AFRL-AFOSR-VA-TR-2022-0608

Theory, Implementations, and Applications of Mean Field Games: the Second Generation

Carmona, Rene
TRUSTEES OF PRINCETON UNIVERSITY
1 NASSAU HALL
PRINCETON, NJ,
US

08/31/2022
Final Technical Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE	**2. REPORT TYPE**	**3. DATES COVERED**	
20220831	Final	**START DATE** 20190615	**END DATE** 20220614

4. TITLE AND SUBTITLE
Theory, Implementations, and Applications of Mean Field Games: the Second Generation

5a. CONTRACT NUMBER	**5b. GRANT NUMBER**	**5c. PROGRAM ELEMENT NUMBER**
	FA9550-19-1-0291	61102F
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Rene Carmona

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	**8. PERFORMING ORGANIZATION REPORT NUMBER**
TRUSTEES OF PRINCETON UNIVERSITY 1 NASSAU HALL PRINCETON, NJ US	

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	**10. SPONSOR/MONITOR'S ACRONYM(S)**	**11. SPONSOR/MONITOR'S REPORT NUMBER(S)**
Air Force Office of Scientific Research 875 N. Randolph St. Room 3112 Arlington, VA 22203	AFRL/AFOSR RTA2	AFRL-AFOSR-VA-TR-2022-0608

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Princeton The first goal of the original proposal was the development of new stochastic and numerical methods for the equilibrium analyses of the theory of Mean Field Games and Mean Field Control. Motivated by a wide spectrum of applications ranging from population dynamics and social networks, to communications and cybersecurity, it will advance the understanding of complex stochastic systems by introducing and developing new mathematical theories and associated computational algorithms. Its ambition is to provide a quantum leap to the state of the current understanding of Mean Field Games. Building on the results of the first generation of investigations, it proposes an ambitious research program whose goal is to define the second generation of equilibrium analyses of these complex stochastic systems. For example, now that we understand the concept of master equation, we plan to fully develop its theory by taking advantage of our newly acquired calculus on spaces of probability measures to produce asymptotic expansions and numerical algorithms, and investigate convergence rates, including models with delay of crucial importance in many logistic applications of interest to the Air Force.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	UU	7
U	U	U		

19a. NAME OF RESPONSIBLE PERSON	**19b. PHONE NUMBER** *(Include area code)*
WARREN ADAMS	00000000

Report Type: STTR/SBIR Final Progress Report
Primary Contact E-mail: rcarmona@princeton.edu
Primary Contact Phone: (609) 258-2310
Organization/Institution Name: Princeton University

Award Information

Title: (MURI) Theory, Implementations, and Applications of Mean Field Games: The Second Generation
Grant Number: FA9550-19-1-0291
Principal Investigator: Rene Carmona
Program Officer: Dr. Warren Adams

Report Information- STTR/SBIR Annual Progress Report

Reporting Start Date: 06/15/2021
Reporting End Date: 06/14/2022

Report Abstract

Princeton
The first goal of the original proposal was the development of new stochastic and numerical methods for the equilibrium analyses of the theory of Mean Field Games and Mean Field Control. Motivated by a wide spectrum of applications ranging from population dynamics and social networks, to communications and cybersecurity, it will advance the understanding of complex stochastic systems by introducing and developing new mathematical theories and associated computational algorithms. Its ambition is to provide a quantum leap to the state of the current understanding of Mean Field Games. Building on the results of the first generation of investigations, it proposes an ambitious research program whose goal is to define the second generation of equilibrium analyses of these complex stochastic systems. For example, now that we understand the concept of master equation, we plan to fully develop its theory by taking advantage of our newly acquired calculus on spaces of probability measures to produce asymptotic expansions and numerical algorithms, and investigate convergence rates, including models with delay of crucial importance in many logistic applications of interest to the Air Force. Finite size effects will also be investigated for models on graphs and networks. Not only will convergence of finite models toward their mean field limits will be studied, but a fine analysis of possible phase transitions and bifurcations will be conducted in search for a classification of the possible equilibrium dynamics.

Columbia
In the past year I have made progress on several projects related to mean field models and their network-based extensions. This includes a general result on the convergence problem for mean field games with common noise, sharp quantitative results on long-time behavior of mean field models, a new and user-friendly formulation of so-called "graphon games," as well as a promising new non-asymptotic perspective on decentralized control theory. Many extensions of the latter two projects remain to be explored.

Technical Highlights and Accomplishments**:**

Princeton:
For the record, and for the ease of reference, we recall the three main objectives of the project.
(1) Development of Machine Learning tools (based on neural network implementations) to compute

numerically Nash equilibria of Mean Field Game models and solutions of optimal Mean Field Control problems

(2) Starting the theory of stochastic dynamic graphon games for the understanding of large network games with interactions given by dense graphs

(3) Development of convergent algorithms for mean field reinforcement learning which could be used for the control of large fleets (e.g. robots)

We disseminated and promoted the results of (1) by a series of lectures (even if most of them were on line). In the process, we revised the two fundamental papers [P7] and [P8] written in collaboration with Mathieu Laurière (postdoc) on the use of Machine Learning algorithms to compute numerical solutions of Mean Field Games. [P7] already appeared in the SIAM Journal on Numerical Analysis (certainly one of the very top journals in the field), while [P8] appeared on line and should soon be scheduled for a printed issue. The confidence gained from the theoretical results of these two papers and the remarkable performance of the numerical algorithms prompted us to use them in all of the papers cited below involving numerical examples of Mean Field Game, Mean Field Control and Reinforcement Learning applications. This is especially the case with papers [P10] and [P11] modeling the impact of a regulator on a large population of individuals or firms which can be in finitely many states and interact in a mean field way. These papers were motivated by applications of great interest: the control of a pandemic and of carbon emissions respectively. Also caused by the exposure from these publications we were asked to write a book chapter on Mean Field Games, Machine Learning and Applications to Finance [PBC2]. This chapter will appear in a Cambridge University Press handbook entitled *Machine Learning in Financial Markets: A guide to contemporary practices,* eds A. Capponi and C.A. LeHalle

Work on point (2) started with the paper [P6] on the static case of graphon games. This paper appeared in Mathematics of Operations Research. We continued the line of research on network games by preparing lecture notes for a short course given in June 2021 at the new NSF institute IMSI in Chicago. We then worked on the extension of this first paper to the dynamic case. This line of research led to the paper [P16] treating the linear quadratic case of the stochastic graphon game. This paper was accepted for publication in Applied Mathematics and Optimization and should appear in print before the end of the year. We are currently reviewing the economic literature on reputation games with the goal of applying the technology of Mean Field Games and large network games to revive the theory of these games in continuous time. In parallel, we are investigating how the analysis of large network games underpinned by random graphs (dense and sparse) can help understand synchronization phenomena observed in neurobiology and social studies. We engaged two new collaborators to pursue these investigations. This led to the recent technical report [PTR4] written in collaboration with Mete Soner and Quentin Cormier (postdoc who just got a position at INRIA in France). This work will be posted on arxiv and submitted for publication in the next few weeks. It provides a first set of results on the stability of mean field game equilibria. Despite the deceptive simplicity of the statements of the main results, the proofs involve highly sophisticated mathematical tools from functional analysis and viscosity solutions of nonlinear partial differential equations.

Investigations linked to bullet point (3) led to the prepublications [P9, P12, P13] and [PTR4]. Some of these papers have been accepted for publication, and other are still under revision. Some of these revisions are extensive. After we posted the original version of [P13], Motte and Pham posted a version of the same problem proposing a framework for the treatment of the open loop version of our model of model free learning for Mean Field reinforcement learning. However, we found their model incomplete and we discovered that in proposing the proper definition for the model, it was possible to prove the result they could not achieve, namely the equality of the open loop and closed loop value functions. We completed a new and improved version of the paper, including original numerical illustrations of the results, and we submitted it for publication in the Annals of Applied Probability.

Beyond the results on the three objectives reviewed above, we completed with PhD student Laura Leal, [P14] using an original optimal control model motivated by prior work on Mean Field Control, a study of the behavior of traders on some high frequency markets. After being posted on SSRN, the paper stayed on the list of the 10 most cited papers for 3 weeks. Since then, these results have been extended in a joint work with a current PhD student, Claire Zeng who introduced a mean field game approach capturing the competition among many traders. This work will be posted in the fall.
I also prepared and revised a chapter [PBC1] for a book published by the American Mathematical Society. This chapter is based on the lectures given during the AMS Short Course on Mean Field Games held in parallel with the 2020 Annual Meeting in Denver CO.

Columbia:

My student A. Soret and I finished a key paper (#3 below) on mean field game type models on large-scale networks. We develop a user-friendly framework for the study of models that have come to be known as "graphon games," in which a large dense network is replaced by a continuous limit. This framework is much simpler than those of prior work, enabling simpler analysis of more general models. My student L. Le Flem and I finished two papers. The first (#1 below) solves the "convergence problem" for mean field games with common noise, a central yet still poorly understood problem in the field. The second (#4 below) obtains optimal rates of convergence for a broad class of mean field interacting diffusions, which were absent from prior literature despite decades of study. My former student J. Zhang (now a postdoc at UC Berkeley IEOR) and I finished a paper (#2 below) giving new stochastic and analytic formulations of stationary solutions of interacting particle systems governed by large sparse (locally tree-like) graphs, which will be a useful building block for future studies of game-theoretic models. Most recently, jointly with my student L.C. Yeung and S. Mukherjee (assistant professor in Columbia Statistics), I co-authored a paper (#5 below) on static mean field approximations which arise in statistical physics, developing a new approach based on log-concavity. Notably, we show how for mean field game and control theory, we showed how our results lead to a new non-asymptotic perspective on mean field control problems, applicable even when the problem does not admit enough symmetry or structure to lend itself to the usual asymptotic methodology.

Publications throughout the period of the grant

Princeton

Peer-reviewed journal articles:

[P1] R. Carmona and C. Graves: Jet Lag Recovery: Synchronization of Circadian Rhythms as a Mean Field Game. *Dynamic Games and Applications,* ***10*** *(2020) 79-99.*
[P2] B. Acciaio, D J. Backho-Veraguas and R. Carmona: Extended Mean Field Control Problems: Stochastic Maximum Principle and Transport Perspective. *SIAM Journal on Optimization and Control* **57** (6) (2019), 3666-3693
http://arxiv.org/1802.05754
[P3] R. Carmona and P.Wang: A Probabilistic Approach to Extended Finite State Mean Field Games. *Mathematics of Operations Research* **46**(2) (2021) 471 - 502.
http://arxiv.org/abs/1808.07635

[P4] R. Carmona and P. Wang: Finite-State Contract Theory with a Principal and a Field of Agents. *Management Science* ***67****(8) (2021) 4725 - 4741*.
http://arxiv.org/abs/1808.07942
[P5] R. Carmona, M. Cerenzia and A.Z. Palmer: The Dyson and Coulomb Games.
*Annales Henri Poincare***, 21** (2020) 2897-2949.
http://arxiv.org/abs/1808.02464
[P6] R. Carmona, D. B. Cooney, C. V. Graves, and M. Lauriere: Stochastic Graphon Games: I. the static case. *Mathematics of Operations Research,* **47** (1) (2022) 750-778.
http://arxiv.org/abs/1911.10664
[P7] R. Carmona and M. Laurière: Convergence Analysis of Machine Learning Algorithms for the Numerical Solution of Mean Field Control and Games: I. The Ergodic Case. *SIAM Journal on Numerical Analysis*, 59 (2021) 1455-1485.
http://arxiv.org/abs/1907.05980
[P8] R. Carmona M. Laurière: Convergence Analysis of Machine Learning Algorithms for the Numerical Solution of Mean Field Control and Games: II. The Finite Horizon Case. *Annals of Applied Probability* (accepted for publication, to appear)
http://arxiv.org/abs/1907.05980
[P9] R. Carmona, K. Hamidouche, M. Laurière, and Z. Tan: Policy Optimization for Linear-Quadratic Zero-Sum Mean-Field Type Games. *59th IEEE Conference on Decision and Control (CDC)* (2020) 1038-1043.
http://arxiv.org/abs/2009.02146
[P10] R. Carmona, G. Dayanikli, and M. Laurière: Mean Field Models to Regulate Carbon Emissions in Electricity Production. *Dynamic Games and Applications (January 2022)*
http://arxiv.org/abs/2102.09434
[P11] A. Aurell, R. Carmona, G. Dayanikli, and M. Laurière: Optimal incentives to mitigate epidemics: a Stackelberg mean field game approach. SIAM Journal of Optimization and Control **60** (2) (2022) 294-322
http://arxiv.org/abs/2011.03105
[P12] R.Carmona, K. Hamidouche, M. Laurière, and Z. Tan: Linear-Quadratic Zero-Sum Mean-Field Type Games: Optimality Conditions and Policy Optimization. *Journal of Dynamics and Games*, **8** (4), (2021) 403 - 443.
http://arxiv.org/abs/2009.00578
[P13] R. Carmona, M. Laurière, and Z. Tan: Model-Free Mean-Field Reinforcement Learning: Mean-Field MDP and Mean-Field Q-Learning.(submitted for publication)
http://arxiv.org/abs/1910.12802
[P14] R. Carmona and L. Leal: Optimal Execution with Quadratic Variation Inventories.
SIAM Journal on Financial Mathematics (to appear 2022)
http://arxiv.org/abs/2104.14615
[P15] A. Aurell, R. Carmona, G. Daynikli, and M. Laurière: Finite State Graphon Games with Applications to Epidemics. *Dynamic Games and Applications (to appear 2022)*
http://arxiv.org/abs/2106.07859
[P16] A. Aurell, R. Carmona and M. Laurière: Stochastic Graphon Games: II. the Linear-Quadratic case. *Applied Mathematics and Optimization* (to appear 2022)
http://arxiv.org/abs/2105.12320
[P17] R. Carmona, and G. Dayanikli) Mean Field Game Model for an Advertising Competition in a Duopoly. *International Game Theory Review, (accepted for publication, to appear in 2022)*
http://arxiv.org/abs/2201.05182

Conference proceedings:

[PC1] A. Angiuli, R. Carmona, C. Graves, H.li, J.F. Chassagneux, and F. Delarue: Numerical Probabilisitc Approach to MFGs. *ESAIM: PROCEEDINGS AND SURVEYS, eds. B. Bouchard, J.-F. Chassagneux, F. Delarue, E. Gobet and J. Lelong, 65, February 2019, 84-113.*
[PC2] R. Carmona, C. Graves and N. Tan: Price of Anarchy for Mean Field Games *ESAIM: PROCEEDINGS AND SURVEYS, eds. B. Bouchard, J.-F. Chassagneux, F. Delarue, E. Gobet and J. Lelong, 65, February 2019, 349-383.*
http://arxiv.org/abs/1802.04644

Book Chapters:

[PBC1] R. Carmona: Applications of Mean Field Games in Financial Engineering and Economic Theory. in *American Mathematical Society*, (2021) Providence RI.
http://arxiv.org/abs/2012.05237
[PBC2] R. Carmona M. Laurière: Deep Learning for Mean Field Games and Mean Field Control with application to Finance, in *Machine Learning in Financial Markets: A guide to contemporary practices,* eds A. Capponi and C.A. LeHalle, Cambridge University Press (2022)
http://arxiv.org/abs/2107.04568

Books:

[PB1] R. Carmona, and F. Delarue: Probabilistic Theory of Mean Field Games: vol. I, Mean Field FBSDEs, Control, and Games. Stochastic Analysis and Applications. Springer Verlag, Feb. 2018 713pp
[PB2] R. Carmona, and F. Delarue: Probabilistic Theory of Mean Field Games: vol. II, Mean Field Games with Common Noise and Master Equations. Stochastic Analysis and Applications. Springer Verlag, Feb. 2018. 697pp

Technical reports:

[PTR1] R. Carmona and K. Webster: Applications of a New Self-Financing Equation
http://arxiv.org/1905.04137
[PTR2] R. Carmona and F. Delarue: Incentivizing Mean Field Game Players and the Price of Instability.
[PTR3] Demographic Mean Field Games
[PTR4] R. Carmona, M. Laurière and Z. Tan: Linear Quadratic Reinforcement Learning: Convergence of Policy Gradient Methods. (submitted for publication)
http://arxiv.org/abs/1910.04295
[PTR4] R. Carmona and X. Yang: GLASSO Model for Electric Load and Wind Power and Monte Carlo Scenario Generation. (submitted for publication)
http://arxiv.org/abs/2111.14628
[PTR5] R. Carmona, Q. Cormier and M. Soner: Synchronization in a Kuramoto Mean Field Game.

%%

Columbia:

[C1] A case study on stochastic games on large graphs in mean field and sparse regimes
With Agathe Soret. To appear in Mathematics of Operations Research.
https://arxiv.org/abs/2005.14102
[C2] Many-player games of optimal consumption and investment under relative performance criteria. With Agathe Soret. Mathematics and Financial Economics 14, 263–281 (2020).
[C3] Inverting the Markovian projection, with an application to local stochastic volatility models. With Mykhaylo Shkolnikov and Jiacheng Zhang. Annals of Probability 48(5), 2189-2211 (2020).
[C4] Hierarchies, entropy, and quantitative propagation of chaos for mean field diffusions.
https://arxiv.org/abs/2105.02983
[C5] Quantitative approximate independence for continuous mean field Gibbs measures.
https://arxiv.org/abs/2105.03238
[C6] Marginal dynamics of interacting diffusions on unimodular Galton-Watson trees. With Kavita Ramanan and Ruoyu Wu.
https://arxiv.org/abs/2009.11667
[C7] Local weak convergence for sparse networks of interacting processes
With Kavita Ramanan and Ruoyu Wu.
https://arxiv.org/abs/1904.02585v3
[C8] A characterization of transportation-information inequalities for Markov processes in terms of dimension-free concentration. With Lane Chun Yeung.
https://arxiv.org/abs/2012.02304
[C9] Closed-loop convergence for mean field games with common noise. With Luc Le Flem.
https://arxiv.org/abs/2107.03273
[C10] Stationary solutions and local equations for interacting diffusions on regular trees. With Jiacheng Zhang.
https://arxiv.org/abs/2111.05416
[C11] A label-state formulation of stochastic graphon games and approximate equilibria on large networks. With Agathe Soret.
https://arxiv.org/abs/2204.09642
[C12] Sharp uniform-in-time propagation of chaos. With Luc Le Flem. Preprint.
https://arxiv.org/abs/2205.12047
[C13] Mean field approximations via log-concavity. With Sumit Mukherjee and Lane Chun Yeung.
https://arxiv.org/abs/2206.01260

Changes in research objectives (if any):

None

Please list personnel associated with the period of performance for this grant (to include any past and present collaborators, students and AFOSR program officers):

Princeton:
Grad Students: Laura Leal, Gokce Daynikli, Zongjun Tan, Claire Zeng, Katerine Donoghue
Post Docs: Mathieu Laurière, Alexander Aurell, Quentin Cormier
Co-authors: A Aurell, D.B. Cooney, Q. Cormier, G. Daynikli, F. Delarue, C. Graves, M. Laurière, L. Leal, M. Soner, Z. Tan, X. Yang

Columbia:
Grad Students: Agathe Soret, Luc Le Flem, Lane Chun Yeung, (co-advised with Ioannis Karatzas), Jiacheng Zhang, Princeton University, co-advised with Mykhaylo Shkolnikov
Co-authors: Sumit Mukherjee (Columbia University), Agathe Soret (Columbia University), Lane Chun Yeung (Columbia University), Jiacheng Zhang (University of California, Berkeley), Kavita Ramanan (Brown University), Ruoyu Wu (Iowa State University)

Extensions Granted or milestones slipped, if any:

An NCE was submitted and is still pending the outcome.