formula shows that, for $P \in \mathcal{P}_f^p$, $G(P)$ reads as a \mathcal{G}_T^n -measurable function from \mathcal{C}^{m_0} to $\mathcal{P}^p(\mathcal{C}^d)$. To prove that $G(P) \in \mathcal{M}_n$, it suffices to check the adaptedness condition in the definition of \mathcal{M}_n (see Section 2.4.3). For our purpose, we can restrict the proof to the case when X is P a.s. piecewise affine as in (2.6). For each $t \in [0, T]$ and $C \in \mathcal{F}_t^X$, we have that $1_C(X) = 1_{C'}(X)$ P a.s. for some $C' \in \mathcal{F}_{\lfloor t \rfloor_n}^X$. Now, $\mathcal{G}_T^n = \mathcal{G}_{\lfloor t \rfloor_n}^n \vee \mathcal{H}$, where $\mathcal{H} \subset \sigma(B_s - B_{\lfloor t \rfloor_n} : s \in [\lfloor t \rfloor_n, T])$. Since \mathcal{H} is P -independent of $\mathcal{F}_{\lfloor t \rfloor_n}^X \vee \mathcal{G}_{\lfloor t \rfloor_n}^n$, we deduce that, P a.s., $P(X \in C | \mathcal{G}_T^n) = P(X \in C' | \mathcal{G}_{\lfloor t \rfloor_n}^n)$. Since the unique event in \mathcal{G}_T^n of null probability under P is the empty set, we deduce that the process $(P(X_{\lfloor t \rfloor_n} \in \cdot | \mathcal{G}_T^n))_{t \in [0, T]}$ is $(\mathcal{G}_t^n)_{t \in [0, T]}$ -adapted. This shows that $G(P) \in \mathcal{M}_n$, and thus $F(\mu) \subset \mathcal{M}_n$.

We will achieve the proof by verifying the hypotheses of the Kakutani–Fan–Glicksberg fixed-point theorem for set-valued functions [2], Corollary 17.55. Namely, we will show that F is upper hemicontinuous with nonempty compact convex values, and we will find a compact convex subset $\mathcal{Q} \subset \mathcal{M}_n$ such that $F(\mu) \subset \mathcal{Q}$ for each $\mu \in \mathcal{Q}$.

First step: Continuity of set-valued functions. For the necessary background on set-valued analysis the reader is referred to [2], Chapter 17. For this paragraph, fix two metric spaces E and F . A set valued function $h : E \rightarrow 2^F$ is *lower hemicontinuous* if, whenever $x_n \rightarrow x$ in E and $y \in h(x)$, there exists $y_{n_k} \in h(x_{n_k})$ such that $y_{n_k} \rightarrow y$. If $h(x)$ is closed for each $x \in E$ then h is called *upper hemicontinuous* if, whenever $x_n \rightarrow x$ in E and $y_n \in h(x_n)$ for each n , the sequence (y_n) has a limit point in $h(x)$. We say h is *continuous* if it is both upper hemicontinuous and lower hemicontinuous. If $h(x)$ is closed for each $x \in E$ and F is compact, then h is upper hemicontinuous if and only if its graph $\{(x, y) : x \in E, y \in h(x)\}$ is closed.

First, we check the continuity of the function

$$\mathcal{P}_f^p \ni P \mapsto P(X \in \cdot | B \in S) \in \mathcal{P}^p(\mathcal{C}^d) \quad \text{for } S \in \Pi_{2^n}^n.$$

This is straightforward, thanks to the finiteness of the conditioning σ -field. Let $\phi : \mathcal{C}^d \rightarrow \mathbb{R}$ be continuous with $|\phi(x)| \leq c(1 + \|x\|_T^p)$ for all $x \in \mathcal{C}^d$, for some $c > 0$. Proposition A.1(3) in the Appendix says that it is enough to prove that $\mathbb{E}^{P_k}[\phi(X) | B \in S] \rightarrow \mathbb{E}^P[\phi(X) | B \in S]$ whenever $P_k \rightarrow P$ in $\mathcal{P}^p(\Omega_f)$. This follows from Lemma A.4, which implies that the following real-valued function is continuous:

$$\mathcal{P}^p(\Omega_f) \ni P \mapsto \mathbb{E}^P[\phi(X) | B \in S] = \mathbb{E}^P[\phi(X) 1_S(B)] / \mathcal{W}^{m_0}(S).$$

Basically, Lemma A.4 handles the discontinuity of the indicator function 1_S together with the fact that ϕ is not bounded. It follows that the function $G : \mathcal{P}_f^p \rightarrow \mathcal{M}_n$ given by (2.8) is continuous. The set-valued function F is simply the composition of G with the set-valued function $\mu \mapsto \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$. Therefore, to prove that F is upper hemicontinuous, it is sufficient to prove that $\mu \mapsto \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ is upper hemicontinuous.

Second step: Analysis of the mapping: $\mu \mapsto \mathcal{R}_f^n(\mu, \mathcal{A}_f)$. Following the first step, the purpose of the second step is to prove continuity of the set-valued function

$$\mathcal{M}_n \ni \mu \mapsto \mathcal{R}_f^n(\mu, \mathcal{A}_f) := \{\mathcal{R}_f^n(\mu, Q) : Q \in \mathcal{A}_f\} \in 2^{\mathcal{P}^p(\Omega_f)}.$$

Since the map $\mathcal{C}^d \ni x \mapsto \hat{x}^n \in \mathcal{C}^d$ is continuous [see (2.6)], it suffices to prove continuity with \mathcal{R}_f^n replaced by \mathcal{R}_f . To do so, we prove first that $\mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)$ is relatively compact by showing that each of the sets of marginal measures is relatively compact; see Lemma A.3. Clearly, $\{P \circ (\xi, B, W)^{-1} : P \in \mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)\} = \{\mathcal{W}_\lambda\}$ is compact in $\mathcal{P}^p(\Omega_0)$. Since A is compact, so is \mathcal{V} , and thus $\{P \circ \Lambda^{-1} : P \in \mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)\}$ is relatively compact in $\mathcal{P}^p(\mathcal{V})$. Since $b, \sigma,$ and σ_0 are bounded, Aldous' criterion (see Proposition B.1 for details) shows that $\{P \circ X^{-1} : P \in \mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)\}$ is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$.

Continuity of the set-valued function $\mathcal{R}_f(\cdot, \mathcal{A}_f)$ will follow from continuity of the single-valued function \mathcal{R}_f . Since the range is relatively compact, it suffices to show that the graph of \mathcal{R}_f is closed. Let $(\mu_k, Q_k) \rightarrow (\mu, Q)$ in $\mathcal{M}_n \times \mathcal{A}_f$ and $P_k := \mathcal{R}_f(\mu_k, Q_k) \rightarrow P$ in $\mathcal{P}^p(\Omega_f)$. It is clear that

$$P \circ (\xi, B, W, \Lambda)^{-1} = \lim_{k \rightarrow \infty} P_k \circ (\xi, B, W, \Lambda)^{-1} = \lim_{k \rightarrow \infty} Q_k = Q.$$

It follows from the results of Kurtz and Protter [36] that the state SDE (2.5) holds under the limiting measure P , since it holds under each P_k . Since $\mathcal{R}_f(\mu, Q)$ is the unique law on Ω_f under which (ξ, B, W, Λ) has law Q and (ξ, B, W, Λ, X) solves (2.5), we deduce that $P = \mathcal{R}_f(\mu, Q)$. We finally conclude that $\mathcal{R}_f(\cdot, \mathcal{A}_f)$, and thus $\mathcal{R}_f^n(\cdot, \mathcal{A}_f)$ are continuous.

Third step: Analysis of the mapping: $\mu \mapsto \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$. As a by-product of the previous analysis, we notice that, for each $\mu \in \mathcal{M}_n, \mathcal{R}_f(\mu, \mathcal{A}_f)$ is closed and relatively compact and thus compact. By continuity of the map $\mathcal{C}^d \ni x \mapsto \hat{x}^n \in \mathcal{C}^d$ [see (2.6)], $\mathcal{R}_f^n(\mu, \mathcal{A}_f)$ is also compact.

Since f and g are continuous in (x, μ, a) and have p -order growth, it can be checked that the reward functional Γ is continuous (although quite elementary, the proof is given in the Appendix; see Lemma A.5). This implies that the expected reward functional

$$\mathcal{M}_n \times \mathcal{P}^p(\Omega_f) \ni (\mu, P) \mapsto J_f(\mu, P) \in \mathbb{R}$$

is also continuous. If Γ is bounded, continuity follows from the fact that $(\mu_k, P_k) \rightarrow (\mu, P)$ implies $P_k \circ (\mu_k(B), \Lambda, X)^{-1} \rightarrow P \circ (\mu(B), \Lambda, X)^{-1}$. In the general case when Γ has p -order growth, it follows from Lemma A.4.

By compactness of $\mathcal{R}_f^n(\mu, \mathcal{A}_f)$ and by continuity of $J_f, \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ is nonempty and compact. Moreover, from a well-known theorem of Berge [2], Theorem 17.31, the set-valued function $\mathcal{R}_f^{*,n} : \mathcal{M}_n \rightarrow 2^{\mathcal{P}(\Omega_f)}$ is upper hemicontinuous.

Fourth step: Convexity of $\mathcal{R}_f^{,n}(\mu, \mathcal{A}_f)$.* We now prove that, for each $\mu \in \mathcal{M}_n$, $\mathcal{R}_f^n(\mu, \mathcal{A}_f)$ is convex. By linearity of the map $\mathcal{C}^d \ni x \mapsto \hat{x}^n \in \mathcal{C}^d$ [see (2.6)], it is sufficient to prove that $\mathcal{R}_f(\mu, \mathcal{A}_f)$ is convex. To this end, we observe first that \mathcal{A}_f is convex. Given $Q_i, i = 1, 2$, in \mathcal{A}_f , and $c \in (0, 1)$, we notice that (B, W) is a Wiener process with respect to $(\mathcal{F}_t^{\xi, B, W, \Lambda})_{t \in [0, T]}$ under $cP^1 + (1 - c)P^2$, where $P^i := \mathcal{R}_f(\mu, Q^i)$ for $i = 1, 2$. [Use the fact that (B, W) is a Wiener process under both P^1 and P^2 .] Moreover, the state equation holds under $cP^1 + (1 - c)P^2$. Since $(cP^1 + (1 - c)P^2) \circ (\xi, B, W, \Lambda)^{-1} = cQ^1 + (1 - c)Q^2$, we deduce that $cP^1 + (1 - c)P^2$ is the unique probability on Ω_f under which (ξ, B, W, Λ) has law $cQ^1 + (1 - c)Q^2$ and (ξ, B, W, Λ, X) solves the state equation. This proves that $cP^1 + (1 - c)P^2 = \mathcal{R}_f(\mu, cQ^1 + (1 - c)Q^2)$.

By linearity of the map $P \mapsto J_f(\mu, P)$, we deduce that the set-valued function $\mathcal{R}_f^{*,n} : \mathcal{M}_n \rightarrow 2^{\mathcal{P}(\Omega_f)}$ has nonempty convex values. (Nonemptiness follows from the previous step.)

Conclusion. Finally, we place ourselves in a convex compact subset of \mathcal{M}_n , by first finding a convex compact set $\mathcal{Q}_0 \subset \mathcal{P}^p(\mathcal{C}^d)$ containing $\{P \circ X^{-1} : P \in \mathcal{R}_f^n(\mathcal{M}_n, \mathcal{A}_f)\}$. To this end, note that the boundedness of (b, σ, σ_0) of assumption (B.1) implies that for each smooth $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with compact support,

$$C_\phi := \sup_{t, x, \mu, a} \left| b(t, x, \mu, a)^\top D\phi(x) + \frac{1}{2} \text{Tr}[(\sigma\sigma^\top + \sigma_0\sigma_0^\top)(t, x, \mu)D^2\phi(x)] \right| < \infty,$$

where D and D^2 denote gradient and Hessian, respectively. Following Lemma 2.4 and again using boundedness of (b, σ, σ_0) , it is standard to show that

$$M := \sup\{\mathbb{E}^P \|X\|_T^{p'} : P \in \mathcal{R}_f^n(\mathcal{M}_n, \mathcal{A}_f)\} < \infty.$$

Now, define \mathcal{Q}_1 to be the set of $P \in \mathcal{P}^p(\mathcal{C}^d)$ satisfying:

- (1) $P \circ X_0^{-1} = \lambda$,
- (2) $\mathbb{E}^P \|X\|_T^{p'} \leq M$,
- (3) for each nonnegative smooth $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with compact support, the process $\phi(X_t) + C_\phi t$ is a P -submartingale.

It is clear that \mathcal{Q}_1 is convex and contains $\{P \circ X^{-1} : P \in \mathcal{R}_f^n(\mathcal{M}_n, \mathcal{A}_f)\}$. Using a well-known tightness criterion of Stroock and Varadhan [43], Theorem 1.4.6, conditions (1) and (3) together imply that \mathcal{Q}_1 is tight, and the p' -moment bound of (2) then ensures that it is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$ (see Proposition A.1). It is straightforward to check that \mathcal{Q}_1 is in fact closed, and thus it is compact. Next, define

$$\mathcal{Q}_2 := \{P \circ (\hat{X}^n)^{-1} : P \in \mathcal{Q}_1\} \subset \mathcal{P}^p(\mathcal{C}^d),$$

and note that \mathcal{Q}_2 is also convex and compact, since $x \mapsto \hat{x}^n$ is continuous and linear.

Recalling the definition of \mathcal{P}_f^p from the first step, let

$$\begin{aligned} \mathcal{Q}_3 &:= \{P \in \mathcal{P}_f^p : P \circ X^{-1} \in \mathcal{Q}_2\} \\ &= \{P \in \mathcal{P}^p(\Omega_f) : P \circ (\xi, B, W, \Lambda)^{-1} \in \mathcal{A}_f, P \circ X^{-1} \in \mathcal{Q}_2\}. \end{aligned}$$

It is easily checked that \mathcal{A}_f is a compact set: closedness is straightforward, and, as in the second step, \mathcal{A}_f is relatively compact since A is compact and the (ξ, B, W) -marginal is fixed. It follows from compactness of \mathcal{A}_f and \mathcal{Q}_2 that \mathcal{Q}_3 is compact (see Lemma A.3). Similarly, it follows from convexity of \mathcal{A}_f and \mathcal{Q}_2 that \mathcal{Q}_3 is convex.

Finally, define $\mathcal{Q} := G(\mathcal{Q}_3)$. Note that $\mathcal{Q} \subset \mathcal{M}_n$, since we saw at the beginning of the proof that indeed $G(P) \in \mathcal{M}_n$ whenever $P \in \mathcal{P}_f^p$ satisfies $P(X = \hat{X}^n) = 1$. As emphasized by (2.8), G is linear and continuous (by the first step). Hence, \mathcal{Q} is convex and compact since \mathcal{Q}_3 is. Moreover, for each $\mu \in \mathcal{M}_n$, $F(\mu) = G(\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f))$ is convex and compact, since $\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ is convex and compact (see the third and fourth steps). Since $F(\mu) \subset \mathcal{Q}$ for each $\mu \in \mathcal{Q}$, the proof is complete. \square

3. Weak limits of discretized MFG. We now aim at passing to the limit in the discretized MFG as the time–space grid is refined, the limit being taken in the weak sense. To do so, we show that any sequence of solutions of the discretized MFG is relatively compact, and we characterize the limits. This requires a lot of precaution, the main reason being that measurability properties are not preserved under weak limits. In particular, we cannot generally ensure that in the limit, the conditional measure μ remains B -measurable in the limit. This motivates the new notion of *weak MFG solution* in the spirit of weak solutions to standard stochastic differential equations. We will thus end up with *weak solutions with weak controls*. Assumption A holds throughout the section.

3.1. *Weak MFG solution with weak control.* Since the conditional measure μ is no longer expected to be measurable with respect to B , we need another space for it. One of the main idea in the sequel is to enlarge the space supporting μ . Namely, instead of considering μ as the conditional distribution of X given some σ -field, we will see μ as the conditional distribution of the whole (W, Λ, X) . This will allow us to describe in a complete way the correlations between the different processes. In other words, μ will be viewed as an element of $\mathcal{P}^p(\mathcal{X})$, with $\mathcal{X} := \mathcal{C}^m \times \mathcal{V} \times \mathcal{C}^d$, and with $\mu^x := \mu(\mathcal{C}^m \times \mathcal{V} \times \cdot)$ denoting the \mathcal{C}^d -marginal.

This brings us to the following definition of a weak MFG solution, the term *weak* referring to the fact that the conditional distribution μ may not be adapted to the noise B :

DEFINITION 3.1 (Weak MFG solution with weak control). A weak MFG solution with weak control (or simply a weak MFG solution) with initial condition λ is a tuple $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$, where $(\Omega, (\mathcal{F}_t)_{t \in [0, T]}, P)$ is a probability space with a complete filtration supporting (B, W, μ, Λ, X) satisfying:

(1) The processes $(B_t)_{t \in [0, T]}$ and $(W_t)_{t \in [0, T]}$ are independent $(\mathcal{F}_t)_{t \in [0, T]}$ Wiener processes of respective dimension m_0 and m , the process $(X_t)_{t \in [0, T]}$ is $(\mathcal{F}_t)_{t \in [0, T]}$ -adapted with values in \mathbb{R}^d , and $P \circ X_0^{-1} = \lambda$. Moreover, μ is a random element of $\mathcal{P}^p(\mathcal{X})$ such that $\mu(C)$ is \mathcal{F}_t -measurable for each $C \in \mathcal{F}_t^{W, \Lambda, X}$ and $t \in [0, T]$.

(2) X_0, W and (B, μ) are independent.

(3) $(\Lambda_t)_{t \in [0, T]}$ is $(\mathcal{F}_t)_{t \in [0, T]}$ -progressively measurable with values in $\mathcal{P}(A)$ and

$$\mathbb{E} \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty.$$

Moreover, $\sigma(\Lambda_s : s \leq t)$ is conditionally independent of $\mathcal{F}_T^{X_0, B, W, \mu}$ given $\mathcal{F}_t^{X_0, B, W, \mu}$ for each $t \in [0, T]$, where

$$\mathcal{F}_t^{X_0, B, W, \mu} = \sigma(X_0, B_s, W_s : s \leq t) \vee \sigma(\mu(C) : C \in \mathcal{F}_t^{W, \Lambda, X}).$$

(4) The state equation holds:

$$(3.1) \quad \begin{aligned} dX_t &= \int_A b(t, X_t, \mu_t^x, a) \Lambda_t(da) dt + \sigma(t, X_t, \mu_t^x) dW_t \\ &\quad + \sigma_0(t, X_t, \mu_t^x) dB_t. \end{aligned}$$

(5) If $(\Omega', (\mathcal{F}'_t)_{t \in [0, T]}, P')$ is another filtered probability space supporting processes $(B', W', \nu, \Lambda', X')$ satisfying (1)–(4) and $P \circ (X_0, B, W, \mu)^{-1} = P' \circ (X'_0, B', W', \nu)^{-1}$, then

$$\mathbb{E}^P[\Gamma(\mu^x, \Lambda, X)] \geq \mathbb{E}^{P'}[\Gamma(\nu^x, \Lambda', X')],$$

where Γ was defined in (2.7).

(6) $\mu = P((W, \Lambda, X) \in \cdot | B, \mu)$ a.s. That is μ is a version of the conditional law of (W, Λ, X) given (B, μ) .

If there exists an A -valued process $(\alpha_t)_{t \in [0, T]}$ such that $P(\Lambda_t = \delta_{\alpha_t} \text{ a.e. } t) = 1$, then we say the tuple is a weak MFG solution with weak strict control. It is said to be a weak MFG solution with strong control if the process $(\alpha_t)_{t \in [0, T]}$ is progressive with respect to the P -completion of $(\mathcal{F}_t^{X_0, B, W, \mu})_{t \in [0, T]}$.

A few comments regarding this definition are in order. The MFG solution is strong (see Definitions 2.1 and 2.3), if μ is B -measurable, and it is weak otherwise. The fact that μ may not be B -measurable justifies the need for enlarging the σ -field in the conditional fixed-point problem (6). The conditioning is now taken

in terms of the smallest σ -field that carries the common sources of randomness, namely B and μ . A crucial fact is that this σ -field is independent of the idiosyncratic source of noise W ; see (2). Similarly, whether or not μ is B -measurable, the control is weak if it is not progressively measurable with respect to the completion of $(\mathcal{F}_t^{X_0, B, W, \mu})_{t \in [0, T]}$. Note finally that assumption (6) in the definition of weak MFG solution with weak control ensures that μ_t^x is \mathcal{F}_t -adapted, as will be seen in Remark 3.5.

Since this notion of “weak control” is unusual, especially the conditional independence requirement in (3), we offer the following interpretation. An agent has full information, in the sense that he observes (in an adapted fashion) the initial state X_0 , the noises B and W , and also the distribution μ of the (infinity of) other agents’ states, controls and noises. That is, the agent has access to $\mathcal{F}_t^{X_0, B, W, \mu}$ at time t . Controls are allowed to be randomized externally to these observations, but such a randomization must be *conditionally independent of future information given current information*. This constraint will be called *compatibility*.

The main result of this section is the following.

THEOREM 3.2. *Under Assumption A, there exists a weak MFG solution with weak control that satisfies (with the notation of Definition 3.1) $\mathbb{E} \int_0^T \int_A |a|^{p'} \times \Lambda_t(da) dt < \infty$.*

3.2. Canonical space. In order to take weak limits of the discretized MFG, which is our purpose, it is convenient to work on a fixed canonical space. This was mentioned in Remark 2.5, although the canonical space introduced here is necessarily somewhat larger than the space Ω_f used in the discretized MFG (with the subscript “ f ” standing for “finite”). Indeed, the notation to follow is used heavily throughout the rest of the paper. As in the previous section, $\Omega_0 := \mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m$ will support the initial condition and the two Wiener processes driving the state equation. We also need the space \mathcal{V} defined in the previous Section 2.3 to handle the relaxed controls and the space \mathcal{C}^d to handle the solution of the state equation. To sum up, we have

$$\begin{aligned} \mathcal{X} &:= \mathcal{C}^m \times \mathcal{V} \times \mathcal{C}^d, & \Omega_0 &:= \mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m, \\ \Omega &:= \mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V} \times \mathcal{C}^d. \end{aligned}$$

The identity map on Ω_0 is still denoted by (ξ, B, W) and the identity map on Ω by $(\xi, B, W, \mu, \Lambda, X)$. The map μ generates the canonical filtration

$$(3.2) \quad \mathcal{F}_t^\mu := \sigma(\mu(C) : C \in \mathcal{F}_t^{W, \Lambda, X}).$$

Recall from (2.3) the definition of the canonical filtration $(\mathcal{F}_t^\Lambda)_{t \in [0, T]}$ on \mathcal{V} , and recall from (2.4) the definition of $\mathcal{W}_\lambda \in \mathcal{P}(\Omega_0)$. We next specify how μ and Λ are allowed to correlate with each other and with the given sources of randomness (ξ, B, W) . We will refer to the conditional independence requirement (3) of Definition 3.1 as *compatibility*, defined a bit more generally as follows:

(1) An element $\rho \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}))$ is said to be in $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ if (ξ, B, W) has law \mathcal{W}_λ under ρ and if B and W are independent $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -Wiener processes under ρ . The subscript c and the symbol \rightsquigarrow in $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ indicate that the extension of the probability measure \mathcal{W}_λ from Ω_0 to $\Omega_0 \times \mathcal{P}^p(\mathcal{X})$ is *compatible*.

(2) For $\rho \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}))$, an element $Q \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$ is said to be in $\mathcal{P}_c^p[(\Omega_0 \times \mathcal{P}^p(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}]$ if (ξ, B, W, μ) has law ρ under Q and $\mathcal{F}_t^{\xi, B, W, \mu}$ and \mathcal{F}_t^Λ are conditionally independent given $\mathcal{F}_t^{\xi, B, W, \mu}$. Again, Q is then *compatible* with ρ in the sense that, given the observation of (ξ, B, W, μ) up until time t , the observation of Λ up until t has no influence on the future of (ξ, B, W, μ) .

REMARK 3.3. These notions of compatibility are special cases of a more general idea, which goes by several names in the literature. It can be viewed as a compatibility of a larger filtration with a smaller one on a single probability space, in which case this is sometimes known as the *H-hypothesis* [7]. Alternatively, this can be seen as a property of an extension of a filtered probability space, known as a *very good extension* [32] or *natural extension* [19]. The term *compatible* is borrowed from Kurtz [35]. The curious reader is referred to [7, 32, 35] for some equivalent definitions, but we will derive the needed results as we go, to keep the paper self-contained.

We now have enough material to describe the optimization problem we will deal with. Given $\rho \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}))$ (that is given the original sources of randomness and a compatible random measure), we denote by $\mathcal{A}(\rho) := \mathcal{P}_c^p[(\Omega_0 \times \mathcal{P}^p(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}]$ [see (2) above] the *set of admissible relaxed controls*.

Observe from (1) and (2) right above that, for $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ and $Q \in \mathcal{A}(\rho)$, the process (B, W) is a Wiener process with respect to the filtration $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda})_{t \in [0, T]}$. Following (1), we will denote by $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}]$ the elements of $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$ under which (B, W) is a Wiener process with respect to the filtration $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda})_{t \in [0, T]}$, so that, if $Q \in \mathcal{A}(\rho)$ with $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$, then $Q \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}]$.

For $Q \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}]$, Λ is p -integrable, that is $\mathbb{E}^Q \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty$. On the completion of the space $(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}, (\mathcal{F}_t^{\xi, B, W, \mu, \Lambda})_{t \in [0, T]}, Q)$ there exists a unique strong solution X of the SDE

$$\begin{aligned}
 (3.3) \quad X_t &= \xi + \int_0^t ds \int_A \Lambda_s(da) b(s, X_s, \mu_s^x, a) + \int_0^t \sigma(s, X_s, \mu_s^x) dW_s \\
 &\quad + \int_0^t \sigma_0(s, X_s, \mu_s^x) dB_s,
 \end{aligned}$$

where we recall that $\mu^x(\cdot) = \mu(C^m \times \mathcal{V} \times \cdot)$ is the marginal law of μ on C^d and $\mu_s^x := \mu^x \circ \pi_s^{-1}$. We then denote by $\mathcal{R}(Q) := Q \circ (\xi, B, W, \mu, \Lambda, X)^{-1} \in \mathcal{P}(\Omega)$

the joint law of the solution. $\mathcal{R}(Q)$ is the unique element P of $\mathcal{P}(\Omega)$ such that $P \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q$ and such that the canonical processes verify the SDE (3.3) under P (again, see footnote 3 on page 3751 for a related discussion about the choice of the filtration). By Lemma 2.4, $\mathcal{R}(Q) \in \mathcal{P}^P(\Omega)$.

For each $\rho \in \mathcal{P}_c^P[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^P(\mathcal{X})]$, define

$$\mathcal{RA}(\rho) := \mathcal{R}(\mathcal{A}(\rho)) = \{\mathcal{R}(Q) : Q \in \mathcal{A}(\rho)\}.$$

Recalling the definition of Γ from (2.7), the expected reward functional $J : \mathcal{P}^P(\Omega) \rightarrow \mathbb{R}$ is defined by

$$(3.4) \quad J(P) := \mathbb{E}^P[\Gamma(\mu^x, \Lambda, X)].$$

The problem of maximizing $J(P)$ over $P \in \mathcal{RA}(\rho)$ is called the *control problem associated to ρ* . Define the set of optimal controls corresponding to ρ by

$$(3.5) \quad \mathcal{A}^*(\rho) := \arg \max_{Q \in \mathcal{A}(\rho)} J(\mathcal{R}(Q)),$$

and note that

$$\mathcal{RA}^*(\rho) := \mathcal{R}(\mathcal{A}^*(\rho)) = \arg \max_{P \in \mathcal{RA}(\rho)} J(P).$$

Pay attention that, a priori, the set $\mathcal{A}^*(\rho)$ may be empty.

3.3. *Relative compactness and MFG pre-solution.* With the terminology introduced above, we make a useful intermediate definition.

DEFINITION 3.4 (MFG pre-solution). Suppose $P \in \mathcal{P}^P(\Omega)$ satisfies the following:

- (1) $(B, \mu), \xi$ and W are independent under P .
- (2) $P \in \mathcal{RA}(\rho)$ where $\rho := P \circ (\xi, B, W, \mu)^{-1}$ is in $\mathcal{P}_c^P[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^P(\mathcal{X})]$.
- (3) $\mu = P((W, \Lambda, X) \in \cdot | B, \mu)$ a.s. That is, μ is a version of the conditional law of (W, Λ, X) given (B, μ) .

Then we say that P is a *MFG pre-solution*.

A *MFG pre-solution* is simply an MFG solution without the requirement that the control be optimal, as is explained in Lemma 3.9 below. Note also that condition (2) implicitly requires that B and W are Wiener processes relative to $(\Omega, (\mathcal{F}_t^{B, W, \mu, \Lambda, X})_{t \in [0, T]}, P)$.

REMARK 3.5. If P is a MFG pre-solution then condition (3) implies that $\mu_t^x = P(X_t \in \cdot | \mathcal{F}_t^{B, \mu^x})$ for each t , where

$$\mathcal{F}_t^{B, \mu^x} := \sigma(B_s, \mu_s^x : s \leq t).$$

Indeed, for any bounded measurable $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$, since $\mathcal{F}_t^{B, \mu^x} \subset \mathcal{F}_T^{B, \mu}$ and μ_t^x is \mathcal{F}_t^{B, μ^x} -measurable, we may condition by \mathcal{F}_t^{B, μ^x} on both sides of the equation $\mathbb{E}[\phi(X_t) | \mathcal{F}_T^{B, \mu}] = \int \phi d\mu_t^x$ to get the desired result. More carefully, this tells us $\mathbb{E}[\phi(X_t) | \mathcal{F}_t^{B, \mu^x}] = \int \phi d\mu_t^x$ a.s. for each ϕ , and by taking ϕ from a countable sequence which is dense in pointwise convergence we conclude that μ_t^x is a version of the regular conditional law of X_t given \mathcal{F}_t^{B, μ^x} .

Definition 3.4 is motivated by the following.

LEMMA 3.6. *Assume that Assumption B holds. For each n , by Theorem 2.6 we may find $\mu^n \in \mathcal{M}_n$ and $P_n \in \mathcal{R}_f^{*,n}(\mu^n, \mathcal{A}_f)$ such that $\mu^n = P_n(X \in \cdot | \mathcal{G}_T^n)$ (both being viewed as random probability measures on \mathcal{C}^d). On \mathcal{X} , define*

$$\bar{\mu}^n := P_n((W, \Lambda, X) \in \cdot | \mathcal{G}_T^n),$$

so that $\bar{\mu}^n$ can be viewed as a map from \mathcal{C}^{m_0} into $\mathcal{P}^p(\mathcal{X})$ and $\bar{\mu}^n(B)$ as a random element of $\mathcal{P}^p(\mathcal{X})$. Then the probability measures

$$\bar{P}_n := P_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1}$$

are relatively compact in $\mathcal{P}^p(\Omega)$, and every limit point is a MFG pre-solution.

PROOF. *First step.* Write $P_n = \mathcal{R}_f^n(\mu^n, Q_n)$, for some $Q_n \in \mathcal{A}_f$, and define $P'_n := \mathcal{R}_f(\mu^n, Q_n)$. Let

$$\bar{P}'_n = P'_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1},$$

so that $\bar{P}_n = \bar{P}'_n \circ (\xi, B, W, \mu, \Lambda, \hat{X}^n)^{-1}$, where \hat{X}^n was defined in (2.6). We first show that \bar{P}'_n are relatively compact in $\mathcal{P}^p(\Omega)$. Clearly, $P'_n \circ (B, W)^{-1}$ are relatively compact, and so are $P'_n \circ \Lambda^{-1}$ by compactness of \mathcal{V} . Since A is compact, the moment bound of Lemma 2.4 yields

$$(3.6) \quad \sup_n \mathbb{E}^{\bar{P}'_n} \int_{\mathcal{C}^d} \|x\|_T^{p'} \mu^x(dx) < \infty.$$

Thus, $P'_n \circ X^{-1}$ are relatively compact, by an application of Aldous' criterion (see Proposition B.1). By Proposition A.2, relative compactness of $P'_n \circ (\bar{\mu}^n(B))^{-1}$ follows from that of the mean measures $P'_n \circ (W, \Lambda, X)^{-1}$ and from the uniform p' -moment bound of Lemma 2.4. Precisely, for a point $\chi_0 \in \mathcal{X}$ and a metric ℓ on \mathcal{X} compatible with the topology,

$$\begin{aligned} \sup_n \int_{\Omega} \left(\int_{\mathcal{X}} \ell^{p'}(\chi_0, \chi) [\bar{\mu}^n(B)](d\chi) \right) dP'_n &= \sup_n \mathbb{E}^{P'_n} [\mathbb{E}^{P'_n} [\ell^{p'}(\chi_0, (W, \Lambda, X)) | \mathcal{G}_T^n]] \\ &= \sup_n \mathbb{E}^{P'_n} [\ell^{p'}(\chi_0, (W, \Lambda, X))] < \infty. \end{aligned}$$

Hence, \overline{P}'_n are relatively compact in $\mathcal{P}^p(\Omega)$.

Second step. Next, we check that $\overline{P}_n = \overline{P}'_n \circ (\xi, B, W, \mu, \Lambda, \hat{X}^n)^{-1}$ are relatively compact and have the same limits as \overline{P}'_n . This will follow essentially from the fact that $\hat{x}^n \rightarrow x$ as $n \rightarrow \infty$ uniformly on compact subsets of \mathcal{C}^d . Indeed, for $t \in [t_i^n, t_{i+1}^n]$, the definition of \hat{x}^n implies

$$|\hat{x}_t^n - x_t| \leq |\hat{x}_{t_i^n}^n - x_{t_i^n}^n| + |x_{t_i^n}^n - x_t| \leq |x_{t_i^n}^n - x_{t_{i-1}^n}^n| + |x_{t_{i-1}^n}^n - x_t|.$$

Since $|t - t_{i-1}^n| \leq 2 \cdot 2^{-n}T$ for $t \in [t_i^n, t_{i+1}^n]$, we get

$$\|\hat{x}^n - x\|_T \leq 2 \sup_{|t-s| \leq 2^{1-n}T} |x_t - x_s| \quad \forall x \in \mathcal{C}^d.$$

If $K \subset \mathcal{C}^d$ is compact, then it is equicontinuous by Arzelà–Ascoli, and the above implies $\sup_{x \in K} \|\hat{x}^n - x\|_T \rightarrow 0$. With this uniform convergence in hand, we check as follows that \overline{P}_n has the same limiting behavior as \overline{P}'_n . By Prohorov’s theorem, for each $\varepsilon > 0$ there exists a compact set $K_\varepsilon \subset \mathcal{C}^d$ such that $\mathbb{E}^{\overline{P}'_n}[\|X\|_T^p 1_{\{X \in K_\varepsilon^c\}}] \leq \varepsilon$ for each n . Using the obvious coupling and the fact that $\|\hat{x}^n\|_T \leq \|x\|_T$ for all $x \in \mathcal{C}^d$,

$$\ell_{\Omega,p}(\overline{P}_n, \overline{P}'_n) \leq \mathbb{E}^{\overline{P}'_n}[\|X - \hat{X}^n\|_T^p]^{1/p} \leq 2\varepsilon^{1/p} + \sup_{x \in K_\varepsilon} \|\hat{x}^n - x\|_T.$$

Send $n \rightarrow \infty$ and then $\varepsilon \downarrow 0$.

Third step. It remains to check that any limit point \overline{P} of \overline{P}_n (and thus of \overline{P}'_n) satisfies the required properties. Note first that (B, μ) , ξ and W are independent under \overline{P} , since $\bar{\mu}^n(B)$ is B -measurable and since B, ξ and W are independent under P_n . Moreover, (B, W) is an $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ Wiener process (of dimension $m_0 + m$) under \overline{P} since it is under P_n . In particular, $\rho := \overline{P} \circ (\xi, B, W, \mu)^{-1} \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}(\mathcal{X})]$. Since $(\bar{\mu}^n(B))^x = \mu^n(B)$, the canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the state equation (3.3) under \overline{P}'_n for each n . Hence, it follows from the results of Kurtz and Protter [36] that (3.3) holds under the limiting measure \overline{P} as well.

We now check that $\mu = \overline{P}((W, \Lambda, X) \in \cdot | \mathcal{F}_T^{B, \mu})$. Let \overline{P}_{n_k} be a subsequence converging to \overline{P} . Fix $n_0 \in \mathbb{N}$ and $S \in \mathcal{G}_T^{n_0}$, and let $\psi : \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ and $\phi : \mathcal{X} \rightarrow \mathbb{R}$ be bounded and continuous. Then, since $\bar{\mu}^n = P_n((W, \Lambda, X) \in \cdot | \mathcal{G}_T^n)$ and $\mathcal{G}_T^{n_0} \subset \mathcal{G}_T^n$ for $n \geq n_0$, we compute (using Lemma A.4 to handle the indicator function)

$$\begin{aligned} \mathbb{E}^{\overline{P}}[1_S(B)\psi(\mu)\phi(W, \Lambda, X)] &= \lim_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}}[1_S(B)\psi(\bar{\mu}^{n_k})\phi(W, \Lambda, X)] \\ &= \lim_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}}\left[1_S(B)\psi(\bar{\mu}^{n_k}) \int \phi d\bar{\mu}^{n_k}\right] \\ &= \mathbb{E}^{\overline{P}}\left[1_S(B)\psi(\mu) \int \phi d\mu\right]. \end{aligned}$$

Conclude by noting that $\sigma(\bigcup_{n=1}^\infty \mathcal{G}_T^n) = \sigma(B)$.

Conclusion. We have checked (1) and (3) in Definition 3.4. Concerning (2), we already know from the beginning of the second step that $\rho = \bar{P} \circ (\xi, B, W, \mu)^{-1} \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$. It thus remains to prove that $Q = P \circ (\xi, B, W, \mu, \Lambda)^{-1}$ is in $\mathcal{A}(\rho)$ (that is the relaxed control is admissible). This follows from the more general Lemma 3.7 right below. \square

The definition of MFG pre-solution requires that ρ is compatible with \mathcal{W}_λ [in the sense of point (1) in Section 3.2], but also the admissibility $P \in \mathcal{RA}(\rho)$ requires that $P \circ (\xi, B, W, \mu, \Lambda)^{-1}$ is compatible with ρ [in the sense of (2) in Section 3.2]. Because the latter compatibility does not behave well under limits, it will be crucial to have an alternative characterization of MFG pre-solutions which allows us to avoid directly checking admissibility. Namely, Lemma 3.7 below shows that admissibility essentially follows automatically from the fixed-point condition (3) of Definition 3.4. In fact, Lemma 3.7 is the main reason we work with the conditional law of (W, Λ, X) , and not just X .

LEMMA 3.7. *Let $P \in \mathcal{P}^p(\Omega)$ such that (B, W) is a Wiener process with respect to the filtration $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ under P , and define $\rho := P \circ (\xi, B, W, \mu)^{-1}$. Suppose that (1) and (3) in Definition 3.4 are satisfied and that $P(X_0 = \xi) = 1$. Then, for $P \circ \mu^{-1}$ -almost every $\nu \in \mathcal{P}^p(\mathcal{X})$, $(W_t)_{t \in [0, T]}$ is an $(\mathcal{F}_t^{W, \Lambda, X})_{t \in [0, T]}$ Wiener process under ν . Moreover, $Q = P \circ (\xi, B, W, \mu, \Lambda)^{-1}$ is in $\mathcal{A}(\rho)$.*

PROOF. *First step.* For $\nu \in \mathcal{P}(\mathcal{X})$, let $\nu^w = \nu \circ W^{-1} \in \mathcal{P}(\mathcal{C}^m)$. To prove the first claim, let $\phi_1 : \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ and $\phi_2 : \mathcal{C}^m \rightarrow \mathbb{R}$ be bounded and measurable. Then, since $P \circ W^{-1} = \mathcal{W}^m$ (with \mathbb{E} denoting expectation under P),

$$\mathbb{E}[\phi_1(\mu)] \int_{\mathcal{C}^m} \phi_2 d\mathcal{W}^m = \mathbb{E}[\phi_1(\mu)\phi_2(W)] = \mathbb{E}\left[\phi_1(\mu) \int_{\mathcal{C}^m} \phi_2 d\mu^w\right].$$

The first equality follows from (1) in Definition 3.4 and the second one from (3) in Definition 3.4.

This holds for all ϕ_1 , and thus $\int \phi_2 d\mu^w = \int \phi_2 d\mathcal{W}^m$ a.s. This holds for all ϕ_2 , and thus $\mu^w = \mathcal{W}^m$ a.s. Now fix $t \in [0, T]$. Suppose $\phi_1 : \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ is bounded and \mathcal{F}_t^μ -measurable, $\phi_2 : \mathcal{C}^m \rightarrow \mathbb{R}$ is bounded and $\sigma(W_s - W_t : s \in [t, T])$ -measurable, and $\phi_3 : \mathcal{X} \rightarrow \mathbb{R}$ is bounded and $\mathcal{F}_t^{W, \Lambda, X}$ -measurable. Then $\phi_2(W)$ and $(\phi_1(\mu), \phi_3(W, \Lambda, X))$ are P -independent [since W is a Wiener process with respect to $(\mathcal{F}_s^{\xi, B, W, \mu, \Lambda})_{s \in [0, T]}$], and so

$$\begin{aligned} &\mathbb{E}\left[\phi_1(\mu) \int_{\mathcal{X}} \phi_3 d\mu\right] \int_{\mathcal{C}^m} \phi_2 d\mathcal{W}^m \\ &= \mathbb{E}[\phi_1(\mu)\phi_3(W, \Lambda, X)] \int_{\mathcal{C}^m} \phi_2 d\mathcal{W}^m \end{aligned}$$

$$\begin{aligned} &= \mathbb{E}[\phi_1(\mu)\phi_2(W)\phi_3(W, \Lambda, X)] \\ &= \mathbb{E}\left[\phi_1(\mu) \int_{\mathcal{X}} \phi_2(w)\phi_3(w, q, x)\mu(dw, dq, dx)\right], \end{aligned}$$

the first and third equalities following from (3) in Definition 3.4. This holds for all ϕ_1 , and thus

$$\begin{aligned} &\int_{\mathcal{C}^m} \phi_2 d\mathcal{W}^m \int_{\mathcal{X}} \phi_3(w, q, x)\mu(dw, dq, dx) \\ &= \int_{\mathcal{X}} \phi_2(w)\phi_3(w, q, x)\mu(dw, dq, dx) \quad \text{a.s.} \end{aligned}$$

This holds for all ϕ_2 and ϕ_3 , and thus it holds P -a.s. that $\sigma(W_s - W_t : s \in [0, T])$ and $\mathcal{F}_t^{W, \Lambda, X}$ are independent under almost every realization of μ .

Second step. We now prove that Q is in $\mathcal{A}(\rho)$ [notice that, by assumption, (B, W) is a Wiener process with respect to the filtration $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ under P]. Fix $t \in [0, T]$. Let $\phi_t : \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$ be $\mathcal{F}_t^{\Lambda, X}$ -measurable, let $\phi_t^w : \mathcal{C}^m \rightarrow \mathbb{R}$ be \mathcal{F}_t^W -measurable, let $\phi_{t+}^w : \mathcal{C}^m \rightarrow \mathbb{R}$ be $\sigma(W_s - W_t : s \in [t, T])$ -measurable, let $\psi_T : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_T^{B, \mu}$ -measurable, and let $\psi_t : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_t^{B, \mu}$ -measurable. Assume all of these functions are bounded. We first compute

$$\begin{aligned} &\mathbb{E}[\psi_T(B, \mu)\phi_{t+}^w(W)\psi_t(B, \mu)\phi_t^w(W)] \\ &= \mathbb{E}[\psi_T(B, \mu)\psi_t(B, \mu)]\mathbb{E}[\phi_{t+}^w(W)]\mathbb{E}[\phi_t^w(W)] \\ &= \mathbb{E}[\mathbb{E}[\psi_T(B, \mu)|\mathcal{F}_t^{B, \mu}]\psi_t(B, \mu)]\mathbb{E}[\phi_{t+}^w(W)]\mathbb{E}[\phi_t^w(W)] \\ &= \mathbb{E}[\mathbb{E}[\psi_T(B, \mu)|\mathcal{F}_t^{B, \mu}]\phi_{t+}^w(W)\psi_t(B, \mu)]\mathbb{E}[\phi_t^w(W)], \end{aligned}$$

the first and third lines following from (1) in Definition 3.4. This shows that

$$(3.7) \quad \mathbb{E}[\psi_T(B, \mu)\phi_{t+}^w(W)|\mathcal{F}_t^{B, W, \mu}] = \mathbb{E}[\psi_T(B, \mu)|\mathcal{F}_t^{B, \mu}] \int_{\mathcal{C}^m} \phi_{t+}^w d\mathcal{W}^m.$$

On the other hand, the first result of this lemma implies that $(W_t)_{t \in [0, T]}$ is an $(\mathcal{F}_t^{W, \Lambda, X})_{t \in [0, T]}$ Wiener process under almost every realization of μ , so that

$$(3.8) \quad \begin{aligned} &\int_{\mathcal{X}} \phi_t(q, x)\phi_{t+}^w(w)\phi_t^w(w)\mu(dw, dq, dx) \\ &= \int_{\mathcal{X}} \phi_t(q, x)\phi_t^w(w)\mu(dw, dq, dx) \int_{\mathcal{C}^m} \phi_{t+}^w d\mathcal{W}^m \quad \text{a.s.} \end{aligned}$$

By (3.2), note also that $\int_{\mathcal{X}} \phi_t(q, x)\phi_t^w(w)\mu(dw, dq, dx)$ is $\mathcal{F}_t^{B, \mu}$ -measurable, since $\phi_t(\Lambda, X)\phi_t^w(W)$ is $\mathcal{F}_t^{W, \Lambda, X}$ -measurable. Putting it together (see right after

the computations for more explanations):

$$\begin{aligned}
 & \mathbb{E}[\phi_t(\Lambda, X)\psi_T(B, \mu)\phi_{t+}^w(W)\psi_t(B, \mu)\phi_t^w(W)] \\
 &= \mathbb{E}\left[\left(\int_{\mathcal{X}} \phi_t(q, x)\phi_{t+}^w(w)\phi_t^w(w)\mu(dw, dq, dx)\right)\psi_T(B, \mu)\psi_t(B, \mu)\right] \\
 &= \mathbb{E}\left[\left(\int_{\mathcal{X}} \phi_t(q, x)\phi_t^w(w)\mu(dw, dq, dx)\right)\psi_T(B, \mu)\psi_t(B, \mu)\right] \\
 &\quad \times \int \phi_{t+}^w d\mathcal{W}^m \\
 &= \mathbb{E}\left[\left(\int_{\mathcal{X}} \phi_t(q, x)\phi_t^w(w)\mu(dw, dq, dx)\right)\mathbb{E}[\psi_T(B, \mu)|\mathcal{F}_t^{B,\mu}]\psi_t(B, \mu)\right] \\
 &\quad \times \int_{\mathcal{C}^m} \phi_{t+}^w d\mathcal{W}^m \\
 &= \mathbb{E}[\phi_t(\Lambda, X)\phi_t^w(W)\mathbb{E}[\psi_T(B, \mu)|\mathcal{F}_t^{B,\mu}]\psi_t(B, \mu)] \int_{\mathcal{C}^m} \phi_{t+}^w d\mathcal{W}^m \\
 &= \mathbb{E}[\mathbb{E}[\phi_t(\Lambda, X)|\mathcal{F}_t^{B,W,\mu}]\mathbb{E}[\psi_T(B, \mu)|\mathcal{F}_t^{B,\mu}]\psi_t(B, \mu)\phi_t^w(W)] \\
 &\quad \times \int_{\mathcal{C}^m} \phi_{t+}^w d\mathcal{W}^m \\
 &= \mathbb{E}[\mathbb{E}[\phi_t(\Lambda, X)|\mathcal{F}_t^{B,W,\mu}]\mathbb{E}[\psi_T(B, \mu)\phi_{t+}^w(W)|\mathcal{F}_t^{B,W,\mu}]\psi_t(B, \mu)\phi_t^w(W)],
 \end{aligned}$$

the first equality following from (3) in Definition 3.4, the second one from (3.8), the third one from the fact that $\int \phi_t(q, x)\phi_t^w(w)\mu(dw, dq, dx)$ is $\mathcal{F}_t^{B,\mu}$ -measurable, the fourth one from (3) in Definition 3.4 and the last one from (3.7).

Replacing $\phi_t^w(W)$ with $\phi_t^w(W)\psi_t^w(W)$, where both ϕ_t^w and ψ_t^w are \mathcal{F}_t^W -measurable, we see that

$$\begin{aligned}
 & \mathbb{E}[\phi_t(\Lambda, X)\psi_T(B, \mu)\phi_{t+}^w(W)\phi_t^w(W)|\mathcal{F}_t^{B,W,\mu}] \\
 &= \mathbb{E}[\phi_t(\Lambda, X)|\mathcal{F}_t^{B,W,\mu}]\mathbb{E}[\psi_T(B, \mu)\phi_{t+}^w(W)\phi_t^w(W)|\mathcal{F}_t^{B,W,\mu}].
 \end{aligned}$$

Since random variables of the form $\phi_t^w(W)\phi_{t+}^w(W)$ generate \mathcal{F}_T^W , this shows that $\mathcal{F}_t^{\Lambda, X}$ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_t^{B,W,\mu}$.

Last step. It now remains to prove that \mathcal{F}_t^Λ is conditionally independent of $\mathcal{F}_T^{\xi, B, W, \mu}$ given $\mathcal{F}_t^{\xi, B, W, \mu}$, which is slightly different from the result of the previous step. To do so, we use the fact that $P(X_0 = \xi) = 1$. Let $\phi_t : \mathcal{V} \rightarrow \mathbb{R}$ be \mathcal{F}_t^Λ -measurable, $\psi_t : \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_t^{B,W,\mu}$ -measurable, $\psi_T : \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_T^{B,W,\mu}$ -measurable and $\zeta_0 : \mathbb{R} \rightarrow \mathbb{R}$ be Borel measurable. Assume all of these functions are bounded. From the previous step, we

deduce that

$$\begin{aligned} & \mathbb{E}[\phi_t(\Lambda)\psi_T(B, W, \mu)\psi_t(B, W, \mu)\zeta_0(\xi)] \\ &= \mathbb{E}[\phi_t(\Lambda)\zeta_0(X_0)\psi_T(B, W, \mu)\psi_t(B, W, \mu)] \\ &= \mathbb{E}[\mathbb{E}[\phi_t(\Lambda)\zeta_0(X_0)|\mathcal{F}_t^{B,W,\mu}]\mathbb{E}[\psi_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}]\psi_t(B, W, \mu)] \\ &= \mathbb{E}[\phi_t(\Lambda)\zeta_0(X_0)\mathbb{E}[\psi_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}]\psi_t(B, W, \mu)] \\ &= \mathbb{E}[\phi_t(\Lambda)\mathbb{E}[\psi_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}]\psi_t(B, W, \mu)\zeta_0(\xi)], \end{aligned}$$

the second equality following from the conditional independence of $\mathcal{F}_t^{\Lambda, X}$ and $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_t^{B,W,\mu}$. In order to complete the proof, notice that $\mathbb{E}[\psi_T(B, W, \mu)|\mathcal{F}_t^{B,W,\mu}] = \mathbb{E}[\psi_T(B, W, \mu)|\mathcal{F}_t^{\xi, B, W, \mu}]$ since ξ and (B, W, μ) are independent under P [see (1) in Definition 3.4]. Therefore, for another bounded Borel measurable function $\zeta'_0 : \mathbb{R} \rightarrow \mathbb{R}$, we get

$$\begin{aligned} & \mathbb{E}[\phi_t(\Lambda)\psi_T(B, W, \mu)\psi_t(B, W, \mu)\zeta_0(\xi)\zeta'_0(\xi)] \\ &= \mathbb{E}[\phi_t(\Lambda)\mathbb{E}[\psi_T(B, W, \mu)|\mathcal{F}_t^{\xi, B, W, \mu}]\psi_t(B, W, \mu)\zeta_0(\xi)\zeta'_0(\xi)] \\ &= \mathbb{E}[\phi_t(\Lambda)\mathbb{E}[\zeta'_0(\xi)\psi_T(B, W, \mu)|\mathcal{F}_t^{\xi, B, W, \mu}]\psi_t(B, W, \mu)\zeta_0(\xi)], \end{aligned}$$

which proves that \mathcal{F}_t^Λ and $\mathcal{F}_T^{\xi, B, W, \mu}$ are conditionally independent given $\mathcal{F}_t^{\xi, B, W, \mu}$. \square

3.4. *Existence of a MFG solution under Assumption B.* The goal of this section is to prove that the limit points constructed in the previous paragraph are not only MFG pre-solutions but are weak MFG solutions.

THEOREM 3.8. *Assume that Assumption B holds and keep the notation of Lemma 3.6. Then every limit point of \bar{P}_n is a weak MFG solution with weak control.*

Generally speaking, it remains to show that any limit point of the sequence of Lemma 3.6 is optimal for the corresponding control problem.

LEMMA 3.9. *Assume that a MFG pre-solution P satisfies $P \in \mathcal{RA}^*(\rho)$, with ρ given by $\rho := P \circ (\xi, B, W, \mu)^{-1}$, then $(\Omega, (\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$ is a weak MFG solution with weak control.*

PROOF. The proof is quite straightforward since the pre-solution properties of P guarantee that the canonical process under P satisfy (1)–(4) and (6) in Definition 3.1. Condition (3) of Definition 3.1 uses $\sigma(\Lambda_s : s \leq t)$, whereas the notion of compatibility in the definition of MFG pre-solutions uses the canonical filtration \mathcal{F}_t^Λ defined by (2.3), but this is no cause for concern in light of the discussion

following (2.3). The additional condition $P \in \mathcal{RA}^*(\rho)$ permits to verify (5) in Definition 3.1 by transferring any $(\Omega', (\mathcal{F}'_t)_{t \in [0, T]}, P', B', W', \nu, \Lambda', X')$ as in (5) onto the canonical space. \square

3.4.1. *Strategy.* In order to check the condition $P \in \mathcal{RA}^*(\rho)$ in Lemma 3.9, the idea is to approximate any alternative MFG control by a sequence of particularly well-behaved controls for the discretized game. The crucial technical device is Lemma 3.11, but we defer its proof to the Appendix. The following definition is rather specific to the setting of compact control space A (we assume that Assumption B holds throughout the section), but it will return in a more general form in Section 3.5:

DEFINITION 3.10. A function $\phi : \Omega_0 \times \mathcal{P}^P(\mathcal{X}) \rightarrow \mathcal{V}$ is said to be *adapted* if $\phi^{-1}(C) \in \mathcal{F}_t^{\xi, B, W, \mu}$ for each $C \in \mathcal{F}_t^\Lambda$ and $t \in [0, T]$. For $\rho \in \mathcal{P}_c^P[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^P(\mathcal{X})]$ [i.e., (ξ, B, W) has law \mathcal{W}_λ under ρ and B and W are independent $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -Wiener processes under ρ], let $\mathcal{A}_a(\rho)$ denote the set of measures of the form

$$(3.9) \quad \rho(d\omega, d\nu)\delta_{\phi(\omega, \nu)}(dq) = \rho \circ (\xi, B, W, \mu, \phi(\xi, B, W, \mu))^{-1},$$

where ϕ is adapted and *continuous*.

LEMMA 3.11. For each $\rho \in \mathcal{P}_c^P[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^P(\mathcal{X})]$, $\mathcal{A}_a(\rho)$ is a dense subset of $\mathcal{A}(\rho)$.

We also need continuity lemmas, the proofs of which are deferred to the end of the subsection. Notice that these lemmas are stated under Assumption A, not B.

LEMMA 3.12. Suppose a set $K \subset \mathcal{P}_c^P[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^P(\mathcal{X}) \times \mathcal{V}]$ satisfies

$$\sup_{P \in K} \mathbb{E}^P \left[\int_{\mathcal{C}^d} \|x\|_T^{p'} \mu^x(dx) + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \right] < \infty.$$

Under Assumption A, the map $\mathcal{R} : K \rightarrow \mathcal{P}^P(\Omega)$ [that maps $Q \in K$ to the law of the solution $(\xi, B, W, \mu, \Lambda, X)$ of (3.3) when $(\xi, B, W, \mu, \Lambda)$ has law Q] is continuous.

LEMMA 3.13. Under Assumption A, the expected reward functional $J : \mathcal{P}^P(\Omega) \rightarrow \mathbb{R}$ given by (3.4) is upper semicontinuous. If also A is compact, then J is continuous.

LEMMA 3.14. Define $\Pi_n : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega)$ by

$$\Pi_n(P) := P \circ (\xi, B, W, \mu, \Lambda, \hat{X}^n)^{-1}.$$

[See (2.6) for the definition of \hat{X}^n .] If $P_n \rightarrow P$ in $\mathcal{P}^P(\Omega)$, then $\Pi_n(P_n) \rightarrow P$ in $\mathcal{P}^P(\Omega)$.

3.4.2. *Proof of Theorem 3.8.* Let $\mu^n, \bar{\mu}^n, P_n$ and \bar{P}_n be as in Lemma 3.6, and let \bar{P} denote any limit point. Relabel the subsequence, and assume that \bar{P}_n itself converges. Let $\rho := \bar{P} \circ (\xi, B, W, \mu)^{-1}$. By Lemma 3.6, $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ and $\bar{P} \in \mathcal{RA}(\rho)$ is a MFG pre-solution, and it remains only to show that \bar{P} is optimal, or $\bar{P} \in \mathcal{RA}^*(\rho)$. Fix $P^* \in \mathcal{RA}(\rho)$ arbitrarily with $J(P^*) > -\infty$. Let

$$Q^* := P^* \circ (\xi, B, W, \mu, \Lambda)^{-1}.$$

By Lemma 3.11, we may find a sequence of $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -adapted continuous functions $\phi_k : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{V}$ such that

$$Q^* = \lim_{k \rightarrow \infty} Q^k \quad \text{where } Q^k := \rho \circ (\xi, B, W, \mu, \phi_k(\xi, B, W, \mu))^{-1}.$$

Define $Q_n^k \in \mathcal{A}_f$ (see Section 2.4.4 for the definition of \mathcal{A}_f) by

$$Q_n^k := \mathcal{W}_\lambda \circ (\xi, B, W, \phi_k(\xi, B, W, \bar{\mu}^n(B)))^{-1}.$$

Note that $\bar{P}_n \rightarrow \bar{P}$ implies

$$\rho = \lim_{n \rightarrow \infty} \bar{P}_n \circ (\xi, B, W, \mu)^{-1} = \lim_{n \rightarrow \infty} \mathcal{W}_\lambda \circ (\xi, B, W, \bar{\mu}^n(B))^{-1},$$

where the second equality comes from the definition of \bar{P}_n in Lemma 3.6. Since ϕ_k is continuous with respect to μ , we deduce from Lemma A.4 (that permits to handle the possible discontinuity of ϕ_k in the other variables)

$$\begin{aligned} & \lim_{n \rightarrow \infty} Q_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda)^{-1} \\ (3.10) \quad & = \lim_{n \rightarrow \infty} \mathcal{W}_\lambda \circ (\xi, B, W, \bar{\mu}^n(B), \phi_k(\xi, B, W, \bar{\mu}^n(B)))^{-1} \\ & = Q^k. \end{aligned}$$

Now let $P_n^k := \mathcal{R}_f^n(\mu^n, Q_n^k)$. Since P_n is optimal for $J_f(\mu^n, \cdot)$,

$$J_f(\mu^n, P_n^k) \leq J_f(\mu^n, P_n).$$

Since A is compact, Lemma 3.13 assures us that J is continuous, and so

$$\lim_{n \rightarrow \infty} J_f(\mu^n, P_n) = \lim_{n \rightarrow \infty} \mathbb{E}^{P_n}[\Gamma(\mu^n(B), \Lambda, X)] = \lim_{n \rightarrow \infty} J(\bar{P}_n) = J(\bar{P}),$$

where the second equality follows simply from the definition of J . We will complete the proof by showing that, on the other hand,

$$(3.11) \quad J(P^*) = \lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} J_f(\mu^n, P_n^k),$$

and both limits exist; indeed, this implies $J(\bar{P}) \geq J(P^*)$, completing the proof since $P^* \in \mathcal{RA}(\rho)$ was arbitrary. Define Π_n as in Lemma 3.14. The trick is to

notice (just applying the basic definition of the different objects) that

$$P_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1} = \Pi_n(\mathcal{R}(Q_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda)^{-1})),$$

$$J_f(\mu^n, P_n^k) = J(P_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1}).$$

Now note that $P_n^k \circ (\mu^x)^{-1} = \mathcal{W}_\lambda \circ (\mu^n)^{-1} = \bar{P}_n \circ (\mu^x)^{-1}$, and thus by (3.6) we have

$$\sup_n \mathbb{E}^{P_n^k} \int_{\mathcal{C}^d} \|y\|_T^{p'} \mu^x(dy) < \infty.$$

Since also A is compact, we may apply Lemma 3.12 (continuity of \mathcal{R}), along with Lemma 3.14 and (3.10), to get

$$\lim_{n \rightarrow \infty} P_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1}$$

$$= \lim_{n \rightarrow \infty} \Pi_n(\mathcal{R}(Q_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda)^{-1})) = \mathcal{R}(Q^k).$$

Thus, again using continuity of \mathcal{R} ,

$$\lim_{k \rightarrow \infty} \lim_{n \rightarrow \infty} P_n^k \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1} = \mathcal{R}(Q^*) = P^*.$$

Finally, (3.11) follows from continuity of J .

3.4.3. *Proof of Lemma 3.12.* Let $Q_n \rightarrow Q$ in K . Note that $\mathcal{R}(Q_n) \circ (X_0, B, W, \mu, \Lambda)^{-1} = Q_n$ are relatively compact in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$. It can be shown using Aldous' criterion (see Proposition B.1) that this implies that $\mathcal{R}(Q_n) \circ X^{-1}$ are relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$, and thus $\mathcal{R}(Q_n)$ are relatively compact in $\mathcal{P}^p(\Omega)$. Let P be any limit point, so $\mathcal{R}(Q_{n_k}) \rightarrow P$ for some n_k . Then

$$P \circ (\xi, B, W, \mu, \Lambda)^{-1} = \lim_{k \rightarrow \infty} \mathcal{R}(Q_{n_k}) \circ (\xi, B, W, \mu, \Lambda)^{-1} = \lim_{k \rightarrow \infty} Q_{n_k} = Q.$$

It follows from the results of Kurtz and Protter [36] that the canonical processes verify the SDE (3.3) under P . Hence, $P = \mathcal{R}(Q)$.

3.4.4. *Proof of Lemma 3.13.* Since f and g are continuous in (x, μ, a) , the upper bounds of f and g [which grow in order p in (x, μ)] along with Lemma A.5 imply both that Γ is upper semicontinuous and then also that J is upper semicontinuous from $\mathcal{P}^p(\Omega)$ to \mathbb{R} . If A is compact, then the same p -order upper bounds of f and g hold for the negative parts as well, and the second part of Lemma A.5 provides the claimed continuity.

3.4.5. *Proof of Lemma 3.14.* This was essentially already proven in the second step of the proof of Lemma 3.6. Note that

$$\ell_{\Omega,p}(\Pi_n(P_n), P) \leq \ell_{\Omega,p}(P_n, P) + \ell_{\Omega,p}(P_n, \Pi_n(P_n)).$$

The first term tends to zero by assumption. Fix $\varepsilon > 0$. Since $\{P_n : n \geq 1\}$ is relatively compact in $\mathcal{P}^p(\Omega)$, by Prohorov’s theorem there exists a compact set $K \subset \mathcal{C}^d$ such that $\mathbb{E}^{P_n}[\|X\|_T^p 1_{\{X \notin K\}}] \leq \varepsilon$ for all n . Use the obvious coupling and the fact that $\|\hat{x}^n\|_T \leq \|x\|_T$ for all $x \in \mathcal{C}^d$ to get

$$\ell_{\Omega,p}(P_n, \Pi_n(P_n)) \leq \mathbb{E}^{P_n}[\|X - \hat{X}^n\|_T^p]^{1/p} \leq 2\varepsilon^{1/p} + \sup_{x \in K} \|x - \hat{x}^n\|_T.$$

We saw in the second step of the proof of Lemma 3.6 that $\hat{x}^n \rightarrow x$ as $n \rightarrow \infty$ uniformly on compact subsets of \mathcal{C}^d , and so the proof is complete.

3.5. *Unbounded coefficients.* Finally, with existence in hand for bounded state coefficients (b, σ, σ_0) and compact control space A , we turn to the general case. The goal is thus to complete the proof of Theorem 3.2 under Assumption A instead of B.

The idea of the proof is to approximate the data (b, σ, σ_0, A) by data satisfying Assumption B. Let $(b^n, \sigma^n, \sigma_0^n)$ denote the projection of (b, σ, σ_0) into the ball centered at the origin with radius n in $\mathbb{R}^d \times \mathbb{R}^{d \times m} \times \mathbb{R}^{d \times m_0}$, respectively. Let A_n denote the intersection of A with the ball centered at the origin with radius n . For sufficiently large n_0 , A_n is nonempty and compact for all $n \geq n_0$, and thus we will always assume $n \geq n_0$ in what follows. Note that the data $(b^n, \sigma^n, \sigma_0^n, f, g, A_n)$ satisfy Assumption B. Moreover, (A.4) and (A.5) hold for each n with the same constants c_1, c_2, c_3 ; this implies that Lemma 2.4 holds with the same constant c_4 for each set of data, that is, independent of n .

Define \mathcal{V}_n as before in terms of A_n , but now view it as a subset of \mathcal{V} . That is, $\mathcal{V}_n := \{q \in \mathcal{V} : q([0, T] \times A_n^c) = 0\}$. Naturally, define $\mathcal{A}_n(\rho)$ to be the set of admissible controls with values in A_n :

$$(3.12) \quad \mathcal{A}_n(\rho) := \{Q \in \mathcal{A}(\rho) : Q(\Lambda \in \mathcal{V}_n) = 1\}.$$

Finally, define $\mathcal{R}_n(Q)$ to be the unique element P of $\mathcal{P}(\Omega)$ such that $P \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q$ and the canonical processes verify the SDE

$$(3.13) \quad \begin{aligned} X_t = X_0 &+ \int_0^t ds \int_A \Lambda_s(da) b^n(s, X_s, \mu_s^x, a) + \int_0^t \sigma^n(s, X_s, \mu_s^x) dW_s \\ &+ \int_0^t \sigma_0^n(s, X_s, \mu_s^x) dB_s. \end{aligned}$$

Define naturally

$$\mathcal{R}_n \mathcal{A}_n^*(\rho) := \arg \max_{P \in \mathcal{R}_n \mathcal{A}_n(\rho)} J(P).$$

By Theorem 3.8, there exists for each n a MFG solution corresponding to the n th truncation of the data. In the present notation, this means there exist $\rho_n \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ and $P_n \in \mathcal{R}_n \mathcal{A}_n^*(\rho_n)$ such that

$$(3.14) \quad \mu = P_n((W, \Lambda, X) \in \cdot | \mathcal{F}_T^{B, \mu}), \quad P^n\text{-a.s.}$$

Once again, the strategy of the proof is to show first that P_n are relatively compact and then that each limit point is a MFG solution.

3.5.1. *Relative compactness.* We start with the following.

LEMMA 3.15. *The measures P_n are relatively compact in $\mathcal{P}^p(\Omega)$. Moreover,*

$$(3.15) \quad \begin{aligned} \sup_n \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt < \infty, \\ \mathbb{E}^{P_n} \int_{\mathcal{C}^d} \|y\|_T^{p'} \mu^x(dy) = \sup_n \mathbb{E}^{P_n} \|X\|_T^{p'} < \infty. \end{aligned}$$

PROOF. Noting that the coefficients $(b^n, \sigma^n, \sigma_0^n)$ satisfy (A.1)–(A.5) with the same constants (independent of n), Lemma 2.4 and (3.14) imply

$$(3.16) \quad \mathbb{E}^{P_n} \int_{\mathcal{C}^d} \|x\|_T^p \mu^x(dx) = \mathbb{E}^{P_n} \|X\|_T^p \leq c_4 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right).$$

Fix $a_0 \in A_{n_0}$. Let R_n denote the unique element of $\mathcal{R}_n \mathcal{A}_n(\rho_n)$ satisfying $R_n(\Lambda_t = \delta_{a_0}$ for a.e. t) = 1. That is R_n is the law of the solution of the state equation arising from the constant control equal to a_0 , in the n th truncation. The first part of Lemma 2.4 implies

$$(3.17) \quad \mathbb{E}^{R_n} \|X\|_T^p \leq c_4 \left(1 + \mathbb{E}^{R_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) + T|a_0|^p \right).$$

Noting that $R_n \circ \mu^{-1} = P_n \circ \mu^{-1}$, we combine (3.17) with (3.16) to get

$$(3.18) \quad \mathbb{E}^{R_n} \|X\|_T^p \leq C_0 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right),$$

where $C_0 > 0$ depends only on c_4, T , and $|a_0|^p$. Use the optimality of P_n , the lower bounds on f and g , and then (3.16) and (3.18) to get

$$(3.19) \quad \begin{aligned} J(P_n) &\geq J(R_n) \\ &\geq -c_2(T + 1) \left(1 + \mathbb{E}^{R_n} \|X\|_T^p + \mathbb{E}^{R_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) + |a_0|^{p'} \right) \\ &\geq -C_1 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right), \end{aligned}$$

where $C_1 > 0$ depends only on $c_2, c_4, T, |a_0|^{p'}$, and C_0 . On the other hand, we may use the upper bounds on f and g along with (3.16) to get

$$\begin{aligned}
 J(P_n) &\leq c_2(T + 1) \left(1 + \mathbb{E}^{P_n} \|X\|_T^p + \mathbb{E}^{P_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) \right) \\
 (3.20) \quad &- c_3 \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \\
 &\leq C_2 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right) - c_3 \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt,
 \end{aligned}$$

where $C_2 > 0$ depends only on c_2, c_3, c_4 , and T . Combining (3.19) and (3.20) and rearranging, we find two constants, $\kappa_1 \in \mathbb{R}$ and $\kappa_2 > 0$, such that

$$\mathbb{E}^{P_n} \int_0^T \int_A (|a|^{p'} + \kappa_1 |a|^p) \Lambda_t(da) dt \leq \kappa_2.$$

[Note that $\mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty$ for each n .] These constants are independent of n , and the first bound in (3.15) follows from the fact that $p' > p$. Combined with Lemma 2.4, this implies the second bound in (3.15).

To show that P_n are relatively compact, we check that each of the sets of marginals is relatively compact; see Lemma A.3. Compactness of $P_n \circ (B, W)^{-1}$ is obvious. Moreover, by (3.15),

$$\sup_n \mathbb{E}^{P_n} \left[\|W\|_T^{p'} + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt + \|X\|_T^{p'} \right] < \infty.$$

Aldous' criterion (Proposition B.1) shows that $P_n \circ (\Lambda, X)^{-1}$ are relatively compact. The mean measures of $P_n \circ \mu^{-1}$ are $P_n \circ (W, \Lambda, X)^{-1}$, which we have shown are relatively compact. Hence, by Proposition A.2, $P_n \circ \mu^{-1}$ are relatively compact in $\mathcal{P}^p(\mathcal{P}^p(\mathcal{X}))$. \square

3.5.2. *Limit points.* Now that we know P_n are relatively compact, we may fix $P \in \mathcal{P}^p(\Omega)$ and a subsequence n_k such that $P_{n_k} \rightarrow P$ in $\mathcal{P}^p(\Omega)$. Define $\rho := P \circ (\xi, B, W, \mu)^{-1}$, and note that $\rho_{n_k} \rightarrow \rho$.

LEMMA 3.16. *The limit point P is a MFG pre-solution and satisfies*

$$\mathbb{E}^P \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \leq \liminf_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt < \infty.$$

PROOF. Fatou's lemma and the first bound in (3.15) imply the stated inequality. We now check (1), (2) and (3) in Definition 3.4. Since $(B, \mu), \xi$ and W are independent under P_n , the same is true under the limit P , which gives (1). We now check (2). The strategy is to apply Lemma 3.7. By passage to the limit, it is well checked that (B, W) is a Wiener process with respect to the filtration

$(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ under P (which implies in particular that $\rho \in \mathcal{P}_c^P[(\Omega_0 \times \mathcal{P}^P(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}]$). Moreover, it must also hold $P(X_0 = \xi) = 1$. Therefore, in order to prove (2), it sufficient to check (3) and to check that the state equation (3.3) is satisfied under P .

We first check (3). If $\psi : \mathcal{C}^{m_0} \times \mathcal{P}^P(\mathcal{X}) \rightarrow \mathbb{R}$ and $\phi : \mathcal{X} \rightarrow \mathbb{R}$ are bounded and continuous, we have

$$\begin{aligned} \mathbb{E}^P[\psi(B, \mu)\phi(W, \Lambda, X)] &= \lim_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}}[\psi(B, \mu)\phi(W, \Lambda, X)] \\ &= \lim_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}}\left[\psi(B, \mu) \int_{\mathcal{X}} \phi d\mu\right] \\ &= \mathbb{E}^P\left[\psi(B, \mu) \int_{\mathcal{X}} \phi d\mu\right]. \end{aligned}$$

Thus $\mu = P((W, \Lambda, X) \in \cdot | \mathcal{F}_T^{B, \mu})$ a.s., which gives (3) in Definition 3.4. Now, to check that the state equation is satisfied, define processes $(Z_t^q)_{t \in [0, T]}$ on Ω by

$$Z_t^q := 1 + |X_t|^q + \left(\int_{\mathbb{R}^d} |y|^p \mu_t^x(dy)\right)^{q/p}, \quad q > 0.$$

Using the growth assumptions on b of (A.4), note that $b(t, y, v, a) \neq b^n(t, y, v, a)$ if and only if

$$(3.21) \quad n < |b(t, y, v, a)| \leq c_1 \left(1 + |y| + \left(\int_{\mathbb{R}^d} |z|^p v(dz)\right)^{1/p} + |a|\right),$$

so that

$$\begin{aligned} &\mathbb{E}^{P_n} \left| \int_0^t ds \int_A \Lambda_s(da) (b^n - b)(s, X_s, \mu_s^x, a) \right| \\ &\leq 2c_1 \mathbb{E}^{P_n} \int_0^t ds \int_A \Lambda_s(da) (Z_s^1 + |a|) 1_{\{c_1(Z_s^1 + |a|) > n\}}. \end{aligned}$$

By Lemma 3.15, this tends to zero as $n \rightarrow \infty$. Similarly, $\sigma(t, y, v) \neq \sigma^n(t, y, v)$ if and only if

$$(3.22) \quad n^2 < |\sigma(t, y, v)|^2 \leq c_1 \left(1 + |y|^{p_\sigma} + \left(\int_{\mathbb{R}^d} |z|^p v(dz)\right)^{p_\sigma/p}\right),$$

so that the Burkholder–Davis–Gundy inequality yields

$$\begin{aligned} &\mathbb{E}^{P_n} \left[\sup_{0 \leq t \leq T} \left| \int_0^t (\sigma^n - \sigma)(s, X_s, \mu_s^x) dW_s \right| \right] \\ &\leq 2c_1^{1/2} \mathbb{E}^{P_n} \left[\left(\int_0^T Z_s^{p_\sigma} 1_{\{c_1 Z_s^{p_\sigma} > n^2\}} ds \right)^{1/2} \right]. \end{aligned}$$

This tends to zero as well, as does $\mathbb{E}^{P_n} |\int_0^t (\sigma_0^n - \sigma_0)(s, X_s, \mu_s^x) dB_s|$. It follows that

$$0 = \lim_{n \rightarrow \infty} \mathbb{E}^{P_n} \sup_{0 \leq t \leq T} \left| X_t - X_0 - \int_0^t ds \int_A \Lambda_s(da) b(s, X_s, \mu_s^x, a) - \int_0^t \sigma(s, X_s, \mu_s^x) dW_s - \int_0^t \sigma_0(s, X_s, \mu_s^x) dB_s \right|.$$

Finally, combine this with the results of Kurtz and Protter [36] to conclude that the SDE (3.3) holds under P . \square

3.5.3. *Optimality.* It remains to show the limit point P in Lemma 3.15 is optimal. Generally speaking, the argument is as follows. Fix $P' \in \mathcal{RA}(\rho)$, with $\rho := P \circ (\xi, B, W, \mu)^{-1}$. If we can prove that there exist $P'_n \in \mathcal{R}_n \mathcal{A}_n(\rho_n)$ such that $J(P'_n) \rightarrow J(P')$, then, by optimality of P_n for each n , it holds that $J(P_n) \geq J(P'_n)$. Since J is upper semicontinuous by Lemma 3.13, we then get

$$J(P) \geq \limsup_{k \rightarrow \infty} J(P_{n_k}) \geq \lim_{k \rightarrow \infty} J(P'_{n_k}) = J(P').$$

Since P' was arbitrary, this implies that P is optimal, or $P \in \mathcal{RA}^*(\rho)$, which completes the proof of Theorem 3.2.

The goal is thus to prove the existence of the sequence $(P'_n)_{n \geq 1}$. For this, we need again to approximate general controls by adapted controls, as in Lemma 3.11. To this end, now that A is noncompact, we generalize the definition of the class $\mathcal{A}_a(\rho)$ in Definition 3.10: Let $\mathcal{A}_a(\rho)$ now denote the set of measures of the form

$$\rho(d\omega, d\nu) \delta_{\phi(\omega, \nu)}(dq) = \rho \circ (\xi, B, W, \mu, \phi(\xi, B, W, \mu))^{-1},$$

where $\phi : \Omega_0 \times \mathcal{P}^P(\mathcal{X}) \rightarrow \mathcal{V}_m$ is adapted and continuous and m is some positive integer [see (3.12) for the definition of \mathcal{V}_m]. In particular, a control $Q \in \mathcal{A}_a(\rho)$ satisfies $Q(\Lambda \in \mathcal{V}_m) = 1$ for some m and renders $(\Lambda_t)_{t \in [0, T]}$ (a.s.-) adapted to $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$. Note that when A is compact this definition specializes to the one provided before. The construction of $(P'_n)_{n \geq 1}$ then follows from the combination of the two next lemmas.

LEMMA 3.17. *For each $P' \in \mathcal{RA}(\rho)$ such that $J(P') > -\infty$, there exist $P'_n \in \mathcal{RA}_a(\rho)$ such that $J(P') = \lim_{n \rightarrow \infty} J(P'_n)$. [As usual $\mathcal{RA}_a(\rho)$ is the image of $\mathcal{A}_a(\rho)$ by \mathcal{R} .]*

LEMMA 3.18. *For each $P' \in \mathcal{RA}_a(\rho)$, there exist $P'_n \in \mathcal{R}_n \mathcal{A}_n(\rho_n)$ such that $J(P') = \lim_{n \rightarrow \infty} J(P'_n)$. [See (3.12) for the definition of $\mathcal{A}_n(\rho_n)$.]*

3.5.4. *Proof of Lemma 3.17. First step.* First, assume $P' \in \mathcal{RA}_m(\rho)$ for some fixed m , so trivially $J(P') > -\infty$. That is, $P'(\Lambda \in \mathcal{V}_m) = 1$. Write $P' = \mathcal{R}(Q')$, where $Q' \in \mathcal{A}_m(\rho)$. By Lemma 3.11, $\mathcal{A}_a(\rho)$ is dense in $\mathcal{A}_m(\rho)$, and there exist $Q'_n \in \mathcal{RA}_a(\rho)$ such that $Q'_n \rightarrow Q'$ in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$. Since A_m is compact, $J \circ \mathcal{R}$ is continuous on $\mathcal{A}_m(\rho)$ by Lemma 3.13, and $J(P') = \lim_{n \rightarrow \infty} J(\mathcal{R}(Q'_n))$.

Second step. Now assume $P' \in \mathcal{RA}(\rho)$ satisfies $J(P') > -\infty$. By the first step, it suffices to show that there exist $P'_n \in \mathcal{RA}_n(\rho)$ such that $J(P') = \lim_{n \rightarrow \infty} J(P'_n)$, since we just showed that each P'_n may be approximated by elements of $\mathcal{RA}_a(\rho)$.

First, the upper bounds of f and g imply

$$-\infty < J(P') \leq c_2(T + 1) \left(1 + \mathbb{E}^{P'} \|X\|_T^p + \mathbb{E}^{P'} \int_{\mathcal{C}^d} \|z\|_T^p \mu(dz) \right) - c_3 \mathbb{E}^{P'} \int_0^T dt \int_A |a|^{p'} \Lambda_t(da).$$

Since $P' \in \mathcal{RA}(\rho)$, it must hold $\mathbb{E}^{P'} \int_{\mathcal{C}^d} \|x\|_T^p \mu(dx) < \infty$ and $\mathbb{E}^{P'} \|X\|_T^p < \infty$, which implies

$$(3.23) \quad \mathbb{E}^{P'} \int_0^T dt \int_A |a|^{p'} \Lambda_t(da) < \infty.$$

Let $\iota_n : A \rightarrow A$ denote any measurable function satisfying $\iota_n(A) \subset A_n$ and $\iota_n(a) = a$ for all $a \in A_n$, so that ι_n converges pointwise to the identity. Let Λ^n denote the image under Λ of the map $(t, a) \mapsto (t, \iota_n(a))$, so that $P'(\Lambda^n \in \mathcal{V}_n) = 1$. Let $Q'_n := P' \circ (\xi, B, W, \mu, \Lambda^n)^{-1}$, which is in $\mathcal{A}_n(\rho)$. Since $\Lambda^n \rightarrow \Lambda$ P' -a.s., it follows that $Q'_n \rightarrow Q'$ in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$, where Q' satisfies $P' = \mathcal{R}(Q')$. By continuity of \mathcal{R} [see Lemma 3.12, which applies thanks to (3.15)], $\mathcal{R}(Q'_n) \rightarrow \mathcal{R}(Q') = P'$ in $\mathcal{P}^p(\Omega)$. Now, since $|\iota_n(a)| \leq |a|$, we have

$$\int_0^T \int_A |a|^{p'} \Lambda_t^n(da) dt \leq \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt,$$

which implies that the sequence

$$\left(\int_0^T \int_A |a|^{p'} \Lambda_t^n(da) dt \right)_{n=1}^\infty$$

is uniformly P' -integrable. By Lemma 2.4, we then have

$$\sup_n \mathbb{E}^{\mathcal{R}(Q'_n)} \left[\|X\|_T^{p'} + \int_{\mathcal{C}^d} \|z\|_T^{p'} \mu(dz) \right] < \infty.$$

The growth assumptions of f and g imply that the rewards are uniformly integrable in the sense that

$$\lim_{r \rightarrow \infty} \sup_n \mathbb{E}^{\mathcal{R}(Q'_n)} [\Gamma(\mu, \Lambda, X) 1_{\{|\Gamma(\mu, \Lambda, X)| > r\}}] = 0.$$

Finally, from the continuity of Γ we conclude that $J(\mathcal{R}(Q'_n)) \rightarrow J(P')$.

3.5.5. *Proof of Lemma 3.18.* Find $Q' \in \mathcal{A}_a(\rho)$ such that $P' = \mathcal{R}(Q')$. There exist m and an adapted function $\phi : \Omega_0 \times \mathcal{P}^P(\mathcal{X}) \rightarrow \mathcal{V}_m$ such that $\phi(\omega, \cdot)$ is continuous for each $\omega \in \Omega_0$ and

$$Q' := \rho \circ (\xi, B, W, \mu, \phi(\xi, B, W, \mu))^{-1}.$$

Recalling the definition of ρ_n from just before (3.14), define

$$Q'_n := \rho_n \circ (\xi, B, W, \mu, \phi(\xi, B, W, \mu))^{-1}.$$

Note that $Q'_n(\Lambda \in \mathcal{V}_m) = 1$. Hence, $Q'_n \in \mathcal{A}_n(\rho_n)$ for $n \geq m$. It follows from boundedness and continuity of ϕ in μ [and Lemma A.4 to handle the fact that ϕ may not be continuous in (ξ, B, W)] that $Q'_n \rightarrow Q'$. The proof will be complete if we can show

$$(3.24) \quad \mathcal{R}_n(Q'_n) \rightarrow P' \quad \text{in } \mathcal{P}^P(\Omega).$$

Indeed, since A_m is compact, we use the continuity of J (see Lemma 3.13) to complete the proof. We prove (3.24) with exactly the same argument as in Lemma 3.12: Since $\mathcal{R}_n(Q_n) \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q_n$ are relatively compact in $\mathcal{P}^P(\Omega_0 \times \mathcal{P}^P(\mathcal{X}) \times \mathcal{V})$, Aldous' criterion (see Proposition B.1 for details) implies that $\mathcal{R}_n(Q'_n) \circ X^{-1}$ are relatively compact in $\mathcal{P}^P(\mathcal{C}^d)$. Thus $\mathcal{R}_n(Q'_n)$ are relatively compact in $\mathcal{P}^P(\Omega)$. Conclude exactly as in the proof of Lemma 3.12 that any limit point must equal P' .

4. Strict and strong controls. This section addresses the question of the existence of strict and strong controls. Recall that $Q \in \mathcal{A}(\rho)$ [resp., $P \in \mathcal{RA}(\rho)$] is a strict control if $Q(\Lambda \in \mathcal{V}_a) = 1$ [resp., $P(\Lambda \in \mathcal{V}_a) = 1$], where

$$(4.1) \quad \mathcal{V}_a := \{q \in \mathcal{V} : q_t = \delta_{\alpha(t)} \text{ for some } \alpha \in L^P([0, T]; A)\}.$$

Recall also that Q is a strong control if there exists an A -valued process $(\alpha_t)_{t \in [0, T]}$, progressively-measurable with respect to the P -completion of $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$, such that $Q(\Lambda = dt\delta_{\alpha_t}(da)) = 1$. Section 4.1 addresses this point under a quite standard condition in control theory. Section 4.2 identifies more specialized assumptions which allow us to find a weak MFG solution with strong control. The idea in each case is the same as in references on relaxed controls: given any weak (relaxed) control, under suitable convexity assumptions, the optional projection of the control onto a suitable sub-filtration will yield an admissible control with a greater value than the original control, without disturbing the joint laws of the other processes.

4.1. *Strict controls.* The following assumption is well known in control theory (dating to Filippov [21]) and permits the construction of a weak MFG solution with weak *strict* control.

ASSUMPTION C. For each $(t, x, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d)$, the following set is convex:

$$K(t, x, \mu) := \{(b(t, x, \mu, a), z) : a \in A, z \leq f(t, x, \mu, a)\} \subset \mathbb{R}^d \times \mathbb{R}.$$

The most obvious examples of Assumption C are the affine drifts b , that is, $b(t, x, \mu, a) = b^1(t, x, \mu)a + b^2(t, x, \mu)$, and objectives $f(t, x, \mu, a)$ which are concave in a . Here is the main result of this subsection.

THEOREM 4.1. *In addition to Assumption A, suppose also that Assumption C holds. Then there exists a weak MFG solution with weak strict control that satisfies $E \int_0^T |\alpha_t|^{p'} dt < \infty$.*

The proof of Theorem 4.1 relies on the following.

PROPOSITION 4.2. *Assume Assumptions A and C hold. Let $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ and $P \in \mathcal{RA}(\rho)$. Then there exists a strict control $P' \in \mathcal{RA}(\rho)$ such that*

$$P' \circ (\xi, B, W, \mu, X)^{-1} := P \circ (\xi, B, W, \mu, X)^{-1}$$

and $J(P') \geq J(P)$.

PROOF. Note that $\int_A \Lambda_t(da)(b, f)(t, X_t, \mu_t^x, a) \in K(t, X_t, \mu_t^x)$. By [28], Theorem A.9, or rather a slight extension thereof in [17], Lemma 3.1, there exist $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ -progressive processes $\hat{\alpha}$ and \hat{z} , taking values in A and $[0, \infty)$, respectively, such that

$$(4.2) \quad \int_A \Lambda_t(da)(b, f)(t, X_t, \mu_t^x, a) = (b, f)(t, X_t, \mu_t^x, \hat{\alpha}_t) - (0, \hat{z}_t).$$

Define $P' \in \mathcal{P}(\Omega)$ by $P' = P \circ (\xi, B, W, \mu, dt\delta_{\hat{\alpha}_t}(da), X)^{-1}$. Clearly, (B, W) is a Wiener process with respect to $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ under P' . Since the state equation (3.3) holds under P , it follows from (4.2) that the state equation holds under P' as well, since the first coordinate of the remainder $(0, \hat{z}_t)$ in (4.2) is zero. Moreover,

$$\begin{aligned} J(P') &= \mathbb{E}^P \left[\int_0^T dt f(t, X_t, \mu_t^x, \hat{\alpha}_t) + g(X_T, \mu_T^x) \right] \\ &\geq \mathbb{E}^P \left[\int_0^T dt \int_A \Lambda_t(da) f(t, X_t, \mu_t^x, a) dt + g(X_T, \mu_T^x) \right] = J(P). \end{aligned}$$

Letting

$$Q' := P' \circ (\xi, B, W, \mu, \Lambda)^{-1} = P \circ (\xi, B, W, \mu, dt\delta_{\hat{\alpha}_t}(da))^{-1},$$

Q' is in $\mathcal{A}(\rho)$. The reason is that $(\hat{\alpha}_t)_{t \in [0, T]}$ is $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ -progressive, so that, for each $t \in [0, T]$ and $C \in \mathcal{B}(A)$, $\int_0^t 1_C(\hat{\alpha}_s) ds$ is $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X}$ -measurable. Since the solution of the state equation (3.1) is strong, $\int_0^t 1_C(\hat{\alpha}_s) ds$ coincides P a.s. with a $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda}$ -measurable random variable. By assumption, \mathcal{F}_t^Λ and $\mathcal{F}_t^{\xi, B, W, \mu}$ are conditionally independent under P given $\mathcal{F}_t^{\xi, B, W, \mu}$. We deduce that $\sigma(\int_0^r 1_C(\hat{\alpha}_s) ds : r \leq t, C \in \mathcal{B}(A))$ and $\mathcal{F}_T^{\xi, B, W, \mu}$ are also conditionally independent under P given $\mathcal{F}_t^{\xi, B, W, \mu}$, which is enough to prove that $Q' \in \mathcal{A}(\rho)$. \square

PROOF OF THEOREM 4.1. Let $P \in \mathcal{P}(\Omega)$ be a MFG solution, whose existence is guaranteed by Theorem 3.2, and set $\rho = P \circ (\xi, B, W, \mu)^{-1}$. By Proposition 4.2, there exists $P' \in \mathcal{RA}(\rho)$ such that $P \circ (\xi, B, W, \mu, X)^{-1} = P' \circ (\xi, B, W, \mu, X)^{-1}$, $J(P) \leq J(P')$, and $P'(\Lambda_t = \delta_{\alpha_t} \text{ a.e. } t) = 1$ for some $(\alpha_t)_{t \in [0, T]}$ -progressive process $(\alpha_t)_{t \in [0, T]}$. But since $P \in \mathcal{RA}^*(\rho)$ (i.e., P is optimal for the control problem corresponding to ρ), it follows that $P' \in \mathcal{RA}^*(\rho)$. It remains to deal with the fixed-point condition. Define

$$\bar{\mu} := P'((W, \Lambda, X) \in \cdot | B, \mu).$$

Conditioning on $(B, \bar{\mu})$ yields $\bar{\mu} := P'((W, \Lambda, X) \in \cdot | B, \bar{\mu})$. Now if $\phi : \mathcal{X} \rightarrow \mathbb{R}$ is $\mathcal{F}_t^{W, \Lambda, X}$ -measurable then

$$\begin{aligned} \int_{\mathcal{X}} \phi d\bar{\mu} &= \mathbb{E}^{P'}[\phi(W, \Lambda, X) | B, \mu] = \mathbb{E}^{P'}[\mathbb{E}^{P'}[\phi(W, \Lambda, X) | \mathcal{F}_T^{\xi, B, W, \mu}] | \mathcal{F}_T^{B, \mu}] \\ &= \mathbb{E}^{P'}[\mathbb{E}^{P'}[\phi(W, \Lambda, X) | \mathcal{F}_t^{\xi, B, W, \mu}] | \mathcal{F}_T^{B, \mu}] \\ &= \mathbb{E}^{P'}[\phi(W, \Lambda, X) | \mathcal{F}_t^{B, \mu}]. \end{aligned}$$

The second equality follows from the conditional independence of $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X}$ and $\mathcal{F}_T^{\xi, B, W, \mu}$ given $\mathcal{F}_t^{\xi, B, W, \mu}$ under P' , which holds because $P' \in \mathcal{RA}(\rho)$, and the last equality follows easily from the independence of (ξ, W) and (B, μ) . This holds for each ϕ , and thus $\mathcal{F}_t^{\bar{\mu}} \subset \mathcal{F}_t^{B, \mu}$ for all t , up to $\rho \circ (B, \mu)^{-1}$ -null sets. It follows that (B, W) is a Wiener process under P' with respect to the filtration generated by $(\xi, B, W, \bar{\mu}, \Lambda, X)$, which is smaller than $(\xi, B, W, \mu, \Lambda, X)$. Moreover, by definition, $\bar{\mu}^x = P'(X \in \cdot | B, \mu)$, and since $P' \circ (B, \mu, X)^{-1} = P \circ (B, \mu, X)^{-1}$ and $\mu^x = P(X \in \cdot | B, \mu)$ imply $\mu^x = P'(X \in \cdot | B, \mu)$, we deduce that $P'(\bar{\mu}^x = \mu^x) = 1$.

Now define $\bar{P} := P' \circ (\xi, B, W, \bar{\mu}, \Lambda, X)^{-1}$; we will show that this is in fact the MFG solution we are looking for. Indeed, from $P'(\bar{\mu}^x = \mu^x) = 1$ it follows that the canonical processes verify the state equation (3.3) under \bar{P} . Hence, in light of the above considerations, we may apply Lemma 3.7 to conclude that \bar{P} is a MFG pre-solution (with weak strict control). In particular, we have $\bar{P} \in \mathcal{RA}(\bar{\rho})$, where $\bar{\rho} := \bar{P} \circ (\xi, B, W, \mu)^{-1} = \rho \circ (\xi, B, W, \bar{\mu})^{-1}$. Moreover, $P' \circ (\mu^x, \Lambda, X)^{-1} =$

$\bar{P} \circ (\mu^x, \Lambda, X)^{-1}$ clearly implies $J(P') = J(\bar{P})$. Although $P' \in \mathcal{RA}^*(\rho)$, this does not immediately imply that $\bar{P} \in \mathcal{RA}^*(\bar{\rho})$, and we must complete the proof carefully.

Fix $\bar{Q} \in \mathcal{A}_a(\bar{\rho})$, where we recall the definition of $\mathcal{A}_a(\rho)$ from Section 3.5.3. That is

$$\bar{Q} = \bar{\rho} \circ (\xi, B, W, \mu, \phi(\xi, B, W, \mu))^{-1} = \rho \circ (\xi, B, W, \bar{\mu}, \phi(\xi, B, W, \bar{\mu}))^{-1}$$

for some adapted function $\phi : \Omega_0 \times \mathcal{P}^P(\mathcal{X}) \rightarrow \mathcal{V}$. Define

$$\bar{Q}' := \rho \circ (\xi, B, W, \mu, \phi(\xi, B, W, \bar{\mu}))^{-1}.$$

Then, since ϕ is adapted and $\mathcal{F}_t^{\bar{\mu}} \subset \mathcal{F}_t^{B, \mu}$ up to null sets, we conclude that \bar{Q}' is compatible, in the sense that $\bar{Q}' \in \mathcal{A}(\rho)$. Since $P'(\bar{\mu}^x = \mu^x) = 1$, we have $\mathcal{R}(\bar{Q}) \circ (\mu^x, \Lambda, X)^{-1} = \mathcal{R}(\bar{Q}') \circ (\mu^x, \Lambda, X)^{-1}$. Thus, $P' \in \mathcal{RA}^*(\rho)$ implies

$$J(\bar{P}) = J(P') \geq J(\mathcal{R}(\bar{Q}')) = J(\mathcal{R}(\bar{Q})).$$

Since this holds for all $\bar{Q} \in \mathcal{A}_a(\bar{\rho})$, we finally conclude that $\bar{P} \in \mathcal{RA}^*(\bar{\rho})$ by combining the density results of Lemmas 3.17 and 3.18. \square

REMARK 4.3. It is possible to strengthen this result slightly to conclude that there exists a relaxed MFG solution with weak strict control α_t adapted to $\mathcal{F}_t^{\xi, B, W, \mu, X}$. Indeed, the argument could proceed along the lines of Proposition 4.4 or by way of martingale problems, as in [19, 28].

4.2. *Strong controls.* A strong but common linearity assumption on the coefficients b, σ and σ_0 allows us to find strong controls.

ASSUMPTION D.

(D.1) A is a convex subset of an Euclidean space, and the state coefficients are affine in (x, a) , in the following form:

$$b(t, x, \mu, a) = b^1(t, \mu)x + b^2(t, \mu)a + b^3(t, \mu),$$

$$\sigma(t, x, \mu) = \sigma^1(t, \mu)x + \sigma^2(t, \mu), \quad \sigma_0(t, x, \mu) = \sigma_0^1(t, \mu)x + \sigma_0^1(t, \mu).$$

(D.2) The objective functions are concave in (x, a) ; that is, the maps $(x, a) \mapsto f(t, x, \mu, a)$ and $x \mapsto g(x, \mu)$ are concave for each (t, μ) .

(D.3) f is strictly concave in (x, a) .

PROPOSITION 4.4. Under Assumptions A and (D.1)–(D.2), then:

(1) For each $\rho \in \mathcal{P}_c^P[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^P(\mathcal{X})]$ there exists a strong optimal control; that is $\mathcal{A}^*(\rho)$ contains a strong control.

If also (D.3) holds, then:

(2) For each ρ the optimal control is unique; that is $\mathcal{A}^*(\rho)$ is a singleton for each ρ .

(3) Every weak MFG solution with weak control is a weak MFG solution with strong control.

PROOF. *Proof of (1).* Let $P \in \mathcal{R}\mathcal{A}^*(\rho)$, which is nonempty by Lemma 3.13. Under Assumption D, the state equation writes as

$$\begin{aligned}
 X_t &= X_0 + \int_0^t (b^1(s, \mu_s^x)X_s + b^2(s, \mu_s^x)\alpha_s + b^3(t, \mu_s^x)) ds \\
 (4.3) \quad &+ \int_0^t (\sigma^1(s, \mu_s^x)X_s + \sigma^2(s, \mu_s^x)) dW_s \\
 &+ \int_0^t (\sigma_0^1(s, \mu_s^x)X_s + \sigma_0^2(s, \mu_s^x)) dB_s,
 \end{aligned}$$

where we have let $\alpha_s := \int_A a \Lambda_s(da)$. Let $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ denote the P -completion of the filtration $(\mathcal{F}_t^{\xi, B, W, \mu})_{t \in [0, T]}$. By optional projection (see [34], Appendix A.3, for a treatment without right-continuity of the filtration), there exist $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -optional (and thus progressive) processes $(\tilde{X}_t)_{t \in [0, T]}$ and $(\tilde{\alpha}_t)_{t \in [0, T]}$ such that such that, for each $t \in [0, T]$,

$$\tilde{X}_t := \mathbb{E}[X_t | \tilde{\mathcal{F}}_t^{\xi, B, W, \mu}], \quad \tilde{\alpha}_t := \mathbb{E}[\alpha_t | \tilde{\mathcal{F}}_t^{\xi, B, W, \mu}] \quad \text{a.s.}$$

In fact, it holds that for each $0 \leq s \leq t \leq T$,

$$(4.4) \quad \tilde{X}_s := \mathbb{E}[X_s | \tilde{\mathcal{F}}_t^{\xi, B, W, \mu}], \quad \tilde{\alpha}_s := \mathbb{E}[\alpha_s | \tilde{\mathcal{F}}_t^{\xi, B, W, \mu}] \quad \text{a.s.}$$

Indeed, since (α_s, X_s) is $\mathcal{F}_s^{\xi, B, W, \mu, \Lambda, X}$ -measurable, and since the solution of the state equation (3.1) is strong, we know that (α_s, X_s) is a.s. $\mathcal{F}_s^{\xi, B, W, \mu, \Lambda}$ -measurable. By compatibility, $\mathcal{F}_t^{\xi, B, W, \mu}$ and \mathcal{F}_s^Λ are conditionally independent given $\mathcal{F}_s^{\xi, B, W, \mu}$, and thus $\tilde{\mathcal{F}}_t^{\xi, B, W, \mu}$ and the completion of \mathcal{F}_s^Λ are conditionally independent given $\tilde{\mathcal{F}}_s^{\xi, B, W, \mu}$. This implies (4.4).

Now, for a given $t \in [0, T]$, take the conditional expectation with respect to $\tilde{\mathcal{F}}_t^{\xi, B, W, \mu}$ in (4.3). Using a conditional version of Fubini’s theorem together with (4.4), we get that for each $t \in [0, T]$ it holds P -a.s. that

$$\begin{aligned}
 \tilde{X}_t &= \xi + \int_0^t (b^1(s, \mu_s^x)\tilde{X}_s + b^2(s, \mu_s^x)\tilde{\alpha}_s + b^3(s, \mu_s^x)) ds \\
 (4.5) \quad &+ \int_0^t (\sigma^1(s, \mu_s^x)\tilde{X}_s + \sigma^2(s, \mu_s^x)) dW_s \\
 &+ \int_0^t (\sigma_0^1(s, \mu_s^x)\tilde{X}_s + \sigma_0^2(s, \mu_s^x)) dB_s.
 \end{aligned}$$

Since the right-hand side is continuous a.s. and the filtration is complete, we replace \tilde{X} with an a.s.-continuous modification, so that (4.5) holds for all $t \in [0, T]$,

P -a.s. That is, the processes on either side of the equation are indistinguishable.

Now define $\tilde{P} := P \circ (\xi, B, W, \mu, dt \delta_{\tilde{\alpha}_t}(da), \tilde{X})^{-1}$. It is clear from (4.5) that $\tilde{P} \in \mathcal{RA}(\rho)$. Jensen’s inequality provides

$$\begin{aligned}
 (4.6) \quad J(P) &\leq \mathbb{E}^P \left[\int_0^T f(t, X_t, \mu_t^x, \alpha_t) dt + g(X_T, \mu_T^x) \right] \\
 &= \mathbb{E}^P \left[\int_0^T \mathbb{E}^P [f(t, X_t, \mu_t^x, \alpha_t) | \tilde{\mathcal{F}}_t^{\xi, B, W, \mu}] dt \right. \\
 &\quad \left. + \mathbb{E}^P [g(X_T, \mu_T^x) | \tilde{\mathcal{F}}_T^{\xi, B, W, \mu}] \right]
 \end{aligned}$$

$$(4.7) \quad \leq \mathbb{E}^P \left[\int_0^T f(t, \tilde{X}_t, \mu_t^x, \tilde{\alpha}_t) dt + g(\tilde{X}_T, \mu_T^x) \right] = J(\tilde{P}).$$

Hence, $\tilde{P} \in \mathcal{RA}^*(\rho)$, and (1) is proven.

Proof of (2) and (3). Now suppose assumption (D.3) holds. We prove only (2), from which (3) follows immediately. Unless Λ is already a strict control, then inequality (4.6) is strict, and unless $\int_A \alpha \Lambda_t(da)$ is already $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -adapted, the inequality (4.7) is strict: $J(\tilde{P}) > J(P)$. This proves that all optimal controls must be strict and $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -adapted. Now suppose we have two strict adapted optimal controls, which without loss of generality we construct on the same space $(\Omega_0 \times \mathcal{P}(\mathcal{X}), (\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}, \rho)$. That is,

$$\begin{aligned}
 X_t^i &= X_0 + \int_0^t (b^1(s, \mu_s^x) X_s^i + b^2(s, \mu_s^x) \alpha_s^i + b^3(s, \mu_s^x)) ds \\
 &\quad + \int_0^t (\sigma^1(s, \mu_s^x) X_s^i + \sigma^2(s, \mu_s^x)) dW_s \\
 &\quad + \int_0^t (\sigma_0^1(s, \mu_s^x) X_s^i + \sigma_0^2(s, \mu_s^x)) dB_s, \quad i = 1, 2,
 \end{aligned}$$

where α^i is $\mathcal{F}_t^{X_0, B, W, \mu}$ -adapted. Define

$$X_t^3 := \frac{1}{2} X_t^1 + \frac{1}{2} X_t^2, \quad \alpha_t^3 := \frac{1}{2} \alpha_t^1 + \frac{1}{2} \alpha_t^2.$$

Again taking advantage of the linearity of the coefficients, it is straightforward to check that (X^3, α^3) also solve the state equation. Unless $\alpha^1 = \alpha^2$ holds $dt \otimes dP$ -a.e., the strict concavity and Jensen’s inequality easily imply that this new control achieves a strictly larger reward than either α^1 or α^2 , which is a contradiction. \square

5. Counter-examples. In this section, simple examples are presented to illustrate two points. First, we demonstrate why we cannot expect existence of a strong MFG solution at the level of generality allowed by Assumption A. Second, by providing an example of a mean field game which fails to admit even a weak solution, we show that the exponent p in both the upper and lower bounds of f and g cannot be relaxed to p' .

5.1. *Nonexistence of strong solutions.* Suppose σ is constant, $g \equiv 0$, $p' = 2$, $p = 1$, $A = \mathbb{R}^d$, and choose the following data:

$$b(t, x, \mu, a) = a, \quad f(t, x, \mu, a) = a^\top \tilde{f}(t, \bar{\mu}) - \frac{1}{2}|a|^2,$$

$$\sigma_0(t, x, \mu) = \tilde{\sigma}_0(t, \bar{\mu}),$$

for some bounded continuous functions $\tilde{f} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\tilde{\sigma}_0 : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m_0}$. Here, we have abbreviated $\bar{\mu} := \int_{\mathbb{R}} z \mu(dz)$ for $\mu \in \mathcal{P}^1(\mathbb{R})$. Proposition 4.4 ensures that there exists a weak MFG solution P with strong control. That is (with the same notation as in Proposition 4.4), there exists an $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -progressive \mathbb{R}^d -valued process $(\alpha_t^*)_{t \in [0, T]}$ such that

$$P(\Lambda = dt \delta_{\alpha_t^*}(da)) = 1, \quad \mathbb{E}^P \int_0^1 |\alpha_t^*|^2 dt < \infty.$$

If $(\alpha_t)_{t \in [0, T]}$ is any bounded $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$ -progressive \mathbb{R}^d -valued processes, then optimality of α^* implies

$$\mathbb{E} \int_0^1 \left((\alpha_t^*)^\top \tilde{f}(t, \bar{\mu}_t^x) - \frac{1}{2} |\alpha_t^*|^2 \right) dt \geq \mathbb{E} \int_0^1 \left(\alpha_t^\top \tilde{f}(t, \bar{\mu}_t^x) - \frac{1}{2} |\alpha_t|^2 \right) dt.$$

Hence, $\alpha_t^* = \tilde{f}(t, \bar{\mu}_t^x)$ holds $dt \otimes dP$ -a.e. The optimally controlled state process is given by

$$dX_t^* = \tilde{f}(t, \bar{\mu}_t^x) dt + \sigma dW_t + \tilde{\sigma}_0(t, \bar{\mu}_t^x) dB_t.$$

Conditioning on (B, μ) and using the fixed-point property $\bar{\mu}_t^x = \mathbb{E}[X_t | B, \mu]$ yields

$$d\bar{\mu}_t^x = \tilde{f}(t, \bar{\mu}_t^x) dt + \tilde{\sigma}_0(t, \bar{\mu}_t^x) dB_t, \quad \bar{\mu}_0^x = \mathbb{E}[X_0].$$

We have only assumed that \tilde{f} and $\tilde{\sigma}_0$ are bounded and continuous. For the punch-line, note that uniqueness in distribution may hold for such a SDE even if it fails to possess a strong solution, in which case $\bar{\mu}_t^x$ cannot be adapted to the completion of \mathcal{F}_t^B and the MFG solution cannot be strong. Such cases are not necessarily pathological; see Barlow [3] for examples in dimension $d = 1$ with $\tilde{f} \equiv 0$ and σ_0 bounded above and below away from zero.

5.2. *Nonexistence of weak solutions.* Unfortunately, Assumption A does not cover linear-quadratic models with quadratic objectives in x or μ . That is, we do not allow

$$f(t, x, \mu, a) = -|a|^2 - c \left| x + c' \int_{\mathbb{R}^d} z \mu(dz) \right|^2, \quad c, c' \in \mathbb{R}.$$

Even when $c > 0$, so that f and g are bounded from above, we cannot expect a general existence result if $p' = p$. This was observed in [12, 38] in the case $\sigma_0 = 0$; the authors showed that only *certain* linear-quadratic mean field games

admit (strong) solutions. The following example reiterates this point in the setting of common noise and weak solutions, extending the example of [38], Section 7.

Consider constant volatilities σ and σ_0 , $d = 1$, $p' = p = 2$, $A = \mathbb{R}$, and the following data:

$$b(t, x, \mu, a) = a, \quad f(t, x, \mu, a) = -a^2, \\ g(x, \mu) = -(x + c\bar{\mu})^2, \quad c \in \mathbb{R},$$

where $\bar{\mu} := \int_{\mathbb{R}^d} z\mu(dz)$ for $\mu \in \mathcal{P}^1(\mathbb{R}^d)$. Choose $T > 0$, $c \in \mathbb{R}$, and $\lambda \in \mathcal{P}^2(\mathbb{R})$ such that

$$c = (1 - T)/T, \quad T \neq 1, \bar{\lambda} \neq 0.$$

Assumptions (A.1)–(A.5) hold with the one exception that the assumption $p' > p$ is violated. Suppose P is a weak MFG solution with weak control and then define $y_t := \mathbb{E}\bar{\mu}_t^x$. Arguing as in [38], we get

$$y_T = \frac{y_0}{1 - T} + y_T,$$

which implies $y_0 = 0$ and which contradicts $\bar{\lambda} \neq 0$ since $y_0 = \mathbb{E}\bar{\mu}_0^x = \bar{\lambda}$. Hence, for this particular choice of data, there is no weak solution.

It would be interesting to find additional structural conditions under which existence of a solution holds in the case $p' = p$. This question has been addressed in [10] when $p' = p = 2$, b is linear, σ is constant, f and g are convex in (x, α) and without common noise. Therein, the strategy consists in solving approximating equations, for which the related p is indeed less than 2, and then in passing to the limit. In order to guarantee the tightness of the approximating solutions, the authors introduce a so-called *weak mean-reverting condition*, which somehow generalizes the classical conditions for handling linear-quadratic MFG. It reads $\langle x, \partial_x g(0, \delta_x) \rangle \leq c(1 + |x|)$ and $\langle x, \partial_x f(t, 0, \delta_x, 0) \rangle \leq c(1 + |x|)$, where δ_x is the Dirac mass at point x . This clearly imposes some restriction on the coefficients as, in full generality (when $p = p' = 2$), $\partial_x g(0, \delta_x)$ and $\partial_x f(t, 0, \delta_x, 0)$ are expected to be of order 1 in x . The *weak mean-reverting condition* assures that the expectations of the approximating solutions remain bounded along the approximation, which actually suffices to prove tightness. We feel that the same strategy could be applied to our setting by considering the conditional expectation given the common noise instead of the expectation itself. Anyhow, in order to limit the length of the paper, we refrain from discussing further this question.

6. Uniqueness. We now discuss uniqueness of solutions. The goal is twofold. Inspired by the Yamada–Watanabe theory for weak and strong solutions to standard stochastic differential equations, we first claim that every weak MFG solution is actually a strong MFG solution provided the MFG solutions are *pathwise unique*. This is a quite important point from the practical point of view as it guarantees that the equilibrium measure μ^x is adapted to the common noise B . As an

illustration, we prove a modest uniqueness result, inspired by the earlier works by Lasry and Lions [39]. When there is no mean field term in the state coefficients, when the optimal controls are unique, and when the monotonicity condition of Lasry and Lions [39] holds, we indeed have a form of *pathwise uniqueness*.

6.1. *Pathwise uniqueness and uniqueness in law.* The starting point of our analysis is to notice that the law of a weak MFG solution is really determined by the law of (B, μ) . Indeed, for an element $\gamma \in \mathcal{P}^p(C^{m_0} \times \mathcal{P}^p(\mathcal{X}))$, we can define $M\gamma \in \mathcal{P}(\Omega)$ by

$$M\gamma(d\xi, d\beta, dw, dv, dq, dx) = \gamma(d\beta, dv)v(dw, dq, dx)\delta_{x_0}(d\xi).$$

We will say γ is a *MFG solution basis* if the distribution $M\gamma$ together with the canonical processes on Ω form a weak MFG solution. We say *uniqueness in law* holds for the MFG if there is at most one MFG solution basis, or equivalently if any two weak MFG solutions induce the same law on Ω . Given two MFG solution bases γ^1 and γ^2 , we say $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ is a *coupling* of γ^1 and γ^2 if:

1. $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q)$ is a probability space with a complete filtration.
2. B is a $(\mathcal{G}_t)_{t \in [0, T]}$ -Wiener process on Θ .
3. For $i = 1, 2$, $\mu^i : \Theta \rightarrow \mathcal{P}^p(\mathcal{X})$ is such that, for each $t \in [0, T]$ and $C \in \mathcal{F}_t^{W, \Lambda, X}$, $\mu^i(C)$ is \mathcal{G}_t -measurable.
4. For $i = 1, 2$, $Q \circ (B, \mu^i)^{-1} = \gamma^i$.
5. μ^1 and μ^2 are conditionally independent given B .

Suppose that for any coupling $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ of any two MFG solution bases γ^1 and γ^2 we have $\mu^1 = \mu^2$ a.s. Then we say *pathwise uniqueness* holds for the mean field game. The following proposition essentially follows from Theorem 1.5 and Lemma 2.10 of [35], but we include the proof since we use slightly different notions of compatibility and of pathwise uniqueness.

PROPOSITION 6.1. *Suppose Assumption A and pathwise uniqueness hold. Then there exists a unique in law weak MFG solution with weak control, and it is in fact a strong MFG solution with weak control.*

PROOF. Let γ^1 and γ^2 be any two MFG solution bases. Let $\Theta = C^{m_0} \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{P}^p(\mathcal{X})$, and let (B, μ^1, μ^2) denote the identity map on Θ . Let Q be the unique probability measure on Θ under which $Q \circ (B, \mu^i)^{-1} = \gamma^i$ for $i = 1, 2$ and also μ^1 and μ^2 are conditionally independent given B . Define the $(\mathcal{G}_t)_{t \in [0, T]}$ to be the Q -completion of the filtration

$$(\sigma(B_s, \mu^1(C), \mu^2(C)) : s \leq t, C \in \mathcal{F}_t^{W, \Lambda, X})_{t \in [0, T]}.$$

Then $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ satisfies conditions (1) and (3)–(5) of the definition of a coupling. We will check that in a moment that (2) necessarily holds

as well. It then follows from pathwise uniqueness that $\mu^1 = \mu^2$ almost surely, which in turn implies $\gamma^1 = \gamma^2$. Conclude in the usual way (as in [35], Theorem 1.5 or [32], Theorem 3.20) that this unique solution is in fact a strong solution.

To see that B is a $(\mathcal{G}_t)_{t \in [0, T]}$ -Wiener process, we need only to check that $\sigma(B_s - B_t : s \in [t, T])$ is independent of \mathcal{G}_t for each $t \in [0, T]$. Fix $t \in [0, T]$ and let Z_t, Z_{t+}, Y_t^1 , and Y_t^2 be bounded random variables, measurable with respect to $\mathcal{F}_t^B := \sigma(B_s : s \leq t)$, $\sigma(B_s - B_t : s \in [t, T])$, $\sigma(\mu^1(C) : C \in \mathcal{F}_t^{W, \Lambda, X})$, and $\sigma(\mu^2(C) : C \in \mathcal{F}_t^{W, \Lambda, X})$, respectively. Since γ^i is a MFG solution basis, we know that B is a Wiener process with respect to the filtration

$$(\sigma(B_u, \mu^i(C) : u \leq s, C \in \mathcal{F}_s^{W, \Lambda, X}))_{s \in [0, T]}.$$

Thus, Y_t^i is independent of $\sigma(B_s - B_t : s \in [t, T])$. Conditional independence of μ^1 and μ^2 implies

$$\mathbb{E}^Q[Y_t^1 Y_t^2 | B] = \mathbb{E}^Q[Y_t^1 | B] \mathbb{E}^Q[Y_t^2 | B] = \mathbb{E}^Q[Y_t^1 | \mathcal{F}_t^B] \mathbb{E}^Q[Y_t^2 | \mathcal{F}_t^B].$$

Thus, since Z_{t+} is independent of \mathcal{F}_t^B ,

$$\begin{aligned} \mathbb{E}^Q[Z_{t+} Z_t Y_t^1 Y_t^2] &= \mathbb{E}^Q[Z_{t+} Z_t \mathbb{E}^Q[Y_t^1 | \mathcal{F}_t^B] \mathbb{E}^Q[Y_t^2 | \mathcal{F}_t^B]] \\ &= \mathbb{E}^Q[Z_{t+}] \mathbb{E}^Q[Z_t \mathbb{E}^Q[Y_t^1 | \mathcal{F}_t^B] \mathbb{E}^Q[Y_t^2 | \mathcal{F}_t^B]] \\ &= \mathbb{E}^Q[Z_{t+}] \mathbb{E}^Q[Z_t Y_t^1 Y_t^2]. \end{aligned}$$

This implies that $\sigma(B_s - B_t : s \in [t, T])$ is independent of \mathcal{G}_t . \square

6.2. *Lasry–Lions monotonicity condition.* An application of Proposition 6.1 is possible under the following.

ASSUMPTION U.

(U.1) b, σ , and σ_0 have no mean field term.

(U.2) f is of the form $f(t, x, \mu, a) = f_1(t, x, a) + f_2(t, x, \mu)$.

(U.3) For all $\mu, \nu \in \mathcal{P}^p(\mathcal{C}^d)$ we have the Lasry–Lions monotonicity condition:

$$(6.1) \quad \int_{\mathcal{C}^d} (\mu - \nu)(dx) \left[g(x_T, \mu_T) - g(x_T, \nu_T) + \int_0^T (f_2(t, x_t, \mu_t) - f_2(t, x_t, \nu_t)) dt \right] \leq 0.$$

(U.4) For any $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ the set $\mathcal{A}^*(\rho)$ is a singleton, which means that the maximization problem *in the environment* ρ has a unique (relaxed) solution. See (3.5) for the definition of $\mathcal{A}^*(\rho)$.

Note that assumptions (D.1)–(D.3) imply (U.4), by Proposition 4.4. We then claim the following.

THEOREM 6.2. *Suppose Assumptions **A** and **U** hold. Then there exists a unique in law weak MFG solution with weak control, and it is in fact a strong MFG solution with weak control. In particular, under Assumption **A**, (D.1)–(D.3), and (U.1)–(U.3), the unique in law weak MFG solution with weak control is in fact a strong MFG solution with strong control.*

PROOF. *First step.* Let γ^1 and γ^2 be two MFG solution bases, and define

$$\rho^i := (M\gamma^i) \circ (\xi, B, W, \mu)^{-1}.$$

Let $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ be any coupling of γ^1 and γ^2 . In view of Proposition 6.1, we will prove $\mu^1 = \mu^2$ a.s. In fact, we may assume without loss of generality that

$$\Theta = \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{P}^p(\mathcal{X}), \quad \mathcal{G}_t = \mathcal{F}_t^B \otimes \mathcal{F}_t^{\mu^1} \otimes \mathcal{F}_t^{\mu^2},$$

and Q is the joint distribution of the canonical processes B, μ^1 , and μ^2 on Θ . For each $i = 1, 2$, there is a kernel

$$\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \ni \omega \mapsto K_\omega^i \in \mathcal{P}(\mathcal{V} \times \mathcal{C}^d)$$

such that

$$M\gamma^i = \rho^i(d\omega)K_\omega^i(dq, dx).$$

The key point is that K^i is necessarily adapted to the completed filtration $(\tilde{\mathcal{F}}_t^{\xi, B, W, \mu})_{t \in [0, T]}$, which means that, for each $t \in [0, T]$ and each $\mathcal{F}_t^{\Lambda, X}$ -measurable $\phi : \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$, the map $\omega \mapsto \int \phi dK_\omega^i$ is $\tilde{\mathcal{F}}_t^{\xi, B, W, \mu}$ -measurable. The proof is as follows. Since $M\gamma^i$ is a weak MFG solution, the σ -fields $\mathcal{F}_T^{\xi, B, W, \mu}$ and \mathcal{F}_t^Λ are conditionally independent under $M\gamma^i$ given $\mathcal{F}_t^{\xi, B, W, \mu}$. Since the solution of the state equation (3.1) is strong, $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X}$ is included in the $M\gamma^i$ -completion of $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda}$, from which we deduce that $\mathcal{F}_T^{\xi, B, W, \mu}$ and $\mathcal{F}_t^{\Lambda, X}$ are conditionally independent under $M\gamma^i$ given $\mathcal{F}_t^{\xi, B, W, \mu}$. Therefore, for each $t \in [0, T]$ and each $\mathcal{F}_t^{\Lambda, X}$ -measurable $\phi : \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$, we have

$$\int \phi dK^i = \mathbb{E}^{M\gamma^i}[\phi(\Lambda, X) | \mathcal{F}_T^{\xi, B, W, \mu}] = \mathbb{E}^{M\gamma^i}[\phi(\Lambda, X) | \mathcal{F}_t^{\xi, B, W, \mu}] \quad \text{a.s.}$$

Second step. Define now the extended probability space

$$\bar{\Omega} := \Theta \times (\mathbb{R}^d \times \mathcal{C}^m) \times (\mathcal{V} \times \mathcal{C}^d)^2, \quad \bar{\mathcal{F}}_t := \mathcal{G}_t \otimes \mathcal{F}_t^{\xi, W} \otimes \mathcal{F}_t^{\Lambda, X} \otimes \mathcal{F}_t^{\Lambda, X},$$

endowed with the probability measure

$$\bar{P} := Q(d\beta, dv^1, dv^2)\lambda(d\xi)\mathcal{W}^m(dw) \prod_{i=1}^2 K_{\xi, \beta, w, v^i}^i(dq^i, dx^i).$$

Let $(B, \mu^1, \mu^2, \xi, W, \Lambda^1, X^1, \Lambda^2, X^2)$ denote the coordinate maps on $\bar{\Omega}$. Let $\mu^{i, x} = (\mu^i)^x$. In words, we have constructed \bar{P} so that the following hold:

- (1) $(B, \mu^1, \mu^2), W,$ and ξ are independent.
- (2) (Λ^1, X^1) and (Λ^2, X^2) are conditionally independent given $(B, \mu^1, \mu^2, \xi, W)$.
- (3) The state equation holds, for each $i = 1, 2,$

$$X_t^i = \xi + \int_0^t ds \int_A \Lambda_s^i(da) b(s, X_s^i, a) ds + \int_0^t \sigma(s, X_s^i) dW_s + \int_0^t \sigma_0(s, X_s^i) dB_s.$$

For $i, j = 1, 2,$ define

$$P^{i,j} := \bar{P} \circ (\xi, B, W, \mu^i, \Lambda^j, X^j)^{-1}.$$

By assumption (U.4), $P^{i,i}$ is the *unique* element of $\mathcal{RA}^*(\rho^i),$ for each $i = 1, 2.$ On the other hand, we will verify that

$$(6.2) \quad P^{1,2} \in \mathcal{RA}(\rho^1) \quad \text{and} \quad P^{2,1} \in \mathcal{RA}(\rho^2).$$

Indeed, defining

$$Q^{1,2} := P^{1,2} \circ (\xi, B, W, \mu, \Lambda)^{-1} = \bar{P} \circ (\xi, B, W, \mu^1, \Lambda^2)^{-1},$$

it is clear that $P^{1,2} = \mathcal{R}(Q^{1,2})$ because of the lack of mean field terms in the state equation [by assumption (U.1)]. It remains only to check that $Q^{1,2}$ is compatible with ρ^1 in the sense of (2) in Section 3.2, or equivalently that, under $\bar{P},$ $\mathcal{F}_T^{\xi, B, W, \mu^1}$ and $\mathcal{F}_t^{\Lambda^2}$ are conditionally independent given $\mathcal{F}_t^{\xi, B, W, \mu^1}.$ Given three bounded real-valued functions ϕ_t^1, ϕ_T^1 and $\psi_t^2,$ where ϕ_t^1 and ϕ_T^1 are both defined on $\Omega_0 \times \mathcal{P}^p(\mathcal{X})$ and are $\mathcal{F}_t^{\xi, B, W, \mu^1}$ -measurable and $\mathcal{F}_T^{\xi, B, W, \mu^1}$ -measurable (respectively), and where ψ_t^2 is defined on \mathcal{V} and is $\mathcal{F}_t^{\Lambda^2}$ -measurable, we have

$$\begin{aligned} & \mathbb{E}^{\bar{P}}[(\phi_t^1 \phi_T^1)(\xi, B, W, \mu^1) \psi_t^2(\Lambda^2)] \\ &= \mathbb{E}^{\bar{P}}\left[(\phi_t^1 \phi_T^1)(\xi, B, W, \mu^1) \int_{\mathcal{V}} \psi_t^2(q) K_{\xi, B, W, \mu^2}^2(dq)\right] \\ &= \mathbb{E}^{\bar{P}}\left[(\phi_t^1 \phi_T^1)(\xi, B, W, \mu^1) \mathbb{E}^{\bar{P}}\left[\int_{\mathcal{V}} \psi_t^2(q) K_{\xi, B, W, \mu^2}^2(dq) \mid \mathcal{F}_t^{\xi, B, W}\right]\right], \end{aligned}$$

where the last equality follows from the fact that μ^1 and μ^2 are conditionally independent given $(\xi, B, W).$ Since (B, W) is an $(\mathcal{F}_t^{\xi, B, W, \mu^2})_{t \in [0, T]}$ -Wiener process and $\int_{\mathcal{V}} \psi_t^2(q) K_{\xi, B, W, \mu^2}^2(dq)$ is $\tilde{\mathcal{F}}_t^{\xi, B, W, \mu^2}$ -measurable by the argument above, the conditioning in the third line can be replaced by a conditioning by $\mathcal{F}_t^{\xi, B, W}.$ Then, using once again the fact that μ^1 and μ^2 are conditionally independent given $(\xi, B, W),$ the conditioning by $\mathcal{F}_t^{\xi, B, W}$ can be replaced by a conditioning by $\mathcal{F}_t^{\xi, B, W, \mu^1},$ which proves the required property of conditional independence. This shows that $Q^{1,2} \in \mathcal{A}(\rho^1),$ and thus $P^{1,2} \in \mathcal{RA}(\rho^1).$ The proof that $P^{2,1} \in \mathcal{RA}(\rho^2)$ is identical.

Third step. Note that (X^i, Λ^i, W) and μ^j are conditionally independent given (B, μ^i) , for $i \neq j$, and thus

$$(6.3) \quad \bar{P}((W, \Lambda^i, X^i) \in \cdot | B, \mu^1, \mu^2) = \bar{P}((W, \Lambda^i, X^i) \in \cdot | B, \mu^i) = \mu^i, \quad i = 1, 2.$$

Now suppose it does not hold that $\mu^1 = \mu^2$ a.s. Suppose that both

$$(6.4) \quad \begin{aligned} P^{1,1} &= P^{1,2} \\ \text{that is } \bar{P} \circ (\xi, B, W, \mu^1, \Lambda^1, X^1)^{-1} &= \bar{P} \circ (\xi, B, W, \mu^1, \Lambda^2, X^2)^{-1}, \end{aligned}$$

$$(6.5) \quad \begin{aligned} P^{2,2} &= P^{2,1} \\ \text{that is } \bar{P} \circ (\xi, B, W, \mu^2, \Lambda^2, X^2)^{-1} &= \bar{P} \circ (\xi, B, W, \mu^2, \Lambda^1, X^1)^{-1}. \end{aligned}$$

It follows that

$$\begin{aligned} \bar{P}((W, \Lambda^2, X^2) \in \cdot | B, \mu^1) &= \bar{P}((W, \Lambda^1, X^1) \in \cdot | B, \mu^1) = \mu^1, \\ \bar{P}((W, \Lambda^1, X^1) \in \cdot | B, \mu^2) &= \bar{P}((W, \Lambda^2, X^2) \in \cdot | B, \mu^2) = \mu^2. \end{aligned}$$

Combined with (6.3), this implies

$$\begin{aligned} \mathbb{E}^{\bar{P}}[\mu^2 | B, \mu^1] &= \mathbb{E}^{\bar{P}}[\bar{P}((W, \Lambda^2, X^2) \in \cdot | B, \mu^1, \mu^2) | B, \mu^1] = \mu^1, \\ \mathbb{E}^{\bar{P}}[\mu^1 | B, \mu^2] &= \mathbb{E}^{\bar{P}}[\bar{P}((W, \Lambda^1, X^1) \in \cdot | B, \mu^1, \mu^2) | B, \mu^2] = \mu^2. \end{aligned}$$

These conditional expectations are understood in terms of mean measures. By conditional independence, $\mathbb{E}^{\bar{P}}[\mu^i | B, \mu^j] = \mathbb{E}^{\bar{P}}[\mu^i | B]$ for $i \neq j$, and thus

$$\mathbb{E}^{\bar{P}}[\mu^2 | B] = \mu^1 \quad \text{and} \quad \mathbb{E}^{\bar{P}}[\mu^1 | B] = \mu^2.$$

Thus, μ^1 and μ^2 are in fact B -measurable and equal, which is a contradiction. Hence, one of the distributional equalities (6.4) or (6.5) must fail. By optimality of $P^{1,1}$ and $P^{2,2}$ and by (6.2), we have the following two inequalities, and assumption (U.4) implies that at least one of them is strict:

$$0 \leq J(P^{2,2}) - J(P^{2,1}) \quad \text{and} \quad 0 \leq J(P^{1,1}) - J(P^{1,2}).$$

Writing out the definition of J and using the special form of f from assumption (U.2),

$$\begin{aligned} 0 \leq & \mathbb{E}^{\bar{P}} \int_0^T dt \left[\int_A \Lambda_t^2(da) f_1(t, X_t^2, a) + f_2(t, X_t^2, \mu_t^{2,x}) \right. \\ & \left. - \int_A \Lambda_t^1(da) f_1(t, X_t^1, a) - f_2(t, X_t^1, \mu_t^{2,x}) \right] \\ & + \mathbb{E}^{\bar{P}} [g(X_T^2, \mu_T^{2,x}) - g(X_T^1, \mu_T^{2,x})], \end{aligned}$$

$$\begin{aligned}
 0 \leq & \mathbb{E}^{\bar{P}} \int_0^T dt \left[\int_A \Lambda_t^1(da) f_1(t, X_t^1, a) + f_2(t, X_t^1, \mu_t^{1,x}) \right. \\
 & \left. - \int_A \Lambda_t^2(da) f_1(t, X_t^2, a) - f_2(t, X_t^2, \mu_t^{1,x}) \right] \\
 & + \mathbb{E}^{\bar{P}} [g(X_T^1, \mu_T^{1,x}) - g(X_T^2, \mu_T^{1,x})],
 \end{aligned}$$

one of the two inequalities being strict. Add these inequalities to get

$$\begin{aligned}
 (6.6) \quad 0 < & \mathbb{E}^{\bar{P}} \left[\int_0^T (f_2(t, X_t^2, \mu_t^{2,x}) - f_2(t, X_t^2, \mu_t^{1,x}) \right. \\
 & \left. + f_2(t, X_t^1, \mu_t^{1,x}) - f_2(t, X_t^1, \mu_t^{2,x})) dt \right] \\
 & + \mathbb{E}^{\bar{P}} [g(X_T^2, \mu_T^{2,x}) - g(X_T^2, \mu_T^{1,x}) + g(X_T^1, \mu_T^{1,x}) - g(X_T^1, \mu_T^{2,x})].
 \end{aligned}$$

Then, conditioning on (B, μ^1, μ^2) inside of (6.6) and applying (6.3) yields

$$\begin{aligned}
 0 < & \mathbb{E}^{\bar{P}} \int_{\mathcal{C}^d} (\mu^{2,x} - \mu^{1,x})(dx) \left[\int_0^T (f_2(t, x_t, \mu_t^{2,x}) - f_2(t, x_t, \mu_t^{1,x})) dt \right. \\
 & \left. + g(x_T, \mu_T^{2,x}) - g(x_T, \mu_T^{1,x}) \right].
 \end{aligned}$$

This contradicts assumption (U.3), and so $\mu^1 = \mu^2$ a.s. \square

APPENDIX A: TOPOLOGY OF WASSERSTEIN SPACES

Recall the definition of the Wasserstein metric from (2.1). For ease of reference, this appendix compiles several known results on Wasserstein spaces.

PROPOSITION A.1 (Theorem 7.12 of [44]). *Let (E, ℓ) be a metric space, and suppose $\mu, \mu_n \in \mathcal{P}^p(E)$. Then the following are equivalent:*

- (1) $\ell_{E,p}(\mu_n, \mu) \rightarrow 0$.
- (2) $\mu_n \rightarrow \mu$ weakly and for some $x_0 \in E$ we have

$$\lim_{r \rightarrow \infty} \sup_n \int_{\{x: \ell^p(x, x_0) \geq r\}} \mu_n(dx) \ell^p(x, x_0) = 0.$$

- (3) $\int \phi d\mu_n \rightarrow \int \phi d\mu$ for all continuous functions $\phi : E \rightarrow \mathbb{R}$ such that there exists $x_0 \in E$ and $c > 0$ for which $|\phi(x)| \leq c(1 + \ell^p(x, x_0))$ for all $x \in E$.

In particular, (2) implies that a sequence $\{\mu_n\} \subset \mathcal{P}^p(E)$ is relatively compact if and only if it is tight [i.e., relatively compact in $\mathcal{P}(E)$] and satisfies the uniform integrability condition (2).

The rest of the results listed here are borrowed from Appendices A and B of [38], but the proofs are straightforward and essentially just extend known results on weak convergence using a homeomorphism between $\mathcal{P}(E)$ and $\mathcal{P}^p(E)$. Indeed, if $x_0 \in E$ is fixed and $\psi(x) := 1 + \ell^p(x, x_0)$, then the map $\mu \mapsto \psi\mu / \int \psi d\mu$ is easily seen to define a homeomorphism from $(\mathcal{P}^p(E), \ell_{E,p})$ to $\mathcal{P}(E)$ with the weak topology, where for each $\mu \in \mathcal{P}^p(E)$ the measure $\psi\mu \in \mathcal{P}(E)$ is defined by $\psi\mu(C) = \int_B C\psi d\mu$ for $C \in \mathcal{B}(E)$. For $P \in \mathcal{P}(\mathcal{P}(E))$, define the mean measure $mP \in \mathcal{P}(E)$ by

$$mP(C) := \int_{\mathcal{P}(E)} \mu(C)P(d\mu), \quad C \in \mathcal{B}(E).$$

PROPOSITION A.2. *Let (E, ℓ) be a complete separable metric space. Suppose $K \subset \mathcal{P}^p(\mathcal{P}^p(E))$ is such that $\{mP : P \in K\} \subset \mathcal{P}(E)$ is tight and*

$$\sup_{P \in K} \int_E mP(dx) \ell^{p'}(x, x_0) < \infty \quad \text{for some } p' > p, x_0 \in E.$$

Then K is relatively compact.

In the next two lemmas, let (E, ℓ_E) and (F, ℓ_F) be two complete separable metric spaces. We equip $E \times F$ with the metric formed by adding the metrics of E and F , given by $((x_1, x_2), (y_1, y_2)) \mapsto \ell_E(x_1, y_1) + \ell_F(x_2, y_2)$, although this choice is inconsequential.

LEMMA A.3. *A set $K \subset \mathcal{P}^p(E \times F)$ is relatively compact if and only if $\{P(\cdot \times \cdot) : P \in K\} \subset \mathcal{P}^p(E)$ and $\{P(E \times \cdot) : P \in K\} \subset \mathcal{P}^p(F)$ are relatively compact.*

LEMMA A.4. *Let $\phi : E \times F \rightarrow \mathbb{R}$ satisfy the following:*

- (1) $\phi(\cdot, y)$ is measurable for each $y \in F$.
- (2) $\phi(x, \cdot)$ is continuous for each $x \in E$.
- (3) There exist $c > 0, x_0 \in E$, and $y_0 \in F$ such that

$$|\phi(x, y)| \leq c(1 + \ell_1^p(x, x_0) + \ell_2^p(y, y_0)) \quad \forall (x, y) \in E \times F.$$

If $P^n \rightarrow P$ in $\mathcal{P}^p(E \times F)$ and $P^n(\cdot \times F) = P(\cdot \times F)$ for all n , then $\int \phi dP^n \rightarrow \int \phi dP$.

The last result we state specialize the above to the space \mathcal{V} , defined in Section 2.3.

LEMMA A.5. *Let (E, ℓ) be a complete separable metric space. Let $\phi : [0, T] \times E \times A \rightarrow \mathbb{R}$ be measurable with $\phi(t, \cdot)$ jointly continuous for each $t \in [0, T]$. Suppose there exist $c > 0$ and $x_0 \in E$ such that*

$$\phi(t, x, a) \leq c(1 + \ell^p(x, x_0) + |a|^p).$$

Then the following map is upper semicontinuous:

$$C([0, T]; E) \times \mathcal{V} \ni (x, q) \mapsto \int q(dt, da)\phi(t, x_t, a).$$

If also $|\phi(t, x, a)| \leq c(1 + \ell^p(x, x_0) + |a|^p)$, then this map is continuous.

APPENDIX B: A COMPACTNESS RESULT FOR ITÔ PROCESSES

Recall from Assumption A that A is a closed subset of a Euclidean space, $p' > p \geq 1 \vee p_\sigma$, $p_\sigma \in [0, 2]$, and $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$. Recall that \mathcal{V} was defined in Section 2.3.

PROPOSITION B.1. *Let d be a positive integer, and fix $c > 0$. Let $\mathcal{Q} \subset \mathcal{P}(\mathcal{V} \times \mathcal{C}^d)$ be the set of laws of $\mathcal{V} \times \mathcal{C}^d$ -valued random variables (Λ, X) defined on some complete filtered probability space $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, P)$ satisfying:*

- (1) $dX_t = \int_A B(t, X_t, a)\Lambda_t(da) dt + \Sigma(t, X_t) dW_t$.
- (2) W is a k -dimensional $(\mathcal{G}_t)_{t \in [0, T]}$ -Wiener process.
- (3) $\Sigma : [0, T] \times \Theta \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times k}$ and $B : [0, T] \times \Theta \times \mathbb{R}^d \times A \rightarrow \mathbb{R}^d$ are jointly measurable, using the $(\mathcal{G}_t)_{t \in [0, T]}$ -progressive σ -field on $[0, T] \times \Theta$.
- (4) X_0 has law λ and is \mathcal{G}_0 -measurable.
- (5) There exists a nonnegative \mathcal{G}_T -measurable random variable Z such that, for each $(t, x, a) \in [0, T] \times \mathbb{R}^d \times A$,

$$|B(t, x, a)| \leq c(1 + |x| + Z + |a|), \quad |\Sigma(t, x)|^2 \leq c(1 + |x|^{p_\sigma} + Z^{p_\sigma})$$

and

$$\mathbb{E}^P \left[|X_0|^{p'} + Z^{p'} + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \right] \leq c.$$

(That is, we vary Σ, B , and the probability space of definition.) Then \mathcal{Q} is a relatively compact subset of $\mathcal{P}^p(\mathcal{V} \times \mathcal{C}^d)$.

PROOF. For each $P \in \mathcal{Q}$ with corresponding probability space $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, P)$ and coefficients B, Σ , standard estimates as in Lemma 2.4 yield

$$\mathbb{E}^P \|X\|_T^{p'} \leq C \mathbb{E}^P \left[1 + |X_0|^{p'} + Z^{p'} + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \right],$$

where $C > 0$ does not depend on P . Hence assumption (5) implies

$$(B.1) \quad \sup_{P \in \mathcal{Q}} \mathbb{E}^P \|X\|_T^{p'} \leq C(1 + c) < \infty.$$

Suppose we can show that $\mathcal{Q}_X := \{P \circ X^{-1} : P \in \mathcal{Q}\} \subset \mathcal{P}(\mathcal{C}^d)$ is tight. Then, from (B.1) (and Proposition A.1) that \mathcal{Q}_X is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$. Moreover,

$$\sup_{P \in \mathcal{Q}} \mathbb{E}^P \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt < \infty$$

implies that $\{P \circ \Lambda^{-1} : P \in \mathcal{Q}\}$ is relatively compact in $\mathcal{P}^p(\mathcal{V})$, by Proposition A.2. Hence, \mathcal{Q} is relatively compact in $\mathcal{P}^p(\mathcal{V} \times \mathcal{C}^d)$, by Lemma A.3. It remains to check that \mathcal{Q}_X is tight, which we will check by verifying Aldous' criterion (see [33], Lemma 16.12) for tightness, or

$$(B.2) \quad \limsup_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}} \sup_{\tau} \mathbb{E}^P [|X_{(\tau+\delta) \wedge T} - X_{\tau}|^p] = 0,$$

where the supremum is over stopping times τ valued in $[0, T]$. The Burkholder–Davis–Gundy inequality implies that there exists a constant $C' > 0$ (which does not depend on P but may change from line to line) such that

$$\begin{aligned} & \mathbb{E}^P [|X_{(\tau+\delta) \wedge T} - X_{\tau}|^p] \\ & \leq C' \mathbb{E}^P \left[\left| \int_{\tau}^{(\tau+\delta) \wedge T} dt \int_A \Lambda_t(da) B(t, X_t, a) \right|^p \right] \\ & \quad + C' \mathbb{E}^P \left[\left(\int_{\tau}^{(\tau+\delta) \wedge T} |\Sigma(t, X_t)|^2 dt \right)^{p/2} \right] \\ & \leq C' \mathbb{E}^P \left[\left| \int_{\tau}^{(\tau+\delta) \wedge T} dt \int_A \Lambda_t(da) c(1 + \|X\|_T + Z + |a|) \right|^p \right] \\ & \quad + C' \mathbb{E}^P \left[\left(\int_{\tau}^{(\tau+\delta) \wedge T} c(1 + \|X\|_T^{p_{\sigma}} + Z^{p_{\sigma}}) dt \right)^{p/2} \right] \\ & \leq C' \mathbb{E}^P \left[(\delta^p + \delta^{p/2})(1 + \|X\|_T^p + Z^p) + \int_{\tau}^{(\tau+\delta) \wedge T} \int_A |a|^p \Lambda_t(da) dt \right]. \end{aligned}$$

The last line simply used Jensen's inequality with $p \geq 1$, and we used also the fact that $p_{\sigma} \leq 2$. Since

$$\sup_{P \in \mathcal{Q}} \mathbb{E}^P [\|X\|_T^p + Z^p] < \infty,$$

it follows that

$$\limsup_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}} \sup_{\tau} \mathbb{E}^P [(\delta^p + \delta^{p/2})(1 + \|X\|_T^p + Z^p)] = 0.$$

By assumption,

$$\sup_{P \in \mathcal{Q}} \mathbb{E}^P \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \leq c < \infty,$$

and since $p < p'$ it follows that

$$\limsup_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}} \sup_{\tau} \mathbb{E}^P \int_{\tau}^{(\tau+\delta) \wedge T} \int_A |a|^p \Lambda_t(da) dt = 0.$$

Putting this together proves (B.2). \square

APPENDIX C: DENSITY OF ADAPTED CONTROLS

The goal of this section is to prove Lemma 3.11, which is essentially an *adapted* analog of the following version of a classical result.

PROPOSITION C.1. *Suppose E and F are complete separable metric spaces and $\mu \in \mathcal{P}(E)$. If μ is nonatomic, then the set*

$$\{\mu(dx)\delta_{\phi(x)}(dy) \in \mathcal{P}(E \times F) : \phi : E \rightarrow F \text{ is measurable}\}$$

is dense in $\mathcal{P}(E, \mu; F) := \{P \in \mathcal{P}(E \times F) : P(\cdot \times F) = \mu\}$. If additionally F is (homeomorphic to) a convex subset of a locally convex space H , then the set

$$\{\mu(dx)\delta_{\phi(x)}(dy) \in \mathcal{P}(E \times F) : \phi : E \rightarrow F \text{ is continuous}\}$$

is also dense in $\mathcal{P}(E, \mu; F)$.

PROOF. This first claim is well known and can be found, for example, in [16], Theorem 2.2.3. To prove the second claim from the first, it suffices to show that any measurable function $\phi : E \rightarrow F$ can be obtained as the μ -a.s. limit of *continuous* functions. By Lusin’s theorem ([6], Theorem 7.1.13), for each $\varepsilon > 0$ we may find a compact $K_\varepsilon \subset E$ such that $\mu(K_\varepsilon^c) \leq \varepsilon$ and the restriction $\phi|_{K_\varepsilon} : K_\varepsilon \rightarrow F$ is continuous. Using a generalization of the Tietze extension theorem due to Dugundji [18], Theorem 4.1, we may find a continuous function $\tilde{\phi}_\varepsilon : E \rightarrow H$ such that $\tilde{\phi}_\varepsilon = \phi$ on K_ε and such that the range $\tilde{\phi}_\varepsilon(E)$ is contained in the convex hull of $\phi|_{K_\varepsilon}(E)$, which is itself contained in the convex set F . We may thus view $\tilde{\phi}_\varepsilon$ as a continuous function from E to F . Since $\mu(\tilde{\phi}_\varepsilon \neq \phi) \leq \mu(K_\varepsilon^c) \leq \varepsilon$, we may find a subsequence of $\tilde{\phi}_\varepsilon$ which converges μ -a.s. to ϕ . \square

As in Lemma 3.11, we work under Assumption B. Recall the definition of an adapted function, given in Definition 3.10.

PROOF OF LEMMA 3.11. It is clear from the definition of an adapted function that $\mathcal{A}_a(\rho) \subset \mathcal{A}(\rho)$. Let $S = (\xi, B, W, \mu)$ abbreviate the identity map on $\Omega_0 \times \mathcal{P}^p(\mathcal{X})$, and let and

$$S^t := (\xi, B_{\cdot \wedge t}, W_{\cdot \wedge t}, \mu^t) \quad \text{where}$$

$$\mu^t := \mu \circ (W_{\cdot \wedge t}, 1_{[0,t]} \Lambda, X_{\cdot \wedge t})^{-1}.$$

On $\Omega' := \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}$, define the filtrations $(\mathcal{F}_t^S)_{t \in [0,T]}$ and $(\mathcal{F}_t^{S,\Lambda})_{t \in [0,T]}$ by $\mathcal{F}_t^S := \sigma(S^t)$ and $\mathcal{F}_t^{S,\Lambda} := \sigma(S^t, 1_{[0,t]} \Lambda)$. Equivalently, our notational conventions allow us to write $\mathcal{F}_t^S = \mathcal{F}_t^{\xi, B, W, \mu}$ and $\mathcal{F}_t^{S,\Lambda} = \mathcal{F}_t^{\xi, B, W, \mu, \Lambda}$.

Fix $Q \in \mathcal{A}(\rho) \subset \mathcal{P}^p(\Omega')$. It is clear that we may approximate elements of \mathcal{V} (in the topology of \mathcal{V}) by piece-wise constant $\mathcal{P}(A)$ -valued paths; that is, we may find a sequence of piece-wise constant $(\mathcal{F}_t^{S,\Lambda})_{t \in [0,T]}$ -adapted $\mathcal{P}(A)$ -valued

processes $(\alpha^k(t))_{t \in [0, T]}$ on Ω' such that $d\alpha^k(t)(da) \rightarrow \Lambda$, Q -a.s., and a fortiori $Q \circ (S, dt\delta_{\alpha^k(t)}(da))^{-1} \rightarrow Q$ weakly. Since A is compact and the S -marginal is fixed, this convergence happens also in $\mathcal{P}^p(\Omega')$, and thus we need not bother to distinguish $\mathcal{P}^p(\Omega')$ -convergence from weak convergence in what follows. Here, a piece-wise constant $\mathcal{F}_t^{S, \Lambda}$ -adapted $\mathcal{P}(A)$ -valued process $(\alpha(t))_{t \in [0, T]}$ is of the form

$$\alpha(t) = a_0 1_{[0, t_0]}(t) + \sum_{i=1}^n a_i 1_{(t_i, t_{i+1}]}(t),$$

where $a_0 \in \mathcal{P}(A)$ is deterministic, a_i is an $\mathcal{F}_{t_i}^{S, \Lambda}$ -measurable $\mathcal{P}(A)$ -valued random variable, and $0 < t_0 < t_1 < \dots < t_{n+1} = T$ for some n . It remains to show that, for any piece-wise constant $(\mathcal{F}_t^{S, \Lambda})_{t \in [0, T]}$ -adapted $\mathcal{P}(A)$ -valued process $(\alpha(t))_{t \in [0, T]}$, there exists a sequence $(\alpha^k(t))_{t \in [0, T]}$ of $(\mathcal{F}_t^S)_{t \in [0, T]}$ -adapted $\mathcal{P}(A)$ -valued processes such that $Q \circ (S, dt\delta_{\alpha^k(t)}(da))^{-1} \rightarrow Q \circ (S, dt\delta_{\alpha(t)}(da))^{-1}$ weakly. The proof is an inductive application of Proposition C.1, the second part of which applies because of the convexity of $\mathcal{P}(A)$.

By the second part of Proposition C.1, there exists a sequence of continuous $\mathcal{F}_{t_1}^S$ -measurable functions $a_1^j : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{P}(A)$ such that $Q \circ (S^{t_1}, a_1^j(S))^{-1} \rightarrow Q \circ (S^{t_1}, a_1)^{-1}$. Since $Q \in \mathcal{A}(\rho)$, $\mathcal{F}_t^{S, \Lambda}$ and \mathcal{F}_T^S are conditionally independent given \mathcal{F}_t^S . In particular, S and (S^{t_1}, a_1) are conditionally independent given S^{t_1} , and so are S and $(S^{t_1}, a_1^j(S))$. Now let $\phi : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be bounded and measurable, and let $\psi : \mathcal{P}(A) \rightarrow \mathbb{R}$ be continuous. Letting \mathbb{E} denote expectation under Q , Lemma A.4 implies

$$\begin{aligned} \lim_{j \rightarrow \infty} \mathbb{E}[\phi(S)\psi(a_1^j(S))] &= \lim_{j \rightarrow \infty} \mathbb{E}[\mathbb{E}[\phi(S)|S^{t_1}]\psi(a_1^j(S))] \\ &= \mathbb{E}[\mathbb{E}[\phi(S)|S^{t_1}]\psi(a_1)] \\ &= \mathbb{E}[\mathbb{E}[\phi(S)|S^{t_1}]\mathbb{E}[\psi(a_1)|S^{t_1}]] \\ &= \mathbb{E}[\mathbb{E}[\phi(S)\psi(a_1)|S^{t_1}]] \\ &= \mathbb{E}[\phi(S)\psi(a_1)]. \end{aligned}$$

This is enough to show that $Q \circ (S, a_1^j(S))^{-1} \rightarrow Q \circ (S, a_1)^{-1}$ (see, e.g., [20], Proposition 3.4.6(b)).

We proceed inductively as follows: suppose we are given $a_1^j, \dots, a_i^j : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{P}(A)$ for some $i \in \{1, \dots, n - 1\}$, where a_k^j is $\mathcal{F}_{t_k}^S$ -measurable for each $k = 1, \dots, i$, and

$$\lim_{j \rightarrow \infty} Q \circ (S, a_1^j(S), \dots, a_i^j(S))^{-1} = Q \circ (S, a_1, \dots, a_i)^{-1}.$$

By Proposition C.1, there exists a sequence of continuous $\mathcal{F}_{t_i+1}^S \otimes \mathcal{B}(\mathcal{P}(A)^i)$ -measurable functions $\hat{a}^k : (\Omega_0 \times \mathcal{P}^p(\mathcal{X})) \times \mathcal{P}(A)^i \rightarrow \mathcal{P}(A)$ such that

$$\begin{aligned} \lim_{k \rightarrow \infty} Q \circ (S^{t_{i+1}}, a_1, \dots, a_i, \hat{a}^k(S, a_1, \dots, a_i))^{-1} \\ = Q \circ (S^{t_{i+1}}, a_1, \dots, a_i, a_{i+1})^{-1}. \end{aligned}$$

It follows as above that in fact

$$\begin{aligned} \lim_{k \rightarrow \infty} Q \circ (S, a_1, \dots, a_i, \hat{a}^k(S, a_1, \dots, a_i))^{-1} \\ = Q \circ (S, a_1, \dots, a_i, a_{i+1})^{-1}. \end{aligned}$$

By continuity of \hat{a}^k , it holds for each k that

$$\begin{aligned} \lim_{j \rightarrow \infty} Q \circ (S, a_1^j(S), \dots, a_i^j(S), \hat{a}^k(S, a_1^j(S), \dots, a_i^j(S)))^{-1} \\ = Q \circ (S, a_1, \dots, a_i, \hat{a}^k(S, a_1, \dots, a_i))^{-1}. \end{aligned}$$

These above two limits imply that there exists a subsequence j_k such that

$$\begin{aligned} \lim_{k \rightarrow \infty} Q \circ (S, a_1^{j_k}(S), \dots, a_i^{j_k}(S), \hat{a}^k(S, a_1^{j_k}(S), \dots, a_i^{j_k}(S)))^{-1} \\ = Q \circ (S, a_1, \dots, a_i, a_{i+1})^{-1}. \end{aligned}$$

Define $a_{i+1}^k(S) := \hat{a}^k(S, a_1^{j_k}(S), \dots, a_i^{j_k}(S))$ to complete the induction.

By the above argument, we construct n sequences $a_i^k : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{P}(A)$, for $i = 1, \dots, n$, where a_i^k is continuous and $\mathcal{F}_{t_i}^S$ -measurable, and

$$\lim_{k \rightarrow \infty} Q \circ (S, a_1^k(S), \dots, a_n^k(S))^{-1} = Q \circ (S, a_1, \dots, a_n)^{-1}.$$

Define

$$\alpha^k(t) = a_0 1_{[0, t_0]}(t) + \sum_{i=1}^n a_i^k(S) 1_{(t_i, t_{i+1}]}(t).$$

The map

$$\mathcal{P}(A)^n \ni (\alpha_1, \dots, \alpha_n) \mapsto dt \left[a_0(da) 1_{[0, t_0]}(t) + \sum_{i=1}^n \alpha_i(da) 1_{(t_i, t_{i+1}]}(t) \right] \in \mathcal{V}$$

is easily seen to be continuous, and thus $Q \circ (S, dt\alpha^k(t)(da))^{-1} \rightarrow Q \circ (S, dt\alpha(t)(da))^{-1}$, completing the proof. \square

REFERENCES

- [1] AHUJA, S. (2014). Wellposedness of mean field games with common noise under a weak monotonicity condition. Preprint. Available at [arXiv:1406.7028](https://arxiv.org/abs/1406.7028).
- [2] ALIPRANTIS, C. and BORDER, K. (2007). *Infinite Dimensional Analysis: A Hitchhiker's Guide*, 3rd ed. Springer, Berlin.
- [3] BARLOW, M. T. (1982). One-dimensional stochastic differential equations with no strong solution. *J. Lond. Math. Soc. (2)* **26** 335–347. [MR0675177](https://doi.org/10.1093/iml/26.3.335)
- [4] BENSOUSSAN, A., FREHSE, J. and YAM, P. (2013). *Mean Field Games and Mean Field Type Control Theory*. Springer, New York. [MR3134900](https://doi.org/10.1007/978-1-4939-9832-7)
- [5] BENSOUSSAN, A., FREHSE, J. and YAM, P. (2014). The master equation in mean field theory. Preprint. Available at [arXiv:1404.4150](https://arxiv.org/abs/1404.4150).
- [6] BOGACHEV, V. I. (2007). *Measure Theory. Vol. I, II*. Springer, Berlin. [MR2267655](https://doi.org/10.1007/978-3-540-34514-0)
- [7] BRÉMAUD, P. and YOR, M. (1978). Changes of filtrations and of probability measures. *Z. Wahrsch. Verw. Gebiete* **45** 269–295. [MR0511775](https://doi.org/10.1007/BF011775)
- [8] CARDALIAGUET, P. (2012). Notes on mean field games.
- [9] CARDALIAGUET, P., GRABER, P. J., PORRETTA, A. and TONON, D. (2015). Second order mean field games with degenerate diffusion and local coupling. *NoDEA Nonlinear Differential Equations Appl.* **22** 1287–1317. [MR3399179](https://doi.org/10.1007/s00033-015-0591-7)
- [10] CARMONA, R. and DELARUE, F. (2013). Probabilistic analysis of mean-field games. *SIAM J. Control Optim.* **51** 2705–2734. [MR3072222](https://doi.org/10.1137/1213001)
- [11] CARMONA, R. and DELARUE, F. (2014). The master equation for large population equilibriums. In *Stochastic Analysis and Applications 2014. Springer Proc. Math. Stat.* **100** 77–128. Springer, Cham. [MR3332710](https://doi.org/10.1007/978-3-319-07710-1_5)
- [12] CARMONA, R., DELARUE, F. and LACHAPPELLE, A. (2013). Control of McKean–Vlasov dynamics versus mean field games. *Math. Financ. Econ.* **7** 131–166. [MR3045029](https://doi.org/10.1007/s11464-013-0302-7)
- [13] CARMONA, R., FOUQUE, J.-P. and SUN, L.-H. (2015). Mean field games and systemic risk. *Commun. Math. Sci.* **13** 911–933. [MR3325083](https://doi.org/10.1142/S175137581550083)
- [14] CARMONA, R. and LACKER, D. (2015). A probabilistic weak formulation of mean field games and applications. *Ann. Appl. Probab.* **25** 1189–1231. [MR3325272](https://doi.org/10.1214/14-AAP1072)
- [15] CARMONA, R. and ZHU, X. (2016). A probabilistic approach to mean field games with major and minor players. *Ann. Appl. Probab.* **26** 1535–1580. [MR3513598](https://doi.org/10.1214/15-AAP1072)
- [16] CASTAING, C., RAYNAUD DE FITTE, P. and VALADIER, M. (2004). *Young Measures on Topological Spaces. Mathematics and Its Applications* **571**. Kluwer Academic, Dordrecht. [MR2102261](https://doi.org/10.1007/978-1-4020-2261-1)
- [17] DUFOUR, F. and STOCKBRIDGE, R. H. (2012). On the existence of strict optimal controls for constrained, controlled Markov processes in continuous time. *Stochastics* **84** 55–78. [MR2876473](https://doi.org/10.1080/00987797.2012.68473)
- [18] DUGUNDJI, J. (1951). An extension of Tietze's theorem. *Pacific J. Math.* **1** 353–367. [MR0044116](https://doi.org/10.1080/00255718.1951.10556116)
- [19] EL KAROUI, N., NGUYEN, D. H. and JEANBLANC-PICQUÉ, M. (1987). Compactification methods in the control of degenerate diffusions: Existence of an optimal control. *Stochastics* **20** 169–219. [MR0878312](https://doi.org/10.1080/009877987088312)
- [20] ETHIER, S. N. and KURTZ, T. G. (2005). *Markov Processes: Characterization and Convergence*, 2nd ed. Wiley, New York.
- [21] FILIPPOV, A. F. (1962). On certain questions in the theory of optimal control. *J. SIAM Control Ser. A* **1** 76–84. [MR0149985](https://doi.org/10.1137/07149985)
- [22] GOMES, D. A., MOHR, J. and SOUZA, R. R. (2010). Discrete time, finite state space mean field games. *J. Math. Pures Appl. (9)* **93** 308–328. [MR2601334](https://doi.org/10.1016/j.jmp.2010.05.004)
- [23] GOMES, D. A., MOHR, J. and SOUZA, R. R. (2013). Continuous time finite state mean field games. *Appl. Math. Optim.* **68** 99–143. [MR3072242](https://doi.org/10.1007/s00245-013-9224-2)

- [24] GOMES, D. A., PATRIZI, S. and VOSKANYAN, V. (2014). On the existence of classical solutions for stationary extended mean field games. *Nonlinear Anal.* **99** 49–79. [MR3160525](#)
- [25] GOMES, D. A., PIMENTEL, E. A. and SÁNCHEZ-MORGADO, H. (2015). Time-dependent mean-field games in the subquadratic case. *Comm. Partial Differential Equations* **40** 40–76. [MR3268921](#)
- [26] GOMES, D. A. and SAÚDE, J. (2014). Mean field games models—A brief survey. *Dyn. Games Appl.* **4** 110–154. [MR3195844](#)
- [27] GUÉANT, O., LASRY, J.-M. and LIONS, P.-L. (2011). Mean field games and applications. In *Paris-Princeton Lectures on Mathematical Finance 2010. Lecture Notes in Math.* **2003** 205–266. Springer, Berlin. [MR2762362](#)
- [28] HAUSSMANN, U. G. and LEPETIER, J.-P. (1990). On the existence of optimal controls. *SIAM J. Control Optim.* **28** 851–902. [MR1051628](#)
- [29] HUANG, M. (2010). Large-population LQG games involving a major player: The Nash certainty equivalence principle. *SIAM J. Control Optim.* **48** 3318–3353. [MR2599921](#)
- [30] HUANG, M., MALHAMÉ, R. P. and CAINES, P. E. (2006). Large population stochastic dynamic games: Closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle. *Commun. Inf. Syst.* **6** 221–251. [MR2346927](#)
- [31] JACOD, J. and MÉMIN, J. (1981). Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité. In *Seminar on Probability, XV (Univ. Strasbourg, Strasbourg, 1979/1980) (French). Lecture Notes in Math.* **850** 529–546. Springer, Berlin. [MR0622586](#)
- [32] JACOD, J. and MÉMIN, J. (1981). Weak and strong solutions of stochastic differential equations: Existence and stability. In *Stochastic Integrals (Proc. Sympos., Univ. Durham, Durham, 1980). Lecture Notes in Math.* **851** 169–212. Springer, Berlin. [MR0620991](#)
- [33] KALLENBERG, O. (2002). *Foundations of Modern Probability*, 2nd ed. Springer, New York. [MR1876169](#)
- [34] KURTZ, T. G. (1998). Martingale problems for conditional distributions of Markov processes. *Electron. J. Probab.* **3** 29 pp. (electronic). [MR1637085](#)
- [35] KURTZ, T. G. (2014). Weak and strong solutions of general stochastic models. *Electron. Commun. Probab.* **19** 16. [MR3254737](#)
- [36] KURTZ, T. G. and PROTTER, P. (1991). Weak limit theorems for stochastic integrals and stochastic differential equations. *Ann. Probab.* **19** 1035–1070. [MR1112406](#)
- [37] LACKER, D. (2014). A general characterization of the mean field limit for stochastic differential games. Preprint. Available at [arXiv:1408.2708](#).
- [38] LACKER, D. (2015). Mean field games via controlled martingale problems: Existence of Markovian equilibria. *Stochastic Process. Appl.* **125** 2856–2894. [MR3332857](#)
- [39] LASRY, J.-M. and LIONS, P.-L. (2007). Mean field games. *Jpn. J. Math.* **2** 229–260. [MR2295621](#)
- [40] LASRY, J.-M. and LIONS, P.-L. and GUÉANT, O. (2008). Application of mean field games to growth theory. Preprint.
- [41] LIONS, P. L. (2007). *Théorie des jeux à champs moyen et applications*. Cours au College de France.
- [42] NOURIAN, M. and CAINES, P. E. (2013). ε -Nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents. *SIAM J. Control Optim.* **51** 3302–3331. [MR3092738](#)
- [43] STROCK, D. W. and VARADHAN, S. R. S. (1979). *Multidimensional Diffusion Processes. Grundlehren der Mathematischen Wissenschaften* **233**. Springer, Berlin. [MR0532498](#)
- [44] VILLANI, C. (2003). *Topics in Optimal Transportation. Graduate Studies in Mathematics* **58**. Amer. Math. Soc., Providence, RI. [MR1964483](#)

- [45] YAMADA, T. and WATANABE, S. (1971). On the uniqueness of solutions of stochastic differential equations. *J. Math. Kyoto Univ.* **11** 155–167. MR0278420

R. CARMONA
ORFE
BENDHEIM CENTER FOR FINANCE
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: rcarmona@princeton.edu

F. DELARUE
LABORATOIRE JEAN-ALEXANDRE DIEUDONNÉ
UNIVERSITÉ DE NICE SOPHIA-ANTIPOLIS
PARC VALROSE
06108 CEDEX 02, NICE
FRANCE
E-MAIL: francois.delarue@unice.fr

D. LACKER
ORFE
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: dlacker@princeton.edu