Mean Field Games with incomplete information

Bruno Ziliotto

CEREMADE, CNRS, Paris Dauphine University

joint work with Eran Shmaya

Northwestern University, Chicago

April 16th, 2019

Standard model (Lasry and Lions and Huang, Caines and Malhamé, 2006)

A Mean Field Game is a continuous-time dynamic game with a continuum of non-atomic players, where :

- at each time instant t, each player in position $X_t \in \mathbb{R}^n$ chooses its velocity $\alpha_t \in \mathbb{R}^n$ and its new position is $X_{t+dt} = X_t + \alpha_t dt + \sigma dB_t$,
- at each time instant t, each player suffers a stage cost that depends on position X_t, velocity α_t and **distribution of players** at time t.
- The total cost is the expected average cost between time 0 and $T > 0$ (alternatives : discounted, long-run average payoff...).

Standard model (Lasry and Lions and Huang, Caines and Malhamé, 2006)

A Mean Field Game is a continuous-time dynamic game with a continuum of non-atomic players, where :

- at each time instant t, each player in position $X_t \in \mathbb{R}^n$ chooses its velocity $\alpha_t \in \mathbb{R}^n$ and its new position is $X_{t+dt} = X_t + \alpha_t dt + \sigma dB_t$,
- at each time instant t, each player suffers a stage cost that depends on position X_t, velocity α_t and **distribution of players** at time t.
- The total cost is the expected average cost between time 0 and $T > 0$ (alternatives : discounted, long-run average payoff...).

Standard model (Lasry and Lions and Huang, Caines and Malhamé, 2006)

A Mean Field Game is a continuous-time dynamic game with a continuum of non-atomic players, where :

- at each time instant t, each player in position $X_t \in \mathbb{R}^n$ chooses its velocity $\alpha_t \in \mathbb{R}^n$ and its new position is $X_{t+dt} = X_t + \alpha_t dt + \sigma dB_t$,
- at each time instant t, each player suffers a stage cost that depends on position X_t, velocity α_t and **distribution of players** at time t.
- The total cost is the expected average cost between time 0 and $T > 0$ (alternatives : discounted, long-run average payoff...).

A few examples

- Bank run (Carmona, Delarue and Lacker 2017)
- Traders market (Cardaliaguet and Lehalle 2017)
- Macro-economic growth model (Achdou et al. 2014)

We introduce a MFG model with incomplete information :

- Cost function depends on a fixed parameter called *state*, that is unknown to players.
- Players get a private stream of signals about the state during the game.

- An (unknown) state $S \in \mathbb{R}^m$, such that $S \sim \mathcal{N}(0, Id)$.

- A signalling process (Z_t) for the representative player, that is a process on $\mathbb{R}^m$ that follows the equation ($\sigma > 0$)

$$dZ_t = Sdt + \sigma dB_t.$$

- The representative player has a position $X_t \in \mathbb{R}^n$, that he observes, and that evolves according to

$$dX_t = \alpha(t, X_t, Z_t)dt + \sqrt{2}\, dB'_t,$$

where (B_t) and (B'_t) are independent Brownians, and α is measurable and chosen by the player.

Model

- An (unknown) state $S \in \mathbb{R}^m$, such that $S \sim \mathcal{N}(0, Id)$.

- A signalling process (Z_t) for the representative player, that is a process on $\mathbb{R}^m$ that follows the equation ($\sigma > 0$)

$$dZ_t = Sdt + \sigma dB_t.$$

- The representative player has a position $X_t \in \mathbb{R}^n$, that he observes, and that evolves according to

$$dX_t = \alpha(t, X_t, Z_t)dt + \sqrt{2}\, dB'_t,$$

where (B_t) and (B'_t) are independent Brownians, and α is measurable and chosen by the player.

Model

- An (unknown) state $S \in \mathbb{R}^m$, such that $S \sim \mathcal{N}(0, Id)$.

- A signalling process (Z_t) for the representative player, that is a process on $\mathbb{R}^m$ that follows the equation ($\sigma > 0$)

$$dZ_t = Sdt + \sigma dB_t.$$

- The representative player has a position $X_t \in \mathbb{R}^n$, that he observes, and that evolves according to

$$dX_t = \alpha(t, X_t, Z_t)dt + \sqrt{2}\, dB'_t,$$

where (B_t) and (B'_t) are independent Brownians, and α is measurable and chosen by the player.

- Given the state s, the density of players that are at time t in position x is denoted by $\rho_s(t,x)$.

- Fix $T > 0$. In the problem starting from t, the player minimizes

$$\mathbb{E}\left(\int_t^T \frac{1}{2}|\alpha(t,X_t,Z_t)|^2 + F(S,X_t,\rho_S(t,X_t))\,dt\right),$$

where F is continuous and bounded.

Optimal control problem of the agent

- Given the state s, the density of players that are at time t in position x is denoted by $\rho_s(t,x)$.

- Fix $T > 0$. In the problem starting from t, the player minimizes

$$\mathbb{E}\left(\int_t^T \frac{1}{2}|\alpha(t,X_t,Z_t)|^2 + F(S,X_t,\rho_S(t,X_t))\,dt\right),$$

where F is continuous and bounded.

For all t, the law of $S|(Z_u)_{0\leq u\leq t}$ is $\mathcal{N}(r_t(Z_t), {\sigma_t}^2)$, where

$$r_t(z) = \frac{z}{t+\sigma^2} \quad \text{and} \quad \sigma_t^2 = \frac{\sigma^2}{\sigma^2+t}.$$

- Given $(Z_u)_{0\leq u\leq t}$, the law of S depends only on t and Z_t.
- The situation of the player at time t can be described by his current position x and current information=signal z.

For all t, the law of $S|(Z_u)_{0\leq u\leq t}$ is $\mathcal{N}(r_t(Z_t), {\sigma_t}^2)$, where

$$r_t(z) = \frac{z}{t+\sigma^2} \quad \text{and} \quad \sigma_t^2 = \frac{\sigma^2}{\sigma^2+t}.$$

- Given $(Z_u)_{0\leq u\leq t}$, the law of S depends only on t and Z_t.
- The situation of the player at time t can be described by his current position x and current information=signal z.

Definition

Let m^0 be an initial player density over $\mathbb{R}^n$. For each $s \in S$, let $m_s : [0, T] \times \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ "regular enough", such that for each t, $m_s(t, .)$ is a probability density over $\mathbb{R}^n \times \mathbb{R}^m$.
The distribution $m = (m_s)_{s \in S}$ is an equilibrium if there exists a control α s.t. :

- α is a minimizer in the control problem $\Gamma(t, x, z)$ with cost

$$\mathbb{E}\left(\int_t^T \frac{1}{2}|\alpha(t, X_t, Z_t)|^2 + F(S, X_t, \rho_S(t, X_t))\, dt\right),$$

  where $\rho_s(t, .)$ is the marginal of $m_s(t, .)$ on $\mathbb{R}^n$.
- For any s and t, the distribution of (X_t, Z_t) under control α, given that $X_0 \sim m^0$ and $S = s$, is $m_s(t, .)$.

"If the initial distribution of players and signals is m^0, and all the other players play α, then it is optimal for a player to play α."

Definition

Let m^0 be an initial player density over $\mathbb{R}^n$. For each $s \in S$, let $m_s : [0,T] \times \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ “regular enough", such that for each t, $m_s(t,.)$ is a probability density over $\mathbb{R}^n \times \mathbb{R}^m$.
The distribution $m = (m_s)_{s \in S}$ is an equilibrium if there exists a control α s.t. :

- α is a minimizer in the control problem $\Gamma(t,x,z)$ with cost
$$\mathbb{E}\left(\int_t^T \frac{1}{2}|\alpha(t,X_t,Z_t)|^2 + F(S,X_t,\rho_S(t,X_t))\,dt\right),$$
where $\rho_s(t,.)$ is the marginal of $m_s(t,.)$ on $\mathbb{R}^n$.
- For any s and t, the distribution of (X_t,Z_t) under control α, given that $X_0 \sim m^0$ and $S = s$, is $m_s(t,.)$.

"If the initial distribution of players and signals is m^0, and all the other players play α, then it is optimal for a player to play α."

Assume for this slide that the state is known (standard case).

Let m be a distribution. Let $u(t,x)$ be the optimal cost in $\Gamma(t,x)$. Then m is an equilibrium if and only if (m,u) satisfies the MFG system

$$-\partial_t u + \frac{1}{2}|D_x u|^2 - \Delta_x u = F(x,m) \qquad \text{(Hamilton-Jacobi equation)}$$

$$\partial_t m - div_x(D_x u\, m) - \Delta_x m = 0 \qquad \text{(Fokker-Planck equation)}$$

- The Hamilton-Jacobi equation corresponds to a dynamic programming principle for $\Gamma(t,x)$.
- The Fokker-Planck equation describes the evolution of the density distribution of the process (X_t), given that $dX_t = \alpha(t,X_t)dt + \sqrt{2}dB_t$, and $\alpha(t,X_t) = -D_x u(t,X_t)$.

Assume for this slide that the state is known (standard case).

Let m be a distribution. Let $u(t,x)$ be the optimal cost in $\Gamma(t,x)$. Then m is an equilibrium if and only if (m,u) satisfies the MFG system

$$-\partial_t u + \frac{1}{2}|D_x u|^2 - \Delta_x u = F(x,m) \qquad \text{(Hamilton-Jacobi equation)}$$

$$\partial_t m - div_x(D_x u\, m) - \Delta_x m = 0 \qquad \text{(Fokker-Planck equation)}$$

- The Hamilton-Jacobi equation corresponds to a dynamic programming principle for $\Gamma(t,x)$.
- The Fokker-Planck equation describes the evolution of the density distribution of the process (X_t), given that $dX_t = \alpha(t,X_t)dt + \sqrt{2}dB_t$, and $\alpha(t,X_t) = -D_x u(t,X_t)$.

Assume for this slide that the state is known (standard case).

Let m be a distribution. Let $u(t,x)$ be the optimal cost in $\Gamma(t,x)$. Then m is an equilibrium if and only if (m,u) satisfies the MFG system

$$-\partial_t u + \frac{1}{2}|D_x u|^2 - \Delta_x u = F(x,m) \qquad \text{(Hamilton-Jacobi equation)}$$

$$\partial_t m - div_x(D_x u\, m) - \Delta_x m = 0 \qquad \text{(Fokker-Planck equation)}$$

- The Hamilton-Jacobi equation corresponds to a dynamic programming principle for $\Gamma(t,x)$.
- The Fokker-Planck equation describes the evolution of the density distribution of the process (X_t), given that $dX_t = \alpha(t,X_t)dt + \sqrt{2}dB_t$, and $\alpha(t,X_t) = -D_x u(t,X_t)$.

For $(t, x, z) \in \mathbb{R}_+ \times \mathbb{R}^n \times \mathbb{R}^m$, and $m = (m_s : [0, T] \times \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R})_s$, let

$$\tilde{F}(t, x, z, m) = \mathbb{E}_{t,z}(F(S, x, \rho_S(t, x)),$$

where the expectation $\mathbb{E}_{t,z}$ is w.r.t. $S \sim \mathcal{N}(r_t(z), \sigma_t^2)$, and $\rho_s(t, .)$ is the marginal of $m_s(t, .)$ on $\mathbb{R}^n$.

Let $m = (m_s)_{s \in S}$ be a distribution. Let $u(t,x,z)$ be the optimal cost in $\Gamma(t,x,z)$. Then m is an equilibrium if and only if (m,u) satisfies the MFG system

$$-\partial_t u + \frac{1}{2}|D_x u|^2 - r_t(z) \cdot D_z u - \Delta_{x,z} u = \widetilde{F}(t,x,z,m)$$

$$\forall s \in S \quad \partial_t m_s - div_x(D_x u \, m_s) + div_z(s \, m_s) - \Delta_{x,z} m_s = 0.$$

The derivation is similar to the standard MFG model. Indeed,

$$dX_t = \alpha(t, X_t, Z_t)dt + \sqrt{2}\, dB'_t,$$

$$dZ_t = Sdt + \sigma dB_t$$

Let $m = (m_s)_{s \in S}$ be a distribution. Let $u(t, x, z)$ be the optimal cost in $\Gamma(t, x, z)$. Then m is an equilibrium if and only if (m, u) satisfies the MFG system

$$-\partial_t u + \frac{1}{2}|D_x u|^2 - r_t(z) \cdot D_z u - \Delta_{x,z} u = \widetilde{F}(t, x, z, m)$$

$$\forall s \in S \quad \partial_t m_s - div_x(D_x u\, m_s) + div_z(s\, m_s) - \Delta_{x,z} m_s = 0.$$

The derivation is similar to the standard MFG model. Indeed,

$$dX_t = \alpha(t, X_t, Z_t)dt + \sqrt{2}\, dB'_t,$$

$$dZ_t = Sdt + \sigma dB_t$$

Proposition (Shmaya and Z.)

For each initial probability distribution, there exists an equilibrium. Moreover, $\alpha^*(t, x, z) = -D_x u(t, x, z)$ is an optimal control.

Proof : same as in the complete information case :

- Given a distribution m, consider the optimal cost u associated to m
- Let m' be the distribution generated by the optimal control α^*.
- Apply a fixed-point argument to the mapping $m \to m'$.

Proposition (Shmaya and Z.)

For each initial probability distribution, there exists an equilibrium. Moreover, $\alpha^*(t, x, z) = -D_x u(t, x, z)$ is an optimal control.

Proof : same as in the complete information case :

- Given a distribution m, consider the optimal cost u associated to m
- Let m' be the distribution generated by the optimal control α^*.
- Apply a fixed-point argument to the mapping $m \to m'$.

Monotonicity assumption

For all $s \in \mathbb{R}^m$, for all $\mu^1 \neq \mu^2 \in \mathcal{P}(\mathbb{R}^n)$,

$$\int_{\mathbb{R}^n} \left(\mu^1(x) - \mu^2(x)\right) \left(F(x, \mu^1, s) - F(x, \mu^2, s)\right) dx > 0.$$

Remark : this holds when $F(x, \mu, s) = \mu(x)$.

Proposition (Shmaya and Z.)

Assume that the monotonicity assumption is satisfied. Let m^1 and m^2 be two equilibria that coincide for $t = 0$. Then $m^1 = m^2$.

Monotonicity assumption

For all $s \in \mathbb{R}^m$, for all $\mu^1 \neq \mu^2 \in \mathcal{P}(\mathbb{R}^n)$,

$$\int_{\mathbb{R}^n} \left(\mu^1(x) - \mu^2(x)\right) \left(F(x,\mu^1,s) - F(x,\mu^2,s)\right) dx > 0.$$

Remark : this holds when $F(x,\mu,s) = \mu(x)$.

Proposition (Shmaya and Z.)

Assume that the monotonicity assumption is satisfied. Let m^1 and m^2 be two equilibria that coincide for $t = 0$. Then $m^1 = m^2$.

Monotonicity assumption

For all $s \in \mathbb{R}^m$, for all $\mu^1 \neq \mu^2 \in \mathcal{P}(\mathbb{R}^n)$,

$$\int_{\mathbb{R}^n} \left(\mu^1(x) - \mu^2(x)\right) \left(F(x, \mu^1, s) - F(x, \mu^2, s)\right) dx > 0.$$

Remark : this holds when $F(x, \mu, s) = \mu(x)$.

Proposition (Shmaya and Z.)

Assume that the monotonicity assumption is satisfied. Let m^1 and m^2 be two equilibria that coincide for $t = 0$. Then $m^1 = m^2$.

- Let m^1 and m^2 be two equilibria such that for all s, $m^1_s(0,.) = m^2_s(0,.) := m^0$. Let $\rho^i_s(t,.)$ be the marginal of $m^i_s(t,.)$ on $\mathbb{R}^n$.

- Denote by $\gamma(i,j)$ the cost of the player when he plays α_i and at each time t, the other players are distributed according to ρ^j, and the initial position of the player is randomized according to m^0.

$$\gamma(i,j) := \mathbb{E}\left[\int_t^T \int_{\mathbb{R}^n} \left[\alpha_i(t,x,z)^2 + F(x,\rho^j_S,S)\right] m^i_S(t,x,z)dt\right].$$

-

$$[\gamma(1,2)-\gamma(2,2)] + [\gamma(2,1)-\gamma(1,1)] = -\mathbb{E}\left[\int_{\mathbb{R}^n} \left[\rho^1_S(t,x)-\rho^2_S(t,x)\right]\left[F(x,\rho^1_S,S)-F(x,\rho^2_S,S)\right] dx\right]$$

Thus, if $\rho^1 \neq \rho^2$, the player has either a profitable deviation in the equilibrium m^1 or in m^2, which is absurd.

Proof

- Let m^1 and m^2 be two equilibria such that for all s, $m^1_s(0,.) = m^2_s(0,.) := m^0$. Let $\rho^i_s(t,.)$ be the marginal of $m^i_s(t,.)$ on $\mathbb{R}^n$.

- Denote by $\gamma(i,j)$ the cost of the player when he plays α_i and at each time t, the other players are distributed according to ρ^j, and the initial position of the player is randomized according to m^0.

$$\gamma(i,j) := \mathbb{E}\left[\int_t^T \int_{\mathbb{R}^n} \left[\alpha_i(t,x,z)^2 + F(x,\rho^j_S,S)\right] m^i_S(t,x,z)dt\right].$$

- $$[\gamma(1,2)-\gamma(2,2)] + [\gamma(2,1)-\gamma(1,1)] = -\mathbb{E}\left[\int_{\mathbb{R}^n} \left[\rho^1_S(t,x)-\rho^2_S(t,x)\right]\left[F(x,\rho^1_S,S)-F(x,\rho^2_S,S)\right]dx\right]$$

  Thus, if $\rho^1 \neq \rho^2$, the player has either a profitable deviation in the equilibrium m^1 or in m^2, which is absurd.

- Let m^1 and m^2 be two equilibria such that for all s, $m^1_s(0,.) = m^2_s(0,.) := m^0$. Let $\rho^i_s(t,.)$ be the marginal of $m^i_s(t,.)$ on $\mathbb{R}^n$.
- Denote by $\gamma(i,j)$ the cost of the player when he plays α_i and at each time t, the other players are distributed according to ρ^j, and the initial position of the player is randomized according to m^0.

$$\gamma(i,j) := \mathbb{E}\left[\int_t^T \int_{\mathbb{R}^n} \left[\alpha_i(t,x,z)^2 + F(x,\rho^j_S,S)\right] m^i_S(t,x,z)dt\right].$$

- $$\begin{aligned}[\gamma(1,2)-\gamma(2,2)] + [\gamma(2,1)-\gamma(1,1)] &= \\ -\mathbb{E}\left[\int_{\mathbb{R}^n} \left[\rho^1_S(t,x) - \rho^2_S(t,x)\right]\left[F(x,\rho^1_S,S) - F(x,\rho^2_S,S)\right] dx\right]\end{aligned}$$

  Thus, if $\rho^1 \neq \rho^2$, the player has either a profitable deviation in the equilibrium m^1 or in m^2, which is absurd.

- A continuum of firms sell the same object,
- The preferences of the consumers are unknown, and firms get signals about them,
- Firms aim at selling an object that reflect the preferences of consumers, but without facing too much competition.

- Position $x \in \mathbb{R}^n$ of a firm=characteristics of the object it sells
- Preferences of the consumer are represented by a state $S \in \mathbb{R}^n$.
- Cost function :

$$F(x, \mu, s) := g(|x - s|) + \int_{\mathbb{R}^n} h(-|x - y|)\mu(y)dy,$$

where g and h are continuous, increasing and integrable.

Value of information

- Does more precise information leads to a better equilibrium ?
- Recall that players receive a stream of signals (Z_t) such that $dZ_t = Sdt + \sigma dB_t$.
- Is the equilibrium cost increasing in σ ?
- Simpler question : is the equilibrium better for $\sigma = 0$ rather than for $\sigma = +\infty$?

Proposition (Shmaya and Z.)

Assume that

$$\mathbb{E}_{S \sim \mathcal{N}(0,Id)}(g(|S|)) - g(0) > \|h\|_{\infty}.$$

Then for large duration of time T, the equilibrium for $\sigma = 0$ is better than the equilibrium for $\sigma = +\infty$.

Thank you for your attention !