Mean Field Games: An Introduction

Xin Guo, UC Berkeley

ESSFM20, September 1, 2020

Overview

- Basics of MFG
- Relation between N-player game and MFG
- Computation and Learning for MFG

Part I: Basics of MFG

Outline

MFG: (comprehensive treatment see Carmona and Delarue (2018) Vol I and II, Bensoussan, J. Frehse, and P. Yam (2013)

- the idea of “mean field” originated from statistical mechanics on weakly interacting particles
- theoretical works pioneered by Lasry and Lions (2007) and Huang, Malhamé and Caines (2006)
- stochastic strategic decision games with very large population of small interacting individuals
- about small interacting individuals, with each player choosing optimal strategy in view of the macroscopic information (mean field)

Main idea of MFG

- Take an N-player game
- When N is large, consider instead the "aggregated" version of the N-player game
- By (f)SLLN, the aggregated version, MFG, becomes an "approximation" of the N-player game, in terms of ϵ-Nash equilibrium

N-player game

$$\inf_{\alpha^i \in \mathcal{A}} E\{\int_0^T f^i(t, X_t^1, \cdots, X_t^N, \alpha_t^i)dt + g^i(T, X_T^1, \cdots, X_T^N, \alpha_T^i)\}$$

$$\textit{subject to} \quad dX_t^i = b^i(t, X_t^1, \cdots, X_t^N, \alpha_t^i)dt + \sigma dW_t^i$$

$$\text{and} \quad X_0^i = x^i$$

- X_t^i the state of player i at time t
- α_t^i the action/control of player i at time t, in an appropriate control set $\mathcal{A}$
- f^i the running cost for player i
- g^i the terminal cost for player i
- b^i the drift term for player i
- σ a volatility term for player i
- W_t^i i.i.d. standard Brownian motions

From N-player game to MFG

Aggregation

$$\inf_{\alpha^i \in \mathcal{A}} E\left\{ \int_0^T \frac{1}{N} \sum_{i=1}^N f^i(t, X_t^1, \cdots, X_t^N, \alpha_t^i) dt \right.$$

$$\left. + \frac{1}{N} \sum_{i=1}^N g^i(T, X_T^1, \cdots, X_T^N, \alpha_T^i) \right\}$$

$$s.t. \quad dX_t^i = \frac{1}{N} \sum_{i=1}^N b^i(t, X_t^1, \cdots, X_t^N, \alpha_t^i) dt + \sigma dW_t^i$$

$$\text{and} \quad X_0^i = x^i$$

As $N \to \infty$,

MFG

$$\inf_{\alpha \in \mathcal{A}} E\{\int_0^T f(t, X_t^i, \mu_t, \alpha_t)dt + g(T, X_T, \mu_T, \alpha_T)\}$$

$$\text{such that} \quad dX_t^i = b(t, X_t^i, \mu_t, \alpha_t)dt + \sigma dW_t^i \quad \text{and} \quad X_0^i = x^i$$

with $\mu_t = \lim_{N\to\infty} \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i}$

Assumptions (symmetry)

Players are indistinguisheable: they are rational, identical, and interchangeable

MFG

- revolutionizes the analysis of game theory
- provides simple approximation method for N-player game
- leads new theoretical development in stochastic analysis
  - Itô lemma for flows of probability measures
  - master equation for characterization of MFG

Toy model (Guéant, Lasry, and Lions (2011))

Question: deciding the starting time of a meeting

- identical $N = 10K$ agents distributed on the negative half-line according to the distribution m_0 ($m_0(x) = 0$ for $x \geq 0$)
- a meeting scheduled to hold at $x = 0$ at time t_0
- actual meeting starting time T depending on the arrivals of participants (e.g., meeting starts when 90% agents arrive at 0)

Toy model for MFG

Assume

- a continuum of agents, as N is large
- agents are rational and interchangable
- X_t^i position of representative agent i at time t
- α_t^i agent i's action at time t
- $X_t^i = \alpha_t^i t + \sigma W_t^i$, with W_t^i i. i. d. standard Brownian motions
- τ_i the time at which agent i would like to arrive
- real arrival time is $\tilde{\tau}^i = \tau^i + \sigma \epsilon^i$ where $\epsilon^i \sim N(0,1)$ i. i. d. with

$$\tilde{\tau}^i = \inf\{t > 0, X_t^i = 0\}$$

Agent i's objective

Minimize the expected cost functional $c^i(t_0, T, \tilde{\tau}^i, \alpha^i)$,

- quadratic running cost with respect to α^i
- terminal cost including
  - $(T - \tilde{\tau}^i)^+$: inconvenience due to waiting time
  - $(\tilde{\tau}^i - t_0)^+$: lateness w.r.t. to scheduled time
  - $(\tilde{\tau}^i - T)^+$: lateness w.r.t. to actual time

Solution approach

- First, Fix T

$$\min_{\alpha} E[c^i(t_0, T, \tilde{\tau}^i, \alpha^i)]$$
$$s.t. \quad X_t^i = \alpha_t^i t + \sigma W_t^i, \quad X_0^i = x_0$$
$$\tilde{\tau}^i = \min\{s : X_s^i = 0\};$$

with α^* the optimal control of this problem

- Based on α^*, get distribution of arrivals μ (all agents are identical). From μ, update T'. Repeat this process until a fixed point.

$$T \to \alpha^* \to T' \to \alpha^{*\prime} \to \cdots \to \text{fixed point}$$

Key insight from the toy example

- game interaction depends on the distribution of individual state
- aggregation simplifies computation
- solution approach is by iteration via Banach fixed point theorem

An MFG on $\mathbb{R}^d$ over a time horizon $[0, T]$,

MFG

$$\inf_{\alpha \in \mathcal{A}} \mathbb{E}\left[\int_0^T f(t, X_t, \mu_t, \alpha_t)dt + g(T, X_T, \mu_T, \alpha_T)\right]$$

subject to

$$dX_t = b(t, X_t, \mu_t, \alpha_t)dt + \sigma(t, X_t, \mu_t, \alpha_t)dW_t$$

$$X_0 \sim \mu^0, \quad \mu_t = Law(X_t) \quad \forall t \in [0, T].$$

Definition: solution to MFG

If there exists a control policy $\{\alpha_t^*\}_t$ and a flow of probability measures $\{\mu_t^*\}_t$ such that

- under $\{\mu_t^*\}_t$, $\{\alpha_t^*\}_t$ solves the optimal control problem

$$\inf_{\alpha \in \mathcal{A}} \mathbb{E}\left[\int_0^T f(t, X_t, \mu_t^*, \alpha_t)dt + g(T, X_T, \mu_T^*, \alpha_T)\right]$$

  subjecto to

$$dX_t = b(t, X_t, \mu_t^*, \alpha_t)dt + \sigma(t, X_t, \mu_t^*, \alpha_t)dW_t, \quad X_0 \sim \mu^0;$$

- under $\{\alpha_t^*\}_t$, the controlled process $\{X_t^*\}_t$ given by

$$dX_t^* = b(t, X_t^*, \mu_t^*, \alpha_t^*)dt + \sigma(t, X_t^*, \mu_t^*, \alpha_t^*)dW_t, \quad X_0^* \sim \mu^0$$

  satisfies $\mu_t^* = Law(X_t^*)$ for all t;

then the pair $(\{\alpha_t^*\}_t, \{\mu_t^*\}_t)$ is called a solution to the MFG.

PDE/control approach for MFG

- (i) Fix a deterministic function $t \in [0, T] \to \mu_t \in \mathcal{P}(\mathbb{R}^d)$
- (ii) Solve the stochastic control problem

$$\inf_{\alpha \in \mathcal{A}} \int_0^T f(t, X_t, \mu_t, \alpha_t) dt$$
$$s.t. \quad dX_t = b(t, X_t, \mu_t, \alpha_t) dt + \sigma dW_t \quad \text{and} \quad X_0 = x$$

- (iii) Update the function $t \in [0, T] \to \mu'_t \in \mathcal{P}(\mathbb{R}^d)$ so that $\mathcal{P}_{X_t} = \mu'_t$
- (iv) Repeat (ii) and (iii). If there exists a fixed point solution μ_t and α_t, then it is a solution for the MFG

First step: HJB equation

Solving the control problem yields

$$\partial_s v(s,x) + H(s,x,\nabla_x v(s,x), \nabla_x^2 v(s,x)) = 0, \qquad \text{(HJB)}$$

with terminal condition $v(T,x) = g(x,\mu_T)$, and the Hamiltonian

$$H(s,x,y,z) = \inf_{\alpha} \left\{ b(s,x,\mu_s,\alpha) \cdot y + \frac{1}{2} Tr(\sigma\sigma^T(s,x,\mu_s,\alpha)z) + f(s,x,\mu_s,\alpha) \right\}$$

Second step: FP equation

Updating $\{\mu_t\}_t$ for the controlled dynamic $\{X_t\}_t$ under the optimal policy $\{\alpha_t^*\}_t$, the density function $m(s,x)$ of $\mu_s = Law(X_s)$ satisfies

$$\begin{aligned}
&\partial_s m(s,x) + div_x\left(b(s,x,\mu_s,\alpha_s^*)m(s,x)\right) \\
&\quad - \frac{1}{2} Tr\left[\nabla_x^2\left(\sigma\sigma^T(s,x,\mu_s,\alpha)m(s,x)\right)\right] = 0, \\
&\int_x m(s,x)dx = 1,\ \forall s; \quad m(0,\cdot) = m^0(\cdot).
\end{aligned} \tag{FP}$$

Third step: iteration

- The coupled HJB-FP system characterizes a mapping Γ

$$\Gamma : \{\mu_t\}_t \xmapsto{\text{HJB}} \{\alpha_t^*\}_t \xmapsto{\text{FP}} \text{updated } \{\mu_t\}_t.$$

- If Γ admits a fixed point $\{\mu_t^*\}_t$, then $(\{\mu_t^*\}_t, \{\alpha_t^*\}_t)$ is a solution to to the (MFG)

Fixed point solution

Schaefer's Fixed Point Theorem

Given a Banach space X. Suppose a mapping $A : X \to X$ is continuous and compact. If the set

$$\{u \in X | u = \lambda A(u) \text{ for some } \lambda \in [0,1]\}$$

is bounded, then A admits a fixed point.

Note: Banach fixed point theorem leads to a unique solution for a contraction mapping.

Existence of MFG solution

Lasry and Lions (2007) considers the class of MFG

- $b(s, x, \mu, \alpha) = -\alpha$
- $\sigma(s, x, \mu, \alpha) \equiv \sigma$ for some given σ;
- $f(s, x, \mu_s, \alpha) = L(x, \alpha) + \tilde{f}(s, x, m(s, x))$; and
- $g(x, \mu_T) = \tilde{g}(x, m(T, x))$.

Existence of MFG solution (Lasry, Lions (2007))

Suppose

- if the density function $m(s,x)$ belongs to $\mathcal{C}([0,T],\mathcal{P}^1(\mathbb{R}^d))$, then $\tilde{f}(s,x,m(s,x))$ and $\tilde{g}(x,m(T,x))$ are both (uniformly) bounded and continuous;
- $\tilde{f}$ is continuous with respect to $m \in \mathcal{C}([0,T],\mathcal{P}^1(\mathbb{R}^d))$;
- if $m \in \mathcal{C}^{k,\alpha}$, then $\tilde{f},\tilde{g} \in \mathcal{C}^{k+1,\alpha}$;
- $\bar{H}(x,y)$ is smooth on $\mathbb{R}^d \times \mathbb{R}^d$ and there exists a constant $c > 0$ such that for any (x,y),

$$|\partial_y \bar{H}(x,y)| \leq C(1+|y|), \quad |\partial_x \bar{H}(x,y)| \leq c(1+|y|).$$

Then the HJB-FP system admits a pair of solutions $(v^*(s,x), m^*(s,x))$, and the corresponding control-mean pair $(\{\alpha_t^*\}_t, \{\mu_t^*\}_t)$ is a solution of the MFG.

Notation

- $\mathcal{C}([0,T],\mathcal{P}^1(\mathbb{R}^d))$: set of continuous mappings from $[0,T]$ to the set of probability measures $\mathcal{P}^1(\mathbb{R}^d)$ where $\int_{\mathbb{R}^d}|x|\mu(dx)<\infty$ for any $\mu\in\mathcal{P}^1(\mathbb{R}^d)$.
- $\mathcal{C}^{k,\alpha}$: set of $\mathcal{C}^k$ functions with bounded and Hölder continuous of exponent α derivatives up to order k, for any integer $k\geq 0$ and $\alpha\in(0,1)$.

Several approaches for MFGs with regular controls

- PDE/control approach:
  backward HJB equation + forward Kolmogorov equation
  Lions and Lasry (2007), Huang, Malhame and Caines (2006), Lions, Lasry and Guánt (2009),
- Probabilistic approach: FBSDEs
  Buckdahn, Li and Peng (2009), Carmona and Delarue (2013)
- Master equation (and verification argument)
  Cardaliaguet, Delarue, Lasry, and Lions (2019)

Existence of MFGs

More general existence results of Lacker (2015)

- controls are not necessarily absolutely continuous
- adopting the notion of relaxed control, with measure selection argument to get strict Markovian MFG solution via convexity condition
- Kakutani-Fan-Glicksberg fixed point theorem on set-valued mapping

Note: this argument could be adapted for MFGs with singular controls.

Singular controls

- more general mathematical framework as controls could be discontinuous
- adding gradient constraint to the HJB
- controlled processes give rise to Skorokhod problems
- stochastic maximum principle fails, and Hamiltonian diverges

Classical fuel follower problem

Single player, dynamics of its position given by

$$X_t = x + W_t + \xi_t,$$

- W_t standard Browian motion
- ξ_t controlled càdlàg process, with finite variation such that

$$\xi_t = \xi_t^+ - \xi_t^-$$

  with ξ_t^+ and ξ_t^- non-decreasing càdlàg processes, $\xi_0^- = \xi_0^+ = 0$
- $\check{\xi}_t = \xi_t^+ + \xi_t^-$ accumulative fuel consumption

The objective is to minimize over all admissible controls $\mathcal{A}$

$$E\int_0^\infty e^{-\alpha t}\{d\check{\xi}_t + h(X_t)dt\}$$

- h convex, $h(-x) = h(x)$, $h''(x)$ decreasing and $0 < k < h''(x) < K$
- admissible control set

$$\begin{aligned}\mathcal{A} \quad &= \{(\xi_t^+, \xi_t^-)|\xi_t^+, \xi_t^-, \ \mathcal{F}^{W_t}\text{-progressively measurable,} \\ &\text{càdlàg non-decreasing, } \mathbb{E}[\int_0^\infty e^{-\alpha t} d\xi_t^+] < \infty, \\ &\mathbb{E}[\int_0^\infty e^{-\alpha t} d\xi_t^-] < \infty, \xi_{0-}^+ = \xi_{0-}^- = 0\}\end{aligned}$$

Solution

- solving the HJB via smooth-fit-principle

$$\min\{\frac{1}{2}v''(x) + h(x) - \alpha v(x), 1 - v'(x), 1 + v'(x)\} = 0$$

- "bang-bang" type optimal control, with $c > 0$ being the unique solution to

$$\frac{1}{\sqrt{2\alpha}} tanh(c\sqrt{2\alpha}) = \frac{p_1^{'}(c) - 1}{p_1^{''}(c)}.$$

with

$$p_1(x) = E[\int_0^{\infty} e^{-\alpha t} h(x + B_t) dt].$$

Benes, Shepp and Witsenhausen (1980) and Karatzas (1983)

Existence of solutions to MFG with singular controls

- For singular control of bounded velocity
  Lacker (2018)
- For singular control of finite variation
  Fu and Horst (2017), Cao, G., and Lee (2019)

Example: cash management

- deterministic model
  Baumol (1952), Miller and Orr (1966)
- stochastic model with fixed and proportional cost
  Eppen and Fama (1968),
  Constantinides and Richard (1978)

Impulse controls

- Quasi-Variational Inequalities:
  Caffarelli and Friedman(1978, 1979), Bensoussan and Lions (1982)
- inventory controls:
  Harrison and Taylor (1978), Harrison, Sellke, and Taylor (1983), Sulem (1986)
- portfolio management with transaction cost:
  Davis and Norman ('90), Korn (1998, 1999), Øksendal and Sulem (2002)
- exchange rates:
  Jeanblanc-Piqué (1993), Cadenillas and Zapatero (1999),
- Insurance models: Candenillas, Choulli, Taksar and Zhang (2006)
- liquidity risk: Ly Vath et al. (2007)
- irreversible investment: Scheinkman and Zariphopoulou (2001)

Cash management

- cash balance as inventory for exchange
- withdrawal from cash balance at the end of each period
- fixed cost for broker's fee and proportional holding cost

Constantinides and Richard (1978)

$X_t \in \mathbb{R}^n$: cash balance at time t

- without intervention,

$$X(t) = \mu t + \sigma W(t), \quad X(0) = x,$$

where $\{W_t\}_{t \geq 0}$ a standard Brownian motion

- with intervention $\varphi = \{\tau_n, \xi_n\}_{n\geq 0} \in \mathcal{A}$

$$dX_t = \mu dt + \sigma dW_t + \sum_{n=1}^{\infty} \delta(t-\tau_n)\xi_n, \quad X_{0-} = x$$

  - $\{\tau_n\}_{n\geq 1}$ a sequence of stopping times w.r.t. $\{\mathcal{F}_t^W\}$ s.t. $\tau_n \uparrow \infty$ a.s.
  - $\xi_n \in \mathcal{F}_{\tau_n}^W$
  - $\mathcal{A}$ the set of all such controls

Objective Function

$$v(x) = \inf_{\varphi \in \mathcal{A}} J(x, \varphi) = \inf_{\varphi \in \mathcal{A}} E\left[\int_0^\infty e^{-rt} C(X(t))dt + \sum_{n=1}^{\infty} e^{-r\tau_n}\phi(\xi_n)\right]$$

- running cost $C(x) = \max\{hx, -px\}$ where $h, p > 0$
- cost of control

$$\phi(x) = \begin{cases} K^+ + k^+x, & x \geq 0 \\ K^- - k^-x, & x < 0 \end{cases} \quad K^+, K^-, k^+, k^- > 0$$

- zero control: $\phi(0) = K^+ > 0$
- discount rate $r > 0$

Quasi-Variational Inequality (QVI)[1]

$$\min\left\{\frac{\sigma^2}{2}v'' - \mu v' - rv + C, \inf_{\xi}\phi(\xi) + v(x+\xi)\right\} = 0$$

- consider $X(t-) = x$, time interval $[t, t+\Delta t]$
- finite action space: either no control or control at time t
- no control:
  $v(X(t-)) \le E\left[\int_t^{t+\Delta t} e^{-r(s-t)}C(X(s))ds\right] + e^{-rdt}v(X(t-)+\Delta X)$
- control at time t: $v(X(t-)) \le \inf_{\xi}\phi(\xi) + v(X(t-)+\xi)$
- sending Δt to 0 and applying Itô's formula

[1] Constantinides and Richard (1978)

Contantinides and Richard (1978)

The (S, s) optimal strategy

- if $X_t \in \mathcal{C} = (q, s)$: no action;
- X_t reaches q (or initial $x \leq q$): raise it to Q immediately
- X_t reaches s (or initial $x \geq s$): lower it to S immediately

Impulse control vs regular control

- (S, s) policy and K-convexity
  Constantinides and Richard (1978)
- QVI vs classical HJB equation: additional *nonlocal* operators
- classical references: Benssousan and Lions (1982)
- how to adapt classical PDEs approach for regularity studies?
  G. and Wu (2010), Davis, G. and Wu (2012)

PDE tools for regularity

- Sobolev imbedding: $W^{2,p} \subset C^{1,\alpha}, \alpha = 1 - n/p$
- Schauder's estimates: $Lu = f \in C^{\alpha} \Rightarrow u \in C^{2,\alpha}$
- Calderon-Zygmund estimates: $Lu = f \in L^{p} \Rightarrow u \in W^{2,p}$

K-convexity and regularity of PDEs

- the nonlocal operator $\mathcal{M}u(x) := \inf_\xi \phi(\xi) + u(x+\xi)$
- $\mathcal{D}(x) := \{\xi \in \mathbb{R}^n : \mathcal{M}v(x) = v(x+\xi) + B(\xi)\}$
- K-convexity is critical for the regularity of the value function

Key property

- For any point $x \in \mathcal{A}$ and $\xi(x) \in \mathcal{D}(x)$, then $x + \xi(x) \in \mathcal{C}$
- $\mathcal{M}u(x) \geq \mathcal{M}(x + \xi(x)) + B(\xi(x)) - K$

This ensures

- a bound for the second difference quotient at x, and
- an appropriate estimate to apply Calderon-Zygmund estimate,
- and eventual $W^{2,p}$ regularity for the value function

Existence of solutions to MFGs with impulse controls

Mostly open

- QVI form for "FP" equation
  Bertucci (2020)
- culprit: the nonlocal operator in QVI
- only known for special cases
  Basei, Cao, G. (2020) on cash management problem

Coming up next: N vs MFG

- MFG as an ϵ-NE approximation to N-player games
- $N \neq$ MFG
  comparisons through several examples
- master equations for convergence of N to MFG, CLT, and large deviations
  Cardaliaguet, Delarue, Lasry, and Lions (2015, 2019), Delarue, Lacker, and Ramanan (2019), Bayraktar and Cohen (2018), Cecchin and Pelino (2018)