

A General Theory for Discrete-Time Mean-Field Games

TAMER BAŞAR

Director, Center for Advanced Study
Dept of ECE and Coordinated Science Laboratory
University of Illinois at Urbana-Champaign

(based on joint work with Naci Saldi and Maxim Raginsky)

IPAM Workshop III on Mean Field Games and Applications

UCLA, CA (Virtual)

4-8 May 2020

Nonzero-sum stochastic dynamic games

- Inherent difficulties in obtaining (and showing existence of) Nash equilibria in stochastic nonzero-sum dynamic games with asymmetric information and a finite number, N , of players
 - ▶ strategic interaction
 - ▶ non-standard dynamic optimization by each player for reaction functions
 - ▶ iterative second guessing
- Mean field game approach provides a way out
- Four recent papers this talk is based on
 - ▶ N. Saldi, T. Başar, M. Raginsky, “Markov-Nash equilibria (MNE) in mean-field games with discounted cost,” *SICON*, 56(6):4256-4287, 2018
 - ▶ — “Approximate Nash Equilibria in partially observed stochastic games with mean-field interactions,” *MOR*, 44(3):1006-1033, 2019.
 - ▶ — “Approximate MNE for DT risk-sensitive MFGs,” *MOR*, 2020.
 - ▶ — “Partially observed DT risk-sensitive MFGs,” 2020 (also CDC’19).

Outline

- **Part I Fully-observed DT mean-field games**
 - ▶ Formulation of the finite-agent game problem
 - ▶ Formulation of the mean-field game problem
 - ▶ Existence of mean-field equilibrium
 - ▶ Existence of approximate Markov-Nash equilibrium
- **Part II Partially-observed DT mean-field games**
 - ▶ Formulation of the finite-agent game problem
 - ▶ Formulation of the mean-field game problem
 - ▶ Existence of mean-field equilibrium
 - ▶ Existence of approximate Nash equilibrium
- **Part III Fully- and partially-observed DT risk-sensitive MFGs**
- **Recap and Conclusions**

Part I

Fully-observed DT mean-field games

Fully-observed DT mean-field game model

- For $i \in \{1, \dots, N\}$, Agent i has the following state dynamics:

$$x_i^N(t+1) = F(x_i^N(t), a_i^N(t), e_t^{(N)}, v_i^N(t)),$$

where the noise process $\{v_i^N(t)\}$ is an i.i.d. sequence and has the same distribution for each i and N .

- $x_i^N(t)$ is the (\mathbf{X} -valued) state variable, $a_i^N(t)$ is the (\mathbf{A} -valued) control action variable, and $e_t^{(N)}$ is the empirical distribution of the state configuration, i.e.

$$e_t^{(N)}(\cdot) := \frac{1}{N} \sum_{i=1}^N \delta_{x_i^N(t)}(\cdot) \in \mathcal{P}(\mathbf{X})$$

- Agent i selects its actions $\{a_i^N(t)\}$ to minimize a functional given in terms of the one-stage cost function $c : \mathbf{X} \times \mathbf{A} \times \mathcal{P}(\mathbf{X}) \rightarrow [0, \infty)$.

Control policies

- History spaces $\mathbf{H}_t = (\mathbf{X} \times \mathbf{A} \times \mathcal{P}(\mathbf{X}))^t \times (\mathbf{X} \times \mathcal{P}(\mathbf{X}))$, $t = 0, 1, 2, \dots$
- A *policy* for a generic agent is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on \mathbf{A} given \mathbf{H}_t .
- A policy is *Markov* if each π_t is a Markov kernel on \mathbf{A} given \mathbf{X} .
- The set of all policies for Agent i is denoted by Π_i and the subset consisting of all Markov policies by \mathbf{M}_i .
- Let $\mathbf{\Pi}^{(N)} = \prod_{i=1}^N \Pi_i$ and $\mathbf{M}^{(N)} = \prod_{i=1}^N \mathbf{M}_i$.

A special class of stochastic NZSDGs

- Dynamics of a generic agent i :

$$\begin{aligned}x_i^N(t+1) &= \frac{1}{N} \sum_{j=1}^N f(x_i^N(t), a_i^N(t), x_j^N(t)) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\ &= \int_{\mathcal{X}} f(x_i^N(t), a_i^N(t), y) e_t^{(N)}(dy) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\ &=: F(x_i^N(t), a_i^N(t), e_t^{(N)}, v_i^N(t)),\end{aligned}$$

where $\{v_i^N(t)\}$ is a sequence of i.i.d. Gaussian random variables.

- One-stage cost function of a generic agent i :

$$\begin{aligned}c(x_i^N(t), a_i^N(t), e_t^{(N)}) &= \frac{1}{N} \sum_{j=1}^N d(x_i^N(t), a_i^N(t), x_j^N(t)) \\ &= \int_{\mathcal{X}} d(x_i^N(t), a_i^N(t), y) e_t^{(N)}(dy)\end{aligned}$$

Cost functions

- For Agent i , the infinite-horizon discounted cost under a policy $\pi^{(N)} \in \Pi^{(N)}$ is given by

$$J_i^{(N)}(\pi^{(N)}) = E^{\pi^{(N)}} \left[\sum_{t=0}^{\infty} \beta^t c(x_i^N(t), a_i^N(t), e_t^{(N)}) \right],$$

where $\beta \in (0, 1)$ is the discount factor.

Nash equilibria

The standard notion of optimality is a player-by-player one:

Definition

A policy $\boldsymbol{\pi}^{(N^*)} = (\pi^{1^*}, \dots, \pi^{N^*})$ constitutes a *Nash equilibrium* if

$$J_i^{(N)}(\boldsymbol{\pi}^{(N^*)}) = \inf_{\pi^i \in \Pi_i} J_i^{(N)}(\boldsymbol{\pi}_{-i}^{(N^*)}, \pi^i)$$

for each $i = 1, \dots, N$, where $\boldsymbol{\pi}_{-i}^{(N^*)} := (\pi^{j^*})_{j \neq i}$.

Nash equilibria

- There are two challenges pertaining to Nash equilibria:
 - ▶ Almost decentralized nature of the information structure of the problem.
 - ▶ *Curse of dimensionality*: the solution of the problem becomes intractable when the numbers of states/actions and/or agents are large.
- Therefore, it is of interest to find an approximate decentralized equilibrium with reduced complexity.
- This is where the MFG approach comes in handy.

(Approximate) Markov-Nash equilibria

To that end, we adopt the following solution concept:

Definition

A policy $\pi^{(N^*)} \in \mathbf{M}^{(N)}$ is a *Markov-Nash equilibrium* if

$$J_i^{(N)}(\pi^{(N^*)}) = \inf_{\pi^i \in \mathbf{M}_i} J_i^{(N)}(\pi_{-i}^{(N^*)}, \pi^i)$$

for each $i = 1, \dots, N$, and an ε -*Markov-Nash equilibrium* (for a given $\varepsilon > 0$) if

$$J_i^{(N)}(\pi^{(N^*)}) \leq \inf_{\pi^i \in \mathbf{M}_i} J_i^{(N)}(\pi_{-i}^{(N^*)}, \pi^i) + \varepsilon$$

for each $i = 1, \dots, N$.

Highlights–Main steps

- A proof of existence of ε -Markov-Nash equilibria for games with sufficiently many agents.
- To this end, we first consider a mean-field game that arises in the infinite-population limit $N \rightarrow \infty$.
- We prove the existence of a mean-field equilibrium for this limiting mean-field game.
- Then, we show that the mean-field equilibrium is an approximate Markov-Nash equilibrium for the original game problem with sufficiently many agents.

Relevant literature

- History and literature on continuous-time mean-field games are well known to this audience.
- There are relatively fewer results on *discrete-time* mean-field games.
- Most earlier works (Gomes et al.'10, Adlakha et al.'15, Elliot et al.'13, Moon and Başar'15, Nourian and Nair'13) consider discrete or linear models.
- Biswas'15 considers average-cost setting with σ -compact Polish state space. But, in this work, agents are only coupled through the cost function.

Mean field game (MFG)

- Mean field game is the infinite-population limit, as $N \rightarrow \infty$, of the original game.
- We have a single agent and model the collective behavior of other agents by an exogenous *state-measure flow* $\boldsymbol{\mu} := (\mu_t)_{t \geq 0} \in \mathcal{P}(\mathbf{X})^\infty$ with a given initial condition μ_0 .
- Given $\boldsymbol{\mu}$, a generic agent has the following state dynamics:

$$x(t+1) = F(x(t), a(t), \mu_t, v(t)).$$

- The generic agent selects its actions $\{a(t)\}$ to minimize a functional given in terms of the one-stage cost function c .

Control policies for MFG

- History spaces $\mathbf{G}_t = (\mathbf{X} \times \mathbf{A})^t \times \mathbf{X}$ for $t = 0, 1, \dots$
- A *policy* is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on \mathbf{A} given \mathbf{G}_t . The set of all policies is denoted by Π .
- A *Markov* policy is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on \mathbf{A} given \mathbf{X} . The set of Markov policies is denoted by \mathbf{M} .

Mean field equilibria

- A policy $\pi^* \in \mathbf{M}$ is optimal for μ if

$$J_{\mu}(\pi^*) = \inf_{\pi \in \Pi} J_{\mu}(\pi),$$

where

$$J_{\mu}(\pi) := E^{\pi} \left[\sum_{t=0}^{\infty} \beta^t c(x(t), a(t), \mu_t) \right].$$

- Let $\mathcal{M} := \{ \mu \in \mathcal{P}(X)^{\infty} : \mu_0 \text{ is fixed} \}$.

Mean field equilibria

- Define the set-valued mapping $\Phi : \mathcal{M} \rightarrow 2^{\mathbf{M}}$ as

$$\Phi(\boldsymbol{\mu}) = \{\pi \in \mathbf{M} : \pi \text{ is optimal for } \boldsymbol{\mu}\}.$$

- Conversely, define a mapping $\Lambda : \mathbf{M} \rightarrow \mathcal{M}$ as follows: given $\pi \in \mathbf{M}$, $\boldsymbol{\mu} := \Lambda(\pi)$ is constructed recursively as

$$\mu_{t+1}(\cdot) = \int_{\mathbf{X} \times \mathbf{A}} p(\cdot | x(t), a(t), \mu_t) \pi_t(da(t) | x(t)) \mu_t(dx(t)).$$

Mean field equilibria

We now introduce the notion of an equilibrium for the mean-field game:

Definition

A pair $(\pi, \mu) \in \mathbf{M} \times \mathcal{M}$ is a *mean-field equilibrium* if $\pi \in \Phi(\mu)$ and $\mu = \Lambda(\pi)$.

Notation

- Let $w : \mathbf{X} \rightarrow \mathbb{R}_+$ be a continuous moment function and $\alpha > 0$.
(We assume $w(x) \geq 1 + d_X(x, x_0)^p$ for some $p \geq 1$ and $x_0 \in \mathbf{X}$; w can be replaced by 1 if c is bounded.)
- For any $g : \mathbf{X} \rightarrow \mathbb{R}$ and $\mu \in \mathcal{P}(\mathbf{X})$, define

$$\|g\|_w := \sup_{x \in \mathbf{X}} \frac{|g(x)|}{w(x)}$$

$$\|\mu\|_w := \sup_{\|g\|_w \leq 1} \left| \int_{\mathbf{X}} g(x) \mu(dx) \right|.$$

- For each $t \geq 0$, define

$$\mathcal{P}_w^t(\mathbf{X}) := \left\{ \mu \in \mathcal{P}(\mathbf{X}) : \int_{\mathbf{X}} w(x) \mu(dx) \leq \alpha^t M \right\},$$

where $\int_{\mathbf{X}} w(x) \mu_0(dx) =: M < \infty$.

Assumption 1

- (a) The cost function c is continuous.
- (b) \mathbf{A} is compact and \mathbf{X} is locally compact.
- (c) The stochastic kernel $p(\cdot|x, a, \mu)$ is weakly continuous.
- (d) We have

$$\sup_{(a, \mu) \in \mathbf{A} \times \mathcal{P}(\mathbf{X})} \int_{\mathbf{X}} w(y) p(dy|x, a, \mu) \leq \alpha w(x).$$

- (e) The function $\int_{\mathbf{X}} w(y) p(dy|x, a, \mu)$ is continuous in (x, a, μ) .
- (f) There exist $\gamma \geq 1$ and a positive real number R such that for each $t \geq 0$, if we define $M_t := \gamma^t R$, then

$$\sup_{(a, \mu) \in \mathbf{A} \times \mathcal{P}_w^t(\mathbf{X})} c(x, a, \mu) \leq M_t w(x).$$

- (g) $\alpha\beta\gamma < 1$.

Revisiting the special class of stochastic NZSDGs

- Dynamics of a generic agent i :

$$\begin{aligned}x_i^N(t+1) &= \frac{1}{N} \sum_{j=1}^N f(x_i^N(t), a_i^N(t), x_j^N(t)) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\ &= \int_{\mathcal{X}} f(x_i^N(t), a_i^N(t), y) e_t^{(N)}(dy) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\ &=: F(x_i^N(t), a_i^N(t), e_t^{(N)}, v_i^N(t)),\end{aligned}$$

where $\{v_i^N(t)\}$ is a sequence of i.i.d. Gaussian random variables.

- One-stage cost function of a generic agent i :

$$\begin{aligned}c(x_i^N(t), a_i^N(t), e_t^{(N)}) &= \frac{1}{N} \sum_{j=1}^N d(x_i^N(t), a_i^N(t), x_j^N(t)) \\ &= \int_{\mathcal{X}} d(x_i^N(t), a_i^N(t), y) e_t^{(N)}(dy)\end{aligned}$$

Assumption 1 vis à vis the special class

Assumption 1 holds for the special class of NZSDGs, with $w(x) = 1 + x^2$, under the following conditions:

- \mathbf{A} is compact.
- g is continuous, and f is bounded and continuous.
- $\sup_{a \in \mathbf{A}} g^2(x, a) \leq Lx^2$ for some $L > 0$.
- $\sup_{(x,a,y) \in K \times \mathbf{A} \times \mathbf{X}} d(x, a, y) < \infty$ for any compact $K \subset \mathbf{X}$.
- $d(x, a, y) \leq R w(x) w(y)$ for some $R > 0$.
- $\omega_d(r) \rightarrow 0$ as $r \rightarrow 0$, where

$$\omega_d(r) = \sup_{y \in \mathbf{X}} \sup_{|x-x'|+|a-a'| \leq r} |d(x, a, y) - d(x', a', y)|$$

- $\alpha^2 \beta < 1$, where $\alpha = \max\{1 + \|f\|, L\}$

Existence of MF equilibrium

Theorem 1

Under Assumption 1, the mean-field game admits a mean-field equilibrium (π, μ) .

Sketch of proof (for the case of bounded c):

For any $\nu \in \mathcal{P}(\mathbf{X} \times \mathbf{A})^\infty$, let $J_{*,t}^\nu$ denote the value function at time t of the nonhomogeneous Markov decision process that a generic agent is faced with:

$$J_{*,t}^\nu(x) := \inf_{\pi \in \mathbf{M}} E^\pi \left[\sum_{k=t}^{\infty} \beta^k c(x(t), a(t), \nu_{t,1}) \mid x(t) = x \right].$$

Hence, using dynamic programming,

$$J_{*,t}^\nu(x) = \min_{a \in \mathbf{A}} \left[c(x, a, \nu_{t,1}) + \beta \int_{\mathbf{X}} J_{*,t+1}^\nu(y) p(dy \mid x, a, \nu_{t,1}) \right] =: T_t^\nu J_{*,t+1}^\nu(x).$$

Existence of MF equilibrium

Adopting the technique of *Jovanovic-Rosental (JME'88)*, introduce a set-valued mapping $\Gamma : \mathcal{P}(\mathbf{X} \times \mathbf{A})^\infty \rightarrow 2^{\mathcal{P}(\mathbf{X} \times \mathbf{A})^\infty}$ by

$$\Gamma(\nu) = C(\nu) \cap B(\nu),$$

where

$$C(\nu) := \left\{ \nu' : \nu'_{0,1} = \mu_0, \nu'_{t+1,1}(\cdot) = \int_{\mathbf{X} \times \mathbf{A}} p(\cdot | x, a, \nu_{t,1}) \nu_t(dx, da) \right\}$$

and

$$B(\nu) := \left\{ \nu' : \forall t \geq 0, \nu'_t \left(\left\{ (x, a) : c(x, a, \nu_{t,1}) + \beta \int_{\mathbf{X}} J_{*,t+1}^\nu p(dy|x, a, \nu_{t,1}) = T_t^\nu J_{*,t+1}^\nu(x) \right\} \right) = 1 \right\}.$$

Existence of MF equilibrium

Suppose that Γ has a fixed point $\nu = (\nu_t)_{t \geq 0}$. Construct a Markov policy $\pi = (\pi_t)_{t \geq 0}$ by disintegrating each ν_t as $\nu_t(dx, da) = \nu_{t,1}(dx)\pi_t(da|x)$, and let $\nu_1 = (\nu_{t,1})_{t \geq 0}$. Then the pair (π, ν_1) is a mean-field equilibrium. Hence, to complete the proof it suffices to prove that Γ has a fixed point.

Let us define

$$\mathcal{P}_w^t(\mathbf{X} \times \mathbf{A}) := \{\nu \in \mathcal{P}(\mathbf{X} \times \mathbf{A}) : \nu_1 \in \mathcal{P}_w^t(\mathbf{X})\}.$$

Note that $\mathcal{P}_w^t(\mathbf{X} \times \mathbf{A})$ is compact. Let $\Xi := \prod_{t=0}^{\infty} \mathcal{P}_w^t(\mathbf{X} \times \mathbf{A})$, which is convex and compact with respect to the product topology. One can prove that for any $\nu \in \Xi$, we have $\Gamma(\nu) \subset \Xi$.

Existence of MF equilibrium

The final piece we need in order to deduce the existence of a fixed point of Γ by an appeal to Kakutani's fixed point theorem is the following: The graph of Γ , i.e., the set

$$\text{Gr}(\Gamma) := \{(\nu, \xi) \in \Xi \times \Xi : \xi \in \Gamma(\nu)\},$$

is closed.

Approximate Markov-Nash equilibria

Now the question is: Given a MFE (π, μ) , and an $\varepsilon > 0$, does there exist $N(\varepsilon)$ such that in an N -player game, with $N \geq N(\varepsilon)$, if all players (agents) use the MFE policy π , they would be in ε -Nash equilibrium in Markov policies; that is, with $\boldsymbol{\pi}^{(N^*)} := (\pi, \pi, \dots, \pi)$,

$$J_i^{(N)}(\boldsymbol{\pi}^{(N^*)}) \leq \inf_{\tilde{\pi}^i \in \mathbf{M}_i} J_i^{(N)}(\boldsymbol{\pi}_{-i}^{(N^*)}, \tilde{\pi}^i) + \varepsilon$$

for each $i = 1, \dots, N$?

Existence of approximate Markov-Nash equilibrium

- Define the following moduli of continuity:

$$\omega_p(r) := \sup_{(x,a) \in \mathbf{X} \times \mathbf{A}} \sup_{\substack{\mu, \nu: \\ \rho_w(\mu, \nu) \leq r}} \|p(\cdot | x, a, \mu) - p(\cdot | x, a, \nu)\|_w$$

$$\omega_c(r) := \sup_{(x,a) \in \mathbf{X} \times \mathbf{A}} \sup_{\substack{\mu, \nu: \\ \rho_w(\mu, \nu) \leq r}} |c(x, a, \mu) - c(x, a, \nu)|.$$

- For any function $g : \mathcal{P}_w(\mathbf{X}) \rightarrow \mathbb{R}$, we define the w -norm of g as follows:

$$\|g\|_w^* := \sup_{\mu \in \mathcal{P}_w(\mathbf{X})} \frac{|g(\mu)|}{\int_{\mathbf{X}} w d\mu}.$$

Assumption 2

- (h) $\omega_p(r) \rightarrow 0$ and $\omega_c(r) \rightarrow 0$ as $r \rightarrow 0$. Moreover, for any $\mu \in \mathcal{P}_w(\mathbf{X})$, the functions

$$\omega_p(\rho_w(\cdot, \mu)) : \mathcal{P}_w(\mathbf{X}) \rightarrow \mathbb{R}$$

$$\omega_c(\rho_w(\cdot, \mu)) : \mathcal{P}_w(\mathbf{X}) \rightarrow \mathbb{R}$$

have finite w -norms.

- (i) There exists a positive real number B such that

$$\sup_{(a, \mu) \in \mathbf{A} \times \mathcal{P}_w(\mathbf{X})} \int_{\mathbf{X}} v^2(y) p(dy|x, a, \mu) \leq Bv^2(x).$$

Revisiting the special class of stochastic NZSDGs

- Dynamics of a generic agent i :

$$\begin{aligned}x_i^N(t+1) &= \frac{1}{N} \sum_{j=1}^N f(x_i^N(t), a_i^N(t), x_j^N(t)) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\&= \int_{\mathcal{X}} f(x_i^N(t), a_i^N(t), y) e_t^{(N)}(dy) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\&=: F(x_i^N(t), a_i^N(t), e_t^{(N)}, v_i^N(t)),\end{aligned}$$

where $\{v_i^N(t)\}$ is a sequence of i.i.d. Gaussian random variables.

- One-stage cost function of a generic agent i :

$$\begin{aligned}c(x_i^N(t), a_i^N(t), e_t^{(N)}) &= \frac{1}{N} \sum_{j=1}^N d(x_i^N(t), a_i^N(t), x_j^N(t)) \\&= \int_{\mathcal{X}} d(x_i^N(t), a_i^N(t), y) e_t^{(N)}(dy)\end{aligned}$$

Assumption 2 vis à vis the special class

Assumption 2 holds for the special class of NZSSDGs under the following conditions:

- $d(x, a, y)$ is (uniformly) Hölder continuous in y with exponent p and Hölder constant K_d .
- $f(x, a, y)$ is (uniformly) Hölder continuous in y with exponent p and Hölder constant K_f .
- g is bounded and $\inf_{(x,a) \in X \times A} |g(x, a)| > 0$.

Main result

Theorem 2

Suppose that Assumptions 1 and 2 hold. Then, for any $\varepsilon > 0$, there exists a positive integer $N(\varepsilon)$, such that, for each $N \geq N(\varepsilon)$, the policy $\boldsymbol{\pi}^{(N)} = \{\pi, \pi, \dots, \pi\}$ is an ε -Markov-Nash equilibrium for the game with N agents.

Logic of the proof

- We first show that as $N \rightarrow \infty$ the empirical distribution of the agents' states at each time t converges to a deterministic limit given by the mean-field equilibrium distribution of the state at time t .
 - ▶ This allows us to deduce that the evolution of the state of a generic agent closely tracks the equilibrium state-measure flow in the infinite-population limit.
- We then show that the infinite population limit is insensitive to individual-agent deviations from the mean-field equilibrium policy.

Sketch of the proof (for bounded c)

Let (π, μ) denote the mean-field equilibrium. One can prove that for all $t \geq 0$, $\mathcal{L}(e_t^{(N)}) \rightarrow \delta_{\mu_t}$ weakly in $\mathcal{P}(\mathcal{P}(\mathbf{X}))$, as $N \rightarrow \infty$. Using this, we obtain

$$\lim_{N \rightarrow \infty} J_1^{(N)}(\pi^{(N)}) = J_\mu(\pi) = \inf_{\pi' \in \Pi} J_\mu(\pi'). \quad (1)$$

Sketch of the proof (*continued*)

Let $\{\tilde{\pi}^{(N)}\}_{N \geq 1} \subset \mathbf{M}_1^c$ be an arbitrary sequence of Markov policies for Agent 1 and let $\tilde{e}_t^{(N)}$ denote the corresponding empirical distribution. Then, we can prove that for each $t \geq 0$, $\mathcal{L}(\tilde{e}_t^{(N)}) \rightarrow \delta_{\mu_t}$ weakly in $\mathcal{P}(\mathcal{P}(\mathbf{X}))$, as $N \rightarrow \infty$.

Using this, we obtain

$$\lim_{N \rightarrow \infty} |J_1^{(N)}(\tilde{\pi}^{(N)}, \pi, \dots, \pi) - J_{\mu}(\tilde{\pi}^{(N)})| = 0. \quad (2)$$

Sketch of the proof (*continued*)

Given $\varepsilon > 0$, for each $N \geq 1$, let $\tilde{\pi}^{(N)} \in M_1^c$ be such that

$$J_1^{(N)}(\tilde{\pi}^{(N)}, \pi, \dots, \pi) < \inf_{\pi' \in M_1^c} J_1^{(N)}(\pi', \pi, \dots, \pi) + \frac{\varepsilon}{3}.$$

Then, by (1) and (2), we have

$$\begin{aligned} \lim_{N \rightarrow \infty} J_1^{(N)}(\tilde{\pi}^{(N)}, \pi, \dots, \pi) &= \lim_{N \rightarrow \infty} J_\mu(\tilde{\pi}^{(N)}) \geq \inf_{\pi'} J_\mu(\pi') = J_\mu(\pi) \\ &= \lim_{N \rightarrow \infty} J_1^{(N)}(\pi, \pi, \dots, \pi), \end{aligned}$$

which implies that there exists $N(\varepsilon)$ such that

$$\begin{aligned} \inf_{\pi' \in M_1^c} J_1^{(N)}(\pi', \pi, \dots, \pi) + \varepsilon &\geq J_1^{(N)}(\tilde{\pi}^{(N)}, \pi, \dots, \pi) + \frac{2\varepsilon}{3} \\ &\geq J_\mu(\pi) + \frac{\varepsilon}{3} \geq J_1^{(N)}(\pi, \pi, \dots, \pi) \end{aligned}$$

for all $N \geq N(\varepsilon)$.

Part II

Partially-observed mean-field games

Partially-observed mean-field game model

- For $i \in \{1, \dots, N\}$, Agent i has the following state and observation dynamics:

$$\begin{aligned}x_i^N(t+1) &= F(x_i^N(t), a_i^N(t), e_i^{(N)}, v_i^N(t)), \\y_i^N(t) &= H(x_i^N(t), e_i^{(N)}, w_i^N(t)),\end{aligned}$$

where the noise processes $\{v_i^N(t)\}$ and $\{w_i^N(t)\}$ are i.i.d. sequences and independent of each other.

- $y_i^N(t)$ is the (\mathbf{Y} -valued) observation variable.
- Agent i selects its actions $\{a_i^N(t)\}$ based on its observations to minimize discounted cost given in terms of the one-stage cost function $c : \mathbf{X} \times \mathbf{A} \times \mathcal{P}(\mathbf{X}) \rightarrow [0, \infty)$.

Control policies

- History spaces $\mathbf{H}_t = (\mathbf{Y} \times \mathbf{A})^t \times \mathbf{Y}$, $t = 0, 1, 2, \dots$
- A *policy* for a generic agent is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on A given \mathbf{H}_t .
- The set of all policies for Agent i is denoted by Π_i .
- Let $\Pi^{(N)} = \prod_{i=1}^N \Pi_i$.

The special class with noisy measurements

- Dynamics of a generic agent i :

$$\begin{aligned}x_i^N(t+1) &= \frac{1}{N} \sum_{j=1}^N f(x_i^N(t), a_i^N(t), x_j^N(t)) + g(x_i^N(t), a_i^N(t)) v_i^N(t) \\ &=: F(x_i^N(t), a_i^N(t), e_t^{(N)}, v_i^N(t)), \\ y_i^N(t) &= \frac{1}{N} \sum_{j=1}^N h(x_i^N(t), x_j^N(t)) + w_i^N(t) \\ &=: H(x_i^N(t), e_t^{(N)}, w_i^N(t)),\end{aligned}$$

where $\{v_i^N(t)\}$ and $\{w_i^N(t)\}$ are sequences of i.i.d. Gaussian random variables.

- One-stage cost function of a generic agent i is same as in fully-observed case.

Nash equilibria

- Establishing the existence of Nash equilibria for PO-MFG is difficult due to:
 - ▶ Almost decentralized and noisy nature of the information structure of the problem.
 - ▶ *Curse of dimensionality*: the solution of the problem becomes intractable when the numbers of states/actions/observations and/or agents are large.
- Therefore, it is of interest to find an approximate decentralized equilibrium with reduced complexity.
- Note that even if the number of players is small, it is all but impossible to show even the existence of approximate Nash equilibria for this general class (it is possible for the LQG case – TB (TAC'78)).

Approximate Nash equilibria

Definition

A policy $\pi^{(N^*)} \in \Pi^{(N)}$ is an ε -Nash equilibrium (for a given $\varepsilon > 0$) if

$$J_i^{(N)}(\pi^{(N^*)}) \leq \inf_{\pi^i \in \Pi_i} J_i^{(N)}(\pi_{-i}^{(N^*)}, \pi^i) + \varepsilon$$

for each $i = 1, \dots, N$.

Relevant literature

- In the literature relatively few results are available on PO-MFGs.
- This one appears to be the first one that studies discrete-time PO-MFGs with such generality.
- Most earlier works (Huang et al.'06, Şen and Caines'14,15,16,16, Tang and Meng'16, Huang and Wang'14) consider continuous time models.

Partially-observed mean field game (MFG)

- Similar to fully-observed case, mean field game is the infinite-population limit, as $N \rightarrow \infty$, of the original partially-observed game.
- We have a single agent and model the collective behavior of other agents by an exogenous *state-measure flow* $\mu := (\mu_t)_{t \geq 0} \in \mathcal{P}(\mathbf{X})^\infty$ with a given initial condition μ_0 .
- Given μ , a generic agent has the following state and observation dynamics:

$$\begin{aligned}x(t+1) &= F(x(t), a(t), \mu_t, v(t)), \\y(t) &= H(x(t), \mu_t, w(t)).\end{aligned}$$

- The generic agent selects its actions $\{a(t)\}$ based on its observations to minimize discounted cost given in terms of the one-stage cost function c .

Control policies for MFG

- History spaces $\mathbf{G}_t = (\mathbf{Y} \times \mathbf{A})^t \times \mathbf{Y}$ for $t = 0, 1, \dots$
- A *policy* is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on \mathbf{A} given \mathbf{G}_t . The set of all policies is denoted by Π .

Mean field equilibria

- A policy $\pi^* \in \Pi$ is optimal for μ if

$$J_{\mu}(\pi^*) = \inf_{\pi \in \Pi} J_{\mu}(\pi).$$

- Let $\mathcal{M} := \{\mu \in \mathcal{P}(\mathcal{X})^{\infty} : \mu_0 \text{ is fixed}\}$.

Mean field equilibria

- Define the set-valued mapping $\Phi : \mathcal{M} \rightarrow 2^\Pi$ as

$$\Phi(\boldsymbol{\mu}) = \{\pi \in \Pi : \pi \text{ is optimal for } \boldsymbol{\mu}\}.$$

- Conversely, define a mapping $\Lambda : \Pi \rightarrow \mathcal{M}$ as follows: given $\pi \in \Pi$, $\boldsymbol{\mu} := \Lambda(\pi)$ is constructed recursively as

$$\mu_{t+1}(\cdot) = \int_{\mathbf{X} \times \mathbf{A}} p(\cdot | x(t), a(t), \mu_t) P^\pi(da(t) | x(t)) \mu_t(dx(t)),$$

where $P^\pi(da(t) | x(t))$ is the conditional distribution of $a(t)$ given $x(t)$, under π .

Mean field equilibria

Definition

A pair $(\pi, \mu) \in \Pi \times \mathcal{M}$ is a *mean-field equilibrium* if $\pi \in \Phi(\mu)$ and $\mu = \Lambda(\pi)$.

Assumption 3

- (a) The cost function c is bounded and continuous.
- (b) \mathbf{A} is compact and \mathbf{X} is locally compact.
- (c) The stochastic kernel p is weakly continuous.
- (d) The observation kernel r is continuous in (x, μ) with respect to total variation distance.
- (e) For some $\alpha > 0$ and a continuous moment function $w : \mathbf{X} \rightarrow [0, \infty)$,

$$\sup_{(a, \mu) \in \mathbf{A} \times \mathcal{P}(\mathbf{X})} \int_{\mathbf{X}} w(y) p(dy|x, a, \mu) \leq \alpha w(x).$$

Assumption 3 vis à vis the special class

Assumption 3 holds for the special case with $w(x) = 1 + x^2$ and $\alpha = \max\{1 + \|f\|^2, L\}$ under the following conditions:

- \mathbf{A} is compact.
- d is continuous and bounded.
- g is continuous, and f is bounded and continuous.
- $\sup_{a \in \mathbf{A}} g^2(x, a) \leq Lx^2$ for some $L > 0$.
- h is continuous and bounded.

Existence of MF equilibrium

Theorem 3

Under Assumption 3, the partially-observed mean-field game admits a mean-field equilibrium (π, μ) .

- To establish the existence of mean-field equilibrium, we use fully observed reduction of partially observed control problems and the dynamic programming principle.

Sketch of proof

- Given any measure flow μ , the optimal control problem for the mean-field game (X, A, Y, p, r, c) reduces to partially observed Markov decision process (POMDP).
- We know that this POMDP can be reduced to a fully observed MDP (belief-state MDP) $(Z, A, \{\eta_t^\mu\}_{t \geq 0}, \{C_t^\mu\}_{t \geq 0})$, where $Z = \mathcal{P}(X)$.
- Hence, we can use the technique developed for fully-observed case to prove the existence of mean-field equilibrium.
- However, there is a crucial difference between this problem and the fully-observed version: We do not have explicit analytical expression describing the relation between μ and η_t^μ from which we can deduce the continuity of η_t^μ with respect to μ .

Sketch of proof

- Define the mapping $\mathbf{B} : \mathcal{P}(\mathbf{Z}) \rightarrow \mathcal{P}(\mathbf{X})$ (barycenter of ν) as follows:

$$\mathbf{B}(\nu)(\cdot) = \int_{\mathbf{Z}} z(\cdot) \nu(dz).$$

- Using this definition, for any $\nu \in \mathcal{P}(\mathbf{Z} \times \mathbf{A})^\infty$, we define the measure flow $\mu^\nu \in \mathcal{P}(\mathbf{X})^\infty$ as follows:

$$\mu^\nu = (\mathbf{B}(\nu_{t,1}))_{t \geq 0}.$$

- For any $\nu \in \mathcal{P}(\mathbf{Z} \times \mathbf{A})^\infty$, let $J_{*,t}^\nu$ denote the value function at time t of the belief-state MDP that a generic agent is faced with for μ^ν .

Sketch of proof

- Similar to the fully-observed case, define the set-valued mapping $\Gamma : \mathcal{P}(\mathbf{Z} \times \mathbf{A})^\infty \rightarrow 2^{\mathcal{P}(\mathbf{Z} \times \mathbf{A})^\infty}$ by

$$\Gamma(\nu) = C(\nu) \cap B(\nu),$$

where

$$C(\nu) := \left\{ \nu' : \nu'_{0,1} = \delta_{\mu_0}, \nu'_{t+1,1}(\cdot) = \int_{\mathbf{Z} \times \mathbf{A}} \eta_t^\nu(\cdot | z, a) \nu_t(dz, da) \right\}$$

and

$$B(\nu) := \left\{ \nu' : \forall t \geq 0, \nu'_t \left(\left\{ (z, a) : C_t^\nu(z, a) + \beta \int_{\mathbf{Z}} J_{*,t+1}^\nu \eta_t^\nu(dy | z, a) = T_t^\nu J_{*,t+1}^\nu(z) \right\} \right) = 1 \right\}.$$

Sketch of proof

- We can prove that η_t^ν is continuous with respect to ν (under the weak topology), which is the key element of the overall proof, different from the full-observation case. A sufficient condition for this is Assumption 3(d).
- The rest of the proof is the same as the proof in fully-observed case.

Existence of approximate Markov-Nash equilibrium

- Let (π', μ) be a mean-field equilibrium.
- Define the following moduli of continuity:

$$\omega_p(r) := \sup_{(x,a) \in X \times A} \sup_{\substack{\mu, \nu: \\ d_{BL}(\mu, \nu) \leq r}} \|p(\cdot | x, a, \mu) - p(\cdot | x, a, \nu)\|_w$$

$$\omega_c(r) := \sup_{(x,a) \in X \times A} \sup_{\substack{\mu, \nu: \\ d_{BL}(\mu, \nu) \leq r}} |c(x, a, \mu) - c(x, a, \nu)|.$$

Assumption 4

- (f) $\omega_p(r) \rightarrow 0$ and $\omega_c(r) \rightarrow 0$ as $r \rightarrow 0$.
- (g) For each $t \geq 0$, $\pi'_t : \mathbf{G}_t \rightarrow \mathcal{P}(\mathbf{A})$ is deterministic; that is,
 $\pi'_t(\cdot | g(t)) = \delta_{f_t(g(t))}(\cdot)$ for some measurable function $f_t : \mathbf{G}_t \rightarrow \mathbf{A}$.
- (h) The observation kernel $r(\cdot | x)$ does not depend on the mean-field term.

Assumption 4 vis à vis the special class

Assumption 4 holds for the special case under the following conditions:

- $d(x, a, y)$ is (uniformly) Lipschitz in y with Lipschitz constant K_d .
- $f(x, a, y)$ is (uniformly) Lipschitz in y with Lipschitz constant K_f .
- g is bounded and $\inf_{(x,a) \in X \times A} |g(x, a)| > 0$.
- h is only a function of x .

Main result

Theorem 4

Suppose that Assumptions 3 and 4 hold. Then, for any $\varepsilon > 0$, there exists a positive integer $N(\varepsilon)$, such that, for each $N \geq N(\varepsilon)$, the policy $\boldsymbol{\pi}^{(N)} = \{\pi, \pi, \dots, \pi\}$ is an ε -Nash equilibrium for the game with N agents.

Logic of the proof

- Since π' is deterministic, we can construct another deterministic policy π , which only uses observations, such that (π, μ) is also a mean-field equilibrium.
- Note that π is not necessarily Markovian.
- We construct an equivalent game model whose states are the state of the original model plus the current and past observations.
- In this new model, π automatically becomes Markov.
- Then, we use the proof technique in fully-observed case to show the existence of an approximate Nash equilibrium.

Part III

Fully- and partially-observed discrete-time risk-sensitive MFGs

Fully- and partially-observed DT risk-sensitive MFGs

Main references:

- N. Saldi, T. Başar, M. Raginsky, “Approximate Markov-Nash equilibria for discrete-time risk-sensitive mean-field games,” *Mathematics of Operations Research*, 2020 (to appear); also arXiv: <http://arxiv.org/abs/1808.03929>.
- N. Saldi, T. Başar, M. Raginsky, “Partially observed discrete-time risk-sensitive mean-field games,” submitted 2020 (also *CDC'19*, pp. 317-322).

Partially-observed DT risk-sensitive MFGs

- For $i \in \{1, \dots, N\}$, Agent i has the following state and observation dynamics:

$$\begin{aligned} s_i^N(t+1) &= F(s_i^N(t), a_i^N(t), d_t^{(N)}, v_i^N(t)) \\ y_i^N(t) &= H(s_i^N(t), w_i^N(t)) \end{aligned}$$

where the noise processes $\{v_i^N(t)\}$ and $\{w_i^N(t)\}$ are i.i.d. sequences.

- $s_i^N(t) \in \mathbf{S}$ state variable, $y_i^N(t) \in \mathbf{Y}$ observation variable, $a_i^N(t) \in \mathbf{A}$ control variable, and $d_t^{(N)} \in \mathcal{P}(\mathbf{S})$ empirical distribution of the states, i.e.

$$d_t^{(N)}(\cdot) := \frac{1}{N} \sum_{i=1}^N \delta_{s_i^N(t)}(\cdot) \in \mathcal{P}(\mathbf{S})$$

Control policies

- A generic agent selects its actions $\{a_i^N(t)\}$ to minimize a functional given in terms of the one-stage cost function $m : \mathbf{S} \times \mathbf{A} \times \mathcal{P}(\mathbf{S}) \rightarrow [0, \infty)$.
- A *policy* for a generic agent is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on \mathbf{A} given $Y^{t+1} = \prod_{i=0}^t Y$.
- The set of all policies for Agent i is denoted by Π_i .
- Let $\Pi^{(N)} = \prod_{i=1}^N \Pi_i$.

Risk-sensitive cost function (finite horizon)

- For Agent i , the finite-horizon risk-sensitive cost under a policy $\pi^{(N)} \in \mathbf{\Pi}^{(N)}$ is given by

$$W_i^{(N)}(\pi^{(N)}) = E^{\pi^{(N)}} \left[e^{\lambda \sum_{t=0}^T \beta^t m(s_i^N(t), a_i^N(t), e_i^{(N)})} \right],$$

where $\beta \in (0, 1)$ is the discount factor and $\lambda > 0$ is the risk factor.

Nash and ε -Nash equilibria

Definition

A policy $\pi^{(N^*)} = (\pi^{1^*}, \dots, \pi^{N^*})$ constitutes a *Nash equilibrium* if

$$W_i^{(N)}(\pi^{(N^*)}) = \inf_{\pi^i \in \Pi_i} W_i^{(N)}(\pi_{-i}^{(N^*)}, \pi^i)$$

for each $i = 1, \dots, N$, where $\pi_{-i}^{(N^*)} := (\pi^{j^*})_{j \neq i}$.

Definition

A policy $\pi^{(N^*)} \in \Pi^{(N)}$ is an ε -*Nash equilibrium* (for a given $\varepsilon > 0$) if

$$W_i^{(N)}(\pi^{(N^*)}) \leq \inf_{\pi^i \in \Pi_i} W_i^{(N)}(\pi_{-i}^{(N^*)}, \pi^i) + \varepsilon$$

for each $i = 1, \dots, N$.

Selected literature on RS MFGs

- Mostly in continuous time, initiated by Tembine-Zhu-Başar (IEEE TAC' 14).
- Others (again in CT): Tembine (Automatica' 15); Moon-Başar (IJC' 16; IEEE TAC' 17)
- PO-MF type control (CT): Djehiche-Tembine (edited volume' 16)
- This one appears to be the first to study discrete-time PO-MFGs, and with such generality.

Equivalent game model

- In the risk-sensitive cost, the one-stage cost functions are in a multiplicative form as opposed to the risk-neutral case.
- To overcome this issue, we first construct an equivalent game model, where the cost can be written in an additive form as in risk-neutral case.
- In this new model, time horizon will be $T + 1$ instead of T .
- Then, using results from the earlier part, the existence of an approximate Nash equilibrium can be established.

Equivalent game model

- In this new model, states are the states of the original model plus the one-stage costs incurred up to that time:

$$x_i^N(t) := \left(s_i^N(t), \sum_{k=0}^{t-1} \beta^k m(s_i^N(k), a_i^N(k), d_k^{(N)}) \right).$$

- Let $x_i^N(t) = (s_i^N(t), m_i^N(t)) \in \mathbf{S} \times [0, L] =: \mathbf{X}$, where L is the maximum risk-neutral cost that can be incurred.

Equivalent game model

- Then, state and observation dynamics are given by

$$\begin{aligned}x_i^N(t+1) &= \left(F(s_i^N(t), a_i^N(t), d_t^{(N)}, v_i^N(t)), m_i^N(t) + \beta^t m(s_i^N(t), a_i^N(t), d_t^{(N)}) \right) \\ &=: T(x_i^N(t), a_i^N(t), e_t^{(N)}, v_i^N(t)) \\ y_i^N(t) &= H(x_i^N(t), w_i^N(t)) = H(s_i^N(t), w_i^N(t))\end{aligned}$$

where

$$e_t^{(N)}(\cdot) := \frac{1}{N} \sum_{i=1}^N \delta_{x_i^N(t)}(\cdot) \in \mathcal{P}(\mathbf{X}).$$

Equivalent game model

- For each $t = 0, \dots, T + 1$, one-stage cost function

$c_t : \mathbf{X} \times \mathbf{A} \times \mathcal{P}(\mathbf{X}) \rightarrow [0, \infty)$ is defined as:

$$c_t(x, a, \mu) = c((s, m), a, \mu) := \begin{cases} 0, & \text{if } t \leq T \\ e^{\lambda m}, & \text{if } t = T + 1. \end{cases}$$

- In this new model, finite-horizon risk-sensitive cost function can be written in an additive-form like in risk-neutral case:

$$J_i^{(N)}(\pi^{(N)}) = E^{\pi^{(N)}} \left[\sum_{t=0}^{T+1} c_t(x_i^N(t), a_i^N(t), e_t^{(N)}) \right].$$

Equivalent game model

- The following proposition establishes the equivalence of this new model and the original model.

Proposition 1

For any $\pi^{(N)} \in \Pi^{(N)}$ and $i = 1, \dots, N$, we have $J_i^{(N)}(\pi^{(N)}) = W_i^{(N)}(\pi^{(N)})$.

- By this result, we can consider the new game model in place of the original one.

Mean field game (MFG)

- Mean field game is the infinite-population limit $N \rightarrow \infty$ of the new game model.
- We have a single agent and model the collective behavior of other agents by an exogenous *state-measure flow* $\mu := (\mu_t)_{t=0}^{T+1} \subset \mathcal{P}(\mathbf{X})$.
- Given μ , a generic agent has the following state and observation dynamics:

$$x(t+1) = T(x(t), a(t), \mu_t, v(t))$$

$$y(t) = H(x(t), w(t)).$$

- The generic agent selects its actions $\{a(t)\}$ to minimize a functional given in terms of the one-stage cost functions $\{c_t\}_{t=0}^{T+1}$.

Control policies for MFG

- History spaces $\mathbf{G}_t = (\mathbf{Y} \times \mathbf{A})^t \times \mathbf{Y}$ for $t = 0, 1, \dots$
- A *policy* is a sequence $\pi = \{\pi_t\}$ of stochastic kernels on \mathbf{A} given \mathbf{G}_t . The set of all policies is denoted by Π .

Mean field equilibria

- A policy $\pi^* \in \Pi$ is optimal for μ if

$$J_{\mu}(\pi^*) = \inf_{\pi \in \Pi} J_{\mu}(\pi),$$

where

$$J_{\mu}(\pi) := E^{\pi} \left[\sum_{t=0}^{T+1} c_t(x(t), a(t), \mu_t) \right].$$

- Let \mathcal{M} denote the set of state-measure flows.

Mean field equilibria

- Define the set-valued mapping $\Phi : \mathcal{M} \rightarrow 2^\Pi$ as

$$\Phi(\boldsymbol{\mu}) = \{\pi \in \Pi : \pi \text{ is optimal for } \boldsymbol{\mu}\}.$$

- Conversely, define a mapping $\Lambda : \Pi \rightarrow \mathcal{M}$ as follows: given $\pi \in \Pi$, $\boldsymbol{\mu} := \Lambda(\pi)$ is the sequence of the distributions of the states under the policy π .

Mean field equilibria

Then, the notion of an equilibrium for the mean-field game is (as before):

Definition

A pair $(\pi, \mu) \in \Pi \times \mathcal{M}$ is a *mean-field equilibrium* if $\pi \in \Phi(\mu)$ and $\mu = \Lambda(\pi)$.

Assumption 5

- (a) The cost function m is bounded and continuous.
- (b) The state dynamics is continuous.
- (c) The observation dynamics is continuous in s and the noise process in observation dynamics has a continuous density.
- (d) \mathbf{A} is compact.
- (e) There exist a constant $\alpha \geq 0$ and a continuous moment function $v : \mathbf{S} \rightarrow [1, \infty)$ such that

$$\sup_{(u, \kappa) \in \mathbf{A} \times \mathcal{P}(\mathbf{S})} \int v(F(s, a, \kappa, v)) \mathbb{P}_v(dv) \leq \alpha v(s).$$

Existence of MF equilibrium

Theorem 5

Under Assumption 5, the mean-field game admits a mean-field equilibrium (π_*, μ_*) .

- This result follows from the proof of the earlier Theorem 3 on risk-neutral version of the same problem (showing that under Assumption 5 the corresponding Assumption 3 of Theorem 3 holds when underlying POMDP is the one corresponding to the transformed cost-additive game, converted to fully-observed MDP, and using again Kakutani's fixed point theorem).

Existence of approximate Nash equilibrium

Define the following moduli of continuity:

$$\omega_q(r) := \sup_{(s,a) \in \mathbf{S} \times \mathbf{A}} \sup_{\substack{\mu, \nu: \\ d_{BL}(\mu, \nu) \leq r}} \|p(\cdot | s, a, \mu) - p(\cdot | s, a, \nu)\|_{\text{TV}}$$

$$\omega_m(r) := \sup_{(s,a) \in \mathbf{S} \times \mathbf{A}} \sup_{\substack{\mu, \nu: \\ d_{BL}(\mu, \nu) \leq r}} |m(s, a, \mu) - m(s, a, \nu)|,$$

where

$$p(\cdot | s, a, \mu) = \mathbb{P}_\nu(F(s, a, \mu, \nu) \in \cdot).$$

Assumption 6:

- (a) $\omega_p(r) \rightarrow 0$ and $\omega_m(r) \rightarrow 0$ as $r \rightarrow 0$.
- (b) For each $t \geq 0$, $\pi_{t,*} : \mathbf{X} \rightarrow \mathcal{P}(\mathbf{A})$ is deterministic and weakly continuous.

Main result

Theorem 6

Under Assumptions 5 and 6, for any $\varepsilon > 0$, there exists a positive integer $N(\varepsilon)$, such that, for each $N \geq N(\varepsilon)$, the policy $\boldsymbol{\pi}^{(N)} = \{\pi_*, \pi_*, \dots, \pi_*\}$ is an ε -Nash equilibrium for the new game model with N agents. Since the new game model is equivalent to the original model, the policy $\boldsymbol{\pi}^{(N)} = \{\pi_*, \pi_*, \dots, \pi_*\}$ is also an ε -Nash equilibrium for the original game.

- This result follows from the earlier Theorem 4 on risk-neutral version of the same problem.

Going from finite (T) to infinite horizon

Using the result [SBR, MOR'20] that any infinite-horizon risk-sensitive cost can be approximated by a finite T-horizon one with error bound $\theta\beta^{T+1}$, for some constant $\theta > 0$ depending on the parameters of the game (but not on policies), we have:

Theorem 7

For any $\varepsilon > 0$, choose T such that $\theta\beta^{T+1} < \varepsilon/3$, and let $N(\varepsilon/3)$ be the constant in Theorem 6 for the finite horizon T . Then, for $N \geq N(\varepsilon/3)$, the policy in Theorem 6, extended arbitrarily from $T + 1$ to ∞ is an ε -Nash equilibrium for the infinite-horizon risk-sensitive game with N agents.

Recap and Conclusion

- Existence of approximate Nash equilibria for finite-population game problems of the mean-field type for both fully-observed and partially-observed MFGs, with discounted regular (risk-neutral) and risk-sensitive costs
 - ▶ First established the existence of NE in the limiting MFG.
 - ▶ Then applied this policy to the finite population game and demonstrated that it constitutes an ε -Nash equilibrium for games with a sufficiently large number of agents.
 - ▶ For RS MFGs, used a transformation to convert to risk-neutral MFGs by expanding the state space.
- Extensions
 - ▶ Non-homogeneous MFGs / Average cost MFGs
 - ▶ MFGs with conditional risk measures
 - ▶ Unspecified models, learning and data-driven approaches