Mean Field Games

Peter E. Caines, Minyi Huang, and Roland P. Malhamé

Abstract

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 asymptotically negligible 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 to the infinite population problem is given by (i) the Hamilton-Jacobi-Bellman (HJB) equation of optimal control for a generic agent and (ii) the Fokker-Planck-Kolmogorov (FPK) equation for that agent, where these equations are linked by the probability distribution of the state of the generic agent, otherwise known as the system's mean field. Moreover, (i) and (ii) have an equivalent expression in terms of the stochastic maximum principle together with a McKean-Vlasov stochastic differential equation, and yet a third characterization is in terms of the so-called master equation. The article first describes problem areas which motivate the development of MFG theory and then presents the theory's basic mathematical formalization. The main results of MFG theory are then presented, namely the existence and uniqueness of infinite population Nash equilibiria, their approximating finite population ε-Nash equilibria, and the associated best response strategies. This is followed by a presentation of the three main

P.E. Caines (✉)
Department of Electrical and Computer Engineering, McGill University, Montreal, QC, Canada
e-mail: peterc@cim.mcgill.ca

M. Huang
School of Mathematics and Statistics, Carleton University, Ottawa, ON, Canada
e-mail: mhuang@math.carleton.ca

R.P. Malhamé
Department of Electrical Engineering, École Polytechnique de Montréal, Montreal, QC, Canada
e-mail: roland.malhame@polymtl.ca

T. Başar, G. Zaccour (eds.), *Handbook of Dynamic Game Theory*,
DOI 10.1007/978-3-319-27335-8_7-1

mathematical methodologies for the derivation of the principal results of the theory. Next, the particular topics of major-minor agent MFG theory and the common noise problem are briefly described and then the final section concisely presents three application areas of MFG theory.

Contents

1 Introduction . . . 2
1.1 The Fundamental Idea of Mean Field Game Theory . . . 2
1.2 Background . . . 3
1.3 Scope . . . 3
2 Problem Areas and Motivating Examples . . . 4
2.1 Engineering . . . 4
2.2 Economics and Finance . . . 4
2.3 Social Phenomena . . . 5
3 Mathematical Framework . . . 5
3.1 Agent Dynamics . . . 5
3.2 Agent Cost Functions . . . 7
3.3 Information Patterns . . . 10
3.4 Solution Concepts: Equilibria and Optima . . . 11
4 Analytic Methods: Existence and Uniqueness of Equilibria . . . 13
4.1 Linear-Quadratic Systems . . . 14
4.2 Nonlinear Systems . . . 14
4.3 PDE Methods and the Master Equation . . . 16
4.4 The Hybrid Approach: PDEs and SDEs, from Infinite to Finite Populations . . . 18
4.5 The Probabilistic Approach . . . 19
4.6 MFG Equilibrium Theory Within the Nonlinear Markov Framework . . . 20
5 Major and Minor Agents . . . 21
6 The Common Noise Problem . . . 22
7 Applications of MFG Theory . . . 23
7.1 Residential Power Storage Control for Integration of Renewables . . . 23
7.2 Communication Networks . . . 24
7.3 Stochastic Growth Models . . . 25
References . . . 25

1 Introduction

1.1 The Fundamental Idea of Mean Field Game Theory

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 asymptotically negligible 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 to the infinite population problem is given by (i) the Hamilton-Jacobi-Bellman (HJB) equation of optimal control for a generic agent and (ii) the Fokker-Planck-Kolmogorov (FPK) equation for that agent, where these equations are linked by the distribution of the state of the generic agent, otherwise known as the system's mean field. Moreover, (i) and (ii) have an equivalent expression in terms of the stochastic maximum principle together with a McKean-Vlasov stochastic differential equation, and yet a third characterization is in terms of the so-called master equation. An important feature of MFG solutions is that they have fixed-point properties regarding the individual responses to and the formation of the mean field which conceptually correspond to equilibrium solutions of the associated games.

1.2 Background

Large population dynamical multi-agent noncooperative 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. Here, large is taken to mean numerically large with respect to some implied normal range or infinite (as a discrete or uncountable set). Analysis of such systems with even a moderate number of agents is regarded as being extremely difficult using the finite population game theoretic methods which were developed over several decades for multi-agent control systems (see, e.g., Basar and Ho 1974; Ho 1980; Basar and Olsder 1999; and Bensoussan and Frehse 1984). In contrast to the dynamical system formulation of multi-agent games, 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 classical transportation analysis (Correa and Stier-Moses 2010; Haurie 1985; Wardrop 1952). However, dynamical (also termed sequential) stochastic games were analyzed in the continuum limit in the work of Jovanovic and Rosenthal (1988) and Bergin and Bernhardt (1992), where a form of the mean field equations can be recognized in a discrete-time dynamic programming equation linked with an evolution equation for the population state distribution.

Subsequently, what is now called MFG theory originated in the equations for dynamical games with (i) large finite populations of asymptotically negligible agents together with (ii) their infinite limits, in the work of (Huang et al. 2003, 2007), Huang et al. (2006) (where the framework was called the Nash certainty equivalence principle; see Caines (2014)) and independently in that of Lasry and Lions (2006a,b, 2007), where the now standard terminology of mean field games (MFGs) was introduced. The closely related notion of oblivious equilibria for large population dynamic games was also independently introduced by Weintraub et al. (2005, 2008) within the framework of discrete-time Markov decision processes (MDP).

1.3 Scope

The theory and methodology of MFG has rapidly developed since its inception and is still advancing. Consequently, the objective of this article is only to present the fundamental conceptual framework of MFG in the continuous time setting and the main techniques that are currently available. Moreover, the important topic of numerical methods will not be included, but it is addressed elsewhere in this volume by other contributors.

2 Problem Areas and Motivating Examples

Topics which motivate MFG theory or form potential areas of applications include the following:

2.1 Engineering

In the domain of power grid network control, an MFG methodology is being applied to create decentralized schemes for power network peak load reduction and compensation of fluctuations originating in renewable sources (see Sect. 7). Vast numbers of individual electric water-heating devices are planned to be coordinated in a decentralized way using an MFG architecture which would limit the required flows of information, such that individual controls give rise to a desired mean consumption.

For cell phone communication networks where coded signals can overlap in the frequency spectrum (called CDMA networks), a degradation of individual reception can occur when multiple users emit in the same frequency band. Compensation for this by users increasing their individual signal powers will shorten battery life and is collectively self-defeating. However, in the resulting dynamic game, a Nash equilibrium is generated when each cellular user controls its transmitted power as specified by MFG theory (see Sect. 7). Other applications include decentralized charging control of large populations of plug-in electric vehicles (Ma et al. 2013).

2.2 Economics and Finance

Human capital growth has been considered in an MFG setting by Guéant et al. (2011) and Lucas and Moll (2014) where the individuals invest resources (such as time and money) for the improvement of personal skills to better position themselves in the labor market when competing with each other.

Chan and Sircar (2015) considered the mean field generalization of Bertrand and Cournot games in the production of exhaustible resources where the price acts as a medium for the producers to interact. Furthermore, an MFG formulation has been used by Carmona et al. (2015) to address systemic risk as characterized by a large number of banks having reached a default threshold by a given time, where interbank loaning and lending is regarded as an instrument of control.

2.3 Social Phenomena

Closely related to the application of MFG theory to economics and finance is its potential application to a whole range of problems in social dynamics. As a short list of current examples, we mention:

2.3.1 Opinion Dynamics

The evolution of the density of the opinions of a mass of agents under hypotheses on the dynamics and stubbornness of the agents is analyzed in an MFG framework in Bauso et al. (2016).

2.3.2 Vaccination Games

When the cost to each individual is represented as a function of (a) the risk of side effects, (b) the benefits of being vaccinated, and (c) the proportion of the population which is vaccinated, as in Bauch and Earn (2004), it is evident that an MFG formulation is relevant, and this has been pursued in the work of Laguzet and Turinici (2015).

2.3.3 Congestion Studies

MFG methodology has been employed in the study of crowds and congested flows in Dogbé (2010) and Lachapelle and Wolfram (2011), where numerical methods reveal the possibility of lane formation.

3 Mathematical Framework

3.1 Agent Dynamics

In MFG theory individual agents are modeled by controlled stochastic systems which may be coupled by their dynamics, their cost functions, and their observation functions.

The principal classes of dynamical models which are used in MFG theory are sketched below; in all of them, the individual agent controls its own state process (invariably denoted here by x_i or x_θ) and is subject to individual and possibly common stochastic disturbances.

Concerning terminology, throughout this article, the term *strategy* of an agent means the functional mapping from an agent's information set to its control actions over time, in other words, the *control law* of that agent.

3.1.1 Diffusion Models

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

$$dx_i = (A_i x_i + B_i u_i)dt + C_i dw_i, \quad 1 \leq i \leq N, \tag{1}$$

on a finite or infinite time interval, where for the ith agent $\mathscr{A}_i$, $x_i \in \mathbb{R}^n$ is the state, $u_i \in \mathbb{R}^m$ the control input, and $w_i \in \mathbb{R}^r$ a standard Wiener process and where $\{w_i, 1 \leq i \leq N\}$ are independent processes. For simplicity, all collections of system initial conditions are taken to be independent and have finite second moment.

A simplified form of the general case is given by the following set of controlled SDEs which for each agent $\mathscr{A}_i$ includes state coupling with all other agents:

$$dx_i(t) = \frac{1}{N}\sum_{j=1}^{N} f(t, x_i(t), u_i(t), x_j(t))dt + \sigma dw_i(t) \tag{2}$$

$$= \int_{R^n} f(t, x_i(t), u_i(t), z)\{\frac{1}{N}\sum_{j=1}^{N}\delta_{x_j}(dz)\}dt + \sigma dw_i(t)$$

$$=: f[t, x_i(t), u_i(t), \{\frac{1}{N}\sum_{j=1}^{N}\delta_{x_j}\}]dt + \sigma dw_i(t) \tag{3}$$

$$= f[t, x_i(t), u_i(t), \mu_t^N]dt + \sigma dw_i(t), \tag{4}$$

where the function $f[\cdot,\cdot,\cdot,\cdot]$, with the empirical measure of the population states $\mu_t^N := \frac{1}{N}\sum_{j=1}^{N}\delta_{x_j}$ at the instant t as its fourth argument, is defined via

$$f[t, x(t), u(t), \nu_t] := \int_{R^n} f(t, x(t), u(t), z)\nu_t(dz), \tag{5}$$

for any measure flow ν_t, as in Cardaliaguet (2012) and Kolokoltsov et al. (2012). For simplicity, we do not consider diffusion coefficients depending on the system state or control.

Equation (2) is defined on a finite or infinite time interval, where, here, for the sake of simplicity, only the uniform (i.e., nonparameterized) 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 McKean-Vlasov equation

$$dx_t = f[x_t, u_t, \mu_t]dt + \sigma dw_t, \quad 1 \leq i \leq N, \;\; 0 \leq t \leq T,$$

where $f[x, u, \mu_t] = \int_{\mathbb{R}} f(x, u, y)\mu_t(dy)$, $\mu_t(\cdot)$ denotes the distribution of the state of the generic agent at $t \in [0, T]$ and the initial condition measure μ_0 is specified. (The dynamics used in Lasry and Lions (2006a,b, 2007) and Cardaliaguet (2012) are of the form $dx_i(t) = u_i(t)dt + dw_i(t)$, where u_i, x_i, w_i are scalar-valued processes.)

It is reasonable to speculate that results described below for the case of system dynamics driven by a Wiener process would hold in the general case of a Wiener process plus a point process and ultimately to the general case of Lévy processses;

indeed, in an operator framework, this generalization is carried out in the work of Kolokoltsov et al. (see below).

3.1.2 Nonlinear Markov Processes

The mean field game dynamic modeling framework has been significantly generalized by Kolokoltsov, Li, and Wei (2012) via the introduction of controlled nonlinear Markov processes where, in this framework, instead of diffusion SDEs, the evolution of a typical agent is described by an integrodifferential generator of Lévy-Khintchine type, where, as in the diffusion models described in the rest of this paper, the coefficients of the dynamical system of each agent, and its associated costs, are permitted to depend upon the empirical measure of the population of agents. As a consequence, by virtue of the Markov property, game theoretic best response problems in this framework can still be solved within the HJB formalism, and moreover the sensitivity analysis of the controls and dynamics with respect to perturbations in the population measure flow is facilitated.

3.1.3 Markov Chains and Other Discrete-Time Processes

The dynamical evolution of the state x_i of the ith agent $\mathscr{A}_i$ is formulated as a discrete-time Markov decision process (MDP). The so-called anonymous sequential games (Bergin and Bernhardt 1992; Jovanovic and Rosenthal 1988) deal with a continuum of agents, where a generic agent's cost function depends on its own state and action, and the joint state-action distribution of the agent population.

In the context of industry dynamics, Weintraub et al. (2005, 2008) adopted a large finite population, where the dynamics may be described by a Markov transition kernel model $P_{t+1} := P(x_i(t+1)|x_i(t), x_{-i}(t), u_i(t))$, where x_{-i} denotes the states of other players; also see Adlakha et al. (2015).

3.1.4 Finite State Models

Within continuous time modeling, Gomes et al. (2013) formulated a mean field game of switching among finite states and determined the equilibrium by a coupled system of ordinary differential equations. Finite state mean field games have applications in social-economic settings and networks (Gomes et al. 2014; Kolokoltsov and Malafeyev 2017; Kolokoltsov and Bensoussan 2016).

3.2 Agent Cost Functions

Throughout this article we shall only refer to cost functions which are the additive (or integral) composition over a finite or infinite time interval of instantaneous (running) costs; in MFG theory these will depend upon the individual state of an agent along with its control and possibly a function of the states of all other agents in the system. As usual in stochastic decision problems, the cost function for any agent will be defined by the expectation of the integrated running costs over all possible sample paths of the system. An important class of such functions is the so-called ergodic cost functions which are defined as the time average of integral cost functions.

3.2.1 Individual Agent Performance Functions in Noncooperative Games

The principal types of games considered in MFG theory are, first, noncooperative games, where each agent seeks to minimize its own loss represented by its cost function. In the most basic finite population linear-quadratic diffusion case, the agent $\mathscr{A}_i$, $1 \leq i \leq N$, possesses a cost function of the form:

$$J_i^N(u_i, u_{-i}) = E \int_0^T \{\|x_i(t) - Hm_N(t)\|_Q^2 + \|u_i(t)\|_R^2\} dt, \tag{6}$$

where $\|\cdot\|_M^2$ denotes the squared (semi-)norm arising from the positive semi-definite matrix M, where we assume the cost-coupling term 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 ith agent, $\overline{x}_N$ 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 Sect. 3.1.1, a corresponding finite population mean field cost function is

$$J_i^N(u_i, u_{-i}) := E \int_0^T \Big((1/N)\sum_{j=1}^N L(x_i(t), u_i(t), x_j(t))\Big) dt, \quad 1 \leq i \leq N, \tag{7}$$

where $L(\cdot)$ is the pairwise cost rate function. Setting the infinite population cost rate $L[x, u, \mu_t] = \int_{\mathbb{R}} L(x, u, y)\mu_t(dy)$, hence the corresponding infinite population expected cost for a generic agent $\mathscr{A}_i$ is given by

$$J_i(u_i, \mu) \ := \ E \int_0^T L[x(t), u_i(t), \mu_t] dt, \tag{8}$$

which is the general expression appearing in Huang et al. (2006) and Nourian and Caines (2013) and which includes those of Lasry and Lions (2006a,b, 2007), Cardaliaguet (2012). $e^{-\rho t}$ discounted costs are employed for infinite time horizon cost functions (Huang et al. 2003, 2007), while the long-run average cost is used for ergodic MFG problems (Bardi 2012; Lasry and Lions 2006a,b, 2007; Li and Zhang 2008).

3.2.2 Risk-Sensitive Performance Functions

This article will solely focus on additive type costs although other forms can be adopted for the individual agents. One important such form is a risk-sensitive cost function:

$$J_i^N(u_i, u_{-i}) = E \exp[\int_0^T (1/N) \sum_{j=1}^N L(x_i(t), u_i(t), x_j(t))dt],$$

which allows the use of dynamic programming to compute the best response. For related analyses in the linear-exponential-quadratic-Gaussian (LEQG) case, see, e.g., Tembine et al. (2014).

3.2.3 Performance Functions in Major-Minor Agent Systems

We start with the most basic finite population linear-quadratic case with a major agent $\mathscr{A}_0$ having state x_0 and N minor agents $\mathscr{A}_i$, $1 \le i \le N$, with states x_i. The SDEs of $\mathscr{A}_0$ and $\mathscr{A}_i$ are given by

$$dx_0 = (A_0 x_0 + B_0 u_0 + F_0 m_N)dt + D_0 dw_0,$$
$$dx_i = (Ax_i + Bu_i + Fm_N + Gx_0)dt + Ddw_i, \quad 1 \le i \le N,$$

where $m_N = \frac{1}{N}\sum_{i=1}^N x_i$ and the initial states are $x_0(0)$ and $x_i(0)$. The major agent has a cost function of the form:

$$J_0^N(u_0, u_{-0}) = E \int_0^T \{\|x_0(t) - H_0 m_N(t)\|^2_{Q_0} + \|u_i(t)\|^2_{R_0}\}dt,$$

and the minor agent $\mathscr{A}_i$ possesses a cost function of the form:

$$J_i^N(u_i, u_{-i}) = E \int_0^T \{\|x_i(t) - H_1 m_N(t) - H_2 x_0(t)\|^2_Q + \|u_i(t)\|^2_R\}dt.$$

Correspondingly, in the nonlinear case with a major agent, the N nonlinear equations in (2) are generalized to include the state of a major agent described by an additional SDE, giving a system described by $N+1$ equations. The cost functions are given by

$$J_0^N(u_0, u_{-0}) \quad := \quad E \int_0^T (1/N) \sum_{j=1}^N L_0(x_0(t), u_0(t), x_j(t))dt,$$

and

$$J_i^N(u_i, u_{-i}) \quad := \quad E \int_0^T (1/N) \sum_{j=1}^N L(x_i(t), u_i(t), x_0(t), x_j(t))dt.$$

Consequently, the infinite population mean field cost functions for the major and minor agents respectively are given by

$$J_0(u_0, \mu) \;:=\; E \int_0^T L_0[x_0(t), u_0(t), \mu_t]dt,$$

and

$$J_i(u_i, \mu) \;:=\; E \int_0^T L[x_i(t), u_i(t), x_0(t), \mu_t]dt,$$

where $L_0[x_0(t), u_0(t), \mu_t]$ and $L[x_i(t), u_i(t), x_0(t), \mu_t]$ correspond to their finite population versions as in the basic minor agent only case.

3.3 Information Patterns

We now introduce the following definitions and characterizations of information patterns in dynamic game theory which shall be used in the rest of this article.

The States of a Set of Agents: A state in dynamic games is taken to be either (i) an individual (agent) state as defined in the Sect. 3.1, in which case it will constitute a component of the global system state, namely, the union of the individual states, or (ii) the global state, which is necessarily sufficient to describe the dynamical evolution of all the agents once the system inputs are specified. We emphasize that in this setting (see, for instance, (2)), only knowledge of the entire system state (i.e., the union of all the individual states) plus all the system inputs would in general permit such an extrapolation.

Moreover, in the infinite population case, the (global) system state may refer to the statistical or probability distribution of the population of individual states, i.e., the *mean field*.

Variety of Information Patterns: Information on dynamical states: For any given agent, this may constitute (i) the initial state, (ii) the partial past history, or (iii) the purely current state values of either (i) that individual agent or (ii) a partial set of all the agents or (iii) the entire set of the agents.

Open-Loop and Closed-Loop Control Laws: The common definition of an open-loop control law for an agent is that it is solely a function of the information set consisting of time and the initial state of that agent or of the whole system (i.e., the global initial state). A closed-loop (i.e., feedback) control law is one which is a function of time and the current state of that agent or the global state of the system subject to the given information pattern constraints, where a particular case of importance is that in which an agent's strategy at any instant depends only upon its current state.

A significant modification of the assertion above must be made in the classical mean field game situation with no common noise or correlating major agent; indeed, in that case all agents in the population will be employing an infinite population-based Nash equilibrium strategy. As a result, the probability distribution of the generic agent, which can be identified with the global state as defined earlier, becomes deterministically predictable for all future times, provided it is known

at the initial time. Consequently, for such MFGs, an adequate characterization of sufficient information is the initial global state. Furthermore, in the MFG framework among others, an agent lacking complete observations on its own current state may employ recursive filtering theory to estimate its own state.
Statistical Information on the Population: For any individual agent, information is typically available on that agent's own dynamics (i.e., the structure and the parameters of its own controlled dynamic system); but it is a distinct assumption that no such information is available to it concerning other individual agents. Furthermore, this information may be available in terms of a distribution (probabilistic or otherwise) over the population of agents and not associated with any identifiable individual agent. This is particularly the case in the MFG framework.
Who Knows What About Whom: A vital aspect of information patterns in game theory is that knowledge concerning (i) other agents' control actions, or, more generally, concerning (ii) their strategies (i.e., their control laws), may or may not be available to any given agent. This is a fundamental issue since the specification of an agent's information pattern in terms of knowledge of other agent's states, system dynamics, cost functions, and parameters leads to different possible methods to solve for different types of equilibrium strategies and even for their existence.

In the MFG setting, if the common assumption is adopted that all agents will compute their best response in reaction to the best responses of all other agents (through the system dynamics), it is then optimal for each agent to solve for its strategy through the solution of the MFG equations. The result is that each agent will know the control strategy of every other agent, but not its control action since the individual state of any other agent is not available to a given agent. Note however that the state distribution of any (i.e., other generic) agent is known by any given generic agent since this is the system's mean field which is generated by the MFG equations.

3.4 Solution Concepts: Equilibria and Optima

In contrast to the situation in classical stochastic control, in the game theoretic context, the notion of an optimal level of performance and associated optimal control for the entire system is in general not meaningful. The fundamental solution concept is that of an equilibrium and here we principally consider the notion of a Nash equilibrium.

3.4.1 Equilibria in Noncooperative Games: Nash Equilibria

For a set of agents $\mathscr{A}_i, 1 \leq i \leq N < \infty$, let $\mathscr{U}^N := \mathscr{U}_1 \times \ldots \times \mathscr{U}_N$ denote the joint admissible strategy space, where each space $\mathscr{U}_i$ consists of a set of strategies (i.e., control laws) u_i which are functions of information specified for $\mathscr{A}_i$ via the underlying information pattern.

The joint strategy $u \triangleq (u_1, \ldots, u_N)$ (sometimes written as $\{u_i, 1 \leq i \leq N\}$) lying in $\mathcal{U}_1 \times \ldots \times \mathcal{U}_N$ will constitute an input for a specific system in one of the classes specified in Sects. 3.1 and 3.2.

The joint strategy (or control law) $u^{\circ,N} \triangleq \{u_i^\circ, 1 \leq i \leq N\} \in \mathcal{U}^N$ is said to generate an *ε-Nash equilibrium*, $\varepsilon \geq 0$, if for each i,

$$J_i^N(u_i^\circ, u_{-i}^\circ) - \varepsilon \leq \inf_{u_i \in \mathcal{U}_i} J_i^N(u_i, u_{-i}^\circ) \leq J_i^N(u_i^\circ, u_{-i}^\circ). \tag{9}$$

In case $\varepsilon = 0$, the equilibrium is called a *Nash equilibrium*.

This celebrated concept has the evident interpretation that when all agents except agent $\mathcal{A}_i$ employ a set of control laws $\{u_j^\circ, j \neq i, 1 \leq j \leq N\}$, any deviation by $\mathcal{A}_i$ from u_i° can yield a cost reduction of at most ε.

In the MFG framework, in its basic noncooperative formulation, the objective of each agent is to find strategies (i.e., control laws) which are compatible with respect to the information pattern and other dynamical constraints and which minimize its individual performance function. Consequently the resulting problem is necessarily game theoretic and the central results of the topic concern the existence of Nash Equilibria and their properties.

For a system of N players, under the hypothesis of closed-loop state information (see Sect. 3.3), we shall define the *set of value functions* $\{V_i(t,x), 1 \leq i \leq N,\}$ *in a Nash equilibrium*, as the set of costs of N agent $\mathcal{A}_i, 1 \leq i \leq N$, with respect to the time and *global state* pair (t, x). The set of value functions and its existence may be characterized by a set of coupled HJB equations.

Under closed-loop information, the Nash equilibrium, if it exists, is *sub-game perfect* in the sense that by restricting to any remaining period of the original game, the set of strategies is still a Nash equilibrium for the resulting sub-game. In this case, the strategy of each agent is determined as a function of time and the current states of the agents and is usually called a *Markov strategy*.

3.4.2 Pareto Optima

A set of strategies yields a *Pareto optimum* if a change of strategies which strictly decreases the cost incurred by one agent strictly increases the cost incurred by at least one other agent.

3.4.3 Social Optima and Welfare Optimization

Within the framework of this article, a *social cost* or *(negative) welfare function* is defined as the sum of the individual cost functions of a set of agents (Huang et al. 2012). As a result a cooperative game may be defined which consists of the agents minimizing the social cost as a cooperative optimal control problem, where the individual strategies will depend upon the information pattern. We observe that a social optimum is necessarily a Pareto optimum with respect to the vector of individual costs since otherwise at least one summand of the social cost function may be strictly reduced without any other agent's cost increasing. The so-called *person-by-person optimality* is the property that at the social optimum, the strategy change of any single agent can yield no improvement of the social cost and so provides a useful necessary condition for social optimality. The exact solution of this problem in general requires a centralized information pattern.

3.4.4 Team Optima

Team problems are distinguished from cooperative game problems by the fact that only one cost function is defined a priori for the entire set of agents while they have access to different sets of information. A necessary condition for a solution to be team optimal is that the person-by-person optimality condition is satisfied (Ho 1980). Team problems do not in general reduce to single agent optimum problems due to the variety of information patterns that are possible for the set of agents.

3.4.5 Mean Field Type Control Optimality

Mean field type control deals with optimal control problems where the mean field of the state process either is involved in the cost functional in a nonlinear manner, such as being associated with the variance of the state, or appears in the system dynamics, or both, and is a function of the single agent's control. Unlike standard stochastic optimal control problems, mean field type control problems do not possess an iterated expectation structure due to the mean field term (i.e., there is time inconsistency), which excludes the direct (i.e., without state extension) application of dynamic programming. In this case, the stochastic maximum principle is an effective tool for characterizing the optimal control; see Andersson and Djehiche (2011) and the monograph of Bensoussan et al. (2013). Carmona et al. (2013) considered a closely related problem termed the control of McKean-Vlasov dynamics. Mean field games dealing with mean field type dynamics and costs and addressing time consistency are considered by Djehiche and Huang (2016).

4 Analytic Methods: Existence and Uniqueness of Equilibria

The objective of each agent in the classes of games under consideration is to find strategies which are admissible with respect to the given dynamic and information constraints and which achieve one of corresponding types of equilibria or optima described in the previous section. In this section we present some of the main analytic methods for establishing the existence, uniqueness, and the nature of the related control laws and their equilibria.

The fundamental feature of MFG theory is the relation between the game theoretic behavior (assumed here to be noncooperative) of finite populations of agents and the infinite population behavior characterized by a small number of equations in the form of PDEs or SDEs.

4.1 Linear-Quadratic Systems

The basic mean field problem in the linear-quadratic case 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 $\mathscr{A}_\theta$ (representing a generic agent $\mathscr{A}_i$ taking parameter θ) with parameter distribution $F(\theta), \theta \in \Theta$ and system parameters identified as $A_\theta = a_\theta, B_\theta = b_\theta, Q := 1$, $R = r > 0$, $H = \gamma$; the extension to the vector case and more general parameter dependence on θ is straightforward. The so-called Nash certainty equivalence (NCE) scheme generating the equilibrium solution takes the form:

$$\rho s_\theta = \frac{ds_\theta}{dt} + a_\theta s_\theta - \frac{b_\theta^2}{r}\Pi_\theta s_\theta - x^*, \tag{10}$$

$$\frac{d\overline{x}_\theta}{dt} = (a_\theta - \frac{b_\theta^2}{r}\Pi_\theta)\overline{x}_\theta - \frac{b_\theta^2}{r}s_\theta, \qquad 0 \le t < \infty, \tag{11}$$

$$\overline{x}(t) = \int_\Theta \overline{x}_\theta(t) dF(\theta), \tag{12}$$

$$x^*(t) = \gamma(\overline{x}(t) + \eta), \tag{13}$$

$$\rho\Pi_\theta = 2a_\theta\Pi_\theta - \frac{b_\theta^2}{r}\Pi_\theta^2 + 1, \quad \Pi_\theta > 0, \quad \text{Riccati Equation} \tag{14}$$

where the control law of the generic parameterized agent $\mathscr{A}_\theta$ has been substituted into the system equation (1) and is given by $u_\theta^0(t) = -\frac{b_\theta}{r}(\Pi_\theta x_\theta(t) + s_\theta(t))$, $0 \le t < \infty$. u_θ^0 is the optimal tracking feedback law with respect to $x^*(t)$ which is an affine function of the mean field term $\overline{x}(t)$, the average with respect to the parameter distribution F of the $\theta \in \Theta$ parameterized state means $x_\theta(t)$ of the agents. Subject to the conditions for the NCE scheme to have a solution, each agent is necessarily in a Nash equilibrium with respect to all full information causal (i.e., non-anticipative) feedback laws with respect to the remainder of agents when these are employing the law u_θ^0 associated with their own parameter.

It is an important feature of the best response control law u_θ^0 that its form depends only on the parametric distribution F of the entire set of agents, and at any instant it is a feedback function of only the state of the agent $\mathscr{A}_\theta$ itself and the deterministic mean field-dependent offset s_θ, and is thus decentralized.

4.2 Nonlinear Systems

For the general nonlinear case, the MFG equations on $[0, T]$ are given by the linked equations for (i) the value function V for each agent in the continuum, (ii) the FPK equation for the 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 agent scalar case, these take the following form:

The mean field game HJB-FPK equations are as follows:

$$[\text{HJB}] \quad -\frac{\partial V(t,x)}{\partial t} = \inf_{u \in U} \left\{ f[x,u,\mu_t]\frac{\partial V(t,x)}{\partial x} + L[x,u,\mu_t] \right\} + \frac{\sigma^2}{2}\frac{\partial^2 V(t,x)}{\partial x^2} \tag{15}$$

$$V(T,x) = 0,$$

$$[\text{FPK}] \quad \frac{\partial p_\mu(t,x)}{\partial t} = -\frac{\partial\{f[x,u^\circ(t,x),\mu_t]p_\mu(t,x)\}}{\partial x} + \frac{\sigma^2}{2}\frac{\partial^2 p_\mu(t,x)}{\partial x^2} \tag{16}$$

$$(t,x) \in [0,T] \times \mathbb{R}$$

$$p_\mu(0,x) = p_{\mu_0}(x),$$

$$[\text{BR}] \quad u^\circ(t,x) = \varphi(t,x|\mu_t), \tag{17}$$

where $p_\mu(t,\cdot)$ is the density of the measure μ_t, which is assumed to exist, and the function $\varphi(t,x|\mu_t)$ is the infimizer in the HJB equation. The (t,x,μ_t)-dependent feedback control gives an optimal control (also known as a best response (BR) strategy) for the generic individual agent with respect to the infinite population-dependent performance function (8) (where the infinite population is represented by the generic agent measure μ).

By the very definition of the solution to FPK equations, the solution μ above will be the state distribution in the process distribution solution pair (x,μ) in

$$[\text{SDE}] \quad dx_t = f[x_t,u_t^\circ,\mu_t]dt + \sigma dw_t, \quad 1 \le i \le N, \quad 0 \le t \le T. \tag{18}$$

This equivalence of the controlled sample path pair (x,μ) solution to the SDE and the corresponding FPK PDE is very important from the point of view of the existence, uniqueness, and game theoretic interpretation of the solution to the system's equation.

A solution to the mean field game equations above may be regarded as an equilibrium solution for an infinite population game in the sense that each BR feedback control (generated by the HJB equation) enters an FPK equation – and hence the corresponding SDE – and so generates a pair (x,μ), where each generic agent in the infinite population with state distribution μ solves the same optimization problem and hence regenerates μ.

In this subsection we briefly review the main methods which are currently available to establish the existence and uniqueness of solutions to various sets of MFG equations. In certain cases the methods are based upon iterative techniques which converge subject to various well-defined conditions. The key feature of the methods is that they yield individual state and mean field-dependent feedback control laws generating ε-Nash equilibria together with an upper bound on the approximation error.

The general nonlinear MFG problem is approached by different routes in the basic sets of papers Huang et al. (2007, 2006), Nourian and Caines (2013), Carmona and Delarue (2013) on one hand, and Lasry and Lions (2006a,b, 2007), Cardaliaguet (2012), Cardaliaguet et al. (2015), Fischer (2014), Carmona and Delarue (2014) on the other. Roughly speaking, the first set uses an infinite to finite population approach (to be called the top-down approach) where the infinite population game equations are first analyzed by fixed-point methods and then ε-Nash equilibrium results are obtained for finite populations by an approximation analysis, while the latter set analyzes the Nash equilibria of the finite population games, with each agent using only individual state feedback, and then proceeds to the infinite population limit (to be called the bottom-up approach).

4.3 PDE Methods and the Master Equation

In Lasry and Lions (2006a,b, 2007), it is proposed to obtain the MFG equation system by a finite N agent to infinite agent (or bottom-up) technique of solving a sequence of games with an increasing number of agents. Each solution would then give a Nash equilibrium for the corresponding finite population game. In this framework there are then two fundamental problems to be tackled: first, the proof of the convergence, in an appropriate sense, of the finite population Nash equilibrium solutions to limits which satisfy the infinite population MFG equations and, second, the demonstration of the existence and uniqueness of solutions to the MFG equations.

In the expository notes of Cardaliaguet (2012), the analytic properties of solutions to the infinite population HJB-FPK PDEs of MFG theory are established for finite time horizon using PDE methods including Schauder fixed-point theory and the theory of viscosity solutions. The relation to finite population games is then derived, that is to say an ε-Nash equilibrium result is established, predicated upon the assumption of strictly individual state feedback for the agents in the sequence of finite games. We observe that the analyses in both cases above will be strongly dependent upon the hypotheses concerning the functional form of the controlled dynamics of the individual agents and their cost functions, each of which may possibly depend upon the mean field measure.

4.3.1 Basic PDE Formulation

In the exposition of the basic analytic MFG theory (Cardaliaguet 2012), agents have the simple dynamics:

$$dx_t^i = u_t^i dt + \sqrt{2} dw_t^i \tag{19}$$

and the cost function of agent i is given in the form:

$$J_i^N(u_i, u_{-i}) = E\int_0^T \left[\frac{1}{2}(u_t^i)^2 + F(x_t^i, \mu_t^{N,-i})\right] dt + EG(x_t^i, \mu_t^{N,-i}),$$

where $\mu_t^{N,-i}$ is the empirical distribution of the states of all other agents. This leads to MFG equations in the simple form:

$$-\partial_t V - \Delta V + \frac{1}{2}|DV|^2 = F(x, m), \quad (x, t) \in \mathbb{R}^d \times (0, T) \tag{20}$$

$$\partial_t m - \Delta m - \text{div}(mDV) = 0, \tag{21}$$

$$m(0) = m_0, \quad V(x, T) = G(x, m(T)), \quad x \in \mathbb{R}^d, \tag{22}$$

where $V(t, x)$ and $m(t, x)$ are the value function and the density of the state distribution, respectively.

The first step is to consider the HJB equation with some fixed measure μ; it is shown by use of the Hopf-Cole transform that a unique Lipschitz continuous solution v to the new HJB equation exists for which a certain number of derivatives are Hölder continuous in space and time and for which the gradient Dv is bounded over R^n.

The second step is to show that the FPK equation with DV appearing in the divergence term has a unique solution function which is as smooth as V. Moreover, as a time-dependent measure, m is Hölder continuous with exponent $\frac{1}{2}$ with respect to the Kantorovich-Rubinstein (KR) metric.

Third, the resulting mapping of μ to V and thence to m, denoted Ψ, is such that Ψ is a continuous map from the (KR) bounded and complete space of measures with finite second moment (hence a compact space) into the same. It follows from Schauder fixed-point theorem that Ψ has a fixed point, which consequently constitutes a solution to the MFG equations with the properties listed above.

The fourth and final step is to show that the Lasry-Lions monotonicity condition (a form of strict passivity condition) on F

$$\int_{\mathbb{R}^d} F(x, m_1) - F(x, m_2) d(m_1 - m_2) > 0, \qquad \forall m_1 \neq m_2,$$

combined with a similar condition for G allowing for equality implies the uniqueness of the solution to the MFG equations.

Within the PDE setting, in-depth regularity investigation of the HJB-FPK equation under different growth and convexity conditions on the Hamiltonian have been developed by Gomes and Saude (2014).

4.3.2 General Theory: The Master Equation Method

The master equation formulation was initially introduced by P-L Lions and has been investigated by various researchers (Bensoussan et al. 2013; Cardaliaguet et al. 2015; Carmona and Delarue 2014).

The program of working from the finite population game equations and their solution to the infinite population MFG equations and their solution has been carried out in Cardaliaguet et al. (2015) for a class of systems with simple dynamics but which, in an extension of the standard MFG theory, include a noise process common to all agents in addition to their individual system noise processes. The basic idea is to reinterpret the value function of a typical player in a game of N players as a function $U(t, x_i, m)$ of time, its own state, and the empirical distribution of the states of all other players.

The analysis using the master equation begins with a set of equations which may be interpreted as the dynamic programming equations for the population. Furthermore, the information set permitted for this optimization is the full N agent system state. The derivation of the master equation is carried out (on an appropriately dimensioned torus) by arguing that the value function of a representative agent i, from a population of N, is a function of time, its own state, and a measure formed by $N-1$ particle states and by taking the limit when $N \to \infty$. In the end, the state space for the master equation is the joint space of a generic agent state and a probability distribution.

The main result of the extensive analysis in Cardaliaguet et al. (2015) is the convergence of the set of Nash value functions $V_i^N(t_0, x)$, $1 \le i \le N$, of the set of agents for the population of size N, with initial condition $x = (x_1, \ldots, x_N)$ at time t_0, to the set of corresponding infinite population value functions U (given as solutions to the master equation), evaluated at the corresponding initial state x_i and the empirical measure m_x^N. This convergence is in the average sense

$$\frac{1}{N}\Sigma|V_i^N(t_0,x) - U(t_0, x_i, m_x^N)| \to 0, \qquad \text{as} \qquad N \to \infty.$$

In Bensoussan et al. (2015), following the derivation of the master equation for mean field type control, the authors apply the standard approach of introducing the system of HJB-FPK equations as an equilibrium solution, and then the Master equation is obtained by decoupling the HJB equation from the Fokker-Planck-Kolmogorov equation. Carmona and Delarue (2014) take a different route by deriving the master equation from a common optimality principle of dynamic programming with constraints. Gangbo and Swiech (2015) analyze the existence and smoothness of the solution for a first-order master equation which corresponds to a mean field game without involving Wiener processes.

4.4 The Hybrid Approach: PDEs and SDEs, from Infinite to Finite Populations

The infinite to finite route is top-down: one does not solve the game of N agents directly. The solution procedure involves four steps. First, one passes directly to the infinite population situation and formulates the dynamical equation and cost function for a single agent interacting with an infinite population possessing a fixed state distribution μ. Second, the stochastic optimization problem for that generic agent is then solved via dynamic programming using the HJB equation and the resulting measure for the optimally controlled agent is generated via the agent's SDE or equivalently FPK equation. Third, one solves the resulting fixed-point problem by the use of various methods (e.g., by employing the Banach, Schauder, or Kakutani fixed-point theorems). Finally, fourth, it is shown that the infinite population Nash equilibrium control laws are ε-Nash equilibrium for finite populations. This formulation was introduced in the sequence of papers Huang et al. (2003, 2006, 2007) and used in Nourian and Caines (2013), Sen and Caines (2016); it corresponds to the "limit first" method employed by Carmona, Delarue, and Lachapelle (2013) for mean field games.

Specifically, subject to Lipschitz and differentiability conditions on the dynamical and cost functions, and adopting a contraction argument methodology, one establishes the existence of a solution to the HJB-FPK equations via the Banach fixed-point theorem; the best response control laws obtained from these MFG equations are necessarily Nash equilibria within all causal feedback laws for the infinite population problem. Since the limiting distribution is equal to the original measure μ, a fixed point is obtained; in other words a consistency condition is satisfied. By construction this must be (i) a self-sustaining population distribution when all agents in the infinite population apply the corresponding feedback law, and (ii) by its construction via the HJB equation, it must be a Nash equilibrium for the infinite population. The resulting equilibrium distribution of a generic agent is called the mean field of the system.

The infinite population solution is then related to the finite population behavior by an ε-Nash equilibrium theorem which states that the cost of any agent can be reduced by at most ε when it changes from the infinite population feedback law to another while all other agents stick to their infinite population-based control strategies. Specifically, it is then shown (Huang et al. 2006) that the set of strategies $\{u_i^\circ(t) = \varphi_i(t, x_i(t)|\mu_t), 1 \leq i \leq N\}$ yields an ε-Nash equilibrium for all ε, i.e., for all $\varepsilon > 0$, there exists $N(\varepsilon)$ such that for all $N \geq N(\varepsilon)$

$$J_i^N(u_i^\circ, u_{-i}^\circ) - \varepsilon \leq \inf_{u_i \in \mathscr{U}_i} J_i^N(u_i, u_{-i}^\circ) \leq J_i^N(u_i^\circ, u_{-i}^\circ). \tag{23}$$

4.5 The Probabilistic Approach

4.5.1 Maximum Principle Solutions Within the Probabilistic Formulation

A different solution framework for the mean field game with nonlinear diffusion dynamics is to take a stochastic maximum principle approach (Carmona and Delarue 2013) for determining the best response of a representative agent. The procedure is carried out in the following steps: (i) A measure flow μ_t is introduced to specify the empirical state distribution associated with an infinite population. (ii) An optimal control problem is solved for that agent by introducing an adjoint process, which then determines the closed-loop system. (iii) The measure flow μ_t is then required to be equal to the law of the closed-loop state processes. This procedure yields a McKean-Vlasov forward-backward stochastic differential equation.

Necessary and sufficient conditions are available to establish the validity of the stochastic maximum (SM) principle approach to MFG theory. In particular, convexity conditions on the dynamics and the cost function, with respect to the state and controls, may be taken as sufficient conditions for the main results characterizing an MFG equilibrium through the solution of the forward-backward stochastic differential equations (FBSDEs), where the forward equation is that of the optimally controlled state dynamics and the backward equation is that of the adjoint process generating the optimal control, where these are linked by the mean field measure process. Furthermore the Lasry-Lions monotonicity condition on the cost function with respect to the mean field forms the principal hypothesis yielding the uniqueness of the solutions.

4.5.2 Weak Solutions Within the Probabilistic Formulation

Under a weak formulation of mean field games (Carmona and Lacker 2015), the stochastic differential equation in the associated optimal control problem is interpreted according to a weak solution. This route is closely related to the weak formulation of stochastic optimal control problems, also known as the martingale approach.

The solution of the mean field game starts by fixing a mean field, as a measure to describe the effect of an infinite number of agents, and a nominal measure for the probability space. Girsanov's transformation is then used to define a new probability measure under which one determines a diffusion process with a controlled drift and a diffusion term. Subsequently, the optimal control problem is solved under this new measure. Finally the consistency condition is introduced such that the distribution of the closed-loop state process agrees with the mean field. Hence the existence and uniqueness of solutions to the MFG equations under the specified conditions are obtained for weak solutions.

The proof of existence under the weak formulation relies on techniques in set-valued analysis and a generalized version of Kakutani's theorem.

4.6 MFG Equilibrium Theory Within the Nonlinear Markov Framework

The mean field game dynamic modeling framework is significantly generalized by Kolokoltsov, Li, and Wei (2012) via the introduction of controlled nonlinear Markov processes where, instead of diffusion SDEs, the evolution of a typical agent is described by an integrodifferential generator of Levy-Khintchine type; as a consequence, by virtue of the Markov property, optimal control problems in this framework can still be solved within the HJB formalism.

Similar to the diffusion models described in the rest of this paper, the coefficients of the dynamical system of each agent, and its associated costs, are permitted to depend upon the empirical measure of the population of agents.

In the formal analysis, again similar to the procedures in the rest of the paper (except for Cardaliaguet et al. 2015), this measure flow is initially fixed at the infinite population limit measure and the agents then optimize their behavior via the corresponding HJB equations.

Finally, invoking the standard consistency requirement of mean field games, the MFG equations are obtained when the probability law of the resulting closed-loop configuration state is set equal to the infinite population distribution (i.e., the limit of the empirical distributions).

Concerning the methods used by Kolokoltsov, Li, and Wei (2012), we observe that there are two methodologies which are employed to ensure the existence of solutions to the kinetic equations (corresponding to the FPK equations in this generalized setup) and the HJB equations: First, (i) the continuity of the mapping from the population measure to the measure generated by the kinetic (i.e., generalized FPK) equations is proven, and (ii) the compactness of the space of measures is established; then (i) and (ii) yield the existence (but not necessarily uniqueness) of a solution measure corresponding to any fixed control law via the Schauder fixed-point theory. Second, an estimate of the sensitivity of the best response mapping (i.e., control law) with respect to an a priori fixed measure flow is proven by an application of the Duhamel principle to the HJB equation. This analysis provides the ingredients which are then used for an existence theory for the solutions to the joint FPK-HJB equations of the MFG. Within this framework an ε-Nash equilibrium theory is then established in a straightforward manner.

5 Major and Minor Agents

The basic structure of mean field games can be remarkably enriched by introducing one or more major agents to interact with a large number of minor agents. A major agent has significant influence, while a minor agent has negligibly small influence on others. Such a differentiation of the strength of agents is well motivated by many practical decision situations, such as a sector consisting of a dominant corporation and many much smaller firms, the financial market with institutional traders and a huge number of small traders. The traditional game theoretic literature has studied such models of mixed populations and coined the name mixed games, but this is only in the context of static cooperative games (Haimanko 2000; Hart 1973; Milnor and Shapley 1978).

Huang (2010) introduced a large population LQG game model with mean field couplings which involves a large number of minor agents and also a major agent. A distinctive feature of the mixed agent MFG problem is that even asymptotically (as the population size N approaches infinity), the noise process of the major agent causes random fluctuation of the mean field behavior of the minor agents. This is in contrast to the situation in the standard MFG models with only minor agents. A state-space augmentation approach for the approximation of the mean field behavior of the minor agents is taken in order to Markovianize the problem and hence to obtain ε-Nash equilibrium strategies. The solution of the mean field game reduces to two local optimal control problems, one for the major agent and the other for a representative minor agent.

Nourian and Caines (2013) extend the LQG model for major and minor (MM) agents (Huang 2010) to the case of a nonlinear MFG systems. The solution to the mean field game problem is decomposed into two nonstandard nonlinear stochastic optimal control problems (SOCPs) with random coefficient processes which yield forward adapted stochastic best response control processes determined from the solution of (backward in time) stochastic Hamilton-Jacobi-Bellman (SHJB) equations. A core idea of the solution is the specification of the conditional distribution of the minor agent's state given the sample path information of the major agent's noise process. The study of mean field games with major-minor agents and nonlinear diffusion dynamics has also been developed in Carmona and Zhu (2016) and Bensoussan et al. (2013) which rely on the machinery of FBSDEs.

An extension of the model in Huang (2010) to the systems of agents with Markov jump parameters in their dynamics and random parameters in their cost functions is studied in Wang and Zhang (2012) for a discrete-time setting.

In MFG problems with purely minor agents, the mean field is deterministic, and this obviates the need for observations on other agents' states so as to determine the mean field. However, a new situation arises for systems with a major agent whose state is partially observed; in this case, best response controls generating equilibria exist which depend upon estimates of the major agent's state (Sen and Caines 2016; Caines and Kizilkale 2017).

6 The Common Noise Problem

An extension of the basic MFG system model occurs when what is called common noise is present in the global system, that is to say there is a common Wiener process whose increments appear on the right-hand side of the SDEs of every agent in the system (Ahuja 2016; Bensoussan et al. 2015; Cardaliaguet et al. 2015; Carmona and Delarue 2014). Clearly this implies that asymptotically in the population size, the individual agents cannot be independent even when each is using local state plus mean field control (which would give rise to independence in the standard case). The study of this case is well motivated by applications such as economics, finance, and, for instance, the presence of common climatic conditions in renewable resource power systems.

There are at least two approaches to this problem. First it may be treated explicitly in an extension of the master equation formulation of the MFG equations as indicated above in Sect. 4 and, second, common noise may be taken to be the state process of a passive (that is to say uncontrolled) major agent whose state process enters each agent's dynamical equation (as in Sect. 3.2.3).

This second approach is significantly more general than the former since (i) the state process of a major agent will typically have nontrivial dynamics, and (ii) the state of the major agent typically enters the cost function of each agent, which is not the case in the simplest common noise problem. An important difference in these treatments is that in the common noise framework, the control of each agent will be a function of (i.e., it is measurable with respect to) its own Wiener process and the common noise, while in the second formulation each minor agent can have complete, partial, or no observations on the state of the major agent which in this case is the common noise.

7 Applications of MFG Theory

As indicated in the Introduction, a key feature of MFG theory is the vast scope of its potential applications of which the following is a sample: Smart grid applications: (i) Dispersed residential energy storage coordinated as a virtual battery for smoothing intermittent renewable sources (Kizilkale and Malhamé 2016); (ii) The recharging control of large populations of plug-in electric vehicles for minimizing system electricity peaks (Ma et al. 2013). Communication systems: (i) Power control in cellular networks to maintain information throughput subject to interference (Aziz and Caines 2017); (ii) Optimization of frequency spectrum utilization in cognitive wireless networks; (iii) Decentralized control for energy conservation in ad hoc environmental sensor networks. Collective dynamics: (i) Crowd dynamics with xenophobia developing between two groups (Lachapelle and Wolfram 2011) and collective choice models (Salhab et al. 2015); (ii) Synchronization of coupled oscillators (Yin et al. 2012). Public health models: Mean field game-based anticipation of individual vaccination strategies (Laguzet and Turinici 2015).

For lack of space and in what follows, we further detail only three examples, respectively, drawn among smart grid applications, communication system applications, and economic applications as follows.

7.1 Residential Power Storage Control for Integration of Renewables

The general objective in this work is to coordinate the loads of potentially millions of dispersed residential energy devices capable of storage, such as electric space heaters, air conditioners, or electric water heaters; these will act as a virtual battery whose storage potential is directed at mitigating the potentially destabilizing effect of the high power system penetration of renewable intermittent energy sources (e.g., solar and wind). A macroscopic level model produces tracking targets for the mean temperature of the controlled loads, and then an application of MFG theory generates microscopic device level decentralized control laws (Kizilkale and Malhamé 2016).

A scalar linear diffusion model with state x_i is used to characterize individual heated space dynamics and includes user activity-generated noise and a heating source. A quadratic cost function is associated with each device which is designed so that (i) pressure is exerted so that devices drift toward z which is set either to their maximum acceptable comfort temperature H if extra energy storage is desired or to their minimum acceptable comfort temperature $L < H$ if load deferral is desired and (ii) average control effort and temperature excursions away from initial temperature are penalized. This leads to the cost function:

$$E\int_0^{\infty} e^{-\delta t}[q_t(x_i - z)^2 + q_{x_0}(x_i - x_i(0))^2 + ru_i^2]dt, \tag{24}$$

where $q_t = \lambda|\int_0^t (\bar{x}(\tau) - y)d\tau|$, $\bar{x}$ is the population average state with $L < \bar{x}(0) < H$, and $L < y < H$ is the population target. In the design the parameter $\lambda > 0$ is adjusted to a suitable level so as to generate a stable population behavior.

7.2 Communication Networks

The so-called CDMA communication networks are such that cell phone signals can interfere by overlapping in the frequency spectrum causing a degradation of individual signal to noise ratios and hence the quality of service. In a basic version of the standard model there are two state variables for each agent: the transmitted power $p \in R^+$ and channel attenuation $\beta \in R$. Conventional power control algorithms in mobile devices use gradient-type algorithms with bounded step size for the transmitted power which may be represented by the so-called rate adjustment model: $dp^i = u_p^i dt + \sigma_p^i dW_p^i$, $u_p^i \leq |u_{max}|$, $1 \leq i \leq N$, where N represents the number of the users in the network, and W_p^i, $1 \leq i \leq N$, independent standard Wiener processes. Further, a standard model for time-varying channel attenuation is the lognormal model, where the channel gain for the ith agent with respect to the base station is given by $e^{\beta^i(t)}$ at the instant t, $0 \leq t \leq T$ and the received power at the base station from the $i$the agent is given by the product $e^{\beta^i} \times p^i$. The channel state, $\beta^i(t)$, evolves according to the power attenuation dynamics: $d\beta^i = -a^i(\beta^i + b^i)dt + \sigma_\beta^i \mathrm{d}W_\beta^i, t \geq 0, 1 \leq i \leq N$. For the generic agent $\mathscr{A}_i$ in the infinite user's case, the cost function $L_i(\beta_i, p_i)$ is given by

$$\begin{aligned}\lim_{N\to\infty} E\left[\int_0^T \left\{-\frac{e^{\beta^i} p^i}{\frac{1}{N}\sum_{j=1}^n p^j e^{\beta^j} + \eta} + p^i\right\} dt\right] \\ = \int_0^T \left\{-\frac{e^{\beta^i} p^i}{\int_{\Omega_\beta \times \Omega_p} e^{\beta} p \mu_t(\beta, p) d\beta dp + \eta} + p^i\right\} dt,\end{aligned}$$

where μ_t denotes the system mean field. As a result, the power control problem may be formulated as a dynamic game between the cellular users whereby each agent's cost function $L_i(\beta, p)$ involves both its individual transmitted power and its signal to noise ratio. The application of MFG theory yields a Nash equilibrium together with the control laws generated by the system's MFG equations (Aziz and Caines 2017). Due to the low dimension of the system state in this formulation, and indeed in that with mobile agents in a planar zone, the MFG PDEs can be solved efficiently.

7.3 Stochastic Growth Models

The model described below is a large population version of the so-called capital accumulation games (Amir 1996). Consider N agents (as economic entities). The capital stock of agent i is x_t^i and modeled by

$$dx_t^i = \left[A(x_t^i)^\alpha - \delta x_t^i\right]dt - C_t^i dt - \sigma x_t^i dw_t^i, \quad t \geq 0, \tag{25}$$

where $A > 0$, $0 < \alpha < 1$, $x_0^i > 0$, $\{w_t^i, 1 \leq i \leq N\}$ are i.i.d. standard Wiener processes. The function $F(x) := Ax^\alpha$ is the Cobb-Douglas production function with capital x and a constant labor size; $(\delta dt + \sigma dw_t^i)$ is the stochastic capital depreciation rate; and C_t is the consumption rate.

The utility functional of agent i takes the form:

$$J_i(C^1, \ldots, C^N) = E\left[\int_0^T e^{-\rho t} U(C_t^i, C_t^{(N,\gamma)})dt + e^{-\rho T} S(X_T)\right], \tag{26}$$

where $C_t^{(N,\gamma)} = \frac{1}{N}\sum_{i=1}^N (C_t^i)^\gamma$ is the population average utility from consumption. The motivation of taking the utility function $U(C_t^i, C_t^{(N,\gamma)})$ is based on relative performance. We take $\gamma \in (0, 1)$ and the utility function (Huang and Nguyen 2016):

$$U(C_t^i, C_t^{(N,\gamma)}) = \frac{1}{\gamma}(C_t^i)^{\gamma(1-\lambda)}\left(\frac{(C_t^i)^\gamma}{C_t^{(N,\gamma)}}\right)^\lambda, \quad \lambda \in [0, 1]. \tag{27}$$

So $U(C_t^i, C_t^{(N,\gamma)})$ may be viewed as a weighted geometric mean of the own utility $U_0 = (C_t^i)^\gamma/\gamma$ and the relative utility $U_1 = (C_t^i)^\gamma/(\gamma C_t^{(N,\gamma)})$. For a given θ, $U(c, \theta)$ is a hyperbolic absolute risk aversion (HARA) utility since $U(c, \theta) = \frac{c^\gamma}{\gamma\theta^\lambda}$, where $1 - \gamma$ is usually called the relative risk aversion coefficient. We further take $S(x) = \frac{\eta x^\gamma}{\gamma}$, where $\eta > 0$ is a constant.

Concerning growth theory in economics, human capital growth has been considered in an MFG setting by Lucas and Moll (2014) and Guéant et al. (2011) where the individuals invest resources (such as time and money) for the improvement of personal skills to better position themselves in the labor market when competing with each other (Guéant et al. 2011).

References

Adlakha S, Johari R, Weintraub GY (2015) Equilibria of dynamic games with many players: existence, approximation, and market structure. J Econ Theory 156:269–316

Ahuja S (2016) Wellposedness of mean gield games with common noise under a weak monotonicity condition. SIAM J Control Optim 54(1):30–48

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

Amir R (1996) Continuous stochastic games of capital accumulation with convex transitions. Games Econ Behav 15:111–131

Andersson D, Djehiche B (2011) A maximum principle for SDEs of mean-field type. Appl Math Optim 63(3):341–356

Aumann RJ, Shapley LS (1974) Values of non-atomic games. Princeton University Press, Princeton

Aziz M, Caines PE (2017) A mean field game computational methodology for decentralized cellular network optimization. IEEE Trans Control Syst Technol 25(2):563–576

Bardi M (2012) Explicit solutions of some linear-quadratic mean field games. Netw Heterog Media 7(2):243–261

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

Bauch CT, Earn DJD (2004) Vaccination and the theory of games. Proc Natl Acad Sci U.S.A. 101:13391–13394

Bauso D, Pesenti R, Tolotti M (2016) Opinion dynamics and stubbornness via multi-population mean-field games. J Optim Theory Appl 170(1):266–293

Bensoussan A, Frehse J (1984) Nonlinear elliptic systems in stochastic game theory. J Reine Angew Math 350:23–67

Bensoussan A, Frehse J, Yam P (2013) Mean field games and mean field type control theory. Springer, New York

Bensoussan A, Frehse J, and Yam SCP (2015) The master equation in mean field theory. J Math Pures Appl 103:1441–1474

Bergin J, Bernhardt D (1992) Anonymous sequential games with aggregate uncertainty. J Math Econ 21:543–562

Caines PE (2014) Mean field games. In: Samad T, Baillieul J (eds) Encyclopedia of systems and control. Springer, Berlin

Caines PE, Kizilkale AC (2017, in press) ε-Nash equilibria for partially observed LQG mean field games with a major player. IEEE Trans Autom Control

Cardaliaguet P (2012) Notes on mean field games. University of Paris, Dauphine

Cardaliaguet P, Delarue F, Lasry J-M, Lions P-L (2015, preprint) The master equation and the convergence problem in mean field games

Carmona R, Delarue F (2013) Probabilistic analysis of mean-field games. SIAM J Control Optim 51(4):2705–2734

Carmona R, Delarue F (2014) The master equation for large population equilibriums. In: Crisan D., Hambly B., Zariphopoulou T (eds) Stochastic analysis and applications. Springer proceedings in mathematics & statistics, vol 100. Springer, Cham

Carmona R, Delarue F, Lachapelle A (2013) Control of McKean-Vlasov dynamics versus mean field games. Math Fin Econ 7(2):131–166

Carmona R, Fouque J-P, Sun L-H (2015) Mean field games and systemic risk. Commun Math Sci 13(4):911–933

Carmona R, Lacker D (2015) A probabilistic weak formulation of mean field games and applications. Ann Appl Probab 25:1189–1231

Carmona R, Zhu X (2016) A probabilistic approach to mean field games with major and minor players. Ann Appl Probab 26(3):1535–1580

Chan P, Sircar R (2015) Bertrand and Cournot mean field games. Appl Math Optim 71:533–569

Correa JR, Stier-Moses NE (2010) Wardrop equilibria. In: Cochran JJ (ed) Wiley encyclopedia of operations research and management science. John Wiley & Sons, Inc, Hoboken

Djehiche B, Huang M (2016) A characterization of sub-game perfect equilibria for SDEs of mean field type. Dyn Games Appl 6(1):55–81

Dogbé C (2010) Modeling crowd dynamics by the mean-field limit approach. Math Comput Model 52(9–10):1506–1520

Fischer M (2014, preprint) On the connection between symmetric N-player games and mean field games. arXiv:1405.1345v1
Gangbo W, Swiech A (2015) Existence of a solution to an equation arising from mean field games. J Differ Equ 259(11):6573–6643
Gomes DA, Mohr J, Souza RR (2013) Continuous time finite state mean field games. Appl Math Optim 68(1):99–143
Gomes DA, Saude J (2014) Mean field games models – a brief survey. Dyn Games Appl 4(2):110–154
Gomes D, Velho RM, Wolfram M-T (2014) Socio-economic applications of finite state mean field games. Phil Trans R Soc A 372:20130405 http://dx.doi.org/10.1098/rsta.2013.0405
Guéant O, Lasry J-M, Lions P-L (2011) Mean field games and applications. In: Paris-Princeton lectures on mathematical finance. Springer, Heidelberg, pp 205–266
Haimanko O (2000) Nonsymmetric values of nonatomic and mixed games. Math Oper Res 25:591–605
Hart S (1973) Values of mixed games. Int J Game Theory 2(1):69–85
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 M (2010) Large-population LQG games involving a major player: the Nash certainty equivalence principle. SIAM J Control Optim 48:3318–3353
Huang M, Caines PE, Malhamé RP (2003) Individual and mass behaviour in large population stochastic wireless power control problems: centralized and Nash equilibrium solutions. In: Proceedings of the 42nd IEEE CDC, Maui, pp 98–103
Huang M, Caines PE, Malhamé RP (2007) Large-population cost-coupled LQG problems with non-uniform agents: individual-mass behavior and decentralized ε-Nash equilibria. IEEE Trans Autom Control 52:1560–1571
Huang M, Caines PE, Malhamé RP (2012) Social optima in mean field LQG control: centralized and decentralized strategies. IEEE Trans Autom Control 57(7):1736–1751
Huang M, Malhamé RP, Caines PE (2006) Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. Commun Inf Syst 6(3):221–251
Huang M, Nguyen SL (2016) Mean field games for stochastic growth with relative utility. Appl Math Optim 74:643–668
Jovanovic B, Rosenthal RW (1988) Anonymous sequential games. J Math Econ 17(1):77–87
Kizilkale AC, Malhamé RP (2016) Collective target tracking mean field control for Markovian jump-driven models of electric water heating loads. In: Vamvoudakis K, Sarangapani J (eds) Control of complex systems: theory and applications. Butterworth-Heinemann/Elsevier, Oxford, pp 559–589
Kolokoltsov VN, Bensoussan A (2016) Mean-field-game model for botnet defense in cyber-security. Appl Math Optim 74(3):669–692
Kolokoltsov VN, Li J, Yang W (2012, preprint) Mean field games and nonlinear Markov processes. Arxiv.org/abs/1112.3744v2
Kolokoltsov VN, Malafeyev OA (2017) Mean-field-game model of corruption. Dyn Games Appl 7(1):34–47. doi: 10.1007/s13235-015-0175-x
Lachapelle A, Wolfram M-T (2011) On a mean field game approach modeling congestion and aversion in pedestrian crowds. Transp Res B 45(10):1572–1589
Laguzet L, Turinici G (2015) Individual vaccination as Nash equilibrium in a SIR model with application to the 2009–2010 influenza A (H1N1) epidemic in France. Bull Math Biol 77(10):1955–1984
Lasry J-M, Lions P-L (2006a) Jeux à champ moyen. I – Le cas stationnaire. C Rendus Math 343(9):619–625
Lasry J-M, Lions P-L (2006b) Jeux à champ moyen. II Horizon fini et controle optimal. C Rendus Math 343(10):679–684
Lasry J-M, Lions P-L (2007). Mean field games. Japan J Math 2:229–260

Li T, Zhang J-F (2008) Asymptotically optimal decentralized control for large population stochastic multiagent systems. IEEE Trans Automat Control 53:1643–1660

Lucas Jr. RE, Moll B (2014) Knowledge growth and the allocation of time. J Political Econ 122(1):1–51

Ma Z, Callaway DS, Hiskens IA (2013) Decentralized charging control of large populations of plug-in electric vehicles. IEEE Trans Control Syst Technol 21(1):67–78

Milnor JW, Shapley LS (1978) Values of large games II: oceanic games. Math Oper Res 3:290–307

Neyman A (2002) Values of games with infinitely many players. In: Aumann RJ, Hart S (eds) Handbook of game theory, vol 3. Elsevier, Amsterdam, pp 2121–2167

Nourian M, Caines PE (2013) ε-Nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents. SIAM J Control Optim 51:3302–3331

Salhab R, Malhamé RP, Le Ny L (2015, preprint) A dynamic game model of collective choice in multi-agent systems. ArXiv:1506.09210

Sen N, Caines PE (2016) Mean field game theory with a partially observed major agent. SIAM J Control Optim 54:3174–3224

Tembine H, Zhu Q, Basar T (2014) Risk-sensitive mean-field games. IEEE Trans Autom Control 59:835–850

Wang BC, Zhang J-F (2012) Distributed control of multi-agent systems with random parameters and a major agent. Automatica 48(9):2093–2106

Wardrop JG (1952) Some theoretical aspects of road traffic research. Proc Inst Civ Eng Part II, 1:325–378

Weintraub GY, Benkard C, Van Roy B (2005) Oblivious equilibrium: a mean field approximation for large-scale dynamic games. Advances in neural information processing systems, MIT Press, Cambridge

Weintraub GY, Benkard CL, Van Roy B (2008) Markov perfect industry dynamics with many firms. Econometrica 76(6):1375–1411

Yin H, Mehta PG, Meyn SP, Shanbhag UV (2012) Synchronization of coupled oscillators is a game. IEEE Trans Autom Control 57:920–935