Mean-Field optimization problems and non-anticipative optimal transport

Beatrice Acciaio
London School of Economics

based on ongoing projects with J. Backhoff,
R. Carmona and P. Wang

Thera Stochastics
A Mathematics Conference in Honor of Ioannis Karatzas
Thera, Santorini, May 31 - June 2, 2017

The story in a nutshell

Given

- a (finite or infinite) set of agents
- who need to choose their actions/strategies
- and face a cost depending on their own type, action, and on the symmetric interaction with each other:

$$\text{cost}(i) = \text{fct}\left(\text{type}(i), \text{action}(i), \text{(empirical) distrib. actions}\right)$$

Aim to

→ find/characterize **equilibria**

→ through connections with **non-anticipative optimal transport**

Outline

1. First setting: hidden/no dynamics
    - Problem formulation
    - Connection with non-anticipative optimal transport
    - Existence and uniqueness results
2. Second setting: state dynamics
    - Problem formulation
    - Connection with non-anticipative optimal transport
    - First results
3. Conclusions

Setting

- time set: $\mathbb{T} = \{0, ..., T\}$, or $\mathbb{T} = [0, T]$
- $\mathbb{X}$: agents **types**

  $\mathcal{X} \subseteq \mathbb{X}^{|\mathbb{T}|}$: agents types evolutions
- $\mathbb{Y}$: agents' **actions**

  $\mathcal{Y} \subseteq \mathbb{Y}^{|\mathbb{T}|}$: agents' actions evolutions
- e.g. $\mathbb{X}=\mathbb{Y}=\mathbb{R}$, and $\mathcal{X}=\mathcal{Y}=\mathbb{R}^{T+1}$ or $\mathcal{X}=\mathcal{Y}=\mathcal{C}([0,T];\mathbb{R})$
- $\eta \in \mathcal{P}(\mathcal{X})$: known a priori distribution over types

→ cost function: $\tilde{c}(x, y, \nu)$ (for each agent)

$\nearrow$	$\uparrow$	$\nwarrow$
type	action	actions' distribution
$x \in \mathcal{X}$	$y \in \mathcal{Y}$	$\nu \in \mathcal{P}(\mathcal{Y})$

Cost function

Separable structure: $\tilde{c}(x, y, \nu) = c(x, y) + V[\nu](y)$

$\nearrow$ idiosyncratic part

$\nwarrow$ mean-field interaction

with $c : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}_+$ l.s.c., $V : \mathcal{P}(\mathcal{Y}) \to \mathcal{B}(\mathcal{Y}; \mathbb{R}_+)$

- **congestion effect**: $V_c[\nu](y) = f\left(y, \frac{d\nu}{dm}(y)\right)$, with $m \in \mathcal{P}(\mathcal{Y})$ reference meas. w.r.t. which congestion measured, $f(y, .) \nearrow$
- **attractive effect**: $V_a[\nu](y) = \int_{\mathcal{Y}} \phi(y, z)\nu(dz)$, with ϕ symmetric, convex, minimal on the diagonal

Static case: Blanchet-Carlier 2015

Pure adapted strategies

pure strategy: all players of type $x \in \mathcal{X}$ choose the same strategy

$$y = A(x) = (A_t(x))_{t \in \mathbb{T}}$$

adapted strategy: $A_t(x) = T^t(x_{0:t})$ for some measurable T^t

Denote by $\mathcal{A}$ the set of pure adapted strategies $A : \mathcal{X} \to \mathcal{Y}$

- type distribution: $\eta \in \mathcal{P}(\mathcal{X})$ (known)
- strategy distribution: $\nu = A_{\#}\eta = T_{\#}\eta \in \mathcal{P}(\mathcal{Y}), \quad T = (T^t)_{t \in \mathbb{T}}$ (will be determined in equilibrium)

Pure equilibrium

Social planner perspective: minimize average cost

For every $\nu \in \mathcal{P}(\mathcal{Y})$, denote

$$P(\nu) := \inf_{A \in \mathcal{A}} \int \left\{ c(x, A(x)) + V[\nu](A(x)) \right\} \eta(dx)$$

Definition

An element $A \in \mathcal{A}$ is called a pure equilibrium if

- A attains $P(\nu)$,
- where $\nu = A_{\#}\eta$.

Cournot-Nash equilibrium

Remark. Let $\mathbb{T} = \{0, 1, ..., T\}$ (analogous in continuous time). Let $c(x, y) = \sum_{t=0}^{T} c_t(x_{0:t}, y_t)$ and $V[\nu](y) = \sum_{t=0}^{T} V_t[\nu_t](y_t)$, then

pure equilibrium for social planner = Cournot-Nash equilibrium

(η-a.s. each agent acts as best response to other agents' actions)

The equilibria are described by the set $\left\{A^{\nu} : A^{\nu}_{\#}\eta = \nu\right\}$, where

$$A^{\nu}_t(x) = T^{\nu}_t(x_{0:t}) := \arg\min_z \left\{c_t(x_{0:t}, z) + V_t[\nu_t](z)\right\}.$$

- This is clearly a specific situation
- Anyway, pure equilibria rarely exists, so we shall consider the natural generalization to mixed-strategy equilibria.

From pure to mixed-strategy equilibrium

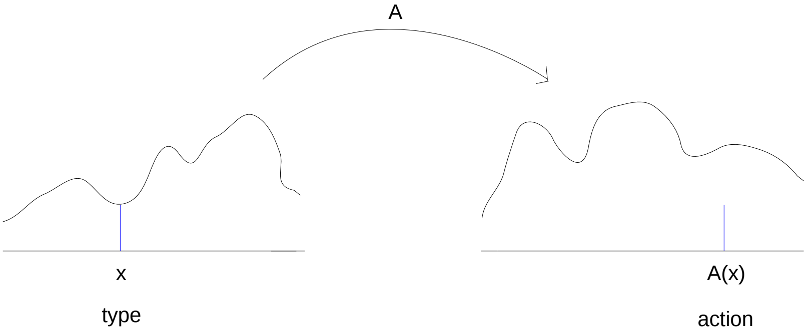

adapted **pure** strategy = adapted **Monge** transport

From pure to mixed-strategy equilibrium

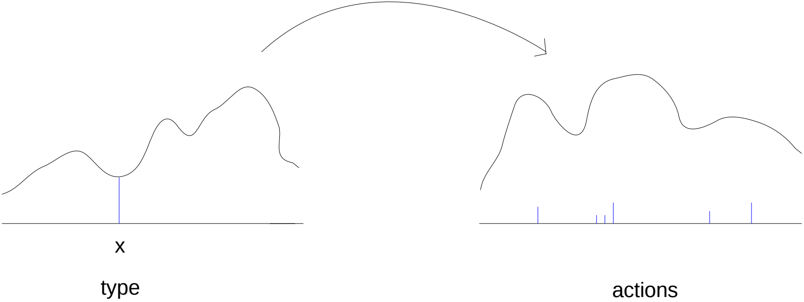

non-anticipative **mixed** strategy = causal **Kantorovich** transport

Mixed non-anticipative strategy

mixed-strategy: players of same type can choose different actions

non-anticipative: $A_t(x) = \text{fct}(x_{0:t}) + \text{sth indep. of } x$

$\downarrow$

Non-anticipative (causal) transport: $\pi \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ s.t. $p_{1\#}\pi = \eta$, and for all t and $D \in \mathcal{F}_t^{\mathcal{Y}}$, the map $\mathcal{X} \ni x \mapsto \pi^x(D)$ is $\mathcal{F}_t^{\mathcal{X}}$-measurable (where $(\mathcal{F}_t^{\mathcal{X}})$,$(\mathcal{F}_t^{\mathcal{Y}})$ canonical filtr. in $\mathcal{X}$,$\mathcal{Y}$, and π^x reg. cond. kernel)

Denote by $\Pi_c(\eta, \nu)$ the set of causal transports between η and ν, and let $\Pi_c(\eta, .) := \bigcup_{\nu \in \mathcal{P}(\mathcal{Y})} \Pi_c(\eta, \nu)$

Note that $\pi = (id, T)_{\#}\eta \in \Pi_c(\eta, .)$ are the pure adapted strategies.

Mixed-strategy equilibrium

For every $\nu \in \mathcal{P}(\mathcal{Y})$, denote

$$M(\nu) := \inf_{\pi \in \Pi_c(\eta,.)} \mathbb{E}^{\pi}\Big[c(x,y) + V[\nu](y)\Big]$$

Definition

An element $\pi \in \Pi_c(\eta,.)$ is called a mixed-strategy equilibrium if

- π attains $M(\nu)$,
- where $\nu = p_{2\#}\pi$, i.e., $\pi \in \Pi_c(\eta,\nu)$.

Remark. Mixed-strategy equilibria are solutions to causal transport problems: if π^* m-s equilibrium, with $p_{2\#}\pi^* = \nu^*$, then it attains

$$\inf_{\pi \in \Pi_c(\eta,\nu^*)} \mathbb{E}^{\pi}[c(x,y)].$$

Analogously, pure equilibria = solutions to CT pbs over Monge maps

Potential games

From the remark, we always have equilibrium $\Longrightarrow$ optimal transport

For **potential games**, we will have "$\Longleftrightarrow$" in some sense

Assumption

There exists $\mathcal{E} : \mathcal{P}(\mathcal{Y}) \to \mathbb{R}$ *such that* V *is the* **first variation** *of* $\mathcal{E}$*:*

$$\lim_{\epsilon \to 0^+} \frac{\mathcal{E}(\nu + \epsilon(\mu - \nu)) - \mathcal{E}(\nu)}{\epsilon} = \int_{\mathcal{Y}} V[\nu](y)(\mu - \nu)(dy), \quad \forall\, \nu, \mu \in \mathcal{P}(\mathcal{Y})$$

E.g. $V = V_c + V_a$ (repulsive+attractive effect) is the first variation of

$$\mathcal{E}(\nu) = \int_{\mathcal{Y}} F\left(y, \frac{d\nu}{dm}(y)\right) m(dy) + \frac{1}{2} \int_{\mathcal{Y} \times \mathcal{Y}} \phi(y, z)\nu(dz)\nu(dy),$$

where $F(y, u) = \int_0^u f(y, s)ds$.

Potential games

Consider the variational problem

$$(VP) \qquad \inf_{\nu \in \mathcal{P}(\mathcal{Y})} \left\{ \underbrace{\inf_{\pi \in \Pi_c(\eta,\nu)} \mathbb{E}^{\pi}[c(x,y)]}_{\mathrm{CT}(\eta,\nu)} + \mathcal{E}[\nu] \right\}$$

Theorem

Let $\mathcal{E}$ be convex, then the following are **equivalent***:*

(i) π^* *is a mixed-strategy equilibrium, with* $p_{2\#}\pi^* = \nu^*$;

(ii) ν^* *solves (VP), and* π^* *solves* $CT(\eta, \nu^*)$.

Remarks. 1. Convexity only needed for “$(i) \Rightarrow (ii)$”

2. Convexity satisfied in the congestion case ($V = V_c$)

3. Alternatively: displacement convexity can be used

Potential games

Corollary (uniqueness)

If $\mathcal{E}$ strictly convex $\Rightarrow$ all m-s equilibria have same second marginal ν^*, i.e., unique optimal distribution of actions.

Indeed, $\nu \mapsto \mathrm{CT}(\eta,\nu)$ convex, hence $\mathcal{E}$ strictly convex implies unique solution ν^* for (VP). Then apply theorem.

Corollary (existence)

For $V = V_c$ and growth condition on f $\Rightarrow$ $\exists$ m-s equilibrium.

Indeed, the growth condition ensures existence of a solution ν^* for (VP), and $\mathrm{CT}(\eta,\nu^*)$ admits a solution π^* since c is bounded below and l.s.c. Then apply theorem.

Example

Let $\mathbb{T} = \{0, 1, ..., T\}$, and $\mathcal{X} = \mathcal{Y} = \mathbb{R}^{T+1}$. If

- η has independent increments, and
- $c(x, y) = c_0(x_0, y_0) + \sum_{t=1}^{T} c_t(x_t - x_{t-1}, y_t - y_{t-1})$, with $c_t(u, v) = k_t(u - v)$ and k_t convex,

Then:

- m-s equilibria (if $\exists$) are determined by the second marginal
- m-s equilibria are the **Knothe-Rosenblatt rearrangements**
- if moreover η has a density, all m-s equilibria are in fact pure

The Knothe-Rosenblatt map

The Knothe-Rosenblatt map

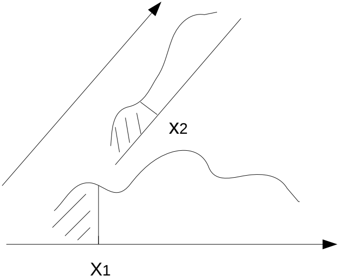

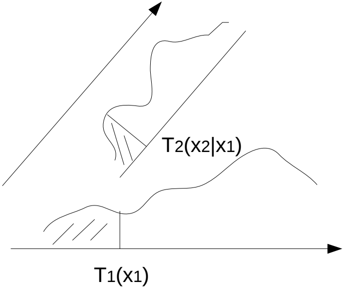

Actions as controls on dynamics

- The previous result describes a specific situation where optimal actions are increasing with the type.
- When these conditions not satisfied, which form of CT/equilibria?

Example. Let actions = controls on dynamics:

$$X_t = (k_t^1 X_{t-1} + k_t^2 \alpha_t) + \epsilon_t, \; t = 1, ..., T, \; X_0 = x_0,$$

with associated cost $f_t(X_t, \alpha_t, \nu_t)$ at time t. As $X_t = fct(\epsilon_i, \alpha_i, i \leq t)$,

$$f_t(X_t, \alpha_t, \nu_t) = c_t(\epsilon_{0:t}, \alpha_{0:t}, \nu_t),$$

hence total cost $= \mathbb{E}[\sum_{t=0}^{T} c_t(\epsilon_{0:t}, \alpha_{0:t}, \nu_t)]$.

$\hookrightarrow$ Fits into previous framework, by reading “noises as types”.

McKean-Vlasov control problem

- With the above example in mind, we will consider

McKean-Vlasov control problem:

$$\inf_{\alpha} \mathbb{E}^{\mathbb{P}}\left[\int_0^T \tilde{f}_t\left(X_t, \alpha_t, \mathbb{P}\circ(X_t, \alpha_t)^{-1}\right) dt + \tilde{g}\left(X_T, \mathbb{P}\circ X_T^{-1}\right)\right]$$

subject to

$$dX_t = b_t\left(X_t, \alpha_t, \mathbb{P}\circ X_t^{-1}\right) dt + dW_t$$

- Let us fist mention connections to large systems of interacting controlled state processes

N-player stochastic differential game

The private state process X^i of player i is given by the solution to

$$dX_t^i = b_t(X_t^i, \alpha_t^i, \bar{\mu}_t^N)dt + dW_t^i$$

- $W^1, ..., W^N$ independent Wiener processes
- $\alpha^1, ..., \alpha^N$ controls of the N players
- $\bar{\mu}_t^N = \frac{1}{N-1}\sum_{j \neq i} \delta_{X_t^j}$ empirical distrib. of states of the other players

The objective of player i is to choose a α^i in order to minimize

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

- $\bar{\nu}_t^N = \frac{1}{N-1}\sum_{j \neq i} \delta_{(X_t^j, \alpha_t^j)}$ empirical joint distrib. of states and controls of the other players

Statistically identical players: same functions $b_t, \tilde{f}_t, \tilde{g}$

From N-player game to McKean-Vlasov control problem

Approximation by **asymptotic arguments**:

- first optimization then limit for $N \to \infty$, or
- viceversa, first limit for $N \to \infty$ and then optimization

$$\begin{tikzcd}
\text{SDE State Dynamics for N players} \arrow[r, "\text{optimization}"] \arrow[dd, "\lim_{N\to\infty}"'] & \text{Nash equilibrium for N players} \arrow[d, "\lim_{N\to\infty}"] \\
& \text{Mean-Field Game} \\
\text{McKean-Vlasov dynamics} \arrow[r, "\text{optimization}"] & \text{controlled McK-V dyn}
\end{tikzcd}$$

(Carmona-Delarue-Lachapelle 2012)

McKean-Vlasov control problem

Back to the McKean-Vlasov control problem.

For simplicity:

- no terminal cost: $\tilde{g} = 0$
- separable costs: $\tilde{f}_t(x, a, \nu) = f_t(x, a) + K_t(\nu)$

Therefore

$$\inf_{\alpha} \mathbb{E}^{\mathbb{P}} \left[\int_0^T \Big\{ f_t(X_t, \alpha_t) + K_t \Big(\mathbb{P} \circ (X_t, \alpha_t)^{-1} \Big) \Big\} dt \right]$$
$$dX_t = b_t \Big(X_t, \alpha_t, \mathbb{P} \circ X_t^{-1} \Big) dt + dW_t,$$

with $f_t : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, $K_t : \mathcal{P}(\mathbb{R} \times \mathbb{R}) \to \mathbb{R}$, $b_t : \mathbb{R} \times \mathbb{R} \times \mathcal{P}(\mathbb{R}) \to \mathbb{R}$

McKean-Vlasov control problem

Definition. A weak solution to the McKean-Vlasov control problem is a tuple $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P}, W, X, \alpha)$ such that:

(i) $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$ supports X and a BM W, α is $\mathcal{F}^X$-progress. measurable and $\mathbb{E}^{\mathbb{P}}\Big[\int_0^T |\alpha_t|^2 \Big] < \infty$

(ii) the state equation $dX_t = b_t\left(X_t, \alpha_t, \mathbb{P}\circ X_t^{-1}\right) dt + dW_t$ holds

(iii) if $(\Omega', (\mathcal{F}'_t)_{t\in[0,T]}, \mathbb{P}', W', X', \alpha')$ is another tuple s.t. (i)-(ii) hold,

$$\mathbb{E}^{\mathbb{P}}\left[\int_0^T\left\{f_t(X_t,\alpha_t)+K_t\left(\mathbb{P}\circ(X_t,\alpha_t)^{-1}\right)\right\}dt\right] \le \mathbb{E}^{\mathbb{P}'}\left[\int_0^T\left\{f_t(X'_t,\alpha'_t)+K_t\left(\mathbb{P}'\circ(X'_t,\alpha'_t)^{-1}\right)\right\}dt\right].$$

Assumptions

→ We need some **technical assumptions**.

→ In the case of linear drift:

$$dX_t = (c_t^1 X_t + c_t^2 \alpha_t + c_t^3 \mathbb{E}[X_t])dt + dW_t,$$

$c_t^i \in \mathbb{R}, c_t^2 > 0$, the assumptions reduce to:

- $f_t(x, .)$ convex (and $f_t(., y)$ at least quadratic growth)
- K_t is $\prec_c$-monotone

Example.

- $f_t(x, a) = d_t^1 x + d_t^2 a + d_t^3 x^2 + d_t^4 a^2, \;\; d_t^i \in \mathbb{R}, d_t^4 > 0$
- $K_t(\zeta) = F_t(\bar{\zeta}_1, \bar{\zeta}_2)$, any F_t, $\bar{\zeta}_i := \int y d(p_{i\#}\zeta)(y)$

Characterization via non-anticipative optimal transport

- formulate a transport problem in the path space $\mathcal{C}([0,T])$
- denote by γ the Wiener measure on $\mathcal{C}([0,T])$
- $(\omega,\overline{\omega})$ generic element on $\mathcal{C}([0,T])\times\mathcal{C}([0,T])$
- "move noises into states"

Theorem

Under the mentioned assumptions, the weak MKV problem is **equivalent** *to the variational problem*

$$\inf_{\mu\ll\gamma}\ \inf_{\pi\in\Pi_{bc}(\gamma,\mu)}\left\{\mathbb{E}^{\pi}\left[\int_0^T f_t\left(\overline{\omega}_t, u_t(\omega,\overline{\omega},\mu)\right)dt\right]+\int_0^T K_t\Big(\big(p_2, u_t(\omega,\overline{\omega},\mu)\big)_{\#}\pi_t\Big)dt\right\}$$

where $u_t(\omega,\overline{\omega},\mu)=b_t^{-1}(\overline{\omega}_t,.,\mu_t)\Big(\big(\overline{\omega}\dot{-}\omega\big)_t\Big).$

$$\Pi_{bc}(\gamma,\mu)=\Big\{\pi\in\Pi_c(\gamma,\mu):\ell_{\#}\pi\in\Pi_c(\mu,\gamma)\Big\},\ \ \text{where } \ell(x,y)=(y,x)$$

Characterization via non-anticipative optimal transport

Remarks.

- The optimization over $\Pi_{bc}(\gamma, \mu)$ is not a standard optimal transport problem ⇒ new analysis for existence/duality.
- When mean-field cost is $K_t(\mathbb{P} \circ X_t^{-1})$ ⇒ standard causal transport problem (A.-Backhoff-Zalashko 2016)

Example.

- state dynamics: $dX_t = \alpha_t dt + dW_t$
- cost: $\mathbb{E}^{\mathbb{P}}\left[\frac{1}{2}\int_0^T \left(X_t^2 + \alpha_t^2\right) dt\right] + \int_0^T K_t(\mathbb{P} \circ X_t^{-1}) dt$

⇒ in the variational problem we have causal optimal transport w.r.t. Cameron-Martin distance:

$$\inf_{\pi \in \Pi_{bc}(\gamma,\mu)} \mathbb{E}^{\pi}[|\overline{\omega} - \omega|_H^2] = \mathcal{H}(\mu|\gamma),$$

hence we are left with

$$\inf_{\mu \ll \gamma} \{\mathcal{H}(\mu|\gamma) + P(\mu)\}, \quad P(\mu) \text{ penalty term}$$

Conclusions

In the case with **hidden/no dynamics**:

- characterization of equilibrium via non-anticipative transport
- existence and uniqueness results
- characterization of causal optimal transport ($\neq$ KR)...

In the case with **state dynamics**:

- characterization of weak McKean-Vlasov solutions via non-anticipative transport
- existence and uniqueness...
- characterization of causal optimal transport...

Bibliography et al.

- Acciaio, B, and Backhoff, J, and Zalashko, A. "Causal optimal transport and its links to enlargement of filtrations and continuous-time stochastic optimization", arXiv:1611.02610, 2016.
- Blanchet, A, and Carlier, G. "Optimal transport and Cournot-Nash equilibria", Mathematics of Operations Research 41, 125-145, 2015.
- Carmona, R, and Delarue, F. "Probabilistic Theory of Mean Field Games with Applications I-II", Springer, 2017.

Thank you for your attention

and

Buon compleanno Ioannis! :)