Mean Field Game & Mean Field Control: Where do we stand 20 years Later?

René Carmona

Department of Operations Research & Financial Engineering
PACM
Princeton University

Van Eenam Lecture, University of Michigan

Acknowledgement of collaborations

Papers & books with

François Delarue (Nice)

a series of papers with

Colleagues and Ph.D. students

J.P. Fouque, A. Lachapelle, M. Lauriere, D. Lacker, P. Wang, G. Zhu, K. Webster, P. Wang, M. Cerenzia, Z. Tan, L. Leal, G. Dayanikli, C. Zeng

Goal of the Lecture

• Review of **Mathematical Theory** developed for the analysis of the behavior of **Large Populations**

minimizing costs or maximizing rewards in a random environment

• **Optimal Control** problems in **Very High Dimension**

when the control decisions are made by a ***single decision maker***

• **Game Theory** (**Multiperson Decision Theory**) for a **Very Large** number of individuals / agents / players

Analysis of situations in which the cost/reward of a decision maker depends

not only on his own actions but also on those of others

What kind of "Mathematical" Model?

- Finite number of **agents**, say N
  - Traders, managers, households, consumers, birds in a flock, pedestrians, countries, . . .
  - We'll eventually try to shoot for $N >> 1$ (large games) !
- Each agent has a quantitative **objective**, say J^i for agent i
  - Expected profit, savings, consumption, energy spent, CO_2 emissions, travel time, trade balance, foraged food, . . .
  - **Each individual is rational and tries to maximize/minimize** their objective
- Each player is **pro-active**
  - Each individual **action** influences the outcome (and the objective costs/rewards)
  - In **non-cooperative** models, each **individual** agent optimizes **selfishly**
- So what happens to the **overall system**?
  - Does it reach an **equilibrium**?
  - What **kind** of equilibrium?
    - reasonable / rational?
    - erratic / chaotic / irrational?
  - Could we end-up with an instance of **rational irrationality**?

Models of Competition: Nash Equilibria

- Say player i takes **action** α^i,
- Their cost J^i depends upon the actions $\alpha^1, \cdots, \alpha^N$ of **ALL** the players

$$J^i = J^i(\alpha^1, \cdots, \alpha^N)$$

- A **strategy profile** $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a **Nash equilibrium** if for every i and feasible action α^i

$$J^i(\hat{\alpha}^1, \cdots, \hat{\alpha}^{i-1}, \hat{\alpha}^i, \hat{\alpha}^{i+1}, \cdots, \hat{\alpha}^N) \leq J^i(\hat{\alpha}^1, \cdots, \hat{\alpha}^{i-1}, \alpha^i, \hat{\alpha}^{i+1}, \cdots, \hat{\alpha}^N)$$

  whatever $i = 1, \cdots, N$ is !
- In other words,

  the system is in a Nash equilibrium if any player trying to deviate from their action cannot end up better off !

- Not traditional minimization (not the typical *steady state found in physics*)
  - **Identification**: what should these equilibria look like?
  - **Exisence**: in fact, they may not exist
  - **Uniqueness**: when they do, they are in large numbers, often a continuum
  - **Computation:** difficult, both mathematically and numerically
  - Why should a system **settle** in a Nash equilibrium? Which one?

Model of Cooperation

Nash Equilibria vs Social Optimality

If agents take actions $\alpha^1, \cdots, \alpha^N$, **Social Cost** is defined as:

$$J^{SC}(\alpha^1, \cdots, \alpha^N) = \frac{1}{N}\left[J^1(\alpha^1, \cdots, \alpha^N) + \cdots + J^N(\alpha^1, \cdots, \alpha^N)\right]$$

- If $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a Nash Equilibrium (NE)

  $$J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$$

  is the (average) cost to the population for settling in the Nash Equilibrium

Model of Cooperation

Nash Equilibria vs Social Optimality

If agents take actions $\alpha^1, \cdots, \alpha^N$, **Social Cost** is defined as:

$$J^{SC}(\alpha^1, \cdots, \alpha^N) = \frac{1}{N}\left[J^1(\alpha^1, \cdots, \alpha^N) + \cdots + J^N(\alpha^1, \cdots, \alpha^N)\right]$$

- If $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a Nash Equilibrium (NE)

  $$J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$$

  is the (average) cost to the population for settling in the Nash Equilibrium
- A **Central Planner** could **minimize** the social cost and find

  $$(\alpha^{1*}, \cdots, \alpha^{N*}) = \arg \inf_{(\alpha^1, \cdots, \alpha^N)} J^{SC}(\alpha^1, \cdots, \alpha^N)$$

  $J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})$ is the minimal social cost ! **Unfortunately**, it is **rarely** a Nash equilibrium

Model of Cooperation

Nash Equilibria vs Social Optimality

If agents take actions $\alpha^1, \cdots, \alpha^N$, **Social Cost** is defined as:

$$J^{SC}(\alpha^1, \cdots, \alpha^N) = \frac{1}{N}\left[J^1(\alpha^1, \cdots, \alpha^N) + \cdots + J^N(\alpha^1, \cdots, \alpha^N)\right]$$

- If $(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$ is a Nash Equilibrium (NE)

  $$J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)$$

  is the (average) cost to the population for settling in the Nash Equilibrium
- A **Central Planner** could **minimize** the social cost and find

  $$(\alpha^{1*}, \cdots, \alpha^{N*}) = \arg \inf_{(\alpha^1, \cdots, \alpha^N)} J^{SC}(\alpha^1, \cdots, \alpha^N)$$

  $J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})$ is the minimal social cost ! **Unfortunately**, it is **rarely** a Nash equilibrium
- How **bad / suboptimal** can Nash equilibria be?

$$\textbf{PoS} = \frac{inf_{(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)\ NE} J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)}{J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})}$$

quantifying how **much worse** the best Nash equilibrium is

Price of Stability

$$\textbf{PoA} = \frac{sup_{(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)\ NE} J^{SC}(\hat{\alpha}^1, \cdots, \hat{\alpha}^N)}{J^{SC}(\alpha^{1*}, \cdots, \alpha^{N*})}$$

quantifying how **much worse** the worst Nash equilibrium is

Price of Anarchy

Computational Issues

Very difficult to compute Nash Equilibria even if N is only reasonably ***large***

- How large is large depends upon the complexity of the model
  - Static vs dynamic models
  - Complexity of the dynamics (for each time period)
  - Finite states vs continuous states
  - Continuous time (systems of ODEs & PDEs)
  - Randomness (systems of SDEs & SPDEs)
- New industry for algorithm development to compute Nash Equilibria
  - Even for **small** (say $N = 2$) **deterministic** and **static** games
- Emergence of a new field: **Computational Game Theory**

Mathematicians are loosing even more ground!

Bird Flocks are Amazing, aren't they?

Engineering Applications: Crowd Motion

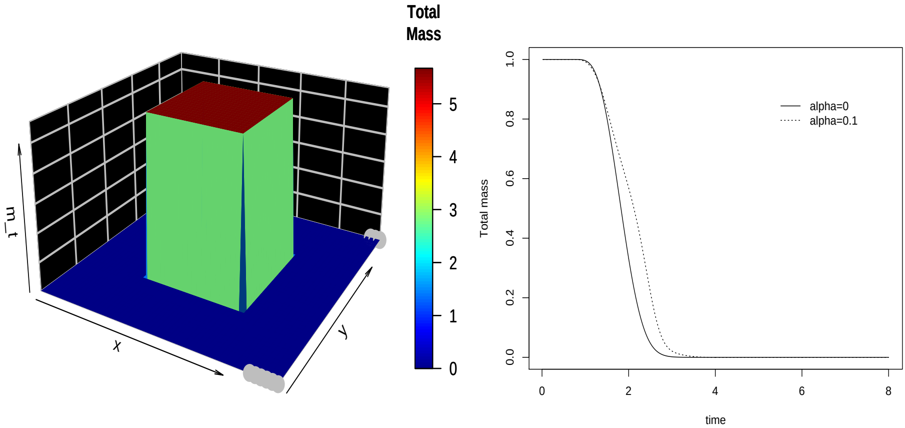

FIGURE: Left: Initial distribution m_0. Right: Time evolution of the total mass of the distribution m_t of the individuals still in the room at time t without congestion (continuous line) and with moderate congestion (dotted line).

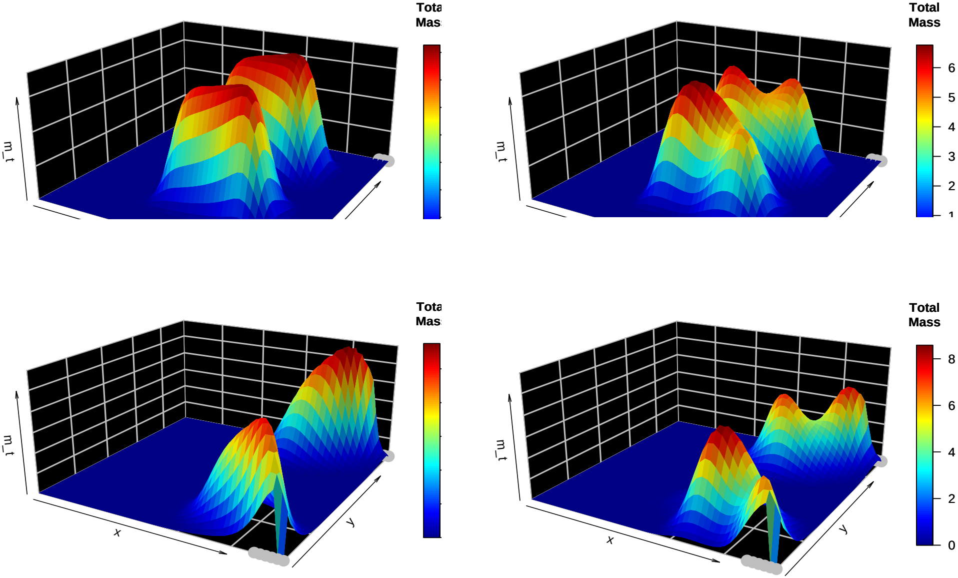
Total Mass
m_t
x
y
1
2
3
4
5
6
0
2
4
6
8

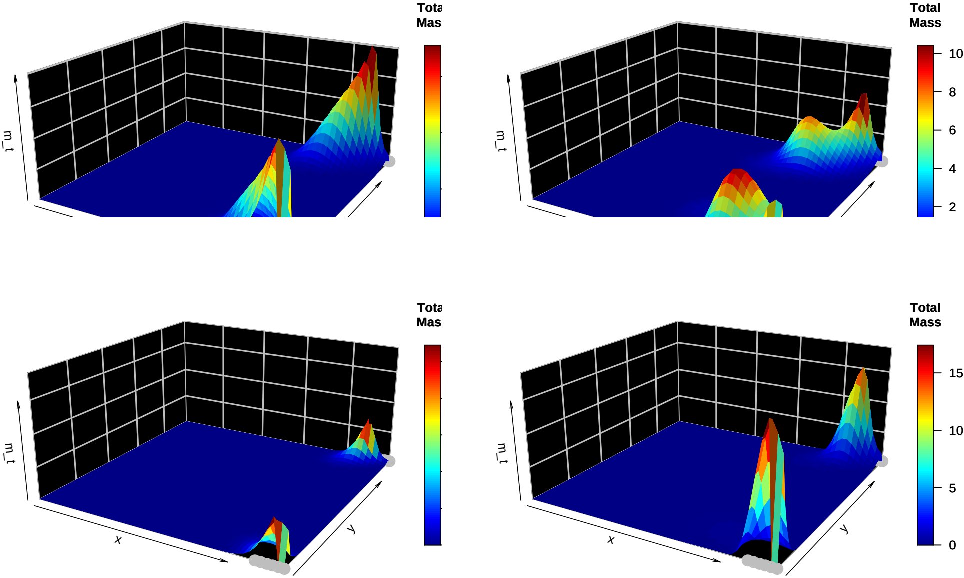
Total Mass
m_t
Total Mass
10
8
6
4
2
m_t
Total Mass
m_t
x
y
Total Mass
15
10
5
0
m_t
x
y

ENGINEERING APPLICATIONS: CYBER SECURITY

Mean Field because the probability that my station will be infected depends upon the proportion of stations already infected.

Ditto for spread of infectious diseases, bank runs, ...

Need for a **Central Decision Maker**? **PoA**?

Engineering Applications: Brain functions

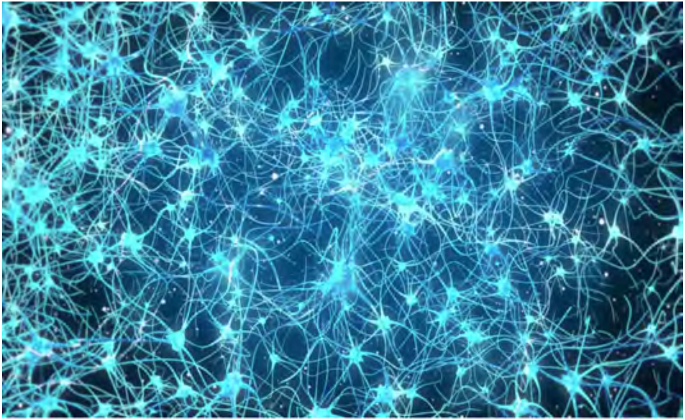

Synchronization of circadian rhythms in SupraChiasmatic Nucleus (SCN)
Nobel Prize for **Hall**, **Rosbash** and **Young**

Can study **Jet Lag Recovery** as MFG!

The Mean Field Game Strategy & the Mean Field Game Problem

N-PLAYER STOCHASTIC DIFFERENTIAL GAMES

Assume Mean Field Interactions (symmetric game)

$$dX_t^{N,i} = b(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N}, \alpha_t^i)dt + \sigma(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N}, \alpha_t^i)dW_t^i \quad i = 1, \cdots, N$$

where $\overline{\mu}^N_{X_t^N}$ is the empirical measure $\overline{\mu}^N_{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^N \delta_{x^i}$

Assume individual i tries to minimize

$$J^i(\boldsymbol{\alpha}^1, \cdots, \boldsymbol{\alpha}^N) = \mathbb{E}\Big[\int_0^T f(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N}, \alpha_t^i)dt + g(X_T, \overline{\mu}^N_{X_T^N})\Big]$$

Search for Nash equilibria

- Very difficult in general, even if N is small
- ϵ-Nash equilibria? Still hard.
- How about in the limit $N \to \infty$?

Mean Field Games 2006 **Lasry - Lions** (MFG) **Caines - Malhamé - Huang** (NCE)

Propagation of Chaos & McKean-Vlasov SDEs

System of N particles $X_t^{N,i}$ at time t with **symmetric (Mean Field)** interactions

$$dX_t^{N,i} = b(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N})dt + \sigma(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N})dW_t^i, \quad i = 1, \cdots, N$$

where $\overline{\mu}^N_{X_t^N}$ is the empirical measure $\overline{\mu}^N_{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^N \delta_{x^i}$

Large population asymptotics ($N \to \infty$)

1. The N processes $\mathbf{X}^{N,i} = (X_t^{N,i})_{0\leq t\leq T}$ become asymptotically **i.i.d.**
2. Each of them is (asymptotically) distributed as the solution of the **McKean-Vlasov** SDE
$$dX_t = b(t, X_t, \mathcal{L}(X_t))dt + \sigma(t, X_t, \mathcal{L}(X_t))dW_t$$

Frequently used notation:

$$\mathcal{L}(X) = \mathbb{P}_X \qquad \text{distribution of the random variable } X.$$

Classical result (see e.g. **Sznitman**), with renewed wave of interest (**Lacker**, **Crowell**, **Bayraktar-Ekren-Zhou**)

MFG PARADIGM

A **typical** agent plays against a **field** of players whose states he/she feels through the statistical distribution **distribution** μ_t of their states at time t

1. For each **Fixed** measure flow $\boldsymbol{\mu} = (\mu_t)$ in $\mathcal{P}(\mathbb{R}^d)$, solve the **standard stochastic control problem**

$$\hat{\boldsymbol{\alpha}} = \arg \inf_{\boldsymbol{\alpha} \in \mathbb{A}} \mathbb{E} \left\{ \int_0^T f(t, X_t, \mu_t, \alpha_t) dt + g(X_T, \mu_T) \right\}$$

subject to

$$dX_t = b(t, X_t, \mu_t, \alpha_t) dt + \sigma(t, X_t, \mu_t, \alpha_t) dW_t$$

2. **Fixed Point Problem**: determine $\boldsymbol{\mu} = (\mu_t)$ so that

$$\forall t \in [0, T], \quad \mathcal{L}(X_t^{\hat{\boldsymbol{\alpha}}}) = \mu_t.$$

$\boldsymbol{\mu}$ or $\hat{\boldsymbol{\alpha}}$ is called a solution of the MFG.

Once this is done one expects that, if $\hat{\alpha}_t = \phi(t, X_t)$,

$$\alpha_t^{j*} = \phi^*(t, X_t^j), \qquad j = 1, \cdots, N$$

form an **approximate Nash equilibrium** for the game with N players.

The Analytic (PDE) Approach to MFGs

For **fixed** $\boldsymbol{\mu} = (\mu_t)_t$, the **value function**

$$V^{\boldsymbol{\mu}}(t,x) = \inf_{(\alpha_s)_{t\le s\le T}} \mathbb{E}\Big[\int_t^T f(s, X_s, \mu_s, \alpha_s)ds + g(X_T, \mu_T)\Big| X_t = x\Big]$$

solves a **HJB (backward) equation**

$$\partial_t V^{\boldsymbol{\mu}}(t,x) + \inf_\alpha [b(t,x,\mu_t,\alpha)\cdot \partial_x V^{\boldsymbol{\mu}}(t,x) + f(t,x,\mu_t,\alpha)]$$
$$\frac{1}{2}\text{trace}[\sigma(t,x)^\dagger \sigma(t,x)\partial^2_{xx} V^{\boldsymbol{\mu}}(t,x)] = 0$$

with terminal condition $V^{\boldsymbol{\mu}}(T,x) = g(x,\mu_T)$

The fixed point step is implemented by requiring that $t \to \mu_t$ solves the **(forward) Fokker-Planck-Kolmogorov** equation

$$\partial_t \mu_t = \frac{1}{2}\text{trace}[\sigma(t,x)^\dagger \sigma(t,x)\partial^2_{xx}\boldsymbol{\mu}_t] + \text{div}[b(t,x,\mu_t,\alpha)\cdot \partial_x V^{\boldsymbol{\mu}}(t,x)\mu_t]$$

This is also a **nonlinear PDE** because μ_t appears in b ......

System of **strongly coupled nonlinear PDEs**! Time goes in **both directions**

Proof of my Commitment to the FIeld

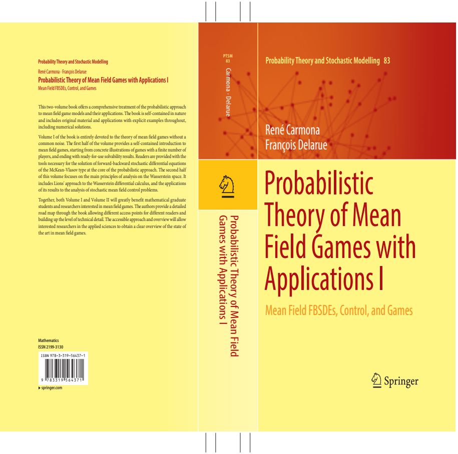

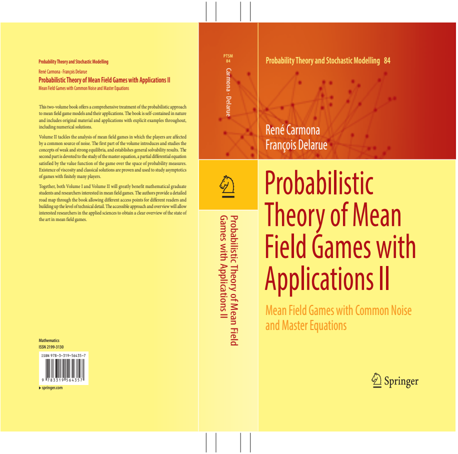

Classical Stochastic Differential Control

$$\inf_{\boldsymbol{\alpha}\in\mathbb{A}} \mathbb{E}\left[\int_0^T f(t, X_t, \alpha_t)dt + g(X_T, \mu_T)\right]$$

$$\text{subject to} \quad dX_t = b(t, X_t, \alpha_t)dt + \sigma(t, X_t, \alpha_t)dW_t; \quad X_0 = x_0.$$

- **Analytic Approach** (by PDEs)
  - HJB equation
- **Probabilistic Approaches** (by FBSDEs)
  1. Represent value function as solution of a BSDE
  2. Represent the gradient of the value function as solution of a FBSDE (Stochastic Maximum Principle)

I. First Probabilistic Approach

Assumptions

- σ is uncontrolled
- σ is invertible

Reduced **Hamitonian**

$$H(t,x,y,\alpha) = b(t,x,\alpha)\cdot y + f(t,x,\alpha)$$

For each control strategy $\boldsymbol{\alpha} = (\alpha_t)_{0\le t\le T}$ with associated state $\mathbf{X}^{\boldsymbol{\alpha}}$, solve **BSDE**

$$dY_t = -H(t, X_t, Z_t\sigma(t,X_t^{\boldsymbol{\alpha}})^{-1}, \alpha_t)dt + Z_t\cdot dW_t, \qquad Y_T = g(X_T^{\boldsymbol{\alpha}}),$$

and denote its solution $(Y^{\boldsymbol{\alpha}}, Z^{\boldsymbol{\alpha}})$. Then

$$Y_0^{\boldsymbol{\alpha}} = J(\boldsymbol{\alpha}) = \mathbb{E}\left[\int_0^T f(t, X_t^{\boldsymbol{\alpha}}, \alpha_t)dt + g(X_T^{\boldsymbol{\alpha}}, \mu_T)\right].$$

So by **comparison theorems** for BSDEs, optimal control $\hat{\boldsymbol{\alpha}} = (\hat\alpha_t)_{0\le t\le T}$ given by:

$$\hat\alpha_t = \hat\alpha(t, X_t^{\boldsymbol{\alpha}}, Z_t^{\boldsymbol{\alpha}}\sigma(t,X_t^{\boldsymbol{\alpha}})^{-1}),$$

with

$$\hat\alpha(t,x,y) \in \text{argmin}_{\alpha\in A} H(t,x,y,\alpha)$$

and the optimum is $Y_0^{\hat{\boldsymbol{\alpha}}} = J(\hat{\boldsymbol{\alpha}})$.

II. Pontryagin Stochastic Maximum Approach

Assumptions

- Coefficients b, σ and f differentiable

Hamitonian

$$H(t,x,y,z,\alpha) = b(t,x,\alpha)\cdot y + \sigma(t,x,\alpha)\cdot z + f(t,x,\alpha)$$

For each control $\boldsymbol{\alpha}$ solve **BSDE** for the adjoint processes $\mathbf{Y} = (Y_t)_t$ and $\mathbf{Z} = (Z_t)_t$

$$dY_t = -\partial_x H(t,X_t,Y_t,Z_t,\alpha_t)dt + Z_t \cdot dW_t, \qquad Y_T = \partial_x g(X_T)$$

Then, optimal control $\hat{\boldsymbol{\alpha}}$ given by:

$$\hat\alpha_t = \hat\alpha(t,X_t,Y_t,Z_t), \quad \text{with} \quad \hat\alpha(t,x,y,z) \in \text{argmin}_{\alpha\in A} H(t,x,y,z,\alpha)$$

II. Pontryagin Stochastic Maximum Approach (cont.)

- **Necessary Condition**
  - If $\hat{\boldsymbol{\alpha}} = (\hat{\alpha}_t)_{0 \le t \le T}$ is an optimal control, then

$$H(t, X_t^{\hat{\boldsymbol{\alpha}}}, Y_t^{\hat{\boldsymbol{\alpha}}}, \hat{\alpha}_t) = \inf_{\alpha \in A} H(t, X_t^{\hat{\boldsymbol{\alpha}}}, Y_t^{\hat{\boldsymbol{\alpha}}}, \alpha)$$

  i.e. $\hat{\boldsymbol{\alpha}}$ minimizes the Hamiltonian along the optimal trajectory.

- **Sufficient Condition**
  - f convex in (x, α) and g convex
  - If $\hat{\boldsymbol{\alpha}} = (\hat{\alpha}_t)_{0 \le t \le T}$ is an admissible control satisfying

$$H(t, X_t^{\hat{\boldsymbol{\alpha}}}, Y_t^{\hat{\boldsymbol{\alpha}}}, \hat{\alpha}_t) = \inf_{\alpha \in A} H(t, X_t^{\hat{\boldsymbol{\alpha}}}, Y_t^{\hat{\boldsymbol{\alpha}}}, \alpha)$$

  then it is optimal.

Summary

In both cases (σ uncontrolled), need to **solve a FBSDE**

$$\begin{cases} dX_t = B(t, X_t, Y_t, Z_t)dt + \Sigma(t, X_t)dW_t, \\ dY_t = F(t, X_t, Y_t, Z_t)dt + Z_t dW_t \end{cases}$$

First Approach

$$\begin{aligned} B(t, x, y, z) &= b\big(t, x, \hat{\alpha}(t, x, z\sigma(t, x)^{-1})\big), \\ F(t, x, y, z) &= -f\big(t, x, \hat{\alpha}(t, x, z\sigma(t, x)^{-1}\big) \\ &\qquad - \big(z\sigma(t, x)^{-1}\big) \cdot b\big(t, x, \hat{\alpha}(t, x, z\sigma(t, x)^{-1})\big). \end{aligned}$$

Second Approach

$$\begin{aligned} B(t, x, y, z) &= b(t, x, \hat{\alpha}(t, x, y)), \\ F(t, x, y, z) &= -\partial_x f(t, x, \hat{\alpha}(t, x, y)) - y \cdot \partial_x b(t, x, \hat{\alpha}(t, x, y)). \end{aligned}$$

FBSDE Decoupling Field

To solve the **standard FBSDE**

$$\begin{cases} dX_t = B(t, X_t, Y_t)dt + \Sigma(t, X_t)dW_t \\ dY_t = -F(t, X_t, Y_t)dt + Z_t dW_t \end{cases}$$

with $X_0 = x_0$ and $Y^T = g(X_T)$,

a standard approach is to look for a solution of the form $Y_t = u(t, X_t)$

- $(t, x) \hookrightarrow u(t, x)$ is called the **decoupling field** of the FBSDE
- If u is smooth,
  - apply Itô's formula to $du(t, X_t)$ using forward equation
  - identify the result with dY_t in backward equation

  $(t, x) \hookrightarrow u(t, x)$ is the solution of a nonlinear PDE

Oh well, So much for the probabilistic approach !

Back to the MFG Problem

- **Forward Dynamics** of state X

$$dX_t = b(t, X_t, \mu_t, \hat{\alpha}(t, X_t, \mu_t, Y_t))dt + \sigma dW_t$$

Back to the MFG Problem

- **Forward Dynamics** of state X

$$dX_t = b(t, X_t, \mu_t, \hat{\alpha}(t, X_t, \mu_t, Y_t))dt + \sigma dW_t$$

- **Backward Dynamics** of the adjoint Y_t

$$dY_t = -\partial_x f(t, X_t, \hat{\alpha}(t, X_t, Y_t)) - Y_t \cdot \partial_x b(t, X_t, \hat{\alpha}(t, X_t, Y_t))dt + Z_t dW_t$$

Back to the MFG Problem

- **Forward Dynamics** of state X

$$dX_t = b(t, X_t, \mu_t, \hat{\alpha}(t, X_t, \mu_t, Y_t))dt + \sigma dW_t$$

- **Backward Dynamics** of the adjoint Y_t

$$dY_t = -\partial_x f(t, X_t, \hat{\alpha}(t, X_t, Y_t)) - Y_t \cdot \partial_x b(t, X_t, \hat{\alpha}(t, X_t, Y_t))dt + Z_t dW_t$$

- **In equilibrium** $\mu_t = \mathbb{P}_{X_t}$

So the FBSDE is of McKean-Vlasov type !

Back to the MFG Problem

- **Forward Dynamics** of state X

$$dX_t = b(t, X_t, \mu_t, \hat{\alpha}(t, X_t, \mu_t, Y_t))dt + \sigma dW_t$$

- **Backward Dynamics** of the adjoint Y_t

$$dY_t = -\partial_x f(t, X_t, \hat{\alpha}(t, X_t, Y_t)) - Y_t \cdot \partial_x b(t, X_t, \hat{\alpha}(t, X_t, Y_t))dt + Z_t dW_t$$

- **In equilibrium** $\mu_t = \mathbb{P}_{X_t}$

So the FBSDE is of McKean-Vlasov type !

- For $\boldsymbol{\mu} = (\mu_t)_t$ **fixed**, assume decoupling field $u^{\boldsymbol{\mu}} : [0, T] \times \mathbb{R}^d \hookrightarrow \mathbb{R}$ exists so that

$$Y_t = u^{\mu}(t, X_t)$$

so in equilibrium

$$Y_t = u^{\mathbb{P}_{X_t}}(t, X_t).$$

Back to the MFG Problem

- **Forward Dynamics** of state X

$$dX_t = b(t, X_t, \mu_t, \hat{\alpha}(t, X_t, \mu_t, Y_t))dt + \sigma dW_t$$

- **Backward Dynamics** of the adjoint Y_t

$$dY_t = -\partial_x f(t, X_t, \hat{\alpha}(t, X_t, Y_t)) - Y_t \cdot \partial_x b(t, X_t, \hat{\alpha}(t, X_t, Y_t))dt + Z_t dW_t$$

- **In equilibrium** $\mu_t = \mathbb{P}_{X_t}$

So the FBSDE is of McKean-Vlasov type !

- For $\boldsymbol{\mu} = (\mu_t)_t$ **fixed**, assume decoupling field $u^{\boldsymbol{\mu}} : [0, T] \times \mathbb{R}^d \hookrightarrow \mathbb{R}$ exists so that

$$Y_t = u^{\mu}(t, X_t)$$

so in equilibrium

$$Y_t = u^{\mathbb{P}_{X_t}}(t, X_t).$$

- **Challenge:** Could the function

$$(t, x, \mu) \hookrightarrow U(t, x, \mu) = u^{\mu}(t, X_t)$$

which contains all the information be the solution of an **infinite dimensional PDE**, with time evolving in **one single direction**?

MASTER EQUATION touted by **P.L. Lions** in his lectures.

Differentiability of Functions of Measures

$\mathcal{M}(\mathbb{R}^d)$ space of **signed** (finite) measures on $\mathbb{R}^d$

- Banach space (dual of a space of continuous functions)
- Classical differential calculus available
- If
$$\mathcal{M}(\mathbb{R}^d) \ni m \hookrightarrow \phi(m) \in \mathbb{R}$$
"ϕ ***is differentiable***" has a meaning
- For $m_0 \in \mathcal{M}(\mathbb{R}^d)$ one can define
$$\frac{\delta \phi(m_0)}{\delta m}(\cdot)$$
as a function on $\mathbb{R}^d$ in **Fréchet** or **Gâteaux** sense

Bensoussan-Frehe-Yam alternative is to work only with measures with **densities** and view ϕ as a function on $L^1(\mathbb{R}^d, dx)$!

Topology on Wasserstein Space

Measures appearing in MFG theory are probability distributions of random variables !!!

Wasserstein space

$$\mathcal{P}_2(\mathbb{R}^d) = \left\{ \mu \in \mathcal{P}(\mathbb{R}^d);\ \int_{\mathbb{R}^d} |x|^2 d\mu(x) < \infty \right\}$$

Metric space for the 2-Wasserstein distance

$$W_2(\mu, \nu) = \inf_{\pi \in \Pi(\mu,\nu)} \left[\int_{\mathbb{R}^d \times \mathbb{R}^d} |x - y|^2 \pi(dx, dy) \right]^{1/2}$$

where $\Pi(\mu, \nu)$ is the set of probability measures coupling μ and ν.
Topological properties of Wasserstein space well understood as following statements are equivalents

- $\mu^N \longrightarrow \mu$ in Wasserstein space
- $\mu^N \longrightarrow \mu$ weakly and $\int |x|^2 \mu^N(dx) \longrightarrow \int |x|^2 \mu(dx)$

Differential Calculus on Wasserstein Space

What does it mean "ϕ ***is differentiable***" or "ϕ ***is convex***" for

$$\mathcal{P}_2(\mathbb{R}^d) \ni \mu \hookrightarrow \phi(\mu) \in \mathbb{R}$$

Wasserstein space $\mathcal{P}_2(\mathbb{R}^d)$ is a **metric space** for W_2

- **Optimal transportation** (**Monge-Ampere-Kantorovich**)
- Curve length and shortest paths (**geodesics**)
- Notion of **convex function** on $\mathcal{P}_2(\mathbb{R}^d)$
- **Tangent spaces** and **differential geometry** on $\mathcal{P}_2(\mathbb{R}^d)$.
- **Differential calculus** on Wasserstein space

Brenier, Benamou, Ambrosio, Gigli, Otto, Caffarelli, Villani, Carlier, ...

DIFFERENTIABILITY IN THE SENSE OF P.L.LIONS

If $\mathcal{P}_2(\mathbb{R}^d) \ni \mu \hookrightarrow \phi(\mu) \in \mathbb{R}$ is "***differentiable***" on **Wasserstein space** what about

$$\mathbb{R}^{dN} \ni (x^1, \cdots, x^N) \mapsto u(x^1, \cdots, x^N) = \phi\left(\frac{1}{N}\sum_{j=1}^{N} \delta_{x^j}\right) \, ?$$

How does $\partial\phi(\mu)$ relate to $\partial_{x^i} u(x^1, \cdots, x^N)$?

Lions' Solution

- **Lift** ϕ up to $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ into $\tilde{\phi}$ defined by $\tilde{\phi}(X) = \phi(\tilde{\mathbb{P}}_X)$
- Use Fréchet differentials on **flat space** L^2

Definition of L-differentiability

ϕ is differentiable at μ_0 if $\tilde{\phi}$ is Fréchet differentiable at X_0 s.t. $\tilde{\mathbb{P}}_{X_0} = \mu_0$

- Check definition is **intrinsic**

Examples of L-differentials

- **Examples**

$$\phi(\mu) = \int_{\mathbb{R}^d} h(x)\mu(dx) \quad \Longrightarrow \quad \partial\phi(\mu)(\cdot) = \partial h(\cdot)$$

$$\phi(\mu) = \int_{\mathbb{R}^d}\int_{\mathbb{R}^d} h(x-y)\mu(dx)\mu(dy) \quad \Longrightarrow \quad \partial\phi(\mu)(\cdot) = [2\partial h(\cdot) * \mu](\cdot)$$

$$\phi(\mu) = \int_{\mathbb{R}^d} \varphi(x,\mu)\mu(dx) \quad \Longrightarrow \quad \partial\phi(\mu)(\cdot) = \partial_x\varphi(\cdot,\mu) + \int_{\mathbb{R}^d} \partial_\mu\varphi(x',\mu)(\cdot)\mu(dx')$$

- **Connection with flate derivative** Assume $\phi : \mathcal{M}_2(\mathbb{R}^d) \mapsto \mathbb{R}$ has a linear functional derivative (at least in a neighborhood of $\mathcal{P}_2(\mathbb{R}^d)$ and that $\mathbb{R}^d \ni x \mapsto [\delta\phi/\delta m](m)(x)$ is differentiable and the derivative

$$\mathcal{M}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (m,x) \mapsto \partial_x\Big[\frac{\delta\phi}{\delta m}\Big](m)(x) \in \mathbb{R}^d$$

is jointly continuous in (m,x) and is of linear growth in x, then ϕ is L-differentiable and

$$\partial_\mu\phi(\mu)(\,\cdot\,) = \partial_x\frac{\delta\phi}{\delta m}(\mu)(\,\cdot\,), \quad \mu \in \mathcal{P}_2(\mathbb{R}^d).$$

- **A sobering counter-example**. If $\mu_0 \in \mathcal{P}_2(E)$ is fixed, the square distance function

$$\mathcal{P}_2(E) \ni \mu \to W_2(\mu_0,\mu)^2 \in \mathbb{R}$$

may not be convex or even L-differentiable!

Itô's Chain Rule

- If u is smooth
- If $d\xi_t = \eta_t dt + \gamma_t dW_t$
- If $dX_t = b_t dt + \sigma_t dW_t$ and $\mu_t = \mathbb{P}_{X_t}$

$$
\begin{aligned}
u(t, \xi_t, \mu_t) &= u(0, \xi_0, \mu_0) + \int_0^t \partial_x u(s, \xi_s, \mu_s) \cdot (\gamma_s dW_s) \\
&+ \int_0^t \Big(\partial_t u(s, \xi_s, \mu_s) + \partial_x u(s, \xi_s, \mu_s) \cdot \eta_s + \frac{1}{2} \mathrm{trace}\big[\partial^2_{xx} u(s, \xi_s, \mu_s) \gamma_s \gamma_s^\dagger\big] \Big) ds \\
&+ \int_0^t \tilde{\mathbb{E}}\big[\partial_\mu u(s, \xi_s, \mu_s)(\tilde{X}_s) \cdot \tilde{b}_s\big] ds + \frac{1}{2} \int_0^t \tilde{\mathbb{E}}\Big[\mathrm{trace}\Big(\partial_v \big[\partial_\mu u(s, \xi_s, \mu_s)\big](\tilde{X}_s) \tilde{\sigma}_s \tilde{\sigma}_s^\dagger\Big)\Big] ds
\end{aligned}
$$

where the process $(\tilde{X}_t, \tilde{b}_t, \tilde{\sigma}_t)_{0 \le t \le T}$ is an **independent copy** of the process $(X_t, b_t, \sigma_t)_{0 \le t \le T}$, on a different probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$

EXAMPLE OF AN ACTUAL MASTER EQUATION

$$
\begin{aligned}
&\partial_t \mathcal{U}(t,x,\mu) + b\big(t,x,\mu,\hat{\alpha}(t,x,\mu,\partial\mathcal{U}(t,x,\mu))\big)\cdot\partial_x\mathcal{U}(t,x,\mu) \\
&\quad + \frac{1}{2}\mathrm{trace}\Big[\partial^2_{xx}\mathcal{U}(t,x,\mu)\Big] + f\big(t,x,\mu,\hat{\alpha}(t,x,\mu,\partial\mathcal{U}(t,x,\mu))\big) \\
&\quad + \int_{\mathbb{R}^d}\Big[b\big(t,x',\mu,\hat{\alpha}(t,x,\mu,\partial\mathcal{U}(t,x,\mu))\big)\cdot\partial_\mu\mathcal{U}(t,x,\mu)(x') \\
&\qquad\qquad + \frac{1}{2}\mathrm{trace}\Big(\partial_{x'}\partial_\mu\mathcal{U}(t,x,\mu)(x')\Big)\Big]d\mu(x') = 0,
\end{aligned}
$$

for $(t,x,\mu)\in[0,T]\times\mathbb{R}^d\times\mathcal{P}_2(\mathbb{R}^d)$, with the **terminal** condition $V(T,x,\mu)=g(x,\mu)$.

Some Results

- Infinite dimensional PDE, rarely solved !
- **Lions**' solution of several first order models in **Cardaliaguet**'s notes
- The master field is a **viscosity solution** of the master equation (not too hard)
- Desirable results:
  - existence of classical solutions
  - convergence of solutions of finite player games to solutions of the MFG
- **Cardaliaguet-Delarue-Lasry-Lions**: major breakthrough (140 pages) despite very restrictive assumptions
- Assuming existence of classical solutions to the master equation
  - Large deviations & rates of convergence **Delarue-Lacker-Ramanan**
  - Analysis of finite MFGs **R.C.-Delarue, Bayraktar-Cohen**
  - Erdos-Renyi graphs **Delarue**

Taking Stock

SDE State Dynamics for N players	*Optimization* $\longrightarrow$	Nash Equilibrium for N players
$\downarrow$ Fixed Point $\lim_{N\to\infty}$		$\downarrow$ Fixed Point $\lim_{N\to\infty}$
McKean Vlasov Dynamics	*Optimization* $\longrightarrow$	**Mean Field Game, MFG** **Controlled McKean-Vlasov SDE, MFC**

The diagram IS NOT commutative

CONTROLLED MCKEAN-VLASOV SDES

$$\inf_{\boldsymbol{\alpha}=(\alpha_t)_{0\le t\le T}} \mathbb{E}\left[\int_0^T f(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dt + g(X_T, \mathbb{P}_{X_T})\right]$$

under dynamical constraint $dX_t = b(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dt + \sigma(t, X_t, \mathbb{P}_{X_t}, \alpha_t)\,dW_t$.

- State $(X_t, \mathbb{P}_{X_t})$ infinite dimensional
- State trajectory $t \mapsto (X_t, \mu_t)$ is a very thin submanifold due to constraint $\mu_t = \mathbb{P}_{X_t}$
- Open loop form: $\boldsymbol{\alpha} = (\alpha_t)_{0\le t\le T}$ adapted
- Closed loop form: $\alpha_t = \phi(t, X_t, \mathbb{P}_{X_t})$

Whether we use

- Infinite dimensional HJB equation
- Pontryagin stochastic maximum principle with Hamiltonian

  $$H(t, x, \mu, y, z, \alpha) = b(t, x, \mu, \alpha) \cdot y + \sigma(t, x, \mu, \alpha) \cdot z + f(t, x, \mu, \alpha)$$

  and introduce the adjoint equations,

WE NEED TO DIFFERENTIATE FUNCTIONS OF MEASURES !

The Adjoint Equations

Given an admissible control $\boldsymbol{\alpha} = (\alpha_t)_{0 \le t \le T}$ and the corresponding controlled state process $\mathbf{X}^\alpha = (X_t^\alpha)_{0 \le t \le T}$, any couple $(Y_t, Z_t)_{0 \le t \le T}$ satisfying:

$$\begin{cases} dY_t = -\partial_x H(t, X_t^\alpha, \mathbb{P}_{X_t^\alpha}, Y_t, \alpha_t)dt + Z_t dW_t \\ \qquad\qquad\qquad\qquad -\tilde{\mathbb{E}}[\partial_\mu H(t, \tilde{X}_t, \mathbb{P}_{X_t^\alpha}, \tilde{Y}_t, \tilde{\alpha}_t)(X_t^\alpha)]dt \\ Y_T = \partial_x g(X_T^\alpha, \mathbb{P}_{X_T^\alpha}) + \tilde{\mathbb{E}}[\partial_\mu g(\tilde{X}_T^\alpha, \mathbb{P}_{X_T^\alpha})(X_T^\alpha)] \end{cases}$$

where $(\tilde{\alpha}, \tilde{X}, \tilde{Y}, \tilde{Z})$ is an independent copy of (α, X^α, Y, Z). $(\mathbf{Y}, \mathbf{Z})$ is called a set of **adjoint processes**

Extra terms in red are the ONLY difference between MFG and Control of McKean-Vlasov dynamics !!!

Pontryagin Maximum Principle (Sufficiency)

Assume

1. Coefficients continuously differentiable with bounded derivatives;
2. Terminal cost function g is convex;
3. $\boldsymbol{\alpha} = (\alpha_t)_{0 \le t \le T}$ admissible control, $\mathbf{X} = (X_t)_{0 \le t \le T}$ corresponding dynamics, $(\mathbf{Y}, \mathbf{Z}) = (Y_t, Z_t)_{0 \le t \le T}$ adjoint processes and

$$(x, \mu, \alpha) \hookrightarrow H(t, x, \mu, Y_t, Z_t, \alpha)$$

is $dt \otimes d\mathbb{P}$ a.e. **convex**,

then, if moreover

$$H(t, X_t, \mathbb{P}_{X_t}, Y_t, Z_t, \alpha_t) = \inf_{\alpha \in A} H(t, X_t, \mathbb{P}_{X_t}, Y_t, \alpha), \quad \text{a.s.}$$

Then α is an optimal control, i.e.

$$J(\boldsymbol{\alpha}) = \inf_{\boldsymbol{\beta} \in \mathbb{A}} J(\boldsymbol{\beta}).$$

Solution of the McKV Control Problem

Assume

- $b(t, x, \mu, \alpha) = b_0(t) \int_{\mathbb{R}^d} x d\mu(x) + b_1(t)x + b_2(t)\alpha$
  with b_0, b_1 and b_2 is $\mathbb{R}^{d\times d}$-valued and are bounded.
- f and g as in MFG problem.

Then there exists a solution $(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = (X_t, Y_t, Z_t)_{0\le t\le T}$ of the McKean-Vlasov FBSDE

$$\begin{cases} dX_t = b_0(t)\mathbb{E}(X_t)dt + b_1(t)X_t dt + b_2(t)\hat{\alpha}(t, X_t, \mathbb{P}_{X_t}, Y_t)dt + \sigma dW_t, \\ dY_t = -\partial_x H(t, X_t, \mathbb{P}_{X_t}, Y_t, \hat{\alpha}_t)\, dt \\ \qquad\qquad - \mathbb{E}\big[\partial_\mu \tilde{H}(t, \tilde{X}_t, X_t, \tilde{Y}_t, \tilde{\hat{\alpha}}_t)\big]\, dt + Z_t dW_t. \end{cases}$$

with $Y_t = u(t, X_t, \mathbb{P}_{X_t})$ for a function

$$u : [0, T] \times \mathbb{R}^d \times \mathcal{P}_1(\mathbb{R}^d) \ni (t, x, \mu) \mapsto u(t, x, \mu)$$

uniformly of Lip-1 and with linear growth in x.

Existing particular case: *Mean Variance Portfolio Optimization* (**Anderson - Djehiche**)

Large Game Asymptotics with Common Noise

If $\mathbf{W}^0 = (W_t^0)_{t\geq 0}$ is an independent Wiener process **independent** of the idiosyncratic $\mathbf{W}^i = (W_t^i)_{t\geq 0}$ for $i \geq 1$,

$$dX_t^{N,i} = b(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N}, \alpha_t^i)dt + \sigma(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N}, \alpha_t^i)dW_t^i + \sigma^0(t, X_t^{N,i}, \overline{\mu}^N_{X_t^N}, \alpha_t^i)dW_t^0$$

In the limit $N \to \infty$, **Conditional Law of Large Numbers**

- If we consider **exchangeable equilibria**,$(\alpha_t^1, \cdots, \alpha_t^N)$, then
  - By **LLN**
    $$\lim_{N\to\infty} \overline{\mu}_t^N = \mathbb{P}_{X_t^1 | \mathcal{F}_t^0}$$
  - **Dynamics** of player 1 (or any other player) becomes
    $$dX_t^1 = b(t, X_t^1, \mu_t, \alpha_t^1)dt + \sigma(t, X_t^1)dW_t + \sigma^0(t, X_t, \mu_t, \alpha_t^1)dW_t^0$$
    with $\mu_t = \mathbb{P}_{X_t^1 | \mathcal{F}_t^0}$.
  - **Cost** to player 1 (or any other player) becomes
    $$\mathbb{E}\left\{\int_0^T f(t, X_t, \mu_t, \alpha_t^1)dt + g(X_T, \mu_T)\right\}$$

MFG Problem with Common Noise

As usual assume σ and σ^0 do not depend upon μ and α.

1. **Fix** a measure valued $(\mathcal{F}_t^0)$-adapted process (μ_t) in $\mathcal{P}(\mathbb{R})$;
2. Solve the standard **stochastic control problem**
$$\hat{\alpha} = \arg\inf_{\alpha} \mathbb{E}\left\{\int_0^T f(t, X_t, \mu_t, \alpha_t)dt + g(X_T, \mu_T)\right\}$$
subject to
$$dX_t = b(t, X_t, \mu_t, \alpha_t)dt + \sigma(t, X_t)dW_t + \sigma^0(t, X_t) \circ dW_t^0;$$
3. **Fixed Point Problem**: determine (μ_t) so that
$$\forall t \in [0, T], \quad \mathbb{P}_{X_t | \mathcal{F}_t^0} = \mu_t \quad a.s.$$

Once this is done, **if** $\hat{\alpha}_t = \phi(t, X_t)$, go back to N player game and show that:
$$\alpha_t^{j*} = \phi^*(t, X_t^j), \qquad j = 1, \cdots, N$$
form an **approximate Nash equilibrium** for the game with N players.

Among the many complications

the master equation is now second order in the measure argument!

Exciting Challenges being Actively Pursued

- **Moral Hazard & Contract Theory**
  - one principal (regulator) offering a contract to incentivize behavior
  - many agents compete and optimize their objectives

Exciting Challenges being Actively Pursued

- **Moral Hazard & Contract Theory**
  - one principal (regulator) offering a contract to incentivize behavior
  - many agents compete and optimize their objectives
- **Games of timing** (bank runs)

EXCITING CHALLENGES BEING ACTIVELY PURSUED

- **Moral Hazard & Contract Theory**
  - one principal (regulator) offering a contract to incentivize behavior
  - many agents compete and optimize their objectives
- **Games of timing** (bank runs)
- **Network Games** on large ($N \to \infty$) **random graphs**
  - **sparse graphs** $\to$ special behaviors (**Delarue**, **Lacker-Soret**, **Fouque et al**)
  - **dense graphs** $\to$ **graphon models** (**Ozdaglar-Parise**, **R.C. et al**, **Caines et al**, **Bayraktar et al**)

Exciting Challenges being Actively Pursued

- **Moral Hazard & Contract Theory**
  - one principal (regulator) offering a contract to incentivize behavior
  - many agents compete and optimize their objectives
- **Games of timing** (bank runs)
- **Network Games** on large ($N \to \infty$) **random graphs**
  - **sparse graphs** $\to$ special behaviors (**Delarue**, **Lacker-Soret**, **Fouque et al**)
  - **dense graphs** $\to$ **graphon models** (**Ozdaglar-Parise**, **R.C. et al**, **Caines et al**, **Bayraktar et al**)
- **Mean Field Reinforcement Learning (RL)**
  - **RL** plays crucial role in **games** (chess,go), **robotics**, **AI** (ChatGPT, DeepSeek)
  - **IMSI workshop** during RL Special Long Program (Spring 2026)

Thank You

Temporary page!

LaTeX was unable to guess the total number of pages correctly. As there w
some unprocessed data that should have been added to the final page thi
extra page has been added to receive it.
If you rerun the document (without altering it) this surplus page will go awa
because LaTeX now knows how many pages to expect for this document.