Long time behavior of finite state (potential) mean-field games via Γ-convergence

Rita Ferreira

Pos-Doc @ Instituto Superior Técnico
&
Centro de Matemática e Aplicações (FCT-UNL)

KAUST, February 12, 2014

Abstract

In this talk we address the study of the long time behavior of continuous time, finite state (potential) mean-field games using Γ-convergence.

The method we propose allows extending existent results in literature as it requires weaker hypotheses than the ones usually assumed.

In collaboration with Diogo Gomes (KAUST)

Problem set-up

- very large number of indistinguishable players
- players can be in a finite number of states, $I_d = \{1, 2, 3, ..., d\}$
- distribution of the players among the different states is given by a probability vector $\theta \in \mathcal{P}(I_d)$, where

$$\begin{cases} \theta^1 + ... + \theta^d = 1\,, \\ \theta^i \geq 0 \quad \forall i \in I_d. \end{cases}$$

- each player only knows its state and the probability θ
- states evolve randomly in time by following a controlled continuous time Markov chain and each player controls its switching rate in order to optimize a certain functional:

running cost c + terminal cost ψ

- Running cost:

$$c : I_d \times \mathcal{P}(I_d) \times (\mathbb{R}_0^+)^d \to \mathbb{R}$$
$$(i, \theta, \alpha) \mapsto c(i, \theta, \alpha)$$

  - i ... state of the player
  - θ ... probability distribution of the population among states
  - α_j ... transition rate the player uses to change from state i to state j

- Terminal Cost:

$$\psi : I_d \to \mathbb{R}$$

- Assume $\alpha \mapsto c(i,\theta,\alpha)$ convex, not depending on α_i
- Let $\Delta_i : \mathbb{R}^d \to \mathbb{R}^d$ denote the difference operator with respect to i:

$$\Delta_i z := (z^1 - z^i, \cdots, z^d - z^i), \quad \text{for all } z = (z^1, \ldots, z^d) \in \mathbb{R}^d.$$

- We define the generalized Legendre transform of the function $c(i,\theta,\cdot)$ as

$$\begin{aligned} h(z,\theta,i) \; &:= \min_{\mu \in (\mathbb{R}_0^+)^d} \left\{ c(i,\theta,\mu) + \sum_{j=1}^{d} \mu_j (z^j - z^i) \right\} \\ &= \min_{\mu \in (\mathbb{R}_0^+)^d} \left\{ c(i,\theta,\mu) + \mu \cdot \Delta_i z \right\}, \end{aligned}$$

which defines a concave function in z

Mean-Field Games

The value function is the unique solution of Hamilton-Jacobi ODE

$$-\dot{v}^i = h(\Delta_i v, \theta, i),$$

satisfying $v(T) = \psi$.

The optimal switching policy for a player in state i is

$$\alpha_i^*(\Delta_j v, \theta, j) = \frac{\partial}{\partial z^j}\big(h(\Delta_i v, \theta, i)\big).$$

The mean-field equilibrium arises when all players use the same optimal switching policy, which gives rise to the system:

$$\begin{cases} \dot{\theta}^i = \sum_{j=1}^{d} \theta^j \alpha_i^*(\Delta_j v, \theta, j), \\ -\dot{v}^i = h(\Delta_i v, \theta, i), \end{cases}$$

together with the initial-terminal conditions

$$\theta(0) = \theta_0, \qquad v(T) = \psi,$$

where θ_0 is the initial distribution of players.

- v is the value function and θ is the probability distribution of the population among states;
- h is known (it is related to the running cost function c) and is concave in the first variable;
- α_i^* is known (provides optimal switching policy for a player in state i);
- $\Delta_i z := (z^1 - z^i, \cdots, z^d - z^i)$ for $z = (z^1, \ldots, z^d) \in \mathbb{R}^d$.

Remark

From the ODE viewpoint these equations are non-standard as some of the variables have initial conditions whereas others have prescribed terminal data.

Study undertaken in [GMS2013] D. Gomes, J. Mohr, and R. Souza, *Continuous time finite state mean field games*, Appl Math Optim (2013) 68:99–143.

Potential Mean-Field Games

Corresponds to the case in which h has the form

$$h(z, \theta, i) = \tilde{h}(z, i) + f(\theta, i),$$

where $\frac{\partial F}{\partial \theta_i}(\cdot) = f(\cdot, i)$ for some convex function $F : \mathbb{R}^d \to \mathbb{R}$.

If F strictly convex, superlinear growth at infinity, and non-increasing in each coordinate, then (v, ϑ) is a solution of the MFG iff

- v is a critical point of the functional

  $$\int_0^T F^*(\dot{v} + \tilde{h}(\Delta.v, \cdot))\,\mathrm{d}t - \theta_0 \cdot v(0),$$

  where we are looking for critical points v that have fixed boundary condition at T, namely $v(T) = \psi$

- $\vartheta(t) = -\nabla F^*(\dot{v}(t) + \tilde{h}(\Delta.v(t), \cdot))$.

The stationary setting

Goal:

Study of long time convergence (trend to equilibrium problem) for finite state mean-field games

A triplet $(\bar{\theta}, \bar{v}, \bar{\lambda}) \in \mathcal{P}(I_d) \times \mathbb{R}^d \times \mathbb{R}$ is called a stationary solution of the mean-field equations if

$$\begin{cases} \sum_{j=1}^{d} \bar{\theta}^j \alpha_i^*(\Delta_j \bar{v}, \bar{\theta}, j) = 0\,, \\ h(\Delta_i \bar{v}, \bar{\theta}, i) = \bar{\lambda}\,, \end{cases}$$

for all $i \in I_d$.

If $(\bar{\theta}, \bar{v}, \bar{\lambda})$ is a stationary solution for the MFG equations, then $(\bar{\theta}, \bar{v} - \bar{\lambda} t \mathbf{1})$, where $\mathbf{1} := (1, ..., 1)$, solves the MFG equations.

Consider

$$\min\left\{\int_0^1 F^*(\tilde{h}(\Delta.v,\cdot)-\lambda\mathbf{1})\,\mathrm{d}t-\lambda\colon\right.$$

$$\left. v:[0,1]\to\mathbb{R}^d \text{ continous}, \sum_{i=1}^d v_i(t)=0,\ \ \lambda\in\mathbb{R}\right\}.$$

Jensen's inequality + F^* convex, componentwise non-increasing function + $\tilde{h}(\cdot,i)$ concave function $\Rightarrow$ it suffices to consider minimizers to this problem in the class of constant functions v. Therefore it is enough to look at minimizers of

$$\min\left\{F^*(\tilde{h}(\Delta.v,\cdot)-\lambda\mathbf{1})-\lambda\colon\ v\in\mathbb{R}^d,\ \sum_{i=1}^d v^i=0,\ \ \lambda\in\mathbb{R}\right\}.$$

In turn, if $(\bar{v},\bar{\lambda})$ is a critical point for the latter, then setting

$$\bar{\theta}^j := -\frac{\partial F^*}{\partial p_j}(\tilde{h}(\Delta.\bar{v},\cdot)-\lambda\mathbf{1}),\quad j\in I_d,$$

we conclude that $(\bar{\theta},\bar{v},\bar{\lambda})$ is a stationary solution of the mean-field eqns.

An important estimate

In $\mathbb{R}^d/\mathbb{R}$ we define the norm

$$\|z\|_\sharp = \frac{\max_{i\in I_d} z^i - \min_{i\in I_d} z^i}{2}, \quad z \in \mathbb{R}^d.$$

Definition

Let $\langle v \rangle := \frac{1}{d}\sum_{j=1}^d v^j$. We say that h is contractive if there exists $M > 0$ such that if $\|v\|_\sharp > M$, then the two following conditions hold for all $\theta \in \mathcal{P}(I_d)$ and $i \in I_d$:

$$(\Delta_i v)^j \leq 0 \text{ for all } j \in I_d \text{ implies } h(\Delta_i v, \theta, i) - \langle h(v, \theta, \cdot)\rangle < 0\,,$$
$$(\Delta_i v)^j \geq 0 \text{ for all } j \in I_d \text{ implies } h(\Delta_i v, \theta, i) - \langle h(v, \theta, \cdot)\rangle > 0\,.$$

Many mean-field games are contractive [GMS2013].

Moreover, for contractive, potential mean-field games:

If $\|v(T)\|_\sharp \leqslant M$, where v is a solution, and M is large enough, then $\|v(t)\|_\sharp \leqslant M$ for all $t \in [0, T]$.

Scaling - study of long time behavior of mean-field games

We introduce a scaled version of the mean-field game, where $\epsilon = \frac{1}{T}$,

$$\begin{cases} \epsilon\dot{\theta}^i_\epsilon = \displaystyle\sum_{j=1}^{d} \theta^j_\epsilon \alpha^*_i(\Delta_j v_\epsilon, \theta_\epsilon, j), \\ -\epsilon\dot{v}^i_\epsilon = h(\Delta_i v_\epsilon, \theta_\epsilon, i), \end{cases}$$

together with the initial-terminal conditions

$$\theta_\epsilon(0) = \theta_0, \qquad v^i_\epsilon(1) = \psi^i.$$

Wlog, $\sum_{i=1}^d \psi^i = 0$.

Important remark:

Scaling in time does not change the bounds mentioned above. Hence, for contractive, potential mean-field games:

$$\sup_{t\in[0,1]} \sup_{\epsilon>0} \|v_\epsilon(t)\|_\sharp < +\infty.$$

In order to rewrite the associated functional in a in a convenient form for the use of Γ-convergence we decompose v_ϵ as follows.

Let $\lambda_\epsilon \in \mathbb{R}$, $u_\epsilon : [0,1] \to \mathbb{R}^d$, and $w_\epsilon : [0,1] \to \mathbb{R}$ be defined by

$$\lambda_\epsilon := \int_0^1 \sum_{i=1}^d h(\Delta_i v_\epsilon, \theta_\epsilon, i)\,\mathrm{d}t\,,$$

$$w_\epsilon(t) := \frac{\epsilon}{d} \sum_{i=1}^d v_\epsilon^i(t) - \lambda_\epsilon(1-t)\,,$$

$$u_\epsilon^i(t) := v_\epsilon^i(t) - \frac{1}{\epsilon} w_\epsilon(t) - \frac{\lambda_\epsilon}{\epsilon}(1-t).$$

Observing that $\Delta_i u_\epsilon = \Delta_i v_\epsilon$ for all $i \in I_d$, the (scaled) mean-field equations become

$$\begin{cases} \epsilon \dot\theta^i_\epsilon = \sum_{j=1}^{d} \theta^j_\epsilon \alpha^*_i(\Delta_j u_\epsilon, \theta_\epsilon, j), \\ \lambda_\epsilon - \dot w_\epsilon - \epsilon \dot u^i_\epsilon = h(\Delta_i u_\epsilon, \theta_\epsilon, i). \end{cases}$$

Moreover, for all $t \in [0,1]$, $\varepsilon > 0$, and $i \in I_d$,

$$\sup_{\varepsilon>0} |\lambda_\varepsilon| < +\infty, \qquad \sup_{t\in[0,1]} \sup_{\epsilon>0} \|u_\epsilon(t)\|_\sharp < +\infty,$$

$$\sum_{i=1}^{d} u^i_\epsilon(t) = 0, \quad u^i_\epsilon(1) = \psi^i, \quad \sup_{\epsilon>0} \|\epsilon \dot u_\epsilon\|_\infty < +\infty,$$

$$w_\epsilon(0) = 0, \;\; w_\epsilon(1) = 0, \quad \sup_{\epsilon>0} \|\dot w_\epsilon\|_\infty < +\infty.$$

From the variational point of view, we look for minimizers of

$$\int_0^1 F^*(\dot{w}\mathbf{1} + \epsilon\dot{u} + \tilde{h}(\Delta.u, \cdot) - \lambda\mathbf{1})\,\mathrm{d}t - \epsilon\theta_0 \cdot u(0) - \lambda$$

over $\lambda \in \mathbb{R}$, $u : [0,1] \to \mathbb{R}^d$, and $w : [0,1] \to \mathbb{R}$ according to the conditions above.

Expect to obtain in the limit:

$$\int_0^1 F^*(\dot{w}\mathbf{1} + \tilde{h}(\Delta.u, \cdot) - \lambda\mathbf{1})\,\mathrm{d}t - \lambda,$$

which corresponds to the *stationary functional* introduced before provided that w does not depend on t.

A Γ-convergence result

- X be a reflexive Banach space endowed with its weak topology
- $\mathcal{F}_n : X \to \overline{\mathbb{R}}$ equi-coercive in the weak topology of X
- exists $\mathcal{F} : X \to \overline{\mathbb{R}}$ such that
    - for every $x \in X$ and for every sequence $\{x_n\}_{n\in\mathbb{N}}$ weakly converging to x in X, one has $\mathcal{F}(x) \leqslant \liminf_{n\to+\infty} \mathcal{F}_n(x_n)$;
    - or every $x \in X$ there exists a sequence $\{x_n\}_{n\in\mathbb{N}}$ weakly converging to x in X such that $\mathcal{F}(x) = \lim_{n\to+\infty} \mathcal{F}_n(x_n)$.

Then

- $\min_{x\in X} \mathcal{F}(x) = \lim_{n\to+\infty} \inf_{x\in X} \mathcal{F}_n(x)$;
- if x_n is a δ_n-minimizer of $\mathcal{F}_n$ in X, where $\delta_n \searrow 0^+$, and x is a cluster point of $\{x_n\}_{n\in\mathbb{N}}$, then x is a minimizer of $\mathcal{F}$ in X, and $\mathcal{F}(x) = \limsup_{n\to+\infty} \mathcal{F}_n(x_n)$;
- if $\{x_n\}_{n\in\mathbb{N}}$ weakly converges to x in X, then x is a minimizer of $\mathcal{F}$ in X, and $\mathcal{F}(x) = \lim_{n\to+\infty} \mathcal{F}_n(x_n)$.

Convergence of functionals, and its minima, associated with mean-field games

The space of continuous functions is not the most appropriate one for this study as it is not a reflexive Banach space.

We then extend the functionals to the product space $L^p \times W_0^{1,p}$, for $p \in (1, +\infty)$, in a natural way.

Assumptions

- $F^* : \mathbb{R}^d \to \mathbb{R}$ is componentwise non-incresing and convex
- $\tilde{h} : \mathbb{R}^d \times I_d \to \mathbb{R}$ is concave in the first variable
- $\tilde{\mathbf{h}} : \mathbb{R}^d \to \mathbb{R}^d$ is the function defined by

$$\tilde{\mathbf{h}}(z) := (\tilde{h}^1(z), \cdots, \tilde{h}^d(z)), \quad \text{with } \tilde{h}^i(z) := \tilde{h}(\Delta_i z, i) \text{ for all } i \in I_d,$$

$\mathcal{F}_\varepsilon : L^p((0,1);\mathbb{R}^d) \times W_0^{1,p}(0,1) \times \mathbb{R} \to \overline{\mathbb{R}}$,

$$\mathcal{F}_\varepsilon(u,w,\lambda) := \begin{cases} \mathcal{G}_\varepsilon(u,w,\lambda) & \text{if } (u,w,\lambda) \in \Phi_\varepsilon, \\ +\infty & \text{otherwise,} \end{cases}$$

$$\mathcal{G}_\varepsilon(u,w,\lambda) := \int_0^1 F^*\big(\dot{w}(t)\mathbf{1} + \varepsilon\dot{u}(t) + \tilde{\mathbf{h}}(u(t)) - \lambda\mathbf{1}\big)\,\mathrm{d}t - \varepsilon\theta_0 \cdot u(0) - \lambda,$$

$$\Phi_\varepsilon := \Bigg\{ (u,w,\lambda) \in W^{1,p}((0,1);\mathbb{R}^d) \times W_0^{1,p}(0,1) \times \mathbb{R} :$$

$$u(1) = \psi,\ \|u(\cdot)\|_\sharp \leqslant \bar{M}_0 \text{ and } \sum_{i=1}^d u^i(\cdot) = 0 \ \mathcal{L}^1\text{-a.e. in } (0,1),$$

$$\max\left\{ \int_0^1 |\varepsilon\dot{u}(t)|^p\,\mathrm{d}t, \|\dot{w}\|_{L^\infty(0,1)} \right\} \leqslant M_0,\ |\lambda| \leqslant R_0 \Bigg\}$$

(some M_0, $\bar{M}_0$, $R_0 \in \mathbb{R}$.)

$$\mathcal{F}_0 : L^p((0,1);\mathbb{R}^d) \times W_0^{1,p}(0,1) \times \mathbb{R} \to \overline{\mathbb{R}},$$

$$\mathcal{F}_0(u,w,\lambda) := \begin{cases} \mathcal{G}_0(u,w,\lambda) & \text{if } (u,w,\lambda) \in \Phi_0, \\ +\infty & \text{otherwise,} \end{cases}$$

$$\mathcal{G}_0(u,w,\lambda) := \int_0^1 F^*\big(\dot{w}(t)\mathbf{1} + \tilde{\mathbf{h}}(u(t)) - \lambda\mathbf{1}\big)\,\mathrm{d}t - \lambda,$$

$$\Phi_0 := \Bigg\{ (u,w,\lambda) \in L^p((0,1);\mathbb{R}^d) \times W_0^{1,p}(0,1) \times \mathbb{R} :$$

$$\|u(\cdot)\|_\sharp \leqslant \bar{M}_0 \text{ and } \sum_{i=1}^d u^i(\cdot) = 0 \ \mathcal{L}^1\text{-a.e. in } (0,1),$$

$$\|\dot{w}\|_{L^\infty(0,1)} \leqslant M_0,\ |\lambda| \leqslant R_0 \Bigg\}.$$

Main Theorem

- $\{\mathcal{F}_\varepsilon\}_{\varepsilon>0}$ Γ-converges as $\varepsilon \to 0^+$ to $\mathcal{F}_0$ with respect to the weak convergence in $L^p((0,1);\mathbb{R}^d) \times W_0^{1,p}(0,1) \times \mathbb{R}$.

- Let $(u_\varepsilon, w_\varepsilon, \lambda_\varepsilon)$ be a δ_ε-minimizer, where $\delta_\varepsilon \searrow 0^+$, of $\mathcal{G}_\varepsilon$ in Φ_ε. Then $\{(u_\varepsilon, w_\varepsilon, \lambda_\varepsilon)\}_{\varepsilon>0}$ is bounded in norm; if (u, w, λ) is a cluster point, then (u, w, λ) is a minimizer of $\mathcal{G}_0$ in Φ_0 and

$$\mathcal{G}_0(u, w, \lambda) = \limsup_{\varepsilon \to 0^+} \mathcal{G}_\varepsilon(u_\varepsilon, w_\varepsilon, \lambda_\varepsilon).$$

- If $\{(u_\varepsilon, w_\varepsilon, \lambda_\varepsilon)\}_{\varepsilon>0}$ weakly converges to (u, w, λ) in $L^p((0,1);\mathbb{R}^d) \times W_0^{1,p}(0,1) \times \mathbb{R}$, then (u, w, λ) is a minimizer of $\mathcal{G}_0$ in Φ_0, and

$$\mathcal{G}_0(u, w, \lambda) = \lim_{\varepsilon \to 0^+} \mathcal{G}_\varepsilon(u_\varepsilon, w_\varepsilon, \lambda_\varepsilon).$$

Furthermore,

Assume that F^* is strictly convex and that $\tilde{h}$ is strictly concave in $\mathbb{R}^d/\mathbb{R}$, that is, for all $0 < \mu < 1$ we have

$$\tilde{h}(\mu\Delta_i u + (1-\mu)\Delta_i v, i) = \mu\tilde{h}(\Delta_i u, i) + (1-\mu)\tilde{h}(\Delta_i v, i)$$

implies $u = v + k\mathbf{1}$, for some $k \in \mathbb{R}$. Using Jensen's inequality, we conclude that a solution $(u, w, \lambda) \in \Phi_0$ to

$$\min\left\{\mathcal{G}_0(u, w, \lambda)\colon\ (u, w, \lambda) \in \Phi_0\right\}$$

is such that (w, u) does not depend on time. Thus, in this setting, we establish in addition convergence of solutions of the mean-field game to stationary solutions without imposing uniform convexity and uniform monotonicity hypotheses

Thanks for your attention!