Probability Theory and Stochastic Modelling 83

René Carmona
François Delarue

Probabilistic Theory of Mean Field Games with Applications I

Mean Field FBSDEs, Control, and Games

Springer

Probability Theory and Stochastic Modelling

Volume 83

The **Probability Theory and Stochastic Modelling** series is a merger and continuation of Springer's two well established series Stochastic Modelling and Applied Probability and Probability and Its Applications series. It publishes research monographs that make a significant contribution to probability theory or an applications domain in which advanced probability methods are fundamental. Books in this series are expected to follow rigorous mathematical standards, while also displaying the expository quality necessary to make them useful and accessible to advanced students as well as researchers. The series covers all aspects of modern probability theory including

- Gaussian processes
- Markov processes
- Random fields, point processes and random sets
- Random matrices
- Statistical mechanics and random media
- Stochastic analysis

as well as applications that include (but are not restricted to):

- Branching processes and other models of population growth
- Communications and processing networks
- Computational methods in probability and stochastic processes, including simulation
- Genetics and other stochastic models in biology and the life sciences
- Information theory, signal processing, and image synthesis
- Mathematical economics and finance
- Statistical methods (e.g. empirical processes, MCMC)
- Statistics for stochastic processes
- Stochastic control
- Stochastic models in operations research and stochastic optimization
- Stochastic models in the physical sciences

More information about this series at http://www.springer.com/series/13205

René Carmona • François Delarue

Probabilistic Theory of Mean Field Games with Applications I

Mean Field FBSDEs, Control, and Games

Springer

René Carmona
ORFE Department
Program in Applied and
Computational Mathematics
Princeton University
Princeton, NJ, USA

François Delarue
Institut Universitaire de France
& Laboratoire J.A. Dieudonné
Université Nice Sophia Antipolis
Nice, France

ISSN 2199-3130 ISSN 2199-3149 (electronic)
Probability Theory and Stochastic Modelling
ISBN 978-3-319-56437-1 ISBN 978-3-319-58920-6 (eBook)
DOI 10.1007/978-3-319-58920-6

Library of Congress Control Number: 2017940988

Mathematics Subject Classification (2010): 91A13, 91A06, 91A55, 49N70, 49N90, 60H10, 60H15, 60H30

Printed on acid-free paper

This Springer imprint is published by Springer Nature
The registered company is Springer International Publishing AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Foreword

Since its inception about a decade ago, the theory of Mean Field Games has rapidly developed into one of the most significant and exciting sources of progress in the study of the dynamical and equilibrium behavior of large systems. The introduction of ideas from statistical physics to identify approximate equilibria for sizeable dynamic games created a new wave of interest in the study of large populations of competitive individuals with "mean field" interactions. This two-volume book grew out of series of lectures and short courses given by the authors over the last few years on the mathematical theory of Mean Field Games and their applications in social sciences, economics, engineering and finance. While this is indeed the object of the book, by taste, background, and expertise, we chose to focus on the probabilistic approach to these game models.

In a trailblazing contribution, Lasry and Lions proposed in 2006 a methodology to produce approximate Nash equilibria for stochastic differential games with symmetric interactions and a large number of players. In their models, a given player *feels* the presence and the behavior of the other players through the empirical distribution of their private states. This type of interaction was extensively studied in the statistical physics literature under the name of *mean field* interaction, hence the terminology Mean Field Game coined by Lasry and Lions. The theory of these new game models was developed in lectures given by Pierre-Louis Lions at the Collège de France which were video-taped and made available on the internet. Simultaneously, Caines, Huang, and Malhamé proposed a similar approach to large games under the name of Nash Certainty Equivalence principle. This terminology fell from grace and the standard reference to these game models is now Mean Field Games.

While slow to pick up momentum, the subject has seen a renewed wave of interest over the last seven years. The mean field game paradigm has evolved from its seminal principles into a full-fledged field attracting theoretically inclined investigators as well as applied mathematicians, engineers, and social scientists. The number of lectures, workshops, conferences, and publications devoted to the subject has grown exponentially, and we thought it was time to provide the applied mathematics community interested in the subject with a textbook presenting the state of the art, as we see it. Because of our personal taste, we chose to focus on what we like to call the probabilistic approach to mean field games. While a significant portion of the text is based on original research by the authors, great care was taken to include models and results contributed by others, whether or not they were aware of the fact they were working with mean field games. So the book should feel and read like a textbook, not a research monograph.

Most of the material and examples found in the text appear for the first time in book form. In fact, a good part of the presentation is original, and the lion's share of the arguments used in the text have been designed especially for the purpose of the book. Our concern for pedagogy justifies (or at least explains) why we chose to divide the material in two volumes and present mean field games without a common noise first. We ease the introduction of the technicalities needed to treat models with a common noise in a crescendo of sophistication in the complexity of the models. Also, we included at the end of each volume four extensive indexes (author index, notation index, subject index, and assumption index) to make navigation throughout the book seamless.

Acknowledgments

First and foremost, we want to thank our wives Debbie and Mélanie for their understanding and unwavering support. The intensity of the research collaboration which led to this two-volume book increased dramatically over the years, invading our academic lives as well as our social lives, pushing us to the brink of sanity at times. We shall never be able to thank them enough for their patience and tolerance. This book project would not have been possible without them: our gratitude is limitless.

Next we would like to thank Pierre-Louis Lions, Jean-Michel Lasry, Peter Caines, Minyi Huang, and Roland Malhamé for their incredible insight in introducing the concept of mean field games. Working independently on both sides of the pond, their original contributions broke the grounds for an entirely new and fertile field of research. Next in line is Pierre Cardaliaguet, not only for numerous private conversations on game theory but also for the invaluable service provided by the notes he wrote from Pierre-Louis Lions' lectures at the Collège de France. Although they were never published in printed form, these notes had a tremendous impact on the mathematical community trying to learn about the subject, especially at a time when writings on mean field games were few and far between.

We also express our gratitude to the organizers of the 2013 and 2015 conferences on mean field games in Padova and Paris: Yves Achdou, Pierre Cardaliaguet, Italo Capuzzo-Dolcetta, Paolo Dai Pra, and Jean-Michel Lasry.

While we like to cast ourselves as proponents of the probabilistic approach to mean field games, it is fair to say that we were far from being the only ones following this path. In fact, some of our papers were posted essentially at the same time as papers of Bensoussan, Frehse, and Yam, addressing similar questions, with the same type of methods. We benefitted greatly from this stimulating and healthy competition.

We also thank our coauthors, especially Jean-François Chasagneux, Dan Crisan, Jean-Pierre Fouque, Daniel Lacker, Peiqi Wang, and Geoffrey Zhu. We used our joint works as the basis for parts of the text which they will recognize easily.

Also, we would like to express our gratitude to the many colleagues and students who gracefully tolerated our relentless promotion of this emerging field of research through courses, seminar, and lecture series. In particular, we would like to thank Jean-François Chasagneux, Rama Cont, Dan Crisan, Romuald Elie, Josselin Garnier, Marcel Nutz, Huyen Pham, and Nizar Touzi for giving us the opportunity to do just that.

Finally, we would like to thank Paul-Eric Chaudru de Raynal, Mark Cerenzia, Christy Graves, Dan Lacker, Victor Marx, and Peiqi Wang for their help in proofreading an earlier version of the manuscript.

The work of R. Carmona was partially supported by NSF DMS-1211928 and DMS-1716673, and by ARO W911NF-17-1-0578. The work of F. Delarue was partially supported by the Institut Universitaire de France and by the Agence Nationale de la Recherche (project MFG, ANR-16-CE40-0015-01).

Princeton, NJ, USA — René Carmona
Nice, France — François Delarue
July 29, 2016

Preface to Volume I

This first volume of the book is entirely devoted to the theory of mean field games in the absence of a source of random shocks common to all the players. We call these models *games without a common noise*. This volume is divided into two main parts. Part I is a self-contained introduction to mean field games, starting from practical applications and concrete illustrations, and ending with ready-for-use solvability results for mean field games without a common noise. While Chapters 1 and 2 are mostly dedicated to games with a finite number of players, the asymptotic formulation which constitutes the core of the book is introduced in Chapter 3. For the exposition to be as pedagogical as possible, we chose to defer some of the technical aspects of this asymptotic formulation to Chapter 4 which provides a complete toolbox for solving forward-backward stochastic differential equations of the McKean-Vlasov type. Part II has a somewhat different scope and focuses on the main principles of analysis on the Wasserstein space of probability measures with a finite second moment, which plays a key role in the study of mean field games and which will be intensively used in the second volume of the book. We present the mathematical theory in Chapter 5, and we implement its results in Chapter 6 with the analysis of stochastic mean field control problems, which are built upon a notion of equilibrium different from the search for Nash equilibria at the root of the definition of mean field games. Extensions, including infinite time horizon models and games with finite state spaces, are discussed in the epilogue of this first volume.

The remainder of this preface expands, chapter by chapter, the short content summary given above. A diagram summarizing the connections between the different components of the book is provided on page xix.

The first chapter sets the stage for the introduction of mean field games with a litany of examples of increasing complexity. Starting with one-period deterministic games with a large number of players, we introduce the mean field game paradigm. We use examples from domains as diverse as finance, macroeconomics, population biology, and social science to motivate the introduction of mean field games in different mathematical settings. Some of these examples were studied in the literature before the introduction of, and without any reference to, mean field games. We chose them because of their powerful illustrative power and the motivation they offer for the introduction of new mathematical models. The examples of *bank runs* modeled as mean field games of timing are a case in point. For pedagogical reasons, we highlight practical applications where the interaction between the players does not necessarily enter the model through the empirical distributions of the states of the players, but via the empirical distributions of the actions of the players, or even the joint empirical distributions of the states and the controls of the players. Most of these examples will be revisited and solved throughout the book.

Chapter 2 offers a crash course on the mathematical theory of stochastic differential games with a finite number of players. The material of this chapter is not often found in book form, and since we make extensive use of its notations and results throughout the book, we thought it was important to present them early for the sake of completeness and future references. We concentrate on what we call the probabilistic approach to the search for Nash equilibria, and we introduce games with mean field interactions as they are the main object of the book. Explicitly solvable models are few and far between. Among them, linear quadratic (LQ for short) models play a very special role because their solutions, when they exist, can be obtained by solving matrix Riccati equations. The last part of the chapter is devoted to a detailed analysis of a couple of linear quadratic models already introduced in Chapter 1, and for which explicit solutions can be derived. To wit, these models do not require the theory of mean field games since their finite player versions can be solved explicitly. However, they provide a testbed for the analysis of the limit of finite player equilibria when the number of players grows without bound, offering an invaluable opportunity to introduce the concept of mean field game and discover some of its essential features.

The probabilistic approach to mean field games is the main thrust of the book. The underpinnings of this approach are presented in Chapter 3. Stochastic control problems and the search for equilibria for stochastic differential games can be tackled by reformulating the optimization and equilibrium problems in terms of backward stochastic differential equations (BSDEs throughout the book) and forward-backward stochastic differential equations (FBSDEs for short). In this chapter, we review the major forms of FBSDEs that may be used to represent the optimal trajectories of a standard optimization problem: the first one is based on a probabilistic representation of the value function, and the second one on the stochastic Pontryagin maximum principle. Combined with the consistency condition issued from the search for Nash equilibria as fixed points of the best response function, this prompts us to introduce a new class of FBSDEs with a distinctive McKean-Vlasov character. This chapter presents a basic existence result for McKean-Vlasov FBSDEs. This result will be extended in Chapter 4. As a by-product, we obtain early solvability results for mean field games by straightforward implementations of the two forms of the probabilistic approach just mentioned. However, since our primary aim in this chapter is to make the presentation as pedagogical as possible, we postpone the most general versions of the existence results for mean field games to Chapter 4, as some of the proofs are rather technical. Instead, we highlight the role of monotonicity, as captured by the so-called Lasry-Lions monotonicity conditions, in the analysis of uniqueness of equilibria. Finally, we specialize the results of this chapter to the case of linear-quadratic mean field games, which can be handled directly via the analysis of Riccati equations. Most of the results of this chapter will be revisited and extended in the second volume to accommodate a *common noise* which is found in many economic and physical applications.

Chapter 4 starts with a stochastic analysis primer on the theory of FBSDEs. As explained above, in the mean field limit of large games, the fixed point step of the search for Nash equilibria turns the standard FBSDEs derived from optimization problems into equations of the McKean-Vlasov type by introducing the distribution of the solution into the coefficients. These FBSDEs characterize the equilibria. Since this new class of FBSDEs was not studied before the advent of mean field games, one of the main objectives of Chapter 4 is to provide a systematic approach to their solution. We show how to use Schauder's fixed point theorem to prove existence of a solution. The chapter closes with the analysis of the so-called *extended mean field games*, in which the players are interacting not only through the distribution of their states but also through the distribution of their controls. Finally, we demonstrate how the methodology developed in the chapter applies to some of the examples presented in the opening chapter.

Although it contains very few results on mean field games, Chapter 5 plays a pivotal role in the book. It contains all the results on spaces of probability measures which we use throughout the book, including the definitions and properties of the Wasserstein distances, the convergence of the empirical measures of a sequence of independent and identically distributed random variables ... and most importantly, a detailed presentation of the differential calculus on the Wasserstein space introduced by Lions in his unpublished lectures at the Collège de France, and by Cardaliaguet in the notes he wrote from Lions' lectures. Even though the use of this differential calculus in the first volume is limited to the ensuing Chapter 6, the differential calculus on the Wasserstein space plays a fundamental role in the study of the master equation for mean field games, whose presentation and analysis will be provided in detail in the second volume. Still, a foretaste of the master equation is given at the end of this chapter. Its derivation is based on an original form of Itô's formula for functionals of the marginal laws of an Itô process, the proof of which is given in full detail. For the sake of completeness, we also provide a thorough and enlightening discussion of the connections between Lions' differential calculus, which we call L-differential calculus throughout the book, and other forms of differential calculus on the space of probability measures, among which the differential calculus used in optimal transportation theory.

One of the remarkable features of the construction of solutions to mean field game problems is the similarity with a natural problem which did not get much attention from analysts and probabilists: the optimal control of (stochastic) differential equations of the McKean-Vlasov type, which could also be called mean field optimal control. The latter is studied in Chapter 6. Both problems can be interpreted as searches of equilibria for large populations, claim which will be substantiated in Chapter 6 in the second volume of the book. Interestingly, the optimal control of McKean-Vlasov stochastic dynamics is intrinsically a stochastic optimization problem while the search for Nash equilibria in mean field games is more of a fixed point problem than an optimization problem. So despite the strong similarities between the two problems, differences do exist, and we highlight them starting with Chapter 6. There, we show that since the problem at hand is a stochastic control problem, the optimal control of McKean-Vlasov stochastic dynamics can be tackled by means of an appropriate version of the Pontryagin stochastic maximum principle. Following this strategy leads to FBSDEs for which the backward part involves the derivative of the Hamiltonian with respect to a measure argument. This novel feature is handled with the tools provided in Chapter 5. We close the chapter with the discussion of an alternative strategy for solving mean field optimal control problems, based on the notion of relaxed controls. Also, we review several crucial examples, among them potential games. These latter models are mean field games for which the solutions can be reduced to the solutions of mean field optimal control problems, and optimal transportation problems.

Chapter 7 is a capstone which we use to revisit some of the examples introduced in Chapter 1, especially those which are not exactly covered by the probabilistic theory of stochastic differential mean field games developed in the first volume. Indeed, Chapter 1 included a considerable amount of applications hinting at mathematical models with distinctive features which are not accommodated by the models and results of the first part of this first volume. We devote this chapter to presentations, even if only informal, of extensions of the Mean Field Game paradigm to these models. They include extensions to several homogenous populations, infinite horizon optimization, and models with finite state spaces. These mean field game models have a great potential for the quantitative analysis of very important practical applications, and we show how the technology developed in the first volume of the book can be brought to bear on their solutions.

Princeton, NJ, USA — René Carmona
Nice, France — François Delarue
July 29, 2016

Organization of the Book: Volume I Organigram

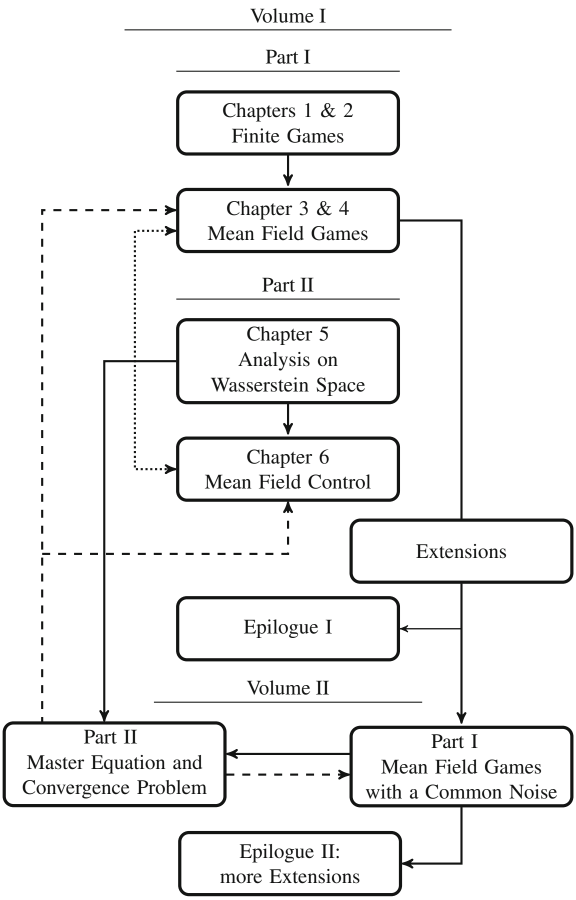

Thick lines indicate the logical order of the chapters. The dotted line between Chapters 3–4 and 6 emphasizes the fact that—in some cases like potential games—mean field games and mean field control problems share the same solutions. Finally, the dashed lines starting from Part II (second volume) point toward the games and the optimization problems for which we can solve approximately the finite-player versions or for which the finite-player equilibria are shown to converge.

References to the second volume appear in the text in the following forms: Chapter (Vol II)-X, Section (Vol II)-$X.x$, Theorem (Vol II)-$X.x$, Proposition (Vol II)-$X.x$, Equation (Vol II)-($X.x$), $\ldots$, where X denotes the corresponding chapter in the second volume and x the corresponding label.

Contents

Volume I