Large deviations, stochastic filtering and trajectorial Gibbs principle in stochastic control

Louis-Pierre Chaintron

Abstract

This dissertation presents the main topics I have been working on for the last two years. I started to work on particle systems and propagation of chaos during my M1 internship with Pierre Degond at Imperial College, in the context of stochastic modelling of biological systems. During my M2 internship, I worked with Philippe Moireau at Inria Saclay on stochastic modelling applied to biology and related estimation problems [CMC19; CCK22]. With the initial motivation of justifying macroscopic laws from microscopic stochastic models, I turned back to particle systems under the supervision of Tony Lelièvre and Julien Reygner at the CERMICS laboratory. The following dissertation presents the mathematical tools we have been using through these works, at the interface between stochastic diffusions and Hamilton-Jacobi-Bellman (HJB) PDEs. The connecting tool from the stochastic setting to the deterministic one is here the large deviation theory, which is briefly described in the first section through the lens of the Freidlin-Wentzell framework [FW98] and its particle system generalisation [DG87]. The second part presents the consequences of the Freidlin-Wentzell setting in stochastic filtering, establishing a first link with HJB PDEs and finite dimensional control. The third part describes the effect of mean-field conditioning on large particle systems; this system can be somehow seen as an infinite dimensional version of the second part, establishing some links with stochastic control and mean-field game systems.

Contents

1 A short introduction to large deviations and Freidlin-Wentzell theory **1**

2 From stochastic filtering to deterministic estimation **4**

3 Gibbs principle on path space and stochastic control **7**

References **14**

1 A short introduction to large deviations and Freidlin-Wentzell theory

This first section is a short introduction to the large deviation theory, starting from the finite dimensional Freidlin-Wentzell theory, before to present an infinite dimensional generalisation using large particle systems. A first example is given by the stochastic diffusion in $\mathbb{R}^d$

$$\mathrm{d}X_t^\varepsilon = b(X_t^\varepsilon)\mathrm{d}t + \sqrt{\varepsilon}\mathrm{d}B_t, \tag{1}$$

for some Lipschitz continuous drift function $b : \mathbb{R}^d \to \mathbb{R}^d$, the stochastic process $(B_t)_{t\geq 0}$ being a $\mathbb{R}^d$-valued Brownian motion. As $\varepsilon \to 0$, the stochastic process $X^\varepsilon_{[0,T]} := (X^\varepsilon_t)_{0\leq t\leq T}$ is expected to converge towards the deterministic flow $(x(t))_{0\leq t\leq T}$ solution of the ODE

$$\frac{\mathrm{d}x(t)}{\mathrm{d}t} = b(x(t)). \tag{2}$$

The large deviation framework which quantifies the deviations away from this limit is known as the Freidlin-Wentzel theory [FW98].

Definition 1 (Large deviation principle)**.** Let E be a Polish space E, $(a_N)_N$ a sequence of positive numbers with $a_N \to +\infty$, and I a non-negative lower-semicontinuous function I on E. A sequence $(\mu^N)_N$ in $\mathcal{P}(E)$ satisfies a Large Deviation Principle (LDP) with speed a_N and rate function I when for any Borel set $A \subset E$, if

$$-\inf_{\mathring{A}} I \leq \liminf_{N\to\infty} a_N^{-1} \log \mu^N(A) \quad \text{and} \quad \limsup_{N\to\infty} a_N^{-1} \log \mu^N(A) \leq -\inf_{\overline{A}} I,$$

where $\mathring{A}$ and $\overline{A}$ denote respectively the interior and closure of A.

Let $(X^N)_{N\geq 1}$ be a sequence of E-valued random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, such that X^N is μ^N-distributed. Roughly speaking, the above definition means that for every x in E,

$$\mathbb{P}\big(X^N \approx x\big) \approx e^{-a_N I(x)} \text{ as } N \to +\infty,$$

capturing deviations from the limit at the exponential scale. When $I(x) > 0$ the above probability goes to 0 as $N \to +\infty$ suggesting that X^N asymptotically belongs to the set $\{x \in E \,/\, I(x) = 0\}$. When this set reduces to a single point $\overline{x}$, the Borel-Cantelli lemma provides the strong convergence

$$X^N \xrightarrow[N\to+\infty]{} \overline{x} \quad \mathbb{P}-\text{a.s.},$$

from the exponential scale estimates. In particular, the sequence $(\mu^N)_N$ weakly converges in $\mathcal{P}(E)$ towards the Dirac measure $\delta_{\overline{x}}$. Once a LDP with rate function I is known, the problem of determining the limit of $(\mu^N)_N$ reduces to the study of minimisers for I.

Let us hark back to Example (1). By sequential characterisation of the limit, Definition (1) can be adapted in a straightforward way to the continuous framework $\varepsilon \to 0$ instead of $N \to +\infty$. In the following result, the speed rate a_N becomes thus ε^{-1}:

Theorem 2 (Freidlin-Wentzel)**.** *The family $(X^\varepsilon_{[0,T]})_{\varepsilon>0}$ satisfies a LDP with speed ε^{-1} and rate function*

$$I((x_t)_{0\leq t\leq T}) = \begin{cases} \frac{1}{2}\int_0^T |\dot{x}_t - b(x_t)|^2 \mathrm{d}t & \text{if } (x_t)_{0\leq t\leq T} \in AC([0,T], \mathbb{R}^d), \\ +\infty & \text{else.} \end{cases} \tag{3}$$

The deterministic flow (2) *is the unique minimiser of I.*

The second part of this dissertation describes the $\varepsilon \to 0$ limit of the stochastic filtering procedure for (1), and how deterministic estimation can be recovered from it [JB88]. It then briefly quotes results that have been proved in the under review paper [Cha+22] in collaboration with A. Gonzales, L. Mertz and P. Moireau (who is now my PhD co-advisor), which study a similar limit when the differential equality (1) is replaced by a sub-differential inclusion (case of reflected processes).

Another example of LDP can be obtained by considering a countable sequence $(X^i_{[0,T]})_{N\geq 1}$ of i.i.d. copies of the process defined in (1) with $\varepsilon = 1$. The $(\mathbb{R}^{dN})$-valued diffusion process

$$\vec{X}^N_{[0,T]} = (X^1_{[0,T]}, \dots X^N_{[0,T]}).$$

will be called *particle system.* Let

$$\Pi(\vec{X}^N_{[0,T]}) := \frac{1}{N}\sum_{i=1}^N \delta_{X^i_{[0,T]}}$$

be the pathwise empirical measure of $\vec{X}^N_{[0,T]}$: it is a $\mathcal{P}(C([0,T],\mathbb{R}^d))$-valued random variable. The $N\to+\infty$ behaviour of $\Pi(\vec{X}^N_{[0,T]})$ is described by the following LDP, where $\nu_{[0,T]}$ denotes the law of $X^1_{[0,T]}$.

Theorem 3 (Sanov)**.** *The sequence* $(\mathrm{Law}(\Pi(\vec{X}^N_{[0,T]})))_{N\geq 1}$ *in* $\mathcal{P}(\mathcal{P}(C([0,T],\mathbb{R}^d))$ *satisfies a LDP with speed* N *and rate function*

$$H(\mu_{[0,T]}|\nu_{[0,T]}) = \begin{cases} \int_{C([0,T],\mathbb{R}^d)} \log\frac{\mathrm{d}\mu_{[0,T]}}{\mathrm{d}\nu_{[0,T]}}\mathrm{d}\mu_{[0,T]} & \text{if } \mu_{[0,T]} \ll \nu_{[0,T]}, \\ +\infty & \text{else,} \end{cases}$$

which is called the relative entropy w.r.t. to $\nu_{[0,T]}$.

The convention $0\log 0 = 0$ is used, and $H(\mu|\nu) = +\infty$ if the function within the integral is not integrable. Since $\nu_{[0,T]}$ is the unique minimiser of $H(\mu_{[0,T]}|\nu_{[0,T]})$, this implies the a.s. weak convergence of the random measure $\mathrm{Law}(\Pi(\vec{X}^N_{[0,T]}))$ towards $\nu_{[0,T]}$, and *this last result is nothing but the classical law of large numbers.*

From $\nu_{[0,T]}$, one can obtain the continuous flow of time-marginals $(\nu_t)_{0\leq t\leq T}$, ν_t being the law of X^1_t at time t. The infinitesimal generator of the Markov process $X^1_{[0,T]}$ is the differential operator

$$L := b\cdot\partial_x + \frac{1}{2}\partial^2_x.$$

Using the notation

$$\mathbb{E}\varphi(X^1_t) = \int_{\mathbb{R}^d}\varphi\mathrm{d}\nu_t =: \langle\nu_t,\varphi\rangle,$$

the flow $(\nu_t)_{0\leq t\leq T}$, given ν_0, is characterised by

$$\forall\varphi\in C^2_b(\mathbb{R}^d), \qquad \frac{\mathrm{d}}{\mathrm{d}t}\langle\nu_t,\varphi\rangle = \langle\nu_t, L\varphi\rangle. \tag{4}$$

Similarly, the flow of marginals for $\Pi(\vec{X}^N_{[0,T]})$ is $(\pi(\vec{X}^N_t))_{0\leq t\leq T}$, the empirical measure $\pi(\vec{X}^N_t)$ at time t being defined as

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

The random flow $(\pi(\vec{X}^N_t))_{0\leq t\leq T}$ is thus a $C([0,T],\mathcal{P}(\mathbb{R}^d))$-valued random variable. The next LDP can be seen as an infinite-dimensional analogous of Theorem 2. Using Varadhan's contraction principle the sequence $(\mathrm{Law}((\pi(\vec{X}^N_t))_{0\leq t\leq T}))_{N\geq 1}$ in $\mathcal{P}(C([0,T],\mathcal{P}(\mathbb{R}^d))$ satisfies a LDP with rate function

$$S((\mu_t)_{0\leq t\leq T}) := \inf_{\substack{\mu'\in\mathcal{P}(C([0,T],\mathbb{R}^d))\\ \forall 0\leq t\leq T,\ \mu'_t=\mu_t}} H(\mu'_{[0,T]}|\nu_{[0,T]}).$$

Theorem 4 (Dawson-Gärtner)**.** *The above rate function rewrites*

$$S((\mu_t)_{0\leq t\leq T}) = \begin{cases} \frac{1}{2}\int_0^T \|\partial_t\mu_t - L^\star\mu_t\|^2_{\mu_t}\mathrm{d}t & \text{if } \mu_{[0,T]}\in AC([0,T],\mathcal{P}(\mathbb{R}^d)) \text{ and } \mu_0=\nu_0, \\ +\infty & \text{else.} \end{cases} \tag{5}$$

where the distribution $L^\star\mu_t$ *is defined against test functions* φ *in* $C^\infty_c(\mathbb{R}^d)$ *by* $\langle L^\star\mu_t,\varphi\rangle = \langle\mu_t, L\varphi\rangle$*, and for any distribution* $\mathcal{T}$ *one defines* $\|\mathcal{T}\|^2_\mu := \sup_{\varphi\in C^\infty_c}\langle\mathcal{T},\varphi\rangle - \frac{1}{2}\langle\mu,|\partial_x\varphi|^2\rangle$.

In particular, the unique minimiser of S is the flow $(\nu_t)_{0\leq t\leq T}$ which solves (4), or equivalently $\partial_t \nu_t = L^\star \nu_t$ in a weak sense. Using Ito's formula with some test function φ in $C_b^2(\mathbb{R}^d)$,

$$\langle \pi(\vec{X}_t^N), \varphi\rangle = \langle \nu_0, \varphi\rangle + \int_0^t \langle \pi(\vec{X}_s^N), L\varphi\rangle \mathrm{d}s + \frac{1}{N}\sum_{i=1}^N \int_0^t \partial_x \varphi(X_s^i)\cdot \mathrm{d}B_s^i.$$

The $X_{[0,T]}^i$ (and the related $B_{[0,T]}^i$) being independent, the standard deviation of the last term is $O(N^{-\frac{1}{2}})$ and the above equation on $\pi(\vec{X}_t^N)$ appears as an infinite dimensional version (in a weak sense) of Equation (1), ε being replaced by N^{-1}.

The third part of this dissertation makes use of these tools to describe the large deviations of $\Pi(\vec{X}_{[0,T]}^N)$ conditioned by the mean-field event

$$\forall t \in [0,T], \qquad \Psi(\pi(\vec{X}_t^N)) \leq 0, \tag{6}$$

for some regular function $\Psi : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$, which can be seen as the "energy" of the system (a physically relevant linear example when $b(x) = -\partial_x V(x)$ is given by $\Psi(\mu) = \langle \mu, V\rangle$). The related minimisation problem on the path space is a work in progress with G. Conforti (CMAP, École Polytechnique) and J. Reygner (who is now my main PhD advisor), and it can be interpreted as a stochastic control problem, which is reminiscent of mean-field game situations.

2 From stochastic filtering to deterministic estimation

This section presents the effect of vanishing noise on the stochastic filtering procedure, showing how deterministic estimation can be recovered from it. Let us assume that the process $X_{[0,T]}^\varepsilon$ given by (1) is observed through the noisy process

$$\mathrm{d}Y_t^\varepsilon = h(X_t^\varepsilon)\mathrm{d}t + \sqrt{\varepsilon}\mathrm{d}B'_t, \tag{7}$$

the Brownian motion $(B'_t)_{t\geq 0}$ being independent from the one driving X_t^ε, and h being some Lipschitz continuous function; these processes are assumed to be built on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$. Stochastic filtering aims to compute an optimal estimator of X_t^ε given the knowledge of Y_s^ε until time t. More rigorously, one wants to compute the conditional law of X_t^ε knowing $(Y_s^\varepsilon)_{0\leq s\leq t}$, i.e. the random probability measure π_t^ε such that

$$\langle \pi_t^\varepsilon, \varphi\rangle = \mathbb{E}\left[\varphi(X_t^\varepsilon) | \sigma\left(Y_s^\varepsilon\right)_{0\leq s\leq t}\right],$$

for any bounded continuous $\varphi : \mathbb{R}^d \to \mathbb{R}$, $\sigma\left(Y_s^\varepsilon\right)_{0\leq s\leq t}$ being the σ-algebra generated by the observation Y_s^ε until time t. In the current setting, one can prove that all the random variables at stakes are square-integrable, and the $\sigma\left(Y_s^\varepsilon\right)_{0\leq s\leq t}$-measurable random variable $\langle \pi_t^\varepsilon, \varphi\rangle$ is optimal in the sense that

$$\mathbb{E}|\langle \pi_t^\varepsilon, \varphi\rangle - \varphi(X_t^\varepsilon)|^2 = \inf_{Z \text{ is } \sigma(Y_s^\varepsilon)_{0\leq s\leq t}\text{-measurable}} \mathbb{E}|Z - \varphi(X_t^\varepsilon)|^2.$$

When the Novikok condition

$$\mathbb{E}\exp\left[\frac{1}{2\varepsilon}\int_0^T |h(X_t^\varepsilon)|^2 \mathrm{d}t\right] < +\infty$$

is verified, a classical approach (see e.g. [BC08]) considers the change of probability measure

$$\left.\frac{\mathrm{d}\tilde{\mathbb{P}}}{\mathrm{d}\mathbb{P}}\right|_{\mathcal{F}_t} = (Z_t^\varepsilon)^{-1},$$

where

$$Z_t^\varepsilon = \exp\left[\frac{1}{\sqrt{\varepsilon}}\int_0^t h(X_s^\varepsilon)\mathrm{d}Y_s^\varepsilon - \frac{1}{2\varepsilon}\int_0^t h^2(X_s^\varepsilon)\mathrm{d}s\right].$$

The Girsanov theorem then shows [Zak69; BC08] that $(\frac{Y_t^\varepsilon}{\sqrt{\varepsilon}})_{t\geq 0}$ is a Brownian motion under $\tilde{\mathbb{P}}$. Define now the non-normalised measure ρ_t by

$$\langle \rho_t^\varepsilon, \varphi\rangle = \tilde{\mathbb{E}}\left[Z_t^\varepsilon \varphi(X_t) | \sigma\left(Y_s\right)_{0\leq s\leq t}\right].$$

This measure can be linked to π_t by the Kallianpur-Striebel formula

$$\langle \pi_t^\varepsilon, \varphi\rangle = \frac{\langle \rho_t^\varepsilon, \varphi\rangle}{\langle \rho_t^\varepsilon, 1\rangle},$$

which is an analogous of the Bayes formula [KS68]. A linear measure-valued stochastic PDE can then be obtained for ρ_t^ε: it is the Zakai equation [Zak69]. Under suitable assumptions, ρ_t^ε has a density $q^\varepsilon(t,x)$ w.r.t. the Lebesgue measure, which solves in a weak sense

$$\mathrm{d}q^\varepsilon(t,\cdot) = \frac{\varepsilon}{2}\partial_x^2 q^\varepsilon(t,\cdot)\mathrm{d}t - \mathrm{div}(q^\varepsilon(t,\cdot)b)\mathrm{d}t + \frac{q^\varepsilon(t,\cdot)}{\varepsilon}\mathrm{d}Y_t^\varepsilon,$$

see [Par82] for a complete presentation. In many areas [CC05], one would like to write q^ε as a continuous function of the observation $(Y_t^\varepsilon)_{0\leq s\leq t}$, in order to perform recursive computations. The method in [Dos77] suggests the transform

$$p^\varepsilon(t,x) := \exp\left[-\frac{Y_t^\varepsilon h(x)}{\varepsilon}\right] q^\varepsilon(t,x),$$

which leads to the *robust Zakai equation*

$$\partial_t p^\varepsilon(t,\cdot) + g^\varepsilon(t,x)\partial_x p^\varepsilon(t,\cdot) + \frac{1}{\varepsilon}V^\varepsilon(t,x)p^\varepsilon(t,\cdot) = \frac{\varepsilon}{2}\partial_{xx}^2 p^\varepsilon(t,\cdot), \tag{8}$$

where the Y_t^ε-dependent coefficients are $g^\varepsilon(t,x) := b(x) - Y_t^\varepsilon \partial_x h(x,t)$ and

$$V^\varepsilon(t,x) := \frac{h^2(x)}{2} + Y_t^\varepsilon L_\varepsilon h(x) - \frac{1}{2}(Y_t^\varepsilon)^2(\partial_x h(x))^2 + \varepsilon \mathrm{div}\left[b(x) - Y_t^\varepsilon \partial_x h(x)\right].$$

In (8), Y_t^ε is a mere parameter, which only appears within the coefficients of the equation. The random variable Y_t^ε being defined as a measurable function $\omega \in \Omega \mapsto Y^\varepsilon(t,\omega)$, (8) can be seen as a family of deterministic PDEs indexed by the parameter ω. At this point, it is only necessary to solve the PDE (8) for a given realisation $(y(s))_{0\leq s\leq t}$ of the continuous trajectory $(Y^\varepsilon(s,\omega))_{0\leq s\leq t}$. The remaining question will then be the measurability of the solution in ω, in order to recover a stochastic process $\omega \mapsto (p^\varepsilon(\omega,s,x))_{0\leq s\leq t}$ from solving (8) for each continuous $(y(s))_{0\leq s\leq t}$. This measurability is proved in [Dos77; Sus78], and the result still holds if (8) is solved only for C^1 trajectories $(y(s))_{0\leq s\leq t}$: p^ε can thus be seen as a deterministic function which depends on a C^1 trajectory $(y(s))_{0\leq s\leq t}$, and p^ε is C^2 using classical results on linear parabolic PDEs.

On the other side, deterministic estimation considers a C^1 function $(y(s))_{0\leq s\leq t}$ which models observation, and looks for some (ζ, w) in $\mathbb{R}^d \times L^2(0,t)$ which minimises the error functional

$$J(\zeta,w,t) := \psi(\zeta) + \int_0^t \frac{1}{2}|w(s)|^2 + \frac{1}{2}|\dot{y}(s) - h(x_w^\zeta(s))|^2 \mathrm{d}s,$$

where the dynamics model is

$$\begin{cases} \dot{x}_w^\zeta(s) = b(x_w^\zeta(s)) + w(s), \\ x_w^\zeta(0) = \zeta. \end{cases} \tag{9}$$

In the above equation analogous to (1), the "control parameter" w acts as a deterministic noise, and minimising $J(\zeta, w, t)$ constrains $h(x^\zeta_w(s))$ to be close to the observation $\dot{y}(s)$, as in (7). Given x in $\mathbb{R}^d$, [Mor68] defines the cost-to-come function

$$V(t,x) = \inf_{(\zeta,w)\ /\ x^\zeta_w(t)=x} J(\zeta, w, t),$$

i.e. the cost to bring the dynamics at x at time t. A good estimator of $x(t)$ knowing the observation $(y(s))_{0\le s\le t}$ is then given by any minimiser of $x \mapsto V(t,x)$. Moreover, using the dynamic programming principle, one can compute V as the unique continuous viscosity solution [Bar94] of the Hamilton-Jacobi-Bellman (HJB) equation

$$\begin{cases} \partial_t V(t,x) + H(x,t,\partial_x V(t,x)) = 0, \\ V(0,x) = \psi(x). \end{cases} \tag{10}$$

the hamiltonian H being defined as

$$H(x,t,p) = \frac{1}{2}|p|^2 + b(x)p - \frac{1}{2}|\dot{y}(t) - h(x)|^2.$$

The following result shows that deterministic estimation selects the maximum likelihood estimator from the stochastic filtering density.

Theorem 5 (Baras-James). *For any (t,x) in $\mathbb{R}_+ \times \mathbb{R}^d$,*

$$\lim_{\varepsilon\to 0} -\varepsilon \log q^\varepsilon(t,x) = V(t,x),$$

provided this holds true at $t = 0$.

Roughly speaking, this means that

$$q^\varepsilon(t,x) \approx_{\varepsilon\to 0} \exp\left[-\frac{V(t,x)}{\varepsilon}\right],$$

and this is strengthened into an LDP for the law of $X^\varepsilon_{[0,t]}$ given the realisation $(y(s))_{0\le s\le t}$ of $(Y^\varepsilon_s)_{0\le s\le t}$ in [Hij84; JB88]. The proof of Theorem 5 relies on the Hopf-Cole transformation [EI85; Fle97]

$$S^\varepsilon = -\varepsilon \log p^\varepsilon - y(t)h(x),$$

which turns the parabolic equation 8 into a viscous HJB equation

$$\partial_t S^\varepsilon(t,x) + H^\varepsilon(x,t,\partial_x S^\varepsilon(t,x)) = \frac{\varepsilon}{2}\partial^2_{xx} S^\varepsilon(t,x),$$

where H^ε uniformly converges towards H on every compact set as ε goes to 0. The maximum principle then provides uniform in ε estimates on $S^\varepsilon(t,x)$ and $\partial_x S^\varepsilon(t,x)$, giving some compactness in the space of continuous functions. By stability of viscosity solutions, S^ε eventually converges towards the unique viscosity solution of (10).

In [Cha+22], we wanted to extend Theorem 5 to the case of the 1D sub-differential dynamics in $\mathbb{R}_+$

$$\mathrm{d}X_t + \partial I_{\mathbb{R}_+}(X_t)(\mathrm{d}t) \ni b(X_t) + \sqrt{\varepsilon}\mathrm{d}B_t, \tag{11}$$

where $I_{\mathbb{R}_+} = \infty \mathbb{1}_{\mathbb{R}_-}$ is the convex indicator function of $\mathbb{R}_+$, so that

$$\partial I_{\mathbb{R}_+}(x) = \begin{cases} \{0\} \text{ if } x > 0, \\ \mathbb{R}_- \text{ if } x = 0. \end{cases}$$

When $b = 0$, the process (11) is well-known as the solution of the Skorokhod problem [Sko61]; the filtering procedure for this kind of reflected process is extensively studied in [Par78]. We managed to establish the equivalent of Theorem 5, but the equivalent cost-to-come in our paper [Cha+22] is defined from the backward in time dynamics,

$$\begin{cases} \dot{z}^{t,x}(s) \in b(z^{t,x}(s)) + w(s) + \partial I_{\mathbb{R}_+}\left(z^{t,x}(s)\right), \quad 0 \le s \le t, \\ z^{t,x}(t) = x. \end{cases} \tag{12}$$

instead of the forward one

$$\dot{y}(s) + b(y(s)) + \partial I_{\mathbb{R}_+}\left(y(s)\right) \ni w(s).$$

This does not provide a recursive way to compute minimisers of $x \mapsto V(t,x)$, and understanding how to circumvent this problem is an ongoing work.

3 Gibbs principle on path space and stochastic control

This section goes back to the infinite dimensional particle setting, studying the effect of mean-field conditioning in the large N limit. Let now $\nu_{[0,T]}$ denote *any* measure in $\mathcal{P}(C([0,T],\mathbb{R}^d))$, and as previously let $\vec{X}^N_{[0,T]}$ be a $\nu^{\otimes N}_{[0,T]}$-distributed random variable. Given a bounded from below lower semi-continuous function $\Psi : \mathcal{P}(\mathbb{R}^d) \to \mathbb{R}$, let us define the stopping time

$$\tau^N_\Psi := \inf\{t > 0 \,|\, \Psi(\pi(\vec{X}^N_t)) > 0\}.$$

We are interested in the $N \to +\infty$ behaviour of the Markov process $\vec{X}^N_{[0,T]}$ conditioned by the mean-field path-dependent event $\{T < \tau^N_\Psi\}$. The conditioned single particle processes are no more independent in law, but exchangeability is preserved (in particular, the conditional law of single particles are identical). From this, the law of the process $(\pi(\vec{X}^N_t))_{0\le t\le T}$ conditioned by $\{T < \tau^N_\Psi\}$ is still a Markov measure. From [SZ91; DZ96], it is expected that the sequence of pathwise conditional laws

$$\mathrm{Law}\big(\Pi(\vec{X}^N_{[0,T]})\big|T < \tau^N_\psi\big) \in \mathcal{P}\left(\mathcal{P}\left(C\left([0,T],\mathbb{R}^d\right)\right)\right), \quad N \ge 1,$$

satisfies a LDP with rate function

$$I^T_\psi\left(\mu_{[0,T]}\right) := \begin{cases} H\big(\mu_{[0,T]}\big|\nu_{[0,T]}\big) - \overline{I}^T_\psi \text{ if } \mu_{[0,T]} \ll \nu_{[0,T]} \text{ and } \forall\, 0 \le t \le T,\, \langle \mu_t, \psi\rangle \le 0, \\ +\infty \text{ else.} \end{cases}$$

where

$$\overline{I}^T_\Psi := \inf_{\substack{\mu_{[0,T]}\in\mathcal{P}(C([0,T],\mathbb{R}^d))\\ \forall 0\le t\le T, \Psi(\mu_t)\le 0}} H\big(\mu_{[0,T]}\big|\nu_{[0,T]}\big). \tag{13}$$

The rigorous proof of this LDP is an ongoing work. When $\Psi(\mu) = \langle \mu, \psi\rangle$ is linear, the following is proven in [DZ96, Theorem 2.1] (it is an adaptation of [Csi84, Theorem 1]):

Theorem 6 (H-convergence)**.** *When $\Psi(\mu) = \langle \mu, \psi\rangle$ is linear, $\overline{I}^T_\psi$ is realised by a unique $\overline{\mu}_{[0,T]}$ such that*

$$H\big(\mathrm{Law}(X^1_{[0,T]}\,|\,T < \tau^N_\psi)|\overline{\mu}_{[0,T]}\big) \le \frac{1}{N} H\big(\mathrm{Law}(\vec{X}^N_{[0,T]}\,|\,T < \tau^N_\psi)|\overline{\mu}^{\otimes N}_{[0,T]}\big)$$
$$\le -\frac{1}{N}\log \mathbb{P}(T < \tau^N_\Psi) - H(\overline{\mu}_{[0,T]}|\nu_{[0,T]}),$$

and the r.h.s. vanishes as $N \to +\infty$ by the Sanov theorem.

This motivates the study of the minimisation problem (13), which was already interesting in itself. Let us start with an easy existence result, whose proof is strongly inspired from [Nut21, Lemma 1.8].

Lemma 7 (Existence of a minimiser in the weak l.s.c. case)**.** *Assume Ψ to be lower semi-continuous (l.s.c.) w.r.t. the weak topology on $\mathcal{P}(\mathbb{R}^d)$. The infimum* (13) *is then achieved by at least one $\overline{\mu}_{[0,T]}$ in $\mathcal{P}(C([0,T],\mathbb{R}^d))$. Uniqueness holds if Ψ is moreover convex.*

Proof. The infimum $\overline{I}_{\Psi}^{T}$ being assumed finite, consider a minimising sequence $(\overline{\mu}_{[0,T]}^{k})_{k\in\mathbb{N}}$ of path measures satisfying the constraint with finite entropy w.r.t. $\nu_{[0,T]}$. Then

$$\sup_{k\in\mathbb{N}} H\big(\mu_{[0,T]}^{k}|\nu_{[0,T]}\big) < +\infty,$$

giving uniform $\nu_{[0,T]}$-integrability for the sequence $(\frac{\mathrm{d}\mu_{[0,T]}^{k}}{\mathrm{d}\nu_{[0,T]}})_{k\in\mathbb{N}}$. The Dunford-Pettis theorem [BR07, Theorem 4.7.18] then shows that this sequence is relatively compact for the weak topology (against L^{∞} functions) on $L^{1}(\nu_{[0,T]})$, and let $f_{[0,T]}$ be any limit point. The set of probability densities being weakly closed, $f_{[0,T]}$ gives rise to a path measure $\mathrm{d}\overline{\mu}_{[0,T]} := f_{[0,T]}\mathrm{d}\nu_{[0,T]}$, which is a limit point of $(\mu_{[0,T]}^{k})_{k\in\mathbb{N}}$ w.r.t. weak convergence along some increasing sub-sequence $(k_i)_{i\in\mathbb{N}}$. The space $C([0,T],\mathbb{R}^d)$ being Polish, the weak lower semi-continuity of $\mu_{[0,T]} \mapsto H(\mu_{[0,T]}|\nu_{[0,T]})$ (see e.g. [Nut21, Lemma 1.3]) gives

$$H\big(\overline{\mu}_{[0,T]}|\nu_{[0,T]}\big) \le \liminf_{i\to+\infty} H\big(\mu_{[0,T]}^{k_i}|\nu_{[0,T]}\big) = \overline{I}_{\Psi}^{T},$$

and the lower semi-continuity of Ψ gives

$$\Psi\big(\overline{\mu}_{[0,T]}\big) \le \liminf_{i\to+\infty} \Psi\big(\mu_{[0,T]}^{k_i}\big) \le 0,$$

which concludes the proof. When Ψ is convex, uniqueness stems from the fact that $\mu_{[0,T]} \mapsto H(\mu_{[0,T]}|\nu_{[0,T]})$ is strictly convex. □

From now on, let us assume that Ψ is convex and C^1 in the sense of [CD+18]: Ψ admits a jointly continuous linear functional derivative $\frac{\delta\Psi}{\delta\mu} : \mathcal{P}(\mathbb{R}^d)\times\mathbb{R}^d \to \mathbb{R}$, i.e.

$$\forall \mu,\mu' \in \mathcal{P}(\mathbb{R}^d),\ \Psi(\mu)-\Psi(\mu') = \int_0^1 \Big\langle \mu-\mu', \frac{\delta\Psi}{\delta\mu}((1-r)\mu'+r\mu,\cdot)\Big\rangle \mathrm{d}r, \tag{14}$$

with the (arbitrary) convention that $\langle \mu, \frac{\delta\Psi}{\delta\mu}(\mu)\rangle = 0$; Ψ is in particular continuous and Lemma 7 applies. Let us eventually assume some compatibility between $\nu_{[0,T]}$ and Ψ (constraint qualification): a threshold $\gamma > 0$ exists such that

$$\nu_{\Psi}^{T}(\gamma) := \inf_{\mu_{[0,T]}\in\mathcal{P}(C([0,T],\mathbb{R}^d))} \nu_{[0,T]}\Big(\Big\{\forall 0\le t\le T,\ \frac{\delta\Psi}{\delta\mu}(\mu_t,x_t) \le -\gamma-\Psi(\mu_t)\Big\}\Big) > 0. \tag{15}$$

Example 8 (Linear case)**.** *When $\Psi(\mu) = \langle\mu,\psi\rangle$, note that $\frac{\delta\Psi}{\delta\mu}(\mu) = \psi - \langle\mu,\psi\rangle$, and the qualification assumption simplifies as*

$$\nu_{\psi}^{T}(\gamma) := \nu_{[0,T]}\big(\big\{\forall 0\le t\le T,\ \psi(x_t)\le -\gamma\big\}\big) > 0. \tag{16}$$

The main result below characterises the minimisers in (13) as Gibbs measures.

Theorem 9 (Gibbs principle)**.** *Under the above assumptions, the minimum in* (13) *is realised by a unique $\overline{\mu}_{[0,T]}$, and a finite non-negative measure $\overline{\lambda}$ over $[0,T]$ exists such that*

$$\mathrm{d}\overline{\mu}_{[0,T]}\big(x_{[0,T]}\big) = \big(Z_T^{\Psi}\big)^{-1} \exp\Big[-\int_0^T \frac{\delta\Psi}{\delta\mu}(\overline{\mu}_t,x_t)\overline{\lambda}(\mathrm{d}t)\Big]\mathrm{d}\nu_{[0,T]}\big(x_{[0,T]}\big), \tag{17}$$

the partition function Z_{Ψ}^{T} being a normalising constant. Moreover, the complementary slackness condition is satisfied:

$$\int_0^T \Psi(\overline{\mu}_t)\overline{\lambda}(\mathrm{d}t) = 0. \tag{18}$$

Since $\Psi(\mu_t) \leq 0$ *for every* t*, this ensures that* $\Psi(\mu_t) = 0$ $\overline{\lambda}$*-a.e. Reciprocally,* $\overline{\mu}_{[0,T]}$ *is characterised by the facts of being admissible and such that* (17)-(18) *holds for some* $\overline{\lambda}$ *in* $\mathcal{M}_+([0,T])$*.*

Let us define the Lagrangian

$$\mathcal{L}(\mu_{[0,T]}, \lambda) := H\big(\mu_{[0,T]} \big| \nu_{[0,T]}\big) + \int_0^T \Psi(\mu_t)\lambda(\mathrm{d}t).$$

To turn the constrained problem (13) into an unconstrained one, introduce the dual problem

$$\sup_{\lambda \in \mathcal{M}_+([0,T])} \inf_{\mu_{[0,T]} \ll \nu_{[0,T]}} \mathcal{L}(\mu_{[0,T]}, \lambda). \tag{19}$$

Define then the functional, for any λ in $\mathcal{M}_+([0,T])$

$$F(\lambda) := \inf_{\mu_{[0,T]} \ll \nu_{[0,T]}} \mathcal{L}(\mu_{[0,T]}, \lambda). \tag{20}$$

A direct adaptation of the proof of Lemma 7 shows that this infimum is realised by a unique path measure $\mu^\lambda_{[0,T]}$. As an infimum of such functions (even linear ones), F is both concave and upper semi-continuous (u.s.c.). The next lemma is a classical linearisation result (see e.g. [BC18, Lemma 3.1] or [Dau20, Proposition 3.1]).

Lemma 10 (Linearised problem)**.** *Given* λ *in* $\mathcal{M}_+([0,T])$*,* $\mu^\lambda_{[0,T]}$ *is the unique minimiser for the linearised problem*

$$\inf_{\mu_{[0,T]} \ll \nu_{[0,T]}} H\big(\mu_{[0,T]} \big| \nu_{[0,T]}\big) + \int_0^T \langle \mu_t, \frac{\delta\Psi}{\delta\mu}(\mu^\lambda_t)\rangle \lambda(\mathrm{d}t). \tag{21}$$

Proof. For any path measure $\mu_{[0,T]}$ and $\varepsilon > 0$, consider the perturbation $\mu^\varepsilon_{[0,T]} := (1-\varepsilon)\mu^\lambda_{[0,T]} + \varepsilon\mu_{[0,T]}$. Since $\mathcal{L}(\mu^\varepsilon_{[0,T]}, \lambda) \geq \mathcal{L}(\mu^\lambda_{[0,T]}, \lambda)$, straightforward computations give the result, using the C^1 regularity of Ψ and the strict convexity of H, before to send ε to 0. □

Lemma 11 (Gibbs measure)**.** *The unique minimiser* $\mu^\lambda_{[0,T]}$ *in* (21) *has the Gibbs shape*

$$\mathrm{d}\mu^\lambda_{[0,T]}\big(x_{[0,T]}\big) = Z^{-1}_\lambda \exp\Big[-\int_0^T \lambda(t)\frac{\delta\Psi}{\delta\mu}(\mu^\lambda_t, x_t)\mathrm{d}t\Big]\mathrm{d}\nu_{[0,T]}\big(x_{[0,T]}\big), \tag{22}$$

Proof. This is a direct consequence of the classical computation:

$$H\big(\mu_{[0,T]}|\nu_{[0,T]}\big) + \int_0^T \langle \mu_t, \frac{\delta\Psi}{\delta\mu}(\mu^\lambda_t)\rangle\lambda(\mathrm{d}t) = H\big(\mu^\lambda_{[0,T]}|\nu_{[0,T]}\big) + \int_0^T \langle \mu^\lambda_t, \frac{\delta\Psi}{\delta\mu}(\mu^\lambda_t)\rangle\lambda(\mathrm{d}t) + H\big(\mu_{[0,T]}|\mu^\lambda_{[0,T]}\big), \tag{23}$$

together with the fact that $H\big(\mu_{[0,T]}|\mu^\lambda_{[0,T]}\big) \geq 0$ with equality if and only if $\mu_{[0,T]} = \mu^\lambda_{[0,T]}$. □

Lemma 12 (Existence of maximisers)**.** *The supremum in* (19) *is realised by at least one* $\overline{\lambda}$ *in* $\mathcal{M}_+([0,T])$*.*

Proof. In the problem (19) of maximising $F(\lambda)$, consider an optimising sequence $(\lambda_k)_{k\in\mathbb{N}}$, and choose a related sequence of optimising path measures $(\mu^{\lambda_k}_{[0,T]})_{k\in\mathbb{N}}$. By the above lemma, the $\mu^{\lambda_k}_{[0,T]}$ have the shape (22), so that

$$\begin{aligned} F(\lambda_k) &= -\log Z_{\lambda_k} - \int_0^T \langle \mu^{\lambda_k}_t, \frac{\delta\Psi}{\delta\mu}(\mu^{\lambda_k}_t)\rangle\lambda^k(\mathrm{d}t) + \int_0^T \Psi(\mu^{\lambda_k}_t)\lambda^k(\mathrm{d}t) \\ &= -\log \mathbb{E}_{\nu_{[0,T]}} \exp\Big[-\int_0^T \frac{\delta\Psi}{\delta\mu}(\mu^{\lambda_k}_t, x_t)\lambda^k(\mathrm{d}t)\Big] + \int_0^T \Psi(\mu^{\lambda_k}_t)\lambda^k(\mathrm{d}t), \end{aligned}$$

recalling the convention $\langle \mu_t^{\lambda_k}, \frac{\delta\Psi}{\delta\mu}(\mu_t^{\lambda_k})\rangle = 0$. Using Assumption (15),

$$\begin{aligned}
\log \mathbb{E}_\nu \exp&\Big[-\int_0^T \frac{\delta\Psi}{\delta\mu}(\mu_t^{\lambda_k}, x_t)\lambda_k(\mathrm{d}t)\Big] \\
&\geq \log \mathbb{E}_\nu \mathbb{1}_{\{\forall 0\leq t\leq T,\ \frac{\delta\Psi}{\delta\mu}(\mu_t^{\lambda_k}, x_t)\lambda_k(\mathrm{d}t)\leq -\gamma-\Psi(\mu_t)\}} \exp\Big[-\int_0^T \frac{\delta\Psi}{\delta\mu}(\mu_t^{\lambda_k}, x_t)\lambda_k(\mathrm{d}t)\Big] \\
&\geq \log \nu_\psi^T(\gamma) + \gamma\lambda_k([0,T]) + \int_0^T \Psi(\mu_t^{\lambda_k})\lambda^k(\mathrm{d}t),
\end{aligned}$$

hence

$$F(\lambda_k) \leq -\gamma\lambda_k([0,T]) - \log \nu_\Psi^T(\gamma).$$

Since $(\lambda_k)_k$ is maximising and $F(0)=0$, $F(\lambda_k)$ must be greater than -1 for k large enough, giving the bound

$$\lambda_k([0,T]) \leq \gamma^{-1}(1-\log \nu_\Psi^T(\gamma)),$$

which is uniform in k. The sequence $(\lambda_k([0,T]))_{k\in\mathbb{N}}$ being bounded, one can assume (up to extracting a sub-sequence) that it converges towards some non-negative real $\tilde{\lambda}$. If $\tilde{\lambda}=0$, there is nothing more to do; otherwise, $(\frac{\lambda_k}{\lambda_k([0,T])})_{k\in\mathbb{N}}$ is a sequence of probability measures over the compact set $[0,T]$, hence a tight sequence by the Prokhorov theorem. Up to extracting once again, this sequence weakly converges towards a probability measure over $[0,T]$, and $(\lambda_k)_{k\in\mathbb{N}}$ weakly converges towards some non-negative measure $\overline{\lambda}$ over $[0,T]$ with finite mass $\tilde{\lambda}$. F being u.s.c., this yields

$$F(\overline{\lambda}) \geq \limsup_{k\to+\infty} F(\lambda_k),$$

proving that $\overline{\lambda}$ achieves the supremum in the dual problem (19). □

Proposition 13 (Primal optimality)**.** *The measure* $\overline{\mu}_{[0,T]} := \mu_{[0,T]}^{\overline{\lambda}}$ *is admissible for the primal problem* (13)*:*

$$\forall t\in[0,T], \qquad \Psi(\overline{\mu}_t)\leq 0.$$

and the complementary slackness condition

$$\int_0^T \Psi(\overline{\mu}_t)\overline{\lambda}(\mathrm{d}t) \geq 0,$$

is verified: $\mu_{[0,T]}^{\overline{\lambda}}$ *is thus the unique solution of* (13)*.*

Proof. Consider $0\leq t_0 < T$ and a sequence $(\varepsilon_k)_{k\in\mathbb{N}}$ in $(0, T-t_0]$ which converges towards 0. Define then $\lambda_k := \overline{\lambda} + \mathbb{1}_{[t_0,t_0+\varepsilon_k]}$. For every k, using the maximality of $\overline{\lambda}$:

$$\mathcal{L}(\mu_{[0,T]}^{\lambda_k}, \lambda_k) = F(\lambda_k) \leq F(\overline{\lambda}) \leq \mathcal{L}(\mu_{[0,T]}^{\lambda_k}, \overline{\lambda}),$$

so that

$$\int_{t_0}^{t_0+\varepsilon_k} \Psi(\mu_t^{\lambda_k})\mathrm{d}t = \int_0^T (\lambda_k(t)-\overline{\lambda}(t))\Psi(\mu_t^{\lambda_k})\mathrm{d}t \leq 0. \tag{24}$$

By the compactness argument used in the proof of Lemma 7, the sequence $(\mu_{[0,T]}^{\lambda_k})_{k\in\mathbb{N}}$ has some limit point $\mu_{[0,T]}^\infty$, and recalling that

$$F(\lambda_k) = \mathcal{L}(\mu_{[0,T]}^{\lambda_k}, \lambda_k) = H\big(\mu_{[0,T]}^{\lambda_k}\big|\nu_{[0,T]}\big) + \int_0^T \lambda_k(t)\Psi(\mu_t^{\lambda_k})\mathrm{d}t,$$

the convergence of λ_k towards $\overline{\lambda}$, the continuity of Ψ and the lower semi-continuity of H imply in turn (up to a sub-sequence) that

$$\mathcal{L}(\mu_{[0,T]}^\infty, \overline{\lambda}) \leq \liminf_{k\to+\infty} \mathcal{L}(\mu_{[0,T]}^{\lambda_k}, \lambda_k) \leq \limsup_{k\to+\infty} F(\lambda_k) \leq F(\overline{\lambda}),$$

the last step using the fact F is u.s.c. By uniqueness of $\overline{\mu}_{[0,T]}$, this gives $\mu^{\infty}_{[0,T]} = \overline{\mu}_{[0,T]}$ and then the convergence of the full sequence $(\mu^{\lambda_k}_{[0,T]})_{k\in\mathbb{N}}$ towards $\overline{\mu}_{[0,T]}$. Dividing now (24) by ε_k and sending k to infinity concludes that $\Psi(\overline{\mu}_{t_0}) \leq 0$, by continuity of Ψ. By continuity, this still holds at $t_0 = T$. The complementary slackness condition follows from the same argument applied to $\lambda_k = \overline{\lambda} - \min(\overline{\lambda}, \varepsilon_k)\mathbb{1}_{\overline{\lambda}>0}\overline{\lambda}$. □

Lemma 14 (Sufficient condition)**.** *A path measure $\tilde{\mu}_{[0,T]}$ of the shape* (17) *for some λ in $\mathcal{M}_+([0,T])$ is the unique minimiser $\mu^{\lambda}_{[0,T]}$ which realises $F(\lambda)$.*

Proof. Using the C^1-regularity of Ψ, the equivalent of (23) now reads for any $\mu_{[0,T]}$

$$H\big(\mu_{[0,T]}|\nu_{[0,T]}\big) + \int_0^T \Psi(\mu_t)\lambda(\mathrm{d}t) = H\big(\tilde{\mu}_{[0,T]}|\nu_{[0,T]}\big) + \int_0^T \Psi(\tilde{\mu}_t)\lambda(\mathrm{d}t)$$

$$+ H\big(\mu_{[0,T]}|\tilde{\mu}_{[0,T]}\big) + \int_0^T \int_0^1 \langle \mu_t - \tilde{\mu}_t, \frac{\delta\Psi}{\delta\mu}((1-r)\mu_t + r\tilde{\mu}_t) - \frac{\delta\Psi}{\delta\mu}(\tilde{\mu}_t)\rangle \mathrm{d}r\lambda(\mathrm{d}t).$$

From the convexity of Ψ, one easily gets that Ψ is above its tangents, so that:

$$\Psi(\mu_t) \geq \Psi(\tilde{\mu}_t) + \langle \mu_t - \tilde{\mu}_t, \frac{\delta\Psi}{\delta\mu}(\tilde{\mu}_t)\rangle.$$

Subtracting $\Psi(\tilde{\mu}_t)$, the C^1 regularity yields

$$\int_0^1 \langle \mu_t - \tilde{\mu}_t, \frac{\delta\Psi}{\delta\mu}((1-r)\mu_t + r\tilde{\mu}_t)\rangle \mathrm{d}r \geq \langle \mu_t - \tilde{\mu}_t, \frac{\delta\Psi}{\delta\mu}(\tilde{\mu}_t)\rangle,$$

so that the double integral two lines above is non-negative, and one can now conclude as in Lemma 10. □

Let us go back to the case $\nu_{[0,T]}$ is the law of the diffusion (1) with $\varepsilon = 1$. Using a Doob transform, $\mathrm{Law}(\vec{X}^N_{[0,T]}|T < \tau^N_{\Psi})$ is the law of the interacting particle system $\vec{Z}^N_{[0,T]} = (Z^{1,N}_{[0,T]}, \dots, Z^{N,N}_{[0,T]})$ described by

$$\mathrm{d}Z^{i,N}_t = b(Z^{i,N}_t)\mathrm{d}t + \partial_x \ln \mathbb{P}_{\vec{Z}^N_t}\big(\tau^N_{\psi} > T - t\big)\mathrm{d}t + \mathrm{d}B^{i,N}_t,$$

the function $\vec{x}^N \mapsto \mathbb{P}_{\vec{x}^N}\big(\tau^N_{\psi} > t\big)$ being the probability that $\tau^N_{\psi} > t$ knowing that $\vec{X}^N_0 = \vec{x}^N$. Particles are no more independent, they interact in an exchangeable way through this additional drift term. Theorem 6 thus appears as a *propagation of chaos* [Szn91; BZ99; CD22a; CD22b] result: it shows that correlations between particles disappear in the $N \to +\infty$ limit, particles being asymptotically independent and $\overline{\mu}_{[0,T]}$-distributed. Let us briefly describe how this can be interpreted in terms of a stochastic control problem. To make things lighter, let us assume that $b = 0$ and that $\nu_{[0,T]}$ is the Wiener measure on $C([0,T,\mathbb{R}^d)$ with an arbitrary initial distribution ν_0 at $t = 0$. Consider the controlled dynamics

$$\mathrm{d}X^{\alpha}_t = \alpha_t(X^{\alpha}_t)\mathrm{d}t + \mathrm{d}B_t, \tag{25}$$

for some deterministic Markov policy $\alpha_{[0,T]} \in L^2_{\mathrm{Law}(X^{\alpha}_t)\otimes\mathrm{d}t}([0,T] \times \mathbb{R}^d)$. A careful analysis of the proofs in [DG87] shows that the functional S defined in (5) satisfies

$$S\big((\mathrm{Law}(X^{\alpha}_t))_{0\leq t\leq T}\big) = \inf_{\substack{\alpha_{[0,T]}\,/\,\mathrm{Law}(X^{\alpha}_0)=\nu_0\,,\\ \forall 0\leq t\leq T,\, \Psi(\mathrm{Law}(X^{\alpha}_t))\leq 0\,.}} \frac{1}{2}\int_0^T \mathbb{E}|\alpha_t(X^{\alpha}_t)|^2\mathrm{d}t. \tag{26}$$

This quadratic problem is non-standard because of the mean-field constraint, which acts on the law of the process itself (one could easily replace $\alpha_t(X_t)$ in (25)-(26) by any adapted process $(\alpha_t)_{0\leq t\leq T}$). The following result is proven in [Dau21; Dau20].

Theorem 15 (Structure of optimal curves). *If $\Psi(\nu_0) < 0$, under additional regularity assumptions on Ψ, the control problem* (25)-(26) *admits at least one solution, and for any such solution $(\mu_{[0,T]}, \alpha_{[0,T]})$, there exist $\varphi_{[0,T]}$ in $W^{1,\infty}([0,T] \times \mathbb{R}^d) \cap C([0,T], C^2(\mathbb{R}^d))$ and λ in $L^\infty([0,T], \mathbb{R}_+)$ such that*

$$\alpha_t(x) = -\partial_x \varphi_t(x),$$

solving the following mean-field game system

$$\begin{cases} -\partial\varphi_t + \frac{1}{2}|\partial_x\varphi_t|^2 - \partial_x^2\varphi_t = \lambda(t)\frac{\delta\Psi}{\delta\mu}(\mu_t), \\ \partial_t\mu_t - L^\star\mu_t - \mathrm{div}(\mu_t\partial_x\varphi_t) = 0, \\ \varphi_T = 0, \\ \mu_0 = \nu_0, \end{cases} \tag{27}$$

together with the complementary slackness condition $\lambda_t\langle\mu_t, \psi\rangle = 0$ for a.e. t. Reciprocally, any solution of (27) *for an adequate multiplier satisfying the above condition is a solution of the control problem* (25)-(26).

The link with Theorem 9 is done by the Girsanov transform, which shows for such a solution $(\mu_{[0,T]}, \alpha_{[0,T]})$ that the pathwise law of the stochastic process $X^\alpha_{[0,T]}$ has the shape

$$\mathrm{d}\mu_{[0,T]}\big(x_{[0,T]}\big) = Z_\alpha^{-1} \exp\Big[-\int_0^T \lambda(t)\frac{\delta\Psi}{\delta\mu}(\mu_t, x_t)\mathrm{d}t\Big]\mathrm{d}\nu_{[0,T]}\big(x_{[0,T]}\big),$$

so that $\mu_{[0,T]} = \overline{\mu}_{[0,T]}$ using the sufficient condition in Theorem 9. This moreover proves the uniqueness in pathwise law for the solution of (26). Our current work aims to extend these results to non-convex Ψ, and to study the $T \to +\infty$ limit of the system, which is expected to correspond to the ergodic stochastic control problem

$$\inf_{\substack{\alpha_{[0,T]}\,/\,\mathrm{Law}(X_0^\alpha)=\nu_0\,, \\ \forall t\geq 0,\, \Psi(\mathrm{Law}(X_t^\alpha))\leq 0\,.}} \limsup_{T\to+\infty} \frac{1}{2T}\int_0^T \mathbb{E}|\alpha_t(X_t^\alpha)|^2\mathrm{d}t.$$

References

[BC08] A. Bain and D. Crisan. *Fundamentals of stochastic filtering*. Vol. 60. Springer Science & Business Media, 2008.

[Bar94] G. Barles. "Solutions de viscosité des équations de Hamilton-Jacobi". In: *Collection SMAI* (1994).

[BZ99] G. Ben Arous and O. Zeitouni. "Increasing propagation of chaos for mean field models". In: *Ann. Inst. Henri Poincaré Probab. Stat.* 35.1 (1999), pp. 85–102.

[BR07] V. I. Bogachev and M. A. S. Ruas. *Measure theory*. Vol. 1. Springer, 2007.

[BC18] A. Briani and P. Cardaliaguet. "Stable solutions in potential mean field game systems". In: *Nonlinear Differential Equations and Applications NoDEA* 25.1 (2018), pp. 1–26.

[CD+18] R. Carmona, F. Delarue, et al. *Probabilistic theory of mean field games with applications I-II*. Springer, 2018.

[CMC19] M. Caruel, P. Moireau, and D. Chapelle. "Stochastic modeling of chemical–mechanical coupling in striated muscles". In: *Biomechanics and modeling in mechanobiology* 18.3 (2019), pp. 563–587.

[CCK22] L.-P. Chaintron, M. Caruel, and F. Kimmig. “Modeling acto-myosin interaction: beyond the Huxley–Hill framework”. In: *preprint hal-03699263* (2022).

[CD22a] L.-P. Chaintron and A. Diez. “Propagation of chaos: A review of models, methods and applications. I. Models and methods”. In: *Kinetic and Related Models* 15.6 (2022), pp. 895–1015.

[CD22b] L.-P. Chaintron and A. Diez. “Propagation of chaos: A review of models, methods and applications. II. Applications”. In: *Kinetic and Related Models* 15.6 (2022), pp. 1017–1173.

[Cha+22] L.-P. Chaintron, Á. Mateos, L. González, and P. Moireau. “Mortensen Observer for a class of variational inequalities-Lost equivalence with stochastic filtering approaches”. In: *arXiv preprint arXiv:2205.02503* (2022).

[CC05] J. Clark and D. Crisan. “On a robust version of the integral representation formula of nonlinear filtering”. In: *Probability theory and related fields* 133.1 (2005), pp. 43–56.

[Csi84] I. Csiszár. “Sanov property, generalized I-projection and a conditional limit theorem”. In: *The Annals of Probability* (1984), pp. 768–793.

[Dau20] S. Daudin. “Optimal Control of Diffusion Processes with Terminal Constraint in Law”. In: *arXiv preprint arXiv:2012.10707* (2020).

[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.

[Dos77] H. Doss. “Liens entre équations différentielles stochastiques et ordinaires”. In: *Annales de l’IHP Probabilités et statistiques.* Vol. 13. 2. 1977, pp. 99–125.

[EI85] L. Evans and H. Ishii. “A PDE approach to some asymptotic problems concerning random differential equations with small noise intensities”. In: *Annales de l’Institut Henri Poincaré C, Analyse non linéaire* 2.1 (1985), pp. 1–20.

[Fle97] W. H. Fleming. “Deterministic nonlinear filtering”. In: *Annali della Scuola Normale Superiore di Pisa-Classe di Scienze* 25.3-4 (1997), pp. 435–454.

[FW98] M. I. Freidlin and A. D. Wentzell. “Random perturbations”. In: *Random perturbations of dynamical systems.* Springer, 1998, pp. 15–43.

[Hij84] O. Hijab. “Asymptotic Bayesian Estimation of a First Order Equation with Small Diffusion”. In: *The Annals of Probability* 12.3 (1984), pp. 890–902.

[JB88] M. R. James and J. S. Baras. “Nonlinear filtering and large deviations:a pde-control theoretic approach”. In: *Stochastics* 23.3 (1988), pp. 391–412.

[KS68] G. Kallianpur and C. Striebel. “Estimation of stochastic systems: Arbitrary system process with additive white noise observation errors”. In: *The Annals of Mathematical Statistics* 39.3 (1968), pp. 785–801.

[Mor68] R. E. Mortensen. “Maximum-likelihood recursive nonlinear filtering”. In: *Journal of Optimization Theory and Applications* 2.6 (1968), pp. 386–394.

[Nut21] M. Nutz. *Introduction to Entropic Optimal Transport.* 2021.

[Par78] E. Pardoux. “Stochastic partial differential equation for the density of the conditional law of a diffusion process with boundary”. In: *Stochastic analysis (Proceeings of International Conference, Northwestern University, Evanston, Ill., 1978).* 1978, pp. 239–269.

[Par82] E. Pardoux. “Équations du filtrage non linéaire de la prédiction et du lissage”. In: *Stochastics* 6.3-4 (1982), pp. 193–231.

[Sko61] A. V. Skorokhod. “Stochastic equations for diffusion processes in a bounded region”. In: *Theory of Probability & Its Applications* 6.3 (1961), pp. 264–274.

[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.

[Sus78] H. J. Sussmann. “On the gap between deterministic and stochastic ordinary differential equations”. In: *The Annals of Probability* (1978), pp. 19–41.

[Szn91] A.-S. Sznitman. “Topics in propagation of chaos”. In: *Éc. Été Probab. St.-Flour XIX—1989.* Springer, 1991, pp. 165–251.

[Zak69] M. Zakai. “On the optimal filtering of diffusion processes”. In: *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 11.3 (1969), pp. 230–243.