

Mean Field Games

Definition

Mean Field Game (MFG) theory studies the existence of Nash equilibria, together with the individual strategies which generate them, in games involving a large number of agents modeled by controlled stochastic dynamical systems. This is achieved by exploiting the relationship between the finite and corresponding infinite limit population problems. The solution of the infinite population problem is given by the fundamental MFG Hamilton-Jacobi-Bellman (HJB) and Fokker-Planck-Kolmogorov (FPK) equations which are linked by the state distribution of a generic agent, otherwise known as the system's mean field.

Introduction

Large-population, dynamical, multi-agent, competitive, and cooperative phenomena occur in a wide range of designed and natural settings such as communication, environmental, epidemiological, transportation, and energy systems, and they underlie much economic and financial behavior. Analysis of such systems is intractable using the finite population game theoretic methods which have been developed for multi-agent control systems (see, e.g., Basar and Ho 1974; Ho 1980; Basar and Olsder 1999; and Bensoussan and Frehse 1984). The continuum population game theoretic models of economics (Aumann and Shapley 1974; Neyman 2002) are static, as, in general, are the large-population models employed in network games (Altman et al. 2002) and transportation analysis (Wardrop 1952; Haurie and Marcotte 1985; Correa and Stier-Moses 2010). However, dynamical (or sequential) stochastic games were analyzed in the continuum limit in the work of Jovanovic and Rosenthal (1988) and Bergin and Bernhardt (1992), where the fundamental mean field equations appear in the form of a discrete time dynamic programming equation and an evolution equation for the population state distribution.

The mean field equations for dynamical games with large but finite populations of asymptotically negligible agents originated in the work of Huang et al. (2003, 2006, 2007) (where the framework was called the Nash Certainty Equivalence Principle) and independently in that of Lasry and Lions (2006a,b, 2007), where the now standard terminology of Mean Field Games (MFGs) was introduced. Independent of both of these, the closely related notion of Oblivious Equilibria for large-population dynamic games was introduced by Weintraub et al. (2005) in the framework of Markov Decision Processes (MDPs).

One of the main results of MFG theory is that in large-population stochastic dynamic games individual feedback strategies exist for which any given agent will be in a Nash equilibrium with respect to the pre-computable behavior of the mass of the other agents; this holds exactly in the asymptotic limit of an infinite population and with increasing accuracy for a finite population of agents using the infinite population feedback laws as the finite population size tends to infinity, a situation which is termed an ϵ -Nash equilibrium. This behavior is described by the solution to the infinite population MFG equations which are fundamental to the theory; they consist of (i) a parameterized family of HJB equations (in the nonuniform parameterized agent case) and (ii) a corresponding family of McKean-Vlasov (MV) FPK PDEs, where (i) and (ii) are linked by the probability distribution of the state of a generic agent, that is to say, the mean field. For each agent, these yield (i) a Nash value of the game, (ii) the best response strategy for the agent, (iii) the agent's stochastic differential equation (SDE) (i.e., the MV-SDE pathwise description), and (iv) the state distribution of such an agent (via the MV FPK for the parameterized individual).

Dynamical Agents

In the diffusion-based models of large-population games, the state evolution of a collection of N agents $A_i, 1 \leq i \leq N < \infty$, is specified by a set N of controlled stochastic differential equations (SDEs) which in the important linear case take the form

$$dx_i(t) = [F_i x_i(t) + G_i u_i(t)] dt + D_i dw_i(t), \quad 1 \leq i \leq N,$$

where $x_i \in \mathbb{R}^n$ is the state, $u_i \in \mathbb{R}^m$ the control input, and w_i the state Wiener process of the i th agent A_i ,

where $\{w_i, 1 \leq i \leq N\}$ is a collection of N independent standard Wiener processes in \mathbb{R}^r independent of all mutually independent initial conditions. For simplicity, throughout this entry, all collections of system initial conditions are taken to be independent, zero mean and have finite second moment.

A simplified form of the general case treated in Huang et al. (2007) and Nourian and Caines (2013) is given by the following set of controlled SDEs which for each agent A_i includes state coupling with all other agents:

$$dx_i(t) = \frac{1}{N} \sum_{j=1}^N f(t, x_i(t), u_i(t), x_j(t)) dt + \sigma dw_i(t), \quad 1 \leq i \leq N,$$

where here, for the sake of simplicity, only the uniform (non-parameterized) generic agent case is presented. The dynamics of a generic agent in the infinite population limit of this system is then described by the following controlled MV stochastic differential equation:

$$dx(t) = f[t, x(t), u(t), \mu_t] dt + \sigma dw(t),$$

where $f[t, x, u, \mu_t] = \int_{\mathbb{R}^n} f(t, x, u, y) \mu_t(dy)$, with the initial condition measure μ_0 specified, where $\mu_t(\cdot)$ denotes the state distribution of the population at $t \in [0, T]$. The dynamics used in the analysis in Lasry and Lions (2006a,b, 2007) and Cardaliaguet (2012) are of the form $dx_i(t) = u_i(t)dt + \sigma dw_i(t)$.

The dynamical evolution of the state x_i of the i th agent A_i in the discrete time Markov Decision Processes (MDP)-based formulation of the so-called anonymous sequential games (Jovanovic and Rosenthal 1988; Bergin and Bernhardt 1992; Weintraub et al. 2005) is described by a Markov state transition function, or kernel, of the form $P_{t+1} = P(x_i(t+1)|x_i(t), x_{-i}(t), u_i(t), P_t)$.

Agent Performance Functions

In the basic finite population linear-quadratic diffusion case, the agent $A_i, 1 \leq i \leq N$, possesses a performance, or loss, function of the form

$$J_i^N(u_i, u_{-i}) = E \int_0^T \{ \|x_i(t) - m_N(t)\|_Q^2 + \|u_i(t)\|_R^2 \} dt,$$

where we assume the cost coupling to be of the form $m_N(t) := (\overline{x_N(t)} + \eta), \eta \in \mathbb{R}^n$, where u_{-i} denotes all

agents' control laws except for that of the i th agent and $\overline{x_N}$ denotes the population average state $(1/N) \sum_{i=1}^N x_i$,

and where here and below the expectation is taken over an underlying sample space which carries all initial conditions and Wiener processes.

For the nonlinear case introduced in the previous section, a corresponding finite population mean field loss function is

$$J_i^N(u_i; u_{-i}) := E \int_0^T \left((1/N) \sum_{j=1}^N L(t, x_i(t), u_i(t), x_j(t)) \right) dt, \quad 1 \leq i \leq N,$$

where L is the nonlinear state cost-coupling function. Setting, by possible abuse of notation,

$L[t, x, u, \mu_t] = \int_{\mathbb{R}^n} L(t, x, u, y) \mu_t(dy)$, the infinite population limit of this cost function for a generic

individual agent A is given by

$$J(u, \mu) := E \int_0^T L[t, x(t), u(t), \mu_t] dt,$$

which is the general expression for the infinite population individual performance functions appearing in Huang et al. (2006) and Nourian and Caines (2013) and which includes those of Lasry and Lions (2006a,b, 2007) and Cardaliaguet (2012). Exponentially discounted costs with discount rate parameter ρ are employed for infinite time horizon performance functions in Huang et al. (2003, 2007), while the sample path limit of the long-range average is used for ergodic MFG problems in Lasry and Lions (2006a, 2007) and Li and Zhang (2008) and in the analysis of adaptive MFG systems (Kizilkale and Caines 2013).

The Existence of Equilibria

The objective of each agent is to find strategies (i.e., control laws) which are admissible with respect to information and other constraints and which minimize its performance function. The resulting problem is necessarily game theoretic and consequently central results of the topic concern the existence of Nash Equilibria and their properties.

The basic linear-quadratic mean field problem has an explicit solution characterizing a Nash equilibrium (see Huang et al. 2003, 2007). Consider the scalar infinite time horizon discounted case, with nonuniform parameterized agents A_θ with

parameter distribution $F(\theta), \theta \in \mathcal{A}$, and dynamical parameters identified as $a_\theta := F_\theta, b_\theta := G_\theta, Q := 1, r := R$; then the so-called Nash Certainty Equivalence (NCE) equation scheme generating the equilibrium solution takes the form

$$\begin{aligned} \rho s_\theta &= \frac{ds_\theta}{dt} + a_\theta s_\theta - \frac{b_\theta^2}{r} \Pi_\theta s_\theta - x^*, \\ \frac{d\bar{x}_\theta}{dt} &= \left(a_\theta - \frac{b_\theta^2}{r} \Pi_\theta \right) \bar{x}_\theta - \frac{b_\theta^2}{r} s_\theta, \quad 0 \leq t < \infty, \\ \bar{x}(t) &= \int_{\mathcal{A}} \bar{x}_\theta(t) dF(\theta), \\ x^*(t) &= \gamma(\bar{x}(t) + \eta), \\ \rho \Pi_\theta &= 2a_\theta \Pi_\theta - \frac{b_\theta^2}{r} \Pi_\theta + 1, \quad \Pi_\theta > 0, \quad \text{Riccati Equation} \end{aligned}$$

where the control action of the generic parameterized agent A_θ is given by

$$u_\theta^0(t) = -\frac{b_\theta}{r} (\Pi_\theta x_\theta(t) + s_\theta(t)), 0 \leq t < \infty. \quad u_\theta^0 \text{ is the optimal tracking feedback law with respect to } x^*(t)$$

which is an affine function of the mean field term $\bar{x}(t)$, the mean with respect to the parameter distribution F of the $\theta \in \mathcal{A}$ parameterized state means of the agents. Subject to the conditions for the NCE scheme to have a solution, each agent is necessarily in a Nash equilibrium in all full information causal (i.e., non-anticipative) feedback laws with respect to the remainder of the agents when these are employing the law u^0 .

It is an important feature of the best response control law u_θ^0 that its form depends only on the parametric data of the entire set of agents, and at any instant it is a feedback function of only the state of the agent A_θ itself and the deterministic mean field-dependent offset s_θ .

For the general nonlinear case, the MFG equations on $[0, T]$ are given by the linked equations for (i) the performance function V for each agent in the continuum, (ii) the FPK for the MV-SDE for that agent, and (iii) the specification of the best response feedback law depending on the mean field measure μ_t and the agent's state $x(t)$. In the uniform agents case, these take the following form.

The Mean Field Game HJB: (MV) FPK Equations

$$\begin{aligned}
 \text{[MV-HJB]} \quad & -\frac{\partial V(t,x)}{\partial t} = \inf_{u \in U} \left\{ f[t, x(t), u(t), \mu_t] \frac{\partial V(t,x)}{\partial x} + L[t, x(t), u(t), t] \right. \\
 & \quad \left. + \frac{\sigma^2}{2} \frac{\partial^2 V(t,x)}{\partial x^2} \right\}, \\
 & V(T, x) = 0, \quad (t, x) \in [0, T] \times \mathbb{R}, \\
 \text{[MV-FPK]} \quad & \frac{\partial \mu(t,x)}{\partial t} = -\frac{\partial \{f[t,x,u(t),\mu_t]\mu(t,x)\}}{\partial x} + \frac{\sigma^2}{2} \frac{\partial^2 \mu(t,x)}{\partial x^2}, \\
 \text{[MV-BR]} \quad & u(t) = \varphi(t, x(t)|\mu_t), \quad (t, x) \in [0, T] \times \mathbb{R}
 \end{aligned}$$

The general nonlinear MFG problem is approached by different routes in Huang et al. (2006) and Nourian and Caines (2013), and Lasry and Lions (2006a,b, 2007) and Cardaliaguet (2012), respectively. In the former, the so-called probabilistic method solves the MFG equations directly. Subject to technical conditions, an iterated contraction argument establishes the existence of a solution to the HJB-(MV) FPK equations; the best response control laws are obtained from these MFG equations, and these are necessarily Nash equilibria within all causal feedback laws for the infinite population problem. In Lasry and Lions (2006a, 2007) the MFG equations on the infinite time interval (i.e., the ergodic case) are obtained as the limit of Nash equilibria for increasing finite populations, while in the expository notes of Cardaliaguet (2012) the analytic properties of solutions to the HJB-FPK equations on the finite interval are analyzed using PDE methods including the theory of viscosity solutions.

In Huang et al. (2003, 2006, 2007), Nourian and Caines (2013), and Cardaliaguet (2012), it is shown that subject to technical conditions, the solutions to the HJB-FPK scheme yield ϵ -Nash solutions for finite population MFGs in that for any $\epsilon > 0$, there exists a population size N_ϵ such that for all larger populations the use of the feedback law given by the MFG infinite population scheme gives each agent a value to its performance function within ϵ of the infinite population Nash value.

A counterintuitive feature of these results is that, asymptotically in population size, observations of the states of rival agents are of no value to any given agent; this is in contrast to the situation in single-agent optimal control theory where the value of observations on an agent's environment is in general positive.

Current Developments and Open Problems

There is now an extensive literature on Mean Field Games, the following being a sample: the mathematical literature has focused on the study of general classes of solutions to the fundamental HJB-FPK equations (see e.g., Cardaliaguet (2013)), while in systems and control, the theory of major-minor agent MFG problems (in economics terminology, atoms and continua) is being developed (Huang 2010; Nguyen and Huang 2012; Nourian and Caines 2013), adaptive control extensions of the LQG theory have been carried out (Kizilkale and Caines 2013), and the risk-sensitive case has been analyzed (Tembine et al. 2012). Much work is now under way in the applications of MFG theory to economics, finance, distributed energy systems, and electrical power markets. Each of these areas has significant open problems, including the application of mathematical transport theory to HJB-FPK equations, the role of MFG theory in portfolio optimization, and the analysis of systems where the presence of partially observed major and minor agent states incurs mean field and agent state estimation.

Bibliography

- Altman E, Basar T, Srikant R (2002) Nash equilibria for combined flow control and routing in networks: asymptotic behavior for a large number of users. *IEEE Trans Autom Control* 47(6):917-930. Special issue on Control Issues in Telecommunication Networks
- Aumann RJ, Shapley LS (1974) *Values of non-atomic games*. Princeton University Press, Princeton
- Basar T, Ho YC (1974) Informational properties of the Nash solutions of two stochastic nonzero-sum games. *J Econ Theory* 7:370-387
- Basar T, Olsder GJ (1999) *Dynamic noncooperative game theory*. SIAM, Philadelphia
- Bensoussan A, Frehse J (1984) Nonlinear elliptic systems in stochastic game theory. *J Reine Angew Math* 350:23-67

- Bergin J, Bernhardt D (1992) Anonymous sequential games with aggregate uncertainty. *J Math Econ* 21:543-562. North-Holland
- Cardaliaguet P (2012) Notes on mean field games. Collège de France
- Cardaliaguet P (2013) Long term average of first order mean field games and work KAM theory. *Dyn Games Appl* 3:473-488
- Correa JR, Stier-Moses NE (2010) In: Cochran JJ (ed) Wardrop equilibria. Wiley encyclopedia of operations research and management science. John Wiley & Sons, Chichester, UK
- Haurie A, Marcotte P (1985) On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* 15(3):295-308
- Ho YC (1980) Team decision theory and information structures. *Proc IEEE* 68(6):15-22
- Huang MY (2010) Large-population LQG games involving a major player: the Nash certainty equivalence principle. *SIAM J Control Optim* 48(5):3318-3353
- Huang MY, Caines PE, Malhamé RP (2003) Individual and mass behaviour in large population stochastic wireless power control problems: centralized and Nash equilibrium solutions. In: IEEE conference on decision and control, Maui, pp 98-103
- Huang MY, Malhamé RP, Caines PE (2006) Large population stochastic dynamic games: closed loop Kean-Vlasov systems and the Nash certainty equivalence principle. *Commun Inf Syst* 6(3):221-252
- Huang MY, Caines PE, Malhamé RP (2007) Large population cost-coupled LQG problems with non-uniform agents: individual-mass behaviour and decentralized ϵ - Nash equilibria. *IEEE Tans Autom Control* 52(9):1560-1571
- Jovanovic B, Rosenthal RW (1988) Anonymous sequential games. *J Math Econ* 17(1):77-87. Elsevier
- Kizilkale AC, Caines PE (2013) Mean field stochastic adaptive control. *IEEE Trans Autom Control* 58(4):905-920
- Lasry JM, Lions PL (2006a) Jeux à champ moyen. I - Le cas stationnaire. *Comptes Rendus Math* 343(9):619-625
- Lasry JM, Lions PL (2006b) Jeux à champ moyen. II - Horizon fini et controle optimal. *Comptes Rendus Math* 343(10):679-684
- Lasry JM, Lions PL (2007) Mean field games. *Jpn J Math* 2:229-260
- Li T, Zhang JF (2008) Asymptotically optimal decentralized control for large population stochastic multiagent systems. *IEEE Tans Autom Control* 53(7):1643-1660
- Neyman A (2002) Values of games with infinitely many players. In: Aumann RJ, Hart S (eds) Handbook of game theory, vol 3. North Holland, Amsterdam, pp 2121-2167
- Nourian M, Caines PE (2013) ϵ -Nash Mean field games theory for nonlinear stochastic dynamical systems with major and minor agents. *SIAM J Control Optim* 50(5):2907-2937
- Nguyen SL, Huang M (2012) Linear-quadratic-Gaussian mixed games with continuum-parametrized minor players. *SIAM J Control Optim* 50(5):2907-2937
- Tembine H, Zhu Q, Basar T (2012) Risk-sensitive mean field games. arXiv:1210.2806
- Wardrop JG (1952) Some theoretical aspects of road traffic research. In: Proceedings of the institute of civil engineers, London, part II, vol 1, pp 325-378
- Weintraub GY, Benkard C, Van Roy B (2005) Oblivious equilibrium: a mean field approximation for large-scale dynamic games. In: Advances in neural information processing systems. MIT, Cambridge

Mean Field Games

Peter Caines Montreal, Canada

DOI: 10.1007/SpringerReference_364780

URL: <http://www.springerreference.com/index/chapterdbid/364780>

Part of: Encyclopedia of Systems and Control

Editors: Dr. Tariq Samad and Professor John Baillieul

PDF created on: June, 20, 2014 14:09

© Springer-Verlag Berlin Heidelberg 2014