Weak Formulation for Mean Field Games and Applications

René Carmona (joint work with D. Lacker)

Bendheim Center for Finance
Department of Operations Research & Financial Engineering
Princeton University

Rennes, May 23, 2013

Example I: Systemic Risk (R.C. - Fouque)

- $X_t^{(i)}, i = 1, \ldots, N$ log-monetary reserves of N banks
- $W_t^{(i)}, i = 0, 1, \ldots, N$ **independent Brownian motions**, $\sigma > 0$
- Model **borrowing and lending** through the drifts:

$$dX_t^i = \left[a(\overline{X}_t - X_t^i) + \alpha_t^i\right] dt + \sigma\left(\sqrt{1-\rho^2}dW_t^i + \rho dW_t^0\right), \quad i = 1, \cdots, N$$

α^i **is the control of bank** i which tries to **minimize**

$$J^i(\alpha^1, \cdots, \alpha^N) = EE\left\{\int_0^T \left[\frac{1}{2q}(\alpha_t^i)^2 - \alpha_t^i(\overline{X}_t - X_t^i)\right] dt\right\}$$

Regulator chooses $q > 0$ to control the cost of borrowing and lending.

- If X_t^i is small (relative to the empirical mean $\overline{X}_t$) then bank i will want to borrow($\alpha_t^i > 0$)
- If X_t^i is large then bank i will want to lend ($\alpha_t^i < 0$)

Approximate Nash MFG-Equilibrium

◇ $\rho = 0$

◇ Simple example (L-Q) of **Mean Field Game (MFG)** *à la Lasry - Lions*

◇ Banks act **independently** of each other

◇ Bank i chooses $\alpha_t^i = q(\overline{X}_t - X_t^i) - \eta_t X_t^i$

$$dX_t^i = \left[(a+q)(\overline{X}_t - X_t^i) - \eta_t X_t^i\right] dt + \sigma dW_t^i$$

for a deterministic function $t \hookrightarrow \eta_t$ solving a Ricatti equation.Therefore

$$d\overline{X}_t = -\eta_t \overline{X}_t dt + \frac{\sigma}{\sqrt{N}} d\overline{W}_t$$

where $\overline{W}_t = \frac{1}{\sqrt{N}} \sum_1^N dW_t^i$.

- **Gaussian system** (OU processes)
- **Large Deviation** estimates for probabilities of **systemic events**

Example II: A Simple Model of Price Impact

(**Almgren-Chriss** '01, **Carlin et al.** '09)

- n brokers trade in the same asset and maximize wealth;
- Brokers $(i = 1, \cdots, n)$ face identical limit order books;
- Broker i trade at *rate* α_t^i at time t
- **Transaction** price = martingale + drift (**price impact**).

Case of Flat Order Book (Quadratic Costs)

- **Asset price**:

$$dS_t = \frac{\gamma}{n} \sum_{i=1}^{n} \alpha_t^i dt + \sigma_0 dB_t$$

- Broker i's **cash** and **volume**:

$$dK_t^i = -(\alpha_t^i S_t + (\alpha_t^i)^2) dt$$
$$dX_t^i = \alpha_t^i dt + \sigma dW_t^i$$

- Broker i's **wealth**: $V_t^i = V_0^i + X_t^i S_t + K_t^i$,

$$dV_t^i = \left(\frac{\gamma}{n} \sum_{j=1}^{n} \alpha_t^j X_t^i - (\alpha_t^i)^2 \right) dt + \sigma S_t dW_t^i + \sigma_0 X_t^i dB_t$$

Risk Neutral Agents

Broker i **maximizes expected wealth** $\mathbb{E}[V_T^i]$:

$$\sup_{\alpha^i} \mathbb{E} \int_0^T \left(\frac{\gamma}{n} \sum_{j=1}^n \alpha_t^j X_t^i - (\alpha_t^i)^2 \right) dt,$$
$$\text{s.t. } dX_t^i = \alpha_t^i dt + \sigma dW_t^i$$

Are there Nash equilibria?

L-Q Mean Field Game

More General Order Books

- Given a **transaction cost curve** $c : \mathbb{R} \to [0, \infty]$ (convex, $c(0) = 0$);
- **Order book shape function** given by Legendre transform γ;
- **Price impact** given by c':
- Optimization of **expected terminal wealth** becomes:

$$\sup_{\alpha^i} \mathbb{E} \int_0^T \left(\frac{\gamma}{n} \sum_{j=1}^{n} c'(\alpha_t^j) X_t^i - c(\alpha_t^i) \right) dt,$$

$$\text{s.t. } dX_t^i = \alpha_t^i dt + \sigma dW_t^i$$

In General

- Adding **benchmark tracking penalties**, **carrying and inventory costs**, ...

$$\sup_{\alpha^i} \mathbb{E}\left[G(X_T^i) + \int_0^T \left(\frac{\gamma}{n} \sum_{j=1}^n c'(\alpha_t^j) X_t^i - c(\alpha_t^i) - F(t, X_t^i) \right) dt \right],$$
$$\text{s.t. } dX_t^i = \alpha_t^i dt + \sigma dW_t^i$$

- Still **MFG** but
  - Brokers' optimization problems **coupled through the empirical distribution of the controls**;
  - Maximizing utility instead of wealth leads to a much harder problem (**common noise would not go away!**)

Search for Nash Equilibriums

- Construct **Best Response Map**
  - for each **strategy profiles** $(\alpha^1, \cdots, \alpha^n)$
  - for each $i \in \{1, \cdots, n\}$
  - find $\hat{\alpha}^i$ maximizing $J^i(\alpha^1, \cdots, \alpha^n)$ over α^i
  - $(\alpha^1, \cdots, \alpha^n) \hookrightarrow (\hat{\alpha}^1, \cdots, \hat{\alpha}^n)$
- Find a **fixed point** for the Best Response map

Can be quite **complex** (**prohibitive** when n is large)

Mean Field Game (MFG) Strategy

- By **symmetry**, interactions depend upon **empirical distributions**
- **ALL** stochastic optimizations should be "**the same**"
- When n is **large**
  - empirical distributions should converge
  - capture interactions with limits of empirical distributions
  - **ONE** standard stochastic control problem **for each possible limit**
- Still need a **fixed point** for choice of the **limit distribution** to be the right one

Lasry - Lions (MFG) **Caines - Malhamé - Huang** (NCE)

The PDE Approach in LaTeX

Motivation (Lasry-Lions, Guéant, La Chapelle, ...)

$$u(t,x) = \sup_{(\alpha_s)_{t\le s\le T}, X_t = x} \mathbb{E}\left[\int_t^T e^{-\rho(s-t)}[g(m(s,X_s)) + h(|\alpha(s,X_s)|)]ds\right]$$

under constraint $dX_t = \alpha(t, X_t)dt + \sigma dW_t$

Formulation (given $m(0,\cdot)$ & $u(T,\cdot)$)

$$\partial_t u + \frac{\sigma^2}{2}\Delta u + H(\nabla u) - \rho u = -g(m) \qquad \text{(Hamilton-Jacobi-Bellman)}$$

$$\partial_t m + \nabla \cdot (mH'(\nabla u)) = \frac{\sigma^2}{2}\Delta m, \qquad \text{(Kolmogorov)}$$

where $m(t,\cdot)$ probability measure, $H(p) = \sup_a(ap - h(a))$.

Stationary Case

$$\frac{\sigma^2}{2}\Delta u + H(\nabla u) - \rho u = -g(m)$$

$$\nabla \cdot (mH'(\nabla u)) = \frac{\sigma^2}{2}\Delta m,$$

Probabilistic Approach

Disclaimer (to PL and the PDE *aficionados*)

> **"Mathematicians (Probabilists) are like Frenchmen: whatever you say to them they translate into their own language and forthwith it is something entirely different."**

Johann Wolfgang von Goethe

Probabilistic Approach

- (Pontryagin) Stochastic Maximum Principle
- **FBSDEs** of **McKean Vlasov** type
  - **R.C. - F. Delarue** (this afternoon - don't miss it)
- **Weak Formulation and BSDEs**

Specific Mean Field Stochastic Differential Games

Player $i \in \{1, \cdots, N\}$ **state process**

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

Objective function

$$J^i(\alpha^1, \cdots, \alpha^N) = \mathbb{E}\left[\int_0^T f(t, X_t^i, \overline{\mu}_t^N, \overline{\nu}_t^N, \alpha_t^i) dt + g(X_T^i, \overline{\mu}_T^N)\right],$$

where

$$\overline{\mu}_t^N := \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}, \quad \overline{\nu}_t^N := \frac{1}{N}\sum_{j=1}^N \delta_{\alpha_t^j}$$

Mean Field Problem

1. **Fix the measure flows** $t \mapsto (\mu_t, \nu_t)$
2. Solve a **standard stochastic control** problem

$$\sup_{\alpha} \mathbb{E}\left[\int_0^T f(t, X_t, \mu_t, \nu_t, \alpha_t)dt + g(X_T, \mu_T)\right], \text{ s.t.}$$
$$dX_t = b\left(t, X_t, \mu_t, \alpha_t\right) dt + \sigma(t, X_t)dW_t,$$

3. Let μ'_t denote the law of the optimally controlled state process at time t and ν'_t the law of the optimal control at time t.
4. Find a **fixed point** $(\mu'_t, \nu'_t) = (\mu_t, \nu_t)$.

First Set of Assumptions

- The *control space* A is a compact convex;
- All progressively measurable A-valued processes are admissible;
- Drift $b : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times A \to \mathbb{R}^d$ progressively measurable, continuous in μ.
- Volatility $\sigma : [0, T] \times \mathcal{C} \to \mathbb{R}^{d \times d}$ progressively measurable.
- There exists a unique strong solution X of the driftless state equation

$$dX_t = \sigma(t, X)dW_t, \qquad X_0 = \xi$$

  such that $\mathbb{E}[\psi^2(X)] < \infty$,
- $\sigma(t, X) > 0$ for all $t \in [0, T]$ almost surely,
- $\sigma^{-1}(t, X)b(t, X, \mu, a)$ is bounded.

Weak Formulation

For each $\mu \in \mathcal{P}_\psi(\mathcal{C})$ and admissible $\alpha \in \mathbb{A}$, define

◇ the probability $\mathbb{P}^{\mu,\alpha}$ on $(\Omega, \mathcal{F}_T)$ by

$$\frac{d\mathbb{P}^{\mu,\alpha}}{d\mathbb{P}} = \exp\left[\int_0^T \sigma^{-1} b(t, X, \mu, \alpha_t) \cdot dW_t - \frac{1}{2}\int_0^T \left|\sigma^{-1} b(t, X, \mu, \alpha_t)\right|^2 dt\right].$$

◇ the process $W^{\mu,\alpha}$ defined by

$$W_t^{\mu,\alpha} := W_t - \int_0^t \sigma^{-1} b(s, X, \mu, \alpha_s)\, ds$$

◇ so that

$$dX_t = b(t, X, \mu, \alpha_t)\, dt + \sigma(t, X) dW_t^{\mu,\alpha}.$$

Weak Formulation (cont.)

- Running objective $f : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times A \to \mathbb{R}$ of the form
$$f(t, x, \mu, q, a) = f_1(t, x, \mu, a) + f_2(t, x, \mu, q).$$
- Terminal objective $g : \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \to \mathbb{R}$ is measurable

$$|g(x, \mu)| + |f(t, x, \mu, q, a)| \le c \left(\psi(x) + \rho \left(\int \psi \, d\mu \right) \right), \quad \forall (t, x, \mu, q, a).$$

for $c > 0$ and an increasing function $\rho : \mathbb{R}_+ \to \mathbb{R}_+$

Problem Statement

Given

- a measure $\mu \in \mathcal{P}(\mathcal{C})$
- a measurable map $[0, T] \ni t \mapsto q_t \in \mathcal{P}(A)$

define the associated conditional expected reward for $\alpha \in \mathbb{A}$ by

$$J_t^{\mu,q}(\alpha) := \mathbb{E}^{\mu,\alpha}\left[\int_t^T f(s, X, \mu, q_s, \alpha_s)ds + g(X, \mu)\middle| \mathcal{F}_t\right]$$

and the conditional value function by

$$V_t^{\mu,q} = \operatorname*{ess\,sup}_{\alpha \in \mathbb{A}} J_t^{\mu,q}(\alpha).$$

Goal: **Find μ and q s.t.**

- **there exists $\hat{\alpha} \in \mathbb{A}$ such that $V_0^{\mu,q} = J_0^{\mu,q}(\hat{\alpha})$,**
- **$P^{\mu,\hat{\alpha}} \circ X^{-1} = \mu$, and $P^{\mu,\hat{\alpha}} \circ \hat{\alpha}_t^{-1} = q_t$ for almost every t**

Existence and uniqueness

Hamiltonian $h : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times \mathbb{R}^d \times A \to \mathbb{R}$,

$$h(t, x, \mu, q, z, a) = f(t, x, \mu, q, a) + z \cdot \sigma^{-1} b(t, x, \mu, a)$$

Maximized Hamiltonian $H : [0, T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times \mathbb{R}^d \to \mathbb{R}$

$$H(t, x, \mu, q, z) := \sup_{a \in A} h(t, x, \mu, q, z, a)$$

Arg-max set

$$A(t, x, \mu, z) := \{a \in A : h(t, x, \mu, q, z, a) = H(t, x, \mu, q, z)\}$$

- $A(t, x, \mu, z)$ does not depend upon q
- $A(t, x, \mu, z)$ is not empty

Finally, a BSDE !

$$Y_t^{\mu,\nu} = g(X,\mu) + \int_t^T H(s,X,\mu,\nu_s,Z_s^{\mu,\nu})ds - \int_t^T Z_s^{\mu,\nu} \cdot dW_s$$

For each $\alpha \in \mathbb{A}$, we may also solve the BSDE

$$\begin{aligned} Y_t^{\mu,\nu,\alpha} &= g(X,\mu) + \int_t^T h(s,X,\mu,\nu_s,Z_s^{\mu,\nu,\alpha},\alpha_s)ds - \int_t^T Z_s^{\mu,\nu,\alpha} \cdot dW_s \\ &= g(X,\mu) + \int_t^T f(s,X,\mu,\nu_s,\alpha_s)ds - \int_t^T Z_s^{\mu,\nu,\alpha} \cdot dW_s^{\mu,\alpha}. \end{aligned}$$

and since $W^{\mu,\alpha}$ is a Wiener process under $P^{\mu,\alpha}$ and $Y^{\mu,\alpha}$ is adapted

$$Y_t^{\mu,\nu,\alpha} = \mathbb{E}^{\mu,\alpha}\left[g(X,\mu) + \int_t^T f(s,X,\mu,\nu,\alpha_s)ds \,\middle|\, \mathcal{F}_t^n \right] = J_t^{\mu,\nu}(\alpha).$$

- By comparison principle $Y_t^{\mu,\nu} \geq V_t^{\mu,\nu}$
- By measurable selection, there exists $\hat{\alpha} : [0,T] \times \mathcal{C} \times \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{P}(A) \times \mathbb{R}^d \to A$

$$H(t,x,\mu,\nu,z) = h(s,x,\mu,\nu,z,\hat{\alpha}(t,x,\mu,z)), \quad \text{for all } (t,x,\mu,\nu,z),$$

The process $\alpha^{\mu,\nu}$

$$\alpha_t^{\mu,\nu} := \hat{\alpha}(t,X,\mu,Z_t^{\mu,\nu})$$

is an optimal control, **but so is any process in the set**

$$\mathbb{A}(\mu,\nu) := \{\alpha \in \mathbb{A} : H(t,X,\mu,\nu_t,y) = h(t,X,\mu,\nu_t,Z_t^{\mu,\nu},\alpha_t) \ dt \times dP - a.e.\}$$

Final Step

Define $\Phi : \mathcal{P}_\psi(\mathcal{C}) \times \mathbb{A} \to \mathcal{P}(\mathcal{C}) \times \mathcal{M}$ by

$$\Phi(\mu, \alpha) := (P^{\mu,\alpha} \circ X^{-1}, \delta_{P^{\mu,\alpha} \circ \alpha_t^{-1}}(dq)dt)$$

The goal now is to find a point $(\mu, \nu) \in \mathcal{P}_\psi(\mathcal{C}) \times \mathcal{M}$ for which there exists $\alpha \in \mathbb{A}(\mu, \nu)$ such that $(\mu, \nu) = \Phi(\mu, \alpha)$. In other words, we seek a fixed point of the set-valued map $(\mu, \nu) \mapsto \Phi(\mu, \mathbb{A}(\mu, \nu)) := \{\Phi(\mu, \alpha) : \alpha \in \mathbb{A}(\mu, \nu)\}$.

McKean-Vlasov FBSDEs: Wishful Thinking !

Main difficulty is the analysis is the adjoint process $Z^{\mu,\nu}$.

For each (μ, ν), $Z_t^{\mu,\nu} = \zeta_{\mu,\nu}(t, X)$ and if $\hat{\alpha}$ is a measurable selection as before, **any solution of**

$$\begin{cases} dX_t = b(t, X, \mu, \hat{\alpha}(t, X, \mu, \zeta_{\mu,\nu}(t, X)))dt + \sigma(t, X)dW_t, \\ X \quad \sim \mu, \ \ \mu \circ (\hat{\alpha}(t, \cdot, \mu, \zeta_{\mu,\nu}(t, \cdot)))^{-1} = \nu_t \ a.e. \end{cases}$$

is a solution of our MFG problem

Can't solve this McKean-Vlasov SDE!

Still, some results (loosely stated)!

Theorem

- *If b, f, g are continuous in (μ, ν, α), the Hamiltonian h is concave in α, some growth conditions hold and $f = f_1(t, x, \mu, a) + f_2(t, x, \mu, \nu)$, then* **there exists a fixed point**.
- *if the Hamiltonian h is strictly concave in α, $f = f_1(t, \mu, \nu) + f_2(t, x, a)$, and $b = b(t, x, a)$, then* **the fixed point is unique**.

Approximate equilibria for the finite-player game

Theorem

If $\alpha = \alpha(t, X_\cdot)$ is an optimal feedback control for the MFG problem, then the strategy profiles $\alpha(t, X_\cdot^i)$ form an **approximate Nash equilibrium** *for the finite-player game (i.e. for some $\epsilon_n \downarrow 0$, no player can increase his expected reward by more than ϵ_n by unilaterally changing strategy).*

Price Impact Model Revisited

Price impact model corresponds to

- $b(t, x, \mu, \alpha) = \alpha$;
- σ constant;
- $g(x, \mu) = G(x)$;
- $f(t, x, \mu, \nu, \alpha) = \gamma x \int c' d\nu - c(\alpha) - F(t, x)$.

Theorem

For a bounded order book, with c' continuous, the mean field price impact model has a solution. Moreover, the errors ϵ_n are $O(1/\sqrt{n})$.

Some Examples of New Types of Interactions

- **Rank Effects**
  - $f(t, x, \mu, q, a)$ contains $G(\mu_t(-\infty, x_t])$
  - Oil production model (Guéant-Lasry-Lions)
- **Quantile Interactions**
  - $f(t, x, \mu, q, a)$ involves the quantile function $y \hookrightarrow F_{\mu_t}^{-1}(y) = \inf\{x \in \mathbb{R};\ \mu_t(-\infty, x] \geq y\}$
- **Functions of the Density of the Population** à la Lasry - Lions
  - ???????????????????