

STOCHASTIC DIFFERENTIAL MEAN FIELD GAME THEORY

DANIEL LACKER

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
OPERATIONS RESEARCH AND FINANCIAL ENGINEERING
ADVISER: RENÉ CARMONA

SEPTEMBER 2015

© Copyright by Daniel Lackner, 2015.
All rights reserved.

Abstract

Mean field game (MFG) theory generalizes classical models of interacting particle systems by replacing the particles with rational agents, making the theory more applicable in economics and other social sciences. Intuitively, (stochastic differential) MFGs are infinite-population or continuum limits of large-population stochastic differential games of a certain symmetric type, and a solution of an MFG is analogous to a Nash equilibrium. This thesis tackles several fundamental problems in MFG theory. First, if (approximate) equilibria exist in the large-population games, to what limits (if any) do they converge as the population size tends to infinity? Second, can the limiting system be used to construct approximate equilibria for the finite-population games? Finally, what can be said about existence and uniqueness of equilibria, for the finite- or infinite-population models?

This thesis presents a complete picture of the limiting behavior of the large-population systems, both with and without common noise, under modest assumptions on the model inputs. Approximate Nash equilibria in the n -player games admit certain weak limits as n tends to infinity, and every limit is a weak solution of the MFG. Conversely, every weak MFG solution can be obtained as the limit of a sequence of approximate Nash equilibria in the n -player games. Even in the setting without common noise, a new solution concept is needed in order to capture all of the possible limits. Interestingly, and in contrast with well known results on related interacting particle systems, empirical state distributions often admit stochastic limits which are not simply randomizations among the deterministic solutions.

With the limit theory in mind, the thesis then develops new existence and uniqueness results. Using controlled martingale problems together with relaxed controls, a general existence theorem is derived by means of Kakutani's fixed point theorem. In the common noise case, a natural notion of weak solution is introduced, and the existence and uniqueness theory is designed in perfect analogy with weak solutions of stochastic differential equations. An existence theorem for weak solutions is proven by a discretization procedure, and a Yamada-Watanabe result is presented and illustrated under some stronger assumptions which ensure pathwise uniqueness.

Acknowledgements

There are so many people who supported me both academically and personally during my time at Princeton, and all of them have influenced my research in some way.

First and foremost, I am infinitely grateful to my advisor, René Carmona, for teaching me not just how to answer questions but how to ask them. His passion for research and genuine curiosity are highly contagious, and he imparted in me a courage to challenge myself and to blaze my own trail. He was always patient when I made wrong turns and wandered down questionable roads, and indeed there were many in my early days. From René I learned how to explore the wilderness of research, and to never doubt that I can make it over the next hill.

I had the pleasure of interacting with a wonderful group of professors at Princeton. I am grateful for the encouragement and stimulating discussions with Ramon van Handel, Ronnie Sircar, Birgit Rudloff, Patrick Cheridito, Erhan Çinlar, and also François Delarue and Shige Peng during their visits to Princeton. Erhan is a well known stickler for notation, and, while I hesitate to make any claims regarding the elegance of the notation of this thesis, it certainly would have been far less readable without his influence. I was fortunate to have the opportunity to work closely with François Delarue, both in Princeton and in Nice, which resulted in some of the material of this thesis. I hope to emulate his incredibly careful work ethic in my own career. Special thanks are due also to Ramon van Handel, who was always willing to take the time to help me out. I've approached Ramon (often unannounced) with such a range of questions, and he is always happy to share his thoughts. His uncanny ability to communicate difficult concepts has undoubtedly shaped the way I think about mathematics.

I will miss working so closely with my office-mates and academic brothers, Mark Cerenzia, Kevin Webster, Geoffrey Zhu, and Michael Li. Many thanks to Kevin, for all of the great discussions over the years, and for teaching me nearly everything I know about the strange and fascinating world of high frequency trading. Kevin was a tough act to follow, but Mark was undoubtedly a worthy successor. Mark shares our advisor's contagious curiosity, and I've learned a great many things from him, ranging from acutely relevant mathematics to the meaning of "turn down for what."

I certainly can't leave out my rock climbing pals and all of the Princeton University Mountaineering Club. Thanks to all of you for making grad school way more fun than it had any right to be. Shout-outs are due especially to Alex, Jono, and Ben: I learned a lot from you guys, from those long road trip discussions of everything from particle physics to gut bacteria, and really life in general. Long live PUMC!

I am amazingly lucky to have such a wonderful family, Jeff, my dad, Lisa, my mom, and Ben, my brother who constantly remind me of what is important in life. They continue to shape me today, and I am so grateful for their unwavering support and encouragement.

Last but absolutely not least, I want to thank Maria, the love of my life, mi media naranja, for so many things, most of all for keeping me human. Thanks to her, I am so much more than a machine for turning coffee into theorems.

To my grandpas, Jack (Zayde) and Bill

Contents

Abstract	iii
Acknowledgements	iv
1 Introduction	1
1.1 From particle systems to mean field games	2
1.1.1 Interacting diffusion models	2
1.1.2 The McKean-Vlasov limit	3
1.1.3 Mean field games	5
1.1.4 Convergence to the mean field game limit	7
1.1.5 Existence	9
1.1.6 Uniqueness	11
1.2 Common noise	11
1.2.1 Interacting diffusion models	12
1.2.2 Mean field games	13
1.3 Extensions of the mean field game framework	14
1.4 Continuum games in economics	16
1.5 Outline of the thesis	17
2 Preliminaries	19
2.1 Spaces of probability measures	19
2.1.1 Wasserstein distances	24
2.1.2 Mean measures and compactness	25
2.1.3 Relaxed controls	27
2.2 McKean-Vlasov limits	30
2.2.1 No common noise	30
2.2.2 Common noise	32
2.2.3 An unorthodox derivation	33
2.3 Static mean field games	37
2.3.1 Elements of set-valued analysis	38
2.3.2 The deterministic case	39
2.3.3 The stochastic case with independent noises	41
2.3.4 The stochastic case with common noise	46
3 Stochastic differential mean field games	48
3.1 MFG solution concepts	48
3.2 Limits of finite games	53

3.2.1	Equilibrium concepts	55
3.2.2	The limit theorem	56
3.2.3	The converse limit theorem	58
3.3	Existence and uniqueness	60
4	Stochastic differential mean field games without common noise	62
4.1	MFG solution concepts	62
4.2	The limit theorem	64
4.3	An illuminating example	65
4.4	Supports of weak solutions	67
4.5	Applications of Theorem 4.4.2	71
4.6	Existence and uniqueness	72
5	Properties of mean field game solutions	74
5.1	Pre-solutions	74
5.2	Canonical space	78
5.3	Continuity results	81
5.4	Proofs of Propositions 3.1.4 and 3.1.5	87
5.4.1	Proof of Proposition 3.1.4	88
5.4.2	Proof of Proposition 3.1.5	90
6	Proofs of the limit theorems	92
6.1	Proof of main limit Theorem 3.2.4	92
6.1.1	Estimates	92
6.1.2	Relative compactness and pre-solutions	96
6.1.3	Modified finite-player games	99
6.1.4	Optimality in the limit	103
6.2	Proof of converse limit Theorem 3.2.10	106
6.2.1	Construction of environments	107
6.2.2	Trajectorial propagation of chaos	108
6.2.3	Proof of Theorem 6.2.1	110
6.2.4	Proof of Proposition 6.2.2	116
6.3	Proof of Theorem 4.2.2	120
6.4	Proofs of Propositions 3.2.2 and 4.2.1	121
6.4.1	Proof of Proposition 3.2.2	121
6.4.2	Proof of Proposition 4.2.1	122
7	Existence, with common noise	123
7.1	Discretized mean field games	123
7.1.1	Discretization procedure	124
7.1.2	Strong solutions	126
7.2	Weak limits of discretized mean field games	130
7.2.1	Existence of a weak solution under Assumption B	132
7.2.2	Unbounded coefficients	134
7.3	Uniqueness	138

7.4	Counterexamples	142
7.4.1	Nonexistence of strong solutions	142
7.4.2	Nonexistence of weak solutions	143
8	Existence, without common noise	147
8.1	Controlled martingale problems	148
8.2	Bounded coefficients	154
8.3	Unbounded coefficients	158
8.3.1	Relative compactness of the approximations	159
8.3.2	Limiting state process dynamics	160
8.3.3	Optimality of the limiting control	163
8.4	A general framework	166
A	Elements of stochastic analysis	169
A.1	Enlargements of filtered probability spaces	169
A.2	Existence and uniqueness for SDEs	171
A.3	Yamada-Watanabe theorems	172
A.4	Martingale measures and martingale problems	174
B	Tightness results for Itô processes	177
	Bibliography	181

Chapter 1

Introduction

Large systems of interacting individuals are central to countless areas of science; the individuals may be people, computers, animals, or particles, and the large systems may be financial markets, networks, flocks, or fluids. Mean field theory (an enormously broad term) was originally developed to study particle systems and has since emerged as the most widespread mathematical foundation for studying a broader class of these systems. The key insight of this approach is that the infinite-population (continuum) limit of the right kind of finite-population model can effectively approximate macroscopic and statistical features of the system as well as the behavior of a typical or average particle. Applications of mean field theory beyond its traditional domain of statistical physics, though plentiful, are often criticized, particularly in the social sciences, for their inability to model individuals as rational. Indeed, the so-called individuals of classical mean field theories behave according to exogenous laws of motion and are thus best understood as particles.

The young theory of *mean field games* directly addresses this criticism and fundamentally generalizes traditional mean field theory by granting individuals choice. Each individual is allowed to optimize some criterion, as an investor maximizes wealth, a manufacturer chooses how much to produce, or a driver avoids traffic. Mean field game (MFG) theory again facilitates succinct descriptions of the behavior of a representative agent as well as the *distribution of states* across the population and over time. The distribution of states could variously represent, for example, the income distribution in a given country or the distribution of fish in a school. In contrast with classical mean field theory, the dynamics of the system emerge *endogenously* in a (typically) competitive equilibrium. Because equilibria of large competitive systems are usually difficult to analyze, MFG theory again seeks more tractable infinite-population limits that retain important statistical features of finite systems.

MFG theory has the potential to advance research on a number of problems of intellectual and practical importance, from financial market stability to the dynamics of the income distribution. However, incorporating choice naturally renders MFG models much more complex than their classical mean field counterparts. As a result, the demands of applications far exceed their as yet underdeveloped theoretical foundations. The original developments around 2006 (see the work of Lasry and Lions [91, 89, 90] and Huang et al. [67, 68]) painted a broad picture of the possibilities of MFG theory and its applications. Subsequent research, however, has focused largely on theoretical questions of existence and uniqueness of solutions for the equations governing the particular class of MFG systems that warrants the more

specific title of *stochastic differential mean field games*. This thesis deepens the analysis of these two questions but also investigates the under-emphasized problem of rigorously justifying the mean field limit. The goal of the thesis is to study the following three problems, in the context of stochastic differential mean field games:

- (1) What is the precise nature of the *mean field game limit*? More specifically, can we identify all of the possible limits of n -player approximate Nash equilibria as $n \rightarrow \infty$ with a sensible *mean field game* equilibrium concept? This “limit” should at least capture the behavior of the distribution of state processes.
- (2) When does there exist an (approximate) Nash equilibria for the finite games, and when does there exist an equilibrium for the mean field game?
- (3) When is the solution of the mean field game unique?

A good equilibrium concept for the mean field game would render the set of MFG equilibria *precisely* equal to the set of limits of approximate Nash equilibria. In this case, the two questions of point (2) are equivalent, and uniqueness results for the MFG immediately imply that there is a unique limit to which any n -player approximate Nash equilibria converge. This thesis presents some solutions to these problems, rigorously proving the correspondence between the finite-population game and a new “weak solution” concept for the mean field game. Moreover, under various assumptions on the model inputs, several existence and uniqueness results for this notion of weak solution are derived, along with their implications for the limit theory.

Throughout the thesis, special attention is paid to models with *common noise*, also known as *aggregate uncertainty*. These models, though largely neglected in the extant literature, are important for applications due to their ability to model certain types of macroscopic sources of noise which persist in the mean field limit.

The rest of this introductory chapter describes mean field game theory in more detail and summarizes the known results on the foundational questions outlined above. Some time is spent first on background material from the theory of McKean-Vlasov limits of interacting diffusions, which can be seen, in contrast with mean field games, as *zero-intelligence models*.

1.1 From particle systems to mean field games

1.1.1 Interacting diffusion models

Let us begin by describing the interacting particle systems on which stochastic differential mean field games are based. These systems are by now well understood and serve as a chief source of intuition when studying their competitive game counterpart.

Imagine n particles are moving in d -dimensional space continuously in time, and the position of particle i at time t is denoted X_t^i . The particle system evolves according to a system of stochastic differential equations (SDEs) of the form

$$\begin{cases} dX_t^i &= b(X_t^i, \hat{\mu}_t^n)dt + \sigma(X_t^i, \hat{\mu}_t^n)dW_t^i, & i = 1, \dots, n, \\ \hat{\mu}_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}. \end{cases} \quad (1.1)$$

There are several important structural features of this model that warrant discussion. First of all, the driving noises W^1, \dots, W^n are independent Wiener processes, and we note that particle i is influenced directly only by W^i . We assume also that the initial positions X_0^1, \dots, X_0^n are i.i.d. The drift and volatility functions b and σ are *the same* for each particle, but for particle i these functions are evaluated at the particle's own position X_t^i as well as the *empirical distribution* of the n particles' positions. To be clear, the arguments of the functions b and σ are a spatial variable and a probability measure. If the Wiener processes W^1, \dots, W^n are correlated, we have a *common noise* model; we postpone the discussion of such models to Section 1.2.

A typical special class of this model arises when the drift (and volatility, though we will typically take this to be constant for simplicity) is of the form

$$b(x, \mu) = \int \tilde{b}(x, y) \mu(dy),$$

for a function \tilde{b} taking two spatial variables as arguments. The SDEs above then takes the form

$$dX_t^i = \frac{1}{n} \sum_{j=1}^n \tilde{b}(X_t^i, X_t^j) dt + \sigma dW_t^i.$$

This model was introduced by McKean in [93], building on some ideas of Kac [72], in an effort to rigorously derive certain reduced equations (e.g., Burger's or Boltzmann's) from finite-particle models. The more general form is discussed in the monograph of Sznitman [104].

One reason these particle systems admit tractable limits as $n \rightarrow \infty$ is their symmetry. Of course, the particles are exchangeable in the sense that the distribution of $(X^{\pi(1)}, \dots, X^{\pi(n)})$ is the same for any choice of permutation π of $\{1, \dots, n\}$, at least when the SDE system is well-posed. Moreover, when n is large, the influence of a single particle on the empirical measure $\hat{\mu}^n$ is small; since this is the only source of coupling or interdependence between the particles, we expect intuitively that some asymptotic independence should arise as n tends to infinity.

Particle systems of this form are quite natural starting points for many scientific models, and understanding their limiting behavior is often an important tool in their analysis. On the other hand, in some applications the limiting system is the starting point (e.g., an idealized physical model), and the finite particle system is used primarily for the purpose of simulation and numerical approximation [21, 84].

1.1.2 The McKean-Vlasov limit

A well known heuristic argument allows us to identify the candidate *mean field limit* of the system (1.1). Suppose for the moment that $(\hat{\mu}_t^n)_{t \in [0, T]}$ converges to a *deterministic* measure flow $(\mu_t)_{t \in [0, T]}$. Then, if b and σ are suitably continuous, the limiting dynamics of a single particle should become

$$d\bar{X}_t^i = b(\bar{X}_t^i, \mu_t) dt + \sigma(\bar{X}_t^i, \mu_t) dW_t^i.$$

That is, X^i should converge to \bar{X}^i in some sense. Then, for continuous functions f , we should have both $\mathbb{E}[f(X_t^i)] \rightarrow \mathbb{E}[f(\bar{X}_t^i)]$ and $\mathbb{E}[f(X_t^i)] = \mathbb{E} \int f d\hat{\mu}_t^n \rightarrow \int f d\mu_t$ for each i , and so μ_t should actually agree with the law of \bar{X}_t^i . In other words, the X^i should converge in a sense to independent copies of the solution of the *McKean-Vlasov equation*

$$dX_t = b(X_t, \text{Law}(X_t))dt + \sigma(X_t, \text{Law}(X_t))dW_t. \quad (1.2)$$

Alternatively, the dynamics of $(\mu_t)_{t \in [0, T]}$ can be described by a nonlinear Kolmogorov forward equation for μ , which may be written (assuming μ_t has a smooth density)

$$\partial_t \mu_t = L_{\mu_t}^* \mu_t = -\partial_x(b(x, \mu_t)\mu_t(x)) + \frac{1}{2}\partial_x^2(\sigma(x, \mu_t)^2\mu_t(x)), \quad (1.3)$$

where $L_{\mu_t}^*$ is the adjoint of the operator L_{μ_t} given by (assuming the particles live in one dimension)

$$L_{\mu_t}\varphi(x) = b(x, \mu_t)\varphi'(x) + \frac{1}{2}\sigma(x, \mu_t)^2\varphi''(x).$$

This partial differential equation (PDE) is also sometimes called the McKean-Vlasov equation, and it may be derived directly from the n -particle model (again heuristically) by applying Itô's formula to $\varphi(X_t^i)$, for smooth φ :

$$d \int \varphi d\hat{\mu}_t^n = \frac{1}{n} \sum_{i=1}^n d\varphi(X_t^i) = \left(\int L_{\mu_t}\varphi d\hat{\mu}_t^n \right) dt + \frac{1}{n} \sum_{i=1}^n \varphi'(X_t^i)\sigma(X_t^i, \mu_t) dW_t^i. \quad (1.4)$$

Since W^i are independent, the last term should vanish as n tends to infinity, and the equation which results is simply a weak form of the equation (1.3).

More precisely, when the McKean-Vlasov equation admits a unique solution, it has been shown rigorously in many settings that the n -particle empirical measures converge in some sense to this unique solution. This type of result is known as *propagation of chaos*, a term coined by Mark Kac. On the other hand, when multiple solutions of the McKean-Vlasov equation exist, then typically all one can say is that $\hat{\mu}^n$ admit limits in distribution, and every such limit is (a stochastic measure flow) concentrated on (i.e., supported by) the set of solutions of the McKean-Vlasov equation.

This McKean-Vlasov limit and many variations have been studied thoroughly in the past several decades, using a wide range of techniques. For the basic form of the model outlined here, there are two dominant strategies for rigorously deriving this limit. The first and more widely applicable technique is weak convergence arguments. By placing the empirical measures $(\hat{\mu}_t^n)_{t \geq 0}$ in a good topological space, proving the relative compactness of this sequence typically requires only modest assumptions on the data b and σ . Either of the above heuristic arguments may then be made rigorous in order to characterize the limit points. See [96, 53, 58] for implementations of this strategy.

A second technique, often called *trajectorial propagation of chaos*, tends to yield stronger convergence results but only under accordingly stronger assumptions (e.g., Lipschitz coefficients). These assumptions also yield uniqueness of the McKean-Vlasov equation. The idea is to construct an explicit coupling between the limiting process and the n -particle models,

by building independent copies of the unique solution X of the McKean-Vlasov equation on the same probability space as the finite-particle system and driven by the same Brownian motions. An advantage of this approach is that it permits good estimates of the rate of convergence to the limit. See [104] for details of this approach.

The non-unique regime, in which the n -particle system admits multiple limits, is emphasized in this thesis for a number of reasons. First of all, uniqueness of Nash equilibria is rare in game theory, and mean field games are no exception. This should not to be seen as a nuisance or a pathology, but rather as a fact of life and a potentially useful modelling tool. For example, the existence of both “good” and “bad” equilibria is a critical feature of the seminal Diamond-Dybvig [44] model of bank runs. Even in the particle models described above, non-uniqueness can be exploited to model phase transitions and quantum tunneling, as in a series of papers of Dawson and Gärtner [42, 41, 43].

1.1.3 Mean field games

Let us now return to the setup of the n -particle model (1.1) and bestow upon the particles some capacity for choice. This will turn the model into a stochastic differential game, and we take care to design the general model so as to preserve the system’s symmetry. To reflect their new-found rationality, we will now refer to the “particles” instead as “agents,” and the process X^i is called the *state process* of agent i . The game takes place on a fixed finite time horizon $T > 0$. Agent i chooses a *control process* $\alpha^i = (\alpha_t^i)_{t \in [0, T]}$, which influences the evolution of the state process according to the following dynamics:

$$\begin{cases} dX_t^i &= b(X_t^i, \hat{\mu}_t^n, \alpha_t^i)dt + \sigma(X_t^i, \hat{\mu}_t^n, \alpha_t^i)dW_t^i, \\ \hat{\mu}_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}. \end{cases} \quad (1.5)$$

This is the same SDE as before, except that now the dynamics of the state process of each agent depends additionally on the agent’s own control. Agent i will seek to maximize a certain objective, of the form

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

Note that the *running objective* f and the *terminal objective* g are the same for each agent. Because the data (b, σ, f, g) depend on the empirical measure $\hat{\mu}^n$, these optimization problems are coupled. The optimal strategy of a single agent i depends through $\hat{\mu}^n$ on which controls the other agents choose, but note that this dependence is *anonymous*, in the sense that agent i does not care *which agent* chooses which controls. This is indeed a very particular class of games, and Section 1.3 will discuss some interesting extensions of this most basic class of model.

To simultaneously resolve these optimization problems, we will look for Nash equilibria. Somewhat more generally, we say that $(\alpha^1, \dots, \alpha^n)$ is an ϵ -Nash equilibrium (or more vaguely an approximate Nash equilibrium) if

$$J_i(\alpha^1, \dots, \alpha^n) + \epsilon \geq J_i(\alpha^1, \dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots, \alpha^n)$$

for each admissible alternative strategy β and each $i = 1, \dots, n$. While there are of course many more refined concepts of equilibrium, the Nash equilibrium is the prototypical *competitive* equilibrium. Working at this stage with a cooperative type of equilibrium (e.g., Pareto) leads to an entirely different mean field theory, which is carefully contrasted with the mean field game paradigm in [34].

As is the case in the analogous particle models, understanding the mean field limit of these games can be useful in different ways. First, it could be the case that a continuum model is convenient or tractable to work with but should still be rigorously grounded in a more realistic or tangible finite-population system; this is much in the spirit of Kac's influential work on the Boltzmann equation [72]. On the other hand, it is clear that n -player stochastic differential of this form tend to be quite intractable, especially when n is large. Closed-form solutions of such n -player games are almost never available, with the possible exception of sufficiently simple linear-quadratic models. Numerical analysis would presumably go through a system of n coupled Hamilton-Jacobi-Bellman (HJB) PDEs (see [12]) which describe the value functions of the n agents, but common finite-difference schemes naturally suffer from the curse of dimensionality.¹ With this in mind, one may hope to find a simpler system by passing to the limit $n \rightarrow \infty$. Although the literature is limited so far [2, 1, 27, 28, 87, 86], some numerical methods have been developed for certain types of mean field games, and this paves the way for approximate solutions of otherwise intractable n -player models.

In fact, it is difficult even to abstractly establish the existence of equilibria for n -player stochastic differential games. While existence theorems abound for two-player zero sum stochastic differential games, there are not nearly as many results for nonzero-sum games or games with more than two players; some work in this direction is by PDE methods [12, 13], BSDE methods [62, 62], and relaxed control arguments [20]. A key point of MFG theory, as we will soon discuss in more detail, is that solutions of the limiting equations may be used to construct approximate equilibria for large-population games, for which existence of equilibria can be hard to prove directly.

Let us now describe the limiting MFG system on an intuitive level. If the number of agents n is large, then a single representative agent has little influence on the empirical measure flow $(\hat{\mu}_t^n)_{t \in [0, T]}$, and this agent expects to lose little in the way of optimality by ignoring her own effect on the empirical measure. If there were a continuum of agents, then each agents' influence on this empirical measure would be null, and the optimization problems of the agents would be decoupled and identical. This line of reasoning leads to the following mean field notion of equilibrium:

For a fixed (deterministic) measure flow $\mu = (\mu_t)_{t \in [0, T]}$, consider the following stochastic optimal control problem:

$$(\mathbf{P}_\mu) \quad \begin{cases} \sup_\alpha \mathbb{E} \left[\int_0^T f(X_t^{\mu, \alpha}, \mu_t, \alpha_t) dt + g(X_T^{\mu, \alpha}, \mu_T) \right], \\ \text{s.t.} \quad dX_t^{\mu, \alpha} = b(X_t^{\mu, \alpha}, \mu_t, \alpha_t) dt + \sigma(X_t^{\mu, \alpha}, \mu_t, \alpha_t) dW_t. \end{cases}$$

¹This approach through HJB equations is preferred for closed-loop equilibria, while a Pontryagin-type maximum principle is more appropriate for open-loop. The latter approach leads to an equally intractable n -dimensional system of forward-backward SDEs. This thesis works exclusively with open-loop equilibria.

A deterministic measure-valued function $t \mapsto \mu_t$ is called an *equilibrium* or a *MFG solution* if $\mu_t = \text{Law}(X_t^{\mu, \alpha^*})$ for each $t \in [0, T]$, for some control α^* which is optimal for the problem (\mathbf{P}_μ) .

Intuitively, the state process of (\mathbf{P}_μ) is that of a single representative agent, and μ_t represents the distribution of an infinity of agents' state processes. The representative agent cannot influence μ_t and thus considers it as fixed when solving the optimization problem. If each agent among the infinity is identical and acts in the same way, then the law of large numbers suggests that the statistical distribution of the representative's optimally controlled state process at time t must agree with μ_t . For this reason, the equation $\mu_t = \text{Law}(X_t^{\mu, \alpha^*})$ is often called the *consistency condition*.

A somewhat more mathematical heuristic argument is as follows. Assume that we are given for each n a Nash equilibrium $(\alpha^1, \dots, \alpha^n)$ for the n -player game, and assume also that there exists a single deterministic function $\hat{\alpha}$, independent of n and i , such that $\alpha_t^i = \hat{\alpha}(t, X_t^i)$ for each $1 \leq i \leq n$. (This is a huge assumption, but the symmetry of the system and the weakness of the coupling for large n suggest that it may not be far from reasonable.) If $\hat{\alpha}$ is sufficiently well-behaved, then the state process empirical measure should converge to a McKean-Vlasov limit (as discussed in Section 1.1.1), i.e. $\hat{\mu}^n \rightarrow \mu$ for some deterministic measure flow μ . The state process X^1 of agent 1, controlled by α^1 , should also converge as $n \rightarrow \infty$ to the solution $X^{\mu, \alpha}$ of the state equation of the mean field problem (\mathbf{P}_μ) with $\alpha_t = \hat{\alpha}(t, X_t^{\mu, \alpha})$. On the other hand, suppose agent 1 in the n -player game chooses to use an alternative control β , while the other agents stick with α^i ; then as $n \rightarrow \infty$ the corresponding empirical measure should be close to the original $\hat{\mu}^n$, since only one player has changed strategy, and thus this new empirical measure should converge to the same μ . Similarly, the new state process of agent 1 (controlled by β) should then converge to the state process $X^{\mu, \beta}$ of the mean field problem (\mathbf{P}_μ) . The Nash equilibrium assumption on $(\alpha^1, \dots, \alpha^n)$ provides an inequality which, when $n \rightarrow \infty$, implies that α is superior to β in the limiting control problem (\mathbf{P}_μ) . Since β was arbitrary, this yields the optimality condition of (\mathbf{P}_μ) . Making this argument rigorous turns out to be a highly nontrivial task, and the interaction between the optimization and the $n \rightarrow \infty$ limit is subtle.

1.1.4 Convergence to the mean field game limit

The first question raised by such an optimistic, informal derivation of the MFG system is, of course: Does the MFG system actually describe the limit as $n \rightarrow \infty$, in some rigorous sense? In the literature, this is most commonly answered by using a solution of the MFG to construct ϵ_n -Nash equilibria for the n -player games, where $\epsilon_n \rightarrow 0$. More specifically, suppose μ is a MFG solution, and the corresponding optimal control may be written in feedback form $\alpha_t^* = \hat{\alpha}(t, X_t^{\mu, \alpha^*})$, for some nice function $\hat{\alpha}$. Then, if each agent i uses the control $\alpha_t^i = \hat{\alpha}(t, X_t^i)$, then the heuristic argument of the previous paragraph can be adapted to prove rigorously that we have an ϵ_n -Nash equilibrium for some $\epsilon_n \rightarrow 0$. Following [67, 32], most of the probabilistic work on MFGs adopts this strategy, it is by now indisputable that MFG solutions are useful in constructing approximate equilibria for n -player games.

Little is known, however, regarding the opposite and arguably more direct convergence problem, and a thorough study of this problem is one of the main contributions of this thesis. Namely, if we are given for each n an approximate Nash equilibrium, then what

can be said about the $n \rightarrow \infty$ limit? Thinking cautiously that the MFG solutions may not be unique, one might guess from the McKean-Vlasov theory that limits of the empirical measures exist and are concentrated on the set of MFG equilibria. We will see, however, that this is not the case, and the full story of MFG limits is more subtle. A genuinely *stochastic* notion of equilibrium is required for a full description of the limits of n -player equilibria, and *these stochastic equilibria are not necessarily just randomizations among the family of deterministic equilibria*. As a consequence, the (deterministic) solution concept considered thus far in the literature on mean field games does not fully capture the limiting dynamics of n -player equilibria. This thesis studies this point in some detail, proving some admittedly difficult-to-apply results which nevertheless shed some light on this phenomenon: The fundamental obstruction is the adaptedness required of controls, which renders the class of admissible controls quite sensitive to whether or not $(\mu_t)_{t \in [0, T]}$ is stochastic. In short, a stochastic equilibrium (or weak MFG solution) requires that the stochastic measure flow μ is independent of the noise W , and the consistency condition reads $\mu_t = \text{Law}(X_t^{\mu, \alpha^*} \mid \mathcal{F}_t^\mu)$, where \mathcal{F}_t^μ is the filtration generated by μ .

The early work of Lasry and Lions [91, 89] first attacked the direct convergence problem rigorously using PDE methods, working with an infinite time horizon and strong simplifying assumptions on the data, and their results were later generalized by Feleqi [49]. Bardi and Priuli [7, 8] justified the MFG limit for certain linear-quadratic problems, and Gomes et al. [54] studied models with finite state space. Substantial progress was made in a recent paper of Fischer [51], which deserves special mention also because both the level of generality and the method of proof are quite similar to ours; we will return to this point shortly.

With the exception of [51], the aforementioned results share the important limitation that the agents have only *partial information*: the control of agent i may depend only on her own state process X^i or Wiener process W^i . The results of this thesis allow for arbitrary full-information strategies, partially resolving a conjecture of Lasry and Lions (stated in Remark x after [91, Theorem 2.3] for the case of infinite time horizon and closed-loop controls). Combined in [91, 89, 49] with the assumption that the state process coefficients (b, σ) do not depend on the empirical measure, the assumption of partial information leads to the immensely useful simplification that the state processes of the n -player games are independent.

Fischer [51], on the other hand, allows for full-information controls but characterizes only the *deterministic* limits of $(\hat{\mu}_t^n)_{t \in [0, T]}$ as MFG equilibria. Assuming that the limit is deterministic implicitly restricts the class of n -player equilibria in question. By characterizing even the stochastic limits of $(\hat{\mu}_t^n)_{t \in [0, T]}$, which we show are in fact quite typical, we impose no such restriction on the equilibrium strategies of the n -player games. This not to say, however, that our results completely subsume those of [51], which work with a more flexible notion of *local* approximate equilibria and which notably include conditions under which the assumption of a deterministic limit can be verified.

The proof of our main limit theorem works by studying the full joint distribution of those processes $(\hat{\mu}^n, W^U, \alpha^U, X^U)$ directly relevant to a representative agent U , with U chosen uniformly at random from $\{1, \dots, n\}$. Randomly selecting the representative agent injects some important symmetry, since equilibrium controls are non necessarily symmetric (see Section 2.3.3 for some discussion of this point). Deriving the correct limiting state process dynamics is fairly routine once adequate estimates on the state processes are established,

and the needed moment bounds on the control processes come from a crucial coercivity assumption on the objective functions. A key technical difficulties is specifying the right class of admissible controls in the limit, and this leads to a notion we call *compatibility*; this notion is introduced in Section 3.1.

1.1.5 Existence

Again, the more common way to justify the MFG system is by using its solution to construct approximate equilibria for the n -player games. When this is possible, the next natural question is how to solve the MFG. This has been done primarily in one of two ways, using either PDEs or FBSDEs.

Numerous conditions are known under which the value function $v(t, x)$ of the optimal control problem (\mathbf{P}_μ) can be expressed as a (viscosity) solution of a Hamilton-Jacobi-Bellman equation, of the form

$$\begin{cases} -\partial_t v(t, x) - \sup_a [L_{\mu_t}^a v(t, x) + f(x, \mu, a)] = 0, & \text{on } (0, T) \times \mathbb{R}^d, \\ v(T, x) = g(x, \mu_T). \end{cases} \quad (1.6)$$

Here we define the generator L_m^a on smooth test functions (again assuming the dimension of the state process is one) by

$$L_m^a \varphi(t, x) = b(x, m, a) \partial_x \varphi(t, x) + \frac{1}{2} \sigma(x, m, a)^2 \partial_x^2 \varphi(t, x).$$

On the other hand, if the optimal control is of the form $\alpha_t^* = \hat{\alpha}(t, X_t^{\mu, \alpha^*})$ for a nice function $\hat{\alpha}$, the fixed point condition $\mu_t = \text{Law}(X_t^{\mu, \alpha^*})$ implies (written assuming μ_t has a smooth density) that the Kolmogorov forward equation for μ is

$$\begin{cases} \partial_t \mu_t(x) &= -\partial_x (b(x, m, \hat{\alpha}(t, x)) \mu_t(x)) + \frac{1}{2} \partial_x^2 (\sigma(x, m, \hat{\alpha}(t, x))^2 \mu_t(x)), \\ \mu_0 &= \text{Law}(X_0). \end{cases}$$

Coupling this Kolmogorov equation with the HJB equation summarizes the important features of the MFG problem. This PDE system has been the subject of much analysis, beginning with the work of Lasry and Lions [91] and surveyed in [24, 55], with the chief difficulty stemming from the opposing directions of time in the two equations. On the one hand, PDE techniques typically require strong simplifying assumptions on the structure of the data; for example, typically $b(x, \mu, a) = a$, σ is constant, and f is of the form $f(x, \mu, a) = f_1(x, a) + f_2(x, \mu)$. On the other hand, PDE methods are powerful in their ability to handle *local* mean field interactions, meaning that coefficients depend on (x, μ) through the density $d\mu/dx(x)$. In this thesis, we only consider *nonlocal smoothing* interactions, meaning the dependence of the coefficients on the measure argument is continuous with respect weak convergence, or more generally a Wasserstein metric.

An alternative approach, pioneered by Carmona and Delarue [32], uses the stochastic (Pontryagin) maximum principle to reduce the MFG problem to a forward-backward SDE

of McKean-Vlasov type as follows: Let

$$H(x, y, m, a) = b(x, m, a)y + f(x, m, a),$$

and suppose $\hat{\alpha}(x, y, m)$ maximizes $H(x, y, m, a)$ over a , for each x, y, m . Fix a measure flow $\mu = (\mu_t)_{t \in [0, T]}$, and suppose (X, Y, Z) solves the FBSDE

$$\begin{cases} dX_t &= b(X_t, Y_t, \hat{\alpha}(X_t, Y_t, \mu_t))dt + \sigma dW_t, \\ dY_t &= -\partial_x H(X_t, Y_t, \mu_t, \hat{\alpha}(X_t, Y_t, \mu_t))dt + Z_t dW_t, \\ X_0 &= \xi, \quad Y_T = \partial_x g(X_T, \mu_T) \end{cases} \quad (1.7)$$

According to the well known maximum principle (see e.g. [98, Theorem 6.4.6]), under appropriate differentiability and convexity assumptions, the control $\hat{\alpha}(X_t, Y_t, \mu_t)$ is optimal for the problem (\mathbf{P}_μ) . Thus, to solve the MFG, it suffices to find μ and a solution (X, Y, Z) to the above FBSDE such that $\mu_t = \text{Law}(X_t)$ for all t .

In a select few cases, the PDE and FBSDE systems above (and thus the MFG problem) can be solved fairly explicitly. While several examples are provided in [61], very few stochastic control problems admit explicit solutions beyond simple linear-quadratic models. A linear-quadratic control problem is one in which b and σ are affine in the state and control arguments (x, a) , and the objectives f and g are quadratic in (x, a) . A linear-quadratic mean field game typically involves an affine function of the mean $\int y\mu(dy)$ in the state dynamics and a quadratic function of the same term in the objectives. Generally, linear-quadratic control problems and MFGs can be reduced to certain Riccati differential equations [15, 34], which occasionally themselves admit explicit solutions [35]. Beyond the linear-quadratic case, explicit solutions of MFGs are typically unavailable, and it is nontrivial to prove that a solution exists. Both the PDE and FBSDE have been the basis for well-posedness studies, and in both cases a key difficulty comes from the forward-backward nature of the problem. When the time horizon T is small, a contraction argument yields existence and uniqueness of solutions under standard Lipschitz assumptions [67], but a restriction on the time horizon is often unsatisfactory. Most existence proofs instead turn on Schauder's fixed point theorem, establishing the requisite compactness and continuity via a priori estimates on the PDE or FBSDE system.

This thesis takes a very different approach to proving existence theorems by working with *relaxed controls*. Rather than characterizing solutions of the control problems (\mathbf{P}_μ) via solutions of PDEs or FBSDEs, a more functional-analytic framework is employed in which admissible controls are described by joint distributions of state-control pairs, i.e. (X, α) . This framework was studied heavily until a couple of decades ago, and it is quite convenient for proving the existence of optimal (Markovian) controls [75, 63] and for weak convergence arguments, such as those arising in numerical methods [85]. For mean field games, there are several advantages to adopting this approach. First, by moving to a compactification of the control space, there is no need for precise analysis of the optimal feedback control. This is often the crux of the PDE and FBSDE approaches, which require strong convexity assumptions to ensure that the continuity of the fixed point map is not lost because of a poorly behaved control. Second, the arguments can be carried out under much less restrictive assumptions on the data (b, σ, f, g) . To the reader familiar with the existence theory for

stochastic optimal controls, it should come at no surprise that this relaxed control approach simultaneously affords both greater generality and a simpler proof. Indeed, the PDE and FBSDE methods are the preferred methods for *constructing* optimal controls and yield much more explicit information about the solution than the more abstract approach of relaxed controls. However, in the study of existence for mean field games, this explicit information is almost always immediately surrendered by an application of an abstract fixed point theorem, usually Schauder's. Hence, purely for the sake of proving existence of equilibria, the use of relaxed controls seems more natural in that it is abstract from the start.

1.1.6 Uniqueness

Uniqueness is much harder to come by than existence in game theory in general, and mean field games are no exception. Existence results vastly outnumber uniqueness results in the literature, and uniqueness generally should not be expected. When the time horizon is small, or alternatively when the product of certain Lipschitz constants is small, contraction arguments are available [67]. Beyond such restrictive assumptions, Lasry and Lions [91] discovered a versatile uniqueness criterion, often called the *monotonicity condition*. More recently, Ahuja [3] showed by FBSDE methods that a weaker monotonicity condition still yields uniqueness for a particular class of MFGs. Finally, various authors have proven uniqueness by specialized PDE methods for certain classes of MFGs with local interactions [25, 26, 57].

The only novelty in the uniqueness results of this thesis comes from the notion of solution. Indeed, the precise uniqueness results rely on the Lasry-Lions monotonicity condition. But this novelty is an important one; the new notion of *weak solution* is developed in this thesis precisely because it captures the limiting behavior of the n -player approximate equilibria completely. Uniqueness of weak solutions thus translates to a unique limit for all n -player approximate equilibria. We will see that the notion of solution in previous literature, which we call a *strong solution*, does not necessarily describe all of the limits of n -player games, and thus a uniqueness result for strong solutions carries no information about uniqueness of limits of n -player games.

1.2 Common noise

One important feature of the interacting diffusion model (1.1) is the independence of the driving noises. This is unrealistic in many applications, especially in economics and finance where *aggregate shocks* should not be ruled out. A model of a large financial market should include, for example, a common set of assets to which all of the agents have access. Mean field games with common noise are naturally more complicated and as such have appeared very little in the literature so far. This thesis seeks to remedy this, providing some first existence and uniqueness results for such MFGs. First, let us explore the corresponding particle systems to get a sense of what changes in the presence of common noise.

1.2.1 Interacting diffusion models

The diffusion model (1.1) can be extended to incorporate common noise, most easily by including an additional independent Brownian motion B and altering the dynamics to

$$\begin{cases} dX_t^i &= b(X_t^i, \hat{\mu}_t^n)dt + \sigma(X_t^i, \hat{\mu}_t^n)dW_t^i + \sigma_0(X_t^i, \hat{\mu}_t^n)dB_t, \\ \hat{\mu}_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}. \end{cases} \quad (1.8)$$

While these systems have not been studied quite as thoroughly as the one without common noise, it is still fairly well understood how to describe the limit $n \rightarrow \infty$. Because of the additional correlations entering the system through the common noise B , one should no longer expect the limiting measure flow to be deterministic, even if it is unique.

The heuristic arguments of Section 1.1.2 adapt naturally to deriving the limiting dynamics of $(\hat{\mu}_t^n)_{t \in [0, T]}$. By again applying Itô's formula to $\varphi(X_t^i)$, for smooth φ , we get an extra term compared to (1.4):

$$\begin{aligned} d \int \varphi d\hat{\mu}_t^n &= \left(\int L_{\hat{\mu}_t^n} \varphi d\hat{\mu}_t^n \right) dt + \frac{1}{n} \sum_{i=1}^n \varphi'(X_t^i) \sigma(x, \hat{\mu}_t^n) dW_t^i \\ &\quad + \left(\int \varphi'(x) \sigma(x, \hat{\mu}_t^n) \hat{\mu}_t^n(dx) \right) dB_t. \end{aligned}$$

The dW term should again vanish in the limit, but now the dB term remains. We arrive at a weak form of the stochastic PDE (SPDE)

$$d\mu_t(x) = L_{\mu_t}^* \mu_t(x) dt - \partial_x (\sigma(x, \mu_t) \mu_t(x)) dB_t \quad (1.9)$$

This is essentially a stochastic form of a Kolmogorov forward equation, and indeed this suggests that μ_t should coincide with the *conditional* law of X_t given B , when $(X_t)_{t \in [0, T]}$ solves the SDE

$$dX_t = b(X_t, \mu_t)dt + \sigma(X_t, \mu_t)dW_t + \sigma_0(X_t, \mu_t)dB_t. \quad (1.10)$$

This is essentially the McKean-Vlasov SDE of (1.2), except that the law appearing in the coefficients is now conditional. Of course, it is a highly nontrivial effort to turn these brief heuristics into a rigorous derivation, but this has been done. Under Lipschitz assumptions on the data, Kurtz and Xiong prove a trajectorial propagation of chaos result in [83]; they show that the descriptions (1.9) and (1.10) are equivalent and characterize the unique limit of $(\hat{\mu}_t^n)_{t \in [0, T]}$.

In the non-unique regime, the literature is much thinner, but some results were obtained by Dawson and Vaillancourt [40]. The SPDE (1.9) is replaced by a martingale problem on the space of probability measures, shown to correspond to a *weak solution* of the SPDE. Here the word “weak” is interpreted not only in the distributional sense but also in the probabilistic sense; the SPDE must be integrated against test functions, and also the solution $(\mu_t)_{t \in [0, T]}$ is *not* required to be adapted to the filtration generated by the driving noise B . Unfortunately, no alternative formulation in the spirit of (1.10) is provided. This gap will be filled in Section

2.2, as a useful first step toward mean field games with common noise. Indeed, rather than requiring μ_t equal the conditional law of X_t given B , the right notion of weak solution requires that we condition on (B, μ) , and again remove the restriction that μ be B -measurable.

1.2.2 Mean field games

When a common noise is introduced into the stochastic differential game system described in Section 1.1.3, one should now expect that the correct MFG limit involves conditional measure flow. To be clear, let us now consider a stochastic differential games with n players, with state processes given by X^1, \dots, X^n satisfying

$$\begin{cases} dX_t^i &= b(X_t^i, \hat{\mu}_t^n, \alpha_t^i)dt + \sigma(X_t^i, \hat{\mu}_t^n, \alpha_t^i)dW_t^i + \sigma_0(X_t^i, \hat{\mu}_t^n, \alpha_t^i)dB_t, \\ \hat{\mu}_t^n &= \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}. \end{cases} \quad (1.11)$$

The description of the system and the objective functions is exactly as in Section 1.1.3, with the only exception of the dB_t term in the dynamics (1.11). The Wiener process B is independent of W^1, \dots, W^n .

Combining our intuition from the McKean-Vlasov setting with common noise and the MFG without common noise leads to the following candidate MFG problem. For a fixed *stochastic* measure flow $\mu = (\mu_t)_{t \in [0, T]}$, consider the following stochastic optimal control problem:

$$(\mathbf{P}'_\mu) \quad \begin{cases} \sup_\alpha \mathbb{E} \left[\int_0^T f(X_t^{\mu, \alpha}, \mu_t, \alpha_t) dt + g(X_T^{\mu, \alpha}, \mu_T) \right], \\ \text{s.t.} \quad dX_t^{\mu, \alpha} = b(X_t^{\mu, \alpha}, \mu_t, \alpha_t)dt + \sigma(X_t^{\mu, \alpha}, \mu_t, \alpha_t)dW_t + \sigma_0(X_t^{\mu, \alpha}, \mu_t, \alpha_t)dB_t. \end{cases}$$

A stochastic measure flow $(\mu_t)_{t \in [0, T]}$ is called an *equilibrium* or a *MFG solution* if $\mu_t = \text{Law}(X_t^{\mu, \alpha^*} \mid \mathcal{F}_t^B)$ for each $t \in [0, T]$, for some control α^* which is optimal for the problem (\mathbf{P}'_μ) . Here $\mathcal{F}_t^B = \sigma(B_s : s \leq t)$ denotes the filtration generated by the common noise.

Although we understand heuristically how to formulate the common noise problem, the literature on such problems is very thin. The papers [61, 35] contain specific common noise models which can be solved explicitly, and the latter paper provides the first and only rigorous justification of this equilibrium concept, i.e. the only limit theorem, by a direct calculation. Otherwise, Ahuja [3] provided the first and only somewhat general existence results, for an essentially linear-quadratic model with common noise but with non-quadratic terminal objective g .

In theory, both the PDE and FBSDE approaches described in Section 1.1.5 can adapt to this setting, but both become far more complicated. The FBSDE system described by (1.7) adapts in the obvious way once the dB_t term is included; the only difference is that the consistency condition now requires $\mu_t = \text{Law}(X_t \mid \mathcal{F}_t^B)$ instead of simply $\mu_t = \text{Law}(X_t)$. The PDE system, on the other hand, is more sensitive. The control problem (\mathbf{P}'_μ) is not Markovian when μ is stochastic, and thus a standard HJB equation is not available. Using Peng's ideas [97], we can derive in its place a *stochastic* HJB equation, which now reads as a backward SPDE. Similarly, there is a Kolmogorov forward equation for the conditional laws $(\mu_t)_{t \in [0, T]}$, but it is now a (forward) stochastic PDE as in (1.9). This forward-backward SPDE

system is difficult to work with, but some interesting ideas appear here nonetheless. One may guess, as in the theory of standard FBSDEs, that the backward part of the solution (the value function) can be written as a deterministic function of time and of the forward part (the measure flow). If so, this deterministic function (the arguments of which are time, a state, and a measure) should itself solve a certain deterministic PDE, which is called the *master equation*. This approach is described thoroughly in [33, 14, 39], and verification theorems are available, but so far it is of limited value in establishing existence theorems.

The randomness of the equilibrium measure flow $(\mu_t)_{t \in [0, T]}$ significantly complicates the analysis of common noise problems. Indeed, existence theorems must study a fixed point problem in the much larger space of *stochastic* measure-valued paths. The compactness issue is resolved by again formulating the problem in a weak sense, in terms of joint laws of the relevant processes (B, W, μ, α, X) , and by working with relaxed controls. A more stubborn problem is posed by the rather discontinuous operation of *conditioning* required in the fixed point procedure. Indeed, if random variables (Z_n, Y) converge in law (with Y independent of n), there is generally no useful sense in which we can say that the conditional law of Z_n given Y converges. However, if the support of Y is finite, then a conditional law given Y is merely a finite vector of (deterministic) probability measures, and it will converge weakly in a natural and useful sense. With this in mind, we approximate the common noise B by a sequence of random walks with finite support and solve a discretized form of the MFG; by then refining the discretization and taking weak limits, we prove existence for the original problem.

Both the existence theory and, of course, the limit theory hinge on weak convergence arguments. In fact, our proof of the convergence of n -player equilibria is not complicated much by the common noise. However, these weak convergence arguments make it much more clear why a notion of *weak solution* arises naturally. For the common noise problem, a weak solution requires only that (B, μ) is independent of W and that $\mu_t = \text{Law}(X_t^{\mu, \alpha^*} \mid \mathcal{F}_t^{B, \mu})$, where the filtration $(\mathcal{F}_t^{B, \mu})_{t \in [0, T]}$ is generated by B and μ . Note that μ is not required to be B -measurable, and if it is we call it a *strong solution*. As is well known in the theory of SDEs, taking weak limits (even of strong solutions) yields weak solutions, in which the solution process is not necessarily adapted to the driving Wiener process; pathwise uniqueness is usually needed in order to prove the limit remains a strong solution. We will develop a theory of strong and weak solutions of MFGs with common noise in exact analogy with the SDE theory, complete with natural notions of pathwise uniqueness and uniqueness in law, as well as an analog of the famous theorem of Yamada and Watanabe [109].

1.3 Extensions of the mean field game framework

This section summarizes several interesting extensions to the mean field game frameworks described above. This thesis will not discuss any of the problems discussed here, instead focusing on deepening the theory of the simple setting first. However, the methods of this thesis seem versatile enough to apply to any of these more general models, and hopefully future work will explore these possibilities.

A first simple extension of the basic MFG setup is to allow for different *types* of agents, as a way of introducing more heterogeneity among agents. One simple way to do this is

by drawing each agent's data (b, σ, f, g) in an i.i.d. fashion at time zero. More precisely, in the n -player game, introduce auxiliary i.i.d. random factors $(\zeta^i)_{i=1}^n$, measurable at time zero; then, for example, assign agent i the drift $b(X_t^i, \hat{\mu}_t^n, \alpha_t^i, \zeta^i)$. To illustrate what this can accomplish, consider a simple example in which ζ^i is 1 with probability $p \in (0, 1)$ or 0 with probability $1 - p$, and set $g(X_T^i, \hat{\mu}_T^n, \zeta^i) = -\zeta^i |X_T^i|^2$ (recall that the agents are maximizers, not minimizers). This means that a fraction of p of the agents are penalized for having a nonzero terminal state X_T^i , while the remaining fraction $1 - p$ have no such penalty. This is not the most common way of incorporating types in MFGs (see e.g. [67]), but it seems the simplest to which to adapt the arguments of this thesis. Incorporating types into this thesis in this way would complicate nothing but the notation.

Another class of MFG models, called *major-minor player* models, supplements the field of small (minor) agents with an additional *major* player. For example, the states X^1, \dots, X^n of the minor agents may evolve according to

$$dX_t^i = b(X_t^i, X_t^0, \hat{\mu}_t^n, \alpha_t^i)dt + \sigma dW_t^i,$$

where again $\hat{\mu}_t^n = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}$ is the empirical distribution of the minor agents' states, while the state X^0 of the major player follows

$$dX_t^0 = b_0(X_t^0, \hat{\mu}_t^n, \alpha_t^0)dt + \sigma_0 dW_t^0.$$

Objective functions should of course be modified similarly to reflect this change. The corresponding mean field game is now intuitively much like a two-player game, in which the major player competes with the mean field of minor players. The curious reader is referred to [94, 95, 11, 37] for more details. It seems natural to cover this model and the common noise models treated in this thesis under the same blanket, by studying this major-minor player model with a common dW_t^0 term added to the dynamics of the minor players, but this is beyond the scope of this thesis.

Another natural extension of this framework is to allow mean field interactions based on the *controls*, and not just the states. That is, one might replace the empirical measure $\hat{\mu}_t^n$ with the joint empirical measure of states and controls, $\frac{1}{n} \sum_{k=1}^n \delta_{(X_t^k, \alpha_t^k)}$. This was studied by the author and Carmona in [36], pointing out that many models arising in the price impact and optimal execution literature naturally incorporate this form of interaction. This type of MFG has also been called an *extended mean field games*; see [56] for a PDE approach.

Finally, ideas resembling MFG theory are discernible in several other recent extensions of classical mean field models. For example, the recent information percolation model of Duffie et al. [45] is in some ways a MFG form of the classical Smoluchowski coagulation equation, which suggests a corresponding n -player game based on the Marcus-Lushnikov model of stochastic coalescence; see [4] for a good survey of the classical mean field theory. Additionally, MFG-like analogs of models of spin systems (e.g., the Ising model) have begun to appear in economics [65, 64], albeit under different names.

1.4 Continuum games in economics

We have seen how mean field games incorporate competition into particle systems, explaining the choice terminology. However, the idea of a continuum of agents is old news to economists, and such models have been immensely popular for over half a century. The work of Aumann [6] in 1964 and Schmeidler [101] in 1973 seem to have popularized continuum-agent models in cooperative and competitive settings, respectively. This section aims to clarify what about *mean field game theory* is novel, apart from its name.

First of all, the economic theory of competitive games with a continuum of agents is dominated by *static* games. The early results of Schmeidler [101] and Mas-Colell [92] laid the foundations and proved some general existence results, and countless subsequent papers studied extensions, variations, and different aspects of these models. Discrete-time dynamics were incorporated in later theoretical work, e.g. [71, 16, 17], although plenty of specific models had appeared previously which involved dynamic games with a continuum of agents. The latter papers of Bergin and Bernhardt [16, 17] notably study models with common noise, which economists refer to as *aggregate uncertainty*. Compared to this line of literature, one novelty of mean field games is that they are typically set in continuous time, which leads to interesting new PDE and SDE systems. However, the term *mean field game*, like *anonymous games* in economics, is general enough that we should avoid tying it to a particular time set; *differential* mean field games will refer specifically to continuous-time models.

A more interesting novelty of modern mean field game theory is its attention to finite-population models. Many economics papers work exclusively with a continuum of agents, justified as a sort of idealization or limit of some corresponding model with a large but finite population. By far the most popular way to justify continuum models in economics is by means of an “exact law of large numbers,” which directly models a continuum of independent agents. The intuitive idea of this approach is that if a continuum of independent coins are flipped, then *exactly* half of them will be heads. To make mathematical sense of this, however, runs into well known measure-theoretic difficulties. Substantial work has gone into making the exact law of large numbers rigorous; see [105, 103] and the references therein.

Scarcely in economics is any attempt made to rigorously expose a continuum model as some kind of limit of finite models. Some of the few papers include [59, 66] as well as a number of papers of G. Carmona and collaborators [29, 9, 31] (not to be confused with R. Carmona, the thesis advisor and author of more recent papers on mean field games). Most if not all such results seem to pertain to *static* games, with no time dynamics. Some of these papers study the limits of sequences of approximate Nash equilibria as the number of agents tend to infinity, while others [100, 30] show how continuum equilibria give rise to approximate equilibria for finite games, in the spirit of much of the mean field game research.

Interestingly, some of the economic literature seems quite aware of a key phenomenon this thesis takes care to elucidate. Namely, even without common noise, there often exist stochastic MFG equilibria which are qualitatively quite different from the deterministic equilibria. Both the zero-intelligence McKean-Vlasov models and mean field static games discussed in Sections 2.2 and 2.3 respectively exhibit a useful simplifying feature: While there may exist stochastic equilibria, they are all mixtures (i.e., randomizations) of the deterministic equilibria. This breaks down in the MFG setting because of the effect of the *aggregate uncertainty* (i.e., stochasticity of the measure flow) on the information available

to agents; it is only the *combination* of dynamics and competition that produces this effect. This point was observed, for example, in the seminal paper of Diamond and Dybvig [44], which analyzes a two-period model of bank runs. Green and Lin [60] studied the stochastic equilibria in more detail, and the papers of Bergin and Bernhardt [16, 17] explored this issue for more general continuum games.

1.5 Outline of the thesis

This thesis consists of eight chapters and two appendices.

Chapter 2 gathers some preliminary material on weak convergence, relaxed controls, McKean-Vlasov equations, and static mean field games. Many but not all of the weak convergence results compiled in the first Section 2.1 are well known, and they are the bread and butter of many of the proofs of this thesis. Proofs are given for those results which are new (or at least for which references are unknown to the author). Strictly speaking, Section 2.2 on McKean-Vlasov limits is not a prerequisite for this thesis, but this known material will shape our intuition for the analogous MFG theory. We do sketch here a novel proof of a known result, because it seems to adapt better to the MFG setting than previously known proofs. Finally, Section 2.3 discusses *static* mean field games, i.e. with no time component, and derives some simple limit theorems; while the presentation is mostly original, the results most likely are not. In a sense, we isolate separately the *dynamic* and *competitive* aspects of mean field theory in these Sections 2.2 and 2.3, respectively, before turning to stochastic differential mean field games, which combine the two.

Chapter 3 begins the discussion of stochastic differential mean field games with common noise by defining the central equilibrium concepts and stating the main theorems on convergence, existence, and uniqueness. The various notions of mean field game solution are carefully defined and discussed first, before turning to the n -player games. The main limit theorem and its converse are stated in Section 3.2, along with some useful corollaries. Finally, the existence and uniqueness theory is summarized in Section 3.3, including statements of the main results and the definitions of pathwise uniqueness and uniqueness in law. This chapter contains almost no proofs, as the proofs of the main theorems are long enough to warrant their own chapters.

Chapter 4 focuses on mean field games *without* common noise, specializing the results of Chapter 3 to this setting. The solution concepts simplify somewhat in this setting, and we elaborate on the implications of the main limit theorem. The gap between weak and strong solutions is discussed in depth here. The new phenomenon is illustrated by an example, in which a MFG without common noise has a stochastic equilibrium (i.e., weak solution), the support of which is disjoint from the set of deterministic equilibria (i.e., strong solutions). An intuitively natural but difficult-to-verify assumption is presented which rules out this possibility. Finally, in Section 4.6, existence and uniqueness results are stated for MFG without common noise.

Chapter 5 begins the road to proofs of the main theorems. The weak solution concept is reformulated on a canonical space, and some first properties are developed. Several technical results are proven here which are useful in both the proof of the limit theorem and the proof

of existence. For example, some first continuity properties are established for certain maps which appear repeatedly in the convergence and existence proofs.

Chapter 6 proves the main limit theorems of Chapters 3 and 4, completing the characterization of weak MFG solutions as the limits of n -player approximate equilibria. Chapter 7 proves the existence and uniqueness results for MFGs with common noise, while Chapter 8 proves existence and uniqueness results in the setting without common noise. The first two of these three, Chapters 6 and 7, work in the framework developed in Chapter 5. Chapter 8, on the other hand, develops a controlled martingale problem framework which is better-suited to studying strong MFG solutions. This more specialized framework is useful in that it accommodates a control-dependent volatility coefficient σ , unlike the framework we adopt for common noise problems.

Appendix A compiles some useful but rather non-standard topics in stochastic analysis. In particular, weak and strong solutions of stochastic differential equations with random coefficients are discussed in some detail. These results are used mostly implicitly throughout the body of the thesis, as they are intuitive and not at all surprising. Nonetheless, this material is worth developing carefully, not only because concise references are difficult to locate: This material serves as a good illustration of two ideas which are also developed in this appendix and which play important roles elsewhere in the thesis: abstract Yamada-Watanabe theorems and compatible extensions of filtrations. Finally, with applications in Chapter 8 in mind, Appendix A closes with a discussion of the stochastic calculus for martingale measures, and some basic results on SDEs and martingale problems driven by them. Lastly, Appendix B proves three tightness results for solutions of stochastic differential equations.

Chapter 2

Preliminaries

This chapter serves first to compile the notation and basic tools that will be used repeatedly in what follows. Weak convergence arguments pervade this thesis, not only in the formalism of the main limit theorem but also in existence proofs, which turn on topological fixed point theorems. The reader is assumed to be familiar with the basics of measure-theoretic probability, though presumably all of the unemphasized background material can be found in the book of Kallenberg [73].

The second goal of this chapter is to explore McKean-Vlasov limits and static MFGs, both as a warm-up and to illustrate some key ideas which will show up in our study of stochastic differential MFGs.

2.1 Spaces of probability measures

All of the topological spaces we will encounter are *Polish spaces*, meaning they are separable and completely metrizable. Given a Polish space E , we let $\mathcal{B}(E)$ denote the Borel sets of E , and let $\mathcal{P}(E)$ denote the set of Borel probability measures on E . We endow $\mathcal{P}(E)$ with the *weak topology*, which is the weakest topology making the maps $\mu \mapsto \int \varphi d\mu$ continuous for each bounded, continuous function $\varphi : E \rightarrow \mathbb{R}$. When E is a Polish space, so is $\mathcal{P}(E)$, and we equip it with its Borel σ -field. Given E -valued random variables X and $(X_n)_{n=1}^\infty$, defined perhaps on different probability spaces, we say that X_n converges in law (or in distribution) to X if the sequence of laws of X_n converges weakly (i.e., in the weak topology) to the law of X .

Several useful facts about weak convergence are outlined in the following theorem:

Theorem 2.1.1 (Portmanteau's theorem). *For probability measures μ_n, μ on a given metric space E , the following are equivalent:*

1. $\mu_n \rightarrow \mu$ weakly.
2. $\limsup_n \int \varphi d\mu_n \leq \int \varphi d\mu$ for all upper semicontinuous φ bounded from above.
3. $\liminf_n \int \varphi d\mu_n \geq \int \varphi d\mu$ for all lower semicontinuous φ bounded from below.

Polish spaces are of particular importance in weak convergence arguments because of *Prokhorov's theorem*, which characterizes compact sets of $\mathcal{P}(E)$. A set $K \subset \mathcal{P}(E)$ is called

tight if for every $\epsilon > 0$ there exists a compact set $S \subset E$ such that $\mu(S) \geq 1 - \epsilon$ for each $\mu \in K$.

Theorem 2.1.2 (Prokhorov's theorem). *Suppose E is a Polish space. A set $K \subset \mathcal{P}(E)$ is relatively compact if and only if it is tight.*

Much of the analysis of this thesis deals with weak convergence on *product spaces*, for which several special results are available. The following result is well known and easy to prove.

Lemma 2.1.3. *Let E and F be Polish spaces. A set $K \subset \mathcal{P}(E \times F)$ is tight if and only if $K_E := \{P(\cdot \times F) : P \in K\} \subset \mathcal{P}(E)$ and $K_F := \{P(E \times \cdot) : P \in K\} \subset \mathcal{P}(F)$ are tight.*

When working with probability measures on a product space, such as $\Omega \times E$, there is a fairly natural alternative to the topology of weak convergence. Namely, the *stable convergence* of a sequence of probabilities P_n on $\Omega \times E$ to P means that $\int \varphi dP_n \rightarrow \int \varphi dP$ for every bounded measurable function $\varphi : \Omega \times E \rightarrow \mathbb{R}$ for which $\varphi(\omega, \cdot)$ is continuous on E for each $\omega \in \Omega$. Here Ω is any measurable space, while E is again a Polish space. We will make no use of this topology, except to note that when Ω too is a Polish space and the Ω -marginals of P_n are fixed, stable convergence and weak convergence of P_n are equivalent. This is stated more precisely in the following theorem, the proof of which can be found in [69, Corollary 2.9]:

Lemma 2.1.4. *Suppose E and F are Polish spaces and $P, P_n \in \mathcal{P}(E \times F)$. Let $\varphi : E \times F \rightarrow \mathbb{R}$ satisfy the following:*

1. $\varphi(\cdot, y)$ is measurable for each $y \in F$.
2. $\varphi(x, \cdot)$ is continuous for each $x \in E$.

If $P_n \rightarrow P$ weakly and $P_n(\cdot \times F) = P(\cdot \times F)$ for all n , then $\int \varphi dP_n \rightarrow \int \varphi dP$.

The following important result is well known, and will be quite useful later.

Proposition 2.1.5. *Suppose E and F are Polish spaces and $\mu \in \mathcal{P}(E)$. Let $S_\mu = \{P \in \mathcal{P}(E \times F) : P(\cdot \times F) = \mu\}$ denote the set of joint laws with first marginal μ .*

1. *A sequence φ_n of measurable functions from E to F converges in μ -measure to φ if and only if $\mu(dx)\delta_{\varphi_n(x)}(dy)$ converges to $\mu(dx)\delta_{\varphi(x)}(dy)$ weakly.*
2. *If μ is nonatomic, then the set*

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

is dense in S_μ .

3. *If μ is nonatomic and F is a convex subset of a locally convex space H (we still assume F is Polish with the induced topology), then the set*

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

is dense in S_μ .

Proof.

1. The first claim is proven in [69, Corollary 3.7], but here is a direct proof using Lemma 2.1.4. One direction is clear, since convergence in probability implies weak convergence. Assume that $\mu(dx)\delta_{\varphi_n(x)}(dy)$ converges to $\mu(dx)\delta_{\varphi(x)}(dy)$ weakly. Then, since $(x, y) \mapsto \varphi(x) - y$ is continuous in y , Lemma 2.1.4 implies

$$\begin{aligned} \int \mu(dx) |\varphi(x) - \varphi_n(x)| \wedge 1 &= \int \mu(dx) \delta_{\varphi_n(x)}(dy) |\varphi(x) - y| \wedge 1 \\ &\rightarrow \int \mu(dx) \delta_{\varphi(x)}(dy) |\varphi(x) - y| \wedge 1 = 0. \end{aligned}$$

2. This is well known, but precise references are hard to find, so for the reader's convenience we provide an adaptation of the proof given in [38, Theorem 2.2.3]. First, it follows from [69, Proposition 2.4] and Lemma 2.1.4 that the topology induced on S_μ by weak convergence admits as a base the system of (open) neighborhoods of the form

$$U = \left\{ P \in S_\mu : \left| \int_{A_i \times F} f_i(y) P(dx, dy) - \int_{A_i \times F} f_i(y) Q(dx, dy) \right| < \epsilon, \forall i = 1, \dots, m \right\},$$

where $m \geq 1$, $Q \in S_\mu$, $\epsilon > 0$, $f_i : F \rightarrow \mathbb{R}$ are bounded continuous functions, and $A_i \subset E$ are Borel sets. In fact, the same set system forms a base even if we require (A_1, \dots, A_m) to form a partition of E . Now fix $P \in S_\mu$. To prove the claim, it suffices to show that for every $m \geq 1$, $\epsilon > 0$, bounded continuous functions $f_i : F \rightarrow \mathbb{R}$, and measurable partition (A_1, \dots, A_m) of E , there exists a measurable map $\varphi : E \rightarrow F$ such that

$$\left| \int_{A_i \times F} f_i(y) P(dx, dy) - \int_{A_i} f_i(\varphi(x)) \mu(dx) \right| < \epsilon, \forall i = 1, \dots, m.$$

As (A_1, \dots, A_m) form a partition, we are free to specify φ separately on each A . Hence, it suffices to show that for every $\epsilon > 0$, every continuous $f : F \rightarrow \mathbb{R}$, and every measurable $A \subset E$ there exists a measurable map $\varphi : A \rightarrow F$ such that

$$\left| \int_{A \times F} f(y) P(dx, dy) - \int_A f(\varphi(x)) \mu(dx) \right| < \epsilon.$$

The set of convex combinations of Dirac measures is dense in $\mathcal{P}(F)$ (see, e.g., [5, Theorem 5.10]), and thus there exist $n \geq 1$, $y_i \in F$, and constants $c_i \geq 0$ such that $\sum_{i=1}^n c_i = P(A \times F)$ and

$$\left| \sum_{i=1}^n c_i f(y_i) - \int_{A \times F} f(y) P(dx, dy) \right| < \epsilon.$$

As μ is nonatomic, we may apply Sierpiński's theorem [5, Theorem 10.52] to find a measurable partition (B_1, \dots, B_n) of A such that $\mu(B_i) = c_i$ for each $i = 1, \dots, n$.

Define $\varphi(x)$ to be y_i for $x \in B_i$, and complete the proof by noting that

$$\int_A f(\varphi(x))\mu(dx) = \sum_{i=1}^n c_i f(y_i).$$

3. In light of the second part, we must only show that any measurable function $\varphi : E \rightarrow F$ can be obtained as the μ -a.s. limit of continuous functions. By a form of Lusin's theorem [19, Theorem 7.1.13], for each $\epsilon > 0$ we may find a compact $K_\epsilon \subset E$ such that $\mu(K_\epsilon^c) \leq \epsilon$. the restriction $\varphi|_{K_\epsilon} : K_\epsilon \rightarrow F$ is continuous. By the Tietze extension theorem, or rather a generalization due to Dugundji [46, Theorem 4.1], we may find a continuous function $\tilde{\varphi}_\epsilon : E \rightarrow F$ such that $\tilde{\varphi}_\epsilon = \varphi$ on K_ϵ and such that the range $\tilde{\varphi}_\epsilon(E)$ is contained in the convex hull of $\varphi|_{K_\epsilon}(E)$, which is itself contained in the convex set F . We may thus view $\tilde{\varphi}_\epsilon$ as a continuous function from E to F . Since $\mu(\tilde{\varphi}_\epsilon \neq \varphi) \leq \mu(K_\epsilon^c) \leq \epsilon$, we may find a subsequence of $\tilde{\varphi}_\epsilon$ which converges μ -a.s. to φ .

□

For instance, that the set of joint laws on a product space $E \times F$ which are concentrated on the graph of a measurable (continuous if F is convex) function from E to F is dense in the set of all joint laws on $E \times F$. More useful to us will be the following *adapted* version of the previous result, presented in discrete time. Various continuous-time analogs readily follow via approximation by simple processes, and one such analog, Proposition 2.1.15, will be important in our analysis of mean field games.

Proposition 2.1.6. *Let (Ω, \mathcal{F}, P) be a probability space supporting two stochastic processes $S = (S_1, \dots, S_N)$ and $X = (X_1, \dots, X_N)$ with values in Polish spaces \mathcal{S} and \mathcal{X} , and let $\mathcal{F}_n^S = \sigma(S_1, \dots, S_n)$ denote the filtration generated by S . Assume \mathcal{X} is a convex subset of a locally convex space. Suppose that the law of S_1 is nonatomic and that (X_1, \dots, X_n) is conditionally independent of \mathcal{F}_N^S given \mathcal{F}_n^S , for each n . Then there exist continuous functions $h_k^j : \mathcal{S}^k \rightarrow \mathcal{X}$, for $k \in \{1, \dots, N\}$ and $j \geq 1$, such that*

$$(S, (h_1^j(S_1), h_2^j(S_1, S_2), \dots, h_N^j(S_1, \dots, S_N))) \rightarrow (S, X)$$

in law in the space $\mathcal{S}^N \times \mathcal{X}^N$.

Proof. The proof is an inductive application of Proposition 2.1.5. First, use Proposition 2.1.5 to find a sequence of continuous functions $h_1^j : \mathcal{S} \rightarrow \mathcal{X}$ such that $(S_1, h_1^j(S_1)) \rightarrow (S_1, X_1)$ as $j \rightarrow \infty$, where convergence is in distribution throughout this proof. Let us show that in fact $(S, h_1^j(S_1))$ converges to (S, X_1) . Let $\varphi : \mathcal{S}^N \rightarrow \mathbb{R}$ be bounded and measurable, and let $\psi : \mathcal{X} \rightarrow \mathbb{R}$ be continuous. Noting that S and (S_1, X_1) are conditionally independent given

S_1 , we use Lemma 2.1.4 to get

$$\begin{aligned}
\lim_{j \rightarrow \infty} \mathbb{E}[\varphi(S)\psi(h_1^j(S_1))] &= \lim_{j \rightarrow \infty} \mathbb{E}[\mathbb{E}[\varphi(S)|S_1]\psi(h_1^j(S_1))] \\
&= \mathbb{E}[\mathbb{E}[\varphi(S)|S_1]\psi(X_1)] \\
&= \mathbb{E}[\mathbb{E}[\varphi(S)|S_1]\mathbb{E}[\psi(X_1)|S_1]] \\
&= \mathbb{E}[\mathbb{E}[\varphi(S)\psi(X_1)|S_1]] \\
&= \mathbb{E}[\varphi(S)\psi(X_1)]
\end{aligned}$$

This is enough to show that $(S, h_1^j(S_1)) \rightarrow (S, X_1)$ (see e.g. [47, Proposition 3.4.6(b)]).

We proceed inductively as follows. Abbreviate $S^n := (S_1, \dots, S_n)$ and $X^n := (X_1, \dots, X_n)$ for each $n = 1, \dots, N$, noting $S^N = S$. Suppose we are given $1 \leq n < N$ and continuous functions $g_k^j : \mathcal{S}^k \rightarrow \mathcal{X}$, for $k \in \{1, \dots, n\}$ and $j \geq 1$, satisfying

$$\lim_{j \rightarrow \infty} (S, g_1^j(S^1), \dots, g_n^j(S^n)) = (S, X_1, \dots, X_n). \quad (2.1)$$

We will show that there exist continuous functions $h_k^i : \mathcal{S}^k \rightarrow \mathcal{X}$ for each $k \in \{1, \dots, n+1\}$ and $i \geq 1$ such that

$$\lim_{i \rightarrow \infty} (S, h_1^i(S^1), \dots, h_{n+1}^i(S^{n+1})) = (S, X_1, \dots, X_{n+1}). \quad (2.2)$$

By Proposition 2.1.5 there exists a sequence of continuous functions $\hat{g}^j : (\mathcal{S}^{n+1} \times \mathcal{X}^n) \rightarrow \mathcal{X}$ such that

$$\lim_{j \rightarrow \infty} (S^{n+1}, X_1, \dots, X_n, \hat{g}^j(S^{n+1}, X^n)) = (S^{n+1}, X_1, \dots, X_n, X_{n+1}).$$

Note that S and (S^n, X^n) are conditionally independent given S_1 . Using the same argument as above, it follows that in fact

$$\lim_{j \rightarrow \infty} (S, X_1, \dots, X_n, \hat{g}^j(S^{n+1}, X^n)) = (S, X_1, \dots, X_n, X_{n+1}). \quad (2.3)$$

By continuity of \hat{g}^j , the limit (2.1) implies that, for each j ,

$$\begin{aligned}
&\lim_{i \rightarrow \infty} (S, g_1^i(S^1), \dots, g_n^i(S^n), \hat{g}^j(S^{n+1}, g_1^i(S^1), \dots, g_n^i(S^n))) \\
&= (S, X_1, \dots, X_n, \hat{g}^j(S^{n+1}, X_1, \dots, X_n)).
\end{aligned} \quad (2.4)$$

Combining the two limits (2.3) and (2.4), we may find a subsequence j_i such that

$$\begin{aligned}
&\lim_{i \rightarrow \infty} (S, g_1^{j_i}(S^1), \dots, g_n^{j_i}(S^n), \hat{g}^{j_i}(S^{n+1}, g_1^{j_i}(S^1), \dots, g_n^{j_i}(S^n))) \\
&= (S, X_1, \dots, X_n, X_{n+1}).
\end{aligned}$$

Define $h_k^i := h_k^{j_i}$ for $k = 1, \dots, n$ and $h_{n+1}^i(S^{n+1}) := \hat{g}^{j_i}(S^{n+1}, g_1^{j_i}(S^1), \dots, g_n^{j_i}(S^n))$ to complete the induction. \square

2.1.1 Wasserstein distances

Given a complete separable metric space (E, d) , and a real number $p \geq 1$, define

$$\mathcal{P}^p(E) := \left\{ \mu \in \mathcal{P}(E) : \int_E d(x_0, x)^p \mu(dx) < \infty \text{ for some } x_0 \in E \right\}. \quad (2.5)$$

Using the triangle inequality, it is easily proven that the words “for some” may be replaced above by “for every.” The p -Wasserstein distance is defined for $\mu, \nu \in \mathcal{P}^p(E)$ by

$$\ell_{E,p}(\mu, \nu) := \inf \left\{ \int_{E \times E} d(x, y)^p \gamma(dx, dy) : \gamma \in \mathcal{P}(E \times E) \text{ has marginals } \mu, \nu \right\}^{1/p}. \quad (2.6)$$

That is, the infimum is over all probability measures γ on $E \times E$ with $\gamma(E \times \cdot) = \nu$ and $\gamma(\cdot \times E) = \mu$. It is well known that $(\mathcal{P}^p(E), \ell_{E,p})$ is a complete, separable metric space. Its convergent sequences are summarized in the following proposition.

Proposition 2.1.7 (Theorem 7.12 of [107]). *Let (E, d) be a metric space, and suppose $\mu, \mu_n \in \mathcal{P}^p(E)$. Then the following are equivalent for $p \geq 1$:*

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

$$\lim_{r \rightarrow \infty} \sup_n \int_{\{x: d(x, x_0)^p \geq r\}} \mu_n(dx) d(x, x_0)^p = 0. \quad (2.7)$$

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

In particular, (2) implies that a sequence $(\mu_n)_n \subset \mathcal{P}^p(E)$ is relatively compact if and only if it is tight and satisfies (2.7).

We may think of the weak topology of $\mathcal{P}(E)$ as the space $\mathcal{P}^p(E)$ with $p = 0$, and the Wasserstein topology simply incorporates some additional moment behavior. Several of the results on product spaces of the previous section on product spaces will now be extended to Wasserstein space. The proofs are all completely straightforward with the aid of a simple homeomorphism: Fix $x_0 \in E$, and define $\psi(x) := 1 + d^p(x, x_0)$. For each $\mu \in \mathcal{P}^p(E)$ define a measure $\psi\mu \in \mathcal{P}(E)$ by $(\psi\mu)(B) = \int_B \psi d\mu$ for all $B \in \mathcal{B}(E)$. Then $\mu \mapsto \psi\mu / \int \psi d\mu$ is easily seen to define a homeomorphism from $(\mathcal{P}^p(E), \ell_{E,p})$ to $\mathcal{P}(E)$ with the weak topology. Indeed, the details of the proofs are omitted, as one simply uses this homeomorphism to transfer the corresponding weak convergence results to Wasserstein space.

In the following two lemmas, let (E, d_E) and (F, d_F) be two complete separable metric spaces. Equip $E \times F$ with the metric formed by adding the metrics of E and F , given by $((x_1, x_2), (y_1, y_2)) \mapsto d_1(x_1, y_1) + d_2(x_2, y_2)$, but this choice is inconsequential.

Lemma 2.1.8. *A set $K \subset \mathcal{P}^p(E \times F)$ is relatively compact if and only if $\{P(\cdot \times F) : P \in K\} \subset \mathcal{P}^p(E)$ and $\{P(E \times \cdot) : P \in K\} \subset \mathcal{P}^p(F)$ are relatively compact.*

Lemma 2.1.9. *Let $\varphi : E \times F \rightarrow \mathbb{R}$ satisfy the following:*

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

$$|\varphi(x, y)| \leq c(1 + d_E^p(x, x_0) + d_F^p(y, y_0)), \quad \forall (x, y) \in E \times F.$$

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

Corollary 2.1.10. *Suppose φ satisfies (1) and (2) of Lemma 2.1.9, and instead*

- 3'. There exist $c > 0$, $x_0 \in E$, and $y_0 \in F$ such that

$$\varphi(x, y) \leq c(1 + \rho_E^p(x, x_0) + \rho_F^p(y, y_0)), \quad \forall (x, y) \in E \times F.$$

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

$$\limsup_{n \rightarrow \infty} \int \varphi dP_n \leq \int \varphi dP.$$

Proof. For each $M < 0$, Lemma 2.1.9 implies

$$\int M \vee \varphi dP = \lim_{n \rightarrow \infty} \int M \vee \varphi dP_n \geq \limsup_{n \rightarrow \infty} \int \varphi dP_n.$$

Send $M \downarrow -\infty$ and use the monotone convergence theorem. □

2.1.2 Mean measures and compactness

Fix a complete separable metric space (E, d) and an exponent $p \geq 1$ throughout the section. It is well known [18, Corollary 7.29.1] that for any real-valued measurable function $\varphi : E \rightarrow \mathbb{R}$, not necessarily continuous, the map $\mu \mapsto \int \varphi d\mu$ is Borel measurable on $\mathcal{P}(E)$, at least if the integral is defined to be $+\infty$ when it is otherwise ill-defined. From this it follows easily that the Borel σ -field of $\mathcal{P}^p(E)$ is nothing but the trace of the Borel σ -field of $\mathcal{P}(E)$. In this section, we study the space $\mathcal{P}^p(\mathcal{P}^p(E))$; recall that we implicitly endow $\mathcal{P}^p(E)$ with the p -Wasserstein metric $\ell_{E,p}$.

For $P \in \mathcal{P}(\mathcal{P}(E))$, define the mean measure $mP \in \mathcal{P}(E)$ by

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

The following Proposition 2.1.11 is contained in Proposition 2.2(ii) of Sznitman [104]. We will omit its proof and instead prove a useful extension to Wasserstein space, using a natural adaptation of Sznitman's proof.

Proposition 2.1.11. *A set K in $\mathcal{P}(\mathcal{P}(E))$ is tight if and only if $\{mP : P \in K\}$ is tight.*

Proposition 2.1.12. *A subset K of $\mathcal{P}^p(\mathcal{P}^p(E))$ is relatively compact if and only if $\{mP : P \in K\}$ is relatively compact in $\mathcal{P}^p(E)$ and*

$$\lim_{r \rightarrow \infty} \sup_{P \in K} \int_{\{\mu: \int_E \mu(dx) d(x, x_0)^p > r\}} P(d\mu) \int_E \mu(dx) d(x, x_0)^p = 0, \quad (2.8)$$

for some $x_0 \in E$.

Proof. Suppose first that K is relatively compact. Note that

$$\ell_{E,p}(\mu, \delta_{x_0})^p = \int_E \mu(dx) d(x, x_0)^p, \quad (2.9)$$

and thus the uniform integrability (2.8) holds by Proposition 2.1.7(2). It is straightforward to show that $m : \mathcal{P}^p(\mathcal{P}^p(E)) \rightarrow \mathcal{P}^p(E)$ is continuous; indeed, suppose $P_n \rightarrow P$ in $\mathcal{P}^p(\mathcal{P}^p(E))$, and $\varphi : E \rightarrow \mathbb{R}$ is continuous with $|\varphi(x)| \leq c(1 + d(x, x_0)^p)$ for some $c \geq 0$. Then

$$\left| \int \varphi d\mu \right| \leq c(1 + \ell_{E,p}(\mu, \delta_{x_0})^p),$$

and thus Proposition 2.1.7(3) implies

$$\int \varphi d[mP_n] = \int P_n(d\mu) \int \varphi d\mu \rightarrow \int P(d\mu) \int \varphi d\mu = \int \varphi d[mP].$$

Continuity of m implies that $\{mP : P \in K\}$ is relatively compact.

Conversely, assume $\{mP : P \in K\}$ is relatively compact and (2.8) holds. Because of (2.9), the uniform integrability assumption rewrites as

$$\lim_{r \rightarrow \infty} \sup_{P \in K} \int_{\{\mu: \ell_{E,p}(\mu, \delta_{x_0})^p \geq r\}} P(d\mu) \ell_{E,p}(\mu, \delta_{x_0})^p = 0,$$

so we need only to show that K is tight, in light of Proposition 2.1.7. Now suppose $P_n \in K$, and let $I_n := mP_n$. Define $\psi(x) := 1 + d(x, x_0)^p$. Relative compactness of I_n in $\mathcal{P}^p(E)$ implies that

$$\lim_{r \rightarrow \infty} \sup_n \int_{\{\psi \geq r\}} \psi dI_n = 0.$$

Thus, for each $\epsilon > 0$ there exist $r(\epsilon) > 0$ and a compact set $K_\epsilon \subset E$ such that

$$\sup_n I_n(K_\epsilon^c) \leq \epsilon/2, \quad \sup_n \int_{\{\psi > r(\epsilon)\}} \psi dI_n \leq \epsilon/2.$$

Now fix $\epsilon > 0$, and for each k define

$$C_k = \left\{ \mu \in \mathcal{P}^p(E) : \mu(K_{\epsilon 2^{-k}/k}^c) \leq 1/k, \text{ and } \int_{\{\psi > r(\epsilon 2^{-k}/k)\}} \psi d\mu \leq 1/k \right\}.$$

Markov's inequality implies

$$\begin{aligned}
P_n(C_k^c) &\leq P_n \left\{ \mu : \mu(K_{\epsilon 2^{-k}/k}^c) > 1/k \right\} + P_n \left\{ \mu : \int_{\{\psi > r(\epsilon 2^{-k}/k)\}} \psi d\mu > 1/k \right\} \\
&\leq k I_n(K_{\epsilon 2^{-k}/k}^c) + k \int_{\{\psi > r(\epsilon 2^{-k}/k)\}} \psi dI_n \\
&\leq 2^{-k} \epsilon,
\end{aligned}$$

and thus $P_n(\bigcup_{k \geq 1} C_k^c) \leq \epsilon$. Since $1_{K_\eta^c}$ and $\psi 1_{\psi > \eta}$ are lower semicontinuous on E for each $\eta > 0$, it follows from Portmanteau's theorem that each C_k is closed. Thus $\bigcap_{k \geq 1} C_k$ is compact, and P_n are tight. \square

Corollary 2.1.13. *Suppose $K \subset \mathcal{P}^p(\mathcal{P}^p(E))$ is such that $\{mP : P \in K\} \subset \mathcal{P}(E)$ is tight and*

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

Then K is relatively compact.

Proof. The assumption along with Jensen's inequality implies

$$\sup_{P \in K} \int P(d\mu) \left(\int d(x, x_0)^p \mu(dx) \right)^{p'/p} < \infty.$$

This in turn implies the uniform integrability condition (2.8) of Proposition 2.1.12. \square

2.1.3 Relaxed controls

This section introduces a useful tool for continuous-time control theory (both deterministic and stochastic) known as *relaxed controls*, which are used heavily in this thesis. The essential idea is that the path space for control processes should generally be the space of (equivalence classes of Lebesgue-a.e. equal) measurable functions from an interval $[0, T]$ to a Polish space A . To use weak convergence arguments we would like to endow this path space with a Polish topology, with the obvious choice being the topology of convergence in Lebesgue measure. However, this topology is far too strong for compactness purposes, and the space of relaxed controls is essentially a completion of this path space under a weaker (metric) topology.

Let A be a Polish space and $p \geq 1$. Let $\mathcal{V}[A]$ denote the set of measures q on $[0, T] \times A$ with first marginal equal to Lebesgue measure (i.e., $q([0, t] \times A) = t$ for $t \in [0, T]$) such that

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

When A is understood, we write simply \mathcal{V} . An element of \mathcal{V} is called a *relaxed control*. Endow \mathcal{V} with the weakest topology making the maps $q \mapsto \int \varphi dq$ continuous for each continuous function $\varphi : [0, T] \times A \rightarrow \mathbb{R}$ satisfying $|\varphi(t, a)| \leq 1 + |a|^p$ for all (t, a) . This topology is metrizable via a natural analog p -Wasserstein metric, which we will not need to define

explicitly. Note that $\{q/T : q \in \mathcal{V}\}$ is a closed subset of $\mathcal{P}^p([0, T] \times A)$, and thus \mathcal{V} is a Polish space. When A is compact, so is \mathcal{V} .

We will frequently identify an element $q \in \mathcal{V}$ with the measurable map $t \mapsto q_t \in \mathcal{P}(A)$ arising from its disintegration $q(dt, da) = dtq_t(da)$, which is unique up to (Lebesgue) almost everywhere equality. There is a one-to-one map between \mathcal{V} and the space of (equivalence classes of a.e. equal) functions from $[0, T]$ to $\mathcal{P}(A)$, but let us check that we may reasonably define a *canonical* $\mathcal{P}(A)$ -valued process on \mathcal{V} , which will be denoted Λ . The natural filtration on \mathcal{V} is $\mathbb{F}^\Lambda = (\mathcal{F}_t^\Lambda)_{t \in [0, T]}$, where for each t \mathcal{F}_t^Λ is generated by the map $q \mapsto 1_{[0, t]}q$, or equivalently by the maps $q \mapsto q(C)$, where C ranges over measurable subsets of $[0, t] \times A$. The following lemma seems to be known and often used implicitly in the literature, but a sketch the proof is included for the reader's convenience, as a precise reference is difficult to locate.

Lemma 2.1.14. *There exists a \mathbb{F}^Λ -predictable process $\Lambda : [0, T] \times \mathcal{V} \rightarrow \mathcal{P}(A)$ such that $\Lambda(t, q) = q_t$ for almost every $t \in [0, T]$, for each $q \in \mathcal{V}$. In particular, $q = \Lambda(t, q)(da)$ for each $q \in \mathcal{V}$.*

Proof. Define $F_\epsilon : [0, T] \times \mathcal{V} \rightarrow \mathcal{P}(A)$ by

$$F_\epsilon(t, q) := \frac{1}{t - (t - \epsilon)^+} \int_{(t - \epsilon)^+}^t q_s ds.$$

Then $F_\epsilon(\cdot, q)$ is continuous for each $q \in \mathcal{V}$, and $F_\epsilon(t, \cdot)$ is \mathcal{F}_t^Λ -measurable for each $t \in [0, T]$. Hence F_ϵ is \mathcal{F}_t^Λ -predictable. Fix arbitrarily $q^0 \in \mathcal{P}(A)$, and for each $(t, q) \in [0, T] \times \mathcal{V}$ define

$$\Lambda(t, q) := \begin{cases} \lim_{\epsilon \downarrow 0} F_\epsilon(t, q) & \text{if the limit exists,} \\ q^0 & \text{otherwise.} \end{cases}$$

Then Λ is predictable, and it follows from Lebesgue's differentiation theorem (arguing with a countable convergence-determining class of functions on A) that $\Lambda(t, q) = q_t$ for almost every $t \in [0, T]$. \square

Write $\Lambda_t = \Lambda(t, \cdot)$ for the canonical process on \mathcal{V} given by Lemma 2.1.14. Then, it is straightforward to check that

$$\mathcal{F}_t^\Lambda = \sigma(\Lambda_s : s \leq t).$$

Of particular interest, of course, are the *strict controls*, which are of the form $q = dt\delta_{\alpha(t)}(da)$ for measurable $\alpha : [0, T] \rightarrow A$. It follows from Proposition 2.1.5 that the set of strict controls is dense in \mathcal{V} . Since we will be working with random elements of \mathcal{V} , the following adapted form of this statement will be useful. As a guiding example of the rather strange-looking assumptions, think of E as the path space $C([0, T]; \mathbb{R}^m)$, W as an m -dimensional Wiener process, and S_t equal to W stopped at time t , i.e. $S_t = W_{\cdot \wedge t}$.

Proposition 2.1.15. *Let $T > 0$, and fix a Polish space E and a closed convex subset A of a Euclidean space. Let (Ω, \mathcal{F}, P) be a probability space supporting stochastic processes $(S_t)_{t \in [0, T]}$ and $\Lambda = (\Lambda_t)_{t \in [0, T]}$, with values in E and $\mathcal{P}(A)$, respectively. Let $\mathcal{F}_t^S = \sigma(S_t)$, and assume that this defines a filtration, i.e. that $\mathcal{F}_u^S \subset \mathcal{F}_t^S$ for $u < t$. Assume that law of S_t*

is nonatomic for all $t > 0$. Finally, assume $\sigma(\Lambda_u : u \leq t)$ is conditionally independent of \mathcal{F}_t^S given \mathcal{F}_t^S , for each $t \in [0, T]$. Then there exists a sequence of $\mathbb{F}^S := (\mathcal{F}_t^S)_{t \in [0, T]}$ -adapted A -valued process $(\alpha_t^k)_{t \in [0, T]}$ satisfying:

1. For each k , α^k is uniformly bounded.
2. $(dt\delta_{\alpha_t^k}(da), S)$ converges in distribution to $(\Lambda = dt\Lambda_t(da), S)$, on the space $\mathcal{V} \times E$.
3. For each k , there exists a continuous function $\varphi_k : E \rightarrow \mathcal{V}$ such that $\varphi_k(S) = dt\delta_{\alpha_t^k}(da)$ a.s.
4. If $P \circ S^{-1} \in \mathcal{P}^p(E)$ for $p \geq 1$ and also $\mathbb{E} \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty$, then $P \circ (dt\delta_{\alpha_t^k}(da), S)^{-1}$ converges to $P \circ (\Lambda, S)^{-1}$ in $\mathcal{P}^p(\mathcal{V} \times E)$.

Proof. First, we may reduce to the case where A is compact as follows. Let A_n denote the intersection of A with the closed ball of radius n centered at the origin. For sufficiently large n_0 , A_n is nonempty for all $n \geq n_0$. Fix $a_0 \in A_{n_0}$ arbitrarily, and define $\iota_n : A \rightarrow A$ by $\iota_n(a) = a$ for $a \in A_n$ and $\iota_n(a) = a_0$ for $a \notin A_n$. Letting $\Lambda_t^n = \Lambda_t \circ \iota_n^{-1}$, it is clear that $\Lambda^n \rightarrow \Lambda$ almost surely (in the topology of \mathcal{V}). Moreover, $\Lambda_t^n(A_n) = 1$ for all t and $n \geq n_0$, almost surely. Finally, since $|\iota_n(a)| \leq |a|$ for $n \geq n_0$, we have the uniform integrability

$$\lim_{r \rightarrow \infty} \sup_n \int_0^T \int_{|a| > r} |a|^p \Lambda_t^n(da) dt \leq \lim_{r \rightarrow \infty} \lim \int_0^T \int_{|a| > r} |a|^p \Lambda_t(da) dt = 0.$$

Thus $P \circ (\Lambda^n, S)^{-1}$ converges to $P \circ (\Lambda, S)^{-1}$ in $\mathcal{P}^p(\mathcal{V} \times E)$.

Now assume that A is compact, and note that it remains only to construct α^k satisfying property (2). Define $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ by

$$\mathcal{F}_t = \sigma(S_t, \Lambda_u : u \leq t).$$

By a version of the so-called *Chattering Lemma* [76, Theorem 2.2(b)], there exists a sequence of *simple* \mathbb{F} -adapted A -valued processes $(\alpha_t^k)_{t \in [0, T]}$ on Ω such that $dt\delta_{\alpha_t^k}(da) \rightarrow \Lambda$ almost surely, and a fortiori $(S, dt\delta_{\alpha_t^k}(da))$ converges to (Λ, S) in distribution. Here, a *simple* \mathbb{F} -adapted A -valued process α is of the form

$$\alpha_t = a_0 1_{[0, t_0]}(t) + \sum_{i=1}^{n-1} a_i 1_{(t_i, t_{i+1}]}(t), \quad (2.10)$$

where $a_0 \in A$ is deterministic, a_i is an \mathcal{F}_{t_i} -measurable A -valued random variable, and $0 < t_0 < t_1 < \dots < t_n = T$ for some n is a fixed, deterministic time grid. It remains to show that, for any simple \mathbb{F} -adapted A -valued process α , there exist \mathbb{F}^S -adapted A -valued processes α^k such that $(S, dt\delta_{\alpha_t^k}(da)) \rightarrow (S, dt\delta_{\alpha_t}(da))$ weakly.

Now let α be of the form (2.10). By Proposition 2.1.6, there exists a sequence of $(\mathcal{F}_{t_i}^S)_{i=1}^n$ -adapted processes $(a_i^k)_{i=1}^n$ such that $(S_{t_1}, \dots, S_{t_n}, a_1^k, \dots, a_n^k)$ converges in law to $(S_{t_1}, \dots, S_{t_n}, a_1, \dots, a_n)$ as $k \rightarrow \infty$. Define

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

The map

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

is easily seen to be continuous, and thus $(S, dt\delta_{\alpha_t^k}(da))^{-1}$ converges in law to $(S, dt\delta_{\alpha_t}(da))^{-1}$, completing the proof. \square

We close the section with an easy but useful compactness criterion for the space \mathcal{V} .

Lemma 2.1.16. *Assume $A \subset \mathbb{R}^k$ for some k , and suppose a set $K \subset \mathcal{V}$ satisfies*

$$\lim_{r \rightarrow \infty} \sup_{q \in K} \int_0^T \int_{|a| > r} |a|^p q_t(da) dt = 0. \quad (2.11)$$

Then K is relatively compact. In particular, (2.11) holds if there exists $p' > p$ such that

$$\sup_K \int_0^T \int_A |a|^{p'} q_t(da) dt < \infty.$$

Proof. Since $p \geq 1$, it follows from Markov's inequality and (2.11) that

$$\sup_{q \in K} q \{(t, a) : |a| > r\} \leq \frac{1}{r} \sup_{q \in K} \int_0^T \int_A |a| q_t(da) dt < \infty,$$

for each $r > 0$. Thus K is tight, and it follows from Proposition 2.1.7 that K is relatively compact. \square

2.2 McKean-Vlasov limits

This section surveys some of the known results on McKean-Vlasov limits of interacting diffusions. The goal is not to be completely formal or general in the statements of results, but rather to present the ideas of the results and a new perspective on their proofs. First, we elaborate on the interacting diffusion model discussed of Section 1.1.1.

2.2.1 No common noise

For the sake of this section, consider bounded continuous (recalling that we implicitly equip $\mathcal{P}(\mathbb{R}^d)$ with weak convergence) functions

$$(b, \sigma) : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d \times \mathbb{R}^{d \times m}.$$

Consider a system of n particles with positions $(X_t^{n,1}, \dots, X_t^{n,n})$ at time t , which evolve according to the dynamics

$$\begin{cases} dX_t^{n,i} &= b(X_t^{n,i}, \hat{\mu}_t^n)dt + \sigma(X_t^{n,i}, \hat{\mu}_t^n)dW_t^i, \\ \hat{\mu}_t^n &= \frac{1}{n} \sum_{i=1}^n \delta_{X_t^{n,i}}, \end{cases} \quad (2.12)$$

where W^1, \dots, W^n are independent standard Wiener processes of dimension m , and the initial positions $X_0^{n,1}, \dots, X_0^{n,n}$ are i.i.d. with initial law $\lambda \in \mathcal{P}(\mathbb{R}^d)$. Here, we view $\hat{\mu}^n$ as a random measure on the path space $\mathcal{C}^d = C([0, \infty); \mathbb{R}^d)$, and its time-marginals are $\hat{\mu}_t^n = \frac{1}{n} \sum_{i=1}^n \delta_{X_t^{n,i}}$. Note that under the stated assumptions, there exists a weak solution of this SDE system, but it may not be unique.

The heuristic argument outlined in the introduction suggests that the limiting dynamics as $n \rightarrow \infty$ could be described by the McKean-Vlasov equation, defined as follows. Let us say that a probability measure μ on \mathcal{C}^d is a *solution of the McKean-Vlasov equation with initial law λ* if $\mu_0 = \lambda$ and there exists a weak solution $(X_t)_{t \geq 0}$ of the SDE

$$dX_t = b(X_t, \mu_t)dt + \sigma(X_t, \mu_t)dW_t,$$

(defined on some filtered probability space supporting an m -dimensional Wiener process W) such that the law of X is μ . Alternative definitions are not uncommon in the literature: First, we may study the (nonlinear) martingale problem associated to the generator which acts on smooth functions φ by

$$L_m \varphi(x) = b(x, m)^\top D\varphi(x) + \frac{1}{2} \text{Tr}[\sigma(x, m)\sigma^\top(x, m)D^2\varphi(x)], \quad (2.13)$$

where D and D^2 denote the gradient and Hessian, respectively. A solution of the martingale problem is a law $\mu \in \mathcal{P}(\mathcal{C}^d)$ under which the process defined on \mathcal{C}^d by

$$[0, \infty) \times \mathcal{C}^d \ni (t, x) \mapsto \varphi(x_t) - \int_0^t L_{\mu_s} \varphi(x_s) ds$$

is a martingale for each smooth φ with compact support. (Here the filtration on \mathcal{C}^d is the canonical one.) Second, the McKean-Vlasov equation is related to the nonlinear Kolmogorov forward equation, formally (assuming μ_t has a smooth density) written as

$$\partial_t \mu_t(x) = L_{\mu_t}^* \mu_t(x), \quad \mu_0 = \lambda$$

where $L_{\mu_t}^*$ is adjoint to L_{μ_t} . For discussions of these various formulations of the McKean-Vlasov equation and connections between them, see [52, 96, 53, 104].

The following limit theorem and its proof are archetypal in the literature on McKean-Vlasov limits [96, 53]. To fix some ideas, endow \mathcal{C}^d with the usual topology of uniform convergence on compacts.

Theorem 2.2.1. *For each n , let $\hat{\mu}^n$ arise from a weak solution of the SDE system (2.12). Then the set $\{\hat{\mu}^n : n \geq 1\}$ of $\mathcal{P}(\mathcal{C}^d)$ -valued random variables is tight, and every limit is supported on the set of solutions of the McKean-Vlasov equation with initial law λ .*

There are two obvious corollaries. First, from existence of the n -particle systems we deduce the existence of a solution of the McKean-Vlasov equation. Second, if it can be shown that the solution of the McKean-Vlasov equation is unique, then it follows that $\hat{\mu}^n$ converge in law to this unique solution, which is a *deterministic* measure! This phenomenon, known as *propagation of chaos*, is discussed in detail in [104].

Soon we will discuss a new proof of Theorem 2.2.1, but let us first mention two interesting ways this is typically proven. First of all, to prove tightness it suffices to show that the mean measures of $\hat{\mu}^n$ are tight (see Proposition 2.1.12), and this typically employs well known tightness criteria on \mathcal{C}^d . The meat of the proof is in characterizing the limits, which is often done by studying the (nonlinear) martingale problem associated with the McKean-Vlasov equation [96, 53].

Another approach works when the coefficients b and σ are Lipschitz (using again a Wasserstein distance for the measure argument), known as a *trajectorial propagation of chaos* argument detailed in [104]. In this case, the McKean-Vlasov solution can be shown to be unique, and a direct coupling argument yields a stronger form of convergence. More precisely, let $(\mu_t)_{t \geq 0}$ denote this unique solution, and fix a common probability space supporting independent Wiener processes W^1, W^2, \dots and i.i.d. random vectors ξ^1, ξ^2, \dots with common law λ . We may solve the SDE (2.12) with $X_0^{n,i} = \xi^i$, and *using the same driving Wiener process* solve also the SDEs

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

Then $(Y_t^i)_{i=1}^\infty$ are i.i.d. with law μ_t , and it can be shown that a much stronger form of convergence holds, namely

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\sup_{0 \leq s \leq t} |X_s^{n,i} - Y_s^i| \wedge 1 \right] = 0.$$

This can be shown to imply the weak convergence of the empirical measures stated in Theorem (2.2.1), and often yields a rate of convergence. See [104] for details of these arguments. This coupling technique will be useful in Section 6.2.

2.2.2 Common noise

Suppose we are given another bounded, continuous (for simplicity) function $\sigma_0 : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^{d \times m_0}$, and the particle system (2.12) is replaced with the following:

$$\begin{cases} dX_t^{n,i} &= b(X_t^{n,i}, \hat{\mu}_t^n)dt + \sigma(X_t^{n,i}, \hat{\mu}_t^n)dW_t^i + \sigma_0(X_t^{n,i}, \hat{\mu}_t^n)dB_t, \\ \hat{\mu}^n &= \frac{1}{n} \sum_{i=1}^n \delta_{X^{n,i}}, \end{cases} \quad (2.14)$$

Here B is an m_0 -dimensional Wiener process, independent of the others. Intuitively, even if the empirical measures $\hat{\mu}^n$ admit a unique limit, we should not expect it to be deterministic, since there is a common source of noise B which persists in the limit $n \rightarrow \infty$. The new limiting dynamics can again be described in a number of ways, either in terms of a McKean-

Vlasov SDE, a martingale problem, or a PDE. The rigorous connections between these three formulations become more subtle in this case, but let us describe them optimistically:

1. First, adapting the notion of McKean-Vlasov equation above, we may look for a *random* measure μ on \mathcal{C}^d such that there exists a solution X of the SDE

$$dX_t = b(X_t, \mu_t)dt + \sigma(X_t, \mu_t)dW_t + \sigma_0(X_t, \mu_t)dB_t, \quad (2.15)$$

where B and W are independent Wiener processes, such that μ equals the conditional law of X given B . It is somewhat a matter of taste to work with a random measure μ on the path space, rather than the flow of marginal laws $(\mu_t)_{t \geq 0}$, but our choice will prove useful later.

2. Under Lipschitz assumptions, Kurtz and Xiong [83] proved existence and uniqueness for this equation and showed that it arises as the limit of $\hat{\mu}^n$ in distribution. They also show rigorously that the solution μ solves the *stochastic* Kolmogorov forward equation:

$$\partial_t \mu_t(x) = L_{\mu_t}^* \mu_t(x) - \sum_{i,j} \partial_{x_i} [(\sigma_0)_{i,j}(x, \mu_t) \mu_t(x)] dB_t^j, \quad \mu_0 = \lambda, \quad (2.16)$$

where L_{μ}^* now contains the term $(\sigma\sigma^\top + \sigma_0\sigma_0^\top)$ in place of $\sigma\sigma^\top$. (No effort will be made here to make rigorous sense of this SPDE.)

3. Alternatively, Dawson and Vaillancourt in [40, 106] prefer a martingale problem approach. Under weaker assumptions than those of [83], the limits of $\hat{\mu}^n$ is now non-unique but can be described by the solutions of a martingale problem on the space $\mathcal{P}(\mathbb{R}^d)$. They show that this corresponds to a *weak solution* of the SPDE (2.16), in the sense that μ need not be B -measurable, and also the equation is required to hold only in a distributional sense, i.e. when μ_t is integrated against smooth compactly supported functions.

In the non-unique regime, the first formulation described above seems less flexible so far as the literature seems to provide no corresponding *weak solution* concept. It turns out the methods developed in this thesis for mean field games with common noise yield a new interpretation of this McKean-Vlasov equation along with a new proof of the limit theorem. The next section describes this approach in some detail, with the goal of illustrating some of the key ideas in a setting much simpler than mean field games. The idea of the argument is to tag a representative particle; this point of view adapts nicely to the MFG setting, where it is quite useful to keep track a representative agent.

2.2.3 An unorthodox derivation

Let us loosely define a *weak solution of the McKean-Vlasov equation starting from λ* to be a random measure μ such that there exists a *weak* solution X of the SDE (2.15) such that μ agrees with the conditional law of X given (B, μ) , and also (B, μ) is assumed to be independent of W . More precisely:

Definition 2.2.2. A weak solution of the McKean-Vlasov equation starting from λ is a tuple $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, X)$, where $(\tilde{\Omega}, \mathcal{F}_T, \mathbb{F}, P)$ is a complete filtered probability space supporting (B, W, μ, X) satisfying

1. B and W are independent \mathbb{F} -Wiener processes of respective dimension m_0 and m , respectively. The processes X is continuous and \mathbb{F} -adapted with values in \mathbb{R}^d , and $P \circ X_0^{-1} = \lambda$. Finally, μ is a random element of $\mathcal{P}(\mathcal{C}^d)$ such that $\mu(C)$ is \mathcal{F}_t -measurable whenever $C \in \mathcal{F}_t^X$ for each $t \geq 0$, where $(\mathcal{F}_t^X)_{t \geq 0}$ denotes the natural filtration on \mathcal{C}^d .
2. X_0 , W , and (B, μ) are independent.
3. The SDE holds

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

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

We may shorten condition (4) to $\mu = P(X \in \cdot \mid B, \mu)$, and note that it implies the arguably more natural condition

$$\mu_t = P(X_t \in \cdot \mid B, \mu) = P(X_t \in \cdot \mid B_s, \mu_s : s \leq t), \text{ a.s.},$$

for each $t \geq 0$. Alternatively, we may define a weak solution in terms of the law of (B, W, X, μ) induced on the canonical space

$$\Omega = C([0, \infty); \mathbb{R}^{m_0} \times \mathbb{R}^m \times \mathbb{R}^d) \times \mathcal{P}(\mathcal{C}^d).$$

It is not too difficult to show that a weak solution in this sense gives rise to a certain type of weak solution of the SPDE (2.16), and even the martingale problems described in [40, 106]. The other direction is not as clear; constructing a weak solution of the McKean-Vlasov equation from a solution of the SPDE (2.16) requires an important uniqueness assumption, as in [83], but we will not go into the details. Even when uniqueness fails for the McKean-Vlasov equation, we have the following limit theorem:

Theorem 2.2.3. Assume the SDE system (2.14) is unique in law for each n , and for each n let $B^n, W^{n,1}, \dots, W^{n,n}$, and $X^{n,i}$ denote a weak solution, defined on some filtered probability space $(\Omega_n, (\mathcal{F}_t^n)_{t \geq 0}, P_n)$. Define $\hat{\mu}^n$ accordingly. Then

$$\{Q_n := P_n \circ (B^n, W^{n,1}, X^{n,1}, \hat{\mu}^n)^{-1} : n \geq 1\}$$

is tight in Ω , and every limit is a weak solution of the McKean-Vlasov equation with initial law λ .

Proof sketch. Note that $Q_n = P_n \circ (B^n, W^{n,k}, X^{n,k}, \hat{\mu}^n)^{-1}$ for each choice of index $k \leq n$, because of the symmetry of the system and the assumption of uniqueness in law (the easy proof can be found in [106]). It is easy to check that $(\Omega_n, (\mathcal{F}_t^n)_{t \geq 0}, P_n, B^n, W^{n,1}, X^{n,1}, \hat{\mu}^n)$ in fact satisfies all of the properties of Definition (2.2.2), with the exception of the second. Indeed, the fourth property follows by symmetry. Moreover, it can be shown that these

properties are *closed*, in the sense that they must hold under any limit point of Q_n . See the proof of Lemma 6.1.5 for more details, along with the simple proof that the independence property (2) appears in the limit. For example, property (4) holds at the limit because if $(B^n, X^{n,1}, \hat{\mu}^n)$ converges in law to (B, X, μ) , then for continuous bounded functions φ and ψ we have, by symmetry,

$$\mathbb{E} \left[\varphi(B, \mu) \left(\psi(X) - \int \psi d\mu \right) \right] = \lim_n \mathbb{E} \left[\varphi(B^n, \hat{\mu}^n) \left(\psi(X^{n,1}) - \int \psi d\hat{\mu}^n \right) \right] = 0.$$

Finally, tightness follows from fairly standard arguments, as will be demonstrated more carefully in Proposition 5.3.2 during our analysis of mean field games. \square

It is worth emphasizing that the more standard arguments of [40, 106] focus solely on the convergence of the empirical measures $\hat{\mu}^n$. On the other hand, the new argument above keeps track of more information, namely the joint law of $(B^n, W^{n,1}, \hat{\mu}^n, X^{n,1})$, which we interpret as the joint law of precisely those processes which are *directly relevant to the first particle* in the sense that they appear in the dynamics (2.15). By symmetry, we may choose any particle to be the *representative*, not just the first. The assumption of uniqueness in law in 2.2.3 was only used for the sake of symmetry. In fact, we can make sense of this argument even if we drop this uniqueness assumption! The proof of the following theorem is essentially identical.

Theorem 2.2.4. *For each n let $B^n, W^{n,1}, \dots, W^{n,n}$, and $X^{n,i}$ denote a weak solution of the SDE system (2.14), defined on some filtered probability space $(\Omega_n, (\mathcal{F}_t^n)_{t \geq 0}, P_n)$. Define $\hat{\mu}^n$ accordingly. Then*

$$\left\{ Q_n := \frac{1}{n} \sum_{k=1}^n P_n \circ (B^n, W^{n,k}, X^{n,k}, \hat{\mu}^n)^{-1} : n \geq 1 \right\} \quad (2.17)$$

is tight in Ω , and every limit is a weak solution of the McKean-Vlasov equation with initial law λ .

Remark 2.2.5. Note that Q_n defined in (2.17) may alternatively be written as

$$Q_n = P_n \circ (B^n, W^{n,U}, X^{n,U}, \hat{\mu}^n)^{-1},$$

where U is a random variable, uniformly distributed on $\{1, \dots, n\}$, drawn independently of the other processes (constructed, if necessary, by enlarging the probability space). The point is that now, without symmetry, the *representative* particle cannot be chosen arbitrarily, but rather it must be chosen *uniformly at random*. Note also that the limiting behavior of $P_n \circ (B^n, \hat{\mu}^n)^{-1}$ can be determined from that of Q_n , simply by projection onto the appropriate marginals.

Theorem 2.2.4 is not stated merely for the sake of generality, but rather because this idea will reappear in our study of mean field game limits. Many symmetric games have asymmetric Nash equilibria, and thus to study the limits of arbitrary Nash equilibria we will use this trick to force exchangeability into the picture.

Now, returning to the setting *without* common noise, it is natural to try to specialize this new approach. Indeed, removing the common noise term everywhere, it can be shown that the limits of $P_n \circ (W^{n,1}, \hat{\mu}^n, X^{n,1})^{-1}$ as in Theorem 2.2.3 (or we may argue more generally as in Theorem 2.2.4) are weak solutions in the following sense:

Definition 2.2.6. A *weak solution of the McKean-Vlasov equation without common noise starting from λ* is a tuple $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P, W, X, \mu)$, where $(\tilde{\Omega}, \mathcal{F}_T, \mathbb{F}, P)$ is a complete filtered probability space supporting (W, μ, X) satisfying

1. W is an \mathbb{F} -Wiener processes of dimension m , respectively. The processes X is continuous and \mathbb{F} -adapted with values in \mathbb{R}^d , and $P \circ X_0^{-1} = \lambda$. Finally, μ is a random element of $\mathcal{P}(\mathcal{C}^d)$ such that $\mu(C)$ is \mathcal{F}_t -measurable whenever $C \in \mathcal{F}_t^X$ for each $t \geq 0$, where $(\mathcal{F}_t^X)_{t \geq 0}$ denotes the natural filtration on \mathcal{C}^d .
2. X_0 , W , and μ are independent.
3. The SDE holds

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

4. $\mu = P(X \in \cdot \mid \mu)$. That is, μ is a version of the conditional law of X given μ .

If it happens that μ is deterministic a.s., then we refer to μ itself as a strong solution. More precisely, if there exists $\nu \in \mathcal{P}(\mathcal{C}^d)$ such that $P(\mu = \nu) = 1$, then we may refer to ν as a strong solution.

The following proposition indicates that there is actually no need for this weak solution concept for describing the limits of finite systems, at least in the absence of common noise.

Proposition 2.2.7. *Let $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P, W, X, \mu)$ be a weak solution of the McKean-Vlasov equation without common noise. Then μ is supported by the set of strong solutions. That is, if $S \subset \mathcal{P}(\mathcal{C}^d)$ is the set of strong solutions, then $P(\mu \in S) = 1$.*

Proof. By independence of X_0 , W , and μ , the conditional law of (X_0, W) given μ is the product of λ with the Wiener measure. Note also that

$$\nu = P(X \in \cdot \mid \mu = \nu), \tag{2.18}$$

for almost every ν . Intuitively, we now just “freeze” $\mu = \nu$ and note that the correct SDE and fixed point condition hold under this conditional measure. This will be made rigorous using martingale problems (see [102] for background). First, define the enlarged filtration

$$\mathcal{G}_t := \mathcal{F}_t \vee \sigma(\mu).$$

Independence of W and μ implies that W remains a Wiener process with respect to $\mathbb{G} := (\mathcal{G}_t)_{t \geq 0}$. Thus, in the equation

$$X_t = X_0 + \int_0^t b(X_s, \mu_s)ds + \int_0^t \sigma(X_s, \mu_s)dW_s,$$

the meaning of stochastic integral is insensitive to the choice of reference filtration \mathbb{F} or \mathbb{G} (see [99, Theorem 2.16]). Hence, defining the generator L_m as in (2.13), the process

$$M_t^\varphi := \varphi(X_t) - \int_0^t L_{\mu_s} \varphi(X_s) ds$$

is a martingale under both \mathbb{F} and \mathbb{G} , for each smooth function φ of compact support. Now fix $t > s \geq 0$, and let Y and Z denote bounded $\sigma(X_u : u \leq s)$ -measurable and $\sigma(\mu)$ -measurable random variables, respectively. Then YZ is \mathcal{G}_s -measurable, and for each φ as above we have

$$0 = \mathbb{E}^P [ZY (M_t^\varphi - M_s^\varphi)].$$

Since this holds for every such choice of Z ,

$$0 = \mathbb{E}^P [Y (M_t^\varphi - M_s^\varphi) | \mu], \text{ a.s.} \quad (2.19)$$

This holds for every bounded $\sigma(X_u : u \leq s)$ -measurable Y , and we conclude that M^φ is a martingale under the conditional measure $P(\cdot | \mu)$, almost surely, for each φ . Although the null set in (2.19) may depend on the choice of Y , it is well known that we need only check this for countably many Y . Similarly, by restricting to a countable dense set of φ , we may interchange the order of the quantifiers once more to conclude that the following holds almost surely: under $P(\cdot | \mu)$, M^φ is a martingale for each φ . Combined with (2.18), we conclude that $P \circ \mu^{-1}$ -almost every $\nu \in \mathcal{P}(\mathcal{C}^d)$ is a strong solution. \square

Remark 2.2.8. A converse of Proposition 2.2.7 is also true: Given any probability measure ρ on $\mathcal{P}(\mathcal{C}^d)$ with $\rho(S) = 1$, there exists a weak solution $(\tilde{\Omega}, \mathbb{F}, P, W, X, \mu)$ of the McKean-Vlasov equation without common noise with $P \circ \mu^{-1} = \rho$. We will not need this fact, and to prove it cleanly requires a notion of martingale problem with random coefficients

Finally, as a corollary of Proposition 2.2.7 and Theorem 2.2.4 (or rather its simple adaptation to the case without common noise), we obtain the following generalization of Theorem 2.2.1.

Corollary 2.2.9. *For each n let $(W^{n,i}, X^{n,i})_{i=1}^n$ denote a weak solution of the SDE system (2.12), defined on some filtered probability space $(\Omega_n, (\mathcal{F}_t^n)_{t \geq 0}, P_n)$. Define $\hat{\mu}^n$ accordingly. Then*

$$\left\{ Q_n := \frac{1}{n} \sum_{k=1}^n P_n \circ (W^{n,k}, X^{n,k}, \hat{\mu}^n)^{-1} : n \geq 1 \right\}$$

is tight in Ω , and every limit is concentrated on the set of strong solutions (without common noise) with initial law λ .

2.3 Static mean field games

Let us now forget about *dynamics* for the moment and see how mean field theory looks for static but *competitive* systems. The goal of this section is to develop some intuition for mean field games by beginning with the technically simpler setting of *static* or *one-shot*

games, before turning to stochastic differential games. The first section treats the simplest conceivable setting of *deterministic* static games, borrowing heavily from Cardaliaguet's notes [24]. The second section introduces stochastic factors into these static games and derives a limit theorem which resembles in many ways the main limit theorem on stochastic differential games but is unencumbered by technical details arising from continuous-time dynamics. First, we summarize a few results we will use from the analysis of set-valued maps, with a view toward fixed point theorems.

2.3.1 Elements of set-valued analysis

The material of this short section is borrowed from [5, Chapter 17], to which the reader is referred for more background. Given two sets X and Y , the *graph* of a set-valued function $\varphi : X \rightarrow 2^Y$ is the subset of $X \times Y$ given by $\{(x, y) : x \in X, y \in \varphi(x)\}$. If Y is a topological space, we say φ has *closed values* if $\varphi(x)$ is a closed set for each $x \in X$. Similarly, if Y is a subset of a vector space, we say φ has *convex values* if $\varphi(x)$ is a convex set for each x . The following version of Kakutani's fixed point theorem is due to K. Fan [48]; see also Corollary 17.55 of [5].

Theorem 2.3.1 (Kakutani's fixed point theorem). *Let K be a nonempty compact convex subset of a locally convex Hausdorff (topological vector) space, and suppose $\varphi : K \rightarrow 2^K$ has a closed graph and nonempty convex compact values. Then φ admits a fixed point; that is, there exists $x \in K$ such that $x \in \varphi(x)$.*

In our applications, this will be applied to set-valued maps of the form

$$\varphi(x) = \arg \max_{y \in \psi(x)} f(x, y) := \left\{ y \in Y : f(x, y) = \sup_{y' \in \psi(x)} f(x, y') \right\}, \quad (2.20)$$

where $\psi : X \rightarrow 2^Y$ and $f : X \times Y \rightarrow \mathbb{R}$. A well known theorem due to Berge allows us to determine that φ has a closed graph when ψ and f have adequate continuity properties. This requires a brief discussion of continuity notions for set-valued maps. Assume X and Y are metric spaces, and let $\varphi : X \rightarrow 2^Y$ be a set-valued map. We say φ is *lower hemicontinuous* if, for every sequence $(x_n)_n$ in X converging to x , and for every $y \in \varphi(x)$, there exist $y_{n_k} \in \varphi(x_{n_k})$ such that $y_{n_k} \rightarrow y$. If φ is closed-valued, we say φ is *upper hemicontinuous* if, whenever $x_n \rightarrow x$ in E and $y_n \in \varphi(x_n)$ for each n , the sequence $(y_n)_n$ has a limit point in $\varphi(x)$. We say φ is *continuous* if it is both upper hemicontinuous and lower hemicontinuous. We may finally state Berge's theorem, quoted from [5, Theorem 17.31].

Theorem 2.3.2 (Berge's Theorem). *Suppose X and Y are metric spaces, $f : X \times Y \rightarrow \mathbb{R}$ is continuous, and $\psi : X \rightarrow 2^Y$ is continuous. Then the map φ defined by (2.20) is upper hemicontinuous and has nonempty compact values.*

In order to apply Kakutani's theorem, Berge's theorem is often used in conjunction with the following version of the closed graph theorem:

Theorem 2.3.3 (Theorem 17.11 of [5]). *Suppose X and Y are metric spaces, with Y compact. Then a closed-valued map $\varphi : X \rightarrow 2^Y$ is upper hemicontinuous if and only if its graph is closed.*

2.3.2 The deterministic case

Imagine we have a large population of n agents, each of whom can choose an action from a common *action set* A . Suppose we are also given *objective functions* $J_i^n : A^n \rightarrow \mathbb{R}$. A vector (a_1, \dots, a_n) in A^n is called a *strategy profile*, meaning the i^{th} agent has chosen the strategy a_i . The objective of agent i is to choose a_i to try to maximize $J_i^n(a_1, \dots, a_n)$. Of course, the optimization problems of each agents are interdependent. In order to resolve them simultaneously we will look for a *Nash equilibrium*. More generally, for $\epsilon > 0$, an ϵ -Nash equilibrium is defined to be any strategy profile (a_1^*, \dots, a_n^*) satisfying

$$J_i^n(a_1^*, \dots, a_n^*) + \epsilon \geq \sup_{a \in A} J_i^n(a_1^*, \dots, a_{i-1}^*, a, a_{i+1}^*, \dots, a_n^*), \text{ for } i = 1, \dots, n.$$

Following Nash's famous argument, one can show under quite modest assumptions (e.g., when A is compact metric and J_i^n are continuous) that a Nash equilibrium exists, at least among *mixed strategies*. A mixed strategy is a vector in $\mathcal{P}(A)^n$ rather than A^n , and Nash equilibrium among mixed strategies is a vector $(\pi_1^*, \dots, \pi_n^*) \in \mathcal{P}(A)^n$ satisfying

$$\bar{J}_i^n(\pi_1^*, \dots, \pi_n^*) \geq \sup_{\pi \in \mathcal{P}(A)} \bar{J}_i^n(\pi_1^*, \dots, \pi_{i-1}^*, \pi, \pi_{i+1}^*, \dots, \pi_n^*), \text{ for } i = 1, \dots, n,$$

where

$$\bar{J}_i^n(\pi_1, \dots, \pi_n) := \int_{A^n} J_i^n(a_1, \dots, a_n) \prod_{k=1}^n \pi_k(da_k).$$

The point of passing to mixed strategies is to acquire *convexity*, both of the action set $\mathcal{P}(A)$ and of the objective functions \bar{J}_i^n , which facilitate an application of Kakutani's fixed point theorem.

When the number of agents n is large, Nash equilibria may be difficult to compute. On the other hand, the desired output of a game-theoretic model is not always a full description of the equilibria. In large-population games, for example, the *distribution* of the equilibrium strategies (or a quantity derived it) is often the main object of interest. In this case, if n is sufficiently large and the game is *symmetric* in a certain sense, then some simplifying analysis is available. By *symmetric* we mean that

$$J_{\pi(i)}^n(a_{\pi(1)}, \dots, a_{\pi(n)}) = J_i^n(a_1, \dots, a_n),$$

for every permutation π of $\{1, \dots, n\}$. More specifically, we will assume there is a single function $F : A \times \mathcal{P}(A) \rightarrow \mathbb{R}$, where $\mathcal{P}(A)$ is the set of probability measures on A , such that

$$J_i^n(a_1, \dots, a_n) = F\left(a_i, \frac{1}{n} \sum_{k=1}^n \delta_{a_k}\right), \text{ for } 1 \leq i \leq n.$$

In an asymptotic sense illustrated more precisely in [24, Theorem 2.1], it is not too restrictive to assume the J_i^n take this form. The main limiting result is the following:

Theorem 2.3.4 (Theorem 2.4 of [24]). *Suppose that for each n we are given an ϵ_n -Nash equilibrium (a_1^n, \dots, a_n^n) , where $\epsilon_n \rightarrow 0$. Let $\hat{\mu}^n = \frac{1}{n} \sum_{i=1}^n \delta_{a_i^n}$ denote the empirical measure of these strategy profiles. Suppose the action space A is compact and metrizable, and the function F is jointly continuous. Then $\{\hat{\mu}^n\} \subset \mathcal{P}(A)$ is tight, and every weak limit μ^* is supported by the set of maximizers of $F(\cdot, \mu^*)$. That is, $\mu^*\{a \in A : F(a, \mu^*) \geq F(b, \mu^*) \ \forall b \in A\} = 1$.*

Proof. Fix any alternative action $b \in A$. The Nash property implies

$$F(a_i^n, \hat{\mu}^n) + \epsilon_n \geq F(b, \hat{\nu}_i^n), \text{ where } \hat{\nu}_i^n = \frac{1}{n} \left(\delta_b + \sum_{k \neq i}^n \delta_{a_k^n} \right).$$

Summing over $i = 1, \dots, n$, we get

$$\int_A F(a, \hat{\mu}^n) \hat{\mu}^n(da) + \epsilon_n \geq \frac{1}{n} \sum_{i=1}^n F(b, \hat{\nu}_i^n) \quad (2.21)$$

If the action space A is a compact metrizable space, then the sequence $\{\hat{\mu}^n\}$ admits weak limits. If $\hat{\mu}^n$ converges to a measure μ^* along a subsequence, then $\hat{\nu}_i^n$ must too converge to the same μ along the same subsequence, uniformly in $i \leq n$. More precisely, recalling the definition of the Wasserstein distance from (2.6), it is not hard to check that

$$\ell_{A,1}(\hat{\nu}^n, \hat{\nu}_i^n) \leq c/n, \text{ for } 1 \leq i \leq n,$$

where we have chosen any compatible metric on A , and $c > 0$ is the diameter of A with respect to this metric, finite thanks to the compactness of A . Since F is uniformly continuous, we may pass to the limit in (2.21) (using the assumption that $\epsilon_n \rightarrow 0$) to get

$$\int_A F(a, \mu^*) \mu^*(da) \geq F(b, \mu^*).$$

This holds for each $b \in A$, and thus μ^* is concentrated on the set of maximizers of $F(\cdot, \mu^*)$. \square

With this result in mind, it is natural to define a *mean field game (MFG) solution* as any probability measure μ on A supported on the set of maximizers of $F(\cdot, \mu)$. Note, for example, that when $F(\cdot, \mu)$ admits a *unique* maximizer for each μ , we conclude that every MFG solution is in fact a point mass!

The definition of a MFG solution can be alternatively expressed as a fixed point condition. For each $\mu \in \mathcal{P}(A)$, define $\varphi(\mu)$ to be the set of probability measures on A supported on the set of maximizers of $F(\cdot, \mu)$. Then, the MFG solutions are exactly the fixed points of this set-valued map, in the sense that $\mu^* \in \varphi(\mu^*)$. Intuitively, given that the distribution of a continuum of adverse agents' strategies is μ , any given representative agent wishes to choose a strategy from the set of maximizers of $F(\cdot, \mu)$. Since all of the agents are identical, the fixed point is a natural expression of consistency, or equilibrium.

With this fixed point formulation in mind, we illustrate how to prove an existence theorem using Kakutani's theorem, in what is by now a well known argument [92, Theorem 1].

(Naturally, this could also be proven by combining Nash's existence theorem with a version of Theorem 2.3.4 for mixed strategy.)

Theorem 2.3.5. *Under the same assumptions of Theorem 2.3.4, the map φ admits a fixed point.*

Proof. Since $\mathcal{P}(A)$ is compact and convex, Kakutani's fixed point theorem will apply if we can show that $\varphi(\mu)$ is convex for each μ and that the graph $S := \{(\mu, \nu) \in \mathcal{P}(A) \times \mathcal{P}(A) : \nu \in \varphi(\mu)\}$ is closed. The convexity is clear. To show S is closed, note that $\nu \in \varphi(\mu)$ if and only if

$$\int_A [F(a, \mu) - F(b, \mu)] \nu(da) \geq 0, \quad \forall b \in A.$$

Since the functional of (μ, ν) on the left-hand side is continuous, it follows easily that S is closed. \square

2.3.3 The stochastic case with independent noises

Now suppose we are given i.i.d. random variables $(W_i)_{i=1}^\infty$ with values in a Polish space E and with common law λ , defined on some common probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Interpret W_1, W_2, \dots as idiosyncratic noises or shocks, with W_i specific to agent i . Suppose now that $F : E \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$. Each agent i in the n -player game chooses an A -valued $\sigma(W_1, \dots, W_n)$ -measurable random variable X_i to try to maximize

$$J_i^n(X_1, \dots, X_n) = \mathbb{E} \left[F \left(W_i, \frac{1}{n} \sum_{k=1}^n \delta_{X_k}, X_i \right) \right].$$

Nash equilibrium is defined in the same manner as before. Again, suppose A is a compact metrizable space and F is bounded and continuous. To prove a limit theorem, the basic idea is to think of $(W_i, \hat{\mu}^n, X_i)$ as the random variables *relevant* to agent i , where $\hat{\mu}^n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ is the empirical measure. Then, we uniformly at random pick one of the n agents to be the *representative agent*, and look at the law of the quadruple of random variables which are relevant to that representative. This leads to the law (2.22) considered in the limit theorem below.

We will state the limit theorem first in its more complicated form, involving a notion of *weak solution*. This is more robust to changes in the specification of the problem, in a sense which will be more clear in the next Section 2.3.4 dealing with common noise. Then, we show that weak solutions admit a simpler description as mixtures of *strong solutions*. First, we need some notation.

Let $\bar{\Omega} := E \times \mathcal{P}(A) \times A$ denote our canonical space, and let (W, μ, X) denote the canonical projections. Given a probability measure ρ on $E \times \mathcal{P}(A)$, let $\mathcal{A}(\rho)$ denote the set of probability measures P on $\bar{\Omega}$ with $E \times \mathcal{P}(A)$ -marginal equal to ρ , i.e. $P \circ (W, \mu)^{-1} = \rho$. The letter \mathcal{A} stands for “admissible,” as this set specifies which joint laws of action X and inputs (W, μ) are admissible for the optimization problems given in the definition below. Finally, let us say that a measure $Q \in \mathcal{P}(\bar{\Omega})$ is a *weak solution of the MFG* (or a *weak MFG solution*) if:

1. Q has E -marginal equal to λ , i.e. $Q \circ W^{-1} = \lambda$.

2. For each $P \in \mathcal{A}(\rho)$ where $\rho := Q \circ (W, \mu)^{-1}$, we have $\int F dQ \geq \int F dP$.
3. W and μ are independent under Q .
4. $Q(X \in \cdot \mid \mu) = \mu$, Q -almost surely. That is, $\mathbb{E}^Q[\varphi(X) \mid \mu] = \int \varphi d\mu$ a.s. for each bounded measurable $\varphi : A \rightarrow \mathbb{R}$.

For an intuitive explanation of the definition of Q^n in the following limit theorem, refer back to Remark 2.2.5.

Theorem 2.3.6. *Suppose for each n we are given an ϵ_n -Nash equilibrium (X_1^n, \dots, X_n^n) , where $\epsilon_n \rightarrow 0$. Let $\hat{\mu}^n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i^n}$ denote the empirical measure of these strategy profiles. Define*

$$Q^n := \frac{1}{n} \sum_{i=1}^n \mathbb{P} \circ (W_i, \hat{\mu}^n, X_i^n)^{-1}. \quad (2.22)$$

Then $(Q^n)_{n=1}^\infty \subset \mathcal{P}(\bar{\Omega})$ is tight, and every weak limit Q is a weak MFG solution.

Proof. Note that the E -marginal of Q^n does not depend on n . Since $\mathcal{P}(A)$ and A are compact, the $\mathcal{P}(A) \times A$ -marginals of Q^n are tight. Hence $(Q^n)_{n=1}^\infty$ is tight (see Lemma 2.1.3). Fix a limit point Q , and let us abuse notation by letting $(Q^n)_{n=1}^\infty$ denote a subsequence convergent to Q . Now we check the four defining properties of a weak MFG solution.

1. Define $\rho := Q \circ (W, \mu)^{-1}$. Note that $Q^n \circ W^{-1} = \mathbb{P} \circ W_1^{-1} = \lambda$, since $(W_i)_{i=1}^\infty$ are i.i.d. with law λ under \mathbb{P} . Thus, passing to the limit, $Q \circ W^{-1} = \lambda$ as well.
2. Fix any continuous function $\varphi : E \times \mathcal{P}(A) \rightarrow A$, and define

$$P := \rho \circ (W, \mu, \varphi(W, \mu))^{-1}.$$

Then P is in $\mathcal{A}(\rho)$, and we will show that $\int F dQ \geq \int F dP$ for this type of P . Indeed, any element of $\mathcal{A}(\rho)$ can be approximated by a P of this form, according to Proposition 2.1.5. Define a random measure $\hat{\mu}_{-i}^n$ on Ω by

$$\hat{\mu}_{-i}^n := \frac{1}{n-1} \sum_{k \neq i}^n \delta_{X_k^n},$$

to represent the empirical measure with the i^{th} agent removed. Set $Y_i^n = \varphi(W_i, \hat{\mu}_{-i}^n)$, and define another random measure $\hat{\mu}_i^n$ on Ω by

$$\hat{\mu}_i^n := \frac{1}{n} \left(\delta_{Y_i^n} + \sum_{k \neq i}^n \delta_{X_k^n} \right).$$

Using the ϵ_n -Nash property, we have

$$\begin{aligned} \int F dQ^n + \epsilon_n &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [F(W_i, X_i^n, \hat{\mu}^n)] + \epsilon_n \\ &\geq \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [F(W_i, Y_i^n, \hat{\mu}_i^n)] \end{aligned}$$

Since $\epsilon_n \rightarrow 0$, we have

$$\begin{aligned} \int F dQ &= \lim_{n \rightarrow \infty} \int F dQ^n \\ &\geq \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [F(W_i, Y_i^n, \hat{\mu}_i^n)] \\ &= \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}} [F(W_i, Y_i^n, \hat{\mu}_{-i}^n)]. \end{aligned} \quad (2.23)$$

The last equality holds because F is bounded and continuous and because $\hat{\mu}^n$ and $\hat{\mu}_i^n$ are close; namely, $\ell_{A,1}(\hat{\mu}^n, \hat{\mu}_i^n) \leq c/n$, where $c > 0$ is the diameter of A . Finally, continuity of φ implies

$$\mathbb{P} \circ (W_i, Y_i^n, \hat{\mu}_{-i}^n)^{-1} = \mathbb{P} \circ (W_i, \varphi(W_i, \hat{\mu}_{-i}^n), \hat{\mu}_{-i}^n)^{-1} \rightarrow P.$$

This holds uniformly in i , and we conclude from (2.23) that $\int F dQ \geq \int F dP$.

3. Let $\varphi : E \rightarrow \mathbb{R}$ and $\psi : \mathcal{P}(A) \rightarrow \mathbb{R}$ be bounded and continuous. Since ψ is bounded, the law of large numbers implies

$$\begin{aligned} \mathbb{E}^Q [(\varphi(W) - \mathbb{E}^Q[\varphi(W)]) \psi(\mu)] &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}} \left[\left(\frac{1}{n} \sum_{i=1}^n \varphi(W_i) - \mathbb{E}^{\mathbb{P}}[\varphi(W_1)] \right) \psi(\hat{\mu}^n) \right] \\ &= 0. \end{aligned}$$

4. Let $\varphi : A \rightarrow \mathbb{R}$ and $\psi : \mathcal{P}(A) \rightarrow \mathbb{R}$ be bounded and continuous. Since $\int \varphi d\hat{\mu}^n = \frac{1}{n} \sum_{i=1}^n \varphi(X_i)$, we have

$$\begin{aligned} \mathbb{E}^Q \left[\left(\varphi(X) - \int \varphi d\mu \right) \psi(\mu) \right] &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}} \left[\left(\frac{1}{n} \sum_{i=1}^n \varphi(X_i) - \int \varphi d\hat{\mu}^n \right) \psi(\hat{\mu}^n) \right] \\ &= 0. \end{aligned}$$

□

Remark 2.3.7. The properties (1) and (4) of a weak MFG solution in fact hold before the limit, i.e. for each Q^n .

There is a nice way to simplify the above result, which requires some alternative terminology. We say that $m \in \mathcal{P}(A)$ is a *strong MFG solution* if there exists $Q \in \mathcal{P}(\bar{\Omega})$ satisfying properties (1-4) of Theorem 2.3.6 as well as $Q(\mu = m) = 1$. Essentially, if we avoid making precise sense of which probability space these objects are defined on, a strong MFG solution is any $m \in \mathcal{P}(A)$ such that there exists an A -valued random variable X satisfying $\mathbb{E}[F(W, m, X)] \geq \mathbb{E}[F(W, m, X')]$ for any other random variable X' .

Proposition 2.3.8. *For any weak MFG solution Q , the measure $Q \circ \mu^{-1}$ is concentrated on the set of strong MFG solutions. In particular, if Q^n is defined as in Theorem 2.3.6, then for every limit point Q of $(Q^n)_{n=1}^\infty$, the measure $Q \circ \mu^{-1}$ is concentrated on the set of strong MFG solutions. Conversely, suppose that for each $(w, m) \in E \times \mathcal{P}(A)$ there exists a unique maximizer $\hat{x}(w, m)$ of the function $x \mapsto F(w, m, x)$. Then, for every probability measure M on $\mathcal{P}(A)$ concentrated on the set of strong MFG solutions, there exists a weak MFG solution Q with $Q \circ \mu^{-1} = M$.*

Proof. Let Q be a weak MFG solution. The function $F(w, m, \cdot)$ admits a maximizer for each (w, m) , by compactness of A . It follows from property (2) that Q is necessarily of the form

$$Q(dw, dm, dx) = \lambda(dw)M(dm)K_{w,m}(dx),$$

where $M := Q \circ \mu^{-1}$, the kernel $E \times \mathcal{P}(A) \ni (w, m) \mapsto K_{w,m} \in \mathcal{P}(A)$ is measurable, and for $\lambda \times M$ -almost every (w, m) the measure $K_{w,m}$ is concentrated on the set of maximizers of $F(w, m, \cdot)$. For each $m \in \mathcal{P}(A)$ let $\bar{K}_m \in \mathcal{P}(A)$ denote the mean measure

$$\bar{K}_m(C) := \int_E \lambda(dw)K_{w,m}(C).$$

Then properties (3) and (4) imply that

$$m = Q(X \in \cdot \mid \mu = m) = \bar{K}_m, \tag{2.24}$$

for M -almost every m in $\mathcal{P}(A)$. Now fix some such $\tilde{m} \in \mathcal{P}(A)$ satisfying (2.24). Define a measure $\tilde{Q} \in \mathcal{P}(\bar{\Omega})$ by

$$\tilde{Q}(dw, dm, dx) = \lambda(dw)\delta_{\tilde{m}}(dm)K_{w,m}(dx).$$

Trivially \tilde{Q} satisfies properties (1) and (3) of a weak MFG solution. Since \tilde{m} satisfies (2.24), it holds that \tilde{Q} satisfies property (4) as well. But \tilde{Q} also clearly satisfies property (2), since $K_{w,m}$ is concentrated on the set of maximizers of $F(w, m, \cdot)$. Thus \tilde{m} is a strong MFG solution.

Conversely, suppose $M \in \mathcal{P}(\mathcal{P}(A))$ is concentrated on the set of strong MFG solutions. First note that the maximizer \hat{x} must be jointly measurable (e.g., by [5, Theorem 18.19]). Consider the measure on $\bar{\Omega}$ given by

$$Q(d\omega, dm, dx) = \lambda(d\omega)M(dm)\delta_{\hat{x}(\omega, m)}(dx).$$

It is clear that $Q \circ W^{-1} = \lambda$ and that W and μ are independent under Q . The fixed point property (4) for strong MFG solutions implies that for M -a.e. $m \in \mathcal{P}(A)$ we have

$m = \lambda \circ \hat{x}(\cdot, m)^{-1}$. Thus

$$\mathbb{E}^Q[\varphi(X) \mid \mu = m] = \int_E \lambda(d\omega) \varphi(\hat{x}(\omega, m)) = \int_A \varphi dm,$$

for every bounded measurable $\varphi : A \rightarrow \mathbb{R}$. This shows $\mu = Q(X \in \cdot \mid \mu)$. Finally, we must check the optimality property (2). Put $\rho = Q \circ (W, \mu)^{-1}$, and fix $P \in \mathcal{A}(\rho)$. By disintegration, we may find a measurable kernel $E \times \mathcal{P}(A) \ni (\omega, m) \mapsto K_{\omega, m} \in \mathcal{P}(A)$ such that

$$P(d\omega, dm, dx) = \lambda(d\omega) M(dm) K_{\omega, m}(dx).$$

For each fixed $m \in \mathcal{P}(A)$, define $P_m \in \mathcal{P}(E \times A)$ by $P_m(d\omega, dx) = \lambda(d\omega) K_{\omega, m}(dx)$. Since M -a.e. $m \in \mathcal{P}(A)$ is a strong MFG solution, it follows that

$$\int_E \lambda(d\omega) F(\omega, m, \hat{x}(\omega, m)) \geq \int_{E \times A} P_m(d\omega, dx) F(\omega, m, x), \text{ for } M - \text{a.e. } m,$$

which implies

$$\begin{aligned} \int_{\Omega} F dQ &= \int_{\mathcal{P}(A)} M(dm) \int_E \lambda(d\omega) F(\omega, m, \hat{x}(\omega, m)) \\ &\geq \int_{\mathcal{P}(A)} M(dm) \int_{E \times A} P_m(d\omega, dx) F(\omega, m, x) \\ &= \int_{\Omega} F dP. \end{aligned}$$

This proves property (2), and so Q is a weak MFG solution. \square

Notice that we haven't shown that our description of the limits of n -player approximate Nash equilibria is sharp. That is, we don't know that *every* weak solution can be reached as the limit of ϵ_n -Nash equilibria with $\epsilon_n \rightarrow 0$. This is indeed an important point, and a good MFG solution concept should be sharp in this sense, at least for a reasonably wide variety of models. In many situations, this weak solution concept will indeed be sharp; in particular, the weak solution concept we use for stochastic differential mean field games has this feature. In any case, sharp or not, Proposition 2.3.8 supports the claim that there is often little need for weak solutions in describing the limits of n -player equilibria, in the static MFG model presented in this section.

As in Section 2.3.2, we may prove existence of a strong solution by solving a fixed point problem. Namely, let \mathcal{A} denote the set of $P \in \mathcal{P}(E \times A)$ with first marginal equal to λ . Define a set-valued map $\varphi : \mathcal{P}(A) \rightarrow 2^{\mathcal{P}(A)}$ by setting

$$\varphi(m) := \left\{ P(E \times \cdot) : \int F(w, m, x) P(dw, dx) = \sup_{Q \in \mathcal{A}} \int F(w, m, x) Q(dw, dx) \right\}.$$

That is, $\varphi(m)$ is the set of A -marginals of the set of maximizers of $\mathcal{A} \ni P \mapsto \int F(w, m, x) P(dw, dx)$. A strong solution is simply a fixed point $m \in \varphi(m)$. Since \mathcal{A} is compact, it is not hard to prove an existence result using Kakutani's theorem along the

lines of Theorem 2.3.5. Again, we emphasize that a direct existence theorem for MFG solutions is not necessarily useful from the perspective of n -player games unless we have a *converse limit theorem* of the type described in the previous paragraph, that is if we can use the MFG solutions to construct approximate Nash equilibria for the n -player games.

2.3.4 The stochastic case with common noise

Now we sketch how the idea of the previous section extends when *common noise* is present. Continue with the same setting and notation, but now suppose an additional independent random variable W_0 is given, with values in another Polish space E_0 and with law λ_0 . Suppose that A is a compact metric space and now that $F : E_0 \times E \times \mathcal{P}(A) \times A \rightarrow \mathbb{R}$ is continuous. Each agent i in the n -player game chooses an A -valued $\sigma(W_0, W_1, \dots, W_n)$ -measurable random variable X_i to try to maximize

$$J_i^n(X_1, \dots, X_n) = \mathbb{E} \left[F \left(W_0, W_i, \frac{1}{n} \sum_{k=1}^n \delta_{X_k}, X_i \right) \right].$$

Now let $\bar{\Omega} := E_0 \times E \times \mathcal{P}(A) \times A$ denote our canonical space, and let (B, W, μ, X) denote the canonical projections. Given $\rho \in \mathcal{P}(E_0 \times E \times \mathcal{P}(A))$, let $\mathcal{A}(\rho)$ denote the set of probability measures P on $\bar{\Omega}$ with $E_0 \times E \times \mathcal{P}(A)$ -marginal equal to ρ , i.e. $P \circ (B, W, \mu)^{-1} = \rho$. A *weak MFG solution* is now defined as follows:

1. Q has $E_0 \times E$ -marginal equal to $\lambda_0 \times \lambda$.
2. For each $P \in \mathcal{A}(\rho)$ where $\rho := Q \circ (B, W, \mu)^{-1}$, we have $\int F dQ \geq \int F dP$.
3. W and (B, μ) are independent under Q .
4. $Q(X \in \cdot \mid B, \mu) = \mu$, Q -almost surely.

The proof of the following limit theorem is analogous to that of Theorem 2.3.6:

Theorem 2.3.9. *Suppose for each n we are given a ϵ_n -Nash equilibrium (X_1^n, \dots, X_n^n) , where $\epsilon_n \rightarrow 0$. Let $\hat{\mu}^n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i^n}$ denote the empirical measure of these strategy profiles. Define*

$$Q^n := \frac{1}{n} \sum_{i=1}^n \mathbb{P} \circ (W_0, W_i, \hat{\mu}^n, X_i^n)^{-1}.$$

Then $(Q^n)_{n=1}^\infty \subset \mathcal{P}(\bar{\Omega})$ is tight, and every weak limit Q is a weak MFG solution.

In this setting it is natural to call a weak MFG solution Q a *strong MFG solution* if μ is a.s. B -measurable under Q , i.e. μ is measurable with respect to the Q -completion of $\sigma(B)$. In this case, the defining property (4) of a weak MFG solution reduces to $Q(X \in \cdot \mid B) = \mu$. Not only can μ fail to be a.s. B -measurable, but also X can fail to be a.s. (B, W, μ) -measurable; two layers of randomization may appear in the limit, which are in a sense *external* to the given sources (B, W) of randomness.

There are now two points which make this setting much more complicated than when common noise was absent. First of all, it is not clear that an analog of Proposition 2.3.8 is available, and thus we must stick with weak solutions in order to describe the limits of finite games. Second, an existence proof is much more complicated here. It is difficult to find the right topological spaces in which to formulate a fixed point problem, and the usual continuity-compactness tradeoff seems unassailable at first. There is a critical failure of continuity due to the operation of conditioning a joint measure. In general, if some random variables (Z_n, Y) converge weakly to (Z, Y) , then there is no reason that the conditional laws of Z_n given Y should converge in any useful sense. However, if the support of Y is *finite*, then there is no problem, as a conditional law given a finite σ -field can then be seen as a finite vector of (deterministic) probability measures. Hence, if λ_0 is nonatomic, we can obtain an existence result via the following procedure:

1. Approximate λ_0 by a sequence of measures with finite support.
2. Prove existence of strong MFG solutions for the approximate systems, exploiting the additional continuity.
3. Take weak limits to obtain a weak solution of the original problem.

This is exactly the strategy employed in Chapter 7 to prove existence of weak solutions of mean field games with common noise.

Chapter 3

Stochastic differential mean field games

This chapter describes in detail the n -player stochastic differential games under consideration as well as the associated mean field game solution concepts. We take now an approach opposite to that of the previous chapter, in the sense that we start immediately with the most general situation, with common noise, before successively refining and specializing the results. Refer back to the introduction, specifically Section 1.2.2, for an informal discussion of the mean field game with common noise. The main limit theorems and existence theorems are stated, but the proofs are postponed to Chapters 6 and 7. The equilibrium concepts for mean field games (MFGs) developed in this chapter are admittedly cumbersome. But they are central to this thesis, so we take time to thoroughly explain all of the moving parts.

3.1 MFG solution concepts

The basic inputs to the model are the following data. We are given a time horizon $T > 0$, three exponents (p', p, p_σ) with $p \geq 1$, a control space A , an initial state distribution $\lambda \in \mathcal{P}(\mathbb{R}^d)$, and the following functions:

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

The state, idiosyncratic noises, and common noise are of dimension d , m , and m_0 , respectively. The following standing assumptions are assumed to hold *throughout the thesis*, although additional assumptions will be imposed for certain results later on:

Assumption A1.

- (A1.1) The control space A is a closed convex subset of a Euclidean space. (More generally, as in [63], a closed convex σ -compact subset of a Banach space would suffice.)
- (A1.2) The exponents satisfy $p' > p \geq 1 \vee p_\sigma$ and $p_\sigma \in [0, 2]$. Moreover, assume $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$.
- (A1.3) The functions b, σ, σ_0, f , and g of (t, x, μ, a) are jointly measurable and are continuous in (x, μ, a) for each t .

(A1.4) The functions (b, σ, σ_0) are uniformly Lipschitz in x . That is, there exists $c_1 > 0$ such that, for all $(t, x, y, \mu, \nu, a) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

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

and

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

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

(A1.5) There exist $c_2, c_3 > 0$ such that, for each $(t, x, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

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

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

$$f(t, x, \mu, a) \geq -c_2 \left(1 + |x|^p + \int_{\mathbb{R}^d} |z|^p \mu(dz) + |a|^{p'} \right).$$

For simplicity, the initial distribution λ on \mathbb{R}^d is fixed throughout, and the state processes of the n -player games are assumed to be independent and identically distributed (i.i.d.) with common law λ . It should be clear how the definitions to follow depend on the choice of initial condition. Remark 3.2.6 will explain how to extend the main limit theorem beyond the i.i.d. setting.

A typical case is $p' = 2$, $p = 1$, and $p_\sigma = 0$ (i.e., σ and σ_0 bounded). Unfortunately, our assumptions do not cover all linear-quadratic models. When the objective f is quadratic in the control a , we are forced to choose $p' = 2$, and the constraint $p < p'$ forces f and g to be strictly subquadratic in x . However, this is not surprising in light of the counterexample discussed in Section 7.4.

For all of the results except for one existence theorem, it is assumed that the volatility coefficients are uncontrolled. This is not so benign an assumption; several arguments involving relaxed controls are simply unavailable when the volatilities are controlled. We assume throughout this chapter that the following assumption holds, in addition to assumption A1.

Assumption A2. The volatilities σ and σ_0 are uncontrolled.

Before defining the various solution concepts for MFGs, let us define some additional canonical spaces. For positive integers k let $\mathcal{C}^k = C([0, T]; \mathbb{R}^k)$ denote the set of continuous functions from $[0, T]$ to \mathbb{R}^k , endowed with the supremum norm and its Borel σ -field. For $\mu \in \mathcal{P}(\mathcal{C}^k)$, let $\mu_t \in \mathcal{P}(\mathbb{R}^k)$ denote the image of μ under the map $x \mapsto x_t$. Let

$$\mathcal{X} := \mathcal{C}^m \times \mathcal{V} \times \mathcal{C}^d. \tag{3.1}$$

This space will house the idiosyncratic noise, the relaxed control, and the state process. Let $(\mathcal{F}_t^\mathcal{X})_{t \in [0, T]}$ denote the canonical filtration on \mathcal{X} , where $\mathcal{F}_t^\mathcal{X}$ is the σ -field generated by the maps

$$\mathcal{X} \ni (w, q, x) \mapsto (w_s, x_s, q(C)) \in \mathbb{R}^m \times \mathbb{R}^d \times \mathbb{R}, \text{ for } s \leq t, C \in \mathcal{B}([0, t] \times A).$$

For $\mu \in \mathcal{P}(\mathcal{X})$, let

$$\mu^x := \mu(\mathcal{C}^m \times \mathcal{V} \times \cdot) \quad (3.2)$$

denote the \mathcal{C}^d -marginal. Finally, for ease of notation let us define the objective functional $\Gamma : \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R} \cup \{-\infty\}$ by

$$\Gamma(\mu, q, x) := \int_0^T \int_A f(t, x_t, \mu_t, a) q_t(da) dt + g(x_T, \mu_T). \quad (3.3)$$

Note that Γ is well-defined, thanks to the growth assumptions [A1.5](#) pertaining to f and g .

Definition 3.1.1. A *weak MFG solution with weak control*, or simply a *weak MFG solution*, is a tuple $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$, where $(\tilde{\Omega}, \mathcal{F}_T, \mathbb{F}, P)$ is a complete filtered probability space supporting (B, W, μ, Λ, X) satisfying

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

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

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

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

4. The state equation holds:

$$dX_t = \int_A b(t, X_t, \mu_t^x, a) \Lambda_t(da) dt + \sigma(t, X_t, \mu_t^x) dW_t + \sigma_0(t, X_t, \mu_t^x) dB_t. \quad (3.4)$$

5. If $(\tilde{\Omega}', \mathbb{F}', P')$ is another filtered probability space supporting $(B', W', \mu', \Lambda', X')$ satisfying (1-4) and $P \circ (B, \mu)^{-1} = P' \circ (B', \mu')^{-1}$, then

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

6. μ is a version of the conditional law of (W, Λ, X) given (B, μ) .

If also there exists an A -valued process α such that $P(\Lambda_t = \delta_{\alpha_t} \text{ a.e. } t) = 1$, then we say the weak MFG solution has *strict control*. If this α is progressively measurable with respect to the completion of $(\mathcal{F}_t^{X_0, B, W, \mu})_{t \in [0, T]}$, we say the weak MFG solution has *strong control*. If μ is a.s. B -measurable, then we have a *strong MFG solution* (with either weak control, strict control, or strong control).

Given a weak MFG solution $(\tilde{\Omega}, \mathbb{F}, P, B, W, \mu, \Lambda, X)$, we view $(X_0, B, W, \mu, \Lambda, X)$ as a random element of the canonical space

$$\Omega := \mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V} \times \mathcal{C}^d. \quad (3.5)$$

A weak MFG solution thus induces a probability measure on Ω , which itself we would like to call a MFG solution, as it is really the object of interest more than the particular probability space. The following definition will be reformulated in Section 5.2 in a more intrinsic manner.

Definition 3.1.2. If $P \in \mathcal{P}(\Omega)$ satisfies $P = P' \circ (X_0, B, W, \mu, \Lambda, X)^{-1}$ for some weak MFG solution $(\Omega', \mathbb{F}', P', B, W, \mu, \Lambda, X)$, then we refer to P itself as a *weak MFG solution*. Naturally, we may also refer to P as a weak MFG solution with strict control or strong control, or as a strong MFG solution, under the analogous additional assumptions.

The definition is to be interpreted as follows. The first conditions (1-3) are largely technical, but we will elaborate notion of compatibility shortly. Remark 3.2.6 will explain why X_0 is independent of μ in point (2). The fourth condition simply says that the given processes verify the correct state equation. The fifth is an optimality condition stated, requiring that the given control Λ achieves a greater reward than any other compatible control; the agent is allowed to change the underlying probability space to construct an alternative control, but the joint distribution of the inputs (X_0, B, W, μ) of the control problem must remain intact. The final condition (6) is the *consistency* condition.

As we saw in Sections 2.2 and 2.3, dealing respectively with interacting diffusions and static mean field games, some care is needed in defining a notion of equilibrium that captures a given mean field limit. This is because weak convergence does not preserve measurability properties, as is strikingly illustrated by Proposition 2.1.5. Even if we require the strategies of the n -player games to be adapted to a given filtration (e.g., generated by the driving Wiener processes), there may appear additional randomness in the $n \rightarrow \infty$ limit. For this reason, we allow both the random measure μ and the control Λ to be randomized externally to the inputs (X_0, B, W) . Following the terminology of weak and strong solutions of SDEs, we call the MFG solution strong if μ happens to be B -adapted and weak otherwise. Similarly, a strong control is adapted to the filtration $(\mathcal{F}_t^{X_0, B, W, \mu})_{t \in [0, T]}$ generated by the inputs to the control problem, whereas a weak control may not be.

This precise notion of “weak control” is unusual. It is not surprising that we need measure-valued controls, which we interpret as a continuous-time form of *mixed strategy*, but the compatibility condition (3) deserves some discussion. Compatibility is actually well known by various names in diverse areas of stochastic analysis, and there are several interesting equivalent definitions. Appendix A.1 elaborates on this, but we will prove all of the needed properties in the body of the thesis, to keep the presentation as self-contained as

possible. Strong controls are more prevalent in the literature, but, as discussed above, weak limits of strong controls are not guaranteed to remain strong controls, as is illustrated by Proposition 2.1.15. On the other extreme, it seems more natural at first to omit compatibility and merely require that Λ is \mathbb{F} -progressively measurable; this class of controls does indeed catch any relevant weak limits of n -player strategies, but requiring *optimality* of a limiting control among this class turns out to be too restrictive. The compatibility requirement falls between these two extremes, and mathematically the best intuition for why compatibility should be relevant comes from Proposition 2.1.15: This says that the set of compatible controls is precisely the *closure* of the set of strong controls, in a certain topological sense.

A more game-theoretic interpretation of compatibility is as follows. An agent has full information, in the sense that she observes (in an adapted fashion) the initial state X_0 , the noises B and W , and also the distribution μ of the (infinity of) other agents' states, controls, and noises. That is, the agent has access to $\mathcal{F}_t^{X_0, B, W, \mu}$ at time t . Controls are allowed to be randomized externally to these observations, but such a randomization must be *conditionally independent of future information given current information*.

Remark 3.1.3. Given a MFG solution, the *consistency condition* (6) implies that $\mu_t^x = P(X_t \in \cdot \mid \mathcal{F}_t^{B, \mu^x})$ for each t , where

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

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

Somewhat less clear is why we must work with the full conditional law of (W, Λ, X) in the consistency condition (6). Indeed, as in Remark (3.1.3), it seems more natural to work simply with the X -marginals, and to require that $\mu^x = P(X \in \cdot \mid B, \mu^x)$. After all, only μ^x appears in the state equation and the objective functions. Mathematically, the compatibility required of the control prevents such a simplification, and even our method of proof depends crucially on the use of the full joint law; see Lemma 5.1.2. Intuitively, the full conditional law μ carries more information than μ^x , particularly pertaining to correlations between the states, controls, and idiosyncratic noises of the other agents. The example of Section 4.3 (specifically Remark 4.3.4) elaborates on this point.

The definition of weak MFG solution is a bit complicated, but additional convexity and concavity assumptions translate into additional structure of weak MFG solutions. For example, we can rule out relaxed controls under the following convexity assumption, which is familiar in control theory ever since Filippov's work [50].

Assumption (Convex). For each $(t, x, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d)$, the subset

$$\{(b(t, x, \mu, a), (\sigma\sigma^\top + \sigma_0\sigma_0^\top)(t, x, \mu, a), z) : a \in A, z \leq f(t, x, \mu, a)\}$$

of $\mathbb{R}^d \times \mathbb{R}^{d \times d} \times \mathbb{R}$ is convex.

Proposition 3.1.4. *Under assumption **(Convex)**, for every weak MFG solution with weak control, there exists a weak MFG solution with strict control which induces the same joint law of (B, W, μ^x, X) . More precisely, if $(\tilde{\Omega}, \mathbb{F}, P, B, W, \mu, \Lambda, X)$ is a weak solution with weak control, there exists a strict control $\tilde{\Lambda}$ and a random element $\tilde{\mu}$ of $\mathcal{P}^p(\mathcal{X})$, both defined on $\tilde{\Omega}$, such that $(\tilde{\Omega}, \mathbb{F}, P, B, W, \tilde{\mu}, \tilde{\Lambda}, X)$ is a weak solution with strict control, and $\tilde{\mu}^x = \mu^x$ a.s.*

Note, for future reference, that when the volatilities are uncontrolled (i.e., when assumption **A2** is in place), the $\mathbb{R}^{d \times d}$ term can be omitted from the set of assumption **(Convex)**. A typical special case of assumption **(Convex)** is detailed in the following stronger assumption, under which a stronger conclusion is possible.

Assumption (Linear-Convex).

1. The state coefficients are affine in (x, a) , in the following form:

$$\begin{aligned} b(t, x, \mu, a) &= b^1(t, \mu)x + b^2(t, \mu)a + b^3(t, \mu), \\ \sigma(t, x, \mu) &= \sigma^1(t, \mu)x + \sigma^2(t, \mu), \quad \sigma_0(t, x, \mu) = \sigma_0^1(t, \mu)x + \sigma_0^1(t, \mu), \end{aligned}$$

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

Proposition 3.1.5. *Under assumption **(Linear-Convex)**, every weak MFG solution with weak control is in fact a weak MFG solution with strong control.*

The proofs of Propositions (3.1.4) and (3.1.5) are deferred to Section 5.4.

3.2 Limits of finite games

Let us now return to defining the n -player games precisely. As before, we assume throughout that assumptions **A1** and **A2** are in force. In this section, however, we also assume the following:

Assumption A3. The exponent p' is at least 2, and (b, σ, σ_0) are uniformly Lipschitz in (x, μ) in the sense that, for all $(t, x, y, \mu, \nu, a) \in [0, T] \times \mathbb{R}^d \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times \mathcal{P}^p(\mathbb{R}^d) \times A$,

$$|(b, \sigma, \sigma_0)(t, x, \mu, a) - (b, \sigma, \sigma_0)(t, y, \nu, a)| \leq c_1 (|x - y| + \ell_{\mathbb{R}^d, p}(\mu, \nu)).$$

An n -player environment is defined to be any tuple

$$\mathcal{E}_n = (\Omega_n, \mathbb{F}^n = (\mathcal{F}_t^n)_{t \in [0, T]}, \mathbb{P}_n, \xi, B, W),$$

where $(\Omega_n, \mathcal{F}_T^n, \mathbb{F}^n, \mathbb{P}_n)$ is a complete filtered probability space supporting an \mathcal{F}_0^n -measurable $(\mathbb{R}^d)^n$ -valued random variable $\xi = (\xi^1, \dots, \xi^n)$ with law $\lambda^{\times n}$, an m_0 -dimensional $(\mathcal{F}_t^n)_{t \in [0, T]}$ -Wiener process B , and a nm -dimensional \mathbb{F}^n -Wiener process $W = (W^1, \dots, W^n)$, independent of B . For simplicity, we consider i.i.d. initial states ξ^1, \dots, ξ^n with common law λ , although it is possible to generalize this (see Remark 3.2.6). Perhaps all of the notation here should be parametrized by \mathcal{E}_n or an additional index for n , but, since we will typically focus

on a fixed sequence of environments $(\mathcal{E}_n)_{n=1}^\infty$, we avoid complicating the notation. Indeed, the subscript n on the measure \mathbb{P}_n will be enough to remind us on which environment we are working at any moment.

Until further notice, we work with a fixed n -player environment \mathcal{E}_n . An *admissible (relaxed) control* is any \mathbb{F}^n -progressively measurable $\mathcal{P}(A)$ -valued process Λ satisfying

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

An *admissible strategy* is a vector of n admissible controls. The set of admissible controls is denoted $\mathcal{A}_n(\mathcal{E}_n)$, and accordingly the set of admissible strategies is the Cartesian product $\mathcal{A}_n^n(\mathcal{E}_n)$. A *strict control* is any control $\Lambda \in \mathcal{A}_n(\mathcal{E}_n)$ such that $\mathbb{P}_n(\Lambda_t = \delta_{\alpha_t}, \text{ a.e. } t) = 1$ for some \mathbb{F}^n -progressively measurable A -valued process α , and a *strict strategy* is any vector of n strict controls. Given an admissible control $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$ define the state processes $X[\Lambda] := (X^1[\Lambda], \dots, X^n[\Lambda])$ by

$$\begin{aligned} dX_t^i[\Lambda] &= \int_A b(t, X_t^i[\Lambda], \hat{\mu}_t^x[\Lambda], a) \Lambda_t^i(da) dt + \sigma(t, X_t^i[\Lambda], \hat{\mu}_t^x[\Lambda]) dW_t^i \\ &\quad + \sigma_0(t, X_t^i[\Lambda], \hat{\mu}_t^x[\Lambda]) dB_t, \quad X_0^i = \xi^i, \\ \hat{\mu}^x[\Lambda] &:= \frac{1}{n} \sum_{k=1}^n \delta_{X^k[\Lambda]}. \end{aligned}$$

Note that we abbreviate

$$\hat{\mu}_t^x[\Lambda] := (\hat{\mu}^x[\Lambda])_t = \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k[\Lambda]}$$

Assumption **A3** ensures that a unique strong solution of this SDE system exists.¹ Indeed, the Lipschitz assumption and the obvious inequality

$$\ell_{\mathbb{R}^d, p} \left(\frac{1}{n} \sum_{i=1}^n \delta_{x_i}, \frac{1}{n} \sum_{i=1}^n \delta_{y_i} \right) \leq \left(\frac{1}{n} \sum_{i=1}^n |x_i - y_i|^p \right)^{1/p}$$

together imply, for example, that the function

$$(\mathbb{R}^d)^n \ni (x_1, \dots, x_n) \mapsto b \left(t, x_1, \frac{1}{n} \sum_{i=1}^n \delta_{x_i}, a \right) \in \mathbb{R}^d$$

is Lipschitz, uniformly in (t, a) .

¹ The filtrations are required to be complete or right-continuous. No problems will arise in the stochastic integration, thanks to the careful treatment of Stroock and Varadhan [102, Section 4.3].

3.2.1 Equilibrium concepts

The value for player i corresponding to a strategy $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$ is defined by

$$J_i(\Lambda) := \mathbb{E}^{\mathbb{P}_n} \left[\int_0^T f(t, X_t^i[\Lambda], \hat{\mu}_t^x[\Lambda], a) \Lambda_t^i(da) dt + g(X_T^i[\Lambda], \hat{\mu}_T^x[\Lambda]) \right].$$

A standard estimate using assumption (A1.4) (proven in Lemma 6.1.1), shows that

$$\mathbb{E}^{\mathbb{P}_n} \left[\sup_{t \in [0, T]} |X_t^i[\Lambda]|^p \right] < \infty$$

for each $\Lambda \in \mathcal{A}_n^n(\mathcal{E}_n)$, $n \geq i \geq 1$. Thus $J_i(\Lambda) < \infty$ is well-defined because of the upper bounds of assumption (A1.5), although it is possible that $J_i(\Lambda) = -\infty$ since we do not require that an admissible control possess a finite moment of order p' . Given a strategy $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$ and a control $\beta \in \mathcal{A}_n(\mathcal{E}_n)$, define a new strategy $(\Lambda^{-i}, \beta) \in \mathcal{A}_n^n(\mathcal{E}_n)$ by

$$(\Lambda^{-i}, \beta) = (\Lambda^1, \dots, \Lambda^{i-1}, \beta, \Lambda^{i+1}, \dots, \Lambda^n).$$

Given $\epsilon = (\epsilon_1, \dots, \epsilon_n) \in [0, \infty)^n$, a *relaxed ϵ -Nash equilibrium in \mathcal{E}_n* is any strategy $\Lambda \in \mathcal{A}_n^n(\mathcal{E}_n)$ satisfying

$$J_i(\Lambda) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{-i}, \beta)) - \epsilon_i, \quad i = 1, \dots, n.$$

Naturally, if $\epsilon_i = 0$ for each $i = 1, \dots, n$, we use the simpler term *Nash equilibrium*, as opposed to *0-Nash equilibrium*. A *strict ϵ -Nash equilibrium in \mathcal{E}_n* is any *strict* strategy $\Lambda \in \mathcal{A}_n^n(\mathcal{E}_n)$ satisfying

$$J_i(\Lambda) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n) \text{ strict}} J_i((\Lambda^{-i}, \beta)) - \epsilon_i, \quad i = 1, \dots, n.$$

Note that the optimality is required only among *strict* controls.

Note that the role of the filtration \mathbb{F}^n in the environment \mathcal{E}_n is mainly to specify the class of admissible controls. We are particularly interested in the sub-filtration generated by the Wiener processes and initial states; define $\mathbb{F}^{s,n} = (\mathcal{F}_t^{s,n})_{t \in [0, T]}$ to be the \mathbb{P}_n -completion of

$$(\sigma(\xi, B_s, W_s : s \leq t))_{t \in [0, T]}.$$

Of course, $\mathcal{F}_t^{s,n} \subset \mathcal{F}_t^n$ for each t . Let us say that $\Lambda \in \mathcal{A}_n(\mathcal{E}_n)$ is a *strong control* if $\mathbb{P}_n(\Lambda_t = \delta_{\alpha_t} \text{ a.e. } t) = 1$ for some $\mathbb{F}^{s,n}$ -progressively measurable A -valued process α . Naturally, a *strong strategy* is a vector of strong controls. A *strong ϵ -Nash equilibrium in \mathcal{E}_n* is any *strong* strategy $\Lambda \in \mathcal{A}_n^n(\mathcal{E}_n)$ such that

$$J_i(\Lambda) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n) \text{ strong}} J_i((\Lambda^{-i}, \beta)) - \epsilon_i, \quad i = 1, \dots, n.$$

Remark 3.2.1. A strong ϵ -Nash equilibrium in $\mathcal{E}_n = (\Omega_n, \mathbb{F}^n, \mathbb{P}_n, \xi, B, W)$ is equivalently a strict ϵ -Nash equilibrium in $\tilde{\mathcal{E}}_n := (\Omega_n, \mathbb{F}^{s,n}, \mathbb{P}_n, \xi, B, W)$.

The most common type of Nash equilibrium considered in the literature is, in our terminology, a strong Nash equilibrium. The next proposition assures us that our equilibrium concept using relaxed controls (and general filtrations) truly generalizes this more standard situation, thus permitting a unified analysis of all of the equilibria described thusfar. The proof is deferred to Appendix 6.4.1.

Proposition 3.2.2. *On any n -player environment \mathcal{E}_n , every strong ϵ -Nash equilibrium is also a strict ϵ -Nash equilibrium, and every strict ϵ -Nash equilibrium is also a relaxed ϵ -Nash equilibrium.*

Remark 3.2.3. Another common type of strategy in dynamic game theory is called *closed-loop*. Whereas our strategies (also called *open-loop*) are specified by *processes*, a closed-loop (strict) strategy is specified by feedback functions $\varphi_i : [0, T] \times (\mathbb{R}^d)^n \rightarrow A$, for $i = 1, \dots, n$, to be evaluated along the path of the state process. In the model of Carmona et al. [35], both the open-loop and closed-loop equilibria are computed explicitly for the n -player games, and they are shown to converge to the same MFG limit. There is no distinction between open-loop and closed-loop in the MFG, and this begs the question of whether or not closed-loop equilibria converge to the same MFG limit that obtained in Theorem 3.2.4. This thesis does not attempt to answer this question.

3.2.2 The limit theorem

We are ready now to state the first main Theorem 3.2.4 and its corollaries. The proof is deferred to Section 6.1. Given an admissible strategy $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$ defined on some n -player environment $\mathcal{E}_n = (\Omega_n, \mathbb{F}^n = (\mathcal{F}_t^n)_{t \in [0, T]}, \mathbb{P}_n, \xi, B, W)$, define (on Ω_n) the random element $\hat{\mu}[\Lambda]$ of $\mathcal{P}^p(\mathcal{X})$ (recalling the definition of \mathcal{X} from (3.1)) by

$$\hat{\mu}[\Lambda] := \frac{1}{n} \sum_{i=1}^n \delta_{(W^i, \Lambda^i, X^i[\Lambda])}.$$

Note that this is consistent with the notation of (3.2), i.e. $\hat{\mu}^x[\Lambda] = (\hat{\mu}[\Lambda])^x$. As usual, we identify a $\mathcal{P}(A)$ -valued process $(\Lambda_t^i)_{t \in [0, T]}$ with the random element $\Lambda^i = dt \Lambda_t^i(da)$ of \mathcal{V} .

Theorem 3.2.4. *Suppose assumptions, A1, A2, and A3 hold. For each n , let $\epsilon^n = (\epsilon_1^n, \dots, \epsilon_n^n) \in [0, \infty)^n$, and let $\mathcal{E}_n = (\Omega_n, (\mathcal{F}_t^n)_{t \in [0, T]}, \mathbb{P}_n, \xi, B, W)$ be any n -player environment. Assume*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \epsilon_i^n = 0. \quad (3.6)$$

Suppose for each n that $\Lambda^n = (\Lambda^{n,1}, \dots, \Lambda^{n,n}) \in \mathcal{A}_n^n(\mathcal{E}_n)$ is a relaxed ϵ^n -Nash equilibrium, and let

$$P_n := \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (\xi^i, B, W^i, \hat{\mu}[\Lambda^n], \Lambda^{n,i}, X^i[\Lambda^n])^{-1}. \quad (3.7)$$

Then $(P_n)_{n=1}^\infty$ is relatively compact in $\mathcal{P}^p(\Omega)$, and each limit point is a weak MFG solution.

Remark 3.2.5. As in Remark 2.2.5, averaging over $i = 1, \dots, n$ in (3.7) circumvents the problem that the strategies $(\Lambda^{n,1}, \dots, \Lambda^{n,n})$ need not be exchangeable. Note also that the limiting behavior of $\mathbb{P}_n \circ (B, \hat{\mu}[\Lambda^n])^{-1}$ can always be recovered from that of P_n . To interpret the definition of P_n , note that we may again write

$$P_n = \mathbb{P}_n \circ (\xi^{U_n}, B, W^{U_n}, \hat{\mu}[\Lambda^n], \Lambda^{n,U_n}, X^{U_n}[\Lambda^n])^{-1},$$

where U_n is a random variable independent of \mathcal{F}_T^n , uniformly distributed among $\{1, \dots, n\}$, constructed by extending the probability space Ω_n . In words, P_n is the joint law of the processes relevant to a *randomly selected representative agent*. Of course, Theorem 3.2.4 specializes when there is exchangeability, in the following sense. For any set E , any element $e = (e^1, \dots, e^n) \in E^n$, and any permutation π of $\{1, \dots, n\}$, let $e_\pi := (e^{\pi(1)}, \dots, e^{\pi(n)})$. If

$$\mathbb{P}_n \circ (\xi_\pi, B, W_\pi, \Lambda_\pi^n)^{-1}$$

is independent of the choice of permutation π , then so is

$$\mathbb{P}_n \circ (\xi_\pi, B, W_\pi, \hat{\mu}[\Lambda_\pi^n], \Lambda_\pi^n, X[\Lambda_\pi^n]_\pi)^{-1}.$$

It then follows that

$$P_n = \mathbb{P}_n \circ (\xi^k, B, W^k, \hat{\mu}[\Lambda^n], \Lambda^{n,k}, X^k[\Lambda^n])^{-1}, \text{ for } n \geq k.$$

Remark 3.2.6. The assumption that the initial states ξ^i are i.i.d. is not strictly necessary. As is common in the literature on McKean-Vlasov limits [53], a form of Theorem 3.2.4 still holds assuming merely that

$$\mathbb{P}_n \circ \left(\frac{1}{n} \sum_{i=1}^n \delta_{\xi^i} \right)^{-1}$$

are tight. The only difference is that in the definition of a weak MFG solution, instead of requiring X_0 , W , and (B, μ) to be independent, we instead require only that W and (X_0, B, μ) are independent; the reason for this becomes clear in the second step of the proof of Lemma 6.1.5. To keep track of this point throughout the thesis would only add unnecessarily to the already heavy notation.

Theorem 3.2.4 is stated in quite a bit of generality, devoid even of standard convexity assumptions on the objective functions f and g . Theorem 3.2.4 includes quite degenerate cases, such as the case of *no objectives*, where $f \equiv g \equiv 0$ and A is compact. In this case, *any strategy profile* whatsoever in the n -player game is a Nash equilibrium, and any weak control can arise in the limit. This explains why such a relaxed solution concept is needed in Theorem 3.2.4. Under additional convexity assumptions, however, we saw in Propositions 3.1.4 and 3.1.5 that weak MFG solutions (and thus limits of n -player equilibria) admit more refined descriptions. Let us see how to apply these results to strengthen the conclusions of Theorem 3.2.4.

Corollary 3.2.7. *Suppose the assumptions of Theorem 3.2.4 hold, as well as assumption **(Convex)**. Then*

$$\left\{ \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (B, W^i, \hat{\mu}^x[\Lambda^n], X^i[\Lambda^n])^{-1} : n \geq 1 \right\}$$

is relatively compact in $\mathcal{P}(\mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}(\mathcal{C}^d) \times \mathcal{C}^d)$, and every limit is of the form $P \circ (B, W, \mu^x, X)^{-1}$, for some weak MFG solution with strict control $(\tilde{\Omega}, \mathbb{F}, P, B, W, \mu, \Lambda, X)$.

Proof. This follows immediately from Theorem 3.2.4 and Proposition 3.1.4. \square

Corollary 3.2.8. *Suppose the assumptions of Theorem 3.2.4 hold, as well as assumption **(Linear-Convex)**. Then $(P_n)_{n=1}^\infty$ is relatively compact in $\mathcal{P}(\Omega)$, and every limit point is a weak MFG solution with strong control.*

Proof. By Proposition 3.1.5, the present assumptions guarantee that every weak MFG solution is a weak MFG solution with strong control. The claim then follows from Theorem 3.2.4. \square

Finally, we provide an example of the satisfying situation in which there is a unique MFG solution. Say that *uniqueness in law* holds for the MFG if any two weak MFG solutions induce the same law on Ω . The following corollary is an immediate consequence of Theorem 3.2.4 and the uniqueness Theorem 3.3.5 to be developed later.

Corollary 3.2.9. *Suppose the assumptions of Corollary 3.2.8 hold, and define P_n as in (3.7). Assume also that*

1. b, σ , and σ_0 have no mean field term, i.e. no μ dependence,
2. f is of the form $f(t, x, \mu, a) = f_1(t, x, a) + f_2(t, x, \mu)$,
3. For each $\mu, \nu \in \mathcal{P}(\mathcal{C}^d)$ we have

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

Then there exists a unique in law weak MFG solution, and it is a strong MFG solution with strong control. In particular, P_n converges in $\mathcal{P}(\Omega)$ to this unique MFG solution.

3.2.3 The converse limit theorem

This section states and discusses a converse to Theorem 3.2.4. For this, we need an additional technical assumption, which we note holds automatically under assumption **A1** in the case that the control space A is compact.

Assumption A4. The function f of (t, x, μ, a) is continuous in (x, μ) , uniformly in a , for each $t \in [0, T]$. That is,

$$\lim_{(x', \mu') \rightarrow (x, \mu)} \sup_{a \in A} |f(t, x', \mu', a) - f(t, x, \mu, a)| = 0, \quad \forall t \in [0, T].$$

Moreover, there exists $c_4 > 0$ such that, for all (t, x, x', μ, μ', a) ,

$$|f(t, x', \mu', a) - f(t, x, \mu, a)| \leq c_4 \left(1 + |x'|^p + |x|^p + \int_{\mathbb{R}^d} |z|^p (\mu' + \mu)(dz) \right).$$

Theorem 3.2.10. *Suppose assumptions **A1**, **A2**, **A3**, and **A4** hold. Let $P \in \mathcal{P}(\Omega)$ be a weak MFG solution, and for each n let $\mathcal{E}_n = (\Omega_n, \mathbb{F}^n, \mathbb{P}_n, \xi, B, W)$ be any n -player environment. Then there exist $\epsilon_n \geq 0$ and a strong $(\epsilon_n, \dots, \epsilon_n)$ -Nash equilibrium $\Lambda^n = (\Lambda^{n,1}, \dots, \Lambda^{n,n})$ on \mathcal{E}_n , such that $\lim_{n \rightarrow \infty} \epsilon_n = 0$ and*

$$P = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (\xi^i, B, W^i, \hat{\mu}[\Lambda^n], \Lambda^{n,i}, X^i[\Lambda^n])^{-1}, \text{ in } \mathcal{P}^p(\Omega). \quad (3.8)$$

Combining Theorems 3.2.4 and 3.2.10 shows that the set of weak MFG solutions is exactly the set of limits of (strong) approximate Nash equilibria. More precisely, the set of weak MFG solutions is exactly the set of limits

$$\lim_{k \rightarrow \infty} \frac{1}{n_k} \sum_{i=1}^{n_k} \mathbb{P}_{n_k} \circ (\xi^i, B, W^i, \hat{\mu}[\Lambda^{n_k}], \Lambda^{n_k,i}, X^i[\Lambda^{n_k}])^{-1},$$

where $\Lambda^n \in \mathcal{A}_n^n(\mathcal{E}_n)$ are strong ϵ^n -Nash equilibria and $\epsilon^n = (\epsilon_1^n, \dots, \epsilon_n^n) \in [0, \infty)^n$ satisfies (3.6). The same statement is true when the word “strong” is replaced by “strict” or “relaxed”, because of Proposition 3.2.2. Similarly, combining Theorem 3.2.10 with Corollaries 3.2.7 and 3.2.8 yields characterizations of the mean field limit without recourse to relaxed controls.

Remark 3.2.11. In light of Remark 3.2.1, the statement of Theorem 3.2.10 is insensitive to the choice of environments \mathcal{E}_n . Without loss of generality, they may all be assumed to satisfy $\mathbb{F}^n = \mathbb{F}^{s,n}$; that is, the filtration may be taken to be the one generated by the process $(\xi, B_t, W_t)_{t \in [0, T]}$.

Remark 3.2.12. It follows from the proofs of Theorems 3.2.4 and 3.2.10 that the *values* converge as well, in the sense that $\frac{1}{n} \sum_{i=1}^n J_i(\Lambda^n)$ converges (along a subsequence in the case of Theorem 3.2.4) to the corresponding optimal value corresponding to the MFG solution.

Remark 3.2.13. Theorem 3.2.10 is admittedly abstract, and not as strong in its conclusion as the typical results of this nature in the literature. Namely, in the setting without common noise, it is usually argued as in [67] that a MFG solution may be used to construct not just any sequence of approximate equilibria, but rather one consisting of *symmetric distributed strategies*, in which the control of agent i is of the form $\hat{\alpha}(t, X_t^i)$ for some function $\hat{\alpha}$ which depends neither on the agent i nor the number of agents n . The techniques of this paper seem too abstract to yield a result of this nature, but in any case this would stray from the objective of the paper. On a somewhat related note, at the level of generality of Theorem 3.2.10 we do not expect to obtain a rate of convergence of ϵ_n , as in [78, 32].

3.3 Existence and uniqueness

This section summarizes the main existence and uniqueness theorems for mean field games with common noise. The proof of the following existence theorem is the main subject of Chapter 7.

Theorem 3.3.1. *Under assumptions **A1** and **A2**, there exists a weak MFG solution (with weak control). If also assumption **(Convex)** holds, there exists a weak MFG solution with strict control.*

Exactly analogous to the theory of SDEs, we will define two notions of uniqueness. The terminology introduced here will appear again only in Sections 4.6 and 7.3. The first definition is quite natural:

Definition 3.3.2. An MFG is *unique in law* if any two weak solutions induce the same law on Ω , i.e. the law of (B, μ, W, Λ, X) .

To define pathwise uniqueness properly requires a bit more care. The starting point is to notice that the law of a weak MFG solution is really determined by the law of (B, μ) . Indeed, for an element $\gamma \in P^p(\mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}))$, we can define $M\gamma \in \mathcal{P}(\Omega)$ by

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

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

1. $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q)$ is a probability space with a complete filtration.
2. B is a $(\mathcal{G}_t)_{t \in [0, T]}$ -Wiener process on Θ .
3. For each $t \in [0, T]$, we have (up to null sets)

$$\mathcal{G}_t = \sigma(B_s, \mu^1(C), \mu^2(C) : s \leq t, C \in \mathcal{F}_t^{\mathcal{X}}).$$

4. For $i = 1, 2$, $Q \circ (B, \mu^i)^{-1} = \gamma^i$.

An *independent coupling* of γ^1 and γ^2 is any coupling of the two satisfying the additional property

5. μ^1 and μ^2 are conditionally independent given B .

Definition 3.3.3. We say *pathwise uniqueness* (resp. independent pathwise uniqueness) holds for the MFG if, for any coupling (resp. independent coupling) $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ of any two MFG solution bases, we have $\mu^1 = \mu^2$ a.s.

The following Proposition is analogous to the famous theorem of Yamada and Watanabe.

Proposition 3.3.4. *Suppose assumptions **A1** and **A2** hold, and suppose independent pathwise uniqueness holds for the MFG. Then uniqueness in law and pathwise uniqueness hold as well, and the unique weak MFG solution is in fact a strong solution with weak control.*

Using this result, we will prove a modest uniqueness result, inspired by the work of Lasry and Lions [91]. When there is no mean field term in the state coefficients, when the optimal controls are unique, and when the monotonicity condition of Lasry and Lions [91] holds, we indeed have a form of *pathwise uniqueness*.

Assumption U.

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

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

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

$$\int_{\mathcal{C}^d} (\mu - \nu)(dx) \left[g(x_T, \mu_T) - g(x_T, \nu_T) + \int_0^T (f_2(t, x_t, \mu_t) - f_2(t, x_t, \nu_t)) dt \right] \leq 0. \quad (3.9)$$

Theorem 3.3.5. *Suppose assumptions **A1**, **(Linear-Convex)**, and **U** hold. Then the MFG is pathwise unique. In particular, there exists a unique in law weak MFG solution with weak control, and it is in fact a strong MFG solution with strong control.*

A somewhat more general form of this theorem will be presented, along with the proofs of Proposition 3.3.4 and Theorem 3.3.5, in Section 7.3.

Chapter 4

Stochastic differential mean field games without common noise

The goal of this section is to specialize the limit theorem to MFGs without common noise. Indeed *we assume that $\sigma_0 \equiv 0$ throughout this section*. The assumption of Theorem 3.2.4 permits degenerate volatility, but when $\sigma_0 \equiv 0$ our general definition of weak MFG solution still involves the common noise B , which in a sense should no longer play any role. To be absolutely clear, we will rewrite the definitions and the two main theorems so that they do not involve a common noise; most notably, the notion of *strong controls* for the finite-player games is refined to *very strong controls*.

The case without common noise is also worth treating separately because better existence results are available. These will be discussed in Section 4.6.

4.1 MFG solution concepts

The proofs of the limit theorems of this chapter are deferred to Section 6.3, where we will see how to deduce almost all of the results *without* common noise from those *with* common noise. Crucially, even without common noise, a weak MFG solution still involves a *random* measure μ , and the consistency condition becomes $\mu = P((W, \Lambda, X) \in \cdot \mid \mu)$. We illustrate by example just how different weak solutions can be from the strong solutions typically considered in the MFG literature, in which μ is deterministic. Finally we close the section by discussing some situations in which weak solutions are concentrated on the family of strong solutions.

First, let us state a simplified definition of MFG solution for the case $\sigma_0 \equiv 0$, which is really just Definition 3.1.1 rewritten without B . Again, the following definition is relative to the initial state distribution λ .

Definition 4.1.1. A *weak MFG solution without common noise* is a tuple $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$, where $(\tilde{\Omega}, \mathcal{F}_T, \mathbb{F}, P)$ is a complete filtered probability space supporting (W, μ, Λ, X) satisfying

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

2. X_0 , W , and μ are independent.
3. Λ is a \mathbb{F} -progressively measurable process with values in $\mathcal{P}(A)$ and

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

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

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

4. The state equation holds:

$$dX_t = \int_A b(t, X_t, \mu_t^x, a) \Lambda_t(da) dt + \sigma(t, X_t, \mu_t^x) dW_t. \quad (4.1)$$

5. For any other $(\tilde{\Omega}', (\mathcal{F}'_t)_{t \in [0, T]}, P', W', \mu', \Lambda', X')$ satisfying (1-4) and also $P' \circ (\mu')^{-1} = P \circ \mu^{-1}$, we have

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

6. μ is a version of the conditional law of (W, Λ, X) given μ .

As in Definition 3.1.2, we may refer to the law $P \circ (W, \mu, \Lambda, X)^{-1}$ itself as a weak MFG solution. Again, if also there exists an A -valued process $(\alpha_t)_{t \in [0, T]}$ such that $P(\Lambda_t = \delta_{\alpha_t} \text{ a.e. } t) = 1$, then we say the MFG solution has *strict control*. If this $(\alpha_t)_{t \in [0, T]}$ is progressively measurable with respect to the completion of $(\mathcal{F}_t^{X_0, W, \mu})_{t \in [0, T]}$, we say the MFG solution has *strong control*. If μ is a.s.-constant, then we have a *strong MFG solution without common noise*. In this case, we may abuse the terminology somewhat by saying that a measure $\tilde{\mu} \in \mathcal{P}^p(\mathcal{X})$ is itself a *strong MFG solution (without common noise)*, if there exists a weak MFG solution $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$ without common noise such that $P(\mu = \tilde{\mu}) = 1$.

Remark 4.1.2. We state Definitions 4.1.1 and 3.1.1 only under the assumption that the volatilities are uncontrolled for a number of reasons. The main reason at this stage is the SDE (4.1). To correctly include a relaxed control in the volatility requires the use of martingale measures, and this would lose us control over the driving Wiener process W . There is no problem, however, in defining weak MFG solutions with *strict control*, and the existence theorem of this chapter, Theorem 4.6.1, allows for controlled volatility. Under only assumption A1, we may define a weak MFG solution with strict control as a tuple $(\tilde{\Omega}, \mathbb{F}, P, W, \mu, \alpha, X)$ satisfying all the properties of Definition 4.1.1 with $\Lambda = dt\delta_{\alpha_t}(da)$, except with property (4) replaced by the state equation

$$dX_t = b(t, X_t, \mu_t^x, \alpha_t) dt + \sigma(t, X_t, \mu_t^x) dW_t.$$

Note that for a strong solution, i.e. when μ is deterministic, the consistency condition (6) is strengthened to $\mu = P \circ (W, \Lambda, X)^{-1}$. The definition of a strong solution may be simplified

somewhat, noting that the X marginal of μ is all that appears in the state equation and objective functions. Indeed, we may alter the definition to require that μ be a (deterministic) element of $\mathcal{P}^p(\mathcal{C}^d)$, rather than $\mathcal{P}^p(\mathcal{X})$, and the fixed point condition (6) may be replaced by $\mu = P \circ X^{-1}$. This loses no information, and there is an obvious one-to-one mapping between the two types of solutions.

For weak solutions, however, no such simplification is available, and there is no redundant information in the full conditional law $\mu = P((W, \Lambda, X) \in \cdot \mid \mu)$ of Definition 4.1.1. While only the X -marginal appears in conditions (4) and (5) of the definition, the filtration generated by (X_0, W, μ) is generally larger than that of (X_0, W, μ^x) , and this is important to the conditional independence of condition (3).

Remark 4.1.3. Our notion of *strong MFG solution without common noise with strong control* corresponds to the usual definition of MFG solution in the literature. It is exactly the definition used in the recent papers [51, 88], and it is a generalization of the more standard definition of MFG solution without common noise found in [67, 32, 15], for example. The latter papers require optimality *only relative to other strong controls*, not among all weak controls as we do in condition (5) of Definition 4.1.1. Under assumption A1, however, optimality among strong controls implies optimality among weak controls, as will be seen in Proposition 5.3.7 (which is just a consequence of Proposition 2.1.15). Thus our definition does include this more standard one. See also [75] or the more recent [77] for further discussion of this point.

4.2 The limit theorem

We continue to work with the definition of the n -player games of Section 3.2. Suppose we are given an n -player environment $\mathcal{E}_n = (\tilde{\Omega}_n, \mathbb{F}^n = (\mathcal{F}_t^n)_{t \in [0, T]}, \mathbb{P}_n, \xi, B, W)$, as was defined in Section 3.2. Let $\mathbb{F}^{vs, n} = (\mathcal{F}_t^{vs, n})_{t \in [0, T]}$ denote the \mathbb{P}_n -completion of $(\sigma(\xi, W_s : s \leq t))_{t \in [0, T]}$, that is the filtration generated by the initial state and the idiosyncratic noises (but not the common noise). Let us say that a control $\Lambda \in \mathcal{A}_n(\mathcal{E}_n)$ is a *very strong control* if $\mathbb{P}_n(\Lambda_t = \delta_{\alpha_t} \text{ a.e. } t) = 1$, for some $\mathbb{F}^{vs, n}$ -progressively measurable A -valued process α . A *very strong strategy* is a vector of strong controls. For $\epsilon = (\epsilon_1, \dots, \epsilon_n) \in [0, \infty)^n$, a *very strong ϵ -Nash equilibrium in \mathcal{E}_n* is any *very strong strategy* $\Lambda \in \mathcal{A}_n^n(\mathcal{E}_n)$ such that

$$J_i(\Lambda) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n) \text{ very strong}} J_i((\Lambda^{-i}, \beta)) - \epsilon_i, \quad i = 1, \dots, n.$$

The *very strong equilibrium* is arguably the most natural notion of equilibrium in the case of no common noise, and it is certainly one of the most common in the literature. The proof of the following Proposition is deferred to Section 6.4.2.

Proposition 4.2.1. *When $\sigma_0 \equiv 0$, every very strong ϵ -Nash equilibrium is also a relaxed ϵ -Nash equilibrium.*

The following Theorem 4.2.2 rewrites Theorems 3.2.4 and 3.2.10 in the setting without common noise. Although this is mostly derived from Theorems 3.2.4 and 3.2.10, the proof is spelled out in Section 6.3, as it is not entirely straightforward.

Theorem 4.2.2. *Suppose $\sigma_0 \equiv 0$. Theorem 3.2.4 remains true if the term “weak MFG solution” is replaced by “weak MFG solution without common noise,” and if P_n is defined instead by*

$$P_n := \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (\xi^i, W^i, \hat{\mu}[\Lambda^n], \Lambda^{n,i}, X^i[\Lambda^n])^{-1}. \quad (4.2)$$

Theorem 3.2.10 remains true if “weak MFG solution” is replaced by “weak MFG solution without common noise,” if P_n is defined by (4.2), and if “strong” is replaced by “very strong.”

Since strong MFG solutions are more familiar in the literature on mean field games and presumably more accessible computationally, it would be nice to have a description of weak solutions in terms of strong solutions. Recall from Sections 2.2 and 2.3.3 that, without common noise, both McKean-Vlasov equations and static mean field game exhibit the property that weak solutions are simply randomizations among the set of strong solutions. We will see that this is not true in general for dynamic mean field games and that the interplay between the *dynamics* and the *optimization* leads to a fundamental difference between stochastic and deterministic equilibria (i.e., weak and strong MFG solutions). More specifically, the adaptiveness requirement renders the class of admissible controls quite sensitive to how random μ is. To highlight this point, Section 4.3 below describes a model possessing weak MFG solutions which are not randomizations of strong MFG solutions. Subsection 4.4 discusses some partial results on when this simplification can occur in the MFG setting.

4.3 An illuminating example

This section describes a deceptively simple example which illustrates the difference between weak and strong solutions. Consider the time horizon $T = 2$, the initial state distribution $\lambda = \delta_0$, and the following data (still with $\sigma_0 \equiv 0$):

$$\begin{aligned} b(t, x, \nu, a) &= a, \quad \sigma \text{ constant}, \quad A = [-1, 1] \\ g(x, \nu) &= x\bar{\nu}, \quad f \equiv 0, \end{aligned}$$

where for $\nu \in \mathcal{P}^1(\mathbb{R})$ we define $\bar{\nu} := \int x\nu(dx)$. Similarly, for $\mu \in \mathcal{P}^1(\mathcal{X})$ write $\bar{\mu}_t^x := \int_{\mathbb{R}} x\mu_t^x(dx)$. Assumption A1 is verified by choosing $p = 2$, $p_\sigma = 0$, and any $p' > 2$. Let us first study the optimization problems arising in the MFG problem. Let $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, 2]}, P, W, \mu, \Lambda, X)$ satisfy (1-5) of Definition 4.1.1. For \mathbb{F} -progressively measurable $\mathcal{P}([-1, 1])$ -valued processes β , define

$$\tilde{J}(\beta) := \mathbb{E} \left[X_2^\beta \bar{\mu}_2^x \right],$$

where

$$X_t^\beta = \int_0^t \int_{[-1, 1]} a\beta_t(da)dt + \sigma W_t, \quad t \in [0, 2].$$

Independence of W and μ implies

$$\tilde{J}(\beta) = \mathbb{E} \left[\int_0^2 \int_{[-1,1]} a \bar{\mu}_2^x \beta_t(da) dt \right] = \mathbb{E} \left[\int_0^2 \int_{[-1,1]} a \mathbb{E}^P[\bar{\mu}_2^x \mid \mathcal{F}_t^\beta] \beta_t(da) dt \right],$$

where $\mathcal{F}_t^\beta := \sigma(\beta_s : s \leq t)$. If it is also required that \mathcal{F}_t^β is conditionally independent of $\mathcal{F}_2^{X_0, W, \mu}$ given $\mathcal{F}_t^{X_0, W, \mu}$, then

$$\mathbb{E}^P[\bar{\mu}_2^x \mid \mathcal{F}_t^\beta] = \mathbb{E}^P[\bar{\mu}_2^x \mid \mathcal{F}_t^{X_0, W, \mu}] = \mathbb{E}^P[\bar{\mu}_2^x \mid \mathcal{F}_t^\mu],$$

where the last equality follows from independence of (X_0, W) and μ , and $\mathcal{F}_t^\mu := \sigma(\mu(C) : C \in \mathcal{F}_t^\mathcal{X})$. Hence

$$\tilde{J}(\beta) = \mathbb{E} \left[\int_0^2 \int_{[-1,1]} a \mathbb{E}^P[\bar{\mu}_2^x \mid \mathcal{F}_t^\mu] \beta_t(da) dt \right]. \quad (4.3)$$

Condition (5) of Definition 4.1.1 implies that Λ maximizes J over all such processes β , which implies that $\Lambda_t(\omega)$ must equal $\delta_{\alpha_t^*(\omega)}$ on the (t, ω) -set $\{\alpha^* \neq 0\}$, where

$$\alpha_t^* := \text{sign}(\mathbb{E}[\bar{\mu}_2^x \mid \mathcal{F}_t^\mu]),$$

and we use the convention $\text{sign}(0) := 0$.

Remark 4.3.1. This already highlights the key point: When μ is deterministic, an optimal control is the constant $\text{sign}(\bar{\mu}_2^x)$, but when μ is random, this control is inadmissible since it is not adapted.

Proposition 4.3.2. *Every strong MFG solution (without common noise) satisfies $\bar{\mu}_2^x \in \{-2, 0, 2\}$ and $\bar{\mu}_t^x = t \text{sign}(\bar{\mu}_2^x)$.*

Proof. Let $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0,2]}, P, W, \mu, \Lambda, X)$ satisfy Definition 4.1.1, with μ deterministic. In this case, $\alpha_t^* = \text{sign}(\bar{\mu}_2^x)$ for all t . Suppose that $\bar{\mu}_2^x \neq 0$. Then $\Lambda_t = \delta_{\alpha_t^*}$ must hold $dt \otimes dP$ -a.e., and thus

$$X_t = t \text{sign}(\bar{\mu}_2^x) + \sigma W_t, \quad t \in [0, 2].$$

The consistency condition (6) of Definition 4.1.1 implies $\bar{\mu}_t^x = \mathbb{E}[X_t] = t \text{sign}(\bar{\mu}_2^x)$. In particular, $\bar{\mu}_2^x = 2 \text{sign}(\bar{\mu}_2^x)$, which implies $\bar{\mu}_2^x = \pm 2$ since we assumed $\bar{\mu}_2^x \neq 0$. \square

Proposition 4.3.3. *There exists a weak MFG solution (without common noise) satisfying $P(\bar{\mu}_2^x = 1) = P(\bar{\mu}_2^x = -1) = 1/2$.*

Proof. Construct on some probability space $(\tilde{\Omega}, \mathcal{F}, P)$ a random variable γ with $P(\gamma = 1) = P(\gamma = -1) = 1/2$ and an independent Wiener process W . Let $\alpha_t^* = \gamma 1_{(1,2]}(t)$ for each t (noticing that this interval is open on the left), and define $(\mathcal{F}_t)_{t \in [0,2]}$ to be the complete filtration generated by $(W_t, \alpha_t^*)_{t \in [0,2]}$. Let

$$X_t := \int_0^t \alpha_s^* ds + \sigma W_t = (t-1)\gamma 1_{(1,2]}(t) + \sigma W_t, \quad t \in [0, 2].$$

Finally, let $\Lambda = dt\delta_{\alpha_t^*}(da)$, and define $\mu := P((W, \Lambda, X) \in \cdot \mid \gamma)$. Clearly μ is γ -measurable. On the other hand, independence of γ and W implies

$$\bar{\mu}_2^x = \mathbb{E}[X_2 \mid \gamma] = \gamma.$$

Thus γ is also μ -measurable, and we conclude that $\mu := P((W, \Lambda, X) \in \cdot \mid \mu)$. It is straightforward to check that

$$\mathcal{F}_t^\mu = \begin{cases} \{\emptyset, \tilde{\Omega}\} & \text{if } t \leq 1 \\ \sigma(\gamma) & \text{if } 1 < t \leq 2 \end{cases}.$$

Thus

$$\mathbb{E}[\bar{\mu}_2^x \mid \mathcal{F}_t^\mu] = \begin{cases} \mathbb{E}[\gamma] = 0 & \text{if } t \leq 1 \\ \mathbb{E}[\gamma \mid \gamma] = \gamma & \text{if } 1 < t \leq 2 \end{cases}.$$

Since $\bar{\mu}_2^x = \gamma = \text{sign}(\gamma)$, we conclude that $\alpha_t^* = \text{sign}(\mathbb{E}[\bar{\mu}_2^x \mid \mathcal{F}_t^\mu])$. It is then readily checked using the previous arguments that $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0,2]}, P, W, \mu, \Lambda, X)$ is a weak MFG solution. \square

To be absolutely clear, the above two propositions imply the following: If $S := \{\nu \in \mathcal{P}(\mathcal{X}) : \bar{\nu}_2^x \in \{-2, 0, 2\}\}$, then every strong MFG solution lies in S , but there exists a weak MFG solution with $P(\mu \in S) = 0$.

Remark 4.3.4. The example of Proposition 4.3.3 can be modified to illustrate another strange phenomenon. The proof of Proposition 4.3.3 has $\alpha_t^* = \gamma$ for $t \in (1, 2]$ and $\alpha_t^* = 0$ for $t \leq 1$. Instead, we could set $\alpha_t^* = \eta_t$ for $t \leq 1$, for any mean-zero $[-1, 1]$ -valued process $(\eta_t)_{t \in [0,1]}$ independent of γ and W . The rest of the proof proceeds unchanged, yielding another weak MFG solution with the same conditional *mean* state $\bar{\mu}^x$, but with different conditional *law* μ^x . (In fact, we could even choose α^* to be any mean-zero *relaxed* control on the time interval $[0, 1]$.) Intuitively, for $t \leq 1$ we have $\mathbb{E}[\bar{\mu}_2^x \mid \mathcal{F}_t^\mu] = 0$, and the choice of control on the time interval $[0, 1]$ does not matter in light of (4.3); the agent then has some freedom to *randomize* her choice of control among the family of non-unique optimal choices. This type of randomization can typically occur when optimal controls are non-unique, and although it is unnatural in some sense, Theorem 3.2.4 indicate that this behavior can indeed arise in the limit from the finite-player games.

4.4 Supports of weak solutions

In this section, we attempt to partially explain what permits the existence of weak solutions which are not randomizations among strong solutions. As was mentioned in Remark 4.3.1, the culprit is the adaptedness required of controls. Indeed, in the example of Section 4.3, very different optimal controls arise depending on whether or not the measure μ is random. If μ is deterministic, then so is the optimal control, and we may write this optimal control as a functional of μ by

$$\hat{\alpha}^D(t, \mu) = \text{sign}(\bar{\mu}_T^x), \quad t \in [0, T].$$

The problem is as follows: for each fixed deterministic μ , the optimal control $(\hat{\alpha}^D(t, \mu))_{t \in [0, T]}$ is deterministic and thus trivially adapted, but when μ is allowed to be random then this

control is no longer adapted and thus no longer admissible. If, for a different MFG problem, it happens that $\hat{\alpha}^D$ is in fact progressively measurable with respect to $(\mathcal{F}_t^\mu)_{t \in [0, T]}$, then this control is still admissible when μ is randomized; moreover, it should be *optimal* when μ is randomized, since it was optimal for each realization of μ .

The following results make these idea precise, but first some terminology will be useful. The discussion of this subsection unfortunately require very cumbersome definitions, but rest assured these definitions are local to this and the following section (Sections 4.4 and 4.5), which can thus be safely skipped with no loss of continuity. As usual we work under assumption **A1** at all times, and the initial state distribution $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$ is fixed.

Definition 4.4.1. We say that a function $\hat{\alpha} : [0, T] \times \mathcal{C}^m \times \mathcal{C}^d \times \mathcal{P}^p(\mathcal{X}) \rightarrow A$ is a *universally admissible control* if:

1. $\hat{\alpha}$ is progressively measurable with respect to the (universal completion of the) natural filtration $(\mathcal{F}_t^{W, X, \mu})_{t \in [0, T]}$ on $\mathcal{C}^m \times \mathcal{C}^d \times \mathcal{P}^p(\mathcal{X})$. Here $\mathcal{F}_t^{W, X, \mu} := \sigma(W_s, X_s, \mu(C) : s \leq t, C \in \mathcal{F}_t^X)$ for each t , where (W, X, μ) denotes the identity map on $\mathcal{C}^m \times \mathcal{C}^d \times \mathcal{P}^p(\mathcal{X})$.
2. For each fixed $\nu \in \mathcal{P}^p(\mathcal{X})$, the SDE

$$dX_t = b(t, X_t, \nu_t^x, \hat{\alpha}(t, W, X, \nu))dt + \sigma(t, X_t, \nu_t^x)dW_t, \quad X_0 \sim \lambda, \quad (4.4)$$

is unique in joint law; that is, if we are given two pairs of processes $(W_t^i, X_t^i)_{t \in [0, T]}$ for $i = 1, 2$, possibly on different filtered probability spaces but with $(W_t^i)_{t \in [0, T]}$ a Wiener process in either case, then (W^1, X^1) and (W^2, X^2) have the same law.

3. Suppose we are given a filtered probability space $(\tilde{\Omega}, (\tilde{\mathcal{F}}_t)_{t \in [0, T]}, \tilde{P})$ supporting an $(\tilde{\mathcal{F}}_t)_{t \in [0, T]}$ -Wiener process \tilde{W} , an $\tilde{\mathcal{F}}_0$ -measurable \mathbb{R}^d -valued random variable $\tilde{\xi}$ with law λ , and a $\mathcal{P}^p(\mathcal{X})$ -valued random variable $\tilde{\mu}$ independent of (ξ, W) such that $\tilde{\mu}(C)$ is $\tilde{\mathcal{F}}_t$ -measurable for each $C \in \mathcal{F}_t^X$ and $t \in [0, T]$. Then there exists a strong solution \tilde{X} of the SDE

$$d\tilde{X}_t = b(t, \tilde{X}_t, \tilde{\mu}_t^x, \hat{\alpha}(t, W, \tilde{X}, \tilde{\mu}))dt + \sigma(t, \tilde{X}_t, \tilde{\mu}_t^x)d\tilde{W}_t, \quad \tilde{X}_0 = \tilde{\xi},$$

and it satisfies $\mathbb{E} \int_0^T |\hat{\alpha}(t, W, \tilde{X}, \tilde{\mu})|^p dt < \infty$.

If $\hat{\alpha}$ is a universally admissible control, we say it is *locally optimal* if for each fixed $\nu \in \mathcal{P}^p(\mathcal{X})$ there exists a complete filtered probability space $(\Omega^{(\nu)}, (\mathcal{F}_t^{(\nu)})_{t \in [0, T]}, P^\nu)$ supporting a Wiener process W^ν and a continuous adapted process X^ν such that (W^ν, X^ν) satisfies the SDE (4.4) and:

- (4) If $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P)$ supports a m -dimensional Wiener process W , a progressive $\mathcal{P}(A)$ -valued process Λ satisfying $\mathbb{E}^P \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty$, and a continuous adapted \mathbb{R}^d -valued process X satisfying

$$dX_t = \int_A b(t, X_t, \nu_t^x, a) \Lambda_t(da) dt + \sigma(t, X_t, \nu_t^x) dW_t, \quad P \circ X_0^{-1} = \lambda,$$

then

$$\mathbb{E}^{P^{(\nu)}} [\Gamma(\nu^x, dt\delta_{\hat{\alpha}(t, W^\nu, X^\nu, \nu)}(da), X^\nu)] \geq \mathbb{E}^P [\Gamma(\nu^x, \Lambda, X)].$$

We need an additional assumption **C**, which simply requires the uniqueness of the optimal controls. A typical example when this holds is for linear-convex coefficients; in the proof of Proposition 3.1.5 in Section 5.4, it is shown that assumption **(Linear-Convex)** implies **C**.

Assumption C. If $(\tilde{\Omega}^i, (\mathcal{F}_t^i)_{t \in [0, T]}, P^i, W^i, \mu^i, \Lambda^i, X^i)$ for $i = 1, 2$ both satisfy (1-5) of Definition 4.1.1 as well as $P^1 \circ (\mu^1)^{-1} = P^2 \circ (\mu^2)^{-1}$, then $P^1 \circ (W^1, \mu^1, \Lambda^1, X^1)^{-1} = P^2 \circ (W^2, \mu^2, \Lambda^2, X^2)^{-1}$.

Note the similarities between the following Theorem 4.4.2 and Proposition 2.3.8, which derived a similar result for static MFGs. In Proposition 2.3.8, it was shown without further assumption that weak solutions concentrate on strong solutions; here, because of the adaptedness issues discussed above, we need to assumption that there exists a universally admissible locally optimal control.

Theorem 4.4.2. Assume **C** holds. Suppose that there exists a universally admissible and locally optimal control $\hat{\alpha} : [0, T] \times \mathcal{C}^m \times \mathcal{C}^d \times \mathcal{P}^p(\mathcal{X}) \rightarrow A$. Then, for every weak MFG solution $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$ (without common noise), $P \circ \mu^{-1}$ is concentrated on the set of strong MFG solutions (without common noise). Conversely, if $\rho \in \mathcal{P}^p(\mathcal{P}^p(\mathcal{X}))$ is concentrated on the set of strong MFG solutions (without common noise), then there exists a weak MFG solution (without common noise) with $P \circ \mu^{-1} = \rho$.

Proof. Let $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$ be a weak MFG solution (without common noise). *Step 1:* We will first show that necessarily $\Lambda_t = \delta_{\hat{\alpha}(t, W, X, \mu)}$ holds $dt \otimes dP$ -a.e. On $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P)$ we may use (3) of Definition 4.4.1 to find a strong solution X' of the SDE

$$dX'_t = b(t, X'_t, \mu_t^x, \hat{\alpha}(t, W, X', \mu))dt + \sigma(t, X'_t, \mu_t^x)dW_t, \quad X'_0 = X_0,$$

with $\mathbb{E}^P \int_0^T |\hat{\alpha}(t, W, X', \mu)|^p dt < \infty$. In particular, X' is adapted to the (completion of the) filtration $\mathcal{F}_t^{X_0, W, \mu} := \sigma(X_0, W_s, \mu(C) : s \leq t, C \in \mathcal{F}_t^{\mathcal{X}})$. Let $\Lambda' := dt\delta_{\hat{\alpha}(t, W, X', \mu)}(da)$. Then it is clear that $(\tilde{\Omega}, (\mathcal{F}_t^{X_0, W, \mu})_{t \in [0, T]}, P, W, \mu, \Lambda', X')$ satisfies conditions (1-4) of Definition 4.1.1. Optimality of P implies

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

On the other hand, for $P \circ \mu^{-1}$ -a.e. $\nu \in \mathcal{P}^p(\mathcal{X})$, the following hold under $P(\cdot \mid \mu = \nu)$:

- W is a $(\mathcal{F}_t)_{t \in [0, T]}$ -Wiener process.
- (W, Λ, X) satisfies

$$dX_t = \int_A b(t, X_t, \nu_t^x, a) \Lambda_t(da) + \sigma(t, X_t, \nu_t^x)dW_t.$$

- (W, X') solves the SDE (4.4).

From the local optimality of $\hat{\alpha}$ we conclude (keeping in mind the uniqueness condition (2) of Definition 4.4.1) that

$$\mathbb{E}^P [\Gamma(\mu^x, \Lambda, X) | \mu] \leq \mathbb{E}^P [\Gamma(\mu^x, \Lambda', X') | \mu].$$

Thus

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

By assumption **C**, there is only one optimal control, and so $\Lambda = \Lambda' = dt\delta_{\hat{\alpha}(t, W, X', \mu)}(da)$, P -a.s. From uniqueness of the SDE solutions we conclude that $X = X'$ a.s. as well, completing the first step. (Note we do not use the assumptions of Definition 4.4.1 for this last conclusion, but only the Lipschitz assumption **A1.4**.)

Step 2: Next, we show that $P \circ \mu^{-1}$ is concentrated on the set of strong MFG solutions. Using (2) and (3) of Definition 4.4.1, we know that for $P \circ \mu^{-1}$ -a.e. $\nu \in \mathcal{P}^p(\mathcal{X})$ there exists on some filtered probability space $(\Omega^{(\nu)}, (\mathcal{F}_t^{(\nu)})_{t \in [0, T]}, P^\nu)$ a weak solution X^ν of the SDE

$$dX_t^\nu = b(t, X_t^\nu, \nu_t^x, \hat{\alpha}(t, W^\nu, X^\nu, \nu))dt + \sigma(t, X_t^\nu, \nu_t^x)dW_t^\nu, \quad P^\nu \circ (X_0^\nu)^{-1} = \lambda,$$

where W^ν is an $(\mathcal{F}_t^{(\nu)})_{t \in [0, T]}$ -Wiener process. From Step 1, on $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P)$ we have

$$dX_t = b(t, X_t, \mu_t^x, \hat{\alpha}(t, W, X, \mu))dt + \sigma(t, X_t, \mu_t^x)dW_t, \quad P \circ X_0^{-1} = \lambda.$$

It follows from the P -independence of μ , X_0 , and W along with the uniqueness in law of condition (2) of Definition 4.4.1 that

$$P((W, \Lambda, X) \in \cdot \mid \mu = \nu) = P^\nu \circ (W^\nu, dt\delta_{\hat{\alpha}(t, W^\nu, X^\nu, \nu)}(da), X^\nu)^{-1}, \quad (4.5)$$

for $P \circ \mu^{-1}$ -a.e. $\nu \in \mathcal{P}^p(\mathcal{X})$. Since $\mu = P((W, \Lambda, X) \in \cdot \mid \mu)$, it follows that

$$\nu = P^\nu \circ (W^\nu, dt\delta_{\hat{\alpha}(t, W^\nu, X^\nu, \nu)}(da), X^\nu)^{-1}, \quad \text{for } P \circ \mu^{-1}\text{-a.e. } \nu \in \mathcal{P}^p(\mathcal{X}). \quad (4.6)$$

We conclude that $P \circ \mu^{-1}$ -a.e. $\nu \in \mathcal{P}^p(\mathcal{X})$ is a strong MFG solution, or more precisely that

$$(\Omega^{(\nu)}, (\mathcal{F}_t^{(\nu)})_{t \in [0, T]}, P^\nu, W^\nu, \nu, dt\delta_{\hat{\alpha}(t, W^\nu, X^\nu, \nu)}(da), X^\nu)$$

is a strong MFG solution. Indeed, we just verified condition (6) of Definition 4.1.1, and conditions (1-4) are obvious. The optimality condition (5) of Definition 4.1.1 is a simple consequence of the local optimality of $\hat{\alpha}$

Step 3: We turn now to the converse. Let $(\tilde{\Omega}, \mathcal{F}, P)$ be any probability space supporting a random variable (ξ, W, μ) with values in $\mathbb{R}^d \times \mathcal{C}^m \times \mathcal{P}^p(\mathcal{X})$ with law $\lambda \times \mathcal{W}^m \times \rho$, where \mathcal{W}^m is Wiener measure on \mathcal{C}^m . Let $(\mathcal{F}_t)_{t \in [0, T]}$ denote the P -completion of $(\sigma(\xi, W_s, \mu(C) : s \leq t, C \in \mathcal{F}_t^{\mathcal{X}}))_{t \in [0, T]}$. Solve strongly on $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P)$ the SDE

$$dX_t = b(t, X_t, \mu_t^x, \hat{\alpha}(t, W, X, \mu))dt + \sigma(t, X_t, \mu_t^x)dW_t, \quad X_0 = \xi.$$

Note that hypothesis (3) makes this possible. Define $\Lambda := dt\delta_{\hat{\alpha}(t, W, X, \mu)}(da)$. Clearly $P \circ \mu^{-1} = \rho$ by construction, and we claim that $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$ is a weak MFG solution.

Using hypothesis (1), it is clear that conditions (1-4) of Definition 4.1.1 hold, and thus we must only check the optimality condition (5) and the fixed point condition (6).

First, let $(\tilde{\Omega}', (\mathcal{F}_t')_{t \in [0, T]}, P', W', \mu', \Lambda', X')$ be an alternative probability space satisfying (1-4) of Definition 4.1.1 and $P' \circ (\mu')^{-1} = P \circ \mu^{-1} = \rho$. The uniqueness in law condition (2) of Definition 4.4.1 implies that $P((W, X) \in \cdot \mid \mu = \nu)$ is exactly the law of the solution of the SDE (4.4), for $P \circ \mu^{-1}$ -a.e. ν . Applying local optimality of $\hat{\alpha}$ for each ν , we conclude that

$$\mathbb{E}^P [\Gamma(\nu^x, \Lambda, X) \mid \mu = \nu] \geq \mathbb{E}^{P'} [\Gamma(\nu^x, \Lambda', X') \mid \mu = \nu], \text{ for } \rho - a.e. \nu.$$

Integrate with respect to ρ on both sides to get $\mathbb{E}^P [\Gamma(\mu^x, \Lambda, X)] \geq \mathbb{E}^{P'} [\Gamma((\mu')^x, \Lambda', X')]$, which verifies condition (5) of Definition 4.1.1. Finally, we check (6) by applying Step 1 to deterministic μ and again using uniqueness of the SDE (4.4) to find that both (4.5) and (4.6) hold for ρ -a.e. ν . \square

4.5 Applications of Theorem 4.4.2

It is admittedly quite difficult to check that there exists a universally admissible, locally optimal control, and we will leave this problem open in all but the simplest cases. Note, however, that conditions (2) and (3) of Definition 4.4.1 hold automatically when $\hat{\alpha}(t, w, x, \nu) = \hat{\alpha}'(t, w, x_0, \nu)$, for some $\hat{\alpha}' : [0, T] \times \mathcal{C}^m \times \mathbb{R}^d \times \mathcal{P}^p(\mathcal{X}) \rightarrow A$.

A simple class of examples

Suppose $A \subset \mathbb{R}^k$ is convex, $g \equiv 0$, and $f = f(t, \mu, a)$ is twice differentiable in a with uniformly negative Hessian. That is, $D_a^2 f(t, \mu, a) \leq -\delta$ for all (t, μ) , for some $\delta > 0$. Suppose as usual that assumption A1 holds. Define

$$\hat{\alpha}(t, w, x, \nu) := \arg \max_{a \in A} f(t, \nu_t^x, a), \text{ for } (t, w, x, \nu) \in [0, T] \times \mathcal{C}^m \times \mathcal{C}^d \times \mathcal{P}^p(\mathcal{X}).$$

It is straightforward to check that assumption C holds and that $\hat{\alpha}$ is a universally admissible and locally optimal control. Of course, this example is quite simple in that the state process does not influence the optimization.

A possible general strategy

The following approach may be more widely applicable. First, for a fixed $\nu \in \mathcal{P}^p(\mathcal{X})$, we may define the value function $V[\nu](t, x)$ of the corresponding optimal control problem in the usual way, and it should solve a Hamilton-Jacobi-Bellman (HJB) PDE of the form

$$\begin{cases} -\partial_t V[\nu](t, x) - H(t, x, \nu_t^x, D_x V[\nu](t, x), D_x^2 V[\nu](t, x)) &= 0, \text{ on } [0, T] \times \mathbb{R}^d, \\ V[\nu](T, x) &= g(x, \nu_T^x) \end{cases},$$

where the Hamiltonian $H : [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$ is defined by

$$H(t, x, \mu, y, z) := \sup_{a \in A} [y^\top b(t, x, \mu, a) + f(t, x, \mu, a)] + \frac{1}{2} \text{Tr} [z \sigma \sigma^\top(t, x, \mu)].$$

Suppose that we can show (as is well known to be possible in very general situations) that for each ν the value function $V[\nu]$ is the unique (viscosity) solution of this HJB equation. Then, an optimal control can be obtained by finding $\hat{\alpha}(t, x_t, \nu)$ which achieves the supremum in

$$H(t, x_t, \nu_t^x, D_x V[\nu](t, x_t), D_x^2 V[\nu](t, x_t)),$$

for each (t, x, ν) . The crux of this approach is to show that the value function $V[\nu](t, x)$ is *adapted with respect to ν* in some sense, which would imply that $\hat{\alpha}$ is universally admissible and locally optimal. A nice special case would be a Markovian dependence, $V[\nu](t, x) = \tilde{V}(t, x, \nu_t^x)$. In short, we must study the dependence of a family of HJB equations on a path-valued parameter.

4.6 Existence and uniqueness

The study of existence of solutions is much simpler in the absence of common noise. The existence of *strong* solutions can be established directly in this case, solely under assumption **A1**; notably, the volatility term σ may be controlled. The following existence theorem is stated under assumption **(Convex)** (defined on page 52), but again without this assumption a form of the theorem still holds with relaxed controls. Because of the presence of control in the volatility term, the precise statement of the general result requires additional technical developments which we postpone to Chapter 8. Although we only defined weak MFG solutions under the assumption **A2** that the volatilities are uncontrolled, note that we may define weak MFG solutions *with strict control* in a natural way without assumption **A2**; see Remark 4.1.2.

Theorem 4.6.1. *Under assumption **A1** and **(Convex)**, the MFG without common noise ($\sigma_0 \equiv 0$) admits a strong solution with strict Markovian control.*

The question of uniqueness requires some new definitions compared to the setting with common noise. The results here are simpler than in Section 3.3 with common noise, but we still need the notions of uniqueness in law and pathwise uniqueness.

Definition 4.6.2. An MFG without common noise is *unique in law* if any two weak solutions induce the same law of (μ, W, Λ, X) .

A weak MFG solution without common noise is determined by the law of μ , just as for MFG with common noise the law of a weak MFG solution is determined by the law of (B, μ) . Indeed, for an element $\gamma \in P^p(\mathcal{P}^p(\mathcal{X}))$, we can define $M\gamma \in \mathcal{P}(\mathcal{C}^m \mathcal{P}^p(\mathcal{X}) \times \mathcal{V} \times \mathcal{C}^d)$ by

$$M\gamma(dw, dm, dq, dx) = \gamma(dm)m(dw, dq, dx).$$

We will say γ is a *MFG solution basis without common noise* if the distribution $M\gamma$ together with the canonical processes and filtrations on $\mathcal{C}^m \mathcal{P}^p(\mathcal{X}) \times \mathcal{V} \times \mathcal{C}^d$ form a weak MFG solution. Given two MFG solution bases (without common noise) γ^1 and γ^2 , we say $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, \mu^1, \mu^2)$ is a *coupling* of γ^1 and γ^2 if:

1. $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q)$ is a probability space with a complete filtration.

2. For each $t \in [0, T]$, we have (up to null sets)

$$\mathcal{G}_t = \sigma(\mu^1(C), \mu^2(C) : s \leq t, C \in \mathcal{F}_t^{\mathcal{X}}).$$

3. For $i = 1, 2$, $Q \circ (\mu^i)^{-1} = \gamma^i$.

4. μ^1 and μ^2 are independent.

Definition 4.6.3. We say *pathwise uniqueness* holds for the MFG without common noise if, for any coupling $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ of any two MFG solution bases (without common noise), we have $\mu^1 = \mu^2$ a.s.

Remark 4.6.4. In a sense, the analog of the Yamada-Watanabe theorem is trivial in this setting, and thus we make only a remark rather than a proposition. If there exists a weak solution, and if pathwise uniqueness holds, then the weak solution is strong and is unique in law. The proof follows simply from the fact that if μ^1 and μ^2 are independent, but $\mu^1 = \mu^2$ almost surely, then μ^1 and μ^2 must be almost surely constant. Such a concept is still useful, however, because of the fact that weak solutions do not necessarily concentrate on the set of strong solutions. If they did, then uniqueness of weak solutions would follow immediately from uniqueness of strong solutions.

The proof of the following uniqueness Theorem 4.6.5 is essentially the same as that of Theorem 3.3.5, which is proven in Section 7.3.

Theorem 4.6.5. Suppose assumptions **A1**, **A2**, **U**, and **(Linear-Convex)** hold. Then the MFG without common noise ($\sigma_0 \equiv 0$) is pathwise unique. In particular, there exists a unique in law weak MFG solution with weak control, and it is in fact a strong MFG solution with strong control.

Chapter 5

Properties of mean field game solutions

We now turn toward the proofs of the main theorems on convergence, existence, and uniqueness. As a first step, this chapter is devoted to the derivation of several useful structural properties of MFG solutions. The first section defines and discusses *MFG pre-solutions*, leading to the critical Lemma 5.1.2, which provides a shortcut around the troublesome compatibility condition. Then, in Section 5.2, we transfer the definitions to a canonical space and set up convenient notation that will be used throughout the proofs. Section 5.3 contains a number of useful topological results, including the useful Proposition 5.3.7 which allows us to check optimality against a dense subclass of admissible controls. Using some of these results, we return to the proofs of Propositions 3.1.4 and 3.1.5. Finally, the chapter closes with a reformulation of the definition of weak MFG solution, solely in terms of the induced law on the canonical space Ω , providing a more intrinsic form of Definition 3.1.2.

5.1 Pre-solutions

The following definition of MFG pre-solution will be useful. It is exactly the same as the definition of weak MFG solution, except that the property of optimality is omitted. Recall the definitions of the canonical spaces \mathcal{X} and Ω from (3.1) and (3.5).

Definition 5.1.1. An *MFG pre-solution* is a tuple $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$ satisfying properties (1-4) and (6) of Definition 4.1.1. Alternatively, we may refer to the law of (B, W, μ, Λ, X) on Ω as an MFG pre-solution.

The definition of MFG solution and pre-solution both require the compatibility of Λ , that is the conditional independence of $\mathcal{F}_t^\Lambda = \sigma(\Lambda_s : s \leq t)$ and $\mathcal{F}_T^{X_0, B, W, \mu}$ given $\mathcal{F}_t^{X_0, B, W, \mu}$, for each t . This property does not behave well under limits, and it will be crucial to have an alternative characterization of MFG pre-solutions which allows us to avoid directly compatibility directly. Namely, Lemma 5.1.2 below shows that compatibility essentially follows automatically from the fixed point condition (3) of Definition 3.1.1. In fact, Lemma 5.1.2 is the main reason we work with the conditional law of (W, Λ, X) , and not just X .

Lemma 5.1.2. *Let $(\tilde{\Omega}, \mathcal{F} = (\mathcal{F}_t)_{t \in [0, T]}, P)$ is a complete filtered probability space supporting (B, W, μ, Λ, X) satisfying:*

1. *B and W are independent \mathbb{F} -Wiener processes of respective dimension m_0 and m , respectively. The process X is \mathcal{F} -adapted with values in \mathbb{R}^d , and $P \circ X_0^{-1} = \lambda$. Moreover, μ is a random element of $\mathcal{P}^p(\mathcal{X})$ such that $\mu(C)$ is \mathcal{F}_t -measurable for each $C \in \mathcal{F}_t^{\mathcal{X}}$ and $t \in [0, T]$. Finally, Λ is \mathcal{F} -progressively measurable with values in $\mathcal{P}^p(A)$.*
2. *X_0, W , and (B, μ) are independent.*
3. *μ is a version of the conditional law of (W, Λ, X) given (B, μ) .*

Then the following statements hold:

1. *Letting (W, Λ, X) denote the canonical processes on $\mathcal{X} = \mathcal{C}^m \times \mathcal{V} \times \mathcal{C}^d$, it holds for $P \circ \mu^{-1}$ -almost every $\nu \in \mathcal{P}^p(\mathcal{X})$ that W is an $\mathcal{F}_t^{\mathcal{X}}$ -Wiener process under ν .*
2. *On $\tilde{\Omega}$, it holds that Λ is compatible, in the sense that \mathcal{F}_t^{Λ} is conditionally independent of $\mathcal{F}_T^{X_0, B, W, \mu}$ given $\mathcal{F}_t^{X_0, B, W, \mu}$, for each $t \in [0, T]$.*

Proof. First step. We begin by proving the first claim. For $\nu \in \mathcal{P}(\mathcal{X})$, let $\nu^w = \nu \circ W^{-1}$ denote the \mathcal{C}^m -marginal. Let also \mathcal{W}^m denote Wiener measure on \mathcal{C}^m . To prove the first claim, let $\varphi_1 : \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ and $\varphi_2 : \mathcal{C}^m \rightarrow \mathbb{R}$ be bounded and measurable. Then, since $P \circ W^{-1} = \mathcal{W}^m$ (with \mathbb{E} denoting expectation under P),

$$\mathbb{E} [\varphi_1(\mu)] \int_{\mathcal{C}^m} \varphi_2 d\mathcal{W}^m = \mathbb{E} [\varphi_1(\mu) \varphi_2(W)] = \mathbb{E} \left[\varphi_1(\mu) \int_{\mathcal{C}^m} \varphi_2 d\mu^w \right].$$

The first equality follows from the independence hypothesis (2), and the second follows from hypothesis (3). This holds for all φ_1 , and thus $\int \varphi_2 d\mu^w = \int \varphi_2 d\mathcal{W}^m$ a.s. This holds for all φ_2 , and thus $\mu^w = \mathcal{W}^m$ a.s.

It remains to check that $\sigma(W_s - W_t : s \in [0, T])$ and $\mathcal{F}_t^{\mathcal{X}}$ are independent under almost every realization of μ . Fix $t \in [0, T]$. Suppose $\varphi_1 : \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ is bounded and \mathcal{F}_t^{μ} -measurable, $\varphi_2 : \mathcal{C}^m \rightarrow \mathbb{R}$ is bounded and $\sigma(W_s - W_t : s \in [t, T])$ -measurable, and $\varphi_3 : \mathcal{X} \rightarrow \mathbb{R}$ is bounded and $\mathcal{F}_t^{\mathcal{X}}$ -measurable. Then, on $\tilde{\Omega}$, $\varphi_2(W)$ and $(\varphi_1(\mu), \varphi_3(W, \Lambda, X))$ are independent (since W is a Wiener process with respect to \mathbb{F}), and so

$$\begin{aligned} \mathbb{E} \left[\varphi_1(\mu) \int_{\mathcal{X}} \varphi_3 d\mu \right] \int_{\mathcal{C}^m} \varphi_2 d\mathcal{W}^m &= \mathbb{E} [\varphi_1(\mu) \varphi_3(W, \Lambda, X)] \int_{\mathcal{C}^m} \varphi_2 d\mathcal{W}^m \\ &= \mathbb{E} [\varphi_1(\mu) \varphi_2(W) \varphi_3(W, \Lambda, X)] \\ &= \mathbb{E} \left[\varphi_1(\mu) \int_{\mathcal{X}} \varphi_2(w) \varphi_3(w, q, x) \mu(dw, dq, dx) \right], \end{aligned}$$

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

$$\int_{\mathcal{C}^m} \varphi_2 d\mathcal{W}^m \int_{\mathcal{X}} \varphi_3(w, q, x) \mu(dw, dq, dx) = \int_{\mathcal{X}} \varphi_2(w) \varphi_3(w, q, x) \mu(dw, dq, dx), \text{ a.s.}$$

Sine this holds for all φ_2 and φ_3 , the proof is completed by arguing with a countable family of (φ_2, φ_3) , dense in the family of bounded measurable functions under pointwise convergence.

Second step. The second claim is proven in two steps. First, we show that \mathcal{F}_t^Λ is conditionally independent of $\mathcal{F}_T^{B,W,\mu}$ given $\mathcal{F}_t^{B,W,\mu}$, for each $t \in [0, T]$. Here we define $(\mathcal{F}_t^{B,W,\mu})_{t \in [0, T]}$ and $(\mathcal{F}_t^{B,\mu})_{t \in [0, T]}$ by

$$\begin{aligned}\mathcal{F}_t^{B,W,\mu} &= \sigma(B_s, W_s, \mu(C) : s \leq t, C \in \mathcal{F}_t^\mathcal{X}) \\ \mathcal{F}_t^{B,\mu} &= \sigma(B_s, \mu(C) : s \leq t, C \in \mathcal{F}_t^\mathcal{X}).\end{aligned}$$

Define also $\mathcal{F}_t^{\Lambda,X}$ and \mathcal{F}_t^W in the natural way. Let $\varphi_t : \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$ be $\mathcal{F}_t^{\Lambda,X}$ -measurable, let $\varphi_t^w : \mathcal{C}^m \rightarrow \mathbb{R}$ be \mathcal{F}_t^W -measurable, let $\varphi_{t+}^w : \mathcal{C}^m \rightarrow \mathbb{R}$ be $\sigma(W_s - W_t : s \in [t, T])$ -measurable, let $\psi_T : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_T^{B,\mu}$ -measurable, and let $\psi_t : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_t^{B,\mu}$ -measurable. Assume all of these functions are bounded. We first compute

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

This shows that

$$\mathbb{E} [\psi_T(B, \mu) \varphi_{t+}^w(W) | \mathcal{F}_t^{B,W,\mu}] = \mathbb{E} [\psi_T(B, \mu) | \mathcal{F}_t^{B,\mu}] \int_{\mathcal{C}^m} \varphi_{t+}^w d\mathcal{W}^m. \quad (5.1)$$

On the other hand, using the first result of this Lemma, we get, almost surely,

$$\int_{\mathcal{X}} \varphi_t \varphi_{t+}^w \varphi_t^w d\mu = \int_{\mathcal{X}} \varphi_t \varphi_t^w d\mu \int_{\mathcal{C}^m} \varphi_{t+}^w d\mathcal{W}^m, \quad (5.2)$$

Note also that $\int_{\mathcal{X}} \varphi_t \varphi_t^w d\mu = \int_{\mathcal{X}} \varphi_t(q, x) \varphi_t^w(w) \mu(dw, dq, dx)$ is $\mathcal{F}_t^{B,\mu}$ -measurable, since $\varphi_t(\Lambda, X) \varphi_t^w(W)$ is $\mathcal{F}_t^{W,\Lambda,X}$ -measurable. Putting it together (further explanations follow the

computations):

$$\begin{aligned}
& \mathbb{E} [\varphi_t(\Lambda, X) \psi_T(B, \mu) \varphi_{t+}^w(W) \psi_t(B, \mu) \varphi_t^w(W)] \\
&= \mathbb{E} \left[\left(\int_{\mathcal{X}} \varphi_t \varphi_{t+}^w \varphi_t^w d\mu \right) \psi_T(B, \mu) \psi_t(B, \mu) \right] \\
&= \mathbb{E} \left[\left(\int_{\mathcal{X}} \varphi_t \varphi_t^w d\mu \right) \psi_T(B, \mu) \psi_t(B, \mu) \right] \int_{\mathcal{C}^m} \varphi_{t+}^w d\mathcal{W}^m \\
&= \mathbb{E} \left[\left(\int_{\mathcal{X}} \varphi_t \varphi_t^w d\mu \right) \mathbb{E} [\psi_T(B, \mu) | \mathcal{F}_t^{B, \mu}] \psi_t(B, \mu) \right] \int_{\mathcal{C}^m} \varphi_{t+}^w d\mathcal{W}^m \\
&= \mathbb{E} [\varphi_t(\Lambda, X) \varphi_t^w(W) \mathbb{E} [\psi_T(B, \mu) | \mathcal{F}_t^{B, \mu}] \psi_t(B, \mu)] \int_{\mathcal{C}^m} \varphi_{t+}^w d\mathcal{W}^m \\
&= \mathbb{E} [\mathbb{E} [\varphi_t(\Lambda, X) | \mathcal{F}_t^{B, W, \mu}] \mathbb{E} [\psi_T(B, \mu) | \mathcal{F}_t^{B, \mu}] \psi_t(B, \mu) \varphi_t^w(W)] \int_{\mathcal{C}^m} \varphi_{t+}^w d\mathcal{W}^m \\
&= \mathbb{E} [\mathbb{E} [\varphi_t(\Lambda, X) | \mathcal{F}_t^{B, W, \mu}] \mathbb{E} [\psi_T(B, \mu) \varphi_{t+}^w(W) | \mathcal{F}_t^{B, W, \mu}] \psi_t(B, \mu) \varphi_t^w(W)],
\end{aligned}$$

the first and fourth equalities following from hypothesis (3), the second from (5.2), the third from the fact that $\int \varphi_t(q, x) \varphi_t^w(w) \mu(dw, dq, dx)$ is $\mathcal{F}_t^{B, \mu}$ -measurable, and the sixth from (5.1).

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

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

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

Last step. It now remains to prove that \mathcal{F}_t^{Λ} is conditionally independent of $\mathcal{F}_T^{X_0, B, W, \mu}$ given $\mathcal{F}_t^{X_0, B, W, \mu}$ for each t , which is slightly different from the result of the previous step. Fix $t \in [0, T]$. Let $\varphi_t : \mathcal{V} \rightarrow \mathbb{R}$ be \mathcal{F}_t^{Λ} -measurable, $\psi_t : \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_t^{B, W, \mu}$ -measurable, $\psi_T : \mathcal{C}^{m_0} \times \mathcal{C}^m \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be $\mathcal{F}_T^{B, W, \mu}$ -measurable and $\zeta_0 : \mathbb{R} \rightarrow \mathbb{R}$ be Borel measurable. Assume all of these functions are bounded. From the previous step, we deduce that

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

where the first equality follows from the conditional independence of $\mathcal{F}_t^{\Lambda, X}$ and $\mathcal{F}_T^{B, W, \mu}$ given $\mathcal{F}_t^{B, W, \mu}$. In order to complete the proof, notice that $\mathbb{E} [\psi_T(B, W, \mu) | \mathcal{F}_t^{B, W, \mu}] = \mathbb{E} [\psi_T(B, W, \mu) | \mathcal{F}_t^{X_0, B, W, \mu}]$ since X_0 and (B, W, μ) are independent by hypothesis (2). There-

fore, for another bounded Borel measurable function $\zeta'_0 : \mathbb{R} \rightarrow \mathbb{R}$, we get

$$\begin{aligned} & \mathbb{E} [\varphi_t(\Lambda) \psi_T(B, W, \mu) \psi_t(B, W, \mu) \zeta_0(X_0) \zeta'_0(X_0)] \\ &= \mathbb{E} \left[\varphi_t(\Lambda) \mathbb{E} \left[\psi_T(B, W, \mu) | \mathcal{F}_t^{X_0, B, W, \mu} \right] \psi_t(B, W, \mu) \zeta_0(X_0) \zeta'_0(X_0) \right] \\ &= \mathbb{E} \left[\varphi_t(\Lambda) \mathbb{E} \left[\zeta'_0(X_0) \psi_T(B, W, \mu) | \mathcal{F}_t^{X_0, B, W, \mu} \right] \psi_t(B, W, \mu) \zeta_0(X_0) \right], \end{aligned}$$

which proves that \mathcal{F}_t^Λ and $\mathcal{F}_t^{X_0, B, W, \mu}$ are conditionally independent given $\mathcal{F}_t^{X_0, B, W, \mu}$. \square

5.2 Canonical space

This section briefly elaborates on the notion of mean field game solution on the canonical space. This will allow us to state simpler conditions by which may check that a measure $P \in \mathcal{P}(\Omega)$ is a weak MFG solution, in the sense of Definition 3.1.2, and also to streamline certain weak continuity results. First, we mention some notational conventions. We will routinely use the same letter φ to denote the natural extension of a function $\varphi : E \rightarrow F$ to any product space $E \times E'$, given by $\varphi(x, y) := \varphi(x)$ for $(x, y) \in E \times E'$. Similarly, we will use the same symbol $(\mathcal{F}_t)_{t \in [0, T]}$ to denote the natural extension of a filtration $(\mathcal{F}_t)_{t \in [0, T]}$ on a space E to any product space $E \times E'$, given by $(\mathcal{F}_t \otimes \{\emptyset, E'\})_{t \in [0, T]}$.

We will make heavy use of the following canonical spaces, two of which have been defined already but are recalled for convenience:

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

From now on, let ξ , B , W , μ , Λ , and X denote the identity maps on \mathbb{R}^d , \mathcal{C}^{m_0} , \mathcal{C}^m , $\mathcal{P}^p(\mathcal{X})$, \mathcal{V} , and \mathcal{C}^d , respectively. Note, for example, that our convention permits W to denote both the identity map on \mathcal{C}^m and the projection from Ω to \mathcal{C}^m .

The canonical processes B , W , and X generate obvious natural filtrations, on \mathcal{C}^{m_0} , \mathcal{C}^m , and \mathcal{C}^d , denoted $\mathbb{F}^B = (\mathcal{F}_t^B)_{t \in [0, T]}$, $\mathbb{F}^W = (\mathcal{F}_t^W)_{t \in [0, T]}$, and $\mathbb{F}^X = (\mathcal{F}_t^X)_{t \in [0, T]}$, respectively. The natural filtration $\mathbb{F}^\Lambda = (\mathcal{F}_t^\Lambda)_{t \in [0, T]}$ is defined by

$$\mathcal{F}_t^\Lambda := \sigma(1_{[0, t]} \Lambda) = \sigma(\Lambda(C) : C \in \mathcal{B}([0, t] \times A)). \quad (5.3)$$

We will abuse notation somewhat by writing Λ both for $(\Lambda_t)_{t \in [0, T]}$, the canonical $\mathcal{P}(A)$ -valued process on \mathcal{V} (see Lemma 2.1.14), and also for $\Lambda(dt, da)$, the identity map on \mathcal{V} . Define the canonical filtration $\mathbb{F}^\mu = (\mathcal{F}_t^\mu)_{t \in [0, T]}$ on $\mathcal{P}^p(\mathcal{X})$ by

$$\mathcal{F}_t^\mu := \sigma(\mu(C) : C \in \mathcal{F}_t^\mathcal{X}).$$

We will frequently work with filtrations generated by several canonical processes, such as $\mathbb{F}^{\xi, B, W} = (\mathcal{F}_t^{\xi, B, W} := \sigma(\xi, B_s, W_s : s \leq t))_{t \in [0, T]}$ defined on Ω_0 , and $\mathbb{F}^{\xi, B, W, \Lambda} = (\mathcal{F}_t^{\xi, B, W, \Lambda} := \mathcal{F}_t^{\xi, B, W} \otimes \mathcal{F}_t^\Lambda)$ defined on $\Omega_0 \times \mathcal{V}$. Our convention on canonical extensions of filtrations to product spaces permits the use of $\mathbb{F}^{\xi, B, W}$ to refer also to the filtration on $\Omega_0 \times \mathcal{V}$ generated by (ξ, B, W) , and it should be clear from context on which space the filtration is defined.

Hence, the filtration $(\mathcal{F}_t^{\mathcal{X}})_{t \in [0, T]}$ defined just before Definition 3.1.1 could alternatively be denoted $\mathcal{F}_t^{\mathcal{X}} = \mathcal{F}_t^{W, \Lambda, X}$, but we stick with the former notation for consistency.

There is somewhat of a conflict in notation, between our use of (ξ, B, W) here as the identity map on $\mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m$ and our previous use (beginning in Section 3.2) of the same letters for random variables with values in $(\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n$, defined on an n -player environment $\mathcal{E}_n = (\Omega_n, (\mathcal{F}_t^n)_{t \in [0, T]}, \mathbb{P}_n, \xi, B, W)$. However, we will almost exclusively discuss the random variables (ξ, B, W) through the lenses of various probability measures, and thus it should be clear from context (i.e., from the nearest notated probability measure) which random variables (ξ, B, W) we are working with at any given moment. For example, given $P \in \mathcal{P}(\Omega)$, the notation $P \circ (\xi, B, W)^{-1}$ refers to a measure on $\mathbb{R}^d \times \mathcal{C}^{m_0} \times \mathcal{C}^m$. On the other hand, \mathbb{P}_n is reserved for the measure on Ω_n in a typical n -player environment, and so $\mathbb{P}_n \circ (\xi, B, W)^{-1}$ refers to a measure on $(\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n$. In any case, this is only be a potential issue in Chapter 6.

We next specify how μ and Λ are allowed to correlate with each other and with the given sources of randomness (ξ, B, W) . It will be useful to fix some notation for the joint law

$$\mathcal{W}_\lambda := \lambda \times \mathcal{W}^{m_0} \times \mathcal{W}^m \in \mathcal{P}^{p'}(\Omega_0), \quad (5.4)$$

where \mathcal{W}^k denotes Wiener measure on \mathcal{C}^k for any positive integer k ; note that p' -integrability follows from the assumption $\lambda \in \mathcal{P}^{p'}(\mathbb{R}^d)$. Recall that the conditional independence requirement in (3) of Definition 3.1.1 is referred to as *compatibility*, and we will require also that μ is compatible in a slightly different sense.

1. An element $\rho \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}))$ is said to be in $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ if (ξ, B, W) has law \mathcal{W}_λ under ρ and if B and W are independent $\mathbb{F}^{\xi, B, W, \mu}$ -Wiener processes under ρ . The subscript c and the symbol \rightsquigarrow in $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ indicate that the extension of the probability measure \mathcal{W}_λ from Ω_0 to $\Omega_0 \times \mathcal{P}^p(\mathcal{X})$ is *compatible*.
2. For $\rho \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}))$, an element $Q \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$ is said to be in $\mathcal{P}_c^p[(\Omega_0 \times \mathcal{P}^p(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}]$ if (ξ, B, W, μ) has law ρ under Q and $\mathcal{F}_T^{\xi, B, W, \mu}$ and \mathcal{F}_t^Λ are conditionally independent given $\mathcal{F}_t^{\xi, B, W, \mu}$. Again, Q is then *compatible* with ρ in the sense that, given the observation of (ξ, B, W, μ) up until time t , the observation of Λ up until t has no influence on the future of (ξ, B, W, μ) .

Remark 5.2.1. These notions of compatibility are both special cases of a more general and well understood idea, which goes by several names in the literature. We will elaborate on this in Appendix A.1, but to keep the presentation self-contained, we will derive the needed results as we go.

We now have enough material to describe the relevant control problems. For a given joint law $\rho \in \mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}))$ of the original sources of randomness (ξ, B, W) and a random measure μ , we denote by

$$\mathcal{A}(\rho) := \mathcal{P}_c^p[(\Omega_0 \times \mathcal{P}^p(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}] \quad (5.5)$$

the *set of admissible relaxed controls* (see (2) above). In words, the controls are specified by joint laws of $(\xi, B, W, \mu, \Lambda)$ that are compatible with the given joint law of (ξ, B, W, μ) , in the above sense.

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

For $Q \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}]$ we have that Λ is p -integrable in the sense that

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

On the completion of the filtered probability space $(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}, \mathbb{F}^{\xi, B, W, \mu, \Lambda}, Q)$ there exists a unique strong solution Y of the SDE

$$Y_t = \xi + \int_0^t ds \int_A \Lambda_s(da) b(s, Y_s, \mu_s^x, a) + \int_0^t \sigma(s, Y_s, \mu_s^x) dW_s + \int_0^t \sigma_0(s, Y_s, \mu_s^x) dB_s.$$

where we recall that $\mu^x(\cdot) = \mu(\mathcal{C}^m \times \mathcal{V} \times \cdot)$ is the marginal law of μ on \mathcal{C}^d and μ_s^x is the time- s marginal. We then denote by

$$\mathcal{R}(Q) := Q \circ (\xi, B, W, \mu, \Lambda, Y)^{-1} \in \mathcal{P}(\Omega) \quad (5.6)$$

the joint law of the solution and the inputs. Equivalently (by Theorem A.2.3), $\mathcal{R}(Q)$ is the unique element P of $\mathcal{P}(\Omega)$ such that $P \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q$ and such that the canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the SDE

$$X_t = \xi + \int_0^t ds \int_A \Lambda_s(da) b(s, X_s, \mu_s^x, a) + \int_0^t \sigma(s, X_s, \mu_s^x) dW_s + \int_0^t \sigma_0(s, X_s, \mu_s^x) dB_s. \quad (5.7)$$

under P . (As before, we do not augment our filtrations; see footnote 1 on page 54 for a related discussion.) A standard estimate (see Lemma 5.3.1 below) will show that this measure is p -integrable, i.e. $\mathcal{R}(Q) \in \mathcal{P}^p(\Omega)$.

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

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

Recalling the definition of Γ from (3.3), the expected reward functional $J : \mathcal{P}^p(\Omega) \rightarrow \mathbb{R}$ is defined by

$$J(P) := \mathbb{E}^P [\Gamma(\mu^x, \Lambda, X)]. \quad (5.8)$$

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

$$\begin{aligned} \mathcal{A}^*(\rho) &:= \arg \max_{Q \in \mathcal{A}(\rho)} J(\mathcal{R}(Q)) \\ &= \left\{ Q \in \mathcal{A}(\rho) : J(\mathcal{R}(Q)) = \sup_{Q' \in \mathcal{A}(\rho)} J(\mathcal{R}(Q')) \right\}, \end{aligned} \quad (5.9)$$

and note that

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

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

We may now reformulate some of the definitions and results from before. First, we simply rewrite the definitions of MFG solution (Definition 3.1.1) and pre-solution (Definition 5.1.1) in the new notation. Lemma 5.2.3 uses Lemma 5.1.2 to provide what will be the most useful description of MFG solutions and pre-solutions.

Definition 5.2.2. A measure $P \in \mathcal{P}^p(\Omega)$ is a MFG pre-solution if and only if it satisfies the following:

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

Similarly, a measure $P \in \mathcal{P}^p(\Omega)$ is a MFG solution if and only if it satisfies (1-4) as well as $P \in \mathcal{RA}^*(\rho)$.

Lemma 5.2.3. Suppose $P \in \mathcal{P}^p(\Omega)$ satisfies the following hold:

1. B and W are independent $\mathbb{F}^{\xi, B, W, \mu, \Lambda, X}$ -Wiener processes, and $P \circ \xi^{-1} = \lambda$.
2. ξ , W , and (B, μ) are independent.
3. $\mu = P((W, \Lambda, X) \in \cdot \mid B, \mu)$, a.s.
4. The canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the state equation (5.7) on Ω .

Then P is a MFG pre-solution.

Proof. The only point that needs to be checked is the compatibility, i.e. that $P \circ (\xi, B, W, \mu, \Lambda)^{-1}$ is in $\mathcal{A}(\rho)$ where $\rho = P \circ (\xi, B, W, \mu)^{-1}$. Equivalently, we need to check that \mathcal{F}_t^Λ is conditionally independent of $\mathcal{F}_T^{\xi, B, W, \mu}$ given $\mathcal{F}_t^{\xi, B, W, \mu}$, for each $t \in [0, T]$. But this follows from Lemma 5.1.2. \square

5.3 Continuity results

Before beginning the proofs of the main theorems, it is useful to study the continuity properties of the maps \mathcal{R} , \mathcal{A} , and J defined in the previous section. Let us begin with a version of a standard integral estimate for the state process. To this end, define the truncated supremum norms $\|\cdot\|_t$ on \mathcal{C}^d by

$$\|x\|_t := \sup_{s \in [0, t]} |x_s|, \quad t \in [0, T]. \quad (5.10)$$

Lemma 5.3.1. *On some filtered probability space $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P)$, suppose B and W are independent \mathbb{F} -Wiener processes of dimension m and m_0 , suppose μ is a $\mathcal{P}(\mathcal{C}^d)$ -valued random variable such that $\mu(C)$ is \mathcal{F}_t -measurable for each $C \in \mathcal{F}_t^X$ and each $t \in [0, T]$, suppose Λ is an \mathbb{F} -progressive $\mathcal{P}^p(A)$ -valued process, and suppose ξ is a \mathcal{F}_0 -measurable random vector with law λ . Assume **A1** and **A2** hold. Then there exists a unique solution X of the state equation*

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

Moreover, for each $\gamma \in [p, p']$, there exists a constant $c_4 > 0$, depending only on γ, λ, T , and the constant c_1 of (A1.4) such that,

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

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

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

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

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

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

then the above becomes

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

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

Next, we state an immensely useful compactness result, the proof of which is deferred to Appendix B. Recall that $p' > p \geq 1$. The result could likely be sharpened with respect to uniform integrability requirements, but this will be good enough for our purposes.

Proposition 5.3.2. *Let d be a positive integer, and fix $c > 0$. Let $\mathcal{Q} \subset \mathcal{P}(\mathcal{V} \times \mathcal{C}^d)$ be the set of laws of $\mathcal{V} \times \mathcal{C}^d$ -valued random variables (Λ, X) defined on some complete filtered probability space $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P)$ satisfying satisfying*

$$dX_t = \int_A B(t, a) \Lambda_t(da) dt + \Sigma(t) dW_t,$$

where the following hold:

1. W is a \mathbb{F} -Wiener process of dimension k , and X is a continuous \mathbb{F} -adapted d -dimensional process with $P \circ X_0^{-1} = \lambda$.
2. $\Sigma : [0, T] \times \tilde{\Omega} \rightarrow \mathbb{R}^{d \times k}$ is progressively measurable, and $B : [0, T] \times \tilde{\Omega} \times A \rightarrow \mathbb{R}^d$ is jointly measurable with respect to the \mathbb{F} -progressive σ -field on $[0, T] \times \tilde{\Omega}$ and the Borel σ -field on A .
3. There exists a nonnegative \mathcal{F}_T -measurable random variable Z such that, for each $(t, \omega, a) \in [0, T] \times \tilde{\Omega} \times A$,

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

and

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

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

Let us put this to use to study the continuity of \mathcal{R} . Recall that topological statements involving a space \mathcal{P}^p are meant with respect to the p -Wasserstein metric. Let us define also some norms on

Lemma 5.3.3. *Suppose $K \subset \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}]$ satisfies*

$$\sup_{Q \in K} \mathbb{E}^Q \left[\int_{\mathcal{C}^d} \|x\|_T^{p'} \mu^x(dx) + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \right] < \infty. \quad (5.11)$$

Then the map $\mathcal{R} : K \rightarrow \mathcal{P}^p(\Omega)$ is continuous.

Proof. Let $Q_n \rightarrow Q$ in $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X}) \times \mathcal{V}]$. Note that $\mathcal{R}(Q_n) \circ (X_0, B, W, \mu, \Lambda)^{-1} = Q_n$ are relatively compact in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$. The integrability hypothesis (5.11) permits an application of Proposition 5.3.2 to conclude that $\mathcal{R}(Q_n) \circ X^{-1}$ are relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$, and thus $\mathcal{R}(Q_n)$ are relatively compact in $\mathcal{P}^p(\Omega)$ (see Proposition 2.1.8). Let P be any limit point, so $\mathcal{R}(Q_{n_k}) \rightarrow P$ for some n_k . Then

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

The canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the same state equation under each of the measures Q_n , and it follows (see Kurtz and Protter [81], or Jacod and Mémmin [70]) that the canonical processes verify the SDE (5.7) under P . Hence, $P = \mathcal{R}(Q)$. \square

Let us turn now to the objective functional J , the continuity properties of which will mostly follow from the preliminary material of Section 2.1.1.

Lemma 5.3.4. *The map $J : \mathcal{P}^p(\Omega) \rightarrow \mathbb{R}$ is upper semicontinuous. Moreover, the restriction of J to a set $K \subset \mathcal{P}^p(\Omega)$ is continuous whenever K satisfies the uniform integrability condition*

$$\lim_{r \rightarrow \infty} \sup_{P \in K} \mathbb{E}^P \left[\int_0^T \int_{\{|a| > r\}} |a|^{p'} \Lambda_t(da) dt \right] = 0. \quad (5.12)$$

In particular, if A is compact, then J is continuous everywhere.

Proof. Since f and g are continuous in (x, μ, a) , the upper bounds of f and g (which grow in order p in (x, μ)) along with Lemma 2.1.10 imply both that Γ is upper semicontinuous and then also that J is upper semicontinuous from $\mathcal{P}^p(\Omega)$ to \mathbb{R} .

Under the additional assumption, the claimed continuity is intuitively quite clear, but a proof should be careful of the two exponents $p' > p$ and the potential discontinuity of f in t . Let $P_n \rightarrow P_\infty$ in $\mathcal{P}^p(\Omega)$ with $P_n \in K$ for each n . The continuity and growth assumptions on g imply that $\mathbb{E}^{P_n}[g(X_T, \mu_T^x)] \rightarrow \mathbb{E}^P[g(X_T, \mu_T^x)]$, and the f term causes the only problems. The convergence $P_n \rightarrow P_\infty$ implies (e.g., by [107, Theorem 7.12])

$$\lim_{r \rightarrow \infty} \sup_n \mathbb{E}^{P_n} \left[\|X\|_T^p 1_{\{\|X\|_T^p > r\}} + \int_{\mathcal{C}^d} \|z\|_T^p \mu^x(dz) 1_{\{\int_{\mathcal{C}^d} \|z\|_T^p \mu^x(dz) > r\}} \right] = 0. \quad (5.13)$$

For $1 \leq n \leq \infty$, define probability measures Q_n on $\tilde{\Omega} := [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$ by

$$Q_n(C) := \frac{1}{T} \mathbb{E}^{P_n} \left[\int_0^T \int_A 1_{\{(t, X_t, \mu_t^x, a) \in C\}} \Lambda_t(da) dt \right], \quad C \in \mathcal{B}(\tilde{\Omega}).$$

Certainly $Q_n \rightarrow Q_\infty$ weakly in $\mathcal{P}(\tilde{\Omega})$. Since the $[0, T]$ -marginal is the same for each Q_n , this implies (by Lemma 2.1.4) $\int \varphi dQ_n \rightarrow \int \varphi dQ_\infty$ for each bounded measurable $\varphi : \tilde{\Omega} \rightarrow \mathbb{R}$ with $\varphi(t, \cdot)$ continuous for each t . Thus $Q_n \circ f^{-1} \rightarrow Q_\infty \circ f^{-1}$ weakly in $\mathcal{P}(\mathbb{R})$, by continuity

of $f(t, \cdot)$ for each t . But it follows from (5.12), (5.13), and the growth assumption of A1.5 that

$$\lim_{r \rightarrow \infty} \sup_n \int_{\{|f| > r\}} f dQ_n = 0,$$

and thus $\int f dQ_n \rightarrow \int f dQ_\infty$. \square

Finally, the following density result is critical, and it is simply a corollary of Proposition 2.1.15.

Definition 5.3.5. For $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$, let $\mathcal{A}_a(\rho) \subset \mathcal{A}(\rho)$ denote the set of measures of the form

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

where the function $\varphi : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{V}$ satisfies the following:

1. φ is *adapted*, in the sense that $\varphi^{-1}(C) \in \mathcal{F}_t^{\xi, B, W, \mu}$ for each $C \in \mathcal{F}_t^\Lambda$ and $t \in [0, T]$.
2. φ is continuous.
3. φ is *bounded*, in the sense that there exists a compact set $K \subset A$ such that $\varphi(\omega, \nu)([0, T] \times K^c) = 0$ for each $(\omega, \nu) \in \Omega_0 \times \mathcal{P}^p(\mathcal{X})$.
4. φ takes values in the set of strict controls, $\{dt\delta_{\alpha_t}(da) : \alpha : [0, T] \rightarrow A \text{ measurable}\}$.

Corollary 5.3.6. For each $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$, $\mathcal{A}_a(\rho)$ is a dense subset of $\mathcal{A}(\rho)$ (in the p -Wasserstein topology).

Proof. This is a consequence of Proposition 2.1.15. For $\nu \in \mathcal{P}(\mathcal{X})$, define $\nu^t := \nu \circ (W_{\cdot \wedge t}, \Lambda_{\cdot \wedge t}, X_{\cdot \wedge t})^{-1}$, where we identify as usual the $\mathcal{P}(A)$ -valued process $(\Lambda_{s \wedge t})_{s \in [0, T]}$ with the random measure $ds\Lambda_{s \wedge t}(da)$. Then set $S_t = (\xi, B_{\cdot \wedge t}, W_{\cdot \wedge t}, \mu^t)$. \square

An important consequence of this result is Proposition 5.3.7 below, which ensures that in any of the control problems we face we need only to check optimality against the dense class of strong controls, not against all admissible relaxed controls.

Proposition 5.3.7. Let $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$, and let $Q \in \mathcal{A}(\rho)$. If $J(\mathcal{R}(Q)) \geq J(\mathcal{R}(Q'))$ for each $Q' \in \mathcal{A}_a(\rho)$, then $Q \in \mathcal{A}^*(\rho)$.

Proposition 5.3.7 follows immediately from Corollary 5.3.6 and the following Lemma 5.3.8, which will be useful again later:

Lemma 5.3.8. Let $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ satisfy

$$\mathbb{E}^\rho \int_{\mathcal{C}^d} \|x\|_T^{p'} \mu^x(dx) < \infty.$$

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

Proof. This would follow immediately from Corollary 5.3.6 if we could ensure that the approximations were sufficiently uniformly integrable; see Lemma 5.3.4. We will proceed somewhat differently.

First step. First, assume $P \in \mathcal{RA}(\rho)$ satisfies

$$P(\Lambda([0, T] \times K^c) = 0) = 1,$$

for some compact set $K \subset A$. That is, Λ is bounded in the sense of Definition 5.3.5. Thanks to the lower bounds on f and g of assumption A1.4, we know $J(P') > -\infty$. Write $P = \mathcal{R}(Q)$, where $Q \in \mathcal{A}(\rho)$. By Corollary 5.3.6, there exist $Q_n \in \mathcal{RA}_a(\rho)$ such that $Q_n \rightarrow Q$ in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$. Since K is compact, $J(\mathcal{R}(Q_n)) \rightarrow J(\mathcal{R}(Q)) = J(P)$ by Lemma 5.3.4.

Second step. To prove the general case, assume now that $P \in \mathcal{RA}(\rho)$ satisfies $J(P) > -\infty$. Because of the first step, it suffices to show that there exist $P_n \in \mathcal{RA}(\rho)$ such that $J(P) = \lim_{n \rightarrow \infty} J(P_n)$ and such that for each n there exists a compact $K_n \subset A$ such that

$$P_n(\Lambda([0, T] \times K_n^c) = 0) = 1. \quad (5.14)$$

First, the upper bounds of f and g of assumption A1.4 imply

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

Thanks to our hypothesis and Lemma 5.3.1, we have $\mathbb{E}^P \int_{\mathcal{C}^d} \|x\|_T^p \mu(dx) < \infty$ and $\mathbb{E}^P \|X\|_T^p < \infty$. This implies

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

Now fix any compact sets $K_1 \subset K_2 \subset \dots$ with $\cup_n K_n = A$. Let $\iota_n : A \rightarrow A$ denote any measurable function satisfying $\iota_n(A) \subset K_n$ and $\iota_n(a) = a$ for all $a \in A_n$, so that ι_n converges pointwise to the identity. Let $\Lambda_t^n = \Lambda_t \circ \iota_n^{-1}$, and as usual identify $\Lambda^n = dt \Lambda_t^n(da)$. Let $Q_n := P \circ (\xi, B, W, \mu, \Lambda^n)^{-1}$ and $P_n = \mathcal{R}(Q_n)$, so that $P_n \in \mathcal{RA}(\rho)$ satisfies (5.14). Since $\Lambda^n \rightarrow \Lambda$ P -a.s., it follows that $Q_n \rightarrow Q$ in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$, where $Q := P \circ (\xi, B, W, \mu, \Lambda)^{-1}$ satisfies $P = \mathcal{R}(Q)$. Now, since $|\iota_n(a)| \leq |a|$, we have

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

which implies that the sequence

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

is uniformly P' -integrable. By continuity of \mathcal{R} and J (see Lemmas 5.3.3 and 5.3.4), $P_n = \mathcal{R}(Q_n) \rightarrow \mathcal{R}(Q) = P$ in $\mathcal{P}^p(\Omega)$ and thus $J(P_n) \rightarrow J(P)$. \square

Remark 5.3.9. At the expense of being overly verbose, it is interesting now to note that an equivalent simplified definition of MFG solution is available which makes no explicit mention of compatibility. Indeed, a tuple $(\tilde{\Omega}, \mathbb{F}, P, B, W, \mu, \Lambda, X)$ is a weak MFG solution if and only if it satisfies properties (1-5) of Definition 3.1.1, with the compatibility requirement (i.e., conditional independence) removed from (3), and in place of (6) the following optimality condition: For each $\mathbb{F}^{X_0, B, W, \mu}$ -progressively measurable A -valued process α , defined on $\tilde{\Omega}$ and satisfying

$$\mathbb{E} \int_0^T |\alpha_t|^p dt < \infty,$$

we have

$$\mathbb{E} [\Gamma(\mu^x, \Lambda, X)] \geq \mathbb{E} [\Gamma(\mu^x, dt\delta_{\alpha_t}(da), Y)],$$

where

$$Y_t = X_0 + \int_0^t b(s, Y_s, \mu_s^x, \alpha_s) ds + \int_0^t \sigma(s, Y_s, \mu_s^x) dW_s + \int_0^t \sigma_0(s, Y_s, \mu_s^x) dB_s.$$

The analogous statement is true in the case without common noise; analogs of Lemma 5.1.2 and 5.2.3 hold with the same proofs.

5.4 Proofs of Propositions 3.1.4 and 3.1.5

This section puts some of these preliminary results to use in proving Propositions 3.1.4 and 3.1.5. First, we quote a measurable selection theorem due to Haussmann and Lepeltier [63]. It will be used to construct optimal strict controls from optimal relaxed controls, both in this section and later in Section 8.1.

Proposition 5.4.1 (Theorem A.9 of [63]). *Let (E, \mathcal{E}) be a measurable space, let A be a closed subset of a Euclidean space, and let n be a positive integer. Consider measurable functions*

$$(\varphi, \psi) : E \times A \rightarrow \mathbb{R}^n \times \mathbb{R}, \quad (\tilde{\varphi}, \tilde{\psi}) : E \rightarrow \mathbb{R}^n \times \mathbb{R},$$

such that $a \mapsto (\varphi, \psi)(x, a)$ is continuous for each $x \in E$. Suppose finally that for each $x \in E$ we have $(\tilde{\varphi}(x), \tilde{\psi}(x)) \in K(x)$, where

$$K(x) := \{(\varphi(x, a), z) : a \in A, z \leq \psi(x, a)\}.$$

Then there exists a measurable function $\hat{\alpha} : E \rightarrow A$ such that

$$\tilde{\varphi}(x) = \varphi(x, \hat{\alpha}(x)), \quad \text{and} \quad \tilde{\psi}(x) \leq \psi(x, \hat{\alpha}(x)), \quad \text{for all } x \in E.$$

5.4.1 Proof of Proposition 3.1.4

With the volatilities uncontrolled, we note that for each (t, x) the set

$$K(t, x, m) := \{(b(t, x, m, a), z) : a \in A, z \leq f(t, x, m, a)\}$$

is convex thanks to assumption **(Convex)**. As in [63, Proposition 3.5], $K(t, x, m)$ is a closed set for each (t, x, m) . Thus, since $(b, f)(t, X_t, \mu_t^x, a) \in K(t, X_t, \mu_t^x)$ for each $a \in A$, we have $\int_A \Lambda_t(da)(b, f)(t, X_t, \mu_t^x, a) \in K(t, X_t, \mu_t^x)$. Consider the completion $\mathbb{F}^{X_0, B, W, \mu, \Lambda}$ of the filtration

$$\sigma(X_0, B_s, W_s, \mu(C), \Lambda_s : s \leq t, C \in \mathcal{F}_t^X), \quad t \in [0, T].$$

Note that the SDE has a unique strong solution, so X is necessarily adapted to $\mathbb{F}^{X_0, B, W, \mu, \Lambda}$. Using the measurable selection result of Proposition 5.4.1 (choosing the space (E, \mathcal{E}) as $[0, T] \times \tilde{\Omega}$ with the $\mathbb{F}^{X_0, B, W, \mu, \Lambda}$ -progressive σ -field), we may find $\mathbb{F}^{X_0, B, W, \mu, \Lambda}$ -progressively measurable processes α and z , taking values in A and $[0, \infty)$, respectively, such that

$$\int_A \Lambda_t(da)(b, f)(t, X_t, \mu_t^x, a) = (b, f)(t, X_t, \mu_t^x, \alpha_t) - (0, z_t). \quad (5.15)$$

This implies that the following SDE holds:

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

Moreover,

$$\begin{aligned} \mathbb{E}[\Gamma(\mu^x, dt\delta_{\alpha_t}(da), X)] &= \mathbb{E}\left[\int_0^T dt f(t, X_t, \mu_t^x, \alpha_t) + g(X_T, \mu_T^x)\right] \\ &\geq \mathbb{E}\left[\int_0^T dt \int_A \Lambda_t(da) f(t, X_t, \mu_t^x, a)dt + g(X_T, \mu_T^x)\right] \\ &= \mathbb{E}[\Gamma(\mu^x, \Lambda, X)] \end{aligned} \quad (5.16)$$

Now let us define $\tilde{\Lambda} = dt\delta_{\alpha_t}(da)$, and finally

$$\tilde{\mu} := P((W, \tilde{\Lambda}, X) \in \cdot \mid B, \mu).$$

Note that by conditioning on $\tilde{\mu}$ on both sides, we get

$$\tilde{\mu} := P((W, \tilde{\Lambda}, X) \in \cdot \mid B, \tilde{\mu}).$$

Moreover, we have of course

$$\tilde{\mu}^x = P(X \in \cdot \mid B, \mu) = \mu^x.$$

It remains to prove that $(\tilde{\Omega}, \mathbb{F}, P, B, W, \tilde{\mu}, \tilde{\Lambda}, X)$ is indeed a weak MFG solution, which is done in several steps. The defining properties (1), (2), (4), and (6) of a weak MFG solution

in Definition 3.1.1 are either obvious or have been verified already. Thanks to Lemma 5.1.2 (see also its corollary, Lemma 5.2.3), the compatibility of the control follows from what we have already proven; that is, property (3) of Definition 3.1.1 also holds. It remains to check the optimality condition (5), which we do in two steps.

Step 1: As a preliminary step, let us check next that for each $t \in [0, T]$, up to null sets, we have

$$\sigma(\tilde{\mu}(C) : C \in \mathcal{F}_t^{\mathcal{X}}) \subset \mathcal{F}_t^{B, \mu} := \sigma(B_s, \mu(C) : s \leq t, C \in \mathcal{F}_t^{\mathcal{X}}).$$

Now if $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ is $\mathcal{F}_t^{\mathcal{X}}$ -measurable then

$$\begin{aligned} \int_{\mathcal{X}} \varphi d\tilde{\mu} &= \mathbb{E} \left[\varphi(W, \tilde{\Lambda}, X) \middle| B, \mu \right] = \mathbb{E} \left[\mathbb{E} \left[\varphi(W, \tilde{\Lambda}, X) \middle| \mathcal{F}_T^{X_0, B, W, \mu} \right] \middle| \mathcal{F}_T^{B, \mu} \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[\varphi(W, \tilde{\Lambda}, X) \middle| \mathcal{F}_t^{X_0, B, W, \mu} \right] \middle| \mathcal{F}_T^{B, \mu} \right] \\ &= \mathbb{E} \left[\varphi(W, \tilde{\Lambda}, X) \middle| \mathcal{F}_t^{B, \mu} \right], \end{aligned}$$

The second equality follows from the conditional independence of \mathcal{F}_t^{Λ} and $\mathcal{F}_T^{X_0, B, W, \mu}$ given $\mathcal{F}_t^{X_0, B, W, \mu}$, and from the adaptedness of $\tilde{\Lambda}$ and X with respect to $\mathbb{F}^{X_0, B, W, \mu, \Lambda}$. The last equality follows easily from the independence of (X_0, W) and (B, μ) . Since this holds for each φ , this step is complete.

Step 2: Thanks to Proposition 5.3.7, we need only check optimality among controls adapted to the filtration $\mathbb{F}^{X_0, B, W, \tilde{\mu}} = (\mathcal{F}_t^{X_0, B, W, \tilde{\mu}})_{t \in [0, T]}$ given by

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

Fix a bounded adapted function $\varphi : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{V}$ (see Definition 5.3.5). Define the relaxed control

$$\beta = \varphi(X_0, B, W, \tilde{\mu}),$$

and the corresponding controlled state process

$$Y_t = X_0 + \int_0^t \int_A b(s, Y_s, \tilde{\mu}_s^x, a) \beta_s(da) ds + \int_0^t \sigma(s, Y_s, \tilde{\mu}_s^x) dW_s + \int_0^t \sigma_0(s, Y_s, \tilde{\mu}_s^x) dB_s.$$

The proof will be complete if we show that

$$\mathbb{E} \left[\Gamma(\tilde{\mu}^x, \tilde{\Lambda}, X) \right] \geq \mathbb{E} \left[\Gamma(\tilde{\mu}^x, \beta, Y) \right].$$

Since $(\beta_t)_{t \in [0, T]}$ is adapted to $\mathbb{F}^{X_0, B, W, \tilde{\mu}}$, it is in fact adapted $\mathbb{F}^{X_0, B, W, \mu}$ as well, thanks to the result of Step 2. But this follows from (5.16) as follows:

$$\begin{aligned} \mathbb{E} \left[\Gamma(\tilde{\mu}^x, \tilde{\Lambda}, X) \right] &= \mathbb{E} \left[\Gamma(\mu^x, \tilde{\Lambda}, X) \right] \geq \mathbb{E} \left[\Gamma(\mu^x, \Lambda, X) \right] \\ &\geq \mathbb{E} \left[\Gamma(\mu^x, \beta, Y) \right] = \mathbb{E} \left[\Gamma(\tilde{\mu}^x, \beta, Y) \right]. \end{aligned}$$

Indeed, the two equalities follow from the fact that $\mu^x = \tilde{\mu}^x$, the first inequality from (5.16), and the second inequality from the optimality property satisfied by the given MFG solution. \square

5.4.2 Proof of Proposition 3.1.5

It suffices to prove the following two statements, for each $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$

- (1) If $\mathcal{A}^*(\rho)$ is nonempty, it contains a strong control.
- (2) $\mathcal{A}^*(\rho)$ contains at most one element.

Indeed, then any MFG solution P must have strong control, since $P \in \mathcal{RA}^*(\rho)$ for some ρ . We will work on the canonical space Ω for this proof. Fix ρ throughout.

Proof of (1). Suppose $P \in \mathcal{RA}^*(\rho)$. Under assumption **(Linear-Convex)**, the state equation writes as

$$\begin{aligned} X_t = X_0 &+ \int_0^t (b^1(s, \mu_s^x)X_s + b^2(s, \mu_s^x)\alpha_s + b^3(t, \mu_s^x)) ds \\ &+ \int_0^t (\sigma^1(s, \mu_s^x)X_s + \sigma^2(s, \mu_s^x)) dW_s + \int_0^t (\sigma_0^1(s, \mu_s^x)X_s + \sigma_0^2(s, \mu_s^x)) dB_s, \end{aligned} \quad (5.17)$$

where we have let $\alpha_s := \int_A a\Lambda_s(da)$. Assume all of the filtrations are completed throughout this proof. By optional projection (see [79, Appendix A.3] for a treatment without right-continuity of the filtration), there exist $\mathbb{F}^{\xi, B, W, \mu}$ -optional (and thus progressive) processes \tilde{X} and $\tilde{\alpha}$ such that such that, for each $t \in [0, T]$,

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

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

$$\tilde{X}_s := \mathbb{E}[X_s | \mathcal{F}_t^{\xi, B, W, \mu}], \quad \tilde{\alpha}_s := \mathbb{E}[\alpha_s | \mathcal{F}_t^{\xi, B, W, \mu}], \quad a.s. \quad (5.18)$$

Indeed, since (α_s, X_s) is $\mathcal{F}_s^{\xi, B, W, \mu, \Lambda, X}$ -measurable, and since the solution of the state equation 3.4 is strong, we know that (α_s, X_s) is a.s. $\mathcal{F}_s^{\xi, B, W, \mu, \Lambda}$ -measurable. By compatibility, $\mathcal{F}_t^{\xi, B, W, \mu}$ and \mathcal{F}_s^Λ are conditionally independent given $\mathcal{F}_s^{\xi, B, W, \mu}$ (this property not altered by completing the filtrations). This implies (5.18).

Now, for a given $t \in [0, T]$, take the conditional expectation with respect to $\mathcal{F}_t^{\xi, B, W, \mu}$ in (5.17). Using a conditional version of Fubini's theorem together with (5.18), we get that for each $t \in [0, T]$ it holds P -a.s. that

$$\begin{aligned} \tilde{X}_t = \xi &+ \int_0^t (b^1(s, \mu_s^x)\tilde{X}_s + b^2(s, \mu_s^x)\tilde{\alpha}_s + b^3(t, \mu_s^x)) ds \\ &+ \int_0^t (\sigma^1(s, \mu_s^x)\tilde{X}_s + \sigma^2(s, \mu_s^x)) dW_s + \int_0^t (\sigma_0^1(s, \mu_s^x)\tilde{X}_s + \sigma_0^2(s, \mu_s^x)) dB_s. \end{aligned} \quad (5.19)$$

Since the right-hand side is continuous a.s. and the filtration is complete, we replace \tilde{X} with an a.s.-continuous modification, so that (5.19) holds for all $t \in [0, T]$, P -a.s. That is, the processes on either side of the equation are indistinguishable.

Now define $\tilde{P} := P \circ (\xi, B, W, \mu, dt\delta_{\tilde{\alpha}_t}(da), \tilde{X})^{-1}$. It is clear from (5.19) that $\tilde{P} \in \mathcal{RA}(\rho)$. Jensen's inequality provides

$$J(P) \leq \mathbb{E}^P \left[\int_0^T f(t, X_t, \mu_t^x, \alpha_t) dt + g(X_T, \mu_T^x) \right] \quad (5.20)$$

$$\begin{aligned} &= \mathbb{E}^P \left[\int_0^T \mathbb{E}^P [f(t, X_t, \mu_t^x, \alpha_t) | \mathcal{F}_t^{\xi, B, W, \mu}] dt + \mathbb{E}^P [g(X_T, \mu_T^x) | \mathcal{F}_T^{\xi, B, W, \mu}] \right] \\ &\leq \mathbb{E}^P \left[\int_0^T f(t, \tilde{X}_t, \mu_t^x, \tilde{\alpha}_t) dt + g(\tilde{X}_T, \mu_T^x) \right] = J(\tilde{P}). \end{aligned} \quad (5.21)$$

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

Proof of (2). Unless (under P) Λ is already a strict control, then inequality (5.20) is strict, and unless $\alpha_t = \int_A a \Lambda_t(da)$ is already $\mathbb{F}^{\xi, B, W, \mu}$ -adapted, the inequality (5.21) is strict: $J(\tilde{P}) > J(P)$. This proves that all optimal controls must be strict and $\mathbb{F}^{\xi, B, W, \mu}$ -adapted. Now suppose we have two strict adapted optimal controls, which without loss of generality we construct on the same space $(\Omega_0 \times \mathcal{P}(\mathcal{X}), \mathbb{F}^{\xi, B, W, \mu}, \rho)$. That is,

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

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

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

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

Remark 5.4.2. It is interesting to note, as is clear from the proofs, that the full force of assumption **A1** is not needed for either of Propositions 3.1.4 or 3.1.5. Notably, the assumption $p' > p$ is not needed. In Proposition 3.1.4 the coefficients (b, σ, f) should be continuous in a to ensure that the set $K(t, x, \mu)$ is closed, but continuity in (x, μ) is unnecessary.

Chapter 6

Proofs of the limit theorems

This chapter is devoted to the proofs of the main convergence results, Theorem 3.2.4, its converse Theorem 3.2.10, and finally the analogous results of Theorem 4.2.2 for the setting without common noise. At the end of the section, we will finally prove Propositions 3.2.2 and 4.2.1, neither of which is used in the proofs of the other three theorems. The notation developed in Section 3.2 for the n -player games now returns to the spotlight.

6.1 Proof of main limit Theorem 3.2.4

Let us begin with the proof of Theorem 3.2.4. Throughout the section, we work with the assumptions and notation of Theorem 3.2.4. In particular, assumptions **A1**, **A2**, and **A3** are in force throughout this section. The strategy is to prove the claimed relative compactness, then that any limit is a MFG pre-solution using the characterization of Lemma 5.2.3, and then finally that any limit corresponds to an optimal control. First, we establish some standard but useful estimates for the n -player systems, very much analogous to Lemma 5.3.1.

6.1.1 Estimates

The first estimate below, Lemma 6.1.1, is fairly standard (and the proof is similar to that of Lemma 5.3.1), but it is important that it is independent of the number of agents n . The second estimate, Lemma 6.1.2, will be used to establish some uniform integrability of the equilibrium controls, and it is precisely where we need the coercivity of the running cost f . Note in the following proofs that the initial states $X_0^i[\Lambda] = X_0^i = \xi^i$ and the initial empirical measure $\hat{\mu}_0^x[\Lambda] = \hat{\mu}_0^x = \frac{1}{n} \sum_{i=1}^n \delta_{\xi^i}$ do not depend on the choice of control. Recall the definition of the truncated supremum norm (5.10).

Lemma 6.1.1. *There exists a constant $c_5 \geq 1$, depending only on p, p', T , and the constant c_1 of assumption **A1.4** such that, for each $\gamma \in [p, p']$, $\beta = (\beta^1, \dots, \beta^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$, and*

$$1 \leq k \leq n,$$

$$\begin{aligned} \mathbb{E}^{\mathbb{P}_n} [\|X^k[\beta]\|_T^\gamma] &\leq c_5 \mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^\gamma + \int_0^T \int_A |a|^\gamma \beta_t^k(da) dt \right. \\ &\quad \left. + \frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^\gamma \beta_t^i(da) dt \right], \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}^{\mathbb{P}_n} \int_{\mathcal{C}^d} \|z\|_T^\gamma \hat{\mu}^x[\beta](dz) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}_n} [\|X^i[\beta]\|_T^\gamma] \\ &\leq c_5 \mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^\gamma + \frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^\gamma \beta_t^i(da) dt \right]. \end{aligned}$$

Proof. We omit $[\beta]$ from the notation throughout the proof, as well as the superscript \mathbb{P}_n which should appear above the expectations. Abbreviate $\Sigma := \sigma\sigma^\top + \sigma_0\sigma_0^\top$. Apply the Burkholder-Davis-Gundy inequality and the growth assumption **A1.4** to find a universal constant $C > 0$ (which will change from line to line) such that, for all $\gamma \in [p, p']$,

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

The last line follows from the bound $(\int \|z\|_s^{p_\sigma} \nu(dz))^{\gamma/p} \leq \int \|z\|_s^\gamma \nu(dz)$ for $\nu \in \mathcal{P}(\mathcal{C}^d)$, which holds because $\gamma \geq p$. To deal with the $\gamma/2$ outside of the time integral, we used the following argument. If $\gamma \geq 2$, we simply use Jensen's inequality to pass $\gamma/2$ inside of the time integral, and then use the inequality $|x|^{p_\sigma \gamma/2} \leq 1 + |x|^\gamma$, which holds because $p_\sigma \leq 2$. The other case is $1 \vee p_\sigma \leq p \leq \gamma < 2$, and we use then the inequalities $|x|^{\gamma/2} \leq 1 + |x|$ and $|x|^{p_\sigma} \leq 1 + |x|^\gamma$. By Gronwall's inequality,

$$\mathbb{E}[\|X^k\|_t^\gamma] \leq C \mathbb{E} \left\{ 1 + |\xi^k|^\gamma + \int_0^t \left[\int_{\mathcal{C}^d} \|z\|_s^\gamma \hat{\mu}^x(dz) + \int_A |a|^\gamma \beta_s^k(da) \right] ds \right\} \quad (6.1)$$

Note that $\mathbb{E}^{\mathbb{P}_n}[\|\xi^k\|^\gamma] = \mathbb{E}^{\mathbb{P}_n}[\|\xi^1\|^\gamma]$ for each k , and average over $k = 1, \dots, n$ to get

$$\begin{aligned} \mathbb{E} \int_{\mathcal{C}^d} \|z\|_t^\gamma \hat{\mu}^x(dz) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[\|X^i\|_t^\gamma] \\ &\leq C \mathbb{E} \left\{ 1 + |\xi^1|^\gamma + \int_0^t \left[\int_{\mathcal{C}^d} \|z\|_s^\gamma \hat{\mu}^x(dz) + \frac{1}{n} \sum_{i=1}^n \int_A |a|^\gamma \beta_s^i(da) \right] ds \right\}. \end{aligned}$$

Apply Gronwall's inequality once again to prove the second claimed inequality. The first claim follows from the second and from (6.1). \square

Lemma 6.1.2. *There exist constants $c_6, c_7 > 0$, depending only p, p', T , and the constants c_1, c_2, c_3 of assumption **A1**, such that for each $\beta = (\beta^1, \dots, \beta^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$, the following hold:*

1. *For each $1 \leq k \leq n$,*

$$\begin{aligned} \mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \beta_t^k(da) dt \\ \leq c_7 \mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^p + \frac{1}{n} \sum_{i \neq k}^n \int_0^T \int_A |a|^p \beta_t^i(da) \right] - c_7 J_k(\beta). \end{aligned}$$

2. *If for some $n \geq k \geq 1$, $\epsilon > 0$, and $\tilde{\beta}^k \in \mathcal{A}_n(\mathcal{E}_n)$ we have*

$$J_k(\tilde{\beta}^k) \geq \sup_{\tilde{\beta} \in \mathcal{A}_n(\mathcal{E}_n)} J_k((\beta^{-k}, \tilde{\beta})) - \epsilon,$$

then

$$\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \tilde{\beta}_t^k(da) dt \leq c_7 \mathbb{E}^{\mathbb{P}_n} \left[1 + \epsilon + |\xi^1|^p + \frac{1}{n} \sum_{i \neq k}^n \int_0^T \int_A |a|^p \beta_t^i(da) \right].$$

3. *If β is an ϵ -Nash equilibrium for some $\epsilon = (\epsilon_1, \dots, \epsilon_n) \in [0, \infty)^n$, then*

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \beta_t^i(da) dt \leq c_7 \left(1 + \mathbb{E}^{\mathbb{P}_n} |\xi^1|^p + \frac{1}{n} \sum_{i=1}^n \epsilon_i \right).$$

Proof. Recall that $\mathbb{E}^{\mathbb{P}_n}[\|\xi^1\|^p] < \infty$ and that every $\tilde{\beta} \in \mathcal{A}_n(\mathcal{E}_n)$ is required to satisfy

$$\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^p \tilde{\beta}_t(da) dt < \infty.$$

Moreover, if $\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^{p'} \tilde{\beta}_t(da) dt = \infty$ then the upper bound on f of assumption **A1.5** implies that $J_k((\beta^{-k}, \tilde{\beta})) = -\infty$, for each $\beta \in \mathcal{A}_n^n(\mathcal{E}_n)$ and $1 \leq k \leq n$.

Proof of (1): First, use the upper bounds of f and g from assumption **A1.5** to get

$$\begin{aligned}
J_k(\beta) &\leq c_2(T+1)\mathbb{E}^{\mathbb{P}_n} \left[1 + \|X^k[\beta]\|_T^p + \int_{\mathcal{C}^d} \|z\|_T^p \hat{\mu}^x[\beta](dz) \right] \\
&\quad - c_3\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^{p'} \beta_t^k(da) dt \\
&\leq 3c_5c_2(T+1)\mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^p + \int_0^T \int_A |a|^p \beta_t^k(da) dt + \frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^p \beta_t^i(da) dt \right] \\
&\quad - c_3\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^{p'} \beta_t^k(da) dt,
\end{aligned}$$

where the last inequality follows from Lemma 6.1.1 (and $c_5 \geq 1$). This proves the first claim, with $c_6 := 6c_5c_2(T+1)/c_3$ and $c_7 := c_6 \vee (1/c_3)$.

Proof of (2): Fix $a_0 \in A$ arbitrarily. Abuse notation somewhat by writing a_0 in place of the constant strict control $(\delta_{a_0})_{t \in [0, T]} \in \mathcal{A}_n(\mathcal{E}_n)$. Lemma 6.1.1 implies

$$\begin{aligned}
&\mathbb{E}^{\mathbb{P}_n} [\|X^k[(\beta^{-k}, a_0)]\|_T^p] \\
&\leq c_5\mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^p + T \left(1 + \frac{1}{n} \right) |a_0|^p + \frac{1}{n} \sum_{i \neq k}^n \int_0^T \int_A |a|^p \beta_t^i(da) dt \right]
\end{aligned}$$

and

$$\begin{aligned}
&\mathbb{E}^{\mathbb{P}_n} \int_{\mathcal{C}^d} \|z\|_T^p \hat{\mu}^x[(\beta^{-k}, a_0)](dz) \\
&\leq c_5\mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^p + \frac{T}{n} |a_0|^p + \frac{1}{n} \sum_{i \neq k}^n \int_0^T \int_A |a|^p \beta_t^i(da) dt \right].
\end{aligned}$$

Use the hypothesis along with the lower bounds on f and g from assumption **A1.5** to get

$$\begin{aligned}
&J_k((\beta^{-k}, \tilde{\beta}^k)) + \epsilon \\
&\geq J_k((\beta^{-k}, a_0)) \\
&\geq -c_2(T+1)\mathbb{E}^{\mathbb{P}_n} \left[1 + \|X^k[(\beta^{-k}, a_0)]\|_T^p + \int_{\mathcal{C}^d} \|z\|_T^p \hat{\mu}^x[(\beta^{-k}, a_0)](dz) + |a_0|^{p'} \right] \\
&\geq -C\mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^p + \frac{1}{n} \sum_{i \neq k}^n \int_0^T \int_A |a|^p \beta_t^i(da) dt \right],
\end{aligned}$$

where $C > 0$ depends only on c_2 , c_5 , T , and $|a_0|^{p'}$. Applying this with the first result with β replaced by $(\beta^{-k}, \tilde{\beta}^k)$ proves (2), replacing c_7 by $c_7(1+C)$.

Proof of (3): If β is an ϵ -Nash equilibrium, then applying (2) with $\tilde{\beta}^k = \beta^k$ gives

$$\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \beta_t^k(da) dt \leq c_7 \mathbb{E}^{\mathbb{P}_n} \left[1 + \epsilon_k + |\xi^1|^p + \frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^p \beta_t^i(da) \right].$$

The proof is completed by averaging over $k = 1, \dots, n$, rearranging terms, and replacing c_6 by $c_6 + c_7$. \square

6.1.2 Relative compactness and pre-solutions

This section proves that $(P_n)_{n=1}^\infty$, defined in (3.7), is relatively compact and that each limit point is a MFG pre-solution. For this we will need a tailor-made tightness result for Itô processes which generalizes our previous result of Proposition 5.3.2. It is essentially an application of Aldous' criterion, but the proof is deferred to Appendix B.

Proposition 6.1.3. *Fix $c > 0$. For each $\kappa \geq 0$, let $\mathcal{Q}_\kappa \subset \mathcal{P}(\mathcal{V} \times \mathcal{C}^d)$ denote the set of laws $P \circ (\Lambda, X)^{-1}$ of $\mathcal{V} \times \mathcal{C}^d$ -valued random variables (Λ, X) defined on some filtered probability space $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P)$ satisfying*

$$dX_t = \int_A B(t, a) \Lambda_t(da) dt + \Sigma(t) dW_t,$$

where the following hold:

1. W is a \mathbb{F} -Wiener process of dimension k , and X is a continuous \mathbb{F} -adapted d -dimensional process with $P \circ X_0^{-1} = \lambda$.
2. $\Sigma : [0, T] \times \tilde{\Omega} \rightarrow \mathbb{R}^{d \times k}$ is progressively measurable, and $B : [0, T] \times \tilde{\Omega} \times A \rightarrow \mathbb{R}^d$ is jointly measurable with respect to the \mathbb{F} -progressive σ -field on $[0, T] \times \tilde{\Omega}$ and the Borel σ -field on A .
3. There exists a nonnegative \mathcal{F}_T -measurable random variable Z such that, for each $(t, \omega, a) \in [0, T] \times \tilde{\Omega} \times A$,

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

and

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

(That is, we vary over Σ, B, Z, k , and the probability space.) Then, for any triangular array $\{\kappa_{n,i} : 1 \leq i \leq n\} \subset [0, \infty)$ with $\sup_n \frac{1}{n} \sum_{i=1}^n \kappa_{n,i} < \infty$, the set

$$\mathcal{Q} := \left\{ \frac{1}{n} \sum_{i=1}^n P_i : n \geq 1, P_i \in \mathcal{Q}_{\kappa_{n,i}} \text{ for } i = 1, \dots, n \right\}$$

is relatively compact in $\mathcal{P}(\mathcal{V} \times \mathcal{C}^d)$.

Lemma 6.1.4. $(P_n)_{n=1}^\infty$ is relatively compact in $\mathcal{P}^p(\Omega)$, and

$$\sup_n \mathbb{E}^{P_n} \left[\|X\|_T^{p'} + \int_{\mathcal{C}^d} \|z\|_T^{p'} \mu(dz) + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \right] < \infty. \quad (6.2)$$

Proof. We first establish (6.2). Since Λ^n is a ϵ^n -Nash equilibrium, part (3) of Lemma 6.1.2 implies

$$\frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \Lambda_t^{n,k}(da) dt \leq c_7 \left(1 + \mathbb{E}^{\mathbb{P}_n} [|\xi^1|^p] + \frac{1}{n} \sum_{k=1}^n \epsilon_k^n \right).$$

The right-hand side above is bounded in n , because of hypothesis (3.6) and because $\mathbb{P}_n \circ (\xi^1)^{-1} = \lambda \in \mathcal{P}^p(\mathbb{R}^d)$ for each n . Since $p' > p$, it follows that

$$\sup_n \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^{p'} \Lambda_t^{n,k}(da) dt < \infty. \quad (6.3)$$

Lemma 6.1.1 implies

$$\mathbb{E}^{\mathbb{P}_n} \int_{\mathcal{C}^d} \|z\|_T^{p'} \hat{\mu}^x[\Lambda^n](dz) \leq c_5 \mathbb{E}^{\mathbb{P}_n} \left[1 + |\xi^1|^{p'} + \frac{1}{n} \sum_{k=1}^n \int_0^T \int_A |a|^{p'} \Lambda_t^{n,k}(da) dt \right] =: \kappa_n.$$

Thus

$$\begin{aligned} & \mathbb{E}^{P_n} \left[\|X\|_T^{p'} + \int_{\mathcal{C}^d} \|z\|_T^{p'} \mu(dz) + \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \right] \\ &= \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} \left[\|X^k[\Lambda^n]\|_T^{p'} + \int_{\mathcal{C}^d} \|z\|_T^{p'} \hat{\mu}^x[\Lambda^n](dz) + \int_0^T \int_A |a|^{p'} \Lambda_t^{n,k}(da) dt \right] \\ &\leq c_5 \mathbb{E}^{\mathbb{P}_n} \left[2 + 2|\xi^1|^{p'} + \frac{3}{n} \sum_{k=1}^n \int_0^T \int_A |a|^{p'} \Lambda_t^{n,k}(da) dt \right] \\ &\leq 3\kappa_n. \end{aligned}$$

Recall in the last line that $c_5 \geq 1$. From (6.3) we conclude that $\sup_n \kappa_n < \infty$, and (6.2) follows.

To prove that $(P_n)_{n=1}^\infty$, it suffices to show that each family of marginals is relatively compact (see Lemma 2.1.8). Since $(P_n \circ (\xi, B, W)^{-1})_{n=1}^\infty$ is a singleton, it is trivially compact. We may apply Proposition 6.1.3 to show that

$$P_n \circ (\Lambda, X)^{-1} = \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (\Lambda^{n,i}, X^{n,i}[\Lambda^n])^{-1}$$

forms a relatively compact sequence. Indeed, in the notation of Proposition 6.1.3, we use $Z = (\int_{\mathcal{C}^d} \|z\|_T^p \hat{\mu}^x[\Lambda^n](dz))^{1/p}$ and $c = c_1$ of assumption A1.4 to check that $\mathbb{P}_n \circ (\Lambda^{n,i}, X^{n,i}[\Lambda^n])^{-1}$

is in $\mathcal{Q}_{\kappa_{n,i}}$ for each $1 \leq i \leq n$, where

$$\kappa_{n,i} = \kappa_n + \mathbb{E}^{\mathbb{P}_n} \left[|\xi^i|^{p'} + \int_0^T \int_A |a|^{p'} \Lambda_t^{n,i}(da) dt \right].$$

Since $c_5 \geq 1$, we have $\frac{1}{n} \sum_{i=1}^n \kappa_{n,i} \leq 2\kappa_n$, and so $\sup_n \frac{1}{n} \sum_{i=1}^n \kappa_{n,i} < \infty$. Thus, Proposition 6.1.3 establishes the relative compactness of $(P_n \circ (\Lambda, X)^{-1})_{n=1}^\infty$. Next, note that $P_n \circ (W, \Lambda, X)^{-1}$ is the mean measure of $P_n \circ \mu^{-1}$ for each n , since for each bounded measurable $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ we have

$$\mathbb{E}^{P_n} [\varphi(W, \Lambda, X)] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}_n} [\varphi(W^i, \Lambda^{n,i}, X^i[\Lambda^n])] = \mathbb{E}^{\mathbb{P}_n} \int_{\mathcal{X}} \varphi d\hat{\mu}[\Lambda^n] = \mathbb{E}^{P_n} \int_{\mathcal{X}} \varphi d\mu.$$

(See Section 2.1.2 for a discussion of mean measures and their use for deriving compactness.) Since also

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

the relative compactness of $(P_n \circ \mu^{-1})_{n=1}^\infty$ in $\mathcal{P}^p(\mathcal{P}^p(\mathcal{X}))$ follows from the relative compactness of $(P_n \circ (W, \Lambda, X)^{-1})_{n=1}^\infty$ in $\mathcal{P}^p(\mathcal{X})$; see Corollary 2.1.13. This completes the proof. \square

Finally, the following Lemma shows that the limits of P_n are MFG pre-solution. It is interesting to note that this proof is new even when specialized to the McKean-Vlasov case, which is of course the special case when the control problem degenerates in the sense that A is a singleton. Refer back to Section 2.2 and Theorem 2.2.3 for more discussion of McKean-Vlasov limits. In particular, the complete proof of Theorem 2.2.3 is nearly identical to that of the following lemma:

Lemma 6.1.5. *Any limit point P of $(P_n)_{n=1}^\infty$ in $\mathcal{P}^p(\Omega)$ is a MFG pre-solution.*

Proof. We abuse notation somewhat by assume that $P_n \rightarrow P$, with the understanding that this is along a subsequence. We check that P satisfies the four conditions of Lemma 5.2.3.

1. Of course,

$$P_n \circ (\xi, B, W)^{-1} = \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (\xi^i, B, W^i)^{-1} = \lambda \times \mathcal{W}^{m_0} \times \mathcal{W}^m,$$

where \mathcal{W}^k denotes Wiener measure on \mathcal{C}^k . Thus $P \circ (\xi, B, W)^{-1} = \lambda \times \mathcal{W}^{m_0} \times \mathcal{W}^m$ as well. On Ω_n , we know $\sigma(W_s^i - W_t^i, B_s - B_t : i = 1, \dots, n, s \in [t, T])$ is \mathbb{P}_n -independent of \mathcal{F}_t^n for each $t \in [0, T]$. It follows that, on Ω , $\sigma(W_s - W_t, B_s - B_t : s \in [t, T])$ is P_n -independent of $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X}$. Hence B and W are Wiener processes on $(\Omega, \mathbb{F}^{\xi, B, W, \mu, \Lambda, X}, P)$.

2. Fix bounded continuous functions $\varphi : \mathbb{R}^d \times \mathcal{C}^m \rightarrow \mathbb{R}$ and $\psi : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$. Since $(\xi^1, W^1), \dots, (\xi^n, W^n)$ are i.i.d. under \mathbb{P}_n with common law $P \circ (\xi, W)^{-1}$ for each n ,

the law of large numbers implies

$$\lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}_n} \left[\left| \frac{1}{n} \sum_{i=1}^n \varphi(\xi^i, W^i) - \mathbb{E}^P[\varphi(\xi, W)] \right| \psi(B, \hat{\mu}[\Lambda^n]) \right] = 0.$$

This implies

$$\begin{aligned} \mathbb{E}^P[\varphi(\xi, W)\psi(B, \mu)] &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}_n}[\varphi(\xi^i, W^i)\psi(B, \hat{\mu}[\Lambda^n])] \\ &= \mathbb{E}^P[\varphi(\xi, W)] \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\mathbb{P}_n}[\psi(B, \hat{\mu}[\Lambda^n])] \\ &= \mathbb{E}^P[\varphi(\xi, W)] \mathbb{E}^P[\psi(B, \mu)]. \end{aligned}$$

This shows (B, μ) is independent of (ξ, W) under P . Since ξ^i and W^i are independent under \mathbb{P}_n , it follows that ξ and W are independent under P_n , for each n . Thus ξ and W are independent under P , and we conclude that ξ , W , and (B, μ) are independent under P .

3. Let $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ and $\psi : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ be bounded and continuous. Then

$$\begin{aligned} \mathbb{E}^P[\psi(B, \mu)\varphi(W, \Lambda, X)] &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}_n} \left[\psi(B, \hat{\mu}[\Lambda^n]) \frac{1}{n} \sum_{i=1}^n \varphi(W^i, \Lambda^{n,i}, X^i[\Lambda^n]) \right] \\ &= \lim_{n \rightarrow \infty} \mathbb{E}^{\mathbb{P}_n} \left[\psi(B, \hat{\mu}[\Lambda^n]) \int_{\mathcal{X}} \varphi d\hat{\mu}[\Lambda^n] \right] \\ &= \mathbb{E}^P \left[\psi(B, \mu) \int_{\mathcal{X}} \varphi d\mu \right]. \end{aligned}$$

4. Since $(\xi^i, B, W^i, \hat{\mu}[\Lambda^n], \Lambda^{n,i}, X^i[\Lambda^n])$ verify the state SDE under \mathbb{P}_n , the canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the state equation (5.7) under each P_n , for each n . This property is preserved in the limit $P_n \rightarrow P$, thanks to well known results on weak convergence of stochastic integrals (e.g., Kurtz and Protter [81]).

□

6.1.3 Modified finite-player games

The last step of the proof, executed in the next Section 6.1.4, is to show that any limit P of P_n is optimal. This step is more involved, and we devote this subsection to studying a useful technical device which we call the *k-modified n-player game*, in which agent k is removed from the empirical measures. Intuitively, if the n -player game is modified so that the empirical measure (present in the state process dynamics and objective functions) no longer includes agent k , then the optimization problem of agent k de-couples from that of the other agents; agent k may then treat the empirical measure of the other $n - 1$ agents

as fixed and thus faces exactly the type of control problem encountered in the MFG. Let us make this idea precise.

For $\beta = (\beta^1, \dots, \beta^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$, define $Y^{-k}[\beta] = (Y^{-k,1}[\beta], \dots, Y^{-k,n}[\beta])$ to be the unique strong solution on $(\Omega_n, \mathbb{F}^n, \mathbb{P}_n)$ of the SDE

$$\begin{aligned} Y_t^{-k,i}[\beta] &= \xi^i + \int_0^t \int_A b(s, Y_s^{-k,i}[\beta], \hat{\mu}_s^{-k,x}[\beta], a) \beta_s^i(da) dt \\ &\quad + \int_0^t \sigma(s, Y_s^{-k,i}[\beta], \hat{\mu}_s^{-k,x}[\beta]) dW_s^i \\ &\quad + \int_0^t \sigma_0(s, Y_s^{-k,i}[\beta], \hat{\mu}_s^{-k,x}[\beta]) dB_s, \\ \hat{\mu}^{-k,x}[\beta] &:= \frac{1}{n-1} \sum_{i \neq k}^n \delta_{Y^{-k,i}[\beta]}. \end{aligned}$$

Define also

$$\hat{\mu}^{-k}[\beta] = \frac{1}{n-1} \sum_{i \neq k}^n \delta_{(W^i, \beta^i, Y^{-k,i}[\beta])}.$$

Intuitively, $Y^{-k,i}$ is agent i 's state process in an analog of the n -player game, in which agent k has been removed from the empirical measure. Naturally, for fixed k , the k -modified state processes $Y^{-k}[\beta]$ should not be far from the true state processes $X[\beta]$ if n is large, and we will quantify this precisely. We will need to be somewhat explicit about the choice of metric on \mathcal{V} , so we define $d_{\mathcal{V}}$ by

$$\begin{aligned} d_{\mathcal{V}}^p(q, q') &:= T \ell_{[0,T] \times A, p}(q/T, q'/T) \\ &= \inf_{\gamma} \int_{[0,T]^2 \times A^2} (|t - t'|^p + |a - a'|^p) \gamma(dt, dt', da, da'), \end{aligned}$$

where the infimum is over measures γ on $[0, T]^2 \times A^2$ with marginals q and q' . By choosing $\gamma = dt \delta_t(dt') q_t(da) q'_t(da')$, we note that

$$d_{\mathcal{V}}^p(q, q') \leq 2^{p-1} \int_0^T \int_A |a|^p q_t(da) dt + 2^{p-1} \int_0^T \int_A |a|^p q'_t(da) dt. \quad (6.4)$$

Define the p' -Wasserstein distance $\ell_{\mathcal{X}, p'}$ on $\mathcal{P}^{p'}(\mathcal{X})$ with respect to the metric on \mathcal{X} given by

$$d_{\mathcal{X}}((w, q, x), (w', q', x')) := \|w - w'\|_T + d_{\mathcal{V}}(q, q') + \|x - x'\|_T. \quad (6.5)$$

Lemma 6.1.6. *There exists a constant $c_8 > 0$ such that, for each $n \geq k \geq 1$ and $\beta = (\beta^1, \dots, \beta^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$, we have*

$$\mathbb{E}^{\mathbb{P}_n} \left[\ell_{\mathcal{X}, p'}^{p'}(\hat{\mu}^{-k}[\beta], \hat{\mu}[\beta]) + \|X^k[\beta] - Y^{-k, k}[\beta]\|_T^{p'} \right] \leq c_8(1 + M[\beta])/n, \text{ where}$$

$$M[\beta] := \mathbb{E}^{\mathbb{P}_n} \left[|\xi^1|^{p'} + \frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^{p'} \beta_t^i(da) dt \right].$$

Proof. Throughout the proof, n is fixed, expected values are all with respect to \mathbb{P}_n , and the notation $[\beta]$ is omitted. Define the truncated p' -Wasserstein distance ℓ_t on $\mathcal{P}^{p'}(\mathcal{C}^d)$ by

$$\ell_t(\mu, \nu) := \inf \left\{ \int_{\mathcal{C}^d \times \mathcal{C}^d} \|x - y\|_t^{p'} \gamma(dx, dy) : \gamma \in \mathcal{P}(\mathcal{C}^d \times \mathcal{C}^d) \text{ has marginals } \mu, \nu \right\}^{1/p'} \quad (6.6)$$

Apply the Doob's maximal inequality and Jensen's inequality (using $p' \geq 2$, which was part of assumption **A3**) to find a constant $C > 0$ (which will change from line to line but depends only on d, p, p', T, c_1 , and c_5) such that

$$\begin{aligned} \mathbb{E} \left[\|X^i - Y^{-k, i}\|_t^{p'} \right] &\leq C \mathbb{E} \int_0^t \int_A |b(s, X_s^i, \hat{\mu}_s^x, a) - b(s, Y_s^{-k, i}, \hat{\mu}_s^{-k, x}, a)|^{p'} \beta_s^i(da) ds \\ &\quad + C \mathbb{E} \int_0^t |\sigma(s, X_s^i, \hat{\mu}_s^x) - \sigma(s, Y_s^{-k, i}, \hat{\mu}_s^{-k, x})|^{p'} ds \\ &\quad + C \mathbb{E} \int_0^t |\sigma_0(s, X_s^i, \hat{\mu}_s^x) - \sigma_0(s, Y_s^{-k, i}, \hat{\mu}_s^{-k, x})|^{p'} ds \\ &\leq C \mathbb{E} \int_0^t \left(\|X^i - Y^{-k, i}\|_s^{p'} + \ell_s^{p'}(\hat{\mu}^x, \hat{\mu}^{-k, x}) \right) ds. \end{aligned}$$

The last line followed from the Lipschitz assumption **A3**, along with the observation that

$$\ell_{\mathbb{R}^d, p}(\nu_s^1, \nu_s^2) \leq \ell_{\mathbb{R}^d, p'}(\nu_s^1, \nu_s^2) \leq \ell_s(\nu^1, \nu^2),$$

for each $\nu^1, \nu^2 \in \mathcal{P}^p(\mathcal{C}^d)$ and $s \in [0, T]$. By Gronwall's inequality (updating the constant C),

$$\mathbb{E} \left[\|X^i - Y^{-k, i}\|_t^{p'} \right] \leq C \mathbb{E} \int_0^t \ell_s^{p'}(\hat{\mu}^x, \hat{\mu}^{-k, x}) ds. \quad (6.7)$$

Now we define a standard coupling of the empirical measures $\hat{\mu}^x$ and $\hat{\mu}^{-k, x}$: first, draw a number j from $\{1, \dots, n\}$ uniformly at random, and consider X^j to be a sample from $\hat{\mu}^x$. If $j \neq k$, choose $Y^{-k, j}$ to be a sample from $\hat{\mu}^{-k, x}$, but if $j = k$, draw another number j' from $\{1, \dots, n\} \setminus \{k\}$ uniformly at random, and choose $Y^{-k, j'}$ to be a sample from $\hat{\mu}^{-k, x}$. This yields

$$\ell_t^{p'}(\hat{\mu}^x, \hat{\mu}^{-k, x}) \leq \frac{1}{n} \sum_{i \neq k} \|X^i - Y^{-k, i}\|_t^{p'} + \frac{1}{n(n-1)} \sum_{i \neq k} \|X^k - Y^{-k, i}\|_t^{p'} \quad (6.8)$$

We know from Lemma 6.1.1 that

$$\frac{1}{n-1} \sum_{i \neq k}^n \mathbb{E}[\|X^i\|_T^{p'}] \leq c_5(1+M),$$

It should be clear that an analog of Lemma 6.1.1 holds for $Y^{-k,i}$ as well, with the same constant. In particular,

$$\frac{1}{n-1} \sum_{i \neq k}^n \mathbb{E}[\|Y^{-k,i}\|_T^{p'}] \leq c_5(1+M).$$

Combine the above four inequalities, averaging (6.7) over $i \neq k$, to get

$$\mathbb{E} \left[\ell_t^{p'}(\hat{\mu}^x, \hat{\mu}^{-k,x}) \right] \leq C \mathbb{E} \int_0^t \ell_s^{p'}(\hat{\mu}^x, \hat{\mu}^{-k,x}) ds + 2^{p'} c_5(1+M)/n.$$

Gronwall's inequality yields a new constant such that

$$\mathbb{E} \left[\ell_T^{p'}(\hat{\mu}^x, \hat{\mu}^{-k,x}) \right] \leq C(1+M)/n.$$

Return to (6.7) to find

$$\mathbb{E} \left[\|X^i - Y^{-k,i}\|_T^{p'} \right] \leq C(1+M)/n, \text{ for } i = 1, \dots, n. \quad (6.9)$$

The same coupling argument leading to (6.8) also yields

$$\begin{aligned} \ell_{\mathcal{X}, p'}^{p'}(\hat{\mu}, \hat{\mu}^{-k}) &\leq \frac{1}{n} \sum_{i \neq k}^n \|X^i - Y^{-k,i}\|_T^{p'} \\ &\quad + \frac{1}{n(n-1)} \sum_{i \neq k}^n d_{\mathcal{X}}^{p'}((W^i, \beta^i, Y^{-k,i}), (W^k, \beta^k, X^k)) \end{aligned} \quad (6.10)$$

Using (6.4), we find yet another constant such that

$$\begin{aligned} &\mathbb{E} \left[d_{\mathcal{X}}^{p'}((W^i, \beta^i, Y^{-k,i}), (W^k, \beta^k, X^k)) \right] \\ &\leq 3^{p'-1} \mathbb{E} \left[\|W^i - W^k\|_T^{p'} + d_{\mathcal{V}}^{p'}(\beta^i, \beta^k) + \|Y^{-k,i} - X^k\|_T^{p'} \right] \\ &\leq C \mathbb{E} \left[\int_0^T \int_A |a|^{p'}(\beta_t^i + \beta_t^k)(da) dt + \|W^1\|_T^{p'} + \|Y^{-k,i}\|_T^{p'} + \|X^k\|_T^{p'} \right] \\ &\leq C \left(2nM + 2nc_5(1+M) + \mathbb{E}[\|W^1\|_T^{p'}] \right). \end{aligned}$$

Thus

$$\frac{1}{n-1} \sum_{i \neq k}^n \mathbb{E} \left[d_{\mathcal{X}}^{p'}((W^i, \beta^i, Y^{-k,i}), (W^k, \beta^k, X^k)) \right] \leq C(1+M).$$

Applying this bound and (6.9) to (6.10) completes the proof. \square

6.1.4 Optimality in the limit

Before we complete the proof, recall the definitions of \mathcal{R} , \mathcal{A} , and \mathcal{A}^* from Section 5.2. The final step is to show that $P \in \mathcal{RA}^*(P \circ (\xi, B, W, \mu)^{-1})$, for any limit P of $(P_n)_{n=1}^\infty$. The idea of the proof is to use the density of adapted controls (see Corollary 5.3.6) to construct nearly optimal controls for the MFG with nice continuity properties. From these controls we build admissible controls for the n -player game, and it must finally be argued that the inequality obtained from the ϵ^n -Nash assumption on Λ^n may be passed to the limit.

Proof of Theorem 3.2.4. Let P be a limit point of $(P_n)_{n=1}^\infty$, which we know exists by Lemma 6.1.4, and again abuse notation by assuming that $P_n \rightarrow P$. Let $\rho := P \circ (\xi, B, W, \mu)^{-1}$. We know from Lemma 6.1.5 that P is a MFG pre-solution, so we need only to check that P is optimal. By Proposition 5.3.7, it suffices to show that $J(P) \geq J(\mathcal{R}(\tilde{Q}))$ for all $\tilde{Q} \in \mathcal{A}_a(\rho)$ (see Definition 5.3.5). Fix arbitrarily some $\tilde{Q} \in \mathcal{A}_a(\rho)$, and recall that $\tilde{Q} \in \mathcal{A}_a(\rho)$ means that there exist bounded, continuous, adapted functions $\tilde{\varphi} : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{V}$ such that

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

For $1 \leq k \leq n$, let

$$\rho_{n,k} := \mathbb{P}_n \circ (\xi^k, B, W^k, \hat{\mu}^{-k}[\Lambda^n])^{-1},$$

and

$$\begin{aligned} Q_{n,k} &:= \rho_{n,k} \circ (\xi, B, W, \mu, \tilde{\varphi}(\xi, B, W, \mu))^{-1} \\ &= \mathbb{P}_n \circ (\xi^k, B, W^k, \hat{\mu}^{-k}[\Lambda^n], \tilde{\varphi}(\xi^k, B, W^k, \hat{\mu}^{-k}[\Lambda^n]))^{-1}. \end{aligned}$$

It follows from Lemma 6.1.6 that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \rho_{n,k} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{P}_n \circ (\xi^k, B, W^k, \hat{\mu}[\Lambda^n])^{-1} = \rho.$$

Since

$$\frac{1}{n} \sum_{k=1}^n \rho_{n,k} \circ (\xi, B, W)^{-1} = P \circ (\xi, B, W)^{-1}$$

does not depend on n , the continuity of $\tilde{\varphi}$ implies

$$\tilde{Q} = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n Q_{n,k}.$$

It is fairly straightforward to check that \mathcal{R} is a linear map, and it is even more straightforward to check that J is linear. Moreover, note that $\tilde{\varphi}$ is bounded (i.e., concentrated on a compact subset of $[0, T] \times A$, as in Definition 5.3.5), and thus Lemma 6.1.6 yields

$$\sup_{n,k} \mathbb{E}^{Q_{n,k}} \int_{\mathcal{C}^d} \|x\|_T^{p'} < \infty.$$

Hence, the continuity of \mathcal{R} and J of Lemmas 5.3.3 and 5.3.4 imply

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n J(\mathcal{R}(Q_{n,k})) = \lim_{n \rightarrow \infty} J \left(\mathcal{R} \left(\frac{1}{n} \sum_{k=1}^n Q_{n,k} \right) \right) = J(\mathcal{R}(\tilde{Q})). \quad (6.11)$$

Now, for $k \leq n$, define $\beta^{n,k} \in \mathcal{A}_n(\mathcal{E}_n)$ by

$$\beta^{n,k} := \tilde{\varphi}(\xi^k, B, W^k, \hat{\mu}^{-k}[\Lambda^n]).$$

For $\beta \in \mathcal{A}_n(\mathcal{E}_n)$, abbreviate $(\Lambda^{n,-k}, \beta) := ((\Lambda^n)^{-k}, \beta)$. Since agent k is removed from the empirical measure in the k -modified system, we have $\hat{\mu}^{-k}[\Lambda^n] = \hat{\mu}^{-k}[(\Lambda^{n,-k}, \beta)]$ for each $\beta \in \mathcal{A}_n(\mathcal{E}_n)$. The key point is that for each $k \leq n$,

$$\mathbb{P}_n \circ (\xi^k, B, W^k, \hat{\mu}^{-k}[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})])^{-1} = \mathcal{R}(Q_{n,k}). \quad (6.12)$$

To prove (6.12), let $P' \in \mathcal{P}(\Omega)$ denote the measure on the left-hand side. Since $\hat{\mu}^{-k}[\Lambda^n] = \hat{\mu}^{-k}[(\Lambda^{n,-k}, \beta^{n,k})]$, we have

$$P' \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q_{n,k}.$$

Since the processes

$$(\xi^k, B, W^k, \hat{\mu}^{-k}[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})])$$

verify the state SDE (5.7) on $(\Omega_n, \mathbb{F}^n, \mathbb{P}_n)$, the canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the state SDE (5.7) under P' . Hence, $P' = \mathcal{R}(Q_{n,k})$. With (6.12) in hand, by definition of J the equation (6.11) then translates to

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} [\Gamma(\hat{\mu}^{-k,x}[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})])] = J(\mathcal{R}(\tilde{Q})). \quad (6.13)$$

One more technical ingredient is needed before we can complete the proof. Namely, we would like to substitute $X^k[(\Lambda^{n,-k}, \beta^{n,k})]$ for $Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})]$ in (6.13), by proving

$$\begin{aligned} 0 &= \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} [\Gamma(\hat{\mu}^{-k,x}[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})]) \\ &\quad - \Gamma(\hat{\mu}^{-k,x}[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, X^k[(\Lambda^{n,-k}, \beta^{n,k})])] . \end{aligned} \quad (6.14)$$

Indeed, it follows from Lemma 6.1.1 (and an obvious analog for the modified state processes Y) that

$$\begin{aligned} Z_{n,k} &:= \mathbb{E}^{\mathbb{P}_n} \left[\|X^k[(\Lambda^{n,-k}, \beta^{n,k})]\|_T^{p'} + \|Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})]\|_T^{p'} \right. \\ &\quad \left. + \int_{\mathcal{C}^d} \|z\|_T^{p'} \hat{\mu}^x[(\Lambda^{n,-k}, \beta^{n,k})](dz) + \int_{\mathcal{C}^d} \|z\|_T^{p'} \hat{\mu}^{-k,x}[(\Lambda^{n,-k}, \beta^{n,k})](dz) \right] \\ &\leq 4c_4 \mathbb{E}^{\mathbb{P}_n} \left[|\xi^1|^{p'} + \frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^{p'} \Lambda_t^{n,i}(da) dt + \int_0^T \int_A |a|^{p'} \beta_t^{n,k}(da) dt \right]. \end{aligned}$$

Lemma 6.1.4 says that

$$\sup_n \mathbb{E}^{\mathbb{P}_n} \left[\frac{1}{n} \sum_{i=1}^n \int_0^T \int_A |a|^{p'} \Lambda_t^{n,i}(da) dt \right] < \infty.$$

Boundedness of $\tilde{\varphi}$ implies that there exists a compact set $K \subset A$ such that $\tilde{\beta}_t^{n,k}(K^c) = 0$ for a.e. $t \in [0, T]$ and all $n \geq k \geq 1$. Thus

$$\sup_n \frac{1}{n} \sum_{k=1}^n Z_{n,k} < \infty,$$

and we have the uniform integrability needed to deduce (6.14) from Lemma 6.1.6 and from the continuity and growth assumptions A1.5 on f and g .

A simple manipulation of the definitions yields $J(P_n) = \frac{1}{n} \sum_{k=1}^n J_k(\Lambda^n)$. Then, since $P_n \rightarrow P$, the upper semicontinuity of J (proven in Lemma 5.3.4) implies

$$J(P) \geq \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n J_k(\Lambda^n).$$

Finally, use the fact that Λ^n is a relaxed ϵ^n -Nash equilibrium to get

$$\begin{aligned} J(P) &\geq \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n [J_k((\Lambda^{n,-k}, \beta^{n,k})) - \epsilon_k^n] \\ &= \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} [\Gamma(\hat{\mu}^x[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, X^k[(\Lambda^{n,-k}, \beta^{n,k})])] \\ &= \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n \mathbb{E}^{\mathbb{P}_n} [\Gamma(\hat{\mu}^{-k,x}[(\Lambda^{n,-k}, \beta^{n,k})], \beta^{n,k}, Y^{-k,k}[(\Lambda^{n,-k}, \beta^{n,k})])] \\ &= J(\mathcal{R}(\tilde{Q})) \end{aligned}$$

The second line is simply writing out the definition of J_k and dropping the ϵ_k^n , which is permitted by hypothesis (3.6). The third line comes from (6.14), and the last is from (6.13). This completes the proof. \square

6.2 Proof of converse limit Theorem 3.2.10

This section is devoted to the proof of Theorem 3.2.10, which we split into two pieces.

Theorem 6.2.1. *Suppose assumptions A1, A2, A3, and A4 hold. Let $P \in \mathcal{P}(\Omega)$ be a weak MFG solution. Then there exist, for each n ,*

1. $\epsilon_n \geq 0$,
2. an n -player environment $\mathcal{E}_n = (\Omega_n, (\mathcal{F}_t^n)_{t \in [0, T]}, \mathbb{P}_n, \xi, B, W)$, and
3. a relaxed $(\epsilon_n, \dots, \epsilon_n)$ -Nash equilibrium $\Lambda^n = (\Lambda^{n,1}, \dots, \Lambda^{n,n})$ on \mathcal{E}_n ,

such that $\lim_{n \rightarrow \infty} \epsilon_n = 0$ and $P_n \rightarrow P$ in $\mathcal{P}^p(\Omega)$, where

$$P_n := \frac{1}{n} \sum_{i=1}^n \mathbb{P}_n \circ (\xi^i, B, W^i, \hat{\mu}[\Lambda^n], \Lambda^{n,i}, X^i[\Lambda^n])^{-1}.$$

Theorem 6.2.1 is nearly the same as Theorem 3.2.10, except that the equilibria Λ^n are now *relaxed* instead of *strong*, and the environments \mathcal{E}_n are now part of the conclusion of the theorem instead of the input. We will prove Theorem 6.2.1 by constructing a convenient sequence of environments \mathcal{E}_n , which all live on the same larger probability space supporting an i.i.d. sequence of state processes corresponding to the given MFG solution. This kind of argument is known as *trajectorial propagation of chaos* in the literature on McKean-Vlasov limits, and the Lipschitz assumption in the measure argument is useful here. The precise choice of environments also facilitates the proof of the following Proposition. Recall the definition of a *strong ϵ -Nash equilibrium* from Remark 3.2.1 and the discussion preceding it.

Proposition 6.2.2. *Let \mathcal{E}_n be the environments defined in the proof of Theorem 6.2.1 (in Section 6.2.1). Let $\Lambda^0 = (\Lambda^{0,1}, \dots, \Lambda^{0,n}) \in \mathcal{A}_n^n(\mathcal{E}_n)$. Then there exist strong strategies $\Lambda^k = (\Lambda^{k,1}, \dots, \Lambda^{k,n}) \in \mathcal{A}_n^n(\mathcal{E}_n)$ such that:*

1. In $\mathcal{P}^p(\mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n)$,

$$\lim_{k \rightarrow \infty} \mathbb{P}_n \circ (B, W, \Lambda^k, X[\Lambda^k])^{-1} = \mathbb{P}_n \circ (B, W, \Lambda^0, X[\Lambda^0])^{-1},$$

2. $\lim_{k \rightarrow \infty} J_i(\Lambda^k) = J_i(\Lambda^0)$, for $i = 1, \dots, n$,

3.

$$\limsup_{k \rightarrow \infty} \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{k,-i}, \beta)) \leq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{0,-i}, \beta)), \text{ for } i = 1, \dots, n.$$

In particular, if Λ^0 is a relaxed $\epsilon^0 = (\epsilon_1^0, \dots, \epsilon_n^0)$ -Nash equilibrium, then Λ^k is a strong $(\epsilon^0 + \epsilon^k)$ -Nash equilibrium, where

$$\epsilon_i^k := \left[\sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{k,-i}, \beta)) - J_i(\Lambda^k) - \epsilon_i^0 \right]^+ \rightarrow 0 \text{ as } k \rightarrow \infty.$$

Proof of Theorem 3.2.10. Recall that strong strategies are insensitive to the choice of n -player environment (see Remark 3.2.11), and so it suffices to prove the theorem on any given sequence of environments, such as those provided by Theorem 6.2.1. By Theorem 6.2.1 we may find $\epsilon_n \rightarrow 0$ and a relaxed $(\epsilon_n, \dots, \epsilon_n)$ -Nash equilibrium Λ^n for the n -player game, with the desired convergence properties. Then, by Proposition 6.2.2, we find for each n each k a strong $\epsilon^{n,k} = (\epsilon_n + \epsilon_1^{n,k}, \dots, \epsilon_n + \epsilon_n^{n,k})$ -Nash equilibrium $\Lambda^{n,k} \in \mathcal{A}_n^n(\mathcal{E}_n)$ with the convergence properties defined in Proposition 6.2.2. For each n , choose k_n large enough to make $\epsilon_i^{n,k_n} \leq 2^{-n}$ for each $i = 1, \dots, n$ and so that the sequences in (1-3) of Proposition 6.2.2 are each within 2^{-n} of their respective limits. \square

6.2.1 Construction of environments

Fix a weak MFG solution P . We will work on the space

$$\bar{\Omega} := [0, 1] \times \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{X}^\infty.$$

Let $(U, B, \mu, (W^i, \Lambda^i, Y^i)_{i=1}^\infty)$ denote the identity map (i.e., coordinate processes) on $\bar{\Omega}$. For $n \in \mathbb{N} \cup \{\infty\}$, consider the complete filtration $\bar{\mathbb{F}}^n = (\bar{\mathcal{F}}_t^n)_{t \in [0, T]}$ generated by U, B, μ , and $(W^i, \Lambda^i, Y^i)_{i=1}^n$, that is the completion of

$$\sigma \left\{ (U, B_s, \mu(C_1), (W_s^i, \Lambda^i(C_2), Y_s^i)_{i=1}^n) : s \leq t, C_1 \in \mathcal{F}_t^{\mathcal{X}}, C_2 \in \mathcal{B}([0, t] \times A) \right\}.$$

Let $P_{B, \mu} := P \circ (B, \mu)^{-1}$, and define the probability measure \mathbb{P} on $(\bar{\Omega}, \bar{\mathcal{F}}_T^\infty)$ by

$$\mathbb{P} := du P_{B, \mu}(d\beta, d\nu) \prod_{i=1}^{\infty} \nu(dw^i, dq^i, dy^i).$$

By construction,

$$\mathbb{P} \circ (Y_0^i, B, W^i, \mu, \Lambda^i, Y^i)^{-1} = P, \text{ for each } i,$$

and $(W^i, \Lambda^i, Y^i)_{i=1}^\infty$ are conditionally i.i.d. with common law μ given (B, μ) . Moreover, U and $(B, \mu, (W^i, \Lambda^i, Y^i)_{i=1}^\infty)$ are independent under \mathbb{P} . We will work with the n -player environments

$$\mathcal{E}_n := (\bar{\Omega}, \bar{\mathbb{F}}^n, \mathbb{P}, (Y_0^1, \dots, Y_0^n), B, (W^1, \dots, W^n)),$$

and we will show that the canonical process $(\Lambda^1, \dots, \Lambda^n)$ is a relaxed $(\epsilon_n, \dots, \epsilon_n)$ -Nash equilibrium for some $\epsilon_n \rightarrow \infty$. Including the seemingly superfluous random variable U makes the class $\mathcal{A}_n(\mathcal{E}_n)$ of admissible controls as rich as possible, in a sense which will be more clear later; until the proof of Proposition 6.2.2, U will be behind the scenes.

Define $X[\beta]$ and $\hat{\mu}[\beta]$ for $\beta \in \mathcal{A}_n^n(\mathcal{E}_n)$ as usual, as in Section 3.2. For each $\bar{\mathbb{F}}^\infty$ -progressive $\mathcal{P}(A)$ -valued process β on $\bar{\Omega}$ and each $i \geq 1$, define $Y^i[\beta]$ to be the unique solution of the SDE

$$dY_t^i[\beta] = \int_A b(t, Y_t^i[\beta], \mu_t^x, a) \beta_t(da) + \sigma(t, Y_t^i[\beta], \mu_t^x) dW_t^i + \sigma_0(t, Y_t^i[\beta], \mu_t^x) dB_t,$$

with $Y_0^i[\beta] = Y_0^i$. Note that if $\beta = (\beta^1, \dots, \beta^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$ then $X^i[\beta]$ differs from $Y^i[\beta^i]$ only in the measure flow which appears in the dynamics; $X^i[\beta]$ depends on the empirical measure flow of $(X^1[\beta], \dots, X^n[\beta])$, whereas $Y^i[\beta^i]$ depends on the random measure μ coming from the MFG solution. Define the canonical n -player strategy profile by

$$\bar{\Lambda}^n = (\bar{\Lambda}^{n,1}, \dots, \bar{\Lambda}^{n,n}) := (\Lambda^1, \dots, \Lambda^n) \in \mathcal{A}_n^n(\mathcal{E}_n).$$

This abbreviation serves in part to indicate which n we are working with at any given moment, so that we can suppress the index n from the rest of the notation. Note that $Y^i[\bar{\Lambda}^{n,i}] = Y^i[\Lambda^i] = Y^i$.

6.2.2 Trajectorial propagation of chaos

Intuition from the theory of propagation of chaos suggests that the state processes (Y^1, \dots, Y^n) and (X^1, \dots, X^n) should be close in some sense, and the purpose of this section is to make this quantitative. For $\beta \in \mathcal{A}_n(\mathcal{E}_n)$, abbreviate

$$(\bar{\Lambda}^{n,-i}, \beta) := ((\bar{\Lambda}^n)^{-i}, \beta) \in \mathcal{A}_n^n(\mathcal{E}_n).$$

Recall the definition of the metric $d_{\mathcal{X}}$ on \mathcal{X} from (6.5), and again define the p' -Wasserstein metric $\ell_{\mathcal{X},p'}$ on $\mathcal{P}^p(\mathcal{X})$ relative to the metric $d_{\mathcal{X}}$.

Lemma 6.2.3. *Fix i and a \mathbb{F}^∞ -progressive $P(A)$ -valued process β , and define*

$$\hat{\nu}^{n,i}[\beta] := \frac{1}{n} \left(\sum_{k \neq i}^n \delta_{(W^k, \Lambda^k, Y^k)} + \delta_{(W^i, \beta, Y^i[\beta])} \right).$$

There exists a sequence $\delta_n > 0$ converging to zero such that

$$\mathbb{E} \left[\ell_{\mathcal{X},p'}^{p'}(\hat{\nu}^{n,i}[\beta], \mu) \right] \leq \delta_n \left(1 + \mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t(da) dt \right).$$

Proof. Expectations are all with respect to \mathbb{P} throughout the proof. For $1 \leq i \leq n$ define

$$\hat{\nu}^n := \frac{1}{n} \sum_{k=1}^n \delta_{(W^k, \Lambda^k, Y^k)}.$$

Using the obvious coupling, we find

$$\ell_{\mathcal{X},p'}^{p'}(\hat{\nu}^{n,i}[\beta], \hat{\nu}^n) \leq \frac{1}{n} d_{\mathcal{X}}^{p'}((W^i, \Lambda^i, Y^i), (W^i, \beta, Y^i[\beta])).$$

Using (6.4), we find a constant $C > 0$, depending only on p, p' , and T , such that

$$\begin{aligned} & \mathbb{E} \left[d_{\mathcal{X}}^{p'}((W^i, \Lambda^i, Y^i), (W^i, \beta, Y^i[\beta])) \right] \\ & \leq C \mathbb{E} \left[\int_0^T \int_A |a|^{p'}(\beta_t + \Lambda_t^i)(da)dt + \|Y^i\|_T^{p'} + \|Y^i[\beta]\|_T^{p'} \right] \end{aligned}$$

Analogously to Lemma 6.1.1, it holds that

$$\mathbb{E}[\|Y^i[\beta]\|_T^{p'}] \leq c_5 \mathbb{E} \left[1 + |Y_0^i|^{p'} + \int_{\mathcal{C}^d} \|z\|_T^{p'} \mu^x(dz) + \int_0^T \int_A |a|^{p'} \beta_t(da)dt \right]. \quad (6.15)$$

Note that $\mathbb{E} \int_{\mathcal{C}^d} \|z\|_T^{p'} \mu^x(dz) < \infty$ and that $\mathbb{E}[|Y_0^i|^{p'}] = \mathbb{E}[|Y_0^1|^{p'}] < \infty$. Apply (6.15) also with $\beta = \Lambda^i$, we find a new constant, still called C and still independent of n , such that

$$\mathbb{E} \left[d_{\mathcal{X}}^{p'}((W^i, \Lambda^i, Y^i), (W^i, \beta, Y^i[\beta])) \right] \leq C \left(1 + \mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t(da)dt \right).$$

Finally, recall that $(W^k, \Lambda^k, Y^k)_{k=1}^\infty$ are conditionally i.i.d. given (B, μ) with common conditional law μ . Since also they are p' -integrable, it follows from the law of large numbers that

$$\lim_{n \rightarrow \infty} \mathbb{E} \left[\ell_{\mathcal{X}, p'}^{p'}(\widehat{\nu}^n, \mu) \right] = 0.$$

Complete the proof by using the triangle inequality to get

$$\mathbb{E} \left[\ell_{\mathcal{X}, p'}^{p'}(\widehat{\nu}^{n,i}[\beta], \mu) \right] \leq \frac{C 2^{p'-1}}{n} \left(1 + \mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t(da)dt \right) + 2^{p'-1} \mathbb{E} \left[\ell_{\mathcal{X}, p'}^{p'}(\widehat{\nu}^n, \mu) \right].$$

□

Lemma 6.2.4. *There is a sequence $\delta_n > 0$ converging to zero such that for each $1 \leq i \leq n$ and each $\beta \in \mathcal{A}_n(\mathcal{E}_n)$,*

$$\begin{aligned} & \mathbb{E}^{\mathbb{P}} \left[\ell_{\mathcal{X}, p'}^{p'}(\hat{\mu}[(\overline{\Lambda}^{n,-i}, \beta)], \mu) + \left\| X^i[(\overline{\Lambda}^{n,-i}, \beta)] - Y^i[\beta] \right\|_T^{p'} \right] \\ & \leq \delta_n \left(1 + \mathbb{E}^{\mathbb{P}} \int_0^T \int_A |a|^{p'} \beta_t(da)dt \right). \end{aligned}$$

Proof. The proof is similar to that of Lemma 6.1.6, and we work again with the truncated p' -Wasserstein distances ℓ_t on \mathcal{C}^d defined in (6.6). Throughout this proof, n and i are fixed, and expectations are all with respect to \mathbb{P} . Abbreviate $\overline{X}^k = X^k[(\overline{\Lambda}^{n,-i}, \beta)]$ and $\hat{\mu} = \hat{\mu}[(\overline{\Lambda}^{n,-i}, \beta)]$ throughout. Define $\overline{Y}^i := Y^i[\beta]$ and $\overline{Y}^k := Y^k$ for $k \neq i$. As in the proof of Lemma 6.1.6, we use the Burkholder-Davis-Gundy inequality followed by Gronwall's inequality to find a constant $C_1 > 0$, depending only on c_1, p' , and T , such that

$$\mathbb{E} \left[\|\overline{X}^k - \overline{Y}^k\|_t^{p'} \right] \leq C_1 \mathbb{E} \int_0^t \ell_s^{p'}(\hat{\mu}^x, \mu^x) ds, \text{ for } 1 \leq k \leq n. \quad (6.16)$$

Define $\widehat{\nu}^{n,i} = \widehat{\nu}^{n,i}[\beta]$ as in Lemma 6.2.3, and write $\widehat{\nu}^{n,i,x} := (\widehat{\nu}^{n,i})^x$ for the empirical distribution of $(\bar{Y}^1, \dots, \bar{Y}^n)$. Use (6.16) and the triangle inequality to get

$$\begin{aligned} \frac{1}{n} \sum_{k=1}^n \mathbb{E} \left[\|\bar{X}^k - \bar{Y}^k\|_t^{p'} \right] &\leq 2^{p'-1} C_1 \mathbb{E} \int_0^t \left(\ell_s^{p'}(\hat{\mu}^x, \widehat{\nu}^{n,i,x}) + \ell_s^{p'}(\widehat{\nu}^{n,i,x}, \mu^x) \right) ds \\ &\leq 2^{p'-1} C_1 \mathbb{E} \int_0^t \left(\frac{1}{n} \sum_{k=1}^n \|\bar{X}^k - \bar{Y}^k\|_s^{p'} + \ell_s^{p'}(\widehat{\nu}^{n,i,x}, \mu^x) \right) ds \end{aligned}$$

By Gronwall's inequality and Lemma 6.2.3, with $C_2 := 2^{p'-1} C_1 e^{2^{p'-1} C_1 T}$ we have

$$\begin{aligned} \frac{1}{n} \sum_{k=1}^n \mathbb{E} \left[\|\bar{X}^k - \bar{Y}^k\|_t^{p'} \right] &\leq C_2 \mathbb{E} \int_0^t \ell_s^{p'}(\widehat{\nu}^{n,i,x}, \mu^x) ds \leq C_2 T \mathbb{E} \left[\ell_{\mathcal{X},p'}^{p'}(\widehat{\nu}^{n,i}, \mu) \right] \\ &\leq C_2 T \delta_n \left(1 + \mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t(da) dt \right). \end{aligned} \quad (6.17)$$

The obvious coupling yields the inequality

$$\ell_{\mathcal{X},p'}^{p'}(\hat{\mu}, \widehat{\nu}^{n,i}) \leq \frac{1}{n} \sum_{k=1}^n \|\bar{X}^k - \bar{Y}^k\|_T^{p'},$$

and then the triangle inequality implies

$$\mathbb{E} \left[\ell_{\mathcal{X},p'}^{p'}(\hat{\mu}, \mu) \right] \leq 2^{p'-1} \frac{1}{n} \sum_{k=1}^n \mathbb{E} \left[\|\bar{X}^k - \bar{Y}^k\|_T^{p'} \right] + 2^{p'-1} \mathbb{E} \left[\ell_{\mathcal{X},p'}^{p'}(\widehat{\nu}^{n,i}, \mu) \right].$$

Conclude from Lemma 6.2.3 and (6.17). \square

6.2.3 Proof of Theorem 6.2.1

With Lemma 6.2.4 in hand, we begin the proof of Theorem 6.2.1. The convergence $P_n \rightarrow P$ follows immediately from Lemma 6.2.4, and it remains only to check that $\bar{\Lambda}^n$ is a relaxed $(\epsilon_n, \dots, \epsilon_n)$ -Nash equilibrium for some $\epsilon_n \rightarrow 0$. Define

$$\begin{aligned} \epsilon_n &:= \max_{i=1}^n \left[\sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\bar{\Lambda}^{n,-i}, \beta)) - J_i(\bar{\Lambda}^n) \right] \\ &= \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_1((\bar{\Lambda}^{n,-1}, \beta)) - J_1(\bar{\Lambda}^n), \end{aligned}$$

where the second equality follows from exchangeability, or more precisely from the fact that (using the notation of Remark 3.2.5) the measure

$$\mathbb{P} \circ (\xi_\pi, B, W_\pi, \hat{\mu}[\bar{\Lambda}_\pi^n], \bar{\Lambda}_\pi^n, X[\bar{\Lambda}_\pi^n]_\pi)^{-1}$$

does not depend on the choice of permutation π . Recall that $P \in \mathcal{P}(\Omega)$ was the given MFG solution, and define $\rho := P \circ (\xi, B, W, \mu)^{-1}$ so that $P \in \mathcal{RA}^*(\rho)$. For each n , find $\beta^n \in \mathcal{A}_n(\mathcal{E}_n)$ such that

$$J_1((\bar{\Lambda}^{n,-1}, \beta^n)) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_1((\bar{\Lambda}^{n,-1}, \beta)) - 1/n. \quad (6.18)$$

To complete the proof, it suffices to prove the following:

$$\lim_{n \rightarrow \infty} J_1(\bar{\Lambda}^n) = \mathbb{E}^\mathbb{P} [\Gamma(\mu^x, \Lambda^1, Y^1)], \quad (6.19)$$

$$\lim_{n \rightarrow \infty} \left| \mathbb{E}^\mathbb{P} [\Gamma(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta^n)], \beta^n, X^1[(\bar{\Lambda}^{n,-1}, \beta^n)]) - \Gamma(\mu^x, \beta^n, Y^1[\beta^n])] \right| = 0. \quad (6.20)$$

To see this, note that $\mathbb{P} \circ (\xi^1, B, W^1, \mu, \Lambda^1, Y^1)^{-1} = P$ holds by construction. Since

$$P'_n := \mathbb{P} \circ (\xi^1, B, W^1, \mu, \beta^n, Y^1[\beta^n])^{-1}$$

is in $\mathcal{RA}(\rho)$ for each n , and since P is in $\mathcal{RA}^*(\rho)$, we have

$$\mathbb{E}^\mathbb{P} [\Gamma(\mu^x, \beta^n, Y^1)] = J(P) \geq J(P'_n) = \mathbb{E}^\mathbb{P} [\Gamma(\mu^x, \beta^n, Y^1[\beta^n])], \text{ for all } n.$$

Thus, from (6.19) and (6.20) it follows that

$$\begin{aligned} \lim_{n \rightarrow \infty} J_1(\bar{\Lambda}^n) &\geq \limsup_{n \rightarrow \infty} \mathbb{E}^\mathbb{P} [\Gamma(\mu^x, \beta^n, Y^1[\beta^n])] \\ &= \limsup_{n \rightarrow \infty} J_1((\bar{\Lambda}^{n,-1}, \beta^n)) \\ &= \limsup_{n \rightarrow \infty} \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_1((\bar{\Lambda}^{n,-1}, \beta)), \end{aligned}$$

where of course in the last step we have used (6.18). Since $\epsilon_n \geq 0$, this shows $\epsilon_n \rightarrow 0$.

Proof of (6.19):

First, apply Lemma 6.2.4 with $\beta = \Lambda^1$ (so that $(\bar{\Lambda}^{n,-1}, \beta) = \bar{\Lambda}^n$) to get

$$\lim_{n \rightarrow \infty} \mathbb{P} \circ (Y_0^1, B, W^1, \hat{\mu}[\bar{\Lambda}^n], \Lambda^1, X^1[\bar{\Lambda}^n])^{-1} = \mathbb{P} \circ (Y_0^1, B, W^1, \mu, \Lambda^1, Y^1)^{-1},$$

where the limit is taken in $\mathcal{P}^p(\Omega)$. Moreover, by Lemma 6.1.2 we must have $\mathbb{E}^\mathbb{P} \int_0^T \int_A |a|^{p'} \Lambda_t^1(da) dt < \infty$, while Lemma 6.1.1 and symmetry imply that

$$\sup_n \mathbb{E}^\mathbb{P} \int_{\mathcal{C}^d} \|x\|_T^{p'} \hat{\mu}^x[\bar{\Lambda}^n](dx) \leq c_5 \left[1 + \mathbb{E}^\mathbb{P} |\xi^1|^{p'} + \mathbb{E}^\mathbb{P} \int_0^T \int_A |a|^{p'} \Lambda_t^1(da) dt \right] < \infty.$$

This verifies the additional integrability conditions required of the continuity result Lemma 5.3.4 for the functional J , and we conclude that

$$\lim_{n \rightarrow \infty} J_1(\bar{\Lambda}^n) = \lim_{n \rightarrow \infty} \mathbb{E}^\mathbb{P} [\Gamma(\hat{\mu}^x[\bar{\Lambda}^n], \Lambda^1, X^1[\bar{\Lambda}^n])] = \mathbb{E}^\mathbb{P} [\Gamma(\mu^x, \Lambda^1, Y^1)].$$

Proof of (6.20):

This step is fairly involved and is thus divided into several steps. The first two steps identify a relative compactness for the laws of the empirical measure and state process pairs, crucial for the third and fourth steps below. Step (3) focuses on the g term, and Step (4) uses the additional assumption A4 to deal with the f term.

Proof of (6.20), Step (1):

We show first that

$$\sup_n \mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t^n(da) dt < \infty. \quad (6.21)$$

Using (6.18) to apply Lemma 6.1.2(2), we get

$$\begin{aligned} \mathbb{E} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \beta_t^n(da) dt &\leq c_7 \mathbb{E} \left[1 + \frac{1}{n} + |\xi^1|^p + \frac{1}{n} \sum_{i=2}^n \int_0^T \int_A |a|^p \Lambda_t^i(da) dt \right] \\ &= c_7 \mathbb{E} \left[1 + \frac{1}{n} + |\xi^1|^p + \frac{n-1}{n} \int_0^T \int_A |a|^p \Lambda_t^1(da) dt \right], \end{aligned}$$

where the second line follows from symmetry. Since $\mathbb{E} \int_0^T \int_A |a|^p \Lambda_t^1(da) dt < \infty$ and $\mathbb{E}[|\xi^1|^p] < \infty$, we have proven (6.21).

Proof of (6.20), Step (2):

Define \mathcal{A}_R for $R > 0$ to be the set of $\bar{\mathbb{F}}^\infty$ -progressive $\mathcal{P}(A)$ -valued processes β such that

$$\mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t(da) dt \leq R.$$

According to (6.21), there exists $R > 0$ such that $\beta^n \in \mathcal{A}_R$ for all n . Define also

$$S_R := \left\{ \mathbb{P} \circ \left(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], X^1[(\bar{\Lambda}^{n,-1}, \beta)] \right)^{-1} : n \geq 1, \beta \in \mathcal{A}_R \right\}.$$

We show next that S_R is relatively compact in $\mathcal{P}^p(\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d)$. Note first that it follows from Lemma 6.1.1 that

$$\sup \left\{ \mathbb{E}^\mathbb{P} \int_{\mathcal{C}^d} \|z\|_T^{p'} \hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)](dz) : n \geq 1, \beta \in \mathcal{A}_R \right\} < \infty. \quad (6.22)$$

By Proposition 6.1.3 the set

$$\left\{ \frac{1}{n} \sum_{k=1}^n \mathbb{P} \circ (X^k[(\bar{\Lambda}^{n,-k}, \beta)])^{-1} : n \geq 1, \beta \in \mathcal{A}_R \right\}$$

is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$. By symmetry, this set is identical to

$$\left\{ \mathbb{P} \circ (X^1[(\bar{\Lambda}^{n,-1}, \beta)])^{-1} : n \geq 1, \beta \in \mathcal{A}_R \right\}.$$

For $\beta \in \mathcal{A}_R$, the mean measure of $\mathbb{P} \circ (\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)])^{-1}$ is exactly

$$\frac{1}{n} \sum_{k=1}^n \mathbb{P} \circ (X^k[(\bar{\Lambda}^{n,-1}, \beta)])^{-1},$$

and it follows again from Proposition 6.1.3 that the family

$$\left\{ \frac{1}{n} \sum_{k=1}^n \mathbb{P} \circ (X^k[(\bar{\Lambda}^{n,-1}, \beta)])^{-1} : n \geq 1, \beta \in \mathcal{A}_R \right\}$$

is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$. From this and (6.22) we conclude that $\mathbb{P} \circ (\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)])^{-1}$ are relatively compact in $\mathcal{P}^p(\mathcal{P}^p(\mathcal{C}^d))$, by Lemma 2.1.13. Since each family of marginals is relatively compact, so is S_R (see Lemma 2.1.8).

Proof of (6.20), Step (3):

Since $\beta^n \in \mathcal{A}_R$ for each n , to prove (6.20) it suffices to show that

$$\sup_{\beta \in \mathcal{A}_R} I_n^\beta \rightarrow 0, \tag{6.23}$$

where

$$\begin{aligned} I_n^\beta &:= \mathbb{E} \left[\Gamma(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], \beta, X^1[(\bar{\Lambda}^{n,-1}, \beta)]) - \Gamma(\mu^x, \beta, Y^1[\beta]) \right] \\ &= \mathbb{E} \left[\int_0^T \int_A \left(f(t, X_t^1[(\bar{\Lambda}^{n,-1}, \beta)], \hat{\mu}_t^x[(\bar{\Lambda}^{n,-1}, \beta)], a) - f(t, Y_t^1[\beta], \mu_t^x, a) \right) \beta_t(da) dt \right] \\ &\quad + \mathbb{E} \left[g(X_T^1[(\bar{\Lambda}^{n,-1}, \beta)], \hat{\mu}_T^x[(\bar{\Lambda}^{n,-1}, \beta)]) - g(Y_T^1[\beta], \mu_T^x) \right]. \end{aligned}$$

We start with the g term. Define

$$\begin{aligned} Q_n^\beta &:= \mathbb{P} \circ (\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], X^1[(\bar{\Lambda}^{n,-1}, \beta)])^{-1}, \\ Q^\beta &:= \mathbb{P} \circ (\mu^x, Y^1[\beta])^{-1}. \end{aligned}$$

Using the metric on $\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d$ given by

$$((\mu, x), (\mu', x')) \mapsto \left[\ell_{\mathcal{C}^d, p}^p(\mu, \mu') + \|x - x'\|_T^p \right]^{1/p},$$

we define the p -Wasserstein metric $\ell_{\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d, p}$ on $\mathcal{P}^p(\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d)$. By Lemma 6.2.4, we have

$$\begin{aligned} \ell_{\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d, p}^{p'}(Q_n^\beta, Q^\beta) &\leq \mathbb{E} \left[\ell_{\mathcal{C}^d, p}^p \left(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], \mu^x \right) + \|X^1[(\bar{\Lambda}^{n,-1}, \beta)] - Y^1[\beta]\|_T^p \right]^{p'/p} \\ &\leq 2^{p'/p-1} \mathbb{E} \left[\ell_{\mathcal{X}, p'}^{p'} \left(\hat{\mu}[(\bar{\Lambda}^{n,-1}, \beta)], \mu \right) + \|X^1[(\bar{\Lambda}^{n,-1}, \beta)] - Y^1[\beta]\|_T^{p'} \right] \\ &\leq 2^{p'/p-1} \delta_n (1 + R), \end{aligned}$$

and thus $Q_n^\beta \rightarrow Q^\beta$ in $\mathcal{P}^p(\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d)$, uniformly in $\beta \in \mathcal{A}_R$. The function

$$\mathcal{P}^p(\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d) \ni Q \mapsto \int Q(d\nu, dx) g(x_T, \nu_T)$$

is continuous, and so its restriction to the closure of S_R is *uniformly* continuous. Thus, since $\{Q_n^\beta : n \geq 1, \beta \in \mathcal{A}_R\} \subset S_R$,

$$\lim_{n \rightarrow \infty} \sup_{\beta \in \mathcal{A}_R} \left| \mathbb{E} \left[g(X_T^1[(\bar{\Lambda}^{n,-1}, \beta)], \hat{\mu}_T^x[(\bar{\Lambda}^{n,-1}, \beta)]) - g(Y_T^1[\beta], \mu_T^x) \right] \right| = 0.$$

Proof of (6.20), Step (4):

To deal with the f term in I_n^β it will be useful to define $G : \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d \rightarrow \mathbb{R}$ by

$$G((\mu^1, x^1), (\mu^2, x^2)) := \int_0^T \sup_{a \in A} |f(t, x_t^1, \mu_t^1, a) - f(t, x_t^2, \mu_t^2, a)| dt$$

With the g term taken care of in Step (3) above, the proof of (6.23) and thus the theorem will be complete if we show that

$$\begin{aligned} 0 &= \lim_{n \rightarrow \infty} \sup_{\beta \in \mathcal{A}_R} \mathbb{E} [Z_\beta^n], \text{ where} \\ Z_\beta^n &:= G \left((\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], X^1[(\bar{\Lambda}^{n,-1}, \beta)]), (\mu^x, Y^1[\beta]) \right). \end{aligned} \tag{6.24}$$

Fix $\eta > 0$, and note that by relative compactness of S_R we may find (e.g., by Proposition 2.1.7) a compact set $K \subset \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{C}^d$ such that, if the event K_β is defined by

$$K_\beta := \left\{ \left(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], X^1[(\bar{\Lambda}^{n,-1}, \beta)] \right) \in K \right\},$$

then

$$\mathbb{E} \left[\left(1 + \int_{\mathcal{C}^d} \|z\|_T^p \hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)](dz) + \|X^1[(\bar{\Lambda}^{n,-1}, \beta)]\|_T^p \right) 1_{K_\beta^c} \right] \leq \eta,$$

for all $n \geq 1$ and $\beta \in \mathcal{A}_R$. Sending $n \rightarrow \infty$, it follows from Lemma 6.2.4 that also

$$\mathbb{E} \left[\left(1 + \int_{\mathcal{C}^d} \|z\|_T^p \mu^x(dz) + \|Y^1[\beta]\|_T^p \right) 1_{K_\beta^c} \right] \leq \eta.$$

Hence, the growth condition on f of assumption A4 implies

$$\mathbb{E} \left[1_{K_\beta^c} Z_\beta^n \right] \leq c_4 \eta, \quad (6.25)$$

for all $n \geq 1$ and $\beta \in \mathcal{A}_R$. Assumption A4 implies that G is continuous, and thus uniformly continuous on $K \times K$. We will check next that $\mathbb{E}[1_{K_\beta^c} Z_\beta^n]$ converges to zero, uniformly in $\beta \in \mathcal{A}_R$. Indeed, by uniform continuity there exists $\eta_0 > 0$ such that if $(\mu^1, x^1), (\mu^2, x^2) \in K$ and $G((\mu^1, x^1), (\mu^2, x^2)) > \eta$ then $\|x^1 - x^2\|_T + \ell_{\mathcal{C}^d, p}(\mu^1, \mu^2) > \eta_0$. Thus, since G is bounded on $K \times K$, say by $C > 0$, we use Markov's inequality and Lemma 6.2.4 to conclude that

$$\begin{aligned} \mathbb{E} [1_{K_\beta^c} Z_\beta^n] &\leq \eta + C \mathbb{P} \left\{ \left\| X^1[(\bar{\Lambda}^{n,-1}, \beta)] - Y^1[\beta] \right\|_T + \ell_{\mathcal{C}^d, p} \left(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], \mu^x \right) > \eta_0 \right\} \\ &\leq \eta + 2^{p'-1} C \eta_0^{-p'} \mathbb{E} \left[\left\| X^1[(\bar{\Lambda}^{n,-1}, \beta)] - Y^1[\beta] \right\|_T^{p'} + \ell_{\mathcal{C}^d, p}^{p'} \left(\hat{\mu}^x[(\bar{\Lambda}^{n,-1}, \beta)], \mu^x \right) \right] \\ &\leq \eta + 2^{p'-1} C \eta_0^{-p'} \delta_n \left(1 + \mathbb{E} \int_0^T \int_A |a|^{p'} \beta_t(da) dt \right) \\ &\leq \eta + 2^{p'-1} C \eta_0^{-p'} \delta_n (1 + R), \end{aligned}$$

whenever $\beta \in \mathcal{A}_R$, where $\delta_n \rightarrow 0$ is from Lemma 6.2.4. Combining this with (6.25), we get

$$\limsup_{n \rightarrow \infty} \sup_{\beta \in \mathcal{A}_R} \mathbb{E} [Z_\beta^n] \leq (1 + c_4) \eta.$$

This holds for each $\eta > 0$, completing the proof of (6.24) and thus of the theorem. \square

6.2.4 Proof of Proposition 6.2.2

Throughout the section, the number of agents n is fixed, and we work on the n -player environment \mathcal{E}_n specified in Section 6.2.1. The proof of Proposition 6.2.2 is split into two main steps. In this first step, we approximate the relaxed strategy Λ^0 by bounded strong strategies, and we check the convergences (1) and (2) claimed in Proposition 6.2.2. The idea behind this approximation is to view the *product* $\prod_i \Lambda_t^{0,i}$ as a relaxed control on A^n , and use the density of strong controls as in Proposition 2.1.15. The second step verifies the somewhat more subtle inequality (3) of Proposition 6.2.2. First, we need the following lemma, which is a simple variant of a standard result:

Lemma 6.2.5. *Suppose $\tilde{\Lambda}^k = (\tilde{\Lambda}^{k,1}, \dots, \tilde{\Lambda}^{k,n}) \in \mathcal{A}_n^n(\mathcal{E}_n)$ is such that*

$$\lim_{k \rightarrow \infty} \mathbb{P} \circ (\xi, B, W, \tilde{\Lambda}^k)^{-1} = \mathbb{P} \circ (\xi, B, W, \Lambda^0)^{-1},$$

with the limit taken in $\mathcal{P}((\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n)$. Then

$$\lim_{k \rightarrow \infty} \mathbb{P} \circ \left(B, W, \tilde{\Lambda}^k, X[\tilde{\Lambda}^k] \right)^{-1} = \mathbb{P} \circ \left(B, W, \Lambda^0, X[\Lambda^0] \right)^{-1},$$

in $\mathcal{P}(\mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n)$.

Proof. This is analogous to the proof of Lemma 5.3.3, which is itself an instance of a standard method proving weak convergence of SDE solutions, so we only sketch the proof. It can be shown as in Proposition 5.3.2 that $\{\mathbb{P} \circ (X[\tilde{\Lambda}^k])^{-1} : k \geq 1\}$ is relatively compact in $\mathcal{P}((\mathcal{C}^d)^n)$, and thus $\{\mathbb{P} \circ \left(B, W, \tilde{\Lambda}^k, X[\tilde{\Lambda}^k] \right)^{-1} : k \geq 1\}$ is relatively compact in $\mathcal{P}(\mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n)$. Using well known results on convergence of stochastic integrals, like those of Kurtz and Protter [81], it is straightforward to check that under any limit point the canonical processes satisfy a certain SDE, and the claimed convergence follows from uniqueness of the SDE solution. \square

Proof of Proposition 6.2.2.

Step 1: Abbreviate $\bar{\mathcal{V}} := \mathcal{V}[A^n]$, the space of relaxed controls on A^n , as in 2.1.3. Define

$$\bar{\Lambda}_t^0(da_1, \dots, da_n) := \prod_{i=1}^n \Lambda_t^{0,i}(da_i),$$

and identify this $\mathcal{P}(A^n)$ -valued process with the random element $\bar{\Lambda}^0 := dt \bar{\Lambda}_t^0(da)$ of $\bar{\mathcal{V}}$. By Proposition 2.1.15, with A replaced by A^n , there exists a sequence of bounded A^n -valued processes $\alpha^k = (\alpha^{k,1}, \dots, \alpha^{k,n})$ such that, if we define

$$\bar{\Lambda}^k := dt \delta_{\alpha_t^k}(da_1, \dots, da_n) = dt \prod_{i=1}^n \delta_{\alpha_t^{k,i}}(da_i),$$

then we have

$$\lim_{r \rightarrow \infty} \sup_k \mathbb{E}^{\mathbb{P}} \left[\int_0^T |\alpha_t^k|^{p'} 1_{\{|\alpha_t^k| > r\}} dt \right] = 0 \quad (6.26)$$

and

$$\lim_{k \rightarrow \infty} \mathbb{P} \circ \left(\xi, B, W, \bar{\Lambda}^k \right)^{-1} = \mathbb{P} \circ \left(\xi, B, W, \bar{\Lambda}^0 \right)^{-1},$$

in $\mathcal{P}((\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \bar{\mathcal{V}})$. Defining $\pi_i : [0, T] \times A^n \rightarrow [0, T] \times A$ by $\pi_i(t, a_1, \dots, a_n) := (t, a_i)$, we note that the map $\bar{\mathcal{V}} \ni q \mapsto q \circ \pi_i^{-1} \in \mathcal{V}$ is continuous. Define $\Lambda_t^{k,i} := \delta_{\alpha_t^{k,i}}$ and $\Lambda^k = (\Lambda^{k,1}, \dots, \Lambda^{k,n})$, and conclude that

$$\lim_{k \rightarrow \infty} \mathbb{P} \circ \left(\xi, B, W, \Lambda^k \right)^{-1} = \mathbb{P} \circ \left(\xi, B, W, \Lambda^0 \right)^{-1},$$

in $\mathcal{P}^p((\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n)$, for each k . This implies, by Lemma 6.2.5, that

$$\lim_{k \rightarrow \infty} \mathbb{P} \circ (B, W, \Lambda^k, X[\Lambda^k])^{-1} = \mathbb{P} \circ (B, W, \Lambda^0, X[\Lambda^0])^{-1},$$

in $\mathcal{P}^p(\mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n)$, verifying the first claim. Thanks to the state SDE estimate 6.1.1, the uniform p' -integrability of α^k implies the p' -moment bound

$$\sup_k \mathbb{E}^{\mathbb{P}} \left[\|X[\Lambda^k]\|_T^{p'} + \int_{\mathcal{C}^d} \|x\|_T^{p'} \hat{\mu}^x[\Lambda^k](dx) \right] < \infty.$$

The validity of the second claim now follows from the continuity of J of Lemma 5.3.4, which ensures that

$$\lim_{k \rightarrow \infty} J_i(\Lambda^k) = J_i(\Lambda^0), \quad i = 1, \dots, n.$$

Step 2: It remains to justify the third claim of Proposition 6.2.2. We prove this only for $i = 1$, since the cases $i = 2, \dots, n$ are identical. For each k find $\beta^k \in \mathcal{A}_n(\mathcal{E}_n)$ such that

$$J_i((\Lambda^{k,-1}, \beta^k)) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{k,-1}, \beta)) - \frac{1}{k}. \quad (6.27)$$

First, use the second part of Lemma 6.1.2 to get

$$\mathbb{E} \int_0^T \int_A (|a|^{p'} - c_6 |a|^p) \beta_t^k(da) dt \leq c_7 \mathbb{E} \left[1 + \frac{1}{k} + |\xi^1|^p + \frac{1}{n} \sum_{i=2}^n \int_0^T \int_A |a|^p \Lambda_t^{k,i}(da) dt \right].$$

Since $\mathbb{E}[|\xi^1|^p] < \infty$, and since

$$\lim_{k \rightarrow \infty} \mathbb{E} \int_0^T \int_A |a|^p \Lambda_t^{k,i}(da) dt = \mathbb{E} \int_0^T \int_A |a|^p \Lambda_t^{0,i}(da) dt < \infty,$$

holds by construction, for $i = 2, \dots, n$, it follows that

$$R := \sup_k \mathbb{E}^{\mathbb{P}} \int_0^T \int_A |a|^{p'} \beta_t^k(da) dt < \infty.$$

It follows from Proposition 5.3.2 that the set

$$\left\{ \mathbb{P} \circ ((\Lambda^{k,-1}, \beta^k), X[(\Lambda^{k,-1}, \beta^k)])^{-1} : k \geq 1 \right\}$$

is relatively compact in $\mathcal{P}^p(\mathcal{V}^n \times (\mathcal{C}^d)^n)$, and so the set

$$\left\{ P_k := \mathbb{P} \circ (B, W, (\Lambda^{k,-1}, \beta^k), X[(\Lambda^{k,-1}, \beta^k)])^{-1} : k \geq 1 \right\} \quad (6.28)$$

is relatively compact in $\mathcal{P}^p(\mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n)$. By the following Lemma 6.2.6, every limit point P of $(P_k)_{k=1}^\infty$ is of the form

$$P = \mathbb{P} \circ (B, W, (\Lambda^{0,-1}, \beta), X[(\Lambda^{0,-1}, \beta)])^{-1}, \text{ for some } \beta \in \mathcal{A}_n(\mathcal{E}_n). \quad (6.29)$$

This and the upper semicontinuity of J imply

$$\limsup_{k \rightarrow \infty} J_i((\Lambda^{k,-1}, \beta^k)) \leq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{0,-1}, \beta)).$$

Because of (6.27), this completes the proof of Proposition 6.2.2. \square

Lemma 6.2.6. *Every limit point P of $(P_k)_{k=1}^\infty$ (defined in (6.28)) is of the form (6.29).*

Proof. Let us abbreviate

$$\Omega^{(n)} := \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n.$$

Let $(B, W = (W^1, \dots, W^n), \Lambda = (\Lambda^1, \dots, \Lambda^n), X = (X^1, \dots, X^n))$ denote the identity map on $\Omega^{(n)}$, and let $(\mathcal{F}_t^{(n)})_{t \in [0, T]}$ denote the natural filtration,

$$\mathcal{F}_t^{(n)} = \sigma((B_s, W_s, \Lambda(C), X_s) : s \leq t, C \in \mathcal{B}([0, t] \times A)).$$

Fix a limit point P of P_k . It is easily verified that P satisfies

$$P \circ (X_0, B, W, (\Lambda^2, \dots, \Lambda^n))^{-1} = \mathbb{P} \circ (X_0, B, W, (\Lambda^{0,2}, \dots, \Lambda^{0,n}))^{-1}. \quad (6.30)$$

Moreover, for each k , we know that B and W are independent $(\mathcal{F}_t^{(n)})_{t \in [0, T]}$ -Wiener processes under P_k , and thus this is true under P as well. Note that $(B, W, (\Lambda^{k,-1}, \beta^k), X[(\Lambda^{k,-1}, \beta^k)])$ satisfy the state SDE under \mathbb{P} , or equivalently under P_k the canonical processes verify the following SDE system, where $i = 1, \dots, n$:

$$\begin{cases} dX_t^i &= \int_A b(t, X_t^i, \hat{\mu}_t^x, a) \Lambda_t^i(da) dt + \sigma(t, X_t^i, \hat{\mu}_t^x) dW_t^i + \sigma_0(t, X_t^i, \hat{\mu}_t^x) dB_t, \\ \hat{\mu}_t^x &= \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k}. \end{cases} \quad (6.31)$$

As is becoming familiar, this property passes to the limit (e.g., by [81]): The canonical processes on $\Omega^{(n)}$ verify the same SDE under P .

It remains only to show that there exists $\beta \in \mathcal{A}_n(\mathcal{E}_n)$ such that

$$\mathbb{P} \circ (X_0, B, W, (\Lambda^{0,-1}, \beta))^{-1} = P \circ (X_0, B, W, \Lambda)^{-1}. \quad (6.32)$$

Indeed, from uniqueness in law of the solution of the SDE (6.31) it will then follow that

$$\mathbb{P} \circ (B, W, (\Lambda^{0,-1}, \beta), X[(\Lambda^{0,-1}, \beta)])^{-1} = P.$$

The independent uniform random variable U built into \mathcal{E}_n now finally comes into play. Using a well known result from measure theory (e.g., [73, Theorem 6.10]) we may find a measurable

function

$$\bar{\beta} = (\bar{\beta}^1, \dots, \bar{\beta}^n) : [0, 1] \times (\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \rightarrow \mathcal{V}^n$$

such that

$$\mathbb{P} \circ (X_0, B, W, \bar{\beta}(U, X_0, B, W))^{-1} = P \circ (X_0, B, W, \Lambda)^{-1}. \quad (6.33)$$

Since B and W are independent $(\mathcal{F}_t^{(n)})_{t \in [0, T]}$ -Wiener processes under P , it follows that

$$(\bar{\beta}(U, X_0, B, W)_s)_{s \in [0, t]} \quad \text{and} \quad \sigma(B_s - B_t, W_s - W_t : s \in [t, T])$$

are independent under \mathbb{P} , for each $t \in [0, T]$. Thus, $(\bar{\beta}(U, X_0, B, W)_t)_{t \in [0, T]}$ is progressively measurable with respect to the \mathbb{P} -completion of the filtration $(\sigma(U, X_0, B_s, W_s : s \leq t))_{t \in [0, T]}$. In particular, $(\bar{\beta}(U, X_0, B, W))_{t \in [0, T]} \in \mathcal{A}_n(\mathcal{E}_n)$ and $\beta := (\bar{\beta}^1(U, X_0, B, W)_t)_{t \in [0, T]}$ is in $\mathcal{A}_n(\mathcal{E}_n)$. Now note that (6.30) and (6.33) together imply

$$\begin{aligned} & \mathbb{P} \circ \left(X_0, B, W, \left(\bar{\beta}^2(U, X_0, B, W), \dots, \bar{\beta}^n(U, X_0, B, W) \right) \right)^{-1} \\ &= P \circ (X_0, B, W, (\Lambda^2, \dots, \Lambda^n))^{-1}. \end{aligned}$$

On the other hand, (6.33) implies that the conditional law under P of Λ^1 given $(X_0, B, W, \Lambda^2, \dots, \Lambda^n)$ is the same as the conditional law under \mathbb{P} of $\bar{\beta}^1(U, X_0, B, W)$ given

$$(X_0, B, W, \bar{\beta}^2(U, X_0, B, W), \dots, \bar{\beta}^n(U, X_0, B, W)).$$

This completes the proof of (6.32). \square

6.3 Proof of Theorem 4.2.2

This section explains the proof of Theorem 4.2.2, which specializes the main results to the setting without common noise essentially by means of the following simple observation. Note that although we assume $\sigma_0 \equiv 0$ throughout the section, *weak MFG solution* has the same meaning as in Definition 3.1.1, distinct from Definition 4.1.1 of *weak MFG solution without common noise*.

Lemma 6.3.1. *Assume $\sigma_0 \equiv 0$. If $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$ is a weak MFG solution, then $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$ is a weak MFG solution without common noise. Conversely, if $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, W, \mu, \Lambda, X)$ is a weak MFG solution without common noise, then we may construct (by enlarging the probability space, if necessary) an m_0 -dimensional Wiener process B independent of (W, μ, Λ, X) such that $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$ is a weak MFG solution.*

Proof. Using the simplified characterization of MFG pre-solution given in Lemma 5.2.3 (see also Remark 5.3.9), this is nearly immediate. First, suppose we are given a weak MFG solution $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$. Since $\mu = P((W, \Lambda, X) \in \cdot \mid B, \mu)$, we may condition on both sides by μ to get $\mu = P((W, \Lambda, X) \in \cdot \mid \mu)$. The other necessary properties of a weak

solution without common noise are easy to check; the optimality condition follows simply from the fact that any $\mathbb{F}^{X_0, W, \mu}$ -progressive control is also $\mathbb{F}^{X_0, B, W, \mu}$ -progressive.

Conversely, given a weak MFG solution without common noise, given that B is independent of (W, μ, Λ, X) , it follows from $\mu = P((W, \Lambda, X) \in \cdot \mid \mu)$ that also $\mu = P((W, \Lambda, X) \in \cdot \mid B, \mu)$. Here, we should check the optimality condition more carefully. Given any $\mathbb{F}^{X_0, B, W, \mu}$ -progressive control $\Lambda' = (\Lambda'_t)_{t \in [0, T]}$, the independence of B and (X_0, W, μ) implies that $\sigma(\Lambda'_s : s \leq t)$ is conditionally independent of $\mathcal{F}_T^{X_0, W, \mu}$ given $\mathcal{F}_t^{X_0, W, \mu}$, for each t . Thus, using the optimality of Λ among compatible controls (property (5) of the Definition 4.1.1), we know that $\mathbb{E}^P[\Gamma(\mu^x, \Lambda, X)] \geq \mathbb{E}^P[\Gamma(\mu^x, \Lambda', X')]$, where X' is the state process controlled by Λ' . This verifies the optimality required of weak MFG solutions (with common noise), thanks to the density of strong (i.e $\mathbb{F}^{X_0, B, W, \mu}$ -progressive) controls in the family of compatible controls; see Proposition 2.1.15. Again, the remaining properties of a weak MFG solution are easy to check. \square

Proof of Theorem 4.2.2. At this point, the proof is mostly straightforward. The first claim, regarding the adaptation of Theorem 3.2.4, follows immediately from Theorem 3.2.4 and the observation of Lemma 6.3.1. The second claim, about adapting Theorem 3.2.10, is not so immediate but requires nothing new. First, notice that Theorem 6.2.1 remains true if we replace “weak MFG solution” by “weak MFG solution without common noise,” and if we define P_n instead by (4.2); this is a consequence of Theorem 6.2.1 and Lemma 6.3.1. Then, we must only check that Proposition 6.2.2 remains true if we replace “strong” by “very strong,” and if we replace the conclusion (1) by

$$(1') \text{ In } \mathcal{P}((\mathcal{C}^m)^n \times \mathcal{V}^n \times (\mathcal{C}^d)^n)$$

$$\lim_{k \rightarrow \infty} \mathbb{P}_n \circ (W, \Lambda^k, X[\Lambda^k])^{-1} = \mathbb{P}_n \circ (W, \Lambda^0, X[\Lambda^0])^{-1}.$$

It is straightforward to check that the proof of Proposition 6.2.2 given in Section 6.2.4 translates mutatis mutandis to this new setting. \square

6.4 Proofs of Propositions 3.2.2 and 4.2.1

6.4.1 Proof of Proposition 3.2.2

Step 1:

We first show that every strong ϵ -Nash equilibrium is also a relaxed ϵ -Nash equilibrium. Suppose $\Lambda = (\Lambda^1, \dots, \Lambda^n) \in \mathcal{A}_n^n(\mathcal{E}_n)$ is a strong ϵ -Nash equilibrium on \mathcal{E}_n . Lemma 6.1.2(3) implies

$$\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^{p'} \Lambda_t^i(da) dt < \infty, \quad i = 1, \dots, n.$$

Let $\delta > 0$, and find $\beta^* \in \mathcal{A}_n(\mathcal{E}_n)$ such that

$$J_i((\Lambda^{-i}, \beta^*)) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{-i}, \beta)) - \delta. \quad (6.34)$$

Lemma 6.1.2(2) implies $\mathbb{E}^{\mathbb{P}_n} \int_0^T \int_A |a|^{p'} \beta_t^*(da) dt < \infty$. Thus, by Corollary 5.3.6, we may find a sequence of $(\mathcal{F}_t^{s,n})_{t \in [0, T]}$ -progressively measurable A -valued processes $(\alpha_t^k)_{t \in [0, T]}$ such that

$$\lim_{r \rightarrow \infty} \sup_k \mathbb{E}^{\mathbb{P}_n} \int_0^T |\alpha_t^k|^{p'} 1_{\{|\alpha_t^k| > r\}} dt = 0, \quad (6.35)$$

and

$$\mathbb{P}_n \circ (\xi, B, W, \beta^*)^{-1} = \lim_{k \rightarrow \infty} \mathbb{P}_n \circ \left(\xi, B, W, dt \delta_{\alpha_t^k}(da) \right)^{-1},$$

in $\mathcal{P}((\mathbb{R}^d)^n \times \mathcal{C}^{m_0} \times (\mathcal{C}^m)^n \times \mathcal{V})$. Abbreviate $\beta^k = dt \delta_{\alpha_t^k}(da)$. Since Λ is a strong strategy, we may write $\Lambda = \widehat{\Lambda}(\xi, B, W)$ for some measurable function $\widehat{\Lambda}$, and it follows (using Lemma 2.1.9 to deal with the potential discontinuity of $\widehat{\Lambda}$) that

$$\mathbb{P}_n \circ (\xi, B, W, (\Lambda^{-i}, \beta^*))^{-1} = \lim_{k \rightarrow \infty} \mathbb{P}_n \circ (\xi, B, W, (\Lambda^{-i}, \beta^k))^{-1}, \quad (6.36)$$

Lemma 6.2.5 gives

$$\mathbb{P}_n \circ (\xi, B, W, (\Lambda^{-i}, \beta^*), X[(\Lambda^{-i}, \beta^*)])^{-1} = \lim_{k \rightarrow \infty} \mathbb{P}_n \circ (\xi, B, W, (\Lambda^{-i}, \beta^k), X[(\Lambda^{-i}, \beta^k)])^{-1}.$$

Hence, the uniform integrability (6.35) and continuity of J of Lemma 5.3.4 thus imply (noting that the required p' -moment bound on μ^x follows easily from the estimates of Lemma 6.1.1)

$$\lim_{k \rightarrow \infty} J_i((\Lambda^{-i}, \beta^k)) = J_i((\Lambda^{-i}, \beta^*)). \quad (6.37)$$

Finally, since Λ is a strong ϵ -Nash equilibrium, it holds for each k that

$$J_i(\Lambda) + \epsilon_i \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n) \text{ strong}} J_i((\Lambda^{-i}, \beta)) \geq J_i((\Lambda^{-i}, \beta^k)).$$

Thus, sending $k \rightarrow \infty$ and applying (6.34) yields

$$J_i(\Lambda) + \epsilon_i \geq J_i((\Lambda^{-i}, \beta^*)) \geq \sup_{\beta \in \mathcal{A}_n(\mathcal{E}_n)} J_i((\Lambda^{-i}, \beta)) - \delta.$$

Sending $\delta \downarrow 0$ shows that Λ is in fact a relaxed ϵ -Nash equilibrium. \square

Step 2:

The proof that every strict ϵ -Nash is a relaxed ϵ -Nash equilibrium follows the same structure; the only difference is that we construct the sequence α^k from β^* a bit differently. First, let $\iota_k : A \rightarrow A$ be a measurable function satisfying $\iota_k(a) = a$ for $|a| \leq k$ and $|\iota_k(a)| \leq k$ for all $a \in A$. Let $\widetilde{\beta}_t^k := \beta_t^* \circ \iota_k^{-1}$, so that $\widetilde{\beta}^k \rightarrow \beta^*$ a.s., and clearly

$$\int_{\{|a| > r\}} |a|^{p'} \widetilde{\beta}_t^k(da) \leq \int_{\{|a| > r\}} |a|^{p'} \beta_t^*(da), \quad r > 0. \quad (6.38)$$

For each k , apply the well known Chattering Lemma [76, Theorem 2.2(b)] to find a sequence of $(\mathcal{F}_t^n)_{t \in [0, T]}$ -progressively measurable A -valued processes $\alpha_t^{k, j}$ such that

$$\tilde{\beta}^k = \lim_{j \rightarrow \infty} dt \delta_{\alpha_t^{k, j}}(da), \text{ a.s.} \quad (6.39)$$

We then find a subsequence j_k such that $\beta^k := dt \delta_{\alpha_t^{k, j_k}}(da)$ converges a.s. to β^* , and (6.36) holds. It follows also from (6.38) and (6.39) that

$$\lim_{r \rightarrow \infty} \sup_{k, j} \mathbb{E}^{\mathbb{P}_n} \int_0^T |\alpha_t^{k, j}|^{p'} 1_{\{|\alpha_t^{k, j}| > r\}} dt = 0,$$

so that (6.37) holds as well. The rest of the proof is as in Step 1.

6.4.2 Proof of Proposition 4.2.1

First, note that when $\sigma_0 \equiv 0$, Lemma 6.2.5 holds true when the common noise B is omitted everywhere it appears. With this in mind, the proof of Proposition 4.2.1 follows exactly Step 1 of the proof of Proposition 3.2.2, except of course with the word “strong” replaced by “very strong,” and with the common noise B removed everywhere it appears.

Chapter 7

Existence, with common noise

This chapter is devoted to the proof of the main existence result, Theorem 3.3.1, and also the proofs of the uniqueness results, Proposition 3.3.4 and Theorem 3.3.5.

The proof of existence follows a strategy that was foreshadowed at the end of Section 2.3.4: First, in Section 7.1 we introduce discretized of the MFG problem, in which the time grid and the range of the common noise are forced into a finite set. This circumvents a crucial lack of continuity which stems from the operation of conditioning a joint distribution, allowing us to prove that strong MFG solutions exist for the discretized problems. Then, in Section 7.2, the discretization is refined, and it is shown that the discretized solutions admit limits which must be weak MFG solutions. This is done first under the additional assumption that the coefficients (b, σ, σ_0) are bounded, and the proof of the original theorem is finally proven by taking weak limits once more.

Quite often, existence of a solution to a mean-field game without common noise is proved by means of Schauder's fixed point theorem. See for instance [24, 32]. Schauder's theorem is then applied on $\mathcal{P}^p(\mathcal{C}^d)$ (with $p = 2$ in most cases), for which compact subsets may be easily described. In the current setting, the presence of the common noise makes things much more complicated. Indeed, an equilibrium, denoted by μ in Definitions 3.1.1, could be viewed in an L^p -space of (equivalence classes of) measurable functions from \mathcal{C}^{m_0} to $\mathcal{P}^p(\mathcal{X})$; this space, however, is much larger, and the difficulty is to identify compact sets which could be stable under the transformations we consider. Moving to a discretized form of the problem, in which the common noise takes only finitely many values, allows us to restrict our attention to the much more pleasant space $\mathcal{P}^p(\mathcal{C}^d)^k$ for finite k .

7.1 Discretized mean field games

This subsection defines precisely the discretized form of the mean field game. We will begin the search for MFG solutions by working under the following additional assumption:

Assumption B. The functions b, σ, σ_0 are uniformly bounded, and the control space A is compact.

Note that assumption **A1.5** together with **B** imply the following bounds on f and g :

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

Note also that \mathcal{V} is compact since A is.

We work with the following canonical spaces, one of which has been defined already:

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

We denote by ξ , B , W , Λ , and X the identity maps on \mathbb{R}^d , \mathcal{C}^{m_0} , \mathcal{C}^m , \mathcal{V} , and \mathcal{C}^d respectively. It is worth now taking a moment to recall the notation and conventions introduced in Section 5.2. For example, with a slight abuse of notation, we will also denote by ξ , B and W the projections from Ω_0 onto \mathbb{R}^d , \mathcal{C}^{m_0} and \mathcal{C}^m respectively, and by ξ , B , W , Λ and X the projections from Ω_f onto \mathbb{R}^d , \mathcal{C}^{m_0} , \mathcal{C}^m , \mathcal{V} and \mathcal{C}^d respectively. The canonical processes B , W , and X generate obvious natural filtrations on Ω_f , denoted \mathbb{F}^B , \mathbb{F}^W , and \mathbb{F}^X . Recall the definition of \mathbb{F}^Λ on \mathcal{V} from (5.3). As in Section 5.2, we will work often with filtrations generated by several canonical processes, without necessarily making explicit mention of the space in consideration. For example, the filtration $\mathcal{F}_t^{\xi, B, W} = \sigma(\xi, B_s, W_s : s \leq t)$ may be defined on Ω_0 or on Ω_f , and this will be clear from context.

7.1.1 Discretization procedure

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

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

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

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

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

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

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

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

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

Measures parameterized by discretized trajectories

The purpose of the discretization procedure described right below is to reduce the complexity of the scenarios upon which an equilibrium μ depends. In a strong MFG solution, the equilibrium μ is B -measurable, and we will now force μ to depend only on the discretization of the canonical process B on \mathcal{C}^{m_0} . This is accomplished by restricting our attention to \mathcal{G}_T^n -measurable functions $\mu : \mathcal{C}^{m_0} \rightarrow \mathcal{P}^p(\mathcal{C}^d)$. In addition, some adaptedness is needed: Let \mathcal{M}_n denote the set of functions $\mu : \mathcal{C}^{m_0} \rightarrow \mathcal{P}^p(\mathcal{C}^d)$ that are \mathcal{G}_T^n -measurable such that for each $t \in [0, T]$ and $C \in \mathcal{F}_t^X$ the map $\beta \mapsto [\mu(\beta)](C)$ is \mathcal{G}_t^n -measurable. In particular, the process $(\mu_t)_{t \in [0, T]}$ is \mathbb{G}^n -adapted and càdlàg (with values in $\mathcal{P}^p(\mathbb{R}^d)$).

Note that any $\mu \in \mathcal{M}_n$ is constant on S for each $S \in \Pi_{2^n}^n$ in the sense that $\beta \mapsto [\mu(\beta)](F)$ (which depends on the discretized trajectory) is constant on S for each Borel subset F of $\mathcal{P}^p(\mathcal{C}^d)$. Endow \mathcal{M}_n with the topology of pointwise convergence, which of course is the same as the topology of uniform convergence since the common domain of each $\mu \in \mathcal{M}_n$ is effectively $\Pi_{2^n}^n$, which is finite. Since $\mathcal{G}_T^n = \sigma(\Pi_{2^n}^n)$ is finite, the space \mathcal{M}_n is homeomorphic to a closed subset of $\mathcal{P}^p(\mathcal{C}^d)^{|\Pi_{2^n}^n|}$. Hence, \mathcal{M}_n is a metrizable closed convex subset of a locally convex topological vector space.

Let us emphasize that here we work with μ taking values in $\mathcal{P}^p(\mathcal{C}^d)$, instead of $\mathcal{P}^p(\mathcal{X})$. For constructing strong MFG solutions as in Definition 3.1.1, we need only to work with the conditional law of X as opposed to the full conditional law of (W, Λ, X) . The reason for this will become more clear as we proceed with the proof.

Control problems

Control problems will be described in terms of measures on $\Omega_0 \times \mathcal{V}$. Recall that \mathcal{W}^k denotes Wiener measure on \mathcal{C}^k for each k , and $\mathcal{W}_\lambda = \lambda \times \mathcal{W}^{m_0} \times \mathcal{W}^m$ was defined in (5.4) to denote the distribution of the given sources of randomness on Ω_0 . The set of *admissible control rules* \mathcal{A}_f is defined to be the set of $Q \in \mathcal{P}(\Omega_0 \times \mathcal{V})$ such that B and W are independent $\mathbb{F}^{\xi, B, W, \Lambda}$ -Wiener processes under Q and $Q \circ (\xi, B, W)^{-1} = \mathcal{W}_\lambda$. Equivalently, $Q \in \mathcal{P}(\Omega_0 \times \mathcal{V})$ is in \mathcal{A}_f if $Q \circ (\xi, B, W)^{-1} = \mathcal{W}_\lambda$ and $(B_t - B_s, W_t - W_s)$ is Q -independent of $\mathcal{F}_s^{\xi, B, W, \Lambda}$ for each $0 \leq s < t \leq T$. Intuitively, this is just the set of “reasonable” joint laws of the control process with the given randomness. It is easy to check that \mathcal{A}_f is closed in the topology of weak convergence.

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

$$\begin{aligned} Y_t = \xi &+ \int_0^t ds \int_A \Lambda_s(da) b(s, Y_s, \mu_s(B), a) \\ &+ \int_0^t \sigma(s, Y_s, \mu_s(B)) dW_s + \int_0^t \sigma_0(s, Y_s, \mu_s(B)) dB_s. \end{aligned} \quad (7.1)$$

Define the law of the solution and the interpolated solution by

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

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

$$\hat{x}_t^n = \frac{2^n}{T} (t - t_i^n) x_{t_i^n} + \frac{2^n}{T} (t_{i+1}^n - t) x_{t_{i+1}^n}, \quad \text{for } t \in [t_i^n, t_{i+1}^n], \quad i = 0, \dots, 2^n - 1. \quad (7.2)$$

The delay ensures that \hat{X}^n is \mathbb{F}^X -adapted. By Lemma 5.3.1 and compactness of A , $\mathcal{R}_f(\mu, Q)$ and $\mathcal{R}_f^n(\mu, Q)$ are in $\mathcal{P}^p(\Omega_f)$. Note that \mathcal{R}_f and \mathcal{R}_f^n are well-defined; by the uniqueness part in Lemma 5.3.1, $\mathcal{R}_f(\mu, Q)$ is the unique element P of $\mathcal{P}(\Omega_f)$ such that $P \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q$ and such that the canonical processes verify the SDE (7.1) under P . Again, as in footnote 1 on page 54, it is no cause for concern that the Q -completion of the canonical filtration $\mathbb{F}^{\xi, B, W, \Lambda}$ may fail to be right-continuous.

The objective of the discretized control problem is as follows. Recall the definition of the reward functional $\Gamma : \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R} \cup \{-\infty\}$ from (3.3), and define the expected reward functional $J_f : \mathcal{M}_n \times \mathcal{P}^p(\Omega_f) \rightarrow \mathbb{R}$ by

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

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

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

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

7.1.2 Strong solutions

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

Theorem 7.1.1. *For each n , there exist $\mu \in \mathcal{M}_n$ and $P \in \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ such that $\mu = P(X \in \cdot | \mathcal{G}_T^n)$ ($P(X \in \cdot | \mathcal{G}_T^n)$ being seen as a map from \mathcal{C}^{m_0} to $\mathcal{P}^p(\mathcal{C}^d)$, constant on each $S \in \Pi_{2^n}^n$.)*

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

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

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

$$G(P) : \mathcal{C}^{m_0} \ni \beta \mapsto \sum_{S \in \Pi_{2^n}^n} P(X \in \cdot | B \in S) 1_S(\beta) = \sum_{S \in \Pi_{2^n}^n} \frac{P(\{X \in \cdot\} \cap \{B \in S\})}{\mathcal{W}^{m_0}(S)} 1_S(\beta). \quad (7.3)$$

The very first step is then to check that $F(\mu) \subset \mathcal{M}_n$ for each $\mu \in \mathcal{M}_n$. The above formula shows that, for $P \in \mathcal{P}_f^p$, $G(P)$ reads as a \mathcal{G}_T^n -measurable function from \mathcal{C}^{m_0} to $\mathcal{P}^p(\mathcal{C}^d)$. To prove that $G(P) \in \mathcal{M}_n$, it suffices to check the adaptedness condition in the definition of \mathcal{M}_n (see Paragraph 7.1.1). For our purpose, we can restrict the proof to the case when X is P a.s. piecewise affine as in (7.2). For each $t \in [0, T]$ and $C \in \mathcal{F}_t^X$, we have that $1_C(X) = 1_{C'}(X)$ P a.s. for some $C' \in \mathcal{F}_{[t]_n}^X$. Now, $\mathcal{G}_T^n = \mathcal{G}_{[t]_n}^n \vee \mathcal{H}$, where $\mathcal{H} \subset \sigma(B_s - B_{[t]_n} : s \in [[t]_n, T])$. Since \mathcal{H} is P -independent of $\mathcal{F}_{[t]_n}^X \vee \mathcal{G}_{[t]_n}^n$, we deduce that, P a.s., $P(X \in C | \mathcal{G}_T^n) = P(X \in C' | \mathcal{G}_{[t]_n}^n)$. Since the unique event in \mathcal{G}_T^n of null probability under P is the empty set, we deduce that the process $(P(X_{[t]_n} \in \cdot | \mathcal{G}_T^n))_{t \in [0, T]}$ is $(\mathcal{G}_t^n)_{t \in [0, T]}$ -adapted. This shows that $G(P) \in \mathcal{M}_n$ and thus $F(\mu) \subset \mathcal{M}_n$.

We will achieve the proof by verifying the hypotheses the Kakutani's fixed point theorem 2.3.1. Namely, we will show that F is upper hemicontinuous with nonempty compact convex values, and we will find a compact convex subset $\mathcal{Q} \subset \mathcal{M}_n$ such that $F(\mu) \subset \mathcal{Q}$ for each $\mu \in \mathcal{Q}$.

Step 1: Continuity of set-valued functions. For the necessary background and definitions of continuity properties for set-valued functions, refer to Section 2.3.1. First, we check the continuity of the function

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

This is straightforward, thanks to the finiteness of the conditioning σ -field. Let $\varphi : \mathcal{C}^d \rightarrow \mathbb{R}$ be continuous with $|\varphi(x)| \leq c(1 + \|x\|_T^p)$ for all $x \in \mathcal{C}^d$, for some $c > 0$. Proposition 2.1.7(3) in Appendix says that it is enough to prove that $\mathbb{E}^{P_k}[\varphi(X) | B \in S] \rightarrow \mathbb{E}^P[\varphi(X) | B \in S]$ whenever $P_k \rightarrow P$ in $\mathcal{P}^p(\Omega_f)$. This follows from Lemma 2.1.9, which implies that the following real-valued function is continuous:

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

Basically, Lemma 2.1.9 handles the discontinuity of the indicator function 1_S together with the fact that φ is not bounded. It follows that the function $G : \mathcal{P}_f^p \rightarrow \mathcal{M}_n$ given by (7.3) is continuous. The set-valued function F is simply the composition of G with the set-valued function $\mu \mapsto \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$. Therefore, to prove that F is upper hemicontinuous, it is sufficient to prove that $\mu \mapsto \mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ is upper hemicontinuous.

Step 2: Analysis of the map $\mu \mapsto \mathcal{R}_f^n(\mu, \mathcal{A}_f)$. Following the first step, the purpose of the second step is to prove continuity of the set-valued function

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

Since the map $\mathcal{C}^d \ni x \mapsto \hat{x}^n \in \mathcal{C}^d$ is continuous (see (7.2)), it suffices to prove continuity with \mathcal{R}_f^n replaced by \mathcal{R}_f . To do so, we prove first that $\mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)$ is relatively compact by showing that each of the sets of marginal measures is relatively compact; see Lemma 2.1.8. Clearly $\{P \circ (\xi, B, W)^{-1} : P \in \mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)\} = \{\mathcal{W}_\lambda\}$ is compact in $\mathcal{P}^p(\Omega_0)$. Since A is compact, so is \mathcal{V} , and thus $\{P \circ \Lambda^{-1} : P \in \mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)\}$ is relatively compact in $\mathcal{P}^p(\mathcal{V})$. Since b , σ , and σ_0 are bounded, it can be shown using Aldous' criterion for tightness (see Proposition 5.3.2 for details) that $\{P \circ X^{-1} : P \in \mathcal{R}_f(\mathcal{M}_n, \mathcal{A}_f)\}$ is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$.

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

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

It follows from the results of Kurtz and Protter [81] that the state SDE (7.1) holds under the limiting measure P , since it holds under each P_k . Since $\mathcal{R}_f(\mu, Q)$ is the unique law on Ω_f under which (ξ, B, W, Λ) has law Q and (ξ, B, W, Λ, X) solves (7.1), we deduce that $P = \mathcal{R}_f(\mu, Q)$. We finally conclude that $\mathcal{R}_f(\cdot, \mathcal{A}_f)$ and thus $\mathcal{R}_f^n(\cdot, \mathcal{A}_f)$ are continuous.

Step 3: Analysis of the map $\mu \mapsto \mathcal{R}_f^{,n}(\mu, \mathcal{A}_f)$.* As a by-product of the previous analysis, we notice that, for each $\mu \in \mathcal{M}_n$, $\mathcal{R}_f(\mu, \mathcal{A}_f)$ is closed and relatively compact and thus compact. By continuity of the map $\mathcal{C}^d \ni x \mapsto \hat{x}^n \in \mathcal{C}^d$ (see (7.2)), $\mathcal{R}_f^n(\mu, \mathcal{A}_f)$ is also compact.

Since f and g are continuous in (x, μ, a) and have p -order growth in (x, μ) , the compactness of A implies that the reward functional Γ is continuous (see Lemma 2.1.9). This implies that the expected reward functional

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

is also continuous. By compactness of $\mathcal{R}_f^n(\mu, \mathcal{A}_f)$ and by continuity of J_f , $\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ is not empty and compact. Moreover, from Berge's theorem 2.3.2, the set-valued function $\mathcal{R}_f^{*,n} : \mathcal{M}_n \rightarrow 2^{\mathcal{P}^p(\Omega_f)}$ is upper hemicontinuous.

Step 4: Convexity of $\mathcal{R}_f^{,n}(\mu, \mathcal{A}_f)$.* We now prove that, for each $\mu \in \mathcal{M}_n$, $\mathcal{R}_f^n(\mu, \mathcal{A}_f)$ is convex. By linearity of the map $\mathcal{C}^d \ni x \mapsto \hat{x}^n \in \mathcal{C}^d$ (defined in (7.2)), it is sufficient to prove that $\mathcal{R}_f(\mu, \mathcal{A}_f)$ is convex. To this end, we observe first that \mathcal{A}_f is convex. Given Q_i , $i = 1, 2$, in \mathcal{A}_f , and $c \in (0, 1)$, we notice that (B, W) is a Wiener process with respect

to $(\mathcal{F}_t^{\xi, B, W, \Lambda})_{t \in [0, T]}$ under $cP^1 + (1 - c)P^2$, where $P^i := \mathcal{R}_f(\mu, Q^i)$ for $i = 1, 2$. (Use the fact that (B, W) is a Wiener process under both P^1 and P^2 .) Moreover, the state equation holds under $cP^1 + (1 - c)P^2$. Since $(cP^1 + (1 - c)P^2) \circ (\xi, B, W, \Lambda)^{-1} = cQ^1 + (1 - c)Q^2$, we deduce that $cP^1 + (1 - c)P^2$ is the unique probability on Ω_f under which (ξ, B, W, Λ) has law $cQ^1 + (1 - c)Q^2$ and (ξ, B, W, Λ, X) solves the state equation. This proves that $cP^1 + (1 - c)P^2 = \mathcal{R}_f(\mu, cQ^1 + (1 - c)Q^2)$.

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

Step 5: Identifying an invariant compact convex set. The last step is to place ourselves in a convex compact subset of \mathcal{M}_n , by first finding a convex compact set $\mathcal{Q}_0 \subset \mathcal{P}^p(\mathcal{C}^d)$ containing $\{P \circ X^{-1} : P \in \mathcal{R}_f^n(\mathcal{M}_n, \mathcal{A}_f)\}$. To this end, note that the boundedness of (b, σ, σ_0) of assumption (B.1) implies that for each smooth $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ with compact support,

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

where D and D^2 denote gradient and Hessian, respectively. Along the lines of the proof of the standard estimate Lemma 5.3.1, using the boundedness of (b, σ, σ_0) it is straightforward to show that

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

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

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

It is clear that \mathcal{Q}_1 is convex and contains $\{P \circ X^{-1} : P \in \mathcal{R}_f^n(\mathcal{M}_n, \mathcal{A}_f)\}$. Using a well known tightness criterion of Stroock and Varadhan [102, Theorem 1.4.6], conditions (1) and (3) together imply that \mathcal{Q}_1 is tight, and the p' -moment bound of (2) then ensures that it is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$ (see Proposition 2.1.7). It is straightforward to check that \mathcal{Q}_1 is in fact closed, and thus it is compact. Next, define

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

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

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

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

It is easily checked that \mathcal{A}_f is a compact set: closedness is straightforward, and, as in the second step, \mathcal{A}_f is relatively compact since A is compact and the (ξ, B, W) -marginal is fixed.

It follows from compactness of \mathcal{A}_f and \mathcal{Q}_2 that \mathcal{Q}_3 is compact (see Lemma 2.1.8). Similarly, it follows from convexity of \mathcal{A}_f and \mathcal{Q}_2 that \mathcal{Q}_3 is convex.

Conclusion of the proof. Finally, define $\mathcal{Q} := G(\mathcal{Q}_3)$. Note that $\mathcal{Q} \subset \mathcal{M}_n$, since we saw at the beginning of the proof that indeed $G(P) \in \mathcal{M}_n$ whenever $P \in \mathcal{P}_f^p$ satisfies $P(X = \hat{X}^n) = 1$. As emphasized by (7.3), G is linear. Hence, \mathcal{Q} is convex and compact since \mathcal{Q}_3 is. Moreover, for each $\mu \in \mathcal{M}_n$, $F(\mu) = G(\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f))$ is convex and compact, since $\mathcal{R}_f^{*,n}(\mu, \mathcal{A}_f)$ is convex and compact (see the third and fourth steps). Since $F(\mu) \subset \mathcal{Q}$ for each $\mu \in \mathcal{Q}$, the proof is complete. \square

7.2 Weak limits of discretized mean field games

We now aim at passing to the limit in the discretized MFG as the time-space grid is refined, the limit being taken in the weak sense. To do so, we show that any sequence of solutions of the discretized MFG is relatively compact, and we characterize the limits. This requires a lot of precaution, the main reason being that measurability properties are not preserved under weak limits. In particular, we cannot generally ensure that in the limit, the conditional measure μ remains B -measurable in the limit. For this reason, the goal is to identify the limits as *weak* MFG solutions. The proof follows the same three-step procedure used in the proof of the main convergence result, Theorem 3.2.4: first we prove relative compactness of the sequence of discretized solutions, then we show the limit is a MFG pre-solution using Lemma 5.2.3, and finally we check the optimality of the limiting control.

Lemma 7.2.1. *Suppose assumption B holds. For each n , by Theorem 7.1.1 we may find $\mu^n \in \mathcal{M}_n$ and $P_n \in \mathcal{R}_f^{*,n}(\mu^n, \mathcal{A}_f)$ such that $\mu^n = P_n(X \in \cdot \mid \mathcal{G}_T^n)$ (both being viewed as random probability measures on \mathcal{C}^d). Define*

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

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

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

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

Proof. Step 1. Write $P_n = \mathcal{R}_f^n(\mu^n, Q_n)$, for some $Q_n \in \mathcal{A}_f$, and define $P'_n := \mathcal{R}_f(\mu^n, Q_n)$. Let

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

so that $\bar{P}_n = \bar{P}'_n \circ (\xi, B, W, \mu, \Lambda, \hat{X}^n)^{-1}$, where \hat{X}^n was defined in (7.2). We first show that \bar{P}'_n are relatively compact in $\mathcal{P}^p(\Omega)$. Clearly $P'_n \circ (B, W)^{-1}$ are relatively compact, and so are $P'_n \circ \Lambda^{-1}$ by compactness of \mathcal{V} . Moreover, $P'_n \circ X^{-1}$ are relatively compact by Proposition 5.3.2. By Corollary 2.1.13, relative compactness of $P'_n \circ (\bar{\mu}^n(B))^{-1}$ follows from that of the mean measures $P'_n \circ (W, \Lambda, X)^{-1}$ and from the p' -moment bound afforded by Lemma 5.3.1, i.e. $\sup_n \mathbb{E}^{P'_n} \|X\|_T^{p'} < \infty$. Precisely, for any point $\chi_0 \in \mathcal{X}$ and a metric $d_{\mathcal{X}}$ on \mathcal{X} compatible

with the topology,

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

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

Step 2. Next, we check that $\bar{P}_n = \bar{P}'_n \circ (\xi, B, W, \mu, \Lambda, \hat{X}^n)^{-1}$ are relatively compact and have the same limits as \bar{P}'_n . This will follow essentially from the fact that $\hat{x}^n \rightarrow x$ as $n \rightarrow \infty$ *uniformly* on compact subsets of \mathcal{C}^d . Indeed, for $t \in [t_i^n, t_{i+1}^n]$, the definition of \hat{x}^n implies

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

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

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

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

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

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

Step 3. It remains to check that any limit point \bar{P} of \bar{P}_n (and thus of \bar{P}'_n) is a MFG pre-solution, and we will do this by checking the four requirements of Lemma 5.2.3. Note first that (B, μ) , ξ , and W are independent under \bar{P} , since $\bar{\mu}^n(B)$ is B -measurable and since B , ξ , and W are independent under P_n . Moreover, (B, W) is an $\mathbb{F}^{\xi, B, W, \mu, \Lambda, X}$ Wiener process (of dimension $m_0 + m$) under \bar{P} since it is under P_n . In particular, $\rho := \bar{P} \circ (\xi, B, W, \mu)^{-1} \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$. Since $(\bar{\mu}^n(B))^x = \mu^n(B)$, the canonical processes $(\xi, B, W, \mu, \Lambda, X)$ verify the state equation 5.7 under \bar{P}'_n for each n . Hence, it follows from the results of Kurtz and Protter [81] that (5.7) holds under the limiting measure \bar{P} as well.

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

(using Lemma 2.1.9 to handle the indicator function)

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

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

7.2.1 Existence of a weak solution under Assumption B

The goal of this section is to prove that the limit points constructed in the previous paragraph are not only MFG pre-solutions but are weak MFG solutions:

Theorem 7.2.2. *Assume that B holds and keep the notation of Lemma 7.2.1. Then, every limit point of $(\bar{P}_n)_{n=1}^{\infty}$ is a weak MFG solution with weak control.*

In order to check the optimality condition at the limit, the idea is to approximate any alternative MFG control by a sequence of particularly well-behaved controls for the discretized game. The crucial technical device is Proposition 2.1.15, or more specifically its corollary Proposition 5.3.7, but the following point is also worth breaking off into its own lemma:

Lemma 7.2.3. *Define $\Pi_n : \mathcal{P}(\Omega) \rightarrow \mathcal{P}(\Omega)$ by*

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

(See (7.2) for the definition of \hat{X}^n .) *If $P_n \rightarrow P$ in $\mathcal{P}^p(\Omega)$, then $\Pi_n(P_n) \rightarrow P$ in $\mathcal{P}^p(\Omega)$.*

Proof. This was essentially already proven in the second step of the proof of Lemma 7.2.1. Note that

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

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

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

We saw in the second step of the proof of Lemma 7.2.1 that $\hat{x}^n \rightarrow x$ as $n \rightarrow \infty$ uniformly on compact subsets of \mathcal{C}^d , and so the proof is complete. □

Proof of Theorem 7.2.2. Let μ^n , $\bar{\mu}^n$, P_n , and \bar{P}_n be as in Lemma 7.2.1, and let \bar{P} denote any limit point. Relabel the subsequence, and assume that \bar{P}_n itself converges. Let $\rho := \bar{P} \circ (\xi, B, W, \mu)^{-1}$. By Lemma 7.2.1, ρ is in $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$, and $\bar{P} \in \mathcal{RA}(\rho)$ is an

MFG pre-solution, and it remains only to show that \bar{P} is optimal, or $\bar{P} \in \mathcal{RA}^*(\rho)$. According to Proposition 5.3.7, it suffices to check that $J(\bar{P}) \geq J(\tilde{P})$ for each \tilde{P} in $\mathcal{RA}_a(\rho)$.

Fix any \tilde{Q} in $\mathcal{A}_a(\rho)$, and note that we may write

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

for some $\mathbb{F}^{\xi, B, W, \mu}$ -adapted continuous function $\varphi : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{V}$ (see Definition 5.3.5). Define $Q_n \in \mathcal{A}_f$ (see Paragraph 7.1.1 for the definition of \mathcal{A}_f) by

$$\tilde{Q}_n := \mathcal{W}_\lambda \circ (\xi, B, W, \varphi(\xi, B, W, \bar{\mu}^n(B)))^{-1}.$$

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

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

where the second equality comes from the definition of \bar{P}_n in Lemma 7.2.1. Since φ is continuous,

$$\begin{aligned} \tilde{Q} &= \lim_{n \rightarrow \infty} \mathcal{W}_\lambda \circ (\xi, B, W, \bar{\mu}^n(B), \varphi(\xi, B, W, \bar{\mu}^n(B)))^{-1} \\ &= \lim_{n \rightarrow \infty} \tilde{Q}_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda)^{-1}. \end{aligned} \tag{7.4}$$

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

$$J_f(\mu^n, \tilde{P}_n) \leq J_f(\mu^n, P_n).$$

Since A is compact, Lemma 5.3.4 assures us that J is continuous, and so

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

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

$$J(\mathcal{R}(\tilde{Q})) = \lim_{n \rightarrow \infty} J_f(\mu^n, \tilde{P}_n), \tag{7.5}$$

and both limits exist. Define Π_n as in Lemma 7.2.3. Applying the basic definition of the different objects, notice that

$$\begin{aligned} \tilde{P}_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1} &= \Pi_n \left(\mathcal{R} \left(\tilde{Q}_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda)^{-1} \right) \right), \\ J_f(\mu^n, \tilde{P}_n) &= J \left(\tilde{P}_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1} \right). \end{aligned}$$

Continuity of \mathcal{R} (see Lemma 5.3.3), Lemma 7.2.3, and (7.4) together yield

$$\begin{aligned} \lim_{n \rightarrow \infty} \tilde{P}_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda, X)^{-1} &= \lim_{n \rightarrow \infty} \Pi_n \left(\mathcal{R} \left(\tilde{Q}_n \circ (\xi, B, W, \bar{\mu}^n(B), \Lambda)^{-1} \right) \right) \\ &= \mathcal{R}(\tilde{Q}). \end{aligned}$$

Finally, (7.5) follows from continuity of J . \square

7.2.2 Unbounded coefficients

Finally, with existence in hand for bounded state coefficients (b, σ, σ_0) and compact control space A , we turn to the general case. The goal is thus to complete the proof of Theorem 3.3.1 under assumptions **A1** and **A2** instead of **B**.

The idea of the proof is to approximate the data (b, σ, σ_0, A) by data satisfying Assumption **B**. Let $(b^n, \sigma^n, \sigma_0^n)$ denote the projection of (b, σ, σ_0) into the ball centered at the origin with radius n in $\mathbb{R}^d \times \mathbb{R}^{d \times m} \times \mathbb{R}^{d \times m_0}$, respectively. Let A_n denote the intersection of A with the ball centered at the origin with radius n . For sufficiently large n_0 , A_n is nonempty and compact for all $n \geq n_0$, and thus we will always assume $n \geq n_0$ in what follows. Note that the data $(b^n, \sigma^n, \sigma_0^n, f, g, A_n)$ satisfy Assumption **B**. Moreover, **A1.4** and **A1.5** hold for each n with the same constants c_1, c_2, c_3 ; this implies that the estimate of Lemma 5.3.1 holds with the same constant c_4 for each set of data, i.e. independent of n .

Define \mathcal{V}_n as before in terms of A_n , but now view it as a subset of \mathcal{V} . That is,

$$\mathcal{V}_n := \{q \in \mathcal{V} : q([0, T] \times A_n^c) = 0\}. \quad (7.6)$$

Then define $\mathcal{A}_n(\rho)$ to be the set of admissible controls with values in A_n :

$$\mathcal{A}_n(\rho) := \{Q \in \mathcal{A}(\rho) : Q(\Lambda \in \mathcal{V}_n) = 1\}. \quad (7.7)$$

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

$$dX_t = \int_A b^n(t, X_t, \mu_t^x, a) \Lambda_t(da) dt + \sigma^n(t, X_t, \mu_t^x) dW_t + \sigma_0^n(t, X_t, \mu_t^x) dB_t. \quad (7.8)$$

Now, of course,

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

By Theorem 7.2.2, there exists for each n an MFG solution corresponding to the n^{th} truncation of the data. In the present notation, this means there exist $\rho_n \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ and $P_n \in \mathcal{R}_n \mathcal{A}_n^*(\rho_n)$ such that

$$\mu = P_n \left((W, \Lambda, X) \in \cdot \mid \mathcal{F}_T^{B, \mu} \right), \quad P^n - a.s. \quad (7.9)$$

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

Relative compactness

We begin with a compactness result:

Lemma 7.2.4. *The measures P_n are relatively compact in $\mathcal{P}^p(\Omega)$. Moreover,*

$$\sup_n \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt < \infty, \quad \sup_n \mathbb{E}^{P_n} \|X\|_T^{p'} < \infty. \quad (7.10)$$

Proof. The proof of the moment estimates is quite similar to Lemma 6.1.2. Noting that the coefficients $(b^n, \sigma^n, \sigma_0^n)$ satisfy A1.1-5 with the same constants (independent of n), the estimates of Lemma 5.3.1 and the fixed point (7.9) together imply

$$\mathbb{E}^{P_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) = \mathbb{E}^{P_n} \|X\|_T^p \leq c_4 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right). \quad (7.11)$$

Fix $a_0 \in A_{n_0}$. Let R_n denote the unique element of $\mathcal{R}_n \mathcal{A}_n(\rho_n)$ satisfying $R_n(\Lambda_t = \delta_{a_0} \text{ for a.e. } t) = 1$. That is R_n is the law of the solution of the state equation arising from the constant control equal to a_0 , in the n^{th} truncation. The first part of Lemma 5.3.1 implies

$$\mathbb{E}^{R_n} \|X\|_T^p \leq c_4 \left(1 + \mathbb{E}^{R_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) + T|a_0|^p \right). \quad (7.12)$$

Noting that $R_n \circ \mu^{-1} = P_n \circ \mu^{-1}$, we combine (7.12) with (7.11) to get

$$\mathbb{E}^{R_n} \|X\|_T^p \leq C_0 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right), \quad (7.13)$$

where $C_0 > 0$ depends only on c_4 , T , and $|a_0|^p$. Use the optimality of P_n , the lower bounds on f and g , and then (7.11) and (7.13) to get

$$\begin{aligned} J(P_n) &\geq J(R_n) \geq -c_2(T+1) \left(1 + \mathbb{E}^{R_n} \|X\|_T^p + \mathbb{E}^{R_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) + |a_0|^{p'} \right) \\ &\geq -C_1 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right), \end{aligned} \quad (7.14)$$

where $C_1 > 0$ depends only on c_2 , c_4 , T , $|a_0|^{p'}$, and C_0 . On the other hand, we may use the upper bounds on f and g along with (7.11) to get

$$\begin{aligned} J(P_n) &\leq c_2(T+1) \left(1 + \mathbb{E}^{P_n} \|X\|_T^p + \mathbb{E}^{P_n} \int_{\mathcal{C}^d} \|y\|_T^p \mu^x(dy) \right) - c_3 \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \\ &\leq C_2 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right) - c_3 \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt, \end{aligned} \quad (7.15)$$

where $C_2 > 0$ depends only on c_2, c_3, c_4 , and T . Combining (7.14) and (7.15) and rearranging, we find two constants, $\kappa_1 \in \mathbb{R}$ and $\kappa_2 > 0$, such that

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

(Note that $\mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty$ for each n .) These constants are independent of n , and the first bound in (7.10) follows from the fact that $p' > p$. Combined with Lemma 5.3.1, this implies the second bound in (7.10).

To show that P_n are relatively compact, we check that each of the sets of marginals is relatively compact; see Lemma 2.1.8. Compactness of $P_n \circ (B, W)^{-1}$ is obvious. Moreover, by (7.10),

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

It follows from Proposition 5.3.2 that $P_n \circ (\Lambda, X)^{-1}$ are relatively compact. The mean measures of $P_n \circ \mu^{-1}$ are $P_n \circ (W, \Lambda, X)^{-1}$, which we have shown are relatively compact. Hence, by Corollary 2.1.13, $P_n \circ \mu^{-1}$ are relatively compact in $\mathcal{P}^p(\mathcal{P}^p(\mathcal{X}))$. \square

Limit points

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

Lemma 7.2.5. *The limit point P is an MFG pre-solution and satisfies*

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

Proof. Fatou's lemma and Lemma 7.2.4 imply the stated inequality. We now check the conditions of Lemma 5.2.3. Since (B, μ) , ξ and W are independent under P_n , the same is true under the limit P , which gives. By passage to the limit, it is well checked that (B, W) is a Wiener process with respect to the filtration $(\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X})_{t \in [0, T]}$ under P (which implies in particular that $\rho \in \mathcal{P}_c^p[(\Omega_0 \times \mathcal{P}^p(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}]$). Moreover, it must also hold $P(X_0 = \xi) = 1$. Let us next check the consistency condition. If $\psi : \mathcal{C}^{m_0} \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathbb{R}$ and $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ are bounded and continuous, we have

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

Thus $\mu = P((W, \Lambda, X) \in \cdot \mid B, \mu)$ a.s.

It remains to check that the state equation is satisfied under P . Define processes $(Z_t^q)_{t \in [0, T]}$ on Ω by

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

Using the growth assumptions on b of [A1.4](#), note that $b(t, y, \nu, a) \neq b^n(t, y, \nu, a)$ if and only if

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

so that

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

By Lemma [7.2.4](#), this tends to zero as $n \rightarrow \infty$. Similarly, $\sigma(t, y, \nu) \neq \sigma^n(t, y, \nu)$ if and only if

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

so that the Burkholder-Davis-Gundy inequality yields

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

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

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

Finally, combine this with the results of Kurtz and Protter [[81](#)] to conclude that the state SDE ([5.7](#)) holds under P . \square

Optimality

It remains to show the limit point P in Lemma [7.2.4](#) is optimal. Let $\rho := P \circ (\xi, B, W, \mu)^{-1}$. Thanks to Proposition [5.3.7](#), it suffices to show that $J(P) \geq J(P')$ for every P' in the dense subclass $\mathcal{RA}_a(\rho)$, where $\mathcal{A}_a(\rho)$ was defined in Definition [5.3.5](#). If we can prove that for each such P' there exist $P'_n \in \mathcal{R}_n \mathcal{A}_n(\rho_n)$ such that $J(P'_n) \rightarrow J(P')$, then, by optimality of P_n for each n , it holds that $J(P_n) \geq J(P'_n)$. Since J is upper semicontinuous by Lemma [5.3.4](#), we then get

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

Since $P' \in \mathcal{RA}_a(\rho)$ was arbitrary, we conclude from Proposition 5.3.7 that P is optimal, or $P \in \mathcal{RA}^*(\rho)$, which completes the proof of Theorem 3.3.1. Hence, it remains to prove the following lemma:

Lemma 7.2.6. *For each $P' \in \mathcal{RA}_a(\rho)$, there exist $P'_n \in \mathcal{RA}_n(\rho_n)$ such that $J(P') = \lim_{n \rightarrow \infty} J(P'_n)$.*

Proof. Find $Q' \in \mathcal{A}_a(\rho)$ such that $P' = \mathcal{R}(Q')$. By definition of $\mathcal{A}_a(\rho)$, there exists a bounded, continuous, adapted function $\varphi : \Omega_0 \times \mathcal{P}^p(\mathcal{X}) \rightarrow \mathcal{V}$ such that

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

Boundedness of φ means precisely that there exists m such that the range of φ is contained in \mathcal{V}_m , which was defined in (7.6). Recalling that $\rho_n = P_n \circ (\xi, B, W, \mu)^{-1}$, define

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

Note that $Q'_n(\Lambda \in \mathcal{V}_m) = 1$. Hence $Q'_n \in \mathcal{A}_n(\rho_n)$ for $n \geq m$. It follows from boundedness and continuity of φ that $Q'_n \rightarrow Q'$. The proof will be complete if we can show

$$\mathcal{R}_n(Q'_n) \rightarrow P', \text{ in } \mathcal{P}^p(\Omega). \quad (7.16)$$

Indeed, since A_m is compact, we may then use the continuity of J (see Lemma 5.3.4) to complete the proof. We prove (7.16) with exactly the same argument as in Lemma 5.3.3: Since $\mathcal{R}_n(Q_n) \circ (\xi, B, W, \mu, \Lambda)^{-1} = Q_n$ are relatively compact in $\mathcal{P}^p(\Omega_0 \times \mathcal{P}^p(\mathcal{X}) \times \mathcal{V})$, we can show (using Proposition 5.3.2) that $\mathcal{R}_n(Q'_n) \circ X^{-1}$ are relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$. Thus $\mathcal{R}_n(Q'_n)$ are relatively compact in $\mathcal{P}^p(\Omega)$. Conclude exactly as in the proof of Lemma 5.3.3 that any limit point must equal P' . \square

7.3 Uniqueness

This section proves both the Yamada-Watanabe type result, Proposition 3.3.4, and the uniqueness result, Theorem 3.3.5. We begin with Proposition 3.3.4, the proof of which will essentially just apply an abstract form of the Yamada-Watanabe theorem that can be found in the papers of Jacod and M  min [69] or Kurtz [80]. This abstract framework is discussed in more detail in Appendix A.3.

Proof of Proposition 3.3.4. Corollary A.1.4 of the appendix tells us that, in the definition of a coupling of MFG solution bases, it is redundant to require that B is a $(\mathcal{G}_t)_{t \in [0, T]}$ -Wiener process. The claim now follows from the abstract form of the Yamada-Watanabe result, quoted in Theorem A.3.1. \square

Let us now turn to the proof of Theorem 3.3.5. Now that more notation is at our disposal, we will instead prove a more general result. Consider

Assumption U'. Assumption U.1-3 hold, along with

(U'.4) For any $\rho \in \mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ the set $\mathcal{A}^*(\rho)$ is a singleton, which means that the maximization problem *in the environment* ρ has a unique (relaxed) solution. See (5.9) for the definition of $\mathcal{A}^*(\rho)$.

Theorem 7.3.1. *Suppose assumptions **A1**, **A2**, and **U** hold. Then there exists a unique in law weak MFG solution with weak control, and it is in fact a strong MFG solution with weak control.*

Remark 7.3.2. In the proof of Proposition 3.1.5 in Section 5.4, it is shown that assumption **(Linear-Convex)** implies (U.4'). Thus, Theorem 7.3.1 and Proposition 3.1.5 together imply Theorem 3.3.5.

Proof of Theorem 7.3.1.

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

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

Let $(\Theta, (\mathcal{G}_t)_{t \in [0, T]}, Q, B, \mu^1, \mu^2)$ be any coupling of γ^1 and γ^2 . In view of Proposition 3.3.4, we will prove $\mu^1 = \mu^2$ a.s. In fact, we may assume without loss of generality that

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

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

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

such that

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

The key point is that K^i is necessarily adapted to the completed filtration $\mathbb{F}^{\xi, B, W, \mu}$, which means that, for each $t \in [0, T]$ and each $\mathcal{F}_t^{\Lambda, X}$ -measurable $\varphi : \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$, the map $\omega \mapsto \int \varphi dK_\omega^i$ is $\mathcal{F}_t^{\xi, B, W, \mu}$ -measurable. The proof is as follows. Since $M\gamma^i$ is a weak MFG solution, the σ -fields $\mathcal{F}_T^{\xi, B, W, \mu}$ and \mathcal{F}_t^Λ are conditionally independent under $M\gamma^i$ given $\mathcal{F}_t^{\xi, B, W, \mu}$. Since the solution of the state equation (3.4) is strong, $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda, X}$ is included in the $M\gamma^i$ -completion of $\mathcal{F}_t^{\xi, B, W, \mu, \Lambda}$, from which we deduce that $\mathcal{F}_T^{\xi, B, W, \mu}$ and $\mathcal{F}_t^{\Lambda, X}$ are conditionally independent under $M\gamma^i$ given $\mathcal{F}_t^{\xi, B, W, \mu}$. Therefore, for each $t \in [0, T]$ and each $\mathcal{F}_t^{\Lambda, X}$ -measurable $\varphi : \mathcal{V} \times \mathcal{C}^d \rightarrow \mathbb{R}$, we have

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

Second step. Define now the extended probability space:

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

endowed with the probability measure:

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

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

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

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

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

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

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

$$P^{1,2} \in \mathcal{RA}(\rho^1) \quad \text{and} \quad P^{2,1} \in \mathcal{RA}(\rho^2). \quad (7.17)$$

Indeed, defining

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

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

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

where the last equality follows from the fact that μ^1 and μ^2 are conditionally independent given (ξ, B, W) . Since (B, W) is an $(\mathcal{F}_t^{\xi, B, W, \mu^2})_{t \in [0, T]}$ -Wiener process and $\int_{\mathcal{V}} \psi_t^2(q) K_{\xi, B, W, \mu^2}^2(dq)$ is $\mathcal{F}_t^{\xi, B, W, \mu^2}$ -measurable by the argument above, the conditioning

in the third line can be replaced by a conditioning by $\mathcal{F}_t^{\xi, B, W}$. Then, using once again the fact that μ^1 and μ^2 are conditionally independent given (ξ, B, W) , the conditioning by $\mathcal{F}_t^{\xi, B, W}$ can be replaced by a conditioning by $\mathcal{F}_t^{\xi, B, W, \mu^1}$, which proves the required property of conditional independence. This shows that $Q^{1,2} \in \mathcal{A}(\rho^1)$ and thus $P^{1,2} \in \mathcal{RA}(\rho^1)$. The proof that $P^{2,1} \in \mathcal{RA}(\rho^2)$ is identical.

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

$$\bar{P}((W, \Lambda^i, X^i) \in \cdot \mid B, \mu^1, \mu^2) = \bar{P}((W, \Lambda^i, X^i) \in \cdot \mid B, \mu^i) = \mu^i, \quad i = 1, 2. \quad (7.18)$$

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

$$P^{1,1} = P^{1,2}, \quad \text{i.e.} \quad \bar{P} \circ (\xi, B, W, \mu^1, \Lambda^1, X^1)^{-1} = \bar{P} \circ (\xi, B, W, \mu^1, \Lambda^2, X^2)^{-1}, \quad (7.19)$$

$$P^{2,2} = P^{2,1}, \quad \text{i.e.} \quad \bar{P} \circ (\xi, B, W, \mu^2, \Lambda^2, X^2)^{-1} = \bar{P} \circ (\xi, B, W, \mu^2, \Lambda^1, X^1)^{-1}. \quad (7.20)$$

It follows that

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

Combined with (7.18), this implies

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

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

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

Thus μ^1 and μ^2 are in fact B -measurable and equal, which is a contradiction. Hence, one of the distributional equalities (7.19) or (7.20) must fail. By optimality of $P^{1,1}$ and $P^{2,2}$ and by (7.17), we have the following two inequalities, and assumption **U'**.4 implies that at least one of them is strict:

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

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

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

where one of the two inequalities is strict. Add these inequalities to get

$$\begin{aligned}
0 &< \mathbb{E}^{\bar{P}} \left[\int_0^T (f_2(t, X_t^2, \mu_t^{2,x}) - f_2(t, X_t^2, \mu_t^{1,x}) + f_2(t, X_t^1, \mu_t^{1,x}) - f_2(t, X_t^1, \mu_t^{2,x})) dt \right] \\
&\quad + \mathbb{E}^{\bar{P}} [g(X_T^2, \mu_T^{2,x}) - g(X_T^2, \mu_T^{1,x}) + g(X_T^1, \mu_T^{1,x}) - g(X_T^1, \mu_T^{2,x})]
\end{aligned} \tag{7.21}$$

Then, conditioning on (B, μ^1, μ^2) inside of (7.21) and applying (7.18) yields

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

This contradicts assumption **U.3**, and so $\mu^1 = \mu^2$ a.s. □

7.4 Counterexamples

In this section, simple examples are presented to illustrate two points. First, we demonstrate why we cannot expect existence of a strong MFG solution at the level of generality allowed by assumption **A1**. Second, by providing an example of a mean field game which fails to admit even a weak solution, we show that the exponent p in both the upper and lower bounds of f and g cannot be relaxed to p' .

7.4.1 Nonexistence of strong solutions

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

$$b(t, x, \mu, a) = a, \quad f(t, x, \mu, a) = a^\top \tilde{f}(t, \bar{\mu}) - \frac{1}{2}|a|^2, \quad \sigma_0(t, x, \mu) = \tilde{\sigma}_0(t, \bar{\mu}),$$

for some bounded continuous functions $\tilde{f} : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\tilde{\sigma}_0 : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m_0}$. Here we have abbreviated $\bar{\mu} := \int_{\mathbb{R}} z \mu(dz)$ for $\mu \in \mathcal{P}^1(\mathbb{R})$. Theorem 3.3.1 and Proposition 3.1.5 ensure that there exists a weak MFG solution P with strong control, $(\tilde{\Omega}, \mathbb{F}, P, B, W, \mu, \Lambda, X)$. In particular, there exists a $\mathbb{F}^{X_0, B, W, \mu}$ -progressive \mathbb{R}^d -valued process α^* such that

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

If α is any bounded $\mathbb{F}^{X_0, B, W, \mu}$ -progressive \mathbb{R}^d -valued processes, then the optimality of α^* implies

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

Hence, $\alpha_t^* = \tilde{f}(t, \bar{\mu}_t^x)$ holds $dt \otimes dP$ -a.e. The state process thus satisfies the SDE

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

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

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

We have only assumed that \tilde{f} and $\tilde{\sigma}_0$ are bounded and continuous. For the punchline, note that uniqueness in distribution may hold for such a SDE even if it fails to possess a strong solution, in which case $\bar{\mu}_t^x$ cannot be adapted to the completion of \mathcal{F}_t^B and the MFG solution cannot be strong. Such cases are not necessarily pathological; see Barlow [10] for examples in dimension $d = 1$ with $\tilde{f} \equiv 0$ and $\tilde{\sigma}_0$ bounded above and below away from zero.

7.4.2 Nonexistence of weak solutions

It is a bit disappointing that assumption A1 excludes linear-quadratic models with objectives which are quadratic in both a and x . That is, we do not allow

$$f(t, x, \mu, a) = -|a|^2 - c|x + c'\bar{\mu}|^2, \quad c, c' \in \mathbb{R},$$

where we have again abbreviated $\bar{\mu} := \int_{\mathbb{R}} z \mu(dz)$ for $\mu \in \mathcal{P}^1(\mathbb{R})$. On the one hand, if $c < 0$ and $|c|$ is large enough, then it may hold for each μ that there exists *no* solution to the corresponding control problem, and obviously non-existence of optimal controls prohibits the existence of MFG solutions. The goal now is to demonstrate that even when f and g are bounded from above, we cannot expect a general existence result if $p' = p$. We are certainly not the first to notice what can go wrong in linear-quadratic mean field games when the constants do not align properly; see, for example, [34, Theorem 3.1]. Of course, the refined analyses of [15, 34] give many positive results on linear-quadratic mean field games, but we simply wish to provide a tractable example of nonexistence to show that this edge case $p' = p$ requires more careful analysis.

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

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

With great foresight, choose $T > 0$, $c \in \mathbb{R}$, and $\lambda \in \mathcal{P}^2(\mathbb{R})$ such that

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

Assumption **A1** and **(Convex)** hold with the one exception that the assumption $p' > p$ is violated. Proposition 3.1.4 still applies (see Remark 5.4.2), and we conclude that if there exists a weak MFG solution (we work here with common noise, but the same argument is valid for solutions without common noise), then there must exist a weak MFG solution with control. Suppose there exists a weak MFG solution $(\tilde{\Omega}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P, B, W, \mu, \Lambda, X)$, and find a \mathbb{F} -progressive real-valued process α^* satisfying $P(\Lambda = dt\delta_{\alpha_t^*}(da)) = 1$ and

$$\mathbb{E} \int_0^1 |\alpha_t^*|^2 dt < \infty,$$

where \mathbb{E} denotes expectation under P . The state equation becomes

$$X_t = X_0 + \int_0^t \alpha_s^* ds + \sigma W_t + \sigma_0 B_t, \quad t \in [0, T]. \quad (7.22)$$

In particular, α^* is the unique minimizer among \mathbb{F} -progressive square-integrable real-valued processes α of

$$J(\alpha) := \mathbb{E} \left[\int_0^T |\alpha_t|^2 dt + (X_T^\alpha + c\bar{\mu}_T)^2 \right],$$

where

$$X_t^\alpha = X_0 + \int_0^t \alpha_s ds + \sigma W_t + \sigma_0 B_t, \quad t \in [0, T].$$

Expand the square

$$(X_T^\alpha + c\bar{\mu}_T)^2 = \left(X_0 + \int_0^T \alpha_t dt + \sigma W_T + \sigma_0 B_T + c\bar{\mu}_T \right)^2$$

and discard the terms which do not involve α to see that minimizing $J(\alpha)$ is equivalent to minimizing

$$\tilde{J}(\alpha) = \mathbb{E} \left[\int_0^T [|\alpha_t|^2 + 2(X_0 + \sigma W_T + \sigma_0 B_T + c\bar{\mu}_T) \alpha_t] dt + \left(\int_0^T \alpha_t dt \right)^2 \right]$$

Since α^* is the unique minimizer, for any other α it holds that

$$\begin{aligned} 0 &= \frac{d}{d\epsilon} \tilde{J}(\alpha^* + \epsilon\alpha) \Big|_{\epsilon=0} \\ &= 2\mathbb{E} \left[\int_0^T [\alpha_t \alpha_t^* + (X_0 + \sigma W_T + \sigma_0 B_T + c\bar{\mu}_T) \alpha_t] dt + \int_0^T \alpha_t dt \int_0^T \alpha_t^* dt \right]. \end{aligned}$$

In particular, if α is deterministic, then

$$0 = \int_0^T \alpha_t \mathbb{E} \left[\alpha_t^* + X_0 + \sigma W_T + \sigma_0 B_T + c\bar{\mu}_T + \int_0^T \alpha_s^* ds \right] dt$$

Since this holds for every deterministic square-integrable α , it follows that

$$0 = \mathbb{E} \left[\alpha_t^* + X_0 + \sigma W_T + \sigma_0 B_T + c\mathbb{E}\bar{\mu}_T + \int_0^T \alpha_s^* ds \right].$$

Noting that $\bar{\mu}_0 = \mathbb{E}X_0$, we get

$$-\mathbb{E}\alpha_t^* = \bar{\mu}_0 + c\mathbb{E}\bar{\mu}_T + \int_0^T \mathbb{E}\alpha_s^* ds.$$

In particular, $\mathbb{E}\alpha_t^*$ is constant in t . Defining $\bar{\alpha} = \mathbb{E}\alpha_t^*$ for all t , we must have

$$\bar{\alpha} = -\frac{\bar{\mu}_0 + c\mathbb{E}\bar{\mu}_T}{1 + T}.$$

Take expectations in (7.22) to get $\mathbb{E}\bar{\mu}_t = \bar{\mu}_0 + \bar{\alpha}t$. But then

$$\begin{aligned} \mathbb{E}\bar{\mu}_T &= \bar{\mu}_0 + \bar{\alpha}T = \bar{\mu}_0 - \frac{\bar{\mu}_0 + c\mathbb{E}\bar{\mu}_T}{1 + T}T \\ &= \frac{\bar{\mu}_0}{1 + T} + \mathbb{E}\bar{\mu}_T, \end{aligned}$$

where in the last line we finally used the particular choice of $c = -(1 + T)/T$. This implies $\bar{\mu}_0 = 0$, which contradicts $\bar{\lambda} \neq 0$ since $\bar{\mu}_0 = \bar{\lambda}$. Hence, for this particular choice of data, there is no solution.

It would be interesting to find additional structural conditions under which existence of a solution holds in the case $p' = p$. This question has been addressed in [32] when $p' = p = 2$, b is linear, σ is constant, f and g are convex in (x, α) and without common noise. Therein, the strategy consists in solving approximating equations, for which the related p is indeed less than 2, and then in passing to the limit. In order to guarantee the tightness of the approximating solutions, the authors introduce a so-called *weak mean-reverting condition*, which reads $\langle x, \partial_x g(0, \delta_x) \rangle \leq c(1 + |x|)$ and $\langle x, \partial_x f(t, 0, \delta_x, 0) \rangle \leq c(1 + |x|)$. This clearly imposes some restriction on the coefficients as, in full generality (when $p = p' = 2$), we should expect $\partial_x g(0, \delta_x)$ and $\partial_x f(t, 0, \delta_x, 0)$ to be of order 1 in x . The *weak mean-reverting condition* provides yields a crucial moment bound on the approximating solutions, which in

turn provides the needed compactness. The same strategy could conceivably be adapted for a proof of existence in the common noise setting, with the help of the discretization argument of Section 7.1, but this is not pursued in this thesis.

Chapter 8

Existence, without common noise

We now develop a framework for studying strong solutions of MFGs without common noise, built on controlled martingale problems. Ultimately, this will culminate in the proof of Theorem 4.6.1. This framework is quite versatile, and the closing Section 8.4 of this chapter will elaborate on some possible extensions. Throughout the chapter, we assume $\sigma_0 \equiv 0$, and strong MFG solutions are understood in the sense of Definition 4.1.1. Moreover, assumption **A1** is in force throughout the chapter.

In many ways, the proof of existence in this chapter is analogous to but simpler than the proof of existence in the previous Chapter 7. However, the approach and some of the notation used in this chapter, where the focus is on *strong* solutions, is rather different from the past three Chapters which dealt primarily with the common noise setting.

It will be useful at times to indicate the dependence of various notation on the choice of data (b, σ, f, g, A) . For example, we may write $\mathcal{V} = \mathcal{V}[A]$ to make it clear which underlying space A is involved in the definition of the relaxed control space \mathcal{V} . This will be useful because the proof of the main existence theorem is done first for bounded coefficients and compact control space A , and the general case is proven by approximation. It will be useful in the latter step to keep track of this dependence. When the choice of data is clear, as it will be before Section 7.2.2, we will typically omit these parameters from the notation. In this chapter, we will work with the canonical space

$$\Xi[A] := \mathcal{V}[A] \times \mathcal{C}^d.$$

The identity maps on $\mathcal{V}[A]$ and \mathcal{C}^d are denoted Λ and X , respectively, and the notational conventions of Chapter 5 are in place. In particular, Λ and X will also denote the projections from $\Xi[A]$ to $\mathcal{V}[A]$ and \mathcal{C}^d , respectively. The filtration on the canonical space $\Xi[A]$ is always $\mathbb{F}^{\Lambda, X} = (\mathcal{F}_t^{\Lambda, X})_{t \in [0, T]}$, given by

$$\mathcal{F}_t^{\Lambda, X} = \sigma(X_s, \Lambda_s : s \leq t) = \sigma(X_s, \Lambda(C) : s \leq t, C \in \mathcal{B}([0, t] \times A)).$$

See Lemma 2.1.14 for some details on the filtration generated by relaxed controls.

8.1 Controlled martingale problems

The goal of this first section is to simplify and reformulate the definition of strong MFG solution. For example, we will work with the consistency condition $\mu = P \circ X^{-1}$, rather than working with the full joint law of (W, Λ, X) . Moreover, in the spirit of martingale problems, by parametrizing our admissible controls merely by joint laws of (X, Λ) , we will avoid keeping track of the driving noise W . Let us begin by describing the controlled martingale problem framework, and then in Lemma 8.1.5 we will connect it with the original Definition 4.1.1 of MFG solution.

The controlled state process will be described by way of its infinitesimal generator. Let $C_0^\infty(\mathbb{R}^d)$ denote the set of infinitely differentiable functions $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ with compact support, and let $D\varphi$ and $D^2\varphi$ denote the gradient and Hessian of φ , respectively. Define the generator $L = L[b, \sigma, A]$ on $\varphi \in C_0^\infty(\mathbb{R}^d)$ by

$$L\varphi(t, x, \mu, a) = b(t, x, \mu, a)^\top D\varphi(x) + \frac{1}{2} \text{Tr} [\sigma \sigma^\top(t, x, \mu, a) D^2\varphi(x)],$$

for $(t, x, \mu, a) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times A$. For $\varphi \in C_0^\infty(\mathbb{R}^d)$ and $\mu \in \mathcal{P}^p(\mathcal{C}^d)$, define $M_t^{\mu, \varphi} = M_t^{\mu, \varphi}[b, \sigma, A] : \Xi \rightarrow \mathbb{R}$ by

$$M_t^{\mu, \varphi}(q, x) := \varphi(x_t) - \int_{[0, t] \times A} q(ds, da) L\varphi(s, x_s, \mu_s, a).$$

Define the objective functional $\Gamma^\mu = \Gamma^\mu[f, g, A] : \Xi \rightarrow \mathbb{R}$ by

$$\Gamma^\mu(q, x) := g(x_T, \mu_T) + \int_{[0, T] \times A} q(dt, da) f(t, x_t, \mu_t, a).$$

Definition 8.1.1. For a measure $\mu \in \mathcal{P}^p(\mathcal{C}^d)$, let $\mathcal{R}[b, \sigma, A](\mu)$ denote the set of $P \in \mathcal{P}(\Xi)$ satisfying the following:

1. $P \circ X_0^{-1} = \lambda$
2. $\mathbb{E}^P \int_0^T \int_A |a|^p \Lambda_t(da) dt < \infty$.
3. For each $\varphi \in C_0^\infty(\mathbb{R}^d)$, the process $M^{\mu, \varphi} = (M_t^{\mu, \varphi})_{t \in [0, T]}$ is a P -martingale.

As before, we abbreviate $\mathcal{R}[b, \sigma, A](\mu)$ to $\mathcal{R}(\mu)$ when the data is clear; this is the set of admissible joint laws of control-state pairs (Λ, X) . Define $J = J[f, g, A] : \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{P}^p(\Xi) \rightarrow \mathbb{R} \cup \{-\infty\}$ and $\mathcal{R}^* = \mathcal{R}^*[b, \sigma, f, g, A] : \mathcal{P}^p(\mathcal{C}^d) \rightarrow 2^{\mathcal{P}(\Xi)}$ by

$$J(\mu, P) := \int_{\Xi} \Gamma^\mu dP,$$

$$\mathcal{R}^*(\mu) := \arg \max_{P \in \mathcal{R}(\mu)} J(\mu, P).$$

Note that when $\mu \in \mathcal{P}^p(\mathcal{C}^d)$ and $P \in \mathcal{P}^p(\Xi)$, the upper bounds on f and g of assumption A1.5 ensure that the positive part of Γ^μ is P -integrable. Hence, J is well-defined. Using the growth

assumptions [A1.4](#) on the coefficients (b, σ) , Lemma [8.2.3](#) below shows that $\mathcal{R}(\mu) \subset \mathcal{P}^p(\Xi)$ for each $\mu \in \mathcal{P}^p(\mathcal{C}^d)$, so that $\mathcal{R}^*(\mu)$ is also well-defined. A priori, $\mathcal{R}^*(\mu)$ may be empty.

Definition 8.1.2. We say $P \in \mathcal{P}^p(\Xi)$ is a *relaxed mean field game (MFG) solution* if $P \in \mathcal{R}^*(P \circ X^{-1})$. We may also refer to the measure $P \circ X^{-1}$ on \mathcal{C}^d itself as a relaxed MFG solution. In other words, a relaxed MFG solution can be seen as a fixed point of the set-valued map

$$\mathcal{P}^p(\mathcal{C}^d) \ni \mu \mapsto \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\} \in 2^{\mathcal{P}^p(\mathcal{C}^d)}.$$

We say a measure $P \in \mathcal{P}(\Xi)$ *corresponds to a strict control* if its \mathcal{V} -marginal is concentrated on the set of strict controls; that is, there exists an $\mathbb{F}^{\Lambda, X}$ -progressively measurable A -valued process α_t on Ξ such that $P(\Lambda = dt\delta_{\alpha_t}) = 1$. On the other hand, P *corresponds to a relaxed Markovian control* if there exists a measurable function $\hat{q} : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$ such that $P(\Lambda = dt\hat{q}(t, X_t)(da)) = 1$. Finally, P *corresponds to a strict Markovian control* if there exists a measurable function $\hat{\alpha} : [0, T] \times \mathbb{R}^d \rightarrow A$ such that $P(\Lambda = dt\delta_{\hat{\alpha}(t, X_t)}(da)) = 1$. If a relaxed MFG solution P corresponds to a relaxed Markovian (resp. strict Markovian) control, then we say P is a *relaxed Markovian MFG solution* (resp. *strict Markovian MFG solution*).

Remark 8.1.3. In fact, the existence theorem for relaxed MFG solutions, Theorem [8.1.6](#), can be extended to include more general objective structures, such as risk-sensitive or mean-variance objectives. See Remark [8.2.7](#). For the sake of simplicity, we stick with the more standard running-terminal objective structure.

It is sometimes more convenient to represent $\mathcal{R}(\mu)$ in terms of stochastic differential equations. To do this in general with control in the volatility requires some use of martingale measures. The few facts about martingale measures we need (namely their stochastic calculus and martingale problems) are summarized in Appendix [A.4](#), the content of which is borrowed from the paper of El Karoui and Méléard [\[74\]](#) and the monograph of Walsh [\[108\]](#). If one is willing to assume σ is uncontrolled, then there is no need for martingale measures, and one may replace $N(da, dt)$ with dW_t in the following proposition. The following proposition is an immediate consequence of Theorem [A.4.3](#), which is itself quoted from Theorem IV-2 of [\[74\]](#).

Proposition 8.1.4. For $\mu \in \mathcal{P}^p(\mathcal{C}^d)$, $\mathcal{R}(\mu)$ is precisely the set of laws $P' \circ (\Lambda, X)^{-1}$, where:

1. $(\Omega', (\mathcal{F}_t)_{t \in [0, T]}, P')$ is a filtered probability space supporting a d -dimensional adapted process X as well as m continuous orthogonal martingale measures $N = (N^1, \dots, N^m)$ on A , each with intensity $\Lambda_t(da)dt$.
2. $P' \circ X_0^{-1} = \lambda$.
3. $\mathbb{E}^{P'} \int_0^T \int_A |a|^p \Lambda_t(da)dt < \infty$.
4. The state equation holds:

$$dX_t = \int_A b(t, X_t, \mu_t, a) \Lambda_t(da)dt + \int_A \sigma(t, X_t, \mu_t, a) N(da, dt). \quad (8.1)$$

Similarly, the subset of $\mathcal{R}(\mu)$ corresponding to strict controls is the set of laws $P' \circ (dt\delta_{\alpha_t}(da), X)^{-1}$,

- 1'. $(\Omega', (\mathcal{F}_t)_{t \in [0, T]}, P')$ is a filtered probability space supporting a d -dimensional adapted process X , an progressively measurable A -valued process α , and a m -dimensional Wiener process W .
- 2'. $P' \circ X_0^{-1} = \lambda$.
- 3'. $\mathbb{E}^{P'} \int_0^T |\alpha_t|^p dt < \infty$.
- 4'. The state equation holds:

$$dX_t = b(t, X_t, \mu_t, \alpha_t)dt + \sigma(t, X_t, \mu_t, \alpha_t)dW_t. \quad (8.2)$$

It is worth noting that on any filtered probability space satisfying (1-3) (resp. (1'-3')), of Proposition 8.1.4, the Lipschitz and growth assumptions of **A1** ensure that there exists a unique strong solution of (8.1) (resp. (8.2)).

Lemma 8.1.5. *Suppose $P \in \mathcal{P}^p(\Xi)$ is a relaxed MFG solution corresponding to a strict control. Then there exists a strong MFG solution (without common noise) with strict control $(\tilde{\Omega}, (\tilde{\mathcal{F}}_t)_{t \in [0, T]}, \tilde{P}, W, \tilde{\mu}, \Lambda, X)$ such that $\tilde{P} \circ (\Lambda, X)^{-1} = \tilde{P}$. Conversely, given any strong MFG solution with strict control $(\tilde{\Omega}, (\tilde{\mathcal{F}}_t)_{t \in [0, T]}, \tilde{P}, W, \tilde{\mu}, \Lambda, X)$, the law $\tilde{P} \circ (\Lambda, X)^{-1}$ is a relaxed MFG solution with strict control.*

Proof. Note that the compatibility requirement (i.e., the conditional independence) required in point (3) of Definition 4.1.1 is vacuous when the measure μ is deterministic, i.e. we are dealing with strong solutions. The proof of the converse is quite straightforward. For the first statement, let P be a relaxed MFG solution with strict control, and apply Proposition 8.1.4 to find a filtered probability space $(\tilde{\Omega}, (\tilde{\mathcal{F}}_t)_{t \in [0, T]}, \tilde{P})$ supporting (W, Λ, X) satisfying properties (1'-4') of Proposition 8.1.4. Now set

$$\tilde{\mu} = \tilde{P} \circ (W, dt\delta_{\alpha_t}(da), X)^{-1}.$$

Then, since P is a relaxed MFG solution, we have

$$\tilde{\mu}^x = P \circ X^{-1} = \mu.$$

Thus the state equation in point (4) of Definition 4.1.1 holds, and the rest of the points of the definition are easy to check by transferring laws from $\tilde{\Omega}$ to the canonical space Ξ . \square

With Lemma 8.1.5 in mind, existence for strong MFG solutions follows from existence of relaxed MFG solutions. The following two theorems are the key results of this chapter, and the main existence Theorem 4.6.1 follows immediately from Corollary 8.1.8. The rest of this section contains the proof of Theorem 8.1.7, while Sections 8.2 and 7.2.2 are devoted to proving Theorem 8.1.6.

Theorem 8.1.6. *Under assumption **A1**, there exists a relaxed MFG solution.*

Theorem 8.1.7. Suppose **A1** holds. Let $\mu \in \mathcal{P}^p(\mathbb{C}^d)$ and $P \in \mathcal{R}(\mu)$. Then there exist a measurable function $\hat{q} : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$ and $P_0 \in \mathcal{R}(\mu)$ such that:

1. $P_0(\Lambda = dt\hat{q}(t, X_t)(da)) = 1$.
2. $J(\mu, P_0) \geq J(\mu, P)$.
3. $P_0 \circ X_t^{-1} = P \circ X_t^{-1}$ for all $t \in [0, T]$.

If also **(Convex)** holds, we can choose \hat{q} of the form $\hat{q}(t, x) = \delta_{\hat{\alpha}(t, x)}$, for some measurable function $\hat{\alpha} : [0, T] \times \mathbb{R}^d \rightarrow A$.

In words, Theorem 8.1.7 says that for any control (1) there exists a Markovian control (2) producing a greater reward (3) without altering the marginal distributions of the state process. When **(Convex)** holds, the new Markovian control can also be taken to be *strict*.

Corollary 8.1.8. Under assumption **A1**, there exists a relaxed Markovian MFG solution. Under assumptions **A1** and **(Convex)**, there exists a strict Markovian MFG solution.

Proof. Let P be a relaxed MFG solution. Let P_0 be as in Theorem 8.1.7. Since $P \in \mathcal{R}^*(\mu)$ and $J(\mu, P_0) \geq J(\mu, P)$, we have $P_0 \in \mathcal{R}^*(\mu)$. Let $\mu^0 := P_0 \circ X^{-1}$. Then $\mu_t^0 = P_0 \circ X_t^{-1} = P \circ X_t^{-1} = \mu_t$ for all $t \in [0, T]$, and it follows that $\mathcal{R}(\mu) = \mathcal{R}(\mu^0)$, $J(\mu^0, \cdot) \equiv J(\mu, \cdot)$, and $\mathcal{R}^*(\mu) = \mathcal{R}^*(\mu^0)$. Thus $P_0 \in \mathcal{R}^*(\mu^0)$. \square

Proof of Theorem 8.1.7. As in [75, Theorem 2.5(a)], we may find \bar{m} and a measurable function $\bar{\sigma} : [0, T] \times \mathbb{R}^d \times \mathcal{P}^p(\mathbb{R}^d) \times \mathcal{P}(A) \rightarrow \mathbb{R}^{d \times \bar{m}}$ such that $\bar{\sigma}(t, x, \mu, q)$ is continuous in (x, μ, q) for each t ,

$$\bar{\sigma} \bar{\sigma}^\top(t, x, \mu, q) = \int_A q(da) \sigma \sigma^\top(t, x, \mu, a),$$

and $\bar{\sigma}(t, x, \mu, \delta_a) = \sigma(t, x, \mu, a)$ for each (t, x, μ, a) ; moreover, we may find a filtered probability space $(\Omega^1, (\mathcal{F}_t^1)_{t \in [0, T]}, Q_1)$ supporting a \bar{m} -dimensional Wiener process W , a \mathbb{R}^d -valued adapted process X^1 , and a progressively measurable $\mathcal{P}(A)$ -valued process Λ_t such that

$$\begin{aligned} dX_t^1 &= \int_A b(t, X_t^1, \mu_t, a) \Lambda_t(da) dt + \bar{\sigma}(t, X_t^1, \mu_t, \Lambda_t) dW_t, \text{ and} \\ P &= Q_1 \circ (dt \Lambda_t(da), X^1)^{-1}. \end{aligned} \tag{8.3}$$

We claim that there exists a (jointly) measurable function $\hat{q} : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$ such that

$$\hat{q}(t, X_t^1) = \mathbb{E}^{Q_1} [\Lambda_t | X_t^1], \quad Q_1 - a.s., \quad a.e. \quad t \in [0, T].$$

More precisely, we mean that for each bounded measurable function $\varphi : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}$, it holds Q_1 -almost surely for a.e. $t \in [0, T]$ that

$$\int_A \varphi(t, X_t, a) \hat{q}(t, X_t^1)(da) = \mathbb{E}^{Q_1} \left[\int_A \varphi(t, X_t^1, a) \Lambda_t(da) \middle| X_t^1 \right]. \tag{8.4}$$

To see this, define a probability measure η on $[0, T] \times \mathbb{R}^d \times A$ by

$$\eta(C) := \frac{1}{T} \mathbb{E}^{Q_1} \left[\int_0^T \int_A 1_C(t, X_t^1, a) \Lambda_t(da) dt \right].$$

Construct \hat{q} by disintegration by writing $\eta(dt, dx, da) = \eta_{1,2}(dt, dx) [\hat{q}(t, x)](da)$, where $\eta_{1,2}$ denotes the $[0, T] \times \mathbb{R}^d$ -marginal of η and $\hat{q} : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{P}(A)$ is measurable. Then, for each bounded measurable $h : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\begin{aligned} & \mathbb{E}^{Q_1} \left[\int_0^T h(t, X_t^1) \int_A \varphi(t, X_t, a) [\hat{q}(t, X_t^1)](da) dt \right] \\ &= T \int_{[0, T] \times \mathbb{R}^d} h(t, x) \int_A \varphi(t, x, a) [\hat{q}(t, x)](da) \eta_{1,2}(dt, dx) \\ &= T \int_{[0, T] \times \mathbb{R}^d \times A} h(t, x) \varphi(t, x, a) \eta(dt, dx, da) \\ &= \mathbb{E}^{Q_1} \left[\int_0^T h(t, X_t^1) \int_A \varphi(t, X_t^1, a) \Lambda_t(da) dt \right]. \end{aligned}$$

This is enough to establish (8.4), thanks to [23, Lemma 5.2].

With \hat{q} in hand, note that

$$\int_A \hat{q}(t, X_t^1)(da) b(t, X_t^1, \mu_t, a) = \mathbb{E}^{Q_1} \left[\int_A \Lambda_t(da) b(t, X_t^1, \mu_t, a) \middle| X_t^1 \right],$$

and

$$\begin{aligned} \bar{\sigma} \bar{\sigma}^\top(t, X_t^1, \mu_t, \hat{q}(t, X_t^1)) &= \int_A \hat{q}(t, X_t^1)(da) \sigma \sigma^\top(t, X_t^1, \mu_t, a) \\ &= \mathbb{E}^{Q_1} \left[\int_A \Lambda_t(da) \sigma \sigma^\top(t, X_t^1, \mu_t, a) \middle| X_t^1 \right]. \end{aligned}$$

The mimicking result of Brunick and Shreve [23, Corollary 3.7] tells us that there exists another filtered probability space $(\Omega^2, (\mathcal{F}_t^2)_{t \in [0, T]}, Q_2)$ supporting a \bar{m} -dimensional Wiener process W^2 and a \mathbb{R}^d -valued adapted process X^2 such that

$$dX_t^2 = \int_A b(t, X_t^2, \mu_t, a) \hat{q}(t, X_t^2)(da) dt + \bar{\sigma}(t, X_t^2, \mu_t, \hat{q}(t, X_t^2)) dW_t^2, \text{ and} \quad (8.5)$$

$$Q_2 \circ (X_t^2)^{-1} = Q_1 \circ (X_t^1)^{-1} = P \circ X_t^{-1}, \text{ for all } t \in [0, T]. \quad (8.6)$$

It follows from Itô's formula that $P_2 := Q_2 \circ (dt\hat{q}(t, X_t^2), X^2)^{-1}$ is in $\mathcal{R}(\mu)$. Finally, compute

$$\begin{aligned} J(\mu, P_2) &= \mathbb{E}^{Q_2} \left[\int_0^T \int_A f(t, X_t^2, \mu_t, a) [\hat{q}(t, X_t^2)](da) dt + g(X_T^2, \mu_T) \right] \\ &= \mathbb{E}^{Q_1} \left[\int_0^T \int_A f(t, X_t^1, \mu_t, a) [\hat{q}(t, X_t^1)](da) dt + g(X_T^1, \mu_T) \right] \\ &= \mathbb{E}^{Q_1} \left[\int_0^T \int_A f(t, X_t^1, \mu_t, a) \Lambda_t(da) dt + g(X_T^1, \mu_T) \right] \\ &= J(\mu, P). \end{aligned}$$

The second line follows from Fubini's theorem and (8.6). The third line follows from Fubini's theorem and the tower property of conditional expectations. This completes the proof of the first part of the theorem; set $P_0 = P_2$, and note that we have in fact proven (2) with *equality*, not *inequality*.

Now suppose assumption **(Convex)** holds. As in [63, Proposition 3.5], $K(t, x, \mu_t)$ is a closed set for each (t, x) . Since $K(t, x, \mu_t)$ is closed and convex with $(b, \sigma\sigma^\top, f)(t, x, \mu_t, a) \in K(t, x, \mu_t)$ for each $a \in A$, we have

$$(b, \sigma\sigma^\top, f)(t, x, \mu_t, \hat{q}(t, x)) = \int_A [\hat{q}(t, x)](da) (b, \sigma\sigma^\top, f)(t, x, \mu_t, a) \in K(t, x, \mu_t),$$

for each $(t, x) \in [0, T] \times \mathbb{R}^d$. Applying the measurable selection result of Proposition 5.4.1 (taking (E, \mathcal{E}) to be $[0, T] \times \mathbb{R}^d$ with its Borel σ -field), there exist measurable functions $\hat{\alpha} : [0, T] \times \mathbb{R}^d \rightarrow A$ and $\hat{z} : [0, T] \times \mathbb{R}^d \rightarrow [0, \infty)$ such that

$$\int_A \hat{q}(t, x)(da) (b, \sigma\sigma^\top, f)(t, x, \mu_t, a) = (b, \sigma\sigma^\top, f)(t, x, \mu_t, \hat{\alpha}(t, x)) - (0, 0, \hat{z}(t, x)), \quad (8.7)$$

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

$$\begin{aligned} b(t, x, \mu_t, \hat{\alpha}(t, x)) &= \int_A \hat{q}(t, x)(da) b(t, x, \mu_t, a), \text{ and} \\ \sigma\sigma^\top(t, x, \mu_t, \hat{\alpha}(t, x)) &= \int_A \hat{q}(t, x)(da) \sigma\sigma^\top(t, x, \mu_t, a) \\ &= \bar{\sigma}\bar{\sigma}^\top(t, x, \mu_t, \hat{q}(t, x)) \end{aligned} \quad (8.8)$$

Now define

$$P_0 := Q_2 \circ (dt\delta_{\hat{\alpha}(t, X_t^2)}(da), X^2)^{-1}.$$

Using the equality (8.8) and Itô's formula, we conclude that P_0 is in $\mathcal{R}(\mu)$. Intuitively, we are exploiting here the fact that the law of the solution of an SDE does not depend on the

choice of square root of the volatility matrix. Finally,

$$\begin{aligned}
J(\mu, P_0) &= \mathbb{E}^{Q_2} \left[\int_0^T f(t, X_t^2, \mu_t, \hat{\alpha}(t, X_t^2)) dt + g(X_T^2, \mu_T) \right] \\
&\geq \mathbb{E}^{Q_2} \left[\int_0^T \int_A f(t, X_t^2, \mu_t, a) \hat{q}(t, X_t^2)(da) dt + g(X_T^2, \mu_T) \right] \\
&= \mathbb{E}^{Q_1} \left[\int_0^T \int_A f(t, X_t^1, \mu_t, a) \hat{q}(t, X_t^1)(da) dt + g(X_T^1, \mu_T) \right] \\
&= \mathbb{E}^{Q_1} \left[\int_0^T \int_A f(t, X_t^1, \mu_t, a) \Lambda_t(da) dt + g(X_T^1, \mu_T) \right] \\
&= J(\mu, P).
\end{aligned}$$

The second line follows from (8.7). The third line comes from Fubini's theorem and $Q_2 \circ (X_t^2)^{-1} = Q_1 \circ (X_t^1)^{-1}$, $t \in [0, T]$. The fourth line follows from Fubini's theorem and the tower property of conditional expectations. The last step is just (8.3). \square

Remark 8.1.9. It should be noted that the control produced by Theorem 8.1.7 is called *Markovian* because of its form $\hat{\alpha}(t, X_t)$, but it does not necessarily render the state process X a Markov process. Although the dynamics appear to be Markovian, the process X is a solution of a potentially ill-posed martingale problem, and it is well-known (see [102, Chapter 12]) that uniqueness in law is required to guarantee the solution is Markovian. It is well known that if the volatility σ is uncontrolled and uniformly nondegenerate, then the martingale problem is indeed well-posed, and thus X is a strong Markov process (a Feller process, in fact).

8.2 Bounded coefficients

In this section, Theorem 8.1.6 is proven in the case that the coefficients are bounded and the control space compact. The general case is proven in Section 7.2.2 by a limiting argument. The general strategy and is the same as the proof of Theorem 3.3.1, existence of solutions *with* common noise. No discretization procedure is needed here, and some care is needed in dealing with the martingale measures that arise because of the controlled volatility, but otherwise many of the arguments are the same. Again we will make use of assumption **B**, which says simply that (b, σ) is bounded and A is compact.

Theorem 8.2.1. *Under assumptions **A1** and **B**, there exists a relaxed MFG solution.*

Remark 8.2.2. In fact, under assumptions **A1** and **B**, we may take $p' = p = 0$ in assumption **A1**, and Theorem 8.2.1 is true with an even simpler proof.

The proof of Theorem 8.2.1 is broken up into several lemmas. First, we state yet another version of a standard estimate.

Lemma 8.2.3. *Assume **A1** holds, and fix $\gamma \in [p, p']$. Then there exists a constant $c_4 > 0$, depending only on γ , $|\lambda|^{p'}$, T , and the constant c_1 of **A1.4** such that for any $\mu \in \mathcal{P}^p(\mathcal{C}^d)$ and*

$P \in \mathcal{R}[b, \sigma, A](\mu)$ we have

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

In particular, $P \in \mathcal{P}^p(\Omega)$. Moreover, if $P \circ X^{-1} = \mu$, then we have

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

Proof. There is a constant $C > 0$ (which will change from line to line) such that

$$\begin{aligned} |X_t|^\gamma &\leq C |X_0|^\gamma + C \int_0^t ds \int_A \Lambda_s(da) |b(s, X_s, \mu_s, a)|^\gamma \\ &\quad + C \left| \int_0^t \int_A \sigma(s, X_s, \mu_s, a) N(da, ds) \right|^\gamma. \end{aligned}$$

The Burkholder-Davis-Gundy inequality and the growth assumptions on b and σ yield

$$\begin{aligned} \mathbb{E}^P \|X\|_t^\gamma &\leq C \mathbb{E}^P \left[|X_0|^\gamma + \int_0^t ds \int_A \Lambda_s(da) \sup_{0 \leq u \leq s} |b(u, X_u, \mu_u, a)|^\gamma \right. \\ &\quad \left. + \left(\int_0^t ds \int_A \Lambda_s(da) \sup_{0 \leq u \leq s} |\sigma(s, X_s, \mu_s, a)|^2 \right)^{\gamma/2} \right] \\ &\leq C \mathbb{E}^P \left[|X_0|^\gamma + \int_0^t ds \int_A \Lambda_s(da) c_1^\gamma \left(1 + \|X\|_s^\gamma + \int_{\mathcal{C}^d} \|x\|_s^\gamma \mu(dx) + |a|^\gamma \right) \right. \\ &\quad \left. + \left(\int_0^t ds \int_A \Lambda_s(da) c_1 \left(1 + \|X\|_s^{p_\sigma} + \left(\int_{\mathcal{C}^d} \|x\|_s^p \mu(dx) \right)^{p_\sigma/p} + |a|^{p_\sigma} \right) \right)^{\gamma/2} \right] \\ &\leq C \mathbb{E}^P \left[1 + |X_0|^\gamma + \int_0^t ds \int_A \Lambda_s(da) (1 + \|X\|_s^\gamma + \int_{\mathcal{C}^d} \|x\|_s^\gamma \mu(dx) + |a|^\gamma) \right] \end{aligned}$$

We used Jensen's inequality for the second line to get $(\int_{\mathcal{C}^d} \|x\|_s^p \mu(dx))^{\gamma/p} \leq \int_{\mathcal{C}^d} \|x\|_s^\gamma \mu(dx)$. If $\gamma \geq 2$, the last line follows from Jensen's inequality and the inequality $|x|^{p_\sigma \gamma/2} \leq 1 + |x|^\gamma$ for $x \in \mathbb{R}$, which holds since $p_\sigma \leq 2$. If $\gamma/2 \leq 1$, the last line follows from the inequality $|x|^{\gamma/2} \leq 1 + |x|$ followed by $|x|^{p_\sigma} \leq 1 + |x|^\gamma$, which holds since $\gamma \geq p_\sigma$. The first claim follows now from Gronwall's inequality. If $P \circ X^{-1} = \mu$, then the above becomes

$$\begin{aligned} \int_{\mathcal{C}^d} \|x\|_t^\gamma \mu(dx) &= \mathbb{E}^P \|X\|_t^\gamma \\ &\leq C \mathbb{E}^P \left[|X_0|^\gamma + \int_0^t \left(1 + 2 \int_{\mathcal{C}^d} \|x\|_s^\gamma \mu(dx) + \int_A |a|^\gamma \Lambda_t(da) \right) ds \right]. \end{aligned}$$

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

We will need a new tightness result as well, and the proof is deferred to Appendix B.

Proposition 8.2.4. *Fix $c > 0$. Let p, p', p_σ and λ be as in assumption A1, and let A be a compact metric space. Let $\mathcal{Q}_c \subset \mathcal{P}(\Xi[A])$ be the set of laws $P \circ (\Lambda, X)^{-1}$ of $\Xi[A]$ -valued random variables (Λ, X) defined on some filtered probability space $(\tilde{\Omega}, (\mathcal{F}_t)_{t \in [0, T]}, P)$ satisfying the SDE*

$$dX_t = \int_A b(t, X_t, a) \Lambda_t(da) dt + \int_A \sigma(t, X_t, a) N(da, dt),$$

where the following hold:

1. $N = (N^1, \dots, N^m)$ are continuous orthogonal martingale measures on A with intensity $\Lambda_t(da)dt$.
2. X is a continuous d -dimensional adapted process with $P \circ X_0^{-1} = \lambda$.
3. $\sigma : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}^{d \times d}$ and $b : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}^d$ are jointly measurable.
4. For each $(t, x, a) \in [0, T] \times \mathbb{R}^d \times A$,

$$\begin{aligned} |b(t, x, a)| &\leq c(1 + |x| + |a|), \\ |\sigma \sigma^\top(t, x, a)| &\leq c(1 + |x|^{p_\sigma} + |a|^{p_\sigma}). \end{aligned}$$

5. Lastly,

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

(That is, we vary over σ , b , and the probability space of definition.) Then \mathcal{Q}_c is relatively compact in $\mathcal{P}(\Xi[A])$.

The proof of Theorem 8.2.1 is again an application of Kakutani's fixed point theorem. See Section 2.3.1 for a statement of this theorem and for the necessary background on set-valued functions.

Lemma 8.2.5. *Under assumptions A1 and B, the range $\mathcal{R}(\mathcal{P}^p(\mathcal{C}^d)) := \{P \in \mathcal{R}(\mu) : \mu \in \mathcal{P}^p(\mathcal{C}^d)\}$ is relatively compact in $\mathcal{P}^p(\Xi)$, and the set-valued function \mathcal{R} is continuous.*

Proof. When A is compact, so is $\mathcal{V} = \mathcal{V}[A]$, and the topology of $\mathcal{P}^p(\mathcal{V})$ is that of weak convergence. Thus $\{P \circ \Lambda^{-1} : P \in \mathcal{R}(\mathcal{P}^p(\mathcal{C}^d))\}$ is relatively compact in $\mathcal{P}^p(\mathcal{V})$. From Proposition 5.3.2 and boundedness of b and σ it follows that $\{P \circ X^{-1} : P \in \mathcal{R}(\mathcal{P}^p(\mathcal{C}^d))\}$ is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$. Thus $\mathcal{R}(\mathcal{P}^p(\mathcal{C}^d))$ is relatively compact in $\mathcal{P}^p(\Xi)$, by Lemma 2.1.8.

To show \mathcal{R} is upper hemicontinuous, it suffices show its graph is closed, since its range is relatively compact. Let $\mu^n \rightarrow \mu$ in $\mathcal{P}^p(\mathcal{C}^d)$ and $P^n \rightarrow P$ in $\mathcal{P}^p(\Omega)$ with $P^n \in \mathcal{R}(\mu^n)$. Clearly $P \circ X_0^{-1} = \lim_n P^n \circ X_0^{-1} = \lambda$. Now fix $s < t$, $\varphi \in C_0^\infty(\mathbb{R}^d)$, and a bounded, continuous, and $\mathcal{F}_s^{\Lambda, X}$ -measurable function $h : \Xi \rightarrow \mathbb{R}$. Note that $(\mu, q, x) \mapsto M_t^{\mu, \varphi}(q, x)$ is bounded and continuous by Lemma 2.1.4. Since $M_t^{\mu^n, \varphi}$ is a P^n -martingale for each n ,

$$\mathbb{E}^P [(M_t^{\mu, \varphi} - M_s^{\mu, \varphi})h] = \lim_{n \rightarrow \infty} \mathbb{E}^{P^n} [(M_t^{\mu^n, \varphi} - M_s^{\mu^n, \varphi})h] = 0.$$

Hence $M_t^{\mu, \varphi}$ is a P -martingale, and so $P \in \mathcal{R}(\mu)$.

To show \mathcal{R} is lower hemicontinuous, let $\mu^n \rightarrow \mu$ and $P \in \mathcal{R}(\mu)$. By Proposition 8.1.4, there exists a filtered probability space $(\Omega', (\mathcal{F}'_t)_{t \in [0, T]}, P')$ supporting a d -dimensional adapted process X as well as m orthogonal martingale measures $N = (N^1, \dots, N^m)$ on A with intensity $\Lambda_t(da)dt$, such that $P' \circ (\Lambda, X)^{-1} = P$ and the state equation (8.1) holds on Ω' . The Lipschitz assumption A1.4 ensures that for each n we may strongly solve the SDE

$$dX_t^n = \int_A b(t, X_t^n, \mu_t^n, a) \Lambda_t(da)dt + \int_A \sigma(t, X_t^n, \mu_t^n, a) N(da, dt), \quad X_0^n = X_0.$$

Let $\gamma \geq 2$. A standard estimate using the Lipschitz assumption and the Burkholder-Davis-Gundy inequality yields a constant $C > 0$ independent of n (which may change from line to line) such that

$$\begin{aligned} \mathbb{E}^{P'} \|X^n - X\|_t^\gamma &\leq C \mathbb{E}^{P'} \int_0^t \int_A |b(s, X_s^n, \mu_s^n, a) - b(s, X_s, \mu_s, a)|^\gamma \Lambda_s(da)ds \\ &\quad + C \mathbb{E}^{P'} \int_0^t \int_A |\sigma(s, X_s^n, \mu_s^n, a) - \sigma(s, X_s, \mu_s, a)|^\gamma \Lambda_s(da)ds \\ &\leq C \mathbb{E}^{P'} \int_0^t \|X^n - X\|_s^\gamma ds \\ &\quad + C \mathbb{E}^{P'} \int_0^t \int_A |b(s, X_s, \mu_s^n, a) - b(s, X_s, \mu_s, a)|^\gamma \Lambda_s(da)ds \\ &\quad + C \mathbb{E}^{P'} \int_0^t \int_A |\sigma(s, X_s, \mu_s^n, a) - \sigma(s, X_s, \mu_s, a)|^\gamma \Lambda_s(da)ds. \end{aligned}$$

Since b and σ are bounded and continuous in μ , Gronwall's inequality and the dominated convergence theorem yield $\mathbb{E}^{P'} \|X^n - X\|_T^\gamma \rightarrow 0$. Let $P^n := P' \circ (\Lambda, X^n)^{-1}$, and check using Itô's formula that $P^n \in \mathcal{R}(\mu^n)$. Choosing $\gamma \geq p$ implies $P^n \rightarrow P$ in $\mathcal{P}^p(\Xi)$, and the proof is complete. \square

Lemma 8.2.6. *Suppose assumption A1 holds. Then J is upper semicontinuous. If also B holds, then J is continuous.*

Proof. It follows from Corollary 2.1.10 and the upper bounds of f and g of assumption A1.5 that $\mathcal{P}^p(\mathcal{C}^d) \times \mathcal{V} \times \mathcal{C}^d \ni (\mu, q, x) \mapsto \Gamma^\mu(q, x)$ is upper semicontinuous. Hence, J is upper semicontinuous. When A is compact, then Γ is continuous by Lemma 2.1.9, and so J is continuous. \square

Proof of Theorem 8.2.1. Since \mathcal{R} is continuous and has nonempty compact values (Lemma 8.2.5), and since J is continuous (Lemma 8.2.6), it follows from Berge's theorem 2.3.2 that \mathcal{R}^* is upper hemicontinuous. It is clear that $\mathcal{R}(\mu)$ is convex for each μ , and it follows from linearity of $P \mapsto J(\mu, P)$ that $\mathcal{R}^*(\mu)$ is convex for each μ . The map $\mathcal{P}^p(\Xi) \ni P \mapsto P \circ X^{-1} \in \mathcal{P}^p(\mathcal{C}^d)$ is linear and continuous, and it follows that the set-valued map

$$\mathcal{P}^p(\mathcal{C}^d) \ni \mu \mapsto F(\mu) := \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\} \subset \mathcal{P}^p(\mathcal{C}^d)$$

is upper hemicontinuous and has nonempty compact convex values. To apply a fixed point theorem, we must place the range $F(\mathcal{P}^p(\mathcal{C}^d))$ inside of a convex compact subset of a nice topological vector space. To this end, define

$$M := \sup \left\{ \int_{\mathcal{C}^d} \|x\|_T^{p'} \mu(dx) : \mu \in F(\mathcal{P}^p(\mathcal{C}^d)) \right\} < \infty.$$

By assumption, b and σ are bounded, so for each $\varphi \in C_0^\infty(\mathbb{R}^d)$ we may find $C_\varphi > 0$ such that

$$|L\varphi(t, x, \mu, a)| \leq C_\varphi,$$

for all (t, x, μ, a) . Moreover, C_φ depends only on $D\varphi$ and $D^2\varphi$. Let \mathcal{Q} denote the set of probability measures P on \mathcal{C}^d satisfying the following:

1. $P \circ X_0^{-1} = \lambda$,
2. $\mathbb{E}^P \|X\|_T^{p'} \leq M$,
3. For each nonnegative $\varphi \in C_0^\infty(\mathbb{R}^d)$, the process $\varphi(X_t) + C_\varphi t$ is a P -submartingale.

It is clear both that \mathcal{Q} is convex and that $F(\mathcal{P}^p(\mathcal{C}^d))$ is contained in \mathcal{Q} . A well known tightness criterion of Stroock and Varadhan [102, Theorem 1.4.6] implies that \mathcal{Q} is tight, and the p' -moment bound (2) ensures that it is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$. In fact, it is straightforward to check that \mathcal{Q} is closed in $\mathcal{P}^p(\mathcal{C}^d)$, and thus it is compact.

Now note that \mathcal{Q} is a subset of the space $\mathcal{M}(\mathcal{C}^d)$ of bounded signed measures on \mathcal{C}^d . When endowed with the topology τ_w of weak convergence, i.e. the topology $\tau_w = \sigma(\mathcal{M}(\mathcal{C}^d), C_b(\mathcal{C}^d))$ induced by bounded continuous functions, $\mathcal{M}(\mathcal{C}^d)$ is a locally convex Hausdorff space. Since \mathcal{Q} is relatively compact in $\mathcal{P}^p(\mathcal{C}^d)$, the p -Wasserstein metric $d_{\mathcal{C}^d, p}$ on $\mathcal{P}^p(\mathcal{C}^d)$ and the topology τ_w on $\mathcal{M}(\mathcal{C}^d)$ both induce the same topology on \mathcal{Q} . Hence, \mathcal{Q} is τ_w -compact. The set-valued function F maps \mathcal{Q} into itself, it is upper hemicontinuous with respect to τ_w (equivalently, its graph is closed), and its values are nonempty, compact, and convex. Existence of a fixed point now follows from Kakutani's theorem 2.3.1. \square

Remark 8.2.7. If one is not interested in Markovian solutions, it is evident from the proofs of this section that a relaxed existence result holds with much more general objective structures, as indicated in Remark 8.1.3. In particular, we only used the fact that $J : \mathcal{P}^p(\mathcal{C}^d) \times \mathcal{P}^p(\Xi) \rightarrow \mathbb{R}$ is continuous and concave in its second argument.

8.3 Unbounded coefficients

This section is devoted to the proof of Theorem 8.1.6, without assuming that b , σ , and A are bounded. The truncation argument is very similar to that of Section 7.2.2, but the details of the proof are quite different since we are working here with martingale problems. Assume throughout this section that assumption A1 holds. Naturally, the idea is to approximate the data (b, σ, A) with truncated versions which satisfy B. Let b_n and σ_n denote the (pointwise) projections of b and σ into the ball centered at the origin with radius n in \mathbb{R}^d and $\mathbb{R}^{d \times m}$,

respectively. Let A_n denote the intersection of A with the ball centered at the origin with radius r_n , where

$$r_n := [n/(2c_1)]^{1/2}. \quad (8.9)$$

(Recall that the constant c_1 comes from assumption **A1.2**.) For sufficiently large n_0 , A_n is nonempty and compact for all $n \geq n_0$, and thus we will always assume $n \geq n_0$ in what follows. Note that the truncated data $(b_n, \sigma_n, f, g, A_n)$ satisfy **B** as well as **A1**. Moreover, **A1.4** and **A1.5** hold *with the same constants* c_1, c_2, c_3 .

By Theorem 8.2.1 there exists for each n a corresponding MFG solution, which is technically a measure on $\Xi[A_n] = \mathcal{V}[A_n] \times \mathcal{C}^d$ but may naturally be viewed as a measure on $\Xi = \Xi[A]$, since $A_n \subset A$. To clarify: Since $A_n \subset A$ there is a natural embedding $\mathcal{V}[A_n] \hookrightarrow \mathcal{V}[A]$. Define $\mathcal{R}_n(\mu)$ to be the set of $P \in \mathcal{P}(\Xi[A])$ satisfying the following:

1. $P(\Lambda([0, T] \times A_n^c) = 0) = 1$.
2. $P \circ X_0^{-1} = \lambda$.
3. $M^{\mu, \varphi}[b_n, \sigma_n, A]$ is a P -martingale for each $\varphi \in C_0^\infty(\mathbb{R}^d)$.

Define

$$\mathcal{R}_n^*(\mu) := \arg \max_{P' \in \mathcal{R}_n(\mu)} J^\mu[f, g, A](P')$$

Then it is clear that $\mathcal{R}_n(\mu)$ (resp. $\mathcal{R}_n^*(\mu)$) is exactly the image of the set $\mathcal{R}[b_n, \sigma_n, f, g, A_n](\mu)$ (resp. $\mathcal{R}^*[b_n, \sigma_n, f, g, A_n](\mu)$) under the natural embedding $\mathcal{P}(\Xi[A_n]) \hookrightarrow \mathcal{P}(\Xi[A])$. Henceforth, we identify these sets. By Theorem 8.2.1, there exist corresponding MFG solutions which may be interpreted as $\mu^n \in \mathcal{P}(\mathcal{C}^d)$ and $P_n \in \mathcal{R}_n^*(\mu^n)$ with $\mu^n = P_n \circ X^{-1}$.

8.3.1 Relative compactness of the approximations

The strategy of the proof is to show that P_n are relatively compact and then characterize the limit points as MFG solutions for the original data $(b, \sigma, f, g, A, \lambda)$. The following Lemma 8.3.1 makes crucial use of the upper bound on f of assumption **A1.5** along with the assumption $p' > p$, in order to establish some uniform integrability of the controls.

Lemma 8.3.1. *The measures P_n are relatively compact in $\mathcal{P}(\Omega[A])$. Moreover,*

$$\sup_n \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt < \infty \quad (8.10)$$

$$\sup_n \mathbb{E}^{P_n} \|X\|_T^{p'} = \sup_n \int_{\mathcal{C}^d} \|x\|_T^{p'} \mu^n(dx) < \infty. \quad (8.11)$$

Proof. Noting that the coefficients (b_n, σ_n) satisfy **A1** with the same constants (independent of n), the second conclusion of Lemma 8.2.3 implies

$$\int_{\mathcal{C}^d} \|x\|_T^p \mu^n(dx) = \mathbb{E}^{P_n} \|X\|_T^p \leq c_4 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right). \quad (8.12)$$

Fix $a_0 \in A_{n_0}$. For $n \geq n_0$, let Q_n denote the unique element of $\mathcal{R}_n(\mu^n)$ satisfying $Q_n(\Lambda_t = \delta_{a_0} \text{ for a.e. } t) = 1$. That is, Q_n is the law of the solution of the state equation arising from the constant control equal to a_0 . The first part of Lemma 8.2.3 implies

$$\mathbb{E}^{Q_n} \|X\|_T^p \leq c_4 \left(1 + \int_{\mathcal{C}^d} \|x\|_T^p \mu^n(dx) + T|a_0|^p \right) \leq C_0 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right), \quad (8.13)$$

where the constant $C_0 > 0$ depends only on c_4 , T , p , and a_0 . Use the optimality of P_n , the lower bounds on f and g , and then (8.12) and (8.13) to get

$$\begin{aligned} J(\mu^n, P_n) &\geq J(\mu^n, Q_n) \geq -c_2(T+1) \left(1 + \mathbb{E}^{Q_n} \|X\|_T^p + \int_{\mathcal{C}^d} \|x\|_T^p \mu^n(dx) + |a_0|^{p'} \right) \\ &\geq -C_1 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right), \end{aligned} \quad (8.14)$$

where $C_1 > 0$ depends only on c_2 , c_4 , T , p , p' , and a_0 . On the other hand, we may use the upper bounds on f and g along with (8.12) to get

$$\begin{aligned} J(\mu^n, P_n) &\leq c_2(T+1) \left(1 + \mathbb{E}^{P_n} \|X\|_T^p + \int_{\mathcal{C}^d} \|x\|_T^p \mu^n(dx) \right) - c_3 \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \\ &\leq C_2 \left(1 + \mathbb{E}^{P_n} \int_0^T \int_A |a|^p \Lambda_t(da) dt \right) - c_3 \mathbb{E}^{P_n} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt, \end{aligned} \quad (8.15)$$

where $C_2 > 0$ depends only on c_2 , c_3 , c_4 , T , p , and a_0 . Combining (8.14) and (8.15) and rearranging, we find two constants, $\kappa_1 \in \mathbb{R}$ and $\kappa_2 > 0$, such that

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

(Note that $\mathbb{E}^{P_n} \int_0^T |\Lambda_t|^p dt < \infty$ for each n .) Crucially, these constants are independent of n . Since $p' > p$, it holds for all sufficiently large x that $x^{p'} + \kappa_1 x^p \geq x^{p'}/2$, and (8.10) follows. Combined with the second conclusion of Lemma 8.2.3, this implies (8.11). Finally, relative compactness of P_n is proven by an application of Aldous' criterion, detailed in Proposition 8.2.4. \square

8.3.2 Limiting state process dynamics

Now that we know P_n are relatively compact, we may fix $P \in \mathcal{P}^p(\Xi[A])$ and a subsequence n_k such that $P_{n_k} \rightarrow P$ in $\mathcal{P}^p(\Xi[A])$. Define $\mu := P \circ X^{-1}$, and note that $\mu^{n_k} \rightarrow \mu$ in $\mathcal{P}^p(\mathcal{C}^d)$. The next result provides a first description of the limit points of P_n and is very much analogous to Lemma 7.2.5, used in the proof of existence with common noise.

Lemma 8.3.2. *The limit point P satisfies $P \in \mathcal{R}[b, \sigma, A](\mu)$, $\mu = P \circ X^{-1}$, and also*

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

Proof. It is immediate that $\mu = \lim_k \mu^{n_k} = \lim_k P_{n_k} \circ X^{-1} = P \circ X^{-1}$, and in particular $P \circ X_0^{-1} = \lambda$. Fatou's lemma and (7.10) imply

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

We must only prove $P \in \mathcal{R}[b, \sigma, A](\mu)$. Fix $\varphi \in C_0^\infty(\mathbb{R}^d)$, and note that $M_t^{\mu^n, \varphi}[b_n, \sigma_n, A_n]$ is a P_n martingale for each n . We must show that $M_t^{\mu, \varphi}[b, \sigma, A]$ is a P -martingale.

Note that $M_t^{\mu^n, \varphi}[b_n, \sigma_n, A]$ may be identified with $M_t^{\mu^n, \varphi}[b_n, \sigma_n, A_n]$, since P_n -almost surely Λ is concentrated on $[0, T] \times A_n$. Letting L_n denote the generator associated to (b_n, σ_n) , we have

$$\begin{aligned} & M_t^{\mu^n, \varphi}[b_n, \sigma_n, A](q, x) - M_t^{\mu^n, \varphi}[b, \sigma, A](q, x) \\ &= \int_0^t ds \int_A \Lambda_s(da) (L_n \varphi(s, x_s, \mu_s^n, a) - L \varphi(s, x_s, \mu_s^n, a)) \\ &= \int_0^t ds \int_A \Lambda_s(da) (b_n(s, x_s, \mu_s^n, a) - b(s, x_s, \mu_s^n, a))^\top D \varphi(x_s) \\ &\quad + \frac{1}{2} \text{Tr} [(\sigma_n \sigma_n^\top(s, x_s, \mu_s^n, a) - \sigma \sigma^\top(s, x_s, \mu_s^n, a)) D^2 \varphi(x_s)]. \end{aligned} \quad (8.16)$$

By construction, $b_n(s, x_s, \mu_s^n, a) \neq b(s, x_s, \mu_s^n, a)$ implies $|b(s, x_s, \mu_s^n, a)| > n$, which by assumption A1.4 implies

$$n < c_1 \left(1 + |x_s| + \left(\int_{\mathbb{R}^d} |z|^p \mu_s^n(dz) \right)^{1/p} + |a| \right). \quad (8.17)$$

Moreover, $|b_n(s, x_s, \mu_s^n, a) - b(s, x_s, \mu_s^n, a)|$ is bounded above by twice the right-hand side of (8.17). For $\gamma \in (0, p']$, denote

$$Z_\gamma := 1 + \|X\|_T^\gamma + \left(\sup_n \int_{\mathbb{C}^d} \|z\|_T^p \mu^n(dz) \right)^{\gamma/p},$$

noting that the supremum is finite by Lemma 8.3.1. Let $C > 0$ bound the first two derivatives of φ . Because of the definition (8.9) of r_n , for $n \geq 2c_1$ and $\gamma \in [0, 2]$ we have

$$\Lambda\{(t, a) : 2c_1 |a|^\gamma > n\} \leq \Lambda\{(t, a) : 2c_1 |a|^2 > n\} = 0, \quad P_n - a.s.$$

Hence

$$\begin{aligned}
& \int_0^t ds \int_A \Lambda_s(da) \left| (b_n(s, X_s, \mu_s^n, a) - b(s, X_s, \mu_s^n, a))^\top D\varphi(X_s) \right| \\
& \leq C \int_0^t ds \int_A \Lambda_s(da) 2c_1(Z_1 + |a|) 1_{\{c_1(Z_1 + |a|) > n\}} \\
& \leq 2Cc_1 \int_0^t ds \int_A \Lambda_s(da) (Z_1 + |a|) (1_{\{2c_1 Z_1 > n\}} + 1_{\{2c_1 |a| > n\}}) \\
& \leq 2Cc_1 \left(tZ_1 + \int_0^t |\Lambda_s| ds \right) 1_{\{2c_1 Z_1 > n\}}, \quad P_n - a.s.
\end{aligned}$$

We have a similar bound for the $\sigma_n \sigma_n^\top - \sigma \sigma^\top$ term:

$$\begin{aligned}
& \int_0^t ds \int_A \Lambda_s(da) \left| \text{Tr} \left[(\sigma_n \sigma_n^\top(s, X_s, \mu_s^n, a) - \sigma \sigma^\top(s, X_s, \mu_s^n, a)) D^2 \varphi(x_s) \right] \right| \\
& \leq C \int_0^t ds \int_A \Lambda_s(da) 2c_1(Z_{p_\sigma} + |a|^{p_\sigma}) 1_{\{c_1(Z_{p_\sigma} + |a|^{p_\sigma}) > n\}} \\
& \leq 2Cc_1 \left(tZ_{p_\sigma} + \int_0^t |\Lambda_s|^{p_\sigma} ds \right) 1_{\{2c_1 Z_{p_\sigma} > n\}}, \quad P_n - a.s.
\end{aligned}$$

Note that (8.11) implies

$$\sup_n \int_{\mathcal{C}^d} \|x\|_T^p \mu^n(dx) < \infty.$$

Returning to (8.16), it holds P_n -a.s. that

$$\begin{aligned}
\left| M_t^{\mu^n, \varphi}[b_n, \sigma_n, A] - M_t^{\mu^n, \varphi}[b, \sigma, A] \right| & \leq 2Cc_1 \left[\left(TZ_1 + \int_0^T \int_A |a| \Lambda_s(da) ds \right) 1_{\{2c_1 Z_1 > n\}} \right. \\
& \quad \left. + \left(TZ_{p_\sigma} + \int_0^T \int_A |a|^{p_\sigma} \Lambda_s(ds) \right) 1_{\{2c_1 Z_{p_\sigma} > n\}} \right]
\end{aligned}$$

for all $t \in [0, T]$. Since $1 \vee p_\sigma \leq p < p'$, and since Lemma 8.3.1 yields

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

we have

$$\lim_{n \rightarrow \infty} \mathbb{E}^{P_n} \left| M_t^{\mu^n, \varphi}[b_n, \sigma_n, A] - M_t^{\mu^n, \varphi}[b, \sigma, A] \right| = 0. \quad (8.18)$$

On the other hand, the map

$$\mathcal{P}^p(\mathcal{C}^d) \times \Xi[A] \ni (\nu, q, x) \mapsto M_t^{\nu, \varphi}[b, \sigma, A](q, x) \in \mathbb{R}$$

is jointly continuous for each t , by Lemma 2.1.9. Fix $s < t$ and a bounded, continuous, and $\mathcal{F}_s^{\Lambda, X}$ -measurable $h : \Xi[A] \rightarrow \mathbb{R}$. Then, since $P_{n_k} \rightarrow P$ in $\mathcal{P}^p(\Xi[A])$, and since $M^{\mu, \varphi}$ grows

with order $1 \vee p_\sigma \leq p$, we have (by Proposition 2.1.7)

$$\begin{aligned} \lim_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}} \left[\left(M_t^{\mu^{n_k}, \varphi}[b, \sigma, A] - M_s^{\mu^{n_k}, \varphi}[b, \sigma, A] \right) h \right] \\ = \mathbb{E}^P \left[(M_t^{\mu, \varphi}[b, \sigma, A] - M_s^{\mu, \varphi}[b, \sigma, A]) h \right]. \end{aligned} \quad (8.19)$$

Since $M_t^{\mu^{n_k}, \varphi}[b_n, \sigma_n, A]$ is a P_n -martingale, combining (8.18) and (8.19) yields

$$\begin{aligned} 0 &= \lim_{k \rightarrow \infty} \mathbb{E}^{P_{n_k}} \left[\left(M_t^{\mu^{n_k}, \varphi}[b_{n_k}, \sigma_{n_k}, A] - M_s^{\mu^{n_k}, \varphi}[b_{n_k}, \sigma_{n_k}, A] \right) h \right] \\ &= \mathbb{E}^P \left[(M_t^{\mu, \varphi}[b, \sigma, A] - M_s^{\mu, \varphi}[b, \sigma, A]) h \right]. \end{aligned}$$

Hence $M_t^{\mu, \varphi}[b, \sigma, A]$ is a P -martingale, and the proof is complete. \square

8.3.3 Optimality of the limiting control

It remains to show that the limit point P is optimal, or $P \in \mathcal{R}^*[b, \sigma, f, g, A](\mu)$. The crucial tool is the following lemma, of course plays the same role here as Lemma 5.3.8 did in the proof of existence with common noise.

Lemma 8.3.3. *For each $P' \in \mathcal{R}[b, \sigma, A](\mu)$ such that $J[f, g, A](\mu, P') > -\infty$, there exists $P'_n \in \mathcal{R}_n(\mu^n)$ such that*

$$J[f, g, A](\mu, P') = \lim_{k \rightarrow \infty} J[f_{n_k}, g_{n_k}, A_{n_k}](\mu^{n_k}, P'_{n_k}). \quad (8.20)$$

Proof. First, the upper bounds of f and g imply

$$\begin{aligned} J[f, g, A](\mu, P') &\leq c_2(T+1) \left(1 + \mathbb{E}^{P'} \|X\|_T^p + \int_{\mathcal{C}^d} \|x\|_T^p \mu(dx) \right) \\ &\quad - c_3 \mathbb{E}^{P'} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt. \end{aligned}$$

Since $\int_{\mathcal{C}^d} \|x\|_T^p \mu(dx) < \infty$ and $\mathbb{E}^{P'} \|X\|_T^p < \infty$, the assumption $J[f, g, A](\mu, P') > -\infty$ implies

$$\mathbb{E}^{P'} \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt < \infty. \quad (8.21)$$

By Proposition 8.1.4, we may find a filtered probability space $(\Omega', (\mathcal{F}'_t)_{t \in [0, T]}, Q')$ supporting a d -dimensional adapted process X as well as m orthogonal martingale measures $N = (N^1, \dots, N^m)$ on A with intensity $\Lambda_t(da)dt$, such that $Q' \circ (\Lambda, X)^{-1} = P'$ and the state equation (8.1) holds. Find a measurable map $\iota_n : A \rightarrow A$ such that $\iota_n(A) \subset A_n$ and $\iota_n(a) = a$ for all $a \in A_n$, so that ι_n converges pointwise to the identity. Let X^n denote the unique strong solution of

$$dX_t^n = \int_A b_n(t, X_t^n, \mu_t^n, \iota_n(a)) \Lambda_t(da) dt + \int_A \sigma_n(t, X_t^n, \mu_t^n, \iota_n(a)) N(da, dt), \quad X_0^n = X_0.$$

Let Λ^n denote the image of Λ under the map $(t, a) \mapsto (t, \iota_n(a))$. Then $Q'(\Lambda^n \in \mathcal{V}[A_n]) = 1$, and it is easy to check that $P'_n := Q' \circ (\Lambda^n, X^n)^{-1}$ is in $\mathcal{R}_n(\mu^n)$. Note that $\Lambda^n \rightarrow \Lambda$ holds Q' -a.s., and we will show also that $\mathbb{E}^{Q'} \|X^{n_k} - X\|_T^p \rightarrow 0$. To this end, note that

$$\begin{aligned} X_t^n - X_t &= \int_{[0,t] \times A} b_n(s, X_s^n, \mu_s^n, \iota_n(a)) - b_n(s, X_s, \mu_s^n, \iota_n(a)) \Lambda_s(da) ds \\ &\quad + \int_{[0,t] \times A} b_n(s, X_s, \mu_s^n, \iota_n(a)) - b(s, X_s, \mu_s, a) \Lambda_s(da) ds \\ &\quad + \int_{[0,t] \times A} \sigma_n(s, X_s^n, \mu_s^n, \iota_n(a)) - \sigma_n(s, X_s, \mu_s^n, \iota_n(a)) N(da, ds) \\ &\quad + \int_{[0,t] \times A} \sigma_n(s, X_s, \mu_s^n, \iota_n(a)) - \sigma(s, X_s, \mu_s, a) N(da, ds). \end{aligned}$$

Use Jensen's inequality, the Lipschitz estimate, and the Burkholder-Davis-Gundy inequality to find a constant $C > 0$, independent of n , such that

$$\mathbb{E}^{Q'} \|X^n - X\|_t^p \leq I_n + II_n + III_n + IV_n,$$

where

$$\begin{aligned} I_n &:= C \mathbb{E}^{Q'} \int_0^t \|X^n - X\|_s^p ds \\ II_n &:= C \mathbb{E}^{Q'} \int_{[0,t] \times A} |b_n(s, X_s, \mu_s^n, \iota_n(a)) - b(s, X_s, \mu_s, a)|^p \Lambda_s(da) ds \\ III_n &:= C \mathbb{E}^{Q'} \left[\left(\int_0^t \|X^n - X\|_s^2 ds \right)^{p/2} \right] \\ IV_n &:= C \mathbb{E}^{Q'} \left[\left(\int_{[0,t] \times A} |\sigma_n(s, X_s, \mu_s^n, \iota_n(a)) - \sigma(s, X_s, \mu_s, a)|^2 \Lambda_s(da) ds \right)^{p/2} \right]. \end{aligned}$$

Recall that $p \geq 1$, by assumption **A1.5**. If $p \geq 2$, note that $III_n \leq CI_n$, for some new constant C . On the other hand, if $p \in [1, 2)$, we use Young's inequality in the form of $|xy| \leq \epsilon^q |x|^q / q + \epsilon^{-q'} |y|^{q'} / q'$, where $q = 2/(2-p)$, $q' = 2/p$, and $\epsilon > 0$. We deduce that

$$\begin{aligned} &\mathbb{E}^{Q'} \left[\left(\int_0^t \|X^n - X\|_s^2 ds \right)^{p/2} \right] \\ &\leq \mathbb{E}^{Q'} \left[\|X^n - X\|_t^{(2-p)p/2} \left(\int_0^t \|X^n - X\|_s^p ds \right)^{p/2} \right] \\ &\leq \epsilon^{\frac{2}{2-p}} \left(1 - \frac{p}{2} \right) \mathbb{E}^{Q'} \|X^n - X\|_t^p + \frac{p}{2\epsilon^{p/2}} \mathbb{E}^{Q'} \int_0^t \|X^n - X\|_s^p ds. \end{aligned}$$

By choosing ϵ sufficiently small, we deduce

$$\mathbb{E}^{Q'} \|X^n - X\|_t^p \leq C(I_n + II_n + IV_n),$$

for a new constant C . Now, once we show that II_{n_k} and IV_{n_k} tend to zero, we may conclude from Gronwall's inequality that $\mathbb{E}^{Q'} \|X^{n_k} - X\|_T^p \rightarrow 0$. Since $|\iota_n(a)| \leq |a|$ for all $a \in A$, there is another constant (again called) C such that

$$\begin{aligned} & \int_0^t ds \int_A \Lambda_s(da) |b_n(t, X_t, \mu_t^n, \iota_n(a)) - b(t, X_t, \mu_t, a)|^p \\ & \leq C \left(1 + \|X\|_T^p + \int_{\mathcal{C}^d} \|x\|_T^p (\mu^n + \mu)(dx) + \int_0^T \int_A |a|^p \Lambda_t(da) dt \right). \end{aligned}$$

A similar bound is available for the term involving σ , using $p_\sigma \leq 2$ and the same argument as in the proof of Lemma 8.2.3. Lemma 8.2.3 implies that the right side above is Q' -integrable, and recall from (8.11) that $\sup_n \int_{\mathcal{C}^d} \|x\|_T^p \mu^n(dx) < \infty$. Since $\mu^{n_k} \rightarrow \mu$ and $\iota_n(a) \rightarrow a$ for each $a \in A$, the dominated convergence theorem shows that II_{n_k} and IV_{n_k} tend to zero.

With the convergence $\mathbb{E}^{Q'} \|X^{n_k} - X\|_T^p \rightarrow 0$ now established, the proof of the Lemma is nearly complete. Note that $\|X^{n_k}\|_T^p$ are uniformly Q' -integrable, and

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

and the latter is Q' -integrable, as in (8.21). The growth assumption A1.5 and (8.11) then imply that the random variables $g(X_T^n, \mu_T^n)$ and

$$\int_0^T dt \int_A \Lambda_t^n(da) f(t, X_t^n, \mu_t^n, a) = \int_0^T dt \int_A \Lambda_t(da) f_n(t, X_t^n, \mu_t^n, \iota_n(a))$$

are uniformly Q' -integrable. Since $\mu^{n_k} \rightarrow \mu$ and $\iota_n(a) \rightarrow a$, continuity of f and g and the convergence of X^{n_k} imply that

$$g(X_T^{n_k}, \mu_T^{n_k}) - g(X_T, \mu_T) + \int_0^T dt \int_A \Lambda_t(da) (f(t, X_t^{n_k}, \mu_t^{n_k}, \iota_{n_k}(a)) - f(t, X_t, \mu_t, a)) \rightarrow 0$$

in Q' -measure. Now (8.20) follows from the dominated convergence theorem, after a transformation to the space $(\Omega', (\mathcal{F}'_t)_{t \in [0, T]}, Q')$:

$$\begin{aligned} J[f, g, A](\mu, P') &= \mathbb{E}^{Q'} \left[g(X_T, \mu_T) + \int_0^T dt \int_A \Lambda_t(da) f(t, X_t, \mu_t, a) \right] \\ &= \lim_{k \rightarrow \infty} \mathbb{E}^{Q'} \left[g(X_T^{n_k}, \mu_T^{n_k}) + \int_0^T dt \int_A \Lambda_t(da) f(t, X_t^{n_k}, \mu_t^{n_k}, \iota_{n_k}(a)) \right] \\ &= \lim_{k \rightarrow \infty} J[f, g, A_{n_k}](\mu^{n_k}, P'_{n_k}). \end{aligned}$$

□

Proof of Theorem 8.1.6. Fix $P' \in \mathcal{R}[b, \sigma, A]$. Find P'_n as in Lemma 8.3.3. Optimality of P_n for each n implies that

$$J[f_n, g_n, A_n](\mu^n, P'_n) \leq J[f_n, g_n, A_n](\mu^n, P_n).$$

Use Lemma 8.3.3 and the upper semicontinuity of J (see Lemma 8.2.6) to get

$$\begin{aligned} J[f, g, A](\mu, P) &\geq \limsup_{k \rightarrow \infty} J[f_{n_k}, g_{n_k}, A_{n_k}](\mu^{n_k}, P_{n_k}) \\ &\geq \lim_{k \rightarrow \infty} J[f_{n_k}, g_{n_k}, A_{n_k}](\mu^{n_k}, P'_{n_k}) \\ &= J[f, g, A](\mu, P'). \end{aligned}$$

Since P' was arbitrary, this implies that P is optimal, or $P \in \mathcal{R}^*[b, \sigma, f, g, A](\mu)$. Since also $P = P \circ X^{-1}$ by Lemma 8.3.2, it follows that P is a relaxed MFG solution. \square

8.4 A general framework

This short stand-alone section outlines a possible generalization of the strategy presented in this section, in an effort to simplify the structure and key ingredients of the argument. Additionally, this is a promising avenue for extensions of the results to jump processes as well as state processes which cannot (as easily) be described by SDEs. The ideas here are very much inspired by the work of Kurtz and Stockbridge [82], who prove the existence of an optimal Markovian control for a large class of stochastic optimal control problems, with Markovian state processes in general state spaces.

Let E and A be compact metric spaces. (Locally compact, complete and separable should be fine, but let us keep the presentation as simple as possible.) Suppose we are given an operator L mapping from a subspace $D(L)$ of $C(E)$ into $C(E \times A \times \mathcal{P}(E))$. We are given also an initial distribution $\lambda \in \mathcal{P}(E)$. This generator L should satisfy some fairly standard assumptions (dense domain, positive maximum principle, etc.) as in [82], but let us keep this discussion informal. Finally, we are given jointly continuous reward functions $f : E \times A \times \mathcal{P}(E) \rightarrow \mathbb{R}$ and $g : E \times \mathcal{P}(E) \rightarrow \mathbb{R}$.

Define \mathcal{V} in the same manner as before, and define also $\mathcal{D} := \mathcal{D}([0, T]; \mathbb{R}^d)$ to be the Skorohod space endowed with the usual metric. Our canonical space will be $\Xi := \mathcal{V} \times \mathcal{D}$, with canonical process (Λ, X) and the natural filtration. For each $\varphi \in D(L)$ and $\mu \in \mathcal{P}(\mathcal{D})$, define a process $M^{\mu, \varphi}$ on Ξ by

$$M_t^{\mu, \varphi}(q, x) := f(x_t) - \int_{[0, t] \times A} L\varphi(x_s, a, \mu_s) q(ds, da).$$

We will say that $P \in \mathcal{P}(\Xi)$ is a solution to the controlled martingale problem for μ if $P \circ X_0^{-1} = \lambda$ and if $M^{\mu, \varphi}$ is a P -martingale for every $\varphi \in D(L)$. Let $\mathcal{R}(\mu)$ denote the set of solutions of the controlled martingale problem for μ . For $\mu \in \mathcal{P}(\mathcal{D})$ and $P \in \mathcal{P}(\Omega)$, let

$$J(\mu, P) := \int_{\Xi} \left[\int_{[0, T] \times A} f(x_s, a, \mu_s) q(ds, da) + g(x_T, \mu_T) \right] P(dq, dx).$$

The set of optimal laws corresponding to μ is denoted

$$\mathcal{R}^*(\mu) := \arg \max_{P \in \mathcal{R}(\mu)} J(\mu, P).$$

A solution of the mean field game is any $\mu \in \mathcal{P}(\mathcal{D})$ satisfying $\mu = P \circ X^{-1}$ for some $P \in \mathcal{R}^*(\mu)$, or equivalently a fixed point of the map $\varphi : \mathcal{P}(\mathcal{D}) \rightarrow 2^{\mathcal{P}(\mathcal{D})}$ given by

$$\varphi(\mu) := \{P \circ X^{-1} : P \in \mathcal{R}^*(\mu)\}.$$

To prove existence of a fixed point, one must show the following:

1. φ maps a compact convex set into itself.
2. The graph of φ is closed.
3. The values $\varphi(\mu)$ are convex compact nonempty.

It is typically straightforward to show that the set $\mathcal{R}(\mu)$ is convex, compact, and nonempty for each μ ; using this along with the continuity and linearity of $J(\mu, \cdot)$ we deduce that (3) holds. Under appropriate boundedness assumptions, the first condition (1) should not be too much of a problem; indeed, we first restricted to bounded coefficients, and then used a truncation argument to remove this restriction. Arguing (2) appears to be the more difficult step, or at least the one requiring the heaviest assumptions. As a consequence of Berge's theorem [5, Theorem 17.31], the following conditions suffice for proving (2):

(2.1) J is jointly continuous.

(2.2) The set-valued map \mathcal{R} has a closed graph.

(2.3) The set-valued map \mathcal{R} is lower hemicontinuous.

Again, (2.1) should be no problem. It is well known that weak limits of solutions of martingale problems remain solutions of martingale problems, under appropriate continuity assumptions, and thus (2.2) should also pose few problems. The technical crux of the argument appears to come from point (2.3), because lower hemicontinuity will require *uniqueness* of the controlled martingale problems in some sense. Recall that lower hemicontinuity means that if $\mu_n \rightarrow \mu$ and $P \in \mathcal{R}(\mu)$, we should be able to find a subsequence μ_{n_k} along with $P_{n_k} \in \mathcal{R}(\mu_{n_k})$ such that $P_{n_k} \rightarrow P$. The natural way to proceed is to construct $P_n \in \mathcal{R}(\mu_n)$ which correspond to *the same choice of control process* as P . If we could formalize and justify the statement *for each choice of control process, there is a unique solution of the controlled martingale problem*, then we could derive this lower semicontinuity. This was quite possible in the setting of stochastic differential equations, which are well understood even when they depend on *random coefficients*. Weak uniqueness results for martingale problems with random coefficients appear to be quite limited in the literature, and this obstructs a satisfying general result in this direction.

As a final remark on this approach, let us note that the Markovian selection argument adapts nicely even in this abstract framework. Theorem 4.1 of Kurtz and Stockbridge [82] plays the role of the Brunick and Shreve [23] mimicking result, and permits the construction

of a *Markovian* (relaxed) MFG solution from a generic relaxed MFG solution, assuming the latter exists.

Appendix A

Elements of stochastic analysis

This section summarizes some supplementary topics in stochastic analysis. First, there is a general discussion of the notion of *compatibility* discussed repeatedly throughout the thesis. Additionally, some results on weak solutions of stochastic differential equations with *random coefficients* are discussed in detail, as they arise naturally in stochastic control problem. These results are often used implicitly throughout the thesis and indeed in many papers, but concise sources can be difficult to locate. The discussion of Yamada-Watanabe-type theorems serves the dual purposes of completing the discussion on SDEs and laying the groundwork for our uniqueness theory for mean field games developed in Section 7.3. Lastly, we compile some results on martingale measures needed in Chapter 8.

A.1 Enlargements of filtered probability spaces

This section elaborates on the notion of *compatibility* defined loosely in Section 3.1. The point of this section is both to demystify this concept and to prepare for the subsequent sections of the appendix, where compatibility plays a role in rigorously defining weak solutions for SDEs with random coefficients. Let us now define a general notion of compatibility, which includes as special cases both of the seemingly distinct notions defined just before Remark 5.2.1.

Definition A.1.1. Let (Ω, \mathcal{F}, P) be a probability space, and let $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ and $\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}$ be two filtrations on Ω , with $\mathcal{F}_t \subset \mathcal{G}_t$ for each t . We say that \mathbb{G} is *compatible* with \mathbb{F} if \mathcal{G}_t is conditionally independent of \mathcal{F}_T given \mathcal{F}_t , for each $t \in [0, T]$.

This is sometimes known as the *H-hypothesis* [22], or that \mathbb{F} is *immersed* in \mathbb{G} . Alternatively, this can be seen as a property of an extension of a filtered probability space, known as a *very good extension* [70] or *natural extension* [75]. The term *compatible* is borrowed from Kurtz [80]. The following Proposition summarizes some equivalent definitions, borrowed from [22, 70, 80].

Lemma A.1.2 (Lemma 2.17 of [70], Theorem 3 of [22]). *Fix a probability space (Ω, \mathcal{F}, P) with two filtrations $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ and $\mathbb{G} = (\mathcal{G}_t)_{t \in [0, T]}$ satisfying $\mathcal{F}_t \subset \mathcal{G}_t$ for each t . The following are equivalent:*

1. \mathbb{G} is compatible with \mathbb{F} .
2. Every \mathbb{F} -martingale is a \mathbb{G} -martingale.
3. For each $t \in [0, T]$ and each bounded \mathcal{F}_T -measurable random variable Z ,

$$\mathbb{E}[Z|\mathcal{G}_t] = \mathbb{E}^Q[Z|\mathcal{F}_t], \text{ a.s.}$$

4. For each $A \in \mathcal{G}_t$ and $t \in [0, T]$, the random variable $P(A|\mathcal{F}_T)$ is measurable with respect to the P -completion of \mathcal{F}_t .

In the case that \mathbb{F} is generated by a \mathbb{R}^k -valued process $(Y_t)_{t \in [0, T]}$ with independent increments, properties (1-4) are equivalent to

5. For each $0 \leq s < t \leq T$, the increment $Y_t - Y_s$ is independent of \mathcal{G}_s .

Let us explain how this notion of compatibility relates to those of Section 5.2, discussed before Remark 5.2.1. The notation $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$ and its variants used in that section were instances of the following, more systematic definition. Fix two filtered spaces $(\Omega^i, \mathbb{F}^i = (\mathcal{F}_t^i)_{t \in [0, T]})$. Given a probability measure P^1 on $(\Omega^1, \mathcal{F}_T^1)$, let $\mathcal{P}_c[(\Omega^1, P^1) \rightsquigarrow \Omega^2]$ denote the set of probability measures P on $\Omega^1 \times \Omega^2$ satisfying:

1. The first marginal of P is P^1 , i.e. $P^1 = P(\cdot \times \Omega^2)$.
2. Under P , the filtration $(\mathcal{F}_t^1 \otimes \mathcal{F}_t^2)_{t \in [0, T]}$ is compatible with \mathbb{F}^1 .

If Ω^1 and Ω^2 are separable metric spaces, and if $p \geq 1$, we may also write

$$\mathcal{P}_c^p[(\Omega^1, P^1) \rightsquigarrow \Omega^2] := \mathcal{P}_c[(\Omega^1, P^1) \rightsquigarrow \Omega^2] \cap \mathcal{P}^p(\Omega^1 \times \Omega^2).$$

It is clear now that this notation is consistent with the definition of $\mathcal{P}_c^p[(\Omega_0 \times \mathcal{P}^p(\mathcal{X}), \rho) \rightsquigarrow \mathcal{V}]$ provided on Page 79. This is also consistent with the definition of $\mathcal{P}_c^p[(\Omega_0, \mathcal{W}_\lambda) \rightsquigarrow \mathcal{P}^p(\mathcal{X})]$, thanks to part (5) of Lemma A.1.2. In this product space setting, the following lemma is useful even though it essentially just rewrites the fourth condition of Lemma A.1.2, and its corollary below will be useful in Section A.3.

Lemma A.1.3 (Lemma 2.17 of [70]). *Suppose $Q \in \mathcal{P}(\Omega^1 \times \Omega^2)$ is of the form $Q(d\omega_1, d\omega_2) = P(d\omega_1)K_{\omega_1}(d\omega_2)$, for some $P \in \mathcal{P}(\Omega^1)$ and some kernel K . Then Q is in $\mathcal{P}_c[(\Omega^1, P) \rightsquigarrow \Omega^2]$ if and only if for each $t \geq 0$ and $C \in \mathcal{F}_t^2$ the map $\omega_1 \mapsto K_{\omega_1}(C)$ is measurable with respect to the P -completion of \mathcal{F}_t^1 .*

Corollary A.1.4 (c.f. Lemma 2.10 of [80]). *Let $(\Omega, (\mathcal{F}_t)_{t \in [0, T]})$ and $(\Omega^i, (\mathcal{F}_t^i)_{t \in [0, T]})$ be filtered spaces for $i = 1, 2$. Suppose $P^i \in \mathcal{P}_c[(\Omega, P) \rightsquigarrow \Omega^i]$ is of the form $P^i(d\omega, d\omega^i) = P(d\omega)K_\omega^i(d\omega^i)$, for $i = 1, 2$. Define*

$$Q(d\omega, d\omega^1, d\omega^2) := P(d\omega)K_\omega^1(d\omega^1)K_\omega^2(d\omega^2).$$

Then Q is in $\mathcal{P}_c[(\Omega, P) \rightsquigarrow \Omega^1 \times \Omega^2]$.

A.2 Existence and uniqueness for SDEs

The basic setup is the following. We are given a probability space (Ω, \mathcal{F}, P) with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ and an m -dimensional \mathbb{F} -Brownian motion W . The basic object of study will be the following stochastic differential equation:

$$dX_t = b(t, X_t)dt + \sigma(t, X_t)dW_t, \quad X_0 = \xi. \quad (\text{A.1})$$

The coefficients b and σ map are specified by

$$\begin{aligned} b &: [0, T] \times \Omega \times \mathbb{R}^d \rightarrow \mathbb{R}^d, \\ \sigma &: [0, T] \times \Omega \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m}. \end{aligned}$$

As usual, we suppress the argument in Ω from the notation. We assume throughout that b and σ satisfy the following hypotheses:

Assumption (SDE). 1. For each $x \in \mathbb{R}$, the process $(b(t, x))_{t \in [0, T]}$ is progressively measurable.

2. There exists $C > 0$ such that, for each $x, y \in \mathbb{R}^d$ and $(t, \omega) \in [0, T] \times \Omega$,

$$|b(t, \omega, x) - b(t, \omega, y)| + |\sigma(t, \omega, x) - \sigma(t, \omega, y)| \leq C|x - y|,$$

3. For some $p \geq 1$, we have

$$\mathbb{E} \left[|\xi|^p + \int_0^T |b(t, 0)|^p dt \right] < \infty.$$

4. ξ is \mathcal{F}_0 -measurable.

The following existence and uniqueness result is well known, at least when $p = 2$. Note that we have not assumed the filtration to be complete or right-continuous, and we refer to Stroock and Varadhan [102] for a careful treatment of stochastic integration without the usual hypotheses.

Theorem A.2.1. *Under assumptions **(SDE)**, there exists a unique continuous \mathbb{F} -adapted process X on Ω satisfying the SDE (A.1) and*

$$\mathbb{E} \left[\sup_{0 \leq t \leq T} |X_t|^p \right] < \infty.$$

The uniqueness in the above theorem is *strong* sense, in the sense that it is specific to the probability space Ω . Frequently, we will work on an enlarged probability space, in which the meaning of a “unique” solution of (A.1) must be clarified. To this end, we must first develop some terminology. The ultimate goal is to define *uniqueness in law* and *pathwise uniqueness* for (A.1), state a suitable form of the celebrated Yamada-Watanabe theorem, and conclude that essentially all kinds of uniqueness hold for the SDE (A.1).

First, let us recall some notational conventions introduced Section 5.2. Any function φ defined on Ω naturally extends to any product space $\Omega \times \Omega'$ by setting $\varphi(\omega, \omega') = \varphi(\omega)$. Similarly, any σ -field \mathcal{G} on Ω is extended naturally to the σ -field $\mathcal{G} \otimes \{\emptyset, \Omega'\}$ on $\Omega \times \Omega'$. Similarly, any filtration on Ω extends canonically to a product space as well. The path space for the solution process will be denoted $\mathcal{C}^d := C([0, T]; \mathbb{R}^d)$, that is the space of continuous functions from $[0, T]$ to \mathbb{R}^d . The path space is endowed with the supremum norm and the corresponding Borel σ -field. Let X denote the canonical process on \mathcal{C}^d , that is $X(x) = x$ and $X_t(x) = x(t)$ for $x \in \mathcal{C}^d$. We will describe a solution of the SDE (A.1) by specifying a compatible probability measure on $\Omega \times \mathcal{C}^d$.

Definition A.2.2. A *solution measure* is an element of $\mathcal{P}_c[(\Omega, P) \rightsquigarrow \mathcal{C}^d]$ (see Section A.1) under which the SDE (A.1) holds almost surely. We say a solution measure \bar{P} is *strong* if it is of the form $\bar{P} = P \circ \varphi^{-1}$, where $\varphi : \Omega \rightarrow \bar{\Omega}$ is given by $\varphi(\omega) = (\omega, \bar{X}(\omega))$ for some continuous \mathbb{F} -adapted process \bar{X} defined on Ω .

The following uniqueness theorem is proven in the next section, after we discuss a suitable form of the Yamada-Watanabe theorem. It also follows from the results of Jacod & Mémmin [70], but again we include the proof both for the sake of completeness and also to illustrate how to apply the abstract form of the Yamada-Watanabe result given in Theorem A.3.1.

Theorem A.2.3. Assume Ω is a Polish space. Under assumption (SDE), there exists a unique solution measure, and it is strong.

A.3 Yamada-Watanabe theorems

The ideas of this section are largely borrowed from a paper of Kurtz [80] and Jacod and Mémmin [69, Theorem 3.20], which distill the essence of the Yamada-Watanabe theorem. We extend the results slightly in a way which will be useful in establishing pathwise uniqueness for mean field games with common noise. Because of this modest extension of the results, we will use slightly different terminology from Kurtz.

Consider the following simple model. Fix two Polish spaces E and F . We are given $\mu \in \mathcal{P}(E)$ and a subset \mathcal{S} of $\mathcal{P}(E \times F)$ satisfying $\gamma(\cdot \times F) = \mu$ for each $\gamma \in \mathcal{S}$. That is, μ is the first marginal of each element of \mathcal{S} . We think of \mathcal{S} as corresponding to some set of constraints. A *weak solution* is any probability measure in \mathcal{S} . A *strong solution* is any solution of the form $\gamma(dx, dy) = \mu(dx)\delta_{F(x)}(dy)$, where $F : E \rightarrow F$ is Borel measurable. We say *uniqueness in law* holds if \mathcal{S} has at most one element.

We will define two slightly different pathwise uniqueness notions, which turn out to be often equivalent. First, *pathwise uniqueness* holds if whenever X is an E -valued random variable and Y_1, Y_2 are F valued random variables defined on a common probability space (Ω, \mathcal{F}, P) , then $P \circ (X, Y_1)^{-1} \in \mathcal{S}$ and $P \circ (X, Y_2)^{-1} \in \mathcal{S}$ together imply $P(Y_1 = Y_2) = 1$. On the other hand, we say *independent pathwise uniqueness* holds if whenever X is an E -valued random variable and Y_1, Y_2 are F valued random variables defined on a common probability space (Ω, \mathcal{F}, P) , and Y_1 and Y_2 are conditionally independent given X , then again $P \circ (X, Y_1)^{-1} \in \mathcal{S}$ and $P \circ (X, Y_2)^{-1} \in \mathcal{S}$ together imply $P(Y_1 = Y_2) = 1$. The only difference between these two notions of uniqueness is the conditional independence restriction.

The following theorem is an abstract form of the Yamada-Watanabe result. The proof can be found in either of the references provided, but since this is a slight extension of these results we include the proof for the sake of completeness.

Theorem A.3.1 (Theorem 1.5 of [80], Theorem 3.20 of [69]). *The following are equivalent:*

1. $\mathcal{S} \neq \emptyset$ and pathwise uniqueness holds.
2. $\mathcal{S} \neq \emptyset$ and independent pathwise uniqueness holds.
3. There exists a strong solution, and uniqueness in law holds.

Proof. Clearly (1) implies (2). Let us check next that (3) implies (1). Assumption (3) immediately implies that \mathcal{S} is a singleton, and its unique element is of the form $\gamma(dx, dy) = \mu(dx)\delta_{F(x)}(dy)$ for a measurable function $F : E \rightarrow F$. Suppose we are given a probability space (Ω, \mathcal{F}, P) supporting an E -valued random variable X and two F valued random variables Y_1, Y_2 , with $P \circ (X, Y_1)^{-1} \in \mathcal{S}$ and $P \circ (X, Y_2)^{-1} \in \mathcal{S}$. Since $\mathcal{S} = \{\gamma\}$, we have

$$P \circ (X, Y_1)^{-1} = P \circ (X, Y_2)^{-1} = \gamma.$$

Thus $Y_1 = F(X) = Y_2$ almost surely, and so pathwise uniqueness holds.

Lastly, we show that (2) implies (3). Fix $\gamma_1, \gamma_2 \in \mathcal{S}$. Disintegrate each γ_i into $\gamma_i(dx, dy) = \mu(dx)K_x^i(dy)$. Consider the space $\Omega = E \times F \times F$, with the product σ -field and the probability measure

$$P(dx, dy_1, dy_2) = \mu(dx)K_x^1(dy_1)K_x^2(dy_2).$$

Define the canonical random variables X, Y_1 , and Y_2 by $X(x, y_1, y_2) = x$ and $Y_i(x, y_1, y_2) = y_i$ for $i = 1, 2$. By construction, we have $P \circ (X, Y_i)^{-1} = \gamma_i$, and thus our assumption of independent pathwise uniqueness implies $P(Y_1 = Y_2) = 1$, and so $K^1 = K^2$. Thus \mathcal{S} is a singleton. Since $Y_1 = Y_2$ a.s. it follows that for μ -a.e. $x \in E$ the measure $K_x^1 \times K_x^1$ is concentrated on the diagonal $\{(y, y) : y \in F\}$. This implies that K_x^1 is a point mass for each such x , which in turn implies that the unique element of \mathcal{S} is a strong solution. \square

While this result is remarkably straightforward, there is often some work to be done in its application. In particular, in the definition of independent pathwise uniqueness, we would prefer to restrict our attention to couplings with even better properties (for example, “joint compatibility” in the language of [80]). Corollary A.1.4 allows us to do this in a sense which is particularly useful for SDEs and for the MFG uniqueness results of Chapter 7.3. Let us illustrate this by example by proving Theorem A.2.3:

Proof of Theorem A.2.3. To apply Theorem A.3.1 we prove that independent pathwise uniqueness holds, identifying $E = \Omega$ and $F = \mathcal{C}^d$. Suppose we are given a probability measure \tilde{P} on $\tilde{\Omega} := \Omega \times \mathcal{C}^d \times \mathcal{C}^d$. Let $(\zeta, X, Y) : \tilde{\Omega} \rightarrow \mathcal{C}^d \times \mathcal{C}^d$ denote the identity map, $(\zeta, X, Y)(\omega, x, y) = (\omega, x, y)$. Assume X and Y are conditionally independent of ζ under \tilde{P} , and assume $\tilde{P} \circ (\zeta, X)^{-1}$ and $\tilde{P} \circ (\zeta, Y)^{-1}$ are both solution measures. Both X and Y must then satisfy the same SDE (A.1), driven by the same Brownian motion, the same initial state, and the same random parameter in the coefficients.

At this point, we would like to use the usual arguments, using the Lipschitz assumption and Gronwall's inequality, to show that $X = Y$ almost surely. But to do this, we need to show that the SDEs are defined *with respect to the same filtration on $\tilde{\Omega}$* . For this it suffices to just check that W (defined originally on Ω) is a Wiener process with respect to the filtration $\mathcal{F}_t \vee \sigma(X_s, Y_s : s \leq t)$, where \mathcal{F}_t was originally defined on Ω but is extended to $\tilde{\Omega}$ in the natural way. That this holds true is exactly the conclusion of Corollary A.1.4, and is also not difficult to prove directly. \square

A.4 Martingale measures and martingale problems

In Chapter 8, we make some use of *martingale measures* in order to represent solutions of certain martingale problems in terms of stochastic differential equations. This appendix summarizes the basic definitions and results on martingale measures. All of the material we need is from [74], but Walsh [108] for a thorough introduction. Fix throughout the section a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, P)$ and a Polish space E with Borel σ -field $\mathcal{B}(E)$. Let \mathcal{A} be a subring of $\mathcal{B}(E)$, meaning that $A \cup B \in \mathcal{A}$ and $A \setminus B \in \mathcal{A}$ whenever $A, B \in \mathcal{A}$.

Definition A.4.1. An L^2 -valued σ -finite measure is a function $U : \mathcal{A} \rightarrow L^2(\Omega, \mathcal{F}, P)$ such that $\mathbb{E}[U(A)^2] < \infty$ for all $A \in \mathcal{A}$ and $U(A \cup B) = U(A) + U(B)$ a.s. for disjoint sets $A, B \in \mathcal{A}$. We say U is σ -finite if there exist $E_1 \subset E_2 \subset \dots$ in \mathcal{E} such that $\cup_n E_n = E$ and, for all n , we have $\mathcal{E}|_{E_n} := \{B \cap E_n : B \in \mathcal{E}\} \subset \mathcal{A}$ and

$$\sup_{A \in \mathcal{E}|_{E_n}} \mathbb{E}[U(A)^2] < \infty.$$

Definition A.4.2. We say $M = \{M_t(A) : t \geq 0, A \in \mathcal{A}\}$ is a \mathbb{F} -martingale measure on E if:

1. $M_0(A) = 0$ for all $A \in \mathcal{A}$.
2. $(M_t(A))_{t \geq 0}$ is a \mathbb{F} -martingale for each $A \in \mathcal{A}$.
3. For each $t \geq 0$, M_t is a L^2 -valued σ -finite measure.
4. For disjoint sets $A, B \in \mathcal{A}$, the martingales $(M_t(A))_{t \geq 0}$ and $(M_t(B))_{t \geq 0}$ are orthogonal, i.e. the product $(M_t(A)M_t(B))_{t \geq 0}$ is a martingale.

We say two martingale measures M and N are *orthogonal* if $(M_t(A)N_t(B))_{t \geq 0}$ is a martingale for each $A, B \in \mathcal{A}$.

It is known [74, Theorem I-4] that martingale measures admit *intensity measures*: Given a martingale measure M , we can associate to it a random σ -finite positive measure μ on $\mathbb{R}_+ \times E$ such that $(\mu([0, t] \times A))_{t \geq 0}$ is the (predictable) quadratic variation of $(M_t(A))_{t \geq 0}$, for each A . We say M is continuous if $t \mapsto M_t(A)$ is a.s.-continuous, for each A ; this is equivalent to saying $\mu(\{t\} \times A) = 0$ a.s. for all A .

Let us finally construct a stochastic integral with respect to M . Consider the set of simple integrands \mathcal{S} , consisting of functions $h : \Omega \times \mathbb{R}_+ \times E \rightarrow \mathbb{R}$ of the form

$$h(\omega, t, x) = \sum_{i=1}^n \varphi_i(\omega) 1_{(t_{i-1}, t_i]}(t) 1_{A_i}(x),$$

where $0 = t_0 < t_1 < \dots < t_n$, φ_i is \mathcal{F}_{t_i} -measurable, and $A_i \in \mathcal{A}$ for each i . The stochastic integral of such an h with respect to M is defined as the martingale measure $h \cdot M$ satisfying

$$(h \cdot M)_t(B) = \sum_{i=1}^n \varphi_i \left(M_{t_{i-1} \wedge t}(B \cap A_i) - M_{t_i \wedge t}(B \cap A_i) \right).$$

It is straightforward to check that \mathcal{S} is dense in the space L_μ^2 , consisting of those functions $h : \Omega \times \mathbb{R}_+ \times E \rightarrow \mathbb{R}$ which are measurable with respect to the product of the predictable σ -field and $\mathcal{B}(E)$ and satisfy

$$\mathbb{E} \int_{\mathbb{R}_+ \times E} |h(t, x)|^2 \mu(dt, dx) < \infty,$$

where we have suppressed the argument ω in both h and μ . Thus (using the isometry result stated more generally in (A.2) below) we may extend the linear map $h \mapsto h \cdot M$ from \mathcal{S} to L_μ^2 in the usual way. For $h \in L_\mu^2$ and $A \subset E$, we may write

$$\int_{[0, t]} \int_E h(s, x) M(ds, dx) := (h \cdot M)_t(E).$$

Given a vector $M = (M^1, \dots, M^m)$ of orthogonal martingale measures, and given a $d \times m$ matrix of elements $(h_{i,k})$ of L_μ^2 , define the stochastic integral $\int_{[0, t]} \int_E h(s, x) M(ds, dx)$ as the d -dimensional process with i^{th} component

$$\sum_{k=1}^m \int_{[0, t]} \int_E h_{i,k}(s, x) M^k(ds, dx).$$

There are two results we use in Chapter 8. First, if $f, g \in L_\mu^2$ and M is a martingale measure, then the quadratic covariation of $(f \cdot M)_{\cdot}(A)$ and $(g \cdot M)_{\cdot}(B)$ at time t is

$$\int_{(0, t]} \int_{A \cap B} f(s, x) g(s, x) \mu(ds, dx).$$

In particular, the following isometry formula holds:

$$\mathbb{E} \left[\left(\int_{[0, t]} \int_E f(s, x) M(ds, dx) \right)^2 \right] = \mathbb{E} \left[\int_{[0, t]} \int_E |f(s, x)|^2 \mu(ds, dx) \right]. \quad (\text{A.2})$$

The second important result is the following theorem of El Karoui and Méléard pertaining to martingale problems, which generalizes the classical results on the correspondence between weak solutions of SDEs and solutions of martingale problems. Suppose we are given measurable functions

$$(b, \sigma) : \mathbb{R}_+ \times \mathbb{R}^d \times E \rightarrow \mathbb{R}^d \times \mathbb{R}^{d \times m}.$$

Define the operator L on functions $\varphi \in C_0^\infty(\mathbb{R}^d)$ by

$$L\varphi(x, e) := b(x, e)^\top D\varphi(x) + \frac{1}{2} \text{Tr} [\sigma \sigma^\top(x, e) D^2 \varphi(x)].$$

Assume, for simplicity, that (b, σ) is bounded, although we can move beyond this case via localization.

Theorem A.4.3 (Theorem IV-2 of [74]). *Suppose we are given a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$ supporting a continuous d -dimensional adapted process X and a $\mathcal{P}(E)$ -valued predictable process Λ . Suppose that for each $\varphi \in C_0^\infty(\mathbb{R}^d)$ the process*

$$\varphi(X_t) - \int_0^t \int_E L\varphi(X_s, e) \Lambda_s(de) ds$$

is a martingale. Then, perhaps on an extension of the probability space, there exist m continuous orthogonal martingale measures $M = (M^1, \dots, M^m)$, each with intensity $dt \Lambda_t(de)$, such that X satisfies the SDE

$$X_t = X_0 + \int_0^t \int_E b(X_s, e) \Lambda_s(de) ds + \int_{[0, t] \times E} \sigma(X_s, e) M(ds, de).$$

The SDE above may alternatively be written in differential form,

$$dX_t = \int_E b(X_t, e) \Lambda_t(de) dt + \int_E \sigma(X_t, e) M(dt, de).$$

As a sanity check, suppose $E = \{e_0\}$ is a singleton. Then to say that M has intensity $dt \delta_{e_0}(de)$ means that $W_t := M([0, t] \times \{e_0\})$ is a martingale with quadratic variation t , which implies that it is a Wiener process thanks to Lévy's characterization. The stochastic integral then reduces to the usual Brownian integral,

$$\int_{[0, t] \times E} h(s, e) M(ds, de) = \int_0^t h(s, e_0) dW_s,$$

and the martingale problem of Theorem A.4.3 is a classical one, of the form studied by Stroock and Varadhan [102].

Appendix B

Tightness results for Itô processes

This section proves the three tightness results for Itô processes stated throughout the body of the thesis: Proposition 5.3.2, the more general Proposition 6.1.3, and finally Proposition 8.2.4. Indeed, the first of these follows quickly from the second, and both work under the assumption of uncontrolled volatility. The third, on the other hand, works with controlled volatility and must thus employ martingale measures. Nonetheless, its proof is quite similar, but some more careful estimates are required.

Proof of Proposition 5.3.2. Simply note that the set \mathcal{Q} defined in Proposition 5.3.2 is contained in the set \mathcal{Q} defined in Proposition 6.1.3 by choosing the constants $\kappa_{n,i} = c$ for each n, i , where c is the constant from the statement of Proposition 5.3.2. \square

Proof of Proposition 6.1.3. For each $1 \leq i \leq n$ and $P \in \mathcal{Q}_{\kappa_{n,i}}$, apply the Burkholder-Davis-Gundy inequality and the growth assumption to find a constant $C > 0$ (which will change from line to line but depends only on c , T , and p') such that

$$\begin{aligned} \mathbb{E}^P[\|X\|_t^{p'}] &\leq C \mathbb{E}^P \left[|X_0|^{p'} + \left(\int_0^t \int_A |B(s, a)| \Lambda_s(da) ds \right)^{p'} + \left(\int_0^t |\Sigma \Sigma^\top(s)| ds \right)^{p'/2} \right] \\ &\leq C \mathbb{E}^P \left\{ 1 + |X_0|^{p'} + Z^{p'} + \int_0^t \left(\|X\|_s^{p'} + \int_A |a|^{p'} \Lambda_s(da) \right) ds \right\}, \end{aligned}$$

where we used also $p' \geq 2$ and Jensen's inequality. By Gronwall's inequality,

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

Thus

$$\begin{aligned} \sup_{P \in \mathcal{Q}} \mathbb{E}^P[\|X\|_T^{p'}] &= \sup_n \sup \left\{ \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{P_i}[\|X\|_T^{p'}] : P_i \in \mathcal{Q}_{\kappa_{n,i}} \text{ for } i = 1, \dots, n \right\} \\ &\leq C \sup_n \frac{1}{n} \sum_{i=1}^n (1 + \kappa_{n,i}) < \infty. \end{aligned} \tag{B.1}$$

By assumption, we have also

$$\sup_{P \in \mathcal{Q}} \mathbb{E}^P \int_0^T \int_A |a|^{p'} \Lambda_t(da) dt \leq \sup_n \frac{1}{n} \sum_{i=1}^n \kappa_{n,i} < \infty. \quad (\text{B.2})$$

It follows from [B.2](#), [Lemma 2.1.16](#), and [Lemma 2.1.13](#) that $\{P \circ \Lambda^{-1} : P \in \mathcal{Q}\}$ is relatively compact. In light of [\(B.1\)](#) and [Proposition 2.1.7](#), it remains to show that $\{P \circ X^{-1} : P \in \mathcal{Q}\} \subset \mathcal{P}(\mathcal{C}^d)$ is tight. To check this, we will verify Aldous' criterion for tightness [[73](#), [Lemma 16.12](#)], or

$$\lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}} \sup_{\tau} \mathbb{E}^P [|X_{(\tau+\delta) \wedge T} - X_{\tau}|^p] = 0, \quad (\text{B.3})$$

where the supremum is over stopping times τ valued in $[0, T]$. The Burkholder-Davis-Gundy inequality implies that there exists a constant $C' > 0$ (which again depends only on c , T , and p and will change from line to line) such that, for any i and any $P \in \mathcal{Q}_{\kappa_{n,i}}$,

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

Since $p' > p$, we have $\mathbb{E}^P[Z^p] \leq \mathbb{E}^P[Z^{p'}]^{p/p'} \leq \kappa_{n,i}^{p/p'}$ for $P \in \mathcal{Q}_{\kappa_{n,i}}$, and thus by assumption

$$\begin{aligned} \sup_{P \in \mathcal{Q}} \mathbb{E}^P[Z^p] &= \sup_n \sup \left\{ \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{P_i}[Z^p] : P_i \in \mathcal{Q}_{\kappa_{n,i}} \text{ for } i = 1, \dots, n \right\} \\ &\leq \sup_n \frac{1}{n} \sum_{i=1}^n \kappa_{n,i}^{p/p'} < \infty. \end{aligned}$$

This and [\(B.1\)](#) imply

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

To control the term with Λ , note that $p' > p$ and (B.2) imply

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

Putting this all together proves (B.3). \square

Proof of Proposition 8.2.4. For each $P \in \mathcal{Q}_c$ with corresponding probability space $(\Omega, \mathcal{F}_t, P)$ and coefficients b, σ , standard estimates as in Lemma 5.3.1 yield

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

where $C > 0$ does not depend on P . Hence assumption (6) implies

$$\sup_{P \in \mathcal{Q}_c} \mathbb{E}^P \|X\|_T^{p'} \leq C(1 + c) < \infty. \quad (\text{B.4})$$

It follows from the hypothesis (6), from Lemma 2.1.16, and from Corollary 2.1.13 that $\{P \circ \Lambda^{-1} : P \in \mathcal{Q}_c\}$ is relatively compact. In light of (B.4) (and Proposition 2.1.7), it remains only to show that $\{P \circ X^{-1} : P \in \mathcal{Q}_c\} \subset \mathcal{P}(\mathcal{C}^d)$ is tight. To check this, we will again verify Aldous' criterion for tightness [73, Theorem 16.10],

$$\lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}_c} \sup_{\tau} \mathbb{E}^P |X_{(\tau+\delta) \wedge T} - X_{\tau}|^p = 0, \quad (\text{B.5})$$

where the innermost supremum is over stopping times τ valued in $[0, T]$. The Burkholder-Davis-Gundy inequality implies that there exists $C' > 0$ such that, for each $P \in \mathcal{Q}_c$ and each τ ,

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

Now note that if $1 \leq p < 2$ and $x, y \geq 0$ then $(x + y)^{p/2} \leq x^{p/2} + y^{p/2}$, and if $p \geq 2$ then $(x + y)^{p/2} \leq 2^{p/2-1}(x^{p/2} + y^{p/2})$. In either case, we find another constant C'' such that,

each $P \in \mathcal{Q}_c$ and each τ ,

$$\begin{aligned} & \mathbb{E}^P |X_{(\tau+\delta) \wedge T} - X_\tau|^p \\ & \leq C'' \mathbb{E}^P \left[|\delta c(1 + \|X\|_T)|^p + c^p \int_\tau^{(\tau+\delta) \wedge T} \int_A |a|^p \Lambda_t(da) dt \right] \\ & \quad + C'' \mathbb{E}^P \left[|\delta c(1 + \|X\|_T^{p_\sigma})|^{p/2} + \left| c \int_\tau^{(\tau+\delta) \wedge T} \int_A |a|^{p_\sigma} \Lambda_t(da) dt \right|^{p/2} \right] \end{aligned} \quad (\text{B.6})$$

The first term of each line poses no problems, in light of (B.4); that is, since $p_\sigma \leq 2$ and $p < p'$,

$$\lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}_c} \sup_{\tau} \mathbb{E}^P \left[|\delta c(1 + \|X\|_T)|^p + |\delta c(1 + \|X\|_T^{p_\sigma})|^{p/2} \right] = 0.$$

On the other hand, note that

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

by assumption. It follows that for any $\gamma \in [0, p')$,

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

and in particular this holds for $\gamma = p$. Hence, if $p \geq 2$ then Jensen's inequality along with $p_\sigma \leq 2$ implies

$$\begin{aligned} & \lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}_c} \sup_{\tau} \mathbb{E}^P \left| \int_\tau^{(\tau+\delta) \wedge T} \int_A |a|^{p_\sigma} \Lambda_t(da) dt \right|^{p/2} \\ & \leq \lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}_c} \sup_{\tau} \mathbb{E}^P \left| \int_\tau^{(\tau+\delta) \wedge T} \int_A |a|^p \Lambda_t(da) dt \right| = 0, \end{aligned}$$

On the other hand, if $p < 2$, then Jensen's inequality in the other direction implies

$$\begin{aligned} & \lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}_c} \sup_{\tau} \mathbb{E}^P \left| \int_\tau^{(\tau+\delta) \wedge T} \int_A |a|^{p_\sigma} \Lambda_t(da) dt \right|^{p/2} \\ & \leq \lim_{\delta \downarrow 0} \sup_{P \in \mathcal{Q}_c} \sup_{\tau} \left(\mathbb{E}^P \int_\tau^{(\tau+\delta) \wedge T} \int_A |a|^{p_\sigma} \Lambda_t(da) dt \right)^{p/2} = 0, \end{aligned}$$

since $p_\sigma \leq p < p'$. Putting this together and returning to (B.6) proves (B.5). \square

Bibliography

- [1] Y. Achdou, F. Camilli, and I. Capuzzo-Dolcetta, *Mean field games: numerical methods for the planning problem*, SIAM Journal on Control and Optimization **50** (2012), no. 1, 77–109.
- [2] Y. Achdou and I. Capuzzo-Dolcetta, *Mean field games: Numerical methods*, SIAM Journal on Numerical Analysis **48** (2010), no. 3, 1136–1162.
- [3] S. Ahuja, *Wellposedness of mean field games with common noise under a weak monotonicity condition*, arXiv preprint arXiv:1406.7028 (2014).
- [4] D. Aldous, *Deterministic and stochastic models for coalescence (aggregation and coagulation): a review of the mean-field theory for probabilists*, Bernoulli (1999), 3–48.
- [5] C. Aliprantis and K. Border, *Infinite dimensional analysis: A hitchhiker’s guide*, 3 ed., Springer, 2007.
- [6] R.J. Aumann, *Markets with a continuum of traders*, Econometrica: Journal of the Econometric Society (1964), 39–50.
- [7] M. Bardi, *Explicit solutions of some linear-quadratic mean field games*, Networks and Heterogeneous Media **7** (2012), no. 2, 243–261.
- [8] M. Bardi and F.S. Priuli, *Linear-quadratic n -person and mean-field games with ergodic cost*, arXiv preprint arXiv:1401.1421 (2014).
- [9] M. Barlo and G. Carmona, *Strategic behavior in non-atomic games*, Available at SSRN 1977242 (2011).
- [10] M.T. Barlow, *One dimensional stochastic differential equations with no strong solution*, Journal of the London Mathematical Society **2** (1982), no. 2, 335–347.
- [11] A. Bensoussan, M. Chau, and P. Yam, *Mean field games with a dominating player*, arXiv preprint arXiv:1404.4148 (2014).
- [12] A. Bensoussan and J. Frehse, *Nonlinear elliptic systems in stochastic game theory*, Universität Bonn. SFB 72. Approximation und Optimierung, 1983.
- [13] ———, *Stochastic games for N players*, Journal of optimization theory and applications **105** (2000), no. 3, 543–565.

- [14] A. Bensoussan, J. Frehse, and P. Yam, *The master equation in mean field theory*, arXiv preprint arXiv:1404.4150 (2014).
- [15] A. Bensoussan, K.C.J. Sung, S.C.P. Yam, and S.P. Yung, *Linear-quadratic mean field games*, Preprint (2011).
- [16] J. Bergin and D. Bernhardt, *Anonymous sequential games with aggregate uncertainty*, Journal of Mathematical Economics **21** (1992), no. 6, 543–562.
- [17] ———, *Anonymous sequential games: existence and characterization of equilibria*, Economic Theory **5** (1995), no. 3, 461–489.
- [18] D. Bertsekas and S. Shreve, *Stochastic optimal control: The discrete time case*, Athena Scientific, 1996.
- [19] V. Bogachev, *Measure theory*, vol. 2, Springer, 2007.
- [20] V.S. Borkar and M.K. Ghosh, *Stochastic differential games: Occupation measure based approach*, Journal of optimization theory and applications **73** (1992), no. 2, 359–385.
- [21] M. Bossy and D. Talay, *A stochastic particle method for the McKean-Vlasov and the Burgers equation*, Mathematics of Computation of the American Mathematical Society **66** (1997), no. 217, 157–192.
- [22] P. Brémaud and M. Yor, *Changes of filtrations and of probability measures*, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete **45** (1978), no. 4, 269–295.
- [23] G. Brunick and S. Shreve, *Mimicking an Itô process by a solution of a stochastic differential equation*, arXiv preprint arXiv:1011.0111 (2010).
- [24] P. Cardaliaguet, *Notes on mean field games*.
- [25] ———, *Weak solutions for first order mean field games with local coupling*, arXiv preprint arXiv:1305.7015 (2013).
- [26] P. Cardaliaguet, P.J. Graber, A. Porretta, and D. Tonon, *Second order mean field games with degenerate diffusion and local coupling*, arXiv preprint arXiv:1407.7024 (2014).
- [27] E. Carlini and F.J. Silva, *A fully discrete semi-Lagrangian scheme for a first order mean field game problem*, SIAM Journal on Numerical Analysis **52** (2014), no. 1, 45–67.
- [28] E. Carlini and F.J. Silva, *A semi-Lagrangian scheme for a degenerate second order mean field game system*, arXiv preprint arXiv:1404.5932 (2014).
- [29] G. Carmona, *Nash equilibria of games with a continuum of players*, (2004).
- [30] ———, *On the purification of Nash equilibria of large games*, Economics Letters **85** (2004), no. 2, 215–219.

- [31] G. Carmona and K. Podczeck, *Approximation and characterization of Nash equilibria of large games*.
- [32] R. Carmona and F. Delarue, *Probabilistic analysis of mean field games*, SIAM Journal on Control and Optimization **51** (2013), 2705–2734.
- [33] ———, *The master equation for large population equilibriums*, arXiv preprint arXiv:1404.4694 (2014).
- [34] R. Carmona, F. Delarue, and A. Lachapelle, *Control of McKean–Vlasov dynamics versus mean field games*, Mathematics and Financial Economics **7** (2013), no. 2, 131–166.
- [35] R. Carmona, J.P. Fouque, and L.H. Sun, *Mean field games and systemic risk*, arXiv preprint arXiv:1308.2172 (2013).
- [36] R. Carmona and D. Lacker, *A probabilistic weak formulation of mean field games and applications*, Ann. Appl. Probab. **25**, no. 3, 1189–1231.
- [37] R. Carmona and X. Zhu, *A probabilistic approach to mean field games with major and minor players*, arXiv preprint arXiv:1409.7141 (2014).
- [38] C. Castaing, P.R. De Fitte, and M. Valadier, *Young measures on topological spaces: with applications in control theory and probability theory*, vol. 571, Springer, 2004.
- [39] J.-F. Chassagneux, D. Crisan, and F. Delarue, *Classical solutions to the master equation for large population equilibria*, arXiv preprint arXiv:1411.3009 (2014).
- [40] D. Dawson and J. Vaillancourt, *Stochastic McKean-Vlasov equations*, Nonlinear Differential Equations and Applications **2** (1995), no. 2, 199–229.
- [41] D.A. Dawson and J. Gärtner, *Large deviations, free energy functional, and quasipotential for a mean field model of interacting diffusions*, Laboratory for Research in Statistics and Probability, Carleton University= Laboratoire de recherche en statistique et probabilités, Carleton University, 1986.
- [42] ———, *Large deviations from the McKean-Vlasov limit for weakly interacting diffusions*, Stochastics: An International Journal of Probability and Stochastic Processes **20** (1987), no. 4, 247–308.
- [43] ———, *Long-time fluctuations of weakly interacting diffusions*, Springer, 1987.
- [44] D.W. Diamond and P.H. Dybvig, *Bank runs, deposit insurance, and liquidity*, The journal of political economy (1983), 401–419.
- [45] D. Duffie, S. Malamud, and G. Manso, *Information percolation with equilibrium search dynamics*, Econometrica **77** (2009), no. 5, 1513–1574.
- [46] J. Dugundji, *An extension of Tietze’s theorem*, Pacific J. Math **1** (1951), no. 3, 353–367.

- [47] S.N. Ethier and T.G. Kurtz, *Markov processes: characterization and convergence*, 2 ed., vol. 282, Wiley-Interscience, 2005.
- [48] K. Fan, *Fixed-point and minimax theorems in locally convex topological linear spaces*, Proceedings of the National Academy of Sciences of the United States of America **38** (1952), no. 2, 121.
- [49] E. Feleqi, *The derivation of ergodic mean field game equations for several populations of players*, Dynamic Games and Applications **3** (2013), no. 4, 523–536.
- [50] A.F. Filippov, *On certain questions in the theory of optimal control*, Journal of the SIAM, Series A: Control **1** (1962), no. 1, 76–84.
- [51] M. Fischer, *On the connection between symmetric n -player games and mean field games*, arXiv preprint arXiv:1405.1345 (2014).
- [52] T. Funaki, *A certain class of diffusion processes associated with nonlinear parabolic equations*, Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete **67** (1984), no. 3, 331–348.
- [53] J. Gärtner, *On the McKean-Vlasov limit for interacting diffusions*, Mathematische Nachrichten **137** (1988), no. 1, 197–248.
- [54] D.A. Gomes, J. Mohr, and R.R. Souza, *Continuous time finite state mean field games*, Applied Mathematics & Optimization **68** (2013), no. 1, 99–143.
- [55] D.A. Gomes and J. Saúde, *Mean field games models-a brief survey*, Dynamic Games and Applications (2013), 1–45.
- [56] D.A. Gomes and V.K. Voskanyan, *Extended mean field games-formulation, existence, uniqueness and examples*, arXiv preprint arXiv:1305.2600 (2013).
- [57] J. Graber, *Weak solutions for mean field games with congestion*, arXiv preprint arXiv:1503.04733 (2015).
- [58] C. Graham and S. Méléard, *Stochastic particle approximations for generalized boltzmann models and convergence estimates*, The Annals of Probability **25** (1997), no. 1, 115–132.
- [59] E.J. Green, *Continuum and finite-player noncooperative models of competition*, Econometrica: Journal of the Econometric Society (1984), 975–993.
- [60] E.J. Green and P. Lin, *Implementing efficient allocations in a model of financial intermediation*, Journal of Economic Theory **109** (2003), no. 1, 1–23.
- [61] O. Guéant, J.M. Lasry, and P.L. Lions, *Mean field games and applications*, Paris-Princeton Lectures on Mathematical Finance 2010, Lecture Notes in Mathematics, vol. 2003, Springer Berlin / Heidelberg, 2011, pp. 205–266.

- [62] S. Hamadène, *Backward-forward SDEs and stochastic differential games*, Stochastic processes and their applications **77** (1998), no. 1, 1–15.
- [63] U.G. Haussmann and J.P. Lepeltier, *On the existence of optimal controls*, SIAM Journal on Control and Optimization **28** (1990), no. 4, 851–902.
- [64] U. Horst, *Dynamic systems of social interactions*, Journal of economic behavior & organization **73** (2010), no. 2, 158–170.
- [65] U. Horst and J.A. Scheinkman, *Equilibria in systems of social interactions*, Journal of Economic Theory **130** (2006), no. 1, 44–77.
- [66] D. Housman, *Infinite player noncooperative games and the continuity of the Nash equilibrium correspondence*, Mathematics of Operations Research **13** (1988), no. 3, 488–496.
- [67] M. Huang, R. Malhamé, and P. Caines, *Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle*, Communications in Information and Systems **6** (2006), no. 3, 221–252.
- [68] ———, *The Nash certainty equivalence principle and McKean-Vlasov systems: An invariance principle and entry adaptation*, Proceedings of the 46th IEEE Conference on Decision and Control, 2007, pp. 121–126.
- [69] J. Jacod and J. Mémin, *Sur un type de convergence intermédiaire entre la convergence en loi et la convergence en probabilité*, Séminaire de probabilités de Strasbourg **15** (1981), 529–546.
- [70] ———, *Weak and strong solutions of stochastic differential equations: existence and stability*, Stochastic integrals, Springer, 1981, pp. 169–212.
- [71] B. Jovanovic and R.W. Rosenthal, *Anonymous sequential games*, Journal of Mathematical Economics **17** (1988), no. 1, 77–87.
- [72] M. Kac, *Foundations of kinetic theory*, Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, vol. 1955, 1954, pp. 171–197.
- [73] O. Kallenberg, *Foundations of modern probability*, Springer, 2002.
- [74] N. El Karoui and S. Méléard, *Martingale measures and stochastic calculus*, Probability Theory and Related Fields **84** (1990), no. 1, 83–101.
- [75] N. El Karoui, D.H. Nguyen, and M. Jeanblanc-Picqué, *Compactification methods in the control of degenerate diffusions: existence of an optimal control*, Stochastics **20** (1987), no. 3, 169–219.
- [76] ———, *Existence of an optimal Markovian filter for the control under partial observations*, SIAM journal on control and optimization **26** (1988), no. 5, 1025–1061.

- [77] N. El Karoui and X. Tan, *Capacities, measurable selection and dynamic programming part ii: Application in stochastic control problems*, arXiv preprint arXiv:1310.3364 (2013).
- [78] V.N. Kolokoltsov, J. Li, and W. Yang, *Mean field games and nonlinear Markov processes*, arXiv preprint arXiv:1112.3744 (2011).
- [79] T.G. Kurtz, *Martingale problems for conditional distributions of Markov processes.*, Electronic J. Probab. **3** (1998), no. 9, 1–29.
- [80] ———, *Weak and strong solutions of general stochastic models*, arXiv preprint arXiv:1305.6747 (2013).
- [81] T.G. Kurtz and P. Protter, *Weak limit theorems for stochastic integrals and stochastic differential equations*, The Annals of Probability (1991), 1035–1070.
- [82] T.G. Kurtz and R.H. Stockbridge, *Existence of Markov controls and characterization of optimal Markov controls*, SIAM Journal on Control and Optimization **36** (1998), no. 2, 609–653.
- [83] T.G. Kurtz and J. Xiong, *Particle representations for a class of nonlinear SPDEs*, Stochastic Processes and their Applications **83** (1999), no. 1, 103–126.
- [84] ———, *Numerical solutions for a class of SPDEs with application to filtering*, Stochastics in Finite and Infinite Dimensions, Springer, 2001, pp. 233–258.
- [85] H.J. Kushner and P. Dupuis, *Numerical methods for stochastic control problems in continuous time*, vol. 24, Springer Science & Business Media, 2001.
- [86] A. Lachapelle, J. Salomon, and G. Turinici, *Computation of mean field equilibria in economics*, Mathematical Models and Methods in Applied Sciences **20** (2010), no. 04, 567–588.
- [87] A. Lachapelle and M.T. Wolfram, *On a mean field game approach modeling congestion and aversion in pedestrian crowds*, Transportation Research Part B: Methodological **45** (2011), no. 10, 1572 – 1589.
- [88] D. Lacker, *Mean field games via controlled martingale problems: existence of Markovian equilibria*, arXiv preprint arXiv:1404.2642 (2014).
- [89] J.-M. Lasry and P.-L. Lions, *Jeux à champ moyen. I–Le cas stationnaire*, Comptes Rendus Mathématique **343** (2006), no. 9, 619–625.
- [90] ———, *Jeux à champ moyen. II–Horizon fini et contrôle optimal*, Comptes Rendus Mathématique **343** (2006), no. 10, 679–684.
- [91] J.M. Lasry and P.L. Lions, *Mean field games*, Japanese Journal of Mathematics **2** (2007), 229–260.

- [92] A. Mas-Colell, *On a theorem of Schmeidler*, Journal of Mathematical Economics **13** (1984), no. 3, 201–206.
- [93] H.P. McKean, *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.
- [94] S.L. Nguyen and M. Huang, *Linear-quadratic-Gaussian mixed games with continuum-parametrized minor players*, SIAM Journal on Control and Optimization **50** (2012), no. 5, 2907–2937.
- [95] M. Nourian and P. Caines, *ϵ -Nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents*, SIAM Journal on Control and Optimization **51** (2013), no. 4, 3302–3331.
- [96] K. Oelschläger, *A martingale approach to the law of large numbers for weakly interacting stochastic processes*, The Annals of Probability (1984), 458–479.
- [97] S. Peng, *Stochastic Hamilton-Jacobi-Bellman equations*, SIAM Journal on Control and Optimization **30** (1992), no. 2, 284–304.
- [98] H. Pham, *Continuous-time stochastic control and optimization with financial applications*, 1 ed., Springer, 2009.
- [99] P. Protter, *Stochastic integration and differential equations: Version 2.1*, vol. 21, Springer Science & Business Media, 2004.
- [100] S. Rashid, *Equilibrium points of non-atomic games: asymptotic results*, Economics Letters **12** (1983), no. 1, 7–10.
- [101] D. Schmeidler, *Equilibrium points of nonatomic games*, Journal of Statistical Physics **7** (1973), no. 4, 295–300.
- [102] D.W. Stroock and S.R.S. Varadhan, *Multidimensional diffusion processes*, 2 ed., Grundlehren Der Mathematischen Wissenschaften, Springer, 1979.
- [103] Y. Sun, *The exact law of large numbers via Fubini extension and characterization of insurable risks*, Journal of Economic Theory **126** (2006), no. 1, 31–69.
- [104] A.S. Sznitman, *Topics in propagation of chaos*, Ecole d’été de Probabilités de Saint-Flour XIX - 1989, Lecture Notes in Mathematics, vol. 1464, Springer Berlin / Heidelberg, 1991, pp. 165–251.
- [105] H. Uhlig, *A law of large numbers for large economies*, Economic Theory **8** (1996), no. 1, 41–50.
- [106] J. Vaillancourt, *On the existence of random McKean-Vlasov limits for triangular arrays of exchangeable diffusions*, Stochastic Analysis and Applications **6** (1988), no. 4, 431–446.

- [107] C. Villani, *Topics in optimal transportation*, Graduate Studies in Mathematics, American Mathematical Society, 2003.
- [108] J.B. Walsh, *An introduction to stochastic partial differential equations*, Springer, 1986.
- [109] T. Yamada and S. Watanabe, *On the uniqueness of solutions of stochastic differential equations*, J. Math. Kyoto Univ **11** (1971), no. 1, 155–167.