McKean-Vlasov control problems and non-anticipative optimal transport

Beatrice Acciaio
London School of Economics

ongoing work with J. Backhoff and R. Carmona

"Advances in Stochastic Analysis for Risk Modeling"
CIRM, Luminy, 13-17 November 2017

Outline

1. Motivation and problem formulation (weak McKean-Vlasov)
2. Our toolkit: causal transport
3. Characterization of MKV solutions via causal transport
4. Conclusions and ongoing research

Motivation

N-player stochastic differential game

→ N players with private state processes evolving as to

$$dX_t^{N,i} = b_t(X_t^{N,i}, \alpha_t^{N,i}, \bar{\nu}_t^{N,-i})dt + dW_t^i, \quad i = 1, ..., N$$

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

N-player stochastic differential game

→ N players with private state processes evolving as to

$$dX_t^{N,i} = b_t(X_t^{N,i}, \alpha_t^{N,i}, \bar{\nu}_t^{N,-i})dt + dW_t^i, \quad i = 1, ..., N$$

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

→ The objective of player i is to choose a control $\alpha^{N,i}$ that minimizes

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

- $\bar{\eta}_t^{N,-i} = \frac{1}{N-1} \sum_{j \neq i} \delta_{(X_t^{N,j}, \alpha_t^{N,j})}$ empirical distrib. of states & controls

→ Statistically identical players: same functions b_t, f_t, g

N-player stochastic differential game

Problems:

- optimal controls rarely exist (whether cooperative or non-cooperative)
- even when they exist, difficult to characterize

N-player stochastic differential game

Problems:

- optimal controls rarely exist (whether cooperative or non-cooperative)
- even when they exist, difficult to characterize

Idea: resort to approximation by **asymptotic arguments:**

N-player game – – – – – – – –– > $N \to \infty$

N-player stochastic differential game

Problems:

- optimal controls rarely exist (whether cooperative or non-cooperative)
- even when they exist, difficult to characterize

Idea: resort to approximation by **asymptotic arguments:**

N-player game	$- - - - - - - -- >$	$N \to \infty$
Nash equilibrium (non-coop)	$- - -- >$	Mean Field Game
Social planner (cooperative)	$- - -- >$	McKean Vlasov

Approximating cooperative equilibria

Main argument:

- all agents adopt the same feedback control: $\alpha_t^{N,i} = \phi(t, X_t^{N,i})$
- in the limit (# players $\to \infty$) the private states of players evolve independently of each other

Approximating cooperative equilibria

Main argument:

- all agents adopt the same feedback control: $\alpha_t^{N,i} = \phi(t, X_t^{N,i})$
- in the limit (# players $\to \infty$) the private states of players evolve independently of each other
- distribution of private state converges toward distribution of the solution to the **McKean-Vlasov control problem:**

$$\inf_\alpha \mathbb{E}\left[\int_0^T f_t\left(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)\right) dt + g\left(X_T, \mathcal{L}(X_T)\right)\right]$$

$$\text{subject to} \quad dX_t = b_t\left(X_t, \alpha_t, \mathcal{L}(X_t)\right) dt + dW_t$$

- under suitable conditions, the optimal feedback controls are ϵ-optimal for large systems of players

Problem Formulation

McKean-Vlasov control problem

We study the following **McKean-Vlasov control problem:**

$$\inf_{\alpha} \mathbb{E}\left[\int_0^T f_t\left(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)\right) dt + g\left(X_T, \mathcal{L}(X_T)\right)\right]$$

subject to $\quad dX_t = b_t\left(X_t, \alpha_t, \mathcal{L}(X_t)\right) dt + dW_t, \;\; X_0 = 0$

McKean-Vlasov control problem

We study the following **McKean-Vlasov control problem:**

$$\inf_{\alpha} \mathbb{E}\left[\int_0^T f_t\left(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)\right) dt + g\left(X_T, \mathcal{L}(X_T)\right)\right]$$

$$\text{subject to} \quad dX_t = b_t\left(X_t, \alpha_t, \mathcal{L}(X_t)\right) dt + dW_t, \;\; X_0 = 0$$

Remark:

An important subclass of MFGs (the so-called **potential games**) can be formulated as MKV control problems.

McKean-Vlasov control problem

Vast literature: Caines, Cardaliaguet, Carmona, Delarue, Huang, Lachapelle, Lacker, Lasry, Lions, Malhamé, Pham, Sznitman, Wei..

Classical approaches:

- analytic (Lasry-Lions): HJB, forward-backward system of PDEs
- probabilistic: Pontryagin maximum principle, adjoint FBSDEs

McKean-Vlasov control problem

Vast literature: Caines, Cardaliaguet, Carmona, Delarue, Huang, Lachapelle, Lacker, Lasry, Lions, Malhamé, Pham, Sznitman, Wei..

Classical approaches:

- analytic (Lasry-Lions): HJB, forward-backward system of PDEs
- probabilistic: Pontryagin maximum principle, adjoint FBSDEs

Our approach: use some “dynamic” optimal transportation

With the aim of giving:

$\hookrightarrow$ different existence results

$\hookrightarrow$ explicit characterization beyond linear-quadratic case

Causal Transport

Classical Monge-Kantorovich optimal transport

Given two Polish probability spaces $(\mathcal{X}, \mu), (\mathcal{Y}, \nu)$, move the mass from μ to ν minimizing the cost of transportation $c : \mathcal{X} \times \mathcal{Y} \to [0, \infty]$:

$$\mathrm{OT}(\mu, \nu, c) := \inf \left\{ \mathbb{E}^{\pi}[c(x, y)] : \pi \in \Pi(\mu, \nu) \right\},$$

$\Pi(\mu, \nu)$: probability measures on $\mathcal{X} \times \mathcal{Y}$ with marginals μ and ν.

Monge transport: all mass sitting on x is transported into $y = T(x)$.

Kantorovich transport: mass can split.

Classical Monge-Kantorovich optimal transport

Given two Polish probability spaces $(\mathcal{X},\mu),(\mathcal{Y},\nu)$, move the mass from μ to ν minimizing the cost of transportation $c:\mathcal{X}\times\mathcal{Y}\to[0,\infty]$:

$$\mathrm{OT}(\mu,\nu,c) := \inf\left\{\mathbb{E}^{\pi}[c(x,y)] : \pi\in\Pi(\mu,\nu)\right\},$$

$\Pi(\mu,\nu)$: probability measures on $\mathcal{X}\times\mathcal{Y}$ with marginals μ and ν.

Monge transport: all mass sitting on x is transported into $y{=}T(x)$.

Kantorovich transport: mass can split.

In a dynamic setting (we have the “time component”): move the mass in a non-anticipative way: what is transported into the 2^{nd} coordinate at time t, depends on the 1^{st} coordinate only up to t

Causal optimal transport

Definition (Causal (non-anticipative) transport plans)

$\pi \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ is a causal transport plan if, $\forall t$ and $D \in \mathcal{F}_t^{\mathcal{Y}}$, the map $\mathcal{X} \ni x \mapsto \pi^x(D)$ is $\mathcal{F}_t^{\mathcal{X}}$-mbl. ($\mathcal{F}^{\mathcal{X}}, \mathcal{F}^{\mathcal{Y}}$ canonical filtrations on $\mathcal{X}$, $\mathcal{Y}$)

The concept goes back to Yamada-Watanabe (1971); see also Jacod (1980), Kurtz (2014), Lassalle (2015), Backhoff et al. (2016).

Causal optimal transport

Definition (Causal (non-anticipative) transport plans)

$\pi \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ is a causal transport plan if, $\forall t$ and $D \in \mathcal{F}_t^{\mathcal{Y}}$, the map $\mathcal{X} \ni x \mapsto \pi^x(D)$ is $\mathcal{F}_t^{\mathcal{X}}$-mbl. ($\mathcal{F}^{\mathcal{X}}, \mathcal{F}^{\mathcal{Y}}$ canonical filtrations on $\mathcal{X}$, $\mathcal{Y}$)

The concept goes back to Yamada-Watanabe (1971); see also Jacod (1980), Kurtz (2014), Lassalle (2015), Backhoff et al. (2016).

Notation:

$\Pi_c(\mu, \nu)$ = set of causal transports with marginals μ and ν,
$\Pi_c(\mu, .) = \bigcup_{\nu \in \mathcal{P}(\mathcal{Y})} \Pi_c(\mu, \nu)$

$$\mathrm{COT}(\mu, \nu, c) := \inf \left\{ \mathbb{E}^\pi[c(X, Y)] : \pi \in \Pi_c(\mu, \nu) \right\}$$

Example: weak-solutions of SDEs

Here $\mathcal{X} = \mathcal{Y} = C_0 := C_0[0,\infty)$ continuous paths starting at zero

Example (Yamada-Watanabe'71)

Assume weak-existence of the solution to the SDE:

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t, \quad b, \sigma \text{ Borel measurable.}$$

Then $\mathcal{L}(B, Y)$ causal transport between $(C_0, \mathcal{L}(B))$ and $(C_0, \mathcal{L}(Y))$.

Example: weak-solutions of SDEs

Here $\mathcal{X} = \mathcal{Y} = C_0 := C_0[0, \infty)$ continuous paths starting at zero

Example (Yamada-Watanabe'71)

Assume weak-existence of the solution to the SDE:

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t, \quad b, \sigma \text{ Borel measurable.}$$

Then $\mathcal{L}(B, Y)$ causal transport between $(C_0, \mathcal{L}(B))$ and $(C_0, \mathcal{L}(Y))$.

- **Transport perspective:** from an observed trajectory of B, the mass can be split at each moment of time into Y only based on the information available up to that time.

Example: weak-solutions of SDEs

Here $\mathcal{X} = \mathcal{Y} = C_0 := C_0[0, \infty)$ continuous paths starting at zero

Example (Yamada-Watanabe'71)

Assume weak-existence of the solution to the SDE:

$$dY_t = b(Y_t)dt + \sigma(Y_t)dB_t, \quad b, \sigma \text{ Borel measurable.}$$

Then $\mathcal{L}(B, Y)$ causal transport between $(C_0, \mathcal{L}(B))$ and $(C_0, \mathcal{L}(Y))$.

- **Transport perspective:** from an observed trajectory of B, the mass can be split at each moment of time into Y only based on the information available up to that time.
- **No splitting of mass:**

  Monge transport $\iff$ **strong solution** $Y = F(B)$.

MKV via Causal Transport

McKean-Vlasov control problem and Causal Transport

→ Recall our McKean-Vlasov control problem:

$$\inf_{\alpha} \mathbb{E}\left[\int_0^T f_t\left(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)\right) dt + g\left(X_T, \mathcal{L}(X_T)\right)\right]$$

subject to

$$dX_t = b_t\left(X_t, \alpha_t, \mathcal{L}(X_t)\right) dt + dW_t, \quad X_0 = 0$$

→ The joint distribution $\mathcal{L}(W, X)$ is a causal transport plan between $(C_0[0,T], \mathcal{L}(W))$ and $(C_0[0,T], \mathcal{L}(X))$

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, BM W, α is $\mathcal{F}$-progress. meas.

(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,

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

Assumptions

$\rightarrow$ We need some **convexity assumptions**.

Assumptions

→ We need some **convexity 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: for all x, a, η,

- f_t is bounded from below uniformly in t
- $f_t(x, ., \eta)$ is convex
- $f_t(x, a, .)$ is $\prec_{\text{conv}}$-monotone

Example: Inter-bank systemic risk model

[Carmona-Fouque-Sun 2013]

- Inter-bank borrowing/lending, where the log-monetary reserve of each bank, asymptotically, is governed by the MKV eq.

$$dX_t = [k(\mathbb{E}[X_t] - X_t) + \alpha_t]dt + dW_t, \ X_0 = 0$$

  $k \geq 0$ rate of m-r in the interaction from b&I between banks

- All banks can control their rate of borrowing/lending to a central bank with the same policy α, to minimize the cost

$$\mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}\alpha_t^2 - q\alpha_t(\mathbb{E}[X_t]-X_t) + \frac{c}{2}(\mathbb{E}[X_t]-X_t)^2\Big)dt + \frac{d}{2}(\mathbb{E}[X_T]-X_T)^2\Big]$$

  $q > 0$ incentive to borrowing ($\alpha_t > 0$) or lending ($\alpha_t < 0$),
  $c, d > 0$ penalize departure from average

Characterization via non-anticipative optimal transport

- we consider transport problems in the path space $C_0[0,T]$
- γ: Wiener measure on $C_0[0,T]$
- $(\omega,\overline{\omega})$: generic element on $C_0[0,T]\times C_0[0,T]$
- here for simplicity control = drift

Characterization via non-anticipative optimal transport

- we consider transport problems in the path space $C_0[0,T]$
- γ: Wiener measure on $C_0[0,T]$
- $(\omega,\overline{\omega})$: generic element on $C_0[0,T]\times C_0[0,T]$
- here for simplicity control = drift

Theorem

The weak MKV problem is **equivalent** *to the variational problem*

$$\inf_{\nu\in\tilde{\mathcal{P}}}\ \inf_{\pi\in\Pi_c(\gamma,\nu)} \mathbb{E}^{\pi}\Big[\int_0^T f_t\Big(\overline{\omega}_t, (\dot{\widehat{\overline{\omega}-\omega}})_t, p_t\big((\overline{\omega}, \dot{\widehat{\overline{\omega}-\omega}})_{\#}\pi\big)\Big)dt + g(\overline{\omega}_T,\nu_T)\Big]$$

where $p_t(\eta)=\eta_t$, $(\dot{\widehat{\overline{\omega}-\omega}})_t=\beta_t$ when $\overline{\omega}-\omega=\int_0^{\cdot}\beta_t dt$, and

$$\tilde{\mathcal{P}}=\{\nu\in\mathcal{P}(C) : \nu\text{-a.s. pathwise quadr.var. } \exists \text{ and } \langle\omega\rangle_t=t\ \forall\ t\}$$

Characterization via non-anticipative optimal transport

'**Equivalence**' means that the above variational problem and

$$\inf \ \mathbb{E}^{\mathbb{P}}\left[\int_0^T f_t\left(X_t, \alpha_t, \mathbb{P}\circ(X_t,\alpha_t)^{-1}\right)dt + g\left(X_T, \mathbb{P}\circ X_T^{-1}\right)\right]$$

have the same value, where the infimum is taken over tuples $(\Omega, (\mathcal{F}_t), \mathbb{P}, W, X, \alpha)$ s.t. $dX_t = b_t\left(X_t, \alpha_t, \mathbb{P}\circ X_t^{-1}\right)dt + dW_t$, and that moreover the optimizers are related via:

- $\nu^* = \mathcal{L}(X^*)$
- $\pi^* \longleftrightarrow \alpha^*$, with $\pi^* = \mathcal{L}(W^*, X^*)$

Characterization via non-anticipative optimal transport

→ Weak solutions of MKV control problem given by infimum over tuples $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P}, W, X, \alpha)$.

Corollary

1. *The infimum can be taken over tuples s.t. α is $\mathcal{F}^X$-measurable (weak closed loop).*
2. *If the infimum is attained, then the optimal α is of weak closed loop form.*

Characterization via non-anticipative optimal transport

→ Weak solutions of MKV control problem given by infimum over tuples $(\Omega, (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P}, W, X, \alpha)$.

Corollary

1. *The infimum can be taken over tuples s.t. α is $\mathcal{F}^X$-measurable (weak closed loop).*
2. *If the infimum is attained, then the optimal α is of weak closed loop form.*

Remark. The outer minimization in VP can be done over $\{\nu \ll \gamma\}$ instead of $\tilde{\mathcal{P}}$, whenever the drift is guaranteed to be square integr. (e.g. drift = control, and $f_t(x, a, \eta) \geq K|a|^2 \;\; \forall \; x, \eta$ and for a large).

Example: separable cost $f_t(x,a) + \tilde{f}_t(\nu_t, x)$

Example: take $k = q = 0$ in the example above, then

- state dynamics: $dX_t = \alpha_t dt + dW_t$
- cost: $\mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}\alpha_t^2 + \frac{c}{2}(\mathbb{E}[X_t] - X_t)^2\Big)dt + \frac{d}{2}(\mathbb{E}[X_T] - X_T)^2\Big]$

$\Rightarrow$ COT w.r.t. Cameron-Martin distance (Lassalle 2015):

$$\inf_{\pi \in \Pi_c(\gamma,\nu)} \mathbb{E}^\pi[|\overline{\omega} - \omega|_H^2] = \mathcal{H}(\nu|\gamma) \Rightarrow \inf_{\nu \ll \gamma} \Big\{\mathcal{H}(\nu|\gamma) + \frac{c}{2}\int_0^T \mathrm{Var}(\nu_t)dt + \frac{d}{2}\mathrm{Var}(\nu_T)\Big\}$$

Example: separable cost $f_t(x,a) + \tilde{f}_t(\nu_t, x)$

Example: take $k = q = 0$ in the example above, then

- state dynamics: $dX_t = \alpha_t dt + dW_t$
- cost: $\mathbb{E}\Big[\int_0^T \Big(\frac{1}{2}\alpha_t^2 + \frac{c}{2}(\mathbb{E}[X_t] - X_t)^2\Big)dt + \frac{d}{2}(\mathbb{E}[X_T] - X_T)^2\Big]$

$\Rightarrow$ COT w.r.t. Cameron-Martin distance (Lassalle 2015):

$$\inf_{\pi \in \Pi_c(\gamma,\nu)} \mathbb{E}^\pi[|\overline{\omega} - \omega|_H^2] = \mathcal{H}(\nu|\gamma) \Rightarrow \inf_{\nu \ll \gamma} \Big\{\mathcal{H}(\nu|\gamma) + \frac{c}{2}\int_0^T \mathrm{Var}(\nu_t)dt + \frac{d}{2}\mathrm{Var}(\nu_T)\Big\}$$

More generally: for running cost $\frac{1}{2}\alpha_t^2 + h_t(X_t, \mathbb{P} \circ X_t^{-1})$,

by Sanov's theorem, we can approximate

$$\inf_{\nu \ll \gamma} \{\mathcal{H}(\nu|\gamma) + F(\nu)\} = \lim_{n\to\infty} -\frac{1}{n} \ln \mathbb{E} e^{nF(\frac{1}{n}\sum_{i=1}^n \delta_{W^i})}, \ \{W_i\} \text{ ind. BMs.}$$

This does not seem to be limited to the entropy case ($\frac{1}{2}\alpha_t^2$).

Concluding remarks

We have provided:

- a connection of McKean-Vlasov control problems to causal transport problems
- a characterization of weak McKean-Vlasov solutions via CT

Concluding remarks

We have provided:

- a connection of McKean-Vlasov control problems to causal transport problems
- a characterization of weak McKean-Vlasov solutions via CT

Work in progress:

- The optimization over $\Pi_c(\gamma, \nu)$ is not a standard causal transport problem $\Rightarrow$ new analysis for existence/duality
- Use our characterization theorem in order to find
  - existence and uniqueness of weak MKV solutions
  - explicit formulation of solutions to MKV control problems
- Time-discretization and numerical scheme

Concluding remarks

We have provided:

- a connection of McKean-Vlasov control problems to causal transport problems
- a characterization of weak McKean-Vlasov solutions via CT

Work in progress:

- The optimization over $\Pi_c(\gamma, \nu)$ is not a standard causal transport problem $\Rightarrow$ new analysis for existence/duality
- Use our characterization theorem in order to find
  - existence and uniqueness of weak MKV solutions
  - explicit formulation of solutions to MKV control problems
- Time-discretization and numerical scheme

Discrete-time setting:

- By the analogy type $\longleftrightarrow$ noise, we can study Cournot-Nash equilibria for heterogeneous agents via causal transport

Thank you for your attention!

From N-player game to McKean-Vlasov control problem

- rarely expect existence of global minimizers
- resort to approximation by **asymptotic arguments:**

From N-player game to McKean-Vlasov control problem

- rarely expect existence of global minimizers
- resort to approximation by **asymptotic arguments:**

SDE State Dynamics for N players	$\xrightarrow{\text{optimization}}$	Nash equilibrium for N players
$\lim_{N\to\infty} \downarrow$		$\downarrow \lim_{N\to\infty}$
		Mean-Field Game
McKean-Vlasov dynamics	$\xrightarrow{\text{optimization}}$	$\boxed{\text{controlled McK-V dyn}}$

From N-player game to McKean-Vlasov control problem

- rarely expect existence of global minimizers
- resort to approximation by **asymptotic arguments:**

SDE State Dynamics for N players	$\xrightarrow{\text{optimization}}$	Nash equilibrium for N players
$\lim_{N\to\infty} \downarrow$		$\downarrow \lim_{N\to\infty}$
		Mean-Field Game
McKean-Vlasov dynamics	$\xrightarrow{\text{optimization}}$	controlled McK-V dyn

Vast literature: Caines, Carmona, Delarue, Huang, Lachapelle, Lacker, Lasry, Lions, Malhamé, Pham, Sznitman, Wei,...

The red path: approximating cooperative equilibria

Main idea:

- all agents adopt the same feedback control: $\alpha_t^{N,i} = \phi(t, X_t^{N,i})$
- in the limit (# players $\to \infty$) the private states of players evolve independently of each other

The red path: approximating cooperative equilibria

Main idea:

- all agents adopt the same feedback control: $\alpha_t^{N,i} = \phi(t, X_t^{N,i})$
- in the limit (# players $\to \infty$) the private states of players evolve independently of each other
- distribution of private state converges toward distribution of the solution to the **McKean-Vlasov control problem:**

$$\inf_\alpha \mathbb{E}\left[\int_0^T f_t\left(X_t, \alpha_t, \mathcal{L}(X_t, \alpha_t)\right) dt + g\left(X_T, \mathcal{L}(X_T)\right)\right]$$

$$\text{subject to} \quad dX_t = b_t\left(X_t, \alpha_t, \mathcal{L}(X_t)\right) dt + dW_t$$

- under suitable conditions, the optimal feedback controls are ϵ-optimal for large systems of players

The blue path: approximating competitive equilibria

Main idea:

- Seek for Nash equilibria for the N-player game

The blue path: approximating competitive equilibria

Main idea:

- Seek for Nash equilibria for the N-player game
- Model behaviour of a representative agent, and solve the **Mean-Field Game problem**:
  1) for every fixed joint law η, with first marginal ν, solve

$$\inf_\alpha \mathbb{E}\Big[\int_0^T f_t\,(X_t, \alpha_t, \eta_t)\,dt + g\,(X_T, \nu_T)\Big]$$
$$\text{s.t. } dX_t = b_t\,(X_t, \alpha_t, \nu_t)\,dt + dW_t$$

  2) fixed point problem: η s.t. for the solution $\mathcal{L}(X, \alpha) = \eta$
- under suitable conditions, the optimal feedback provides an approximate Nash equilibrium for large system of players

The blue path: approximating competitive equilibria

Main idea:

- Seek for Nash equilibria for the N-player game
- Model behaviour of a representative agent, and solve the **Mean-Field Game problem**:
  1) for every fixed joint law η, with first marginal ν, solve

$$\inf_\alpha \mathbb{E}\Big[\int_0^T f_t\,(X_t, \alpha_t, \eta_t)\, dt + g\,(X_T, \nu_T) \Big]$$
$$\text{s.t. } dX_t = b_t\,(X_t, \alpha_t, \nu_t)\, dt + dW_t$$

  2) fixed point problem: η s.t. for the solution $\mathcal{L}(X, \alpha) = \eta$
- under suitable conditions, the optimal feedback provides an approximate Nash equilibrium for large system of players
- for potential games, MFG can be formulated as MKV

General case

Assumptions. For all $x, a \in \mathbb{R}, m \in \mathcal{P}(\mathbb{R}), \eta \in \mathcal{P}(\mathbb{R} \times \mathbb{R})$:

(A1) $b_t(x, ., m)$ injective and convex

(A2) f_t bdd below unif. in t, and $f_t(x, b_t^{-1}(x, ., m)(y), \eta)$ convex in y

(A3) $f_t(x, a, .)$ is $\prec_{cm}$-monotone (resp. $\prec_{conv}$-monotone if b is linear)

($\prec_{cm}$ (resp. $\prec_{conv}$) denotes the conv/monotone (resp. conv) order)

Pathwise quadratic variation. For $\omega \in C$, $n \in \mathbb{N}$, define

$$\sigma_0^n(\omega) := 0, \quad \sigma_{k+1}^n(\omega) := \inf\{t > \sigma_k^n(\omega) : |\omega(t) - \omega(\sigma_k^n)| \geq 2^{-n}\}, \quad k \in \mathbb{N}$$

We say that ω has quadratic variation if

$$V_n(\omega)(t) := \sum_{k=0}^{\infty} (\omega(\sigma_{k+1}^n \wedge t) - \omega(\sigma_k^n \wedge t))^2 \to_u \; =: \langle \omega \rangle_t \in C$$

Notation. $\tilde{\mathcal{P}} = \{\nu \in \mathcal{P}(C) : \langle \omega \rangle \exists \; \nu\text{-a.s., with } \langle \omega \rangle_t = t \text{ for all } t\}$

General case

Under the above assumptions, the following characterization of weak McKean-Vlasov solutions via causal transport holds.

Theorem

The weak MKV problem is **equivalent** *to the following problem*

$$\inf_{\nu \in \tilde{\mathcal{P}}} \inf_{\pi \in \Pi_c(\gamma,\nu)} \mathbb{E}^\pi \left[\int_0^T f_t\left(\overline{\omega}_t, u_t^\nu(\omega,\overline{\omega}), p_t\big((\overline{\omega}, u^\nu)_{\#}\pi\big)\right) dt + g(\overline{\omega}_T, \nu_T) \right]$$

where $u_t^\nu(\omega,\overline{\omega}) = b_t^{-1}(\overline{\omega}_t, ., \nu_t)\big((\dot{\widehat{\overline{\omega} - \omega}})_t\big)$ *and* $p_t(\eta) = \eta_t$.