Entropic Optimal Planning

Nizar Touzi

Ecole Polytechnique, France

November 17, 2021

Stochastic control of interacting dynamics

Agents $i = 1, \dots, N$

Each agent i controls the state processes

$dX_t^i = \alpha_t^i dt + dW_t^i$, W^i indep Brownian motions

Denote $\alpha^{(-i)} := (\alpha^1, \dots, \alpha^{i-1}, \alpha^{i+1}, \dots \alpha^n)$

Nash Equilibrium

- Agent i's individual optimization :
$$V^i(\alpha^{(-i)}) \quad := \quad \inf_{\alpha^i} \mathbb{E}\Big[\int_0^T \Big\{\frac{1}{2}|\alpha_t^i|^2 - f\Big(X_t^i, \frac{1}{N}\sum_{j=1}^N \delta_{X_t^j}\Big)\Big\}dt + g(X_T^i)\Big]$$
$\Longrightarrow$ optimal response $\hat\alpha^i(\alpha^{(-i)})$
- A Nash equilibrium is $(\alpha^{*,1}, \dots, \alpha^{*,N})$ such that $\alpha^{*,i} \in \hat\alpha^i(\alpha^{*,(-i)})$

Coupled system of HJB equations

Value function of Agent i characterized by means of HJB equation

$$\partial_t v^i + \frac{1}{2}\Delta v^i + \sum_{j\neq i}\alpha^j\partial_{x_j}v^i + \underbrace{\inf_{\alpha^i}\left\{\alpha^i\partial_{x_i}v^i + \frac{1}{2}|\alpha^i|^2\right\}}_{=-\frac{1}{2}|\partial_{x_i}v^i|^2} = f\big(x^i,\bar{\mu}_N(x)\big)$$

$$v^i(T,.) = g$$

with optimal control $\hat{\alpha}^i(\alpha^{(-j)}) = -\partial_{x_i}v^i$, where $v^i = v^i(t,x;\alpha^{(-i)})$

Nash equilibrium : $\alpha^* = (\alpha^{*,1},\ldots,\alpha^{*,N})$ such that

$$\alpha^{*,i} = -\partial_{x_i}v^i(t,x;\alpha^{*,(-i)}) \quad \text{for all} \quad i = 1,\ldots,N$$

... raises many technical difficulties !

SDE and Focker-Planck equation

To simplify, send $N \to \infty \implies$ interaction through marginal distribution

$$\bar{\mu}_N(X_t) := \frac{1}{N} \sum_{i=1}^N \delta_{X_t^i} \longrightarrow \text{ Law of } X_t$$

If X solution of SDE $dX_t = b(t, X_t)dt + \sigma(t, X_t)dW_t$, its marginal law $m(t, dx) := \mathbb{P} \circ X_t^{-1}$ characterized by the Fokker-Planck equation

$$\partial_t m + \text{div}[bm] - \frac{1}{2} \sum_{i,j} \partial^2_{i,j} \{(\sigma\sigma^\intercal)_{ij} m\} = 0, \quad m|_{t=0} = \delta_{\{X_0\}}$$

SDE at Nash equilibrium is

$$dX_t = -\partial_x v(t, X_t, m(t, X_t))dt + dW_t$$

and the corresponding density characterized by the FP equation

$$\partial_t m + \text{div}[-\partial_x v\, m] - \frac{1}{2} \Delta m = 0, \quad m\Big|_{t=0} = \mu_0$$

Mean Field Games

MFG, Huang, Malhamé & Caines '06 and Lasry & Lions '06

Lasry & Lions' formulation $\partial_t v + \frac{1}{2}\Delta v - \frac{1}{2}|Dv|^2 = f(x,m),\ v\big|_{t=T} = g \quad (HJB)$

$$\partial_t m - \frac{1}{2}\Delta m - \operatorname{div}\big[Dv\ m\big] = 0,\ m\big|_{t=0} = \mu_0 \quad (FP)$$

(HJB) : Representative agent optimization problem, parametrized by m

$$\inf_{\alpha} \mathbb{E}\Big[\int_0^T \{\frac{1}{2}|\alpha_t|^2 - f(X_t, m(t,X_t))\}dt + g(X_T)\Big], \quad dX_t = \alpha_t dt + dW_t$$

Hamiltonian $H(z) = \inf_a \Big\{az - \frac{|a|^2}{2}\Big\} = \frac{|z|^2}{2} \Longrightarrow$ opt. cont. $\alpha_t^* = Dv(t,X_t)$

(FP) : $m =$marginal distribution of diffusion with transport coefficient given by optimal control of the HJB : $dX_t = Dv(t,X_t)dt + dW_t$

$\Longrightarrow$ **Nash equilibrium**

P.L. Lions' Planning Problem

Planning Problem

μ_0, μ_T given probability measures on $\mathbb{R}^d$, solve

$$\partial_t v + \frac{1}{2}\Delta v - \frac{1}{2}|Dv|^2 = f(x,m)$$

$$\partial_t m - \frac{1}{2}\Delta m - \mathrm{Div}\big[Dv\, m\big] = 0 \qquad m\big|_{t=0} = \mu_0 \;\text{ and }\; m\big|_{t=T} = \mu_T$$

(in particular $g = v\big|_{t=T}$ to be determined)

Unique solution exists for any pair (m_0, m_T)...

Lions '10, Achdou, Camilli & Capuzzo Dolcetta '12
Porretta '14, Orrieri, Porretta & Savaré '18
Graber, Mészáros & Tonon '18, Benamou
Carlier, Di Marino & Nena '18

Figure 1: EU energy mix evolution (55 percent lower emissions in 2030 compared to 1990 and climate neutrality in 2050)

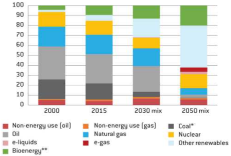

ÉCOLE POLYTECHNIQUE
UNIVERSITÉ PARIS-SACLAY

Path-dependent stochastic control

$\Omega = C^0(\mathbb{R}_+, \mathbb{R}^d)$, canonical process $X_t(\omega) = \omega(t)$, $t \geq 0$
$\mathbb{P}_0$: Wiener measure on Ω

Croud of agents defined by probability distribution m

for fixed m, solve the representative Agent problem :

$$V_0(m, \xi) := \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}^{\mathbb{P}}\Big[\xi(X) - \int_0^T c_t(m(t), \nu_t)dt\Big]$$

where, for some control ν valued in U, $\mathbb{P} \in \mathcal{P}$ is weak solution of

$\mathbb{P} \circ (X_0)^{-1} = m_0$, and $dX_t = \sigma_t(X, \nu_t)\big[\lambda_t(X, \nu_t)dt + dW_t^{\mathbb{P}}\big]$, $\mathbb{P}-$a.s.

Two additional features :

- control affects both drift (transport) and diffusion
- all coefficients are possibly path dependent
- **path dependence of $\xi = \xi(X_{\wedge T})$ is crucial !**

Agents in (Path-dependent) mean field Nash equilibrium

Denote $\hat{\mathcal{P}}(m,\xi) := \{\text{solutions of } V_0(m,\xi)\}$

Definition (Mean field game equilibrium)

$\hat{m}$ is a MFG equilibrium if there exists

$$\hat{\mathbb{P}} \in \hat{\mathcal{P}}(\hat{m},\xi) \quad \text{such that} \quad \hat{\mathbb{P}} \circ (X_t)^{-1} = \hat{m}(t), \quad \text{for all} \quad t \le T$$

Denote MFG(ξ)=$\{$solutions of MFG equilibrium$\}$

Path-dependent formulation of the Planning Problem

Path-dependent Planning Problem

Given μ_0, μ_T probability measures on $\mathbb{R}^d$,

find $\xi \in \mathbb{L}^0(\mathcal{F}_T)$ and $\hat{m} \in \mathrm{MFG}(\xi)$ such that $\hat{m}(0) = \mu_0, \ \hat{m}(T) = \mu_T$

- ξ may be interpreted as the incentive regulation so as to optimally move the population from μ_0 to μ_T
- More freedom than the original Lions' planning problem where $\xi(X) = g(X_T)$
- Multiple solutions, in general...
- Relation with contract theory : 1 Principal facing a crowd of Agents in Nash equilibrium

Forward description of MFG equilibria

Hamiltonian of the representative agent problem (with $b := \sigma\lambda$) :

$$H_t(z,\gamma,\mu) := \sup_{u\in U}\Big\{b_t(u)\cdot z + \frac{1}{2}\sigma_t(u)^2 : \gamma - c_t(\mu,u)\Big\}$$

Proposition (Ren, Tan & NT '21)

Let $p>1$ and $\xi \in \mathcal{L}^p(\mathcal{P})$ be such that $\mathrm{MFG}(\xi) \neq \emptyset$. Then, we may find $Y_0 \in \mathbb{R}$, $Z \in \mathcal{H}^p(\mathcal{P})$, and $\mathbb{F}-$prog.meas. Γ such that

$$\mathrm{MFG}(\xi) = \mathrm{MFG}\big(Y_T^{Z,\Gamma}\big), \quad \text{where}$$

- $Y_T^{Z,\Gamma} := \int_0^T Z_t\cdot dX_t + \frac{1}{2}\Gamma_t : d\langle X\rangle_t - H_t\big(Z_t,\Gamma_t,\mu(t)\big)dt$, $\mathcal{P}-$q.s.
- $\mu(t) = \mathbb{P}\circ X_t^{-1}$ is defined by the McKean-Vlasov controlled dynamics

$$dX_t = \nabla_z H_t(Z_t,\Gamma_t,\mu(t))dt + \sqrt{2\nabla_\gamma H_t(Z_t,\Gamma_t,\mu(t))}\, dW_t, \quad X_0 \sim \mu_0$$

2nd order backward SDEs : Soner, NT & Zhang, Possamaï, Tan & Zhou
Principal-Agent problem : Cvitanić, Passamaï & NT and Elie & Possamaï

Here, we extend to the quadratic setting

Intuitions

- Uncontrolled diffusion case : solve the representative Agent problem

$$V_0(m,\xi) := \sup_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\Big[\xi(X) - \int_0^T c_t(m(t),\nu_t)dt\Big]$$

where, for some control ν valued in U, $\mathbb{P}\in\mathcal{P}$ is weak solution of

$$\mathbb{P}\circ(X_0)^{-1} = m_0, \text{ and } dX_t = \sigma_t(X)\big[\lambda_t(X,\nu_t)dt + dW_t^{\mathbb{P}}\big], \ \mathbb{P}-\text{a.s.}$$

Then the dynamic version of the problems satisfies the BSDE with final $Y_T = \xi$ (El Karoui, Peng & Quenez '97). Notice m is fixed !

- With diffusion control
  - same type of representation for all fixed diffusion, martingale optimality principle induces an non-decreasing process A
  - A vanishes on the support of optimal measure
  - our representation follows by approximating A by an appropriate sequence of a.c. nondecreasing processes

Implication for the Lions planning problem

Definition

Given μ_0, μ_T, Plan(μ_0, μ_T) is the set of all processes Z, Γ such that

- There is a solution to the McKean-Vlasov SDE

$$dX_t = \nabla_z H_t(Z_t, \Gamma_t, \mu(t))dt + \sqrt{2\nabla_\gamma H_t(Z_t, \Gamma_t, \mu(t))}\, dW_t, \quad X_t \sim \mu(t)$$

- $\mu(0) = \mu_0$ and $\mu(T) = \mu_T$

ξ is a solution of the Lions' (path dependent) planning problem **Iff**

$$\xi := Y_T^{Z,\Gamma} := \int_0^T Z_t \cdot dX_t + \frac{1}{2}\Gamma_t : d\langle X\rangle_t - H_t\big(Z_t, \Gamma_t, \mu(t)\big)dt$$

for some $(Z, \Gamma) \in$ Plan(μ_0, μ_T)

Extensions

- Multi-marginal planning problem
- Discount factor in the control problem $\Longrightarrow$ H depends on Y $\Longrightarrow$ Solve a McKean-Vlasov SDE for (X, Y)
- Elliptic setting $\Longrightarrow$ random horizon backward SDEs

Optimal planning

Select a solution of the planning problem

$$\sup_{(Z,\Gamma)\in \mathrm{Plan}(\mu_0,\mu_T)} \mathbb{E}^{\mathbb{P}^{Z,\Gamma}}\Big[\int_0^T \ell(X_t, \mathbb{P}\circ X_t^{-1})dt + L(X_T, \mathbb{P}\circ X_T^{-1}, Y_T)\Big]$$

Back to the Lions purely quadratic setting

Representative agent problem :

$$\inf_{\alpha} \mathbb{E}^{\mathbb{P}^\alpha}\Big[\int_0^T \Big\{ \frac{1}{2}|\alpha_t|^2 - f(X_t, m(t,X_t)) \Big\} dt + \xi \Big], \; dX_t = \alpha_t dt + dW_t, \; \mathbb{P}^\alpha - \text{a.s.}$$

Let $H(x,z,\gamma,m) = H^0(x,z,m) + \frac{1}{2}\mathrm{Tr}[\gamma]$, where

$$H^0(x,z,m) := \inf_a \Big\{ az + \frac{1}{2}a^2 - f(x,m) \Big\} = -\frac{1}{2}z^2 - f(x,m)$$

As $\nabla_z H = -z$ and $2\nabla_\gamma H = I_d$ (independent of m), we have

All solutions of the Lions planning problem are of the form

$$\xi := \int_0^T Z_t dX_t - H^0(X_t, Z_t, m(t))dt$$

for some $Z \in \mathrm{Plan}(\mu_0, \mu_T)$, i.e. $Z \in \mathbb{H}^0$ and

$$dX_t = -Z_t dt + dW_t, \; \mathbb{P}^{-Z} - \text{a.s.} \; \mathbb{P}^{-Z} \circ X_0^{-1} = \mu_0, \; \text{and } \mathbb{P}^{-Z} \circ X_T^{-1} = \mu_T$$

Entropic optimal planning in the purely quadratic MFG

$\mathbf{Q}, \mathbf{R} \in \mathrm{Prob}(\Omega)$, **entropy** of $\mathbf{Q}$ wrt $\mathbf{R}$:

$$H\big(\mathbf{Q}|\mathbf{R}\big) := \mathbb{E}^{\mathbf{Q}}\Big[\ln \frac{d\mathbf{Q}}{d\mathbf{R}}\Big] = \int \ln \frac{d\mathbf{Q}}{d\mathbf{R}} d\mathbf{Q}, \quad \text{for all} \quad \mathbf{Q} \ll \mathbf{R}, \ (\infty, \ \text{otherwise})$$

For $Z \in \mathrm{Plan}(\mu_0, \mu_T)$, denote $\mathbb{P}^Z$ the probability on Ω defined by

$$dX_t = -Z_t dt + dW_t, \quad X_0 \sim \mu_0, \quad \text{and} \quad X_T \sim \mu_T$$

Minimum entropy optimal planning

Given μ_0, μ_T probability measure on dbR^d, solve

$$\min_{Z \in \mathrm{Plan}(\mu_{\mathbf{0}}, \mu_T)} H\big(\mathbb{P}^Z | \mathbb{P}_0\big)$$

If Z^* is a solution, then $\mathbb{P}^* := \mathbb{P}^{Z^*}$ is a minimum entropy optimal planning from μ_0 to μ_T

Explicit solution of the minimum entropy planning problem

Proposition

Denote $m_T := \mu_0 * \mathcal{N}(0, TI_d)$, and assume

$$\mu_T \sim \text{Leb}_{\mathbb{R}^d}, \quad \text{and} \int \ln\left(\frac{d\mu_T}{dm_T}\right)^- dm_T + \int \left(\frac{d\mu_T}{dm_T}\right)^2 dm_T < \infty$$

Then the minimum entropy planning problem has a unique solution $Z^* = -\theta$ defined by

$$\frac{d\mu_T}{dm_T}(X_T) = e^{\int_0^T \theta_t \cdot dX_t - \frac{1}{2}\int_0^T |\theta_t|^2 dt}, \mathbb{P}_0 - a.s.$$

Consequently, $\xi^* := \int_0^T Z_t^* dX_t - H^0(X_t, Z_t^*, m(t))dt$ is a (path-dependent) solution of the Lions optimal planning problem

Hamiltonian with superlinear growth in the gradient

Proposition

Assume that

$$\frac{\nabla_z H_t^0(\omega, y, z, m)}{f(|z|)} = O(1) \quad \text{as} \quad |z| \to \infty$$

for some continuous f with $f(0) = 0$

Then, under the conditions of the previous theorem, the minimum entropy planning problem has a unique solution Z^* defined by

$$\nabla_z H_t^0(Y_t, Z_t^*, \mu_t) \;=\; \theta_t, \;\mathbb{P}_0 - \text{a.s.}$$

Diffusion control, uncontrolled drift

for fixed m, solve the representative Agent problem :

$$V_0(m,\xi) := \sup_{\mathbb{P}\in\mathcal{P}} \mathbb{E}^{\mathbb{P}}\Big[\xi(X) - \int_0^T c_t(m(t),\nu_t)dt\Big]$$

where, for some control ν valued in U, $\mathbb{P}\in\mathcal{P}$ is weak solution of

$$\mathbb{P}\circ(X_0)^{-1} = m_0, \text{ and } dX_t = \sigma_t(X,\nu_t)dW_t^{\mathbb{P}}, \ \mathbb{P}-\text{a.s.}$$

Marginal distributions are in convex order :

$$t \longmapsto \int f(x)\mu_t(dx) \ \text{ nondecreasing}$$

Hamiltonian of the representative agent problem :

$$H_t(\gamma,\mu) := \sup_{u\in U}\Big\{\frac{1}{2}\sigma_t^2 : \gamma - c_t(\mu,u)\Big\}$$

Hamiltonian with superlinear growth in the Hessian

Ongoing work...

Main result

Let $\mu_T \in \text{Prob}(\mathbb{R})$ with $\mu_0 \preceq \mu_T$ in convex order, and assume that

- $\nabla_z H \equiv 0$, i.e. uncontrolled drift
- $\sigma_t(x, U) = \mathbb{S}_d^+$ and ... as $|\gamma| \to \infty$

Then, there exists an MFG equilibrium transporting μ_0 to μ_T

Idea of proof : Start from a solution of the Skorohod Embedding problem

$$\Longrightarrow dX_t = \sigma_t(X) dW_t$$

- $d = 1$: τ stopp time s.t. $B_0 \sim \mu$ and $B_\tau \sim \mu_T$, then $X_t := X_0 + B_{\frac{t}{T-t} \wedge \tau}$ defines a continuous martingale that we may represent as above by the Dubins-Schwarz theorem
- $d \geq 2$...

Hamiltonian with superlinear growth in the Hessian

Denote $\mu(t) := \mathbb{P} \circ X_t^{-1}$ and define Γ_t^* by

$$\sqrt{2\nabla_\gamma H_t\big(X, \Gamma_t^*, \mu(t)\big)} \quad = \quad \sigma_t(X)$$

A solution of the Lions' (path dependent) planning problem is given by

$$\xi^* := \frac{1}{2}\Gamma_t : d\langle X\rangle_t - H_t\big(\Gamma_t^*, \mu(t)\big)dt$$

for some $(Z, \Gamma) \in \mathrm{Plan}(\mu_0, \mu_T)$

• Any solution of the Skorohod embedding problem induces a solution of the optimal planning problem

• Optimality ? ?