SIAG-FME interview with
René Carmona

Paul Wythes ’55 Professor of Engineering and Finance
Operations Research and Financial Engineering (ORFE)
Princeton University

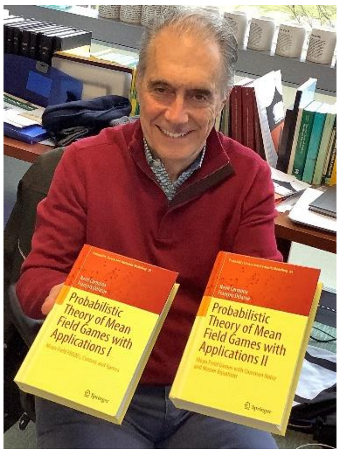

SIAG-FME*: Thank you, René, for agreeing to an interview for the newsletter of the SIAM Activity Group in Financial Mathematics. First of all congratulations on the 2020 Joseph L. Doob Prize, awarded by AMS, for the two-volume set Probabilistic Theory of Mean Field Games with Applications, co-authored with Francois Delarue, and published in 2018 by Springer.*

Can you give a very quick explanation what mean field games are for somebody who never heard about it before?

René Carmona: The theory of MFGs is an attempt to mathematically understand the behavior of large stochastic systems. Fifteen years ago, a group of mathematicians and economists in France, and a group of electrical engineers and mathematicians in Canada, simultaneously proposed a paradigm to understand the behavior of large populations of individuals selfishly minimizing their own costs or maximizing their rewards. What should this system look like? Should this rational behavior of individuals minimizing costs at the microscopic level, lead to rationality at the macroscopic level? In other words, should the system end up in a stable equilibrium? Or will the individualism of the individuals lead to irrationality, in the sense that the whole system would be chaotic, unable to reach an equilibrium.

Equilibrium may have a different meaning for physicists and researchers working in dynamical systems, than for people working in game theory. The most popular notion of equilibrium in game theory is the notion of Nash equilibrium. It was introduced by John Nash to identify strategies leading to stable group behaviors. This notion of stability had to be characterized in an original way as games typically involve the simultaneous optimization of several objective functions. In a Nash equilibrium, no individual has any interest or incentive to change his or her own strategy. Unfortunately, as soon as the number of players exceeds a small threshold, the search for equilibria becomes intractable, both mathematically and computationally. The paradigm of MFG suggests to consider the limit when the number of individuals grows without bound, and search for algorithms within this limiting regime, even if the mathematics become more sophisticated. This will hopefully lead to solutions which will provide approximate equilibria for finite, realistic systems.

SIAG-FME*: Which thought process have you followed for the development of the book. Are their topics or branches which you opted not to cover in detail or other content which you have considered fundamental?*

Carmona: The early works on MFGs, both in France and in Canada, were based on classical partial differential equations, typically Hamilton-Jacobi-Bellman (HJB) equations. Looking at the applications at the core of the mathematical models, one realizes that they all include random behaviors and stochastic environments. In fact, most models are specified via stochastic differential equations. I thought, that there ought to be a way to understand both the game models and the MFG paradigm using probabilistic tools. In mathematics, providing alternative formulations of problems, often leads to a better understanding of the challenges, and sometimes, to the discovery of new solutions and new problems. So, for the last 10 years, I have been working on a probabilistic approach to MFGs, inspired by the standard classical partial differential equations approach, and in search for new solutions. This explains the title of the book. This also explain that we did not cover some of the topics originally treated with partial differential equations when we could not find a natural way to recast them in the probabilistic approach.

SIAG-FME: *which part of the book did you enjoy most to write?*

Carmona: I can honestly tell you the points that I did not enjoy -- *the extremely technical parts which I let my younger co-author take care of*! However, even though models with what we call a common noise are extremely technical and difficult to solve, and even though complete analyses of such models are few and far between, it is important to set the record straight: they are crucial, especially for applications in engineering and economics.

Now back to your question, one task that I really enjoyed was to identify early works in diverse fields of applications, and try to present them as natural applications of mean field games. Even though I did not need to be convinced, that part cemented my conviction that mean field games were an unusually powerful modeling tool with a very broad spectrum of applications ranging from macro-economic to finance, from population biology to system engineering, from sociology to public policy, and I could go on and on.

I also greatly enjoyed writing the chapter where we presented the stochastic calculus on the spaces of probability measures. Calculus and differential geometric analysis on the Wasserstein space of probability measures were studied by many authors, and several excellent books on the subject already exist. Still, I found exciting and rewarding to put together all sorts new material brought to light by the probabilistic approach to MFGs, and to provide self-contained proofs of results scattered in the recent literature. This excitement led to a 100 pages chapter. This may sound insane but it is clean mathematics in accessible form, which I think can be useful in many parts of analysis, mathematics, statistics, and machine learning. For example, many recent theoretical analyses in machine learning rely on calculus on the Wasserstein space.

The chapter devoted to the theory of what we call the Control of McKean-Vlasov Dynamics, and which contains what is known as mean field control, is presumably more original. If we try to provide an intuitive explanation of the differences with standard MFG equilibria, one should think of a central planner doing the optimization for all the individuals and imposing the choice of a common control strategy to the individuals in the population. The optimization now concerns a single objective function: one does not solve a game model, one really solves a control problem, but the dynamics of the state of the system depend upon the statistical distribution of the states of all the individuals. It seems a very natural stochastic control problem, but it was rarely considered before the advent of MFGs. We developed tools for that new class of problems and we wrote a long chapter on this specific topic. I think this part of the books is a rather original and something that was very rewarding to achieve.

SIAG-FME: *With respect to the time when you outlined the book chapters, do you think there were in the meantime new results that you would like to add if you would write it anew? Which of these topics should be known to the community?*

Carmona: Francois and I were relentless in our efforts to include in the submitted manuscript all the works we could get a hold of by the time of the submission. Also, we used the long Notes & Complements sections at the end of each chapter to give extra references, and our own perspective on the history of the discoveries, as well as to discuss connected results when we thought they were relevant to the topic of the chapter in question.

Now a lot has happened since the publication of the book, and the vibrant research activity in mean field games we are witnessing these days is definitely exciting, especially when we realize how many young researchers are getting into the fray.

If I had to speak directly to your question, I would address applications in finance, where I see at least two significant developments worth mentioning. The first one concerns models in which we have a small number of major players and a field of minor players. These models are presented and discussed in the book, but today, the state of understanding of these models is much improved. It was in its infancy when we wrote the book. It was the first attempt to break the symmetry which is required among all the players in a mean field game model. One simple form of global symmetry breaking is to have groups of players and having mean field in each of these groups. But there is another more interesting extension, when you have a small number of major players influencing a field of minor players. This has a clear application in finance, where, if you consider the network of banks, you face many small financial institutions and the SIFIs[1], the Morgan Stanley, the Goldman Sachs and Deutsche Banks of this world, that do not behave in the same way as the small local banks. Thus, you cannot apply a model requiring full symmetry among all the players. These major-minor players models were presented in the book, but new methods and models have been developed since the publication of the book, for example in contract theory, where one can imagine running an office and hiring a large number of managers to whom you give a contract and who will interact in different ways, whether they like it or not. So, there is one major player and a field of minor players, and one could hope that these minor players may reach a mean field Nash equilibrium. The interaction between major player and the field of minor players depends upon the way you set it up. You can apply this to modeling bird flocking or schools of fish when you have a leader. There is currently a lot of interest in studying the Master Equation for these major-minor player models. This is crucial in understanding some phenomena such as large deviations in the setup of MFGs especially with common noise.

The second set of applications where the symmetry is again broken, and the MFG paradigm cannot be applied directly, is when the players interact through a network. Consider already established relationships between players that prevent the symmetry assumption to be satisfied. This is the case if we consider a simple competitive model, in which firms compete with each other, but not in a symmetric way. Pepsi would presumably compete with Coke, but may be not with Tesla; on the other hand, Tesla would compete with Ford, but maybe not with Pepsi. For each player, the set of players with whom they interact is given by a network or a link or a contract. This breaks the symmetry and it makes it impossible to apply the results of MFGs. But hopefully, the ideas developed to study mean field games will be adjusted, adapted, extended so that we are still able to study these new models. These new models will have a big impact on our understanding of social networks and more realistic models of behavior of large populations for which the symmetry assumption is not appropriate. It is a good first approximation. So far, we picked the low hanging fruits, and now we have to reach higher, for something more difficult, definitely more challenging, and most likely, more realistic.

SIAG-FME*: Trying to read your book, what would you recommend a reader how to approach this quite massive book? What is the best way to engage with it?*

Carmona.: Before you start, make sure you have some medication handy, or maybe a good bottle of wine. (laughs) I believe that, the fourteen hundred pages are organized in such a way so that the difficulties and the excitement come in crescendo. The chapters are written so that whatever you learn in one chapter will be used in the following ones. You can stop at any time, and still benefit from your efforts, whether you are interested in game theory, or in stochastic systems, or applications of a specific probabilistic tool, and even in the development of new probabilistic tools. Models with common noise, which are highly technical and very difficult to study, are presented at the end of the book, as a culmination. Nevertheless, they should not be discounted as they are crucial for many applications, not the least, macro-economics and

[1] SIFI stands for Systemically Important Financial Institution.

finance.

SIAG-FME*: In which field of applications do you think MFGs can be the most revolutionary? Where can MFGs overcome currently existing barriers and make significant advances beyond what is possible with existing methods?*

Carmona: As I mentioned, if we want to study a situation where individuals interact in an anonymous way, as soon as the number of individuals is significant - depending on the application that could be just a few individuals, or hundreds or hundreds of thousands – the computations basically become impossible; computing Nash equilibria is very difficult. To wit, I should add that computer scientists invented an entire field, Computational Game Theory, just to develop algorithms to compute Nash equilibria, even for static deterministic games! The MFG paradigm allows us to analyze the system by studying a generic individual interacting with a field of individuals, the field being modeled by a flow of probability measures. This simplification allows to find solutions which give approximate solutions to the original finite game models. This should be of interest to scientists from engineering, biology, social science, political science, and again finance. Let me try to mention a few applications in finance which have not been tackled yet, and which clearly should eventually to be solved with the use of MFG techniques. We are well aware of the interaction between a market maker and a large number of traders through an order book. They compete and their behaviors determine the time evolution of the order book. There are fundamental papers modeling the interaction of traders trading on a single stock, however, these models are relatively simplistic when compared to what is happening in the actual high frequency markets. MFG models have the potential to be more realistic. Another example which comes to mind is the pricing of digital goods and IPOs. It is well known that the major banks running the IPO pricing shows are not always doing well, and the examples of mispriced initial offerings abound. IPOs are auctions, and auctions can be viewed as games. IPOs tours involve a large number of possible investors that the underwriters are visiting and asking for their advice, with several layers of auctions in a given IPO. I think the presence of these several layers should be a motivation for new MFG models, and these MFG-inspired models, if and when they are solved, should provide a better understanding of the pricing of IPOs. Recently, the mathematical technology developed for the analysis of MFGs was touted as a possible way to shed light on the mysterious success of some Artificial Intelligence applications. Significant progress on the theoretical understanding of deep learning and deep neural networks has involved mean field limits where the width of a layer goes to infinity. Naturally, instead of working with the empirical measure of parameters, one works with probability distributions of parameters, and write the stochastic gradient descent used to fit and train neural networks in terms of a gradient flow on Wasserstein space. So even if one does not harness the full MFG paradigm, by which I mean even if one does not use the game component, the technology which has been developed, used, and promoted for the analysis of MFGs comes into play and can be extremely helpful in understanding these new phenomena which we do not fully understand, and which we would really like to understand.

SIAG-FME*: Nothing about the circadian rhythm?*

Carmona: I am glad that you asked. Among the many applications I worked on and found interesting, even if from the mathematical point of view they are not earth shattering, understanding jetlag and the synchronization of the circadian rhythms is dear to my heart. As you know, I have been commuting weekly between the West Coast and the East Coast for the last 25 years, and I am not immune to jet lag. So, I invited a graduate student to work on this problem. It is interesting in two ways. First, we could find a result which is consistent with the common belief that traveling eastwards is more difficult to recover from than traveling westwards. Second, the mathematical geek in me was happy to see a model naturally set up on the torus. Over the years, I was frustrated to see that some of my colleagues working on MFGs, would write a system of partial differential equations for the solution of the MFG and automatically would put the system on the torus without much justification, except for convenience obviously. I understand that periodic boundary conditions mean that there is no real boundary behavior to worry about, and when you look at large time behavior, the state space being compact makes things much easier. In addition, if you can look at an ergodic version of the problem you get rid of the time variable. So, models on the torus provide a lot of simplifications. I was frustrated to see that this mathematical convenience was rarely justified from a practical point of view. So when I started working on the synchronization of the oscillators in the brain to adjust to the light and recover from jetlag, I was happy to see that the model was naturally set on the torus.

Most importantly this model raised the issue of the Nash equilibrium. Once a system is in a Nash equilibrium, there is no incentive to change: no individual has any reason to change their behavior. After all, this is the definition of a Nash equilibrium. So, once we are there, we do not want to change. But how, why and when do we get there? There are many natural phenomena which are suggesting that individuals can evolve and reach a Nash equilibrium. Here is an example that was brought to my attention by a colleague: if you go to the opera, at the end of the show people stand up and start clapping. The clapping starts being completely chaotic, but after a few seconds there is a natural synchronization that happens, and people clap at a very specific frequency and in synchronicity. They reach a Nash equilibrium starting from a chaotic behavior. Many phenomena of that kind would benefit from being better understood mathematically.

SIAG-FME: *If you would make a prediction which topics in mathematical finance will attract the most attention over the next decade, what is the next big topic that MFGs in finance will play a role in?*

Carmona: As I already mentioned problems in high frequency trading on an order book, pricing digital goods and IPOs should make the list. If we can build a theory when the interactions are not symmetric, systemic risk will be the next obvious application which can be tackled. In its current state systemic risk can only be studied for a large number of symmetric institutions and possibly a small number of major institutions, and the detailed network of interactions between institutions cannot be captured by MFG models. Hopefully, we will be able to make significant progress in that direction in the near future.

I am sure there will be many more applications in economic theory. Prominent economists are currently revisiting many of the equilibrium models in macroeconomics and in growth theory to recast them in continuous time and continuous space, offering in so doing, a smooth transition between macro-economic and finance. Indeed, many of the early stochastic growth models basically appear as MFG models. Populations of households or firms compute aggregate quantities, and solve optimization problems once these aggregate quantities are fixed. De facto, they compute their best response to aggregate quantities. General equilibrium then requires that clearing conditions be satisfied. This is nothing more than requiring that a fixed point exists, and this is the second step of the MFG paradigm. I recently looked at a certain number of papers in macroeconomics which I presented in the introduction of the book -- the Ayagari model, the Krusell-Smith model -- and when I search for the word "Nash" I could not find it. These macroeconomic models are not searching for a Nash equilibrium, at least by name, but the general equilibrium that they identify is exactly a Nash equilibrium for a MFG model. Once this is realized, that will not change much to what the economists have done so well for the last 20 or 30 years, but it may open doors to solving many more problems, not necessarily directly in finance, but at least in economics. I believe that the number of applications is going to grow systematically once it is realized that all these general equilibrium models are in fact equilibriums for MFGs. Beyond this realization (I would not want to be criticized for acting like Monsieur Jourdain realizing that he has been speaking prose all along) I believe that the technology developed for the analysis of MFGs, and especially the numerical methods developed for their understanding, will be of great help.

SIAG-FME: *And on the side of theory development of MFGs, what are the upcoming big topic there?*

Carmona: I would bet on a systematic study of the controlled McKean-Vlasov equations and dynamics.

We wrote an entire chapter on this subject in the book, but we did not dwell on the far reaching practical applications of this theory to control and reinforcement learning. For example, imagine trying to control a large number of robots to accomplish a task, for example cleaning a minefield. If these robots could communicate, whether it is broadcasting their position, (or transmitting the number of mines they have already found) to a central unit which is going to collect the information about the individual states, possibly perform an optimization, and dispatch behavior policies, this would be a typical central planner problem, a typical mean field control problem, namely the control of McKean-Vlasov dynamics, because the empirical distribution of the states of the robots is going to enter the optimization of the central planner. This is something we understand well now. What is the next step? Many applications of reinforcement learning, involve Markov decision models for which the environment is not perfectly known, and the optimizer is torn between two obvious strategies. One is to maximize the immediate rewards and the other one is to explore the environment to gain a better grasp of the coefficients of the model. This is a theory which is extremely popular because of its spectacular successes, (beating the best player at Go being the claim of fame of Deep Mind). Next, we could imagine having reinforcement learning for a field of individuals instead of one individual. This is an exciting avenue of research, and with a couple of students and post-docs we started working in this direction. We are making progress, and while it is not earth shattering yet, we are hopeful that we shall be able to use the technology we developed for mean field control problems to our benefit. Stay tuned.

SIAG-FME*: Any last message you want to give the readers about mean field games?*

Carmona: I hope that anybody trying to read anything about MFGs will be as excited as I was when I first discovered the subject ten years ago through lectures of Pierre-Louis Lions and Jean-Michel Lasry. I also hope that they will be smarter than I was. At the beginning, I thought that MFGs and what we call mean field control, namely control of McKean-Vlasov equations, was the same thing. It is very easy to write a diagram, and any mathematician who sees a diagram like this is tempted to think that it has to be commutative: taking the limit when the number of players goes to infinity and then doing the optimization, or doing the optimization first and then taking the limit as the size of the system tends to infinity should give the same result. And it does not. In a sense, I spent a good six months trying to prove something which was wrong, and grossly wrong, because very simple counter-examples can be constructed. But, that gave me a great understanding of the differences between the two, and made me realize that one side of the coin, namely the control of McKean-Vlasov stochastic equations, was a virgin field which had to be unearthed. That was extremely exciting. Younger and brighter applied mathematicians face broad spectra of applications, and they should be able to take advantage of the technology behind the MFG paradigm, the game-theoretical way of thinking, the analysis of asymptotic regimes, McKean-Vlasov equations, propagation of chaos, and more. All these tools can be brought to bear on the solutions of many problems that, we did not even uncover yet.

SIAG-FME*: OK. Thank you very much Rene for the interview.*

Carmona: You are most welcome, that was pretty hard. (laughs) Thanks.