

A general characterization of the mean field limit for stochastic differential games

Daniel Lacker

ORFE, *Princeton University*

January 24, 2015

Section 1

Introduction

Mean field theories beyond physics

Mean field theory in a nutshell: Approximate statistical features of a n -particle system by a ∞ -particle system.

Applications outside of physics:

- ▶ economics & finance (systemic risk, income distribution...)
- ▶ biology (flocking...)
- ▶ sociology (crowd dynamics, voter models...)
- ▶ electrical engineering (telecommunications...)

Mean field theories beyond physics

Mean field theory in a nutshell: Approximate statistical features of a n -particle system by a ∞ -particle system.

Applications outside of physics:

- ▶ economics & finance (systemic risk, income distribution...)
- ▶ biology (flocking...)
- ▶ sociology (crowd dynamics, voter models...)
- ▶ electrical engineering (telecommunications...)

Why mean field game theory?

It replaces **particles** with **rational agents**. Laws of motion emerge in equilibrium and need not be prescribed exogenously.

Mean field theories beyond physics

Main novelties of MFG theory: continuous time (PDEs, SDEs) and **rigorous connection to finite-population models**

Most-studied so far: MFG analogs of McKean-Vlasov interacting diffusion models. \rightsquigarrow **Stochastic differential MFGs**

Some recent literature: MFG analogs of

- ▶ Spin systems (Horst/Scheinkman)
- ▶ Stochastic coalescence (Duffie/Malamud/Manso)

Section 2

A prototypical MFG model

A model of systemic risk: Carmona/Fouque/Sun '13

Mean field model:

n banks with log-monetary reserves $(X_t^i)_{t \in [0, T]}$,

$$dX_t^i = a(\bar{X}_t - X_t^i)dt + \sigma\rho dW_t^i + \sigma\sqrt{1 - \rho^2}dB_t,$$

$$\bar{X}_t = \frac{1}{n} \sum_{k=1}^n X_t^k$$

Rate of borrowing/lending between banks: $a > 0$

A model of systemic risk: Carmona/Fouque/Sun '13

Mean field model:

n banks with log-monetary reserves $(X_t^i)_{t \in [0, T]}$,

$$dX_t^i = a(\bar{X}_t - X_t^i)dt + \sigma\rho dW_t^i + \sigma\sqrt{1 - \rho^2}dB_t,$$

$$\bar{X}_t = \frac{1}{n} \sum_{k=1}^n X_t^k$$

Rate of borrowing/lending between banks: $a > 0$

Goal: Find probabilities of **systemic events** of the form

$$\left\{ \min_{0 \leq t \leq T} \bar{X}_t \leq D \right\}, \quad D = \text{default level.}$$

A model of systemic risk: Carmona/Fouque/Sun '13

Mean field game model:

n banks with log-monetary reserves $(X_t^i)_{t \in [0, T]}$,

$$dX_t^i = [a(\bar{X}_t - X_t^i) + \alpha_t^i] dt + \sigma \rho dW_t^i + \sigma \sqrt{1 - \rho^2} dB_t,$$

$$\bar{X}_t = \frac{1}{n} \sum_{k=1}^n X_t^k$$

Bank i chooses to borrow/lend from a central bank at rate α_t^i ,
to minimize some cost

$$\mathbb{E} \left[\int_0^T f(X_t^i, \bar{X}_t, \alpha_t^i) dt + g(X_T^i, \bar{X}_T) \right].$$

A model of systemic risk: Carmona/Fouque/Sun '13

Mean field game model:

n banks with log-monetary reserves $(X_t^i)_{t \in [0, T]}$,

$$dX_t^i = [a(\bar{X}_t - X_t^i) + \alpha_t^i] dt + \sigma \rho dW_t^i + \sigma \sqrt{1 - \rho^2} dB_t,$$

$$\bar{X}_t = \frac{1}{n} \sum_{k=1}^n X_t^k$$

Bank i chooses to borrow/lend from a central bank at rate α_t^i , to minimize some cost

$$\mathbb{E} \left[\int_0^T f(X_t^i, \bar{X}_t, \alpha_t^i) dt + g(X_T^i, \bar{X}_T) \right].$$

Goal: Find systemic event probabilities in **Nash equilibrium**.

Section 3

Stochastic differential mean field games

Stochastic differential games

Agents $i = 1, \dots, n$ have state process dynamics

$$dX_t^i = b(X_t^i, \bar{\mu}_t^n, \alpha_t^i)dt + \sigma dW_t^i + \sigma_0 dB_t,$$

$$\bar{\mu}_t^n := \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

with B, W^1, \dots, W^n independent, (X_0^1, \dots, X_0^n) i.i.d.

Stochastic differential games

Agents $i = 1, \dots, n$ have state process dynamics

$$dX_t^i = b(X_t^i, \bar{\mu}_t^n, \alpha_t^i)dt + \sigma dW_t^i + \sigma_0 dB_t,$$

$$\bar{\mu}_t^n := \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

with B, W^1, \dots, W^n independent, (X_0^1, \dots, X_0^n) i.i.d.

Agent i chooses α^i to minimize

$$J_i^n(\alpha^1, \dots, \alpha^n) := \mathbb{E} \left[\int_0^T f(X_t^i, \bar{\mu}_t^n, \alpha_t^i) dt + g(X_T^i, \bar{\mu}_T^n) \right].$$

Stochastic differential games

Agents $i = 1, \dots, n$ have state process dynamics

$$dX_t^i = b(X_t^i, \bar{\mu}_t^n, \alpha_t^i)dt + \sigma dW_t^i + \sigma_0 dB_t,$$

$$\bar{\mu}_t^n := \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

with B, W^1, \dots, W^n independent, (X_0^1, \dots, X_0^n) i.i.d.

Agent i chooses α^i to minimize

$$J_i^n(\alpha^1, \dots, \alpha^n) := \mathbb{E} \left[\int_0^T f(X_t^i, \bar{\mu}_t^n, \alpha_t^i) dt + g(X_T^i, \bar{\mu}_T^n) \right].$$

Say $(\alpha^1, \dots, \alpha^n)$ form an ϵ -Nash equilibrium if $\forall i = 1, \dots, n$

$$J_i^n(\alpha^1, \dots, \alpha^n) \leq \epsilon + \inf_{\beta} J_i^n(\dots, \alpha^{i-1}, \beta, \alpha^{i+1}, \dots).$$

Mean field limit $n \rightarrow \infty$?

The problem

Given for each n an ϵ_n -Nash equilibrium $(\alpha^{n,1}, \dots, \alpha^{n,n})$, with $\epsilon_n \rightarrow 0$, can we characterize the possible limits of $\bar{\mu}_t^n$?
Limiting behavior of a representative agent?

Mean field limit $n \rightarrow \infty$?

The problem

Given for each n an ϵ_n -Nash equilibrium $(\alpha^{n,1}, \dots, \alpha^{n,n})$, with $\epsilon_n \rightarrow 0$, can we characterize the possible limits of $\bar{\mu}_t^n$?
Limiting behavior of a representative agent?

Previous results

Lasry/ Lions '06, Bardi '11, Feleqi '13, Gomes '13,
Carmona/Fouque/Sun '13, Fischer '14

Mean field limit $n \rightarrow \infty$?

The problem

Given for each n an ϵ_n -Nash equilibrium $(\alpha^{n,1}, \dots, \alpha^{n,n})$, with $\epsilon_n \rightarrow 0$, can we characterize the possible limits of $\bar{\mu}_t^n$?
Limiting behavior of a representative agent?

Previous results

Lasry/ Lions '06, Bardi '11, Feleqi '13, Gomes '13,
Carmona/Fouque/Sun '13, Fischer '14

A related, better-understood problem

Find a mean field game solution directly, and use it to construct an ϵ_n -Nash equilibrium for the n -player game.

See **Huang/Malhamé/Caines '06** & many others.

Proposed mean field game limit, without common noise

Intuition and the existing literature suggest that $\bar{\mu}^n$ may converge to a **mean field game** (MFG) limit, a process μ satisfying:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^{\alpha}, \mu_t, \alpha_t) dt + g(X_T^{\alpha}, \mu_T) \right], \\ dX_t^{\alpha} & = b(X_t^{\alpha}, \mu_t, \alpha_t) dt + \sigma dW_t, \end{cases}$$

Proposed mean field game limit, without common noise

Intuition and the existing literature suggest that $\bar{\mu}^n$ may converge to a **mean field game** (MFG) limit, a process μ satisfying:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^{\alpha}, \mu_t, \alpha_t) dt + g(X_T^{\alpha}, \mu_T) \right], \\ dX_t^{\alpha} & = b(X_t^{\alpha}, \mu_t, \alpha_t) dt + \sigma dW_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*}). \end{cases}$$

Proposed mean field game limit, with common noise

Intuition and the existing literature suggest that $\bar{\mu}^n$ may converge to a **mean field game** (MFG) limit, a process μ satisfying:

$$\left\{ \begin{array}{l} \alpha^* \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^\alpha, \mu_t, \alpha_t) dt + g(X_T^\alpha, \mu_T) \right], \\ dX_t^\alpha = b(X_t^\alpha, \mu_t, \alpha_t) dt + \sigma dW_t + \sigma_0 dB_t, \end{array} \right.$$

Proposed mean field game limit, with common noise

Intuition and the existing literature suggest that $\bar{\mu}^n$ may converge to a **mean field game** (MFG) limit, a process μ satisfying:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^\alpha, \mu_t, \alpha_t) dt + g(X_T^\alpha, \mu_T) \right], \\ dX_t^\alpha & = b(X_t^\alpha, \mu_t, \alpha_t) dt + \sigma dW_t + \sigma_0 dB_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*} \mid \mathcal{F}_t^B), \quad \mathcal{F}_t^B := \sigma(B_s : s \leq t). \end{cases}$$

Proposed mean field game limit, with common noise

Intuition and the existing literature suggest that $\bar{\mu}^n$ may converge to a **mean field game** (MFG) limit, a process μ satisfying:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^\alpha, \mu_t, \alpha_t) dt + g(X_T^\alpha, \mu_T) \right], \\ dX_t^\alpha & = b(X_t^\alpha, \mu_t, \alpha_t) dt + \sigma dW_t + \sigma_0 dB_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*} \mid \mathcal{F}_t^B), \quad \mathcal{F}_t^B := \sigma(B_s : s \leq t). \end{cases}$$

We call this a **strong MFG solution**, since μ_t is \mathcal{F}_t^B -adapted.

Proposed mean field game limit, with common noise

Intuition and the existing literature suggest that $\bar{\mu}^n$ may converge to a **mean field game** (MFG) limit, a process μ satisfying:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^\alpha, \mu_t, \alpha_t) dt + g(X_T^\alpha, \mu_T) \right], \\ dX_t^\alpha & = b(X_t^\alpha, \mu_t, \alpha_t) dt + \sigma dW_t + \sigma_0 dB_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*} \mid \mathcal{F}_t^B), \quad \mathcal{F}_t^B := \sigma(B_s : s \leq t). \end{cases}$$

We call this a **strong MFG solution**, since μ_t is \mathcal{F}_t^B -adapted. Without some kind of uniqueness (hard to come by!), we should expect only a **weak solution**:

$$\mu_t = \text{Law}(X_t^{\alpha^*} \mid \mathcal{F}_t^{\mu, B}), \text{ with } X_0, (\mu, B), W \text{ independent.}$$

Standing assumptions

Admissible controls for n -player game

Any \mathcal{F}_t^n -adapted process, where

$$\mathcal{F}_t^n \supset \sigma(X_0^1, \dots, X_0^n, W_s^1, \dots, W_s^n, B_s : s \leq t).$$

Technicalities

b, f, g continuous, control space $A \subset \mathbb{R}^k$ closed, b Lipschitz in (x, μ) , growth assumptions...

Main results

Theorem (Mean field limit)

Given for each n an ϵ_n -Nash equilibrium with $\epsilon_n \rightarrow 0$, the sequence $(\bar{\mu}^n)_{n=1}^\infty$ is tight, and **every limit is a weak MFG solution**.

Conversely, every weak MFG solution can be obtained as a limit in this way.

Main results

Theorem (Mean field limit)

*Given for each n an ϵ_n -Nash equilibrium with $\epsilon_n \rightarrow 0$, the sequence $(\bar{\mu}^n)_{n=1}^\infty$ is tight, and **every limit is a weak MFG solution**.*

Conversely, every weak MFG solution can be obtained as a limit in this way.

Theorem (Existence, with R. Carmona & F. Delarue)

There exists a weak MFG solution.

Main results

Theorem (Mean field limit)

Given for each n an ϵ_n -Nash equilibrium with $\epsilon_n \rightarrow 0$, the sequence $(\bar{\mu}^n)_{n=1}^\infty$ is tight, and **every limit is a weak MFG solution**.

Conversely, every weak MFG solution can be obtained as a limit in this way.

Theorem (Existence, with R. Carmona & F. Delarue)

There exists a weak MFG solution.

Theorem (Uniqueness, with R. Carmona & F. Delarue)

A Yamada-Watanabe-type theorem holds for MFGs, and under strong additional assumptions we have pathwise uniqueness.

Main results

Theorem (Mean field limit)

Given for each n an ϵ_n -Nash equilibrium with $\epsilon_n \rightarrow 0$, the sequence $(\bar{\mu}^n)_{n=1}^\infty$ is tight, and **every limit is a weak MFG solution**.
Conversely, every weak MFG solution can be obtained as a limit in this way.

Theorem (Existence, with R. Carmona & F. Delarue)

There exists a weak MFG solution.

Theorem (Uniqueness, with R. Carmona & F. Delarue)

A Yamada-Watanabe-type theorem holds for MFGs, and under strong additional assumptions we have pathwise uniqueness.

Theorem (Existence, with K. Webster)

“Translation invariant” MFGs admit strong solutions.

Section 4

A surprise in the case of no common noise

Interacting particle system without common noise

Particles $i = 1, \dots, n$ have dynamics

$$dX_t^i = b(X_t^i, \bar{\mu}_t^n) dt + \sigma dW_t^i,$$

$$\bar{\mu}_t^n := \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

with W^1, \dots, W^n independent, (X_0^1, \dots, X_0^n) i.i.d.

Mean field limit $n \rightarrow \infty$

Theorem

1. $\bar{\mu}^n$ are tight in $C([0, T]; \mathcal{P}(\mathbb{R}))$.

Mean field limit $n \rightarrow \infty$

Theorem

1. $\bar{\mu}^n$ are tight in $C([0, T]; \mathcal{P}(\mathbb{R}))$.
2. Every weak limit μ is such that a.e. realization $\nu \in C([0, T]; \mathcal{P}(\mathbb{R}))$ satisfies the **McKean-Vlasov (MV) equation**:

$$\begin{cases} dX_t = b(X_t, \nu_t)dt + \sigma dW_t, \\ \nu_t = \text{Law}(X_t). \end{cases}$$

See: Oelschläger '84, Gärtner '88.

McKean-Vlasov equations (MFG without control)

Strong McKean-Vlasov solution: A **deterministic** μ s.t.:

$$\begin{cases} dX_t &= b(X_t, \mu_t)dt + \sigma dW_t, \\ \mu_t &= \text{Law}(X_t). \end{cases}$$

McKean-Vlasov equations (MFG without control)

Strong McKean-Vlasov solution: A **deterministic** μ s.t.:

$$\begin{cases} dX_t &= b(X_t, \mu_t)dt + \sigma dW_t, \\ \mu_t &= \text{Law}(X_t). \end{cases}$$

Weak McKean-Vlasov solution: A **stochastic** μ s.t.:

$$\begin{cases} dX_t &= b(X_t, \mu_t)dt + \sigma dW_t, \\ \mu_t &= \text{Law}(X_t \mid \mathcal{F}_t^\mu), \quad \text{with } X_0, W, \mu \text{ independent.} \end{cases}$$

McKean-Vlasov equations (MFG without control)

Strong McKean-Vlasov solution: A **deterministic** μ s.t.:

$$\begin{cases} dX_t &= b(X_t, \mu_t)dt + \sigma dW_t, \\ \mu_t &= \text{Law}(X_t). \end{cases}$$

Weak McKean-Vlasov solution: A **stochastic** μ s.t.:

$$\begin{cases} dX_t &= b(X_t, \mu_t)dt + \sigma dW_t, \\ \mu_t &= \text{Law}(X_t \mid \mathcal{F}_t^\mu), \quad \text{with } X_0, W, \mu \text{ independent.} \end{cases}$$

Theorem

A random measure μ is a weak solution if and only if it is concentrated on the set of strong solutions, that is a.e. realization is a strong solution.

MFG solutions without common noise

Strong MFG solution: A **deterministic** μ s.t.:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^{\alpha}, \mu_t, \alpha_t) dt + g(X_T^{\alpha}, \mu_T) \right], \\ dX_t^{\alpha} & = b(X_t^{\alpha}, \mu_t, \alpha_t) dt + \sigma dW_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*}). \end{cases}$$

MFG solutions without common noise

Strong MFG solution: A **deterministic** μ s.t.:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^{\alpha}, \mu_t, \alpha_t) dt + g(X_T^{\alpha}, \mu_T) \right], \\ dX_t^{\alpha} & = b(X_t^{\alpha}, \mu_t, \alpha_t) dt + \sigma dW_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*}). \end{cases}$$

Weak MFG solution: A **stochastic** μ s.t.:

$$\begin{cases} \alpha^* & \in \arg \min_{\alpha} \mathbb{E} \left[\int_0^T f(X_t^{\alpha}, \mu_t, \alpha_t) dt + g(X_T^{\alpha}, \mu_T) \right], \\ dX_t^{\alpha} & = b(X_t^{\alpha}, \mu_t, \alpha_t) dt + \sigma dW_t, \\ \mu_t & = \text{Law}(X_t^{\alpha^*} \mid \mathcal{F}_t^{\mu}), \quad \text{with } X_0, W, \mu \text{ independent.} \end{cases}$$

Weak vs strong MFG solutions

Until now, the MFG literature only considered strong solutions:

A natural question:

Are weak MFG solutions concentrated on the set of strong MFG solutions? In other words, is a.e. realization of a weak MFG solution a strong MFG solution?

Weak vs strong MFG solutions

Until now, the MFG literature only considered strong solutions:

A natural question:

Are weak MFG solutions concentrated on the set of strong MFG solutions? In other words, is a.e. realization of a weak MFG solution a strong MFG solution?

Answer

NO.

Weak vs strong MFG solutions

Until now, the MFG literature only considered strong solutions:

A natural question:

Are weak MFG solutions concentrated on the set of strong MFG solutions? In other words, is a.e. realization of a weak MFG solution a strong MFG solution?

Answer

NO.

Conclusion

Strong solutions are not enough to describe mean field limits.

Weak vs strong MFG solutions

Until now, the MFG literature only considered strong solutions:

A natural question:

Are weak MFG solutions concentrated on the set of strong MFG solutions? In other words, is a.e. realization of a weak MFG solution a strong MFG solution?

Answer

NO.

Conclusion

Strong solutions are not enough to describe mean field limits.

The obstruction

When μ is deterministic, the control α can anticipate μ .

A resolution

A sufficient condition

For each deterministic $\mu = (\mu_t)_{t \in [0, T]}$, find an optimal control $\alpha^*[\mu] = (\alpha^*[\mu]_t)_{t \in [0, T]}$. Suppose

$$\alpha^*[\mu]_t = \alpha^*[\mu_{\cdot \wedge t}]_t, \quad \text{for all } t, \mu.$$

Then every weak solution is concentrated on the set of strong solutions.

Open problem

For a family of optimal control problems parametrized by paths $(\mu_t)_{t \in [0, T]}$, under what conditions is the dependence of the optimal control on the parameter **adapted**?

Section 5

MFG limit proof outline

Interacting particle system with common noise

Particles $i = 1, \dots, n$ have dynamics

$$dX_t^i = b(X_t^i, \bar{\mu}_t^n)dt + \sigma dW_t^i + \sigma_0 dB_t,$$

$$\bar{\mu}_t^n := \frac{1}{n} \sum_{k=1}^n \delta_{X_t^k},$$

with B, W^1, \dots, W^n independent, (X_0^1, \dots, X_0^n) i.i.d.

Mean field limit $n \rightarrow \infty$, an unorthodox approach

Theorem

1. $(\bar{\mu}^n, B, W^1, X^1)$ are tight in $C([0, T]; \mathcal{P}(\mathbb{R}) \times \mathbb{R}^3)$.

Mean field limit $n \rightarrow \infty$, an unorthodox approach

Theorem

1. $(\bar{\mu}^n, B, W^1, X^1)$ are tight in $C([0, T]; \mathcal{P}(\mathbb{R}) \times \mathbb{R}^3)$.
2. Every weak limit (μ, B, W, X) solves the *conditional McKean-Vlasov (CMV) equation*:

$$\begin{cases} dX_t = b(X_t, \mu_t)dt + \sigma dW_t + \sigma_0 dB_t, \\ \mu_t = \text{Law}(X \mid \mathcal{F}_t^{\mu, B}), \text{ with } X_0, (\mu, B), W \text{ independent.} \end{cases}$$

Mean field limit $n \rightarrow \infty$, an unorthodox approach

Theorem

1. $(\bar{\mu}^n, B, W^1, X^1)$ are tight in $C([0, T]; \mathcal{P}(\mathbb{R}) \times \mathbb{R}^3)$.
2. Every weak limit (μ, B, W, X) solves the *conditional McKean-Vlasov (CMV) equation*:

$$\begin{cases} dX_t = b(X_t, \mu_t)dt + \sigma dW_t + \sigma_0 dB_t, \\ \mu_t = \text{Law}(X \mid \mathcal{F}_t^{\mu, B}), \text{ with } X_0, (\mu, B), W \text{ independent.} \end{cases}$$

This approach **keeps track of a representative particle** and thus adapts well to the MFG setting.

Proving the MFG limit

Theorem

Given for each n an ϵ_n -Nash equilibrium $(\alpha^{n,1}, \dots, \alpha^{n,n})$ with $\epsilon_n \rightarrow 0$, the sequence $(\bar{\mu}^n)_{n=1}^\infty$ is tight, and *every limit is a weak MFG solution*.

Proof outline

1. Deal with lack of exchangeability.
2. *Control the controls.*
3. Prove tightness.
4. Check dynamics and fixed point condition at limit.
5. *Prove optimality of limits.*

A pipe dream

If $\alpha_t^{n,i} = \hat{\alpha}(t, X_t^i, \bar{\mu}_t^n)$ for some nice function $\hat{\alpha}$, $\forall 1 \leq i \leq n$, then reduce to the particle system case.

Step 1: Exchangeability

Naive idea

Study the joint law of $(\bar{\mu}^n, B, W^1, \alpha^{n,1}, X^1)$.

Step 1: Exchangeability

Naive idea

Study the joint law of $(\bar{\mu}^n, B, W^1, \alpha^{n,1}, X^1)$.

Problem

No reason to expect $(\alpha^{n,1}, \dots, \alpha^{n,n})$ or (X^1, \dots, X^n) to be **exchangeable**.

Step 1: Exchangeability

Naive idea

Study the joint law of $(\bar{\mu}^n, B, W^1, \alpha^{n,1}, X^1)$.

Problem

No reason to expect $(\alpha^{n,1}, \dots, \alpha^{n,n})$ or (X^1, \dots, X^n) to be **exchangeable**.

Solution

Randomly select representative agent. Study the laws

$$\begin{aligned} Q_n &= \frac{1}{n} \sum_{i=1}^n \text{Law}(\bar{\mu}^n, B, W^i, X^i, \alpha^{n,i}) \\ &= \text{Law}(\bar{\mu}^n, B, W^U, X^U, \alpha^{n,U}), \end{aligned}$$

where $U \sim \text{Uniform}\{1, \dots, n\}$ is independent of everything.

Step 2: Control the controls

Problem

Find a good space for the controls, $\alpha^{n,i}$. **Compactness** is difficult in $L^0([0, T]; A)$, with topology of convergence in measure.

Step 2: Control the controls

Problem

Find a good space for the controls, $\alpha^{n,i}$. **Compactness** is difficult in $L^0([0, T]; A)$, with topology of convergence in measure.

Solution

Use **relaxed controls**,

$$\begin{aligned} \mathcal{V} &:= \text{weak closure } \{ dt \delta_{\alpha(t)}(da) : \alpha \in L^0([0, T]; A) \} \\ &\cong (L^0([0, T]; \mathcal{P}(A)), \tau_{\text{relaxed}}). \end{aligned}$$

Drift with a relaxed control Λ is $\int_A b(X_t, \mu_t, a) \Lambda_t(da)$.

Step 2: Control the controls

Problem

Find a good space for the controls, $\alpha^{n,i}$. **Compactness** is difficult in $L^0([0, T]; A)$, with topology of convergence in measure.

Solution

Use **relaxed controls**,

$$\begin{aligned} \mathcal{V} &:= \text{weak closure } \{ dt \delta_{\alpha(t)}(da) : \alpha \in L^0([0, T]; A) \} \\ &\cong (L^0([0, T]; \mathcal{P}(A)), \tau_{\text{relaxed}}). \end{aligned}$$

Drift with a relaxed control Λ is $\int_A b(X_t, \mu_t, a) \Lambda_t(da)$.

An extreme case

Suppose $g \equiv f \equiv 0$. Then **any strategies** $(\alpha^{n,1}, \dots, \alpha^{n,n})$ are Nash, and any relaxed control can arise in the limit.

Step 5: Optimality

Problem

What is the right class of **admissible (relaxed) controls** Λ for the MFG?

Natural but bad choice #1

Require Λ adapted to the filtration $\mathcal{F}_t^{X_0, \mu, B, W}$ generated by (X_0, μ, B, W) , the **given sources of randomness** for the control problems. This class is **too small**, and does not necessarily contain our limit.

Step 5: Optimality

Problem

What is the right class of **admissible (relaxed) controls** Λ for the MFG?

Natural but bad choice #1

Require Λ adapted to the filtration $\mathcal{F}_t^{X_0, \mu, B, W}$ generated by (X_0, μ, B, W) , the **given sources of randomness** for the control problems. This class is **too small**, and does not necessarily contain our limit.

Natural but bad choice #2

Require that B and W remain Wiener processes with respect to the filtration generated by $(X_0, \mu, B, W, \Lambda)$. This class is **too large**, and our limit may not be optimal in this class.

Step 5: Optimality

The right choice

Require Λ to be **compatible**, meaning that \mathcal{F}_t^Λ is conditionally independent of $\mathcal{F}_T^{X_0, \mu, B, W}$ given $\mathcal{F}_t^{X_0, \mu, B, W}$, for each t .

Step 5: Optimality

The right choice

Require Λ to be **compatible**, meaning that \mathcal{F}_t^Λ is conditionally independent of $\mathcal{F}_T^{X_0, \mu, B, W}$ given $\mathcal{F}_t^{X_0, \mu, B, W}$, for each t .

Lemma

Under any weak limit, the relaxed control Λ is compatible.

Step 5: Optimality

The right choice

Require Λ to be **compatible**, meaning that \mathcal{F}_t^Λ is conditionally independent of $\mathcal{F}_T^{X_0, \mu, B, W}$ given $\mathcal{F}_t^{X_0, \mu, B, W}$, for each t .

Lemma

Under any weak limit, the relaxed control Λ is compatible.

Lemma

A relaxed control Λ is compatible if and only if there exists a sequence of $\mathcal{F}_t^{X_0, \mu, B, W}$ -adapted strict controls

$\hat{\alpha}_k = \hat{\alpha}_k(t, X_0, \mu, B, W)$, continuous in μ , such that

$$(X_0, \mu, B, W, \hat{\alpha}_k(t, X_0, \mu, B, W)) \Rightarrow (X_0, \mu, B, W, \Lambda).$$

Step 5: Optimality

Fix a weak limit (μ, B, W, Λ, X) . Show Λ optimal among compatible controls:

1. Consider first a $\mathcal{F}_t^{X_0, \mu, B, W}$ -adapted strict control $\hat{\alpha}(t, X_0, \mu, B, W)$, with $\hat{\alpha}$ continuous in μ .
2. Construct an admissible strategy for the n -player game via $\beta_t^{n,i} = \hat{\alpha}(t, X_0^i, \bar{\mu}^n, B, W^i)$.
3. By ϵ_n -Nash property in n -player game, $\alpha^{n,i}$ is nearly superior to $\beta^{n,i}$ for agent i .
4. Passing the inequality to the limit (using continuity of $\hat{\alpha}$ in μ), Λ is superior to $\hat{\alpha}(t, X_0, \mu, B, W)$.
5. Conclude by approximating general compatible controls by such $\hat{\alpha}(t, X_0, \mu, B, W)$.

Section 6

Refinements

Strict controls

We are more interested in MFG solutions with **strict controls**, meaning $\Lambda_t = \delta_{\alpha_t}$ for some A -valued process α .

Theorem

Suppose for each (x, μ) the set

$$\{(b(x, \mu, a), z) : a \in A, z \geq f(x, \mu, a)\}$$

*is convex. Then for every weak MFG solution there exists another weak MFG solution with strict control **with the same Law** (μ, B, W, X) .*

Strong controls

We are even more interested in MFG solutions with **strong controls**, meaning $\Lambda_t = \delta_{\alpha_t}$ for some $\mathcal{F}_t^{X_0, \mu, B, W}$ -progressive A -valued process α .

Theorem

Suppose b is affine in (x, a) , f is strictly convex in (x, a) , and g is convex in x . Then every weak MFG solution necessarily has strong control.

\Rightarrow Can state MFG limit theorem without reference to relaxed controls or compatibility

Uniqueness

We are even more interested in **strong MFG solutions**, meaning the control is strong and also η is B -measurable, so

$$\mu_t = \text{Law}(X_t \in \cdot \mid \mathcal{F}_t^B).$$

Theorem

Suppose $b = b(x, a)$ is affine in (x, a) and independent of the mean field, f is strictly convex in (x, a) , g is convex in x , $f = f_1(t, x, \mu) + f_2(t, x, a)$, and **monotonicity** holds: $\forall \mu, \nu$,

$$\int [f_1(t, x, \mu) - f_1(t, x, \nu) + g(x, \mu) - g(x, \nu)] (\mu - \nu)(dx) \geq 0.$$

Then “pathwise uniqueness” holds, and the unique weak MFG solution is strong. In particular, for every sequence of ϵ_n -Nash equilibria converges to the unique MFG solution.