Correlated equilibria and mean field games

– Joint work with Markus Fischer (Padova Uni) –

Luciano Campi

London School of Economics and Political Science

1 Introduction

2 Correlated N-player and mean field games

3 Convergence of correlated equilibria in restricted strategies

4 Construction of approximate correlated equilibria

5 Conclusions

Mean field games and N-player games

Mean field games (MFGs), introduced by [Huang et al., 2006] and [Lasry & Lions, 2007], arise as limit systems for certain symmetric stochastic differential N-player games with mean field interaction as the number of players $N \to \infty$.

Mean field games and N-player games

Mean field games (MFGs), introduced by [Huang et al., 2006] and [Lasry & Lions, 2007], arise as limit systems for certain symmetric stochastic differential N-player games with mean field interaction as the number of players $N \to \infty$.

Mean field interaction: each player interacts with her competitors only via the empirical distribution of their positions (closely related to anonymous games).

Mean field games and N-player games

Mean field games (MFGs), introduced by [Huang et al., 2006] and [Lasry & Lions, 2007], arise as limit systems for certain symmetric stochastic differential N-player games with mean field interaction as the number of players $N \to \infty$.

Mean field interaction: each player interacts with her competitors only via the empirical distribution of their positions (closely related to anonymous games).

When $N \to \infty$, one expects the empirical distribution to converge to the law of the "representative player" (Law of Large Numbers aka Propagation of Chaos).

Mean field games and N-player games

Mean field games (MFGs), introduced by [Huang et al., 2006] and [Lasry & Lions, 2007], arise as limit systems for certain symmetric stochastic differential N-player games with mean field interaction as the number of players $N \to \infty$.

Mean field interaction: each player interacts with her competitors only via the empirical distribution of their positions (closely related to anonymous games).

When $N \to \infty$, one expects the empirical distribution to converge to the law of the "representative player" (Law of Large Numbers aka Propagation of Chaos).

In MFGs the "representative player" reacts optimally to the behaviour of the population, which in turn should arise (at equilibrium) by aggregation of all (identical) players best responses.

Connection between MFG and N-player games

Rigorous connection between MFGs and underlying N-player games can be established in two directions:

1. Construction of approximate Nash equilibria for N-player games starting from a solution to the MFG

   (for instance, [Huang et al., 2006], [Carmona & Delarue, 2013], [Gomes et al., 2013],...).

2. Convergence to solutions of the MFG of (approximate) N-player Nash equilibria, as $N \to \infty$.

Crucial, especially in second direction, is the choice of admissible strategies in definition of N-player Nash eq. Standard choices:

- stochastic open-loop;
- feedback or closed-loop Markov over
  - full system state,
  - only individual players' states (restricted or decentralized strategies, "Markov open-loop").

Convergence to the MFG limit

Difficult, especially for non-stationary finite horizon problems:
Convergence of full state Markov feedback Nash equilibria.

- Breakthrough in [Cardaliaguet et al., 2015]: convergence of Nash equilibria through master equation if well-posed, thus under uniqueness of MFG solutions.
- In this situation, also CLT and LDP from the MFG limit: [Cecchin & Pelino, 2017] and [Bayraktar & Cohen, 2017] for finite state games, [Delarue et al., 2018a] in the diffusion setting.

Convergence to the MFG limit

Difficult, especially for non-stationary finite horizon problems: Convergence of full state Markov feedback Nash equilibria.

- Breakthrough in [Cardaliaguet et al., 2015]: convergence of Nash equilibria through master equation if well-posed, thus under uniqueness of MFG solutions.
- In this situation, also CLT and LDP from the MFG limit: [Cecchin & Pelino, 2017] and [Bayraktar & Cohen, 2017] for finite state games, [Delarue et al., 2018a] in the diffusion setting.
- Recent case studies under non-uniqueness: [Nutz et al., 2018] for optimal stopping problems, [Delarue & Foguen Tchuendom, 2018] example of restoration of uniqueness through common noise, [Cecchin et al., 2018], [Bayraktar & Zhang, 2019] for two-state models.

Convergence to the MFG limit

Difficult, especially for non-stationary finite horizon problems: Convergence of full state Markov feedback Nash equilibria.

- Breakthrough in [Cardaliaguet et al., 2015]: convergence of Nash equilibria through master equation if well-posed, thus under uniqueness of MFG solutions.
- In this situation, also CLT and LDP from the MFG limit: [Cecchin & Pelino, 2017] and [Bayraktar & Cohen, 2017] for finite state games, [Delarue et al., 2018a] in the diffusion setting.
- Recent case studies under non-uniqueness: [Nutz et al., 2018] for optimal stopping problems, [Delarue & Foguen Tchuendom, 2018] example of restoration of uniqueness through common noise, [Cecchin et al., 2018], [Bayraktar & Zhang, 2019] for two-state models.
- In [Lacker, 2018], general convergence result in the non-degenerate diffusion setting, but to weak solutions of the MFG; in particular, stochastic flow of measures (even without common noise).

Correlated equilibria: a two-player textbook example

Generalization of Nash eq that allows for correlation between players' strategies due to R. Aumann [Aumann, 1974, Aumann, 1987].

For extension to discrete time stochastic games see, e.g., [Solan, 2001].

Huge literature on correlated equilibria and its generalisations (e.g. communication equilibria): F. Forges, S. Hart, E. Lehrer, R. Myerson ...

Game of "chicken" (Rebel without a cause): D=Dare, S=Swerve

	S	D
S	(6,6)	(2,7)
D	(7,2)	(0,0)

Two pure Nash: (S,D) and (D,S), with resp. payoffs $(2,7)$ and $(7,2)$.

One mixed Nash with payoff $(4\frac{2}{3}, 4\frac{2}{3})$: the players independently of each other select S with prob $2/3$, D with prob $1/3$.

Correlated equilibria: example cont.

Aumann's idea: a mediator or correlation device randomly selects a strategy profile according to some publicly known distribution, then recommends each player in private a strategy according to the profile.

A probability distribution on the space of strategy profiles is a correlated equilibrium (CE) if no player has an incentive to unilaterally deviate from the mediator's recommendation.

Correlated equilibria: example cont.

Aumann's idea: a mediator or correlation device randomly selects a strategy profile according to some publicly known distribution, then recommends each player in private a strategy according to the profile.

A probability distribution on the space of strategy profiles is a correlated equilibrium (CE) if no player has an incentive to unilaterally deviate from the mediator's recommendation.

Nash equilibria, pure or mixed, are correlated equilibria. In the example:

Nash (S,D)

	S	D
S	0	1
D	0	0

Nash (D,S)

	S	D
S	0	0
D	1	0

Mixed Nash

	S	D
S	$4/9$	$2/9$
D	$2/9$	$1/9$

Correlated equilibria: example cont.

Aumann's idea: a mediator or correlation device randomly selects a strategy profile according to some publicly known distribution, then recommends each player in private a strategy according to the profile.

A probability distribution on the space of strategy profiles is a correlated equilibrium (CE) if no player has an incentive to unilaterally deviate from the mediator's recommendation.

Nash equilibria, pure or mixed, are correlated equilibria. In the example:

Nash (S,D)

	S	D
S	0	1
D	0	0

Nash (D,S)

	S	D
S	0	0
D	1	0

Mixed Nash

	S	D
S	$4/9$	$2/9$
D	$2/9$	$1/9$

But also new equilibria:

Convex combinations of NE

	S	D
S	0	α
D	$1-\alpha$	0

Outside convex hull of NE

	S	D
S	$1/3$	$1/3$
D	$1/3$	0

Aim and scope

Consider correlated equilibria (CE) for a simple class of symmetric finite horizon N-player games. Find definition of CE for the limiting MFG ($N \to \infty$).

Justify definition in two ways:

- by showing convergence of N-player CE to the MFG limit
- by constructing approximate CE starting from a solution of the MFG.

Simplifying assumptions:

- dynamics in discrete time over finite horizon
- finite state space, finite set of control actions
- equilibria in restricted strategies.

1. Introduction
2. Correlated N-player and mean field games
3. Convergence of correlated equilibria in restricted strategies
4. Construction of approximate correlated equilibria
5. Conclusions

Set-up

Let $\mathcal{X}$, Γ be finite sets, the set of individual states and control actions, respectively.

Let $T \in \mathbb{N}$ be the time horizon. The dynamics are determined by a meas. system function

$$\Psi \colon \{0, \ldots, T-1\} \times \mathcal{X} \times \mathcal{P}(\mathcal{X}) \times \Gamma \times \mathcal{Z} \to \mathcal{X},$$

where $\mathcal{Z} \doteq [0, 1]$ is the space of noise states, equipped with $\nu \doteq \text{Uniform}([0, 1])$.

Let $\mathcal{R}$ denote the set of admissible individual (restricted) strategies:

$$\mathcal{R} \doteq \{\varphi \colon \{0, \ldots, T-1\} \times \mathcal{X} \to \Gamma\}.$$

Strategies only depend on time and players' own positions: restricted strategies (decentralized Markov strategies).

The N-player game: dynamics I

Let $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$, $\gamma \in \mathcal{P}(\mathcal{R}^N)$, and let $u : \mathcal{R} \to \mathcal{R}$; $\mathfrak{m}_0$ is called an initial distribution, γ a correlated profile, and u a strategy modification.

The N-player game: dynamics I

Let $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$, $\gamma \in \mathcal{P}(\mathcal{R}^N)$, and let $u : \mathcal{R} \to \mathcal{R}$; $\mathfrak{m}_0$ is called an initial distribution, γ a correlated profile, and u a strategy modification.

Let

$$((\Omega, \mathcal{F}, \mathrm{P}), \Phi_1^N, \dots, \Phi_N^N, X^{N,1}, \dots, X^{N,N}, \xi^{N,1}, \dots, \xi^{N,N})$$

be a realization of the triple $(\mathfrak{m}_0, \gamma, u)$ for player i: for all $j \in \{1, \dots, N\}$ the r.v.s

$$\Phi_j^N, \quad X_0^{N,j}, \dots, X_T^{N,j}, \quad \xi_1^{N,j}, \dots, \xi_T^{N,j}$$

takes values in $\mathcal{R}$, $\mathcal{X}$, and $\mathcal{Z}$, respectively, s.t.

- $X_0^{N,1}, \dots, X_0^{N,N}$ are i.i.d. $\sim \mathfrak{m}_0$;
- $\mathrm{P} \circ (\Phi_1^N, \dots, \Phi_N^N)^{-1} = \gamma$;
- $\xi_t^{N,j}$, $j = 1, \dots, N$, $t = 1, \dots, T$, are i.i.d. $\sim \nu$;

Moreover ... (see next slide)

The N-player game: dynamics II

... we also assume:

- $(\xi^{N,j})_{j=1}^N$, $(X_0^{N,j})_{j=1}^N$, and $(\Phi_j^N)_{j=1}^N$ are independent;
- Dynamics: P-a.s., for every $t \in \{0, \ldots, T-1\}$,

$$X_{t+1}^{N,i} = \Psi\left(t, X_t^{N,i}, \mu_t^{N,i}, u \circ \Phi_i^N(t, X_t^{N,i}), \xi_{t+1}^{N,i}\right),$$

$$X_{t+1}^{N,j} = \Psi\left(t, X_t^{N,j}, \mu_t^{N,j}, \Phi_j^N(t, X_t^{N,j}), \xi_{t+1}^{N,j}\right), \quad j \neq i,$$

where $\mu_t^{N,k} \doteq \frac{1}{N-1} \sum_{l \neq k} \delta_{X_t^{N,l}}$.

The N-player game: costs

The costs for player i associated with $(\mathfrak{m}_0, \gamma, u)$ are

$$J_i^N(\mathfrak{m}_0; \gamma, u) \doteq \mathbf{E}\left[\sum_{t=0}^{T-1} f\left(t, X_t^{N,i}, \mu_t^{N,i}, u \circ \Phi_i^N(t, X_t^{N,i})\right) + F\left(X_T^{N,i}, \mu_T^{N,i}\right)\right],$$

where

- $f \colon \{0, \dots, T-1\} \times \mathcal{X} \times \mathcal{P}(\mathcal{X}) \times \Gamma \to \mathbb{R}$ is the running cost,
- $F \colon \mathcal{X} \times \mathcal{P}(\mathcal{X}) \to \mathbb{R}$ is the terminal costs,
- and the expected value **E** is determined according to a realization of $(\mathfrak{m}_0, \gamma, u)$ for player i.

Interpretation: In player i costs, all other players follow the mediator recommendation (values of Φ_j^N, $j \neq i$), while player i applies modified strategy ($u \circ \Phi_i^N$ instead of Φ_i^N).

Player i accepts the mediator's suggestion if $u = \mathrm{Id}$.

The N-player game: correlated equilibria (CE)

Definition 1.

Let $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$, $\varepsilon \geq 0$. A correlated profile $\gamma \in \mathcal{P}(\mathcal{R}^N)$ is called an ε-CE if for every $i \in \{1, \dots, N\}$, every strategy modification u on $\mathcal{R}$,

$$J_i^N(\mathfrak{m}_0; \gamma, \mathrm{Id}) \leq J_i^N(\mathfrak{m}_0; \gamma, u) + \varepsilon.$$

When $\varepsilon = 0$, we say that γ is a CE.

Observations:

1. Here, CE is defined with respect to restricted strategies; analogous definition for full feedback strategies.
2. When γ is a Dirac distribution ($\Phi_1^N, \dots, \Phi_N^N$ constant), then Def 1 reduces to that of a Nash eq in pure (restricted) strategies.
3. When γ has product form ($\Phi_1^N, \dots, \Phi_N^N$ independent), then Def 1 corresponds to Nash in mixed (restricted) strategies.

Proposition.

Let $\mathfrak{m}^N \in \mathcal{P}(\mathcal{X}^N)$ be exchangeable. Then there exists a symmetric CE with initial distribution $\mathfrak{m}^N$.

Sketch of the proof (based on Hart & Schmeidler (1989)):

- Auxiliary 2-player zero-sum game with player I min over symmetric γ, while player II max over modifications u
- Player II pays to player I the amount

$$\sum_{\varphi \in \mathcal{R}^N} \gamma(\varphi) \sum_{u \in \mathcal{U}} \theta(u) \left(J_1^N(\mathfrak{m}^N, \delta_\varphi, u) - J_1^N(\mathfrak{m}^N, \delta_\varphi, \mathrm{Id}) \right).$$

- A symmetric CE is a symmetric strategy for player I that gives a payoff ≥ 0 for any player II strategy.
- Use Minimax Theorem to finds it.

The mean field game: dynamics

Any $\rho \in \mathcal{P}(\mathcal{R} \times \mathcal{P}(\mathcal{X})^{T+1})$ is called correlated flow. Let u be a strategy modification on $\mathcal{R}$.

Let $((\Omega, \mathcal{F}, \mathrm{P}), \Phi, X, \mu, \xi)$ be a realization of the triple $(\mathfrak{m}_0, \rho, u)$:

$$\Phi, \quad X_0, \dots, X_T, \quad \mu_0, \dots, \mu_T, \quad \xi_1, \dots, \xi_T$$

are r.v.s with values in $\mathcal{R}$, $\mathcal{X}$, $\mathcal{P}(\mathcal{X})$, and $\mathcal{Z}$, respectively, such that

- $\mathrm{P} \circ (X_0)^{-1} = \mathfrak{m}_0$
- $\mathrm{P} \circ (\Phi, \mu_0, \dots, \mu_T)^{-1} = \rho$
- ξ_t, $t = 1, \dots, T$, are i.i.d. $\sim \nu$
- ξ, X_0, and (Φ, μ) are independent
- P-a.s., for every $t = 0, \dots, T-1$,

(1) $$X_{t+1} = \Psi\left(t, X_t, \mu_t, u \circ \Phi\left(t, X_t\right), \xi_{t+1}\right).$$

The mean field game: costs

The costs associated with $(\mathfrak{m}_0, \rho, u)$ are given by

$$J(\mathfrak{m}_0; \rho, u) \doteq \mathbf{E}\left[\sum_{t=0}^{T-1} f\left(t, X_t, \mu_t, u \circ \Phi\left(t, X_t\right)\right) + F\left(X_T, \mu_T\right)\right],$$

where the expected value is determined according to a realization of $(\mathfrak{m}_0, \rho, u)$.

The cost functional is well defined, since any two realizations of $(\mathfrak{m}_0, \rho, u)$ generate the same expected value in the definition of $J(\mathfrak{m}_0; \rho, u)$.

The mean field game: correlated solutions

Definition 2.

Let $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$. A correlated flow $\rho \in \mathcal{P}(\mathcal{R} \times \mathcal{P}(\mathcal{X})^{T+1})$ is called a correlated solution of the MFG (in restricted strategies) if:

i *Optimality: For every strategy modification u on $\mathcal{R}$,*

$$J(\mathfrak{m}_0; \rho, \mathrm{Id}) \leq J(\mathfrak{m}_0; \rho, u).$$

ii *Consistency: If $((\Omega, \mathcal{F}, \mathrm{P}), \Phi, X, \mu, \xi)$ is a realization of the triple $(\mathfrak{m}_0, \rho, \mathrm{Id})$, then for every t,*

$$\mu_t(\cdot) = \mathrm{P}\left[X_t \in \cdot \mid \mathcal{F}^{\mu}_{T}\right],$$

where $\mathcal{F}^{\mu}_{T} \doteq \sigma(\mu_s : s = 0, \ldots, T)$.

Correlated solutions and weak solutions

Definition from [Lacker, 2018]; compatible with open-loop formulation from [Lacker, 2016] and work by R. Carmona and F. Delarue:

Definition 2.5. A *weak semi-Markov mean field equilibrium* (or simply a *weak MFE*) is a tuple $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}, W, \alpha^*, X^*, \mu)$, where $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ is a complete filtered probability space and:

(1) μ is a continuous $\mathbb{F}$-adapted $\mathcal{P}(\mathbb{R}^d)$-valued process, W is a $\mathbb{F}$-Brownian motion, and X^* is a continuous $\mathbb{R}^d$-valued $\mathbb{F}$-adapted process with $\mathbb{P} \circ (X_0^*)^{-1} = \lambda$.

(2) $\alpha^* : [0,T] \times \mathbb{R}^d \times C([0,T]; \mathcal{P}(\mathbb{R}^d)) \to A$ is semi-Markov.

(3) X_0^*, μ, and W are independent.

(4) The state equation holds:

$$dX_t^* = b(t, X_t^*, \mu_t, \alpha^*(t, X_t^*, \mu))dt + dW_t.$$

(5) For every alternative semi-Markov $\alpha : [0,T] \times \mathbb{R}^d \times C([0,T]; \mathcal{P}(\mathbb{R}^d)) \to A$ we have

$$\begin{aligned} &\mathbb{E}\left[\int_0^T f(t, X_t^*, \mu_t, \alpha^*(t, X_t^*, \mu)dt + g(X_T^*, \mu_T)\right] \\ &\quad \geq \mathbb{E}\left[\int_0^T f(t, X_t, \mu_t, \alpha(t, X_t, \mu))dt + g(X_T, \mu_T)\right], \end{aligned}$$

where X is the solution (see Remark 2.6 below) of

$$dX_t = b(t, X_t, \mu_t, \alpha(t, X_t, \mu))dt + dW_t, \quad X_0 = X_0^*. \tag{2.2}$$

(6) The consistency condition holds: $\mu_t = \mathbb{P}(X_t^* \in \cdot \mid \mathcal{F}_t^\mu)$ a.s. for each $t \in [0,T]$, where $\mathcal{F}_t^\mu = \sigma(\mu_s : s \leq t)$.

1 Introduction

2 Correlated N-player and mean field games

3 Convergence of correlated equilibria in restricted strategies

4 Construction of approximate correlated equilibria

5 Conclusions

Assumptions

For $N \in \mathbb{N}$, let $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$, $\gamma^N \in \mathcal{P}(\mathcal{R}^N)$, $\varepsilon_N \geq 0$, and $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$. Assume:

A1 Continuity of Ψ: $\exists$ $\boldsymbol{w}\colon [0,\infty) \to [0,1]$ measurable s.t. $\boldsymbol{w}(s) \overset{s\to 0+}{\longrightarrow} 0$ and, for every $(t,x,a) \in \{0,\dots,T-1\} \times \mathcal{X} \times \Gamma$

$$\int_{\mathcal{Z}} \mathbf{1}_{\Psi(t,x,m,a,z)\neq\Psi(t,x,\widetilde{m},a,z)} \nu(dz) \leq \boldsymbol{w}\big(\mathrm{dist}(m,\widetilde{m})\big) \text{ for all } m,\widetilde{m} \in \mathcal{P}(\mathcal{X}).$$

Moreover, for every $t \in \{0,\dots,T-1\}$, every $\tau \in \mathcal{P}(\mathcal{X} \times \Gamma \times \mathcal{P}(\mathcal{X}))$, $\Psi(t,.)$ is $\tau \otimes \nu$-a.e. continuous.

A2 The costs f, F are continuous.

A3 γ^N is a symmetric probability measures, for all $N \in \mathbb{N}$.

A4 Each γ^N is an ε_N-CE in restricted strategies with $\varepsilon_N \to 0$ as $N \to \infty$.

Convergence to the mean field game limit

For $N \in \mathbb{N} \setminus \{1\}$, let

$$\left((\Omega_N, \mathcal{F}_N, \mathrm{P}_N), \Phi_1^N, \ldots, \Phi_N^N, X^{N,1}, \ldots, X^{N,N}, \xi^{N,1}, \ldots, \xi^{N,N}\right)$$

be a realization of the triple $(\mathfrak{m}_0, \gamma^N, \mathrm{Id})$, and set

$$\rho^N \doteq \mathrm{P}_N \circ \left(\Phi_1^N, \mu_0^{N,1}, \ldots, \mu_T^{N,1}\right)^{-1},$$

where $\mu_t^{N,1} = \frac{1}{N-1} \sum_{j=2}^N \delta_{X_t^{N,j}}$. Then:

Theorem 1.

Grant (A1)-(A4). Then $(\rho^N)_{N \in \mathbb{N}}$ is relatively compact in $\mathcal{P}(\mathcal{R} \times \mathcal{P}(\mathcal{X})^{T+1})$, and any limit point is a correlated solution of the MFG with initial law $\mathfrak{m}_0$.

1 Introduction

2 Correlated N-player and mean field games

3 Convergence of correlated equilibria in restricted strategies

4 Construction of approximate correlated equilibria

5 Conclusions

Construction of approximate correlated equilibria

Idea: disintegrate a correlated MFG solution $\rho \in \mathcal{P}(\mathcal{R} \times \mathcal{P}(\mathcal{X})^{T+1})$ as

$$\rho(d\varphi, dm) = \rho_1(\varphi|m)\rho_2(dm).$$

The mediator generates a flow $m \sim \rho_2$, which she uses as correlation device to recommend i.i.d. strategies $\varphi_1, \dots, \varphi_N$ from $\rho_1(\cdot|m)$.

Theorem 2.

Let $\mathfrak{m}_0 \in \mathcal{P}(\mathcal{X})$. Grant (A1) and (A2). Suppose that $\rho \in \mathcal{P}(\mathcal{R} \times \mathcal{P}(\mathcal{X})^{T+1})$ is a correlated solution of the MFG. For $N \in \mathbb{N}$, define $\gamma^N \in \mathcal{P}(\mathcal{R}^N)$ as

$$\gamma^N(C_1 \times \ldots \times C_N) \doteq \int_{\mathcal{P}(\mathcal{X})^{T+1}} \prod_{i=1}^{N} \rho_1(C_i|m)\, \rho_2(dm).$$

Then there exists $\varepsilon_N \geq 0$ such that γ^N is an ε_N-CE with $\varepsilon_N \to 0$ as $N \to \infty$.

1 Introduction

2 Correlated N-player and mean field games

3 Convergence of correlated equilibria in restricted strategies

4 Construction of approximate correlated equilibria

5 Conclusions

Conclusions

- Definition of correlated equilibrium for a simple class of MFGs; extends in some sense concept of weak solution.
- Justification of definition
  - through convergence of N-player correlated equilibria to the MFG limit;
  - through construction of approximate N-player correlated equilibria starting from the MFG.

  Limitation: equilibria only in restricted strategies.

Some open questions:

1. Extension to more general classes, in particular, with infinite set of strategy modifications.
2. Examples.

Bibliography I

R.J. Aumann.
Subjectivity and correlation in randomized strategies.
Journal of Mathematical Economics, 1:67–96, 1974.

R.J. Aumann.
Correlated equilibrium as an expression of Bayesian rationality.
Econometrica, 55:1–18, 1987.

E. Bayraktar and A. Cohen.
Analysis of a finite state many player game using its master equation.
SIAM Journal on Control and Optimization, 56(5): 3538-3568, 2018.

E. Bayraktar and X. Zhang.
On non-uniqueness in mean field games.
arXiv:1908.06207, August 2019.

P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions.
Mean field games.
arXiv:1509.02505, September 2015.

A. Cecchin, P. Dai Pra, M. Fischer, and G. Pelino.
On the convergence problem in mean field games: a two state model without uniqueness.
SIAM J. Control Optim., 57(4):2443–2466, 2019.

A. Cecchin and G. Pelino.
Convergence, fluctuations and large deviations for finite state mean field games via the master equation.
Stochastic Processes and their Applications, pre-published online, `doi.org/10.1016/j.spa.2018.12.002`, December 2018.

Bibliography II

R. Carmona and F. Delarue.
Probabilistic analysis of mean-field games.
SIAM J. Control Optim., 51(4):2705–2734, 2013.

F. Delarue, D. Lacker, and K. Ramanan.
From the master equation to mean field game limit theory: a central limit theorem.
`arXiv:1804.08542 [math.PR]`, April 2018.

F. Delarue, D. Lacker, and K. Ramanan.
From the master equation to mean field game limit theory: large deviations and concentration of measure.
`arXiv:1804.08550 [math.PR]`, April 2018.

F. Delarue and R. Foguen Tchuendom.
Selection of equilibria in a linear quadratic mean field game.
`arXiv:1808.09137 [math.PR]`, August 2018.

M. Fischer.
On the connection between symmetric N-player games and mean field games.
Ann. Appl. Probab., 27(2):757–810, 2017.

D. Gomes, J. Mohr, and R.R. Souza.
Continuous time finite state mean field games.
Appl. Math. Optim., 68:99–143, 2013.

M. Huang, R.P. Malhamé, and P.E. Caines.
Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle.
Commun. Inf. Syst., 6(3):221–252, 2006.

Bibliography III

D. Lacker.
A general characterization of the mean field limit for stochastic differential games.
Probab. Theory Related Fields, 165(3):581–648, 2016.

D. Lacker.
On the convergence of closed-loop Nash equilibria to the mean field game limit.
`arXiv:1808.02745 [math.PR]`, August 2018.

J.-M. Lasry and P.-L. Lions.
Mean field games.
Japan. J. Math., 2(1):229–260, 2007.

M. Nutz, J. San Martin, and X. Tan.
Convergence to the mean field game limit: a case study.
`arXiv:1806.00817 [math.OC]`, June 2018.

E. Solan.
Characterization of correlated equilibria in stochastic games.
Int. J. Game Theory, 30:259–277, 2001.