
The Master Equation and the Convergence Problem in Mean Field Games

Pierre Cardaliaguet
François Delarue
Jean-Michel Lasry
Pierre-Louis Lions

Annals of Mathematics Studies
Number 201

The Master Equation and the Convergence Problem in Mean Field Games

Pierre Cardaliaguet
François Delarue
Jean-Michel Lasry
Pierre-Louis Lions

PRINCETON UNIVERSITY PRESS

PRINCETON AND OXFORD

2019

Copyright © 2019 by Princeton University Press

Requests for permission to reproduce material from this work
should be sent to permissions@press.princeton.edu

Published by Princeton University Press
41 William Street, Princeton, New Jersey 08540
6 Oxford Street, Woodstock, Oxfordshire OX20 1TR

press.princeton.edu

All Rights Reserved

LCCN: 2019936796

ISBN 978-0-691-19070-9

ISBN (pbk.) 978-0-691-19071-6

British Library Cataloging-in-Publication Data is available

Editorial: Vickie Kearn, Susannah Shoemaker, and Lauren Bucca

Production Editorial: Nathan Carr

Production: Erin Suydam

Publicity: Matthew Taylor and Kathryn Stevens

Copyeditor: Theresa Kornak

This book has been composed in L^AT_EX

Printed on acid-free paper. ∞

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Contents

Preface	vii
1 Introduction	1
1.1 From the Nash System to the Master Equation	1
1.2 Informal Derivation of the Master Equation	19
2 Presentation of the Main Results	28
2.1 Notations	28
2.2 Derivatives	30
2.3 Assumptions	36
2.4 Statement of the Main Results	38
3 A Starter: The First-Order Master Equation	48
3.1 Space Regularity of U	49
3.2 Lipschitz Continuity of U	53
3.3 Estimates on a Linear System	60
3.4 Differentiability of U with Respect to the Measure	67
3.5 Proof of the Solvability of the First-Order Master Equation	72
3.6 Lipschitz Continuity of $\frac{\delta U}{\delta m}$ with Respect to m	75
3.7 Link with the Optimal Control of the Fokker–Planck Equation	77
4 Mean Field Game System with a Common Noise	85
4.1 Stochastic Fokker–Planck/Hamilton–Jacobi–Bellman System	86
4.2 Probabilistic Setup	89
4.3 Solvability of the Stochastic Fokker–Planck/Hamilton–Jacobi–Bellman System	89
4.4 Linearization	112
5 The Second-Order Master Equation	128
5.1 Construction of the Solution	128
5.2 First-Order Differentiability	132
5.3 Second-Order Differentiability	141
5.4 Derivation of the Master Equation	150
5.5 Well-Posedness of the Stochastic MFG System	155

6	Convergence of the Nash System	159
6.1	Finite Dimensional Projections of U	160
6.2	Convergence	166
6.3	Propagation of Chaos	172
A	Appendix	175
A.1	Link with the Derivative on the Space of Random Variables	175
A.2	Technical Remarks on Derivatives	183
A.3	Various Forms of Itô's Formula	189
	References	203
	Index	211

Preface

THE PURPOSE OF THIS SHORT MONOGRAPH is to address recent advances in the theory of mean field games (MFGs), which has met an amazing success since the simultaneous pioneering works by Lasry and Lions and by Caines, Huang, and Malhamé more than ten years ago. While earlier developments in the theory have been largely spread out over the last decade, issues that are addressed in this book require a new step forward in the analysis. The text evolved with the objective to provide a self-contained study of the so-called *master equation* and to answer the *convergence problem*, which has remained mainly open so far. As the writing progressed, the manuscript became longer and longer and, in the end, it turned out to be more relevant to publish the whole as a book.

There might be several reasons to explain the growing interest for MFGs. From the technical point of view, the underpinning stakes fall within several mathematical areas, including partial differential equations, probability theory, stochastic analysis, optimal control, and optimal transportation. In particular, several issues raised by the analysis of MFGs may be tackled by analytical or probabilistic tools; sometimes, they even require a subtle implementation of mixed arguments, which is precisely the case in this book. As a matter of fact, researchers from different disciplines have developed an interest in the subject, which has grown very quickly. Another explanation for the interest in the theory is the wide range of applications that it offers. While they were originally inspired by works in economics on heterogeneous agents, MFG models now appear under various forms in several domains, which include, for instance, mathematical finance, study of crowd phenomena, epidemiology, and cybersecurity.

Mean field games should be understood as games with a continuum of players, each of them interacting with the whole statistical distribution of the population. In this regard, they are expected to provide an asymptotic formulation for games with finitely many players with mean field interaction. In most of the related works, the connection between finite games and MFGs is addressed in the following manner: It is usually shown that solutions of the asymptotic problem generate an *almost equilibrium*, understood in the sense of Nash, to the corresponding finite game, the accuracy of the equilibrium getting stronger and stronger as the number of players in the finite game tends to infinity. The main purpose of this book is to focus on the converse problem, which may be formulated as follows: Do the equilibria of the finite games (if they exist) converge to a solution of the corresponding MFG as the number of players becomes

very large? Surprisingly, answering this question turns out to be much more difficult than proving that any asymptotic solution generates an almost equilibrium. Though several works addressed this problem in specific cases, including the case when the equilibria of the finite player game are taken over open loop strategies, the general case when the agents play strategies in closed (Markovian) feedback form has remained open so far. The objective here is to exhibit a quite large class of MFGs for which the answer is positive and, to do so, to implement a method that is robust enough to accommodate other sets of assumption.

The intrinsic difficulty in proving the convergence of finite player equilibria may be explained as follows. When taken over strategies in closed Markovian form, Nash equilibria of a stochastic differential game with N players in a state of dimension d may be described through a system of N quasilinear parabolic partial differential equations in dimension $N \times d$, which we refer to as the *Nash system* throughout the monograph. As N becomes larger and larger, the system obviously becomes more and more intricate. In particular, it seems especially difficult to get any a priori estimate that could be helpful for passing to the limit by means of a compactness argument. The strategy developed in the book is thus to bypass any detailed study of the Nash system. Instead, we use a short cut and focus directly on the expected limiting form of the Nash system. This limiting form is precisely what we call *the master equation* in the title of the book. As a result of the symmetry inherent in the mean field structure, this limiting form is no longer a system of equations but reduces to one equation only, which makes it simpler than the Nash system. It describes the equilibrium cost to one representative player in a continuum of players. Actually, to account for the continuum of players underpinning the game, the master equation has to be set over the Euclidean space times the space of probability measures; the state variable is thus a pair that encodes both the private state of a single representative player together with the statistical distribution of the states of all the players. Most of the book is thus dedicated to the analysis of this master equation. One of the key results in the book is to show that, under appropriate conditions on the coefficients, the master equation is uniquely solvable in the classical sense for an appropriate notion of differential calculus on the space of probability measures. Among the assumptions we require, we assume the coefficients to be monotone in the direction of the measure; as demonstrated earlier by Lasry and Lions, this forces uniqueness of the solution to the corresponding MFG.

Smoothness of the solution then plays a crucial role in our program. It is indeed the precise property we use next for proving the convergence of the finite player equilibria to the solution of the limiting MFG. The key step is indeed to expand the solution of the master equation along the “equilibrium trajectories” of the finite player games, which requires enough regularity. As indicated earlier, this methodology seems to be quite sharp and should certainly work under different sets of assumptions.

Actually, the master equation was already introduced by Lions in his lectures at Collège de France. It provides an alternative formulation to MFGs, different

from the earlier (and most popular) one based on a coupled forward–backward system, known as the *MFG system*, which is made of a backward Hamilton–Jacobi equation and a forward Kolmogorov equation. Part of our program in the book is then achieved by exploiting quite systematically the connection between the MFG system and the master equation: In short, the MFG system plays the role of characteristics for the master equation. In the text, we use this correspondence heavily to establish, by means of a flow method, the smoothness of the solution to the master equation.

Though the MFG system was extensively studied in earlier works, we provide in the book a detailed analysis of it in the case when players in the finite game are subject to a so-called common noise: Under the action of this common noise, both the backward and forward equations in the MFG system become stochastic, which makes it more complicated; as a result, we devote a whole chapter to addressing the solvability of the MFG system under the presence of such a common noise. Together with the study of the convergence problem, this perspective is completely new in the literature.

The book is organized in six chapters, which include a detailed introduction and are followed by an appendix. The guideline follows the aforementioned steps: Part of the book is dedicated to the analysis of the master equation, including the study of the MFG system with a common noise, and the rest concerns the convergence problem. The main results obtained in the book are collected in Chapter 2. Chapter 3 is a sort of warm-up, as it contains a preliminary analysis of the master equation in the simpler case when there is no common noise. In Chapter 4, we study the MFG system in the presence of a common noise, and the corresponding analysis of the master equation is performed in Chapter 5. The convergence problem is addressed in Chapter 6. We suggest the reader start with the Introduction, which contains in particular a formal derivation of the master equation, and then to carry on with Chapters 2 and 3. Possibly, the reader who is interested only in MFGs without common noise may skip Chapters 4 and 5 and switch directly to Chapter 6. In such a case, she/he has to set the parameter β , which stands for the intensity of the common noise throughout the book, equal to 0. The Appendix contains several results on the differential calculus on the space of probability measures together with an Itô’s formula for functionals of a process taking values in the space of probability measures. We emphasize that, for simplicity, most of the analysis provided in the book is on the torus, but, as already explained, we feel that the method is robust enough to accommodate the nonperiodic setting.

To conclude, we would like to thank our colleagues from our field for all the stimulating discussions and work sessions we have shared with them. Some of them have formulated very useful comments and suggestions on the preliminary draft of the book. They include in particular Yves Achdou, Martino Bardi, Alain Bensoussan, René Carmona, Jean-François Chassagneux, Markus Fischer, Wilfrid Gangbo, Christine Grün, Daniel Lacker and Alessio Porretta. We are also very grateful to the anonymous referees who examined the various versions of the manuscript. Their suggestions helped us greatly in improving the text.

We also thank the French National Research Agency, which supported part of this work under the grant ANR MFG (ANR-16-CE40-0015-01).

Pierre Cardaliaguet, Paris
François Delarue, Nice
Jean-Michel Lasry, Paris
Pierre-Louis Lions, Paris

The Master Equation and the Convergence Problem in Mean Field Games

Chapter One

Introduction

1.1 FROM THE NASH SYSTEM TO THE MASTER EQUATION

Game theory formalizes interactions between “rational” decision makers. Its applications are numerous and range from economics and biology to computer science. In this monograph we are interested mainly in noncooperative games, that is, in games in which there is no global planner: each player pursues his or her own interests, which are partly conflicting with those of others.

In noncooperative game theory, the key concept is that of Nash equilibria, introduced by Nash in [82]. A Nash equilibrium is a choice of strategies for the players such that no player can benefit by changing strategies while the other players keep theirs unchanged. This notion has proved to be particularly relevant and tractable in games with a small number of players and action sets. However, as soon as the number of players becomes large, it seems difficult to implement in practice, because it requires that each player knows the strategies the other players will use. Besides, for some games, the set of Nash equilibria is huge and it seems difficult for the players to decide which equilibrium they are going to play: for instance, in repeated games, the Folk theorem states that the set of Nash equilibria coincides with the set of feasible and individually rational payoffs in the one-shot game, which is a large set in general (see [93]).

In view of these difficulties, one can look for configurations in which the notion of Nash equilibria simplifies. As noticed by Von Neumann and Morgenstern [96], one can expect that this is the case when the number of players becomes large and each player individually has a negligible influence on the other players: it “is a well known phenomenon in many branches of the exact and physical sciences that very great numbers are often easier to handle than those of medium size [. . .]. This is of course due to the excellent possibility of applying the laws of statistics and probabilities in the first case” (p. 14). Such *nonatomic games* were analyzed in particular by Aumann [10] in the framework of cooperative games. Schmeidler [91] (see also Mas-Colell [78]) extended the notion of Nash equilibria to that setting and proved the existence of pure Nash equilibria.

In the book we are interested in games with a continuum of players, in continuous time and continuous state space. Continuous time, continuous space games are often called *differential games*. They appear in optimal control problems in which the system is controlled by several agents. Such problems (for a

finite number of players) were introduced at about the same time by Isaacs [59] and Pontryagin [87]. Pontryagin derived optimality conditions for these games. Isaacs, working on specific examples of two-player zero-sum differential games, computed explicitly the solution of these games and established the formal connection with the Hamilton–Jacobi equations. The rigorous justification of Isaacs ideas for general systems took some time. The main difficulty arose from the set of strategies (or from the dependence on the cost of the players with respect to these strategies), which is much more complex than for classical games: indeed, the players have to observe the actions taken by the other players in continuous time and choose their instantaneous actions accordingly. For two-player, zero-sum differential games, the first general existence result of a Nash equilibrium was established by Fleming [39]: in this case the Nash equilibrium is unique and is called the value function (it is a function of time and space). The link between this value function and the Hamilton–Jacobi equations was made possible by the introduction of viscosity solutions by Crandall and Lions [32] (see also [33] for a general presentation of viscosity solutions). The application to zero-sum differential games are due to Evans and Souganidis [35] (for determinist problems) and Fleming and Souganidis [40] (for stochastic ones).

For non-zero-sum differential games, the situation is more complicated. One can show the existence of general Nash equilibria thanks to an adaptation of the Folk theorem: see Kononenko [64] (for differential games of first order) and Buckdahn, Cardaliaguet, and Rainer [23] (for differential games with diffusion). However, this notion of solution does not allow for dynamic programming: it lacks time consistency in general. The existence of time-consistent Nash equilibria, based on dynamic programming, requires the solvability of a strongly coupled system of Hamilton–Jacobi equations. This system, which plays a key role in this book, is here called *the Nash system*. For problems without diffusions, Bressan and Shen explain in [21, 22] that the Nash system is ill-posed in general. However, for systems with diffusions, the Nash system becomes a uniformly parabolic system of partial differential equations. Typically, for a game with N players and with uncontrolled diffusions, this backward in time system takes the form

$$\begin{cases} -\partial_t v^i(t, \mathbf{x}) - \text{tr}(a^i(t, \mathbf{x}) D^2 v^N(t, \mathbf{x})) + \mathcal{H}^i(t, \mathbf{x}, Dv^1(t, \mathbf{x}), \dots, Dv^N(t, \mathbf{x})) = 0 \\ \quad \text{in } [0, T] \times (\mathbb{R}^d)^N, \quad i \in \{1, \dots, N\}, \\ v^i(T, \mathbf{x}) = G^i(\mathbf{x}) \quad \text{in } (\mathbb{R}^d)^N. \end{cases} \quad (1.1)$$

The foregoing system describes the evolution in time of the value function v^i of agent i ($i \in \{1, \dots, N\}$). This value function depends on the positions of all the players $\mathbf{x} = (x_1, \dots, x_N)$, x_i being the position of the state of player i . The second-order terms $\text{tr}(a^i(t, \mathbf{x}) D^2 v^N(t, \mathbf{x}))$ formalize the noises affecting the dynamics of agent i . The Hamiltonian $\mathcal{H}^i$ encodes the cost player i has to pay

to control her state and reaching some goal. This cost depends on the positions of the other players and on their strategies.

The relevance of such a system for differential games has been discussed by Star and Ho [94] and Case [30] (for first-order systems) and by Friedman [43] (1972) (for second-order systems); see also the monograph by Başar and Olsder [11] and the references therein. The well-posedness of this system has been established under some restrictions on the regularity and the growth of the Hamiltonians: See in particular the monograph by Ladyženskaja, Solonnikov, and Ural'ceva [70] and the paper by Bensoussan and Frehse [14].

As for classical games, it is natural to investigate the limit of differential games as the number of players tends to infinity. The hope is that in this limit configuration the Nash system simplifies. This notion makes sense only for time-consistent Nash equilibria, because no simplification occurs in the framework of Folk's theorem, where the player who deviates is punished by all the other players.

Games in continuous space with infinitely many players were first introduced in the economic literature (in discrete time) under the terminology of heterogeneous models. The aim was to formalize dynamic general equilibria in macroeconomics by taking into account not only aggregate variables—GDP, employment, the general price level, for example—but also the distributions of variables, say the joint distribution of income and wealth or the size distribution of firms, and to try to understand how these variables interact. We refer in particular to the pioneering works by Aiyagari [6], Huggett [58], and Krusell and Smith [65], as well as the presentation of the continuous-time counterpart of these models in [5].

In the mathematical literature, the theory of differential games with infinitely many players, known as mean field games (MFGs), started with the works of Lasry and Lions [71, 72, 74]; Huang, Caines, and Malhamé [53–57] presented similar models under the name of the certainty equivalence principle. Since then the literature has grown very quickly, not only for the theoretical aspects, but also for the numerical methods and the applications: we refer to the monographs [16, 48] or the survey paper [49].

This book focuses mainly on the derivation of the MFG models from games with a finite number of players. In classical game theory, the rigorous link between the nonatomic games and games with a large but finite number of agents is quite well-understood: one can show (1) that limits of Nash equilibria as the number of agents tends to infinity is a Nash equilibrium of the nonatomic game (Green [50]), and (2) that any optimal strategy in the nonatomic game provides an ϵ -Nash equilibrium in the game with finitely many players, provided the number of players is sufficiently large (Rashid [90]).

For MFGs, the situation is completely different. If the equivalent of question (2) is pretty well understood, problem (1) turns out to be surprisingly difficult. Indeed, passing from the MFG equilibria to the differential game with finitely many players relies mostly on known techniques in mean field theory: this has been developed since the beginning of the theory in [54] and well studied since then (see also, for instance, [25, 62]). On the contrary, when one considers

a sequence of solutions to the Nash systems with N players and one wants to let N tend to infinity, the problem becomes extremely intricate. The main reason is that, in classical game theory, this convergence comes from compactness properties of the problem; this compactness is completely absent for differential games. This issue is related to the difficulty of building time-consistent solutions for these games. A less technical way to see this is to note that there is a change of nature between the Nash system and its conjectured limit, the MFG. In the Nash system, the players observe each other, and the deviation of a single player could a priori change entirely the outcome of the game. On the contrary, in the MFG, players react only to the evolving population density and therefore the deviation of a single player has no impact at all on the system. The main purpose of this book is to explain why this limit holds despite this change of nature.

1.1.1 Statement of the Problem

To explain our result further, we first need to specify the Nash system we are considering. We assume that players control their own state and interact only through their cost function. Then the Nash system (1.1) takes the more specific form:

$$\begin{cases} -\partial_t v^{N,i}(t, \mathbf{x}) - \sum_{j=1}^N \Delta_{x_j} v^{N,i}(t, \mathbf{x}) - \beta \sum_{j,k=1}^N \text{Tr} D_{x_j, x_k}^2 v^{N,i}(t, \mathbf{x}) \\ \quad + H(x_i, D_{x_i} v^{N,i}(t, \mathbf{x})) + \sum_{j \neq i} D_p H(x_j, D_{x_j} v^{N,j}(t, \mathbf{x})) \cdot D_{x_j} v^{N,i}(t, \mathbf{x}) \\ \quad = F^{N,i}(\mathbf{x}) & \text{in } [0, T] \times (\mathbb{R}^d)^N, \\ v^{N,i}(T, \mathbf{x}) = G^{N,i}(\mathbf{x}) & \text{in } (\mathbb{R}^d)^N. \end{cases} \quad (1.2)$$

As before, the above system is stated in $[0, T] \times (\mathbb{R}^d)^N$, where a typical element is denoted by $(t, \mathbf{x})$ with $\mathbf{x} = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N$. The unknowns are the N maps $(v^{N,i})_{i \in \{1, \dots, N\}}$ (the value functions). The data are the Hamiltonian $H : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, the maps $F^{N,i}, G^{N,i} : (\mathbb{R}^d)^N \rightarrow \mathbb{R}$, the nonnegative parameter β , and the horizon $T \geq 0$. In the second line, the symbol $\cdot$ denotes the inner product in $\mathbb{R}^d$.

System (1.2) describes the Nash equilibria of an N -player differential game (see Section 1.2 for a short description). In this game, the set of “optimal trajectories” solves a system of N coupled stochastic differential equations (SDEs):

$$\begin{aligned} dX_{i,t} &= -D_p H(X_{i,t}, Dv^{N,i}(t, \mathbf{X}_t))dt + \sqrt{2} dB_t^i + \sqrt{2\beta} dW_t, \\ t &\in [0, T], \quad i \in \{1, \dots, N\}, \end{aligned} \quad (1.3)$$

where $v^{N,i}$ is the solution to (1.2) and the $((B_t^i)_{t \in [0, T]})_{i=1, \dots, N}$ and $(W_t)_{t \in [0, T]}$ are d -dimensional independent Brownian motions. The Brownian motions $((B_t^i)_{t \in [0, T]})_{i=1, \dots, N}$ correspond to the *individual noises*, while the Brownian

motion $(W_t)_{t \in [0, T]}$ is the same for all the equations and, for this reason, is called the *common noise*. Under such a probabilistic point of view, the collection of random processes $((X_{i,t})_{t \in [0, T]})_{i=1, \dots, N}$ forms a dynamical system of interacting particles.

The aim of this book is to understand the behavior, as N tends to infinity, of the value functions $v^{N,i}$. Another, but closely related, objective of our book is to study the mean field limit of the $((X_{i,t})_{t \in [0, T]})_{i=1, \dots, N}$ as N tends to infinity.

1.1.2 Link with the Mean Field Theory

Of course, there is no chance to observe a mean field limit for (1.3) under a general choice of the coefficients in (1.2). Asking for a mean field limit certainly requires that the system has a specific symmetric structure in such a way that the players in the differential game are somewhat exchangeable (when in equilibrium). For this purpose, we suppose that, for each $i \in \{1, \dots, N\}$, the maps $(\mathbb{R}^d)^N \ni \mathbf{x} \mapsto F^{N,i}(\mathbf{x})$ and $(\mathbb{R}^d)^N \ni \mathbf{x} \mapsto G^{N,i}(\mathbf{x})$ depend only on x_i and on the empirical distribution of the variables $(x_j)_{j \neq i}$:

$$F^{N,i}(\mathbf{x}) = F(x_i, m_{\mathbf{x}}^{N,i}) \quad \text{and} \quad G^{N,i}(\mathbf{x}) = G(x_i, m_{\mathbf{x}}^{N,i}), \quad (1.4)$$

where $m_{\mathbf{x}}^{N,i} = \frac{1}{N-1} \sum_{j \neq i} \delta_{x_j}$ is the empirical distribution of the $(x_j)_{j \neq i}$ and where $F, G : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ are given functions, $\mathcal{P}(\mathbb{R}^d)$ being the set of Borel probability measures on $\mathbb{R}^d$. Under this assumption, the solution of the Nash system indeed enjoys strong symmetry properties, which imply in particular the required exchangeability property. Namely, $v^{N,i}$ can be written in a form similar to (1.4):

$$v^{N,i}(t, \mathbf{x}) = v^N(t, x_i, m_{\mathbf{x}}^{N,i}), \quad t \in [0, T], \quad \mathbf{x} \in (\mathbb{R}^d)^N, \quad (1.5)$$

for a function $v^N(t, \cdot, \cdot)$ taking as arguments a state in $\mathbb{R}^d$ and an empirical distribution of size $N - 1$ over $\mathbb{R}^d$.

In any case, even under the foregoing symmetry assumptions, it is by no means clear whether the system (1.3) can exhibit a mean field limit. The reason is that the dynamics of the particles $(X_{1,t}, \dots, X_{N,t})_{t \in [0, T]}$ are coupled through the unknown solutions $v^{N,1}, \dots, v^{N,N}$ to the Nash system (1.2), whose symmetry properties (1.5) may not suffice to apply standard results from the theory of propagation of chaos. Obviously, the difficulty is that the function v^N on the right-hand side of (1.5) precisely depends on N . Part of the challenge in the text is thus to show that the interaction terms in (1.3) get closer and closer, as N tends to the infinity, to some interaction terms with a much more tractable and much more explicit shape.

To get a picture of the ideal case under which the mean-field limit can be taken, one can choose for a while $\beta = 0$ in (1.3) and then assume that the function v^N in the right-hand side of (1.5) is independent of N . Equivalently, one can replace in (1.3) the interaction function $(\mathbb{R}^d)^N \ni \mathbf{x} \mapsto D_p H(x_i, v^{N,i}(t, \mathbf{x}))$

by $(\mathbb{R}^d)^N \ni \mathbf{x} \mapsto b(x_i, m_{\mathbf{x}}^{N,i})$, for a map $b : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \mapsto \mathbb{R}^d$. In such a case, the coupled system of SDEs (1.3) turns into

$$dX_{i,t} = b\left(X_{i,t}, \frac{1}{N-1} \sum_{j \neq i} \delta_{X_{j,t}}\right) dt + \sqrt{2} dB_t^i, \quad t \in [0, T], \quad i \in \{1, \dots, N\}, \quad (1.6)$$

the second argument in b being nothing but the empirical measure of the particle system at time t . Under suitable assumptions on b (e.g., if b is bounded and Lipschitz continuous in both variables, the space of probability measures being equipped with the Wasserstein distance) and on the initial distribution of the $((X_{i,t})_{i=1,\dots,N})_{t \in [0,T]}$, both the marginal law of $(X_t^1)_{t \in [0,T]}$ (or of any other player) and the empirical distribution of the whole system converge to the solution of the McKean–Vlasov equation:

$$\partial_t m - \Delta m + \operatorname{div}(m b(\cdot, m)) = 0.$$

(see, among many other references, McKean [77], Sznitman [92], Méléard [79]). The standard strategy for establishing the convergence consists in a coupling argument. Precisely, if one introduces the system of N independent equations

$$dY_{i,t} = b(Y_{i,t}, \mathcal{L}(Y_{i,t})) dt + \sqrt{2} dB_t^i, \quad t \in [0, T], \quad i \in \{1, \dots, N\},$$

(where $\mathcal{L}(Y_{i,t})$ is the law of $Y_{i,t}$ with the same (chaotic) initial condition as that of the processes $((X_{i,t})_{t \in [0,T]})_{i=1,\dots,N}$, then it is known that (under appropriate integrability conditions; see Fournier and Guillin [42])

$$\sup_{t \in [0,T]} \mathbb{E} [|X_{1,t} - Y_{1,t}|] \leqslant C N^{-\frac{1}{\max(2,d)}} (\mathbf{1}_{\{d \neq 2\}} + \ln(1+N) \mathbf{1}_{\{d=2\}}).$$

In comparison with (1.6), all the equations in (1.3) are subject to the common noise $(W_t)_{t \in [0,T]}$, at least when $\beta \neq 0$. This makes a first difference between our limit problem and the above McKean–Vlasov example of interacting diffusions, but, for the time being, it is not clear how deeply this may affect the analysis. Indeed, the presence of a common noise does not constitute a real challenge in the study of McKean–Vlasov equations, the foregoing coupling argument working in that case as well, provided that the distribution of Y is replaced by its conditional distribution given the realization of the common noise. However, the key point here is precisely that our problem is not formulated as a McKean–Vlasov equation, as the drifts in (1.3) are not of the same explicit mean field structure as they are in (1.6) because of the additional dependence on N in the right-hand side of (1.5): obviously this is the second main difference between (1.3) and (1.6). This makes rather difficult any attempt to guess the precise impact of the common noise on the analysis. Certainly, as we already pointed out, the major issue in analyzing (1.3) stems from the complex nature of the underlying interactions. As the equations depend on one another through the

nonlinear system (1.2), the evolution with N of the coupling between all of them is indeed much more intricate than in (1.6). And once again, on the top of that, the common noise adds another layer of difficulty. For these reasons, the convergence of both (1.2) and (1.3) has been an open question since Lasry and Lions' initial papers on MFGs [71, 72].

1.1.3 The Mean Field Game System

If one tries, at least in the simpler case $\beta = 0$, to describe—in a heuristic way—the structure of a differential game with infinitely many indistinguishable players, i.e., a “nonatomic differential game,” one finds a problem in which each (infinitesimal) player optimizes his payoff, depending on the collective behavior of the others, and, meanwhile, the resulting optimal state of each of them is exactly distributed according to the state of the population. This is the “mean field game system” (MFG system):

$$\begin{cases} -\partial_t u - \Delta u + H(x, D_x u) = F(x, m(t)) & \text{in } [0, T] \times \mathbb{R}^d, \\ \partial_t m - \Delta m - \operatorname{div}(m D_p H(x, D_x u)) = 0 & \text{in } [0, T] \times \mathbb{R}^d, \\ u(T, x) = G(x, m(T)), \quad m(0, \cdot) = m_{(0)} & \text{in } \mathbb{R}^d, \end{cases} \quad (1.7)$$

where $m_{(0)}$ denotes the initial state of the population. The system consists in a coupling between a (backward) Hamilton–Jacobi equation, describing the dynamics of the value function of any of the players, and a (forward) Kolmogorov equation, describing the dynamics of the distribution of the population. In that framework, H reads as a Hamiltonian, F is understood as a running cost, and G as a terminal cost. Since its simultaneous introduction by Lasry and Lions [74] and by Huang, Caines, and Malhamé [53], this system has been thoroughly investigated: its existence, under various assumptions, can be found in [15, 25, 54–56, 62, 74, 76]. Concerning uniqueness of the solution, two regimes were identified in [74]. Uniqueness holds under Lipschitz type conditions when the time horizon T is short (or, equivalently, when H , F , and G are “small”), but, as for finite-dimensional two-point boundary value problems, it may fail when the system is set over a time interval of arbitrary length. Over long time intervals, uniqueness is guaranteed under the quite fascinating condition that F and G are monotone; i.e., if, for any measures m, m' , the following holds:

$$\begin{aligned} & \int_{\mathbb{R}^d} (F(x, m) - F(x, m')) d(m - m')(x) \geq 0 \\ \text{and} \quad & \int_{\mathbb{R}^d} (G(x, m) - G(x, m')) d(m - m')(x) \geq 0. \end{aligned} \quad (1.8)$$

The interpretation of the monotonicity condition is that the players dislike congested areas and favor configurations in which they are more scattered; see Remark 2.3.1 for an example. Generally speaking, condition (1.8) plays a key

role throughout the text, as it guarantees not only uniqueness but also stability of the solutions to (1.7).

As observed, a solution to the MFG system (1.7) can indeed be interpreted as a Nash equilibrium for a differential game with infinitely many players: in that framework, it plays the role of the Schmeidler noncooperative equilibrium. A standard strategy to make the connection between (1.7) and differential games consists in inserting the optimal strategies from the Hamilton–Jacobi equation in (1.7) into finitely many player games in order to construct approximate Nash equilibria: see [54], as well as [25, 55, 56, 62]. However, although it establishes the interpretation of the system (1.7) as a differential game with infinitely many players, this says nothing about the convergence of (1.2) and (1.3).

When β is positive, the system describing Nash equilibria within a population of infinitely many players subject to the same common noise of intensity β cannot be described by a deterministic system of the same form as (1.7). Owing to the theory of propagation of chaos for systems of interacting particles (see the short remark earlier), the unknown m in the forward equation is then expected to represent the conditional law of the optimal state of any player given the realization of the common noise. In particular, it must be random. This turns the forward Kolmogorov equation into a forward stochastic Kolmogorov equation. As the Hamilton–Jacobi equation depends on m , it renders u random as well. At any rate, a key fact from the theory of stochastic processes is that the solution to an SDE must be adapted to the underlying observation, as its values at some time t cannot anticipate the future of the noise after t . At first sight, it seems to be very demanding, as u is also required to match, at time T , $G(\cdot, m(T))$, which depends on the whole realization of the noise up until T . The correct formulation to accommodate both constraints is given by the theory of backward SDEs, which suggests penalizing the backward dynamics by a martingale in order to guarantee that the solution is indeed adapted. We refer the reader to the monograph [84] for a complete account on the finite dimensional theory and to the paper [85] for an insight into the infinite dimensional case. Denoting by W “the common noise” (here, a d -dimensional Brownian motion) and by $m_{(0)}$ the initial distribution of the players at time t_0 , the MFG system with common noise then takes the form (in which the unknowns are now (u_t, m_t, v_t))

$$\begin{cases} d_t u_t = [-(1 + \beta)\Delta u_t + H(x, D_x u_t) - F(x, m_t) - \sqrt{2\beta} \operatorname{div}(v_t)] dt \\ \quad + v_t \cdot dW_t, & \text{in } [0, T] \times \mathbb{R}^d, \\ d_t m_t = [(1 + \beta)\Delta m_t + \operatorname{div}(m_t D_p H(x, D_x u_t))] dt \\ \quad - \operatorname{div}(m_t \sqrt{2\beta} dW_t), & \text{in } [0, T] \times \mathbb{R}^d, \\ u_T(x) = G(x, m_T), \quad m_0 = m_{(0)}, & \text{in } \mathbb{R}^d \end{cases} \quad (1.9)$$

where we used the standard convention from the theory of stochastic processes that consists in indicating the time parameter as an index in random functions. As suggested immediately above, the map v_t is a random vector field that forces

the solution u_t of the backward equation to be adapted to the filtration generated by $(W_t)_{t \in [0, T]}$. As far as we know, the system (1.9) has never been investigated and part of this book will be dedicated to its analysis (see, however, [27] for an informal discussion). Below, we call the system (1.9) the *MFG system with common noise*.

Note that the aggregate equations (1.7) and (1.9) (see also the master equation (1.10)) are the continuous-time analogues of equations that appear in the analysis of dynamic stochastic general equilibria in heterogeneous agent models (Aiyagari [6], Bewley [19], and Huggett [58]). In this setting, the factor β describes the intensity of “aggregate shocks,” as discussed by Krusell and Smith in the seminal paper [65]. In some sense, the limit problem studied in the text is an attempt to deduce the macroeconomic models, describing the dynamics of a typical (but heterogeneous) agent in an equilibrium configuration, from the microeconomic ones (the Nash equilibria).

1.1.4 The Master Equation

Although the MFG system has been widely studied since its introduction in [74] and [53], it has become increasingly clear that this system was not sufficient to take into account the entire complexity of dynamic games with infinitely many players. A case in point is that the original system (1.7) becomes much more complex in the presence of a common noise (i.e., when $\beta > 0$); see the stochastic version (1.9). In the same spirit, we may notice that the original MFG system (1.7) does not accommodate MFGs with a major player and infinitely many small players; see [52]. And, last but not least, the main limitation is that, so far, the formulation based on the system (1.7) (or (1.9) when $\beta > 0$) has not allowed establishment of a clear connection with the Nash system (1.2).

These issues led Lasry and Lions [76] to introduce an infinite dimensional equation—the so-called “master equation”—that directly describes, at least formally, the limit of the Nash system (1.2) and encompasses the foregoing complex situations. Before writing down this equation, let us explain its main features. One of the key observations has to do with the symmetry properties, to which we already alluded, that are satisfied by the solution of the Nash system (1.2). Under the standing symmetry assumptions (1.4) on the $(F^{N,i})_{i=1,\dots,N}$ and $(G^{N,i})_{i=1,\dots,N}$, (1.5) says that the $(v^{N,i})_{i=1,\dots,N}$ can be written into a form similar to (1.4), namely $v^{N,i}(t, \mathbf{x}) = v^N(t, x_i, m_{\mathbf{x}}^{N,i})$ (where the empirical measures $m_{\mathbf{x}}^{N,i}$ are defined as in (1.4)), but with the obvious but major restriction that the function v^N that appears on the right-hand side of the equality now depends on N . With such a formulation, the value function to player i reads as a function of the private state of player i and of the empirical distribution formed by the others. Then, one may guess, at least under the additional assumption that such a structure is preserved as $N \rightarrow +\infty$, that the unknown in the limit problem takes the form $U = U(t, x, m)$, where x is the position of the (typical) small player at time t and m is the distribution of the (infinitely many) other agents.

The question is then to write down the dynamics of U . Plugging $U = U(t, x_i, m_x^{N,i})$ into the Nash system (1.2), one obtains—at least formally—an equation stated in the space of measures (see Section 1.2 for a heuristic discussion). This is the so-called master equation. It takes the form

$$\begin{cases} -\partial_t U - (1 + \beta)\Delta_x U + H(x, D_x U) - (1 + \beta) \int_{\mathbb{R}^d} \operatorname{div}_y [D_m U] dm(y) \\ \quad + \int_{\mathbb{R}^d} D_m U \cdot D_p H(y, D_x U(\cdot, y, \cdot)) dm(y) \\ - 2\beta \int_{\mathbb{R}^d} \operatorname{div}_x [D_m U] dm(y) - \beta \int_{\mathbb{R}^{2d}} \operatorname{Tr} [D_{mm}^2 U] dm^{\otimes 2}(y, y') \\ = F(x, m) & \text{in } [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d), \\ U(t, x, m) = G(x, m) & \text{in } \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d), \end{cases} \quad (1.10)$$

where $\partial_t U$, $D_x U$, and $\Delta_x U$ are understood as $\partial_t U(t, x, m)$, $D_x U(t, x, m)$, and $\Delta_x U(t, x, m)$; $D_x U(\cdot, y, \cdot)$ is understood as $D_x U(t, y, m)$; and $D_m U$ and $D_{mm}^2 U$ are understood as $D_m(t, x, m, y)$ and $D_{mm}^2 U(t, x, m, y, y')$.

In Eq. (1.10), $\partial_t U$, $D_x U$, and $\Delta_x U$ stand for the usual time derivative, space derivatives, and Laplacian with respect to the local variables (t, x) of the unknown U , while $D_m U$ and $D_{mm}^2 U$ are the first- and second-order derivatives with respect to the measure m . The precise definition of these derivatives is postponed to Chapter 2. For the time being, let us just note that it is related to the derivatives in the space of probability measures described, for instance, by Ambrosio, Gigli, and Savaré in [7] and by Lions in [76]. It is worth mentioning that the master equation (1.10) is not the first example of an equation studied in the space of measures—by far: for instance, Otto [83] gave an interpretation of the porous medium equation as an evolution equation in the space of measures, and Jordan, Kinderlehrer, and Otto [60] showed that the heat equation was also a gradient flow in that framework; notice also that the analysis of Hamilton–Jacobi equations in metric spaces is partly motivated by the specific case in which the underlying metric space is the space of measures (see in particular [8, 36] and the references therein). The master equation is, however, the first one to combine at the same time the issue of being nonlocal, nonlinear, and of second order and, moreover, without maximum principle.

Besides the discussion in [76], the importance of the master equation (1.10) has been acknowledged by several contributions: see, for instance, the monograph [16] and the companion papers [17] and [18], in which Bensoussan, Frehse, and Yam generalize this equation to mean field type control problems and reformulate it as a partial differential equation (PDE) set on an L^2 space, and [27], where Carmona and Delarue interpret this equation as a decoupling field of forward–backward SDE in infinite dimension.

If the master equation has been discussed and manipulated thoroughly in the aforementioned references, it is mostly at a formal level: the well-posedness of the master equation has remained, to a large extent, open until now. Besides,