I. Restatement of the main result and its proof

We restate Theorem 1 below which now includes the pde formulation as well.

Theorem 1. *Consider the variational problem (11)-(14).*

(i) *For the probabilistic form (11) of the variational problem, the optimal control $U_t^* = e^{\alpha_t - \gamma_t} Y_t$, where the optimal trajectory $\{(X_t, Y_t)\}_{t \geq 0}$ evolves according to the Hamilton's odes:*

$$\frac{\mathrm{d}X_t}{\mathrm{d}t} = U_t^* = e^{\alpha_t - \gamma_t} Y_t, \quad X_0 \sim \rho_0 \tag{I.1a}$$

$$\frac{\mathrm{d}Y_t}{\mathrm{d}t} = -e^{\alpha_t + \beta_t + \gamma_t} \nabla_\rho \mathsf{F}(\rho_t)(X_t), \quad Y_0 = \nabla \phi_0(X_0) \tag{I.1b}$$

where ϕ_0 is a convex function, and $\rho_t = Law(X_t)$.

(ii) *For the pde form (14) of the variational problem, the optimal control is $u_t^* = e^{\alpha_t - \gamma_t} \nabla \phi_t(x)$, where the optimal trajectory $\{(\rho_t, \phi_t)\}_{t \geq 0}$ evolves according to the Hamilton's pdes:*

$$\frac{\partial \rho_t}{\partial t} = -\nabla \cdot (\rho_t \underbrace{e^{\alpha_t - \gamma_t} \nabla \phi_t}_{u_t^*}), \quad \textit{initial condn. } \rho_0 \tag{I.2a}$$

$$\frac{\partial \phi_t}{\partial t} = -e^{\alpha_t - \gamma_t} \frac{|\nabla \phi_t|^2}{2} - e^{\alpha_t + \gamma_t + \beta_t} \nabla_\rho \mathsf{F}(\rho) \tag{I.2b}$$

(iii) *The solutions of the two forms are equivalent in the following sense:*

$$Law(X_t) = \rho_t, \quad U_t = u_t(X_t), \quad Y_t = \nabla \phi_t(X_t)$$

(iv) *Suppose additionally that the functional F is displacement convex and ρ_∞ is its minimizer. Define*

$$V(t) = \frac{1}{2} \mathsf{E}(|X_t + e^{-\gamma_t} Y_t - T_{\rho_t}^{\rho_\infty}(X_t)|^2) + e^{\beta_t}(\mathsf{F}(\rho) - \mathsf{F}(\rho_\infty)) \tag{I.3}$$

where the map $T_{\rho_t}^{\rho_\infty} : \mathbb{R}^d \to \mathbb{R}^d$ is the optimal transport map from ρ_t to ρ_∞. Assume the dimension $d = 1$. Consequently, the following rate of convergence is obtained along the optimal trajectory

$$\mathsf{F}(\rho_t) - \mathsf{F}(\rho_\infty) \leq O(e^{-\beta_t}), \quad \forall t \geq 0$$

Proof. (i) The Hamiltonian function defined in (12) is equal to

$$\mathsf{H}(t, x, \rho, y, u) = y \cdot u - e^{\gamma_t - \alpha_t} \frac{1}{2} |u|^2 + e^{\alpha_t + \gamma_t \beta_t} \tilde{F}(\rho, x)$$

after inserting the formula for the Lagrangian. According to the maximum principle in probabilistic form for (mean-field) optimal control problems (see [1, Sec. 6.2.3]), the optimal control law $U_t^* = \arg\min_v \mathsf{H}(t, X_t, \rho_t, Y_t, v) = e^{\alpha_t - \gamma_t} Y_t$ and the Hamilton's equations are

$$\begin{aligned}
\frac{\mathrm{d}X_t}{\mathrm{d}t} &= +\nabla_y \mathsf{H}(t, X_t, \rho_t, Y_t, U_t^*) = U_t^* = e^{\alpha_t - \gamma_t} Y_t \\
\frac{\mathrm{d}Y_t}{\mathrm{d}t} &= -\nabla_x \mathsf{H}(t, X_t, \rho_t, Y_t, U_t^*) - \tilde{\mathsf{E}}[\nabla_\rho \mathsf{H}(t, \tilde{X}_t, \rho_t, \tilde{Y}_t, \tilde{U}_t^*)(X_t)]
\end{aligned}$$

where $\tilde{X}_t, \tilde{Y}_t, \tilde{U}_t^*$ are independent copies of X_t, Y_t, U_t^*. The derivatives

$$\nabla_x \mathsf{H}(t, x, \rho, y, u) = e^{\alpha_t + \beta_t + \gamma_t} \nabla_x \tilde{F}(\rho, x)$$

$$\nabla_\rho \mathsf{H}(t, x, \rho, y, u) = e^{\alpha_t + \beta_t + \gamma_t} \nabla_\rho \tilde{F}(\rho, x)$$

It follows that

$$\frac{\mathrm{d}Y_t}{\mathrm{d}t} = -e^{\alpha_t + \beta_t + \gamma_t} \left(\nabla_x \tilde{F}(\rho_t, X_t) + \tilde{\mathsf{E}}[\nabla_\rho \tilde{F}(\rho_t, \tilde{X}_t)(X_t)] \right) = -e^{\alpha_t + \beta_t + \gamma_t} \nabla_\rho \mathsf{F}(\rho)(X_t)$$

where we used the definition $\mathsf{F}(\rho) = \int \tilde{F}(x, \rho) \rho(x) \, \mathrm{d}x$ and the identity [1, Sec. 5.2.2 Example 3]

$$\nabla_\rho \mathsf{F}(\rho)(x) = \nabla_x \tilde{F}(\rho, x) + \int \nabla_\rho \tilde{F}(\rho, \tilde{x})(x) \rho(\tilde{x}) \, \mathrm{d}\tilde{x}$$

(ii) The Hamiltonian function defined in (15) is equal to

$$\mathscr{H}(t, \rho, \phi, u) = \int \left[\nabla \phi(x) \cdot u(x) - \frac{1}{2} e^{\gamma_t - \alpha_t} |u(x)|^2 \right] \rho(x) \, \mathrm{d}x + e^{\alpha_t + \gamma_t + \beta_t} \mathsf{F}(\rho)$$

after inserting the formula for the Lagrangian. According to the maximum principle for pde formulation of mean-field optimal control problems (see [1, Sec. 6.2.4]) the optimal control vector field is $u_t^* = \arg\min_v \mathscr{H}(t, \rho_t, \phi_t, v) = e^{\alpha_t - \gamma_t} \nabla \phi_t$ and the Hamilton's equations are:

$$\begin{aligned}
\frac{\partial \rho_t}{\partial t} &= +\frac{\partial \mathscr{H}}{\partial \phi}(t, \rho_t, \phi_t, u_t) = -\nabla \cdot (\rho_t \nabla u_t^*) \\
\frac{\partial \phi_t}{\partial t} &= -\frac{\partial \mathscr{H}}{\partial \rho}(t, \rho_t, \phi_t, u_t) = -(\nabla \phi \cdot u^* - e^{\gamma_t - \alpha_t} \frac{1}{2} |u_t^*|^2 + e^{\alpha_t + \gamma_t + \beta_t} \frac{\partial \mathsf{F}}{\partial \rho}(\rho_t))
\end{aligned}$$

inserting the formula $u_t^* = e^{\alpha_t - \gamma_t} \nabla \phi_t$ concludes the result.

(iii) Consider the (ρ_t, ϕ_t) defined from (I.2). The distribution ρ_t is identified with a stochastic process $\tilde{X}_t$ such that $\frac{d\tilde{X}_t}{dt} = e^{\alpha_t - \gamma_t}\nabla\phi_t(\tilde{X}_t)$ and $\mathrm{Law}(\tilde{X}_t) = \rho_t$. Then define $\tilde{Y}_t = \nabla\phi_t(\tilde{X}_t)$. Taking the time derivative shows that

$$\begin{aligned}\frac{d\tilde{Y}_t}{dt} &= \frac{d}{dt}\nabla\phi_t(\tilde{X}_t) = \nabla^2\phi_t(\tilde{X}_t)\frac{d\tilde{X}_t}{dt} + \nabla\frac{\partial\phi_t}{\partial t}(X_t)\\ &= e^{\alpha_t-\gamma_t}\nabla^2\phi_t(\tilde{X}_t)\nabla\phi_t(\tilde{X}_t) - e^{\alpha_t-\gamma_t}\nabla^2\phi_t(\tilde{X}_t)\nabla\phi_t(X_t) - e^{\alpha_t+\beta_t+\gamma_t}\nabla\frac{\partial \mathsf{F}}{\partial\rho}(\rho_t)(\tilde{X}_t)\\ &= -e^{\alpha_t+\beta_t+\gamma_t}\nabla\frac{\partial \mathsf{F}}{\partial\rho}(\rho_t)(\tilde{X}_t)\\ &= -e^{\alpha_t+\beta_t+\gamma_t}\nabla_\rho \mathsf{F}(\rho_t)(\tilde{X}_t)\end{aligned}$$

with the initial condition $\tilde{Y}_0 = \nabla\phi_0(\tilde{X}_0)$, where we used the identity $\nabla_x\frac{\partial \mathsf{F}}{\partial\rho}(\rho) = \nabla_\rho \mathsf{F}(\rho)$ [1, Prop. 5.48]. Therefore the equations for $\tilde{X}_t$ and $\tilde{Y}_t$ are identical. Hence one can identify (X_t, Y_t) with $(\tilde{X}_t, \tilde{Y}_t)$.

(iv) The energy functional

$$V(t) = \underbrace{\frac{1}{2}\mathsf{E}\left[|X_t + e^{-\gamma_t}Y_t - T_{\rho_t}^{\rho_\infty}(X_t)|^2\right]}_{\text{first term}} + \underbrace{e^{\beta_t}(\mathsf{F}(\rho) - \mathsf{F}(\rho_\infty))}_{\text{second term}}$$

Then the derivative of the first term is

$$\mathsf{E}\left[(X_t + e^{-\gamma_t}Y_t - T_{\rho_t}^{\rho_\infty}(X_t))\cdot(e^{\alpha_t-\gamma_t}Y_t - \dot{\gamma}_t e^{-\gamma_t}Y_t - e^{\alpha_t+\beta_t}\nabla_\rho \mathsf{F}(\rho_t)(X_t) + \xi(T_{\rho_t}^{\rho_\infty}(X_t)))\right]$$

where $\xi(T_{\rho_t}^{\rho_\infty}(X_t)) := \frac{d}{dt}T_{\rho_t}^{\rho_\infty}(X_t)$. Using the scaling condition $\dot{\gamma}_t = e^{\alpha_t}$ the derivative of the first term simplifies to

$$\mathsf{E}\left[(X_t + e^{-\gamma_t}Y_t - T_{\rho_t}^{\rho_\infty}(X_t))\cdot(-e^{\alpha_t+\beta_t}\nabla_\rho \mathsf{F}(\rho_t)(X_t) + \xi(T_{\rho_t}^{\rho_\infty}(X_t)))\right]$$

We claim that when the dimension $d = 1$, the expectation

$$\mathsf{E}[(X_t + e^{-\gamma_t}Y_t - T_{\rho_t}^{\rho_\infty}(X_t))\cdot\xi(T_{\rho_t}^{\rho_\infty}(X_t))] = 0 \tag{I.4}$$

We present the proof for the claim at the end. Assuming that the claim is true, the derivative of the first term simplifies to

$$\mathsf{E}\left[(X_t + e^{-\gamma_t}Y_t - T_{\rho_t}^{\rho_\infty}(X_t))\cdot(-e^{\alpha_t+\beta_t}\nabla_\rho F(\rho_t)(X_t))\right]$$

The derivative of the second term is

$$\begin{aligned}\frac{d}{dt}(\text{second term}) &= \dot{\beta}_t e^{\beta_t}(\mathsf{F}(\rho_t) - \mathsf{F}(\rho_\infty)) + e^{\beta_t}\frac{d}{dt}\mathsf{F}(\rho_t)\\ &= e^{\alpha_t+\beta_t}(\mathsf{F}(\rho_t) - \mathsf{F}(\rho_\infty)) + e^{\beta_t}\mathsf{E}[\nabla_\rho \mathsf{F}(\rho_t)(X_t)e^{\alpha_t-\gamma_t}Y_t]\end{aligned}$$

where we used the scaling condition $\dot{\beta}_t = e^{\alpha_t}$ and the chain-rule for the Wasserstein gradient [2, Ch. 10, E. Chain rule]. Adding the derivative of the first and second term yields:

$$\frac{\mathrm{d}V}{\mathrm{d}t}(t) = e^{\alpha_t+\beta_t}\left(\mathsf{F}(\rho_t) - \mathsf{F}(\rho_\infty) - \mathsf{E}\left[(X_t - T_{\rho_t}^{\rho_\infty}(X_t))\cdot\nabla_\rho \mathsf{F}(\rho_t)(X_t)\right]\right)$$

which is negative by variational inequality characterization of the displacement convex function $\mathsf{F}(\rho)$ [2, Eq. 10.1.7].

We now present the proof of the claim (I.4) under the assumption that $d = 1$. According to Brenier theorem [3], there exists a convex function ψ_t such that $T_{\rho_t}^{\rho_\infty}(x) = \nabla\psi_t(x)$ and $T_{\rho_\infty}^{\rho_t}(x) = \nabla\psi_t^\star(x)$ where $\psi_t^\star$ is the convex conjugate of ψ_t. Because ρ_∞ is the push-forward of ρ_t under the map $\nabla\psi_t$, we have

$$\mathsf{E}[g(\nabla\psi_t(X_t))] = \int g(x)\rho_\infty(x)\,\mathrm{d}x,$$

for all measurable functions g. Upon taking the derivative with respect to time,

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathsf{E}[g(\nabla\psi_t(X_t))] = \frac{\mathrm{d}}{\mathrm{d}t}\int g(x)\rho_\infty(x)\,\mathrm{d}x = 0$$

Hence by application of the dominated convergence theorem (DCT) and interchanging the expectation and the derivative,

$$\mathsf{E}[\frac{\mathrm{d}}{\mathrm{d}t}g(\nabla\psi_t(X_t))] = \mathsf{E}[\nabla g(\nabla\psi_t(X_t))\cdot\xi(\nabla\psi_t(X_t))] = 0 \tag{I.5}$$

Letting $g(x) = \psi^*(x) - e^{-\gamma_t}\int_{-\infty}^{x}\nabla\phi_t(\nabla\psi^\star(z))\,\mathrm{d}z - \frac{1}{2}|x|^2$ where ϕ_t is defined in part-(ii) of the theorem 1 concludes

$$\begin{aligned}0 = \mathsf{E}[\nabla g(\nabla\psi_t(X_t))\cdot\xi(\nabla\psi_t(X_t)))] &= \mathsf{E}[X_t - e^{-\gamma_t}\nabla\phi_t(X_t) - \nabla\psi_t(X_t))\cdot\xi(\nabla\psi_t(X_t)))]\\ &= \mathsf{E}[X_t - e^{-\gamma_t}Y_t - \nabla\psi_t(X_t))\cdot\xi(\nabla\psi_t(X_t)))]\end{aligned}$$

where we used $Y_t = \nabla\phi_t(X_t)$ from part-(iii) of Theorem 1. This concludes the proof of the claim. Note that the application of DCT in (I.5) follows from smoothness of $g(x)$ and assuming $T_{\rho_t}^{\rho_\infty}(x)$ is differentiable with respect to time. Showing $T_{\rho_t}^{\rho_\infty}(x)$ is differentiable with respect to time is technical out of the scope of this work.

□

II. Wasserstein gradient and Gâteaux derivative

This section contains definitions of the Wasserstein gradient and Gâteaux derivative [2], [1].

Let $\mathsf{F}:\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)\to\mathbb{R}$ be a (smooth) functional on the space of probability distributions.

Gâteaux derivative: The Gâteaux derivative of F at $\rho\in\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$ is a real-valued function on $\mathbb{R}^d$ denoted as $\frac{\partial \mathsf{F}}{\partial\rho}(\rho):\mathbb{R}^d\to\mathbb{R}$. It is defined as a function that satisfies the identity

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathsf{F}(\rho_t)\Big|_{t=0}=\int_{\mathbb{R}^d}\frac{\partial \mathsf{F}}{\partial\rho}(\rho)(x)(-\nabla\cdot(\rho(x)u(x)))\,\mathrm{d}x$$

for all path ρ_t in $\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$ such that $\frac{\partial\rho_t}{\partial t}=-\nabla\cdot(\rho_t u)$ with $\rho_0=\rho\in\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$.

Wasserstein gradient: The Wasserstein gradient of F at ρ is a vector-field on $\mathbb{R}^d$ denoted as $\nabla_\rho\mathsf{F}(\rho):\mathbb{R}^d\to\mathbb{R}^d$. It is defined as a vector-field that satisfies the identity

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathsf{F}(\rho_t)\Big|_{t=0}=\int_{\mathbb{R}^d}\nabla_\rho\mathsf{F}(\rho)(x)\cdot u(x)\;\rho(x)\,\mathrm{d}x$$

for all path ρ_t in $\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$ such that $\frac{\partial\rho_t}{\partial t}=-\nabla\cdot(\rho_t u)$ with $\rho_0=\rho\in\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$.

The two definitions imply the following relationship [1, Prop. 5.48]:

$$\nabla_\rho\mathsf{F}(\rho)(\cdot)=\nabla_x\frac{\partial\mathsf{F}}{\partial\rho}(\rho)(\cdot)$$

Example: Let $\mathsf{F}(\rho)=\int\log(\frac{\rho(x)}{\rho_\infty(x)})\rho(x)\,\mathrm{d}x$ be the relative entropy functional. Consider a path ρ_t in $\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$ such that $\frac{\partial\rho_t}{\partial t}=-\nabla\cdot(\rho_t u)$ with $\rho_0=\rho\in\mathscr{P}_{\text{ac},2}(\mathbb{R}^d)$. Then

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\mathsf{F}(\rho_t)&=\int\log(\frac{\rho_t(x)}{\rho_\infty(x)})\frac{\partial\rho_t}{\partial t}(x)\,\mathrm{d}x+\int\frac{\partial\rho_t}{\partial t}(x)\,\mathrm{d}x\\&=-\int\log(\frac{\rho_t(x)}{\rho_\infty(x)})\nabla\cdot(\rho_t(x)u(x))\,\mathrm{d}x\\&=\int\nabla_x\log(\frac{\rho_t(x)}{\rho_\infty(x)})\cdot u(x)\;\rho_t(x)\,\mathrm{d}x\end{aligned}$$

where the divergence theorem is used in the last step. The definitions of the Gâteaux derivative and Wasserstein gradient imply

$$\frac{\partial\mathsf{F}}{\partial\rho}(\rho)(x)=\log(\frac{\rho(x)}{\rho_\infty(x)})$$

$$\nabla_\rho\mathsf{F}(\rho)(x)=\nabla_x\log(\frac{\rho(x)}{\rho_\infty(x)})$$

III. RELATIONSHIP WITH THE UNDER-DAMPED LANGEVIN EQUATION

A basic form of the under-damped (or second order) Langevin equation is given in [4]

$$\begin{aligned} \mathrm{d}X_t &= v_t\,\mathrm{d}t \\ \mathrm{d}v_t &= -\gamma v_t\,\mathrm{d}t - \nabla f(X_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}B_t \end{aligned} \tag{III.1}$$

where $\{B_t\}_{t\geq 0}$ is the standard Brownian motion.

Consider next, the the accelerated flow (19). Denote $v_t := e^{\alpha_t - \gamma_t} Y_t$. Then, with an appropriate choice of scaling parameters (e.g. $\alpha_t = 0$, $\beta_t = 0$ and $\gamma_t = -\gamma t$):

$$\begin{aligned} \mathrm{d}X_t &= v_t\,\mathrm{d}t \\ \mathrm{d}v_t &= -\gamma v_t\,\mathrm{d}t - \nabla f(X_t)\,\mathrm{d}t - \nabla_x \log(\rho_t(X_t)) \end{aligned} \tag{III.2}$$

The scaling parameters are chosen here for the sake of comparison and do not satisfy the ideal scaling conditions of [5].

The sdes (III.1) and (III.2) are similar except that the stochastic term $\sqrt{2}\,\mathrm{d}B_t$ in (III.1) is replaced with a deterministic term $-\nabla_x \log(\rho_t(X_t))$ in (III.2). Because of this difference, the resulting distributions are different. Let $p_t(x,v)$ denote the joint distribution on (X_t, v_t) of (III.1) and let $q_t(x,v)$ denote the joint distribution on (X_t, v_t) of (III.2). Then the corresponding Fokker-Planck equations are:

$$\begin{aligned} \frac{\partial p}{\partial t}(x,v) &= -\nabla_x \cdot (p_t(x,v)v) + \nabla_v \cdot (p_t(x,v)(\gamma v + \nabla f(x))) + \Delta_v p_t(x,v) \\ \frac{\partial q}{\partial t}(x,v) &= -\nabla_x \cdot (q_t(x,v)v) + \nabla_v \cdot (q_t(x,v)(\gamma v + \nabla f(x))) + \nabla_v \cdot (q_t(x,y)\nabla_x \log(\rho_t(x))) \end{aligned}$$

where $\rho_t(x) = \int q_t(x,v)\,\mathrm{d}v$ is the marginal of $q_t(x,y)$ on x. The final term in the Fokker-Planck equations are clearly different. The joint distributions are different as well.

The situation is in contrast to the first order Langevin equation, where the stochastic term $\sqrt{2}\,\mathrm{d}B_t$ and $-\nabla \log(\rho_t(X_t))$ are equivalent, in the sense that the resulting distributions have the same marginal distribution as a function of time. To illustrate this point, consider the following two forms of the Langevin equation:

$$\mathrm{d}X_t = -\nabla f(X_t)\,\mathrm{d}t + \sqrt{2}\,\mathrm{d}B_t \tag{III.3}$$

$$\mathrm{d}X_t = -\nabla f(X_t)\,\mathrm{d}t - \nabla \log(\rho_t(X_t)) \tag{III.4}$$

Let $p_t(x)$ denote the distribution of X_t of (III.3) and let $q_t(x)$ denote the distribution of X_t of (III.4). The corresponding Fokker-Planck equations are as follows

$$\begin{aligned}
\frac{\partial p}{\partial t}(x) &= -\nabla\cdot(p_t(x)\nabla f(x)) + \Delta p_t(x) \\
\frac{\partial q}{\partial t}(x) &= -\nabla\cdot(q_t(x)\nabla f(x)) + \nabla\cdot(q_t(x)\nabla\log(\rho_t(x))) \\
&= -\nabla\cdot(q_t(x)\nabla f(x)) + \nabla\cdot(q_t(x)\nabla\log(q_t(x))) \\
&= -\nabla\cdot(q_t(x)\nabla f(x)) + \Delta q_t(x)
\end{aligned}$$

where we used $\rho_t(x) = q_t(x)$. In particular, this implies that the marginal probability distribution of the stochastic process X_t are the same for first order Langevin sde (III.3) and (III.4) .

REFERENCES

[1] R. Carmona and F. Delarue, Probabilistic Theory of Mean Field Games with Applications I-II. Springer, 2017.

[2] L. Ambrosio, N. Gigli, and G. Savaré, Gradient flows: in metric spaces and in the space of probability measures. Springer Science & Business Media, 2008.

[3] C. Villani, Topics in optimal transportation. American Mathematical Soc., 2003, no. 58.

[4] X. Cheng, N. S. Chatterji, P. L. Bartlett, and M. I. Jordan, "Underdamped langevin mcmc: A non-asymptotic analysis," arXiv preprint arXiv:1707.03663, 2017.

[5] A. Wibisono, A. C. Wilson, and M. I. Jordan, "A variational perspective on accelerated methods in optimization," Proceedings of the National Academy of Sciences, p. 201614734, 2016.