Coarse correlated equilibria in mean field games

Markus Fischer

University of Padua

Joint work with Luciano Campi and Federico Cannerozzi

“Robust Optimization and Simulation of Complex Stochastic Systems” in honor of Paul Dupuis

ICERM Providence, September 13–15, 2024

Mean field games and N-player games

Mean field games (MFG), as introduced by [Huang, Malhamé, Caines, 2006] and [Lasry & Lions, 2007], arise as limit systems for symmetric non-zero-sum non-cooperative stochastic N-player games as the number of players $N \to \infty$.

Passage to the limit analogous to McKean-Vlasov limit for weakly interacting particle systems; symmetry in the sense of statistically indistinguishable components (exchangeable joint laws).

For games, controlled systems; notion of solution at prelimit level usually (approximate) Nash equilibrium (NE).

Mean field games and N-player games

Mean field games (MFG), as introduced by [Huang, Malhamé, Caines, 2006] and [Lasry & Lions, 2007], arise as limit systems for symmetric non-zero-sum non-cooperative stochastic N-player games as the number of players $N \to \infty$.

Passage to the limit analogous to McKean-Vlasov limit for weakly interacting particle systems; symmetry in the sense of statistically indistinguishable components (exchangeable joint laws).

For games, controlled systems; notion of solution at prelimit level usually (approximate) Nash equilibrium (NE).

Rigorous connection can be established in two directions:

1. Construction of approximate N-player equilibria from MFG.
2. Convergence of N-player equilibria to MFG solutions.

Crucial: choice of admissible strategies.

Mean field systems and N-particle systems

McKean-Vlasov limit for uncontrolled weakly interacting particle systems corresponds to law of large numbers for associated empirical measures.

Large deviations from McKean-Vlasov limit: in Itô diffusion setting, classical results by [Tanaka, 1984], [Dawson & Gärtner, 1987].

Based on Dupuis-Ellis weak convergence approach and [Boué & Dupuis, 1998, Budhiraja & Dupuis, 2000], Laplace principle in [Budhiraja, Dupuis, F., 2012] through controlled martingale problems; connection with mean field type control.

Mean field systems and N-particle systems

McKean-Vlasov limit for uncontrolled weakly interacting particle systems corresponds to law of large numbers for associated empirical measures.

Large deviations from McKean-Vlasov limit: in Itô diffusion setting, classical results by [Tanaka, 1984], [Dawson & Gärtner, 1987].

Based on Dupuis-Ellis weak convergence approach and [Boué & Dupuis, 1998, Budhiraja & Dupuis, 2000], Laplace principle in [Budhiraja, Dupuis, F., 2012] through controlled martingale problems; connection with mean field type control.

Technique useful also for convergence to MFG limit of N-player open-loop Nash equilibria [F., 2017, Lacker, 2016].

MFG convergence problem for closed-loop NE more difficult; for instance: [Cardaliaguet, Delarue, Lasry, Lions, 2019], [Lacker, 2020], [Lacker & Le Flem, 2023], [Djete, 2023].

Large deviations of closed-loop NE from MFG limit: [Delarue, Lacker, Ramanan, 2020], [Cecchin & Pelino, 2019].

Mean field systems and N-particle systems

McKean-Vlasov limit for uncontrolled weakly interacting particle systems corresponds to law of large numbers for associated empirical measures.

Large deviations from McKean-Vlasov limit: in Itô diffusion setting, classical results by [Tanaka, 1984], [Dawson & Gärtner, 1987].

Based on Dupuis-Ellis weak convergence approach and [Boué & Dupuis, 1998, Budhiraja & Dupuis, 2000], Laplace principle in [Budhiraja, Dupuis, F., 2012] through controlled martingale problems; connection with mean field type control.

Technique useful also for convergence to MFG limit of N-player open-loop Nash equilibria [F., 2017, Lacker, 2016].

MFG convergence problem for closed-loop NE more difficult; for instance: [Cardaliaguet, Delarue, Lasry, Lions, 2019], [Lacker, 2020], [Lacker & Le Flem, 2023], [Djete, 2023].

Large deviations of closed-loop NE from MFG limit: [Delarue, Lacker, Ramanan, 2020], [Cecchin & Pelino, 2019].

On connection with mean field type control: [Bezemek & Heldman, 2024].

LDP for mean field multi-scale systems: [Bezemek & Spiliopoulos, 2023].

Correlated equilibria

Basis for MFGs are Nash equilibria in underlying N-player games.

Generalization of Nash equilibrium that allows for correlation between players' strategies without cooperation due to Robert Aumann [Aumann, 1974, Aumann, 1987].

Aumann's idea: introduce mediator or correlation device that randomly selects a strategy profile according to some publicly known distribution, then tells each player in private which strategy she should play.
A correlated equilibrium (CE) is a probability distribution on the space of strategy profiles such that no player has an incentive to unilaterally deviate from the mediator's recommendation after having received it.

Prototypical example: traffic lights.

Correlated equilibria

Basis for MFGs are Nash equilibria in underlying N-player games.

Generalization of Nash equilibrium that allows for correlation between players' strategies without cooperation due to Robert Aumann [Aumann, 1974, Aumann, 1987].

Aumann's idea: introduce mediator or correlation device that randomly selects a strategy profile according to some publicly known distribution, then tells each player in private which strategy she should play.
A correlated equilibrium (CE) is a probability distribution on the space of strategy profiles such that no player has an incentive to unilaterally deviate from the mediator's recommendation after having received it.

Prototypical example: traffic lights.

Generalization [Moulin & Vial, 1978]: Players decide whether to follow the mediator's recommendations before receiving them. This leads to the notion of coarse correlated equilibrium (CCE).

Observe:

$$\text{NE} \subset \text{CE} \subset \text{CCE}.$$

A textbook example [Roughgarden, 2016]

Four player routing game. Each player has to choose exactly one of six edges:

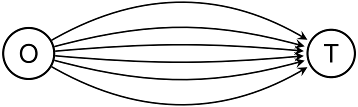

Cost to player i equal to number of players on edge chosen by i: $J_i(\boldsymbol{a}) \doteq \sum_{j=1}^{4} \mathbf{1}_{a_i=a_j}$.

- Pure NE: any set of four distinct edges. Cost 1 to each player.
- Mixed NE: Edges chosen uniformly and independently. Cost $3/2$ to each player.
- CE (not NE): Uniform distribution over those strategy profiles that have exactly three edges occupied. Cost $3/2$ to each player.
- CCE (not CE): Uniform distribution over those strategy profiles that have exactly three edges occupied and occupied edges are either the lower or the upper three. Cost $3/2$ to each player.

(Coarse) correlated equilibria

Advantages of CE and CCE:

- Possibly lower expected costs (higher efficiency).
- Easier to compute [Papadimitriou & Roughgarden, 2008].
- Easier to justify via learning algorithms [Hart & Mas-Colell, 2003].

Applications in economics or computer science; for instance:
[Moulin, Ray, Sen Gupta, 2014] (duopoly game),
[Roughgarden, 2016] (routing games).

Literature on (coarse) correlated equilibria mostly for discrete models, but [Nowak, 1993], [Averboukh, 2019] for continuous time deterministic two-player non-zero-sum differential games.

(Coarse) correlated equilibria

Advantages of CE and CCE:

- Possibly lower expected costs (higher efficiency).
- Easier to compute [Papadimitriou & Roughgarden, 2008].
- Easier to justify via learning algorithms [Hart & Mas-Colell, 2003].

Applications in economics or computer science; for instance:
[Moulin, Ray, Sen Gupta, 2014] (duopoly game),
[Roughgarden, 2016] (routing games).

Literature on (coarse) correlated equilibria mostly for discrete models, but [Nowak, 1993], [Averboukh, 2019] for continuous time deterministic two-player non-zero-sum differential games.

Little for MFGs: [Campi & F., 2022], [Bonesini, Campi, F., 2024], [Muller et al., 2022, Muller et al., 2022+], [Zhao et al., 2024+], all in discrete time with finite state space. More closely related:
[Campi, Cannerozzi, Cartelier, 2023+], [Cannerozzi & Ferrari, 2024+].

Aim and scope

Start from a class of continuous time finite horizon N-player differential games with stochastic dynamics driven by additive Wiener noise.

Consider coarse correlated equilibria in stochastic open-loop strategies.

- Definition of coarse correlated solution for corresponding mean field game; existence of solutions.
- Example of a mean field game possessing explicit non-trivial coarse correlated solutions.
- Construction of approximate N-player CCE from a coarse correlated MFG solution.

Setting

- Finite time horizon: $T > 0$.
- Individual player control actions: $A \subset \mathbb{R}^l$ compact.
- Individual player drift coefficient: $b\colon [0,T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \to \mathbb{R}^d$.
- Individual player cost coefficients: $f\colon [0,T] \times \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \times A \to \mathbb{R}$, $g\colon \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$.
- Initial states and idiosyncratic noise: $(\xi^i)_{i\geq 1}$ i.i.d. $\mathbb{R}^d$-valued random variables with common distribution ν, $(W^i)_{i\geq 1}$ independent d-dimensional standard Wiener processes, all defined on the canonical probability space $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$.
- N-player filtration: $\mathbb{F}^{1,N}$ is the (completed) filtration generated by $(W^1, \ldots, W^N)$ and $(\xi)_{i=1}^N$.
- Admissible N-player strategy profiles: $\mathcal{A}_N^N \doteq \times^N \mathcal{A}_N$ where $\mathcal{A}_N$ space of all $\mathbb{F}^{1,N}$-progressively measurable A-valued processes.

Admissible recommendations

In a pre-game phase, the mediator chooses a probability distribution γ on $\mathcal{A}_N^N$ such that $\gamma = \mathbb{P}^0 \circ \boldsymbol{\Lambda}^{-1}$ for some admissible recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \boldsymbol{\Lambda})$:

- $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is a Polish probability space;
- $\boldsymbol{\Lambda} = (\Lambda^1, \ldots, \Lambda^N)$ is an $\mathcal{A}_N^N$-valued random variable with law $\mathbb{P}^0 \circ \boldsymbol{\Lambda}^{-1} = \gamma$;
- setting $(\Omega, \mathcal{F}, \mathbb{P}) \doteq (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) \otimes (\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$, the A^N-valued process $\boldsymbol{\lambda} = (\lambda_t^1, \ldots, \lambda_t^N)_{t\in[0,T]}$ given by

$$\lambda_t^i(\omega_0, \omega_1) \doteq \Lambda_t^i(\omega_0)(\omega_1), \quad i \in \{1, \ldots, N\},$$

  is progressivly measurable w.r.t. the $\mathbb{P}$-augmented filtration $\mathcal{F}^{0-} \otimes \mathbb{F}^{1,N} = (\mathcal{F}^{0-} \otimes \mathcal{F}_t^{1,N})_{t\in[0,T]}$.

Initial states $\xi^1, \xi^2, \ldots$, Wiener processes $W^1, W^2, \ldots$, and admissible strategies $\beta \in \mathcal{A}_N$ also live on (filtered) product space $(\Omega, \mathcal{F}, \mathbb{P})$.

N-player dynamics and costs

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \mathbf{\Lambda})$ be an admissible recommendation profile.

If all players precommit to mediator's recommendations, then costs for player i given by

$$J_i^N(\mathbf{\Lambda}) \doteq \mathbb{E}_{\mathbb{P}}\left[\int_0^T f(t, X_t^{N,i}, \bar{\mu}_t^N, \lambda_t^i)dt + g(X_T^{N,i}, \bar{\mu}_T^N)\right]$$

with $\bar{\mu}_t^N \doteq \frac{1}{N}\sum_{j=1}^N \delta_{X_t^{N,j}}$ subject to

$$dX_t^{N,j} = b(t, X_t^{N,j}, \bar{\mu}_t^N, \lambda_t^j)dt + dW_t^j, \quad X_0^{N,j} = \xi^j, \quad j \in \{1, \ldots, N\}.$$

N-player dynamics and costs

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \boldsymbol{\Lambda})$ be an admissible recommendation profile.

If all players precommit to mediator's recommendations, then costs for player i given by

$$J_i^N(\boldsymbol{\Lambda}) \doteq \mathbb{E}_{\mathbb{P}}\left[\int_0^T f(t, X_t^{N,i}, \bar{\mu}_t^N, \lambda_t^i)dt + g(X_T^{N,i}, \bar{\mu}_T^N)\right]$$

with $\bar{\mu}_t^N \doteq \frac{1}{N}\sum_{j=1}^N \delta_{X_t^{N,j}}$ subject to

$$dX_t^{N,j} = b(t, X_t^{N,j}, \bar{\mu}_t^N, \lambda_t^j)dt + dW_t^j, \quad X_0^{N,j} = \xi^j, \quad j \in \{1, \ldots, N\}.$$

If player i deviates using $\beta \in \mathcal{A}_N$ while the others precommit, then

$$J_i^N(\beta, \boldsymbol{\Lambda}^{-i}) \doteq \mathbb{E}_{\mathbb{P}}\left[\int_0^T f(t, X_t^{N,i}, \bar{\mu}_t^N, \beta_t)dt + g(X_T^{N,i}, \bar{\mu}_T^N)\right]$$

with $\bar{\mu}_t^N \doteq \frac{1}{N}\sum_{j=1}^N \delta_{X_t^{N,j}}$ subject to

$$dX_t^{N,i} = b(t, X_t^{N,j}, \bar{\mu}_t^N, \beta_t)dt + dW_t^i, \qquad X_0^{N,i} = \xi^i,$$
$$dX_t^{N,j} = b(t, X_t^{N,j}, \bar{\mu}_t^N, \lambda_t^j)dt + dW_t^j, \qquad X_0^{N,j} = \xi^j, \qquad j \neq i.$$

N-player coarse correlated equilibrium

Definition 1.

Let $\varepsilon \geq 0$. An admissible recommendation profile $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \boldsymbol{\Lambda})$ is called an ε-coarse correlated equilibrium (with initial distribution $\otimes^N \nu$) if for every $i \in \{1, \ldots, N\}$, every strategy $\beta \in \mathcal{A}_N$,

$$J_i^N(\boldsymbol{\Lambda}) \leq J_i^N\left(\beta, \boldsymbol{\Lambda}^{-i}\right) + \varepsilon.$$

When $\varepsilon = 0$, we say that $\boldsymbol{\Lambda}$ is a coarse correlated equilibrium.

Observations:

1. Deviations and recommendations in stochastic open-loop strategies.
2. When $\boldsymbol{\Lambda} = (\Lambda^1, \ldots, \Lambda^N)$ is a vector of independent r.v.s, then Definition 1 corresponds to Nash equilibrium in mixed strategies.
3. In a non-coarse correlated equilibrium, deviating player i would choose β in function of Λ_i.

Admissible recommendations

Aim: Limit model for CCE in N-player games when $N \to \infty$.

Let $(\Omega^*, \mathcal{F}^*, \mathbb{F}^*, \mathbb{P}^*)$ be the filtered canonical space for an initial state ξ with distribution ν and a standard d-dimensional Wiener process W. Let $\mathcal{A}$ be the set of admissible strategies, i.e., $\mathbb{F}^*$-progressively measurable A-valued processes.

An admissible recommendation is a pair $((\Omega^0, \mathcal{F}^{0-}, \mathcal{P}^0), \Lambda)$ such that

- $(\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ is a Polish probability space;
- Λ is an $\mathcal{A}$-valued random variable;
- setting $(\Omega, \mathcal{F}, \mathbb{P}) \doteq (\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0) \otimes (\Omega^*, \mathcal{F}^*, \mathbb{P}^*)$, the A-valued process $\lambda = (\lambda_t)_{t\in[0,T]}$ given by

$$\lambda_t(\omega_0, \omega_*) \doteq \Lambda_t(\omega_0)(\omega_*)$$

  is progressivly measurable w.r.t. the $\mathbb{P}$-augmented filtration $\mathbb{F} \doteq \mathcal{F}^{0-} \otimes \mathbb{F}^* = (\mathcal{F}^{0-} \otimes \mathcal{F}^*_t)_{t\in[0,T]}$.

Mean field game dynamics and costs

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ be a correlated flow:

- $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is an admissible recommendation;
- μ is a $\mathbf{C}([0,T], \mathcal{P}_2(\mathbb{R}^d))$-valued random variable on $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ (random flow of measures).

If the representative player precommits to mediator's recommendations, then costs given by

$$J(\Lambda, \mu) \doteq \mathbb{E}_{\mathbb{P}} \left[\int_0^T f(t, X_t, \mu_t, \lambda_t) dt + g(X_T, \mu_T) \right]$$

subject to $dX_t = b(t, X_t, \mu_t, \lambda_t) dt + dW_t$ with $X_0 = \xi$.

Mean field game dynamics and costs

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ be a correlated flow:

- $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda)$ is an admissible recommendation;
- μ is a $\mathbf{C}([0, T], \mathcal{P}_2(\mathbb{R}^d))$-valued random variable on $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0)$ (random flow of measures).

If the representative player precommits to mediator's recommendations, then costs given by

$$J(\Lambda, \mu) \doteq \mathbb{E}_{\mathbb{P}} \left[\int_0^T f(t, X_t, \mu_t, \lambda_t) dt + g(X_T, \mu_T) \right]$$

subject to $dX_t = b(t, X_t, \mu_t, \lambda_t)dt + dW_t$ with $X_0 = \xi$.

If the representative player deviates using $\beta \in \mathcal{A}$, then

$$J(\beta, \mu) \doteq \mathbb{E}_{\mathbb{P}} \left[\int_0^T f(t, X_t, \mu_t, \beta_t) dt + g(X_T, \mu_T) \right]$$

subject to $dX_t = b(t, X_t, \bar{\mu}_t, \beta_t)dt + dW_t$ with $X_0 = \xi$.

Coarse correlated solutions

Definition 2.

A correlated flow $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \Lambda, \mu)$ *is called a coarse correlated solution of the mean field game (with initial distribution* ν*) if the following two conditions hold:*

1. *Optimality: for every strategy* $\beta \in \mathcal{A}$*,*

$$J(\Lambda, \mu) \leq J(\beta, \mu).$$

2. *Consistency: with* X *state process under* Λ*,*

$$\mu_t(\cdot) = \mathbb{P}\left[X_t \in \cdot \mid \mathcal{F}^{\mu}_{T}\right] \quad \text{for all } t \in [0, T].$$

Observations:

1. Consistency condition implies $\mu(t)(\cdot) = \mathrm{P}\left[X(t) \in \cdot \mid \mathcal{F}^{\mu}_{t}\right]$.
2. Classical MFG solution if (Λ, μ) deterministic (Dirac distribution).
3. Flow of measures μ in general non-deterministic although no explicit common noise.

Assumptions

- A1 $A \subset \mathbb{R}^l$ compact;
- A2 initial distribution $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ some $p > 4$;
- A3 b, f, g jointly measurable in all their variables;
- A4 $(x, m, a) \mapsto b(t, x, m, a)$ Lipschitz uniformly in t;
- A5 $t \mapsto (b, f)(t, 0, \delta_0, a_0)$ bounded for some $a_0 \in A$;
- A6 f, g locally Lipschitz of sub-quadratic growth in (x, m, a) (unif. in t).

Extra-assumption for the existence of coarse correlated solutions:

- B For all $(t, x, m) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$,

$$K(t, x, m) \doteq \{(b(t, x, m, a), z) : a \in A, f(t, x, m, a) \leq z\} \subset \mathbb{R}^d \times \mathbb{R}$$

  is closed and convex.

Existence of coarse correlated solutions

Theorem 1.

Grant assumptions A1–A6 and B, then there exists a coarse correlated solution of the mean field game.

Proof based on observation by [Hart & Schmeidler, 1989]:
Let (Λ, μ) be a coarse correlated solution. Then optimality condition:

$$J(\Lambda, \mu) \leq \inf_{\beta \in \mathcal{A}} J(\beta, \mu) \iff \inf_{\beta \in \mathcal{A}} \{J(\beta, \mu) - J(\Lambda, \mu)\} \geq 0.$$

Using a minimax theorem (due to K. Fan), we check that the auxiliary zero-sum game

$$\sup_{(\Lambda,\mu)\text{: admissible and consistent}} \inf_{\beta \in \mathcal{A}} \{J(\beta, \mu) - J(\Lambda, \mu)\}$$

has a value and that $\inf_\beta \sup_{(\Lambda,\mu)} \ldots \geq 0$, thus getting existence of a coarse correlated solution.

Relaxed controls are used to compactify the spaces of strategies in the zero-sum game. Assumption B allows to return to ordinary strategies.

Example

Similar to [Lacker, 2016], [Bardi & F., 2019]: Choose

$$d = 1, \qquad A \doteq [a, b] \text{ with } a < 0 < b, \qquad b(t, x, m, \gamma) = \gamma,$$

$$f \equiv 0, \qquad g(x, m) = -x \cdot \int_{\mathbb{R}} y\, m(dy), \qquad \nu = \delta_0.$$

Costs associated with an admissible recommendation (Λ, μ) thus given by

$$J(\Lambda, \mu) \doteq \mathbb{E}_{\mathbb{P}} \left[-X_T \cdot \int_{\mathbb{R}} y\, \mu_T(dy) \right],$$

subject to

$$dX_t = \Lambda_t dt + dW_t, \quad X_0 = 0.$$

There exist two classical MFG solutions (in open loop strategies):

$$u_t^+ = b, \quad \mu_t^+ = \mathcal{L}(tb + W_t); \qquad u_t^- = a, \quad \mu_t^- = \mathcal{L}(ta + W_t).$$

Example (cont.)

Now, define the following "mixtures":

$$\mu^1 \doteq w_1\mu^+ + (1 - w_1)\mu^-, \qquad \mu^2 \doteq w_2\mu^+ + (1 - w_2)\mu^-,$$

with weights $w_1, w_2 \in [0, 1]$, and set

$$(\Lambda, \mu)((i,j)) = \begin{cases} (u^+, \mu^1) & \text{if } (i,j) = (1,1), \\ (u^+, \mu^2) & \text{if } (i,j) = (1,2), \\ (u^-, \mu^1) & \text{if } (i,j) = (2,1), \\ (u^-, \mu^2) & \text{if } (i,j) = (2,2), \end{cases}$$

$(i,j) \in \Omega^0 \doteq \{1,2\}^2$.

For any $T > 0$, there exist a probability density $(p_{i,j})_{i,j=1,2}$ and weights w_1, w_2 such that (Λ, μ) is a coarse correlated solution.

Consistency condition gives

$$w_1 = \frac{p_{1,1}}{p_{1,1} + p_{2,1}}, \quad w_2 = \frac{p_{1,2}}{p_{1,2} + p_{2,2}}.$$

Example (cont.): infinitely many solutions

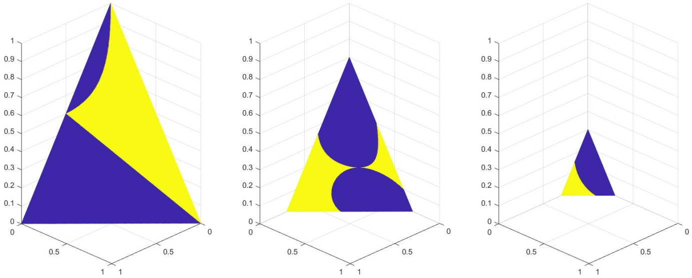

Figure: Case $[a, b] = [-1, 1]$, $T = 2$. Colored points indicate values of $(p_{1,1}, p_{1,2}, p_{2,2})$ such that $p_{1,1} + p_{1,2} + p_{2,2} = 1 - p_{2,1}$ for $p_{2,1} = 0.0$ (left), $p_{2,1} = 0.3$ (center), and $p_{2,1} = 0.7$ (right), respectively. Yellow points correspond to coarse correlated solutions of the mean field game.

Approximate CCE: construction

Aim: Use coarse correlated MFG solution to construct approximate CCE in the N-player game (N large enough).

Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \hat{\Lambda}, \hat{\mu})$ be a coarse correlated MFG solution.

(a) Disintegrate the joint law of $(\hat{\Lambda}, \hat{\mu})$ as

$$\mathbb{P}^0((\hat{\Lambda}, \hat{\mu}) \in C \times B) = \int_B \kappa(C, m)\rho(dm)$$

for some stochastic kernel κ.

(b) Define $\bar{\Omega} = (\times_1^\infty \Omega^0) \times \mathcal{C}([0, T]; \mathcal{P}_2(\mathbb{R}^d))$, $\bar{\mathcal{F}}$ the product σ-field and

$$\bar{\mathbb{P}}(A_1 \times \cdots \times A_N \times B) = \int_B \prod_{i=1}^N K(A_i, m)\rho(dm), \quad N \geq 2,$$

where K is the regular conditional probability of $\mathbb{P}^0$ given $\hat{\mu}$.

Approximate CCE: construction (cont.)

(c) Build random variables $\Lambda^i \colon (\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}}) \to \mathcal{A}(\mathbb{F}^{\xi^i, W^i})$, $i \in \mathbb{N}$, $\mu \colon (\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}}) \to C([0,T]; \mathcal{P}_2(\mathbb{R}^d))$ such that, for all $N \geq 1$,

$$\bar{\mathbb{P}}(\Lambda^1 \in C_1, \dots, \Lambda^N \in C_N, \mu \in B) = \int_B \prod_{i=1}^N \kappa(C_i, m)\rho(dm).$$

Thus, $\bar{\mathbb{P}} \circ (\Lambda^i, \mu)^{-1} = \mathbb{P}^0 \circ (\hat{\Lambda}, \hat{\mu})^{-1}$ for all $i = 1, \dots, N$, and $\Lambda^1, \dots, \Lambda^N$ are conditionally i.i.d. given μ.

Theorem 2.

Assume A1–A6. Let $((\Omega^0, \mathcal{F}^{0-}, \mathbb{P}^0), \hat{\Lambda}, \hat{\mu})$ *be a coarse correlated solution.*

Then there exist admissible recommendations $((\Omega^{0,N}, \mathcal{F}^{0-,N}, \mathbb{P}^{0,N}), \boldsymbol{\Lambda}^N)$ *such that* $\boldsymbol{\Lambda}^N = (\Lambda^1, \dots, \Lambda^N)$ *is an* ε_N*-CCE for the N-player game and* $\varepsilon_N \to 0$ *as* $N \to \infty$*.*

Sketch of proof

By symmetry, focus on deviations of player $i = 1$. Set

$$\varepsilon_N \doteq J_1^N(\mathbf{\Lambda}^N) - \inf_{\beta \in \mathcal{A}_N} J_1^N(\mathbf{\Lambda}^{N,-1}, \beta).$$

Then $\mathbf{\Lambda}^N$ is an ε_N-coarse correlated equilibrium for every $N \geq 2$.
Show that $\varepsilon_N \to 0$.

Sufficient to prove the following:

(1a) $$J(\Lambda^1, \mu) \leq J(\beta^N, \mu),$$

(1b) $$\lim_{N\to\infty} J_1^N(\mathbf{\Lambda}^N) = J(\Lambda^1, \mu),$$

(1c) $$\lim_{N\to\infty} |J_1^N(\beta^N, \mathbf{\Lambda}^{N,-1}) - J(\beta^N, \mu)| = 0.$$

The limits (1b), (1c) can be proved by continuity of the cost functions and propagation of chaos arguments.

Sketch of proof (cont.)

To prove (1a): cannot exploit limit optimality as β^N could depend on all Wiener processes and initial conditions, not only on (ξ^1, W^1).

But conditionally on the paths of $(\xi^i, W^i)_{i=2}^N$ we can compare $J(\beta^N, \mu)$ to $J(\Lambda^1, \mu)$ and conclude:

$$
\begin{aligned}
&J(\Lambda^1, \mu) - J(\beta^N, \mu) \\
&\quad = \mathbb{E}\Bigg[\mathbb{E}\left[\int_0^T f(s, X_s, \mu_s, \lambda_s^1)ds + g(X_T, \mu_T) \mid (\xi^i, W^i)_{i=2}^N\right] \\
&\qquad - \mathbb{E}\left[\int_0^T f(s, Z_s, \mu_s, \beta_s^N)ds + g(Z_T, \mu_T) \mid (\xi^i, W^i)_{i=2}^N\right]\Bigg] \\
&\qquad\qquad = \int \left(J(\hat{\Lambda}, \hat{\mu}) - J(\tilde{\beta}(\mathbf{x}, \mathbf{w}), \hat{\mu})\right) P_\nu(d\mathbf{x}, d\mathbf{w}) \leq 0,
\end{aligned}
$$

where P_ν is the law of $(\xi^i, W^i)_{i=2}^N$ and $\tilde{\beta}(\mathbf{x}, \mathbf{w}) \in \mathcal{A}$.

Conclusions

Thank you.

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 Baysian rationality.
Journal of Mathematical Economics, 1:67–96, 1974.

Y. Averboukh.
Approximate public-signal correlated equilibria for nonzero-sum differential games.
SIAM J. Control Optim., 57(1):743–772, 2019.

M. Bardi and M. Fischer.
On non-uniqueness and uniqueness of solutions in finite-horizon Mean Field Games
ESAIM: COCV, 25: 44, 33 pages, 2019.

Z.W. Bezemek and M. Heldman.
Importance sampling for the empirical measure of weakly interacting diffusions.
Appl. Math. Optim., 89: 7, 59 pages, 2024.

Z.W. Bezemek and K. Spiliopoulos
Large deviations for interacting multiscale particle systems.
Stochastic Process. Appl., 155:27–108, 2023.

O. Bonesini, L. Campi, and M. Fischer.
Correlated equilibria for mean field games with progressive strategies.
Math. Oper. Res., published on-line, 40 pages, April 2024.

Bibliography II

M. Boué and P. Dupuis.
A variational representation for certain functionals of Brownian motion.
Ann. Probab., 26(4):1641–1659, 1998.

A. Budhiraja and P. Dupuis.
A variational representation for positive functionals of infinite dimensional Brownian motion.
Probab. Math. Statist., 20(1):39–61, 2000.

A. Budhiraja, P. Dupuis, and M. Fischer.
Large deviation properties of weakly interacting processes via weak convergence methods.
Ann. Probab., 40(1):74–102, 2012.

L. Campi, F. Cannerozzi, F. Cartelier.
Coarse correlated equilibria in linear quadratic mean field games and application to an emission abatement game
`arXiv:2311.04162 [math.OC]`, December 2023.

L. Campi and M. Fischer.
Correlated equilibria and mean field games: a simple model.
Math. Oper. Res., 47(3):2240–2259, 2022.

F. Cannerozzi and G. Ferrari.
Cooperation, Correlation and Competition in Ergodic N-player Games and Mean-field Games of Singular Controls: A Case Study.
`arXiv:2404.15079 [math.OC]`, April 2024.

P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions.
The Master Equation and the Convergence Problem in Mean Field Games.
Princeton University Press, 2019.

Bibliography III

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 Process. Appl., 129:4510–4555, 2019.

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

D.A. Dawson and J. Gärtner.
Large deviations from the McKean-Vlasov limit for weakly interacting diffusions.
Stochastics, 20(4):247–308, 1987.

F. Delarue, D. Lacker, and K. Ramanan.
From the Master equation to mean field game limit theory: large deviations and concentration of measure.
Ann. Probab., 48(1):211–263, 2020.

M.F. Djete.
Large population games with interactions through controls and common noise: convergence results and equivalence between open-loop and closed-loop controls
ESAIM: COCV, 29: 39, 42 pages, 2023.

T. Dokka, H. Moulin, I. Ray, S. Sen Gupta.
Equilibrium design in an *n*-player quadratic game
Rev Econ Design (2022). https://doi.org/10.1007/s10058-022-00299-2

Bibliography IV

P. Dupuis and R. S. Ellis.
A Weak Convergence Approach to the Theory of Large Deviations.
Wiley Series in Probability and Statistics. John Wiley & Sons, New York, 1997.

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 field2107.03273 games.
Appl. Math. Optim., 68:99–143, 2013.

S. Hart and D. Schmeidler.
Existence of correlated equilibria.
Math. Oper. Res., 14(1):18–25, 1989.

S. Hart and A. Mas-Colell.
Regret-based continuous-time dynamics.
Games and Economic Behavior, 45:375–394, 2003.

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.

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

Bibliography V

D. Lacker.
On the convergence of closed-loop Nash equilibria to the mean field game limit.
Ann. Appl. Probab., 30(4):1693–1761, 2020.

D. Lacker and L. Le Flem.
Closed-loop convergence for the mean field games with common noise.
Ann. Probab., 33(4):2681–2733, 2023.

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

M. Maschler, E. Solan, and S. Zamir.
Game Theory.
Cambridge University Press, 2013.

H. Moulin, I. Ray, and S. Sen Gupta.
Improving Nash by coarse correlation.
Journal of Economic Theory, 150:852–865, 2014.

H. Moulin and J.-P. Vial.
Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon.
Internat. J. Game Theory, 7(3-4):201–221, 1978.

P. Muller, M. Rowland, R. Elie, G. Piliouras, J. Pérolat, M. Laurière, R. Marinier, O. Pietquin, and K. Tuyls.
Learning Equilibria in Mean-Field Games: Introducing Mean-Field PSRO.
Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022), P. Faliszewski, V. Mascardi, C. Pelachaud, M.E. Taylor (eds.), May 9–13, 2022, Online; pp. 926–934.

Bibliography VI

P. Muller, R. Elie, M. Rowland, M. Laurière, J. Pérolat, S. Perrin, M. Geist G. Piliouras, O. Pietquin, and K. Tuyls.
Learning Correlated Equilibria in Mean-Field Games.
arXiv:2208.10138, August 2022.

A.S. Nowak.
Correlated equilibria in nonzero-sum differential games.
J. Mat. Anal. Appl., 174:539–549, 1993.

C.H. Papadimitriou and T. Roughgarden
Computing correlated equilibria in multi-player games.
J. ACM, 55(3): article 14, 29 pages, 2008.

T. Roughgarden.
Twenty Lectures on Algorithmic Game Theory.
Cambridge University Press, 2016.

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

E. Solan.
Correlated equilibrium in stochastic games.
Games Econom. Behav., 38:(2):362–399, 2002.

H. Tanaka.
Limit theorems for certain diffusion processes with interaction.
In K. Itô, editor, *Stochastic Analysis: Proc. Taniguchi Symposium, Katata & Kyoto, 1982*, volume 32 of *North-Holland Mathematical Library*, pages 469–488. Elsevier, Amsterdam, 1984.

Bibliography VII

Z. Zhao, N. Yang, X. Yan, H. Zhang, J. Wang, and Y. Yang.
Correlated Mean Field Imitation Learning
arXiv:2404.09324 [cs.MA], April 2024.

Examples of admissible N-player recommendations

- If Λ takes only $k \geq 1$ values $\alpha^1, \ldots, \alpha^k \in \mathcal{A}_N^N$ with probabilities $\mathbb{P}^0(\Lambda = \alpha^j) = p_j$, then it is admissible. In this case

$$\lambda_t(\omega_0, \omega_1) = \sum_{j=1}^{k} \mathbf{1}_{\{\Lambda = \alpha^j\}}(\omega_0) \alpha_t^j(\omega_1).$$

- More generally, if Λ takes at most countably many values, then it is admissible. We have

$$\lambda_t(\omega_0, \omega_1) = \sum_{j=1}^{\infty} \mathbf{1}_{\{\Lambda = \alpha^j\}}(\omega_0) \alpha_t^j(\omega_1).$$

- Let $(\lambda_t)_{t \in [0,T]}$ be an A-valued progressively measurable as before, define

$$\Lambda(\omega_0) = (\lambda_t(\omega_0, \cdot))_{t \in [0,T]} \mathbf{1}_{N^c}(\omega_0) + a_0 \mathbf{1}_N(\omega_0)$$

for $a_0 \in A$ and N a $\mathbb{P}^0$-null subset of Ω^0. Then Λ is an admissible recommendation