NOTES FOR THE AMS SHORT COURSE ON MEAN FIELD GAMES

DANIEL LACKER

These notes are to supplement the AMS short course on mean field games on January 13–14, 2020. The aim of of these notes in particular is to develop some elements of the limit theory, connecting n-player models to their mean field game limits.

Some of the material is borrowed from my lectures notes from a PhD course I taught at Columbia University,[1] which was an overview of interacting diffusion models, stochastic control, stochastic games, mean field games and the master equation. Those notes, for a semester-long course, are naturally much more comprehensive than these. Both are nowhere near as comprehensive or rigorous in their treatment of the subject as the recent two-volume textbook of Carmona-Delarue [7]. On my website[2] you can also find a set of lecture notes devoted to the history and application of the "compactification methods" which are developed briefly in Section 4 below. This is to mention only some recent probabilistic accounts of mean field game theory, and the notes of Cardaliaguet [4] remain an excellent introduction to the topic from a PDE perspective. Additional surveys and perspectives may be found in [1, 12]. It is safe to say that mean field game theory now boasts a wide range of reading materials tailored to readers of various backgrounds and attention spans.

Contents

[1] http://www.columbia.edu/~dl3133/MFGSpring2018.pdf

[2] http://www.columbia.edu/~dl3133/IPAM-MFGCompactnessMethods.pdf

1. Notation and references for background material

For a metric space $\mathcal{X}$, we write $\mathcal{P}(\mathcal{X})$ for the set of Borel probability measures on $\mathcal{X}$, and $C_b(\mathcal{X})$ for the set of bounded continuous functions on $\mathcal{X}$. We write $\mathcal{L}(Z)$ for the law or distribution of a random variable Z. For $x \in \mathcal{X}$ we write δ_x for the Dirac mass at x, i.e., the probability measure defined by $\delta_x(A) = 1$ if $x \in A$ and 0 otherwise. For $n \in \mathbb{N}$ and a vector $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathcal{X}^n$, we write $m^n_{\boldsymbol{x}}$ for the empirical measure, i.e.,

$$m^n_{\boldsymbol{x}} = \frac{1}{n} \sum_{i=1}^n \delta_{x_i}.$$

We always equip $\mathcal{P}(\mathcal{X})$ with the topology of weak convergence, by which $\mu_n \to \mu$ in $\mathcal{P}(\mathcal{X})$ if and only if $\int \varphi \, d\mu_n \to \int \varphi \, d\mu$ for all $\varphi \in C_b(\mathcal{X})$. We use also the standard abbreviations

$$\langle \mu, \varphi \rangle = \int \varphi \, d\mu = \int_{\mathcal{X}} \varphi(x) \, \mu(dx),$$

for any $\mu \in \mathcal{P}(\mathcal{X})$ and any μ-integrable function φ on $\mathcal{X}$. We refer to appendix A for a summary of some basic facts on weak convergence of probability measures and further references on the topic.

The study of mean field game theory from the probabilistic perspective requires some basic knowledge of stochastic calculus and stochastic optimal control. For background on stochastic calculus, the reader is referred to the classic book of Karatzas and Shreve [15]. The book of Pham [26] is a good introductory text on stochastic control theory, and the first chapter serves as a good review of stochastic calculus as well. The book of Fleming-Soner [10] is also a classic reference, with more emphasis on PDE methods.

2. Interacting diffusions and McKean-Vlasov equations

A first step toward understanding mean field games is to understand their uncontrolled counterparts, which we study in this section. Behind the scenes throughout this section is a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, with $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ satisfying the usual conditions, assumed rich enough to support whatever we ask of it, in particular independent d-dimensional $\mathbb{F}$-Brownian motions.

The main object of study will be a system of n interacting particles $\boldsymbol{X}^n_t = (X^{n,1}_t, \ldots, X^{n,n}_t)$, driven by stochastic differential equations (SDEs) of the form

$$dX^{n,i}_t = b(X^{n,i}_t, m^n_{\boldsymbol{X}^n_t})dt + dW^i_t. \tag{1}$$

Driving this SDE system are n independent Brownian motions, $W^1, \ldots, W^n$, and we typically but not always assume the initial states $X^{n,i}_0$ are i.i.d. We think of $X^{n,i}_t$ as the position of particle i at time t, in a Euclidean space $\mathbb{R}^d$. We could include a state- and measure-dependent diffusion coefficient $\sigma(X^{n,i}_t, m^n_{\boldsymbol{X}^n_t})$, but we choose not to for simplicity of notation.

We think of the number n of particles as very large, and ultimately we will send it to infinity. There is a key structural feature that makes this system amenable to *mean field* analysis: The coefficients b and σ are the same for each particle, and the only dependence of particle i on the rest of the particles $k \neq i$ is through the empirical measure $m^n_{\boldsymbol{X}^n_t}$. Let us build some intuition with a simple example:

2.1. A first example.

In this section we study in more detail a warm-up model mentioned in the introduction. Consider the SDE system (1), with $d = 1$-dimensional particles, with the drift

$$b(x, m) = a \left(x - \int_{\mathbb{R}} y \, m(dy) \right).$$

Here $a > 0$, and we can write more explicitly

$$\begin{aligned} dX_t^{n,i} &= a\left(\overline{X}_t^n - X_t^{n,i}\right) dt + dW_t^i, \quad i = 1, 2, \ldots, n, \\ \overline{X}_t^n &= \frac{1}{n}\sum_{k=1}^n X_t^{n,k}. \end{aligned} \tag{2}$$

The drift pushes each particle toward the empirical average $\overline{X}_t^n$. This is like an Ornstein-Uhlenbeck equation, but the target of mean-reversion is dynamic. To understand how the system behaves for large n, a good way to start is by noticing that if we average the n particles we get very simple dynamics for $\overline{X}_t^n$:

$$d\overline{X}_t^n = \frac{1}{n}\sum_{i=1}^n dW_t^i.$$

In particular,

$$\overline{X}_t^n = \overline{X}_0^n + \frac{1}{n}\sum_{i=1}^n W_t^i.$$

Sending $n \to \infty$, the average of the n Brownian motions vanishes thanks to the law of large numbers. Moreover, if the initial states $(X_0^{n,i})_{n\ge i\ge 1}$ are i.i.d., the empirical average $\overline{X}_0^n$ converges to the true mean, say $\bar{m} \in \mathbb{R}$. Hence, when $n \to \infty$, the empirical average becomes $\lim_n \overline{X}_t^n = \bar{m}$ for all t, almost surely. Plugging this back in to the original equation (2), we find that $\lim_n X^{n,i} = Y^i$ a.s., where

$$dY_t^i = a\left(\bar{m} - Y_t^i\right) dt + dW_t^i, \quad Y_0^i = X_0^{n,i}. \tag{3}$$

We can solve this equation pretty easily. First, writing it in integral form, we have

$$Y_t^i = X_0^{n,i} + a\int_0^t \left(\bar{m} - Y_s^i\right) ds + W_t^i.$$

Take expectations to get

$$\mathbb{E}[Y_t^i] = \bar{m} + a\int_0^t \left(\bar{m} - \mathbb{E}[Y_s^i]\right) ds.$$

Differentiate in t to get

$$\frac{d}{dt}\mathbb{E}[Y_t^i] = a\left(\bar{m} - \mathbb{E}[Y_t^i]\right), \qquad \mathbb{E}[Y_0^i] = \bar{m}.$$

This shows that the function $t \mapsto \mathbb{E}[Y_t^i]$ solves a very simple ordinary differential equation, the solution of which is constant, $\mathbb{E}[Y_t^i] = \bar{m}$. Hence, we may rewrite (3) as

$$dY_t^i = a\left(\mathbb{E}[Y_t^i] - Y_t^i\right) dt + dW_t^i, \quad Y_0^i = X_0^{n,i}. \tag{4}$$

This is an example of a *McKean-Vlasov* equation, an SDE in which the coefficients depend on the law of the solution.

It important to observe that the resulting limiting processes $(Y^i)_{i\in\mathbb{N}}$ are i.i.d. They solve the same SDEs, driven by i.i.d. Brownian motions and i.i.d. initial states. This is a general phenomenon with McKean-Vlasov limits; the particles become asymptotically i.i.d. as $n \to \infty$, in a sense we will soon make precise.

2.2. **Deriving the McKean-Vlasov limit.** While the example of Section 2.1 was simple enough to allow explicit computation, we now set to work on a general understanding of the $n \to \infty$ behavior of systems of the form (1). First, we should specify assumptions on the coefficients to let us ensure at the very least that the n-particle SDE system (1) is well-posed. In this section, we work with the following set of nice assumptions, recalling from (25) the notation for Wasserstein metrics:

Assumption 2.1. Assume the initial states $(X_0^{n,i})_{n \ge i \ge 1}$ are i.i.d. with $\mathbb{E}[|X_0^{n,i}|^2] < \infty$. The Brownian motions $(W^i)_{i \in \mathbb{N}}$ are independent and m-dimensional. Assume $b : \mathbb{R}^d \times \mathcal{P}^2(\mathbb{R}^d) \to \mathbb{R}^d$ is Lipschitz, in the sense that there exists a constant $L > 0$ such that

$$|b(x,m) - b(x',m')| \le L(|x - x'| + \mathcal{W}_2(m,m')). \tag{5}$$

Note that we always write $|\cdot|$ for the Euclidean norm on $\mathbb{R}^d$ and $|\cdot|$ for the Frobenius norm on $\mathbb{R}^{d\times m}$.

This assumption immediately lets us check that the SDE system (1) is well-posed. We make heavy use of the following well known well-posedness result:

Lemma 2.2. *Under Assumption 2.1, the n-particle SDE system* (1) *admits a unique strong solution, for each n.*

Proof. We will fall back on Itô's classical existence and uniqueness result for Lipschitz SDEs. Define the $\mathbb{R}^{nd}$-valued process $\boldsymbol{X}_t = (X_t^{n,1}, \ldots, X_t^{n,n})$, and similar define the nm-dimensional Brownian motion $\boldsymbol{W}_t = (W_t^1, \ldots, W_t^n)$. We may write

$$d\boldsymbol{X}_t = B(\boldsymbol{X}_t)dt + d\boldsymbol{W}_t, \tag{6}$$

if we define also $B : \mathbb{R}^{nd} \to \mathbb{R}^{nd}$ for $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathbb{R}^{nd}$ by

$$B(\boldsymbol{x}) = \begin{pmatrix} b(x_1, m^n_{\boldsymbol{x}}) \\ b(x_2, m^n_{\boldsymbol{x}}) \\ \vdots \\ b(x_n, m^n_{\boldsymbol{x}}) \end{pmatrix}.$$

Then, for $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{nd}$, we have

$$\begin{aligned} |B(\boldsymbol{x}) - B(\boldsymbol{y})|^2 &= \sum_{i=1}^n |b(x_i, m^n_{\boldsymbol{x}}) - b(y_i, m^n_{\boldsymbol{y}})|^2 \\ &\le 2L^2 \sum_{i=1}^n \left(|x_i - y_i|^2 + \mathcal{W}_2^2(m^n_{\boldsymbol{x}}, m^n_{\boldsymbol{y}}) \right) \\ &= 2L^2 |\boldsymbol{x} - \boldsymbol{y}|^2 + 2L^2 n \mathcal{W}_2^2(m^n_{\boldsymbol{x}}, m^n_{\boldsymbol{y}}), \end{aligned}$$

where the second line used the elementary inequality $(a+b)^2 \le 2a^2 + 2b^2$. Next, to bound the Wasserstein distance between empirical measures, we use the most natural coupling; namely, the empirical measure

$$\pi = \frac{1}{n} \sum_{i=1}^n \delta_{(x_i, y_i)}$$

is a coupling of $m^n_{\boldsymbol{x}}$ and $m^n_{\boldsymbol{y}}$. Hence,

$$\begin{aligned}\mathcal{W}_2^2(m^n_{\boldsymbol{x}}, m^n_{\boldsymbol{y}}) &\le \int_{\mathbb{R}^d\times\mathbb{R}^d} |x-y|^2\pi(dx,dy) \\ &= \frac{1}{n}\sum_{i=1}^n |x_i - y_i|^2 \\ &= \frac{1}{n}|\boldsymbol{x}-\boldsymbol{y}|^2.\end{aligned}$$

We conclude that $|B(\boldsymbol{x}) - B(\boldsymbol{y})|^2 \le 4L^2|\boldsymbol{x}-\boldsymbol{y}|^2$, which means that B is $2L$-Lipschitz. Conclude that the SDE (6), which was simply a rewriting of our original SDE system, has a unique strong solution. □

Now, to send $n \to \infty$ in our particle system, we start by "guessing" what the limit should look like. As $n \to \infty$, the interaction becomes weaker and weaker in the sense that the contribution of a given particle i is of order $1/n$, intuitively. Hence, as $n \to \infty$, we expect the interaction to vanish, in the sense that a given particle i does not appear in the measure term anymore. If particles do not affect the measure flow, then the particles should be i.i.d., as they have the same coefficients b and σ and are driven by independent Brownian motions. For a better, more mathematical derivation, see Section 2.4.

For now, the above guess leads us to expect that if $m^n_{\boldsymbol{X}^n_t} \to \mu_t$ in some sense, where μ_t is a *deterministic* measure flow (i.e., a function $\mathbb{R}_+ \ni t \mapsto \mu_t \in \mathcal{P}^2(\mathbb{R}^d)$), then the dynamics of any particle particle will look like

$$dY^i_t = b(Y^i_t, \mu_t)dt + dW^i_t.$$

These particles are i.i.d., but μ_t should still somehow represent their distribution. We know that the empirical measure of i.i.d. samples converges to the true distribution (see Theorem A.7), so we should expect that μ_t is actually the law of Y^i_t, for any i. In other words, the law of the solution shows up in the coefficients of the SDE! We call this a *McKean-Vlasov equation*, after the seminal work of McKean [21, 22].

To formulate the McKean-Vlasov equation more precisely, it is convenient and often more general to lift the discussion to the path space, in the following sense. For convenience, we will fix a time horizon $T > 0$, but it is straightforward to extend much of the discussion to follow to the infinite time interval. Let $\mathcal{C}^d = C([0,T];\mathbb{R}^d)$ denote the set of continuous $\mathbb{R}^d$-valued functions of time, equipped with the supremum norm

$$\|x\| = \sup_{t\in[0,T]} |x_t|$$

and the corresponding Borel σ-field. Rather than work with measure flows, as elements of $C([0,T];\mathcal{P}^2(\mathbb{R}^d))$, we will work with probability measures on $\mathcal{C}^d$. There is a natural surjection

$$\mathcal{P}^2(\mathcal{C}^d) \ni \mu \to (\mu_t)_{t\in[0,T]} \in C([0,T];\mathcal{P}^2(\mathbb{R}^d)),$$

where μ_t is defined as the image of the measure μ through the map $\mathcal{C}^d \ni x \mapsto x_t \in \mathbb{R}^d$. In other words, if $\mu = \mathcal{L}(X)$ is the law of a $\mathcal{C}^d$-valued random variable, then $\mu_t = \mathcal{L}(X_t)$ is the time-t marginal. Note that this surjection is continuous and, in fact, 1-Lipschitz, if $C([0,T];\mathcal{P}^2(\mathbb{R}^d))$ is endowed with the supremum distance induced by the metric $\mathcal{W}_2$ on $\mathbb{R}^d$; explicitly, we have

$$\sup_{t\in[0,T]} \mathcal{W}_{\mathbb{R}^d,2}(\mu_t,\nu_t) \le \mathcal{W}_{\mathcal{C}^d,2}(\mu,\nu) \tag{7}$$

for every $\mu, \nu \in \mathcal{P}^2(\mathcal{C}^d)$. Viewing $\boldsymbol{X}^n = (X^{n,1}, \ldots, X^{n,n})$ as a vector of random elements of $\mathcal{C}^d$, we will in fact show that $m^n_{\boldsymbol{X}^n}$ converges in probability in $\mathcal{P}^2(\mathcal{C}^d)$ to a non-random limit, which is characterized as the law of the unique solution of the McKean-Vlasov equation.

The McKean-Vlasov equation is defined precisely as follows:

$$\begin{aligned} dY_t &= b(Y_t, \mu_t)dt + dW_t, \quad t \in [0,T], \quad Y_0 = \xi, \\ \mu &= \mathcal{L}(Y). \end{aligned} \tag{8}$$

Here we write $\mathcal{L}(Z)$ for the law or distribution of a random variable Z. Here W is a Brownian motion, ξ is an $\mathbb{R}^d$-valued random variable with the same law as $X^{n,i}_0$ (which we recall are i.i.d.), and both are (say) supported on the same filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$. Of course, ξ should be $\mathcal{F}_0$-measurable, and W should be an $\mathbb{F}$-Brownian motion. A strong solution of the McKean-Vlasov equation is a pair (Y, μ), where Y is a continuous $\mathbb{F}$-adapted $\mathbb{R}^d$-valued process (alternatively, a $C([0,T];\mathbb{R}^d)$-valued random variable), and μ is a probability measure on $C([0,T];\mathbb{R}^d)$, such that both equations (8) hold simultaneously. More compactly, one could simply write

$$dY_t = b(Y_t, \mathcal{L}(Y_t))dt + dW_t, \quad t \in [0,T], \quad Y_0 = \xi.$$

Under the Lipschitz assumption, we can always uniquely solve for Y if we know μ, and so we sometimes refer to μ itself (instead of the pair (Y, μ)) as the solution of the McKean-Vlasov equation.

We formalize all of this in the following theorem. We will make use of the 2-Wasserstein distance $\mathcal{W}_2 = \mathcal{W}_{\mathcal{C}^d,2}$, which was defined in (25).

Theorem 2.3. *Suppose Assumption 2.1 holds. There exists a unique strong solution of the McKean-Vlasov equation* (8)*. Moreover, the n-particle system converges in the following two senses. First,*

$$\lim_{n\to\infty} \mathbb{E}\left[\mathcal{W}^2_{\mathcal{C}^d,2}(m^n_{\boldsymbol{X}^n}, \mu)\right] = 0. \tag{9}$$

Second, for a fixed $k \in \mathbb{N}$, we have the weak convergence

$$\mathcal{L}(X^{n,1}, \ldots, X^{n,k}) \to \mathcal{L}(Y^1, \ldots, Y^k), \tag{10}$$

where $Y^1, \ldots, Y^k$ are independent copies of the solution of the McKean-Vlasov equation.

We interpret the second form of the limit (10) as saying that the particles $X^{n,i}$ become asymptotically i.i.d. as $n \to \infty$. Indeed, for any fixed k, the first k particles converge in distribution to i.i.d. random variables. The choice of the "first k" particles here is inconsequential, in light of the fact that the n-particle system (1) is *exchangeable*, in the sense that

$$(X^{n,\pi(1)}, \ldots, X^{n,\pi(n)}) \stackrel{d}{=} (X^{n,1}, \ldots, X^{n,n})$$

for any permutation π of $\{1, \ldots, n\}$. (Here $\stackrel{d}{=}$ means equal in distribution.) It is left as an exercise for the reader to prove this claimed exchangeability, using the uniqueness of the solution of the SDE system (1).

Remark 2.4. The two kinds of limits described in Theorem 2.3 are sometimes referred to as *propagation of chaos*, though this terminology has somewhat lost its original meaning. One would say that propagation of chaos holds for the interacting particle system (1) if the following holds: For any $m_0 \in \mathcal{P}(\mathbb{R}^d)$ and any choice of deterministic initial states $(X^{n,i}_0)$ satisfying

$$\frac{1}{n}\sum_{i=1}^n \delta_{X^{n,i}_0} \to m_0, \tag{11}$$

we have the weak limit $\frac{1}{n}\sum_{i=1}^n \delta_{X_t^{n,i}} \rightarrow \mu_t$ in probability in $\mathcal{P}(\mathbb{R}^d)$, where (Y,μ) is the solution of the McKean-Vlasov equation (8) with initial state $\xi \sim m_0$. Initial states which converge weakly as in (11) are called "m_0-chaotic," and the term *propagation of chaos* means that the "chaoticity" of the initial distributions "propagates" to later times $t > 0$.

We break up the proof of Theorem 2.3 into a two major steps. First, we show existence and uniqueness:

Proof of existence and uniqueness. Define the truncated supremum norm

$$\|x\|_t := \sup_{0\le s\le t} |x_s|$$

for $x \in \mathcal{C}^d$ and $t \in [0,T]$. Using this, define the truncated Wasserstein distance on $\mathcal{P}^2(\mathcal{C}^d)$ by (recalling the notation $\Pi(\mu,\nu)$ for couplings introduced just before (25))

$$d_t^2(\mu,\nu) = \inf_{\pi\in\Pi(\mu,\nu)} \int_{\mathcal{C}^d\times\mathcal{C}^d} \|x-y\|_t^2\, \pi(dx,dy). \tag{12}$$

Notice that for any fixed $\mu \in \mathcal{P}^2(\mathcal{C}^d)$ we have the Lipschitz condition

$$|b(x,\mu_t) - b(y,\mu_t)| \le L|x-y|,$$

and thus classical theory ensures that there exists a unique (square-integrable) solution Y^μ of the SDE

$$dY_t^\mu = b(Y_t^\mu,\mu_t)dt + dW_t, \quad Y_0^\mu = \xi.$$

Define a map $\Phi : \mathcal{P}^2(\mathcal{C}^d) \to \mathcal{P}^2(\mathcal{C}^d)$ by setting $\Phi(\mu) = \mathcal{L}(Y^\mu)$. That is, for a fixed μ, we solve the above SDE, and set $\Phi(\mu)$ to be the law of the solution. Note that (Y,μ) is a solution of the McKean-Vlasov equation if and only if $Y = Y^\mu$ and $\mu = \Phi(\mu)$. That is, fixed points of Φ are precisely solutions of the McKean-Vlasov equation.

Let $\mu,\nu \in \mathcal{P}^2(\mathcal{C}^d)$, and use Jensen's inequality to get, for $t \in [0,T]$,

$$|Y_t^\mu - Y_t^\nu|^2 \le t\int_0^t |b(Y_r^\mu,\mu_r) - b(Y_r^\nu,\nu_r)|^2\, dr.$$

Take the supremum and use the Lipschitz assumption to find

$$\begin{aligned}
\mathbb{E}\left[\|Y^\mu - Y^\nu\|_t^2\right] &\le t\mathbb{E}\left[\int_0^t |b(Y_r^\mu,\mu_r) - b(Y_r^\nu,\nu_r)|^2\, dr\right] \\
&\le t\mathbb{E}\left[\int_0^t |b(Y_r^\mu,\mu_r) - b(Y_r^\nu,\nu_r)|^2\, dr\right] \\
&\le 2tL^2\mathbb{E}\left[\int_0^t \left(|Y_r^\mu - Y_r^\nu|^2 + \mathcal{W}_2^2(\mu_r,\nu_r)\right) dr\right] \\
&\le 2tL^2\mathbb{E}\left[\int_0^t \left(\|Y^\mu - Y^\nu\|_r^2 + \mathcal{W}_2^2(\mu_r,\nu_r)\right) dr\right].
\end{aligned}$$

Use Fubini's theorem and Gronwall's inequality to conclude that

$$\mathbb{E}\left[\|Y^\mu - Y^\nu\|_t^2\right] \le C\int_0^t \mathcal{W}_2^2(\mu_r,\nu_r)dr,$$

for $t \in [0,T]$, where $C = 2TL^2\exp(2TL^2)$. Use a form of the inequality (7) to get

$$\mathbb{E}\left[\|Y^\mu - Y^\nu\|_t^2\right] \le C\int_0^t d_r^2(\mu,\nu)dr.$$

Finally, recall the definition of d_t^2, and notice that the joint distribution of (Y^μ, Y^ν) is a coupling of $(\Phi(\mu), \Phi(\nu))$. Hence,

$$d_t^2(\Phi(\mu), \Phi(\nu)) \le \mathbb{E}\left[\|Y^\mu - Y^\nu\|_t^2\right] \le C\int_0^t d_r^2(\mu,\nu)dr.$$

The proof of existence and uniqueness now follows from the usual Picard iteration argument. In particular, uniqueness follows from the above inequality and another application of Gronwall's inequality, whereas existence is derived by choosing arbitrarily $\mu^0 \in \mathcal{P}^2(\mathcal{C}^d)$, setting $\mu^{k+1} = \Phi(\mu^k)$ for $k \ge 0$, and showing using the above inequality that (μ^k) forms a Cauchy sequence whose limit must be a fixed point of Φ. □

Proof of McKean-Vlasov limit. We now move to the second part of Theorem 2.3, proving the claimed limit theorem. The idea is what is by a coupling argument, where we construct i.i.d. copies of the unique solution of the McKean-Vlasov equation which has desirable joint-distributional properties with the original particle system. To do this, let μ denote the unique solution of the McKean-Vlasov equation (8). Using the same Brownian motions (W^i) and initial states (ξ^i) as our original particle system (1), define Y^i as the solution of the SDE

$$dY_t^i = b(Y_t^i, \mu_t)dt + dW_t^i, \quad Y_0^i = \xi^i.$$

Because the initial states and Brownian motions are i.i.d., so are (Y^i). We now want to estimate the difference $|X_t^{n,i} - Y_t^i|$. To do this, we proceed as in the previous step, starting with a fixed i and n:

$$|X_t^{n,i} - Y_t^i|^2 \le t\int_0^t \left|b(X_r^{n,i}, m^n_{\boldsymbol{X}_r^n}) - b(Y_r^i, \mu_r)\right|^2 dr.$$

Take the supremum and use the Lipschitz assumption to find

$$\begin{aligned}
\mathbb{E}\left[\|X^{n,i} - Y^i\|_t^2\right] &\le t\mathbb{E}\left[\int_0^t \left|b(X_r^{n,i}, m^n_{\boldsymbol{X}_r^n}) - b(Y_r^i, \mu_r)\right|^2 dr\right] \\
&\le t\mathbb{E}\left[\int_0^t \left|b(X_r^{n,i}, m^n_{\boldsymbol{X}_r^n}) - b(Y_r^i, \mu_r)\right|^2 dr\right] \\
&\le 2tL^2\mathbb{E}\left[\int_0^t \left(\|X^{n,i} - Y^i\|_r^2 + \mathcal{W}_2^2(m^n_{\boldsymbol{X}_r^n}, \mu_r)\right) dr\right].
\end{aligned}$$

Use Fubini's theorem, Gronwall's inequality, and a form of the inequality (7) to get

$$\begin{aligned}
\mathbb{E}\left[\|X^{n,i} - Y^i\|_t^2\right] &\le C\mathbb{E}\left[\int_0^t \mathcal{W}_2^2(m^n_{\boldsymbol{X}_r^n}, \mu_r)dr\right] \\
&\le C\mathbb{E}\left[\int_0^t d_r^2(m^n_{\boldsymbol{X}^n}, \mu)dr\right],
\end{aligned} \tag{13}$$

for $t \in [0,T]$, where $C = 2TL^2\exp(2TL^2)$, and where we recall the definition of the truncated Wasserstein distance from (12).

Next, note that the empirical measure $\frac{1}{n}\sum_{i=1}^n \delta_{(X^{n,i}, Y^i)}$ is a coupling of the (random) empirical measures $m^n_{\boldsymbol{X}^n}$ and $m^n_{\boldsymbol{Y}^n}$, where $\boldsymbol{Y}^n = (Y^1, \ldots, Y^n)$. Thus

$$d_t^2(m^n_{\boldsymbol{X}^n}, m^n_{\boldsymbol{Y}^n}) \le \frac{1}{n}\sum_{i=1}^n \|X^{n,i} - Y^i\|_t^2, \quad a.s.$$

Combine this with (13) to get

$$\mathbb{E}[d_t^2(m^n_{\boldsymbol{X}^n}, m^n_{\boldsymbol{Y}^n})] \le C\mathbb{E}\left[\int_0^t d_r^2(m^n_{\boldsymbol{X}^n}, \mu)dr\right].$$

Use the triangle inequality and the previous inequality to get

$$\begin{aligned}\mathbb{E}[d_t^2(m^n_{\boldsymbol{X}^n},\mu)] &\le 2\mathbb{E}[d_t^2(m^n_{\boldsymbol{X}^n},m^n_{\boldsymbol{Y}^n})]+2\mathbb{E}[d_t^2(m^n_{\boldsymbol{Y}^n},\mu)]\\ &\le 2C\mathbb{E}\left[\int_0^t d_r^2(m^n_{\boldsymbol{X}^n},\mu)dr\right]+2\mathbb{E}[d_t^2(m^n_{\boldsymbol{Y}^n},\mu)].\end{aligned}$$

Apply Gronwall's inequality once again to get

$$\mathbb{E}[d_t^2(m^n_{\boldsymbol{X}^n},\mu)] \le 2e^{2CT}\mathbb{E}\left[d_t^2(m^n_{\boldsymbol{Y}^n},\mu)\right].$$

In particular, setting $t=T$, we have

$$\mathbb{E}[\mathcal{W}_2^2(m^n_{\boldsymbol{X}^n},\mu)] \le 2e^{2CT}\mathbb{E}[\mathcal{W}_2^2(m^n_{\boldsymbol{Y}^n},\mu)].$$

But $m^n_{\boldsymbol{Y}^n}$ are the empirical measures of i.i.d. samples from the law μ. Hence, by the law of large numbers in the form of Theorem A.7, the right-hand side converges to zero, and the claimed limit (9) follows.

Finally, to prove the second claimed limit (10), we use (13) to find

$$\begin{aligned}\mathbb{E}\left[\max_{i=1,\dots,k}\|X^{n,i}-Y^i\|_t^2\right] &\le \sum_{i=1}^k \mathbb{E}\left[\|X^{n,i}-Y^i\|_t^2\right]\\ &\le Ck\mathbb{E}\left[\int_0^t d_r^2(m^n_{\boldsymbol{X}^n},\mu)dr\right]\\ &\le CkT\mathbb{E}[\mathcal{W}_2^2(m^n_{\boldsymbol{X}^n},\mu)].\end{aligned}$$

We just showed that this converges to zero, and the claimed convergence in distribution follows.

2.3. **A PDE formulation.** In a first course on stochastic calculus, one encounters the Fokker-Planck equation, a linear PDE which describes the behavior of the distribution of the solution of an SDE. A McKean-Vlasov equation like (8) can similarly be described by a Fokker-Planck equation, but it is markedly different in that it is both *nonlinear* and *nonlocal*, in a sense we will soon make clear.

Suppose (Y,μ) solves the McKean-Vlasov equation (8). Apply Itô's formula to $\varphi(Y_t)$, where φ is a smooth function with compact support, to get

$$d\varphi(Y_t)=\left(b(Y_t,\mu_t)\nabla\varphi(Y_t)+\frac{1}{2}\Delta\varphi(Y_t)\right)dt+\nabla\varphi(Y_t)dW_t,$$

where ∇ and Δ denote the gradient and Laplacian operators, respectively, and we understand $b(x,m)\varphi(x)$ as the inner product of the vectors $b(x,m)$ and $\nabla\varphi(x)$. Integrating this equation, taking expectations to kill the martingale term, and applying Fubini's theorem to exchange the expectation and time integral, we find

$$\mathbb{E}[\varphi(Y_t)]=\mathbb{E}[\varphi(Y_0)]+\int_0^t \mathbb{E}\left[b(Y_s,\mu_s)\nabla\varphi(Y_s)+\frac{1}{2}\Delta(Y_s)\right]ds, \qquad \forall t\ge 0.$$

Now, we know that $Y_t\sim\mu_t$. Hence, we may rewrite the above as

$$\langle\mu_t,\varphi\rangle=\langle\mu_0,\varphi\rangle+\int_0^t\left\langle\mu_s, b(\cdot,\mu_s)\nabla\varphi(\cdot)+\frac{1}{2}\Delta\varphi\right\rangle ds \tag{14}$$

(Recall that $\langle\mu_t,\varphi\rangle=\int\varphi\,d\mu_t$.) The equation (14) holds for every $t\ge 0$ and every smooth function φ of compact support, and this can be interpreted as saying that $(\mu_t)_{t\ge 0}$ is a weak solution of a certain PDE. To write this this PDE in strong form, suppose now that, for each $t\ge 0$, μ_t has a density (with respect to Lebesgue measure), which we write as

$\mu(t,x)$. Assume in addition that it admits one continuous derivative in t and two in x. Differentiating the above equation in t yields

$$\begin{aligned}\frac{d}{dt}\int_{\mathbb{R}^d}\varphi(x)\mu(t,x)dx &= \int_{\mathbb{R}^d}\left(b(x,\mu_t)\nabla\varphi(x)+\frac{1}{2}\Delta\varphi(x)\right)\mu(t,x)dx \\ &= \int_{\mathbb{R}^d}\left(-\mathrm{div}_x(b(x,\mu_t)\mu(t,x))+\frac{1}{2}\Delta\mu(t,x)\right)\varphi(x)dx.\end{aligned}$$

Because this must hold for every test function φ, we conclude that $\mu(t,x)$ solves the PDE

$$\partial_t\mu(t,x) = -\mathrm{div}_x(b(x,\mu_t)\mu(t,x)) + \frac{1}{2}\Delta\mu(t,x). \tag{15}$$

Notice that if b and σ did not depend on μ, this would be a usual (linear) Fokker-Planck equation. Note in addition that the nonlinear dependence on μ_t is also typically *nonlocal*, in the sense that $b = b(x,\mu_t)$ is a function not of the value $\mu(t,x)$ but of the entire distribution $(\mu(t,x))_{x\in\mathbb{R}^d}$.

Even if the density of μ_t does not exist or is not sufficiently smooth, is it still clear that (μ_t) is a weak solution of the PDE (15), in the sense that it holds integrated against test functions. That is, the first equation of (14) is always valid, for every smooth test function φ of compact support. What we have shown with the above argument is that if μ is a solution of the McKean-Vlasov equation (8), then it is also a weak solution of the PDE (15) in a suitable sense. The converse can be shown as well, using recent results of Figalli [9] and Trevisan [28] on the correspondence between SDEs and Fokker-Planck equations. That is, solutions of the McKean-Vlasov equation (8) are in one-to-one correspondence with weak solutions of the PDE (14), under minimal assumptions on the drift b. For this reason the weak PDE (14) is often itself the main object of study, instead of the McKean-Vlasov SDE (8).

2.4. **An alternative derivation of the McKean-Vlasov limit.** In addition to characterizing solutions of the McKean-Vlasov equation (8) as discussed above, the weak PDE (14) can be used as the basis for studying the $n\to\infty$ limit of the n-particle system. This approach requires some machinery from weak convergence theory for stochastic processes, which we will not develop. Instead, the argument is merely sketched, with some warnings when the arguments become hand-wavy. Assume throughout this section that $b = b(x,m)$ is bounded and jointly continuous, using the topology of weak convergence for the measure variable.

An effective way to study the $\mathcal{P}(\mathbb{R}^d)$-valued stochastic process $(m^n_{\boldsymbol{X}^n_t})_{t\ge 0}$ is through its behavior against test functions. Fix a smooth function φ on $\mathbb{R}^d$ with compact support. To identify the behavior of

$$\langle m^n_{\boldsymbol{X}^n_t},\varphi\rangle = \frac{1}{n}\sum_{i=1}^n\varphi(X^{n,i}_t),$$

we first use Itô's formula to write, for each $i=1,\ldots,n$,

$$d\varphi(X^{n,i}_t) = \left(b(X^{n,i}_t, m^n_{\boldsymbol{X}^n_t})\nabla\varphi(X^{n,i}_t)+\frac{1}{2}\Delta\varphi(X^{n,i}_t)\right)dt + \nabla\varphi(X^{n,i}_t)dW^i_t.$$

It is convenient now to define the infinitesimal generator by setting

$$L_m\varphi(x) = b(x,m)\nabla\varphi(x)+\frac{1}{2}\Delta\varphi(x),$$

for each $m\in\mathcal{P}(\mathbb{R}^d)$ and $x\in\mathbb{R}^d$. We may then write

$$d\varphi(X^{n,i}_t) = L_{m^n_{\boldsymbol{X}^n_t}}\varphi(X^{n,i}_t)dt + \nabla\varphi(X^{n,i}_t)dW^i_t.$$

Average over $i = 1, \dots, n$ to get

$$\begin{aligned} d\langle m^n_{\boldsymbol{X}^n_t}, \varphi\rangle &= \frac{1}{n}\sum_{i=1}^n d\varphi(X^{n,i}_t) \\ &= \left\langle m^n_{\boldsymbol{X}^n_t}, L_{m^n_{\boldsymbol{X}^n_t}}\varphi\right\rangle dt + \frac{1}{n}\sum_{i=1}^n \nabla\varphi(X^{n,i}_t)dW^i_t \\ &=: \left\langle m^n_{\boldsymbol{X}^n_t}, L_{m^n_{\boldsymbol{X}^n_t}}\varphi\right\rangle dt + dM^n_t, \end{aligned}$$

where the last line defines the local martingale M^n. Since $\nabla\varphi$ is bounded, M^n is a martingale with quadratic variation

$$[M^n]_t = \frac{1}{n^2}\sum_{i=1}^n \int_0^t |\nabla\varphi(X^{n,i}_s)|^2 ds \le \frac{t}{n}\|\nabla\varphi\|^2_\infty.$$

In particular, this implies $\mathbb{E}[(M^n_t)^2] \le t\|\nabla\varphi\|^2_\infty/n$.

At this point we begin skipping some crucial steps. One should show next that $(m^n_{\boldsymbol{X}^n_t})_{t\ge0}$ is tight as a sequence of random elements of $C(\mathbb{R}_+;\mathcal{P}(\mathbb{R}^d))$, endowed with an appropriate topology. Then, by Prokhorov's theorem, it admits a subsequential limit point, and we would like to describe all such limit points. So suppose the process $(m^n_{\boldsymbol{X}^n_t})_{t\ge0}$ converges along a subsequence to another (measure-valued) process $(\nu_t)_{t\ge0}$. Then, since the martingale term M^n vanishes as $n\to\infty$, we should have

$$d\langle \nu_t, \varphi\rangle = \langle \nu_t, L_{\nu_t}\varphi\rangle\, dt. \tag{16}$$

The identity (16) holds a.s. for all smooth φ of compact support, and with some work one can switch the order of quantifiers to show that in fact, with probability 1, the equation holds (16) for all φ and all $t \ge 0$. (That is, the null set does not depend on t or φ.) This shows that the limit point ν is, with probability 1, a solution of the weak PDE (14).

If we knew, via different arguments, that the weak solution of the PDE (15) is unique, then we would be in good shape. Let us write μ for the unique solution. Every subsequential limit in distribution ν of $(m^n_{\boldsymbol{X}^n_t})_{t\ge0}$ is almost surely a weak solution, and thus $\nu = \mu$ a.s. This implies that the full sequence $(m^n_{\boldsymbol{X}^n_t})_{t\ge0}$ converges in distribution to the deterministic μ.

In fact, this line of argument reveals, interestingly, that even if the PDE (15) is not unique, we can still say that every subsequential limit point of $(m^n_{\boldsymbol{X}^n_t})_{t\ge0}$ is supported on the set of solutions of PDE (15).

3. Mean field and n-player games

This will be covered in more detail in other parts of the short course, so we merely summarize the basics of mean field games (MFGs) here. We start with n-player stochastic differential games of mean field type, and then we move on to MFGs.

We **assume throughout the rest of these notes** that we are given a finite time horizon $T > 0$, a compact metric space A called the *control space*, and bounded continuous drift and cost functions:

$$\begin{aligned} b &: \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \to \mathbb{R}^d \\ f &: \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \to \mathbb{R} \\ g &: \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}. \end{aligned}$$

The boundedness assumption is certainly too strong, but we impose it to avoid encountering integrability issues. This is an introductory course, after all.

3.1. **Stochastic differential games of mean field type.** Fix a time horizon $T > 0$. Fix i.i.d. initial states $(\xi^i)_{i=1}^{\infty}$. At each time $t \in [0,T]$, each player will be able to choose an action from the (compact metric) action space A. Suppose there are n players, labeled $i = 1,\dots,n$, each of which controls a d-dimensional state process X^i, governed by the dynamics

$$dX_t^i = b(X_t^i, m^n_{\boldsymbol{X}_t}, \alpha^i(t,\boldsymbol{X}_t))dt + dW_t^i, \qquad X_0^i = \xi^i. \tag{17}$$

Here $\boldsymbol{X} = (X^1,\dots,X^n)$ denotes the vector of state processes. In parallel with the structure of interacting diffusions of McKean-Vlasov type, studied in Section 2, we have a common drift function b. An *admissible (Markovian) control* is a measurable function from $[0,T] \times (\mathbb{R}^d)^n \to A$, and we write $\mathcal{A}$ for the set of admissible controls. The Cartesian product $\mathcal{A}^n$ then denotes the set of *strategy profiles*, meaning a choice of control for each player. It is a well known result of Veretennikov [29] (see also [16, Theorem 2.1]) that SDEs with bounded measurable drift and constant non-degenerate diffusion coefficient are well-posed in the strong sense, so for any choice of strategy profile the SDE (17) is well-posed.

The goal of player i will be to choose α^i to maximize

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

These optimization problems are interdependent, and we resolve them using the notion of *Nash equilibrium*, or more generally *ϵ-Nash equilibrium*. For $\epsilon \ge 0$, an ϵ-Nash equilibrium is a strategy profile $\boldsymbol{\alpha} = (\alpha^1,\dots,\alpha^n) \in \mathcal{A}^n$ such that

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

If $\epsilon = 0$, we say simply *Nash equilibrium* in place of *0-Nash equilibrium*.

Remark 3.1. We work here with Markovian (feedback) controls, and we will not discuss the common alternative choice of *open-loop* controls, which are often considered less realistic.

3.2. **A PDE formulation.** A standard method for finding Markovian equilibria is by studying an associated PDE system. A one-player stochastic optimal control problem can be studied through a Hamilton-Jacobi-Bellman (HJB) equation, a nonlinear PDE representing the value function (see [26] or [10] for background). Similarly, an n-player stochastic differential game can be studied through a system of n coupled HJB equations. Each of these n HJB equations represents the value function of a single player.

Suppose players adopt the strategy profile $(\alpha^1,\dots,\alpha^n)$. Let us fix a player $i \in \{1,\dots,n\}$ and study the best response problem for this player, i.e., the problem of optimizing over $J_i(\dots,\alpha^{i-1},\beta,\alpha^{i+1},\dots)$ over $\beta \in \mathcal{A}$. Treating $(\alpha^k)_{k\neq i}$ as fixed, player i faces a standard stochastic optimal control problem. Define the Hamiltonian for this control problem by setting, for each $(x,m,y) \in \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d$,

$$H(x,m,y) = \sup_{a\in A}\left(y\cdot b(x,m,a) + f(x,m,a)\right),$$

and let $\widehat{\alpha}(x,m,y)$ denote the maximizer, which we assume to be unique for simplicity. The HJB equation for the value function $v_i^n : [0,T] \times (\mathbb{R}^d)^n \to \mathbb{R}$ of this control problem is

$$\begin{aligned} 0 &= \partial_t v_i^n(t,\boldsymbol{x}) + \sum_{k\neq i} \nabla_k v_i^n(t,\boldsymbol{x}) \cdot b(x_k, m^n_{\boldsymbol{x}}, \alpha^k(t,\boldsymbol{x})) \\ &\quad + H(x_i, m^n_{\boldsymbol{x}}, \nabla_i v_i^n(t,\boldsymbol{x})) + \frac{1}{2}\Delta v_i^n(t,\boldsymbol{x}), \\ v_i^n(T,\boldsymbol{x}) &= g(x_i, m^n_{\boldsymbol{x}}). \end{aligned}$$

The gradient operator ∇_k acts on the k^{th} variable of $\boldsymbol{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$, and the Laplacian Δ acts on all of the dn variables. We highlight in red where the controls of the other players $k \neq i$ appear. If we can solve this PDE, then the find the best response of player i according to

$$\widetilde{\alpha}^i(t, \boldsymbol{x}) := \widehat{\alpha}(x_i, m^n_{\boldsymbol{x}}, \nabla_i v^n_i(t, \boldsymbol{x})). \tag{18}$$

Finally, to find a Nash equilibrium, we need each player to be solving this best response problem. So the PDE system we should solve is

$$\begin{aligned} 0 = \partial_t v^n_i(t, \boldsymbol{x}) + \sum_{k \neq i} \nabla_k v^n_i(t, \boldsymbol{x}) \cdot b(x_k, m^n_{\boldsymbol{x}}, \widehat{\alpha}(x_k, m^n_{\boldsymbol{x}}, \nabla_k v^n_k(t, \boldsymbol{x}))) \\ + H(x_i, m^n_{\boldsymbol{x}}, \nabla_i v^n_i(t, \boldsymbol{x})) + \frac{1}{2}\Delta v^n_i(t, \boldsymbol{x}), \end{aligned} \tag{19}$$

for $i = 1, \dots, n$, with boundary conditions $v^n_i(T, \boldsymbol{x}) = g(x_i, m^n_{\boldsymbol{x}})$. This PDE system is often called the **Nash system**. A standard verification argument yields the following: If $(v^n_i)^n_{i=1}$ is a classical solution of the Nash system (19), then $(\widetilde{\alpha}^i)^n_{i=1}$ given by (18) forms a Markovian Nash equilibrium.

3.3. **Deriving the master equation.** In this section we give a *heuristic* derivation of the so-called *master equation* from sending $n \to \infty$ in the Nash system (19). We emphasize the word *heuristic* here, as many arguments in this section are particularly non-rigorous.

Suppose the n-Nash system (19) is well-posed for each $n \in \mathbb{N}$, and let us first rewrite slightly Nash system (19) slightly as

$$\begin{aligned} 0 = \partial_t v^n_i(t, \boldsymbol{x}) + H(x_i, m^n_{\boldsymbol{x}}, \nabla_i v^n_i(t, \boldsymbol{x})) + \frac{1}{2}\sum_{k=1}^n \Delta_k v^n_i(t, \boldsymbol{x}) \\ + \sum_{k \neq i} \nabla_k v^n_i(t, \boldsymbol{x}) \cdot b(x_k, m^n_{\boldsymbol{x}}, \widehat{\alpha}(x_k, m^n_{\boldsymbol{x}}, \nabla_k v^n_k(t, \boldsymbol{x}))), \end{aligned}$$

for $i = 1, \dots, n$, with boundary conditions $v^n_i(T, \boldsymbol{x}) = g(x_i, m^n_{\boldsymbol{x}})$. We have simply split the Laplacian term into n constituents. Well-posedness implies that the solution is symmetric, in the sense that $v^n_i(t, \boldsymbol{x})$ is symmetric in the $n-1$ variables $(x_k)_{k \neq i}$. In particular, there exists a function $V_n : [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ such that $v^n_i(t, \boldsymbol{x}) = V_n(t, x_i, m^n_{\boldsymbol{x}})$. Note of course that this function V_n depends on the number of players. The limit as $n \to \infty$ should be the solution of the master equation, as we will now formally show by assuming that V_n is as smooth as we need, in particular has bounded derivatives of all orders.

To derive the master equation, we use the calculus rules on $\mathcal{P}(\mathbb{R}^d)$ summarized in Appendix B, particularly Proposition B.6, to express derivatives of v^n_i in terms of $\nabla_x V_n$ and $D_m V_n$. First, we clearly have $\partial_t v^n_i(t, \boldsymbol{x}) = \partial_t V_n(t, x_i, m^n_{\boldsymbol{x}})$. Second, use Proposition B.6 to get

$$\begin{aligned} \nabla_j v^n_i(t, \boldsymbol{x}) &= \nabla_j (V_n(t, x_i, m^n_{\boldsymbol{x}})) \\ &= \frac{1}{n} D_m V_n(t, x_i, m^n_{\boldsymbol{x}}, x_j) + \delta_{ij} \nabla_x V_n(t, x_i, m^n_{\boldsymbol{x}}) \\ &= O(1/n) + \delta_{ij} \nabla_x V_n(t, x_i, m^n_{\boldsymbol{x}}), \end{aligned}$$

where $\delta_{ij} = 1$ if $i = j$ and 0 if $i \neq j$. Lastly, writing $D_m V_n = D_m V_n(t, x, m, v)$, if $j \neq i$ we have

$$\Delta_j v^n_i(t, \boldsymbol{x}) = \frac{1}{n} \mathrm{div}_v D_m V_n(t, x_i, m^n_{\boldsymbol{x}}, x_j) + \frac{1}{n^2} \mathrm{Tr}[D^2_m V_n(t, x_i, m^n_{\boldsymbol{x}}, x_j, x_j)].$$

For $j = i$, this is instead

$$\begin{aligned}\Delta_i v_i^n(t,\boldsymbol{x}) =& \frac{1}{n}\mathrm{div}_v D_m V_n(t,x_i,m^n_{\boldsymbol{x}},x_i) + \frac{1}{n^2}\mathrm{Tr}[D_m^2 V_n(t,x_i,m^n_{\boldsymbol{x}},x_i,x_i)] \\ &+ \Delta_x V_n(t,x_i,m^n_{\boldsymbol{x}}) + \frac{1}{n}\mathrm{Tr}[D_m \nabla_x V_n(t,x_i,m^n_{\boldsymbol{x}},x_i)].\end{aligned}$$

Plug this into the Nash system to get an equation in terms of V_n, for each i:

$$\begin{aligned}0 &= \partial_t V_n(t,x_i,m^n_{\boldsymbol{x}}) + H\big(x_i, m^n_{\boldsymbol{x}}, O(1/n) + \nabla_x V_n(t,x_i,m^n_{\boldsymbol{x}})\big) \\ &\quad + \frac{1}{2n}\sum_{k=1}^n \mathrm{div}_v D_m V_n(t,x_i,m^n_{\boldsymbol{x}},x_k) + \Delta_x V_n(t,x_i,m^n_{\boldsymbol{x}}) + O(1/n) \\ &\quad + \frac{1}{n}\sum_{k\neq i} D_m V_n(t,x_i,m^n_{\boldsymbol{x}},x_k)\cdot b\Big(x_k, m^n_{\boldsymbol{x}}, \widehat{\alpha}\big(x_k, m^n_{\boldsymbol{x}}, O(1/n) + \nabla_x V_n(t,x_k,m^n_{\boldsymbol{x}})\big)\Big) \\ &= O(1/n) + \partial_t V_n(t,x_i,m^n_{\boldsymbol{x}}) + H\big(x_i, m^n_{\boldsymbol{x}}, \nabla_x V_n(t,x_i,m^n_{\boldsymbol{x}})\big) \\ &\quad + \frac{1}{2}\int_{\mathbb{R}^d} \mathrm{div}_v D_m V_n(t,x_i,m^n_{\boldsymbol{x}},v)\, m^n_{\boldsymbol{x}}(dv) + \Delta_x V_n(t,x_i,m^n_{\boldsymbol{x}}) \\ &\quad + \int_{\mathbb{R}^d} D_m V_n(t,x_i,m^n_{\boldsymbol{x}},v)\cdot b\Big(v, m^n_{\boldsymbol{x}}, \widehat{\alpha}\big(v, m^n_{\boldsymbol{x}}, \nabla_x V_n(t,v,m^n_{\boldsymbol{x}})\big)\Big)\, m^n_{\boldsymbol{x}}(dv).\end{aligned}$$

Indeed, we used the fact that the term $(1/n^2)\mathrm{Tr}[D_m^2 V_n]$ in the expression for $\Delta_j v_i^n$ is $O(1/n^2)$ and thus becomes $O(1/n)$ after averaging over $j = 1,\ldots,n$. Similarly, the term $(1/n)\mathrm{Tr}[D_m\nabla_x V_n]$ in the expression for $\Delta_i v_i^n$ is $O(1/n)$. We deduce that, if $V_n \to U$ in some sense, then U should satisfy

$$\begin{aligned}0 &= \partial_t U(t,x,m) + H\big(x,m,\nabla_x U(t,x,m)\big) \\ &\quad + \frac{1}{2}\int_{\mathbb{R}^d} \mathrm{div}_v D_m U(t,x,m,v)\, m(dv) + \Delta_x U(t,x,m) \\ &\quad + \int_{\mathbb{R}^d} D_m U(t,x,m,v)\cdot b\Big(v,m,\widehat{\alpha}\big(v,m,\nabla_x U(t,v,m)\big)\Big)\, m(dv),\end{aligned} \tag{20}$$

with terminal condition $U(T,x,m) = g(x,m)$. This is the master equation.

The master equation allows one to identify the equilibrium control as $\widehat{\alpha}(x,m,\nabla_x U(t,x,m))$. That is, by solving the McKean-Vlasov SDE

$$dX_t = b(X_t, \mathcal{L}(X_t), \widehat{\alpha}(X_t, \mathcal{L}(X_t), \nabla_x U(t,x,\mathcal{L}(X_t)))dt + dW_t,$$

one can show that the measure flow $(\mathcal{L}(X_t))_{t\in[0,T]}$ is a mean field equilibrium, in the sense described in the next section. We refer to the lecture of F. Delarue for more details on the analysis of the master equation and its precise connection to mean field equilibria.

The above derivation of the master equation is made rigorous in the seminal paper of Cardaliaguet-Delarue-Lasry-Lions [5]. In fact, the best way to make it rigorous is by reversing the argument. Rather than showing that $v_i^n(t,\boldsymbol{x}) = V_n(t,x_i,m^n_{\boldsymbol{x}})$ nearly solves the master equation, the argument in [5] is to show that the functions $(u_i^n)_{i=1}^n$ defined by $u_i^n(t,\boldsymbol{x}) := U(t,x_i,m^n_{\boldsymbol{x}})$ nearly solves the n-player Nash system (19). The argument is very similar to that presented above, using the calculus on $\mathcal{P}(\mathbb{R}^d)$ via Proposition B.6. Before implementing this approach, of course, an important and challenging first step is to establish the existence and uniqueness of a sufficiently smooth classical solution U of the master equation.

3.4. **The MFG fixed point problem.** Here we give first a heuristic derivation and second a precise formulation of the mean field game (MFG) fixed point problem.

Working heuristically for now, let us start from the n-player game described in Section 3.1, and let us avoid placing any specific assumptions on b. Suppose the initial states $X_0^{n,i}$ are i.i.d. with some given initial law λ. Suppose $(\alpha^{n,1},\dots,\alpha^{n,n})$ is a Nash equilibrium for each n. Given the symmetry of the game, it is not unreasonable to suspect that for each n we can find a single function $\widehat{\alpha}_n$ such that $\alpha^{n,i}(t,\boldsymbol{x}) = \widehat{\alpha}(t,x_i,m^n_{\boldsymbol{x}})$. It is far less reasonable to expect that $\widehat{\alpha}_n = \widehat{\alpha}$ does not depend on n, but let us make this assumption anyway, as we are doing heuristics. Then, the state processes become

$$dX_t^{n,i} = b(X_t^{n,i}, m^n_{\boldsymbol{X}_t^n}, \widehat{\alpha}(t, X_t^{n,i}, m^n_{\boldsymbol{X}_t^n}))dt + dW_t^i.$$

where as usual $m^n_{\boldsymbol{X}_t^n}$ is the empirical measure of the vector of n state processes $\boldsymbol{X}_t^n = (X_t^{n,1},\dots,X_t^{n,n})$. This is precisely an equation of McKean-Vlasov type. Recalling Theorem 2.3, it should hold (if the coefficients and, notably, $\widehat{\alpha}$ are nice enough) that $(X^{n,1}, m^n_{\boldsymbol{X}^n})$ converges in law to (X,μ), where $\mu = (\mu_t)_{t\in[0,T]}$ is a deterministic measure flow and X^* solves

$$dX_t^* = b(X_t^*, \mu_t, \widehat{\alpha}(t, X_t^*, \mu_t))dt + dW_t,$$

with $\mu_t = \mathcal{L}(X_t^*)$ for all t. On the other hand, because $(\alpha^{n,1},\dots,\alpha^{n,n})$ is a Nash equilibrium, for any other control β we have (suppressing the dependence of the controls on the state processes, for convenience)

$$\begin{aligned}\mathbb{E}&\left[\int_0^T f(X_t^{n,1}, m^n_{\boldsymbol{X}_t^n}, \alpha_t^{n,1})dt + g(X_T^1, \mu_T^n)\right] \\ &\geq \mathbb{E}\left[\int_0^T f(\widetilde{X}_t^{n,1}, m^n_{\widetilde{\boldsymbol{X}}_t^n}, \beta_t)dt + g(\widetilde{X}_T^{n,1}, m^n_{\widetilde{\boldsymbol{X}}_T^n})\right]. \end{aligned} \tag{21}$$

Here we mean that players 1's state process $\widetilde{X}_t^{n,1}$ on the right-hand side is controlled by β, not by $\alpha^{n,1}$, but the other $n-1$ players $i \neq 1$ stick with the same controls $\alpha^{n,i}$. When only player 1 changes strategy and the remaining $n-1$ players stick with the same $\alpha^{n,i}$, the empirical measure should not change too much and should thus converge to the same limit μ. Then, the process $\widetilde{X}^{n,1}$ controlled by β should converge to the solution X^β of the SDE

$$dX_t^\beta = b(X_t^\beta, \mu_t, \beta_t)dt + dW_t,$$

with the same measure flow μ *as before.* Sending $n\to\infty$ on both sides of (21), we find

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

This should hold for every choice of β, and we conclude that the control $\widehat{\alpha}$ is optimal for the control problem

$$\begin{array}{ll}\sup_\alpha & \mathbb{E}\left[\int_0^T f(X_t,\mu_t,\alpha_t)dt + g(X_T,\mu_T)\right] \\ dX_t & = b(X_t,\mu_t,\alpha_t)dt + dW_t.\end{array}$$

Moreover, μ is precisely the law of the optimally controlled state process X^* above!

In summary the above heuristic leads us to the following fixed point problem:

Definition 3.2. Define a map $\Phi : C([0,T],\mathcal{P}(\mathbb{R}^d)) \to C([0,T],\mathcal{P}(\mathbb{R}^d))$ as follows:

(1) Fix a (deterministic) measure flow $\mu = (\mu_t)_{t\in[0,T]} \in C([0,T],\mathcal{P}(\mathbb{R}^d))$, to represent a continuum of agents' state processes.

(2) Solve the control problem faced by a *typical* or *representative agent*:

$$(P_\mu) \qquad \begin{cases} \sup_\alpha & \mathbb{E}\left[\int_0^T f(X_t^{\mu,\alpha}, \mu_t, \alpha(t, X_t))dt + g(X_T^{\mu,\alpha}, \mu_T)\right] \\ dX_t^{\mu,\alpha} & = b(X_t^{\mu,\alpha}, \mu_t, \alpha(t, X_t))dt + dW_t. \end{cases}$$

Here the law of the initial state $X_0^{\mu,\alpha}$ equals the given initial law λ, and the optimization is over the set of all measurable functions $\alpha : [0,T] \times \mathbb{R}^d \to A$.

(3) Let α^* be the optimally controlled state process, which we assume is unique, and define $\Phi(\mu) = (\mathcal{L}(X_t^{\mu,\alpha^*}))_{t\in[0,T]}$.

We say that $\mu \in C([0,T], \mathcal{P}(\mathbb{R}^d))$ is a *mean field equilibrium* (MFE) if it is a fixed point of Φ, or $\mu = \Phi(\mu)$.

Remark 3.3. Definition 3.2 is central to the rest of the course, so we take the time to clarify it with some comments:

- While our definition, strictly speaking, refers only to the measure flow μ, it is often useful to include the optimal control. That is, we may refer more descriptively to the pair (μ, α^*) as an MFE, where α^* is optimal for the control problem (P_μ).
- In general, the optimal control α^* for (P_μ) need not be unique. We could then could define
$$\Phi(\mu) = \left\{(\mathcal{L}(X_t^*))_{t\in[0,T]} : X^* \text{ is optimal for } (P_\mu)\right\},$$
and then try to find a fixed point for this set-valued map, meaning $\mu \in \Phi(\mu)$. In this case, it becomes even more appropriate to include the control α^* in the definition of an MFE, as in the previous bullet point.
- It cannot be stressed enough that the measure flow $\mu = (\mu_t)_{t\in[0,T]}$ must be seen as *fixed* when solving the control problem (P_μ). The control problem (P_μ) is a completely standard stochastic optimal control problem, as the measure flow should be seen just as a time-dependence in the coefficient.
- The fixed point can instead be formulated on the control process itself. That is, suppose we start with a control α, and then solve the McKean-Vlasov equation
$$dY_t = b(Y_t, \mu_t, \alpha(t, X_t))dt + dW_t, \quad \mu_t = \mathcal{L}(Y_t), \ t \in [0,T].$$
Note that $Y = X^{\mu,\alpha}$. Using this measure flow, we then find an optimal control α^* for the control problem (P_μ) relative to this measure flow μ. If α^* agrees with the α we started with, then μ is a mean field equilibrium.
- We assume the control α is Markovian, but a common relaxation is to work with the family of open-loop controls, which are processes adapted to the minimal filtration $(\sigma(X_0, W_s : s \le t))_{t\in[0,T]}$. This relaxation typically does not make much of a difference; see [17] for general results on this.

We refer to the other lectures in this short course for a more detailed discussion of the mean field game problem, examples, extensions, etc. We limit our discussion in the rest of these notes to the question of rigorously justifying the convergence of the n-player game to the MFG, for which we have now seen two different heuristic arguments.

3.5. **Constructing approximate Nash equilibria.** Most of the rest of these notes are focused on the question of proving that a given sequence of Nash equilibria $(\alpha^{n,1}, \ldots, \alpha^{n,n})$ "converges" to the MFG, typically in the sense that the associated empirical measure flows $(m^n_{\boldsymbol{X}_t^n})_{t\in[0,T]}$ converge to a mean field equilibrium. A much easier problem, still of great practical value, is a sort of converse: Given a mean field equilibrium $\mu = (\mu_t)_{t\in[0,T]}$, one can construct a sequence ϵ_n converging to zero along with ϵ_n-Nash equilibria $(\alpha^{n,1}, \ldots, \alpha^{n,n})$ such that the associated empirical measure flows converge to μ. In fact, if the MFE μ is associated with the control $\alpha^* = \alpha^*(t,x)$, then the approximate equilibria can be constructed explicitly by $\alpha^{n,i}(t,\boldsymbol{x}) := \alpha^*(t,x_i)$. These controls have the highly desirable property of being *distributed*, in the sense that each player's control is based solely on her own state process, not the other players, not even through the empirical measure.

The proof (under reasonable technical assumptions on the coefficients) that this choice of $\alpha^{n,i}$ does indeed provide an approximate Nash equilibrium is similar to the heuristic argument given in Section 3.4, but we will not go into details. This idea is implemented in many papers on MFGs, including but not limited to [13, 20, 6, 8, 18, 19], and in textbook form in [7, Volume II, Chapter 6].

4. Convergence to the MFG limit via compactness methods

A critical challenge in formalizing the heuristic derivation of the MFG in 3.4 is that it is very hard to get much of a handle on the equilibrium controls $(\alpha^{n,1},\ldots,\alpha^{n,n})$. For this reason, it is quite useful to find a way to embed these controls in a compact space, so that we can extract a limit point.

To be clear, our controls for the n-player game are measurable functions $\alpha^{n,i} : [0,T]\times(\mathbb{R}^d)^n \to A$, where we assume the control space A is compact. But just knowing that $(\alpha^{n,n},\ldots,\alpha^{n,n})$ is a Nash equilibrium does not give us much in the way of regularity or estimates on the functions $\alpha^{n,i}$. If we knew the Nash system (19) had a classical solution $(v_i^n)_{i=1}^n$, and our Nash equilibrium is the one determined by this classical solution as in (18), then we may be able to get an estimate on the $\alpha^{n,i}$ if we had some estimate on the gradients $\nabla_i v_i^n$ which were *uniform in* n. But even if this worked out, there is still the more fundamental problem of having functions $\alpha^{n,i} = \alpha^{n,i}(t,x_1,\ldots,x_n)$ of $n+1$ variables; what does it mean for a sequence of increasingly many variables to converge?

An alternative idea is to view the controls as *processes* rather than functions. That is, suppose $\alpha^{n,i}$ are given controls and $\boldsymbol{X}^n = (X^{n,1},\ldots,X^{n,n})$ solves the corresponding controlled SDE,

$$dX_t^{n,i} = b(X_t^{n,i}, m^n_{\boldsymbol{X}_t^n}, \alpha^{n,i}(t,\boldsymbol{X}_t^n))dt + dW_t^i. \tag{22}$$

We can then define a process $\tilde{\alpha}_t^{n,i}(\omega) = \alpha^{n,i}(t,\boldsymbol{X}_t^n(\omega))$, which is adapted to the filtration generated by the Brownian motions, and hope that these processes converge in some sense. The space of adapted processes with values in A (modulo $dt\times d\mathbb{P}$ equivalence) is separable and completely metrizeable when topologized by convergence in measure, but there is no hope of showing pre-compactness of $\{\tilde{\alpha}^{n,i}\}$ in this space. Convergence in distribution is more hopeful, if we view $\tilde{\alpha}^{n,i}$ as a random element of the space $L^0([0,T];A)$ of measurable functions (modulo a.e. equality) from $[0,T]$ to A. Equipped with the topology of convergence in Lebesgue measure, this is again too difficult for compactness purposes. But there is fortunately a very convenient way to compactify this space.

4.1. **Relaxed controls.** We embed our controls into a space of measures, for which compactness criteria are straightforward. Precisely, let $\mathcal{V}$ denote the set of positive Borel measures on $[0,T]\times A$ with first marginal equal to Lebesgue measure. That is, $q\in\mathcal{V}$ satisfies $q([s,t]\times A) = t-s$ for $0\le s<t\le T$. Note that every $q\in\mathcal{V}$ thus has total mass T. We may endow $\mathcal{V}$ with the topology of weak convergence, which means that $q_n\to q$ if and only if $\int\varphi\,dq_n \to \int\varphi\,dq$ for every bounded continuous function $\varphi:[0,T]\times A\to\mathbb{R}$. Two incredibly useful lemmas are the following:

Lemma 4.1. *Suppose $q, q_n\in\mathcal{V}$ with $q_n\to q$. Then, for any bounded measurable function $\varphi:[0,T]\times A\to\mathbb{R}$ such that $\varphi(t,\cdot)$ is continuous on A for each $t\in[0,T]$, we have $\int\varphi\,dq_n\to\int\varphi\,dq$.*

Exercise 4.2. Prove Lemma 4.1. **Hint:** Show that the map $\Phi:[0,T]\to C(A)$ defined by $\Phi(t)=\varphi(t,\cdot)$ is Borel measurable. Then, by Lusin's theorem, there exists for each $\delta>0$

a *continuous* function $\Phi_\delta : [0,T] \to C(A)$ such that $\Phi_\delta = \Phi$ except on a set of Lebesgue measure less than δ.

Lemma 4.3. *The space $\mathcal{V}$ is compact and metrizeable.*

Proof. Assume $T = 1$ for simplicity, as the general case amounts to no more than a rescaling. Because $[0,1] \times A$ is a compact metric space, the space $\mathcal{P}([0,1] \times A)$ is compact and metrizeable when equipped with the topology of weak convergence (see Theorem A.3). We also note that $\mathcal{V}$ is a closed subset of $\mathcal{P}([0,1] \times A)$: A measure $\mu \in \mathcal{P}([0,1] \times A)$ belongs to $\mathcal{V}$ if and only if its first marginal is the uniform (Lebesgue) measure. The map $[0,1] \times A \ni (t,a) \mapsto t \in [0,1]$ is continuous, and so if $\mu_n \to \mu$ in $\mathcal{P}([0,1] \times A)$ then the first marginal must converge, thanks to the continuous mapping Theorem A.2.

Alternatively, instead of referring to Theorem A.3, we can note that the Banach space $C([0,T] \times A)$ of continuous functions on $[0,T] \times A$ is separable. Hence, by Banach-Alaoglu, closed bounded sets in the dual $C([0,T] \times A)^*$ are weak*-compact and metrizeable. Clearly $\mathcal{V}$ is bounded, and closedness is argued as above. $\square$

By the disintegration theorem, note that every $q \in \mathcal{V}$ can be written as

$$q(dt, da) = dt q_t(da),$$

for some measurable family of probability measures $(q_t)_{t\in[0,T]}$ on A, uniquely defined up to almost-everywhere equality. Any (open-loop) control $\alpha \in L^0([0,T]; A)$,[3] now called a *strict control*, can be embedded in $\mathcal{V}$ by identifying it with the measure

$$q(dt, da) = dt\delta_{\alpha_t}(da).$$

We may write $\mathcal{V}_0$ for the set of $q \in \mathcal{V}$ of the above form. Elements of $\mathcal{V}_0$ are called *strict controls*, while general elements of $\mathcal{V}$ are called *relaxed controls*. The following remarkable fact is quite important, but we leave it as a (tricky) exercise, given that we will not use it in this course:

Exercise 4.4. Show that $\mathcal{V}_0$ is dense in $\mathcal{V}$. This is sometimes known in the literature as the *chattering lemma*. **Hint:** You may use the following implication of the Borel Isomorphism Theorem: Suppose $\mathcal{X}$ and $\mathcal{Y}$ are complete, separable metric spaces. Suppose μ and ν are nonnegative measures on $\mathcal{X}$ and $\mathcal{Y}$, respectively, with the same total mass, i.e., $\mu(\mathcal{X}) = \nu(\mathcal{Y})$. If μ is nonatomic, then there exists a Borel-measurable function $\varphi : \mathcal{X} \to \mathcal{Y}$ such that $\mu \circ \varphi^{-1} = \nu$.

Now, a stochastic control problem is often originally formulated as the problem of choosing a control process $(\alpha_t)_{t\in[0,T]}$ and then solving the controlled SDE

$$dX_t = b(X_t, \alpha_t)dt + dW_t.$$

On the other hand, a random element $\Lambda = \Lambda(dt, da) = dt\Lambda_t(da)$ of $\mathcal{V}$ can be interpreted as a measure-valued control, where for each (t,ω) the controller chooses a probability measure over the control space instead choosing a single action. The controlled SDE then becomes

$$dX_t = \left(\int_A b(X_t, a)\Lambda_t(da) \right) dt + dW_t.$$

[3]We write $L^0([0,T]; A)$ for the set of (equivalence classes of a.e. equal) measurable functions from $[0,T]$ to A.

4.2. **Relaxed mean field game.** We next reformulate the MFG problem discussed in Section 3.4, using relaxed controls along with the new name of *strong equilibrium.* We work again with the standing assumption that (b, f, g) are bounded and continuous, and the initial states from the n-player game defined in Section 3.1 are i.i.d. with some given law $\lambda \in \mathcal{P}(\mathbb{R}^d)$ with finite moments of every order.

Definition 4.5. We say that $(m_t)_{t\in[0,T]} \in C([0,T];\mathcal{P}(\mathbb{R}^d))$ is a *strong mean field equilibrium (MFE)* if there exists measurable function $\Lambda^* : [0,T] \times \mathbb{R}^d \to \mathcal{P}(A)$ such that the unique in law solution of the SDE

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

satisfies the following:

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

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

where X is the unique in law solution of

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

We use the term *strong equilibrium* because it turns out there is a more general notion, called a *weak equilibrium* which is needed in order to adequately describe the $n \to \infty$ limit theory. To define a weak equilibrium, we need a bit more notation. For a metric space E, we will by default equip $C([0,T];E)$ with the canonical filtration, i.e., the filtration generated by the coordinate process. The notation $\mathrm{Prog}(E)$ with refer to the progressive σ-field on $[0,T] \times C([0,T];E)$, which is generated by all progressively measurable real-valued processes $[0,T] \times C([0,T];E) \to \mathbb{R}$. In the following, let $\mathcal{B}(E)$ denote the Borel σ-field of a metric space E.

Definition 4.6. A *weak mean field equilibrium (MFE)* 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 an $\mathbb{F}$-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) $\Lambda^* : [0,T] \times \mathbb{R}^d \times C([0,T];\mathcal{P}(\mathbb{R}^d)) \to \mathcal{P}(A)$ is $\mathcal{B}(\mathbb{R}^d) \otimes \mathrm{Prog}(\mathcal{P}(\mathbb{R}^d))$-measurable.
(3) X_0^*, μ, and W are independent.
(4) The state equation holds:

$$dX_t^* = \int_A b(X_t^*, \mu_t, a)\Lambda^*(X_t^*, \mu)(da)dt + dW_t. \tag{23}$$

(5) For every alternative $\mathcal{B}(\mathbb{R}^d) \otimes \mathrm{Prog}(\mathcal{P}(\mathbb{R}^d))$-measurable function $\Lambda : [0,T] \times \mathbb{R}^d \times C([0,T];\mathcal{P}(\mathbb{R}^d)) \to \mathcal{P}(A)$, we have

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

where X is the solution (on an extension of the probability space) of

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

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

The law of the measure-valued process μ is a probability measure on $C([0,T];\mathcal{P}(\mathbb{R}^d))$, and with some abuse of notation we may refer to this probability measure itself as a weak MFE.

The key difference between a weak and strong equilibrium, of course, is that the measure flow $(\mu_t)_{t\in[0,T]}$ is stochastic in a weak equilibrium. Accordingly, the (relaxed) control Λ in (2) and (5) in Definition 4.6 should not be allowed to anticipate the measure flow, hence the progressive measurability assumption (2). Notice that a weak MFE $\mu = (\mu_t)_{t\in[0,T]}$ is a strong MFE if and only if it is determinsitic. The following limit theorem from [19]:

Theorem 4.7. *Under some additional technical assumptions: Let $\epsilon_n \ge 0$ with $\epsilon_n \to 0$. Suppose that for each n we are given an ϵ_n-Nash equilibrium $(\alpha^1, \dots, \alpha^n) \in \mathcal{A}_n^n$ for the n-player game, and let $(m^n_{\boldsymbol{X}_t^n})_{t\in[0,T]}$ denote the corresponding empirical measure flow from (17). Then the sequence $\{\mathcal{L}(m^n_{\boldsymbol{X}^n}) : n \in \mathbb{N}\} \subset \mathcal{P}(C([0,T];\mathcal{P}(\mathbb{R}^d)))$ is tight, and every (weak) limit point is a weak MFE in the sense of Definition 4.6.*

The nicest case, of course, is when there is a unique weak MFE, in which case Theorem 4.7 implies that any sequence of Nash equilibria for the n-player games converges to this unique equilibrium.

4.3. **Weak mean field equilibria.** Before discussing ideas of the proof, it is worth dwelling on the notion of weak MFE, introduced in Definition 4.6. There is an interesting phenomenon here: Suppose $S \subset C([0,T];\mathcal{P}(\mathbb{R}^d))$ denotes the set of strong MFE, as defined in Definition 4.5. Then *there can exist weak MFE $\mu = (\mu_t)_{t\in[0,T]}$ which are not supported on the set S*. This is in stark contrast with two simpler situations:

(1) Static or one-shot mean field games, in which there is a game but no dynamics.
(2) McKean-Vlasov particle systems, in which there are dynamics but no game (i.e., dynamics are uncontrolled).

In either of these cases, we can define similar notions of strong and weak equilibrium. And, in either case, it turns out that *weak MFE are always supported on* (or mixtures/randomizations among) *the set of strong MFE*. It is only in *dynamic games*, where optimization and dynamics interact, that this new phenomenon emerges, in which weak MFE can be completely different from strong MFE. See [18, 19] for an explicit example of a weak MFE which is not supported on the set of strong MFE.

4.4. **Limit theorem proof ideas.** We will sketch some of the ideas of the proof of Theorem 4.7, which we divide into three key steps. First we prove the claimed tightness. Second, we argue that for any limit point $M \in \mathcal{P}(C([0,T];\mathcal{P}(\mathbb{R}^d)))$, we may construct a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X^*, \mu)$ such that $\mathbb{P} \circ \mu^{-1} = M$ and such that properties (1-4) and (6) of Definition 4.6 hold. The final step is to check the optimality condition (5), which is the hardest part. We refer to [19] for the complete proof, with full details.

4.4.1. *Tightness.* Proving the tightness claimed in Theorem 4.7 is relatively straightforward in the sense that it simply applies some fairly well known criteria. In particular, one may use the following nice result of Sznitman [27, Proposition 2.2(ii)]:

Proposition 4.8. *Suppose $\mathcal{X}$ is a complete separable metric space, and let $S \subset \mathcal{P}(\mathcal{P}(\mathcal{X}))$. For $M \in \mathcal{P}(\mathcal{P}(\mathcal{X}))$, define the* mean measure $\overline{M} \in \mathcal{P}(\mathcal{X})$ *by*

$$\overline{M}(B) = \int_{\mathcal{P}(\mathcal{X})} m(B)\, M(dm), \quad \text{for} \quad B \subset \mathcal{X} \quad \text{Borel}.$$

Then S is tight if and only if $\{\overline{M} : M \in S\} \subset \mathcal{P}(\mathcal{X})$ is tight.

With this in hand, consider the "lifted" empirical measure

$$m^n_{\boldsymbol{X}^n} = \frac{1}{n}\sum_{k=1}^n \delta_{X^{n,k}},$$

as a random element of $\mathcal{P}(\mathcal{C}^d)$. Let $M_n = \mathcal{L}(m^n_{\boldsymbol{X}^n}) \in \mathcal{P}(\mathcal{P}(\mathcal{C}^d))$. Then the mean measure is given by

$$\overline{M}_n(B) = \mathbb{E}[m^n_{\boldsymbol{X}^n}(B)] = \mathbb{E}\left[\frac{1}{n}\sum_{k=1}^n 1_{\{X^{n,k}\in B\}}\right] = \frac{1}{n}\sum_{k=1}^n \mathbb{P}(X^{n,k} \in B).$$

One may then show that $\{\overline{M}_n : n \in \mathbb{N}\} \subset \mathcal{P}(\mathcal{C}^d)$ is tight by using standard tightness criteria for the continuous processes $X^{n,k}$. Conclude from Proposition 4.8 that $\{M_n = \mathcal{L}(m^n_{\boldsymbol{X}^n}) : n \in \mathbb{N}\} \subset \mathcal{P}(\mathcal{P}(\mathcal{C}^d))$ is tight. Finally, it is straightforward to check that the map

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

is continuous, and we conclude from the continuous mapping Theorem A.2 that $\{\mathcal{L}(m^n_{\boldsymbol{X}^n}) : n \in \mathbb{N}\} \subset \mathcal{P}(C([0,T];\mathcal{P}(\mathbb{R}^d)))$ is tight.

4.4.2. *Identifying the limiting dynamics.* The next step is to prove that, for any limit point $M \in \mathcal{P}(C([0,T];\mathcal{P}(\mathbb{R}^d)))$ of $\{\mathcal{L}(m^n_{\boldsymbol{X}^n}) : n \in \mathbb{N}\}$, we may construct a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \Lambda^*, X^*, \mu)$ such that $\mathbb{P}\circ\mu^{-1} = M$ and such that properties (1-4) and (6) of Definition 4.6 hold. This argument is an embellishment of a well-established argument for deriving the McKean-Vlasov limit for interacting diffusions, developed for instance in [24, 11, 23].

First, we view each control $\alpha^{n,i}$ as inducing a random relaxed control $\Lambda^{n,i}(da, da) = dt\Lambda^{n,i}_t(da) = dt\delta_{\alpha^{n,i}(t,\boldsymbol{X}^n)}(da)$. We then focus on the extended empirical measure

$$m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n} = \frac{1}{n}\sum_{k=1}^n \delta_{(X^{n,k},\Lambda^{n,k})},$$

which is a random element of $\mathcal{P}(\mathcal{C}^d \times \mathcal{V})$. Since $\mathcal{V}$ is compact we may easily deduce from the previous section that $\{\mathcal{L}(m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n}) : n \in \mathbb{N}\} \subset \mathcal{P}(\mathcal{P}(\mathcal{C}^d \times \mathcal{V}))$ is tight (e.g., using Lemma A.5). By applying Itô's formula to $\varphi(X^{n,k}_t)$ and then averaging over $k = 1, \ldots, n$, we find

$$\begin{aligned}
d\int_{\mathcal{C}^d\times\mathcal{V}} \varphi(x_t)\, m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n}(dx, dq) &= \frac{1}{n}\sum_{k=1}^n d\varphi(X^{n,k}_t) \\
&= \frac{1}{n}\sum_{k=1}^n \left(b(, X^{n,k}_t, m^n_{\boldsymbol{X}^n_t}, \alpha^{n,k}(t, \boldsymbol{X}^n_t)) \cdot \nabla\varphi(X^{n,k}_t) + \frac{1}{2}\Delta\varphi(X^{n,k}_t)\right) dt \\
&\quad + \frac{1}{n}\sum_{k=1}^n \nabla\varphi(X^{n,k}_t)\cdot dW^k_t \\
&= \frac{1}{n}\sum_{k=1}^n \int_A \left(b(X^{n,k}_t, m^n_{\boldsymbol{X}^n_t}, a)\cdot\nabla\varphi(X^{n,k}_t) + \frac{1}{2}\Delta\varphi(X^{n,k}_t)\right) \Lambda^{n,k}_t(da)dt \\
&\quad + dM^{n,\varphi}_t,
\end{aligned}$$

where we define the martingale $M^{n,\varphi}$ by setting $M_0^{n,\varphi}=0$ and $dM_t^{n,\varphi}=\frac{1}{n}\sum_{k=1}^n \nabla\varphi(X_t^{n,k})\cdot dW_t^k$. It then holds almost surely that

$$\begin{aligned} M_t^{n,\varphi} = \int_{\mathcal{C}^d\times\mathcal{V}} \Bigg[& \varphi(x_t)-\varphi(x_0)) \\ & - \int_{[0,T]\times A}\left(b(x_s, m^n_{\boldsymbol{X}^n_s}, a)\cdot\nabla\varphi(x_s)+\frac{1}{2}\Delta\varphi(x_s)\right)q(ds,da)\Bigg] m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n}(dx,dq). \end{aligned}$$

Doob's inequality, Itô's isometry, and the independence of the Brownian motions $W^1,\ldots,W^n$ yield

$$\begin{aligned} \mathbb{E}\left[\sup_{t\in[0,T]}|M_t^{n,\varphi}|^2\right] \le 4\mathbb{E}[|M_T^{n,\varphi}|^2] &= \frac{4}{n^2}\sum_{k=1}^n \mathbb{E}\int_0^T |\nabla\varphi(X_t^{n,k})|^2dt \\ &\le \frac{4T\|\nabla\varphi\|_\infty}{n}. \end{aligned}$$

Hence, $M^{n,\varphi}\to 0$, and we pass to the limit above (skipping the details of the weak convergence arguments) to find

$$\begin{aligned} 0 = \int_{\mathcal{C}^d\times\mathcal{V}} \Bigg[& \varphi(x_t)-\varphi(x_0)) \\ & - \int_{[0,T]\times A}\left(b(x_s, \widetilde{\mu}^x_s, a)\cdot\nabla\varphi(x_s)+\frac{1}{2}\Delta\varphi(x_s)\right)q(ds,da)\Bigg] \widetilde{\mu}(dx,dq), \end{aligned}$$

where the $\mathcal{P}(\mathcal{C}^d\times\mathcal{V})$-valued random variable $\widetilde{\mu}$ is any weak limit of $m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n}$, and $\widetilde{\mu}^x$ denotes the $\mathcal{C}^d$-marginal of $\widetilde{\mu}$. This holds for each $\varphi\in C_c^\infty(\mathbb{R}^d)$ and each $t>s\ge 0$.

In some sense, the above integral equation is a randomized version of a Fokker-Planck (or Kolmogorov forward) equation, but with an additional dependence on a non-Markovian control. By carefully applying a *Markovian projection* argument, it can be argued that there exists a tuple $(\Omega,\mathcal{F},\mathbb{F},\mathbb{P},W,\Lambda^*,X^*,\mu)$ such that $\mathbb{P}\circ\mu^{-1}=\mathcal{L}((\widetilde{\mu}^x_t)_{t\in[0,T]})$ and properties (1-4) and (6) hold. This is by no means straightforward and requires some nontrivial results connecting Fokker-Planck equations to solution of martingale problems, but we omit the remaining details of this step.

4.4.3. *Proving optimality at the limit.* Suppose now that we have proven the claimed tightness of Theorem 4.7 as well as the fact that for any limit point $M\in\mathcal{P}(C([0,T];\mathcal{P}(\mathbb{R}^d)))$ we may construct a tuple $(\Omega,\mathcal{F},\mathbb{F},\mathbb{P},W,\Lambda^*,X^*,\mu)$ such that $\mathbb{P}\circ\mu^{-1}=M$ and such that properties (1-4) and (6) of Definition 4.6 hold. The final and most difficult step is to check that this tuple satisfies the optimality condition (5) of Definition 4.6.

The principle here[4] is to fix an arbitrary competitor (strict) control $\widehat{\alpha}:[0,T]\times\mathbb{R}^d\times C([0,T];\mathcal{P}(\mathbb{R}^d))\to A$ and "give it" to each of the players in the n-player game. Precisely, for each n and each $k=1,\ldots,n$, define the n state processes $\boldsymbol{Y}^{n,k}=(Y^{n,k,1},\ldots,Y^{n,k,n})$ by

$$\begin{aligned} dY_t^{n,k,k} &= b(Y_t^{n,k,k}, m^n_{\boldsymbol{Y}_t^{n,k}}, \widehat{\alpha}(t,Y_t^{n,k,k}, m^n_{\boldsymbol{Y}^{n,k}}))dt + dW_t^k, \\ dY_t^{n,k,i} &= b(Y_t^{n,k,i}, m^n_{\boldsymbol{Y}_t^{n,k}}, \alpha^{n,i}(t,\boldsymbol{Y}^{n,k}))dt + dW_t^i, \quad i\ne k, \end{aligned}$$

[4]This is reminiscent of Gamma-convergence arguments.

where as usual the initial states $Y_0^{n,k,1}, \ldots, Y_0^{n,k,n}$ are i.i.d. with law λ. The state process $\boldsymbol{Y}^{n,k}$ would be exactly the same as the state Nash equilibrium process $\boldsymbol{X}^n$ governed by $(\alpha^{n,1}, \ldots, \alpha^{n,n})$, except that we switch player k's control from $\alpha^{n,k}$ to $\widehat{\alpha}$.

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

$$\begin{aligned} &\frac{1}{n}\sum_{k=1}^n \mathbb{E}\left[\int_0^T f(X_t^{n,k}, m^n_{\boldsymbol{X}^n_t}, \alpha^{n,k}(t, \boldsymbol{X}^n))dt + g(X_T^{n,k}, m^n_{\boldsymbol{X}^n_t})\right] \\ &\quad \geq -\epsilon + \frac{1}{n}\sum_{k=1}^n \mathbb{E}\left[\int_0^T f(Y_t^{n,k,k}, m^n_{\boldsymbol{Y}^{n,k}_t}, \widehat{\alpha}(t, Y_t^{n,k,k}, m^n_{\boldsymbol{Y}^{n,k}}))dt + g(Y_T^{n,k,k}, m^n_{\boldsymbol{Y}^{n,k}_T})\right]. \end{aligned} \tag{24}$$

We may rewrite the left-hand side as

$$\int_{\mathcal{C}^d\times\mathcal{V}}\left[\int_{[0,T]\times A} f(x_t, m^n_{\boldsymbol{X}^n_t}, a)q(dt,da) + g(x_T, m^n_{\boldsymbol{X}^n_T})\right] m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n}(dx,dq).$$

Note that along a subsequence for which $m^n_{\boldsymbol{X}^n,\boldsymbol{\Lambda}^n}$ converges to some $\widetilde{\mu}$, we can easily take the limit of this expression to get precisely the left-hand side of the inequality in (5) of Definition 4.6. What remains is to show that the right-hand side of (24) along the same subsequence is precisely the right-hand side of the inequality in (5) of Definition 4.6.

This last point is the technical crux of the argument. It is far from obvious at first how to proceed, because we know very little about the controls $\alpha^{n,1}, \ldots, \alpha^{n,n}$. Intuitively, we would like to claim that $m^n_{\boldsymbol{Y}^{n,k}}$ should be "close" in some sense to $m^n_{\boldsymbol{X}^n}$ and thus shares the same limiting behavior, because we have only switched one single agent's control. 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 $\alpha^{n,1}, \ldots, \alpha^{n,n}$ depend very heavily on, say, player k's state process, in which case a deviation from this player k would have a strong influence on the empirical measure. We may (loosely) refer to such a player k as *influential*. The key point is that this cannot possibly be the case that every (or even "most") players are simultaneously influential. More precisely, while we cannot show that $m^n_{\boldsymbol{Y}^{n,k}}$ and $m^n_{\boldsymbol{X}^n}$ have the same limiting behavior for each k, it can be shown that $\mathcal{L}(m^n_{\boldsymbol{X}^n})$ and

$$\frac{1}{n}\sum_{k=1}^n \mathcal{L}(m^n_{\boldsymbol{Y}^{n,k}})$$

have the same limiting behavior, in the sense that the distance between these two measures converges to zero as $n \to \infty$. There is more to do still once this is rigorously proven, but this is the most important insight in this step.

Appendix A. Weak convergence of probability measures

This short section collects some basic and standard facts about weak convergence of probability measures on metric spaces. This material will not be discussed in the lectures and is included in the notes as a reference for the less probabilistically oriented reader. This machinery forms the foundation for many statements and arguments of these notes, and, in order to get to the meat of the course, we will cover this material far too quickly. For more details, refer to the classic textbook of Billingsley [2]. Weak convergence of processes is treated concisely in Kallenberg's tome [14, Chapter 14], and the book of Parthasarathy [25] is a nice reference for a more topological perspective.

Throughout the section, let $(\mathcal{X}, d)$ denote a metric space. We always equip $\mathcal{X}$ with the Borel σ-field, meaning the σ-field generated by the open sets of $\mathcal{X}$. We will write $\mathcal{P}(\mathcal{X})$ for the set of (Borel) probability measures on X. For a measurable function φ from

$\mathcal{X}$ to another metric space $\mathcal{Y}$, we define the image measure $\mu \circ \varphi^{-1} \in \mathcal{P}(\mathcal{Y})$ by setting $\mu \circ \varphi^{-1}(A) := \mu(\varphi \in A) = \mu\{x \in \mathcal{X} : \varphi(x) \in A\}$. Let $C_b(\mathcal{X})$ denote the set of bounded continuous real-valued functions on X. The fundamental definition is the following, which we state in two equivalent forms, one measure-theoretic and one probabilistic:

Definition A.1. Given $\mu, \nu_n \in \mathcal{P}(\mathcal{X})$, we say that μ_n *converges weakly to* μ, or $\mu_n \to \mu$, if

$$\lim_{n\to\infty} \int_{\mathcal{X}} f \, d\mu_n = \int_{\mathcal{X}} f \, d\mu, \text{ for every } f \in C_b(\mathcal{X}).$$

Alternatively, given a sequence of $\mathcal{X}$-valued random variables (X_n), we say that X_n *converges weakly (or in distribution)* to another $\mathcal{X}$-valued random variable X (often denoted $X_n \Rightarrow X$) if $\mathcal{L}(X_n) \to \mathcal{L}(X)$ weakly, i.e.,

$$\lim_{n\to\infty} \mathbb{E}[f(X_n)] = \mathbb{E}[f(X)], \text{ for every } f \in C_b(\mathcal{X}).$$

The following theorem will often be used implicitly and is completely trivial to prove:

Theorem A.2 (Continuous mapping theorem)**.** *Suppose $\mathcal{X}$ and $\mathcal{Y}$ are metric spaces, and (X_n) is a sequence of $\mathcal{X}$-valued random variables converging in distribution to another $\mathcal{X}$-valued random variable X. Suppose $g : \mathcal{X} \to \mathcal{Y}$ is a continuous function. Then $g(X_n) \Rightarrow g(X)$.*

It is an important fact that weak convergence of probability measures corresponds to a metric topology on $\mathcal{P}(\mathcal{X})$, at least when the underlying metric space $(\mathcal{X}, d)$ is separable. The weak convergence topology on $\mathcal{P}(\mathcal{X})$ is of course generated by the basic open sets of the form

$$\{\nu \in \mathcal{P}(\mathcal{X}) : |\langle \nu, f_i \rangle - \langle \mu, f_i \rangle| < \epsilon_i, \quad i = 1, \ldots, k\},$$

where $k \in \mathbb{N}$, $\mu \in \mathcal{P}(\mathcal{X})$, and $f_1, \ldots, f_k \in C_b(\mathcal{X})$. The following theorem is well known and can be found, for instance, in [25, Section II.6].

Theorem A.3. *If $(\mathcal{X}, d)$ is a separable metric space, then $\mathcal{P}(\mathcal{X})$ can be metrized as a separable metric space. If $(\mathcal{X}, d)$ is a complete and separable metric space, then $\mathcal{P}(\mathcal{X})$ can be metrized as a complete and separable metric space. If $(\mathcal{X}, d)$ is a compact metric space, then $\mathcal{P}(\mathcal{X})$ can be metrized as a compact metric space.*

There are several popular metrics, including the Levy-Prokhorov metric, the bounded-Lipschitz metric, or the family of Wasserstein metrics discussed in Section A.1 below.

Characterizing compact sets of $\mathcal{P}(\mathcal{X})$ is, unsurprisingly, quite useful in our study of the convergence problem via compactification methods in Section 4. The classical theorem of Prokhorov accomplishes this. Given a set $S \subset \mathcal{P}(\mathcal{X})$, we say that the family S of probability measures is *tight* if for all ϵ there exists a compact set $K \subset X$ such that

$$\sup_{\mu\in S} \mu(K^c) \le \epsilon.$$

The following can be found in [2, Theorem 6.1, 6.2] and [14, Theorem 16.3]

Theorem A.4 (Prokhorov's theorem)**.** *Suppose $S \subset \mathcal{P}(\mathcal{X})$. If S is tight, then it is pre-compact in $\mathcal{P}(\mathcal{X})$. Conversely, if S is pre-compact, and if the metric space $(\mathcal{X}, d)$ is separable and complete, then S is tight.*

As a first application, we note that tightness lends itself nicely to working on product spaces. The proof is left as an exercise.

Lemma A.5. *Suppose $\mathcal{X}_1$ and $\mathcal{X}_2$ are complete, separable metric spaces, and endow $\mathcal{X}_1 \times \mathcal{X}_2$ with any metric compatible with the product topology. Define the projections $\pi_i : \mathcal{X}_1 \times \mathcal{X}_2 \to \mathcal{X}_i$, for $i = 1, 2$. Then a set $S \subset \mathcal{P}(\mathcal{X}_1 \times \mathcal{X}_2)$ is tight if and only if the sets $S_1 = \{\mu \circ \pi_1^{-1} : \mu \in S\}$ and $S_2 = \{\mu \circ \pi_2^{-1} : \mu \in S\}$ are tight in $\mathcal{P}(\mathcal{X}_1)$ and $\mathcal{P}(\mathcal{X}_2)$, respectively.*

Another useful lemma says essentially that in order to deduce convergence of integrals of an *unbounded* function from weak convergence of probability measures, one needs a certain uniform integrability. The proof is left as an exercise.

Lemma A.6. *Let $\mathcal{X}$ be a complete separable metric space, and suppose $\mu, \mu_n \in \mathcal{P}(\mathcal{X})$ are such that $\mu_n \to \mu$ weakly. Suppose $f : \mathcal{X} \to \mathbb{R}$ satisfies*

$$\lim_{r\to\infty} \sup_n \int_{\{|f|\geq r\}} |f| \, d\mu_n = 0.$$

Then $\lim_{n\to\infty} \int_{\mathcal{X}} f \, d\mu_n = \int_{\mathcal{X}} f \, d\mu$.

A.1. **Wasserstein metrics.** It will at times be convenient to work with Wasserstein metrics. The definition of the Wasserstein metric is based on the idea of a *coupling*. Assume hereafter that the metric space $(\mathcal{X}, d)$ is separable. For $\mu, \nu \in \mathcal{P}(\mathcal{X})$, we write $\Pi(\mu, \nu)$ to denote the set of Borel probability measures π on $\mathcal{X} \times \mathcal{X}$ with first marginal μ and second marginal ν. Precisely, $\pi(A \times \mathcal{X}) = \mu(A)$ and $\pi(\mathcal{X} \times A) = \nu(A)$ for every Borel set $A \subset \mathcal{X}$. For $p \geq 1$, define $\mathcal{P}^p(\mathcal{X})$ to be the set of probability measures $\mu \in \mathcal{P}(\mathcal{X})$ satisfying

$$\int_{\mathcal{X}} d^p(x, x_0)\mu(dx) < \infty,$$

where $x_0 \in \mathcal{X}$ is an arbitrary reference point. (By the triangle inequality, the choice of x_0 is inconsequential.) The p-Wasserstein metric on $\mathcal{P}^p(\mathcal{X})$ is defined by

$$\mathcal{W}_{\mathcal{X},p}(\mu, \nu) = \left(\inf_{\pi\in\Pi(\mu,\nu)} \int_{\mathcal{X}\times\mathcal{X}} d(x, y)^p \, \pi(dx, dy) \right)^{1/p}. \tag{25}$$

If the space $\mathcal{X}$ is understood, we write simply $\mathcal{W}_p$ instead of $\mathcal{W}_{\mathcal{X},p}$. It is well known that $(\mathcal{P}^p(\mathcal{X}), \mathcal{W}_{\mathcal{X},p})$ is a complete and separable metric space whenever $(\mathcal{X}, d)$ is complete and separable; see [30, Theorem 7.3] or the note of Bolley [3]. The following are equivalent, for $\mu, \mu_n \in \mathcal{P}^p(\mathcal{X})$:

(i) $\lim_{n\to\infty} \mathcal{W}_{\mathcal{X},p}(\mu_n, \mu) = 0$.

(ii) $\mu_n \to \mu$ weakly and, for some $x_0 \in \mathcal{X}$,

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

(iii) $\mu_n \to \mu$ weakl and, for some $x_0 \in \mathcal{X}$,

$$\lim_{n\to\infty} \int d^p(x, x_0) \, \mu_n(dx) \to \int d^p(x, x_0) \, \mu(dx).$$

(iv) $\mu_n \to \mu$ weakl and, for each $x_0 \in \mathcal{X}$ and each continuous function $f : \mathcal{X} \to \mathbb{R}$ satisfying $\sup_{x\in\mathcal{X}} |f(x)|/(1 + d(x, x_0)) < \infty$, we have $\langle \mu_n, f \rangle \to \langle \mu, f \rangle$.

If the metric space $(\mathcal{X}, d)$ is bounded (meaning d is bounded), then $\mathcal{W}_{\mathcal{X},p}$ metrizes the usual topology of weak convergence. Note that any metric space $(\mathcal{X}, d)$ admits a bounded metric which generates the same topology, such as $\min\{d, 1\}$, and under this metric we have $\mathcal{P}^p(\mathcal{X}) = \mathcal{P}(\mathcal{X})$ and $\mu_n \to \mu$ weakly if and only if $\mathcal{W}_{\mathcal{X},p}(\mu_n, \mu) \to 0$. This gives an explicit example of a metric for Theorem A.3.

A.2. **Empirical measures.** In the text we make frequent use of the following well known forms of the law of large numbers.

Theorem A.7. *Suppose $\mathcal{X}$ is a separable metric space. Let $(X_i)_{i\in\mathbb{N}}$ be i.i.d. $\mathcal{X}$-valued random variables with common law $\mu \in \mathcal{P}(\mathcal{X})$. Then*

$$\mathbb{P}\left(\frac{1}{n} \sum_{i=1}^{n} \delta_{X_i} \to \mu \right) = 1,$$

where as usual convergence in $\mathcal{P}(\mathcal{X})$ is with respect to the topology of weak convergence. If $\mu \in \mathcal{P}^p(\mathcal{X})$ for some $p > 1$, then we also have

$$\mathbb{P}\left(\lim_{n\to\infty} \mathcal{W}_{\mathcal{X},p}\left(\frac{1}{n}\sum_{i=1}^n \delta_{X_i}, \mu\right) = 0\right) = 1,$$

$$\lim_{n\to\infty} \mathbb{E}\left[\mathcal{W}^p_{\mathcal{X},p}\left(\frac{1}{n}\sum_{i=1}^n \delta_{X_i}, \mu\right)\right] = 0.$$

Exercise A.8. Prove Theorem A.7 using the strong law of large numbers and the characterizations of $\mathcal{W}_{\mathcal{X},p}$ convergence given in Section A.1.

Appendix B. Calculus on $\mathcal{P}(\mathbb{R}^d)$

Before we can define the master equation, which is a PDE set on $[0,T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$, we must first specify how to differentiate a function of a probability measure. Let us say that a function $U : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ is $\mathscr{C}^1$ if there exists a bounded continuous function $\frac{\delta U}{\delta m} : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}$ such that, for every $m, m' \in \mathcal{P}(\mathbb{R}^d)$,

$$U(m') - U(m) = \int_0^1 \int_{\mathbb{R}^d} \frac{\delta U}{\delta m}((1-t)m + tm', v)\, (m' - m)(dv)\, dt. \tag{26}$$

Equivalently,

$$\frac{d}{dh}\Big|_{h=0} U(m + h(m' - m)) = \int_{\mathbb{R}^d} \frac{\delta U}{\delta m}(m, v)\, (m' - m)(dv). \tag{27}$$

Only one function $\frac{\delta U}{\delta m}$ can satisfy (26), up to a constant shift; that is, if $\frac{\delta U}{\delta m}$ satisfies (26) then so does $(m, v) \mapsto \frac{\delta U}{\delta m}(m, v) + c(m)$ for any function $c : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$. For concreteness we always choose the shift $c(m)$ to ensure

$$\int_{\mathbb{R}^d} \frac{\delta U}{\delta m}(m, v)\, m(dv) = 0. \tag{28}$$

If $\frac{\delta U}{\delta m}(m, v)$ is continuously differentiable in v, we define its *intrinsic derivative* $D_m U : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}^d$ by

$$D_m U(m, v) = D_v\left(\frac{\delta U}{\delta m}(m, v)\right), \tag{29}$$

where we use the notation D_v for the gradient in v.

Remark B.1. One may develop a similar theory of differentiation for functions on $\mathcal{P}^p(\mathbb{R}^d)$, for any exponent $p \geq 1$, as long as one is careful to require that the derivative $\delta U/\delta m$, if unbounded, satisfies some kind of growth conditions to ensure that the integral on the right-hand side of (26) is well-defined.

Before developing any further theory, we mention a few examples:

Example B.2. Suppose $U(m) = \langle m, \varphi\rangle$ for $\varphi : \mathbb{R}^d \to \mathbb{R}$ bounded and continuous, where $\langle m, \varphi\rangle = \int \varphi\, dm$. Then, for any $m, m' \in \mathcal{P}(\mathbb{R}^d)$,

$$\frac{d}{dh}\Big|_{h=0} U(m + h(m' - m)) = \frac{d}{dh}\Big|_{h=0} \left(\langle m, \varphi\rangle + h\langle m' - m, \varphi\rangle\right) = \int \varphi\, d(m' - m).$$

Using (27), this shows that

$$\frac{\delta U}{\delta m}(m, v) = \varphi(v).$$

If φ is continuously differentiable, then

$$D_m U(m,v) = D_v \varphi(v).$$

Example B.3. Suppose $U(m) = F(\langle m, \varphi\rangle)$ for smooth bounded functions $\varphi : \mathbb{R}^d \to \mathbb{R}$ and $F : \mathbb{R} \to \mathbb{R}$. Using the chain rule and the calculation of the previous example, we find

$$\frac{\delta U}{\delta m}(m,v) = F(m)\varphi(v), \qquad D_m U(m,v) = F(m) D_v \varphi(v).$$

Example B.4. Suppose $U(m) = \iint \phi(x,y) m(dx) m(dy)$ for some smooth bounded function ϕ on $\mathbb{R}^d \times \mathbb{R}^d$. First observe

$$\begin{aligned} U(m + h(\tilde{m} - m)) = h^2 \iint \phi d(\tilde{m} - m)^2 + h \iint \phi(x,y) m(dx)(\tilde{m} - m)(dx, dy) \\ + h \iint \phi(x,y)(\tilde{m} - m)(dx) m(dy) + \iint \phi dm^2 \end{aligned}$$

and therefore

$$\frac{d}{dh}|_{h=0} U(m + h(\tilde{m} - m)) = \iint_{\mathbb{R}^d \times \mathbb{R}^d} [\phi(x,y) + \phi(y,x)] m(dy)(\tilde{m} - m)(dx).$$

From this we conclude

$$\frac{\delta U}{\delta m}(m,v) = \int [\phi(x,v) + \phi(v,x)] m(dx).$$

We also make use of second derivatives of functions on $\mathcal{P}(\mathbb{R}^d)$. If, for each $v \in \mathbb{R}^d$, the map $m \mapsto \frac{\delta U}{\delta m}(m,v)$ is $\mathscr{C}^1$, then we say that U is $\mathscr{C}^2$ and let $\frac{\delta^2 U}{\delta m^2}$ denote its derivative, or more explicitly,

$$\frac{\delta^2 U}{\delta m^2}(m,v,v') = \frac{\delta}{\delta m}\left(\frac{\delta U}{\delta m}(\cdot, v)\right)(m,v').$$

We will also make some use of the derivative

$$D_v D_m U(m,v) = D_v[D_m U(m,v)],$$

when it exists, and we note that $D_v D_m U$ takes values in $\mathbb{R}^{d\times d}$. Finally, if U is $\mathscr{C}^2$ and if $\frac{\delta^2 U}{\delta m^2}(m,v,v')$ is twice continuously differentiable in (v,v'), we let

$$D_m^2 U(m,v,v') = D^2_{v,v'} \frac{\delta^2 U}{\delta m^2}(m,v,v')$$

denote the $d \times d$ matrix of partial derivatives $(\partial_{v_i}\partial_{v'_j}[\delta^2 U/\delta m^2](m,v,v'))_{i,j}$. We have the following lemma, whose proof we leave as an exercise:

Lemma B.5. *If U is $\mathscr{C}^2$, then we have*

$$D_m^2 U(m,v,v') = D_m^2 U(m,v',v).$$

Moreover, we may write

$$D_m^2 U(m,v,v') = D_m[D_m U(\cdot, v)](m,v').$$

That is, if we fix v and apply the operator D_m to the function $m \mapsto D_m U(m,v)$, then the resulting function is $D_m^2 U$.

The most important result for our purposes is the following, which shows how these derivatives interact with empirical measures.

Proposition B.6. *Given $U : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$, define $u_n : (\mathbb{R}^d)^n \to \mathbb{R}$ by $u_n(\boldsymbol{x}) = U(m^n_{\boldsymbol{x}})$ for some fixed $n \geq 1$.*

(i) *If U is $\mathscr{C}^1$ and if D_mU exists and is bounded and jointly continuous, then u_n is continuously differentiable, and*

$$D_{x_j}u_n(\boldsymbol{x}) = \frac{1}{n}D_mU(m^n_{\boldsymbol{x}}, x_j), \text{ for } j = 1, \ldots, n. \tag{30}$$

(ii) *If U is $\mathscr{C}^2$ and if D^2_mU exists and is bounded and jointly continuous, then u is twice continuously differentiable, and*

$$D_{x_k}D_{x_j}u_n(\boldsymbol{x}) = \frac{1}{n^2}D^2_mU(m^n_{\boldsymbol{x}}, x_j, x_k) + \delta_{j,k}\frac{1}{n}D_vD_mU(m^n_{\boldsymbol{x}}, x_j),$$

where $\delta_{j,k} = 1$ if $j = k$ and $\delta_{j,k} = 0$ if $j \neq k$.

Proof. Let $m \in \mathcal{P}(\mathbb{R}^d)$ and $\boldsymbol{x} = (x_1, \ldots, x_n) \in (\mathbb{R}^d)^n$. By continuity, it suffices to prove the claims assuming the points $x_1, \ldots, x_n \in \mathbb{R}^d$ are distinct. Fix an index $j \in \{1, \ldots, n\}$ and a bounded continuous function $\phi : \mathbb{R}^d \to \mathbb{R}^d$, to be specified later. We claim that, under the assumptions of part (i),

$$\lim_{h \downarrow 0} \frac{U(m \circ (\mathrm{Id} + h\phi)^{-1}) - U(m)}{h} = \int_{\mathbb{R}^d} D_mU(m, v) \cdot \phi(v)m(dv) \tag{31}$$

holds, where Id denotes the identity map on $\mathbb{R}^d$. Once (31) is proven, we complete the proof as follows. For a fixed vector $v \in \mathbb{R}^d$ we may choose ϕ such that $\phi(x_j) = v$ while $\phi(x_i) = 0$ for $i \neq j$. Let $\boldsymbol{v} \in (\mathbb{R}^d)^n$ have j^{th} coordinate equal to v and i^{th} coordinate zero for $i \neq j$. Then $u_n(\boldsymbol{x}) = U(m^n_{\boldsymbol{x}})$ satisfies

$$\begin{aligned}
\lim_{h \downarrow 0} \frac{u_n(\boldsymbol{x} + h\boldsymbol{v}) - u_n(\boldsymbol{x})}{h} &= \lim_{h \downarrow 0} \frac{U(m^n_{\boldsymbol{x}} \circ (\mathrm{Id} + h\phi)^{-1}) - U(m^n_{\boldsymbol{x}})}{h} \\
&= \frac{1}{n}\sum_{k=1}^n D_mU(m^n_{\boldsymbol{x}}, x_k) \cdot \phi(x_k) \\
&= \frac{1}{n}D_mU(m^n_{\boldsymbol{x}}, x_j) \cdot v.
\end{aligned}$$

This proves (i). Under the additional assumptions, (ii) follows by applying (i) again.

It remains to prove (31). For $h > 0$, $t \in [0, 1]$, and $m \in \mathcal{P}(\mathbb{R}^d)$, let $m_{h,t} = tm \circ (\mathrm{Id} + h\phi)^{-1} + (1 - t)m$. Then, using (26) and (29), respectively, in the first and third equalities below, we obtain

$$\begin{aligned}
U(m \circ (\mathrm{Id} + h\phi)^{-1} - U(m) &= \int_0^1 \int_{\mathbb{R}^d} \frac{\delta U}{\delta m}(m_{h,t}, v)\, (m \circ (\mathrm{Id} + h\phi)^{-1} - m)(dv)\, dt \\
&= \int_0^1 \int_{\mathbb{R}^d} \left(\frac{\delta U}{\delta m}(m_{h,t}, v + h\phi(v)) - \frac{\delta U}{\delta m}(m_{h,t}, v) \right) m(dv)\, dt \\
&= h \int_0^1 \int_{\mathbb{R}^d} \int_0^1 D_mU(m_{h,t}, v + sh\phi(v)) \cdot \phi(v)\, ds\, m(dv)\, dt.
\end{aligned}$$

As $h \downarrow 0$ we have $m_{h,t} \to m$ and $sh\phi(v) \to 0$, and we deduce (31) from the bounded convergence theorem and continuity of D_m. ☐

Lastly, and in part for more practice working with this notion of derivative, we prove the following technical result, which will be used implicitly. It simply says that, for a function $U = U(x, m)$ on $\mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$, the derivatives in x and in m commute. The proof is left as an exercise for the reader.

Lemma B.7. *For any function $U = U(x, m) : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$, we have*

$$D_xD_mU(x, m, v) = D_mD_xU(x, m, v), \qquad \text{for all } v \in \mathbb{R}^d,$$

as long as the derivatives on both sides exist and are bounded and jointly continuous.

References

1. A. Bensoussan, J. Frehse, and P. Yam, *Mean field games and mean field type control theory*, vol. 101, Springer, 2013.
2. P. Billingsley, *Convergence of probability measures*, John Wiley & Sons, 2013.
3. F. Bolley, *Separability and completeness for the Wasserstein distance*, Séminaire de probabilités XLI, Springer, 2008, pp. 371–377.
4. P. Cardaliaguet, *Notes on mean field games*, Tech. report, Technical report, 2010.
5. P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions, *The master equation and the convergence problem in mean field games*, arXiv preprint arXiv:1509.02505 (2015).
6. R. Carmona and F. Delarue, *Probabilistic analysis of mean-field games*, SIAM Journal on Control and Optimization **51** (2013), no. 4, 2705–2734.
7. ______, *Probabilistic theory of mean field games with applications i-ii*, Springer, 2017.
8. R. Carmona and D. Lacker, *A probabilistic weak formulation of mean field games and applications*, The Annals of Applied Probability **25** (2015), no. 3, 1189–1231.
9. A. Figalli, *Existence and uniqueness of martingale solutions for SDEs with rough or degenerate coefficients*, Journal of Functional Analysis **254** (2008), no. 1, 109–153.
10. W.H. Fleming and H.M. Soner, *Controlled Markov processes and viscosity solutions*, vol. 25, Springer Science & Business Media, 2006.
11. J. Gärtner, *On the McKean-Vlasov limit for interacting diffusions*, Mathematische Nachrichten **137** (1988), no. 1, 197–248.
12. D.A. Gomes and J. Saúde, *Mean field games models—a brief survey*, Dynamic Games and Applications **4** (2014), no. 2, 110–154.
13. M. Huang, R. Malhamé, and P. Caines, *Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle*, Communications in Information & Systems **6** (2006), no. 3, 221–252.
14. O. Kallenberg, *Foundations of modern probability*, Springer Science & Business Media, 2006.
15. I. Karatzas and S. Shreve, *Brownian motion and stochastic calculus*, vol. 113, Springer Science & Business Media, 2012.
16. N.V. Krylov and M. Roeckner, *Strong solutions of stochastic equations with singular time dependent drift*, Probability theory and related fields **131** (2005), no. 2, 154–196.
17. D. Lacker, *Mean field games via controlled martingale problems: existence of Markovian equilibria*, Stochastic Processes and their Applications **125** (2015), no. 7, 2856–2894.
18. ______, *A general characterization of the mean field limit for stochastic differential games*, Probability Theory and Related Fields **165** (2016), no. 3-4, 581–648.
19. ______, *On the convergence of closed-loop Nash equilibria to the mean field game limit*, arXiv preprint arXiv:1808.02745 (2018).
20. Huang M, R. Malhamé, and P. Caines, *Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior and decentralized ε-Nash equilibria*, IEEE transactions on automatic control **52** (2007), no. 9, 1560–1571.
21. H.P. McKean, *A class of Markov processes associated with nonlinear parabolic equations*, Proceedings of the National Academy of Sciences **56** (1966), no. 6, 1907–1911.
22. ______, *Propagation of chaos for a class of non-linear parabolic equations*, Stochastic Differential Equations (Lecture Series in Differential Equations, Session 7, Catholic Univ., 1967) (1967), 41–57.
23. S. Méléard, *Asymptotic behaviour of some interacting particle systems; McKean-Vlasov and Boltzmann models*, Probabilistic models for nonlinear partial differential equations, Springer, 1996, pp. 42–95.
24. K. Oelschläger, *A martingale approach to the law of large numbers for weakly interacting stochastic processes*, The Annals of Probability (1984), 458–479.
25. K.R. Parthasarathy, *Probability measures on metric spaces*, vol. 352, American Mathematical Soc., 2005.
26. H. Pham, *Continuous-time stochastic control and optimization with financial applications*, vol. 61, Springer Science & Business Media, 2009.
27. A.S. Sznitman, *Topics in propagation of chaos*, Ecole d'eté de Probabilités de Saint-Flour XIX - 1989, Lecture Notes in Mathematics, vol. 1464, Springer Berlin / Heidelberg, 1991, pp. 165–251.
28. D. Trevisan, *Well-posedness of multidimensional diffusion processes with weakly differentiable coefficients*, Electronic Journal of Probability **21** (2016).
29. A. Veretennikov, *On strong solutions and explicit formulas for solutions of stochastic integral equations*, Sbornik: Mathematics **39** (1981), 387–403.
30. C. Villani, *Topics in optimal transportation*, no. 58, American Mathematical Soc., 2003.