Gibbs principle on path space and relations with stochastic control

Louis-Pierre Chaintron

DMA ENS-PSL, CERMICS ENPC, Inria MΞDISIM

Joint work with G. Conforti (*CMAP, École Polytechnique*)
and J. Reygner (*CERMICS, École des Ponts et Chaussées*)

A concrete example

From "Modern Times", Charlie Chaplin, 1936, 13^{th} *minute.*

$\Rightarrow$ *The Tramp* works.

Let us model *The Tramp*'s work

Three possible states:

The Tramp is on time.

The Tramp takes a break.

The Tramp is late.

Random disturbance

Deterministic evolution:

$$\mathrm{Tramp}_{t+1} = F(t, \mathrm{Tramp}_t).$$

$\hookrightarrow$ It is sufficient to know $\mathrm{Tramp}_{t=0}$.

Random disturbance

Deterministic evolution:

$$\mathrm{Tramp}_{t+1} = F(t, \mathrm{Tramp}_t).$$

$\hookrightarrow$ It is sufficient to know $\mathrm{Tramp}_{t=0}$.

But...

$\Rightarrow$ F does not account for interaction.

Probabilistic model

Let us a consider a prior distribution

$$\nu \in \mathcal{P}(\{\text{on time}, \text{late}, \text{break}\}),$$

for instance

$$\nu(\text{on time}) = \nu(\text{late}) = \nu(\text{break}) = \tfrac{1}{3}.$$

Probabilistic model

Let us a consider a prior distribution

$$\nu \in \mathcal{P}(\{\text{on time}, \text{late}, \text{break}\}),$$

for instance

$$\nu(\text{on time}) = \nu(\text{late}) = \nu(\text{break}) = \frac{1}{3}.$$

This is not accurate...

...because *The Tramp* is never on time!

Conditioning

$$\nu_{\text{on time}} := \nu(\cdot|\text{Tramp}\neq\text{on time}) = Z^{-1}\mathbb{1}_{\{\text{Tramp}\neq\text{on time}\}}\,\nu$$

↪ $\nu_{\text{on time}}$ matches observation: $\nu_{\text{on time}}(\text{on time}) = 0$.

↪ Is this choice of $\nu_{\text{on time}}$ optimal?...

Conditioning

$$\nu_{\text{on time}} := \nu(\cdot|\text{Tramp}\neq\text{on time}) = Z^{-1}\mathbb{1}_{\{\text{Tramp}\neq\text{on time}\}}\,\nu$$

$\hookrightarrow$ $\nu_{\text{on time}}$ matches observation: $\nu_{\text{on time}}(\text{on time}) = 0$.

$\hookrightarrow$ Is this choice of $\nu_{\text{on time}}$ optimal?...

...Yes, it is!

$$\nu_{\text{on time}} = \underset{\substack{\mu\in\mathcal{P} \\ \mu(\text{on time})=0}}{\text{argmin}} \; H(\mu|\nu),$$

where $H(\mu|\nu) \geq 0$ and $H(\mu|\nu) \Leftrightarrow \mu = \nu$.

$\Rightarrow$ Natural distance w.r.t. conditioning.

Relative entropy

Given $\nu \in \mathcal{P}(\mathbb{R}^d)$, the relative entropy is

$$H(\mu|\nu) := \begin{cases} \int_E \log \frac{\mathrm{d}\mu}{\mathrm{d}\nu} \mathrm{d}\mu & \text{if } \mu \ll \nu, \\ +\infty & \text{otherwise.} \end{cases}$$

Given $A \subset \mathbb{R}^d$ with $\nu(A) > 0$,

$$\underset{\substack{\mu \in \mathcal{P}(E) \\ \mu(A)=0}}{\operatorname{argmin}} H(\mu|\nu) = Z^{-1} \mathbb{1}_{A^c} \nu =: \nu_{A^c}$$

Relative entropy

Given $\nu \in \mathcal{P}(\mathbb{R}^d)$, the relative entropy is

$$H(\mu|\nu) := \begin{cases} \int_E \log \frac{\mathrm{d}\mu}{\mathrm{d}\nu} \mathrm{d}\mu & \text{if } \mu \ll \nu, \\ +\infty & \text{otherwise.} \end{cases}$$

Given $A \subset \mathbb{R}^d$ with $\nu(A) > 0$,

$$\boxed{\operatorname*{argmin}_{\substack{\mu \in \mathcal{P}(E) \\ \mu(A)=0}} H(\mu|\nu) = Z^{-1} \mathbb{1}_{A^c} \nu =: \nu_{A^c}}$$

Proof.

$$H(\mu|\nu) + \mu(A) = H(\nu_{A^c}|\nu) + \nu(A) + \underbrace{H(\mu|\nu_{A^c})}_{\geq 0}.$$

□

Adding statistical information

Defining

$$\psi(\text{Tramp}) := \mathbb{1}_{\{\text{Tramp}=\text{on time}\}} - \frac{1}{42},$$

we now measure that

$$\mathbb{E}\,\psi(\text{Tramp}) \leq 0.$$

$\hookrightarrow$ How can we correct μ to account for this?

Adding statistical information

Defining

$$\psi(\text{Tramp}) := \mathbb{1}_{\{\text{Tramp}=\text{on time}\}} - \frac{1}{42},$$

we now measure that

$$\mathbb{E}\,\psi(\text{Tramp}) \leq 0.$$

$\hookrightarrow$ How can we correct μ to account for this?

A natural candidate is

$$\underset{\substack{\mu\in\mathcal{P} \\ \langle\mu,\psi\rangle\leq 0}}{\text{argmin}}\, H(\mu|\nu),$$

where $\langle\mu,\psi\rangle := \int \psi \,\mathrm{d}\mu$.

Gibbs variational principle

Theorem (Gibbs variational principle)

Given ν in $\mathcal{P}(\mathbb{R}^d)$ and $\psi : \mathbb{R}^d \to \mathbb{R}$ continuous, bounded from below,

$$\inf_{\substack{\mu \in \mathcal{P}(\mathbb{R}^d) \\ \langle \mu, \psi \rangle \leq 0}} H(\mu|\nu)$$

is realised by a unique measure $\mu_{\overline{\beta}}$ for some $\overline{\beta} \in \mathbb{R}$, where

$$\frac{\mathrm{d}\mu_{\overline{\beta}}}{\mathrm{d}\nu}(x) = Z_{\overline{\beta}}^{-1} e^{-\overline{\beta}\psi(x)}.$$

$\hookrightarrow$ Single linear constraint [SZ91; DZ96].

$\hookrightarrow$ Lagrange multiplier $\overline{\beta}$.

Proof.

The Gibbs free energy

$$G(\beta) := \inf_{\mu \in \mathcal{P}(\mathbb{R}^d)} H(\mu|\nu) + \beta\langle \mu, \psi \rangle,$$

is uniquely realised by $\mu = \mu_\beta$, because

$$H(\mu|\nu) + \beta\langle \mu, \psi \rangle = H(\mu_\beta|\nu) + \beta\langle \mu_\beta, \psi \rangle + \overbrace{H(\mu|\mu_\beta)}^{\geq 0},$$

hence $G(\beta) = -\log Z_\beta$.

Proof.

The Gibbs free energy

$$G(\beta) := \inf_{\mu \in \mathcal{P}(\mathbb{R}^d)} H(\mu|\nu) + \beta\langle \mu, \psi \rangle,$$

is uniquely realised by $\mu = \mu_\beta$, because

$$H(\mu|\nu) + \beta\langle \mu, \psi \rangle = H(\mu_\beta|\nu) + \beta\langle \mu_\beta, \psi \rangle + \overbrace{H(\mu|\mu_\beta)}^{\geq 0},$$

hence $G(\beta) = -\log Z_\beta$. Moreover,

$$\inf_{\substack{\mu \in \mathcal{P}(\mathbb{R}^d) \\ \langle \mu, \psi \rangle \leq 0}} H(\mu|\nu) = \inf_{\substack{\mu \in \mathcal{P}(\mathbb{R}^d) \\ \langle \mu, \psi \rangle \leq 0}} \sup_{\beta \geq 0} H(\mu|\nu) + \beta\langle \mu, \psi \rangle$$
$$\geq \sup_{\beta \in \mathbb{R}} G(\beta),$$

and $\beta \mapsto \langle \mu_\beta, \psi \rangle$ is continuous and decreasing... □

An alternative approach to entropy

Law of large numbers

For $(X^i)_{i \geq 1}$ independent ν-distributed variables,

$$\frac{1}{N} \sum_{i=1}^{N} \varphi(X^i) \xrightarrow[N \to +\infty]{\text{a.s.}} \mathbb{E}\varphi(X^1) = \langle \nu, \varphi \rangle$$

for every φ, so that

$$\pi(\vec{X}^N) := \frac{1}{N} \sum_{i=1}^{N} \delta_{X^i} \xrightarrow[N \to +\infty]{\text{weak}} \nu.$$

An alternative approach to entropy

Law of large numbers

For $(X^i)_{i \geq 1}$ independent ν-distributed variables,

$$\frac{1}{N} \sum_{i=1}^{N} \varphi(X^i) \xrightarrow[N \to +\infty]{\text{a.s.}} \mathbb{E}\varphi(X^1) = \langle \nu, \varphi \rangle$$

for every φ, so that

$$\pi(\vec{X}^N) := \frac{1}{N} \sum_{i=1}^{N} \delta_{X^i} \xrightarrow[N \to +\infty]{\text{weak}} \nu.$$

$\hookrightarrow$ What about fluctuations?

$$\mathbb{P}(\pi(\vec{X}^N) = \mu) = ?$$

Law of large numbers

For $(X^i)_{i\geq 1}$ independent ν-distributed variables,

$$\frac{1}{N}\sum_{i=1}^{N}\varphi(X^i)\xrightarrow[N\to+\infty]{\text{a.s.}}\mathbb{E}\varphi(X^1)=\langle\nu,\varphi\rangle$$

for every φ, so that

$$\pi(\vec{X}^N):=\frac{1}{N}\sum_{i=1}^{N}\delta_{X^i}\xrightarrow[N\to+\infty]{\text{weak}}\nu.$$

$\hookrightarrow$ What about fluctuations?

$$\boxed{\mathbb{P}(\pi(\vec{X}^N)\simeq\mu)\simeq e^{-NH(\mu|\nu)}}$$

$\Rightarrow$ H quantifies deviations from the LLN.

From large deviations,

$$\mathbb{P}\big(\,\pi(\vec{X}^N) \in A\,\big) \simeq \exp\Big[-N \inf_{\mu \in A} H(\mu|\nu)\Big],$$

so that

$$\mathbb{P}\big(\,\langle \pi(\vec{X}^N), \psi \rangle \leq 0\,\big) \simeq \exp\Big[-N \inf_{\mu, \langle \mu, \psi \rangle \leq 0} H(\mu|\nu)\Big].$$

From large deviations,

$$\mathbb{P}\big(\, \pi(\vec{X}^N) \in A \,\big) \simeq \exp\left[-N \inf_{\mu \in A} H(\mu|\nu) \right],$$

so that

$$\mathbb{P}\big(\, \langle \pi(\vec{X}^N), \psi \rangle \leq 0 \,\big) \simeq \exp\left[-N \inf_{\mu, \langle \mu, \psi \rangle \leq 0} H(\mu|\nu) \right].$$

Consequently,

$$\mathbb{P}\big(\, \langle \pi(\vec{X}^N) \in A \,|\, \langle \pi(\vec{X}^N), \psi \rangle \leq 0 \,\big) \simeq$$
$$\exp\left[-N \left(\inf_{\mu \in A, \langle \mu, \psi \rangle \leq 0} H(\mu|\nu) - \inf_{\mu, \langle \mu, \psi \rangle \leq 0} H(\mu|\nu) \right) \right].$$

Canonical ensemble

For the kinetic energy $\psi(v) := \frac{1}{2}mv^2$ in $\mathbb{R}^3$,

$$\left(\pi(\vec{V}^N) \in A \,|\, \langle \pi(\vec{V}^N), \psi \rangle = \overline{\psi}\right) \xrightarrow[N\to+\infty]{\text{law}} \underset{\mu \in \mathcal{P}(\mathbb{R}^3),\, \langle \mu, \psi \rangle = \overline{\psi}}{\text{argmax}} -\int_{\mathbb{R}^3} \mu(v) \log \mu(v) \mathrm{d}v,$$

Canonical ensemble

For the kinetic energy $\psi(v) := \frac{1}{2}mv^2$ in $\mathbb{R}^3$,

$$\left(\pi(\vec{V}^N) \in A \,\middle|\, \langle \pi(\vec{V}^N), \psi \rangle = \overline{\psi}\right) \xrightarrow[N\to+\infty]{\text{law}} \operatorname*{argmax}_{\mu \in \mathcal{P}(\mathbb{R}^3),\ \langle \mu, \psi \rangle = \overline{\psi}} - \int_{\mathbb{R}^3} \mu(v) \log \mu(v) \mathrm{d}v,$$

giving the Maxwell-Boltzmann distribution

$$\overline{\mu}(v) := \left(\frac{m}{2\pi k_B T} \right)^{3/2} \exp\left[- \frac{mv^2}{2k_B T} \right],$$

where $\overline{\psi} = \frac{3}{2} k_B T = \frac{3}{2\beta}$.

Infinitely many constraints

Example of i.i.d. diffusion processes in $\mathbb{R}^d$

$$\mathrm{d}X_t^{i,N} = -\nabla V(X_t^{i,N})\mathrm{d}t + \mathrm{d}B_t^{i,N}, \quad 1 \leq i \leq N,$$

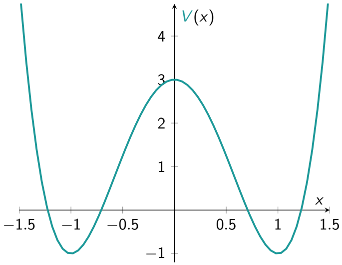

Confinement potential V

with mean-field conditioning:

$$\forall t \in [0, T], \quad \frac{1}{N}\sum_{i=1}^{N} V(X_t^{i,N}) \leq 0.$$

General setting

Let $(X^i_{[0,T]})_{i \geq 1}$ be independent $\nu_{[0,T]}$-distributed $C([0,T],\mathbb{R}^d)$-valued variables, and let

$$\pi(\vec{X}^N_t) := \frac{1}{N} \sum_{i=1}^{N} \delta_{X^i_t},$$

denote the empirical measure at time t.

We condition by

$$\forall t \in [0,T], \quad \Psi(\pi(\vec{X}^N_t)) \leq 0. \tag{1}$$

General setting

Let $(X^i_{[0,T]})_{i \geq 1}$ be independent $\nu_{[0,T]}$-distributed $C([0,T],\mathbb{R}^d)$-valued variables, and let

$$\pi(\vec{X}^N_t) := \frac{1}{N} \sum_{i=1}^{N} \delta_{X^i_t},$$

denote the empirical measure at time t.

We condition by

$$\forall t \in [0,T], \quad \Psi(\pi(\vec{X}^N_t)) \leq 0. \tag{1}$$

As $N \to +\infty$, $\mathrm{Law}(\, X^1_{[0,T]} \,|\, (1)\,)$ converges towards

$$\underset{\substack{\mu_{[0,T]} \in \mathcal{P}(C([0,T],\mathbb{R}^d)) \\ \forall t \in [0,T],\, \Psi(\mu_t) \leq 0}}{\operatorname{argmin}} H(\mu_{[0,T]} | \nu_{[0,T]}).$$

Regularity assumptions

The map $\Psi : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ is C^1 in the following sense: for any $\mu, \mu' \in \mathcal{P}(\mathbb{R}^d)$,

$$\frac{\mathrm{d}}{\mathrm{d}\varepsilon}\Big|_{\varepsilon=0} \Psi(\mu + \varepsilon(\mu' - \mu)) = \langle \mu - \mu', \tfrac{\delta \Psi}{\delta \mu}(\mu) \rangle,$$

for a continuous $\frac{\delta \Psi}{\delta \mu} : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}$.

Regularity assumptions

The map $\Psi : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$ is C^1 in the following sense: for any $\mu, \mu' \in \mathcal{P}(\mathbb{R}^d)$,

$$\frac{\mathrm{d}}{\mathrm{d}\varepsilon}\Big|_{\varepsilon=0} \Psi(\mu + \varepsilon(\mu' - \mu)) = \langle \mu - \mu', \tfrac{\delta \Psi}{\delta \mu}(\mu) \rangle,$$

for a continuous $\frac{\delta \Psi}{\delta \mu} : \mathcal{P}(\mathbb{R}^d) \times \mathbb{R}^d \to \mathbb{R}$.

Example

The constraints

$$\Psi_V(\mu) = \langle \mu, V \rangle \quad \text{and} \quad \Psi_W(\mu) = \langle \mu, W \star \mu \rangle$$

are respectively linear and non-convex in μ. For even W,

$$\frac{\delta \Psi_V}{\delta \mu}(\mu) = V \quad \text{and} \quad \frac{\delta \Psi_W}{\delta \mu}(\mu) = 2W \star \mu.$$

$\hookrightarrow$ For $W(x) = x^2$, we obtain $\Psi_W = \mathrm{Var}$.

Theorem (Gibbs path measure)

For any minimiser $\overline{\mu}_{[0,T]}$ *of*

$$\inf_{\substack{\mu \in \mathcal{P}(C([0,T],\mathbb{R}^d)) \\ \forall t \in [0,T],\, \Psi(\mu_t) \leq 0}} H(\mu_{[0,T]} | \nu_{[0,T]}).$$

some $\overline{\lambda}$ *in* $\mathcal{M}_+([0,T])$ *exists s.t.*

$$\frac{\mathrm{d}\overline{\mu}_{[0,T]}}{\mathrm{d}\nu_{[0,T]}}(x_{[0,T]}) = \left(Z_T^{\Psi}\right)^{-1} \exp\left[- \int_0^T \frac{\delta \Psi}{\delta \mu}(\overline{\mu}_t, x_t) \overline{\lambda}(\mathrm{d}t) \right],$$

with $\Psi(\overline{\mu}_t) = 0$ $\overline{\lambda}$*-a.e. Sufficient condition in the convex case.*

From conditioning to control

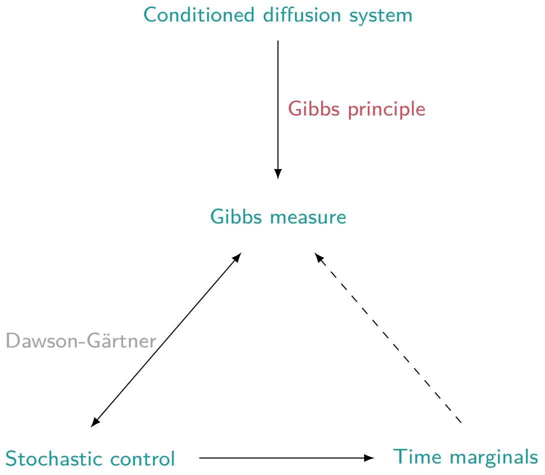

Controlled process

$$\begin{cases} \mathrm{d}X_s^{t,\mu,\alpha} = -\nabla V(X_s^{t,\mu,\alpha})\mathrm{d}s + \alpha_s \mathrm{d}s + \mathrm{d}B_s, \quad t \leq s \leq T, \\ X_t^{t,\mu,\alpha} = X_t^{t,\mu}, \quad X_t^{t,\mu} \sim \mu, \end{cases}$$

for some adapted process $\alpha = (\alpha_s)_{t \leq s \leq T}$.

Value function

$$V(t,\mu) := \inf_{\substack{(\alpha_s)_{t\leq s\leq T} \\ \forall s\in[t,T],\ \Psi(\mathrm{Law}(X_s^{t,\mu,\alpha}))\leq 0}} \mathbb{E}\int_t^T \frac{1}{2}|\alpha_s|^2 \mathrm{d}s.$$

$\hookrightarrow$ One looks for $\overline{\alpha}$ which realises $V(0,\overline{\mu}_0)$.

$\Rightarrow$ Control problem with law constraints [Dau21].

Pathwise law

Theorem

There exists an optimal $\overline{\alpha}_t = \varphi(t, X_t^{\overline{\alpha}})$ *such that*

$$\mathrm{Law}(X_{[0,T]}^{\overline{\alpha}}) = \overline{\mu}_{[0,T]},$$

and

$$\boxed{\inf_{\substack{\mu \in \mathcal{P}(C([0,T],\mathbb{R}^d)) \\ \forall t \in [0,T],\, \Psi(\mu_t) \leq 0}} H(\mu_{[0,T]} | \nu_{[0,T]}) = H(\overline{\mu}_0 | \nu_0) + V(0, \overline{\mu}_0).}$$

$\Rightarrow$ *This characterises time marginals.*

Potential mean-field game structure

$(\varphi, (\overline{\mu}_t)_t)$ is a solution of the MFG system

$$\begin{cases} \partial_t \varphi - \nabla V \cdot \nabla \varphi + \frac{1}{2}\Delta\varphi - \frac{1}{2}|\nabla\varphi|^2 = -\overline{\lambda}\frac{\delta\Psi}{\delta\mu}(\overline{\mu}_t), \\ \partial_t \overline{\mu}_t - \mathrm{div}(\overline{\mu}_t \nabla V + \overline{\mu}_t \nabla\varphi(t,\cdot) + \frac{1}{2}\nabla\overline{\mu}_t) = 0, \\ \varphi(T,\cdot) = 0. \end{cases}$$

As $N \to +\infty$, this describes a Nash equilibrium for the game

$$\inf_{\alpha^1,\dots,\alpha^N} \sum_{i=1}^N \mathbb{E} \int_0^T \frac{1}{2}|\alpha^i_t|^2 \mathrm{d}t + \int_0^T \frac{\delta\Psi}{\delta\mu}(\pi(\vec{X}^N_t), X^i_t)\overline{\lambda}(\mathrm{d}t).$$

Summary

$$\mathrm{Law}\Big(\Pi\Big(\vec{X}^N_{[0,T]}\Big)\,\Big|\,\forall t\in[0,T],\ \Psi(\pi(\vec{X}^N_t))\leq 0\Big)$$

↓ Gibbs principle

$$\operatorname*{argmin}_{\substack{\mu_{[0,T]} \\ \forall t,\ \Psi(\mu_t)\leq 0}} H(\mu_{[0,T]}|\nu_{[0,T]})$$

Dawson-Gärtner (↔) — Girsanov (↑)

$$\inf_{\substack{(\alpha_t)_t \\ \forall t,\ \Psi(\mathrm{Law}(X^\alpha_t))\leq 0}} \frac{1}{2}\int_0^T \mathbb{E}|\alpha_t|^2 \mathrm{d}t$$

→ HJB equation →

MFG system on $(\overline{\mu}_t, V(t,\cdot))$

Mean-field interaction:

$$\mathrm{d}X_t^{i,N} = -\nabla V(X_t^{i,N})\mathrm{d}t - \frac{1}{N}\sum_{j=1}^{N} \nabla W(X_t^{i,N} - X_t^{j,N})\mathrm{d}t + \mathrm{d}B_t^{i,N}$$

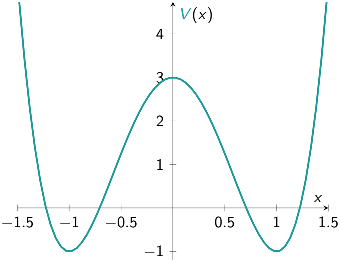

Confinement potential V

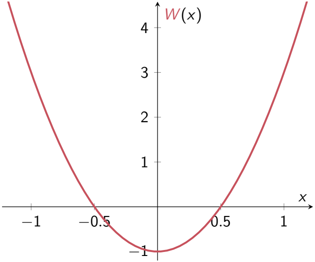

Interaction potential W

Conditioning:

$$\forall t \in [0, T], \quad \Psi(\pi(\vec{X}_t^N)) \le 0.$$

Mean-field stochastic control

Controlled process

$$\begin{cases} \mathrm{d}X_s^{t,\mu,\alpha} = -\nabla V(X_s^{t,\mu,\alpha})\mathrm{d}s - \nabla(W \star \mathrm{Law}(X_s^{t,\mu,\alpha}))(X_s^{t,\mu,\alpha})\mathrm{d}s \\ \qquad\qquad + \alpha_s \mathrm{d}s + \mathrm{d}B_s, \quad t \leq s \leq T, \\ X_t^{t,\mu,\alpha} = X_t^{t,\mu}, \quad X_t^{t,\mu} \sim \mu, \end{cases}$$

for some adapted process $\alpha = (\alpha_s)_{t \leq s \leq T}$.

Value function

$$\mathcal{V}(t,\mu) := \inf_{\substack{(\alpha_s)_{t\leq s\leq T} \\ \forall s \in [t,T],\ \Psi(\mathrm{Law}(X_s^{t,\mu,\alpha})) \leq 0}} \mathbb{E} \int_t^T \frac{1}{2} |\alpha_s|^2 \mathrm{d}s.$$

$\Rightarrow$ McKean-Vlasov control problem with law constraints.

Theorem

There exists an optimal control such that $\mathrm{Law}(X^{\overline{\alpha}}_{[0,T]}) = \overline{\mu}_{[0,T]}$ *and*

$$\inf_{\substack{\mu \in \mathcal{P}(C([0,T],\mathbb{R}^d)) \\ \forall t \in [0,T],\, \Psi(\mu_t) \leq 0}} H(\mu_{[0,T]} | \Gamma(\mu_{[0,T]})) = H(\overline{\mu}_0 | \nu_0) + \mathcal{V}(0, \overline{\mu}_0).$$

Time-marginals

$$\begin{cases} \partial_t \varphi - \nabla V \cdot \nabla \varphi - 2\nabla(W \star \overline{\mu}_t) \cdot \nabla \varphi + \frac{1}{2}\Delta\varphi \\ \qquad\qquad - \frac{1}{2}|\nabla\varphi|^2 = -\overline{\lambda}\frac{\delta\Psi}{\delta\mu}(\overline{\mu}_t), \\ \partial_t \overline{\mu}_t - \mathrm{div}(\overline{\mu}_t \nabla V + \nabla W \star \overline{\mu}_t + \overline{\mu}_t \nabla\varphi(t,\cdot) + \frac{1}{2}\nabla\overline{\mu}_t) = 0, \\ \varphi(T,\cdot) = 0. \end{cases}$$

Thank you!

References

[Dau21] S. Daudin. "Optimal control of the Fokker-Planck equation under state constraints in the Wasserstein space". In: *arXiv preprint arXiv:2109.14978* (2021).

[DG87] D. A. Dawson and J. Gärtner. "Large deviations from the McKean-Vlasov limit for weakly interacting diffusions". In: *Stochastics: An International Journal of Probability and Stochastic Processes* 20.4 (1987), pp. 247–308.

[DZ96] A. Dembo and O. Zeitouni. "Refinements of the Gibbs conditioning principle". In: *Probability theory and related fields* 104.1 (1996), pp. 1–14.

[SZ91] D. W. Stroock and O. Zeitouni. "Microcanonical distributions, Gibbs states, and the equivalence of ensembles". In: *Random walks, Brownian motion, and interacting particle systems*. Springer, 1991, pp. 399–424.