Stochastic Differential Games in Finite and Infinite Population Regimes

Part II: *Going from finite to infinite population and back: Mean field games under Nash equilibrium*

Tamer Başar

Coordinated Science Laboratory
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign
Urbana, IL - 61801

June 19-20, 2018

Outline

- Introduction to Mean Field Games
- Risk-Sensitive Mean Field Games with Coupled PDEs
- LQ Risk-Sensitive Mean Field Games
- Risk-Sensitive Mean Field Games via the SMP
- Conclusions

Outline

1. Introduction

Introduction: Mean Field Games

- Decentralized or distributed optimal decision making problem for a large number of *weakly* coupled players through a mean field
- Mean field: mass behavior, average behavior of all agents (e.g. $\frac{1}{N}\sum_{i=1}^{N} x_i(t)$)
- Initiated by P. L. Lions, 2007 (PDE approach)
- Applications: optimal decentralized control of large-scale multi-agent systems
  - smart grid, traffic networks, social networks, economics, biology, etc.

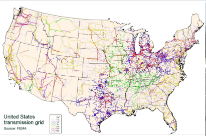



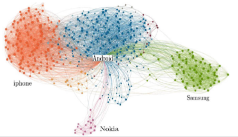

Introduction: Mean Field Games

- Applications
  - finance (Lasry and Lions, Jap. Journal of Math, 2007)
  - economics with a large number of firms (Weintraub et al., Econometrica, 2008)
  - consensus and social networks (Nourian et al., TAC, 2013)
  - control of a large number of oscillators (Yin et al., TAC, 2012)
  - wireless power control (Huang et al., CDC, 2003)
  - the planning problem (Achdou, et al., SICON, 2012)
  - electric vehicle charging (Ma et al., CST, 2013)
  - smart grid (Chen et al., TAC, 2017)
  - biology (Zhu et al., CDC, 2011)

Introduction: Mean Field Games

- Major challenges faced in stochastic DGs with NE
  - Existence (and uniqueness) of Nash equilibria
  - Centralized optimization: It is hard to find a Nash equilibrium when N is very large (*curse of dimensionality*, R. Bellman)
  - Informational constraint: Each player needs to know or deduce state information of all other players (informational restrictions)
  - Strategic interactions and difficulties due to nonlinearity
  - Challenges of distributed control (Blondel and Tsitsiklis, SICON, 1997)

Introduction: Mean Field Games

- (Risk-sensitive) Mean field analysis
  - Consider the situation when $N \to \infty$
  - The empirical distribution ν_N converges to some probability measure (needs to be justified)
  - The impact of each player on other players becomes negligible when N is large
  - The problem reduces to solving the (risk-sensitive) stochastic optimal control problem and the associated fixed point problem
- Obtain an *efficient* algorithm for risk-sensitive mean field games
- Characterize a suboptimal (approximated) decentralized Nash equilibrium
  - Individual Nash control u_i^*: decentralized (function of the local state information)

Introduction: Mean Field Games

- Some observations on games with a large number of players: Theory of Games and Economic Behavior (von Neumann and Morgenstern, 1944, pp.13-14)
- "When the number of participants becomes really great....... a more conventional theory becomes possible."
- Very great numbers are often easier to handle than those of medium size...... This is of course, due to the excellent possibility of applying the laws of statistics and probabilities in the first case

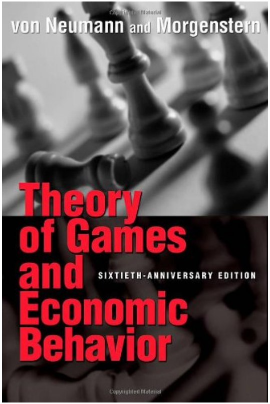

Outline

2. Risk-Neutral Mean Field Games

Risk-Neutral Mean Field Games

- Risk-neutral MFGs: Huang et al (2007), Lasry and Lions (2007), Tao and Zhang (2008)
- Stochastic differential equation for player i (linear dynamics)

$$dx_i(t) = (A_i x(t) + B_i u_i(t))dt + D_i dW_i(t)$$

- The (quadratic) risk-neutral cost function for agent i

$$J_i^N(u_i, u_{-i}) = \mathbb{E} \int_0^T \left\| x_i(t) - \frac{1}{N} \sum_{i=1}^{N} x_i(t) \right\|_Q^2 + \|u_i(t)\|_R^2 dt$$

- Agents are coupled with each other through the mean field term
- Agents have access to local state information only, but with finite N obtaining NE is a (very) challenging problem

Risk-Neutral Mean Field Games

- ϵ-Nash equilibrium

$$J_i^N(u_i^*, u_{-i}^*) \leq J_i^N(u_i, u_{-i}^*) + \epsilon$$

- u_i^*: decentralized ϵ-Nash strategy
  - u_i^* is a function of local state information
- Obtain MF equilibrium for the infinite population, and use it for the N-player game $\Rightarrow$ $\epsilon(N)$ NE

Outline

3. Risk-Sensitive Mean Field Games with Coupled PDEs

Risk-Sensitive Mean Field Games with Coupled PDEs

- Stochastic differential equation of player i, $1 \leq i \leq N$

$$dx_i(t) = f(t, x_i, \nu_N, u_i)dt + \sigma(t)dW_i(t)$$

- ν_N: The empirical distribution (mean field) of the N players

$$\nu_N(t) = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i(t)}(dx)$$

- Risk-sensitive objective function for player i ($\gamma > 0$)

$$J_i(u_i, u_{-i}) = \gamma \log \mathbb{E}\Big[\exp\{\frac{1}{\gamma}\int_0^T l(t, x_i, \nu_N, u_i)dt + \frac{1}{\gamma}m(T, x_i, \nu_N)\}\Big]$$

Tembine, Zhu and Başar, IEEE-TAC, 2014

Risk-Sensitive Mean Field Games with Coupled PDEs

- Consider the symmetric problem
- By taking $N \to \infty$, the empirical measure ν_N converges to some probability measure due to the law of large numbers if the following coupled PDEs have the solution
- Hamilton-Jacobi-Bellman and Fokker-Plank-Kolmogorov coupled PDEs (scalar case)

$$\frac{\partial}{\partial t}V(t,x) + \frac{\sigma^2(t)}{2}\frac{\partial^2}{\partial x^2}V(t,x) + H(t,x,\nu,\frac{\partial}{\partial x}V(t,x),u^*,\gamma) = 0$$
$$\frac{\partial}{\partial t}\nu(t) - \frac{\sigma^2(t)}{2}\frac{\partial^2}{\partial x^2}\nu(t) + \frac{\partial}{\partial x}(f(t,x,\nu,u^*)\cdot\nu(t)) = 0$$

- Initial and boundary conditions: $V(T,x) = m(T,x,\nu)$, $\nu(0) = \delta_{x(0)}$

Risk-Sensitive Mean Field Games with Coupled PDEs

- Hamilton-Jacobi-Bellman and Fokker-Plank-Kolmogorov coupled PDEs (scalar case)

$$\frac{\partial}{\partial t}V(t,x)+\frac{\sigma^2(t)}{2}\frac{\partial^2}{\partial x^2}V(t,x)+H(t,x,\nu,\frac{\partial}{\partial x}V(t,x),u^*,\gamma)=0$$
$$\frac{\partial}{\partial t}\nu(t)-\frac{\sigma^2(t)}{2}\frac{\partial^2}{\partial x^2}\nu(t)+\frac{\partial}{\partial x}(f(t,x,\nu,u^*)\cdot\nu(t))=0$$

- Initial and boundary conditions: $V(T,x)=m(T,x,\nu)$, $\nu(0)=\delta_{x(0)}$

- HJB PDE: stochastic optimal control (backward)
- FPK PDE: evolution of the state probability distribution (forward)
- Existence and uniqueness of the solution of the couple PDEs
  $\Rightarrow$ Very challenging question (especially for risk-sensitive mean field games)
- The solution (u^*,ν) is known as a *mean field equilibrium*

Outline

4. Linear-Quadratic Risk-Sensitive Mean Field Games

Problem Formulation (P1)

- Stochastic differential equation (SDE) for agent i, $1 \leq i \leq N$

$$dx_i(t) = (A(\theta_i)x_i(t) + B(\theta_i)u_i(t))dt + \sqrt{\mu}D(\theta_i)dW_i(t)$$

- The risk-sensitive cost function for agent i with $\delta > 0$

$$J_i^N(u_i, u_{-i}) = \limsup_{T\to\infty} \frac{\delta}{T} \log \mathbb{E}\Big\{\exp[\frac{1}{\delta}\phi_i(x, f_N, u)]\Big\}$$

$$\phi_i(x, f_N, u) := \int_0^T \Big\|x_i(t) - \frac{1}{N}\sum_{i=1}^{N} x_i(t)\Big\|_Q^2 + \|u_i(t)\|_R^2 dt$$

- $f_N(t) := \frac{1}{N}\sum_{i=1}^{N} x_i(t)$: mean field (mass behavior, empirical distribution)
- Agents are coupled with each other through the mean field term
- Individual agents want to *follow* (track through their states) the mean field f_N
- The heterogeneous case $\Rightarrow$ need the prior distribution $F(\theta)$
- The limit $\delta \to \infty$ captures the risk-neutral case [Note: θ in Lecture I is $1/\delta$ here, ϵ there is μ here, and W_i is standard Wiener process]

Moon and Başar, CDC (2014), IEEE-TAC (62(3), 2017), IJC (89(7), 2016)

Problem Formulation (**P1**)

- Motivation of the risk-sensitive cost function
  - The risk-sensitive cost function entails a weighted sum of all the moments of the integral cost $\phi_i \Rightarrow H^\infty$ robustness and performance

$$J_i^N(u_i, u_{-i}) = \limsup_{T\to\infty} \frac{1}{T}\Big[\mathbb{E}\{\phi_i\} + \frac{1}{2\delta}\mathrm{Var}\{\phi_i\} + o(\frac{1}{\delta})\Big]$$

  - $T \to \infty$: mean-square stability
  - Large deviation limit w.r.t. μ (will not be covered in this talk)

Objective

- Obtain $\{u_i^*, 1 \le i \le N\}$ that minimize the cost functions $\Rightarrow$ Characterize a (centralized) Nash equilibrium, $\{u_i^*, 1 \le i \le N\}$,

$$J_i^N(u_i^*, u_{-i}^*) \le \inf_{u_i} J_i^N(u_i, u_{-i}^*),\ 1 \le i \le N$$

LQ Robust MFGs, Problem 2 (**P2**)

- Stochastic differential equation (SDE) for agent i, $1 \leq i \leq N$

$$dx_i(t) = (A(\theta_i)x_i(t) + B(\theta_i)u_i(t) + D(\theta_i)v_i(t))dt + \sqrt{\mu}D(\theta_i)db_i(t)$$

- The worst-case risk-neutral cost function for agent i

$$\mathbf{P2}: \quad J_{2,i}^N(u_i, u_{-i}) = \sup_{v_i \in \mathcal{V}_i} \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}\{\phi_i^2(x, f_N, u, v)\}$$

$$\phi_i^2(x, f_N, u, v) := \int_0^T \|x_i(t) - \frac{1}{N}\sum_{i=1}^N x_i(t)\|_Q^2 + \|u_i(t)\|_R^2 - \gamma^2\|v_i(t)\|^2 dt$$

- v_i can be viewed as a fictitious player (or adversary) of agent i, which strives for a worst-case cost function for agent i
- Agents are coupled with each other through the mean field term

LQ Robust MFGs, Problem 2 (**P2**)

- Stochastic differential equation (SDE) for agent i, $1 \leq i \leq N$

$$dx_i(t) = (A(\theta_i)x_i(t) + B(\theta_i)u_i(t) + D(\theta_i)v_i(t))dt + \sqrt{\mu}D(\theta_i)db_i(t)$$

- The worst-case risk-neutral cost function for agent i

$$\textbf{P2}: \quad J_{2,i}^N(u_i, u_{-i}) = \sup_{v_i \in \mathcal{V}_i} \limsup_{T \to \infty} \frac{1}{T} \mathbb{E}\{\phi_i^2(x, f_N, u, v)\}$$

$$\phi_i^2(x, f_N, u, v) := \int_0^T \|x_i(t) - \frac{1}{N}\sum_{i=1}^N x_i(t)\|_Q^2 + \|u_i(t)\|_R^2 - \gamma^2\|v_i(t)\|^2 dt$$

- v_i can be viewed as a fictitious player (or adversary) of agent i, which strives for a worst-case cost function for agent i
- Agents are coupled with each other through the mean field term

LQ Robust MFGs, Problem 2 (**P2**)

- Stochastic differential equation (SDE) for agent i, $1 \leq i \leq N$

$$dx_i(t) = (A(\theta_i)x_i(t) + B(\theta_i)u_i(t) + D(\theta_i)v_i(t))dt + \sqrt{\mu}D(\theta_i)db_i(t)$$

- The worst-case risk-neutral cost function for agent i

P2: $$J_{2,i}^N(u_i, u_{-i}) = \sup_{v_i \in \mathcal{V}_i} \limsup_{T \to \infty} \frac{1}{T}\mathbb{E}\{\phi_i^2(x, f_N, u, v)\}$$

$$\phi_i^2(x, f_N, u, v) := \int_0^T \|x_i(t) - \frac{1}{N}\sum_{i=1}^N x_i(t)\|_Q^2 + \|u_i(t)\|_R^2 - \gamma^2\|v_i(t)\|^2 dt$$

- v_i can be viewed as a fictitious player (or adversary) of agent i, which strives for a worst-case cost function for agent i
- Agents are coupled with each other through the mean field term

LQ Robust MFGs, Problem 2 (**P2**)

- Stochastic differential equation (SDE) for agent i, $1 \leq i \leq N$

$$dx_i(t) = (A(\theta_i)x_i(t) + B(\theta_i)u_i(t) + D(\theta_i)v_i(t))dt + \sqrt{\mu}D(\theta_i)db_i(t)$$

- The worst-case risk-neutral cost function for agent i

P2: $$J_{2,i}^N(u_i, u_{-i}) = \sup_{v_i \in \mathcal{V}_i} \limsup_{T \to \infty} \frac{1}{T}\mathbb{E}\{\phi_i^2(x, f_N, u, v)\}$$

$$\phi_i^2(x, f_N, u, v) := \int_0^T \|x_i(t) - \frac{1}{N}\sum_{i=1}^N x_i(t)\|_Q^2 + \|u_i(t)\|_R^2 - \gamma^2\|v_i(t)\|^2 dt$$

- v_i can be viewed as a fictitious player (or adversary) of agent i, which strives for a worst-case cost function for agent i
- Agents are coupled with each other through the mean field term

Mean Field Analysis for **P1** and **P2**

- Solve the individual local robust control problem with g instead of f_N

$$\textbf{P1:}\ \bar{J}_1(u,g) = \limsup_{T\to\infty} \frac{\delta}{T} \log \mathbb{E}\{\exp[\frac{1}{\delta}\int_0^T \|x(t)-g(t)\|_Q^2 + \|u(t)\|_R^2 dt]\}$$

$$\textbf{P2:}\ \bar{J}_2(u,v,g) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\{\int_0^T \|x(t)-g(t)\|_Q^2 + \|u(t)\|_R^2 - \gamma^2\|v(t)\|^2 dt\}$$

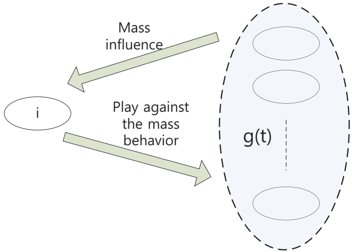

- Characterize g^* that is a *best estimate* of the mean field f_N
  - need to construct a *mean field system* $\mathcal{T}(g)(t)$
  - obtain a fixed point of $\mathcal{T}(g)(t)$, i.e., $g^* = \mathcal{T}(g^*)$

Mean Field Analysis for **P1** and **P2**

- Solve the individual local robust control problem with g instead of f_N

$$\textbf{P1: } \bar{J}_1(u,g) = \limsup_{T\to\infty} \frac{\delta}{T} \log \mathbb{E}\{\exp[\frac{1}{\delta}\int_0^T \|x(t)-g(t)\|_Q^2 + \|u(t)\|_R^2 dt]\}$$

$$\textbf{P2: } \bar{J}_2(u,v,g) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\{\int_0^T \|x(t)-g(t)\|_Q^2 + \|u(t)\|_R^2 - \gamma^2\|v(t)\|^2 dt\}$$

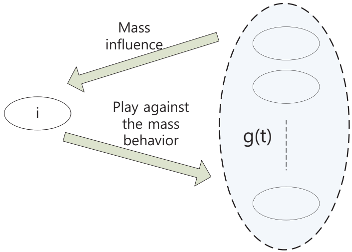

- Characterize g^* that is a *best estimate* of the mean field f_N
  - need to construct a *mean field system* $\mathcal{T}(g)(t)$
  - obtain a fixed point of $\mathcal{T}(g)(t)$, i.e., $g^* = \mathcal{T}(g^*)$

Robust Tracking Control for **P1** and **P2**

Proposition: Individual robust control problems for **P1** and **P2**

Suppose that (A, B) is stabilizable and $(A, Q^{1/2})$ is detectable. Suppose that for a fixed $\gamma = \sqrt{\delta/2\mu} > 0$, there is a matrix $P \geq 0$ that solves the following GARE

$$A^T P + PA + Q - P(BR^{-1}B^T - \frac{1}{\gamma^2}DD^T)P = 0$$

Then **[Rec]**

- $H := A - BR^{-1}B^T P + \frac{1}{\gamma^2}DD^T P$ and $G := A - BR^{-1}B^T P$ are Hurwitz
- The robust decentralized controller: $\bar{u}(t) = -R^{-1}B^T Px(t) - R^{-1}B^T s(t)$ where $\frac{ds(t)}{dt} = -H^T s(t) + Qg(t)$
- The worst-case disturbance (**P2**): $\bar{v}(t) = \gamma^{-2}D^T Px(t) + \gamma^{-2}D^T s(t)$
- $s(t)$ has a unique solution in $\mathcal{C}_n^b$: $s(t) = -\int_t^{\infty} e^{-H^T(t-s)}Qg(s)ds$

Remark

- The two robust tracking problems are identical
- Related to the robust (H^{∞}) control problem w.r.t. γ

Robust Tracking Control for **P1** and **P2**

Proposition: Individual robust control problems for **P1** and **P2**

Suppose that (A, B) is stabilizable and $(A, Q^{1/2})$ is detectable. Suppose that for a fixed $\gamma = \sqrt{\delta/2\mu} > 0$, there is a matrix $P \geq 0$ that solves the following GARE

$$A^T P + PA + Q - P(BR^{-1}B^T - \frac{1}{\gamma^2} DD^T)P = 0$$

Then **[Rec]**

- $H := A - BR^{-1}B^T P + \frac{1}{\gamma^2} DD^T P$ and $G := A - BR^{-1}B^T P$ are Hurwitz
- The robust decentralized controller: $\bar{u}(t) = -R^{-1}B^T Px(t) - R^{-1}B^T s(t)$ where $\frac{ds(t)}{dt} = -H^T s(t) + Qg(t)$
- The worst-case disturbance (**P2**): $\bar{v}(t) = \gamma^{-2} D^T Px(t) + \gamma^{-2} D^T s(t)$
- $s(t)$ has a unique solution in $\mathcal{C}_n^b$: $s(t) = -\int_t^\infty e^{-H^T(t-s)} Qg(s) ds$

Remark

- The two robust tracking problems are identical
- Related to the robust (H^∞) control problem w.r.t. γ

Mean Field Analysis for **P1** and **P2**

- $\bar{x}_\theta(t) = \mathbb{E}\{x_\theta(t)\}$ and we use $h \in \mathcal{C}_n^b$ for **P2**
- Mean field system for **P1** (with the robust decentralized controller) [Rec]

$$\mathcal{T}(g)(t) := \int_{\theta\in\Theta, x\in X} \bar{x}_\theta(t) dF(\theta, x)$$

$$\bar{x}_\theta(t) = e^{G(\theta)t}x + \int_0^t e^{G(\theta)(t-\tau)} B(\theta) R^{-1} B^T(\theta) \times \left(\int_\tau^\infty e^{-H(\theta)^T(\tau-s)} Q g(s) ds \right) d\tau$$

- Mean field system for **P2** (with the robust decentralized controller and the worst-case disturbance) [Rec]

$$\mathcal{L}(h)(t) := \int_{\theta\in\Theta, x\in X} \bar{x}_\theta(t) dF(\theta, x)$$

$$\bar{x}_\theta(t) = e^{H(\theta)t}x + \int_0^t e^{H(\theta)(t-\tau)} \left(B(\theta) R^{-1} B^T(\theta) - \gamma^{-2} D(\theta) D(\theta)^T \right) \times \left(\int_\tau^\infty e^{-H^T(\theta)(\tau-s)} Q h(s) ds \right) d\tau$$

Mean Field Analysis for **P1** and **P2**

- $\bar{x}_\theta(t) = \mathbb{E}\{x_\theta(t)\}$ and we use $h \in \mathcal{C}_n^b$ for **P2**
- Mean field system for **P1** (with the robust decentralized controller) **[Rec]**

$$\mathcal{T}(g)(t) := \int_{\theta\in\Theta, x\in X} \bar{x}_\theta(t) dF(\theta, x)$$

$$\bar{x}_\theta(t) = e^{G(\theta)t}x + \int_0^t e^{G(\theta)(t-\tau)} B(\theta) R^{-1} B^T(\theta)$$

$$\times \left(\int_\tau^\infty e^{-H(\theta)^T(\tau-s)} Qg(s) ds \right) d\tau$$

- Mean field system for **P2** (with the robust decentralized controller and the worst-case disturbance) **[Rec]**

$$\mathcal{L}(h)(t) := \int_{\theta\in\Theta, x\in X} \bar{x}_\theta(t) dF(\theta, x)$$

$$\bar{x}_\theta(t) = e^{H(\theta)t}x + \int_0^t e^{H(\theta)(t-\tau)} \left(B(\theta) R^{-1} B^T(\theta) - \gamma^{-2} D(\theta) D(\theta)^T \right)$$

$$\times \left(\int_\tau^\infty e^{-H^T(\theta)(\tau-s)} Qh(s) ds \right) d\tau$$

Mean Field Analysis for **P1** and **P2**

- $\mathcal{T}(g)(t)$ and $\mathcal{L}(h)(t)$ capture the mass behavior when N is large
- Simplest case

$$\lim_{N\to\infty} f_N(t) = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} x_i(t) = \mathbb{E}\{x_i(t)\} = \mathcal{T}(g)(t), \quad \text{SLLN}$$

- We need to seek g^* and h^* such that $g^* = \mathcal{T}(g^*)$ and $h^* = \mathcal{L}(h^*)$
- Sufficient condition (due to the contraction mapping theorem) **[Rec]**

$$\textbf{P1}: \ \|R^{-1}\|\|Q\| \int_{\theta\in\Theta} \|B(\theta)\|^2 \Big(\int_0^\infty \|e^{G(\theta)\tau}\| d\tau\Big)\Big(\int_0^\infty \|e^{H(\theta)\tau}\| d\tau\Big) dF(\theta) < 1$$

$$\textbf{P2}: \ \int_{\theta\in\Theta} \Big(\int_0^\infty \|e^{H(\theta)t}\|^2 dt\Big)^2 \Big(\|B(\theta)\|^2\|R^{-1}\| + \gamma^{-2}\|D(\theta)\|^2\Big) dF(\theta) < 1$$

- $\lim_{k\to\infty} \mathcal{T}^k(g_0) = g^*$ for any $g_0 \in \mathcal{C}_n^b$
- $g^*(t)$ and $h^*(t)$ are best estimates of $f_N(t)$ when N is large
- Generally $g^* \neq h^*$. But when $\gamma \to \infty$, $g^* \equiv h^*$

Mean Field Analysis for **P1** and **P2**

- $\mathcal{T}(g)(t)$ and $\mathcal{L}(h)(t)$ capture the mass behavior when N is large
- Simplest case

$$\lim_{N\to\infty} f_N(t) = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} x_i(t) = \mathbb{E}\{x_i(t)\} = \mathcal{T}(g)(t), \quad \text{SLLN}$$

- We need to seek g^* and h^* such that $g^* = \mathcal{T}(g^*)$ and $h^* = \mathcal{L}(h^*)$
- Sufficient condition (due to the contraction mapping theorem) **[Rec]**

$$\textbf{P1}: \ \|R^{-1}\|\|Q\| \int_{\theta\in\Theta} \|B(\theta)\|^2 \Big(\int_0^\infty \|e^{G(\theta)\tau}\| d\tau\Big)\Big(\int_0^\infty \|e^{H(\theta)\tau}\| d\tau\Big) dF(\theta) < 1$$

$$\textbf{P2}: \ \int_{\theta\in\Theta} \Big(\int_0^\infty \|e^{H(\theta)t}\|^2 dt\Big)^2 \Big(\|B(\theta)\|^2\|R^{-1}\| + \gamma^{-2}\|D(\theta)\|^2\Big) dF(\theta) < 1$$

- $\lim_{k\to\infty} \mathcal{T}^k(g_0) = g^*$ for any $g_0 \in \mathcal{C}_n^b$
- $g^*(t)$ and $h^*(t)$ are best estimates of $f_N(t)$ when N is large
- Generally $g^* \neq h^*$. But when $\gamma \to \infty$, $g^* \equiv h^*$

Mean Field Analysis for **P1** and **P2**

- $\mathcal{T}(g)(t)$ and $\mathcal{L}(h)(t)$ capture the mass behavior when N is large
- Simplest case

$$\lim_{N\to\infty} f_N(t) = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} x_i(t) = \mathbb{E}\{x_i(t)\} = \mathcal{T}(g)(t), \quad \text{SLLN}$$

- We need to seek g^* and h^* such that $g^* = \mathcal{T}(g^*)$ and $h^* = \mathcal{L}(h^*)$
- Sufficient condition (due to the contraction mapping theorem) **[Rec]**

$$\mathbf{P1}: \; \|R^{-1}\|\|Q\| \int_{\theta\in\Theta} \|B(\theta)\|^2 \Big(\int_0^\infty \|e^{G(\theta)\tau}\| d\tau\Big)\Big(\int_0^\infty \|e^{H(\theta)\tau}\| d\tau\Big) dF(\theta) < 1$$

$$\mathbf{P2}: \; \int_{\theta\in\Theta} \Big(\int_0^\infty \|e^{H(\theta)t}\|^2 dt\Big)^2 \Big(\|B(\theta)\|^2\|R^{-1}\| + \gamma^{-2}\|D(\theta)\|^2\Big) dF(\theta) < 1$$

- $\lim_{k\to\infty} \mathcal{T}^k(g_0) = g^*$ for any $g_0 \in \mathcal{C}_n^b$
- $g^*(t)$ and $h^*(t)$ are best estimates of $f_N(t)$ when N is large
- Generally $g^* \neq h^*$. But when $\gamma \to \infty$, $g^* \equiv h^*$

Mean Field Analysis for **P1** and **P2**

- $\mathcal{T}(g)(t)$ and $\mathcal{L}(h)(t)$ capture the mass behavior when N is large
- Simplest case

$$\lim_{N\to\infty} f_N(t) = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} x_i(t) = \mathbb{E}\{x_i(t)\} = \mathcal{T}(g)(t), \quad \text{SLLN}$$

- We need to seek g^* and h^* such that $g^* = \mathcal{T}(g^*)$ and $h^* = \mathcal{L}(h^*)$
- Sufficient condition (due to the contraction mapping theorem) **[Rec]**

$$\textbf{P1}: \; \|R^{-1}\|\|Q\| \int_{\theta\in\Theta} \|B(\theta)\|^2 \Big(\int_0^\infty \|e^{G(\theta)\tau}\|d\tau\Big)\Big(\int_0^\infty \|e^{H(\theta)\tau}\|d\tau\Big) dF(\theta) < 1$$

$$\textbf{P2}: \; \int_{\theta\in\Theta} \Big(\int_0^\infty \|e^{H(\theta)t}\|^2 dt\Big)^2 \Big(\|B(\theta)\|^2\|R^{-1}\| + \gamma^{-2}\|D(\theta)\|^2\Big) dF(\theta) < 1$$

- $\lim_{k\to\infty} \mathcal{T}^k(g_0) = g^*$ for any $g_0 \in \mathcal{C}_n^b$
- $g^*(t)$ and $h^*(t)$ are best estimates of $f_N(t)$ when N is large
- Generally $g^* \neq h^*$. But when $\gamma \to \infty$, $g^* \equiv h^*$

Mean Field Analysis for **P1** and **P2**

- $\mathcal{T}(g)(t)$ and $\mathcal{L}(h)(t)$ capture the mass behavior when N is large
- Simplest case

$$\lim_{N\to\infty} f_N(t) = \lim_{N\to\infty} \frac{1}{N}\sum_{i=1}^{N} x_i(t) = \mathbb{E}\{x_i(t)\} = \mathcal{T}(g)(t), \quad \text{SLLN}$$

- We need to seek g^* and h^* such that $g^* = \mathcal{T}(g^*)$ and $h^* = \mathcal{L}(h^*)$
- Sufficient condition (due to the contraction mapping theorem) **[Rec]**

$$\textbf{P1}: \ \|R^{-1}\|\|Q\| \int_{\theta\in\Theta} \|B(\theta)\|^2 \Big(\int_0^\infty \|e^{G(\theta)\tau}\| d\tau\Big)\Big(\int_0^\infty \|e^{H(\theta)\tau}\| d\tau\Big) dF(\theta) < 1$$

$$\textbf{P2}: \ \int_{\theta\in\Theta} \Big(\int_0^\infty \|e^{H(\theta)t}\|^2 dt\Big)^2 \Big(\|B(\theta)\|^2\|R^{-1}\| + \gamma^{-2}\|D(\theta)\|^2\Big) dF(\theta) < 1$$

- $\lim_{k\to\infty} \mathcal{T}^k(g_0) = g^*$ for any $g_0 \in \mathcal{C}_n^b$
- $g^*(t)$ and $h^*(t)$ are best estimates of $f_N(t)$ when N is large
- Generally $g^* \not\equiv h^*$. But when $\gamma \to \infty$, $g^* \equiv h^*$

Main Results for **P1** and **P2**

Existence and Characterization of an ϵ-Nash equilibrium

- There exists an ϵ-Nash equilibrium with g^* (**P1**), i.e., there exist $\{u_i^*, 1 \leq i \leq N\}$ and $\epsilon_N \geq 0$ such that

$$J_{1,i}^N(u_i^*, u_{-i}^*) \leq \inf_{u_i \in \mathcal{U}_i^c} J_{1,i}^N(u_i, u_{-i}^*) + \epsilon_N,$$

  where $\epsilon_N \to 0$ as $N \to \infty$. For the uniform agent case, $\epsilon_N = O(1/\sqrt{N})$
- The ϵ-Nash strategy u_i^* is decentralized, i.e., u_i^* is a function of x_i and g^*
- Law of Large Numbers: g^* satisfies **[Rec]**

$$\lim_{N\to\infty} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \ \forall T \geq 0, \ \text{a.s.}$$

$$\lim_{N\to\infty} \limsup_{T\to\infty} \frac{1}{T} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \ \text{a.s.}$$

  g^*: deterministic function and can be computed offline
- The same results also hold for **P2** with the worst-case disturbance

Main Results for **P1** and **P2**

Existence and Characterization of an ϵ-Nash equilibrium

- There exists an ϵ-Nash equilibrium with g^* (**P1**), i.e., there exist $\{u_i^*, 1 \leq i \leq N\}$ and $\epsilon_N \geq 0$ such that

$$J_{1,i}^N(u_i^*, u_{-i}^*) \leq \inf_{u_i \in \mathcal{U}_i^c} J_{1,i}^N(u_i, u_{-i}^*) + \epsilon_N,$$

  where $\epsilon_N \to 0$ as $N \to \infty$. For the uniform agent case, $\epsilon_N = O(1/\sqrt{N})$

- The ϵ-Nash strategy u_i^* is decentralized, i.e., u_i^* is a function of x_i and g^*
- Law of Large Numbers: g^* satisfies **[Rec]**

$$\lim_{N\to\infty} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \ \forall T \geq 0, \text{ a.s.}$$

$$\lim_{N\to\infty} \limsup_{T\to\infty} \frac{1}{T} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \text{ a.s.}$$

  g^*: deterministic function and can be computed offline
- The same results also hold for **P2** with the worst-case disturbance

Main Results for **P1** and **P2**

Existence and Characterization of an ϵ-Nash equilibrium

- There exists an ϵ-Nash equilibrium with g^* (**P1**), i.e., there exist $\{u_i^*, 1 \leq i \leq N\}$ and $\epsilon_N \geq 0$ such that

$$J_{1,i}^N(u_i^*, u_{-i}^*) \leq \inf_{u_i \in \mathcal{U}_i^c} J_{1,i}^N(u_i, u_{-i}^*) + \epsilon_N,$$

  where $\epsilon_N \to 0$ as $N \to \infty$. For the uniform agent case, $\epsilon_N = O(1/\sqrt{N})$
- The ϵ-Nash strategy u_i^* is decentralized, i.e., u_i^* is a function of x_i and g^*
- Law of Large Numbers: g^* satisfies **[Rec]**

$$\lim_{N\to\infty} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \ \forall T \geq 0, \text{ a.s.}$$

$$\lim_{N\to\infty} \limsup_{T\to\infty} \frac{1}{T} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \text{ a.s.}$$

  g^*: deterministic function and can be computed offline
- The same results also hold for **P2** with the worst-case disturbance

Main Results for **P1** and **P2**

Existence and Characterization of an ϵ-Nash equilibrium

- There exists an ϵ-Nash equilibrium with g^* (**P1**), i.e., there exist $\{u_i^*, 1 \leq i \leq N\}$ and $\epsilon_N \geq 0$ such that

$$J_{1,i}^N(u_i^*, u_{-i}^*) \leq \inf_{u_i \in \mathcal{U}_i^c} J_{1,i}^N(u_i, u_{-i}^*) + \epsilon_N,$$

  where $\epsilon_N \to 0$ as $N \to \infty$. For the uniform agent case, $\epsilon_N = O(1/\sqrt{N})$
- The ϵ-Nash strategy u_i^* is decentralized, i.e., u_i^* is a function of x_i and g^*
- Law of Large Numbers: g^* satisfies **[Rec]**

$$\lim_{N\to\infty} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \ \forall T \geq 0, \ \text{a.s.}$$

$$\lim_{N\to\infty} \limsup_{T\to\infty} \frac{1}{T} \int_0^T \left\| \frac{1}{N} \sum_{i=1}^N x_i^*(t) - g^*(t) \right\|^2 dt = 0, \ \text{a.s.}$$

  g^*: deterministic function and can be computed offline
- The same results also hold for **P2** with the worst-case disturbance

Main Results for **P1** and **P2**

Proof (sketch): Law of large numbers (first part)

$$\int_0^T \|f_N^*(t) - g^*(t)\|^2 dt \leq 2\int_0^T \left\|\frac{1}{N}\sum_{i=1}^{N}(x_i^*(t) - \mathbb{E}\{x_i^*(t)\})\right\|^2 dt$$
$$+ 2T \sup_{t\geq 0}\left\|\mathbb{E}\{x_i^*(t)\} - g^*(t)\right\|^2$$

- The second part is zero (due to the fixed-point theorem)
- $e_i^*(t) = x_i^*(t) - \mathbb{E}\{x_i^*(t)\}$ is a mutually orthogonal random vector with $\mathbb{E}\{e_i^*(t)\} = 0$ and $\mathbb{E}\{\|e_i^*(t)\|^2\} < \infty$ for all i and $t \geq 0$
- Strong law of large numbers $\Rightarrow \lim_{N\to\infty} \|(1/N)\sum_{i=1}^N e_i^*(t)\| = 0$ for all $t \in [0, T]$
- $\|(1/N)\sum_{i=1}^N e_i^*(t)\|^2$, $N \geq 1$, is uniformly integrable on $[0, T]$ for all $T \geq 0$, we have the desired result

Main Results for **P1** and **P2**

Proof (sketch): Law of large numbers (first part)

$$\int_0^T \|f_N^*(t) - g^*(t)\|^2 dt \leq 2\int_0^T \Big\| \frac{1}{N}\sum_{i=1}^N (x_i^*(t) - \mathbb{E}\{x_i^*(t)\}) \Big\|^2 dt$$
$$+ 2T \sup_{t\geq 0} \Big\| \mathbb{E}\{x_i^*(t)\} - g^*(t) \Big\|^2$$

- The second part is zero (due to the fixed-point theorem)
- $e_i^*(t) = x_i^*(t) - \mathbb{E}\{x_i^*(t)\}$ is a mutually orthogonal random vector with $\mathbb{E}\{e_i^*(t)\} = 0$ and $\mathbb{E}\{\|e_i^*(t)\|^2\} < \infty$ for all i and $t \geq 0$
- Strong law of large numbers $\Rightarrow \lim_{N\to\infty} \|(1/N)\sum_{i=1}^N e_i^*(t)\| = 0$ for all $t \in [0, T]$
- $\|(1/N)\sum_{i=1}^N e_i^*(t)\|^2$, $N \geq 1$, is uniformly integrable on $[0, T]$ for all $T \geq 0$, we have the desired result

Main Results for **P1** and **P2**

Proof (sketch): Law of large numbers (first part)

$$\int_0^T \|f_N^*(t) - g^*(t)\|^2 dt \leq 2\int_0^T \Big\| \frac{1}{N}\sum_{i=1}^{N}(x_i^*(t) - \mathbb{E}\{x_i^*(t)\})\Big\|^2 dt$$
$$+ 2T \sup_{t\geq 0}\Big\|\mathbb{E}\{x_i^*(t)\} - g^*(t)\Big\|^2$$

- The second part is zero (due to the fixed-point theorem)
- $e_i^*(t) = x_i^*(t) - \mathbb{E}\{x_i^*(t)\}$ is a mutually orthogonal random vector with $\mathbb{E}\{e_i^*(t)\} = 0$ and $\mathbb{E}\{\|e_i^*(t)\|^2\} < \infty$ for all i and $t \geq 0$
- Strong law of large numbers $\Rightarrow \lim_{N\to\infty} \|(1/N)\sum_{i=1}^{N} e_i^*(t)\| = 0$ for all $t \in [0, T]$
- $\|(1/N)\sum_{i=1}^{N} e_i^*(t)\|^2$, $N \geq 1$, is uniformly integrable on $[0, T]$ for all $T \geq 0$, we have the desired result

Main Results for **P1** and **P2**

Proof (sketch): Law of large numbers (first part)

$$\int_0^T \|f_N^*(t) - g^*(t)\|^2 dt \leq 2\int_0^T \left\| \frac{1}{N} \sum_{i=1}^N (x_i^*(t) - \mathbb{E}\{x_i^*(t)\}) \right\|^2 dt$$

$$+ 2T \sup_{t \geq 0} \left\| \mathbb{E}\{x_i^*(t)\} - g^*(t) \right\|^2$$

- The second part is zero (due to the fixed-point theorem)
- $e_i^*(t) = x_i^*(t) - \mathbb{E}\{x_i^*(t)\}$ is a mutually orthogonal random vector with $\mathbb{E}\{e_i^*(t)\} = 0$ and $\mathbb{E}\{\|e_i^*(t)\|^2\} < \infty$ for all i and $t \geq 0$
- Strong law of large numbers $\Rightarrow \lim_{N\to\infty} \|(1/N)\sum_{i=1}^N e_i^*(t)\| = 0$ for all $t \in [0, T]$
- $\|(1/N)\sum_{i=1}^N e_i^*(t)\|^2$, $N \geq 1$, is uniformly integrable on $[0, T]$ for all $T \geq 0$, we have the desired result

Partial Equivalence and Limiting Behaviors of **P1** and **P2**

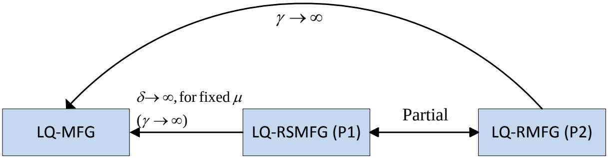

- **P1** and **P2** share the same robust decentralized controller
- Partial equivalence: the mean field systems (and their fixed points) are different
- Limiting behaviors
  - Large deviation (small noise) limit ($\mu, \delta \to 0$ with $\gamma = \sqrt{\delta/2\mu} > 0$): The same results hold under this limit (SDE $\Rightarrow$ ODE)
  - Risk-neutral limit ($\gamma \to \infty$): The results are identical to that of the (risk-neutral) LQ mean field game ($g^* \equiv h^*$)

Partial Equivalence and Limiting Behaviors of **P1** and **P2**

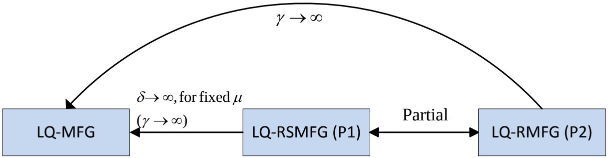

- **P1** and **P2** share the same robust decentralized controller
- Partial equivalence: the mean field systems (and their fixed points) are different
- Limiting behaviors
  - Large deviation (small noise) limit ($\mu, \delta \to 0$ with $\gamma = \sqrt{\delta/2\mu} > 0$): The same results hold under this limit (SDE $\Rightarrow$ ODE)
  - Risk-neutral limit ($\gamma \to \infty$): The results are identical to that of the (risk-neutral) LQ mean field game ($g^* \equiv h^*$)

Partial Equivalence and Limiting Behaviors of **P1** and **P2**

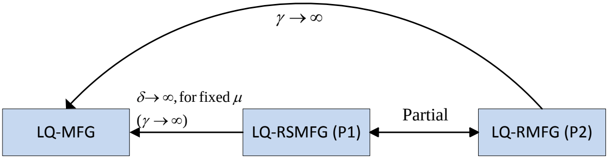

- **P1** and **P2** share the same robust decentralized controller
- Partial equivalence: the mean field systems (and their fixed points) are different
- Limiting behaviors
  - Large deviation (small noise) limit ($\mu, \delta \to 0$ with $\gamma = \sqrt{\delta/2\mu} > 0$): The same results hold under this limit (SDE $\Rightarrow$ ODE)
  - Risk-neutral limit ($\gamma \to \infty$): The results are identical to that of the (risk-neutral) LQ mean field game ($g^* \equiv h^*$)

Simulations ($N = 500$)

- $A_i = \theta_i$ is an i.i.d. uniform random variable with the interval $[2, 5]$, $B = D = Q = R = 1$, $\mu = 2 \Rightarrow \gamma_\theta^* = \gamma^* = 1$, $g^*(t) = 5.086e^{-8.49t}$, $h^*(t) = 5.1e^{-3.37t}$
- $\epsilon^2(N) := \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\int_0^T \|f_N^*(t) - g^*(t)\|^2 dt$
- γ determines robustness of the equilibrium (due to the individual robust control problems)

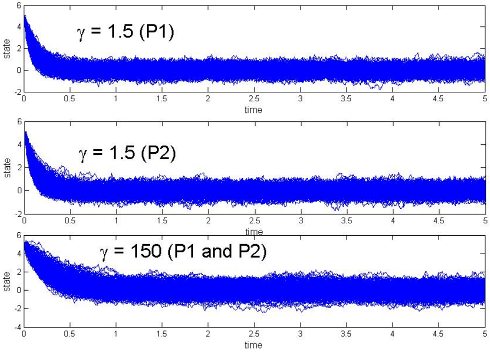

Simulations ($N = 500$)

- $A_i = \theta_i$ is an i.i.d. uniform random variable with the interval $[2, 5]$, $B = D = Q = R = 1$, $\mu = 2 \Rightarrow \gamma_\theta^* = \gamma^* = 1$, $g^*(t) = 5.086e^{-8.49t}$, $h^*(t) = 5.1e^{-3.37t}$
- $\epsilon^2(N) := \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\int_0^T \|f_N^*(t) - g^*(t)\|^2 dt$
- γ determines robustness of the equilibrium (due to the individual robust control problems)

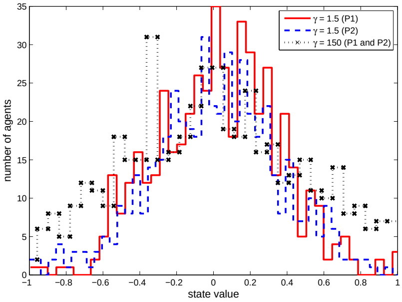

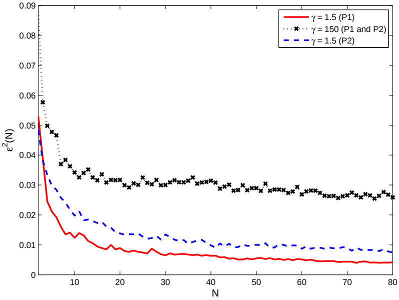

Simulation

- A large number of second-order damping systems: each system follows the approximated mean field behavior

Summary for this part

- LQ risk-sensitive mean field games
- The decentralized ϵ-Nash equilibrium for LQ risk-sensitive mean field games
- ϵ-Nash equilibrium: suboptimal Nash equilibrium, and $\epsilon \to 0$ as $N \to \infty$
- The equilibrium features H^{∞} performance and robustness due to the local risk-sensitive control problem with respect to γ
- Computational complexity
  - Riccati equation P_i: semidefinite programming (efficient!!!)
  - Fixed points g^* and h^*: Picard iteration (efficient!!!)

Outline

5. Risk-Sensitive Mean Field Games via the SMP

Risk-Sensitive Mean Field Games via the SMP

- General nonlinear risk-sensitive mean field games
- Stochastic differential equation of player i, $1 \leq i \leq N$

$$dx_i(t) = f(t, x_i, \nu_N, u_i)dt + \sigma(t)dB_i(t)$$

- ν_N: Empirical distribution (mean field) of the N players

$$\nu_N(t) = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i(t)}(dx)$$

- Risk-sensitive cost function for player i ($\gamma > 0$: robustness)

$$J_i^N(u_i, u_{-i}) = \gamma \log \mathbb{E}\Big[\exp\{\frac{1}{\gamma}\int_0^T l(t, x_i, \nu_N, u_i)dt + \frac{1}{\gamma}m(T, x_i, \nu_N)\}\Big]$$

- Players are coupled with each other through the mean field term

Moon and Başar, DGAA (2018, under review)

Risk-Sensitive Mean Field Games via the SMP

(i) Local risk-sensitive optimization for a fixed probability measure μ

$$\bar{J}(u,\mu) = \gamma \log \mathbb{E}\Big[\exp\{\frac{1}{\gamma}\int_0^T l(t,x,\mu,u)dt + \frac{1}{\gamma}m(T,x,\mu)\}\Big]$$
$$dx(t) = f(t,x,\mu,u)dt + \sigma(t)dW(t)$$

- Solve (i) via the stochastic maximum principle
- Need to analyze the forward-backward SDE (FBSDE)

The corresponding FBSDE **[Rec]**

$$\begin{cases} dx(t) = f(t,x,\mu,u^*)dt + \sigma(t)dW(t) \\ dp(t) = -\Big[f_x^\top(t,x,\mu,u^*)p(t) + l_x(t,x,\mu,u^*) + \frac{1}{\gamma}q(t)\sigma^\top(t)p(t)\Big]dt + q(t)dE \\ x(0) = x_0,\ \ p(T) = m_x(T,x,\mu) \end{cases}$$

- u^* optimal control for $\bar{J}$
- State process x depends on the probability measure μ and the optimal control u^*

Risk-Sensitive Mean Field Games via the SMP

(ii) Fixed point problem: Determine a fixed point of $\mathbb{P}_{x_\mu}$, i.e., $\mu^* = \mathbb{P}_{x_{\mu^*}} = T\mu^*$

- $\mathbb{P}_{x_\mu} = T\mu$: Law of the state process x with the optimal solution determined from step (i)
- Existence of the fixed point: Schauder's fixed point theorem under the 1-Wasserstein metric

$$
\begin{aligned}
W_1(\mu_1, \mu_2) &:= \inf_{\pi \in \Pi(\mu_1, \mu_2)} \int \|x - y\| \pi(dx, dy) \\
&= \sup_{f \in \text{1-Lip}} \left\{ \int f(x) d\mu_1(x) - \int f(x) d\mu_2(x) \right\}
\end{aligned}
$$

- Schauder's fixed point theorem: *Let X be a nonempty closed and bounded convex subset of a normed space S. Let $T : X \to K \subset X$ be continuous, where K is compact. Then T has a fixed point.*

Main Result: ϵ-Nash Equilibrium

The ϵ-Nash Equilibrium

- The set of N optimal distributed controls, $\{u_1^*, \ldots, u_N^*\}$, constitutes an ϵ-Nash equilibrium of the risk-sensitive mean field game. That is, for any u_i,

$$J_i^N(u_1^*, \ldots, u_i^*, \ldots, u_N^*) \leq J_i^N(u_1^*, \ldots, u_i, \ldots, u_N^*) + \epsilon_N, \quad 1 \leq i \leq N$$

where $\epsilon_N \to 0$ as $N \to \infty$ with the convergence rate of $O(\frac{1}{N^{1/(n+4)}})$

Main Result: ϵ-Nash Equilibrium

- ϵ-Nash control u_i^* from (i) and (ii): decentralized (function of x_i and μ^*)
- The equilibrium features robustness under stochastic uncertainty
- The proof follows from **[Rec]**
  - convergence of the 2-Wasserstein distance between the empirical distribution and the fixed point probability measure

$$\mathbb{E}[W_2^2(\nu_N^*(t), \mu^*(t))] = O\left(\frac{1}{N^{2/(n+4)}}\right), \ \forall t \in [0, T]$$

  - the asymptotic analysis for all i:

$$\left| \underbrace{J_i^N(u_i^*, u_{-i}^*)}_{\text{individual NE cost}} - \underbrace{\bar{J}_i(u_i^*, \mu^*)}_{\text{local optimal cost}} \right| \to 0 \ \text{ as } N \to \infty$$

Outline

6 Conclusions

Conclusions

- Risk-sensitive mean field games
- Robust decision analysis in a large population regime
- Local optimization and fixed point problems
- Three different approaches
  - Coupled PDEs (HJB and FPK)
  - Linear-quadratic problem
  - Risk-sensitive maximum principle