Mean field games vs. best reply strategy

Matt Barker (Imperial), Pierre Degond (Imperial) and Marie-Therese Wolfram (Warwick)

Virtual workshop on Mean field games

WARWICK
THE UNIVERSITY OF WARWICK

N-player stochastic differential game

Consider N players, each adjusting its state to maximise/minimise a given cost criterion over a finite or infinite time horizon.

- Mean field game (MFG): agents know the current and future distribution of all others - anticipate the future.[a]
- Best reply strategy (BRS): agents determine their locally best action by minimising their cost over a short receding time horizon; also known as model predictive (MPC) or receding horizon control.[b]

$$
\begin{tikzcd}
\text{N-player stochastic differential game} \arrow[r, "\text{MPC}"] \arrow[d, "N \to \infty"] & \text{BRS for the N-player game} \arrow[d, "N \to \infty"] \\
\text{MFG} \arrow[r, "\text{MPC}"] & \text{BRS for MFG}
\end{tikzcd}
$$

[a] Lasry and Lions, CR Math 343 2006 and Jap. J. Math. 2, 2007; Huang, Caines and Malhame CIS 6 2016 and J. Sys. Sci. Comp. 20, 2017

[b] Degond, Herty and Liu, CMS 15, 2017; Barker, J. Dyn. Games 6, 2019

Focus of the talk

We will focus on (MFG) and (BRS) problems in case of a quadratic Hamiltonian as well as congestion cost and a confinement potential on a bounded domain.

- How do solutions of BRS and MFG models compare ?
- Under which conditions are solutions close ?
- When should we prefer one to another ?

Mean field games: the setting

Consider N interacting agents on the torus $\mathbb{T}^d$, each agent with state X_t is trying to optimise

$$J(\alpha; m) = \mathbb{E}\left[\int_0^T \left(\frac{\alpha_s^2}{2} + h(X_s, m(X_s))\right)\, ds\right].$$

subject to the constraint that

$$dX_t = \alpha_t dt + \sigma dB_t \text{ with initial condition } \mathcal{L}(X_0) = m_0.$$

Functions and parameters:

- $m : \mathbb{T}^d \to [0, \infty)$ is the density of the agent distribution
- $h : \mathbb{T}^d \times (0, \infty) \to \mathbb{R}$ is a density dependent cost
- $\alpha_t : [0, T] \to \mathbb{T}^d$ is a control chosen by a representative agent
- $\sigma > 0$ the amplitude of the noise and dB_t a d-dimensional Wiener process.

Then the optimal cost trajectory $u(x, t)$ is

$$u(x, t) = \inf_{\alpha \in \mathcal{A}} \mathbb{E}\left[\int_t^T \left(\frac{\alpha_s^2}{2} + h(X_s, m(X_s))\right)\, ds \,\middle|\, X_t = x\right].$$

and the optimal control by $\alpha_t^* = -\nabla u(X_t, t)$.

The limit $N \to \infty$:

Then the distribution of agents $m(x,t)$ and the optimal cost trajectory $u = u(x,t)$ solve the MFG

$$
\begin{aligned}
\partial_t u &= \frac{|\nabla u|^2}{2} - h(x,m) - \frac{\sigma^2}{2}\nabla^2 u\\
\partial_t m &= \nabla \cdot [m \nabla u] + \frac{\sigma^2}{2}\nabla^2 m\\
m(x,0) &= m_0\\
u(x,T) &= 0\,.
\end{aligned}
$$

Best reply strategy: the idea

Agents consider a rescaled cost functional over a short time

$$J^{\Delta t}(\alpha; m) = \mathbb{E}\left[\int_t^{t+\Delta t}\left(\frac{\alpha_s^2}{2} + \frac{1}{\Delta t}h(X_s, m(X_s))\right) ds \middle| X_t = x\right].$$

Taking the limit $\Delta t \to 0$, yields the optimal control

$$\alpha_t = -\left.[\nabla h(x, m(x))]\right|_{x=X_t}$$

Then the distribution of agents evolves according to the Fokker–Planck equation

$$\partial_t m = \nabla \cdot [(\nabla h(x, m(x)))\, m] + \frac{\sigma^2}{2}\nabla^2 m$$
$$m(x,0) = m_0 .$$

The stationary problems

- **MFG:** triple $(m, u, \lambda) \in C^2(\mathbb{T}^d) \times C^2(\mathbb{T}^d) \times \mathbb{R}$ solves

$$
\begin{aligned}
-\frac{\sigma^2}{2}\nabla^2 m - \nabla \cdot (m \nabla u) &= 0 \\
-\frac{\sigma^2}{2}\nabla^2 u + \frac{|\nabla u|^2}{2} - h(x, m) + \lambda &= 0 \\
\int_{\mathbb{T}^d} m \; dx &= 1 \\
\int_{\mathbb{T}^d} u \; dx &= 0 \, .
\end{aligned}
$$

- **BRS:** setting $\partial_t m = 0$ gives

$$
\begin{aligned}
\nabla \cdot [(\nabla h(x, m(x))) \, m] + \frac{\sigma^2}{2}\nabla^2 m &= 0 \\
\int_{\mathbb{T}^d} m \; dx &= 1 \, .
\end{aligned}
$$

We will consider the stationary MFG and BRS on Ω with **no flux boundary conditions** in the following.

Problem setting and assumptions:

General assumptions:

(A$_1$) $\Omega \subset \mathbb{R}^d$ is an open bounded set with a $C^{2,\alpha}$ boundary, for some $\alpha \in (0,1)$ and $d \geq 1$.

(A$_2$) $h(x,\cdot)$ is an increasing function for every $x \in \Omega$.

(A$_3$) There exists a continuous function $g : (0,\infty) \to [0,\infty)$ such that

$$\sup_{x\in\Omega} |h(x,m)| \leq g(m) \text{ for every } m \in (0,\infty)$$

BRS assumptions:

(BRS$_1$) $h \in C^2(\Omega \times (0,\infty)) \cap C^1(\bar{\Omega} \times (0,\infty))$.

(BRS$_2$) There exists a continuous function $f : (0,\infty) \to [0,\infty)$ such that

$$\sup_{x\in\Omega} |\nabla_x h(x,m)| \leq f(m) \text{ for every } m \in (0,\infty).$$

MFG assumptions:

(MFG$_1$) $h \in C(\Omega \times (0,\infty))$

(MFG$_2$) $\lim_{m\to 0} \sup_{x\in\Omega} h(x,m) < \inf_{x\in\Omega} h\left(x, \frac{1}{|\Omega|}\right)$.

(MFG$_3$) $\sup_{x\in\Omega} h\left(x, \frac{1}{|\Omega|}\right) < \lim_{m\to\infty} \inf_{x\in\Omega} h(x,m)$.

BRS: existence of stationary solutions

Definition of classical BRS solutions: Let assumptions (A1)-(A3) and (BRS1)-(BRS2) be satisfied. Then a function $m : \Omega \to (0, \infty)$ is a classical solution if it satisfies

$$
\begin{aligned}
m \in C^2(\Omega) \cap C^1(\bar{\Omega}) \\
-\frac{\sigma^2}{2}\nabla^2 m - \nabla \cdot (m \nabla [h(x,m)]) = 0\,, \quad & x \in \Omega \\
-\frac{\sigma^2}{2}\nabla m \cdot \nu - m \nabla [h(x,m)] \cdot \nu = 0\,, \quad & x \in \partial\Omega \\
\int_\Omega m\, dx = 1\,.
\end{aligned}
$$

Theorem:

There exists a unique solution $m : \Omega \to (0, \infty)$ to the stationary BRS.

MFG: existence of stationary solutions

Definition of classical MFG solutions: Let assumptions (A1)-(A3) and (MFG1)-(MFG3) be satisfied. Then the triple $m : \Omega \to (0, \infty)$, $u : \Omega \to \mathbb{R}$ and $\lambda \in \mathbb{R}$ is a classical solution if it satisfies

$$\begin{aligned}
m \in C^2(\Omega) \cap C^1\left(\bar{\Omega}\right) \text{ and } u \in C^2(\Omega) \cap C^1\left(\bar{\Omega}\right)& \\
-\frac{\sigma^2}{2}\nabla^2 m - \nabla \cdot (m \nabla u) = 0\,, &\quad x \in \Omega \\
-\frac{\sigma^2}{2}\nabla^2 u + \frac{|\nabla u|^2}{2} - h(x, m) + \lambda = 0\,, &\quad x \in \Omega \\
-\frac{\sigma^2}{2}\nabla m \cdot \nu = 0 \text{ and } \nabla u \cdot \nu = 0\,, &\quad x \in \partial\Omega \\
\int_\Omega m \; dx = 1 \text{ and } \int_\Omega u \; dx = 0.&
\end{aligned}$$

Theorem:

There exists a unique solution (m, u, λ) to the MFG system.

Idea of the proof

Introduce $Z = \int_\Omega e^{-\frac{2}{\sigma^2} u}\, dx$, then $m = \frac{1}{Z} e^{-\frac{2}{\sigma^2} u}$ and finding a MFG solution is equivalent to finding a unique solution $(u, \lambda, Z) \in \left[C^2(\Omega) \cap C^1(\bar{\Omega})\right] \times \mathbb{R} \times (0, \infty)$ to

$$
\begin{aligned}
u \in C^2(\Omega) \cap C^1(\bar{\Omega}) \\
-\frac{\sigma^2}{2} \nabla^2 u + \frac{|\nabla u|^2}{2} - h\left(x, \frac{1}{Z} e^{-\frac{2}{\sigma^2} u}\right) + \lambda = 0\,, \quad x \in \Omega \\
-\nabla u \cdot \nu = 0\,, \quad x \in \partial\Omega \\
\int_\Omega \frac{1}{Z} e^{-\frac{2}{\sigma^2} u}\, dx = 1, \text{ and } \int_\Omega u\, dx = 0\,.
\end{aligned}
$$

MFG vs. BRS in case of a specific example

Specific example with quadratic potential and logarithmic congestion: [a]

$$h(x, m) = \beta x^2 + \log m, \text{ with } \beta \leq 0.$$

Then the **stationary MFG** on the real line

$$\begin{aligned} \frac{\sigma^2}{2}\partial_{xx}^2 m + \partial_x\left(m\partial_x u\right) &= 0\,, \quad x \in \mathbb{R}\,, \\ -\frac{|\partial_x u|^2}{2} + \log m + \beta x^2 + \frac{\sigma^2}{2}\partial_{xx}^2 u + \lambda &= 0\,, \quad x \in \mathbb{R}\,, \end{aligned}$$

with $\int_{\mathbb{R}} m\, dx = 1$ has the explicit solution

$$m(x) = \left(\frac{a}{\pi}\right)^{1/2} e^{-ax^2},\ u(x) = bx^2 \text{ and } \lambda = \log\left(\frac{\pi}{a}\right) - \sigma^2 b\,,$$

Stationary BRS:

$$\partial_x\left(m\partial_x(\log m + \beta x^2)\right) + \frac{\sigma^2}{2}\partial_{xx}^2 m = 0\,, \quad x \in \mathbb{R}$$

with $\int_{\mathbb{R}} m\, dx = 1$ has the explicit solution $m(x) = \left(\frac{2\beta}{(2+\sigma^2)\pi}\right)^{1/2} e^{-\frac{2\beta}{(2+\sigma^2)}x^2}$.

[a] Gomes, Pimentel and Voskanyan, Springer Briefs in Mathematics, 2016; Gueant, J. Math. Pures Appl. 92, 2009

Comparison:

For $\sigma > 0$, the stationary distributions of the MFG system and the BRS are given by normal distributions, with mean 0 and variances a_1 and a_2 respectively, where

$$a_1 = \frac{\sigma^4}{-2 + 2(1 + 2\sigma^4\beta)^{1/2}} \text{ and } a_2 = \frac{2 + \sigma^2}{4\beta}.$$

Then, for fixed $\beta \geq 0$

$$\lim_{\sigma^2 \to 0} \frac{a_2}{a_1} = 1 \text{ and } \lim_{\sigma^2 \to \infty} \frac{a_2}{a_1} = \frac{1}{(2\beta)^{1/2}}.$$

While, for fixed $\sigma > 0$

$$\lim_{\beta \to 0} \frac{a_2}{a_1} = 1 + \frac{\sigma^2}{2} \text{ and } \lim_{\beta \to \infty} (2\beta)^{1/2} \frac{a_2}{a_1} = \frac{2 + \sigma^2}{\sigma^2}.$$

Simulation results

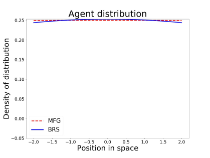

(a) $\beta = 0.1,\ \frac{\sigma^2}{2} = 10$

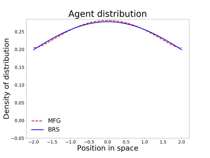

(b) $\beta = 0.1,\ \frac{\sigma^2}{2} = 0.2$

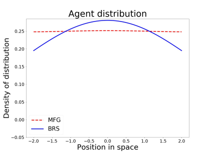

(c) $\beta = 1,\ \frac{\sigma^2}{2} = 10$

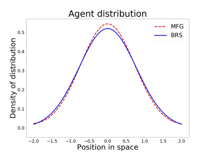

(d) $\beta = 1,\ \frac{\sigma^2}{2} = 0.2$

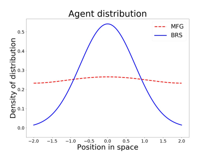

(e) $\beta = 10,\ \frac{\sigma^2}{2} = 10$

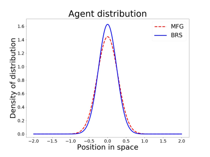

(f) $\beta = 10,\ \frac{\sigma^2}{2} = 0.2$

Trapped in wells:

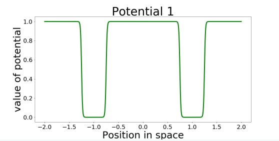

(a) Potential $F(x)$ when $h(x, m) = F(x) + \log(m)$

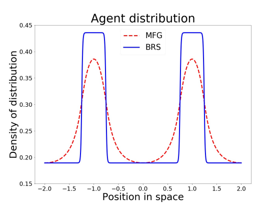

(b) $\frac{\sigma^2}{2} = 0.2$

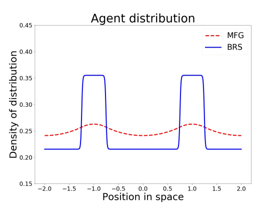

(c) $\frac{\sigma^2}{2} = 1$

Trapped in wells:

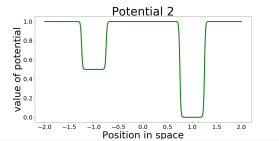

(a) Potential $F(x)$ when $h(x, m) = F(x) + \log(m)$

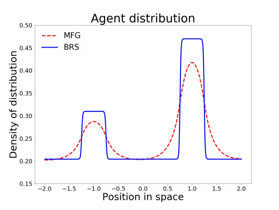

(b) $\frac{\sigma^2}{2} = 0.2$

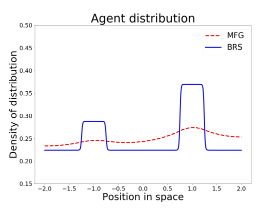

(c) $\frac{\sigma^2}{2} = 1$

Trapped in wells:

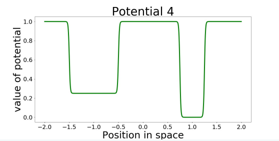

(a) Potential $F(x)$ when $h(x, m) = F(x) + \log(m)$

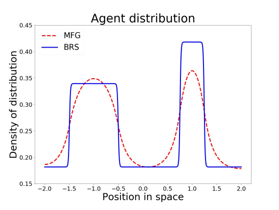

(b) $\frac{\sigma^2}{2} = 0.2$

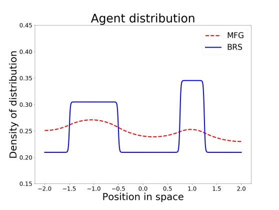

(c) $\frac{\sigma^2}{2} = 1$

Trapped in wells:

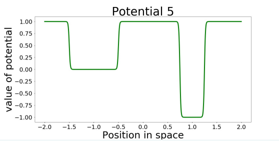

(a) Potential $F(x)$ when $h(x, m) = F(x) + \log(m)$

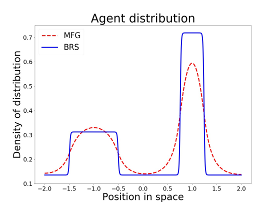

(b) $\frac{\sigma^2}{2} = 0.2$

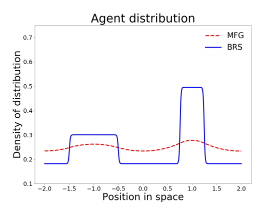

(c) $\frac{\sigma^2}{2} = 1$

Conclusions:

Comparison of stationary (MFG) and (BRS): first insights in concrete example, but no general statements.

Dynamics (MFG) vs. (BRS): to come....

Thank you very much for your attention !

References:

M. Barker, P. Degond and MTW, *Comparing the best reply strategy and mean field games: the stationary case*, Arxiv, 2019.