



Probabilistic Theory of Mean Field Games, Volumes I & II

Reviewed by Jean-Pierre Fouque

Probabilistic Theory of Mean Field Games with Applications: Volume I, Mean Field FBSDEs, Control, and Games; Volume II, Mean Field Games with Common Noise and Master Equations

By René Carmona and François Delarue

Systems evolving in continuous time (t in R_+) and in continuous space (x in R^d) that are subject to noise are modeled by diffusion processes which are obtained as solutions of stochastic differential equations of Itô type. These are differential equations containing a stochastic integral with respect to the celebrated Brownian motion (or Wiener process W_t) which models the noise whose time increments are commonly called white noise. In differential form, they are written as $dX_t = b(t, X_t)dt + a(t, X_t)dW_t$, where $b(t, x)$ is the drift and $a(t, x)$ the diffusion coefficient (or volatility in financial mathematics). The pillars of Itô's theory are stochastic integrals, with respect to the Brownian motion which is not of bounded variations, and the chain rule (Itô's formula) involving a second-order derivative term. The solution of an Itô's stochastic differential equation (SDE) is a continuous Markov process associated to a semigroup and an infinitesimal generator which is a second-order linear differential operator. The evolution of the density of the process is described by the forward Fokker–Planck–Kolmogorov equation involving the adjoint of the infinitesimal generator. A nice introduction to these objects and various applications, such as filtering, optimal stopping, stochastic

control, or mathematical finance, can be found in the popular graduate text by Øksendal [1].

In the context of stochastic control, one single agent controls the coefficients (the drift and eventually the volatility) of a diffusion process (as described above), in order to optimize an objective function.

The theory and applications of stochastic control for diffusion processes and the relation with nonlinear second-order parabolic partial differential equations have been vastly studied and now form a classical topic in probability theory (see the book by Fleming and Soner [2]).

The situation with N -player stochastic differential games is quite different. As opposed to the case of stochastic control with one agent, typically, the dynamics of the states of the N agents are given by a coupled system of stochastic differential equations of the Itô type driven by Brownian motions possibly correlated as in the case with a common noise. Each player, through her action affecting drift and possibly volatility terms, tries to minimize/maximize her objective function made of running and terminal costs/rewards. The information structure of the set of actions/strategies needs to be carefully defined (open-loop, closed-loop, ...) in order to specify the problem. There are several notions of equilibria, each of them lacking uniqueness in general and difficult to compute in particular when N is large.

The books under review are concerned with Nash equilibria. A Nash equilibrium is a set of strategies such that if any player tries to change her action, the others being fixed, then, she cannot be strictly better off. This equilibrium can also be understood as a best-response of the system (see Chapter 5 of the book by Carmona [3]).

In 2006, Lasry and Lions (see [4]) proposed a methodology to obtain approximate Nash equilibria for stochastic differential games between a large number of players

Jean-Pierre Fouque is a Distinguished Professor in the Department of Statistics and Applied Probability at the University of California Santa Barbara. His email address is fouque@pstat.ucsb.edu.

Communicated by Notices Book Review Editor Stephan Ramon Garcia.

For permission to reprint this article, please contact: reprint-permission@ams.org.

DOI: <https://dx.doi.org/10.1090/noti2089>

interacting in a symmetric way such that a given player is influenced by the others through the empirical distribution of their states. In the vast statistical physics literature, this type of interaction is known as mean field interaction. In the limit of infinite number of particles, the law of large numbers takes place, and particles become independent (a phenomenon known as propagation of chaos), and one given particle interacts with its own distribution leading to nonlinear Fokker–Planck equations called McKean–Vlasov equations (see the book by Sznitman [5]).

Lasry and Lions’s analytical approach, named mean field games (MFG), consists in solving a system of coupled nonlinear partial differential equations (PDEs): one forward in time, the Fokker–Planck–Kolmogorov (FPK) PDE, describing the evolution of the distribution of the state of one agent; and the other backward in time, the Hamilton–Jacobi–Bellman (HJB) PDE, characterizing the value function of this single agent in interaction with its own distribution. Simultaneously, Caines, Huang, and Malhamé (see [6]) proposed a similar approach to large games under the name of Nash certainty equivalence principle.

During the last decade, René Carmona and François Delarue have developed a probabilistic approach to MFGs which they present in these two impressive volumes based on their papers, lecture notes, and numerous results that are published here for the first time. I would like to state up front that they did a remarkable job in providing motivating examples, guiding the reader through a huge amount of complicated mathematical concepts and analyses, and providing timely applications, particularly in the area of financial mathematics. We will come back to these in more detail, but, given the size of this contribution (1440 pages!), it is legitimate to ask first the following general questions:

1. What does the probabilistic approach bring to the table when compared to the PDE approach?
2. Why should the presence of a common noise affecting all the players justify a second volume?

The first question can be answered by noting that the probabilistic approach to MFGs has been developed as an alternative to the PDE original approach in the spirit of the time-honored dual PDE-probability approach to problems like the solution of the heat equation and more general second-order linear PDEs for which Itô’s calculus and the martingale theory of diffusion processes have provided an alternative (see the book by Stroock and Varadhan [7]), or control theory for which backward stochastic differential equations (BSDEs) offered a substitute to HJB equations (see [8] and [9]). Historically, this dual (PDE-probability) approach has proven to enhance the understanding of these models, so the extension to stochastic differential games is natural.

The search for Nash equilibria has two clearly delineated steps: 1) an optimization problem whose solution provides the “best response” of a given player to the other players’

strategies; 2) a fixed point problem whose solution guarantees the consistency of the best responses identified in the first step. The original PDE approach to the search for Nash equilibria for MFGs encapsulated these two steps in a system of nonlinear PDEs, an HJB equation coupled with a form of an FPK equation, involving both directions of time. The perfect parallel in the probabilistic approach is a forward-backward stochastic differential equation (FBSDE) of the McKean–Vlasov type. The search for the best responses leads to an FBSDE while the fixed point step forces the distributions of the processes to be injected into the FBSDE, giving it its McKean–Vlasov character. While the probabilistic theory of standard FBSDEs was well developed, the analysis of this new type of forward-backward equations suggested by this approach was suspiciously ignored, and the first volume fills that gap (among other things).


The systematic development of the theory of optimal control of stochastic dynamics of McKean–Vlasov type is another new line of research triggered by the probabilistic approach to MFGs. It is called mean field control (MFC) to offer a parallel terminology to MFG. The special case of potential games was the subject of early investigations with PDE techniques. Its analysis was underpinned by the known fact from classical game theory that the search for Nash equilibria for potential games reduces to a single optimization problem. Building on this, the probabilistic approach identified this optimization as an optimal control problem of the McKean–Vlasov type, calling for the systematic development of its own theory. The 100 pages of Chapter 6 of Volume I give a self-contained account of this new theory of FBSDEs of McKean–Vlasov type. In this formulation, commonalities and differences with MFGs can easily be identified and better understood. In particular, the parallel between these theories and their treatments sheds new light on the differences between centralized and decentralized decision making, and opens up new avenues of research pointing toward progress in reinforcement learning for large populations of agents.

As for the second question, adding common noise to MFG and MFC models is needed in many applications. This common noise models the presence of a systematic noise affecting all the agents as opposed to individual noises. Unfortunately, but not surprisingly, it increases dramatically the levels of technicality and difficulty. These models are the object of Volume II, in which many of the ideas and tools introduced in Volume I are expanded to cover the new models. The natural extension of the PDE approach involves a system of stochastic partial differential equations (SPDEs) generalizing the HJB and the FPK equations. This already sets the challenges in the probabilistic framework, and it is possibly the reason why Lions introduced the master equation, which can capture models with common noise, still with pure PDE methods (see [10]). The notions of master field and master equation, including remarkable convergence

results as the number of players goes to infinity, are covered in detail in the second part of the second volume.

It is noteworthy that for the models with common noise, a good number of results have been obtained within the probabilistic approach, but so far they were not given a PDE analogue.

Now we come back to the books. Each chapter ends with a very helpful section of notes and references. In this review, we chose not to discuss bibliographic and authorship considerations.



Chapter 1 of Volume I offers a variety of motivating examples for the MFG paradigm. They come from different domains including finance, macroeconomics, biology, and social science. They are revisited and eventually solved throughout the book. Of course, in this review we cannot go over all these examples and their details, but as personal advice, I recommend following the example of a toy model of systemic risk

introduced in this first chapter. This is not only because I contributed to it (!) in [11], but also because it can be solved explicitly, and it helps to illustrate the concepts introduced in the general theory, which can be technically involved at times. It is a simple linear-quadratic (LQ) model of a finite number of agents represented by continuous-time diffusions where the drift of one component is controlled by the corresponding agent (its rate of borrowing and lending for instance), and the noises are independent Brownian motions (without a common noise to start with). Each agent tries to minimize her objective function comprising quadratic running and terminal costs so that the agent tries to stay close to the ensemble average (empirical mean). This example is solved in Chapter 2, that is Nash equilibria (open-loop and closed-loop) are computed using the stochastic maximum principle, an essential tool throughout the book. Markovian Nash equilibria are also obtained by PDE methods.

In my opinion, Chapter 3 is crucial for understanding the probabilistic approach to MFGs. It introduces the problem, the various ways to approach it, the different sets of strategies depending on the information structure leading to eventually different Nash equilibria, and the types of FBSDEs characterizing them, in particular the FBSDEs of McKean–Vlasov (MKV) type, which play a central role in the books. The stochastic maximum principle for this type of FBSDE is introduced as well as the notions of Lasry–Lions monotonicity and L -monotonicity crucial for uniqueness.

Chapter 4 goes into the details of FBSDEs and solutions to MFGs. In particular it introduces the fundamental concepts of a decoupling field for an FBSDE and of a master field for a MKV FBSDE. Several approaches are presented:

Schauder's fixed point theorem, probabilistic representation of the value function, and stochastic Pontryagin principle. A new type of MFGs, called extended MFGs, is also introduced. These are quite interesting, arising naturally in many applications where the interaction is through the empirical distribution of the controls, and they are the subject of ongoing research.

Part II of Volume I is devoted to spaces of probability measures (Chapter 5) and to optimal control of MKV SDEs (Chapter 6). Differential calculus over spaces of probability measures plays an essential role in the treatment of stochastic control of MKV dynamics (Chapter 6) and for the concept of the master equation in Part II of Volume II.

The analytic approach to MFGs relies heavily on a form of differential calculus on the Wasserstein space $P_2(R^d)$ of probability measures of order 2 on R^d as developed in the theory of optimal transportation. In his lectures at the Collège de France, Pierre-Louis Lions offered an original alternative to this calculus by lifting probability measures to random variables in an L^2 space. This form of calculus was tailor-made for functions of empirical measures as they appear in game models with mean field interactions. It is also particularly well suited to the control of the limit when the number of players or agents grows without bound and to the introduction of the master equation. The latter is a PDE on the product of the state space of the model and the space of probability measures on the state space. This PDE, typically in infinite dimensions, was touted by Lions as the proper tool to characterize equilibria.

This differential calculus is the object of Chapter 5 of Volume I, where the notion of differentiability based on liftings is called L -differentiability, short for Lions-differentiability. Chapter 5 offers a self-contained introduction to the topology of spaces of probability measures equipped with a Wasserstein distance, to the definition and properties of L -derivatives of functions of probability measures and to L -convexity and monotonicity. Section 5.6 provides a chain rule for differentiation along a flow of probability measures. This section explains very well the two possible approaches to derive such an Itô formula: either the classical approach of applying Taylor's formula on small time intervals or by approximating the flow of probability measures by a flow of empirical distributions of finite-dimensional systems and applying the classical Itô formula. The point that the second approach requires less regularity is well made.

Chapter 6 is about the new problem of optimal control of stochastic MKV dynamics or MFC as called by the authors who introduced it. Of course, the reader will wonder about the difference between MFGs and MFCs. The answer is given in Section 6.1. The starting point is the N -player differential game. For MFGs, one first optimizes to obtain a Nash equilibrium for the N players, and then passes to the limit as N goes to infinity to introduce the corresponding MFG. Alternatively, one can first pass to the limit using a fixed

point argument, and then optimize to obtain controlled MKV dynamics (MFC). However, the N large limit and the optimization steps do not commute in general, and the authors analyze carefully the subtle differences in this chapter. We reproduce their illuminating diagram in Figure 1.




Figure 1.

This new problem calls for revisiting its probabilistic and analytic formulations, and in particular the stochastic Pontryagin maximum principle. Existence and uniqueness of the corresponding FBSDE are obtained under suitable assumptions, in particular joint convexity in the state, its law, and the control. Then, several types of master equations are naturally introduced. It is very important to understand if one deals with a master equation for the value function or for its derivative. It is also interesting to observe that only first derivatives with respect to the probability distribution are involved in the case without common noise, as we differentiate along a deterministic flow of probability measures. This will be a major difference with the case with common noise studied in Volume II.

The first part of Volume II is devoted to games with a common noise. Typically, for an N -player differential game with a common noise, a common Brownian motion W^0 will appear in the dynamics of the states of each player. Intuitively, this common noise will not be averaged in the limiting empirical distribution, and it will also appear in the limiting controlled dynamics. Well, naively, one could say, let's condition on

this Brownian motion and repeat the analysis carried out in Volume I. This is in fact the right idea except that there are many difficult technical problems. They are resolved by the authors and presented in book form here for the first time. For the probabilistic approach to MFGs, the MKV FBSDEs to be solved now depend on this common noise which can be seen as a random environment where the flow of probability measures is now stochastic.

The first three sections of Chapter 1 provide an in-depth analysis of the new concept of a forward-backward system in a random environment. This is rather technical but necessary, and the authors should be thanked for the considerable effort they put into it. These tools are then used to address stochastic control problems in a random environment. The value function of the optimization

problem is now the solution of a stochastic HJB equation. Alternatively, the probabilistic approach leads to revisiting the Pontryagin stochastic maximum principle in the framework of random environment.


Then, the probabilistic framework developed in Chapter 1 is applied to MFGs with a common noise in Chapter 2. This starts with N -player differential games with a common noise represented by a Brownian motion W^0 independent of the (independent) idiosyncratic noises W^i , $i = 1, \dots, N$. Using the "naïve" idea of conditioning on W^0 , but equipped with the tools developed in Chapter 1, the MKV theory is extended to a conditional MKV theory associated to conditional MKV SDEs and to the corresponding non-linear stochastic FPK equations describing the evolution of the conditional limiting distribution. Conditional propagation of chaos is derived as well. Then, turning to control problems, strong and weak solutions to MFG problems are identified as well as a form of a Yamada–Watanabe theorem. A formulation in terms of systems of forward-backward SPDEs is also given.

Existence of weak solutions to MFGs with a common noise is treated in Chapter 3. Extending the fixed-point argument based on Schauder's theorem to the case with a common noise is tempting but intractable due to the fact that equilibria are now stochastic, making compactness much more difficult to establish than in the deterministic case without common noise. The strategy is to use an adequate topology to deal with compactness. This is very technical and is performed with admirable care.

Strong uniqueness under the Lasry–Lions monotonicity condition on the cost functions and interesting examples are presented, showing that the common noise may have a regularizing effect by restoring strong uniqueness for models where it fails in the absence of a common noise.

Part II of Volume II is about the master equation and approximations of N -player games equilibria. The notion of master field presented at the beginning of Chapter 4 is intuitively the generalization of the decoupling field introduced in Chapter 3 of Volume I in the case without common noise. Since we are now in a random environment produced by the common noise, the master field will be a function of time, the state variable, and the conditional law of the state given the common noise. Still, it is expected to "decouple" the equations in the FBSDEs in the random environment. It is also expected that the optimal control will be the gradient with respect to the state variable of the master field. In order to derive an equation for the master field, an Itô formula along flows of conditional probability measures is established. The master equation satisfied by the master field now involves a second derivative with respect to the probability measure as expected due to the diffusion term created by the common noise. In that sense, the equation is of second order in contrast with the case without common noise, as seen in Volume I where the master equation was of first order in the measure. To be

Springer International Publishing, Probability Theory and Stochastic Modelling Series, Volume 64, 712 pages, 2018.



rigorous, one needs to specify in which sense the master field is a solution of the master equation. This is done by a careful extension of the notion of viscosity solution to this type of (infinite-dimensional) equations. Existence and uniqueness for the master equation is then obtained under the conditions presented in Chapter 3. Examples presented at the beginning of Volume I are revisited from the point of view of the master equation. In particular, our favorite LQ toy model of systemic risk leads to an explicit solution of the master equation and a convergence result (as the number of players goes to infinity) which follows directly from explicit formulas.

One of the great achievements of the master equation is that it provides approximations to the corresponding finite-player game Nash equilibria. Start with the case without common noise and look for a closed-loop Markovian Nash equilibrium. The value functions (one for each player) satisfy a system of HJB equations which can be difficult to solve, in particular in high dimension. It can be approximated by injecting the state variables and the empirical distribution of the other states in the solution of the corresponding master equation (assuming existence and uniqueness of a classical solution), and the error of approximation can be shown to be of order $1/N$. Alternatively, one can consider the optimal path where the drift for player i involves the derivative with respect to the state variable x_i of the i th value function (for simplicity think about constant volatilities). This system of SDEs can be compared to the McKean–Vlasov system obtained by using in the drift the gradient of the solution to the corresponding master equation evaluated at the state of the i th player and at the empirical distribution of the states of the other players. This type of convergence is in the regime of the Law of Large Numbers and very recent remarkable results by Lacker, Delarue, and Ramanan show that the master equation is also key to handle central limit theorem, large deviations, and concentration inequalities.

The book ends with some extensions of the theory in Chapter 7. They include MFGs with major and minor players and MFGs of timing, both being very active domains of research.

In conclusion, the probabilistic approach to MFGs with and without a common noise has generated many interesting and challenging new problems requiring the development of new sophisticated tools. The authors have done a fantastic job in arranging all of them in these two volumes, which will stay as a major reference for generations of graduate students in stochastic analysis entering the subject of mean field games.

Well done René et François!

René Carmona and François Delarue received the 2020 Joseph L. Doob Prize for these volumes. See the prize announcement at https://www.ams.org/news?news_id=5596.

References

- [1] B. Øksendal, *Stochastic differential equations* (sixth edition), Universitext, Springer-Verlag, Berlin, 2003. MR2001996
- [2] W. H. Fleming and M. Soner, *Controlled Markov processes and viscosity solutions* (second edition), Stochastic Modeling and Applied Probability, vol. 25, Springer, New York, 2006. MR2179357
- [3] R. Carmona, *Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications*, Financial Mathematics, vol. 1, SIAM, 2016. MR3629171
- [4] J.-M. Lasry and P.-L. Lions, *Mean field games*, Jpn. J. Math. 2 (2007), no. 1, 229–260. MR2295621
- [5] A.-S. Sznitman, *Topics in propagation of chaos*, École d'Été de Probabilités de Saint-Flour XIX–1989, 165–251, Lecture Notes in Math., vol. 1464, Springer, Berlin, 1991. MR1108185
- [6] M. Huang, R. P. Malhamé, and P. E. Caines, *Large population stochastic dynamic games: closed-loop McKean–Vlasov systems and the Nash certainty equivalence principle*, Commun. Inf. Syst. 6 (2006), no. 3, 221–251. MR2346927
- [7] D. W. Stroock and S. R. S. Varadhan, *Multidimensional diffusion processes*, Classics in Mathematics, Springer-Verlag, Berlin, 2006. MR2190038
- [8] É. Pardoux and S. Peng, *Backward stochastic differential equations and quasilinear parabolic partial differential equations*, Stochastic partial differential equations and their applications (Charlotte, NC, 1991), 200–217, Lect. Notes Control Inf. Sci., vol. 176, Springer, Berlin, 1992. MR1176785
- [9] N. El Karoui, S. Peng, and M. C. Quenez, *Backward stochastic differential equations in finance*, Math. Finance 7 (1997), no. 1, 1–71. MR1434407
- [10] P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions, *The master equation and the convergence problem in mean field games*, Annals of Mathematics Studies, vol. 201, Princeton University Press, Princeton, NJ, 2019. MR3967062
- [11] R. Carmona, J.-P. Fouque, and L.-H. Sun, *Mean field games and systemic risk: a toy model*, Commun. Math. Sci. 13 (2015), no. 4, 911–933. MR3325083



Jean-Pierre Fouque

Credits

Covers of volumes I and II are courtesy of Springer Nature. Figure 1 is adapted by permission from Springer Nature: eBook, *Probabilistic Theory of Mean Field Games with Applications*, volume I, by René Carmona and François Delarue, 2018.

Author photo is courtesy of the author.