

ON THE CONVERGENCE OF CLOSED-LOOP NASH EQUILIBRIA TO THE MEAN FIELD GAME LIMIT

BY DANIEL LACKER

Department of Industrial Engineering & Operations Research, Columbia University, daniel.lacker@columbia.edu

This paper continues the study of the mean field game (MFG) convergence problem: In what sense do the Nash equilibria of n -player stochastic differential games converge to the mean field game as $n \rightarrow \infty$? Previous work on this problem took two forms. First, when the n -player equilibria are open-loop, compactness arguments permit a characterization of all limit points of n -player equilibria as weak MFG equilibria, which contain additional randomness compared to the standard (strong) equilibrium concept. On the other hand, when the n -player equilibria are closed-loop, the convergence to the MFG equilibrium is known only when the MFG equilibrium is unique and the associated “master equation” is solvable and sufficiently smooth. This paper adapts the compactness arguments to the closed-loop case, proving a convergence theorem that holds even when the MFG equilibrium is nonunique. Every limit point of n -player equilibria is shown to be the same kind of weak MFG equilibrium as in the open-loop case. Some partial results and examples are discussed for the converse question, regarding which of the weak MFG equilibria can arise as the limit of n -player (approximate) equilibria.

CONTENTS

1. Introduction	1694
2. Setup and main results	1697
2.1. The n -player games	1698
2.2. The mean field game	1699
2.3. Limit theorems	1701
2.4. A partial converse to the main limit theorem	1702
2.5. Closed-loop versus open-loop equilibria	1704
2.6. Ideas of the proof of the main limit theorem	1705
2.6.1. Limiting dynamics	1706
2.6.2. Optimality at the limit	1707
3. Relaxed equilibria	1708
3.1. Relaxed n -player games	1708
3.2. Relaxed mean field equilibria	1710
3.3. Extensions of the limit theorems	1711
4. Relating the various equilibrium concepts	1712
4.1. Proof of Proposition 3.2	1712
4.2. Proof of Proposition 3.3	1713
4.3. Proof of Proposition 2.2	1714
4.4. Proof of Proposition 3.4	1715
4.5. Proof of Proposition 3.7	1716
5. Proof of the main limit theorem	1717
5.1. Tightness	1718
5.2. Relaxed controls	1718
5.3. Projection lemmas	1719
5.4. Identification of limiting dynamics	1721

Received August 2018; revised June 2019.

MSC2020 subject classifications. 60H30, 93E20, 91A06.

Key words and phrases. Mean field games, stochastic differential games, closed-loop controls, McKean–Vlasov equations, relaxed controls.

5.5. Optimality 1723
 5.6. Proof of Theorem 2.9 1732
 6. Closed-loop versus open-loop 1736
 6.1. Proof of Theorem 2.13 1739
 7. Constructing n -player equilibria from mean field equilibria 1740
 7.1. Proof of Theorem 3.10 1740
 7.2. Uncontrolled models and ill-posed ODEs 1742
 7.3. A game-theoretic example 1743
 Appendix A: SDEs with random coefficients 1748
 A.1. Deterministic well-posedness 1749
 A.2. Stochastic well-posedness 1750
 A.3. Stability 1753
 A.4. Forward equations 1754
 Appendix B: Joint measurability of regular conditional laws 1755
 Appendix C: Conditional means of random measures 1757
 References 1758

1. Introduction. The goal of this paper is to deepen the study of the $n \rightarrow \infty$ limit theory for n -player stochastic differential games of mean field type. To briefly summarize the problem, specified in full detail in Section 2, suppose n players have private state processes $\mathbf{X} = (X^1, \dots, X^n)$ governed by the stochastic differential equation (SDE) system

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha^i(t, \mathbf{X}_t)) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

where W^1, \dots, W^n are independent Brownian motions, and X_0^1, \dots, X_0^n are i.i.d. Note that the drift function b is the same for each player, but the dynamics of player i 's state process depend only on X^i itself, the empirical probability measure μ_t^n of all players' states, and the control α^i of player i . Each player chooses α^i from the set \mathcal{A} of measurable functions from $[0, T] \times (\mathbb{R}^d)^n$ to the set A of admissible actions. That is, each player's control is chosen as a (deterministic) function of time and the current states of all players. The goal of player i is to maximize the expected payoff

$$J_i^n(\alpha^1, \dots, \alpha^n) = \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^n, \alpha^i(t, \mathbf{X}_t)) dt + g(X_T^i, \mu_T^n) \right],$$

which takes the same symmetric form as the drift. The primary object of study is a *closed-loop Markovian Nash equilibrium*, defined as any vector $(\alpha^1, \dots, \alpha^n) \in \mathcal{A}^n$ such that

$$J_i^n(\alpha^1, \dots, \alpha^n) \geq \sup_{\beta \in \mathcal{A}} J_i^n(\alpha^1, \dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots, \alpha^n).$$

As is well known, Markovian Nash equilibria can be constructed by solving a parabolic PDE system, representing the value functions of each of the n players, under suitable assumptions on the coefficients (b, f, g) ; see [16], Section 2.1.4. For a more thorough introduction to stochastic differential games and mean field games, refer to the recent books [16, 17].

A fundamental problem in mean field game (MFG) theory is to characterize the limiting behavior of Nash equilibrium as $n \rightarrow \infty$. More specifically, if $(\alpha^{n,1}, \dots, \alpha^{n,n})$ is a Nash equilibrium for each n , how does the associated empirical measure process $\mu^n = (\mu_t^n)_{t \in [0, T]}$ behave as $n \rightarrow \infty$? The heuristic put forth in the foundational work of [53–55] and [44, 45] suggests that the limiting behavior should be captured by what we call in this paper the *strong mean field equilibria*. A strong mean field equilibrium (or *strong MFE*, defined precisely in Definition 2.3) is a flow of probability measures $m = (m_t)_{t \in [0, T]}$ such that $m_t = \text{Law}(X_t^*)$

for all $t \in [0, T]$, where X^* is the optimal state process for the following stochastic control problem, in which m is treated as fixed:

$$(1.1) \quad \sup_{\alpha} \mathbb{E} \left[\int_0^T f(t, X_t, m_t, \alpha(t, X_t)) dt + g(X_T, m_T) \right],$$

$$dX_t = b(t, X_t, m_t, \alpha(t, X_t)) dt + dW_t.$$

The majority of the MFG literature focuses on questions of existence and uniqueness of equilibria, though there is by now a decent understanding of this convergence problem. Early results [30, 53, 55] confirmed the MFE as the relevant limiting concept but imposed strong restrictions on the controls $\alpha^{n,i}$, requiring them to be of the distributed form $\alpha^{n,i}(t, x_1, \dots, x_n) = \widehat{\alpha}^{n,i}(t, x_i)$.

The first comprehensive results came in [33, 51], but notably working with the distinct (and typically simpler to analyze) concept of *open-loop* controls. In an open-loop equilibrium, each player specifies a control as a function of the noises (W^1, \dots, W^n) rather than the states (X^1, \dots, X^n) , and this results in completely different equilibria. See [16], pp. 72–76, for a careful discussion of the differences between the open-loop and closed-loop regimes, which we will review briefly in Section 2.5. For open-loop equilibria, the results of [51] give a rather complete picture of the $n \rightarrow \infty$ behavior: Even when the MFG equilibrium is nonunique, we can still characterize all subsequential limits of μ^n as MFG equilibria, as long as we work with a suitable *weak equilibrium* concept. Conversely, each of these weak equilibria can arise as the limit of μ^n , for a suitable choice of approximate n -player open-loop equilibria.

Our understanding of the $n \rightarrow \infty$ behavior of closed-loop equilibria is much less complete, but a major breakthrough came with the work of Cardaliaguet et al. [11] on the *master equation*, an infinite-dimensional PDE that describes the value function of the mean field game. It was shown in [11], Section 6, how to use a smooth solution of the *master equation* to prove that μ^n converges to the (unique, in their setting) MFG equilibrium $\mu = (\mu_t)_{t \in [0, T]}$; see also [17], Section 6.3. More recently, these ideas were refined in [25, 26] to derive a central limit theorem and a large deviation principle for μ^n , as well as nonasymptotic bounds on various distances between μ^n and its limit μ . The idea of using the master equation to prove limit theorems has proven to be powerful and quite versatile, subsequently adapted to models with finite state space [2, 21], local interactions [10], interactions governed by random graphs [23], and models with a major player [9].

The master equation approach, however, is limited in several ways. The most fundamental shortcoming is that it requires the MFG equilibrium to be unique. In game theory, uniqueness is of course the exception, not the rule, and the aforementioned papers leave open the intriguing question of how to describe the limiting behavior of μ^n when there are multiple MFG equilibria. Furthermore, it is very challenging to produce a classical solution of the master equation, and this has been accomplished so far only in quite restricted settings [11, 22, 39].

This paper fills the gap between the open-loop and closed-loop regimes by proving (in Theorem 2.7) a limit theorem for closed-loop equilibrium which is general enough to accommodate nonunique MFG equilibrium. Under suitable assumptions, we show that the sequence of empirical measure flows (μ^n) is tight in a suitable space, and every limit in distribution is what we call a *weak semi-Markov mean field equilibrium*, or *weak MFE* for short. This equilibrium concept (given precisely in Definition 2.5) differs from the standard strong MFE described above in three key respects:

- The deterministic measure flow $(m_t)_{t \in [0, T]}$ is replaced by a stochastic one $\mu = (\mu_t)_{t \in [0, T]}$.
- The controls α in (1.1) are *semi-Markov*, meaning $\alpha = \alpha(t, X_t, \mu)$, where the dependence on the path $\mu = (\mu_t)_{t \in [0, T]}$ is nonanticipative.
- The consistency condition $m_t = \text{Law}(X_t)$ becomes conditional, $\mu_t = \text{Law}(X_t \mid (\mu_s)_{s \leq t})$.

The philosophy behind the weak MFE is no different from strong MFE: Each individual player treats the mean field μ as given, describing the distribution of states among an infinite (continuum) population of competing players. Each player reacts optimally to μ , in a way that is consistent with (i.e., reproduces) the mean field μ when aggregated over the infinity of players. The stochastic measure flow μ should be thought of as an endogenous common noise, in the sense that its randomness is felt equally by all of the players. An economist might refer to this as *aggregate uncertainty* [4], as opposed to (exogeneous) *aggregate shocks*, which one might produce by allowing correlations between the driving Brownian motions W^1, \dots, W^n as in [11, 18].

In fact, our main limit theorem applies also to *closed-loop path-dependent equilibria*, in which each player can choose a control $\alpha^i = \alpha^i(t, (X_s)_{s \leq t})$ depending on the entire history of the n state processes. Under modest convexity assumptions, we show that every Markovian equilibrium for the n -player game is also a path-dependent equilibrium, which allows us to study the two simultaneously. This seems to be the first MFG limit theorem for path-dependent equilibria.

In the special case where the weak MFE is unique, our main limit theorem becomes a proper convergence result: Every sequence of n -player (closed-loop) Nash equilibria converges to the unique weak MFE. In particular, we show that the well-known monotonicity condition of Lasry–Lions [55] is sufficient (Corollary 2.10), and this recovers and generalizes the aspects of the limit theorems of [11] pertaining to empirical measures.

The proof of our main limit theorem is based on probabilistic weak convergence and compactness arguments, as well as judicious use of Markovian projection arguments which allow one to “mimick” the time- t marginal laws of a general Itô process by a Markovian diffusion (see Theorem 2.14, quoted from [8, 41]). In a sense, the techniques build on those developed in [51] for the open-loop regime, but the adaptation to the closed-loop case is highly nontrivial. The central difficulty of the closed-loop regime comes from the fact that a single player’s change in strategy can have an outsized impact on the empirical measure due to the feedback through the other controls. If we heuristically consider such a player to be *influential*, the key idea behind our proof is that, in a certain averaged sense, *not too many players can be simultaneously influential*. See Section 2.6 for an informal discussion of the proof.

Our notion of *weak MFE* turns out to be equivalent in a certain sense to the notion of *weak MFG solution* introduced in [18, 51]. In particular, we encounter here the same interesting phenomenon explored in [51], Section 3, which is that *not all weak MFE are mixtures of strong MFE*. That is, if we let S denote the set of strong MFE $m = (m_t)_{t \in [0, T]}$ in the usual sense described in (1.1) above, then there can exist weak MFE μ such that $\mathbb{P}(\mu \in S) < 1$. As a consequence, in settings with multiple MFG equilibria, the usual strong MFE concept is inadequate for describing the limiting behavior of n -player equilibria. Notably, this phenomenon does not appear in McKean–Vlasov systems (i.e., uncontrolled systems) or in static mean field games (i.e., one-shot games with no time component).

The converse to our main limit theorem turns out to be challenging to address, and we have only partial results. We show in Theorem 2.12 that, under reasonable assumptions, every *strong MFE* does indeed arise as the $n \rightarrow \infty$ limit of some sequence of approximate n -player (closed-loop) equilibria (see also [17], Section 6.1). The question of if or when this is true for *weak MFE* remains open. We give some examples of weak MFE which are not strong but which do arise as limits of approximate n -player equilibria by exploiting an interesting connection with the *regularization-by-noise* or *Peano phenomenon*, which can be described as follows: Adding ϵdW_t can turn a nonunique ODE into a well-posed SDE, and in certain cases the limit in distribution of the SDE solution as $\epsilon \downarrow 0$ is a particular mixture of the ODE solutions [1]. We mention also the three recent case studies [20, 24, 57], to be discussed further in Section 2.4, explore this equilibrium-selection problem for various specific models.

On the other hand, in the open-loop regime, it is by now well known that strong MFE do arise as the limit of n -player equilibria [15, 44], and it was shown in [51] that the same is true for any weak MFE. In this sense our results support the folklore that open-loop and closed-loop should “converge together” as $n \rightarrow \infty$.¹ This is reminiscent of the study in discrete time in [38], though the precise use of terminology therein is different.

The paper is organized as follows. We begin in Section 2 by specifying notation and assumptions, precisely defining the n -player and mean field games, and stating clearly the main results. Notably, Section 2.6 sketches the key ideas of the proof of the main result, Theorem 2.7. Section 3 makes a first step toward proving the main theorems by relaxing the notions of equilibrium, leading to somewhat more general forms of the main theorems which are interesting in their own right. The heart of the paper is Section 5, devoted to the proof of Theorem 2.7. The comparison between open-loop and closed-loop equilibria is developed further in Section 6, with proofs of several statements from Section 2.5. Lastly, Section 7 contains some proofs and examples surrounding the partial converse to the main limit theorem, discussed in the previous paragraph. The rather long-winded Appendix irons out some necessary but less exciting technical points pertaining to SDEs with random coefficients (Appendix A) as well as measurability properties of regular conditional laws (Appendix B) and conditional expectations of random measures (Appendix C) with respect to the underlying probability law.

2. Setup and main results. We begin by fixing some commonly used notation. We are given a time horizon $T > 0$ and a dimension $d \in \mathbb{N}$, and we write C^d for the space of continuous paths,

$$C^d := C([0, T]; \mathbb{R}^d),$$

equipped with the sup-norm. We use boldface for vectors, such as $\mathbf{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$.

For a complete separable metric space (E, d) , let $\mathcal{P}(E)$ denote the set of Borel probability measures. We always endow $\mathcal{P}(E)$ with the topology of weak convergence and its corresponding Borel σ -field. Although we will not explicitly use it, to fix ideas we suppose throughout the paper that $\mathcal{P}(E)$ is equipped with the Wasserstein metric

$$(2.1) \quad (m, m') \mapsto \inf_{\gamma} \int_{E \times E} 1 \wedge d(x, y) \gamma(dx, dy),$$

where the infimum is over all $\gamma \in \mathcal{P}(E \times E)$ with marginals m and m' . This is known to (completely) metrize weak convergence [67], Theorem 7.12. In particular, we will make frequent use of the space $C([0, T]; \mathcal{P}(\mathbb{R}^d))$, implicitly equipped with the sup-metric.

For any random variable X we write $\mathcal{L}(X)$ for its law, or $\mathcal{L}(X | Y)$ for a version of the conditional law of X given another random variable Y , which is always well-defined up to almost sure equality when the random variables take values in Polish spaces. We write $X \stackrel{d}{=} Y$ when two random variables have the same law, and we write $X \sim \lambda$ to mean that $\mathcal{L}(X) = \lambda$.

We are given a time horizon $T > 0$, a control space A , an initial state distribution $\lambda \in \mathcal{P}(\mathbb{R}^d)$, and the following functions:

$$(b, f) : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \rightarrow \mathbb{R}^d \times \mathbb{R},$$

$$g : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}.$$

The following assumption is in force throughout the paper:

¹Interestingly, however, this connection can break down if the interactions are not sufficiently continuous; the forthcoming [14] studies an explicitly solvable MFG with singular interactions à la Dyson Brownian motion in which the $n \rightarrow \infty$ limits of the open-loop and closed-loop equilibria of the n -player games are different.

ASSUMPTION A.

- (A.1) A is a compact convex subset of normed vector space.
- (A.2) The functions b, f and g are bounded and jointly continuous.

Occasionally, we will also need the following convexity assumption, which dates back to the work of Filippov [32] and Roxin [60]. It holds, for example, if $b = b(t, x, m, a)$ is affine in a and $f = f(t, x, m, a)$ concave in a , for each (t, x, m) .

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

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

Before proceeding, let us comment on these admittedly strong assumptions. The boundedness of b is unfortunate from the perspective of applications but is the most essential for our arguments, which rely heavily on Girsanov’s theorem; to allow unbounded b may be possible but would significantly complicate the already long proofs. On the other hand, it should be straightforward to allow f and g to grow polynomially in x and perhaps even in the mean of the measure m , as long as the initial state distribution λ is granted certain finite moments. It is conceivable also that we could allow a noncompact control space A , under a coercivity assumption on f . Lastly, Assumption B can be dropped at the price of working with relaxed controls, and this is written explicitly in Section 3.

2.1. *The n -player games.* Let $n \in \mathbb{N}$. In the n -player game, an *admissible control* is a progressively measurable function $\alpha : [0, T] \times (\mathcal{C}^d)^n \rightarrow A$.² Let \mathcal{A}_n denote the set of admissible controls. A *Markovian control* is an admissible control $\alpha \in \mathcal{A}_n$ of the form $\alpha(t, x) = \tilde{\alpha}(t, x_t)$, where $\tilde{\alpha} : [0, T] \times (\mathbb{R}^d)^n \rightarrow A$ is Borel measurable. Let $\mathcal{AM}_n \subset \mathcal{A}_n$ denote the set of Markovian controls. Accepting a mild abuse of notation, we will identify \mathcal{AM}_n with the set of Borel measurable functions from $[0, T] \times (\mathbb{R}^d)^n$ to A .

The state processes in the n -player game are described as follows. For any $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathcal{A}_n$, by Girsanov’s theorem the following SDE system has a unique in law solution $X = (X^1, \dots, X^n)$:

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha^i(t, X)) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

where W^1, \dots, W^n are independent d -dimensional Brownian motions, and X_0^1, \dots, X_0^n are i.i.d. with law λ , independent of (W^1, \dots, W^n) .³ We may write $X[\alpha] = (X^1[\alpha], \dots, X^n[\alpha])$ in place of $X = (X^1, \dots, X^n)$ to stress which controls are being applied, and similarly $\mu^n[\alpha] = \mu^n$.

When $(\alpha^1, \dots, \alpha^n)$ are Markovian, the solution of the above SDE is *strong*, thanks to a result of Veretennikov [66] (see also Krylov–Röckner [49], Theorem 2.1). In particular, we can in that case assume the solution processes are all defined on the same probability space. In general, however, we work with weak solutions of SDEs, and keep in mind that for a different control we may need to construct the state process X on a different probability space.

²Here, we may define progressive measurability simply to mean that α is Borel measurable and satisfies $\alpha(t, x) = \alpha(t, x')$ whenever $t \in [0, T]$ and $x, x' \in (\mathcal{C}^d)^n$ satisfy $x_s = x'_s$ for all $s \leq t$.

³We could allow a constant invertible volatility coefficient $\sigma \in \mathbb{R}^{d \times d}$, but by redefining the state variables there is no loss of generality in taking σ to be the identity matrix. Nondegeneracy of the noise is important in this paper, as we will make heavy use of Girsanov’s theorem and well-posedness results for nondegenerate SDEs.

Player $i \in \{1, \dots, n\}$ chooses α^i to try to maximize

$$J_i^n(\alpha^1, \dots, \alpha^n) := \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^n, \alpha^i(t, \mathbf{X})) dt + g(X_T^i, \mu_T^n) \right].$$

DEFINITION 2.1. Let $\epsilon \geq 0$. A closed-loop (path-dependent) ϵ -Nash equilibrium is a tuple $(\alpha^1, \dots, \alpha^n) \in \mathcal{A}_n^n$ such that

$$J_i^n(\alpha^1, \dots, \alpha^n) \geq \sup_{\beta \in \mathcal{A}_n} J_i^n(\alpha^1, \dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots, \alpha^n) - \epsilon \quad \text{for } i = 1, \dots, n.$$

A Markovian ϵ -Nash equilibrium is a tuple $(\alpha^1, \dots, \alpha^n) \in \mathcal{AM}_n^n$ such that⁴

$$J_i^n(\alpha^1, \dots, \alpha^n) \geq \sup_{\beta \in \mathcal{AM}_n} J_i^n(\alpha^1, \dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots, \alpha^n) - \epsilon \quad \text{for } i = 1, \dots, n.$$

Note that the notion of Markovian Nash equilibrium involves a supremum only over \mathcal{AM}_n , so a priori there is no clear relationship between these two equilibrium concepts. Nonetheless, using Assumption B we prove in Section 4 that Markovian equilibria form a subset of closed-loop equilibria, which allows us to study both types of equilibrium simultaneously:

PROPOSITION 2.2. Suppose Assumptions A and B hold, and let $\epsilon \geq 0$. Then any Markovian ϵ -Nash equilibrium is also a closed-loop ϵ -Nash equilibrium.

It is well known that a Markovian Nash equilibrium for an n -player game can be constructed from a classical solution (when one exists) of the corresponding Nash system, a system of n parabolic PDEs representing the value functions of the n players [36]. For a general discussion see [16], Section 2.1.4, or see [11], Section 1.1, for a good summary in a setting closer to ours. Alternatively, Markovian equilibria can be constructed using a form of the stochastic maximum principle [16], Section 2.2.2.

2.2. The mean field game. We next define the limiting (mean field) game, beginning with the usual notion of equilibrium, which we call a strong equilibrium. In the following, for $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ and measurable functions $\alpha : [0, T] \times \mathbb{R}^d \rightarrow A$ we will encounter SDEs which we will write in the form

$$dX_t = b(t, X_t, m_t, \alpha(t, X_t)) dt + dW_t, \quad X_0 \sim \lambda.$$

When we say “ X is the unique solution” of this SDE, we mean implicitly that $X = (X_t)_{t \in [0, T]}$ and $W = (W_t)_{t \in [0, T]}$ are continuous stochastic processes defined on some common filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ on which X is \mathbb{F} -adapted, W is an \mathbb{F} -Brownian motion, the initial state X_0 has law $\mathbb{P} \circ X_0^{-1} = \lambda$ and is independent of W , and the above SDE is satisfied. We avoid making explicit mention of the probability space, as we work exclusively with distributional properties of X . Recall in the following that we write $\mathcal{L}(Y)$ for the law of a random variable Y .

DEFINITION 2.3. We say that $m = (m_t)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is a strong mean field equilibrium (MFE) if there exists a measurable function $\alpha^* : [0, T] \times \mathbb{R}^d \rightarrow A$ such that the unique solution of the SDE

$$dX_t^* = b(t, X_t^*, m_t, \alpha^*(t, X_t^*)) dt + dW_t, \quad X_0^* \sim \lambda$$

satisfies the following:

⁴Alternative terminology is common in the engineering literature: Instead of “closed-loop path-dependent” and “Markovian” one sometimes encounters “closed loop perfect state” and “feedback perfect state,” respectively.

- (1) The consistency condition holds: $m_t = \mathcal{L}(X_t)$ for all $t \in [0, T]$.
- (2) For any measurable function $\alpha : [0, T] \times \mathbb{R}^d \rightarrow A$, we have

$$\begin{aligned} & \mathbb{E} \left[\int_0^T f(t, X_t^*, m_t, \alpha^*(t, X_t^*)) dt + g(X_T^*, m_T) \right] \\ & \geq \mathbb{E} \left[\int_0^T f(t, X_t, m_t, \alpha(t, X_t)) dt + g(X_T, m_T) \right], \end{aligned}$$

where X is the unique solution of

$$dX_t = b(t, X_t, m_t, \alpha(t, X_t)) dt + dW_t, \quad X_0 \sim \lambda.$$

It was shown in [50], Theorem 6.2, that a strong MFE exists under Assumptions **A** and **B**, but we will not make use of this fact. We next define our weak equilibrium concept, after first introducing a useful terminology:

DEFINITION 2.4. For a Polish space E , we say a function $F : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow E$ is semi-Markov if it is Borel measurable and satisfies $F(t, x, m) = F(t, x, m')$ whenever $(t, x) \in [0, T] \times \mathbb{R}^d$ and $m, m' \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ satisfy $m_s = m'_s$ for all $s \leq t$.

We use the term *semi-Markov* because the control $\alpha^*(t, x, m)$ depends on the state process only at its current time (Markovian) but on the entire history of the measure flow (non-Markovian). It is important to notice that the dependence on m is *nonanticipative*.

DEFINITION 2.5. A *weak semi-Markov mean field equilibrium* (or simply a *weak MFE*) is a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^*, X^*, \mu)$, where $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a complete filtered probability space and:

- (1) μ is a continuous \mathbb{F} -adapted $\mathcal{P}(\mathbb{R}^d)$ -valued process, W is a \mathbb{F} -Brownian motion, and X^* is a continuous \mathbb{R}^d -valued \mathbb{F} -adapted process with $\mathbb{P} \circ (X_0^*)^{-1} = \lambda$.
- (2) $\alpha^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$ is semi-Markov.
- (3) X_0^* , μ , and W are independent.
- (4) The state equation holds:

$$dX_t^* = b(t, X_t^*, \mu_t, \alpha^*(t, X_t^*, \mu)) dt + dW_t.$$

- (5) For every alternative semi-Markov $\alpha : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$ we have

$$\begin{aligned} & \mathbb{E} \left[\int_0^T f(t, X_t^*, \mu_t, \alpha^*(t, X_t^*, \mu)) dt + g(X_T^*, \mu_T) \right] \\ & \geq \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \alpha(t, X_t, \mu)) dt + g(X_T, \mu_T) \right], \end{aligned}$$

where X is the solution (see Remark 2.6 below) of

$$(2.2) \quad dX_t = b(t, X_t, \mu_t, \alpha(t, X_t, \mu)) dt + dW_t, \quad X_0 = X_0^*.$$

- (6) The consistency condition holds: $\mu_t = \mathbb{P}(X_t^* \in \cdot \mid \mathcal{F}_t^\mu)$ a.s. for each $t \in [0, T]$, where $\mathcal{F}_t^\mu = \sigma(\mu_s : s \leq t)$.

We refer also to the $\mathcal{P}(\mathbb{R}^d)$ -valued process μ itself as a weak (semi-Markov) MFE. In this way, if μ is deterministic, then it is a weak MFE if and only if it is a strong MFE. In other words, a strong MFE is always a weak MFE.

REMARK 2.6. The SDE (2.2) admits a unique strong solution (in particular, defined on the same probability space Ω), as we discuss in detail in Appendix A. When μ is deterministic, this follows immediately from the main results of [49, 66]. Appendix A extends this to cover stochastic μ , as long as X_0 , W , and μ are independent. In particular, the solution X of (2.2) is adapted to the complete filtration generated by the process $(X_0^*, W_s, \mu_s)_{s \leq t}$, and so is X^* .

2.3. *Limit theorems.* The following is the main result of the paper:

THEOREM 2.7. *Suppose Assumptions A and B hold. Fix a sequence $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$. For each n , suppose $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n}) \in \mathcal{A}_n^n$ is a closed-loop ϵ_n -Nash equilibrium. Then the associated empirical measure flow sequence $\mu^n = \mu^n[\alpha^n]$ is tight as a family of $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ -valued random variables, and every limit in distribution is a weak MFE.*

The proof is given between Sections 4 and 5. Recalling Proposition 2.2, we immediately deduce that Theorem 2.7 remains true if instead α^n is a Markovian ϵ_n -Nash equilibrium.

REMARK 2.8. The proof of Theorem 2.7 will show that we do not need the full strength of the ϵ_n -Nash equilibrium property. In fact, it suffices to assume merely that $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n})$ satisfies the much weaker inequality

$$\frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,n}) + \epsilon_n \geq \sup_{\beta^1, \dots, \beta^n \in \mathcal{A}_n} \frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,k-1}, \beta^k, \alpha^{n,k+1}, \dots, \alpha^{n,n}).$$

This can be interpreted as an averaged form of approximate Nash equilibrium, which interestingly permits a vanishing fraction of players to choose highly suboptimal controls. See also Section 2.6.2 below, particularly inequality (2.6), for an idea of why this is good enough.

It is well known that stronger assumptions, including a suitable monotonicity condition on the payoff functions, ensures that the mean field equilibrium is unique, and we adapt these ideas to our weaker equilibrium concept. The following theorem, inspired by the early uniqueness result of Lasry–Lions [55], is proven in Section 5.6.

THEOREM 2.9. *Suppose Assumption A holds, along with the following:*

- (i) $b(t, x, m, a) = b(t, x, a)$ has no mean field term.
- (ii) $f(t, x, m, a) = f_1(t, x, m) + f_2(t, x, a)$, for some measurable functions f_1 and f_2 .
- (iii) The action space A is a convex, compact subset of \mathbb{R}^k for some k .
- (iv) For each $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$, $b(t, x, a)$ is affine in a , and $f_2(t, x, m, a)$ is strictly concave in a .
- (v) The monotonicity condition holds: For each $m_1, m_2 \in \mathcal{P}(\mathbb{R}^d)$, we have

$$\int_{\mathbb{R}^d} (f_1(t, x, m_1) - f_1(t, x, m_2))(m_1 - m_2)(dx) \leq 0,$$

$$\int_{\mathbb{R}^d} (g(x, m_1) - g(x, m_2))(m_1 - m_2)(dx) \leq 0.$$

Then there exists a unique in law weak MFE, and it is in fact a strong MFE.

Noting that condition (iv) of Theorem 2.9 implies Assumption B, we may combine the Theorems 2.9 and 2.7 to get the following limit theorem:

COROLLARY 2.10. *Suppose the assumptions of Theorem 2.9 hold. For each n , suppose $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n}) \in \mathcal{A}_n^n$ is a closed-loop ϵ_n -Nash equilibrium. Then $\mu^n = \mu^n[\alpha^n]$ converges in probability in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ to the unique strong MFE.*

Corollary 2.10 is worth comparing to the results of [11], Section 2.4.4, the only previous limit theorem for closed-loop n -player equilibria. Assuming a unique (strong) MFE and a smooth solution of the master equation, a comparable limit theorem for μ^n follows from [11], Theorem 2.15, though it is not stated explicitly. Aside from the fact that they treat common noise, our Corollary 2.10 holds under much weaker assumptions. Moreover, our main result, Theorem 2.7, holds even when the MFE is nonunique, which seems completely out of reach of the techniques of [11]. Of course, the smooth regime they work with affords a more refined and quantitative description of the limit theorem including convergence of value functions; see also [25, 26].

REMARK 2.11. Instead of Definition 2.5, one might propose a more natural *weak fully-Markov* equilibrium concept, in which the control is of the form $\alpha^*(t, X_t, \mu_t)$, depending only on the present value of the measure flow. It is not clear if this smaller class of equilibria is sufficient to catch all limit points of n -player equilibria, and we suspect not. The issue is likely the mode of convergence, and the method of proof suggests the following conjecture: In the setting of Theorem 2.7, every limit point of the pre-compact sequence $(\mathcal{L}(\mu_t^n))_{t \in [0, T]}$ in $C([0, T]; \mathcal{P}(\mathcal{P}(\mathbb{R}^d)))$ can be written as $(\mathcal{L}(\mu_t))_{t \in [0, T]}$ for some weak fully-Markov equilibrium, in the sense just described. To prove this would likely require a Markovian projection argument for measure-valued processes, and such technology does not seem to be available at this time.

2.4. A partial converse to the main limit theorem. Theorem 2.7 ensures that all subsequential limits of closed-loop n -player approximate equilibria are weak MFE. The natural followup question is: Are all weak MFE subsequential limits of closed-loop n -player approximate equilibria? This remains unclear in general, but this section discusses a partial result and a sketch of how to build interesting examples, carried out in more detail in Section 7. (Notably, if the n -player equilibria are open-loop rather than closed-loop, then the results of [51] provide an affirmative answer to this question, and we will return to this point in Section 2.5.)

ASSUMPTION C. The drift b is Lipschitz with respect to total variation, in the following sense: There exists $c > 0$ such that, for each $(t, x, a) \in [0, T] \times \mathbb{R}^d \times A$ and $m, m' \in \mathcal{P}(\mathbb{R}^d)$, we have

$$\frac{1}{c} |b(t, x, m, a) - b(t, x, m', a)| \leq \|m - m'\|_{\text{TV}} := \sup_f \int_{\mathbb{R}^d} f d(m - m'),$$

where the supremum is over all measurable functions $f : \mathbb{R}^d \rightarrow [-1, 1]$.

Note that the metric $\|m - m'\|_{\text{TV}}$ dominates the Wasserstein metric defined in (2.1), and thus Assumption C is weaker in a sense than the Wasserstein–Lipschitz assumptions that appear more often in the literature.

We prove the following in Section 7.3, which shows that every strong MFE arises as the limit of n -player approximate equilibria. The only prior result of this nature seems to be the recent [17], Theorem 6.9, which operates under different and mostly stronger assumptions. The same conclusion is also implicit in [11], Proposition 6.3, under even heavier assumptions.

THEOREM 2.12. *Suppose Assumptions A, B and C hold. Suppose $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is a strong MFE. Then there exist $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$ and, for each n , a Markovian ϵ_n -Nash equilibrium $\alpha^n \in \mathcal{AM}_n^n$ such that $\mu^n[\alpha^n]$ converges in law to m in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$.*

The strategy in proving this is standard: Let $\alpha^*(t, x)$ be the corresponding optimal control from Definition 2.3. The state process X in Definition 2.3 is then the solution of

$$(2.3) \quad dX_t = b(t, X_t, m_t, \alpha^*(t, X_t)) dt + dW_t, \quad m_t = \mathcal{L}(X_t), \forall t \in [0, T].$$

Then, we tell each player in the n -player game to adopt the control $\alpha^*(t, X_t^i)$. This results in the n -particle system

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha^*(t, X_t^i)) dt + dW_t^i.$$

We expect from McKean–Vlasov limit theory that μ^n converges in law to m . The inequality of the optimality condition (2) of Definition 2.3 should then translate to the approximate Nash property in the pre-limit. The precise form of Assumption C is inspired from the recent [52], which proves a strong form of propagation of chaos, in total variation norm, which allows us to avoid imposing continuity assumptions on the control α^* .⁵

It is not clear when we can expect Theorem 2.12 to extend to weak MFE. To explain what can go wrong, suppose that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^*, X^*, \mu)$ is a weak MFE in the sense of Definition 2.5. We then have

$$(2.4) \quad dX_t^* = b(t, X_t^*, \mu_t, \alpha^*(t, X_t^*, \mu)) dt + dW_t, \quad \mu_t = \mathcal{L}(X_t^* | \mathcal{F}_t^\mu), t \in [0, T].$$

Because X_0^* , W , and μ are independent, the law of (X_0^*, W) remains unchanged if we condition on μ ; it is then intuitively clear (and follows from Lemma A.2) that the $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ -valued random variable μ belongs almost surely to the set S^* , consisting of those $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ which solve the McKean–Vlasov equation deterministically,

$$(2.5) \quad dX_t^m = b(t, X_t^m, m_t, \alpha^*(t, X_t^m, m)) dt + dW_t, \quad m_t = \mathcal{L}(X_t^m), t \in [0, T].$$

The key point is that if μ is a weak but not strong MFE, then this McKean–Vlasov equation (2.5) is necessarily nonunique; that is, S^* is not a singleton. In other words, a weak MFE can always be expressed as a mixture of solutions of a nonunique McKean–Vlasov equation. As a consequence, we cannot expect propagation of chaos to hold for the corresponding particle system. That is, if we proceed as before by letting the players in the n -player game use the (path-dependent) controls $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n}) \in \mathcal{A}_n^n$ given by

$$\alpha^{n,i}(t, \mathbf{x}) = \alpha^* \left(t, x_t^i, \frac{1}{n} \sum_{k=1}^n \delta_{x^k} \right), \quad \mathbf{x} = (x^1, \dots, x^n) \in (\mathcal{C}^d)^n,$$

then there is no way to know if $\mu^n[\alpha^n]$ converges to the given μ . For nonunique McKean–Vlasov equations, one can often show that the sequence $\mu^n[\alpha^n]$ is tight and that every limit point is supported on S^* . But when S^* is not a singleton, there is no way in general to know which mixture(s) will be “picked out” by the limit $n \rightarrow \infty$.

We will discuss these ideas further in Section 7, which includes examples of weak MFE which are not strong MFE but which do arise as the limits of n -player (approximate) Nash equilibria. Section 7.3, in particular, gives an example of an interesting kind of weak MFE, discussed also in [51], Section 3: If $S \subset C([0, T]; \mathcal{P}(\mathbb{R}^d))$ denotes the set of strong MFE,

⁵Aside from [52], some other papers have developed propagation of chaos results in total variation, such as [3, 46]. But, unlike [52], the results of [3, 46] are quantitative in nature and accordingly require structural assumptions on the coefficients that are not natural in the present context.

then there can exist weak semi-Markov MFE μ with $\mathbb{P}(\mu \in S) < 1$. But we do not address an intriguing open problem: Can one construct a weak MFE μ satisfying $\mathbb{P}(\mu \in S) < 1$ which arises as the limit of n -player approximate equilibria? In the examples we give in Section 7 of weak MFE which arise as the limits of n -player approximate equilibria, the weak MFE are always mixtures of strong MFE; that is, they satisfy $\mathbb{P}(\mu \in S) = 1$. Note, on the other hand, that it is known that all weak MFE do indeed arise as limits of *open-loop* n -player approximate equilibria; see Theorem 2.13 below, essentially quoted from [51].

Lastly, we mention three recent papers [20, 24, 57] which provide enlightening case studies on the equilibrium-selection problem. In the MFG of optimal stopping in [57], certain MFE are shown to arise as the limits of n -player equilibria, while others do not. In the linear-quadratic model of [24], the unique n -player equilibria converge to a mixture of two strong MFE, and there exists a third MFE which is not selected. Lastly, in the two-state model of [20], there are up to three strong MFE, and one of them is selected as the limit of n -player equilibria. Interestingly, both [20, 24] highlight the role of a unique but nonclassical solution of the master equation. Notably, these three papers work with true n -player Nash equilibria, not with ϵ -equilibria like we do. With this in mind, Theorem 2.12 surely fails if we require $\epsilon_n = 0$. This fits well with known results for static games [37], in which all equilibria of a continuum of agents arise as the limit of n -player approximate equilibria, but this fails if the word “approximate” is omitted.

2.5. Closed-loop versus open-loop equilibria. The parallel limit theory for open-loop n -player equilibria is better understood, and we highlight it in this short section for the sake of comparison. First, we recall the definition of open-loop equilibrium. In this section, we impose stronger continuity assumptions on b and f , so that we may apply the results of [51]:

ASSUMPTION D. There exist $c > 0$ such that $\int_{\mathbb{R}^d} |x|^2 \lambda(dx) < \infty$ and, for each $t \in [0, T]$, $a \in A$, $x, x' \in \mathbb{R}^d$, and $m, m' \in \mathcal{P}(\mathbb{R}^d)$, we have

$$|b(t, x, m, a) - b(t, x', m', a)| \leq c(|x - x'| + \mathcal{W}_1(m, m')),$$

where \mathcal{W}_1 denotes the Wasserstein metric, defined by $\mathcal{W}_1(m, m') = \inf_{\gamma} \int_{\mathbb{R}^d \times \mathbb{R}^d} |x - y| \gamma(dx, dy)$, where the infimum is over all probability measures γ on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals m and m' . Moreover, the objective function $f = f(t, x, m, a)$ satisfies the uniform continuity condition

$$\lim_{(x', m') \rightarrow (x, m)} \sup_{a \in A} |f(t, x', m', a) - f(t, x, m, a)| = 0,$$

for all $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$.

The open-loop n -player game is defined on a fixed filtered probability space $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n, \mathbb{P}^n)$, supporting independent \mathbb{F}^n -Brownian motions and i.i.d. \mathcal{F}_0^n -measurable initial states (X_0^1, \dots, X_0^n) with law λ .⁶ Let \mathbb{A}_n denote the set of \mathbb{F}^n -adapted A -valued processes. For $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathbb{A}_n^n$, define the expected payoff

$$J_i^n(\alpha) = \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^n, \alpha_t^i) dt + g(X_T^i, \mu_T^n) \right],$$

where (X^1, \dots, X^n) is the unique strong solution (recalling Assumption D) of the SDE

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha_t^i) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}.$$

⁶The filtration \mathbb{F}^n does not need to be the minimal one generated by the initial states and Brownian motions.

We may again write $\mu^n = \mu^n[\alpha]$ to emphasize the dependence on the choice of control. For $\epsilon \geq 0$, an *open-loop ϵ -equilibrium* is a tuple $\alpha^n = (\alpha^1, \dots, \alpha^n) \in \mathbb{A}_n^n$ such that

$$J_i^n(\alpha) \geq \sup_{\beta \in \mathbb{A}_n} J_i^n(\alpha^1, \dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots, \alpha^n) - \epsilon.$$

It cannot be stressed enough that open-loop and closed-loop equilibria can be very different. See [19] for an example of an n -player game in which the unique (and explicit) open-loop and closed-loop equilibria are distinct, although they converge to the same limit as $n \rightarrow \infty$. Open-loop equilibria are most often found using the stochastic maximum principle [16], Section 2.2.1.

In the open-loop case, at least under the additional Assumption D, a limit theorem analogous to Theorem 2.7 was developed in [51] along with with a full converse. That is, every weak MFE arises as the limit of some sequence of approximate equilibria of the n -player games. The following is essentially [51], Theorem 3.4:

THEOREM 2.13. *Suppose Assumptions A, B and D hold. If, for each n , we are given an open-loop ϵ_n -Nash equilibrium $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n}) \in \mathbb{A}_n^n$ for some $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$, then $\mu^n[\alpha^n]$ is tight in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$, and every limit in distribution is a weak MFE. Conversely, for every weak MFE μ , we may find, for each n , $\epsilon_n \geq 0$ and an open-loop ϵ_n -Nash equilibrium $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n})$ such that $\epsilon_n \rightarrow 0$ and $\mu^n[\alpha^n]$ converges in law to μ in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$.*

The only difference between Theorem 2.13 and [51], Theorem 3.4, lies in the notion of weak MFE. We will prove at the end of Section 6 a correspondence between our notion of weak MFE and the equilibrium concept used in [51], and Theorem 2.13 will then readily follow.

In summary, the main gap remaining between the open-loop and closed-loop regimes is that we have only a *partial* converse for the latter, as discussed in depth in Section 2.4. However, there is also a discrepancy in the assumptions required for each result. The second claim in Theorem 2.13 relies on the additional Assumption D, much as Theorem 2.12 requires Assumption C. Notably, both are Lipschitz assumptions, used in the proofs to ensure uniqueness and propagation of chaos for certain (controlled) McKean–Vlasov equations.

2.6. Ideas of the proof of the main limit theorem. In this section we informally explain some of the main ideas of the rather lengthy proof of Theorem 2.7, which comes in Section 5. Tightness is straightforward here and fairly standard, so we mostly focus on the two bigger challenges of identifying the dynamics at the limit (properties (1–4) and (6) of Definition 2.5) and proving the optimality of the limiting control (property (5) of Definition 2.5).

A key tool in identifying the limiting dynamics and in proving Proposition 2.2 is (a special case of) the Markovian projection theorem, due originally to Gyöngy [41], Theorem 4.6, and later generalized in [8], Corollary 3.7:

THEOREM 2.14 (Markovian projection). *Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space supporting an \mathbb{F} -adapted continuous process X and an \mathbb{F} -Brownian motion W . Suppose $b = (b_t)_{t \in [0, T]}$ is a bounded \mathbb{F} -progressively measurable process such that, almost surely,*

$$X_t = X_0 + \int_0^t b_s ds + W_t, \quad t \in [0, T].$$

Then there exists a bounded measurable function $\hat{b} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that

$$\hat{b}(t, X_t) = \mathbb{E}[b_t | X_t] \quad \text{a.s., } t \in [0, T],$$

and, moreover, the unique strong solution of the SDE

$$dY_t = \widehat{b}(t, Y_t) dt + dW_t, \quad Y_0 = X_0,$$

satisfies $Y_t \stackrel{d}{=} X_t$ for each $t \in [0, T]$.

2.6.1. *Limiting dynamics.* We want to show that, for any weak limit μ of (μ^n) , we may construct a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^*, X^*, \mu)$ and such that properties (1–4) and (6) of Definition 2.5 hold. Much of this argument is an embellishment of a well-established martingale approach for deriving the McKean–Vlasov limit for interacting diffusions, developed for instance in [40, 58]. A first difference is that here we work with the extended empirical measure

$$\overline{\mu}^n = \frac{1}{n} \sum_{k=1}^n \delta_{(X^k, \alpha^{n,k})}.$$

Here we view X^k as a \mathcal{C}^d -valued random variable and $\alpha^{n,k} = \alpha^{n,k}(t, X)$ as a random variable taking values in the space \mathcal{V} of relaxed or measure-valued controls, defined in Section 5.2; the space \mathcal{V} is essentially a convenient compactification of the space $L^0([0, T]; A)$ of measurable A -valued paths. First, we show that every weak limit $\overline{\mu}$ of $(\overline{\mu}^n)$ satisfies

$$\int_{\mathbb{R}^d} \varphi d(\mu_t - \mu_0) = \int_{\mathcal{C}^d \times \mathcal{V}} \int_0^t \int_A \left(b(s, x_s, \mu_s, a) \cdot \nabla \varphi(x_s) + \frac{1}{2} \Delta \varphi(x_s) \right) q_s(da) ds \overline{\mu}(dx, dq)$$

almost surely, for each smooth test function φ on \mathbb{R}^d , where $\mu_t = \overline{\mu} \circ [(x, q) \mapsto x_t]^{-1}$ is the marginal flow associated to the x variable.

The above integral equation closely resembles the weak or integrated form of a Fokker–Planck equation. Instead of an integral $\int_0^t \int_{\mathbb{R}^d} \dots \mu_s(dx) ds$ appearing on the right-hand side, we have a more complicated expression involving the integral with the respect to $\overline{\mu}$. Drawing intuition from the Markovian Projection Theorem 2.14, we would like to condition on the marginal flow $(\mu_t)_{t \in [0, T]}$, in order to “project away the extra randomness” in some sense. Ultimately, we build (cf. Lemma 5.3) a semi-Markov control $\alpha^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$ such that

$$\int_{\mathbb{R}^d} \varphi d(\mu_t - \mu_0) = \int_0^t \int_{\mathbb{R}^d} \left(b(s, x, \mu_s, \alpha^*(t, x, \mu)) \cdot \nabla \varphi(x) + \frac{1}{2} \Delta \varphi(x) \right) \mu_s(dx) ds,$$

almost surely, for each φ , and such that the expected value of objective function is preserved in a suitable sense. This now says that $(\mu_t)_{t \in [0, T]}$ almost surely solves a Fokker–Planck equation, which we can identify with the solution of an SDE. In fact, this SDE is of McKean–Vlasov type, because μ itself appears nonlinearly in the coefficients b and α^* , and this line of reasoning eventually leads us to properties (1–4) and (6) of Definition 2.5.

REMARK 2.15. Throughout the proofs, we repeatedly use projection arguments like Theorem 2.14 to allow us to modify the control into a more convenient form (e.g., semi-Markov), while preserving the (time-marginal) laws of the state process and measure flow. Because of this, our approach does not readily adapt to models in which the functions b or f depend on the distribution the controls in addition to that of the states, which are often known as *extended mean field games*. See [16], Chapter I.4.6–8, for further discussion and references on such models.

2.6.2. *Optimality at the limit.* Suppose now that we have proven the claimed tightness of Theorem 2.7 and also that for any limit point μ of (μ^n) we may construct a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^*, X^*, \mu)$ such that properties (1–4) and (6) of Definition 2.5 hold. The final and most difficult step is to check that this tuple satisfies the optimality property (5) of Definition 2.5. In the following, we work with a relabeled convergent subsequence and assume μ^n converges in law to μ .

The general strategy, reminiscent of Gamma-convergence arguments, is to choose an arbitrary alternative control $\alpha : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$, and give it to each of the players in the n -player game. Precisely, for each n and each $k = 1, \dots, n$, define the n state processes $\mathbf{Y}^k = (Y^{k,1}, \dots, Y^{k,n})$ by

$$dY_t^{k,k} = b(t, Y_t^{k,k}, \mu_t^{n,k}, \alpha(t, Y_t^{k,k}, \mu^{n,k})) dt + dW_t^k, \quad \mu^{n,k} = \frac{1}{n} \sum_{j=1}^n \delta_{Y^{n,k,j}},$$

$$dY_t^{k,i} = b(t, Y_t^{k,i}, \mu_t^{n,k}, \alpha^{n,i}(t, \mathbf{Y}^k)) dt + dW_t^i, \quad i \neq k,$$

with initial states $Y_0^{k,i} = X_0^i$. The state process \mathbf{Y}^k differs from the equilibrium state process $\mathbf{X}[(\alpha^{n,1}, \dots, \alpha^{n,n})]$ only in that we switched player k 's control from $\alpha^{n,k}$ to α .

The assumed ϵ_n -Nash equilibrium property of $(\alpha^{n,1}, \dots, \alpha^{n,n})$ then implies that

$$(2.6) \quad \frac{1}{n} \sum_{k=1}^n \mathbb{E} \left[\int_0^T f(t, X_t^k, \mu_t^n, \alpha^{n,k}(t, \mathbf{X})) dt + g(X_T^k, \mu_T^n) \right] \geq -\epsilon_n + \frac{1}{n} \sum_{k=1}^n \mathbb{E} \left[\int_0^T f(t, Y_t^{k,k}, \mu_t^{n,k}, \alpha(t, Y_t^{k,k}, \mu^{n,k})) dt + g(Y_T^{k,k}, \mu_T^{n,k}) \right].$$

(In fact, this is the only use of the ϵ_n -Nash property in the entire proof, and this explains Remark 2.8.) We then wish to take limits on both sides. First, the arguments of Section 2.6.1 allow us to identify the limit of the left-hand side of (2.6) as precisely the left-hand side of the inequality in (5) of Definition 2.5. What remains is to show that the right-hand side of (2.6) along the same subsequence converges to the right-hand side of the inequality in (5) of Definition 2.5.

This last point is the technical crux of the argument. It is not obvious at first how to approach this, because we know very little about the controls $\alpha^{n,1}, \dots, \alpha^{n,n}$. Intuitively, one is tempted to claim that, because we have only switched one single agent's control, $\mu^{n,k}$ should be close in some sense to μ^n , for each k . The challenge comes from the closed-loop nature of the controls; if one player switches controls, then all of the other players controls react to the change in the state process. It could be the case that all of the controls $\alpha^{n,1}, \dots, \alpha^{n,n}$ depend very heavily on, say, player 1's state process, in which case a change in control from this player 1 would have a strong influence on the empirical measure.

While we cannot show that $\mu^{n,k}$ and μ^n have the same limiting behavior for each k , we are able to show that $\mathcal{L}(\mu^n)$ and $\frac{1}{n} \sum_{k=1}^n \mathcal{L}(\mu^{n,k})$ have the same limiting behavior, in the sense that the total variation distance between these two measures converges to zero as $n \rightarrow \infty$. Indeed, supposing the state process \mathbf{X} is defined on the probability space $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n, \mathbb{P}^n)$, we may define an equivalent probability measure $\mathbb{Q}^{n,k}$ by setting $d\mathbb{Q}^{n,k}/d\mathbb{P}^n := \zeta_T^{n,k}$, where the positive martingale $(\zeta_t^{n,k})_{t \in [0, T]}$ is given as the unique solution of the SDE

$$d\zeta_t^{n,k} = \zeta_t^{n,k} (b(t, X_t^k, \mu_t^n, \alpha(t, X_t^k, \mu^n)) - b(t, X_t^k, \mu_t^n, \alpha^{n,k}(t, \mathbf{X}))) \cdot dW_t^k, \quad \zeta_0^{n,k} = 1.$$

By Girsanov's theorem and uniqueness of the SDEs, we have $\mathcal{L}(\mathbf{Y}^k) = \mathbb{Q}^{n,k} \circ \mathbf{X}^{-1}$. Hence, for any bounded measurable function h ,

$$(2.7) \quad \frac{1}{n} \sum_{k=1}^n \mathbb{E}[h(\mu^{n,k})] = \frac{1}{n} \sum_{k=1}^n \mathbb{E}[\zeta_T^{n,k} h(\mu^n)].$$

Because the Brownian motions W^k are independent, the process $\frac{1}{n} \sum_{k=1}^n \zeta_t^{n,k}$ is a martingale with quadratic variation up to time s given by

$$\frac{1}{n^2} \sum_{k=1}^n \int_0^s |\zeta_s^{n,k}|^2 |b(t, X_t^k, \mu_t^n, \alpha(t, X_t^k, \mu^n)) - b(t, X_t^{n,k}, \mu_t^n, \alpha^{n,k}(t, \mathbf{X}))|^2 dt,$$

which is of order $1/n$ in expectation because b is bounded. Hence, $\frac{1}{n} \sum_{k=1}^n \zeta_T^{n,k} \rightarrow 1$ in probability, and from (2.7) we deduce that

$$(2.8) \quad \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}[h(\mu^{n,k})] - \mathbb{E}[h(\mu^n)] = 0.$$

Most of the intuition behind this proof is contained in this argument that $\mathcal{L}(\mu^n)$ and $\frac{1}{n} \sum_{k=1}^n \mathcal{L}(\mu^{n,k})$ have the same limiting behavior, but one important additional point is worth mentioning: The right-hand side of (2.6) can be written as the integral of a fixed (n -independent) function with respect to the measure $\frac{1}{n} \sum_{k=1}^n \mathcal{L}(Y^{k,k}, \mu^{n,k})$, and it is this measure whose limiting behavior we should identify, not just $\frac{1}{n} \sum_{k=1}^n \mathcal{L}(\mu^{n,k})$. To this end, for any bounded measurable function h , write

$$\frac{1}{n} \sum_{k=1}^n \mathbb{E}[h(Y^{k,k}, \mu^{n,k})] = \frac{1}{n} \sum_{k=1}^n \mathbb{E}[\zeta_T^{n,k} h(X^k, \mu^n)].$$

The limiting behavior of this expression can be identified by studying the $(d + 1)$ -dimensional particle system $(X^k, \zeta^{n,k})_{k=1}^n$, adapting the classical martingale approach for McKean–Vlasov systems mentioned in Section 2.6.1.

3. Relaxed equilibria. Our proofs will make heavy use of *relaxed* or *randomized controls*, essentially replacing A -valued controls with $\mathcal{P}(A)$ -valued controls, which by now have a long history in stochastic optimal control theory [28, 34] for their useful compactness properties. Relaxed controls were employed in an MFG context [18, 50, 51], and we will use them in the same way. It is worth noting, however, that while they are certainly mathematically convenient, relaxed controls also admit a natural interpretation in a game-theoretic context as *mixed strategies*.

3.1. Relaxed n -player games. We begin by extending the equilibrium concepts for n -player games of Section 2.1. Write \mathcal{R}_n for the set of progressively measurable functions $\Lambda : [0, T] \times (\mathbb{C}^d)^n \rightarrow \mathcal{P}(A)$, and let \mathcal{RM}_n denote the subset of functions of the form $\Lambda(t, x) = \tilde{\Lambda}(t, x_t)$ for some measurable function $\tilde{\Lambda} : [0, T] \times (\mathbb{R}^d)^n \rightarrow \mathcal{P}(A)$. Via the embedding $A \ni a \mapsto \delta_a \in \mathcal{P}(A)$, we may view \mathcal{A}_n and \mathcal{AM}_n as subsets of \mathcal{R}_n , and we have the following natural inclusions:

$$\mathcal{AM}_n \subset \mathcal{A}_n \subset \mathcal{R}_n, \quad \mathcal{AM}_n \subset \mathcal{RM}_n \subset \mathcal{R}_n.$$

The state process and objective functions are defined for relaxed controls $\mathbf{\Lambda} = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{R}_n$ as follows:

$$dX_t^i = \int_A b(t, X_t^i, \mu_t^n, a) \Lambda^i(t, \mathbf{X})(da) dt + dW_t^i,$$

$$J_i^n(\Lambda^1, \dots, \Lambda^n) = \mathbb{E} \left[\int_0^T \int_A f(t, X_t^i, \mu_t^n, a) \Lambda^i(t, \mathbf{X})(da) dt + g(X_T^i, \mu_T^n) \right].$$

We may write $\mathbf{X}[\mathbf{\Lambda}] = (X^1[\mathbf{\Lambda}], \dots, X^n[\mathbf{\Lambda}])$ in place of $\mathbf{X} = (X^1, \dots, X^n)$ to stress which controls are being applied, and similarly we may write $\mu^n[\mathbf{\Lambda}]$ in place of μ^n .

DEFINITION 3.1. Let $\epsilon \geq 0$. A *relaxed closed-loop ϵ -Nash equilibrium* is a tuple $(\Lambda^1, \dots, \Lambda^n) \in \mathcal{R}_n^n$ such that

$$J_i^n(\Lambda^1, \dots, \Lambda^n) \geq \sup_{\beta \in \mathcal{R}_n} J_i^n(\Lambda^1, \dots, \Lambda^{i-1}, \beta, \Lambda^{i+1}, \dots, \Lambda^n) - \epsilon \quad \text{for } i = 1, \dots, n.$$

A *relaxed Markovian ϵ -Nash equilibrium* is a tuple $(\Lambda^1, \dots, \Lambda^n) \in \mathcal{RM}_n^n$ such that

$$J_i^n(\Lambda^1, \dots, \Lambda^n) \geq \sup_{\beta \in \mathcal{RM}_n} J_i^n(\Lambda^1, \dots, \Lambda^{i-1}, \beta, \Lambda^{i+1}, \dots, \Lambda^n) - \epsilon \quad \text{for } i = 1, \dots, n.$$

The following trio of propositions, along with Proposition 2.2, will show that the four equilibrium concepts described in Definitions 2.1 and 3.1 are roughly equivalent, if we accept both assumptions A and B. The proofs are given in Section 4.

PROPOSITION 3.2. *Suppose Assumption A holds, and let $\epsilon \geq 0$. Then any relaxed Markovian ϵ -Nash equilibrium is also a relaxed closed-loop ϵ -Nash equilibrium.*

PROPOSITION 3.3. *Suppose Assumptions A and B hold, and let $\epsilon \geq 0$. Then:*

- (a) *Any Markovian ϵ -Nash equilibrium is also a relaxed Markovian ϵ -Nash equilibrium.*
- (b) *Any closed-loop ϵ -Nash equilibrium is also a relaxed closed-loop ϵ -Nash equilibrium.*

PROPOSITION 3.4. *Suppose Assumptions A and B hold, and let $\epsilon \geq 0$. Then:*

- (a) *For any relaxed Markovian ϵ -Nash equilibrium $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{RM}_n^n$, there exists a Markovian ϵ -Nash equilibrium $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathcal{AM}_n^n$ such that $X[\Lambda] \stackrel{d}{=} X[\alpha]$.*
- (b) *For any relaxed closed-loop ϵ -Nash equilibrium $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{R}_n^n$, there exists a closed-loop ϵ -Nash equilibrium $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathcal{A}_n^n$ such that $X[\Lambda] \stackrel{d}{=} X[\alpha]$.*

Some notation helps to summarize the above propositions. Fix $\epsilon \geq 0$, let $\mathcal{A}_n^{*,\epsilon} \subset \mathcal{A}_n^n$ denote the set of closed-loop ϵ -Nash equilibria. Similarly, define $\mathcal{AM}_n^{*,\epsilon}$, $\mathcal{R}_n^{*,\epsilon}$, and $\mathcal{RM}_n^{*,\epsilon}$ respectively as the sets of Markovian, relaxed closed-loop, and relaxed Markovian ϵ -Nash equilibria. We may summarize the relations of Propositions 2.2, 3.2 and 3.3 by writing

$$\mathcal{AM}_n^{*,\epsilon} \subset \mathcal{A}_n^{*,\epsilon} \subset \mathcal{R}_n^{*,\epsilon}, \quad \mathcal{AM}_n^{*,\epsilon} \subset \mathcal{RM}_n^{*,\epsilon} \subset \mathcal{R}_n^{*,\epsilon}.$$

Moreover, we can think of Proposition 3.4(a) (resp. (b)) as reducing $\mathcal{AM}_n^{*,\epsilon} \subset \mathcal{RM}_n^{*,\epsilon}$ (resp. $\mathcal{A}_n^{*,\epsilon} \subset \mathcal{R}_n^{*,\epsilon}$) to equality, if we are content to focus only on the law of the state process X . Precisely, under Assumptions A and B, we have the following relationships between subsets of $\mathcal{P}((\mathbb{C}^d)^n)$:

$$\begin{aligned} \{ \mathcal{L}(X[\Lambda]) : \Lambda \in \mathcal{AM}_n^{*,\epsilon} \} &= \{ \mathcal{L}(X[\Lambda]) : \Lambda \in \mathcal{RM}_n^{*,\epsilon} \} \\ (3.1) \qquad \qquad \qquad &\subset \{ \mathcal{L}(X[\Lambda]) : \Lambda \in \mathcal{R}_n^{*,\epsilon} \} \\ &= \{ \mathcal{L}(X[\Lambda]) : \Lambda \in \mathcal{A}_n^{*,\epsilon} \}. \end{aligned}$$

Recall that our main result, Theorem 2.7, involves only the law of the state process X . Thanks to the above propositions, we may simultaneously cover all four of these possibilities by focusing solely on the laws of closed-loop Markovian equilibria, that is, $\mathcal{A}_n^{*,\epsilon}$.

We will make no claims throughout the paper regarding existence of equilibria for n -player games, but we provide some references. As we have mentioned, Markovian Nash equilibria (the set $\mathcal{AM}_n^{*,0}$, in the notation of the previous paragraph) are by far the most commonly studied in the literature and can be found by solving a system of n Hamilton–Jacobi–Bellman (HJB) equations. Relaxed Markovian equilibria are far less common, but the notable paper of

Borkar and Ghosh [6] has several theorems on existence (i.e., $\mathcal{RM}_n^{*,0} \neq \emptyset$). While their discussion of finite horizon problems is limited to the final sentence of the paper, it is clear that the techniques they develop for infinite-horizon problems can be easily adapted. Closed-loop path-dependent equilibria have appeared with some frequency in the literature on two-player stochastic differential games [12, 42]. They arise quite naturally in the BSDE-based weak formulation of Hamadene–Lepeltier [42], which reduces the existence of Nash equilibria to the solution of a BSDE (which is nothing but the stochastic representation of the corresponding HJB equation). The extension to the n -player setting is written in the lecture notes [13], Section 5.3.2, but be careful that our notion of closed-loop equilibrium is called “open-loop” therein. Lastly, we are unaware of any discussion of *relaxed closed-loop* equilibria in prior literature, but it is useful at the very least as an intermediary in establishing the relations in (3.1).

3.2. Relaxed mean field equilibria. We next extend the MFG equilibrium concepts (Definitions 2.3 and 2.5) of Section 2.2 to the relaxed setting:

DEFINITION 3.5. We say that $m = (m_t)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is a *strong relaxed mean field equilibrium* (or simply a *strong RMFE*) if there exists a measurable function $\Lambda^* : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$ such that the unique solution of the SDE

$$dX_t^* = \int_A b(t, X_t^*, m_t, a) \Lambda^*(t, X_t^*)(da) dt + dW_t, \quad X_0^* \sim \lambda$$

satisfies the following:

- (1) The consistency condition holds: $m_t = \mathcal{L}(X_t)$ for all $t \in [0, T]$.
- (2) For any measurable function $\Lambda : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$, we have

$$\begin{aligned} & \mathbb{E} \left[\int_0^T \int_A f(t, X_t^*, m_t, a) \Lambda^*(t, X_t^*)(da) dt + g(X_T^*, m_T) \right] \\ & \geq \mathbb{E} \left[\int_0^T \int_A f(t, X_t, m_t, a) \Lambda(t, X_t)(da) dt + g(X_T, m_T) \right], \end{aligned}$$

where X is the unique solution of

$$dX_t = \int_A b(t, X_t, m_t, a) \Lambda(t, X_t)(da) dt + dW_t, \quad X_0 \sim \lambda.$$

It was shown in [50], Theorem 6.2, that a strong MFE exists under Assumption A, though we will not need this fact. Recall from Definition 2.4 the notion of a *semi-Markov function*.

DEFINITION 3.6. A *weak semi-Markov relaxed mean field equilibrium* (or simply a *weak RMFE*) is a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X^*, \mu)$, where $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a complete filtered probability space and:

- (1) μ is a continuous \mathbb{F} -adapted $\mathcal{P}(\mathbb{R}^d)$ -valued process, W is a \mathbb{F} -Brownian motion, and X^* is a continuous \mathbb{R}^d -valued \mathbb{F} -adapted process with $\mathbb{P} \circ (X_0^*)^{-1} = \lambda$.
- (2) $\Lambda^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ is semi-Markov.
- (3) X_0^* , μ , and W are independent.
- (4) The state equation holds:

$$(3.2) \quad dX_t^* = \int_A b(t, X_t^*, \mu_t, a) \Lambda^*(t, X_t^*, \mu)(da) dt + dW_t.$$

(5) For every alternative $\Lambda : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ satisfying (2), we have

$$\begin{aligned} & \mathbb{E} \left[\int_0^T \int_A f(t, X_t^*, \mu_t, a) \Lambda^*(t, X_t^*, \mu)(da) dt + g(X_T^*, \mu_T) \right] \\ & \geq \mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda(t, X_t, \mu)(da) dt + g(X_T, \mu_T) \right], \end{aligned}$$

where X is the solution (recall Remark 2.6) of

$$dX_t = \int_A b(t, X_t, \mu_t, a) \Lambda(t, X_t, \mu)(da) dt + dW_t, \quad X_0 \sim \lambda.$$

(6) The consistency condition holds: $\mu_t = \mathbb{P}(X_t^* \in \cdot \mid \mathcal{F}_t^\mu)$ a.s. for each $t \in [0, T]$, where $\mathcal{F}_t^\mu = \sigma(\mu_s : s \leq t)$.

We refer also to the process μ itself as a weak RMFE. More precisely, we say a continuous $\mathcal{P}(\mathbb{R}^d)$ -valued process μ is a weak RMFE if it is defined within some tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X^*, \mu)$ which meets the above requirements.

Similar to the relationships of Section 3.1, under Assumptions A and B we prove in Section 4 that MFE and relaxed MFE induce the same measure flows:

PROPOSITION 3.7. *Suppose Assumptions A and B hold. Then every strong RMFE is a strong MFE, and every strong MFE is a strong RFME. Similarly, a continuous $\mathcal{P}(\mathbb{R}^d)$ -valued processes μ is a weak MFE if and only if it is a weak RMFE.*

REMARK 3.8. Recall from Remark 2.6 that the SDEs in (3.2) admit unique strong solutions, and in particular X^* is necessarily adapted to the complete filtration generated by the process $(X_0^*, W_s, \mu_s)_{s \leq t}$. With this and property (3) of Definition 3.6, we easily deduce that $\mathcal{L}(X_t^* \mid \mathcal{F}_t^\mu) = \mathcal{L}(X_t^* \mid \mu)$ a.s., for each $t \in [0, T]$, which will be useful later.

3.3. Extensions of the limit theorems. This section collects some generalizations of the main results announced in Section 2, which do not require the convexity Assumption B. The results of the previous two subsections show how the various equilibrium concepts related to each other if we impose Assumption B, and this is how we will deduce the results of Section 2 from those announced here.

THEOREM 3.9. *Suppose Assumption A holds. Fix a sequence $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$. For each n , suppose $\alpha^n = (\alpha^{n,1}, \dots, \alpha^{n,n}) \in \mathcal{A}_n^n$ is a closed-loop ϵ_n -Nash equilibrium. Then the associated empirical measure flow sequence $\mu^n = \mu^n[\alpha^n]$ is tight as a family of $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ -valued random variables, and every limit in distribution is a weak RMFE.*

If we impose both Assumption A and B, then Proposition 3.7 tells us that weak RMFE and weak MFE are one and the same. Thus, our main result, Theorem 2.7, follows from Theorem 3.9. Recall also the relations summarized in (3.1). Under Assumptions A and B, we deduce that Theorem 3.9 remains valid when α^n is instead assumed to be any of the four types of equilibrium described in Definitions 2.1 and 3.1.

Similarly, we may deduce the converse Theorem 2.12 from Proposition 3.4(a) the following generalization to relaxed equilibria:

THEOREM 3.10. *Suppose Assumptions A and C hold. Suppose $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is a strong RMFE. Then there exist $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$ and, for each n , a relaxed Markovian ϵ_n -Nash equilibrium $\Lambda^n \in \mathcal{RM}_n^n$ such that $\mu^n[\Lambda^n]$ converges in law to m in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$.*

Sections 5 and 7.1 are devoted to the proofs of Theorems 3.9 and 3.10, respectively. We might lastly state a form of Theorem 2.13 without Assumption B, as long as we use weak RMFE instead of weak MFE, but we opt not to write this out explicitly.

4. Relating the various equilibrium concepts. This section proves the various relationships between different equilibrium concepts of Definitions 2.1 and 3.1, announced in Propositions 2.2, 3.2, 3.3 and 3.4. We also prove Proposition 3.7, which relates MFE to RMFE.

4.1. *Proof of Proposition 3.2.* Let $\epsilon \geq 0$, and fix a relaxed Markovian ϵ -Nash equilibrium $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{RM}_n^n$. The goal is to show that Λ is also a relaxed closed-loop ϵ -Nash equilibrium. The state processes $X = X[\Lambda]$ solve the SDE system

$$dX_t^i = \int_A b(t, X_t^i, \mu_t^n, a) \Lambda^i(t, X_t)(da) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}.$$

Let $\beta \in \mathcal{R}_n$ be an alternative relaxed closed-loop control. We will focus on player 1, showing that

$$(4.1) \quad J_1^n(\Lambda^1, \dots, \Lambda^n) \geq J_1^n(\beta, \Lambda^2, \dots, \Lambda^n) - \epsilon.$$

The argument for other players $i \neq 1$ is identical. To proceed, let $Y = (Y^1, \dots, Y^n) := X[\beta, \Lambda^2, \dots, \Lambda^n]$ be the state processes in which players $i \neq 2$ still use Λ^i :

$$\begin{aligned} dY_t^1 &= \int_A b(t, Y_t^1, v_t^n, a) \beta(t, Y)(da) dt + dW_t^1, \\ dY_t^i &= \int_A b(t, Y_t^i, v_t^n, a) \Lambda^i(t, Y_t)(da) dt + dW_t^i, \quad i \neq 1, \\ v_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{Y_t^k}. \end{aligned}$$

Define $\tilde{\beta} \in \mathcal{RM}_n$ by setting $\tilde{\beta}(t, \mathbf{x}) = \mathbb{E}[\beta(t, Y) \mid Y_t = \mathbf{x}]$, noting that the existence of a jointly measurable version of this conditional mean measure is demonstrated by Lemma C.2. More precisely, this defines a Borel measurable function $\tilde{\beta} : [0, T] \times (\mathbb{R}^d)^n \rightarrow \mathcal{P}(A)$ such that

$$(4.2) \quad \int_A \varphi(t, Y_t, a) \tilde{\beta}(t, Y_t)(da) = \mathbb{E} \left[\int_A \varphi(t, Y_t, a) \beta(t, Y)(da) \mid Y_t \right] \quad \text{a.s., a.e. } t \in [0, T].$$

By Theorem 2.14, the unique solution $\tilde{Y} = (\tilde{Y}^1, \dots, \tilde{Y}^n)$ of the SDE system

$$\begin{aligned} d\tilde{Y}_t^1 &= \int_A b(t, \tilde{Y}_t^1, \tilde{v}_t^n, a) \tilde{\beta}(t, \tilde{Y}_t)(da) dt + dW_t^1, \\ d\tilde{Y}_t^i &= \int_A b(t, \tilde{Y}_t^i, \tilde{v}_t^n, a) \Lambda^i(t, \tilde{Y}_t)(da) dt + dW_t^i, \quad i \neq 1, \\ \tilde{v}_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{\tilde{Y}_t^k}, \end{aligned}$$

satisfies $\tilde{Y}_t \stackrel{d}{=} Y_t$ for all $t \in [0, T]$. Using Fubini's theorem and (4.2), we find

$$J_1^n(\beta, \Lambda^2, \dots, \Lambda^n) = \mathbb{E} \left[\int_0^T \int_A f(t, Y_t^1, v_t^n, a) \beta(t, Y)(da) dt + g(Y_T^1, v_T^n) \right]$$

$$\begin{aligned}
 &= \mathbb{E} \left[\int_0^T \int_A f(t, Y_t^1, v_t^n, a) \tilde{\beta}(t, \mathbf{Y}_t)(da) dt + g(Y_T^1, v_T^n) \right] \\
 &= \mathbb{E} \left[\int_0^T \int_A f(t, \tilde{Y}_t^1, \tilde{v}_t^n, a) \tilde{\beta}(t, \tilde{\mathbf{Y}}_t)(da) dt + g(\tilde{Y}_T^1, \tilde{v}_T^n) \right] \\
 &= J_1^n(\tilde{\beta}, \Lambda^2, \dots, \Lambda^n) \\
 &\leq J_1^n(\Lambda^1, \dots, \Lambda^n) + \epsilon.
 \end{aligned}$$

Indeed, the last inequality follows from the assumption that $(\Lambda^1, \dots, \Lambda^n)$ is a relaxed Markovian ϵ -Nash equilibrium.

4.2. *Proof of Proposition 3.3. Proof of (a):* Let $\epsilon \geq 0$, and fix a Markovian ϵ -Nash equilibrium $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathcal{AM}_n^n$. The goal is to show that α is also a relaxed Markovian ϵ -Nash equilibrium. The state processes $X = X[\alpha]$ solve the SDE system

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha^i(t, \mathbf{X}_t)) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}.$$

Let $\beta \in \mathcal{RM}_n$ be an alternative relaxed control. We will focus on player 1, showing that

$$(4.3) \quad J_1^n(\alpha^1, \dots, \alpha^n) \geq J_1^n(\beta, \alpha^2, \dots, \alpha^n) - \epsilon.$$

The argument for other players $i \neq 1$ is identical. To proceed, let $\mathbf{Y} = (Y^1, \dots, Y^n) = X[\beta, \alpha^2, \dots, \alpha^n]$ be the state processes in which players $i \neq 2$ still use α^i :

$$\begin{aligned}
 dY_t^1 &= \int_A b(t, Y_t^1, v_t^n, a) \beta(t, \mathbf{Y}_t)(da) dt + dW_t^1, \\
 dY_t^i &= b(t, Y_t^i, v_t^n, \alpha^i(t, \mathbf{Y}_t)) dt + dW_t^i, \quad i \neq 1, \\
 v_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{Y_t^k}.
 \end{aligned}$$

Recalling the definition of the convex set $K(t, x, m)$ from Assumption B, we have

$$\int_A (b(t, Y_t^1, v_t^n, a), f(t, Y_t^1, v_t^n, a)) \beta(t, \mathbf{Y}_t)(da) \in K(t, Y_t^1, v_t^n).$$

Let $L_n : (\mathbb{R}^d)^n \rightarrow \mathcal{P}(\mathbb{R}^d)$ denote the empirical measure map, $L_n(\mathbf{x}) = \frac{1}{n} \sum_{k=1}^n \delta_{x_k}$. Using a measurable selection theorem [43], Theorem A.9, we may find a measurable function $\tilde{\alpha} : [0, T] \times (\mathbb{R}^d)^n \rightarrow A$ such that, for each $\mathbf{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$,

$$(4.4) \quad b(t, x_1, L_n(\mathbf{x}), \tilde{\alpha}(t, \mathbf{x})) = \int_A b(t, x_1, L_n(\mathbf{x}), a) \beta(t, \mathbf{x})(da)$$

and

$$\int_A f(t, x_1, L_n(\mathbf{x}), a) \beta(t, \mathbf{x})(da) \leq f(t, x_1, L_n(\mathbf{x}), \tilde{\alpha}(t, \mathbf{x})).$$

The first of these identities implies that in fact

$$dY_t^1 = b(t, Y_t^1, v_t^n, \tilde{\alpha}(t, \mathbf{Y}_t)) dt + dW_t^1,$$

and in particular $\mathbf{Y} \stackrel{d}{=} X[\tilde{\alpha}, \alpha^2, \dots, \alpha^n]$, while the second implies

$$J_1^n(\beta, \alpha^2, \dots, \alpha^n) = \mathbb{E} \left[\int_0^T \int_A f(t, Y_t^1, v_t^n, a) \beta(t, \mathbf{Y}_t)(da) dt + g(Y_T^1, v_T^n) \right]$$

$$\begin{aligned} &\leq \mathbb{E} \left[\int_0^T f(t, Y_t^1, v_t^n, \tilde{\alpha}(t, Y_t)) dt + g(Y_T^1, v_T^n) \right] \\ &= J_1^n(\tilde{\alpha}, \alpha^2, \dots, \alpha^n) \\ &\leq J_1^n(\alpha^1, \dots, \alpha^n) + \epsilon. \end{aligned}$$

Indeed, the last inequality follows from the assumption that $(\alpha^1, \dots, \alpha^n)$ is a Markovian ϵ -Nash equilibrium.

Proof of (b): This proof is identical to that of part (a), except that all of the controls involved (namely, α^i and β) are closed-loop (path-dependent) instead of Markovian.

4.3. *Proof of Proposition 2.2.* This combines ideas of both of the previous proofs. Let $\epsilon \geq 0$, and fix a Markovian ϵ -Nash equilibrium $\alpha = (\alpha^1, \dots, \alpha^n) \in \mathcal{AM}_n^n$. The goal is to show that α is a closed-loop ϵ -Nash equilibrium. The state processes $X = X[\alpha]$ solve the SDE system

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha^i(t, X_t)) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}.$$

Let $\beta \in \mathcal{A}_n$ be an alternative closed-loop control. We will focus on player 1, showing that

$$(4.5) \quad J_1^n(\alpha^1, \dots, \alpha^n) \geq J_1^n(\beta, \alpha^2, \dots, \alpha^n) - \epsilon.$$

Let $Y = (Y^1, \dots, Y^n) = X[\beta, \alpha^2, \dots, \alpha^n]$ be the state processes in which players $i \neq 1$ still use α^i :

$$\begin{aligned} dY_t^1 &= b(t, Y_t^1, v_t^n, \beta(t, Y)) dt + dW_t^1, \\ dY_t^i &= b(t, Y_t^i, v_t^n, \alpha^i(t, Y_t)) dt + dW_t^i, \quad i \neq 1, \\ v_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{Y_t^k}. \end{aligned}$$

Recalling the definition of the convex set $K(t, x, m)$ from Assumption B, notice that

$$\mathbb{E}[b(t, Y_t^1, v_t^n, \beta(t, Y)) \mid Y_t] \in K(t, Y_t^1, v_t^n) \quad \text{a.s.},$$

for each $t \in [0, T]$. Using a measurable selection theorem [43], Theorem A.9, we may find a measurable function $\tilde{\beta} : [0, T] \times (\mathbb{R}^d)^n \rightarrow A$ such that

$$(4.6) \quad b(t, Y_t^1, v_t^n, \tilde{\beta}(t, Y_t)) = \mathbb{E}[b(t, Y_t^1, v_t^n, \beta(t, Y)) \mid Y_t] \quad \text{a.s.},$$

$$(4.7) \quad f(t, Y_t^1, v_t^n, \tilde{\beta}(t, Y_t)) \geq \mathbb{E}[f(t, Y_t^1, v_t^n, \beta(t, Y)) \mid Y_t] \quad \text{a.s.},$$

for each $t \in [0, T]$. By Theorem 2.14, the unique solution $\tilde{Y} = (\tilde{Y}^1, \dots, \tilde{Y}^n)$ of the SDE system

$$\begin{aligned} d\tilde{Y}_t^1 &= b(t, \tilde{Y}_t^1, \tilde{v}_t^n, \tilde{\beta}(t, \tilde{Y}_t)) dt + dW_t^1, \\ d\tilde{Y}_t^i &= b(t, \tilde{Y}_t^i, \tilde{v}_t^n, \alpha^i(t, \tilde{Y}_t)) dt + dW_t^i, \quad i \neq 1, \\ \tilde{v}_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{\tilde{Y}_t^k}, \end{aligned}$$

satisfies $\tilde{Y}_t \stackrel{d}{=} Y_t$ for all $t \in [0, T]$. Note that $\tilde{Y} \stackrel{d}{=} X[\tilde{\beta}, \alpha^2, \dots, \alpha^n]$. Use Fubini's theorem and (4.7) to get

$$\begin{aligned} J_1^n(\beta, \alpha^2, \dots, \alpha^n) &= \mathbb{E} \left[\int_0^T f(t, Y_t^1, v_t^n, \beta(t, Y)) dt + g(Y_T^1, v_T^n) \right] \\ &\leq \mathbb{E} \left[\int_0^T f(t, Y_t^1, v_t^n, \tilde{\beta}(t, Y_t)) dt + g(Y_T^1, v_T^n) \right] \\ &= \mathbb{E} \left[\int_0^T f(t, \tilde{Y}_t^1, \tilde{v}_t^n, \tilde{\beta}(t, \tilde{Y}_t)) dt + g(\tilde{Y}_T^1, \tilde{v}_T^n) \right] \\ &= J_1^n(\tilde{\beta}, \alpha^2, \dots, \alpha^n) \\ &\leq J_1^n(\alpha^1, \dots, \alpha^n) + \epsilon. \end{aligned}$$

Indeed, the last inequality follows from the assumption that $(\alpha^1, \dots, \alpha^n)$ is a Markovian ϵ -Nash equilibrium.

4.4. *Proof of Proposition 3.4. Proof of (a):* Let $\epsilon \geq 0$, and fix a relaxed Markovian ϵ -Nash equilibrium $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{RM}_n^n$. The state processes $X = X[\Lambda]$ solve the SDE system

$$dX_t^i = \int_A b(t, X_t^i, \mu_t^n, a) \Lambda^i(t, X_t)(da) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}.$$

Let $L_n : (\mathbb{R}^d)^n \rightarrow \mathcal{P}(\mathbb{R}^d)$ denote the empirical measure map, $L_n(\mathbf{x}) = \frac{1}{n} \sum_{k=1}^n \delta_{x_k}$. Recalling the definition of $K(t, x, m)$ from Assumption B, it holds for each $t \in [0, T]$, $\mathbf{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$, and $i \in \{1, \dots, n\}$ that

$$\int_A (b(t, x_i, L_n(\mathbf{x}), a), f(t, x_i, L_n(\mathbf{x}), a)) \Lambda^i(t, \mathbf{x})(da) \in K(t, x_i, L_n(\mathbf{x})).$$

Using a measurable selection theorem [43], Theorem A.9, we may find a measurable function $\alpha^i : [0, T] \times (\mathbb{R}^d)^n \rightarrow A$ such that, for each $\mathbf{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$ and $t \in [0, T]$,

$$(4.8) \quad b(t, x_i, L_n(\mathbf{x}), \alpha^i(t, \mathbf{x})) = \int_A b(t, x_i, L_n(\mathbf{x}), a) \Lambda^i(t, \mathbf{x})(da)$$

and

$$\int_A f(t, x_i, L_n(\mathbf{x}), a) \Lambda^i(t, \mathbf{x})(da) \leq f(t, x_i, L_n(\mathbf{x}), \alpha^i(t, \mathbf{x})).$$

The first of these identities implies that in fact X solves the SDE system

$$dX_t^i = b(t, X_t^i, \mu_t^n, \alpha^i(t, X_t)) dt + dW_t^i,$$

that is, $X = X[\alpha^1, \dots, \alpha^n]$, while the second implies

$$\begin{aligned} J_i^n(\Lambda^1, \dots, \Lambda^n) &= \mathbb{E} \left[\int_0^T \int_A f(t, X_t^i, \mu_t^n, a) \Lambda^i(t, X_t)(da) dt + g(X_T^i, \mu_T^n) \right] \\ (4.9) \quad &\leq \mathbb{E} \left[\int_0^T f(t, X_t^i, \mu_t^n, \alpha^i(t, X_t)) dt + g(X_T^i, \mu_T^n) \right] \\ &= J_i^n(\alpha^1, \dots, \alpha^n). \end{aligned}$$

Now, let us show that $(\alpha^1, \dots, \alpha^n)$ is an ϵ -Nash equilibrium. Fix an alternative Markovian control $\beta \in \mathcal{AM}_n$. We will focus on player 1, showing that

$$(4.10) \quad J_1^n(\alpha^1, \dots, \alpha^n) \geq J_1^n(\beta, \alpha^2, \dots, \alpha^n) - \epsilon.$$

To proceed, let $Y = (Y^1, \dots, Y^n) = X[\beta, \Lambda^2, \dots, \Lambda^n]$ be the state processes in which players $i \neq 2$ still use the relaxed controls Λ^i :

$$\begin{aligned} dY_t^1 &= b(t, Y_t^1, v_t^n, \beta(t, Y_t)) dt + dW_t^1, \\ dY_t^i &= \int_A b(t, Y_t^i, v_t^n, a) \Lambda^i(t, Y_t)(da) dt + dW_t^i, \quad i \neq 1, \\ v_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{Y_t^k}. \end{aligned}$$

Using (4.8) we may write

$$dY_t^i = b(t, Y_t^i, v_t^n, \alpha^i(t, Y_t)) dt + dW_t^i, \quad i \neq 1,$$

that is, $Y \stackrel{d}{=} X[\beta, \alpha^2, \dots, \alpha^n]$. Hence, since $(\Lambda^1, \dots, \Lambda^n)$ is ϵ -Nash, we may use (4.9) to get

$$\begin{aligned} J_1^n(\beta, \alpha^2, \dots, \alpha^n) &= J_1^n(\beta, \Lambda^2, \dots, \Lambda^n) \leq J_1^n(\Lambda^1, \dots, \Lambda^n) + \epsilon \\ &\leq J_1^n(\alpha^1, \dots, \alpha^n) + \epsilon. \end{aligned}$$

Proof of (b): This proof is identical to that of part (a), except that all of the controls involved (namely, Λ^i, α^i and β) are closed-loop (path-dependent) instead of Markovian.

4.5. *Proof of Proposition 3.7.* We prove the claims only for weak MFE, as the strong MFE is a special case of a deterministic weak MFE. We begin with a preparatory argument. Let $\Lambda : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ be any semi-Markov function (recall Definition 2.4). Recalling the definition of the convex set $K(t, x, m)$ from Assumption B, note that for $(t, x, m) \in [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d))$ we have

$$\int_A (b(t, x, m_t, a), f(t, x, m_t, a)) \Lambda(t, x, m)(da) \in K(t, x, m_t).$$

Using a measurable selection theorem [43], Theorem A.9, we may find a semi-Markov function $\alpha^\Lambda : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$ such that, for each (t, x, m) ,

$$b(t, x, m_t, \alpha^\Lambda(t, x, m)) = \int_A b(t, x, m, a) \Lambda(t, x, m)(da)$$

and

$$\int_A f(t, x, m, a) \Lambda(t, x, m)(da) \leq f(t, x, m, \alpha^\Lambda(t, x, m)).$$

In particular, if X solves the SDE

$$(4.11) \quad dX_t = \int_A b(t, X_t, \mu_t, a) \Lambda(t, X_t, \mu)(da) dt + dW_t,$$

where W is a Brownian motion, μ is a continuous $\mathcal{P}(\mathbb{R}^d)$ -valued process, and (X_0, μ, W) are independent, then X also solves the SDE

$$dX_t = b(t, X_t, \mu_t, \alpha^\Lambda(t, X_t, \mu)) dt + dW_t,$$

and we have the inequality

$$(4.12) \quad \begin{aligned} &\mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda(t, X_t, \mu)(da) dt + g(X_T, \mu_T) \right] \\ &\leq \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \alpha^\Lambda(t, X_t, \mu)) dt + g(X_T, \mu_T) \right]. \end{aligned}$$

With this construction the proof is straightforward. We first show that a weak RMFE is a weak MFE. Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X^*, \mu)$ be a weak RMFE. It is then easy to check using the above facts that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^{\Lambda^*}, X^*, \mu)$ is a weak MFE. Conversely, let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^*, X^*, \mu)$ be a weak MFE. Define $\Lambda^*(t, x, m) := \delta_{\alpha^*(t, x, m)}$. It is clear that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X^*, \mu)$ satisfies properties (1–4) and (6) of Definition 3.6. To prove (5), let $\Lambda : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ denote any semi-Markov function, and let X solve the corresponding SDE (4.11). Combine property (5) of the definition of weak MFE (Definition 2.5) with (4.12) to get

$$\begin{aligned} & \mathbb{E} \left[\int_0^T \int_A f(t, X_t^*, \mu_t, a) \Lambda^*(t, X_t^*, \mu) (da) dt + g(X_T^*, \mu_T) \right] \\ &= \mathbb{E} \left[\int_0^T f(t, X_t^*, \mu_t, \alpha^*(t, X_t^*, \mu)) dt + g(X_T^*, \mu_T) \right] \\ &\geq \mathbb{E} \left[\int_0^T f(t, X_t, \mu_t, \alpha^\Lambda(t, X_t, \mu)) dt + g(X_T, \mu_T) \right] \\ &\geq \mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda(t, X_t, \mu) (da) dt + g(X_T, \mu_T) \right]. \end{aligned}$$

This completes the proof.

5. Proof of the main limit theorem. This section is devoted to the proof of Theorem 3.9, from which Theorem 2.7 follows (see Section 3.3). We break this up into three major steps. First, we show tightness, which is straightforward in the present context. Next, we identify the limiting dynamics, in the sense that we prove that properties (1–4) and (6) of Definition 3.6 hold at the limit. Lastly, we address the optimality condition (5).

In fact, before we prove Theorem 3.9, we will carry out the bulk of the analysis without using the fact that the n -player controls are given as ϵ_n -Nash equilibria. That is, much of the work of characterizing the limiting behavior can and should be done independently of the Nash property. Only at the end will we use the Nash property to produce an inequality, which is then passed to the limit to obtain the desired optimality condition.

In the following, we work with an arbitrary sequence of controls $(\alpha^{n,1}, \dots, \alpha^{n,n}) \in \mathcal{A}_n^n$. We write $\mathbf{X}^n = (X^{n,1}, \dots, X^{n,n})$ to denote the corresponding state process in the n -player game, which we now index by n for clarity, and which is determined as the unique in law solution of the SDE

$$(5.1) \quad dX_t^{n,i} = b(t, X_t^{n,i}, \mu_t^n, \alpha^{n,i}(t, \mathbf{X}^n)) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^{n,k}}.$$

As usual, $X_0^{n,1}, \dots, X_0^{n,n}$ are i.i.d. with law λ .

It is convenient in the following to work with probability measures on the path space rather the flows of probability measures on \mathbb{R}^d . To distinguish between the two, we will reserve bold font for the former. For $\mathbf{m} \in \mathcal{P}(\mathcal{C}^d)$, define for each $t \in [0, T]$ the marginal law $m_t = \mathbf{m} \circ [x \mapsto x_t]^{-1}$, and note that the map

$$\mathcal{P}(\mathcal{C}^d) \ni \mathbf{m} \mapsto m := (m_t)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$$

is continuous. Given $\mathbf{m} \in \mathcal{P}(\mathcal{C}^d)$, we will refer to this $m = (m_t)_{t \in [0, T]}$ as the *induced or corresponding measure flow*. To keep track of notation, we stick to the following rules:

- We use the Latin m for a deterministic measure and the Greek μ for a random measure.
- We use boldface for a measure on path space, $\mathbf{m} \in \mathcal{P}(\mathcal{C}^d)$, to distinguish it from a measure flow, written as $m = (m_t)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$.

- Starting in Section 5.4, we will encounter probability measures on the extended path space $\mathcal{C}^d \times \mathcal{V}$, with \mathcal{V} defined in Section 5.2. We denote such measures as $\overline{m} \in \mathcal{P}(\mathcal{C}^d \times \mathcal{V})$.

Define μ^n , a random element of $\mathcal{P}(\mathcal{C}^d)$, by

$$\mu^n = \frac{1}{n} \sum_{k=1}^n \delta_{X^{n,k}}.$$

In light of the previous discussion, if μ^n converges in law to some $\mathcal{P}(\mathcal{C}^d)$ -valued random variable μ , then the marginal flow $\mu^n = (\mu_t^n)_{t \in [0, T]}$ converges in law in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ to the corresponding marginal flow $\mu = (\mu_t)_{t \in [0, T]}$.

Let $C_c^\infty(\mathbb{R}^d)$ denote the set of smooth functions of compact support. We define the infinitesimal generator of the controlled process as follows: For $\varphi \in C_c^\infty(\mathbb{R}^d)$, define

$$(5.2) \quad L\varphi(t, x, m, a) := b(t, x, m, a) \cdot \nabla\varphi(x) + \frac{1}{2} \Delta\varphi(x),$$

for $(t, x, m, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A$.

5.1. *Tightness.* We first prove that (μ^n) is tight.

LEMMA 5.1. *The sequence (μ^n) is a tight family of $\mathcal{P}(\mathcal{C}^d)$ -valued random variables.*

PROOF. According to [63], (2.5), it suffices to show that the sequence of mean measures $(m^n) \subset \mathcal{P}(\mathcal{C}^d)$ is tight, where we define m^n for Borel sets $B \subset \mathcal{C}^d$ by

$$m^n(B) = \mathbb{E}[\mu^n(B)] = \frac{1}{n} \sum_{k=1}^n \mathbb{P}(X^{n,k} \in B).$$

Letting $\|b\|_\infty$ denote the minimal uniform bound on $|b|$, note that $|L\varphi|$ is pointwise bounded by the constant

$$C_\varphi := \|b\|_\infty \|\nabla\varphi\|_\infty + \frac{1}{2} \|\Delta\varphi\|_\infty.$$

By Itô’s formula, for every $\varphi \in C_c^\infty(\mathbb{R}^d)$ the process $(\varphi(X_t^k) + C_\varphi t)_{t \in [0, T]}$ is a submartingale. It follows from [62], Theorem 1.4.6, that $\{X^{n,k} : n \in \mathbb{N}, k = 1, \dots, n\}$ is a tight family of \mathcal{C}^d -valued random variables. Hence, (m^n) is tight. \square

5.2. *Relaxed controls.* Before we proceed to identify the dynamics of limit points of (μ^n) , we must first discuss a convenient topological space in which to view the controls. Let \mathcal{V} denote the set of measures q on $[0, T] \times A$ with first marginal equal to Lebesgue measure. Equip \mathcal{V} with the topology of weak convergence, and note that \mathcal{V} is a compact metric space because A is. See [51], Appendix A, for a summary of basic facts about this space and references.

Each $q \in \mathcal{V}$ may be identified with a measurable function $[0, T] \ni t \mapsto q_t \in \mathcal{P}(A)$, determined uniquely (up to a.e. equality) by $dtq_t(da) = q(dt, da)$. Similarly, a measurable $\mathcal{P}(A)$ -valued process $(\Lambda_t)_{t \in [0, T]}$ can be identified with the random element $\Lambda = dt\Lambda_t(da)$ of \mathcal{V} . It is known that one can construct a measurable version of the *canonical process* on \mathcal{V} . More precisely, suppose $\mathbb{F}^\mathcal{V} = (\mathcal{F}_t^\mathcal{V})_{t \in [0, T]}$ denotes the natural filtration, where for each $t \in [0, T]$ we define $\mathcal{F}_t^\mathcal{V}$ as the σ -field generated by the functions $\mathcal{V} \ni q \mapsto q(B) \in \mathbb{R}$, for Borel sets $B \subset [0, t] \times A$. Then there exists (see [50], Lemma 3.2) an $\mathbb{F}^\mathcal{V}$ -predictable process

$$(5.3) \quad \widehat{q} : [0, T] \times \mathcal{V} \rightarrow \mathcal{P}(A) \quad \text{such that} \quad \widehat{q}(t, q) = q_t, \text{ a.e. } t, \forall q \in \mathcal{V}.$$

In particular, the filtration generated by the process $(\widehat{q}(t, \cdot))_{t \in [0, T]}$ is precisely $\mathbb{F}^\mathcal{V}$. With this in mind, we are free to identify $\mathcal{P}(A)$ -valued processes and \mathcal{V} -valued random variables.

5.3. *Projection lemmas.* As a preparation for the next step of identifying the dynamics of the limiting measure flows, we begin with two projection arguments that will be useful again in later sections. The first is straightforward but worth summarizing, while the second hides some delicate measurability questions which are largely outsourced to the appendix. In the following, it is convenient to use the usual duality notation for integration:

$$\langle m, \varphi \rangle = \int \varphi dm.$$

LEMMA 5.2. *Suppose $(Y_t)_{t \in [0, T]}$ is a continuous stochastic process taking values in a Polish space E and defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Suppose $h : E \rightarrow \mathbb{R}$ is continuous, and suppose it holds that*

$$h(Y_t) = h(Y_0) + \int_0^t a_s ds \quad \text{a.s., for a.e. } t \in [0, T],$$

where $(a_t)_{t \in [0, T]}$ is some bounded measurable real-valued process. Suppose $\hat{a} : [0, T] \times C([0, T]; E) \rightarrow \mathbb{R}$ is a progressively measurable function satisfying

$$\hat{a}(t, Y) = \mathbb{E}[a_t \mid \mathcal{F}_t^Y] \quad \text{a.s., for a.e. } t \in [0, T],$$

where $\mathcal{F}_t^Y = \sigma(Y_s : s \leq t)$. Then

$$h(Y_t) = h(Y_0) + \int_0^t \hat{a}(s, Y) ds \quad \text{for all } t \in [0, T] \text{ a.s.}$$

PROOF. By continuity, we have

$$h(Y_t) = h(Y_0) + \int_0^t a_s ds \quad \text{for all } t \in [0, T] \text{ a.s.}$$

Hence, it holds a.s. that for almost every $t \in (0, T)$ we have

$$a_t = \lim_{\delta \downarrow 0} \frac{1}{\delta} (h(Y_t) - h(Y_{t-\delta})).$$

In particular, a_t is measurable with respect to the completion of \mathcal{F}_t^Y , which implies $a_t = \mathbb{E}[a_t \mid \mathcal{F}_t^Y] = \hat{a}(t, Y)$, a.s., for a.e. $t \in [0, T]$. Complete the proof by integrating this identity and using continuity of h and Y to interchange the order of quantifiers as needed. \square

In the following lemma, we show that a solution of a certain kind of randomized Fokker–Planck equation can be realized as the conditional law of the state process under a semi-Markov control. Recall the notion of *semi-Markov function* from Definition 2.4.

LEMMA 5.3. *Suppose $\bar{\mu}$ is a $\mathcal{P}(C^d \times \mathcal{V})$ -valued random variable, and let $\mu = (\mu_t)_{t \in [0, T]}$ denote the corresponding measure flow.⁷ Suppose it holds with probability 1 that for every $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$ we have*

$$(5.4) \quad \langle \mu_t, \varphi \rangle = \langle \mu_0, \varphi \rangle + \int_{C^d \times \mathcal{V}} \left[\int_0^t \int_A L\varphi(s, x_s, \mu_s, a) q_s(da) ds \right] \bar{\mu}(dx, dq).$$

Then there exists a semi-Markov function $\Lambda^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ such that the following hold:

⁷That is, the $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ -valued random variable μ is defined by $\mu_t := \bar{\mu} \circ [(x, q) \mapsto x_t]^{-1}$, for $t \in [0, T]$.

(a) It holds with probability 1 that for every $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$ we have

$$\langle \mu_t, \varphi \rangle = \langle \mu_0, \varphi \rangle + \int_0^t \left\langle \mu_s, \int_A L\varphi(s, \cdot, \mu_s, a) \Lambda^*(t, \cdot, \mu)(da) \right\rangle ds.$$

(b) For each bounded measurable function ψ on $[0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A$, we have

$$\begin{aligned} & \mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} \int_0^T \int_A \psi(t, x_t, \mu_t, a) q_t(da) dt \bar{\mu}(dx, dq) \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} \int_A \psi(t, x, \mu_t, a) \Lambda^*(t, x, \mu)(da) \mu_t(dx) dt \right]. \end{aligned}$$

(c) By enlarging the probability space, we may construct continuous d -dimensional processes X and W such that:

- (i) X_0, W and μ are independent.
- (ii) W is a Brownian motion with respect to the complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ generated by the process $(X_0, \mu_t, W_t)_{t \in [0, T]}$.
- (iii) X is a continuous process with $X_0 \sim \lambda$, adapted to the completion of \mathbb{F} .
- (iv) The state equation holds,

$$dX_t = \int_A b(t, X_t, \mu_t, a) \Lambda^*(t, X_t, \mu)(da) dt + dW_t.$$

(v) For each t , it holds a.s. that $\mu_t = \mathcal{L}(X_t \mid \mathcal{F}_t^\mu)$, where $\mathcal{F}_t^\mu = \sigma(\mu_s : s \leq t)$.

PROOF. We first justify (c), assuming we have already found Λ^* such that (a) and (b) hold. In fact, the claimed processes X and W come from the observation that property (a) is simply a randomized version of a Fokker–Planck equation. Corollary A.7 works out the details and shows that we can construct X and W satisfying properties (i–v).

To construct Λ^* satisfying (a) and (b), we note first that (5.4) rewrites as

$$\langle \mu_t, \varphi \rangle = \langle \mu_0, \varphi \rangle + \int_{\mathcal{C}^d \times \mathcal{V}} \left[\int_0^t \int_A L\varphi(s, x_s, \mu_s, a) \widehat{q}(s, q)(da) ds \right] \bar{\mu}(dx, dq),$$

where \widehat{q} is the “nice version” of the process $[0, T] \times \mathcal{V} \ni (t, q) \mapsto q_t \in \mathcal{P}(A)$ described in (5.3).

Suppose for concreteness that the random variable $\bar{\mu}$ is defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and we may assume without loss of generality that Ω is a Polish space and \mathcal{F} its Borel σ -field. We now use Lemma B.2 to construct a jointly measurable version of the regular condition law of (x, q) given x_t under the random probability measure $\bar{\mu}(dx, dq)$; precisely, there exists a jointly measurable map $[0, T] \times \mathbb{R}^d \times \Omega \ni (t, x, \omega) \mapsto \bar{\mu}_{t,x}(\omega) \in \mathcal{P}(\mathcal{C}^d \times \mathcal{V})$ such that it holds a.s. that for every bounded measurable function $h : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ and $F : [0, T] \times \mathcal{C}^d \times \mathcal{V} \rightarrow \mathbb{R}$ we have

$$\begin{aligned} & \int_{\mathcal{C}^d \times \mathcal{V}} \int_0^T h(t, x_t) F(t, x, q) dt \bar{\mu}(dx, dq) \\ &= \int_{\mathcal{C}^d \times \mathcal{V}} \int_0^T h(t, x_t) \langle \bar{\mu}_{t,x_t}, F(t, \cdot) \rangle dt \bar{\mu}(dx, dq) \quad \text{a.s.} \end{aligned}$$

Using Fubini’s theorem and a change of variables, we may rewrite this as

$$\begin{aligned} & \int_{\mathcal{C}^d \times \mathcal{V}} \int_0^T h(t, x_t) F(t, x, q) dt \bar{\mu}(dx, dq) \\ (5.5) \quad &= \int_0^T \int_{\mathbb{R}^d} h(t, x) \langle \bar{\mu}_{t,x}, F(t, \cdot) \rangle \mu_t(dx) dt \quad \text{a.s.} \end{aligned}$$

Applying this with $F(t, x, q) = \int_A L\varphi(t, x, \mu_t, a)\widehat{q}(t, q)(da)$, we may write (5.4) as

$$(5.6) \quad \begin{aligned} \langle \mu_t, \varphi \rangle &= \langle \mu_0, \varphi \rangle \\ &+ \int_0^t \int_{\mathbb{R}^d} \int_{\mathcal{C}^d \times \mathcal{V}} \int_A L\varphi(s, x, \mu_s, a)\widehat{q}(s, q)(da)\overline{\mu}_{t,x}(d\tilde{x}, dq)\mu_s(dx) ds \quad \text{a.s.} \end{aligned}$$

Next, recall that $\mathbb{F}^\mu = (\mathcal{F}_t^\mu)_{t \in [0, T]}$ is the filtration generated by the process $(\mu_t)_{t \in [0, T]}$. We may find (using Corollary C.3) a semi-Markov function $\Lambda^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ such that, for every bounded measurable function $\psi : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \rightarrow \mathbb{R}$ and every $(t, x) \in [0, T] \times \mathbb{R}^d$, we have

$$(5.7) \quad \begin{aligned} &\int_A \psi(t, x, \mu_t, \cdot) d\Lambda^*(t, x, \mu) \\ &= \mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} \int_A \psi(t, x, \mu_t, a)\widehat{q}(s, q)(da)\overline{\mu}_{t,x}(d\tilde{x}, dq) \mid \mathcal{F}_t^\mu \right]. \end{aligned}$$

Applying (5.7) with $\psi = L\varphi$, and using (5.6) and Lemma 5.2, we get

$$\langle \mu_t, \varphi \rangle = \langle \mu_0, \varphi \rangle + \int_0^t \int_{\mathbb{R}^d} \int_A L\varphi(s, x, \mu_s, a)\Lambda^*(t, x, \mu)(da)\mu_s(dx) ds,$$

for all $t \in [0, T]$, almost surely, for each $\varphi \in C_c^\infty(\mathbb{R}^d)$. This is exactly (a), once we interchange the order of the quantifiers “almost surely” and “for each $\varphi \in C_c^\infty(\mathbb{R}^d)$.” This is easily justified by working with a countable dense family of such φ .

Finally, to prove (b), fix ψ , and simply use (5.5) and (5.7) along with Fubini’s theorem:

$$\begin{aligned} &\mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} \int_0^T \int_A \psi(t, x_t, \mu_t, a)q_t(da) dt \overline{\mu}(dx, dq) \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} \int_{\mathcal{C}^d \times \mathcal{V}} \int_A \psi(t, x, \mu_t, a)\widehat{q}(t, q)(da)\overline{\mu}_{t,x}(d\tilde{x}, dq)\mu_t(dx) dt \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} \int_A \psi(t, x, \mu_t, a)\Lambda^*(t, x, \mu)(da)\mu_t(dx) dt \right]. \quad \square \end{aligned}$$

5.4. *Identification of limiting dynamics.* We next provide a first description of the dynamics of subsequential limit points of $\mu^{n_k} = (\mu_t^{n_k})_{t \in [0, T]}$.

THEOREM 5.4. *Suppose a subsequence $(\mu_t^{n_k})_{t \in [0, T]}$ converges in law in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ to $(\mu_t)_{t \in [0, T]}$. Then there exists a semi-Markov function $\Lambda^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ such that, by extending the probability space if needed, we may construct continuous d -dimensional processes X and W such that:*

- (i) X_0, W and μ are independent.
- (ii) W is a Brownian motion with respect to the complete filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ generated by the process $(X_0, \mu_t, W_t)_{t \in [0, T]}$.
- (iii) X is \mathbb{F} -adapted, with $X_0 \sim \lambda$.
- (iv) The following SDE holds:

$$dX_t = \int_A b(t, X_t, \mu_t, a)\Lambda^*(t, X_t, \mu)(da) dt + dW_t.$$

- (v) For each t , it holds a.s. that $\mu_t = \mathcal{L}(X_t \mid \mathcal{F}_t^\mu)$, where $\mathcal{F}_t^\mu = \sigma(\mu_s : s \leq t)$.

Moreover,

$$\begin{aligned}
 (5.8) \quad & \lim_{k \rightarrow \infty} \frac{1}{n_k} \sum_{i=1}^{n_k} J_i^{n_k}(\alpha^{n_k,1}, \dots, \alpha^{n_k,n_k}) \\
 &= \mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda^*(t, X_t, \mu)(da) dt + g(X_T, \mu_T) \right].
 \end{aligned}$$

PROOF. Let us view each control as a random element of \mathcal{V} , by defining

$$\Lambda^{n,i}(dt, da) = dt \delta_{\alpha^{n,i}(t, X^n)}(da),$$

and define the extended empirical measure $\bar{\mu}^n$, a $\mathcal{P}(\mathcal{C}^d \times \mathcal{V})$ -valued random variable, by

$$\bar{\mu}^n = \frac{1}{n} \sum_{k=1}^n \delta_{(X^{n,k}, \Lambda^{n,k})}.$$

Because the \mathcal{C}^d -marginal μ^n is tight by Lemma 5.1 and \mathcal{V} is compact, the sequence of random measures $\bar{\mu}^n$ is tight. We may then pass to a further subsequence and assume that $\bar{\mu}^n$ converges in law to some random element $\bar{\mu}$ of $\mathcal{P}(\mathcal{C}^d \times \mathcal{V})$ whose \mathcal{C}^d -marginal is μ .

Step 1: We first show that $\bar{\mu}$ must satisfy the hypothesis (5.4) of Lemma 5.3. Recall that $X^n = (X^{n,1}, \dots, X^{n,n})$ is the vector of state processes; see (5.1). Begin by applying Itô’s formula to $\varphi(X_t^{n,k})$ and averaging over $k = 1, \dots, n$ to get

$$\begin{aligned}
 d\langle \mu_t^n, \varphi \rangle &= \frac{1}{n} \sum_{k=1}^n L\varphi(t, X_t^{n,k}, \mu_t^n, \alpha^{n,k}(t, X^n)) dt + dM_t^{n,\varphi} \\
 &= \frac{1}{n} \sum_{k=1}^n \int_A L\varphi(t, X_t^{n,k}, \mu_t^n, a) \Lambda_t^{n,k}(da) dt + dM_t^{n,\varphi},
 \end{aligned}$$

where we define the martingale

$$M_t^{n,\varphi} = \frac{1}{n} \sum_{k=1}^n \nabla \varphi(X_s^{n,k}) \cdot dW_s^k.$$

Notice that the quadratic variation of this martingale is

$$(5.9) \quad [M^{n,\varphi}]_t = \frac{1}{n^2} \sum_{k=1}^n \int_0^t |\nabla \varphi(X_s^{n,k})|^2 ds \leq \frac{\|\nabla \varphi\|_\infty^2}{n}.$$

For a measure $\bar{m} \in \mathcal{P}(\mathcal{C}^d \times \mathcal{V})$, let $m = (m_t)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ denote the associated measure flow, and define $F_t : \mathcal{P}(\mathcal{C}^d \times \mathcal{V}) \rightarrow \mathbb{R}$ by

$$F_t^\varphi(\bar{m}) = \int_{\mathcal{C}^d \times \mathcal{V}} \left[\varphi(x_t) - \varphi(x_0) - \int_0^t \int_A L\varphi(s, x_s, m_s, a) q_s(da) ds \right] \bar{m}(dx, dq).$$

We may then write

$$M_t^{n,\varphi} = F_t^\varphi(\bar{\mu}^n).$$

It can be shown that F_t^φ is a bounded continuous function (see [50], Appendix A, for details). Because $\bar{\mu}^n$ converges in law to $\bar{\mu}$, we conclude from the continuous mapping theorem that $F_t^\varphi(\bar{\mu}^n)$ converges in law to $F_t^\varphi(\bar{\mu})$ (with convergence understood in both cases to be along the same subsequence as before). But (5.9) implies that $F_t^\varphi(\bar{\mu}^n) = M_t^{n,\varphi}$ converges in probability to zero. Hence,

$$F_t^\varphi(\bar{\mu}) = 0 \quad \text{almost surely, for each } t \in [0, T], \varphi \in C_c^\infty(\mathbb{R}^d).$$

For each $\bar{m} \in \mathcal{P}(\mathcal{C}^d \times \mathcal{V})$, it is clear that $\lim_n F_{t_n}^{\varphi_n}(\bar{m}) = F_t^\varphi(\bar{m})$ whenever $t_n \rightarrow t$ and $(\varphi_n, \nabla\varphi_n, \Delta\varphi_n) \rightarrow (\varphi, \nabla\varphi, \Delta\varphi)$ uniformly. Hence, working with a countable dense family, we may interchange the order of quantifiers and conclude that

$$F_t^\varphi(\bar{\mu}) = 0 \quad \text{for each } t \in [0, T], \varphi \in C_c^\infty(\mathbb{R}^d), \text{ almost surely.}$$

This shows that $\bar{\mu}$ satisfies (5.4).

Step 2. We now construct the processes X and W . Thanks to Step 1, we may apply Lemma 5.3 to find a semi-Markov function $\Lambda^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$ such that (a), (b) and (c) of Lemma 5.3 hold.

Step 3. To complete the proof, we address the final claim about convergence of value. Notice that

$$\begin{aligned} J_n &:= \frac{1}{n} \sum_{i=1}^n J_i^n(\alpha^{n,1}, \dots, \alpha^{n,n}) \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left[\int_0^T f(t, X_t^{n,i}, \mu_t^n, \alpha^{n,i}(t, \mathbf{X}^n)) dt + g(X_T^i, \mu_T^n) \right] \\ &= \mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} \left(\int_0^T \int_A f(t, x_t, \mu_t^n, a) q_t(da) dt + g(x_T, \mu_T^n) \right) \bar{\mu}^n(dx, dq) \right]. \end{aligned}$$

Recall that any subsequence contains a further subsequence along which $\bar{\mu}^n$ converges to some $\bar{\mu}$. Along such a subsequence, by boundedness and continuity of f and g , we find that J_n converges to

$$\mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} \left(\int_0^T \int_A f(t, x_t, \mu_t, a) q_t(da) dt + g(x_T, \mu_T) \right) \bar{\mu}(dx, dq) \right].$$

We claim that this is equal to the right-hand side of (5.8). Recalling that $(\mu_t)_{t \in [0, T]}$ is the marginal flow associated with $\bar{\mu}$ and also that $\mu_t = \mathcal{L}(X_t | \mathcal{F}_t^\mu)$ for each t , we may write the second term as

$$\mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} g(x_T, \mu_T) \bar{\mu}(dx, dq) \right] = \mathbb{E} \left[\int_{\mathbb{R}^d} g(x, \mu_T) \mu_T(dx) \right] = \mathbb{E}[g(X_T, \mu_T)].$$

To handle the first term, we use part (b) of Lemma 5.3 along with Fubini’s theorem and the identity $\mu_t = \mathcal{L}(X_t | \mathcal{F}_t^\mu)$ to write

$$\begin{aligned} &\mathbb{E} \left[\int_{\mathcal{C}^d \times \mathcal{V}} \int_0^T \int_A f(t, x_t, \mu_t, a) q_t(da) dt \bar{\mu}(dx, dq) \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} \int_A f(t, x, \mu_t, a) \Lambda^*(t, x, \mu)(da) \mu_t(dx) dt \right] \\ &= \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} \int_A f(t, X_t, \mu_t, a) \Lambda^*(t, X_t, \mu)(da) dt \right]. \end{aligned}$$

This completes the proof. \square

5.5. Optimality. The analysis carried out so far will allow us to check all of the properties of Definition 3.6 at the limit except for the optimality condition (5), and this section will complete this last task. Using Theorem 5.4, we work with a fixed weak limit $\mu = (\mu_t)_{t \in [0, T]}$, and we abuse notation by relabeling the subsequence with the same notation, so that $\mu^n = (\mu_t^n)_{t \in [0, T]} \rightarrow \mu$ weakly in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$. It is crucial to keep in mind that *for the rest of this section we are working with this particular limit point and this particular convergent subsequence.*

By Theorem 5.4, we may assume that μ is defined on a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, which supports d -dimensional processes X and W which satisfy properties (i–v) of Theorem 5.4 and equation (5.8), for some semi-Markov function $\Lambda^* : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$. Throughout this section, the notation $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ of this paragraph will stand.

Relative to this fixed random measure flow μ , we define on $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ the family of all possible alternative strategy choices. Let us write $\mathcal{R}_{\text{semi}}$ for the set of semi-Markov functions from $[0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d))$ to $\mathcal{P}(A)$. For any $\Lambda \in \mathcal{R}_{\text{semi}}$, let $X[\Lambda] = (X_t[\Lambda])_{t \in [0, T]}$ denote the unique strong solution (see Lemmas A.2 and A.3) of the SDE⁸

$$(5.10) \quad dX_t[\Lambda] = \int_A b(t, X_t[\Lambda], \mu_t, a) \Lambda(t, X_t[\Lambda], \mu)(da) dt + dW_t, \quad X_0[\Lambda] = X_0.$$

In this notation, note that $X[\Lambda^*] = X$. Define

$$J(\Lambda) := \mathbb{E} \left[\int_0^T \int_A f(t, X_t[\Lambda], \mu_t, a) \Lambda(t, X_t[\Lambda], \mu)(da) dt + g(X_T[\Lambda], \mu_T) \right].$$

The proof of Theorem 3.9 will be complete if we can show that

$$(5.11) \quad \sup_{\Lambda \in \mathcal{R}_{\text{semi}}} J(\Lambda) = J(\Lambda^*).$$

We accomplish this in two steps. The first and more straightforward step is to reduce the supremum to a nicer subset of $\mathcal{R}_{\text{semi}}$. Precisely, we will show

$$(5.12) \quad \sup_{\beta \in \mathcal{A}_{\text{semi}}^c} J(\beta) = \sup_{\Lambda \in \mathcal{R}_{\text{semi}}} J(\Lambda),$$

where we define $\mathcal{A}_{\text{semi}}^c$ to be the set of *continuous* semi-Markov functions $\beta : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$, which we view as a subset of $\mathcal{R}_{\text{semi}}$ by means of the usual embedding $A \ni a \mapsto \delta_a \in \mathcal{P}(A)$. Indeed, (5.12) follows from:

LEMMA 5.5. *For any $\Lambda \in \mathcal{R}_{\text{semi}}$, there exists a sequence $\beta^n \in \mathcal{A}_{\text{semi}}^c$ such that $(\mu, X[\Lambda^n])$ converges in law to $(\mu, X[\Lambda])$ and $J(\beta^n) \rightarrow J(\Lambda)$.*

Lastly, for each “nice” alternative control $\beta \in \mathcal{A}_{\text{semi}}^c$, we show that $J(\beta)$ is the limit of the average value of some sequence of admissible n -player controls, which is accomplished using the following crucial proposition:

PROPOSITION 5.6. *Let $\beta \in \mathcal{A}_{\text{semi}}^c$. For each n and each $k = 1, \dots, n$, define $\beta^{n,k} \in \mathcal{A}_n$ by*

$$(5.13) \quad \beta^{n,k}(t, \mathbf{x}) = \beta \left(t, x_t^k, \frac{1}{n} \sum_{j=1}^n \delta_{x^j} \right) \quad \text{for } t \in [0, T], \mathbf{x} = (x^1, \dots, x^n) \in (\mathcal{C}^d)^n.$$

Then (taking limits along the same subsequence described above)

$$(5.14) \quad \lim_n \frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,k-1}, \beta^{n,k}, \alpha^{n,k+1}, \dots, \alpha^{n,n}) = J(\beta).$$

⁸It is not important here that we are working with strong solutions, but it is notationally convenient to construct everything on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$.

With Proposition 5.6 in hand, let us see how to complete the proof of Theorem 3.9:

PROOF OF THEOREM 3.9. Using (5.12), for an arbitrary $\delta > 0$ we may find $\beta \in \mathcal{A}_{\text{semi}}^c$ such that

$$\sup_{\Lambda \in \mathcal{R}_{\text{semi}}} J(\Lambda) \leq J(\beta) + \delta.$$

To prove (5.11) it now suffices to show that $J(\beta) \leq J(\Lambda^*)$. Recall (5.8) from Theorem 5.4, which says

$$J(\Lambda^*) = \lim_n \frac{1}{n} \sum_{i=1}^n J_i^n(\alpha^{n,1}, \dots, \alpha^{n,n}),$$

where the limit is taken along the appropriate subsequence. On the other hand, defining $\beta^{n,k}$ as in Proposition 5.6, we have (5.14). Finally using the fact that $(\alpha^{n,1}, \dots, \alpha^{n,n})$ is a closed-loop ϵ_n -Nash equilibrium with $\epsilon_n \rightarrow 0$, we conclude that, along the same convergent subsequence,

$$\begin{aligned} J(\beta) &= \lim_n \frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,k-1}, \beta^{n,k}, \alpha^{n,k+1}, \dots, \alpha^{n,n}) \\ &\leq \lim_n \frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,n}) + \epsilon_n \\ &= J(\Lambda^*). \end{aligned}$$

The proof of Theorem 3.9 is thus complete. \square

PROOF OF LEMMA 5.5. *Step 1.* Before constructing the approximations, we show how to derive the claimed limits. Suppose $\Lambda^n \in \mathcal{R}_{\text{semi}}$, and assume that it holds for almost every $x \in \mathbb{R}^d$ and $\mathcal{L}(\mu)$ -almost every $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ that, for every bounded continuous function $\varphi : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}$,

$$\lim_{n \rightarrow \infty} \int_0^T \int_A \varphi(t, x, a) \Lambda^n(t, x, m)(da) dx dt = \int_0^T \int_A \varphi(t, x, a) \Lambda(t, x, m)(da) dt.$$

Consider the coefficients

$$\begin{aligned} B_n(t, x, m) &:= \int_A b(t, x, m_t, a) \Lambda^n(t, x, m)(da), \\ B(t, x, m) &= \int_A b(t, x, m_t, a) \Lambda(t, x, m)(da). \end{aligned}$$

For continuous $\varphi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with compact support, we have

$$\lim_{n \rightarrow \infty} \int_0^T \int_{\mathbb{R}^d} (B_n(t, x, m) - B(t, x, m)) \cdot \varphi(t, x) dx dt = 0,$$

for $\mathcal{L}(\mu)$ -almost every $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$. Using Lemma A.4, we conclude that $(\mu, X[\Lambda^n])$ converges in law to $(\mu, X[\Lambda])$. To conclude that $J(\Lambda^n) \rightarrow J(\Lambda)$ we would like to simply use the fact that f and g are bounded and continuous, but we must be careful about the fact that Λ and Λ^n may be discontinuous. Begin by writing

$$J(\Lambda^n) = \mathbb{E}[F_n(\mu, X[\Lambda^n])],$$

where we define $F_n : C([0, T]; \mathcal{P}(\mathbb{R}^d)) \times \mathcal{C}^d \rightarrow \mathbb{R}$ by

$$F_n(m, x) = \int_0^T \int_A f(t, x_t, m_t, a) \Lambda^n(t, x_t, m)(da) dt + g(x_T, m_T).$$

Define $F(x, m)$ similarly, with Λ in place of Λ^n , so that $J(\Lambda) = \mathbb{E}[F(\mu, X[\Lambda])]$. We know from Lemma A.4 that $\mathbb{E}[h(\mu, X[\Lambda^n])] \rightarrow \mathbb{E}[h(\mu, X[\Lambda])]$ for every bounded measurable function $h : C([0, T]; \mathcal{P}(\mathbb{R}^d)) \times \mathcal{C}^d \rightarrow \mathbb{R}$. On the other hand, we know by assumption that $F_n \rightarrow F$ pointwise. We may use a form of the dominated convergence theorem [61], Proposition 11.4.18, to conclude that $\mathbb{E}[F_n(\mu, X[\Lambda^n])] \rightarrow \mathbb{E}[F(\mu, X[\Lambda])]$.

Step 2. Next, we construct the desired approximations. Apply the well known ‘‘chattering lemma’’ (see, e.g., [29], Theorem 2.2, or [35], Theorem 4) to find a sequence of semi-Markov functions $\beta^n : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow A$ such that

$$dt \delta_{\beta^n(t,x,m)}(da) \rightarrow dt \Lambda(t, x, m)(da)$$

weakly (i.e., in \mathcal{V}) for each (x, m) . Hence, we may assume Λ is already of the form $\Lambda(t, x, m) = \delta_{\beta(t,x,m)}$ for some semi-Markov function β .

To complete the proof we use the fact that, since A is compact and convex, any measurable function from a Polish probability space into A is the a.e. limit of continuous functions (see, e.g., [18], Proposition C.1). By ‘‘Polish probability space’’ we mean a Polish space E equipped with a Borel probability measure. The only hurdle is that the Borel σ -field of the space $\Theta := [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is strictly larger than the one generated by semi-Markov functions, but this is not difficult to work around. Equip Θ with the probability measure Q defined for Borel sets $S \subset \Theta$ by

$$Q(S) = \frac{1}{T} \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} 1_S(t, x, \mu_t) \Phi_d(x) dx dt \right],$$

where Φ_d is the density of a standard d -dimensional Gaussian random variable. Define the map $\Pi : \Theta \rightarrow \Theta$ by

$$\Pi(t, x, m) = (t, x, m_{\cdot \wedge t}),$$

where $m_{\cdot \wedge t}$ denotes the path which follows m up to time t and is constant thereafter. Then Π is continuous, and the image $\Pi(\Theta)$ is closed. Moreover, the σ -field generated by Θ is precisely the one generated by the semi-Markov functions, and so any semi-Markov function $F : \Theta \rightarrow A$ factorizes through Π , in the sense that $F = F \circ \Pi$. The space $\Pi(\Theta)$ is a Polish space with the induced topology. Hence, as mentioend above, $\beta = \beta \circ \Pi$ is the $Q \circ \Pi^{-1}$ -a.e. limit of a sequence of continuous functions $\tilde{\beta}^n : \Pi(\Theta) \rightarrow A$. Define $\beta^n : \Theta \rightarrow A$ by $\beta^n = \tilde{\beta}^n \circ \Pi$. Then, β^n is continuous for each n , and $\beta^n \rightarrow \beta$ holds Q -a.e. \square

PROOF OF PROPOSITION 5.6. Recall that X^n solves the SDE (5.1). Define the state process

$$Y^{n,k} = (Y^{n,k,1}, \dots, Y^{n,k,n}) := X^n[(\alpha^{n,1}, \dots, \alpha^{n,k-1}, \beta^{n,k}, \alpha^{n,k+1}, \dots, \alpha^{n,n})].$$

Note that $Y^{n,k}$ follows the dynamics

$$\begin{aligned} dY_t^{n,k,k} &= b(t, Y_t^{n,k,k}, \mu_t^{n,k}, \beta(t, Y_t^{n,k,k}, \mu_t^{n,k})) dt + dW_t^k, \\ dY_t^{n,k,i} &= b(t, Y_t^{n,k,i}, \mu_t^{n,k}, \alpha^{n,i}(t, Y^{n,k})) dt + dW_t^i, \quad i \neq k, \\ \mu_t^{n,k} &= \frac{1}{n} \sum_{j=1}^n \delta_{Y^{n,k,j}}. \end{aligned}$$

Assume that X^n is defined on some filtered probability space $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n, \mathbb{P}^n)$, and we of course assume that the Brownian motions W^k from (5.1) are in fact \mathbb{F}^n -Brownian motions. Note Y^n may live on a different probability space which, to avoid complicating notation, we will not give a name. Recall from the second paragraph of Section 5.5 that we are working throughout this proof with a given (relabelled) subsequence along which $\mathcal{L}(\mu^n) = \mathbb{P}^n \circ (\mu^n)^{-1}$ converges in $\mathcal{P}(C([0, T]; \mathcal{P}(\mathbb{R}^d)))$ to $\mathcal{L}(\mu)$.

Step 1. It is convenient in this proof to work on a suitable canonical space, and the first step is simply to set up notation. Define an equivalent probability measure $\mathbb{Q}^{n,k}$ on $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n)$ by setting

$$\begin{aligned} \frac{d\mathbb{Q}^{n,k}}{d\mathbb{P}^n} &= \exp\left(\int_0^T (b(t, X_t^{n,k}, \mu_t^n, \beta(t, X_t^{n,k}, \mu^n)) - b(t, X_t^{n,k}, \mu_t^n, \alpha^{n,k}(t, X^n))) \cdot dW_t^k \right. \\ &\quad \left. - \frac{1}{2} \int_0^T |b(t, X_t^{n,k}, \mu_t^n, \beta(t, X_t^{n,k}, \mu^n)) - b(t, X_t^{n,k}, \mu_t^n, \alpha^{n,k}(t, X^n))|^2 dt\right). \end{aligned}$$

By Girsanov’s theorem and uniqueness of the SDEs, we have $\mathbb{Q}^{n,k} \circ (X^n)^{-1} = \mathcal{L}(Y^{n,k})$, and thus $\mathbb{Q}^{n,k} \circ (\mu^n, X^n)^{-1} = \mathcal{L}(\mu^{n,k}, Y^{n,k})$. Note also that we may write $d\mathbb{Q}^{n,k}/d\mathbb{P}^n = \zeta_T^{n,k}$, where we define $\zeta^{n,k}$ as the unique solution of the SDE

$$\begin{aligned} (5.15) \quad d\zeta_t^{n,k} &= \zeta_t^{n,k} \Xi_t^{n,k} \cdot dW_t^k, \quad \zeta_0^{n,k} = 1, \\ \Xi_t^{n,k} &:= b(t, X_t^{n,k}, \mu_t^n, \beta(t, X_t^{n,k}, \mu^n)) - b(t, X_t^{n,k}, \mu_t^n, \alpha^{n,k}(t, X^n)). \end{aligned}$$

We note for future use that boundedness of b easily yields the estimate

$$(5.16) \quad \sup_{n \in \mathbb{N}} \max_{k=1, \dots, n} \mathbb{E}[|d\mathbb{Q}^{n,k}/d\mathbb{P}^n|^p] = \sup_{n \in \mathbb{N}} \max_{k=1, \dots, n} \mathbb{E}[|\zeta_T^{n,k}|^p] < \infty,$$

for any $p > 0$. Moreover, we may write

$$\begin{aligned} (5.17) \quad &\frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,k-1}, \beta^{n,k}, \alpha^{n,k+1}, \dots, \alpha^{n,n}) \\ &= \frac{1}{n} \sum_{k=1}^n \mathbb{E} \left[\int_0^T f(t, Y_t^{n,k,k}, \mu_t^{n,k}, \beta(t, Y_t^{n,k,k}, \mu^{n,k})) dt + g(Y_T^{n,k,k}, \mu_T^{n,k}) \right] \\ &= \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{Q}^{n,k}} \left[\int_0^T f(t, X_t^{n,k}, \mu_t^n, \beta(t, X_t^{n,k}, \mu^n)) dt + g(X_T^{n,k}, \mu_T^n) \right] \\ (5.18) \quad &= \mathbb{E}^{\mathbb{P}^n} \left[\frac{1}{n} \sum_{k=1}^n \zeta_T^{n,k} \left(\int_0^T f(t, X_t^{n,k}, \mu_t^n, \beta(t, X_t^{n,k}, \mu^n)) dt + g(X_T^{n,k}, \mu_T^n) \right) \right]. \end{aligned}$$

We would like to show that the measure $\frac{1}{n} \sum_{k=1}^n \mathcal{L}(Y^{n,k,k}, \mu^{n,k})$ converges to $\mathcal{L}(X[\beta], \mu)$, along the same subsequence for which $\mathcal{L}(\mu^n)$ converges to $\mathcal{L}(\mu)$. Indeed, we could then pass to the limit directly in (5.17). The change of measure allows us to transform the expression into one involving the original μ^n and the particles $X^{n,k}$, as well as the new auxiliary particles $\zeta^{n,k}$. We will ultimately analyze the limiting behavior of the empirical measure of $(X^{n,k}, \zeta^{n,k}, W^k)_{k=1}^n$, as it is convenient to include the Brownian motion W^k as well.

Precisely, we proceed as follows. Define the $\mathcal{P}(A)$ -valued processes $\Lambda_t^{n,k} = \delta_{\alpha^{n,k}(t, X^n)}$, and view $\Lambda^{n,k}$ as a \mathcal{V} -valued random variable. Consider the extended empirical measure

$$R^n := \frac{1}{n} \sum_{k=1}^n \delta_{(X^{n,k}, \zeta^{n,k}, W^k, \Lambda^{n,k})},$$

viewed as a random variable with values in $\mathcal{P}(\overline{\Omega})$, where $\overline{\Omega} := \mathcal{C}^d \times \mathcal{C}_+^1 \times \mathcal{C}^d \times \mathcal{V}$. Here, $\mathcal{C}_+^1 := C([0, T]; \mathbb{R}_+)$ is the space of nonnegative one-dimensional continuous paths.

Step 2. We first show that the sequence $\{\mathbb{P}^n \circ (\mathbf{R}^n)^{-1} : n \in \mathbb{N}\} \subset \mathcal{P}(\mathcal{P}(\overline{\Omega}))$ is tight. According to [63], (2.5), it suffices to show that the sequence $\{M_n : n \in \mathbb{N}\} \subset \mathcal{P}(\overline{\Omega})$ of mean measures is tight, where the mean measure M_n is defined on Borel sets $S \subset \overline{\Omega}$ by

$$(5.19) \quad M_n(S) = \mathbb{E}^{\mathbb{P}^n}[\mathbf{R}^n(S)] = \frac{1}{n} \sum_{k=1}^n \mathbb{P}^n((X^{n,k}, \zeta^{n,k}, W^k, \Lambda^{n,k}) \in S).$$

To do this, it suffices to show that each marginal sequence is tight. Since \mathcal{V} is compact, the \mathcal{V} -marginal sequence is clearly tight. The third marginal of M_n is precisely Wiener measure; this sequence is constant and therefore tight. We saw in the proof of Lemma 5.1 that the sequence of first marginals

$$\frac{1}{n} \sum_{k=1}^n \mathbb{P}^n \circ (X^{n,k})^{-1}$$

is tight. Finally, we must check that the second marginal sequence

$$\frac{1}{n} \sum_{k=1}^n \mathbb{P}^n \circ (\zeta^{n,k})^{-1}$$

is tight. This is accomplished using Aldous’ criterion for tightness [48], Lemma 16.12. First, note that the estimate (5.16) implies by Doob’s inequality

$$(5.20) \quad \sup_{n \in \mathbb{N}} \max_{k=1, \dots, n} \mathbb{E}^{\mathbb{P}^n} \left[\sup_{t \in [0, T]} |\zeta_t^{n,k}|^p \right] < \infty.$$

Recalling that b is uniformly bounded, we have $|\Xi^{n,k}| \leq 2\|b\|_\infty$, where we recall the notation $\Xi^{n,k}$ from (5.15). For any $\delta > 0$ and any $[0, T - \delta]$ -valued stopping time, Itô’s isometry yields

$$\begin{aligned} \mathbb{E}^{\mathbb{P}^n} [|\zeta_{\tau+\delta}^{n,k} - \zeta_\tau^{n,k}|^2] &= \mathbb{E}^{\mathbb{P}^n} \left[\left| \int_\tau^{\tau+\delta} \zeta_t^{n,k} \Xi_t^{n,k} \cdot dW_t^k \right|^2 \right] = \mathbb{E}^{\mathbb{P}^n} \left[\int_\tau^{\tau+\delta} |\zeta_t^k|^2 |\Xi_t^{n,k}|^2 dt \right] \\ &\leq 4\delta \|b\|_\infty^2 \mathbb{E}^{\mathbb{P}^n} \left[\sup_{t \in [0, T]} |\zeta_t^k|^2 \right]. \end{aligned}$$

This converges to zero as $\delta \rightarrow 0$, uniformly in n, k and τ . This is enough to apply Aldous’ criterion and conclude that the second marginal sequence of M_n is tight, thus completing the proof that \mathbf{R}^n is a tight sequence of $\mathcal{P}(\overline{\Omega})$ -valued random variables.

Step 3. As a first step toward identifying the limit points of \mathbf{R}^n , by first showing that all limit points are supported on the set of solutions of a certain martingale problem. For the moment, fix $n \in \mathbb{N}$ and $k \in \{1, \dots, n\}$. For any $\varphi = \varphi(x, y, w) \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d)$, Itô’s formula yields

$$\begin{aligned} &d\varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) \\ &= \nabla_x \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) \cdot b(t, X_t^{n,k}, \mu_t^n, \alpha^{n,k}(s, \mathbf{X}^n)) dt + \frac{1}{2} \Delta_x \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) dt \\ &\quad + \frac{1}{2} \partial_{yy} \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) |\zeta_t^{n,k}|^2 |\Xi_t^{n,k}|^2 dt + \frac{1}{2} \Delta_w \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) dt \\ &\quad + \zeta_t^{n,k} (\nabla_x + \nabla_w) \partial_y \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) \cdot \Xi_t^{n,k} dt \\ &\quad + (\nabla_w \cdot \nabla_x) \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) dt + \partial_y \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) \zeta_t^{n,k} \Xi_t^{n,k} \cdot dW_t^k \\ &\quad + \nabla_x \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) \cdot dW_t^k + \nabla_w \varphi(X_t^{n,k}, \zeta_t^{n,k}, W_t^k) \cdot dW_t^k. \end{aligned}$$

Here we write $(\nabla_w \cdot \nabla_x)$ for the operator $\sum_{i=1}^d \partial_{w_i} \partial_{x_i}$. For $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$, $t \in [0, T]$, and $\varphi \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d)$, define a random variable $M_t[m, \varphi] : \overline{\Omega} \rightarrow \mathbb{R}$ by

$$M_t[m, \varphi](x, y, w, q) = \varphi(x_t, y_t, w_t) - \int_0^t \int_A \widehat{M}[m, \varphi](u, x_u, y_u, w_u, a) q_u(da) du,$$

where, for $(t, x, y, w, a) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d \times A$, we define

$$\begin{aligned} &\widehat{M}[m, \varphi](t, x, y, w, a) \\ &= \nabla_x \varphi(x, y, w) \cdot b(t, x, m_t, a) + \frac{1}{2} \Delta_x \varphi(x, y, w) \\ &\quad + \frac{1}{2} \partial_{yy} \varphi(x, y, w) |y|^2 |b(t, x, m_t, \beta(t, x, m)) - b(t, x, m_t, a)|^2 \\ &\quad + y(\nabla_x + \nabla_w) \partial_y \varphi(x, y, w) \cdot [b(t, x, m_t, \beta(t, x, m)) - b(t, x, m_t, a)] \\ &\quad + \frac{1}{2} \Delta_w \varphi(x, y, w) + (\nabla_w \cdot \nabla_x) \varphi(x, y, w). \end{aligned}$$

Under \mathbb{P}^n , the above calculation shows that the process

$$M_t^{n,k,\varphi} := M_t[\mu^n, \varphi](X^{n,k}, \zeta^{n,k}, W^k, \Lambda^{n,k})$$

is a martingale. Moreover, the cross-variation $[M^{n,k,\varphi}, M^{n,j,\varphi}]$ vanishes for $j \neq k$.

To completely specify a martingale problem, we equip $\overline{\Omega}$ with a canonical filtration $\overline{\mathbb{F}} = (\overline{\mathcal{F}}_t)_{t \in [0, T]}$. Precisely, this is defined by letting $\overline{\mathcal{F}}_t$ be the σ -field generated by the maps $\overline{\Omega} \ni (x, y, w, q) \mapsto (x_s, y_s, w_s, \widehat{q}(s, \cdot)) \in \mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d \times \mathcal{P}(A)$, for $s \leq t$, where \widehat{q} is the version of the canonical $\mathcal{P}(A)$ -valued process on \mathcal{V} described in (5.3).

For $s < t$ and any continuous $\overline{\mathcal{F}}_s$ -measurable function $h : \overline{\Omega} \rightarrow \mathbb{R}$ bounded in absolute value by 1, define $F[h, \varphi, s, t] : \mathcal{P}(\overline{\Omega}) \rightarrow \mathbb{R}$ by

$$F[h, \varphi, s, t](\mathbf{R}) = |(\mathbf{R}, (M_t[R^x, \varphi] - M_s[R^x, \varphi])h)|^2,$$

where, for $\mathbf{R} \in \mathcal{P}(\overline{\Omega}) = \mathcal{P}(C^d \times C_+^1 \times C^d \times \mathcal{V})$, we write R^x to denote the induced measure flow $R^x = (R_t^x)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ induced by the first C^d -marginal of \mathbf{R} . That is, $R_t^x = \mathbf{R} \circ [(x, y, w, q) \mapsto x_t]^{-1}$. Because b and β are continuous by assumption, the map

$$C([0, T]; \mathcal{P}(\mathbb{R}^d)) \times \overline{\Omega} \ni (m, x, y, w, q) \mapsto \int_0^t \int_A \widehat{M}[m, \varphi](u, x_u, y_u, w_u, a) q_u(da) du$$

is continuous for each t and φ (see, e.g., [50], Appendix A, for details). We would immediately deduce that $F[h, \varphi, s, t]$ is continuous on $\mathcal{P}(\overline{\Omega})$, except that \widehat{M} is unbounded due to the multiplication by $|y|^2$. To deal with this, abbreviate $F = F[h, \varphi, s, t]$, and define for $r > 0$

$$F^r[h, \varphi, s, t](\mathbf{R}) = |(\mathbf{R}, (M_t^r[R^x, \varphi] - M_s^r[R^x, \varphi])h)|^2,$$

where $M_t^r[m, \varphi](x, y, w, q) := M_t[m, \varphi](x, y \wedge r, w, q)$. Then $M^r[m, \varphi]$ is uniformly bounded for each r and φ . Using (5.20), it is straightforward to check that

$$(5.21) \quad \lim_{r \rightarrow \infty} \sup_{n \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^n} [|F[h, \varphi, s, t](\mathbf{R}^n) - F^r[h, \varphi, s, t](\mathbf{R}^n)|^2] = 0.$$

Note that $F^r[h, \varphi, s, t]$ is bounded and continuous on $\mathcal{P}(\overline{\Omega})$.

Now, recalling that the sequence \mathbf{R}^n is tight by Step 2, we may suppose that it converges in law (along a subsequence) to some $\mathcal{P}(\overline{\Omega})$ -valued random variable \mathbf{R} . Use (5.21) to conclude that $\mathbb{P}^n \circ (F[h, \varphi, s, t](\mathbf{R}^n))^{-1}$ converges to $\mathcal{L}(F[h, \varphi, s, t](\mathbf{R}))$. Then, using Fatou's lemma, the fact that $M^{n,k,\varphi}$ and $M^{n,j,\varphi}$ define orthogonal martingales, and the fact that h is

$\overline{\mathcal{F}}_s$ -measurable and bounded in absolute value by 1, we find (taking limits along the same subsequence)

$$\begin{aligned} \mathbb{E}[F[h, \varphi, s, t](\mathbf{R})] &\leq \liminf_n \mathbb{E}^{\mathbb{P}^n} [F[h, \varphi, s, t](\mathbf{R}^n)] \\ &= \liminf_n \mathbb{E}^{\mathbb{P}^n} \left[\left| \frac{1}{n} \sum_{k=1}^n (M_t^{n,k,\varphi} - M_s^{n,k,\varphi}) h(X^{n,k}, \zeta^{n,k}, W^k, \Lambda^{n,k}) \right|^2 \right] \\ &= \liminf_n \frac{1}{n^2} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}^n} [|M_t^{n,k,\varphi} - M_s^{n,k,\varphi}|^2 h^2(X^{n,k}, \zeta^{n,k}, W^k, \Lambda^{n,k})] \\ &\leq \liminf_n \frac{1}{n^2} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}^n} [|M_t^{n,k,\varphi} - M_s^{n,k,\varphi}|^2]. \end{aligned}$$

Finally, noting that

$$\begin{aligned} &\mathbb{E}^{\mathbb{P}^n} [|M_t^{n,k,\varphi} - M_s^{n,k,\varphi}|^2] \\ &= \mathbb{E}^{\mathbb{P}^n} \int_s^t |(\nabla_w + \nabla_x)\varphi(X_u^{n,k}, \zeta_u^{n,k}, W_u^k) + \partial_y\varphi(X_u^{n,k}, \zeta_u^{n,k}, W_u^k)\zeta_u^{n,k} \Xi_u^k|^2 du, \end{aligned}$$

we use (5.20) and boundedness of b to get

$$\sup_{n \in \mathbb{N}} \max_{k=1, \dots, n} \mathbb{E}^{\mathbb{P}^n} [|M_t^{n,k,\varphi} - M_s^{n,k,\varphi}|^2] < \infty.$$

Hence,

$$\mathbb{E}[F[h, \varphi, s, t](\mathbf{R})] = 0.$$

In particular, $F[h, \varphi, s, t](\mathbf{R}) = 0$ a.s. for each h, φ, s, t .

By working with a countably dense family (as in the end of Step 1 of the proof of Theorem 5.4), we may switch the order of quantifiers to conclude that it holds with probability 1 that, for all (h, φ, s, t) , $F[h, \varphi, s, t](\mathbf{R}) = 0$. Recalling the definition of $F[h, \varphi, s, t]$, this means that \mathbf{R} is supported on the set $\mathcal{L} \subset \mathcal{P}(\overline{\Omega})$ consisting of those probability measures R such that:

- $(M_t[R^x, \varphi])_{t \in [0, T]}$ is an R -martingale, for each $\varphi \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d)$, where $R^x = (R_t^x)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ denotes the measure flow associated with the first marginal.
- $R \circ [(x, y, w, q) \mapsto (x_0, y_0, w_0)]^{-1} = \lambda \times \delta_1 \times \delta_0$.

Step 4. We now establish a key identity satisfied by the measures $R \in \mathcal{L}$ identified in the previous step. For $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ let $P^m \in \mathcal{P}(C^d)$ denote the law of the unique solution X^m of the SDE

$$dX_t^m = b(t, X_t^m, m_t, \beta(t, X_t^m, m)) dt + dB_t, \quad X_0^m \sim \lambda.$$

This defines a universally measurable map $C([0, T]; \mathcal{P}(\mathbb{R}^d)) \ni m \mapsto P^m \in \mathcal{P}(C^d)$, by Lemma A.1. Recall that universal measurability means here that, for any Borel set S in $\mathcal{P}(C^d)$, the set $\{m : P^m \in S\}$ is measurable with respect to the completion of the Borel σ -field of $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ by any Borel probability measure; see [5], Definition 7.18 and Chapter 7.7, for background. We claim that every $R \in \mathcal{L}$ satisfies

$$(5.22) \quad \int_{\overline{\Omega}} h(x) y_T R(dx, dy, dw, dq) = \langle P^{R^x}, h \rangle,$$

for bounded measurable functions h on C^d .

Fix $R \in \mathcal{L}$. We can construct, on some filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, an $\overline{\Omega}$ -valued random variable (X, ζ, W, Λ) with law R such that the process

$$\begin{aligned} &\varphi(X_t, \zeta_t, W_t) - \int_0^t \int_A \left[\nabla_x \varphi(X_s, \zeta_s, W_s) \cdot b(s, X_s, R_s^x, a) + \frac{1}{2} \Delta_x \varphi(X_s, \zeta_s, W_s) \right. \\ &\quad + \frac{1}{2} \partial_{yy} \varphi(X_s, \zeta_s, W_s) |y|^2 |b(s, X_s, R_s^x, \beta(s, X_s, R^x)) - b(s, X_s, R_s^x, a)|^2 \\ &\quad + y(\nabla_x + \nabla_w) \partial_y \varphi(X_s, \zeta_s, W_s) \cdot [b(t, X_s, R_s^x, \beta(s, X_s, R^x)) - b(s, X_s, R_s^x, a)] \\ &\quad \left. + \frac{1}{2} \Delta_w \varphi(X_s, \zeta_s, W_s) + (\nabla_w \cdot \nabla_x) \varphi(X_s, \zeta_s, W_s) \right] \Lambda_s(da) ds \end{aligned}$$

is a martingale for each $\varphi \in C_c^\infty(\mathbb{R}^d \times \mathbb{R}_+ \times \mathbb{R}^d)$. This is a bit different from the usual martingale problem framework because of the integration with respect to $\Lambda_s(da)$, so standard theory does not immediately tell us how to represent (X, ζ, W) as the solution of an SDE. But the work of El Karoui and Méléard [27] covers this situation by making use of the notion of *martingale measures*, in the sense of Walsh [68], and the reader is referred to either reference for precise definitions. According to [27], Theorem IV-2, by extending the probability space if needed, we may find a vector $M = (M^1, \dots, M^d)$ of orthogonal martingale measures $M^i = M^i(da, dt)$ on $A \times [0, T]$, each with intensity measure $\Lambda_t(da) dt$, such that the following hold, for $t \in [0, T]$:

$$\begin{aligned} dX_t &= \int_A b(t, X_t, R_t^x, a) \Lambda_t(da) dt + dW_t \\ dW_t &= \int_A M(da, dt) \quad \text{i.e.,} \quad W_t = \int_{A \times [0, t]} M(da, ds) = M(A \times [0, t]), \\ d\zeta_t &= \zeta_t dN_t, \end{aligned}$$

where we define the martingale N by

$$N_t = \int_{A \times [0, t]} (b(s, X_s, R_s^x, \beta(s, X_s, R^x)) - b(s, X_s, R_s^x, a)) \cdot M(da, ds).$$

The only fact we need to know about martingale measures is the following: For any bounded jointly functions $\varphi, \psi : [0, T] \times A \times \Omega \rightarrow \mathbb{R}^d$ (using the Borel σ -field on A and the \mathbb{F} -progressive σ -field on $[0, T] \times \Omega$), the processes $t \mapsto \int_{A \times [0, t]} \varphi(s, a) \cdot M(da, ds)$ and $t \mapsto \int_{A \times [0, t]} \psi(s, a) \cdot M(da, ds)$ are martingales with covariation process $\int_0^t \int_A \varphi(s, a) \cdot \psi(s, a) \Lambda_s(da) ds$. In particular, using this and Lévy’s characterization, we deduce that W is a Brownian motion.

Continuing to work on the same probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, define a change of measure by

$$\frac{d\mathbb{Q}}{d\mathbb{P}} := \zeta_T = \exp\left(N_T - \frac{1}{2}[N]_T\right).$$

By Girsanov’s theorem (e.g., in the general form of [59], Theorem III.39), the process $B = W - [W, N]$ is a \mathbb{Q} -Brownian motion, and we compute

$$B_t = W_t - \int_0^t \int_A (b(s, X_s, R_s^x, \beta(s, X_s, R^x)) - b(s, X_s, R_s^x, a)) \Lambda_s(da) ds.$$

Substitute this into the equation for X to get

$$(5.23) \quad dX_t = b(t, X_t, R_t^x, \beta(t, X_t, R^x)) dt + dB_t,$$

still with initial distribution $\mathbb{Q} \circ X_0^{-1} = \mathbb{P} \circ X_0^{-1} = \lambda$.

The SDE (5.23) has a unique in law solution, and its law is precisely $\mathbb{Q} \circ X^{-1} = P^{R^x}$, where P^m was defined for $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ at the beginning of this step. It then holds, for any bounded measurable $h : \mathcal{C}^d \rightarrow \mathbb{R}$, that

$$\int_{\Omega} h(x) y_T R(dx, dy, dw, dq) = \mathbb{E}^{\mathbb{P}}[h(X)\zeta_T] = \mathbb{E}^{\mathbb{Q}}[h(X)] = \langle P^{R^x}, h \rangle,$$

which establishes (5.22).

Step 5. We are finally ready to take limits. Recalling from Step 2 that R^n is a tight sequence, let R denote any weak limit. From Step 3 we know that R belongs almost surely to \mathfrak{L} . Recalling the identifications of Step 1, we may pass to the limit along the same subsequence along which R^n converges in law to R to get, using (5.22),

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}^n} [\zeta_T^{n,k} h(\mu^n, X^{n,k})] &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}^n} \left[\int_{\Omega} y_T h(\mu^n, x) R^n(dx, dy, dw, dq) \right] \\ &= \mathbb{E} \left[\int_{\Omega} y_T h(R^x, x) R(dx, dy, dw, dq) \right] \\ &= \mathbb{E}[\langle P^{R^x}, h(R^x, \cdot) \rangle], \end{aligned}$$

for any bounded continuous function h on $C([0, T]; \mathcal{P}(\mathbb{R}^d)) \times \mathcal{C}^d$. Recall that μ^n converges in law to μ , which implies that $R^x \stackrel{d}{=} \mu$. Hence,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}^n} [\zeta_T^{n,k} h(\mu^n, X^{n,k})] = \mathbb{E}[\langle P^{\mu}, h(\mu, \cdot) \rangle].$$

Recalling the notation from before the statement the Proposition, the process $X[\beta]$ solves the SDE

$$dX_t[\beta] = b(t, X_t[\beta], \mu_t, \beta(t, X_t[\beta], \mu)) dt + dW_t,$$

where $X_0 \sim \lambda$, W , and μ are independent. Lemma A.2 ensures that the conditional law of $X[\beta]$ given μ is precisely P^{μ} . In particular, $\mathbb{E}[\langle P^{\mu}, h(\mu, \cdot) \rangle] = \mathbb{E}[h(\mu, X[\beta])]$, and we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}^n} [\zeta_T^{n,k} h(\mu^n, X^{n,k})] = \mathbb{E}[h(\mu, X[\beta])].$$

Finally, recalling that f , g and β are continuous by assumption, we may finally return to (5.18) from Step 1 to complete the proof:

$$\begin{aligned} &\frac{1}{n} \sum_{k=1}^n J_k^n(\alpha^{n,1}, \dots, \alpha^{n,k-1}, \beta^{n,k}, \alpha^{n,k+1}, \dots, \alpha^{n,n}) \\ &= \mathbb{E}^{\mathbb{P}^n} \left[\frac{1}{n} \sum_{k=1}^n \zeta_T^{n,k} \left(\int_0^T f(t, X_t^{n,k}, \mu_t^n, \beta(t, X_t^{n,k}, \mu^n)) dt + g(X_T^{n,k}, \mu_T^n) \right) \right] \\ &= \mathbb{E} \left[\int_0^T f(t, X_t[\beta], \mu_t, \beta(t, X_t[\beta], \mu)) dt + g(X_T[\beta], \mu_T) \right] \\ &= J(\beta). \end{aligned} \quad \square$$

5.6. Proof of Theorem 2.9. In this section we prove Theorem 2.9, which asserts uniqueness of weak MFE. The key monotonicity assumption (v) is of course inspired by the work of Lasry and Lions [55], which argued via PDEs. The probabilistic proof below is quite similar

to that of [18], which worked with a different equilibrium concept, though we will see in Theorems 6.2 and 6.3 that our weak MFE concept coincides in a sense with the notion of weak MFG solution used in [18, 51]. We give a complete proof below in part to keep the paper self-contained. In addition, [18], Theorem 6.2, makes the stronger assumption that $b(t, x, m, a)$ and $f(t, x, m, a)$ are respectively affine and strictly concave in (x, a) jointly, which we relax here to be just in a .

For $i = 1, 2$, suppose $(\Omega^i, \mathcal{F}^i, \mathbb{F}^i, \mathbb{P}^i, W^i, \alpha^i, X^i, \mu^i)$ are weak MFE in the sense of Definition 2.5. Noting that (X_0^i, W^i, μ^i) are independent for each $i = 1, 2$, we can without loss of generality couple these weak MFE by assuming that they are defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, with the same Brownian motion and initial state. Precisely, W is an \mathbb{F} -Brownian motion, X_0 is \mathcal{F}_0 -measurable with law λ , and the random variables X_0, W, μ^1 , and μ^2 are independent, with X^i given as the unique strong solution of

$$dX_t^i = b(t, X_t^i, \alpha^i(t, X_t^i, \mu^i)) dt + dW_t, \quad X_0^i = X_0, i = 1, 2.$$

The MFG fixed point property entails that

$$\mu_t^i = \mathcal{L}(X_t^i | \mathcal{F}_t^{\mu^i}) = \mathcal{L}(X_t^i | \mu^i) \quad \text{a.s., } t \in [0, T].$$

(See Lemma A.1 for the latter identity.) Because (X^1, μ^1) and μ^2 are independent, and similarly with the indices 1 and 2 reversed, we deduce that

$$(5.24) \quad \mu_t^i = \mathcal{L}(X_t^i | \mu^1, \mu^2) \quad \text{a.s., } t \in [0, T], i = 1, 2.$$

Let us define also the auxiliary state processes

$$dY_t^{i,m} = b(t, Y_t^{i,m}, \alpha^i(t, Y_t^{i,m}, m)) dt + dW_t, \quad Y_0^{i,m} = X_0,$$

for $i = 1, 2$ and $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$. Now, for a fixed $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$, the function $(t, x, m') \mapsto \alpha^i(t, x, m)$ defines an admissible semi-Markov control, which happens to not depend on the measure argument; we deduce from the optimality of α^1 that

$$(5.25) \quad \mathbb{E} \left[\int_0^T (f_1(t, X_t^1, \mu_t^1) + f_2(t, X_t^1, \alpha^1(t, X_t^1, \mu_t^1))) dt + g(X_T^1, \mu_T^1) \right] \\ \geq \mathbb{E} \left[\int_0^T (f_1(t, Y_t^{2,m}, \mu_t^1) + f_2(t, Y_t^{2,m}, \alpha^2(t, Y_t^{2,m}, m))) dt + g(Y_T^{2,m}, \mu_T^1) \right].$$

According to Lemma A.2, if $M^i := \mathcal{L}(\mu^i)$, then for any bounded measurable function φ on $\mathcal{C}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d))$ we have

$$\mathbb{E}[\varphi(X^2, \mu^2)] = \mathbb{E}[\varphi(Y^{2,\mu^2}, \mu^2)] = \int \mathbb{E}[\varphi(Y^{2,m}, m)] M^2(dm).$$

In other words, the law of $Y^{i,m}$ is a version of the conditional law of X^i given $\mu^i = m$. Because X_0, W, μ^1 , and μ^2 are independent, we similarly have

$$(5.26) \quad \mathbb{E}[\varphi(X^2, \mu^2, \mu^1)] = \int \mathbb{E}[\varphi(Y^{2,m}, m, \mu^1)] M^2(dm),$$

for bounded measurable functions φ on $\mathcal{C}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d))^2$. Hence, integrating (5.25) with respect to $M^2(dm)$, we find

$$(5.27) \quad \mathbb{E} \left[\int_0^T (f_1(t, X_t^1, \mu_t^1) + f_2(t, X_t^1, \alpha^1(t, X_t^1, \mu_t^1))) dt + g(X_T^1, \mu_T^1) \right] \\ \geq \mathbb{E} \left[\int_0^T (f_1(t, X_t^2, \mu_t^1) + f_2(t, X_t^2, \alpha^2(t, X_t^2, \mu_t^2))) dt + g(X_T^2, \mu_T^1) \right].$$

Repeat the argument with the indices 1 and 2 reversed to get

$$\begin{aligned}
 (5.28) \quad & \mathbb{E} \left[\int_0^T (f_1(t, X_t^2, \mu_t^2) + f_2(t, X_t^2, \alpha^2(t, X_t^2, \mu^2))) dt + g(X_T^2, \mu_T^2) \right] \\
 & \geq \mathbb{E} \left[\int_0^T (f_1(t, X_t^1, \mu_t^2) + f_2(t, X_t^1, \alpha^2(t, X_t^1, \mu^1))) dt + g(X_T^1, \mu_T^2) \right].
 \end{aligned}$$

Add these two inequalities, apply (5.24), and rearrange to find

$$\begin{aligned}
 0 \leq & \mathbb{E} \left[\int_0^T \int_{\mathbb{R}^d} (f_1(t, x, \mu_t^1) - f_1(t, x, \mu_t^2)) (\mu_t^1 - \mu_t^2) (dx) dt \right] \\
 & + \mathbb{E} \left[\int_{\mathbb{R}^d} (g(x, \mu_T^1) - g(x, \mu_T^2)) (\mu_T^1 - \mu_T^2) (dx) \right].
 \end{aligned}$$

According to the monotonicity assumption (v), this must hold with equality. Hence, both (5.27) and (5.28) must hold with equality.

To complete the proof, it will suffice to show that the optimal semi-Markov control α^1 is essentially unique; that is, we will show that unless α^1 and α^2 agree in a certain sense then we can construct another control which achieves a strictly greater reward. To do this, we apply a projection argument. Fix $m^1, m^2 \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$. Let U be a random variable, independent of everything else, which takes the value 1 or 2 with equal probability. Then

$$dY_t^{U, m^U} = b(t, Y_t^{U, m^U}, \alpha^U(t, Y_t^{U, m^U}, m^U)) dt + dW_t, \quad Y_0^{U, m^U} = X_0.$$

Suppose we define

$$(5.29) \quad \widehat{\alpha}^{m^1, m^2}(t, x) = \mathbb{E}[\alpha^U(t, Y_t^{U, m^U}, m^U) \mid Y_t^{U, m^U} = x].$$

By Lemmas A.1 and B.2, $\widehat{\alpha}^{m^1, m^2}(t, x)$ is jointly semi-Markov, in the sense that it is jointly measurable and $\widehat{\alpha}^{\widetilde{m}^1, \widetilde{m}^2}(t, x) = \widehat{\alpha}^{m^1, m^2}(t, x)$ whenever $\widetilde{m}_s^i = m_s^i$ for all $s \leq t$ and $i = 1, 2$. By Markovian projection (Theorem 2.14), and because $b(t, x, a)$ is affine in a , we have $\mathcal{L}(Y_t^{U, m^U}) = \mathcal{L}(\widehat{X}_t^{m^1, m^2})$ for all t , where \widehat{X}^{m^1, m^2} is the unique strong solution of

$$d\widehat{X}_t^{m^1, m^2} = b(t, \widehat{X}_t^{m^1, m^2}, \widehat{\alpha}^{m^1, m^2}(t, \widehat{X}_t^{m^1, m^2})) dt + dW_t, \quad \widehat{X}_0^{m^1, m^2} = X_0.$$

By Fubini's theorem and Jensen's inequality, using the assumption that $f(t, x, m, a)$ is concave in a it follows that

$$\begin{aligned}
 (5.30) \quad & \mathbb{E} \left[\int_0^T f(t, \widehat{X}_t^{m^1, m^2}, m_t^1, \widehat{\alpha}^{m^1, m^2}(t, \widehat{X}_t^{m^1, m^2})) dt + g(\widehat{X}_T^{m^1, m^2}, m_T^1) \right] \\
 & = \mathbb{E} \left[\int_0^T f(t, Y_t^{U, m^U}, m_t^1, \widehat{\alpha}^{m^1, m^2}(t, Y_t^{U, m^U})) dt + g(Y_T^{U, m^U}, m_T^1) \right] \\
 & \geq \mathbb{E} \left[\int_0^T f(t, Y_t^{U, m^U}, m_t^1, \alpha^U(t, Y_t^{U, m^U}, m^U)) dt + g(Y_T^{U, m^U}, m_T^1) \right] \\
 & = \frac{1}{2} \mathbb{E} \left[\int_0^T f(t, Y_t^{1, m^1}, m_t^1, \alpha^1(t, Y_t^{1, m^1}, m^1)) dt + g(Y_T^{1, m^1}, m_T^1) \right] \\
 & \quad + \frac{1}{2} \mathbb{E} \left[\int_0^T f(t, Y_t^{2, m^2}, m_t^1, \alpha^2(t, Y_t^{2, m^2}, m^2)) dt + g(Y_T^{2, m^2}, m_T^1) \right].
 \end{aligned}$$

Integrate on both sides with respect to $M^1(dm^1)M^2(dm^2)$, recalling (5.26), to get

$$\begin{aligned} & \mathbb{E} \left[\int_0^T f(t, \widehat{X}_t^{\mu^1, \mu^2}, \mu_t^1, \widehat{\alpha}^{\mu^1, \mu^2}(t, \widehat{X}_t^{\mu^1, \mu^2})) dt + g(\widehat{X}_T^{\mu^1, \mu^2}, \mu_T^1) \right] \\ & \geq \frac{1}{2} \mathbb{E} \left[\int_0^T f(t, X_t^1, \mu_t^1, \alpha^1(t, X_t^1, \mu^1)) dt + g(X_T^1, \mu_T^1) \right] \\ & \quad + \frac{1}{2} \mathbb{E} \left[\int_0^T (f_1(t, X_t^2, \mu_t^1) + f_2(t, X_t^2, \alpha^2(t, X_t^2, \mu^2))) dt + g(X_T^2, \mu_T^1) \right]. \end{aligned}$$

Recalling that equality holds in (5.27), we deduce

$$(5.31) \quad \begin{aligned} & \mathbb{E} \left[\int_0^T f(t, \widehat{X}_t^{\mu^1, \mu^2}, \mu_t^1, \widehat{\alpha}^{\mu^1, \mu^2}(t, \widehat{X}_t^{\mu^1, \mu^2})) dt + g(\widehat{X}_T^{\mu^1, \mu^2}, \mu_T^1) \right] \\ & \geq \mathbb{E} \left[\int_0^T f(t, X_t^1, \mu_t^1, \alpha^1(t, X_t^1, \mu^1)) dt + g(X_T^1, \mu_T^1) \right]. \end{aligned}$$

On the other hand, because \widehat{X}^{μ^1, m^2} solves the SDE

$$d\widehat{X}_t^{\mu^1, m^2} = b(t, \widehat{X}_t^{\mu^1, m^2}, \widehat{\alpha}^{\mu^1, m^2}(t, \widehat{X}_t^{\mu^1, m^2})) dt + dW_t, \quad \widehat{X}_0^{\mu^1, m^2} = X_0,$$

and because $(t, x, m^1) \mapsto \widehat{\alpha}^{m^1, m^2}(t, x)$ is semi-Markov for each m^2 , we deduce from the optimality of α^1 that

$$\begin{aligned} & \mathbb{E} \left[\int_0^T f(t, \widehat{X}_t^{\mu^1, \mu^2}, \mu_t^1, \widehat{\alpha}^{\mu^1, m^2}(t, \widehat{X}_t^{\mu^1, m^2})) dt + g(\widehat{X}_T^{\mu^1, m^2}, \mu_T^1) \right] \\ & \leq \mathbb{E} \left[\int_0^T f(t, X_t^1, \mu_t^1, \alpha^1(t, X_t^1, \mu^1)) dt + g(X_T^1, \mu_T^1) \right]. \end{aligned}$$

Integrate with respect to $M^2(dm^2)$ to deduce the opposite inequality of (5.31), which thus becomes an equality. It follows that the application of Jensen’s inequality in (5.30) to the strictly concave function f must be an equality for $M^1 \times M^2$ -a.e. (m^1, m^2) , which implies that the random variable inside the conditional expectation in (5.29) must in fact be Y_t^{U, m^U} -measurable for a.e. t . That is,

$$\widehat{\alpha}^{m^1, m^2}(t, Y_t^{U, m^U}) = \alpha^U(t, Y_t^{U, m^U}, m^U),$$

a.s., for a.e. t , for $M^1 \times M^2$ -a.e. (m^1, m^2) . Recalling the definition of U , this is equivalent to

$$\widehat{\alpha}^{m^1, m^2}(t, Y_t^{1, m^1}) = \alpha^1(t, Y_t^{1, m^1}, m^1) \quad \text{and} \quad \widehat{\alpha}^{m^1, m^2}(t, Y_t^{2, m^2}) = \alpha^2(t, Y_t^{2, m^2}, m^2),$$

a.s., for a.e. t , for $M^1 \times M^2$ -a.e. (m^1, m^2) . Because the drift b is bounded, the random variable Y_t^{i, m^i} has full support for each $t > 0$, and we deduce that

$$\alpha^1(t, x, m^1) = \alpha^2(t, x, m^2),$$

for (Lebesgue) almost every (t, x) , and for $M^1 \times M^2$ -a.e. (m^1, m^2) . This in turn implies $X^1 \equiv X^2$ a.s., which implies $\mu^1 = \mu^2$ a.s. In summary, we have shown that for any two weak MFE μ^1 and μ^2 , we can couple the two in such a way that they are independent and also satisfy $\mu^1 = \mu^2$ a.s. This shows not only that μ^1 and μ^2 must have the same law but also that they must be deterministic, hence strong MFE.

6. Closed-loop versus open-loop. This section compares our notion of weak semi-Markov RMFE (Definition 3.6) with the notion of *weak MFG solution* of [51], Definition 3.1, which itself is a specialization of [18], Definition 3.1, to the case without common noise. The relevance of the latter definition is that it characterizes the limits of n -player approximate equilibria in open-loop regime [51], Theorem 3.4. Our goal is to show that these two definitions are largely equivalent. To state the definition of a weak MFG solution, we first need a bit of notation.

Recall from Section 5.2 the definition of the space \mathcal{V} of relaxed controls. Define $\mathcal{X} := \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d$, and equip this space with the filtration $\mathbb{F}^{\mathcal{X}} = (\mathcal{F}_t^{\mathcal{X}})_{t \in [0, T]}$, where $\mathcal{F}_t^{\mathcal{X}}$ is the σ -field generated by the maps $\mathcal{X} \ni (w, q, x) \mapsto (w_s, q(S), x_s) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d$, where $s \leq t$ and S is a Borel subset of $[0, t] \times A$. As usual, we identify a $\mathcal{P}(A)$ -valued process $\Lambda = (\Lambda_t)_{t \in [0, T]}$ with the random element of \mathcal{V} given by $dt \Lambda_t(da)$. For a measure $\tilde{m} \in \mathcal{P}(\mathcal{X})$, we write $\tilde{m}^x = (\tilde{m}_t^x)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ for the measure flow associated with the third marginal, that is, $\tilde{m}_t^x = \tilde{m} \circ [(w, q, x) \mapsto x_t]^{-1}$.

DEFINITION 6.1. A *weak MFG solution* is a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \tilde{\mu}, \Lambda, X)$, where:

(1) $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a complete filtered probability space. Also, W is an \mathbb{F} -Brownian motion of dimension d , X is an \mathbb{F} -adapted d -dimensional process with $\mathbb{P} \circ X_0^{-1} = \lambda$, and Λ is a $\mathcal{P}(A)$ -valued \mathbb{F} -progressively measurable process. Lastly, $\tilde{\mu}$ is a $\mathcal{P}(\mathcal{X})$ -valued random variable such that $\tilde{\mu}(S)$ is \mathcal{F}_t -measurable whenever $S \in \mathcal{F}_t^{\mathcal{X}}$ and $t \in [0, T]$.

- (2) $\tilde{\mu}, X_0$ and W are independent.
- (3) The state equation holds,

$$dX_t = \int_A b(t, X_t, \tilde{\mu}_t^x, a) \Lambda_t(da) dt + dW_t.$$

(4) The control Λ is *compatible*, in the sense that $\sigma(\Lambda_s : s \leq t)$ is conditionally independent of $\mathcal{F}_T^{X_0, W, \tilde{\mu}}$ given $\mathcal{F}_t^{X_0, W, \tilde{\mu}}$, for each $t \in [0, T]$, where

$$\mathcal{F}_t^{X_0, W, \tilde{\mu}} := \sigma(X_0, W_s, \tilde{\mu}(S) : s \leq t, S \in \mathcal{F}_t^{\mathcal{X}}).$$

(5) The control Λ is *optimal*, in the sense that if $(\Omega', \mathcal{F}', \mathbb{F}', \mathbb{P}', W', \tilde{\mu}', \Lambda', X')$ satisfies (1-4) and $\mathbb{P}' \circ (\tilde{\mu}')^{-1} = \mathbb{P} \circ \tilde{\mu}^{-1}$, then we have

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

- (6) The consistency condition holds: $\tilde{\mu} = \mathbb{P}((W, \Lambda, X) \in \cdot \mid \tilde{\mu})$ a.s.

We may abuse notation somewhat by referring to $\tilde{\mu}$ itself as a weak MFG solution. This is reasonable because we can recover the full joint law of $(\tilde{\mu}, W, \Lambda, X)$ from that of $\tilde{\mu}$ by using the consistency condition (6).

THEOREM 6.2. Suppose $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X, \mu)$ is a weak semi-Markov RMFE. Let $\Lambda_t = \Lambda^*(t, X_t, \mu)$, and set $\tilde{\mu} = \mathbb{P}((W, \Lambda, X) \in \cdot \mid \mu)$. Then $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \tilde{\mu}, \Lambda, X)$ is a weak MFG solution.

PROOF. First, define $\mathbb{F}^{\tilde{\mu}} = (\mathcal{F}_t^{\tilde{\mu}})_{t \in [0, T]}$ as the filtration generated by $\tilde{\mu}$, namely, $\mathcal{F}_t^{\tilde{\mu}} = \sigma(\tilde{\mu}(S) : S \in \mathcal{F}_t^{\mathcal{X}})$. As usual, let $\mathcal{F}_t^{\mu} = \sigma(\mu_s : s \leq t)$. We claim first that $\mathcal{F}_t^{\tilde{\mu}} = \mathcal{F}_t^{\mu}$ for each t .

Recall from Remark 3.8 that $\mathcal{L}(X_t | \mu) = \mathcal{L}(X_t | \mathcal{F}_t^\mu) = \mu_t$ a.s. for each t . It follows immediately that $\mathcal{F}_t^\mu \subset \mathcal{F}_t^{\tilde{\mu}}$, because

$$\mu_t = \mathbb{P}(X_t \in \cdot | \mu) = \tilde{\mu}_t^x \quad \text{a.s.}$$

For the reverse, fix a bounded $\mathcal{F}_t^{\mathcal{X}}$ -measurable function $h : \mathcal{X} \rightarrow \mathbb{R}$. Note that X is necessarily $\mathbb{F}^{X_0, W, \mu}$ -adapted by Lemma A.2, and thus so is Λ , where $\mathbb{F}^{X_0, W, \mu} = (\mathcal{F}_t^{X_0, W, \mu})_{t \in [0, T]}$ is defined by $\mathcal{F}_t^{X_0, W, \mu} = \sigma(X_0, W_s, \mu_s : s \leq t)$. Hence, we may find a bounded $\mathcal{F}_t^{X_0, W, \mu}$ -measurable random variable $\psi(X_0, W, \mu)$ such that $h(W, \Lambda, X) = \psi(X_0, W, \mu)$ a.s. Then,

$$\begin{aligned} \langle \tilde{\mu}, h \rangle &= \mathbb{E}[h(W, \Lambda, X) | \mu] = \mathbb{E}[\psi(X_0, W, \mu) | \mu] \\ &= \langle \lambda \times \mathcal{W}, \psi(\cdot, \cdot, \mu) \rangle, \end{aligned}$$

where \mathcal{W} denotes Wiener measure on \mathcal{C}^d , and the last identity follows from the independence of X_0, W and μ . Because $\psi(X_0, W, \mu)$ is $\mathcal{F}_t^{X_0, W, \mu}$ -measurable, this shows that $\langle \tilde{\mu}, h \rangle$ is \mathcal{F}_t^μ -measurable. Hence, $\mathcal{F}_t^\mu \supset \mathcal{F}_t^{\tilde{\mu}}$.

Properties (1–3) and (6) of Definition 6.1 are straightforward to check now that we have shown $\mathcal{F}_t^{\tilde{\mu}} = \mathcal{F}_t^\mu$ for each t . The compatibility property (4) follows easily from the fact that Λ is $\mathbb{F}^{X_0, W, \mu} = \mathbb{F}^{X_0, W, \tilde{\mu}}$ -adapted.

It remains to check properly the optimality property (5). According to [18], Lemma 3.11, (see also [51], Lemma 4.7), it suffices to check (5) only for alternative controls Λ' which are adapted to the filtration $\mathbb{F}^{X_0, W, \tilde{\mu}}$, because such controls are dense in a joint distributional sense. Precisely, (5) is equivalent to the following:

(5') For each $\mathbb{F}^{X_0, W, \tilde{\mu}}$ -progressively measurable $\mathcal{P}(A)$ -valued process $\Lambda' = (\Lambda'_t)_{t \in [0, T]}$, we have

$$\begin{aligned} (6.1) \quad & \mathbb{E} \left[\int_0^T \int_A f(t, X_t, \tilde{\mu}_t^x, a) \Lambda_t(da) dt + g(X_T, \tilde{\mu}_T^x) \right] \\ & \geq \mathbb{E} \left[\int_0^T \int_A f(t, X'_t, \tilde{\mu}_t^x, a) \Lambda'_t(da) dt + g(X'_T, \tilde{\mu}_T^x) \right], \end{aligned}$$

where X' is the unique strong solution of the SDE

$$dX'_t = \int_A b(t, X'_t, \tilde{\mu}_t^x, a) \Lambda'_t(da) dt + dW_t, \quad X'_0 = X_0.$$

Let (X', Λ') be as in (5'). Recall that $\mu_t = \tilde{\mu}_t^x$ for all $t \in [0, T]$, and so (6.1) is equivalent to

$$\begin{aligned} (6.2) \quad & \mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda^*(t, X_t, \mu)(da) dt + g(X_T, \mu_T) \right] \\ & \geq \mathbb{E} \left[\int_0^T \int_A f(t, X'_t, \mu_t, a) \Lambda'_t(da) dt + g(X'_T, \mu_T) \right]. \end{aligned}$$

We showed also that $\mathcal{F}_t^\mu = \mathcal{F}_t^{\tilde{\mu}}$ for each $t \in [0, T]$, and thus $\mathcal{F}_t^{X_0, W, \tilde{\mu}} = \mathcal{F}_t^{X_0, W, \mu} := \sigma(X_0, W_s, \mu_s : s \leq t)$. Then Λ' is $\mathbb{F}^{X_0, W, \mu}$ -progressively measurable, and we may write $\Lambda'_t = \Lambda'(t, X_0, W, \mu)$.

Because μ is a weak RMFE, we know that Λ^* is optimal when compared to alternative *semi-Markov* controls. To check that it is optimal over $\mathbb{F}^{X_0, W, \tilde{\mu}}$ controls, we proceed by a projection argument reminiscent of those of Section 4. For $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$, let $\mathbb{P}^m = \mathbb{P}(\cdot | \mu = m)$ denote a version of the regular conditional law given μ . The statements in the rest of this paragraph hold for $\mathbb{P} \circ \mu^{-1}$ -almost every $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$. Since X_0, W and

μ are independent, we have $\mathbb{P}^m \circ (X_0, W)^{-1} = \mathbb{P} \circ (X_0, W)^{-1} = \lambda \times \mathcal{W}$, where \mathcal{W} denotes Wiener measure. Moreover, under \mathbb{P}^m , the SDE still holds, which we may write as

$$dX'_t = \int_A b(t, X'_t, m_t, a) \Lambda'(t, X_0, W, m)(da) dt + dW_t.$$

We wish to apply Theorem 2.14 under this measure \mathbb{P}^m . To do so, we first find a Borel measurable function $\widehat{\Lambda} : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$ such that

$$(6.3) \quad \widehat{\Lambda}(t, X'_t, m) = \mathbb{E}^{\mathbb{P}^m} [\Lambda'(t, X_0, W, m) \mid X'_t], \quad \mathbb{P}^m\text{-a.s.}, \forall t \in [0, T],$$

where these expectations are in the sense of mean measure; see Lemma C.2. The point of this definition is that the unique strong solution X^m (on $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$) of the SDE

$$dX^m_t = \int_A b(t, X^m_t, m_t, a) \widehat{\Lambda}(t, X^m_t, m)(da) dt + dW_t, \quad X^m_0 = X_0,$$

satisfies $\mathbb{P} \circ (X^m_t)^{-1} = \mathbb{P}^m \circ (X'_t)^{-1}$ for each $t \in [0, T]$, by Theorem 2.14.

At this point we would like to re-introduce the random measure flow by replacing m by μ and treating $\widehat{\Lambda}(t, X'_t, \mu)$ as a semi-Markov control. For this to work, we must check that $\widehat{\Lambda}$ is not merely Borel measurable but rather semi-Markov. Note that X' is a strong solution, so it is $\mathbb{F}^{X_0, W, \mu}$ -adapted, and we can write $X'_t = X'(t, X_0, W, \mu)$. We may then write (6.3) as

$$\widehat{\Lambda}(t, X'_t, m) = \mathbb{E}^{\mathbb{P}^m} [\Lambda'(t, X_0, W, m) \mid X'(t, X_0, W, m)].$$

Recall that $\mathbb{P}^m \circ (X_0, W)^{-1} = \lambda \times \mathcal{W}$ and that Λ' and X' are progressive, which implies in particular that $\Lambda'(t, X_0, W, m) = \Lambda'(t, X_0, W, \widetilde{m})$ and $X'(t, X_0, W, m) = X'(t, X_0, W, \widetilde{m})$ a.s., whenever $t \in [0, T]$ and $m_s = \widetilde{m}_s$ for $s \in [0, t]$. From these facts we deduce that $\widehat{\Lambda}(t, X'_t, m) = \widehat{\Lambda}(t, X'_t, \widetilde{m})$ a.s., whenever $m_s = \widetilde{m}_s$ for $s \in [0, t]$.

Finally, returning to the unconditional measure \mathbb{P} , define \widehat{X} to be the unique strong solution (on $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$) of the SDE

$$(6.4) \quad d\widehat{X}_t = \int_A b(t, \widehat{X}_t, \mu_t, a) \widehat{\Lambda}(t, \widehat{X}_t, \mu)(da) dt + dW_t, \quad \widehat{X}_0 = X_0,$$

and note that \widehat{X} is adapted to $\mathbb{F}^{X_0, W, \mu}$. Indeed, see Lemma A.2 and A.3 for well-posedness of this SDE, despite the fact that $\widehat{\Lambda}$ may be discontinuous. In addition, as we check carefully in the same two lemmas, the conditional law of \widehat{X} given $\mu = m$ is precisely $\mathbb{P} \circ (X^m)^{-1}$. In particular, we find

$$\mathbb{P}^m \circ \widehat{X}_t^{-1} = \mathbb{P} \circ (X^m_t)^{-1} = \mathbb{P}^m \circ (X'_t)^{-1},$$

for almost every m and for each t . Equivalently, plugging in the random μ , we have $\mathbb{P}^\mu = \mathbb{P}^\mu \circ (X'_t)^{-1}$ a.s. for each t . Using this and the definition of $\widehat{\Lambda}$, we finally use Fubini's theorem and the tower property of conditional expectation to get

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

Recalling the form of the SDE (6.4) for \widehat{X} , we may finally use the defining property (5) of a weak RMFE (Definition 3.6) to conclude that this expectation is dominated by

$$\mathbb{E} \left[\int_0^T \int_A f(t, X_t, \mu_t, a) \Lambda^*(t, X_t, \mu)(da) dt + g(X_T, \mu_T) \right],$$

which proves (6.2). \square

THEOREM 6.3. *Suppose $\tilde{\mu}$ is a weak MFG solution. Then there exists a weak semi-Markov RMFE μ such that $\mu \stackrel{d}{=} \tilde{\mu}^x$.*

PROOF. Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \tilde{\mu}, \Lambda, X)$ be a weak MFG solution. Recalling that $\tilde{\mu}$ is a random measure on $\mathcal{X} = \mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d$, let $\bar{\mu}$ denote the image under under the map $\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d \ni (w, q, x) \mapsto (x, q) \in \mathcal{C}^d \times \mathcal{V}$. It is straightforward to check using the properties of Definition 6.1 and Itô’s formula that $\bar{\mu}$ satisfies the identity (5.4). Hence Lemma 5.3 applies, in particular part (c), and (enlarging the probability space if necessary) we may define Λ^* and X^* as therein. It is immediate from Lemma 5.3 to check that properties (1–4) and (6) of Definition 3.6 are valid. It remains to check the optimality property (5).

First, from part (b) of Lemma 5.3, note that

$$\begin{aligned} (6.5) \quad & \mathbb{E}^{\mathbb{P}} \left[\int_0^T \int_A f(t, X_t, \tilde{\mu}_t^x, a) \Lambda_t(da) dt + g(X_T, \tilde{\mu}_T^x) \right] \\ & = \mathbb{E}^{\mathbb{P}} \left[\int_0^T \int_A f(t, X_t^*, \mu_t, a) \Lambda^*(t, X_t^*, \mu)(da) dt + g(X_T^*, \mu_T) \right]. \end{aligned}$$

Fix any semi-Markov function $\Lambda' : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathcal{P}(A)$, and let X' denote the unique strong solution (see Lemmas A.2 and A.3) of the SDE

$$dX'_t = \int_A b(t, X'_t, \mu_t, a) \Lambda'(t, X'_t, \mu)(da) dt + dW_t, \quad X_0 \sim \lambda.$$

Note that X' is adapted to the complete filtration generated by the process $(X_0, W_t, \mu_t)_{t \in [0, T]}$. Define the $\mathcal{P}(A)$ -valued process $\tilde{\Lambda}_t = \Lambda'(t, X'_t, \mu)$. One checks easily that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \tilde{\mu}, \tilde{\Lambda}, X')$ satisfies properties (1–4) of Definition 6.1. Hence, using property (5) therein along with (6.5), we find

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

This is valid for any choice of Λ' , and we conclude that property (5) of Definition 3.6 holds. \square

Now that we have shown that weak MFE and weak MFG solutions coincide, we can give a quick proof of Theorem 2.13:

6.1. *Proof of Theorem 2.13.* It was shown in [51], Theorem 3.4, that both claims are true if “weak MFE” is replaced by “weak MFG solution” in the statements. We saw in Theorem 6.3 that a weak MFG solution is a weak RMFE and in Proposition 3.7 that a weak RMFE is a weak MFE under Assumptions A and B.

7. Constructing n -player equilibria from mean field equilibria. This section continues the discussion of Section 2.4 on the question of which weak MFE can arise as the limit of n -player (approximate) Nash equilibria. We begin in Section 7.1 by proving Theorem 3.10, which states that every strong RMFE arises as the limit of n -player approximate equilibria.

The rest of the section is devoted to examples: We warm up in Section 7.2 with some observations on the case where the game-theoretic aspect of the problem degenerates in the sense that A is a singleton. In this uncontrolled regime, we are simply left with the study of McKean–Vlasov limits, which already reveals of some of the range of possible behaviors.

However, much richer behavior is possible when the game-theoretic aspect does not trivialize. Section 7.3 discusses such an example, in which there exist weak MFE which are not mixtures of strong MFE.

7.1. *Proof of Theorem 3.10.* Let (m, Λ^*) be a strong RMFE, in the sense of Definition 3.5. Let X^* denote the corresponding state process,

$$(7.1) \quad dX_t^* = \int_A b(t, X_t^*, m_t, a) \Lambda^*(t, X_t^*)(da) dt + dW_t, \quad X_0^* \sim \lambda.$$

Now, for the n -player game, define $\Lambda^{n,i} \in \mathcal{RM}_n$ by setting

$$\Lambda^{n,i}(t, \mathbf{x}) = \Lambda^*(t, x_i) \quad \text{for } \mathbf{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n.$$

Define

$$\epsilon_n := \sup_{\beta \in \mathcal{RM}_n} J_1^n(\beta, \Lambda^{n,2}, \dots, \Lambda^{n,n}) - J_1^n(\Lambda^{n,1}, \dots, \Lambda^{n,n}).$$

Note that $\epsilon_n \geq 0$, and by symmetry it holds for any $i \in \{1, \dots, n\}$ that

$$\epsilon_n = \sup_{\beta \in \mathcal{RM}_n} J_i^n(\Lambda^{n,1}, \dots, \Lambda^{n,i-1}, \beta, \Lambda^{n,i+1}, \dots, \Lambda^{n,n}) - J_i^n(\Lambda^{n,1}, \dots, \Lambda^{n,n}).$$

Hence, $\Lambda^n = (\Lambda^{n,1}, \dots, \Lambda^{n,n})$ is an ϵ_n -Nash equilibrium. Assumption C lets us apply the result of [52], Theorem 2.5(2), (or more specifically Remark 2.7 therein), a strong form of propagation of chaos, to conclude that $\mu^n \rightarrow m$ in law in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$. Moreover, for any $t \in [0, T]$ and any bounded measurable (not necessarily continuous) function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$, we have

$$(7.2) \quad \int_{\mathbb{R}^d} \varphi d\mu_t^n[\Lambda^n] \rightarrow \int_{\mathbb{R}^d} \varphi dm_t,$$

in probability.

It remains to show that $\epsilon_n \rightarrow 0$. Fix arbitrarily a sequence $\beta^n \in \mathcal{RM}_n$ such that

$$(7.3) \quad J_1^n(\beta^n, \Lambda^{n,2}, \dots, \Lambda^{n,n}) \geq \sup_{\beta \in \mathcal{RM}_n} J_1^n(\beta, \Lambda^{n,2}, \dots, \Lambda^{n,n}) - \frac{1}{n}.$$

Abbreviate $X^n = (X^{n,1}, \dots, X^{n,n}) = X[\Lambda^n]$ and $\mu^n = \mu^n[\Lambda^n]$, as well as

$$Y^n = (Y^{n,1}, \dots, Y^{n,n}) = X[(\beta^n, \Lambda^{n,2}, \dots, \Lambda^{n,n})],$$

$$v^n = \mu^n[(\beta^n, \Lambda^{n,2}, \dots, \Lambda^{n,n})].$$

In particular, the state process X^n follows the SDEs

$$dX_t^{n,i} = \int_A b(t, X_t^{n,i}, \mu_t^n, a) \Lambda^*(t, X_t^{n,i})(da) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^{n,k}},$$

whereas Y^n follows the SDEs

$$\begin{aligned} dY_t^{n,1} &= \int_A b(t, Y_t^{n,1}, v_t^n, a) \beta^n(t, Y^n)(da) dt + dW_t^1, \\ dY_t^{n,k} &= \int_A b(t, Y_t^{n,k}, v_t^n, a) \Lambda^*(t, Y_t^{n,k})(da) dt + dW_t^k, \quad i \neq 1, \\ v_t^n &= \frac{1}{n} \sum_{j=1}^n \delta_{Y_t^{n,j}}. \end{aligned}$$

Suppose that X^n is defined on a filtered probability space $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n, \mathbb{P}^n)$, where W^k are of course assumed to be \mathbb{F}^n -Brownian motions. (We will avoid giving a name to whatever probability space Y^n is defined on, which may be different.) Define a probability measure \mathbb{Q}^n on $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n)$ by

$$\begin{aligned} \frac{d\mathbb{Q}^n}{d\mathbb{P}^n} &= \exp\left(\int_0^T \int_A b(t, X_t^{n,1}, \mu_t^n, a) (\beta^n(t, X^n) - \Lambda^*(t, X_t^{n,1}))(da) dW_t^1 \right. \\ &\quad \left. - \frac{1}{2} \int_0^T \left| \int_A b(t, X_t^{n,1}, \mu_t^n, a) (\beta^n(t, X^n) - \Lambda^*(t, X_t^{n,1}))(da) \right|^2 dt \right). \end{aligned}$$

By Girsanov’s theorem and uniqueness of the SDEs, we have $\mathbb{Q}^n \circ (X^n)^{-1} = \mathcal{L}(Y^n)$. Boundedness of b implies that

$$\sup_{n \in \mathbb{N}} \mathbb{E}^{\mathbb{P}^n} \left[\left| \frac{d\mathbb{Q}^n}{d\mathbb{P}^n} \right|^p \right] < \infty,$$

for all $p \geq 1$. Hence, because μ^n converges in probability to m under \mathbb{P}^n (in the sense that $\lim_{n \rightarrow \infty} \mathbb{P}^n(\mu^n \notin U) = 0$ for any open neighborhood U of m in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$), it also converges in probability to m under \mathbb{Q}^n . But $\mathbb{Q}^n \circ (\mu^n)^{-1} = \mathcal{L}(v^n)$, and so $v^n \rightarrow m$ in probability; this uses the fact that m is deterministic.⁹

Now, view $(Y^{n,1}, \beta^n(\cdot, Y^n), W^1)$ as a random element of $\mathcal{C}^d \times \mathcal{V} \times \mathcal{C}^d$, where the space \mathcal{V} of relaxed controls was defined in Section 5.2. Recalling that \mathcal{V} is compact, it is straightforward to check that this sequence is tight. Letting (Y, β, W) denote any subsequential limit point, one readily checks using continuity of b that W is a Brownian motion with respect to the filtration $(\sigma(Y_s, \beta_s, W_s : s \leq t))_{t \in [0, T]}$, that $Y_0 \sim \lambda$, and that the SDE holds,

$$dY_t = \int_A b(t, Y_t, m_t, a) \beta_t(da) dt + dW_t.$$

Use Lemma C.2 to find a measurable function $\Lambda : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$ such that

$$\Lambda(t, Y_t) = \mathbb{E}[\beta_t | Y_t] \quad \text{a.s., a.e. } t,$$

in the sense of mean measures. Apply Theorem 2.14 to find that $Y_t \stackrel{d}{=} Z_t$ for all $t \in [0, T]$, where Z is the unique strong solution of the SDE

$$dZ_t = \int_A b(t, Z_t, m_t, a) \Lambda(t, Z_t)(da) dt + dW_t.$$

Using the assumption that f and g are bounded and continuous, we conclude that, along the same convergent subsequence for which $(Y^{n,1}, \beta^n(\cdot, Y^n), W^1)$ converges to (Y, β, W) , we

⁹For a metric space (E, d) , a point $e_0 \in E$, and a sequence ξ_n of E -valued random variables, perhaps defined on different probability spaces, recall that $\mathcal{L}(\xi_n) \rightarrow \delta_{e_0}$ weakly if and only if $\xi_n \rightarrow e_0$ in probability, which means $\lim_{n \rightarrow \infty} \mathbb{P}(d(\xi_n, e_0) > \epsilon) = 0$ for all $\epsilon > 0$.

have

$$\begin{aligned} & \lim_n \mathbb{E} \left[\int_0^T \int_A f(t, Y_t^{n,1}, v_t^n, a) \beta^n(t, \mathbf{Y}^n)(da) dt + g(Y_T^{n,1}, v_T^n) \right] \\ &= \mathbb{E} \left[\int_0^T \int_A f(t, Y_t, m_t, a) \beta_t(da) dt + g(Y_T, m_T) \right] \\ &= \mathbb{E} \left[\int_0^T \int_A f(t, Y_t, m_t, a) \Lambda(t, Y_t)(da) dt + g(Y_T, m_T) \right] \\ &= \mathbb{E} \left[\int_0^T \int_A f(t, Z_t, m_t, a) \Lambda(t, Z_t)(da) dt + g(Z_T, m_T) \right] \\ &\leq \mathbb{E} \left[\int_0^T \int_A f(t, X_t^*, m_t, a) \Lambda^*(t, X_t^*)(da) dt + g(X_T^*, m_T) \right], \end{aligned}$$

where the last inequality is from the optimality part of the assumption that m is a strong MFE. This inequality holds for any convergent subsequence of the tight sequence $(Y^{n,1}, \beta^n(t, \mathbf{Y}^n), W^1)$, and we conclude that

$$\begin{aligned} & \limsup_{n \rightarrow \infty} J_1^n(\beta^n, \Lambda^{n,2}, \dots, \Lambda^{n,n}) \\ &= \limsup_{n \rightarrow \infty} \mathbb{E} \left[\int_0^T \int_A f(t, Y_t^{n,1}, v_t^n, a) \beta^n(t, \mathbf{Y}^n)(da) dt + g(Y_T^{n,1}, v_T^n) \right] \\ &\leq \mathbb{E} \left[\int_0^T \int_A f(t, X_t^*, m_t, a) \Lambda^*(t, X_t^*)(da) dt + g(X_T^*, m_T) \right]. \end{aligned}$$

On the other hand, notice that the convergence $\mu^n \rightarrow m$ implies

$$\begin{aligned} & \lim_{n \rightarrow \infty} J_1^n(\Lambda^n) \\ &= \lim_{n \rightarrow \infty} \mathbb{E} \left[\int_0^T \int_A f(t, X_t^{n,1}, \mu_t^n, a) \Lambda^*(t, X_t^{n,1})(da) dt + g(X_T^{n,1}, \mu_T^n) \right] \\ &= \lim_{n \rightarrow \infty} \mathbb{E} \int_0^T \int_{\mathbb{R}^d} \int_A f(t, x, \mu_t^n, a) \Lambda^*(t, x)(da) \mu_t^n(dx) dt + \mathbb{E} \int_{\mathbb{R}^d} g(x, \mu_T^n) \mu_T^n(dx) \\ &= \int_0^T \int_{\mathbb{R}^d} \int_A f(t, x, m_t, a) \Lambda^*(t, x)(da) m_t(dx) dt + \int_{\mathbb{R}^d} g(x, m_T) m_T(dx) \\ &= \mathbb{E} \left[\int_0^T \int_A f(t, X_t^*, m_t, a) \Lambda^*(t, X_t^*)(da) dt + g(X_T^*, m_T) \right], \end{aligned}$$

where the second line used symmetry and the third used (7.2) to deal with the fact that Λ^* may be discontinuous. Recalling the previous inequality and (7.3), we conclude that $\epsilon_n \rightarrow 0$.

7.2. Uncontrolled models and ill-posed ODEs. Weak MFE are easy to construct by building degenerate control problems into ill-posed McKean–Vlasov equations or ODEs, as illustrated in this section. Suppose the drift function is the trivial

$$b(t, x, m, a) = B(\bar{m}),$$

for some bounded continuous function $B : \mathbb{R}^d \rightarrow \mathbb{R}^d$, where we again denote by \bar{m} the mean of a measure $m \in \mathcal{P}(\mathbb{R}^d)$, if it exists. The state process (X^1, \dots, X^n) of the n -player game are then un-controlled, and we do not even need to specify objective functions (f, g) or an

action space A . The dimension d is arbitrary. The state processes then evolve according to

$$(7.4) \quad dX_t^i = B(\bar{\mu}_t^n) dt + dW_t^i, \quad \mu_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

with i.i.d. initial states given by λ .

This is the unique n -player equilibrium, and the above SDE system (7.4) is unique in law. But a broad range of $n \rightarrow \infty$ limiting behavior is possible here, and there are potentially multiple (weak) MFE. Averaging (7.4) over $i = 1, \dots, n$, the empirical mean is seen to follow

$$d\bar{\mu}_t^n = B(\bar{\mu}_t^n) dt + \frac{1}{\sqrt{n}} d\bar{W}_t,$$

where $\bar{W} := \frac{1}{\sqrt{n}} \sum_{k=1}^n W^k$ is a Brownian motion. The sequence of real-valued processes $(\bar{\mu}_t^n)_{t \in [0, T]}$ is easily seen to be tight (using, e.g., Aldous' criterion for tightness [48], Lemma 16.12), and it is straightforward to check that every weak limit is supported on the set $S_{\text{ODE}} \subset C([0, T]; \mathbb{R})$ consisting of those functions $x = x(t)$ satisfying the integral equation

$$(7.5) \quad x(t) = \bar{\lambda} + \int_0^t B(x(s)) ds \quad \forall t \in [0, T].$$

It can be checked that a $\mathcal{P}(\mathbb{R}^d)$ -valued process $\mu = (\mu_t)_{t \in [0, T]}$ is a weak MFE if and only if $(\bar{\mu}_t)_{t \in [0, T]}$ belongs almost surely to S_{ODE} and μ_t is precisely

$$\mu_t = \mathcal{N}_d\left(\bar{\lambda} + \int_0^t B(\bar{\mu}_s) ds, tI\right),$$

where $\mathcal{N}_d(m, \Sigma)$ denotes the d -dimensional Gaussian law with mean vector m and covariance matrix Σ . In particular, weak MFE are parametrized by mixtures of solutions of the ODE (7.5).

Of course, in some cases, such as if B is Lipschitz, this ODE has a unique solution. In this case, there is a unique MFE, and the n -player equilibrium converges to it. But without uniqueness for (7.5), anything could happen. The vanishing noise limit $n \rightarrow \infty$ may select one particular solution, or it may fail to converge at all. See [1, 64] for examples of this phenomenon.

7.3. A game-theoretic example. We now turn to a more interesting example, in which the nonuniqueness of the MFE comes from the game-theoretic aspect rather than from ill-posed state process dynamics. In particular, this example admits many weak MFE which are not mixtures of strong MFE. Consider the $d = 1$ -dimensional mean field game described by the coefficients

$$b(t, x, m, a) = a, \quad f \equiv 0, \quad g(x, m) = x\bar{m}, \quad A = [-1, 1], \quad \lambda = \delta_0,$$

where $\bar{m} = \int_{\mathbb{R}} ym(dy)$. This example was analyzed in [51], Section 3.3. It was shown in Proposition 3.6 therein that there are precisely three strong MFE, m^{-1} , m^0 , and m^1 , defined by

$$(7.6) \quad m_t^c = \mathcal{L}(ct + W_t) \quad \text{for } c \in \{-1, 0, 1\}.$$

On the other hand, there are infinitely many weak MFE, many of which are *not mixtures* of these three strong MFE. In [51], Proposition 3.7, one such weak MFE was constructed explicitly, and we elaborate somewhat on this construction below. Note that [51] works with *weak MFG solutions* in the sense of Definition 6.1 instead of our notion of *weak semi-Markov MFE* (Definition 3.6), but we saw in Section 6 that the two are equivalent in a sense.

To construct a family of weak MFE, let $t_0 \in [0, T]$, and let $(\Omega, \mathcal{F}, \mathbb{P})$ be any probability space supporting a Brownian motion W and an independent random variable γ with $\mathbb{P}(\gamma = 1) = \mathbb{P}(\gamma = -1) = 1/2$. Define a $\mathcal{P}(\mathbb{R}^d)$ -valued process $\mu = (\mu_t)_{t \in [0, T]}$ by

$$(7.7) \quad \mu_t = \begin{cases} \mathcal{L}(W_t) & \text{if } t \leq t_0, \\ \mathcal{L}(W_t + \gamma(t - t_0) \mid \gamma) & \text{if } t \in (t_0, T], \end{cases}$$

and note that the mean of μ_t is

$$(7.8) \quad \bar{\mu}_t = \gamma(t - t_0)^+.$$

Suppose $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ is the complete filtration generated by the processes (W, μ) . In particular, $\mathcal{F}_t = \sigma(W_s : s \leq t)$ for $t \leq t_0$, and $\mathcal{F}_t = \sigma(W_s, \gamma : s \leq t)$ for $t \in (t_0, T]$. Define the state process

$$dX_t^* = \gamma 1_{(t_0, T]}(t) dt + dW_t, \quad X_0^* = 0,$$

and define a control $\alpha_{t_0}^* : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ by

$$\alpha_{t_0}^*(t, x) = \text{sgn}(x) 1_{(t_0, T]}(t),$$

where

$$\text{sgn}(x) := \begin{cases} 1 & \text{if } x > 0, \\ -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0. \end{cases}$$

Then $\gamma = \text{sgn}(\bar{\mu}_t)$ for $t > t_0$, and we can rewrite the dynamics of X^* as

$$(7.9) \quad dX_t^* = \alpha_{t_0}^*(t, \bar{\mu}_t) dt + dW_t.$$

We claim that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha_{t_0}^*, X^*, \mu)$ is a weak MFE in the sense of Definition 2.5. To check that the consistency condition $\mu_t = \mathcal{L}(X_t^* \mid \mathcal{F}_t^\mu)$ holds, note first that the σ -field $\mathcal{F}_t^\mu := \sigma(\mu_s : s \leq t)$ is trivial if $t \leq t_0$ and is equal to $\sigma(\gamma)$ if $t \in (t_0, T]$. Hence,

$$\mathcal{L}(X_t^* \mid \mathcal{F}_t^\mu) = \mathcal{L}\left(W_t + \int_0^t \alpha_{t_0}^*(s, \bar{\mu}_s) ds \mid \mathcal{F}_t^\mu\right) = \mathcal{L}(W_t + \gamma(t - t_0)^+ \mid \mathcal{F}_t^\mu) = \mu_t.$$

We must lastly check that the control $\alpha_{t_0}^*$ defined above is optimal. Fix an alternative semi-Markov control $\alpha = \alpha(t, x, m)$, and define the state process

$$dX'_t = \alpha(t, X'_t, \mu) dt + dW_t, \quad X'_0 = 0.$$

The corresponding reward, using the fact that μ and W are independent, is

$$\begin{aligned} J(\alpha) &:= \mathbb{E}[X'_T \bar{\mu}_T] = \mathbb{E}\left[\int_0^T \alpha(t, X'_t, \mu) dt\right] = \mathbb{E}\left[\int_0^T \alpha(t, X'_t, \mu) \mathbb{E}[\bar{\mu}_T \mid \mathcal{F}_t] dt\right] \\ &= \mathbb{E}\left[\int_0^T \alpha(t, X'_t, \mu) \text{sgn}(\bar{\mu}_t) dt\right]. \end{aligned}$$

Indeed, the last step follows from the independence of W and μ , which yields

$$\begin{aligned} \mathbb{E}[\bar{\mu}_T \mid \mathcal{F}_t] &= \begin{cases} \bar{\mu}_T = \text{sgn}(\bar{\mu}_T) = \text{sgn}(\bar{\mu}_t) & \text{if } t \in (t_0, T], \\ \mathbb{E}[\bar{\mu}_T] = 0 & \text{if } t \leq t_0 \end{cases} \\ &= \alpha_{t_0}^*(t, \bar{\mu}_t). \end{aligned}$$

The optimizers of $J(\alpha)$ over α are precisely those α which satisfy

$$\alpha(t, X'_t, \mu) = \alpha_{t_0}^*(t, \bar{\mu}_t) \quad \text{for } t \in (1, T].$$

In particular, the control $\alpha_{t_0}^*$ itself above is optimal, and we conclude that $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha_{t_0}^*, X^*, \mu)$ is a weak semi-Markov MFE.

REMARK 7.1. This example notably illustrates weak MFE which are not mixtures of strong MFE. Indeed, recall from (7.6) that the three strong MFE are m^{-1}, m^0, m^1 . The weak MFE μ constructed above satisfies in particular $\mathbb{P}(|\bar{\mu}_T| = T - t_0) = 1$. Hence, unless $t_0 = 0$ or $t_0 = T$, this weak MFE is not a mixture of strong MFE.

The McKean–Vlasov equation in (7.9) is ill-posed (by design), which renders this example difficult to analyze. Indeed, consider the set $S_{t_0}^*$ of $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ such that there exists a solution of

$$dX_t = \alpha_{t_0}^*(t, \bar{m}_t) dt + dW_t, \quad X_0 = 0, \mathcal{L}(X_t) = m_t, \forall t \in [0, T].$$

Taking expectations, we find

$$(7.10) \quad d\bar{m}_t = \alpha_{t_0}^*(t, \bar{m}_t) dt, \quad \bar{m}_0 = 0.$$

This is an ill-posed ODE, and its solutions (on the time interval $[0, T]$) are precisely the functions $\{H_s^\pm : s \in [t_0, T]\}$, where

$$(7.11) \quad H_s^\pm(t) = \pm(t - s)^+,$$

noting that $H_T^\pm \equiv 0$. Note then that $S_{t_0}^*$ consists of precisely the measure flows of the form $(\mathcal{L}(W_t + H_s^\pm(t)))_{t \in [0, T]}$, for $s \in [t_0, T]$.

On the other hand, suppose we construct the natural n -particle system

$$dX_t^i = \alpha_{t_0}^*(t, \bar{\mu}_t^n) dt + dW_t^i, \quad X_0^i = 0, \mu_0^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}.$$

Averaging over $i = 1, \dots, n$, we find that the empirical mean satisfies

$$(7.12) \quad d\bar{\mu}_t^n = \alpha_{t_0}^*(t, \bar{\mu}_t^n) dt + \frac{1}{\sqrt{n}} d\bar{W}_t, \quad \bar{\mu}_t^n = 0,$$

where $\bar{W}_t = \frac{1}{\sqrt{n}} \sum_{k=1}^n W_t^k$ is a Brownian motion. One would expect that as $n \rightarrow \infty$ the limit points of $(\bar{\mu}_t^n)_{t \in [0, T]}$ are supported on solutions of the ODE (7.10). But, in fact, this is a well understood example of the “regularization by noise” phenomenon, and a particular mixture is picked out in the limit $n \rightarrow \infty$. Indeed, the law of $(\bar{\mu}_t^n)_{t \in [0, T]}$ converges to the mixture $\frac{1}{2} \delta_{H_0^+} + \frac{1}{2} \delta_{H_0^-}$; this was proven in [64] in the case $t_0 = 0$, and the extension to general t_0 is straightforward. In addition, one can deduce from this that the full measure flow μ^n , not just its mean, converges in law in $C([0, T]; \mathcal{P}(\mathbb{R}^d))$ to μ defined in (7.7).

In light of this discussion, and after studying the proof of Theorem 2.12, it is natural to guess that

$$(7.13) \quad \alpha_{t_0}^{n,i}(t, \mathbf{x}) := \alpha_{t_0}^*\left(t, \frac{1}{n} \sum_{k=1}^n x_k\right)$$

defines an approximate (Markovian) Nash equilibrium for the n -player game, for any $t_0 \in [0, T]$. For $t_0 = T$ this is true and follows from Theorem 2.12, because the MFE $\mu_t = \mathcal{L}(W_t)$ is strong in this case. For general $t_0 \in [0, T)$ it is not as clear, and we have resolved only the $t_0 = 0$ case:

PROPOSITION 7.2. Let $\alpha^n = (\alpha_0^{n,1}, \dots, \alpha_0^{n,n})$, where $\alpha_0^{n,i}$ are defined as in (7.13) with $t_0 = 0$. Then there exists $\epsilon_n \geq 0$ with $\epsilon_n \rightarrow 0$ such that α^n is a Markovian ϵ_n -Nash equilibrium for each n . Moreover, the law of the C^d -valued random variable $(\bar{\mu}_t^n[\alpha^n])_{t \in [0, T]}$ converges to $\frac{1}{2} \delta_{H_0^+} + \frac{1}{2} \delta_{H_0^-}$.

REMARK 7.3. On the other hand, suppose instead that we take γ to be 1, -1 , or 0 with $\mathbb{P}(\gamma = 1) = \mathbb{P}(\gamma = -1) = p < 1/2$ so that $\mathbb{E}\gamma = 0$ and $\mathbb{P}(\gamma = 0) > 0$. Carrying out the exact same construction as above, we arrive at another weak MFE in which (μ, X^*) once again obeys the dynamics

$$dX_t^* = \alpha_{t_0}^*(t, \bar{\mu}_t) dt + dW_t, \quad X_0^* = 0, \mu_t = \mathcal{L}(X_t | \mathcal{F}_t^\mu), \text{ a.s., } \forall t \in [0, T],$$

and again with μ satisfying both (7.8) and (7.7). The point is that in this case the law of $(\bar{\mu}_t)_{t \in [0, T]}$ is given by the mixture $p\delta_{H_0^+} + p\delta_{H_0^-} + (1 - 2p)\delta_0$. This is not the mixture picked out in the limit from the n -particle system (7.12), in which we saw that the law of $\bar{\mu}^n$ converges to $\frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$. In this case, it is not clear if this particular weak MFE can arise as the limit of n -player approximate equilibria, but the naive construction certainly fails.

PROOF OF PROPOSITION 7.2. Recall that our weak MFE μ satisfies $\bar{\mu}_t = \gamma t$, where $\mathbb{P}(\gamma = 1) = \mathbb{P}(\gamma = -1) = 1/2$. The final claim of the Proposition, that the law of $\bar{\mu}^n[\alpha^n]$ converges to $\frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$, was shown in [64].

Define

$$\epsilon_n := \sup_{\beta \in \mathcal{AM}_n} J_1^n(\beta, \alpha_0^{n,2}, \dots, \alpha_0^{n,n}) - J_1^n(\alpha^n).$$

Note that $\epsilon_n \geq 0$, and by symmetry it holds for any $k \in \{1, \dots, n\}$ that

$$\epsilon_n = \sup_{\beta \in \mathcal{AM}_n} J_k^n(\alpha_0^{n,1}, \dots, \alpha_0^{n,k-1}, \beta, \alpha_0^{n,k+1}, \dots, \alpha_0^{n,n}) - J_k^n(\alpha^n).$$

Hence, $\alpha^n = (\alpha_0^{n,1}, \dots, \alpha_0^{n,n})$ is an ϵ_n -Nash equilibrium. It remains to show that $\epsilon_n \rightarrow 0$. A direct calculation, using symmetry and the fact that $|\bar{\mu}_t^n[\alpha^n]| \rightarrow t$ in law, shows that

$$J_1^n(\alpha_0^{n,1}, \dots, \alpha_0^{n,n}) = \mathbb{E}[X_T^1[\alpha^n] \bar{\mu}_T^n[\alpha^n]] = \mathbb{E}[|\bar{\mu}_T^n[\alpha^n]|^2] \rightarrow T^2$$

as $n \rightarrow \infty$. Hence, to show that $\epsilon_n \rightarrow 0$, it suffices to show that

$$(7.14) \quad \limsup_{n \rightarrow \infty} \sup_{\beta \in \mathcal{AM}_n} J_1^n(\beta, \alpha_0^{n,2}, \dots, \alpha_0^{n,n}) \leq T^2.$$

To this end, for each n find $\beta^n \in \mathcal{AM}_n$ such that

$$(7.15) \quad \sup_{\beta \in \mathcal{AM}_n} J_1^n(\beta, \alpha_0^{n,2}, \dots, \alpha_0^{n,n}) \leq J_1^n(\beta^n, \alpha_0^{n,2}, \dots, \alpha_0^{n,n}) + \frac{1}{n}.$$

Abbreviate $X^n = X^1[(\beta^n, \alpha_0^{n,2}, \dots, \alpha_0^{n,n})]$, $Y^n = \bar{\mu}^n[(\beta^n, \alpha_0^{n,2}, \dots, \alpha_0^{n,n})]$. Abuse notation by writing $\beta_t^n = \beta^n(t, X_t[(\beta^n, \alpha_0^{n,2}, \dots, \alpha_0^{n,n})])$. Then

$$dX_t^n = \beta_t^n dt + dW_t^1, \\ dY_t^n = \left(\frac{1}{n} \beta_t^n + \frac{n-1}{n} \text{sgn}(Y_t^n) \right) dt + \frac{1}{n} \sum_{k=1}^n dW_t^k.$$

Finally, we view β^n as a random variable with values in $L_1^2 := L^2([0, T]; [-1, 1])$. Equip L_1^2 with the subspace topology inherited from the weak topology of the Hilbert space $L^2([0, T]; \mathbb{R})$, and note that L_1^2 is then compact and metrizable.

LEMMA 7.4. The sequence (X^n, Y^n, W^1, β^n) of $\mathcal{C} \times \mathcal{C} \times \mathcal{C} \times L_1^2$ -valued random variables is tight, and every weak limit (X, Y, W, β) satisfies:

- (i) $\mathcal{L}(Y) = \frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$, where H_0^\pm are defined in (7.11).
- (ii) The following equations hold, for $t \in [0, T]$:

$$X_t = \int_0^t \beta_s ds + W_t, \quad Y_t = \int_0^t \operatorname{sgn}(Y_s) ds.$$

- (iii) W is a Brownian motion with respect to the filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ defined by $\mathcal{F}_t = \sigma(X_s, Y_s, W_s, \beta_s : s \leq t)$.
- (iv) Y and W are independent.

PROOF. Tightness follows from standard arguments. Let (X, Y, W, β) denote any limit point. Clearly (iii) holds. We first check that (i) holds by showing that the law of Y^n converges weakly to $\frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$. Suppose that $(X^n, Y^n, W^1, \dots, W^n, \beta^n)$ are defined on the filtered probability space $(\Omega^n, \mathcal{F}^n, \mathbb{P}^n)$. Define the Brownian motion $\overline{W}_t = \frac{1}{\sqrt{n}} \sum_{i=1}^n W_t^i$. On this space, let Z^n denote the unique strong solution of the SDE

$$dZ_t^n = \operatorname{sgn}(Z_t^n) dt + \frac{1}{\sqrt{n}} d\overline{W}_t.$$

We know from [64] that $\mathbb{P}^n \circ (Z^n)^{-1} \rightarrow \frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$. Define an equivalent probability measure \mathbb{Q}^n by setting

$$\frac{d\mathbb{Q}^n}{d\mathbb{P}^n} = \exp\left(\frac{1}{\sqrt{n}} \int_0^T (\beta_t^n - \operatorname{sgn}(Z_t^n)) d\overline{W}_t - \frac{1}{2n} \int_0^T (\beta_t - \operatorname{sgn}(Z_t^n))^2 dt\right).$$

By Girsanov’s theorem and uniqueness in law of the SDEs, we have $\mathbb{Q}^n \circ (Z^n)^{-1} = \mathbb{P}^n \circ (Y^n)^{-1}$. This yields the following bound on relative entropy:

$$\mathbb{E}^{\mathbb{P}^n} \left[\frac{d\mathbb{Q}^n}{d\mathbb{P}^n} \log \frac{d\mathbb{Q}^n}{d\mathbb{P}^n} \right] = -\mathbb{E}^{\mathbb{Q}^n} \left[\log \frac{d\mathbb{P}^n}{d\mathbb{Q}^n} \right] = \frac{1}{2n} \mathbb{E}^{\mathbb{Q}^n} \int_0^T (\beta_t - \operatorname{sgn}(Z_t^n))^2 dt \leq \frac{2T}{n}.$$

By Pinsker’s inequality, the total variation norm of $\mathbb{Q}^n - \mathbb{P}^n$ converges to zero. Because $\mathbb{P}^n \circ (Z^n)^{-1} \rightarrow \frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$, we conclude that also $\mathbb{Q}^n \circ (Z^n)^{-1} \rightarrow \frac{1}{2}\delta_{H_0^+} + \frac{1}{2}\delta_{H_0^-}$. Recalling that $\mathbb{Q}^n \circ (Z^n)^{-1} = \mathbb{P}^n \circ (Y^n)^{-1}$, this completes the proof of (i).

With (i) now established, we prove (ii). It is clear that $X_t = \int_0^t \beta_s ds + W_t$ holds, because $X_t^n = \int_0^t \beta_s^n ds + W_t^n$ for each n and because $L_1^2 \ni q \mapsto \int_0^t q_s ds \in \mathbb{R}$ is (weakly) continuous for each t . Finally, note that (i) implies that $Y_t = \int_0^t \operatorname{sgn}(Y_s) ds$ for all t .

To check property (iii), note that the law of W is clearly equal to Wiener measure, so we must only show that $W_t - W_s$ is independent of \mathcal{F}_s for each $t > s \geq 0$. This argument is straightforward and thus omitted.

We finally show that (iv) follows from the other claims. Because W is \mathbb{F} -Brownian, it is also \mathbb{F}_+ -Brownian, where $\mathbb{F}_+ = (\mathcal{F}_{t+})_{t \in [0, T]}$ denotes the right-continuous augmentation, defined by $\mathcal{F}_t = \bigcap_{\epsilon > 0} \mathcal{F}_{t+\epsilon}$. In particular, W is independent of \mathcal{F}_{0+} . Now, from (i) we may write $Y_t = t \operatorname{sgn}(Y_T)$ a.s., from which we conclude that the entire process Y is a.s.-measurable with respect to \mathcal{F}_{0+} . Hence, Y and W are independent. \square

With this Lemma in hand, we now complete the proof of Proposition 7.2. Working with a subsequence of (X^n, Y^n, W^1, β^n) and its limit (X, Y, W, β) , we have

$$\begin{aligned} \lim_n J_1^n(\beta^n, \alpha_0^{n,2}, \dots, \alpha_0^{n,n}) &= \lim_n \mathbb{E}[X_T^n Y_T^n] \\ &= \mathbb{E}[X_T Y_T] = \mathbb{E}\left[Y_T \int_0^T \beta_t dt + Y_T W_T\right] = \mathbb{E}\left[Y_T \int_0^T \beta_t dt\right] \\ &\leq T \mathbb{E}|Y_T| = T^2, \end{aligned}$$

with the limit taken along the appropriate subsequence. Note that the second equality is valid in light of the simple estimate $\sup_{n \in \mathbb{N}} \mathbb{E}[|X_T^n Y_T^n|^p] < \infty$ for any $p > 1$, which provides the uniform integrability needed to pass to the limit. Finally, because this holds for each convergent subsequence, we conclude finally from

$$\limsup_{n \rightarrow \infty} J_1^n(\beta^n, \alpha_0^{n,2}, \dots, \alpha_0^{n,n}) \leq T^2.$$

Recalling (7.14) and (7.15), this completes the proof. \square

APPENDIX A: SDES WITH RANDOM COEFFICIENTS

This section develops some intuitively clear but somewhat delicate technical points regarding SDEs with random coefficients. It will be useful to write $\mathbb{F}^E = (\mathcal{F}_t^E)_{t \in [0, T]}$ for the canonical filtration on the path space $C([0, T]; E)$, defined for any Polish space E .

For the rest of the section, fix a complete separable metric space E (which in applications in this paper will be $E = \mathcal{P}(\mathbb{R}^d)$). As in Definition 2.4, let us say that a function $B : [0, T] \times \mathbb{R}^d \times C([0, T]; E) \rightarrow \mathbb{R}^d$ is semi-Markov if it is Borel measurable and satisfies $F(t, x, e) = F(t, x, e')$ whenever $(t, x) \in [0, T] \times \mathbb{R}^d$ and $e, e' \in C([0, T]; E)$ satisfy $e_s = e'_s$ for all $s \leq t$. Fix throughout the section one such semi-Markov function B , which we assume is bounded. Equip $C([0, T]; E)$ with the supremum distance. We fix also a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supporting a d -dimensional \mathbb{F} -Brownian motion W as well as an \mathcal{F}_0 -measurable \mathbb{R}^d -valued random variable ξ with law λ .

The goal of this section is to justify the following points:

(1) *Deterministic well-posedness*: For a deterministic $e \in C([0, T]; E)$, there is a unique strong solution of the SDE

$$(A.1) \quad dX_t^e = B(t, X_t^e, e) dt + dW_t, \quad X_0^e = \xi.$$

Let $P^e \in \mathcal{P}(\mathcal{C}^d)$ denote its law. By “strong solution” here we mean X^e is adapted to the complete filtration generated by the process $(\xi, W_t)_{t \in [0, T]}$.

(2) *Stochastic well-posedness*: If η is a $C([0, T]; E)$ -valued random variable with law M , independent of (ξ, W) , then there is a unique strong solution of the SDE

$$(A.2) \quad dX_t = B(t, X_t, \eta) dt + dW_t, \quad X_0 = \xi.$$

By “strong solution” we mean X is adapted to the complete filtration generated by the process $(\xi, W_t, \eta_t)_{t \in [0, T]}$.

(3) *Consistency*: The map $C([0, T]; E) \ni e \mapsto P^e \in \mathcal{P}(\mathcal{C}^d)$ is universally measurable and, in the notation of part (2), provides a version of the conditional law of X given η . That is, for each bounded measurable function φ on $C([0, T]; E) \times \mathcal{C}^d$, we have

$$\mathbb{E}[\varphi(\eta, X)] = \int_{C([0, T]; E)} M(de) \int_{\mathcal{C}^d} P^e(dx) \varphi(e, x).$$

(4) *Stability*: Given a uniformly bounded sequence of semi-Markov functions $B_n : [0, T] \times \mathbb{R}^d \times C([0, T]; E) \rightarrow \mathbb{R}^d$ satisfying $B_n(t, x, e) \rightarrow B(t, x, e)$ for M -a.e. e and Lebesgue-a.e. (t, x) , we have

$$\lim_{n \rightarrow \infty} \mathbb{E}[\varphi(\eta, X^n)] = \mathbb{E}[\varphi(\eta, X)]$$

for each bounded measurable function $\varphi : C([0, T]; E) \times \mathcal{C}^d \rightarrow \mathbb{R}$, where X^n is the unique strong solution of

$$(A.3) \quad dX_t^n = B_n(t, X_t^n, \eta) dt + dW_t, \quad X_0^n = \xi.$$

(5) *Equivalence to forward equations:* Suppose a continuous $\mathcal{P}(\mathbb{R}^d)$ -valued process $\mu = (\mu_t)_{t \in [0, T]}$ is a weak solution of the randomized Fokker–Planck equation associated to (A.2). Precisely, suppose μ is adapted to the filtration generated by η , and it holds almost surely that, for all $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$,

$$\langle \mu_t, \varphi \rangle = \langle \lambda, \varphi \rangle + \int_0^t \left\langle \mu_s, B(s, \cdot, \eta) \cdot \nabla \varphi(\cdot) + \frac{1}{2} \Delta \varphi(\cdot) \right\rangle ds.$$

Then $\mu_t = \mathcal{L}(X_t \mid (\eta_s)_{s \leq t})$ a.s., for each t , where X is as in (A.2).

These results are applied in the text in the particular case $E = \mathcal{P}(\mathbb{R}^d)$, and with $\mu = \eta$ in step (5), but we find it clearer and perhaps useful on its own to work in this more general setting.

A.1. Deterministic well-posedness. Part (1) of the program follows from the result of Veretennikov [66] (see also [49], Theorem 2.1). That is, for each $e \in C([0, T]; E)$ there exists a unique strong solution X^e of the SDE (A.1). Let $P^e = \mathbb{P} \circ (X^e)^{-1}$. In particular, pathwise uniqueness holds for this SDE, in the following sense: Suppose our probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supports two continuous \mathbb{F} -adapted processes X^1, X^2 which both satisfy

$$dX_t^i = B(t, X_t^i, e) dt + dW_t, \quad X_0^i = \xi, \quad i = 1, 2,$$

and also as usual the process W is an \mathbb{F} -Brownian motion independent of ξ . Then $X^1 = X^2$ a.s., and the law of X^1 is precisely P^e .

We would like to be able to construct a version of $(t, \omega, e) \mapsto X_t^e(\omega)$ which is jointly measurable and which depends in an adapted fashion on e , but it is not clear how to do this. Uniqueness of the strong solution X^e easily yields $\mathbb{P}(X_s^e = X_s^{\tilde{e}}, \forall s \leq t) = 1$ whenever $t \in [0, T]$ and $e, \tilde{e} \in C([0, T]; E)$ satisfy $e_s = \tilde{e}_s$ for all $s \leq t$. But the null set depends on (t, e, \tilde{e}) , and we thus face a continuum of null sets. There is no continuity in e to exploit, as we have made no continuity assumptions on B , and this is the main technical impediment to our program (1–5). Instead, we work with the law P^e instead of the process X^e itself.

In the following, for each $t \in [0, T]$ we let $\overline{\mathcal{F}}_t^E$ denote the universal completion of \mathcal{F}_t^E . That is, if $\mathcal{N}_t^M := \{S \subset C([0, T]; E) : \exists S' \in \mathcal{F}_t^E, S \subset S', M(S') = 0\}$ denotes the set of M -null sets of \mathcal{F}_t^E , then

$$\overline{\mathcal{F}}_t^E := \bigcap_{M \in \mathcal{P}(C([0, T]; E))} \sigma(\mathcal{F}_t^E \cup \mathcal{N}_t^M).$$

We will make implicit use of the well known fact that the following σ -fields coincide, for any Polish space E' (see [5], Corollary 7.29.1):

- The Borel σ -field on $\mathcal{P}(E')$ induced by the topology of weak convergence.
- The σ -field generated by the maps $\mathcal{P}(E') \ni m \mapsto \langle m, \varphi \rangle$, where φ is a bounded continuous function on E' .
- The σ -field generated by the maps $\mathcal{P}(E') \ni m \mapsto m(S)$, where S is a Borel set of E' .

LEMMA A.1. *The map $C([0, T]; E) \ni e \mapsto P^e \in \mathcal{P}(C^d)$ is universally measurable. Moreover, this map is adapted in the sense that, for every $t \in [0, T]$ and every $S \in \mathcal{F}_t^{\mathbb{R}^d}$, the map $e \mapsto P^e(S)$ is $\overline{\mathcal{F}}_t^E$ -measurable.*

PROOF. We first prove the claims under the additional assumption that $B(t, x, e)$ is continuous in e for each (t, x) . We claim that then $e \mapsto P^e$ is continuous. To see this, suppose $e^n \rightarrow e$ in $C([0, T]; E)$. It then holds for bounded continuous function $\varphi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with compact support that

$$\lim_{n \rightarrow \infty} \int_0^T \int_{\mathbb{R}^d} B(t, x, e^n) \cdot \varphi(t, x) dx dt = \int_0^T \int_{\mathbb{R}^d} B(t, x, e) \cdot \varphi(t, x) dx dt.$$

It follows from [62], Theorem 11.3.3, that $P^{e^n} \rightarrow P^e$. Moreover, for $t \in [0, T]$, uniqueness of the SDE ensures that if $e_s = \tilde{e}_s$ for $s \leq t$ then $X_s^e = X_s^{\tilde{e}}$ for all $s \leq t$, a.s. Hence, if $S \in \mathcal{F}_t^{\mathbb{R}^d}$, then $P^e(S) = \mathbb{P}(X^e \in S) = \mathbb{P}(X^{\tilde{e}} \in S) = P^{\tilde{e}}(S)$, and so $e \mapsto P^e(S)$ is \mathcal{F}_t^E -measurable.

We now address general B by an approximation argument. Fix a probability measure $M \in \mathcal{P}(C([0, T]; E))$. Define the finite measure Q on $[0, T] \times \mathbb{R}^d \times C([0, T]; E)$ by setting, for Borel sets S ,

$$Q(S) = \int_{C([0, T]; E)} \int_{\mathbb{R}^d} \int_0^T 1_S(t, x, e) \exp(-|x|^2) dt dx M(de).$$

We may then find a sequence of continuous semi-Markov functions B_n which converges Q -almost everywhere to B . Define P_n^e as the law of the corresponding SDE solution, that is, $P_n^e = \mathbb{P} \circ (X^{n,e})^{-1}$ where $X^{n,e}$ is given by

$$dX_t^{n,e} = B_n(t, X_t^{n,e}, e) dt + dW_t, \quad X_0^{n,e} = \xi.$$

As argued in the previous paragraph, $e \mapsto P_n^e$ is continuous for each n . Moreover, it holds for M -almost every $e \in C([0, T]; E)$ that

$$\lim_{n \rightarrow \infty} \int_0^T \int_{\mathbb{R}^d} B_n(t, x, e) \cdot \varphi(t, x) dx dt = \int_0^T \int_{\mathbb{R}^d} B(t, x, e) \cdot \varphi(t, x) dx dt$$

for each bounded continuous function $\varphi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with compact support. It follows again from [62], Theorem 11.3.3, that $P_n^e \rightarrow P^e$ for M -a.e. $e \in C([0, T]; E)$. Hence, the map $e \mapsto P^e$ agrees M -a.e. with a Borel measurable function, so it is measurable with respect to the M -completion of the Borel σ -field of $C([0, T]; E)$. This holds for every choice of $M \in \mathcal{P}(C([0, T]; E))$, which proves the first claim. The second claim is proven similarly: For $t \in [0, T]$, we know from above that $e \mapsto P_n^e(S)$ is \mathcal{F}_t^e -measurable for $S \in \mathcal{F}_t^{\mathbb{R}^d}$. Hence, for any bounded continuous $\mathcal{F}_t^{\mathbb{R}^d}$ -measurable function $\varphi : \mathcal{C}^d \rightarrow \mathbb{R}$, we have $\lim_n \langle P_n^e, \varphi \rangle = \langle P^e, \varphi \rangle$ for M -a.e. e , and we deduce that $e \mapsto \langle P^e, \varphi \rangle$ is $\overline{\mathcal{F}}_t^E$ -measurable. \square

A.2. Stochastic well-posedness. We now turn to steps (2) and (3) of the program outlined at the beginning of the section, by proving weak existence and pathwise uniqueness for the SDE (A.2) and then identifying the law of the unique solution as $\mathcal{L}(\eta, X) = M(de)P^e(dx)$. As the SDE (A.2) has random coefficients, the original form of the Yamada–Watanabe theorem does not apply, and we instead use the generalization due to Jacod–Mémmin [47] to conclude, as usual, that weak existence and pathwise uniqueness are together equivalent to uniqueness in law and existence of a strong solution. The first lemma checks that the SDE (A.2) is pathwise unique and identifies its law (rather, it satisfies *very good pathwise uniqueness* in the language of [47], Definition 2.24).

LEMMA A.2. *Let $M \in \mathcal{P}(C([0, T]; E))$. Suppose our filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supports an \mathbb{F} -adapted continuous E -valued process η with law M , independent of (ξ, W) , as well as two d -dimensional \mathbb{F} -adapted processes (X^1, X^2) satisfying*

$$dX_t^i = B(t, X_t^i, \eta) dt + dW_t, \quad X_0^i = \xi, i = 1, 2.$$

Define $\mathcal{F}_t^{\xi, \eta, W} := \sigma(\xi, \eta_s, W_s : s \leq t)$, and assume that $(X_s^1, X_s^2)_{s \in [0, t]}$ is conditionally independent of $\mathcal{F}_t^{\xi, \eta, W}$ given $\mathcal{F}_t^{\xi, \eta, W}$, for each $t \in [0, T]$. Then $X^1 = X^2$ a.s., and $\mathbb{P} \circ (\eta, X^i)^{-1} = M(de)P^e(dx)$ for each $i = 1, 2$.

PROOF. By assumption, $\mathbb{P} \circ \eta^{-1} = M$. Let us show that $e \mapsto P^e$ is a version of the conditional law $\mathbb{P}(X^i \in \cdot \mid \eta = e)$. Define the regular conditional law $C([0, T]; E) \ni e \mapsto$

$\mathbb{Q}_e = \mathbb{P}(\cdot \mid \eta = e) \in \mathcal{P}(\Omega)$. Because W , ξ and η are independent, we have $\mathbb{Q}_e \circ (\xi, W)^{-1} = \mathbb{P} \circ (\xi, W)^{-1}$ for M -a.e. e . Moreover, the SDE

$$X_t^i = \xi + \int_0^t B(s, X_s^i, e) ds + W_t \quad \forall t \in [0, T],$$

holds almost surely under \mathbb{Q}_e , for M -a.e. e .

We would like to conclude from pathwise uniqueness (see the first paragraph of Section A.1) that $\mathbb{Q}_e(X^1 = X^2) = 1$ and $\mathbb{Q}_e \circ (X^i)^{-1} = P^e$ for M -a.e. e . To do so we need only to show that W is an $\widetilde{\mathbb{F}}$ -Brownian motion under \mathbb{Q}_e , for M -a.e. e , where $\widetilde{\mathbb{F}} = (\widetilde{\mathcal{F}}_t)_{t \in [0, T]}$ denotes the filtration (on Ω) generated by X^1 , X^2 and W , that is, $\widetilde{\mathcal{F}}_t = \sigma(X_s, W_s : s \leq t)$. This amounts to proving that for each $t \in [0, T]$, each $\sigma(X_s^1, X_s^2 : s \leq t)$ -measurable random variable $\varphi_t(X)$, each $\sigma(W_s : s \leq t)$ -measurable random variable $h_t(W)$, and each $\sigma(W_s - W_t : s \in [t, T])$ -measurable random variable $h_{t+}(W)$, we have

$$\mathbb{E}^{\mathbb{Q}_e}[\varphi_t(X)h_t(W)h_{t+}(W)] = \mathbb{E}^{\mathbb{Q}_e}[\varphi_t(X)h_t(W)]\mathbb{E}^{\mathbb{Q}_e}[h_{t+}(W)].$$

To prove this, notice that if $\varphi : C([0, T]; E) \rightarrow \mathbb{R}$ is any bounded measurable function, then (taking expectations under \mathbb{P})

$$\begin{aligned} \mathbb{E}[\varphi(\eta)\varphi_t(X)h_t(W)h_{t+}(W)] &= \mathbb{E}[\mathbb{E}[\varphi_t(X) \mid \mathcal{F}_T^{\xi, \eta, W}] \varphi(\eta)h_t(W)h_{t+}(W)] \\ &= \mathbb{E}[\mathbb{E}[\varphi_t(X) \mid \mathcal{F}_t^{\xi, \eta, W}] \varphi(\eta)h_t(W)h_{t+}(W)] \\ &= \mathbb{E}[\mathbb{E}[\varphi_t(X) \mid \mathcal{F}_t^{\xi, \eta, W}] \varphi(\eta)h_t(W)]\mathbb{E}[h_{t+}(W)] \\ &= \mathbb{E}[\varphi(\eta)\varphi_t(X)h_t(W)]\mathbb{E}[h_{t+}(W)]. \end{aligned}$$

Indeed, the second and final lines follow from the assumed conditional independence of $(X_s^1, X_s^2)_{s \leq t}$ and $\mathcal{F}_T^{\xi, \eta, W}$ given $\mathcal{F}_t^{\xi, \eta, W}$, while the second to last identity follows from the fact that $(\xi, \eta, (W_s)_{s \leq t})$ and $(W_s - W_t)_{s \geq t}$ are independent, which is an easy consequence of the independence of ξ , η and W . We conclude that

$$\mathbb{E}[\varphi_t(X)h_t(W)h_{t+}(W) \mid \eta] = \mathbb{E}[\varphi_t(X)h_t(W) \mid \eta]\mathbb{E}[h_{t+}(W)] \quad \text{a.s.,}$$

which completes the proof. \square

Now that we have checked pathwise uniqueness, we turn to the problem of existence. The following lemma shows that we can construct (η, X) with law $M(de)P^e(dx)$ so that the SDE (A.2) does indeed hold, as well as the conditional independence property of Lemma A.2. This will be enough to deduce strong existence, using a form of the Yamada–Watanabe theorem [47], Theorem 2.25.

LEMMA A.3. *Let $M \in \mathcal{P}(C([0, T]; E))$. Suppose our filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supports an \mathbb{F} -adapted continuous E -valued process η with law M , independent of (ξ, W) . Then there exists a continuous \mathbb{F} -adapted process X solving*

$$dX_t = B(t, X_t, \eta) dt + dW_t, \quad X_0 = \xi,$$

such that $\mathbb{P} \circ (\eta, X)^{-1} = M(de)P^e(dx)$. In particular, X is adapted to the complete filtration generated by the process $(\xi, \eta_t, W_t)_{t \in [0, T]}$.

PROOF. Following the strategy described above, we begin by building a weak solution. We work on the canonical space $\overline{\Omega} = C([0, T]; E) \times \mathcal{C}^d$. Let (η, X) denote the canonical (coordinate) processes, and let $\overline{\mathbb{F}} = (\overline{\mathcal{F}}_t)_{t \in [0, T]}$ denote the filtration they generate, which can

be written as $\overline{\mathcal{F}}_t = \mathcal{F}_t^E \otimes \mathcal{F}_t^{\mathbb{R}^d}$. Define $\overline{\mathbb{P}}(de, dx) = M(de)P^e(dx)$. For each $t \in [0, T]$, define $W_t : \overline{\Omega} \rightarrow \mathbb{R}$ by

$$W_t(e, x) = x_t - x_0 - \int_0^t B(s, x_s, e) ds,$$

and define $\xi := X_0$. The process $W = (W_t)_{t \in [0, T]}$ is $\overline{\mathbb{P}}$ -progressively measurable with respect to the canonical filtration. Note that $W(e, \cdot)$ is a Brownian motion on $(\mathcal{C}^d, \mathbb{F}^{\mathbb{R}^d}, P^e)$, for each $e \in C([0, T]; E)$, by definition of P^e . It follows easily that W is an $\overline{\mathbb{P}}$ -Brownian motion under $\overline{\mathbb{P}}$. Moreover, ξ, η and W are independent. By construction, the SDE holds,

$$dX_t = B(t, X_t, \eta) dt + dW_t, \quad X_0 = \xi.$$

We will show that $(X_s)_{s \leq t}$ is conditionally independent of $\mathcal{F}_T^{\xi, \eta, W}$ given $\mathcal{F}_t^{\xi, \eta, W}$ for each t , where $\mathcal{F}_t^{\xi, \eta, W} := \sigma(\xi, \eta_s, W_s : s \leq t)$. To prove this, fix $t \in [0, T]$ as well as random variables $h_t(W), h_{t+}(W), \varphi_0(\xi), \varphi_t(X), \psi_t(\eta)$, and $\psi_T(\eta)$, measurable with respect to $(W_s)_{s \leq t}, (W_s - W_t)_{s \in [t, T]}, \xi, (X_s)_{s \leq t}, (\eta_s)_{s \leq t}$, and η , respectively. Then, by definition of $\overline{\mathbb{P}}$,

$$\begin{aligned} & \mathbb{E}^{\overline{\mathbb{P}}}[h_t(W)h_{t+}(W)\varphi_0(\xi)\varphi_t(X)\psi_t(\eta)\psi_T(\eta)] \\ (A.4) \quad &= \int_{C([0, T]; E)} M(de)\psi_t(e)\psi_T(e) \int_{\mathcal{C}^d} P^e(dx)\varphi_0(x_0)\varphi_t(x)h_t(W(e, x))h_{t+}(W(e, x)) \\ &= (\mathcal{W}, h_{t+}) \int_{C([0, T]; E)} M(de)\psi_t(e)\psi_T(e) \int_{\mathcal{C}^d} P^e(dx)\varphi_0(x_0)\varphi_t(x)h_t(W(e, x)), \end{aligned}$$

where \mathcal{W} denotes Wiener measure on \mathcal{C}^d , and where the last line used the fact that $W(e, \cdot)$ is a Brownian motion on $(\mathcal{C}^d, \mathbb{F}^{\mathbb{R}^d}, P^e)$, mentioned above. Now, the independence of ξ, η and W easily implies

$$\mathbb{E}^{\overline{\mathbb{P}}}[\psi_T(\eta) \mid \mathcal{F}_t^{\xi, \eta, W}] = \mathbb{E}^{\overline{\mathbb{P}}}[\psi_T(\eta) \mid \mathcal{F}_t^\eta] =: \tilde{\psi}_t(\eta).$$

Moreover, because the function $\mathcal{C}^d \ni x \mapsto \varphi_0(x_0)\varphi_t(x)h_t(W(e, x))$ is $\mathcal{F}_t^{\mathbb{R}^d}$ -measurable for each fixed e , the adaptedness of $e \mapsto P^e$ proven in Lemma A.1 implies that the function

$$C([0, T]; E) \ni e \mapsto \int_{\mathcal{C}^d} P^e(dx)\varphi_0(x_0)\varphi_t(x)h_t(W(e, x))$$

is $\overline{\mathcal{F}}_t^E$ -measurable, and in particular it agrees M -a.e. with an \mathcal{F}_t^E -measurable function. Hence, we may condition in the right-hand side of (A.4) to get

$$(\mathcal{W}, h_{t+}) \int_{C([0, T]; E)} M(de)\psi_t(e)\tilde{\psi}_t(e) \int_{\mathcal{C}^d} P^e(dx)\varphi_0(x_0)\varphi_t(x)h_t(W(e, x)),$$

and we deduce from (A.4) that

$$\begin{aligned} & \mathbb{E}^{\overline{\mathbb{P}}}[h_t(W)h_{t+}(W)\varphi_0(\xi)\varphi_t(X)\psi_t(\eta)\psi_T(\eta)] \\ &= \mathbb{E}^{\overline{\mathbb{P}}}[h_{t+}(W)]\mathbb{E}^{\overline{\mathbb{P}}}[h_t(W)\varphi_0(\xi)\varphi_t(X)\psi_t(\eta)\tilde{\psi}_t(\eta)]. \end{aligned}$$

Finally, recall that

$$\tilde{\psi}_t(\eta)\mathbb{E}^{\overline{\mathbb{P}}}[h_{t+}(W)] = \mathbb{E}^{\overline{\mathbb{P}}}[\psi_T(\eta) \mid \mathcal{F}_t^{\xi, \eta, W}]\mathbb{E}^{\overline{\mathbb{P}}}[h_{t+}(W) \mid \mathcal{F}_t^{\xi, \eta, W}],$$

which yields

$$\begin{aligned} & \mathbb{E}^{\overline{\mathbb{P}}}[h_t(W)h_{t+}(W)\varphi_0(\xi)\varphi_t(X)\psi_t(\eta)\psi_T(\eta)] \\ &= \mathbb{E}^{\overline{\mathbb{P}}}[h_t(W)\varphi_0(\xi)\varphi_t(X)\psi_t(\eta)\mathbb{E}^{\overline{\mathbb{P}}}[\psi_T(\eta) \mid \mathcal{F}_t^{\xi, \eta, W}]\mathbb{E}^{\overline{\mathbb{P}}}[h_{t+}(W) \mid \mathcal{F}_t^{\xi, \eta, W}]] \\ &= \mathbb{E}^{\overline{\mathbb{P}}}[h_t(W)\psi_t(\eta)\varphi_0(\xi)\mathbb{E}^{\overline{\mathbb{P}}}[\varphi_t(X) \mid \mathcal{F}_t^{\xi, \eta, W}]\mathbb{E}^{\overline{\mathbb{P}}}[\psi_T(\eta) \mid \mathcal{F}_t^{\xi, \eta, W}]\mathbb{E}^{\overline{\mathbb{P}}}[h_{t+}(W) \mid \mathcal{F}_t^{\xi, \eta, W}]]. \end{aligned}$$

This proves the desired conditional independence (and in fact a bit more).

Finally, we complete the proof in the manner announced before the statement of the lemma. It is well known (see, e.g., [7], Theorem 3) that the following are equivalent:

- (1) Every $\mathbb{F}^{\xi, \eta, W}$ -martingale is an $\mathbb{F}^{\xi, \eta, W, X}$ -martingale, where $\mathbb{F}^{\xi, \eta, W, X} = (\mathcal{F}_t^{\xi, \eta, W, X})_{t \in [0, T]}$ is defined by $\mathcal{F}_t^{\xi, \eta, W, X} = \sigma(\xi, \eta_s, W_s, X_s : s \leq t)$.
- (2) $(X_s)_{s \leq t}$ is conditionally independent of $\mathcal{F}_t^{\xi, \eta, W}$ given $\mathcal{F}_t^{\xi, \eta, W}$ for each $t \in [0, T]$.

This shows that our conditional independence property is in fact equivalent to the notion of *very good solution measure* in [47], Definition 1.7. Thus, by [47], Theorem 2.25, we conclude that the solution measure is in fact strong, which means in our context that X must be adapted with respect to the $\overline{\mathbb{P}}$ -completion of $\mathbb{F}^{\xi, \eta, W}$. \square

A.3. Stability. We turn next to part (4) of the outline from the beginning of the section. Suppose $B_n : [0, T] \times \mathbb{R}^d \times C([0, T]; E) \rightarrow \mathbb{R}^d$ is semi-Markov, for each n . Assume that B_n are uniformly bounded and that $B_n(t, x, e) \rightarrow B(t, x, e)$, for M -a.e. $e \in C([0, T]; E)$ and Lebesgue-a.e. $(t, x) \in [0, T] \times \mathbb{R}^d$. Thanks to the work of the previous section, we may define X^n as the unique strong solution of the SDE (A.3) corresponding to coefficient B_n . That is,

$$dX_t^n = B_n(t, X_t^n, \eta) dt + dW_t, \quad X_0^n = \xi.$$

LEMMA A.4. *Let $M \in \mathcal{P}(C([0, T]; E))$. Suppose our filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supports an \mathbb{F} -adapted continuous E -valued process η with law M , independent of (ξ, W) . Then, for every bounded measurable function $h : C([0, T]; E) \times \mathcal{C}^d \rightarrow \mathbb{R}$, we have*

$$(A.5) \quad \lim_{n \rightarrow \infty} \mathbb{E}[\varphi(\eta, X^n)] = \mathbb{E}[\varphi(\eta, X)].$$

In particular, (η, X^n) converges in law to (η, X) .

PROOF. From Lemmas A.2 and A.3, we know that $\mathbb{P} \circ (\eta, X^n)^{-1} = M(de)P_n^e(dx)$, where we define $P_n^e := \mathbb{P} \circ (X^n, e)^{-1}$ as the law of the unique strong solution of the SDE

$$dX_t^{n,e} = B_n(t, X_t^{n,e}, e) dt + dW_t, \quad X_0^{n,e} = \xi.$$

By assumption, for any bounded continuous function $\varphi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ with compact support, we have

$$\lim_{n \rightarrow \infty} \int_0^T \int_{\mathbb{R}^d} B_n(t, x, e) \cdot \varphi(t, x) dx dt = \int_0^T \int_{\mathbb{R}^d} B(t, x, e) \cdot \varphi(t, x) dx dt,$$

for M -a.e. e . It follows from [62], Theorem 11.3.3, that $P_n^e \rightarrow P^e$ for M -a.e. e . It follows immediately that $M(de)P_n^e(dx) \rightarrow M(de)P^e(dx)$ weakly. To prove that convergence holds for bounded measurable test functions, we need a bit more. Let $\mathcal{W}_\lambda \in \mathcal{P}(\mathcal{C}^d)$ denote the law of a d Brownian motion started from initial law λ . Then

$$\frac{dP_n^e}{d\mathcal{W}_\lambda}(w) = \exp\left(\int_0^T B_n(t, w_t, e) dw_t - \frac{1}{2} \int_0^T |B_n(t, w_t, e)|^2 dt\right).$$

Because B_n is uniformly bounded, it is straightforward to show that

$$\sup_{e \in C([0, T]; E)} \sup_{n \in \mathbb{N}} \int_{\mathcal{C}^d} \left| \frac{dP_n^e}{d\mathcal{W}_\lambda} \right|^2 d\mathcal{W}_\lambda < \infty.$$

This implies that the family $\{dP_n^e/d\mathcal{W}_\lambda : n \in \mathbb{N}, e \in C([0, T]; E)\}$ is precompact in $L^2(\mathcal{W}_\lambda)$ with the weak topology, and this is enough to let us upgrade the convergence. Indeed, we conclude that

$$\lim_{n \rightarrow \infty} \int_{\mathcal{C}^d} h dP_n^e = \int_{\mathcal{C}^d} h dP^e \quad \forall e \in C([0, T]; E),$$

not only for bounded continuous functions $h : \mathcal{C}^d \rightarrow \mathbb{R}$ but also for bounded measurable functions. Finally, if $h : C([0, T]; E) \times \mathcal{C}^d \rightarrow \mathbb{R}$ is bounded and measurable, we conclude from dominated convergence that

$$\lim_{n \rightarrow \infty} \int_{C([0, T]; E)} \int_{\mathcal{C}^d} h(e, x) P_n^e(dx) M(de) = \int_{C([0, T]; E)} \int_{\mathcal{C}^d} h(e, x) P^e(dx) M(de).$$

This is equivalent to the claimed (A.5). \square

A.4. Forward equations. Let E and B be as in the previous section. Let λ denote the law of the initial state ξ . Consider the problem of finding $(m_t)_{t \in [0, T]} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ such that

$$(A.6) \quad \langle m_t, \varphi \rangle = \langle \lambda, \varphi \rangle + \int_0^t \left\langle m_s, B(s, \cdot, e) \cdot \nabla \varphi(\cdot) + \frac{1}{2} \Delta \varphi(\cdot) \right\rangle ds,$$

for all $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$. This is nothing but the Fokker–Planck equation associated with the SDE (A.1). One solution is provided by the marginal flow $(P_t^e = \mathcal{L}(X_t^e))_{t \in [0, T]}$, and the following gives uniqueness.

LEMMA A.5. *Fix $e \in C([0, T]; E)$, and suppose $m \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ satisfies (A.6) for every $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$. Then $m_t = P_t^e$ for all $t \in [0, T]$.*

PROOF. It is well known that the solution of a Fokker–Planck equation, in very general settings, can be represented as the marginal laws of a solution of the corresponding martingale problem. See, for example, [31], Theorem 2.6, or [65], Theorem 2.5. In our context, the martingale problem has a unique solution given by P^e , and the claim follows. \square

LEMMA A.6. *Suppose our filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supports an \mathbb{F} -adapted continuous E -valued process η , independent of (ξ, W) , as well as a continuous $\mathcal{P}(\mathbb{R}^d)$ -valued process μ which is adapted to the filtration generated by η . Suppose it holds almost surely that, for all $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$,*

$$(A.7) \quad \langle \mu_t, \varphi \rangle = \langle \lambda, \varphi \rangle + \int_0^t \left\langle \mu_s, B(s, \cdot, \eta) \cdot \nabla \varphi(\cdot) + \frac{1}{2} \Delta \varphi(\cdot) \right\rangle ds.$$

Then $\mu_t = P_t^\eta$ for all $t \in [0, T]$ a.s. Moreover, we may find a continuous process X , adapted to the complete filtration generated by the process $(\xi, W_t, \eta_t)_{t \in [0, T]}$, such that

$$(A.8) \quad dX_t = B(t, X_t, \eta) dt + dW_t, \quad X_0 = \xi,$$

and also $\mathcal{L}(X_t | \eta) = \mathcal{L}(X_t | (\eta_s)_{s \leq t}) = \mu_t$ a.s. for each t .

PROOF. Let $M \in \mathcal{P}(C([0, T]; E))$ denote the law of η . As we assumed μ is adapted to the filtration of η , we may write $\mu = \widehat{\mu}(\eta)$ a.s., where $\widehat{\mu} : C([0, T]; E) \rightarrow C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is an adapted map in the sense that $\widehat{\mu}^{-1}(S) \in \mathcal{F}_t^{\mathcal{P}(\mathbb{R}^d)}$ for each $S \in \mathcal{F}_t^E$ and each $t \in [0, T]$. Then because of (A.7), for M -a.e. $e \in C([0, T]; E)$ and every $\varphi \in C_c^\infty(\mathbb{R}^d)$ and $t \in [0, T]$ it holds that

$$\langle \widehat{\mu}_t(e), \varphi \rangle = \langle \lambda, \varphi \rangle + \int_0^t \left\langle \widehat{\mu}_s(e), B(s, \cdot, e) \cdot \nabla \varphi(\cdot) + \frac{1}{2} \Delta \varphi(\cdot) \right\rangle ds.$$

From Lemma A.5 we conclude that $\widehat{\mu}_t(e) = P_t^e$ for each $t \in [0, T]$. As this holds for almost every e , we deduce the first claim: $\mu_t = P_t^\eta$ for all t , a.s.

Now, using Lemmas A.2 and A.3, we may safely define X to be the unique strong solution of the SDE (A.8), and we know that $\mathbb{P} \circ (\eta, X)^{-1} = M(de)P^e(dx)$. This last identity is

equivalent to $\mathcal{L}(X | \eta) = P^\eta$ a.s. Marginalizing at time t and using the conclusion of the previous paragraph, we find $\mathcal{L}(X_t | \eta) = P_t^\eta = \mu_t$ for all t , a.s.

Lastly, to deduce that $\mathcal{L}(X_t | \eta) = \mathcal{L}(X_t | (\eta_s)_{s \leq t})$, note that for any bounded measurable $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ we have $\mathbb{E}[\varphi(X_t) | \eta] = \langle \mu_t, \varphi \rangle$. As μ_t is $(\eta_s)_{s \leq t}$ -measurable, we may condition on $(\eta_s)_{s \leq t}$ to get $\mathbb{E}[\varphi(X_t) | (\eta_s)_{s \leq t}] = \langle \mu_t, \varphi \rangle$. \square

We finally note how Lemma A.6 specializes in the most important situation for this paper, where $E = \mathcal{P}(\mathbb{R}^d)$ and the processes η and μ are identical. Assume now that $B : [0, T] \times \mathbb{R}^d \times C([0, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow \mathbb{R}^d$ is a given bounded semi-Markov function (in the sense of Definition 2.4).

COROLLARY A.7. *Suppose our filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ supports an \mathbb{F} -adapted continuous $\mathcal{P}(\mathbb{R}^d)$ -valued process μ , independent of (ξ, W) . Suppose it holds almost surely that, for all $t \in [0, T]$ and $\varphi \in C_c^\infty(\mathbb{R}^d)$,*

$$(A.9) \quad \langle \mu_t, \varphi \rangle = \langle \lambda, \varphi \rangle + \int_0^t \left\langle \mu_s, B(s, \cdot, \mu) \cdot \nabla \varphi(\cdot) + \frac{1}{2} \Delta \varphi(\cdot) \right\rangle ds.$$

Then we may find a continuous process X , adapted to the complete filtration generated by the process $(\xi, W_t, \mu_t)_{t \in [0, T]}$, such that

$$dX_t = B(t, X_t, \mu) dt + dW_t, \quad X_0 = \xi,$$

and also $\mathcal{L}(X_t | \mu) = \mathcal{L}(X_t | \mathcal{F}_t^\mu) = \mu_t$ a.s. for each t , where $\mathcal{F}_t^\mu = \sigma(\mu_s : s \leq t)$.

APPENDIX B: JOINT MEASURABILITY OF REGULAR CONDITIONAL LAWS

This section provides the details of a technical point used in various places in the paper, notably in the proof of Lemma 5.3. Therein, we wanted to define a regular conditional law in a way that is jointly measurable with respect to the underlying probability law. The first lemma in this direction is likely known, but we include a proof. Recall that for a Polish space E we always equip $\mathcal{P}(E)$ with the topology of weak convergence and the corresponding Borel σ -field.

LEMMA B.1. *Let E and E' be Polish spaces, and let $\pi : E \rightarrow E'$ be continuous. Then there exists a measurable map $\Gamma : \mathcal{P}(E) \times E' \rightarrow \mathcal{P}(E)$ such that*

$$\int_E F(x) h(\pi(x)) m(dx) = \int_E \left(\int_E F d\Gamma(m, \pi(x)) \right) h(\pi(x)) m(dx),$$

for all bounded measurable $F : E \rightarrow \mathbb{R}$ and $h : E' \rightarrow \mathbb{R}$.

PROOF. To write this in a more probabilistic notation, let $X : E \rightarrow E$ denote the identity map. What we must find is a version of the regular conditional law $m(X \in \cdot | \pi(X) = x')$ which is jointly measurable as a function of $(x', m) \in E' \times \mathcal{P}(E)$. We borrow a construction of [56], Lemma 3.1. Because E' is Polish, we may find a refining sequence of finite Borel partitions (A_1^n, \dots, A_n^n) of E' such that $\bigcup_n \sigma(A_1^n, \dots, A_n^n)$ generates the Borel σ -field. For each n , define $\Gamma_n : \mathcal{P}(E) \times E' \rightarrow \mathcal{P}(E)$ by

$$\Gamma_n(m, x')(\cdot) = \sum_{k=1}^n \frac{m(\cdot \cap \pi^{-1}(A_k^n))}{m(\pi^{-1}(A_k^n))} 1_{A_k^n}(x'),$$

where we adopt the convention $0/0 := 0$. As E is Polish, we may find a countable sequence (φ_k) of bounded continuous functions such that $\mathcal{P}(E) \ni m \mapsto (\langle m, \varphi_k \rangle)_{k \in \mathbb{N}} \in \mathbb{R}^\mathbb{N}$

is a homeomorphism to its image. Because the σ -algebras $\sigma(A_1^n, \dots, A_n^n)$ increase in n by design, the supermartingale convergence theorem ensures that for each $m \in \mathcal{P}(E)$ the $\lim_n \int \varphi_k d\Gamma_n(m, x')$ exists for $m \circ \pi^{-1}$ -almost every x' and is a version of the conditional expectation $\mathbb{E}^m[\varphi_k(X) \mid \pi(X) = x']$. Now, fixing $x_0 \in E$ arbitrarily, we may set

$$\Gamma(m, x') := \begin{cases} \lim_n \Gamma_n(m, x') & \text{if the limit exists,} \\ \delta_{x_0} & \text{otherwise,} \end{cases}$$

where the limit is in the sense of weak convergence. Then, with the help of the sequence (φ_k) from above, we deduce that for each $m \in \mathcal{P}(E)$ the map $x' \mapsto \Gamma(m, x')$ is a version of the regular conditional law $m(X \in \cdot \mid \pi(X) = x')$. As Γ_n is jointly measurable for each n , so too is Γ . \square

We now turn to the real purpose of this section. In the following, let Ω be a Polish space. Let E be a complete and separable metric space, and let $X = (X_t)_{t \in [0, T]}$ be a measurable process. By *measurable* here we mean that the function $X : [0, T] \times \Omega \rightarrow E$ is jointly Borel-measurable. We show next how to construct a version of $P(\cdot \mid X_t = x)$ which is jointly measurable in t, x , and the underlying probability measure P . Let $\mathbb{E}^P[\cdot]$ denote expectation with respect to a probability measure $P \in \mathcal{P}(\Omega)$.

LEMMA B.2. *There exists a jointly measurable function $\Gamma : [0, T] \times E \times \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega)$ such that, for every bounded measurable function $F : [0, T] \times \Omega \rightarrow \mathbb{R}$ and each $P \in \mathcal{P}(\Omega)$, we have*

$$\mathbb{E}^P[F(t, \cdot) \mid X_t] = \int_{\Omega} F(t, \cdot) d\Gamma(t, X_t, P), \quad P\text{-a.s., a.e. } t \in [0, T].$$

PROOF. Consider the measurable space $\overline{\Omega} = [0, T] \times \Omega$ and $\overline{E} = [0, T] \times E$, and define $\pi : \overline{\Omega} \rightarrow \overline{E}$ by $\pi(t, \omega) = (t, X_t(\omega))$. Apply Lemma B.1 to find a measurable function $\overline{\Gamma} : \mathcal{P}(\overline{\Omega}) \times \overline{E} \rightarrow \mathcal{P}(\overline{\Omega})$ such that for each $\overline{P} \in \mathcal{P}(\overline{\Omega})$ it holds that $\overline{\Gamma}(\overline{P}, \cdot)$ is a version of the conditional law $\overline{P}(\cdot \mid \pi)$. Let U denote the uniform probability measure on $[0, T]$. Then, for $P \in \mathcal{P}(\Omega)$ and bounded measurable functions $F : \overline{\Omega} \rightarrow \mathbb{R}$ and $h : \overline{E} \rightarrow \mathbb{R}$, we have

$$\begin{aligned} & \frac{1}{T} \mathbb{E}^P \int_0^T F(t, \cdot) h(t, X_t) dt \\ &= \int_{\Omega} \int_0^T F(t, \omega) h(\pi(t, \omega)) U(dt) P(d\omega) \\ &= \int_{[0, T] \times \Omega} h(\pi(t, \omega)) \left(\int_{\overline{\Omega}} F d\overline{\Gamma}(U \times P, \pi(t, \omega)) \right) (U \times P)(dt, d\omega) \\ &= \int_{[0, T] \times E} h(t, x) \left(\int_{\overline{\Omega}} F d\overline{\Gamma}(U \times P, (t, x)) \right) (U \times P) \circ \pi^{-1}(dt, dx) \\ &= \int_{[0, T] \times E} h(t, x) \left(\int_{\overline{\Omega}} F d\Gamma(P, t, x) \right) (U \times P) \circ \pi^{-1}(dt, dx), \end{aligned}$$

where we define $\tilde{\Gamma} : \mathcal{P}(\Omega) \times \overline{E} \rightarrow \mathcal{P}(\overline{\Omega})$ by setting $\tilde{\Gamma}(P, t, x) := \overline{\Gamma}(U \times P, (t, x))$. Note next that

$$\int_{[0, T] \times E} g(t, x) (P \times U) \circ \pi^{-1}(dt, dx) = \frac{1}{T} \mathbb{E}^P \int_0^T g(t, X_t) dt,$$

for any bounded measurable $g : \overline{E} \rightarrow \mathbb{R}$. Hence, the above becomes

$$\frac{1}{T} \mathbb{E}^P \int_0^T F(t, \cdot) h(t, X_t) dt = \frac{1}{T} \mathbb{E}^P \int_0^T h(t, X_t) \left(\int_{\overline{\Omega}} F d\tilde{\Gamma}(P, t, X_t) \right) dt.$$

For a measure $\bar{P} \in \mathcal{P}(\bar{\Omega}) = \mathcal{P}([0, T] \times \Omega)$, let \bar{P}^T and \bar{P}^Ω denote the $[0, T]$ and Ω marginals, respectively. By choosing F depending only on t , we find that $\tilde{\Gamma}(P, t, X_t)^T = \delta_t$ a.s. for a.e. $t \in [0, T]$. Finally, define $\Gamma : [0, T] \times E \times \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega)$ by marginalizing, for example, setting

$$\Gamma(t, x, P) := \tilde{\Gamma}(P, t, x)^\Omega.$$

Then

$$\tilde{\Gamma}(P, t, X_t) = \Gamma(t, X_t, P) \times \delta_t \quad \text{a.s., a.e. } t \in [0, T],$$

and we find

$$\frac{1}{T} \mathbb{E}^P \int_0^T F(t, \cdot) h(t, X_t) dt = \frac{1}{T} \mathbb{E}^P \int_0^T h(t, X_t) \left(\int_\Omega F(t, \cdot) d\Gamma(t, X_t, P) \right) dt.$$

This is enough to complete the proof (see [8], Lemma 5.2). \square

APPENDIX C: CONDITIONAL MEANS OF RANDOM MEASURES

This section gives some details regarding one additional technical point, relevant in various applications of the Markovian projection Theorem 2.14 in settings involving relaxed controls, which is to construct a measurable version of the conditional mean of a random measure. This is formalized in the following lemma, stated in a setting abstract enough to allow for the various applications we have in mind. This is implicit in many references on relaxed controls, but we include the short proofs in the spirit of full transparency, as in the previous sections.

LEMMA C.1. *Suppose $\Gamma : E \mapsto \mathcal{P}(\Omega)$ is a measurable map. Suppose also that $K : \Omega \rightarrow \mathcal{P}(A)$ is measurable. Then there exists a measurable function $\Lambda : E \rightarrow \mathcal{P}(A)$ such that, for every bounded measurable function $\varphi : E \times A \rightarrow \mathbb{R}$ and every $x \in E$, we have*

$$\int_A \varphi(x, a) \Lambda(x)(da) = \int_\Omega \left(\int_A \varphi(x, a) dK(\omega)(da) \right) \Gamma(x)(d\omega).$$

PROOF. It is useful here to recall the well-known equivalence stated just before Lemma A.1. Define $\Lambda(x)(S)$ for Borel sets $S \subset A$ and $x \in E$ by setting

$$\Lambda(x)(S) = \int_\Omega K(\omega)(S) \Gamma(x)(d\omega).$$

For each $x \in E$, it is clear that $\Lambda(x)(\cdot)$ is a probability measure on A . On the other hand, for each Borel set $S \subset A$, the map $\omega \mapsto K(\omega)(S)$ is Borel measurable, and thus so is $x \mapsto \Lambda(x)(S)$. We conclude that Λ defines a measurable map from E to $\mathcal{P}(A)$, and the claimed identity holds whenever $\varphi(x, a) = 1_S(a)$ for a Borel set $S \subset A$. It is straightforward to extend this to any φ of the form $\varphi(x, a) = \psi(a)$, for $\psi : A \rightarrow \mathbb{R}$ bounded and measurable. Because the identity holds pointwise, for each $x \in E$, we can then extend to general $\varphi = \varphi(x, a)$. \square

One of the main purposes of the abstract considerations of Sections B and C is the following:

LEMMA C.2. *Suppose we are given, on some Polish probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a measurable $\mathcal{P}(A)$ -valued process $\beta = (\beta_t)_{t \in [0, T]}$ as well as a measurable E -valued process $X = (X_t)_{t \in [0, T]}$. Then there exists a jointly measurable function $\hat{\beta} : [0, T] \times E \rightarrow \mathcal{P}(A)$ such that, for each bounded measurable function $\varphi : [0, T] \times E \times A \rightarrow \mathbb{R}$, we have*

$$\int_A \varphi(t, X_t, a) \hat{\beta}(t, X_t)(da) = \mathbb{E} \left[\int_A \varphi(t, X_t, a) \beta_t(da) \mid X_t \right] \quad \text{a.s., a.e. } t \in [0, T].$$

We may write this with the suggestive mean measure notation,

$$\widehat{\beta}(t, X_t) = \mathbb{E}[\beta_t \mid X_t].$$

PROOF. By Lemma B.2, a version of the regular conditional law $\Gamma(t, x) = \mathbb{P}(\cdot \mid X_t = x) \in \mathcal{P}(\Omega)$ can be constructed which is jointly measurable in (t, x) . Then simply apply Lemma C.1. \square

Lastly, in the proof of Lemma 5.3, we need the following:

LEMMA C.3. Suppose we are given, on some Polish probability space $(\Omega, \mathcal{F}, \mathbb{P})$, a jointly measurable function $\beta : [0, T] \times \mathbb{R}^d \times \Omega \rightarrow \mathcal{P}(A)$, as well as a continuous E -valued process $\eta = (\eta_t)_{t \in [0, T]}$. Let $\mathcal{F}_t^\eta = \sigma(\eta_s : s \leq t)$. Then there exists a semi-Markov function $\Lambda : [0, T] \times \mathbb{R}^d \times C([0, T]; E) \rightarrow \mathcal{P}(A)$ such that

$$\Lambda(t, X_t, \eta) = \mathbb{E}[\beta(t, X_t, \cdot) \mid \mathcal{F}_t^\eta] \quad \text{a.s., for each } t \in [0, T].$$

That is, for each bounded measurable function $\varphi : [0, T] \times \mathbb{R}^d \times C([0, T]; E) \times A \rightarrow \mathbb{R}$, we have

$$\begin{aligned} & \int_A \varphi(t, X_t, \eta, \cdot) d\Lambda(t, X_t, \eta) \\ &= \mathbb{E} \left[\int_A \varphi(t, X_t, \eta, \cdot) d\beta(t, X_t, \cdot) \mid \mathcal{F}_t^\eta \right] \quad \text{a.s., for each } t \in [0, T]. \end{aligned}$$

PROOF. This will follow immediately from Lemma C.1 once we can construct a version of the conditional law $\mathbb{P}(\cdot \mid \mathcal{F}_t^\eta)$ which is progressively measurable in (t, η) . In fact, noting that $\mathcal{F}_t^\eta = \sigma(\eta_{\cdot \wedge t})$, this follows from Lemma B.2 applied with $X = (X_t)_{t \in [0, T]}$ therein given by the stopped process $(\eta_{\cdot \wedge t})_{t \in [0, T]}$. \square

REFERENCES

- [1] BAFICO, R. and BALDI, P. (1981/82). Small random perturbations of Peano phenomena. *Stochastics* **6** 279–292. MR0665404 <https://doi.org/10.1080/17442508208833208>
- [2] BAYRAKTAR, E. and COHEN, A. (2018). Analysis of a finite state many player game using its master equation. *SIAM J. Control Optim.* **56** 3538–3568. MR3860894 <https://doi.org/10.1137/17M113887X>
- [3] BEN AROUS, G. and ZEITOUNI, O. (1999). Increasing propagation of chaos for mean field models. *Ann. Inst. Henri Poincaré Probab. Stat.* **35** 85–102. MR1669916 [https://doi.org/10.1016/S0246-0203\(99\)80006-5](https://doi.org/10.1016/S0246-0203(99)80006-5)
- [4] BERGIN, J. and BERNHARDT, D. (1992). Anonymous sequential games with aggregate uncertainty. *J. Math. Econom.* **21** 543–562. MR1195676 [https://doi.org/10.1016/0304-4068\(92\)90026-4](https://doi.org/10.1016/0304-4068(92)90026-4)
- [5] BERTSEKAS, D. and SHREVE, S. (1996). *Stochastic Optimal Control: The Discrete Time Case*. Athena Scientific.
- [6] BORKAR, V. S. and GHOSH, M. K. (1992). Stochastic differential games: Occupation measure based approach. *J. Optim. Theory Appl.* **73** 359–385. MR1163266 <https://doi.org/10.1007/BF00940187>
- [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/BF00537538>
- [8] BRUNICK, G. and SHREVE, S. (2013). Mimicking an Itô process by a solution of a stochastic differential equation. *Ann. Appl. Probab.* **23** 1584–1628. MR3098443 <https://doi.org/10.1214/12-aap881>
- [9] CARDALIAGUET, P., CIRANT, M. and PORRETTA, A. (2018). Remarks on Nash equilibria in mean field game models with a major player. arXiv preprint [arXiv:1811.02811](https://arxiv.org/abs/1811.02811).
- [10] CARDALIAGUET, P. (2017). The convergence problem in mean field games with local coupling. *Appl. Math. Optim.* **76** 177–215. MR3679342 <https://doi.org/10.1007/s00245-017-9434-0>
- [11] CARDALIAGUET, P., DELARUE, F., LASRY, J.-M. and LIONS, P.-L. (2019). *The Master Equation and the Convergence Problem in Mean Field Games*. *Annals of Mathematics Studies* **201**. Princeton Univ. Press, Princeton, NJ. MR3967062

- [12] CARDALIAGUET, P. and RAINER, C. (2009). Stochastic differential games with asymmetric information. *Appl. Math. Optim.* **59** 1–36. MR2465706 <https://doi.org/10.1007/s00245-008-9042-0>
- [13] CARMONA, R. (2016). *Lectures on BSDEs, Stochastic Control, and Stochastic Differential Games with Financial Applications. Financial Mathematics 1*. SIAM, Philadelphia, PA. MR3629171
- [14] CARMONA, R., CERENZIA, M. and PALMER, A. Z. (2018). The Dyson game. arXiv preprint [arXiv:1808.02464](https://arxiv.org/abs/1808.02464).
- [15] CARMONA, R. and DELARUE, F. (2013). Probabilistic analysis of mean-field games. *SIAM J. Control Optim.* **51** 2705–2734. MR3072222 <https://doi.org/10.1137/120883499>
- [16] CARMONA, R. and DELARUE, F. (2017). *Probabilistic Theory of Mean Field Games: Vol. I, Mean Field FBSDEs, Control, and Games. Stochastic Analysis and Applications*. Springer Verlag.
- [17] CARMONA, R. and DELARUE, F. (2018). *Probabilistic Theory of Mean Field Games with Applications. II: Mean Field Games with Common Noise and Master Equations. Probability Theory and Stochastic Modelling 84*. Springer, Cham. MR3753660
- [18] CARMONA, R., DELARUE, F. and LACKER, D. (2016). Mean field games with common noise. *Ann. Probab.* **44** 3740–3803. MR3572323 <https://doi.org/10.1214/15-AOP1060>
- [19] 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/10.4310/CMS.2015.v13.n4.a4>
- [20] CECCHIN, A., DAI PRA, P., FISCHER, M. and PELINO, G. (2019). On the convergence problem in mean field games: A two state model without uniqueness. *SIAM J. Control Optim.* **57** 2443–2466. MR3981375 <https://doi.org/10.1137/18M1222454>
- [21] CECCHIN, A. and PELINO, G. (2019). Convergence, fluctuations and large deviations for finite state mean field games via the master equation. *Stochastic Process. Appl.* **129** 4510–4555. MR4013871 <https://doi.org/10.1016/j.spa.2018.12.002>
- [22] CHASSAGNEUX, J.-F., CRISAN, D. and DELARUE, F. (2014). A probabilistic approach to classical solutions of the master equation for large population equilibria. arXiv preprint [arXiv:1411.3009](https://arxiv.org/abs/1411.3009).
- [23] DELARUE, F. (2017). Mean field games: A toy model on an Erdős–Rényi graph. In *Journées MAS 2016 de la SMAI—Phénomènes Complexes et Hétérogènes. ESAIM Proc. Surveys 60* 1–26. EDP Sci., Les Ulis. MR3772471
- [24] DELARUE, F. and FOGUEN TCHUENDOM, R. (2020). Selection of equilibria in a linear quadratic mean-field game. *Stochastic Process. Appl.* **130** 1000–1040. MR4046528 <https://doi.org/10.1016/j.spa.2019.04.005>
- [25] DELARUE, F., LACKER, D. and RAMANAN, K. (2020). From the master equation to mean field game limit theory: Large deviations and concentration of measure. *Ann. Probab.* **48** 211–263. MR4079435 <http://dx.doi.org/10.1214/19-AOP1359>
- [26] DELARUE, F., LACKER, D. and RAMANAN, K. (2019). From the master equation to mean field game limit theory: A central limit theorem. *Electron. J. Probab.* **24** Paper No. 51, 54. MR3954791 <https://doi.org/10.1214/19-EJP298>
- [27] EL KAROUI, N. and MÉLÉARD, S. (1990). Martingale measures and stochastic calculus. *Probab. Theory Related Fields 84* 83–101. MR1027822 <https://doi.org/10.1007/BF01288560>
- [28] EL KAROUI, N., NGUYEN, D. 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/10.1080/17442508708833443>
- [29] EL KAROUI, N., NGUYEN, D. H. and JEANBLANC-PICQUÉ, M. (1988). Existence of an optimal Markovian filter for the control under partial observations. *SIAM J. Control Optim.* **26** 1025–1061. MR0957652 <https://doi.org/10.1137/0326057>
- [30] FELEQI, E. (2013). The derivation of ergodic mean field game equations for several populations of players. *Dyn. Games Appl.* **3** 523–536. MR3127148 <https://doi.org/10.1007/s13235-013-0088-5>
- [31] FIGALLI, A. (2008). Existence and uniqueness of martingale solutions for SDEs with rough or degenerate coefficients. *J. Funct. Anal.* **254** 109–153. MR2375067 <https://doi.org/10.1016/j.jfa.2007.09.020>
- [32] FILIPPOV, A. F. (1962). On certain questions in the theory of optimal control. *J. SIAM Control Ser. A 1* 76–84. MR0149985
- [33] FISCHER, M. (2017). On the connection between symmetric N -player games and mean field games. *Ann. Appl. Probab.* **27** 757–810. MR3655853 <https://doi.org/10.1214/16-AAP1215>
- [34] FLEMING, W. H. (1977). Generalized solutions in optimal stochastic control. In *Differential Games and Control Theory, II (Proc. 2nd Conf., Univ. Rhode Island, Kingston, R.I., 1976) Lecture Notes in Pure and Appl. Math. 30* 147–165. MR0688675
- [35] FLEMING, W. H. and NISIO, M. (1984). On stochastic relaxed control for partially observed diffusions. *Nagoya Math. J.* **93** 71–108. MR0738919 <https://doi.org/10.1017/S0027763000020742>
- [36] FRIEDMAN, A. (1972). Stochastic differential games. *J. Differential Equations 11* 79–108. MR0292528 [https://doi.org/10.1016/0022-0396\(72\)90082-4](https://doi.org/10.1016/0022-0396(72)90082-4)

- [37] FUDENBERG, D. and LEVINE, D. (1986). Limit games and limit equilibria. *J. Econom. Theory* **38** 261–279. MR0841698 [https://doi.org/10.1016/0022-0531\(86\)90118-3](https://doi.org/10.1016/0022-0531(86)90118-3)
- [38] FUDENBERG, D. and LEVINE, D. (2009). Open-loop and closed-loop equilibria in dynamic games with many players. In *A Long-Run Collaboration on Long-Run Games* 41–58. World Scientific.
- [39] GANGBO, W. and SWIECH, A. (2015). Existence of a solution to an equation arising from the theory of mean field games. *J. Differential Equations* **259** 6573–6643. MR3397332 <https://doi.org/10.1016/j.jde.2015.08.001>
- [40] GÄRTNER, J. (1988). On the McKean–Vlasov limit for interacting diffusions. *Math. Nachr.* **137** 197–248. MR0968996 <https://doi.org/10.1002/mana.19881370116>
- [41] GYÖNGY, I. (1986). Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probab. Theory Related Fields* **71** 501–516. MR0833267 <https://doi.org/10.1007/BF00699039>
- [42] HAMADENE, S. and LEPELTIER, J.-P. (1995). Zero-sum stochastic differential games and backward equations. *Systems Control Lett.* **24** 259–263. MR1321134 [https://doi.org/10.1016/0167-6911\(94\)00011-J](https://doi.org/10.1016/0167-6911(94)00011-J)
- [43] HAUSSMANN, U. G. and LEPELTIER, J.-P. (1990). On the existence of optimal controls. *SIAM J. Control Optim.* **28** 851–902. MR1051628 <https://doi.org/10.1137/0328049>
- [44] HUANG, M., CAINES, P. E. and MALHAMÉ, R. P. (2007). Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ϵ -Nash equilibria. *IEEE Trans. Automat. Control* **52** 1560–1571. MR2352434 <https://doi.org/10.1109/TAC.2007.904450>
- [45] 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
- [46] JABIN, P.-E. and WANG, Z. (2016). Mean field limit and propagation of chaos for Vlasov systems with bounded forces. *J. Funct. Anal.* **271** 3588–3627. MR3558251 <https://doi.org/10.1016/j.jfa.2016.09.014>
- [47] 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
- [48] KALLENBERG, O. (2002). *Foundations of Modern Probability*, 2nd ed. *Probability and Its Applications (New York)*. Springer, New York. MR1876169 <https://doi.org/10.1007/978-1-4757-4015-8>
- [49] KRYLOV, N. V. and RÖCKNER, M. (2005). Strong solutions of stochastic equations with singular time dependent drift. *Probab. Theory Related Fields* **131** 154–196. MR2117951 <https://doi.org/10.1007/s00440-004-0361-z>
- [50] LACKER, D. (2015). Mean field games via controlled martingale problems: Existence of Markovian equilibria. *Stochastic Process. Appl.* **125** 2856–2894. MR3332857 <https://doi.org/10.1016/j.spa.2015.02.006>
- [51] LACKER, D. (2016). A general characterization of the mean field limit for stochastic differential games. *Probab. Theory Related Fields* **165** 581–648. MR3520014 <https://doi.org/10.1007/s00440-015-0641-9>
- [52] LACKER, D. (2018). On a strong form of propagation of chaos for McKean–Vlasov equations. *Electron. Commun. Probab.* **23** Paper No. 45, 11. MR3841406 <https://doi.org/10.1214/18-ECP150>
- [53] LASRY, J.-M. and LIONS, P.-L. (2006). Jeux à champ moyen. I. Le cas stationnaire. *C. R. Math. Acad. Sci. Paris* **343** 619–625. MR2269875 <https://doi.org/10.1016/j.crma.2006.09.019>
- [54] LASRY, J.-M. and LIONS, P.-L. (2006). Jeux à champ moyen. II. Horizon fini et contrôle optimal. *C. R. Math. Acad. Sci. Paris* **343** 679–684. MR2271747 <https://doi.org/10.1016/j.crma.2006.09.018>
- [55] LASRY, J.-M. and LIONS, P.-L. (2007). Mean field games. *Jpn. J. Math.* **2** 229–260. MR2295621 <https://doi.org/10.1007/s11537-007-0657-8>
- [56] NEUFELD, A. and NUTZ, M. (2014). Measurability of semimartingale characteristics with respect to the probability law. *Stochastic Process. Appl.* **124** 3819–3845. MR3249357 <https://doi.org/10.1016/j.spa.2014.07.006>
- [57] NUTZ, M., SAN MARTIN, J. and TAN, X. (2018). Convergence to the mean field game limit: A case study. arXiv preprint [arXiv:1806.00817](https://arxiv.org/abs/1806.00817).
- [58] OELSCHLÄGER, K. (1984). A martingale approach to the law of large numbers for weakly interacting stochastic processes. *Ann. Probab.* **12** 458–479. MR0735849
- [59] PROTTER, P. E. (2005). *Stochastic Integration and Differential Equations*, 2nd ed. *Stochastic Modelling and Applied Probability* **21**. Springer, Berlin. Version 2.1, Corrected third printing. MR2273672 <https://doi.org/10.1007/978-3-662-10061-5>
- [60] ROXIN, E. (1962). The existence of optimal controls. *Michigan Math. J.* **9** 109–119. MR0136844
- [61] ROYDEN, H. L. (1968). *Real Analysis*, 2nd ed. Macmillan. MR0151555
- [62] STROOCK, D. W. and VARADHAN, S. R. S. (2006). *Multidimensional Diffusion Processes. Classics in Mathematics*. Springer, Berlin. Reprint of the 1997 edition. MR2190038
- [63] SZNITMAN, A.-S. (1991). Topics in propagation of chaos. In *École D’Été de Probabilités de Saint-Flour XIX—1989. Lecture Notes in Math.* **1464** 165–251. Springer, Berlin. MR1108185 <https://doi.org/10.1007/BFb0085169>

- [64] TREVISAN, D. (2013). Zero noise limits using local times. *Electron. Commun. Probab.* **18** no. 31, 7. MR3056068 <https://doi.org/10.1214/ECP.v18-2587>
- [65] TREVISAN, D. (2016). Well-posedness of multidimensional diffusion processes with weakly differentiable coefficients. *Electron. J. Probab.* **21** Paper No. 22, 41. MR3485364 <https://doi.org/10.1214/16-EJP4453>
- [66] VERETENNIKOV, A. Y. (1981). On strong solutions and explicit formulas for solutions of stochastic integral equations. *Math. USSR, Sb.* **39** 387.
- [67] VILLANI, C. (2003). *Topics in Optimal Transportation. Graduate Studies in Mathematics* **58**. Amer. Math. Soc., Providence, RI. MR1964483 <https://doi.org/10.1090/gsm/058>
- [68] WALSH, J. B. (1986). An introduction to stochastic partial differential equations. In *École D'été de Probabilités de Saint-Flour, XIV—1984. Lecture Notes in Math.* **1180** 265–439. Springer, Berlin. MR0876085 <https://doi.org/10.1007/BFb0074920>