Numerical Methods for Mean Field Games

Lecture 1
Introduction to MFGs: Definitions and Equilibrium Conditions

Mathieu Laurière

New York University Shanghai

UM6P Vanguard Center, Université Cadi AYYAD,
University Côte d'Azur, & GE2MI
Open Doctoral Lectures
July 5 – 7, 2023

About this course

- 6 lectures, 15 hours: $3 \times (3+2)$
- Objectives
  1. Introducing Mean Field Games
  2. Presenting the main ideas behind several numerical methods
  3. Providing sample codes to experiment with
- Interdisciplinary topic
- Feel free to ask questions
- The slides are available on my webpage: `https://mlauriere.github.io/#teaching`
- Based on the work of many contributors
- Feel free to reach out: `mathieu.lauriere@nyu.edu`

Outline

Many agent systems

Systems with many agents are ubiquitous in today's interconnected world

Flocking

Crowd motion

Traffic flow

Collective AI

[Image credits: Unsplash, Wikimedia Commons (Kilobots)]

More examples

Economics & finance, energy management, telecommunications, networks, ...

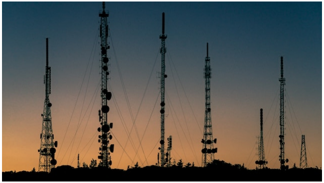

Groups of animals

- Flocking, schooling, herding, . . . have been extensively studied
- Predator-prey models
- Ex.: Cucker-Smale model of flocking, Lotka-Volterra system
- . . .

Groups of human beings

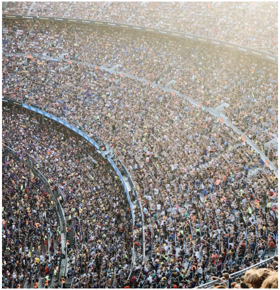

Some existing approaches *("What MFGs are* ***not****")*

- Dynamical systems:
  - describe the dynamics of one or many agents, sometimes mean field
  - but usually **no rationality** (optimization)
- Agent based models (ABM):
  - "Agent-based models are a kind of **microscale model** that simulate the simultaneous operations and interactions of multiple agents in an attempt to re-create and predict the appearance of complex phenomena. "
  - "Individual agents are typically characterized as **boundedly rational**, presumed to be acting in what they perceive as their own interests, such as reproduction, economic benefit, or social status, using heuristics or simple decision-making rules." (Wikipedia)
- Game theory
  - optimization aspects
  - notion of Nash equilibrium, social optimum, . . .
  - but usually limited to a **finite (small) number of agents**
- Evolutionary game theory (EGT)
  - "application of game theory to evolving populations in biology"
  - "an evolutionary version of game theory **does not require players to act rationally** – only that they have a strategy" (Wikipedia)
- Non-atomic anonymous games
  - continuum of rational players; each player has her **own index** and own strategy
  - mostly limited to static games; difficulties for dynamic, stochastic games

MFG paradigm in a nutshell

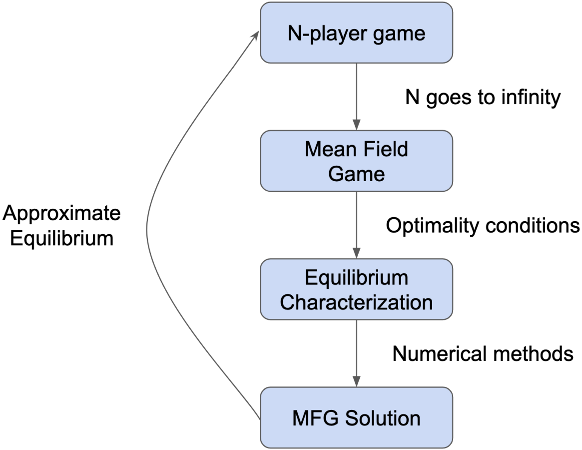

Goal for this lecture: discuss the other aspects and motivate numerical methods

Following lectures: focus on numerical methods

Outline of this course

- Lecture 1: Introduction
  - MFG models in the static setting
  - Dynamic setting and optimality conditions
- Lectures 2 & 3: “Classical” numerical methods (Parts I & II)
- Lectures 4 & 5: Deep learning numerical methods (Parts I & II)
- Lecture 6: Reinforcement learning methods

Some References

- Introduction to Mean Field Games:
  - Pierre-Louis Lions' lectures at Collège de France (`https://www.college-de-france.fr/`)
  - Pierre Cardaliaguet's notes (2013): `https://www.ceremade.dauphine.fr/ cardaliaguet/MFG20130420.pdf`
  - Gomes, D. A., & Saúde, J. (2014). Mean field games models—a brief survey. Dynamic Games and Applications, 4, 110-154.
  - Cardaliaguet, P., & Porretta, A. (2020). An Introduction to Mean Field Game Theory. In *Mean Field Games* (pp. 1-158). Springer, Cham.
  - Carmona, Delarue, Graves, Lacker, Laurière, Malhamé & Ramanan: Lecture notes of the 2020 AMS Short Course on Mean Field Games (American Mathematical Society), organized by François Delarue
  - Achdou, Y., Cardaliaguet, P., Delarue, F., Porretta, A., & Santambrogio, F. (2021). Mean Field Games: Cetraro, Italy 2019 (Vol. 2281). Springer Nature.
  - Delarue, F. (Ed.). (2021). Mean Field Games (Vol. 78). American Mathematical Society.
- Monographs on Mean Field Games and Mean Field Control:
  - Bensoussan, A., Frehse, J., & Yam, P. (2013). *Mean field games and mean field type control theory* (Vol. 101). New York: Springer.
  - Gomes, D. A., Pimentel, E. A., & Voskanyan, V. (2016). *Regularity theory for mean-field game systems. New York: Springer.*
  - Carmona, R., & Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games* (Vol. 83). Springer.
  - Carmona, R., & Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations* (Vol. 84). Springer.
- Surveys about numerical methods for MFGs:
  - Achdou, Y. (2013). Finite difference methods for mean field games. In *Hamilton-Jacobi equations: approximations, numerical analysis and applications* (pp. 1-47). Springer, Berlin, Heidelberg.
  - Achdou, Y., & Laurière, M. (2020). Mean Field Games and Applications: Numerical Aspects. *Mean Field Games: Cetraro, Italy 2019, 2281*, 249.
  - Laurière, M. (2021). Numerical Methods for Mean Field Games and Mean Field Type Control. Lecture notes for the AMS'20 short course. arXiv preprint arXiv:2106.06231.
  - Carmona, R., & Laurière, M. (2021). Deep Learning for Mean Field Games and Mean Field Control with Applications to Finance. arXiv preprint arXiv:2107.04568.
  - Hu, R., & Laurière, M. (2023). Recent developments in machine learning methods for stochastic control and games. arXiv preprint arXiv:2303.10257.
  - Laurière, M., Perrin, S., Geist, M., & Pietquin, O. (2022). Learning mean field games: A survey. arXiv preprint arXiv:2205.12944.

Outline

Outline

Notations

- N players $[N] = \{1, \ldots, N\}$
- action space $\mathcal{A}$ (finite for simplcity)
- each player $i \in [N]$ selects an action $a^i \in \mathcal{A}$
- it induces a population profile of actions $\underline{a} = (a^1, \ldots, a^N) \in \mathcal{A}^N$
- each player pays a cost $f^i(\underline{a})$, where $f^i : \mathcal{A}^N \to \mathbb{R}$
- goal of each player: minimize her own cost $\min_{a^i} f^i(\underline{a})$

Notations

Question: Is there a “stable configuration” in which all the players are “satisfied”?

Intuition: Strategy profile such that no player is interested in deviating by herself

Definition (Nash equilibrium (NE))

Notations

Question: Is there a “stable configuration” in which all the players are “satisfied”?

Intuition: Strategy profile such that no player is interested in deviating by herself

Definition (Nash equilibrium (NE))

$\underline{\hat{a}} = (\hat{a}^1, \ldots, \hat{a}^N) \in \mathcal{A}^N$ is a Nash equilibrium if: for every $i \in [N]$, for every $a^i \in \mathcal{A}$

$$f^i(\underline{\hat{a}}) \leq f^i(\hat{a}^1, \ldots, \hat{a}^{i-1}, a^i, \hat{a}^{i+1}, \ldots, \hat{a}^N)$$

Notations

Question: Is there a “stable configuration” in which all the players are “satisfied”?

Intuition: Strategy profile such that no player is interested in deviating by herself

Definition (Nash equilibrium (NE))

$\underline{\hat{a}} = (\hat{a}^1, \ldots, \hat{a}^N) \in \mathcal{A}^N$ is a Nash equilibrium if: for every $i \in [N]$, for every $a^i \in \mathcal{A}$

$$f^i(\underline{\hat{a}}) \leq f^i(\hat{a}^1, \ldots, \hat{a}^{i-1}, a^i, \hat{a}^{i+1}, \ldots, \hat{a}^N)$$

Convenient notation: $(a^i, \underline{\hat{a}}^{-i}) = (\hat{a}^1, \ldots, \hat{a}^{i-1}, a^i, \hat{a}^{i+1}, \ldots, \hat{a}^N)$.

The above condition rewrites: $f^i(\hat{a}^i, \underline{\hat{a}}^{-i}) \leq f^i(a^i, \underline{\hat{a}}^{-i})$

Example: Population distribution

Example (Target position; no interactions)

Cost:

$$f^i(\underline{a}) = -|a^i - a_{target}|^2$$

Nash equilibrium:

Example: Population distribution

Example (Attraction to the group; interaction through the mean)

Cost:

$$f^i(\underline{a}) = |a^i - \frac{1}{N}\sum_{j=1}^{N} a^j|$$

Nash equilibrium:

Example: Population distribution

Example (Group aversion)

For simplicity, assume $\mathcal{A} = \{1, 2, \ldots, d\}$ is a finite set.

Assume there are $N = k \times d$ players for some integer k.

Cost:

$$f^i(\underline{a}) = \sum_{j=1}^{N} 1_{\{a^j = a^i\}}$$

which is the number of players who choose the same action.

Nash equilibrium?

Example: Population distribution

Example (Group aversion)

For simplicity, assume $\mathcal{A} = \{1, 2, \ldots, d\}$ is a finite set.

Assume there are $N = k \times d$ players for some integer k.

Cost:

$$f^i(\underline{a}) = \sum_{j=1}^{N} 1_{\{a^j = a^i\}}$$

which is the number of players who choose the same action.

Nash equilibrium?

$a^i = \lfloor i/k \rfloor + 1$ form a Nash equilibrium with uniform distribution over actions.

Example (Group aversion)

For simplicity, assume $\mathcal{A} = \{1, 2, \ldots, d\}$ is a finite set.

Assume there are $N = k \times d$ players for some integer k.

Cost:

$$f^i(\underline{a}) = \sum_{j=1}^{N} 1_{\{a^j = a^i\}}$$

which is the number of players who choose the same action.

Nash equilibrium?

$a^i = \lfloor i/k \rfloor + 1$ form a Nash equilibrium with uniform distribution over actions.

Remark: Player i does not know the other players' actions before choosing her actions

$\Rightarrow$ Need to anticipate

Example: Population distribution

Exercise

Consider the following example with spatial preferences + group aversion.

Cost:

$$f^i(\underline{a}) = -|a^i - a_{target}| + \sum_{j=1}^{N} 1_{\{a^j = a^i\}}$$

Nash equilibrium?

Example (Example without NE)

Rock-Paper-Scissor game.

Number of player: $N = 2$.

Action set: $\mathcal{A} = \{R, P, S\}$.

Cost:

$$f^i(\underline{a}) = \begin{cases} 1 & \text{if } (a^i, a^{-i}) \in \{(P,R),(R,S),(S,P)\} \\ 0 & \text{if } a^i = a^{-i} \\ -1 & \text{otherwise} \end{cases}$$

No Nash equilibrium (in the sense defined previously).

Nash Theorem

Warning: a pure NE does not always exist

Nash theorem: existence of mixed NE

THEOREM 1. *Every finite game has an equilibrium point.*

Von Neumann and Morgenstern have developed a very fruitful theory of two-person zero-sum games in their book *Theory of Games and Economic Behavior*. This book also contains a theory of n-person games of a type which we would call cooperative. This theory is based on an analysis of the interrelationships of the various coalitions which can be formed by the players of the game.

Our theory, in contradistinction, is based on the *absence* of coalitions in that it is assumed that each participant acts independently, without collaboration or communication with any of the others.

The notion of an *equilibrium point* is the basic ingredient in our theory. This notion yields a generalization of the concept of the solution of a two-person zero-sum game. It turns out that the set of equilibrium points of a two-person zero-sum game is simply the set of all pairs of opposing "good strategies."

In the immediately following sections we shall define equilibrium points and prove that a finite non-cooperative game always has at least one equilibrium point.

Source: [Nash, 1951]

Mixed strategies

- N players $[N] = \{1, \dots, N\}$
- action space $\mathcal{A}$ (finite for simplicity)
- each player chooses a mixed strategy $\pi^i \in \mathcal{P}(\mathcal{A}) =$ probability measures on $\mathcal{A}$
- each player (independently) picks an action according to her strategy $a^i \sim \pi^i$
- it induces population profiles of strategies $\underline{\pi}$ and actions $\underline{a}$
- each player pays a cost $f^i(\underline{a})$
- goal for each player: minimize her own *expected* cost

$$J^i(\underline{\pi}) = \mathbb{E}_{a^j \sim \pi^j, j=1,\dots,N}[f^i(\underline{a})]$$

Mixed strategies

- N players $[N] = \{1, \ldots, N\}$
- action space $\mathcal{A}$ (finite for simplicity)
- each player chooses a mixed strategy $\pi^i \in \mathcal{P}(\mathcal{A})$ = probability measures on $\mathcal{A}$
- each player (independently) picks an action according to her strategy $a^i \sim \pi^i$
- it induces population profiles of strategies $\underline{\pi}$ and actions $\underline{a}$
- each player pays a cost $f^i(\underline{a})$
- goal for each player: minimize her own *expected* cost

$$J^i(\underline{\pi}) = \mathbb{E}_{a^j \sim \pi^j, j=1,\ldots,N}[f^i(\underline{a})]$$

- Remark: the distribution $\frac{1}{N}\sum_j \delta_{a_j}$ is *random*. But less and less as $N \to +\infty$.

Mixed Nash equilibrium

Definition (Mixed Nash equilibrium in N-player game)

$\underline{\hat{\pi}} \in \mathcal{P}(\mathcal{A})^N$ is a mixed Nash equilibrium for the N-player game if:

$$J^i(\hat{\pi}^i, \underline{\hat{\pi}}^{-i}) \leq J^i(\pi^i, \underline{\hat{\pi}}^{-i}), \qquad \forall i, \forall \pi^i$$

Exercise

Revisit the examples which had a solution.

For each example, compute the mixed Nash equilibria.

Question: What happens to the Rock-Paper-Scissor example without a solution?

Answer: $(\pi^1, \pi^2) = (\hat{\pi}, \hat{\pi})$, with $\hat{\pi} = (1/3, 1/3, 1/3)$ is a Nash equilibrium.

Large populations

Question: *What if N is very large?*

- In many applications, the number of players is extremely large
- Intuitively,
  - each player has a **negligible impact** on the rest of the population
  - the population distribution of actions becomes **deterministic**
- This should simplify the analysis
- Can we formalize this intuition?

Large populations

Question: *What if N is very large?*

- In many applications, the number of players is extremely large
- Intuitively,
  - each player has a **negligible impact** on the rest of the population
  - the population distribution of actions becomes **deterministic**
- This should simplify the analysis
- Can we formalize this intuition?
- Idea: let N go to infinity and study the problem we obtain in the limit
- Key assumptions: homogeneity and anonymity
- "Mean field game" paradigm [Lasry, Lions; Caines, Huang, Malhamé 2006]

Outline

1. Motivations

2. MFG Models: Static Setting
   - Finite Population Game
   - Mean Field Games
   - Social Optimum

3. MFG Models: Dynamic setting

4. Optimality & Equilibrium Conditions

5. Conclusion

Notations

We assume homogeneity and anonymity:

$$f^i(\underline{a}) = f\Big(a^i, \frac{1}{N}\sum_{j=1}^{N}\delta_{a^j}\Big), \qquad i = 1, \ldots, N$$

Notations

We assume homogeneity and anonymity:

$$f^i(\underline{a}) = f\Big(a^i, \frac{1}{N}\sum_{j=1}^N \delta_{a^j}\Big), \qquad i = 1, \dots, N$$

Passing to the limit (formally) as $N \to +\infty$, we have the following setting:

- "Infinitely many" players
- representative player chooses a (mixed) strategy $\pi \in \mathcal{P}(\mathcal{A})$
- player picks an action according to the strategy $a \sim \pi$
- the empirical distribution $\frac{1}{N}\sum_{j=1}^N \delta_{a^j}$ converges to a population distribution of actions: $\pi' \in \mathcal{P}(\mathcal{A})$
- representative player pays a cost $f(a, \pi')$
- goal for each player: minimize her own average cost $J(\pi, \pi') = \mathbb{E}_{a\sim\pi}[f(a, \pi')]$

Notations

We assume homogeneity and anonymity:

$$f^i(\underline{a}) = f\left(a^i, \frac{1}{N}\sum_{j=1}^{N} \delta_{a^j}\right), \qquad i = 1, \dots, N$$

Passing to the limit (formally) as $N \to +\infty$, we have the following setting:

- "Infinitely many" players
- representative player chooses a (mixed) strategy $\pi \in \mathcal{P}(\mathcal{A})$
- player picks an action according to the strategy $a \sim \pi$
- the empirical distribution $\frac{1}{N}\sum_{j=1}^{N} \delta_{a^j}$ converges to a population distribution of actions: $\pi' \in \mathcal{P}(\mathcal{A})$
- representative player pays a cost $f(a, \pi')$
- goal for each player: minimize her own average cost $J(\pi, \pi') = \mathbb{E}_{a\sim\pi}[f(a, \pi')]$

Key points:

- it is enough to understand the behavior of one representative player
- each player has no influence on the rest of the population π'
- "mean-field interactions" is more general than "interactions through the mean"

Mean field Nash equilibrium

The notion of solution in a MFG is:

Definition (Mean field Nash equilibrium (MFNE))

$\hat{\pi} \in \mathcal{P}(\mathcal{A})$ is a mean field Nash equilibrium strategy if:

- $\hat{\pi}$ is an optimal strategy (best response) for a representative player, given the population distribution
- and the population distribution is $\hat{\pi}$

Mean field Nash equilibrium

The notion of solution in a MFG is:

Definition (Mean field Nash equilibrium (MFNE))

$\hat{\pi} \in \mathcal{P}(\mathcal{A})$ is a mean field Nash equilibrium strategy if:

- $\hat{\pi}$ is an optimal strategy (best response) for a representative player, given the population distribution
- and the population distribution is $\hat{\pi}$

Fixed point formulation:

$$\hat{\pi} \in \operatorname*{argmin}_{\pi} J(\pi, \hat{\pi}) = \mathrm{BR}(\hat{\pi})$$

Mean field Nash equilibrium

The notion of solution in a MFG is:

Definition (Mean field Nash equilibrium (MFNE))

$\hat{\pi} \in \mathcal{P}(\mathcal{A})$ is a mean field Nash equilibrium strategy if:

- $\hat{\pi}$ is an optimal strategy (best response) for a representative player, given the population distribution
- and the population distribution is $\hat{\pi}$

Fixed point formulation:

$$\hat{\pi} \in \operatorname*{argmin}_{\pi} J(\pi, \hat{\pi}) = \mathrm{BR}(\hat{\pi})$$

This yields a first algorithm: fixed point iterations $\pi^k \mapsto \pi^{k+1}$.

Simple to implement, but fails to converge on many examples. (More details later.)

Example: Population distribution

Exercise

Revisit the previous finite-population examples in the MFG setting.

- Attraction to the group
- Aversion to the group

From MFG to finite player game

Key motivation for MFG

An MFG equilibrium strategy provides an **approximate** (and usually **decentralized**) Nash equilibrium in the corresponding finite-population game.

From MFG to finite player game

Key motivation for MFG

An MFG equilibrium strategy provides an **approximate** (and usually **decentralized**) Nash equilibrium in the corresponding finite-population game.

Definition (Approximate Nash equilibrium in N-player game)

Let $\epsilon > 0$. $\underline{\hat{\pi}} \in \mathcal{P}(\mathcal{A})^N$ is an ϵ-Nash equilibrium for the N-player game if:

$$J^i(\hat{\pi}^i, \underline{\hat{\pi}}^{-i}) \leq J^i(\pi^i, \underline{\hat{\pi}}^{-i}) + \epsilon, \qquad \forall i, \forall \pi^i$$

Theorem (Informal statement)

Consider an N-player game and the corresponding MFG.

Let $\hat{\pi}$ be a mean field NE.

Then, in the N-player game, $\hat{\pi}$ is an ϵ-NE, with $\epsilon \to 0$ as $N \to +\infty$.

Interpretation: If everyone was using a MFG equilibrium policy, then anyone could be better off by at most ϵ by unilateral deviations.

Approximate NE: Example

Example (Interaction through the mean)

Consider a cost:

$$f(a, \nu) = \varphi(a, \bar{\nu}), \qquad \bar{\nu} = \mathbb{E}_{a' \sim \nu}[a']$$

with φ Lipschitz in ν uniformly in a.

Assumption: Mean field Nash equilibrium property: $\hat{\pi}$ such that

$$J(\hat{\pi}, \hat{\pi}) \leq J(\pi, \hat{\pi}), \qquad \forall \pi$$

Goal: ϵ-Nash equilibrium for N-player game:

$$J^i(\hat{\pi}, \hat{\pi}^{-i}) \leq J^i(\pi, \hat{\pi}^{-i}) + \epsilon, \qquad \forall \pi$$

where $\hat{\pi}^{-i} = (\hat{\pi}, \ldots, \hat{\pi}) \in \mathcal{P}(\mathcal{A})^{N-1}$.

Approximate NE: Example

Proof sketch:

- Idea: compare N-player cost with MF cost:

$$J^i(\hat{\pi}, \hat{\pi}^{-i}) - J^i(\pi, \hat{\pi}^{-i})$$

Approximate NE: Example

Proof sketch:

- Idea: compare N-player cost with MF cost:

$$
\begin{aligned}
&J^i(\hat{\pi}, \hat{\pi}^{-i}) - J^i(\pi, \hat{\pi}^{-i}) \\
&= \underbrace{J^i(\hat{\pi}, \hat{\pi}^{-i}) - J(\hat{\pi}, \hat{\pi})}_{\textbf{(i)}} + \underbrace{J(\hat{\pi}, \hat{\pi}) - J(\pi, \hat{\pi})}_{\leq 0} + \underbrace{J(\pi, \hat{\pi}) - J^i(\pi, \hat{\pi}^{-i})}_{\textbf{(ii)}}
\end{aligned}
$$

Approximate NE: Example

Proof sketch:

- Idea: compare N-player cost with MF cost:

$$
\begin{aligned}
&J^i(\hat\pi,\hat\pi^{-i}) - J^i(\pi,\hat\pi^{-i}) \\
&= \underbrace{J^i(\hat\pi,\hat\pi^{-i}) - J(\hat\pi,\hat\pi)}_{\textbf{(i)}} + \underbrace{J(\hat\pi,\hat\pi) - J(\pi,\hat\pi)}_{\leq 0} + \underbrace{J(\pi,\hat\pi) - J^i(\pi,\hat\pi^{-i})}_{\textbf{(ii)}}
\end{aligned}
$$

- $J^i(\hat\pi,\hat\pi^{-i}) = \mathbb{E}_{\hat a^i\sim\hat\pi}\mathbb{E}_{\hat a^j\sim\hat\pi, j\neq i}\left[\varphi(\hat a^i, \bar{\hat a})\right]$, where $\bar{\hat a} := \frac{1}{N}\sum_{j=1}^N \hat a^j$

- and $J(\hat\pi,\hat\pi) = \mathbb{E}_{\hat a\sim\hat\pi}\left[\varphi(\hat a, \bar{\hat\pi})\right]$, where $\bar{\hat\pi} := \mathbb{E}_{a\sim\hat\pi}[a]$

Approximate NE: Example

Proof sketch:

- Idea: compare N-player cost with MF cost:

$$J^i(\hat\pi, \hat\pi^{-i}) - J^i(\pi, \hat\pi^{-i})$$
$$= \underbrace{J^i(\hat\pi, \hat\pi^{-i}) - J(\hat\pi, \hat\pi)}_{(\mathbf{i})} + \underbrace{J(\hat\pi, \hat\pi) - J(\pi, \hat\pi)}_{\leq 0} + \underbrace{J(\pi, \hat\pi) - J^i(\pi, \hat\pi^{-i})}_{(\mathbf{ii})}$$

- $J^i(\hat\pi, \hat\pi^{-i}) = \mathbb{E}_{\hat a^i \sim \hat\pi} \mathbb{E}_{\hat a^j \sim \hat\pi, j \neq i} \left[\varphi(\hat a^i, \bar{\hat a})\right]$, where $\bar{\hat a} := \frac{1}{N} \sum_{j=1}^N \hat a^j$

- and $J(\hat\pi, \hat\pi) = \mathbb{E}_{\hat a \sim \hat\pi} \left[\varphi(\hat a, \bar{\hat\pi})\right]$, where $\bar{\hat\pi} := \mathbb{E}_{a \sim \hat\pi}[a]$

- We have φ is Lipschitz and $\bar{\hat a} \approx \bar{\hat a}^{-i} := \frac{1}{N} \sum_{j \neq i} \hat a^j \approx \bar{\hat\pi}$, so:

$$\left|\varphi(\hat a^i, \bar{\hat a}) - \varphi(\hat a^i, \bar{\hat\pi})\right| \leq C|\bar{\hat a} - \bar{\hat\pi}| \leq C \underbrace{|\bar{\hat a} - \bar{\hat a}^{-i}|}_{= \frac{1}{N}|\hat a^i|} + C|\bar{\hat a}^{-i} - \bar{\hat\pi}|$$

- Hence:

$$(\mathbf{i}) \leq \frac{C}{N} \mathbb{E}_{\hat a^i \sim \hat\pi} |\hat a^i| + C \underbrace{\mathbb{E}_{\hat a^j \sim \hat\pi, j \neq i} |\bar{\hat a}^{-i} - \bar{\hat\pi}|}_{\to 0 \text{ as } N \to \infty \text{ by LLN}}$$

- Similarly for $(\mathbf{ii})$

Outline

Definition

- Goal: minimize the social cost = average cost for the agents in the population

- N-agent social cost:

$$J^{soc}(\underline{\pi}) = \frac{1}{N}\sum_{i=1}^{N} J^i(\pi, \underline{\pi}^{-i})$$

- Mean field social cost:

$$J^{soc}(\pi) = J(\pi, \pi)$$

- Optimization problem and not fixed point problem anymore

Nash Equilibrium vs Social Optimum

- In general the two notions are different
- i.e., the socially optimal strategy is different from the Nash equilibrium policy
- Price of Anarchy [Koutsoupias & Papadimitriou, 1999]:

ABSTRACT

In a system where noncooperative agents share a common resource, we propose the price of anarchy, which is the ratio between the worst possible Nash equilibrium and the social optimum, as a measure of the effectiveness of the system. Deriving upper and lower bounds for this ratio in a model where several agents share a very simple network leads to some interesting mathematics, results, and open problems.[2]

- More on this later (see lectures 2 and 3, LQ setting and crowd motion)

Nash Equilibrium vs Social Optimum

- In *some cases*, the two notions coincide.
- Example: Potential MFG with cost: $f(a, \nu) = \nabla F(\nu)(a)$, $\mathcal{A}$ finite for simplicity
- The average cost is: $J(\pi, \nu) = \mathbb{E}_{a \sim \pi}[f(a, \nu)] = \sum_a \pi(a) \nabla F(\nu)(a) = \pi \cdot \nabla F(\nu)$
- Assuming the potential F convex, we have the equivalence:

$$
\begin{aligned}
\hat{\pi} \text{ is a NE} &\Leftrightarrow J(\pi, \hat{\pi}) - J(\hat{\pi}, \hat{\pi}) \geq 0, \quad \forall \pi \\
&\Leftrightarrow (\pi - \hat{\pi}) \cdot \nabla F(\hat{\pi}) \geq 0, \quad \forall \pi \\
&\Leftrightarrow \nabla F(\hat{\pi}) = 0 \\
&\Leftrightarrow \hat{\pi} \text{ is a minimizer of } F \\
&\Leftrightarrow \hat{\pi} \text{ is a SO}
\end{aligned}
$$

- Example: entropy: $F(\nu) = \sum_a \nu(a) \log(\nu(a))$
- More on this later (see lecture 3, optimization methods for variational MFGs)

Exercises

Exercise

Design a static MFG with exactly 2 pure NE. How many mixed NE are there?

Exercise

Design a static MFG with exactly 2 mixed social optima.

Exercise

Design a static MFG with a unique mixed NE and a unique mixed SO, such that their values are different (price of anarchy different from 1).

Same question with “such that their values are the same”.

Some references

- As far as I know, the static MFGs have not been studied extensively
- In fact a static MFG can be recast as a dynamic MFG with a single time step and and a single state
- Normal form MFGs (Coop/Betray/Punish, Rock/Paper/Scissor) [Muller et al., 2022b, Muller et al., 2022a]
- Static MFGs Section 2.1 in the survey [Laurière et al., 2022a]

Outline

Overview

Space \ Time	Discrete	Continuous
Discrete		
Continuous		

Overview

Space \ Time	Discrete	Continuous
Discrete		
Continuous		

For simplicity of presentation, we start with the “discrete & discrete” case.

Outline

Notations

For simplicity of presentation, we start with discrete time & discrete (finite) space.

- Time $T < +\infty$, $t \in [T] = \{0, 1, \dots, T\}$
- State space $\mathcal{X}$ finite (for now)
- Action space $\mathcal{A}$ finite (for now)
- Player's state $X_t^i \in \mathcal{X}$
- Population's state $m_t^N = \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i} \in \mathcal{P}(\mathcal{X})$
- One-step strategy (deterministic or mixed)
- "Control" or "policy": let us for now focus on the Markovian case:
  - Control (deterministic): $\alpha^i : [T] \times \mathcal{X} \to \mathcal{A}$
  - Policy (mixed): $\pi^i : [T] \times \mathcal{X} \to \mathcal{P}(\mathcal{A})$
  - Other choices (open-loop, ...)

Notations

- We assume homogeneity and anonymity, meaning:
  - same transition rules and same cost functions
  - interactions only through aggregate quantities

- Player's dynamics:

$$X^i_{t+1} \sim P(\cdot | X^i_t, A^i_t, m^N_t)$$

  where $A^i_t = \alpha^i(t, X^i_t)$ or $A^i_t \sim \pi^i(t, X^i_t)$

- For instance:

$$X^i_{t+1} = F(X^i_t, A^i_t, m_t) + \epsilon^i_{t+1}$$

  where ϵ^i_{t+1} is a random perturbation

- Population's dynamics:

$$m^N_{t+1} = \frac{1}{N} \sum_{i=1}^{N} \delta_{X^i_t}$$

  which is *random* if the players' states are.

Notations and Definition

- Running cost $f : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$

- Terminal cost $g : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$

- Total cost when player i uses policy π^i and the rest of the population uses $\underline{\pi}^{-i}$:

$$J^i(\pi^i, \underline{\pi}^{-i}) = \mathbb{E}\left[\sum_{t=0}^{T-1} f(X_t^i, A_t^i, m_t^N) + g(X_T^i, m_T^N)\right]$$

Notations and Definition

- Running cost $f : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$
- Terminal cost $g : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$
- Total cost when player i uses policy π^i and the rest of the population uses $\underline{\pi}^{-i}$:

$$J^i(\pi^i, \underline{\pi}^{-i}) = \mathbb{E}\left[\sum_{t=0}^{T-1} f(X_t^i, A_t^i, m_t^N) + g(X_T^i, m_T^N)\right]$$

Definition (Nash equilibrium in dynamic N-player game)

$\hat{\underline{\pi}} \in \mathcal{P}(\mathcal{A})^N$ is a Nash equilibrium for the N-player game if:

$$J^i(\hat{\pi}^i, \hat{\underline{\pi}}^{-i}) \leq J^i(\pi^i, \hat{\underline{\pi}}^{-i}), \qquad \forall i, \forall \pi^i$$

The definition of Nash equilibrium is exactly the same, but the definition of J^i is more involved than in the static case.

Outline

Notations

- Time $T < +\infty$, $t \in [T] = \{0, 1, \dots, T\}$
- State space $\mathcal{X}$ finite (for now)
- Action space $\mathcal{A}$ finite (for now)
- Player's state $X_t \in \mathcal{X}$
- Population's state $m_t \in \mathcal{P}(\mathcal{X})$, identified with a vector of dimension $|\mathcal{X}|$
- One-step strategy (deterministic or mixed)
- "Control" or "policy": let us for now focus on the Markovian case:
  - Policy (mixed): $\pi : [T] \times \mathcal{X} \to \mathcal{P}(\mathcal{A})$
  - Control (deterministic): $\alpha : [T] \times \mathcal{X} \to \mathcal{A}$
  - Deterministic is a special case of mixed: $\pi_t(a|x) = \delta_{\alpha_t(x)}(a)$
  - Other choices (open-loop, ...)

Notations

- Time $T < +\infty$, $t \in [T] = \{0, 1, \dots, T\}$
- State space $\mathcal{X}$ finite (for now)
- Action space $\mathcal{A}$ finite (for now)
- Player's state $X_t \in \mathcal{X}$
- Population's state $m_t \in \mathcal{P}(\mathcal{X})$, identified with a vector of dimension $|\mathcal{X}|$
- One-step strategy (deterministic or mixed)
- "Control" or "policy": let us for now focus on the Markovian case:
  - Policy (mixed): $\pi : [T] \times \mathcal{X} \to \mathcal{P}(\mathcal{A})$
  - Control (deterministic): $\alpha : [T] \times \mathcal{X} \to \mathcal{A}$
  - Deterministic is a special case of mixed: $\pi_t(a|x) = \delta_{\alpha_t(x)}(a)$
  - Other choices (open-loop, ...)
- Remark: decentralized policies/controls are enough (in this setting at least) i.e., we do not need to add m_t as an input to π_t or α_t

Notations

- Player's dynamics given a mean-field sequence $m = (m_t)_{t=0,\dots,T}$:

$$X_{t+1} \sim P(\cdot|X_t, A_t, m_t)$$

  where $A_t = \alpha(t, X_t)$ or $A_t \sim \pi(t, X_t)$

- For instance:

$$X_{t+1} = F(X_t, A_t, m_t) + \epsilon_{t+1}$$

- Population distribution dynamics associated to a policy π:

$$m^\pi_{t+1} = P^\pi_t m^\pi_t$$

  where P^π_t is the transition matrix

$$P^\pi_t(x, x') = \sum_{a \in \mathcal{A}} \pi_t(a|x) P(x'|x, a, m^\pi_t)$$

- P^π_t depends implicitly on m^π_t $\Rightarrow$ non linear Markov chain

Notations

- Running cost $f : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$
- Terminal cost $g : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$
- Total cost for a representative player using policy π and the rest of the population uses π':

$$J(\pi, \pi') = \mathbb{E}\left[\sum_{t=0}^{T-1} f(X_t, A_t, m_t^{\pi'}) + g(X_T, m_T^{\pi'})\right]$$

Notations

- Running cost $f : \mathcal{X} \times \mathcal{A} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$
- Terminal cost $g : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$
- Total cost for a representative player using policy π and the rest of the population uses π':
$$J(\pi, \pi') = \mathbb{E}\left[\sum_{t=0}^{T-1} f(X_t, A_t, m_t^{\pi'}) + g(X_T, m_T^{\pi'})\right]$$
- Actually $J(\pi, \pi') = J(\pi, m^{\pi'})$, and we can define $J(\pi, m)$ more generally
- Best response (BR) to a mean field $m \in [T] \times \mathcal{P}(\mathcal{X})$:
$$\mathrm{BR}(m) = \operatorname*{argmin}_{\pi} J(\pi, m)$$
- In general, $\mathrm{BR}(m)$ is a set.

Example: Crowd motion in a grid world

- State space: grid world
- Dynamics: move to a neighboring cell
- Running cost:
  - cost to move:
  - discomfort if crowded:
- Terminal cost:
  - spatial preference:
- Illustrations: [Geist et al., 2022, Laurière et al., 2022b]

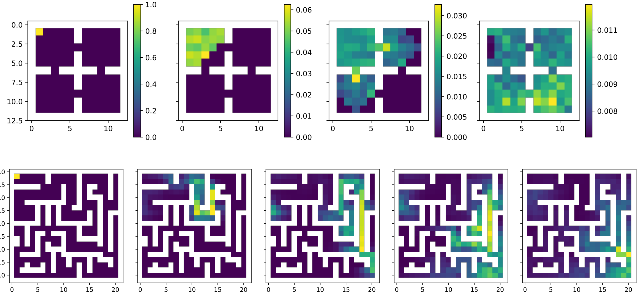

The notion of solution in a MFG is:

Definition (Mean field Nash equilibrium (MFNE))

$\hat{\pi}$ is a mean field Nash equilibrium policy if:

- $\hat{\pi}$ is an optimal policy for a representative player, given the population distribution
- and the population distribution is *generated* by $\hat{\pi}$

Mean field Nash equilibrium (dynamic)

The notion of solution in a MFG is:

Definition (Mean field Nash equilibrium (MFNE))

$\hat{\pi}$ is a mean field Nash equilibrium policy if:

- $\hat{\pi}$ is an optimal policy for a representative player, given the population distribution
- and the population distribution is *generated* by $\hat{\pi}$

Fixed point formulation:

$$\hat{\pi} \in \operatorname*{argmin}_{\pi} J(\pi, m^{\hat{\pi}})$$

Mean field Nash equilibrium (dynamic)

The notion of solution in a MFG is:

Definition (Mean field Nash equilibrium (MFNE))

$\hat{\pi}$ is a mean field Nash equilibrium policy if:

- $\hat{\pi}$ is an optimal policy for a representative player, given the population distribution
- and the population distribution is *generated* by $\hat{\pi}$

Fixed point formulation:

$$\hat{\pi} \in \operatorname*{argmin}_{\pi} J(\pi, m^{\hat{\pi}})$$

This yields a first algorithm: fixed point iterations $\pi^k \mapsto m^k \mapsto \pi^{k+1}$.

Simple to implement, but fails to converge on many examples. (See lectures 2 and 6.)

Existence

Existence of an equilibrium is generally based on the fixed point formulation

Typically:

- Banach/Picard fixed point theorem (requires strict contraction)

- Brouwer/Schauder fixed point theorem (requires only continuity)

Uniqueness

Uniqueness is typically ensured through two types of arguments:

- **Strict contractivity:** uniqueness is obtained as a consequence of Banach fixed point theorem
- **Monotonicity:** uniqueness is a consequence of the monotonicity of the cost
  - Typical setting: $b(x,a,m) = b(x,a)$, $f(x,a,m) = \tilde{f}(x,a) + V(x,m)$
  - V is monotone in L^2 if: $m_1, m_2 \in L^2(\mathbb{R}^d)$,

$$\int (V(x,m_1) - V(x,m_2))(m_1(x) - m_2(x))dx \geq 0$$

  - Example: crowd aversion

Example of existence proof

Sketch of existence proof: look for a fixed point of $\Phi : \pi \xrightarrow{\text{MF}} \tilde{m} \xrightarrow{\text{BR}} \tilde{\pi}$

A simple model:

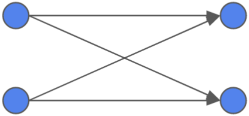

- $\mathcal{X} = \{0, 1\}$, $\mathcal{A} = \{-1, 0, 1\}$, $T = 1$
- $X_{t+1} = X_t + A_t$ with walls at $x = -1, 2$
- $f(x, a, m) = 0, g(x, m) = |x - \bar{m}|, \bar{m} =$ mean of m
- $m_0 = (\frac{1}{2}, \frac{1}{2})$

Example of existence proof

Step 1: Convexity and compactness

Step 2: Continuity of Φ

Step 2.a: Continuity of MF?

Step 2.b: Continuity of BR?

Exercises

Exercise

Complete the proof of existence in the previous example (2 states, 1 time step).

Exercise

Define a dynamic MFG such that the following two conditions hold:

1. there is a unique NE
2. given the equilibrium mean field sequence m, $\mathrm{BR}(m)$ is not a singleton (there are multiple optimal policies)

Some references

Some references on discrete time, finite state space MFGs:

- [Gomes et al., 2010]
- Link with continuous MFGs [Hadikhanloo and Silva, 2019]
- Reinforcement learning for MFGs is often studied in this setting: [Chen et al., 2022], [Guo et al., 2019], [Subramanian and Mahajan, 2019], [Elie et al., 2020b], ...
- See lecture 6 for more references on RL for MFGs; survey: [Laurière et al., 2022a]

Outline

Motivations

Why do we care about continuous time & space?

- More natural for many applications
- Discretizing a continuous time/space process is not trivial
- We can use calculus

Example: Epidemics (continuous time, discrete space)

Example (SIR model)

3 possible states: Susceptible of infection, Infected, Recovered. Mean field dynamics:

$$\begin{cases} \dot{S}(t) = -\beta I(t)S(t) \\ \dot{I}(t) = \beta I(t)S(t) - \gamma I(t) \\ \dot{R}(t) = \gamma I(t) \end{cases}$$

Basic reproduction number: $R_0 = \beta/\gamma$.

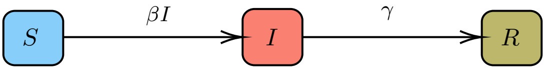

Example (SIR model)

3 possible states: Susceptible of infection, Infected, Recovered. Mean field dynamics:

$$\begin{cases} \dot{S}(t) = -\beta I(t) S(t) \\ \dot{I}(t) = \beta I(t) S(t) - \gamma I(t) \\ \dot{R}(t) = \gamma I(t) \end{cases}$$

Basic reproduction number: $R_0 = \beta/\gamma$.

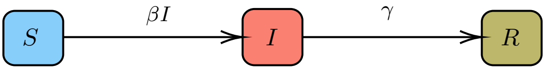

MFG for epidemics: [Laguzet and Turinici, 2015], [Hubert and Turinici, 2018], [Elie et al., 2020a], [Lee et al., 2021], [Olmez et al., 2022], [Aurell et al., 2022b], [Doncel et al., 2022], [Aurell et al., 2022a] ...

Example: MFG for epidemics (continuous time, discrete space)

Example (MFG extension of SIR; borrowed from [Aurell et al., 2022b])

- Action: α_t = "contact factor"
- Action and state distribution $\rho(a, x)$
- Individual's transition rate from S to I: $\beta\alpha_t \int a\rho_t(a, I)da$

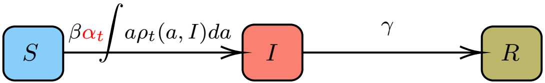

Example: MFG for epidemics (continuous time, discrete space)

Example (MFG extension of SIR; borrowed from [Aurell et al., 2022b])

- Action: $\alpha_t =$ "contact factor"
- Action and state distribution $\rho(a, x)$
- Individual's transition rate from S to I: $\beta \alpha_t \int a \rho_t(a, I) da$

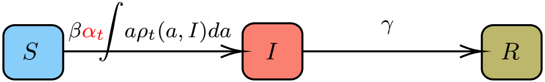

Exercise

Write the mean field dynamics corresponding to the above model, for a fixed state-action distribution flow $\rho = (\rho_t)_{t \in [0,T]}, \rho_t \in \mathcal{P}(\mathcal{A} \times \mathcal{X})$.

Example: MFG for epidemics (continuous time, discrete space)

Example (MFG extension of SIR; borrowed from [Aurell et al., 2022b])

- Action: α_t = "contact factor"
- Action and state distribution $\rho(a, x)$
- Individual's transition rate from S to I: $\beta\alpha_t \int a\rho_t(a, I)da$

$$S \xrightarrow{\beta\alpha_t \int a\rho_t(a, I)da} I \xrightarrow{\gamma} R$$

Exercise

Write the mean field dynamics corresponding to the above model, for a fixed state-action distribution flow $\rho = (\rho_t)_{t\in[0,T]}, \rho_t \in \mathcal{P}(\mathcal{A} \times \mathcal{X})$.

Missing ingredient: to have game, we need to define a cost function.

Other MFG models in discrete space & continuous time: [Gomes et al., 2013, Kolokoltsov and Bensoussan, 2016, Bayraktar et al., 2021], ...

Example: Flocking (continuous time, continuous space)

Example (Cucker-Smale model [Cucker and Smale, 2007])

Position and velocity:

$$\begin{cases} \dot{x}^i(t) = v^i(t) \\ \dot{v}^i(t) = \displaystyle\sum_{j=1}^{N} \frac{v_j(t) - v^i(t)}{(\epsilon + |x^j - x^i|)^\beta} \end{cases}$$

Example: Flocking (continuous time, continuous space)

Example (Cucker-Smale model [Cucker and Smale, 2007])

Position and velocity:

$$\begin{cases} \dot{x}^i(t) = v^i(t) \\ \dot{v}^i(t) = \displaystyle\sum_{j=1}^{N} \frac{v_j(t) - v^i(t)}{(\epsilon + |x^j - x^i|)^\beta} \end{cases}$$

MFG for flocking & acceleration control: [Nourian et al., 2011], [Grover et al., 2018], [Achdou et al., 2020], [Bardi and Cardaliaguet, 2021], [Santambrogio and Shim, 2021], [Perrin et al., 2021], ...

Example: Flocking (continuous time, continuous space)

Example (MFG model of flocking [Nourian et al., 2011])

- Action: $\alpha_t =$ acceleration. Dynamics:

$$\begin{cases} \dot{x}(t) = v(t) \\ \dot{v}(t) = \alpha(t) \end{cases}$$

- Running cost: penalizes deviation from neighbors' velocity:

$$f((x,v),m) = \left| \int \frac{v' - v}{(\epsilon + |x' - x|)^{\beta}} m(dx', dv') \right|^2$$

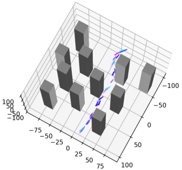

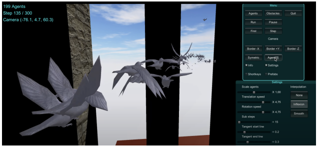

[Perrin et al., 2021], video: `https://www.youtube.com/watch?v=TdXysW_FA3k`

Interacting particle systems

Diffusion (source: Wikipedia):

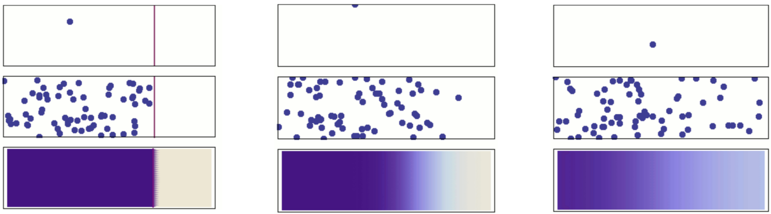

- Particle's dynamics: $dX_t = \sigma dW_t$, with W a Brownian motion
- Macroscopic distribution dynamics: $\partial_t m(t,x) - \frac{\sigma^2}{2}\Delta m(t,x) = 0$
- Link with N-particle system: propagation of chaos [Kac, 1956, Sznitman, 1991]

Interacting particle systems

Diffusion (source: Wikipedia):

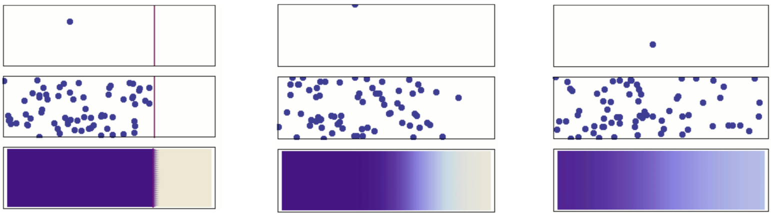

- Particle's dynamics: $dX_t = \sigma dW_t$, with W a Brownian motion
- Macroscopic distribution dynamics: $\partial_t m(t,x) - \frac{\sigma^2}{2}\Delta m(t,x) = 0$
- Link with N-particle system: propagation of chaos [Kac, 1956, Sznitman, 1991]
- Note: We can also add a transport term (convection–diffusion equation):
  - $dX_t = b(t, X_t)dt + \sigma dW_t$
  - $\partial_t m(t,x) - \frac{\sigma^2}{2}\Delta m(t,x) + \mathrm{div}(b(t,x)m(t,x)) = 0$

Continuous time, continuous space MFG

- Player i's state $X_t^i \in \mathbb{R}^d$
- with dynamics:

$$dX_t^i = b(t, X_t^i, \alpha^i(t, X_t^i), m_t^N)dt + \sigma dW_t^i, \qquad X_0^i \sim m^0$$

- W^i is an idiosyncratic (individual) noise, independent from other W^j's
- W is a noise for the representative player
- The population empirical distribution is:

$$m_t^N = \frac{1}{N} \sum_{j=1}^{N} \delta_{X_t^j}$$

- Here again, it is stochastic . . . but less and less as $N \to +\infty$
- Passing rigorously to the limit in the MFG framework: see e.g. [Cardaliaguet et al., 2019], Lacker's lecture notes [Lacker, 2018, Delarue, 2021] and the references therein

Continuous time, continuous space MFG

- Time horizon $T < +\infty$, $t \in [0, T]$
- Player's state $X_t \in \mathbb{R}^d$
- Player's control (deterministic) α_t, typically:
  - most often focus on deterministic controls
  - closed-loop Markovian: $\alpha_t = \alpha(t, X_t)$
  - open-loop: $\alpha_t = \alpha(t, \omega)$ progressively measurable
- Player's dynamics:

$$dX_t = b(t, X_t, \alpha(t, X_t), m_t)dt + \sigma dW_t, \qquad X_0 \sim m^0$$

- Population dynamics: Kolmogorov-Fokker-Planck equation

$$\partial_t m(t,x) - \frac{\sigma^2}{2}\Delta m(t,x) + \mathrm{div}(b(t,x,\alpha(t,x))m(t,x)) = 0, \qquad m_{|t=0} = m^0$$

Continuous time, continuous space MFG

Cost: dependence on the mean field

- non-local (typically “regularizing” operator)

$$f(t, X_t, \alpha_t, m_t)$$

- local (if the population distribution has a density, still denoted by m)

$$f(t, X_t, \alpha_t, m(t, X_t))$$

Next steps

Main question for the rest of this course:

How can we characterize and compute mean field Nash equilibria?

Outline

Literature

Many papers on MFGs start like this:

INTRODUCTION

This paper is devoted to the analysis of second order mean field games systems with a local coupling. The general form of these systems is:

$$\left\{\begin{array}{ll} (i) & \boxed{-\partial_t \phi - A_{ij}\partial_{ij}\phi + H(x, D\phi) = f(x, m(x,t))} \\ (ii) & \partial_t m - \partial_{ij}(A_{ij}m) - \operatorname{div}(mD_pH(x, D\phi)) = 0 \\ (iii) & m(0) = m_0,\ \phi(x,T) = \phi_T(x) \end{array}\right. \quad (1)$$

Source: Cardaliaguet, P., Graber, P.J., Porretta, A. and Tonon, D., 2015. Second order mean field games with degenerate diffusion and local coupling. Nonlinear Differential Equations and Applications NoDEA, 22(5), pp.1287-1317.

What do these equations mean?

In a nutshell, the probabilistic approach to the solution of the mean-field game problem results in the solution of a FBSDE of the McKean–Vlasov type

$$(3.1) \qquad \begin{cases} dX_t = b\big(t, X_t, \mathbb{P}_{X_t}, \hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t)\big)dt + \sigma dW_t, \\ \boxed{dY_t = -\partial_x H\big(t, X_t, \mathbb{P}_{X_t}, Y_t, \hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t)\big)dt + Z_t dW_t,} \end{cases}$$

with the initial condition $X_0 = x_0 \in \mathbb{R}^d$, and terminal condition $Y_T = \partial_x g(X_T, \mathbb{P}_{X_T})$.

Source: Carmona, R. and Delarue, F., 2013. Probabilistic analysis of mean-field games. SIAM Journal on Control and Optimization, 51(4), pp.2705-2734.

Value function

- Value of a state = sum of future costs, when starting from this state

- Value function of a representative player given a mean field sequence

Bellman equation

- Bellman equation for the value function (Dynamic Programming Principle):
  - Terminal time: $V_T(x) = g(X_T, m_T)$
  - Backward induction:
    $V_t(x) = \min_a \mathbb{E}\left[f(X_t, A_t, m_t) + V_{t+1}(X_{t+1}) | X_t = x, A_t = a\right]$
- Recovering the optimal control from the value function: using argmin

Forward-backward system

Coupled system:

- Forward equation for the mean field:

$$\begin{cases} m_{t+1}(x) = \sum_{x'} m_t(x') \sum_a \pi_t(a|x) p(x|x', a, m_t), \\ m_0 \text{ given} \end{cases}$$

- Backward equation for the value function:

$$\begin{cases} V_t(x) = \min_a \mathbb{E}\left[f(X_t, A_t, m_t) + V_{t+1}(X_{t+1}) | X_t = x, A_t = a \right], \\ V_T(x) = g(X_T, m_T) \end{cases}$$

- Equilibrium policy: π satisfies: (1) is optimal against m and (2) generates m

Challenge: We cannot (fully) solve one equation before the other!

Outline

Value function

- Value of a state = sum of future costs, when starting from this state
- Value function of a representative player given a mean field flow
- Dynamic Programming Principle?

HJB equation

- Hamiltonian:

$$H(x,m,p) = \max_a -L(x,a,m,p), \quad L(x,a,m,p) = f(x,a,m) + b(x,a,m)\cdot p$$

- Hamilton-Jacobi-Bellman equation, given the mean field flow:

$$\begin{cases} -\partial_t u(t,x) - \frac{\sigma^2}{2}\Delta u(t,x) + H(x, m(t), \nabla u(t,x))) = 0, \\ u(T,x) = g(x, m(T)) \end{cases}$$

- Recovering the optimal control: optimizer of the Hamiltonian

- Unique action minimizes H under strict convexity assumptions

HJB equation

- Hamiltonian:

$$H(x,m,p) = \max_a -L(x,a,m,p), \quad L(x,a,m,p) = f(x,a,m) + b(x,a,m) \cdot p$$

- Hamilton-Jacobi-Bellman equation, given the mean field flow:

$$\begin{cases} -\partial_t u(t,x) - \frac{\sigma^2}{2}\Delta u(t,x) + H(x, m(t), \nabla u(t,x))) = 0, \\ u(T,x) = g(x, m(T)) \end{cases}$$

- Recovering the optimal control: optimizer of the Hamiltonian
- Unique action minimizes H under strict convexity assumptions
- **Warning:** Another convention: $H(x,m,p) = \min_a L(x,a,m,p) \Rightarrow -H$ in HJB.

Forward-backward PDE system for MFG

The equilibrium control minimizes the Hamiltonian:

$$\hat{\alpha}(t,x) = \operatorname*{argmax}_a -L(x, a, m(t), \nabla u(t,x))$$

where (m, u) solve the forward-backward PDE system:

- Forward equation for the mean field:

$$\begin{cases} \partial_t m(t,x) - \frac{\sigma^2}{2}\Delta m(t,x) + \operatorname{div}(m(t,x) H_p(x, m(t), \nabla u(t,x))) = 0, \\ m(0,x) = m_0(x) \end{cases}$$

- Backward equation for the value function:

$$\begin{cases} -\partial_t u(t,x) - \frac{\sigma^2}{2}\Delta u(t,x) + H(x, m(t), \nabla u(t,x))) = 0, \\ u(T,x) = g(x, m(T)) \end{cases}$$

Challenge: We cannot (fully) solve one equation before the other!

Existence and uniqueness of MFNE

- **Existence:** generally obtained by applying a fixed point theorem, such as:
  - Banach fixed point theorem: typically applicable under “smallness” conditions (small time or small Lipschitz constants); gives uniqueness too
  - Schauder fixed point theorem: applicable more generally; does not yield uniqueness
  - Compactness can be challenging
- **Uniqueness:**
  - Contractivity (application of Banach fixed point theorem; “smallness” assumptions)
  - Monotonicity condition (Lasry & Lions; “structural” assumption)

Remark: Variational MFGs

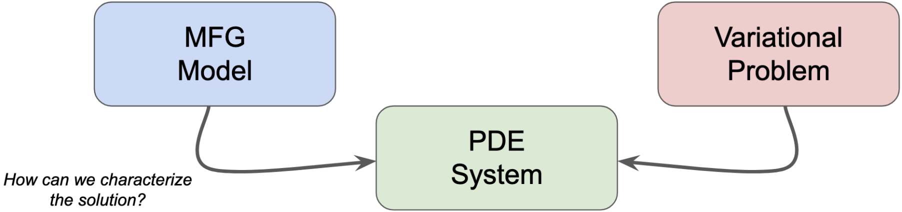

In some cases, the MFG PDE system can be interpreted as the optimality conditions for a variational problem

See e.g. [Lasry and Lions, 2007], [Cardaliaguet and Graber, 2015], . . .

This can also inspire numerical methods. (More on this in lecture 3.)

Remark: MFG with common noise

- Common noise: randomness affecting the whole population
- Example: extra Brownian motion common to all the players
- Then the two PDEs become stochastic PDEs

$$
\begin{cases}
d_t u_t = \left\{ -(1+\beta)\Delta u_t + H(x, Du_t) - F(x, m_t) - \sqrt{2\beta}\,\mathrm{div}(v_t) \right\} dt + v_t \cdot \sqrt{2\beta} dW_t \\
\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{in } [0,T] \times \mathbb{T}^d, \\
d_t m_t = \left[(1+\beta)\Delta m_t + \mathrm{div}\big(m_t D_p H(m_t, Du_t) \big) \right] dt - \mathrm{div}\big(m_t \sqrt{2\beta} dW_t \big), \\
\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{in } [0,T] \times \mathbb{T}^d, \\
u_T(x) = G(x, m_T), \; m_0 = m_{(0)}, \qquad \text{in } \mathbb{T}^d
\end{cases}
$$

Source: [Cardaliaguet et al., 2019]

Remark: Master equation

- Common noise: randomness affecting the whole population
- Example: extra Brownian motion common to all the players
- Convergence analysis (as $N \to \infty$) based on the Master equation

$$
\begin{cases}
-\partial_t U - (1+\beta)\Delta_x U + H(x, D_x U) \\
\qquad -(1+\beta)\displaystyle\int_{\mathbb{R}^d} \operatorname{div}_y \left[D_m U\right] \, dm(y) + \int_{\mathbb{R}^d} D_m U \cdot D_p H(y, D_x U) \, dm(y) \\
\qquad -2\beta \displaystyle\int_{\mathbb{R}^d} \operatorname{div}_x \left[D_m U\right] dm(y) - \beta \int_{\mathbb{R}^{2d}} \operatorname{Tr}\left[D^2_{mm} U\right] \, dm \otimes dm = F(x, m) \\
\qquad\qquad \text{in } [0, T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \\
U(T, x, m) = G(x, m) \qquad \text{in } \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)
\end{cases}
$$

Source: [Cardaliaguet et al., 2019]

Exercise

For the following drift and running cost functions ($d = 1$ to simplicity), write the KFP equation, the Hamiltonian and the HJB equation:

- Linear-quadratic (LQ):

$$b(x,a,m) = Ax + Ba + \bar{A}\bar{m}^2,\, f(x,a,m) = Qx^2 + Ra^2 + \bar{Q}\bar{m}^2,\, g(x,m) = Q_T x^2 + \bar{Q}_T \bar{m}^2$$

  with $\bar{m} = \int \xi m(\xi) d\xi$
- Congestion: $b(x,a,m) = a$, $f(x,a,m) = m(x)|a|^2$
- Aversion: $b(x,a,m) = a$, $f(x,a,m) = |a|^2 + m(x)$

Exercise

Derive optimality conditions for the social optimum problem.

Exercises

Exercise [Bogachev, Krylov, Röckner, Shaposhnikov; Thm 9.8.41]

Consider the MFG PDE system:

$$\begin{cases} -\partial_t u - \Delta u + \frac{1}{2}|\nabla u|^2 = F(x, m_t), & \mathbb{R}^d \times [0,T) \\ \partial_t m - \Delta m - \operatorname{div}(\nabla u m) = 0, & \mathbb{R}^d \times (0,T] \end{cases}$$

with $u(T,x) = G(x, m(T))$ and $m_0 = \nu$.

Part 1: Write the player's dynamics and the cost function.

Part 2: Show existence of a classical solution, assuming:

- ν is a probability distribution on $\mathbb{R}^d$ with finite second moment
- $F, G : \mathbb{R}^d \times \mathcal{P}_1(\mathbb{R}^d) \to \mathbb{R}$ are bounded and Lipschitz

Source: [Bogachev et al., 2022]

Outline

From Bellman equation

- Value function of a representative player given a mean field sequence

- Hamilton-Jacobi-Bellman equation, given equilibrium mean field flow $\hat{m}_t$:

$$-\partial_t u(t,x) - \frac{\sigma^2}{2}\Delta u(t,x) + H(x, \hat{m}(t), \nabla u(t,x))) = 0, \qquad u(T,x) = g(x, m(T))$$

From Bellman equation

- Value function of a representative player given a mean field sequence
- Hamilton-Jacobi-Bellman equation, given equilibrium mean field flow $\hat{m}_t$:

$$-\partial_t u(t,x) - \frac{\sigma^2}{2}\Delta u(t,x) + H(x, \hat{m}(t), \nabla u(t,x))) = 0, \qquad u(T,x) = g(x, m(T))$$

- Actually in practice, we do not really need to know u *everywhere*
- Motivation for probabilistic numerical methods (see lectures 3, 4, 5)
- Two approaches, based respectively on Bellman & Pontryagin principles

From Bellman equation

- We want to know u and the control along the path of X
- Introduce $Y_t = u(t, X_t)$ where u is the value function, solution to HJB
- Dynamics of Y:

$$
\begin{aligned}
dY_t &= \frac{d}{dt} u(t, X_t) \\
&= \left[\partial_t u(t, X_t) + \frac{\sigma^2}{2} \Delta u(t, X_t) \right] dt + \nabla u(t, X_t) dX_t && \text{(by Itô's lemma)} \\
&= -f(X_t, A_t^*, \hat{m}_t) dt + Z_t dW_t && \text{(by HJB equation)}
\end{aligned}
$$

where A_t^* is the optimal action when at state X_t

Pontryagin's maximum principle: ODE

- Assume X has a deterministic evolution:

$$\dot{x}_t = b(x_t, a_t), \quad x_0 \text{ given}$$

- Hamiltonian:

$$H(x,p) = \max_a -L(x,a,p), \quad L(x,a,p) = f(x,a) + b(x,a) \cdot p$$

- Pontryagin's maximum principle:

$$\begin{cases} \dot{x}_t^* = b(x_t^*, a_t^*), & x_0^* \text{ given} \\ \dot{y}_t^* = -\nabla L(x_t^*, a_t^*, y_t^*), & y_T^* = \nabla g(x_T^*) \\ a_t^* = \operatorname{argmax}_a -L(x_t^*, a, y_t^*) & \end{cases}$$

- In fact, y_t^* can be interpreted as $\nabla u(t, x_t^*)$

Pontryagin's maximum principle: SDE

- If X satisfies an SDE:

$$dX_t = b(X_t, A_t)dt + \sigma dW_t, \quad X_0 \sim m_0$$

- Hamiltonian:

$$H(x,p) = \max_a -L(x,a,p), \quad L(x,a,p) = f(x,a) + b(x,a) \cdot p$$

- Stochastic Pontryagin maximum principle:

$$\begin{cases} dX_t^* = b(X_t^*, A_t^*)dt + \sigma dW_t, & X_0^* \sim m_0 \\ dY_t^* = -\nabla L(X_t^*, A_t^*, Y_t^*)dt + Z_t^* dW_t, & Y_T^* = \nabla g(X_T^*) \\ A_t^* = \operatorname{argmax}_a -L(X_t^*, a, Y_t^*) \end{cases}$$

- In fact, Y_t^* can be interpreted as $\nabla u(t, X_t^*)$

Pontryagin's maximum principle: MKV SDE

- If X satisfies a mean field SDE:

$$dX_t = b(X_t, A_t, \hat{m}_t)dt + \sigma dW_t, \quad X_0 \sim m_0$$

- Hamiltonian:

$$H(x,p) = \max_a -L(x,a,m,p), \quad L(x,a,m,p) = f(x,a,m) + b(x,a,m) \cdot p$$

- Stochastic Pontryagin maximum principle with mean field interactions:

$$\begin{cases} dX_t^* = b(X_t^*, A_t^*, \hat{m}_t)dt + \sigma dW_t, & X_0^* \sim m_0 \\ dY_t^* = -\nabla L(X_t^*, A_t^*, \hat{m}_t, Y_t^*)dt + Z_t^* dW_t, & Y_T^* = \nabla g(X_T^*, \hat{m}_T) \\ A_t^* = \operatorname{argmax}_a -L(X_t^*, a, \hat{m}_t, Y_t^*) \end{cases}$$

- In fact, Y_t^* can be interpreted as $\nabla u(t, X_t^*)$
- For the equilibrium, we need to include the consistency condition for the MF
- $\hat{m}_t = \mathcal{L}(X_t^*)$

MKV FBSDE systems

In both cases (from Bellman or Pontryagin's principles), we get an instance of a McKean-Vlasov forward-backward SDEs (MKV-FBSDE):

$$\begin{cases} dX_t = B(X_t, Y_t, Z_t, m_t)dt + \sigma dW_t, & X_0 \sim m_0 \\ dY_t = F(X_t, Y_t, Z_t, m_t)dt + Z_t dW_t, & Y_T = G(X_T, m_T) \\ m_t = \mathcal{L}(X_t) \end{cases}$$

- Analysis: existence, uniqueness, ...
- Extensions (common noise, ...)
- Link with Master equation
- See book [Carmona and Delarue, 2018a, Carmona and Delarue, 2018b] for (many) more details

Exercise

For the following drift and running cost functions ($d = 1$ to simplicity), write the MKV FBSDE system:

- Linear-quadratic (LQ):

$$b(x,a,m) = Ax + Ba + \bar{A}\bar{m}^2, f(x,a,m) = Qx^2 + Ra^2 + \bar{Q}\bar{m}^2, g(x,m) = Q_T x^2 + \bar{Q}_T \bar{m}^2$$

  with $\bar{m} = \int \xi m(\xi) d\xi$
- Congestion: $b(x,a,m) = a, f(x,a,m) = m(x)|a|^2$
- Aversion: $b(x,a,m) = a, f(x,a,m) = |a|^2 + m(x)$

Exercise

Derive an FBSDE system for the social optimum problem.

Outline

Summary

- N-player games
- Mean field games
- Connection in two directions
- Several settings (static, dynamics discrete/continuous)
- Optimality conditions

Extensions

From the modeling viewpoint, many possible extensions:

- More settings, e.g. MFG with **ergodic** cost [Cardaliaguet et al., 2012], [Feleqi, 2013], [Bardi and Priuli, 2014], [Arapostathis et al., 2017], [Anahtarci et al., 2023], ...
- Interactions through the **action distribution** ("extended MFGs", "MFGs of controls", ...): [Gomes et al., 2014], [Gomes and Voskanyan, 2016], [Cardaliaguet and Lehalle, 2018], [Achdou and Kobeissi, 2020], [Laurière and Tangpi, 2022], [Kobeissi, 2022], ...
- **Common noise**: in the continuous space case see [Carmona and Delarue, 2018b] and references therein; in the finite state case, see e.g. [Bertucci et al., 2019], [Bayraktar et al., 2021], ...
- **Several populations** MFGs: [Huang et al., 2006], [Feleqi, 2013], [Cirant, 2015], [Achdou et al., 2017], [Bensoussan et al., 2018], ...
- **Mean field type games**: [Djehiche et al., 2017], [Barreiro-Gomez and Tembine, 2021] and references therein; [Miller and Pham, 2019], [Cosso and Pham, 2019], [Carmona et al., 2019], ...
- **Mean field control games**: [Angiuli et al., 2022b], [Angiuli et al., 2022a]

Extensions

- **Major player**: [Carmona and Zhu, 2016], [Caines and Kizilkale, 2016], [Carmona and Wang, 2017], [Lasry and Lions, 2018], [Cardaliaguet et al., 2020], [Carmona and Dayanıklı, 2021], [Carmona et al., 2022b], ...
- **Stackelberg** MFGs [Bensoussan et al., 2015], [Moon and Başar, 2018], [Elie et al., 2019], [Firoozi et al., 2021], [Aurell et al., 2022b], [Vasal and Berry, 2022], [Guo et al., 2022], [Dayanikli and Lauriere, 2023],...
- **Graphon** games [Parise and Ozdaglar, 2019], [Caines and Huang, 2019], [Caines and Huang, 2021], [Lacker and Soret, 2022], [Gao et al., 2020], [Vasal et al., 2021], [Carmona et al., 2022a], [Aurell et al., 2022c], [Aurell et al., 2022a], [Bayraktar et al., 2023], ...
- ...

For simplicity, in these lectures, we will mostly focus on “plain” MFGs, although many ideas can be extended.

Thank you for your attention

Questions?

Feel free to reach out: `mathieu.lauriere@nyu.edu`

References I

[Achdou et al., 2017] Achdou, Y., Bardi, M., and Cirant, M. (2017).
Mean field games models of segregation.
Mathematical Models and Methods in Applied Sciences, 27(01):75–113.

[Achdou and Kobeissi, 2020] Achdou, Y. and Kobeissi, Z. (2020).
Mean field games of controls: Finite difference approximations.
arXiv preprint arXiv:2003.03968.

[Achdou et al., 2020] Achdou, Y., Mannucci, P., Marchi, C., and Tchou, N. (2020).
Deterministic mean field games with control on the acceleration.
Nonlinear Differential Equations and Applications NoDEA, 27:1–32.

[Anahtarci et al., 2023] Anahtarci, B., Kariksiz, C. D., and Saldi, N. (2023).
Learning mean-field games with discounted and average costs.
Journal of Machine Learning Research, 24(17):1–59.

[Angiuli et al., 2022a] Angiuli, A., Detering, N., Fouque, J.-P., Lauriere, M., and Lin, J. (2022a).
Reinforcement learning algorithm for mixed mean field control games.
arXiv preprint arXiv:2205.02330.

[Angiuli et al., 2022b] Angiuli, A., Detering, N., Fouque, J.-P., Laurière, M., and Lin, J. (2022b).
Reinforcement learning for intra-and-inter-bank borrowing and lending mean field control game.
In *Proceedings of the Third ACM International Conference on AI in Finance*, pages 369–376.

References II

[Arapostathis et al., 2017] Arapostathis, A., Biswas, A., and Carroll, J. (2017).
On solutions of mean field games with ergodic cost.
Journal de Mathématiques Pures et Appliquées, 107(2):205–251.

[Aurell et al., 2022a] Aurell, A., Carmona, R., Dayanıklı, G., and Laurière, M. (2022a).
Finite state graphon games with applications to epidemics.
Dynamic Games and Applications, 12(1):49–81.

[Aurell et al., 2022b] Aurell, A., Carmona, R., Dayanikli, G., and Lauriere, M. (2022b).
Optimal incentives to mitigate epidemics: a stackelberg mean field game approach.
SIAM Journal on Control and Optimization, 60(2):S294–S322.

[Aurell et al., 2022c] Aurell, A., Carmona, R., and Lauriere, M. (2022c).
Stochastic graphon games: Ii. the linear-quadratic case.
Applied Mathematics & Optimization, 85(3):39.

[Bardi and Cardaliaguet, 2021] Bardi, M. and Cardaliaguet, P. (2021).
Convergence of some mean field games systems to aggregation and flocking models.
Nonlinear Analysis, 204:112199.

[Bardi and Priuli, 2014] Bardi, M. and Priuli, F. S. (2014).
Linear-quadratic n-person and mean-field games with ergodic cost.
SIAM Journal on Control and Optimization, 52(5):3022–3052.

[Barreiro-Gomez and Tembine, 2021] Barreiro-Gomez, J. and Tembine, H. (2021).
Mean-field-type Games for Engineers.
CRC Press.

[Bayraktar et al., 2021] Bayraktar, E., Cecchin, A., Cohen, A., and Delarue, F. (2021).
Finite state mean field games with wright–fisher common noise.
Journal de Mathématiques Pures et Appliquées, 147:98–162.

[Bayraktar et al., 2023] Bayraktar, E., Wu, R., and Zhang, X. (2023).
Propagation of chaos of forward–backward stochastic differential equations with graphon interactions.
Applied Mathematics & Optimization, 88(1):25.

[Bensoussan et al., 2015] Bensoussan, A., Chau, M. H., and Yam, S. C. P. (2015).
Mean field stackelberg games: Aggregation of delayed instructions.
SIAM Journal on Control and Optimization, 53(4):2237–2266.

[Bensoussan et al., 2018] Bensoussan, A., Huang, T., and Laurière, M. (2018).
Mean field control and mean field game models with several populations.
Minimax Theory and its Applications, 3(2):173–209.

[Bertucci et al., 2019] Bertucci, C., Lasry, J.-M., and Lions, P.-L. (2019).
Some remarks on mean field games.
Communications in Partial Differential Equations, 44(3):205–227.

References IV

[Bogachev et al., 2022] Bogachev, V. I., Krylov, N. V., Röckner, M., and Shaposhnikov, S. V. (2022).
Fokker–Planck–Kolmogorov Equations, volume 207.
American Mathematical Society.

[Caines and Huang, 2019] Caines, P. E. and Huang, M. (2019).
Graphon mean field games and the gmfg equations: ε-nash equilibria.
In *2019 IEEE 58th conference on decision and control (CDC)*, pages 286–292. IEEE.

[Caines and Huang, 2021] Caines, P. E. and Huang, M. (2021).
Graphon mean field games and their equations.
SIAM Journal on Control and Optimization, 59(6):4373–4399.

[Caines and Kizilkale, 2016] Caines, P. E. and Kizilkale, A. C. (2016).
ϵ-nash equilibria for partially observed lqg mean field games with a major player.
IEEE Transactions on Automatic Control, 62(7):3225–3234.

[Cardaliaguet et al., 2020] Cardaliaguet, P., Cirant, M., and Porretta, A. (2020).
Remarks on nash equilibria in mean field game models with a major player.
Proceedings of the American Mathematical Society, 148(10):4241–4255.

[Cardaliaguet et al., 2019] Cardaliaguet, P., Delarue, F., Lasry, J.-M., and Lions, P.-L. (2019).
The master equation and the convergence problem in mean field games, volume 201 of *Annals of Mathematics Studies*.
Princeton University Press, Princeton, NJ.

References V

[Cardaliaguet and Graber, 2015] Cardaliaguet, P. and Graber, P. J. (2015).
Mean field games systems of first order.
ESAIM Control Optim. Calc. Var., 21(3):690–722.

[Cardaliaguet et al., 2012] Cardaliaguet, P., Lasry, J.-M., Lions, P.-L., and Porretta, A. (2012).
Long time average of mean field games.
Networks & Heterogeneous Media, 7(2).

[Cardaliaguet and Lehalle, 2018] Cardaliaguet, P. and Lehalle, C.-A. (2018).
Mean field game of controls and an application to trade crowding.
Mathematics and Financial Economics, 12:335–363.

[Carmona et al., 2022a] Carmona, R., Cooney, D. B., Graves, C. V., and Lauriere, M. (2022a).
Stochastic graphon games: I. the static case.
Mathematics of Operations Research, 47(1):750–778.

[Carmona and Dayanıklı, 2021] Carmona, R. and Dayanıklı, G. (2021).
Mean field game model for an advertising competition in a duopoly.
International Game Theory Review, 23(04):2150024.

[Carmona et al., 2022b] Carmona, R., Dayanıklı, G., and Laurière, M. (2022b).
Mean field models to regulate carbon emissions in electricity production.
Dynamic Games and Applications, 12(3):897–928.

References VI

[Carmona and Delarue, 2018a] Carmona, R. and Delarue, F. (2018a).
Probabilistic theory of mean field games with applications. I, volume 83 of *Probability Theory and Stochastic Modelling*.
Springer, Cham.
Mean field FBSDEs, control, and games.

[Carmona and Delarue, 2018b] Carmona, R. and Delarue, F. (2018b).
Probabilistic theory of mean field games with applications. II, volume 84 of *Probability Theory and Stochastic Modelling*.
Springer, Cham.
Mean field games with common noise and master equations.

[Carmona et al., 2019] Carmona, R., Laurière, M., and Tan, Z. (2019).
Linear-quadratic mean-field reinforcement learning: convergence of policy gradient methods.
arXiv preprint arXiv:1910.04295.

[Carmona and Wang, 2017] Carmona, R. and Wang, P. (2017).
An alternative approach to mean field game with major and minor players, and applications to herders impacts.
Applied Mathematics & Optimization, 76:5–27.

[Carmona and Zhu, 2016] Carmona, R. and Zhu, X. (2016).
A probabilistic approach to mean field games with major and minor players.
Annals of applied probability: an official journal of the Institute of Mathematical Statistics, 26(3):1535–1580.

References VII

[Chen et al., 2022] Chen, Y., Zhang, L., Liu, J., and Hu, S. (2022).
Individual-level inverse reinforcement learning for mean field games.
In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 253–262.

[Cirant, 2015] Cirant, M. (2015).
Multi-population mean field games systems with neumann boundary conditions.
Journal de Mathématiques Pures et Appliquées, 103(5):1294–1315.

[Cosso and Pham, 2019] Cosso, A. and Pham, H. (2019).
Zero-sum stochastic differential games of generalized mckean–vlasov type.
Journal de Mathématiques Pures et Appliquées, 129:180–212.

[Cucker and Smale, 2007] Cucker, F. and Smale, S. (2007).
Emergent behavior in flocks.
IEEE Transactions on automatic control, 52(5):852–862.

[Dayanikli and Lauriere, 2023] Dayanikli, G. and Lauriere, M. (2023).
A machine learning method for stackelberg mean field games.
arXiv preprint arXiv:2302.10440.

[Delarue, 2021] Delarue, F. (2021).
Mean Field Games, volume 78.
American Mathematical Society.

[Djehiche et al., 2017] Djehiche, B., Tcheukam, A., and Tembine, H. (2017).
Mean-field-type games in engineering.
AIMS Electronics and Electrical Engineering, 1(1):18–73.

[Doncel et al., 2022] Doncel, J., Gast, N., and Gaujal, B. (2022).
A mean field game analysis of sir dynamics with vaccination.
Probability in the Engineering and Informational Sciences, 36(2):482–499.

[Elie et al., 2020a] Elie, R., Hubert, E., and Turinici, G. (2020a).
Contact rate epidemic control of covid-19: an equilibrium view.
Mathematical Modelling of Natural Phenomena, 15:35.

[Elie et al., 2019] Elie, R., Mastrolia, T., and Possamaï, D. (2019).
A tale of a principal and many, many agents.
Mathematics of Operations Research, 44(2):440–467.

[Elie et al., 2020b] Elie, R., Perolat, J., Laurière, M., Geist, M., and Pietquin, O. (2020b).
On the convergence of model free learning in mean field games.
In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7143–7150.

[Feleqi, 2013] Feleqi, E. (2013).
The derivation of ergodic mean field game equations for several populations of players.
Dynamic Games and Applications, 3:523–536.

References IX

[Firoozi et al., 2021] Firoozi, D., Shrivats, A. V., and Jaimungal, S. (2021).
Principal agent mean field games in rec markets.
arXiv preprint arXiv:2112.11963.

[Gao et al., 2020] Gao, S., Tchuendom, R. F., and Caines, P. E. (2020).
Linear quadratic graphon field games.
arXiv preprint arXiv:2006.03964.

[Geist et al., 2022] Geist, M., Pérolat, J., Laurière, M., Elie, R., Perrin, S., Bachem, O., Munos, R., and Pietquin, O. (2022).
Concave utility reinforcement learning: The mean-field game viewpoint.
In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 489–497.

[Gomes et al., 2010] Gomes, D. A., Mohr, J., and Souza, R. R. (2010).
Discrete time, finite state space mean field games.
Journal de mathématiques pures et appliquées, 93(3):308–328.

[Gomes et al., 2013] Gomes, D. A., Mohr, J., and Souza, R. R. (2013).
Continuous time finite state mean field games.
Applied Mathematics & Optimization, 68(1):99–143.

[Gomes et al., 2014] Gomes, D. A., Patrizi, S., and Voskanyan, V. (2014).
On the existence of classical solutions for stationary extended mean field games.
Nonlinear Analysis: Theory, Methods & Applications, 99:49–79.

[Gomes and Voskanyan, 2016] Gomes, D. A. and Voskanyan, V. K. (2016).
Extended deterministic mean-field games.
SIAM Journal on Control and Optimization, 54(2):1030–1055.

[Grover et al., 2018] Grover, P., Bakshi, K., and Theodorou, E. A. (2018).
A mean-field game model for homogeneous flocking.
Chaos: An Interdisciplinary Journal of Nonlinear Science, 28(6).

[Guo et al., 2019] Guo, X., Hu, A., Xu, R., and Zhang, J. (2019).
Learning mean-field games.
Advances in Neural Information Processing Systems, 32:4966–4976.

[Guo et al., 2022] Guo, X., Hu, A., and Zhang, J. (2022).
Optimization frameworks and sensitivity analysis of stackelberg mean-field games.
arXiv preprint arXiv:2210.04110.

[Hadikhanloo and Silva, 2019] Hadikhanloo, S. and Silva, F. J. (2019).
Finite mean field games: fictitious play and convergence to a first order continuous mean field game.
Journal de Mathématiques Pures et Appliquées, 132:369–397.

[Huang et al., 2006] Huang, M., Malhamé, R. P., Caines, P. E., et al. (2006).
Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle.
Communications in Information & Systems, 6(3):221–252.

References XI

[Hubert and Turinici, 2018] Hubert, E. and Turinici, G. (2018).
Nash-mfg equilibrium in a sir model with time dependent newborn vaccination.
Ricerche di matematica, 67:227–246.

[Kac, 1956] Kac, M. (1956).
Foundations of kinetic theory.
In *Proceedings of The third Berkeley symposium on mathematical statistics and probability*, volume 3, pages 171–197.

[Kobeissi, 2022] Kobeissi, Z. (2022).
Mean field games with monotonous interactions through the law of states and controls of the agents.
Nonlinear Differential Equations and Applications NoDEA, 29(5):52.

[Kolokoltsov and Bensoussan, 2016] Kolokoltsov, V. N. and Bensoussan, A. (2016).
Mean-field-game model for botnet defense in cyber-security.
Appl. Math. Optim., 74(3):669–692.

[Lacker, 2018] Lacker, D. (2018).
Mean field games and interacting particle systems.
preprint.

References XII

[Lacker and Soret, 2022] Lacker, D. and Soret, A. (2022).
A label-state formulation of stochastic graphon games and approximate equilibria on large networks.
Mathematics of Operations Research.

[Laguzet and Turinici, 2015] Laguzet, L. and Turinici, G. (2015).
Individual vaccination as nash equilibrium in a sir model with application to the 2009–2010 influenza a (h1n1) epidemic in france.
Bulletin of Mathematical Biology, 77:1955–1984.

[Lasry and Lions, 2007] Lasry, J.-M. and Lions, P.-L. (2007).
Mean field games.
Jpn. J. Math., 2(1):229–260.

[Lasry and Lions, 2018] Lasry, J.-M. and Lions, P.-L. (2018).
Mean-field games with a major player.
Comptes Rendus Mathematique, 356(8):886–890.

[Laurière et al., 2022a] Laurière, M., Perrin, S., Geist, M., and Pietquin, O. (2022a).
Learning mean field games: A survey.
arXiv preprint arXiv:2205.12944.

[Laurière et al., 2022b] Laurière, M., Perrin, S., Girgin, S., Muller, P., Jain, A., Cabannes, T., Piliouras, G., Pérolat, J., Elie, R., Pietquin, O., et al. (2022b).
Scalable deep reinforcement learning algorithms for mean field games.
In *International Conference on Machine Learning*, pages 12078–12095. PMLR.

[Laurière and Tangpi, 2022] Laurière, M. and Tangpi, L. (2022).
Convergence of large population games to mean field games with interaction through the controls.
SIAM Journal on Mathematical Analysis, 54(3):3535–3574.

[Lee et al., 2021] Lee, W., Liu, S., Tembine, H., Li, W., and Osher, S. (2021).
Controlling propagation of epidemics via mean-field control.
SIAM Journal on Applied Mathematics, 81(1):190–207.

[Miller and Pham, 2019] Miller, E. and Pham, H. (2019).
Linear-quadratic mckean-vlasov stochastic differential games.
Modeling, Stochastic Control, Optimization, and Applications, pages 451–481.

[Moon and Başar, 2018] Moon, J. and Başar, T. (2018).
Linear quadratic mean field stackelberg differential games.
Automatica, 97:200–213.

References XIV

[Muller et al., 2022a] Muller, P., Elie, R., Rowland, M., Lauriere, M., Perolat, J., Perrin, S., Geist, M., Piliouras, G., Pietquin, O., and Tuyls, K. (2022a).
Learning correlated equilibria in mean-field games.
arXiv preprint arXiv:2208.10138.

[Muller et al., 2022b] Muller, P., Rowland, M., Elie, R., Piliouras, G., Perolat, J., Lauriere, M., Marinier, R., Pietquin, O., and Tuyls, K. (2022b).
Learning equilibria in mean-field games: Introducing mean-field psro.
In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 926–934.

[Nash, 1951] Nash, J. (1951).
Non-cooperative games.
Annals of mathematics, pages 286–295.

[Nourian et al., 2011] Nourian, M., Caines, P. E., and Malhamé, R. P. (2011).
Mean field analysis of controlled cucker-smale type flocking: Linear analysis and perturbation equations.
IFAC Proceedings Volumes, 44(1):4471–4476.

[Olmez et al., 2022] Olmez, S. Y., Aggarwal, S., Kim, J. W., Miehling, E., Başar, T., West, M., and Mehta, P. G. (2022).
Modeling presymptomatic spread in epidemics via mean-field games.
In *2022 American Control Conference (ACC)*, pages 3648–3655. IEEE.

[Parise and Ozdaglar, 2019] Parise, F. and Ozdaglar, A. (2019).
Graphon games.
In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 457–458.

[Perrin et al., 2021] Perrin, S., Laurière, M., Pérolat, J., Geist, M., Élie, R., and Pietquin, O. (2021).
Mean field games flock! the reinforcement learning way.
In *IJCAI*.

[Santambrogio and Shim, 2021] Santambrogio, F. and Shim, W. (2021).
A cucker–smale inspired deterministic mean field game with velocity interactions.
SIAM Journal on Control and Optimization, 59(6):4155–4187.

[Subramanian and Mahajan, 2019] Subramanian, J. and Mahajan, A. (2019).
Reinforcement learning in stationary mean-field games.
In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 251–259.

[Sznitman, 1991] Sznitman, A.-S. (1991).
Topics in propagation of chaos.
In *Ecole d'été de probabilités de Saint-Flour XIX—1989*, pages 165–251. Springer.

[Vasal and Berry, 2022] Vasal, D. and Berry, R. (2022).
Master equation for discrete-time stackelberg mean field games with a single leader.
In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5529–5535. IEEE.

[Vasal et al., 2021] Vasal, D., Mishra, R., and Vishwanath, S. (2021).
Sequential decomposition of graphon mean field games.
In *2021 American Control Conference (ACC)*, pages 730–736. IEEE.