Thermodynamic limit and phase transitions in non-cooperative games: some mean-field examples

Paolo Dai Pra

Università di Padova

Padova, May 24, 2018

These researches have involved:

Alekos Cecchin, Markus Fischer, Guglielmo Pelino, Elena Sartori, Marco Tolotti.

Complex stochastic models motivated, or inspired, by Statistical Mechanics are often introduced and studied on the basis of the following program.

Complex stochastic models motivated, or inspired, by Statistical Mechanics are often introduced and studied on the basis of the following program.

- Define stochastic (Markovian) dynamics for a system of N interacting units $x_1^N(t), x_2^N(t), \ldots, x_N^N(t)$, modeling the microscopic evolution of the system at hand.

Complex stochastic models motivated, or inspired, by Statistical Mechanics are often introduced and studied on the basis of the following program.

- Define stochastic (Markovian) dynamics for a system of N interacting units $x_1^N(t), x_2^N(t), \ldots, x_N^N(t)$, modeling the microscopic evolution of the system at hand.
- Show that the limit as $N \to +\infty$ is well defined, producing the evolution $(x_i(t))_{i=1}^{+\infty}$ of countably many units. This should reveal the collective, or macroscopic, behavior of the system.

Complex stochastic models motivated, or inspired, by Statistical Mechanics are often introduced and studied on the basis of the following program.

- Define stochastic (Markovian) dynamics for a system of N interacting units $x_1^N(t), x_2^N(t), \ldots, x_N^N(t)$, modeling the microscopic evolution of the system at hand.
- Show that the limit as $N \to +\infty$ is well defined, producing the evolution $(x_i(t))_{i=1}^{+\infty}$ of countably many units. This should reveal the collective, or macroscopic, behavior of the system.
- Study the qualitative behavior (e.g. fixed points, attractors, stability...) of the limit ($N \to +\infty$) dynamics. Whenever this behavior has sudden changes as some parameter of the model crosses a critical value, we say there is a *phase transition*.

Example: the Ising model

Let (V_N, E_N) be a sequence of connected graphs, $|V_N| = N$, for which a limit as $N \to +\infty$ can be defined.

Example: the Ising model

Let (V_N, E_N) be a sequence of connected graphs, $|V_N| = N$, for which a limit as $N \to +\infty$ can be defined.

$x_i^N(t) \in \{-1, 1\}$, for $i \in V_N$, evolve according to the following rule:

the rate at which $x_i^N(t)$ switches to the opposite of its current value is

$$e^{-\beta x_i^N(t) m_N^i(t)}$$

Example: the Ising model

Let (V_N, E_N) be a sequence of connected graphs, $|V_N| = N$, for which a limit as $N \to +\infty$ can be defined.

$x_i^N(t) \in \{-1, 1\}$, for $i \in V_N$, evolve according to the following rule:

the rate at which $x_i^N(t)$ switches to the opposite of its current value is

$$e^{-\beta x_i^N(t) m_N^i(t)}$$

where

$$m_N^i(t) := \frac{1}{d(i)} \sum_{j \sim i} x_j^N(t),$$

$j \sim i$ indicates that j and i are neighbors, $d(i)$ is the number of neighbors (degree) of i, $\beta > 0$ is the parameter tuning the interaction (inverse temperature).

Example: the Ising model

The existence of the macroscopic limit can be proved for many graph sequences.

Example: the Ising model

The existence of the macroscopic limit can be proved for many graph sequences.

Moreover in many cases (e.g. $V = \mathbb{Z}^d$, $d \geq 2$) there exists $\beta_c > 0$ such that for $\beta < \beta_c$ the limit dynamics has a unique stationary distribution, while multiple stationary distributions emerge as $\beta > \beta_c$.

Example: the Ising model in the complete graph

A particularly simple example is that of the complete graph or the mean-field case: all pairs of vertices are connected, so

$$m_N^i = \frac{1}{N-1} \sum_{j \neq i} x_j^N(t).$$

Example: the Ising model in the complete graph

A particularly simple example is that of the complete graph or the mean-field case: all pairs of vertices are connected, so

$$m_N^i = \frac{1}{N-1} \sum_{j \neq i} x_j^N(t).$$

Assuming that the initial states $(x_i(0))_{i=1}^N$ are i.i.d., then the processes $x_i^N(t)$ converge, as $N \to +\infty$, to the i.i.d. processes $x_i(t)$ with switch rates

$$e^{-\beta x_i(t) m(t)},$$

Example: the Ising model in the complete graph

A particularly simple example is that of the complete graph or the mean-field case: all pairs of vertices are connected, so

$$m_N^i = \frac{1}{N-1} \sum_{j \neq i} x_j^N(t).$$

Assuming that the initial states $(x_i(0))_{i=1}^N$ are i.i.d., then the processes $x_i^N(t)$ converge, as $N \to +\infty$, to the i.i.d. processes $x_i(t)$ with switch rates

$$e^{-\beta x_i(t) m(t)},$$

where $m(t) := \mathbb{E}(x_i(t))$ can be computed by solving

$$\dot{m} = -2\sinh(\beta m) + 2m\sinh(\beta m).$$

$$\dot{m} = -2\sinh(\beta m) + 2m\sinh(\beta m).$$

$$\dot{m} = -2\sinh(\beta m) + 2m\sinh(\beta m).$$

Note that this last equation has $m = 0$ as global attractor for $\beta \leq 1$, which bifurcates into two stable attractors as $\beta > 1$.

$$\dot{m} = -2\sinh(\beta m) + 2m\sinh(\beta m).$$

Note that this last equation has $m = 0$ as global attractor for $\beta \leq 1$, which bifurcates into two stable attractors as $\beta > 1$.

The limit processes $x_i(t)$, $i \geq 1$, are independent: this property is called propagation of chaos.

In the example above, the interaction is coded in the jump rate of the state of each unit.

In the example above, the interaction is coded in the jump rate of the state of each unit.

It is expressed in terms of a given function of the state of the system.

In the example above, the interaction is coded in the jump rate of the state of each unit.

It is expressed in terms of a given function of the state of the system.

This is satisfactory for systems driven by fundamental physical laws. In various applications, however, e.g. in social sciences, interaction may be the result of optimization strategies, where each unit acts non-cooperatively to maximize her/his own interest.

Our program now reads as follows:

- Illustrate examples of N-players game for which the macroscopic limit $N \to +\infty$ can be dealt with;

Our program now reads as follows:

- Illustrate examples of N-players game for which the macroscopic limit $N \to +\infty$ can be dealt with;
- detect phase transitions in the macroscopic strategic behavior of the community.

Our program now reads as follows:

- Illustrate examples of N-players game for which the macroscopic limit $N \to +\infty$ can be dealt with;
- detect phase transitions in the macroscopic strategic behavior of the community.

The simplest case is the mean-field case, where any two players are neighbors.

Finite state mean field games

Mean Field Games were introduced as limit models for symmetric non-zero-sum non-cooperative N-player dynamic games when the number N of players tends to infinity.

(J.M. Lasry and P.L. Lions '06; M. Huang, R. P. Mallamé, P. E. Caines '06)

Finite state mean field games

Mean Field Games were introduced as limit models for symmetric non-zero-sum non-cooperative N-player dynamic games when the number N of players tends to infinity.

(J.M. Lasry and P.L. Lions '06; M. Huang, R. P. Mallamé, P. E. Caines '06)

Despite of the deep study of the limit model, the question of convergence of the N-player game to the corresponding mean-field game is still open to a large extent.

Finite state mean field games

Mean Field Games were introduced as limit models for symmetric non-zero-sum non-cooperative N-player dynamic games when the number N of players tends to infinity.

(J.M. Lasry and P.L. Lions '06; M. Huang, R. P. Mallamé, P. E. Caines '06)

Despite of the deep study of the limit model, the question of convergence of the N-player game to the corresponding mean-field game is still open to a large extent.

We have considered the problem in the case of finite state space.

Finite state mean field games

Let $\Sigma := \{1, 2, \dots, d\}$ and

$$\mathcal{P}(\Sigma) := \{m \in \mathbb{R}^d : \; m_j \geq 0, \; m_1 + m_2 + \cdots + m_d = 1\}$$

Finite state mean field games

Let $\Sigma := \{1, 2, \dots, d\}$ and

$$\mathcal{P}(\Sigma) := \{m \in \mathbb{R}^d : \ m_j \geq 0, \ m_1 + m_2 + \cdots + m_d = 1\}$$

N players control their (continuous-time) dynamics on Σ. We denote by $X_i(t)$ the state of the i-th player at time t, and $\mathbf{X}(t) = (X_1(t), X_2(t), \dots, X_N(t))$.

Finite state mean field games

Let $\Sigma := \{1, 2, \ldots, d\}$ and

$$\mathcal{P}(\Sigma) := \{m \in \mathbb{R}^d : \; m_j \geq 0, \; m_1 + m_2 + \cdots + m_d = 1\}$$

N players control their (continuous-time) dynamics on Σ. We denote by $X_i(t)$ the state of the i-th player at time t, and $\mathbf{X}(t) = (X_1(t), X_2(t), \ldots, X_N(t))$.

Each player i is allowed to control the rate $a^i_y(t)$ of jumping to y at time t. We restrict to feedback strategies:

$$a^i_y(t) = \alpha^i_y(t, \mathbf{X}(t)).$$

Finite state mean field games

Each player i aims at minimizing an index of the form

$$J_i^N(\alpha) := \mathbb{E}\left[\int_0^T \left(L(X_i(t), \alpha^i(t, \mathbf{X}(t))) + F(X_i(t), m_{\mathbf{X(t)}}^{N,i})\right) dt + G(X_i(T), m_{\mathbf{X(T)}}^{N,i})\right]$$

where for $\mathbf{x} \in \Sigma^N$,

$$m_{\mathbf{x}}^{N,i} = \frac{1}{N-1}\sum_{j \neq i} \delta_{x_j} \in \mathcal{P}(\Sigma).$$

Finite state mean field games

Each player i aims at minimizing an index of the form

$$J_i^N(\alpha) := \mathbb{E}\left[\int_0^T \left(L(X_i(t), \alpha^i(t, \mathbf{X}(t))) + F(X_i(t), m^{N,i}_{\mathbf{X(t)}})\right) dt + G(X_i(T), m^{N,i}_{\mathbf{X(T)}})\right]$$

where for $\mathbf{x} \in \Sigma^N$,

$$m^{N,i}_{\mathbf{x}} = \frac{1}{N-1}\sum_{j \neq i} \delta_{x_j} \in \mathcal{P}(\Sigma).$$

We work under the assumption that $L(x, \alpha)$ is smooth and uniformly convex in α, that guarantee existence and uniqueness of the Nash equilibrium.

Finite state mean field games

If we define the value function for player i

$$v^{N,i}(t,\mathbf{x}) := \mathbb{E}_{x,t}\left[\int_t^T \left(L(X_i(s),\overline{\alpha}^i(t,\mathbf{X}(s))) + F(X_i(s), m^{N,i}_{\mathbf{X(s)}})\right) dt + G(X_i(T), m^{N,i}_{\mathbf{X(T)}})\right]$$

it can be determined, as well as the control $\overline{\alpha}$ corresponding to the Nash equilibrium, by solving a system of differential equation, called Hamilton-Jacobi-Bellman (HJB) equation.

The mean-field game

We now let $N \to +\infty$ and consider a "representative" player i.

The mean-field game

We now let $N \to +\infty$ and consider a "representative" player i.

Heuristically,

$$m^{N,i}_{\mathbf{X(t)}} \to m_t \in \mathcal{P}(\Sigma)$$

where m_t is deterministic.

The mean-field game

We now let $N \to +\infty$ and consider a "representative" player i.

Heuristically,

$$m^{N,i}_{\mathbf{X(t)}} \to m_t \in \mathcal{P}(\Sigma)$$

where m_t is deterministic.

Then the player, whose state is denoted by $X(t)$, asymptotically aims at minimizing

$$J(\alpha) := \mathbb{E}\left[\int_0^T \left(L(X(t), \alpha(t, X(t))) + F(X(t), m_t)\right) dt + G(X(T), m_T)\right]$$

The mean-field game

We now let $N \to +\infty$ and consider a "representative" player i.

Heuristically,

$$m^{N,i}_{\mathbf{X(t)}} \to m_t \in \mathcal{P}(\Sigma)$$

where m_t is deterministic.

Then the player, whose state is denoted by $X(t)$, asymptotically aims at minimizing

$$J(\alpha) := \mathbb{E}\left[\int_0^T (L(X(t), \alpha(t, X(t))) + F(X(t), m_t))\, dt + G(X(T), m_T)\right]$$

where $\alpha_y(t, X(t))$ is the rate the player jumps to $y \in \Sigma$.

The mean-field game

We now let $N \to +\infty$ and consider a "representative" player i.

Heuristically,

$$m^{N,i}_{\mathbf{X(t)}} \to m_t \in \mathcal{P}(\Sigma)$$

where m_t is deterministic.

Then the player, whose state is denoted by $X(t)$, asymptotically aims at minimizing

$$J(\alpha) := \mathbb{E}\left[\int_0^T (L(X(t), \alpha(t, X(t))) + F(X(t), m_t))\, dt + G(X(T), m_T)\right]$$

where $\alpha_y(t, X(t))$ is the rate the player jumps to $y \in \Sigma$. Moreover, the optimal process $\overline{X}(t)$ must satisfy the consistency relation

$$m_t = Law(\overline{X}(t))$$

for $t \in [0, T]$.

The mean-field game equation

Solving the mean-field game amounts to:

The mean-field game equation

Solving the mean-field game amounts to:

- for a given flow $(m_t)_{t \in [0,T]}$, find the optimal strategy $\overline{\alpha}$;

The mean-field game equation

Solving the mean-field game amounts to:

- for a given flow $(m_t)_{t\in[0,T]}$, find the optimal strategy $\overline{\alpha}$;
- compute the flow $(\overline{m}_t)_{t\in[0,T]}$ of the marginal distributions of the optimal process;

The mean-field game equation

Solving the mean-field game amounts to:

- for a given flow $(m_t)_{t \in [0,T]}$, find the optimal strategy $\overline{\alpha}$;
- compute the flow $(\overline{m}_t)_{t \in [0,T]}$ of the marginal distributions of the optimal process;
- find $(m_t)_{t \in [0,T]}$ so that $m_t = \overline{m}_t$ for all $t \in [0, T]$.

The mean-field game equation

This can be translated into a pair of coupled equations:

The mean-field game equation

This can be translated into a pair of coupled equations:

an HJB equation

$$\begin{cases} -\dfrac{d}{dt}u(t,x) = -H(x,\nabla u(t,x)) + F(x,m_t) \\ \quad u(T,x) = G(x,m_T) \end{cases}$$

with

$$H(x,p) := \sup_{\alpha}\left[-\alpha \cdot p - L(x,\alpha)\right]$$

The mean-field game equation

This can be translated into a pair of coupled equations:

an HJB equation

$$\begin{cases} -\dfrac{d}{dt}u(t,x) = -H(x, \nabla u(t,x)) + F(x, m_t) \\ \quad u(T,x) = G(x, m_T) \end{cases}$$

with

$$H(x,p) := \sup_{\alpha} \left[-\alpha \cdot p - L(x,\alpha)\right] = -\alpha^*(x,p) \cdot p - L(x, \alpha^*(x,p))$$

The mean-field game equation

This can be translated into a pair of coupled equations:

an HJB equation

$$\begin{cases} -\dfrac{d}{dt}u(t,x) = -H(x, \nabla u(t,x)) + F(x, m_t) \\ u(T,x) = G(x, m_T) \end{cases}$$

with

$$H(x,p) := \sup_{\alpha} \left[-\alpha \cdot p - L(x,\alpha)\right] = -\alpha^*(x,p) \cdot p - L(x, \alpha^*(x,p))$$

and a Kolmogorov equation

$$\begin{cases} \dfrac{d}{dt}m_t(x) = \displaystyle\sum_y m_t(y)\alpha^*(y, \nabla u(t,y)) \\ m_0 = \mu_0 \end{cases}$$

The mean-field game equation

This can be translated into a pair of coupled equations:

an HJB equation

$$\begin{cases} -\dfrac{d}{dt}u(t,x) = -H(x, \nabla u(t,x)) + F(x, m_t) \\ u(T,x) = G(x, m_T) \end{cases}$$

with

$$H(x,p) := \sup_{\alpha} \left[-\alpha \cdot p - L(x,\alpha) \right] = -\alpha^*(x,p) \cdot p - L(x, \alpha^*(x,p))$$

and a Kolmogorov equation

$$\begin{cases} \dfrac{d}{dt}m_t(x) = \sum_y m_t(y)\alpha^*(y, \nabla u(t,y)) \\ m_0 = \mu_0 \end{cases}$$

Note: $(\nabla u(t,x))_y := u(t,y) - u(t,x)$.

The Master equation

Recently an alternative description of the mean-field game has been introduced.

The Master equation

Recently an alternative description of the mean-field game has been introduced.

We start from the ansatz that the value $v^{N,i}(t, \mathbf{x})$ of the game for player i in the N-player game is of the form

$$v^{N,i}(t, \mathbf{x}) = U^N(t, x_i, m_{\mathbf{x}}^{N,i}).$$

The Master equation

Recently an alternative description of the mean-field game has been introduced.

We start from the ansatz that the value $v^{N,i}(t, \mathbf{x})$ of the game for player i in the N-player game is of the form

$$v^{N,i}(t, \mathbf{x}) = U^N(t, x_i, m_{\mathbf{x}}^{N,i}).$$

Pretending U^N has a limit U as $N \to +\infty$, one derives an equation for U, called the Master equation:

$$\begin{cases} -\frac{d}{dt} U(t,x,m) = -H(x, \nabla_x U(t,x,m)) \\ \qquad\qquad + \sum_y m(y)\alpha^*(y, \nabla_y U(t,y,m)) \cdot \nabla_m U(t,y,m) F(x,m) \\ U(T,x,m) = G(x,m) \end{cases}$$

Convergence of the N-players game

But in which sense the N-players game converges to the mean-field game?

Convergence of the N-players game

But in which sense the N-players game converges to the mean-field game?

Note that standard conditions that guarantee uniqueness on the Nash equilibrium of the N-players game, do not in general imply uniqueness of the mean-field game equation.

Convergence of the N-players game

But in which sense the N-players game converges to the mean-field game?

Note that standard conditions that guarantee uniqueness on the Nash equilibrium of the N-players game, do not in general imply uniqueness of the mean-field game equation.

The following remarkable result has been proved by A. Cecchin and G. Pelino (see P. Cardaliaguet, F. Delarue, J.M. Lasry, P.L. Lions '15 for results in the continuous setting).

Convergence of the N-players game

Theorem 1.

Suppose the Master equation has a classical (smooth) solution U. Then

Convergence of the N-players game

Theorem 1.

Suppose the Master equation has a classical (smooth) solution U. Then

- *This solution is unique.*

Convergence of the N-players game

Theorem 1.

Suppose the Master equation has a classical (smooth) solution U. Then

- *This solution is unique.*
- *The mean-field equation has a unique solution.*

Convergence of the N-players game

Theorem 1.

Suppose the Master equation has a classical (smooth) solution U. Then

- *This solution is unique.*
- *The mean-field equation has a unique solution.*
- *Let $(\mathbf{X}^{\mathbf{N},*}(t))_{t\in[0,T]}$ be the dynamics of the N-players game corresponding to the unique Nash equilibrium, and $(X^*(t))_{t\in[0,T]}$ be the optimal process in the mean-field game. Then for each i*

$$X_i^{N,*} \to X^*$$

in distribution.

Convergence of the N-players game

Theorem 1.

Suppose the Master equation has a classical (smooth) solution U. Then

- *This solution is unique.*
- *The mean-field equation has a unique solution.*
- *Let $(\mathbf{X}^{\mathbf{N},*}(t))_{t\in[0,T]}$ be the dynamics of the N-players game corresponding to the unique Nash equilibrium, and $(X^*(t))_{t\in[0,T]}$ be the optimal process in the mean-field game. Then for each i*

$$X_i^{N,*} \to X^*$$

  in distribution.
- *Propagation of chaos holds: for $i \neq j$, $X_i^{N,*}$ and $X_j^{N,*}$ converge to independent copies of X^*.*

Convergence of the N-players game

Sufficient conditions for having a classical solution of the Master Equation are known. It is the case if the cost functions $F(x, m)$ and $G(x, m)$ satisfy the following monotonicity condition: for every $m, m' \in \mathcal{P}(\Sigma)$

$$\sum_x [F(x, m) - F(x, m')][m(x) - m'(x)] \geq 0$$

and the same for G.

Convergence of the N-players game

Sufficient conditions for having a classical solution of the Master Equation are known. It is the case if the cost functions $F(x, m)$ and $G(x, m)$ satisfy the following monotonicity condition: for every $m, m' \in \mathcal{P}(\Sigma)$

$$\sum_x [F(x, m) - F(x, m')][m(x) - m'(x)] \geq 0$$

and the same for G.

But what happens otherwise, in particular when the mean-field equation has multiple solutions?

A non-uniqueness example

Assume: $\Sigma = \{-1, 1\}$. So $m \in \mathcal{P}(\Sigma)$ can be identified with its mean.

$$L(x, \alpha) = \frac{\alpha^2}{2}, \quad F(x, m) \equiv 0, \quad G(x, m) = -xm.$$

So each player, at time T, aims at aligning with the others.

A non-uniqueness example

The mean-field game equation

$$\begin{cases} \frac{d}{dt}u(x,t) = H(\nabla u(x,t)) \\ \frac{d}{dt}m_t = -m_t\,|\nabla u(1,t)| + \nabla u(1,t) \\ u(x,T) = -x m_T \\ m_0 \ \text{given} \end{cases}$$

has:

- a unique solution for $T < T(m_0)$;
- three solutions for $T > T(m_0)$,

A non-uniqueness example

The mean-field game equation

$$\begin{cases} \frac{d}{dt}u(x,t) = H(\nabla u(x,t)) \\ \frac{d}{dt}m_t = -m_t\,|\nabla u(1,t)| + \nabla u(1,t) \\ u(x,T) = -x m_T \\ m_0 \ \text{given} \end{cases}$$

has:

- a unique solution for $T < T(m_0)$;
- three solutions for $T > T(m_0)$,

with $T(m_0)$ given implicitly by

$$|m_0| = \frac{(2T-1)^2(T+4)}{27\,T}.$$

A non-uniqueness example

The master equation for $U(t,x,m)$ is more conveniently written in terms of

$$Z(t,m) = \nabla U(t,1,m),$$

and reads

$$\begin{cases} \dfrac{d}{dt} Z = -\dfrac{d}{dm}\left(m\dfrac{Z|Z|}{2} - \dfrac{Z^2}{2} \right) \\ Z(T,m) = 2m \end{cases}$$

A non-uniqueness example

The master equation for $U(t, x, m)$ is more conveniently written in terms of

$$Z(t, m) = \nabla U(t, 1, m),$$

and reads

$$\begin{cases} \dfrac{d}{dt} Z = -\dfrac{d}{dm}\left(m\dfrac{Z|Z|}{2} - \dfrac{Z^2}{2} \right) \\ Z(T, m) = 2m \end{cases}$$

This is a scalar conservation law.

A non-uniqueness example

Theorem 2.

For $T > \frac{1}{2}$ the master equation has no classical solution. It has many solutions in the weak sense, but a unique entropy solution $Z(t, m)$, i.e. such that

- *it is a classical solution whenever it is continuous;*
- *if $Z(t, m)$ is discontinuous in m^* then*

$$-\lim_{m \uparrow m^*} Z(t, m) = \lim_{m \downarrow m^*} Z(t, m) > 0.$$

A non-uniqueness example

Theorem 2.

For $T > \frac{1}{2}$ the master equation has no classical solution. It has many solutions in the weak sense, but a unique entropy solution $Z(t,m)$, i.e. such that

- *it is a classical solution whenever it is continuous;*
- *if $Z(t,m)$ is discontinuous in m^* then*

$$-\lim_{m \uparrow m^*} Z(t,m) = \lim_{m \downarrow m^*} Z(t,m) > 0.$$

The entropy solution of the master equation corresponds to one particular solution of the mean-field game equation for $m_0 \neq 0$, and to a randomization of two solutions for $m_0 = 0$. We call this the entropy solution of the mean-field game equation.

A non-uniqueness example

Theorem 3.

(Cecchin, D.P., Fischer, Pelino) Let $(\mathbf{X}^{\mathbf{N},*}(t))_{t\in[0,T]}$ *be the dynamics of the N-players game corresponding to the unique Nash equilibrium, and* $(X^*(t))_{t\in[0,T]}$ *be the optimal process in the mean-field game corresponding to the entropy solution. Then for each i*

$$X_i^{N,*} \to X^*$$

in distribution. Moreover, propagation of chaos holds if $m_0 \neq 0$ *and does not hold for* $m_0 = 0$.

A non-uniqueness example

Thus, the optimal process for N players "selects", in the limit $N \to +\infty$, one particular solution of the Mean Field Game equation.

A non-uniqueness example

Thus, the optimal process for N players "selects", in the limit $N \to +\infty$, one particular solution of the Mean Field Game equation.

What is the meaning of the other solutions?

A non-uniqueness example

We recall that, for a N player game, a strategy is ϵ-Nash if any player, changing its own strategy letting unchanged that of the others, can improve his index by at most ϵ.

A non-uniqueness example

We recall that, for a N player game, a strategy is ϵ-Nash if any player, changing its own strategy letting unchanged that of the others, can improve his index by at most ϵ.

Theorem 4.

(Cecchin, Fisher) Let α^* be the optimal control for the mean-field game corresponding to any solution of the mean-field game equation. The same feedback, if used by all players in the N-player game, is ϵ_N-Nash, where

$$\epsilon_N \to 0 \quad \textit{as} \quad N \to +\infty.$$

An example in infinite horizon

For games with infinite time horizon the behavior may become even richer.

An example in infinite horizon

For games with infinite time horizon the behavior may become even richer.

Consider, as before, models on $\Sigma := \{-1, 1\}$ in which player control their jump rates $a_y^i(t) = \alpha^i(t, \mathbf{X}(t)))$, with cost function:

$$J_i^N(\alpha) := \mathbb{E}\left[\int_0^{+\infty} e^{-\lambda t} L\left(X_i(t), \alpha^i(t, \mathbf{X}(t)), m_{\mathbf{X(t)}}^{N,i}\right) dt\right]$$

where

$$L(x, \alpha, m) = \frac{1}{\mu(1 + \varepsilon x m)}\alpha^2 - xm.$$

with $\lambda, \mu > 0, \varepsilon \in (0, 1]$.

For λ, ε fixed we observe three regimes for the mean-field game (D.P., Sartori, Tolotti '18):

For λ, ε fixed we observe three regimes for the mean-field game (D.P., Sartori, Tolotti '18):

- μ small (low mobility): the equilibrium control is unique, and leads to consensus: $\lim_{t \to +\infty} m(t) \in \{-1, 1\}$.

For λ, ε fixed we observe three regimes for the mean-field game (D.P., Sartori, Tolotti '18):

- μ small (low mobility): the equilibrium control is unique, and leads to consensus: $\lim_{t \to +\infty} m(t) \in \{-1, 1\}$.
- μ intermediate (moderate mobility): for some initial conditions there may be many equilibrium controls, but they all lead to consensus.

For λ, ε fixed we observe three regimes for the mean-field game (D.P., Sartori, Tolotti '18):

- μ small (low mobility): the equilibrium control is unique, and leads to consensus: $\lim_{t \to +\infty} m(t) \in \{-1, 1\}$.
- μ intermediate (moderate mobility): for some initial conditions there may be many equilibrium controls, but they all lead to consensus.
- μ large (high mobility): there are equilibrium controls leading to periodic behavior of $m(t)$, in particular consensus is not obtained.

For λ, ε fixed we observe three regimes for the mean-field game (D.P., Sartori, Tolotti '18):

- μ small (low mobility): the equilibrium control is unique, and leads to consensus: $\lim_{t \to +\infty} m(t) \in \{-1, 1\}$.
- μ intermediate (moderate mobility): for some initial conditions there may be many equilibrium controls, but they all lead to consensus.
- μ large (high mobility): there are equilibrium controls leading to periodic behavior of $m(t)$, in particular consensus is not obtained.

Convergence of the N-players game to the mean-field game is still open.

Thanks