Mean Field Games with Absorption

– Joint works with M. Ghio, M. Fischer, G. Livieri –

Luciano Campi

London School of Economics, Department of Statistics

Innovative Research in Mathematical Finance, Luminy 2018
dedicated to the 70^{th} birthday of Yuri Kabanov

Contents

Introduction

N-player games

Mean field games

Approximate Nash equilibria

A counterexample

Dependence on past absorptions

Conclusions

Introduction: mean-field games (MFG)

- Introduced by Lasry-Lions (2006) and Huang-Caines-Malhamé (2006).
- Main idea:
  - consider a symmetric N-player game where the players interact through their average behaviour
  - pass to the limit $N \to \infty$ making the effect of any single player on the others more and more negligible,
  - hence reducing to a one-player game interacting with the theoretical distribution of the population.
- The simplification at the limit is due to a “propagation of chaos” phenomenon (LLN) emerging as $N \to \infty$.
- Further difficulty wrt particles systems due to the fact that here particles are optimizing agents.
- Game theory: MFG are closely related to anonymous and non-atomic games (mainly within the context of static or sequential games).

Introduction: literature

- Intuition suggests the following recipe:
  - Pass to the limit MFG first;
  - study the equilibria in the limit problem;
  - use those equilibria as approximation of the equilibria in the pre-limit problems (N large).
- Many approaches in the literature:
  - PDE: Lasry, Lions, Cardaliaguet, Bensoussan, Achdou, Guéant, Gomes, Porretta, Bardi ...
  - Probability via FBSDE or compactification methods: Carmona, Delarue, Lacker, Kolokoltsov, Yam ...
- Very rich and flexible framework: various applications in social sciences (economics, finance, crowd dynamics ...) and engineering.
- See Cardaliaguet LN, Lions' recordings at Collège de France or the very recent Carmona-Delarue book (probabilistic approach!).

Introduction: main results

We consider a system of weakly interacting (controlled) diffusions, stopped as soon as they exit a given set.

Under some technical assumptions (e.g. non-degeneracy of the diffusion coefficient) we have:

- Existence of a feedback solution for the MFG.
- The limit solution is nearly Nash for the pre-limit N-player games.
- Counter-example in the degenerate case: limit solutions might not be nearly Nash for the pre-limit games.

Related works (none of them are games except the last ones):

- Giesecke et al. (2014, 2015), Hambly-Ledger (2016) on McKean-Vlasov with absorption (large portfolios of credit derivatives);
- Delarue-Rubenthaler-Tanré (applications to neurology);
- Chan-Sircar (2015), Graber-Bensoussan (2016) (Bertrand & Cournot, fracking and renewables).

Introduction: main results

We consider a system of weakly interacting (controlled) diffusions, stopped as soon as they exit a given set.

Under some technical assumptions (e.g. non-degeneracy of the diffusion coefficient) we have:

- Existence of a feedback solution for the MFG.
- The limit solution is nearly Nash for the pre-limit N-player games.
- Counter-example in the degenerate case: limit solutions might not be nearly Nash for the pre-limit games.

Related works (none of them are games except the last ones):

- Giesecke et al. (2014, 2015), Hambly-Ledger (2016) on McKean-Vlasov with absorption (large portfolios of credit derivatives);
- Delarue-Rubenthaler-Tanré (applications to neurology);
- Chan-Sircar (2015), Graber-Bensoussan (2016) (Bertrand & Cournot, fracking and renewables).

N-player dynamics

Let $T > 0$ be the finite horizon, $\Gamma \subset \mathbb{R}^d$ the set of control actions, O the set of non-absorbing states.
Given a vector $\boldsymbol{u} = (u_1, \ldots, u_N)$ of Γ-valued progressive feedback strategies, the players' states evolve as

$$X_t^{N,i} = X_0^{N,i} + \int_0^t \left(u_i(s, \boldsymbol{X}^N) + \bar{b}\left(s, X_s^{N,i}, \int_{\mathbb{R}^d} w(y)\pi_s^N(dy)\right)\right) ds + \sigma W_t^{N,i}$$

where $W^{N,1}, \ldots, W^{N,N}$ are ind. Wiener processes, σ $d \times d$-matrix, $\bar{b}$, w given functions, and $\pi^N(t, \cdot)$ the conditional empirical measure of the states of the players still in O at time t:

$$\pi_t^N(\cdot) \doteq \begin{cases} \frac{1}{\bar{N}_t^N} \sum_{j=1}^N \mathbf{1}_{[0,\tau^{X^{N,j}})}(t) \cdot \delta_{X_t^{N,j}}(\cdot) & \text{if } \bar{N}_t^N > 0, \\ \delta_0(\cdot) & \text{if } \bar{N}_t^N = 0, \end{cases}$$

$\bar{N}_t^N \doteq \sum_{j=1}^N \mathbf{1}_{[0,\tau^{X^{N,j}})}(t)$ and $\tau^{X^{N,j}} \doteq \inf\{t \geq 0 : X_t^{N,j} \notin O\}$.

Initial distribution $\nu_N \doteq \text{Law}(X_0^{N,1}, \ldots, X_0^{N,N})$ fixed and symmetric.

N-player costs

Let $\mathcal{U}_N$ be the set of all Γ-valued progr. meas. processes. Let $\mathcal{U}_{fb}^N$ be the set of all strategy vectors $\boldsymbol{u} \in \times^N \mathcal{U}_N$ such that Eq. (16) under $\boldsymbol{u}$ with initial law ν_N has a solution unique in distribution.

Player i evaluates $\boldsymbol{u} = (u_1, \ldots, u_N) \in \mathcal{U}_{fb}^N$ so as to minimize

$$J^{N,i}(\boldsymbol{u}) \doteq \mathbf{E}\left[\int_0^{\tau^{N,i}} f\left(s, X_s^{N,i}, \int_{\mathbb{R}^d} w(y)\pi_s^N(dy), u_i\left(s, \boldsymbol{X}^N\right)\right) ds + F\left(\tau^{N,i}, X_{\tau^{N,i}}^{N,i}\right)\right],$$

where

- $\boldsymbol{X}^N = (X^{N,1}, \ldots, X^{N,N})$ is a solution of Eq. (16) under $\boldsymbol{u}$ with initial distribution ν_N,
- $\tau^{N,i} \doteq \tau^{X^{N,i}} \wedge T$, the random time horizon for player i,
- $\pi^N(\cdot)$ is the conditional empirical measure process induced by $(X^{N,1}, \ldots, X^{N,N})$.

Assumptions

(H1) Boundedness and measurability: w, $\bar{b}$, f, F are Borel measurable functions uniformly bounded by some constant $K > 0$.

(H2) Continuity: w, F are continuous, and $f(t, \cdot)$ is continuous uniformly in t.

(H3) Lipschitz continuity: $\bar{b}(t, \cdot)$ Lipschitz with constant L uniformly in t.

(H4) Action space: $\Gamma \subset \mathbb{R}^d$ is compact (and non-empty).

(H5) State space: $O \subset \mathbb{R}^d$ is non-empty, open, and bounded such that ∂O is a C^2-manifold.

For main results, additional non-degeneracy assumption:

- σ is a matrix of full rank.

Under non-degeneracy assumption, $\mathcal{U}_{fb}^N = \times^N \mathcal{U}_N$.

Assumptions

- (H1) Boundedness and measurability: w, $\bar{b}$, f, F are Borel measurable functions uniformly bounded by some constant $K > 0$.
- (H2) Continuity: w, F are continuous, and $f(t, \cdot)$ is continuous uniformly in t.
- (H3) Lipschitz continuity: $\bar{b}(t, \cdot)$ Lipschitz with constant L uniformly in t.
- (H4) Action space: $\Gamma \subset \mathbb{R}^d$ is compact (and non-empty).
- (H5) State space: $O \subset \mathbb{R}^d$ is non-empty, open, and bounded such that ∂O is a C^2-manifold.

For main results, additional non-degeneracy assumption:

- σ is a matrix of full rank.

Under non-degeneracy assumption, $\mathcal{U}_{fb}^N = \times^N \mathcal{U}_N$.

Nash equilibria

Given a strategy vector $\boldsymbol{u} = (u_1, \ldots, u_N)$ and an individual strategy $v \in \mathcal{U}_N$, indicate by

$$[\boldsymbol{u}^{-i}, v] \doteq (u_1, \ldots, u_{i-1}, v, u_{i+1}, \ldots, u_N)$$

the strategy vector obtained from $\boldsymbol{u}$ by replacing u_i with v.

Definition

Let $\varepsilon \geq 0$. A strategy vector $\boldsymbol{u} = (u_1, \ldots, u_N) \in \mathcal{U}_{fb}^N$ is called an ε-*Nash equilibrium* for the N-player game if for every $i \in \{1, \ldots, N\}$, every $v \in \mathcal{U}_N$ such that $[\boldsymbol{u}^{-i}, v] \in \mathcal{U}_{fb}^N$,

$$J^{N,i}(\boldsymbol{u}) \leq J^{N,i}\left([\boldsymbol{u}^{-i}, v]\right) + \varepsilon.$$

If $\boldsymbol{u}$ is an ε-Nash equilibrium with $\varepsilon = 0$, then $\boldsymbol{u}$ is called a *Nash equilibrium*.

Nash equilibria in full information feedback strategies!

Limit dynamics

Mean field limit suggests to consider the equation

$$X_t = X_0 + \int_0^t \left(u(s, X) + \bar{b} \left(s, X_s, \int_{\mathbb{R}^d} w(y) \mathfrak{p}_s(dy) \right) \right) ds + \sigma W_t \tag{3.1}$$

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

- $\mathfrak{p} \in \mathcal{M} \doteq \mathbf{M}([0, T], \mathcal{P}(\mathbb{R}^d))$ is a flow of probability measures,
- $u \in \mathcal{U}_1$ a Γ-valued progressive feedback strategy,
- and W a d-dimensional Wiener process.

In view of N-player game, $\mathfrak{p}$ should correspond to a flow of conditional probabilities!

Limit costs

Let $\mathcal{U}_{fb}$ denote the set of all feedback strategies $u \in \mathcal{U}_1$ such that Eq. (3.1) has a solution unique in law given any initial distribution with support in O.

Under non-degeneracy assumption, $\mathcal{U}_{fb} = \mathcal{U}_1$.

Costs associated with a strategy $u \in \mathcal{U}_{fb}$, a flow of measures $\mathfrak{p} \in \mathcal{M}$, and an initial distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O are given by

$$J(\nu, u; \mathfrak{p}) \doteq \mathbf{E}\left[\int_0^{\tau} f\left(s, X_s, \int_{\mathbb{R}^d} w(y)\mathfrak{p}_s(dy), u(s, X)\right) ds + F\left(\tau, X_\tau\right)\right]$$

where

- X is a solution of Eq. (3.1) under u with initial distribution ν,
- and $\tau \doteq \tau^X \wedge T$ the random time horizon.

Mean field game (MFG)

Definition

A feedback solution of the MFG with absorption is a triple $(\nu, u, \mathfrak{p})$ s.t.

(i) $\nu \in \mathcal{P}(\mathbb{R}^d)$ with $\operatorname{supp}(\nu) \subset O$, $u \in \mathcal{U}_{fb}$, and $\mathfrak{p} \in \mathcal{M}$;

(ii) optimality property: strategy u is optimal for $\mathfrak{p}$ and initial distribution ν in the sense that

$$J(\nu, u; \mathfrak{p}) = V(\nu; \mathfrak{p});$$

(iii) conditional mean field property: if X is a solution of Eq. (3.1) with flow of measures $\mathfrak{p}$, strategy u, and initial distribution ν, then

$$\mathfrak{p}(t, \cdot) = \mathbf{P}\left(X_t \in \cdot \mid \tau^X > t\right)$$

for every $t \in [0, T]$ such that $\mathbf{P}\left(\tau^X > t\right) > 0$.

Approximate Nash equilibria from the mean field game

Set $\mathcal{X} \doteq \mathbf{C}([0,T],\mathbb{R}^d)$. For $\nu \in \mathcal{P}(\mathbb{R}^d)$, let $\Theta_\nu \in \mathcal{P}(\mathcal{X})$ denote the law of $X_t = \xi + \sigma W_t$, $t \in [0,T]$, where $\mathrm{Law}(\xi) = \nu$.

Theorem

Assume non-degeneracy in addition to (A1) – (A5). Suppose $(\nu_N)_{N\in\mathbb{N}}$ is ν-chaotic for some $\nu \in \mathcal{P}(\mathbb{R}^d)$ with support in O.
If $(\nu, u, \mathfrak{p})$ is a feedback solution of the MFG regular in the sense that

$$\Theta_\nu\left(\{\varphi \in \mathcal{X} \;:\; u(t,\cdot) \text{ is discontinuous at } \varphi\}\right) = 0, \;\; a.e.\; t \in [0,T],$$

then $(\boldsymbol{u}^N)_{N\in\mathbb{N}} \subset \mathcal{U}^N_{fb}$ with $\boldsymbol{u}^N = (u^N_1,\ldots,u^N_N)$ defined by

$$u^N_i(t,\boldsymbol{\varphi}) \doteq u(t,\varphi_i), \quad (t,\boldsymbol{\varphi}) \in [0,T]\times\mathcal{X}^N,$$

yields a sequence of approximate Nash equilibria: $\forall\varepsilon > 0$, $\exists N_0(\varepsilon) \in \mathbb{N}$ such that $\boldsymbol{u}^N$ is an ε-Nash equilibrium for the N-player game whenever $N \geq N_0(\varepsilon)$.

Existence of regular solutions to the MFG

Theorem

In addition to the hypotheses of Theorem 1, assume that

$$\Gamma \ni \gamma \mapsto f(t, x, \int w(y)\mu(dy), \gamma) + z \cdot \sigma^{-1} b(t, x, \int w(y)\mu(dy), \gamma)$$

has a unique minimizer for all (x, z, μ) (plus some mild technical assumptions such as $\bar{b}, f$ jointly continuous in all their variables).

Then there exists a feedback solution of the MFG $(\nu, u, \mathfrak{p})$ such that

$$u(t, \varphi) = \tilde{u}(t, \varphi(t))$$

for some continuous function $\tilde{u} \colon [0, T] \times \mathbb{R}^d \to \Gamma$.

Counterexample: intuition

It's possible to construct an example with the following properties:

- for some initial condition, X_i^N cannot leave O before the terminal time T,
- while this is possible in the MFG, where leaving before T is actually preferable.
- Hence implementing MFG opt strategy in the N-player game will push X_i^N towards ∂O but will fail to make it exit.
- This will produce costs higher than those of an alternative strategy.
- Hence MFG solution cannot provide a nearly Nash in the N-player game.
- The example have deterministic state dynamics ($\sigma = 0$), except random initial conditions (singular wrt Lebesgue).

See the joint paper with M. Fischer for details (AAP, 2018).

Dependence on past absorptions

Goal: we want to keep track of the fraction of absorbed players in the dynamics (joint work in progress with M. Ghio and G. Livieri).

We start from the *N*-player game with player *i*'s state variables

$$X_t^{N,i} = X_0^{N,i} + \int_0^t \left(u_i(s, \boldsymbol{X}^N) + \bar{b}\left(s, X_s^{N,i}, L_s^N, \int_{\mathbb{R}^d} w(y)\mu_s^N(dy) \right) \right) ds + \sigma W_t^{N,i},$$

where

- L_s^N is the % of absorbed players before time s;
- $\mu_s(dy) = N^{-1} \sum_i \mathbf{1}_{\{\tau^{X^{N,i}} < s\}} \delta_{X_s^{N,i}}(dy)$.

Moreover player *i* minimizes

$$J^{N,i}(\boldsymbol{u}) \doteq \mathbf{E}\left[\int_0^{\tau^{N,i}} f\left(s, X_s^{N,i}, L_s^N, \int_{\mathbb{R}^d} w(y)\mu_s^N(dy), u_i\left(s, \boldsymbol{X}^N \right) \right) ds + F\left(\tau^{N,i}, X_{\tau^{N,i}}^{N,i} \right) \right].$$

Dependence on past absorptions

When (formally) passing to the limit as $N \to \infty$ we obtain the dynamics

$$X_t = X_0 + \int_0^t \left(u(s,X) + \bar{b}\left(s, X_s, \overline{m}_s(\mathcal{O}), \int_{\mathbb{R}^d} w(y) m_s(dy) \right) \right) ds + \sigma W_t$$

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

- $m \in \mathcal{M}_{\leq 1} \doteq \mathbf{M}([0,T], \mathcal{P}_{\leq 1}(\mathbb{R}^d))$ is a flow of sub-probability measures;
- $\overline{m}_s(\cdot) = 1 - m_s(\cdot)$, $s \in [0,T]$.

Given m, the (representative) player minimizes

$$J(\nu, u; \mathfrak{p}) \doteq \mathbf{E}\left[\int_0^\tau f\left(s, X_s, \overline{m}_s(\mathcal{O}), \int_{\mathbb{R}^d} w(y) m_s(dy), u(s,X)\right) ds + F(\tau, X_\tau) \right]$$

Mean-field condition (fixed point): $m_s(dy) = \mathbf{P}(X_s \in dy, \tau^X > s)$, $\forall s$.

Dependence on past absorptions

Our results in a nutshell:

- Under weaker assumptions as before, e.g. coefficients with polynomial growth, $\mathcal{O}$ unbounded with smooth boundary,
- we have existence of a MFG solution in feedback form.
- The proof has two steps:
  - truncate the coefficients by some threshold $K > 0$ and show existence in the bounded case as before;
  - pass to the limit ($K \to \infty$) using some compactification methods for controls.
- We also have (in progress!) that any of those MFG solutions provides an approximate Nash equilibrium in the N-player game with N large enough (under non-degeneracy for σ).

Conclusions

- We studied N-player games and MFG with an absorbing set.
- Warning: the limit solution might not be approximately Nash for the pre-limit games.
- Under non-degeneracy, everything goes well as usual (existence, approximation)
- In progress: adding dependence on past absorptions in the drift.

Future directions:

- Applications to economics/finance, e.g. interaction between defaultable firms, MFG for corporate finance models (in progress).
- Singular controls (as recent papers by Fu-Horst and Guo-Lee).
- Cascade effect (as recent papers by Nadtochiy-Shkolnikov).

Conclusions

- We studied N-player games and MFG with an absorbing set.
- Warning: the limit solution might not be approximately Nash for the pre-limit games.
- Under non-degeneracy, everything goes well as usual (existence, approximation)
- In progress: adding dependence on past absorptions in the drift.

Future directions:

- Applications to economics/finance, e.g. interaction between defaultable firms, MFG for corporate finance models (in progress).
- Singular controls (as recent papers by Fu-Horst and Guo-Lee).
- Cascade effect (as recent papers by Nadtochiy-Shkolnikov).

Thanks very much for your attention!