Sinkhorn Flow as Mirror Flow: A Continuous-Time Framework for Generalizing the Sinkhorn Algorithm

Mohammad Reza Karimi* ETH Zürich

Ya-Ping Hsieh* ETH Zürich

Andreas Krause ETH Zürich

Abstract

Many problems in machine learning can be formulated as solving entropy-regularized optimal transport on the space of probability measures. The canonical approach involves the *Sinkhorn iterates*, renowned for their rich mathematical properties. Recently, the Sinkhorn algorithm has been recast within the *mirror descent* framework, thus benefiting from classical optimization theory insights. Here, we build upon this result by introducing a *continuous-time* analogue of the Sinkhorn algorithm. This perspective allows us to derive novel variants of Sinkhorn schemes that are robust to noise and bias. Moreover, our continuous-time dynamics offers a unified perspective on several recently discovered dynamics in machine learning and mathematics, such as the "Wasserstein mirror flow" of Deb et al. (2023) or the "mean-field Schrödinger equation" of Claisse et al. (2023).

1 INTRODUCTION

Many modern machine learning tasks can be reframed as solving an entropy-regularized optimal transport (OT) problem over the space of probability measures. One particularly noteworthy instance that has attracted significant attention is the *Schrödinger bridge* (SB) problem, whose primary objective is to dynamically transform a given measure into another measure. Consequently, SB has found widespread application in diverse domains that require an understanding of complex continuous-time systems, spanning applications such as sampling (Bernton et al., 2019; Huang et al., 2021), generative modeling (Chen et al., 2022; De Bortoli et al., 2021; Wang et al., 2021), molecular biology (Holdijk et al., 2022), single-cell dynamics (Bunne et al., 2023; Pariset et al., 2023; Somnath et al., 2023), and mean-field games (Liu et al., 2022).

The canonical approach for solving entropy-regularized OT is the *Sinkhorn* algorithm (Sinkhorn and Knopp, 1967) or the closely related *Iterative Proportional Fitting* (IPF) procedure (Chen et al., 2021; Fortet, 1940; Kullback, 1968). Traditionally, these algorithms have been framed as an *alternating projection* procedure and extensively studied in this regard (Chen et al., 2016; Cuturi, 2013; Ghosal and Nutz, 2022; Peyré and Cuturi, 2019). Conversely, recent research (Ballu and Berthet, 2023; Mensch and Peyré, 2020; Mishchenko, 2019) sheds new light on Sinkhorn for discrete probability distributions by associating it to the *mirror descent* (MD) scheme (Beck and Teboulle, 2003; Nemirovsky and Yudin, 1983), thereby opening up new avenues for understanding Sinkhorn through the lens of classical optimization theory. Furthermore, Aubin-Frankowski et al. (2022); Léger (2021) generalize this insight to the space of continuous probability measures. To be more precise, these studies reveal that the Sinkhorn iterates can be perceived as MD steps, specifically implemented *with step-size 1.*

Here, we advance this mirror descent perspective by introducing a novel, continuous-time variant of the Sinkhorn algorithm on the space of probability measures. Our objectives are twofold. Firstly, we deepen our comprehension of Sinkhorn iterates by emphasizing that the key components for establishing convergence are not limited to the previous focus of the specific step-size 1, but are inherently linked to the choice of the mirror map, objective function, and constraints. Building upon this insight, we harness stochastic MD analysis to yield novel Sinkhorn variations that, unlike the traditional approach, maintain convergence even in the presence of noise and bias. Secondly, we demonstrate that our continuous-time dynamics opens up a

*Equal contributions.

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238.

unified perspective on various intriguing dynamics recently explored in machine learning and mathematics.

To summarize, we make the following contributions:

- We derive a specific formula for a discrete-time mirror descent scheme that encompasses Sinkhorn iterates as a special case by setting the step-size to 1. By driving the step-sizes to 0, we obtain a novel continuous-time version of the Sinkhorn iterates, comprising a pair of "primal" and "dual" systems referred to as the Sinkhorn$_\varepsilon$ and Schrödinger$_\varepsilon$ flows, respectively.
- Importantly, both our novel discrete-time and continuous-time schemes preserve the essential properties of Sinkhorn needed for ensuring its convergence. This insight deepens our understanding by highlighting that Sinkhorn's convergence is primarily a result of the mirror descent algorithm with a specific mirror map, as opposed to the conventional interpretation centered around alternating projection. Leveraging this fresh perspective, we further devise new schemes accompanied by robust asymptotic and non-asymptotic guarantees.
- With the aid of Otto calculus (Otto, 2001), we demonstrate that our dynamics formally encompass the "Wasserstein Mirror Flow" introduced by Deb et al. (2023) as well as the evolution proposed by Claisse et al. (2023). The former is motivated by modeling the behavior of Sinkhorn for *unregularized* optimal transport (i.e., when $\varepsilon \to 0$ in (OT_ε) below), while the latter seeks to relax the conditions for establishing the convergence of neural network training in the mean-field limit. Until now, these dynamics were treated in isolation, primarily due to the lack of a continuous-time mirror descent perspective.
- We contextualize these findings in the Schrödinger bridge setting and provide a mirror descent interpretation of the Iterative Proportional Fitting procedure widely adopted in the machine learning community (Chen et al., 2022; De Bortoli et al., 2021; Vargas et al., 2021). Additionally, we establish that these new mirror descent iterations can be expressed as stochastic differential equations with an explicit drift formula.

2 BACKGROUND ON MIRROR DESCENT

In this section, we revisit the fundamental components of the classical mirror descent (MD) scheme (Beck and Teboulle, 2003; Nemirovsky and Yudin, 1983).

§ MD as a "minimizing movement" scheme. Let $F : \mathbb{R}^d \to \mathbb{R}$ be a differentiable objective function, $\varphi : \mathbb{R}^d \to \mathbb{R}$ be strictly convex and differentiable called the *Bregman potential*, and $\mathcal{C} \subseteq \mathbb{R}^d$ be a convex constraint. The MD algorithm aims to find $\min_{x\in\mathcal{C}} F(x)$ by following the iterations

$$x_{n+1} = \arg\min_{x\in\mathcal{C}} \left\{ \langle \nabla F(x_n), x - x_n\rangle + \frac{D_\varphi(x \,\|\, x_n)}{\gamma_n} \right\}. \tag{1}$$

Here, γ_n is a sequence of step-sizes, and $D_\varphi(x \,\|\, x_n)$ is the Bregman divergence associated with φ:

$$D_\varphi(x' \,\|\, x) \coloneqq \varphi(x') - \varphi(x) - \langle \nabla\varphi(x), x' - x\rangle. \tag{2}$$

This is the *minimizing movement* interpretation of MD: At each iteration, linearize the objective F and minimize it while staying "close" to the previous iterate, where the measure of closeness is determined by the Bregman divergence.

§ The dual perspective of MD. A particularly insightful approach for studying MD employs the concept of *convex duality* (Rockafellar, 1997). Recall that for a convex set $\mathcal{C}$, the convex indicator function $\mathrm{I}_\mathcal{C}$ is defined as $\mathrm{I}_\mathcal{C}(x) = 0$ if $x \in \mathcal{C}$ and $+\infty$ otherwise. Let φ^* be the *Fenchel conjugate* of $\varphi + \mathrm{I}_\mathcal{C}$:

$$\begin{aligned} \varphi^*(y) &= \sup_{x\in\mathcal{X}} \{\langle x, y\rangle - (\varphi + \mathrm{I}_\mathcal{C})(x)\} \\ &= \sup_{x\in\mathcal{C}} \{\langle x, y\rangle - \varphi(x)\}. \end{aligned} \tag{3}$$

As $\varphi + \mathrm{I}_\mathcal{C}$ is strictly convex, it holds that φ^* is essentially differentiable (Rockafellar, 1997), and the Danskin's theorem implies

$$\nabla\varphi^*(y) = \arg\max_{x\in\mathcal{C}} \{\langle x, y\rangle - \varphi(x)\}.$$

Thus, $\nabla\varphi^*(\nabla\varphi(x')) = \arg\min_{x\in\mathcal{C}} D_\varphi(x \,\|\, x')$. In particular, for all $x \in \mathcal{C}$, $\nabla\varphi^*(\nabla\varphi(x)) = x$.

We shall call any $x \in \mathcal{C}$ a *primal point* and any $y \in \mathrm{dom}(\nabla\varphi^*)$ a *dual point*. Notice that a dual point y uniquely identifies a primal point $x = \nabla\varphi^*(y)$, but several dual points might correspond to the same primal point x, one of which is $\nabla\varphi(x)$.[1]

Now, let $y_0 = \nabla\varphi(x_0)$ and consider the following dual iterations:

$$\begin{cases} y_{n+1} = y_n - \gamma_n \nabla F(x_n), \\ x_{n+1} = \nabla\varphi^*(y_{n+1}), \end{cases} \tag{4}$$

which can also be written solely in terms of the dual variables as

$$y_{n+1} = y_n - \gamma_n \nabla F(\nabla\varphi^*(y_n)).$$

[1] For a fixed primal point $\bar{x} \in \mathcal{C}$, any $x \in \mathbb{R}^d$ such that $\arg\min_{x'\in\mathcal{C}} D_\varphi(x' \,\|\, x) = \bar{x}$ satisfies $\nabla\varphi^*(\nabla\varphi(x)) = \bar{x}$.

Under certain mild conditions which hold in our setting, (4) coincides with (1) (Beck and Teboulle, 2003). Driving the step-size $\gamma_n \to 0$, we obtain the *continuous-time* limit of the dual iterations (Krichene et al., 2015; Tzen et al., 2023), called the *mirror flow*:

$$\begin{cases} \frac{\mathrm{d}}{\mathrm{d}t} y(t) = -\nabla F(x(t)), \\ x(t) = \nabla\varphi^*(y(t)). \end{cases} \tag{5}$$

Alternatively, we can write the mirror flow as

$$\frac{\mathrm{d}}{\mathrm{d}t} y(t) = -(\nabla F \circ \nabla\varphi^*)(y(t)). \tag{6}$$

§ MD on the space of probability measures. While the discussion has thus far been limited to the Euclidean spaces, the extension of MD to the infinite-dimensional space of *probability measures* presents no conceptual difficulty: One can simply replace the inner product in (2) with the *duality pairing* $\langle p, f\rangle \coloneqq \mathbb{E}_p f$ (Halmos, 2013), and the gradient operator with the *first variation* as defined in Appendix A (Bauschke et al., 2003; Hsieh et al., 2019; Léger, 2021).

3 CONTINUOUS-TIME SINKHORN FLOWS

This section constitutes our core contribution: The introduction of a pair of "primal" and "dual" dynamics, which can be viewed as the analogue of the Sinkhorn algorithm in continuous time. Our derivation builds upon a well-established MD interpretation of Sinkhorn, reviewed in detail in Section 3.1. In Section 3.2, we present a simple yet pivotal extension of the MD interpretation, serving as the foundation for our main result in Section 3.3. We then showcase in Sections 3.4–3.5 how our novel dynamics allows unifying existing dynamics as well as motivating new schemes with advantageous properties compared to Sinkhorn. All proofs can be found in Appendix B.

3.1 Background: Entropic Optimal Transport and the Sinkhorn Algorithm

We first recall the central properties of the entropy-regularized OT; the materials are classical, and presented e.g., by Peyré and Cuturi (2019). Let μ and ν be two given probability measures on $\mathcal{X}$ and $\mathcal{Y}$, respectively. Consider a cost function $c : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ and a regularization parameter $\varepsilon > 0$. The *entropy-regularized OT* is the minimization problem

$$\mathrm{OT}_\varepsilon(\mu,\nu) \coloneqq \min_{\pi\in\Gamma(\mu,\nu)} \mathbb{E}_\pi[c] + \varepsilon H(\pi \,\|\, \mu\otimes\nu) \tag{OT_ε}$$

where $\Gamma(\mu,\nu)$ is the set of all couplings between μ and ν, and $H(\cdot\,\|\,\cdot)$ is the relative entropy. One can alternatively rewrite (OT_ε) as the *static Schrödinger problem* (see Section 4 for an explanation of the terminology):

$$\min_{\pi\in\Gamma(\mu,\nu)} H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}}), \tag{7}$$

where the reference measure $\pi_\varepsilon^{\mathrm{ref}}$ is defined as $\mathrm{d}\pi_\varepsilon^{\mathrm{ref}} \propto \exp(-c/\varepsilon)\,\mathrm{d}(\mu\otimes\nu)$ and encodes all the information about ε and the cost c. Without loss of generality, we assume that c is normalized s.t. $\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}/\mathrm{d}(\mu\otimes\nu) = \exp(-c/\varepsilon)$.

The optimal solution $\pi_\varepsilon^{\mathrm{opt}}$ of (OT_ε) admits the following dual representation: There exists potential functions $f : \mathcal{X} \to \mathbb{R}$ and $g : \mathcal{Y} \to \mathbb{R}$, unique up to constants, such that

$$\mathrm{d}\pi_\varepsilon^{\mathrm{opt}} = e^{f\oplus g - \frac{c}{\varepsilon}}\,\mathrm{d}(\mu\otimes\nu) = \exp(f\oplus g)\,\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}. \tag{8}$$

Here, we use the notation $(f\oplus g)(x,y) = f(x)+g(y)$, and call f and g the *Schrödinger potentials* of $\pi_\varepsilon^{\mathrm{opt}}$.[2] Moreover, the reverse direction also holds: If a coupling $\pi\in\Gamma(\mu,\nu)$ has the form (8), then it is the optimal solution of $\mathrm{OT}_\varepsilon(\mu,\nu)$.

§ The Sinkhorn Algorithm. A popular method for solving (OT_ε) is the Sinkhorn algorithm (Cuturi, 2013; Sinkhorn and Knopp, 1967): Starting from $\pi^0 \coloneqq \pi_\varepsilon^{\mathrm{ref}}$, the algorithm iterates as

$$\begin{aligned} \pi^{n+1/2} &\coloneqq \arg\min\{H(\pi \,\|\, \pi^n) : \pi_{\mathcal{Y}} = \nu\}, \\ \pi^{n+1} &\coloneqq \arg\min\{H(\pi \,\|\, \pi^{n+1/2}) : \pi_{\mathcal{X}} = \mu\}. \end{aligned} \tag{Sink_1}$$

Here, $\pi_{\mathcal{X}}$ denotes the $\mathcal{X}$-marginal of π. We use the notation $\pi^{n+1} = \mathrm{S}_1[\pi^n]$ to represent a full (Sink_1) iteration. A key attribute of the Sinkhorn algorithm is that all the information concerning the cost c and the regularization parameter ε is embedded in the initialization of π^0; the operator S_1 itself is independent of c and ε.

The special structure of algorithm (Sink_1) guarantees that the iterations π^n admit the form $\mathrm{d}\pi^n = \exp(f_\varepsilon^n \oplus g_\varepsilon^n)\,\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}$ for some potentials $f_\varepsilon^n, g_\varepsilon^n$. Furthermore, since the $\mathcal{X}$-marginal of each successive π^n is always μ, we can determine f_ε^n from g_ε^n as:

$$f_\varepsilon^n(x) = -\log\int \exp\!\left(g_\varepsilon^n(y) - \frac{c(x,y)}{\varepsilon}\right)\nu(\mathrm{d}y). \tag{9}$$

This also implies that π^n can be recovered solely from g_ε^n. In fact, in Lemma 3 below, we will demonstrate how to retrieve π^n from g^n through the application of the *dual* mirror map.

§ Mirror Descent interpretation. It is recently shown in (Aubin-Frankowski et al., 2022) that the

[2] Some authors (such as Nutz and Wiesel, 2022) prefer writing $\mathrm{d}\pi_\varepsilon^{\mathrm{opt}}$ as $\exp\big(\frac{1}{\varepsilon}(f\oplus g - c)\big)\,\mathrm{d}(\mu\otimes\nu)$, and call these f and g the Schrödinger potentials.

Sinkhorn algorithm can be viewed as an MD iteration in the space of probability measures. Specifically, by defining the objective function $F(\pi) \coloneqq H(\pi_{\mathcal{Y}} \,\|\, \nu)$, the Bregman potential $\varphi(\pi) \coloneqq H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}})$, and the constraint set $\mathcal{C} \coloneqq \{\pi : \pi_{\mathcal{X}} = \mu\}$, the Sinkhorn update ($\mathrm{Sink}_1$) corresponds to the MD update

$$\pi^{n+1} = \arg\min_{\pi\in\mathcal{C}} \{\langle \delta F(\pi^n), \pi - \pi^n\rangle + D_\varphi(\pi \,\|\, \pi^n)\}, \tag{10}$$

where $\delta F(\pi)$ is the first variation of F at π (see Appendix A for the related background).

3.2 Sinkhorn with *Arbitrary* Step-sizes

Comparing (10) with the MD update rule (1), we see that the classical Sinkhorn corresponds to MD *with constant step-size 1*. However, once the connection to MD is established, we can use arbitrary step-sizes γ_n to get the γ-Sinkhorn iteration, defined as follows:

Definition 1 (γ-Sinkhorn iteration)**.** Let $F(\pi) = H(\pi_{\mathcal{Y}} \,\|\, \nu)$, $\varphi(\pi) = H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}})$, and $\mathcal{C} = \{\pi : \pi_{\mathcal{X}} = \mu\}$. Starting from $\pi_\gamma^0 \coloneqq \pi_\varepsilon^{\mathrm{ref}}$, we define the iterates π_γ^n as

$$\pi_\gamma^{n+1} = \arg\min_{\pi\in\mathcal{C}} \left\{ \langle \delta F(\pi_\gamma^n), \pi - \pi_\gamma^n\rangle + \frac{D_\varphi(\pi \,\|\, \pi_\gamma^n)}{\gamma_n} \right\}$$

and write $\pi_\gamma^{n+1} = \mathrm{S}_{\gamma_n}[\pi_\gamma^n]$.

The following lemma is a simple observation that lies at the heart of all forthcoming derivations:

Lemma 1. *The γ-Sinkhorn iterates π_γ^n defined as in Definition 1 correspond to the update rule:*

$$\begin{aligned} \pi_\gamma^{n+1/2} &\coloneqq \arg\min\{H(\pi \,\|\, \pi_\gamma^n) : \pi_{\mathcal{Y}} = \nu\}, \\ \pi_\gamma^{n+1} &\coloneqq \arg\min\{\gamma_n H(\pi \,\|\, \pi_\gamma^{n+1/2}) \\ &\qquad + (1-\gamma_n) H(\pi \,\|\, \pi_\gamma^n) : \pi_{\mathcal{X}} = \mu\}. \end{aligned} \tag{Sink_γ}$$

Observe that in (Sink_γ) the half-iteration updates are exactly the same as in the classical Sinkhorn (Sink_1). However, unlike (Sink_1), integer iterations π_γ^{n+1}'s are now computed based on both $\pi_\gamma^{n+1/2}$ *and* π_γ^n, hence losing the interpretation of being a KL-projection step from $\pi_\gamma^{n+1/2}$. Nevertheless, the next lemma establishes that *the existence of Schrödinger potentials is retained* for (Sink_γ):

Lemma 2. *The γ-Sinkhorn iterates π_γ^n in (Sink_γ) admit the representation*

$$\mathrm{d}\pi_\gamma^n = \exp\big(f_\gamma^n \oplus g_\gamma^n\big)\, \mathrm{d}\pi_\varepsilon^{\mathrm{ref}}. \tag{11}$$

Moreover, the potentials g_γ^n satisfy the recursion

$$g_\gamma^{n+1} = g_\gamma^n - \gamma_n \log \frac{\mathrm{d}(\pi_\gamma^n)_{\mathcal{Y}}}{\mathrm{d}\nu},$$

and f_γ^{n+1} is computed from g_γ^{n+1} as in (9).

In Section 3.4, we will establish the convergence of the generalized iterates in (Sink_γ). In essence, this means that the dual representation provided in (11), which arises from the choice of the objective function F, the Bregman potential φ, and the constraint set $\mathcal{C}$ in Definition 1, is sufficient for ensuring convergence without resorting to the conventional alternating projection interpretation.

3.3 Sinkhorn$_\varepsilon$ and Schrödinger$_\varepsilon$ Flows

Analyzing an algorithm's continuous-time limits often provides a more manageable analytical perspective than for its discrete-time counterpart. Further, continuous-time dynamics are beneficial for implementing *stochastic approximation* techniques (see Section 3.4), which are essential for determining the algorithm's convergence. The following proposition characterizes the limiting behavior of operator S_γ as $\gamma \to 0$.

Proposition 1. *Fix a coupling $\pi \in \{\pi : \pi_{\mathcal{X}} = \mu\}$. For any $\gamma > 0$, let $\pi^\gamma = \mathrm{S}_\gamma[\pi]$. Then,*

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0} \log \pi^\gamma(x,y) = -\log \frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}(y) + \mathbb{E}_{\pi(\cdot|x)}\left[\log \frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}\right].$$

Moreover, if $\mathrm{d}\pi = \exp(f \oplus g)\, \mathrm{d}\pi_\varepsilon^{\mathrm{ref}}$, then for all $\gamma > 0$, $\mathrm{d}\pi^\gamma = \exp(f^\gamma \oplus g^\gamma)\, \mathrm{d}\pi_\varepsilon^{\mathrm{ref}}$, and f^γ and g^γ satisfy

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0} g^\gamma(y) &= -\log \frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}(y), \\ \frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0} f^\gamma(x) &= \mathbb{E}_{\pi(\cdot|x)}\left[\log \frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}\right]. \end{aligned}$$

In view of Proposition 1, we are now ready to define the *Sinkhorn$_\varepsilon$* and the *Schrödinger$_\varepsilon$ flows.*

Definition 2. Consider the set of all joint distributions on $\mathcal{X} \times \mathcal{Y}$ that solve (OT_ε) with cost function c for their own marginals:

$$\Pi_{c,\varepsilon} \coloneqq \{\pi : \pi \text{ solves } \mathrm{OT}_\varepsilon(\pi_{\mathcal{X}}, \pi_{\mathcal{Y}})\}. \tag{12}$$

For any $\pi_\varepsilon^0 \in \{\pi : \pi_{\mathcal{X}} = \mu\} \cap \Pi_{c,\varepsilon}$, we construct a curve $(\pi_\varepsilon^t)_{t\geq 0}$ whose velocity is determined by

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t} \log \frac{\mathrm{d}\pi_\varepsilon^t}{\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}}(x,y) \\ &\qquad = -\log \frac{\mathrm{d}(\pi_\varepsilon^t)_{\mathcal{Y}}}{\mathrm{d}\nu}(y) + \mathbb{E}_{\pi_\varepsilon^t(\cdot|x)}\left[\log \frac{\mathrm{d}(\pi_\varepsilon^t)_{\mathcal{Y}}}{\mathrm{d}\nu}\right]. \end{aligned} \tag{13}$$

We call the mapping $(\pi_\varepsilon^0, t) \mapsto \pi_\varepsilon^t$ the *Sinkhorn$_\varepsilon$ flow*. Similarly, we call the mapping $(g_\varepsilon^0, t) \mapsto g_\varepsilon^t$ the *Schrödinger$_\varepsilon$ flow*, which describes the evolution of the Schrödinger potential corresponding to π_ε^t:

$$\frac{\mathrm{d}}{\mathrm{d}t} g_\varepsilon^t = -\log \frac{\mathrm{d}(\pi_\varepsilon^t)_{\mathcal{Y}}}{\mathrm{d}\nu} = -\delta F(\pi_\varepsilon^t). \tag{14}$$

Remark. Using $\pi_\varepsilon^t = \delta\varphi^*(g_\varepsilon^t)$, the Schrödinger$_\varepsilon$ flow (14) can be written solely in terms of g_ε^t. The other direction also holds: if we integrate (14) to get g_ε^t, we can recover the Sinkhorn$_\varepsilon$ flow (13). In short, the two flows (13) and (14) are equivalent.

§ Mirror Flow interpretation. As the Sinkhorn$_\varepsilon$ and Schrödinger$_\varepsilon$ flows emerge from driving the step-size of an MD iterate to zero, it is natural to anticipate a *mirror flow* interpretation in the sense of (5). In this section, we make this link precise.

Consider the primal space $\mathcal{P}$ to be the space of probability measures over $\mathcal{X}\times\mathcal{Y}$ having smooth densities with respect to the Lebesgue measure, and the dual space $\mathcal{D} \coloneqq L^1(\mathcal{X}\times\mathcal{Y})$ to be the space of integrable functions. The first variation of the Bregman potential $\varphi(\pi) = H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}})$ gives a link from $\mathcal{P}$ to $\mathcal{D}$:

$$\delta\varphi(\pi) = \log\frac{\mathrm{d}\pi}{\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}} \in \mathcal{D}, \quad \forall \pi\in\mathcal{P},\, \pi \ll \pi_\varepsilon^{\mathrm{ref}}. \tag{15}$$

Moreover, the Fenchel conjugate φ^* of $\varphi + \mathrm{I}_{\mathcal{C}}$, defined in (3), gives a link from $\mathcal{D}$ to $\mathcal{C} = \{\pi : \pi_{\mathcal{X}} = \mu\} \subset \mathcal{P}$ through its first variation.

Crucially, unlike prior studies such as (Aubin-Frankowski et al., 2022; Ballu and Berthet, 2023; Mishchenko, 2019), we incorporate the $\mathrm{I}_{\mathcal{C}}$ in our definition of φ^*. This plays a vital role in preventing the emergence of a differential *inclusion*, as opposed to an equation, in the dual iteration, which is pivotal in our interpretation of mirror flow and the Otto calculus.

Lemma 3. *The Fenchel conjugate φ^* of $\varphi + \mathrm{I}_{\mathcal{C}}$ evaluated at $h\in\mathcal{D}$ is given by $\varphi^*(h) = \langle\hat\pi, h\rangle - H(\hat\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}})$, where*

$$\hat\pi(x,y) \coloneqq \frac{\pi_\varepsilon^{\mathrm{ref}}(x,y)\, e^{h(x,y)}}{\int \pi_\varepsilon^{\mathrm{ref}}(x,y')\, e^{h(x,y')}\, dy'}\,\mu(x) \in \mathcal{C}. \tag{16}$$

Moreover, one has $\delta\varphi^*(h) = \hat\pi$ where $\hat\pi$ is defined in (16).

Remark. We have seen that, when $\pi = \exp(f\oplus g)\pi_\varepsilon^{\mathrm{ref}}$, the recovery of π from g can be accomplished using the relationship expressed in (9). Lemma 3 further demonstrates that this operation is none other than an application of the dual mirror map: $\pi = \delta\varphi^*(g)$.

Using these formulas for $\delta\varphi$ and $\delta\varphi^*$, we can then define an infinite-dimensional *mirror flow* as follows. Fix $\hat\pi_\varepsilon^0 \in \Pi_{c,\varepsilon}\cap\{\pi : \pi_{\mathcal{X}} = \mu\}$ and $h_\varepsilon^0 \coloneqq \delta\varphi(\hat\pi_\varepsilon^0) \in \mathcal{D}$, and consider

$$\begin{cases} \frac{\mathrm{d}}{\mathrm{d}t} h_\varepsilon^t = -\delta F(\hat\pi_\varepsilon^t), \\ \hat\pi_\varepsilon^t = \delta\varphi^*(h_\varepsilon^t), \end{cases} \tag{17}$$

which can be equivalently written as

$$\frac{\mathrm{d}}{\mathrm{d}t} h_\varepsilon^t = -(\delta F \circ \delta\varphi^*)(h_\varepsilon^t). \tag{18}$$

We then have:

Theorem 3.1. *The dynamics* (17) *or* (18) *coincide with the Schrödinger$_\varepsilon$ flow* (14)*, and the corresponding $(\hat\pi_\varepsilon^t)_{t\geq 0}$ solves the Sinkhorn$_\varepsilon$ flow* (13) *starting at $\hat\pi_\varepsilon^0$.*

Comparing Eqs. (17)–(18) with Eqs. (5)–(6), we thus see that Sinkhorn$_\varepsilon$ and Schrödinger$_\varepsilon$ flows can be seen as the analogue of the mirror flow in the space of probability measures.

3.4 Convergence of Sinkhorn$_\varepsilon$ Flow and γ-Sinkhorn Iterates

The flexibility of the variable step-sizes in γ-Sinkhorn schemes offers a straightforward framework for improving the traditional Sinkhorn algorithm. Furthermore, the continuous-time Sinkhorn$_\varepsilon$ and Schrödinger$_\varepsilon$ flows pave the way for integrating the powerful machinery of *stochastic approximation* techniques (Karimi et al., 2022a,b; Mertikopoulos et al., 2023). In this section, we illustrate how to leverage our theory to enhance the convergence of Sinkhorn schemes in scenarios involving *noisy gradients*.

§ Rate analysis for the Sinkhorn$_\varepsilon$ flow. Before we proceed to present improved Sinkhorn schemes, we first establish the convergence rate of the continuous-time Sinkhorn$_\varepsilon$ flow:

Theorem 3.2. *Starting from $\pi_\varepsilon^0 \in \Pi_{c,\varepsilon}\cap\{\pi : \pi_{\mathcal{X}} = \mu\}$, consider the Sinkhorn$_\varepsilon$ flow π_ε^t and the corresponding Schrödinger$_\varepsilon$ flow g_ε^t. Then,*

$$F(\pi_\varepsilon^t) \leq \frac{D_{\varphi^*}(g_\varepsilon^0 \,\|\, g_\varepsilon^{\mathrm{opt}})}{t} = \mathcal{O}(t^{-1}),$$

where $g_\varepsilon^{\mathrm{opt}}$ is the Schrödinger potential of the optimal coupling for (OT$_\varepsilon$)*. That is, the $\mathcal{Y}$-marginal of π_ε^t converges (in relative entropy) to ν with the rate $1/t$.*

While our proof of Theorem 3.2 is guided by the mirror flow formalism presented in the previous section, it does *not* follow directly from existing results for mirror flows such as (Krichene et al., 2015; Tzen et al., 2023). This is primarily due to the presence of the additional constraint $\mathcal{C}$, which is absent in conventional mirror descent analyses.

§ Convergent Sinkhorn under noise. In entropy-regularized OT, neural networks (NNs) are commonly used to parameterize the transport plans. Typically, the Sinkhorn iterations (Sink$_1$) are employed, requiring solving an infinite-dimensional optimization problem, approximated via multiple stochastic gradient steps over NNs. However, inherent stochasticity in computations can *prevent convergence* when δF in (10) is replaced by a noisy estimate $\tilde\delta F$, necessitating a remedy (Hanzely and Richtárik, 2021).

In this section, we introduce two improvements. First, Theorem 3.3 shows that using our variable step-size method (Sink_γ) with $\gamma_n = \mathcal{O}(n^{-1/2})$, one maintains a convergence rate of $\mathcal{O}(n^{-1/2})$, when the "stochastic gradients" remain unbiased with finite variance. Second, Theorem 3.4 establishes asymptotic *last-iterate* convergence when one employs stochastic *and* biased gradient estimates.

Theorem 3.3. *Suppose that we have a stochastic estimate $\tilde{\delta}F$ of δF such that $\mathbb{E}[\tilde{\delta}F(\pi)] = \delta F(\pi)$ and $\mathbb{E}[\|\tilde{\delta}F(\pi)\|_\infty^2] \le \sigma^2 < \infty$ for all π. Consider the iterations π_γ^n generated by (Sink_γ) using $\tilde{\delta}F$ and a fixed step-size γ. Then we have, with $\bar{\pi}_\gamma^n := \frac{1}{n}\sum_{k=0}^n \pi_\gamma^k$,*

$$\mathbb{E}\big[H((\bar{\pi}_\gamma^n)_\mathcal{Y} \,\|\, \nu)\big] \le \frac{H(\pi^{\text{opt}} \,\|\, \pi_\varepsilon^{\text{ref}})}{\gamma n} + \gamma\sigma^2. \tag{19}$$

The proof of Theorem 3.3 is established by combining our framework with a classical analysis of stochastic Bregman schemes (Dragomir et al., 2021; Hanzely and Richtárik, 2021).

While (19) immediately yields an $\mathcal{O}(n^{-1/2})$ convergence rate, there are two significant drawbacks in Theorem 3.3. First, since the stochastic estimate $\tilde{\delta}F$ aims to capture noise introduced during the intermediate optimization procedures for NNs, the unbiasedness assumption is rather restrictive. Second, even if $\tilde{\delta}F$ is unbiased, we are still required to produce an *ergodic* iterate $\bar{\pi}_\gamma^n$, whereas in practice, the last iterate π_γ^n is often the most utilized. To address these issues, we leverage stochastic approximation analysis, which relies on the *continuous-time* convergence in Theorem 3.2 to prove the following (see Theorem B.1 for details):

Theorem 3.4 (Informal)**.** *Let π^n be the sequence of measures generated by (Sink_γ) using noisy and biased gradients $\tilde{\delta}F$, along with a step-size rule γ_n such that $\sum \gamma_n = \infty$ and $\sum \gamma_n^2 < \infty$. Then $\lim_{n\to\infty} \pi^n = \pi_\varepsilon^{\text{opt}}$ almost surely, if the biases vanish asympotically and the noises have uniformly bounded variance.*

Theorem 3.4 offers two advantages over Theorem 3.3. First, it replaces ergodic convergence with the more desirable *last-iterate* convergence. Secondly, if we consider the bias as the error during the optimization of the NN at each step of (Sink_γ), then Theorem 3.4 allows for a level of flexibility where the precision of the intermediate steps may progressively improve, instead of always requiring perfect optimization as stipulated by the unbiased assumption in Theorem 3.3. However, we acknowledge that these advantages come at the cost of losing a non-asymptotic rate.

3.5 Otto Calculus for $\text{Sinkhorn}_\varepsilon$ and $\text{Schrödinger}_\varepsilon$ Flows

The Otto calculus, a groundbreaking development in 21st-century mathematics, has found far-reaching implications in machine learning applications (Ambrosio et al., 2005; Otto, 2001; Villani, 2008). Leveraging this powerful framework, we showcase how our $\text{Sinkhorn}_\varepsilon$ and $\text{Schrödinger}_\varepsilon$ flows offer a unified template for various recently discovered dynamics that hold significant relevance to machine learning (Claisse et al., 2023; Deb et al., 2023).

§ Wasserstein Mirror Flow. Consider the time-dilated Schrödinger potential as $\tilde{g}_\varepsilon^t := g_\varepsilon^{t/\varepsilon}$. Recall that $\tilde{g}_\varepsilon^t \equiv \tilde{g}_\varepsilon^t(y)$ is a function on $\mathcal{Y}$. By taking the gradient with respect to the y variable in (14), we obtain:

$$\frac{\mathrm{d}}{\mathrm{d}t}\nabla(\varepsilon\tilde{g}_\varepsilon^t) = \frac{\mathrm{d}}{\mathrm{d}t}\nabla_{\mathbb{W}_2}\mathrm{OT}_\varepsilon(\mu,\cdot) = -\nabla_{\mathbb{W}_2}F(\pi_\varepsilon^t), \tag{20}$$

where $\nabla_{\mathbb{W}_2}F(\pi) := \nabla\delta F(\pi)$ is the *Wasserstein gradient* of F (Otto, 2001) and similar for $\nabla_{\mathbb{W}_2}\mathrm{OT}_\varepsilon$ (Benamou et al., 2023), which is evaluated at $(\delta\varphi^*(\tilde{g}_\varepsilon^t))_\mathcal{Y}$.

Now, let $c(x,y) = \frac{1}{2}|x-y|^2$ in (OT_ε). In this case, it is well-known that, as $\varepsilon \to 0$, $\pi_\varepsilon^{\text{opt}}$ converges to the optimal $\mathbb{W}_2$ coupling (denoted by ϖ) and $\varepsilon g_\varepsilon^{\text{opt}}$ converges to the corresponding Kantorovich potential (denoted by g_0) (Chiarini et al., 2023; Pooladian and Niles-Weed, 2021). Thus, by driving $\varepsilon \to 0$ in (20), we formally get:

$$\frac{\mathrm{d}}{\mathrm{d}t}\nabla g_0^t = \frac{\mathrm{d}}{\mathrm{d}t}\nabla_{\mathbb{W}_2}\mathrm{OT}_0(\mu,\varpi_\mathcal{Y}^t) = -\nabla_{\mathbb{W}_2}F(\varpi^t) \tag{21}$$

where ϖ^t is the optimal $\mathbb{W}_2$ coupling of μ and $\varpi_\mathcal{Y}^t$, and g_0^t is the corresponding Kantorovich potential. This equation is exactly the "Wasserstein Mirror Flow" of Deb et al. (2023, equation (2.3)) proposed to study Sinkhorn for the unregularized OT. Therefore, the dynamics discovered by Deb et al. (2023) can be seen as a limiting case of the more general $\text{Sinkhorn}_\varepsilon$ and $\text{Schrödinger}_\varepsilon$ flows.

§ JKO Flow with Relative Entropy. Recently, motivated by the mean-field limit of neural network training, Claisse et al. (2023) study a variation of the JKO flow of a functional F (Jordan et al., 1998), replacing the $\mathbb{W}_2$ distance with the relative entropy:

$$p_\gamma^{n+1} := \arg\min_p\big\{F(p) + \gamma^{-1}H(p \,\|\, p_\gamma^n)\big\}, \tag{22}$$

where the minimization is over the set of all distributions with regular densities. In the limit $\gamma \to 0$, they show that such as a scheme converges to the flow

$$\frac{\mathrm{d}}{\mathrm{d}t}\log p_t = -\delta F(p_t) \tag{23}$$

which formally resembles (14). However, there are two important differences. Firstly, (22) constitutes an *unconstrained* minimization problem, whereas our mirror descent scheme seeks minimizers within the constraint set $\mathcal{C} = \{\pi : \pi_{\mathcal{X}} = \mu\}$, giving rise to the central notion of Schrödinger potentials. Secondly, the update (22) involves the original objective function F in the minimization, whereas our scheme employs the linearized objective. While the second difference becomes negligible as $\gamma \to 0$ (intuitively, there would be no distinction between F and its linearization), the impact of the constraint set $\mathcal{C}$ persists. Consequently, one can view (23) as an *unconstrained* mirror flow for the objective F with the entropy as the Bregman potential.

4 EXTENSION TO SCHRÖDINGER BRIDGES

In Section 3, we established the continuous-time variant of the Sinkhorn iterates for (OT_ε), which pertains to the "static" entropy-regularized OT. Motivated by the strong connections to diffusion models, in this section, we broaden our scope to encompass the *dynamical* scenario, often referred to as the *Schrödinger bridge* (SB) problem. Beyond adapting the results in Section 3 to the SB, we provide additional insights by demonstrating that each time point in the continuous-time SB flow can be characterized as a *stochastic differential equation* with a well-defined drift formula. All proofs can be found in Appendix C.

4.1 Review of Schrödinger Bridges

Given two probability measures μ_0, μ_T on $\mathbb{R}^d$, the SB refers to the following entropy minimization problem over the space of all *stochastic processes* over $[0, T]$:

$$\min_{\mathsf{P}} \big\{ H(\mathsf{P} \,\|\, \mathsf{P}^{\mathrm{ref}}) : \mathsf{P}_0 = \mu_0,\, \mathsf{P}_T = \mu_T \big\}, \tag{SB}$$

where $\mathsf{P}^{\mathrm{ref}}$ is a given path measure induced by the solutions of the stochastic differential equation

$$\mathrm{d}X_t = b_t^{\mathrm{ref}}(X_t)\,\mathrm{d}t + \sigma \mathrm{d}W_t, \tag{24}$$

and P_t is the marginal of P at time t. It turns out that solving (SB) is intimately related to solving the *static* Schrödinger problem (Léonard, 2014):

$$\min_{\pi \in \Gamma(\mu_0, \mu_T)} H(\pi \,\|\, \mathsf{P}^{\mathrm{ref}}_{0,T}). \tag{25}$$

§ Connection to Entropy-Regularized OT. In the case where $\mathsf{P}^{\mathrm{ref}}$ is the law of a *reversible* Brownian motion on $[0, 1]$ with diffusion parameter σ (Léonard, 2014), the joint distribution of the end time points satisfies $\mathsf{P}^{\mathrm{ref}}_{0,1}(\mathrm{d}x, \mathrm{d}y) \propto \exp(-|x-y|^2/2\sigma^2)$. Therefore, (25) becomes an instance of entropy-regularized OT (7) with the cost $c(x, y) = \frac{1}{2}|x-y|^2$ and $\varepsilon = \sigma^2$. Thus, (SB) can be viewed as the *dynamic* formulation of (OT_ε) where, instead of merely seeking an optimal coupling $\pi_\varepsilon^{\mathrm{opt}}$, one solves for an entire stochastic process that transforms μ_0 into μ_1.

§ Iterative Proportional Fitting. The classical algorithm for solving (SB) is the *Iterative Proportional Fitting* (IPF) procedure, which can be seen as the *dynamic* version of the Sinkhorn scheme: Starting from $\mathsf{P}^0 = \mathsf{P}^{\mathrm{ref}}$, define for $n \geq 0$,

$$\begin{aligned} \mathsf{P}^{n+1/2} &= \arg\min\{H(\mathsf{P} \,\|\, \mathsf{P}^n) : \mathsf{P}_T = \mu_T\}, \\ \mathsf{P}^{n+1} &= \arg\min\{H(\mathsf{P} \,\|\, \mathsf{P}^{n+1/2}) : \mathsf{P}_0 = \mu_0\}. \end{aligned} \tag{IPF}$$

Notice that, in complete analogy to Sinkhorn, we have $\mathsf{P}^n_0 = \mu_0$ and $\mathsf{P}^{n+1/2}_T = \mu_T$ for all $n \geq 0$.

With this background, we now present an approach to IPF for (SB) that parallels the Sinkhorn for (OT_ε).

4.2 IPF as Mirror Descent

Similar to the Sinkhorn algorithm, we show that IPF can be interpreted through the lens of MD. This finding serves as the dynamic counterpart to (Aubin-Frankowski et al., 2022, **Proposition 5**).

Proposition 2. *The iterations P^n of (IPF) satisfy*

$$\mathsf{P}^{n+1} = \arg\min_{\mathsf{P} \in \mathcal{C}} \{\langle \delta F(\mathsf{P}^n), \mathsf{P} - \mathsf{P}^n \rangle + D_\varphi(\mathsf{P} \,\|\, \mathsf{P}^n)\}, \tag{26}$$

with $F(\mathsf{P}) \coloneqq H(\mathsf{P}_T \,\|\, \mu_T)$, $\varphi(\mathsf{P}) \coloneqq H(\mathsf{P} \,\|\, \mathsf{P}^{\mathrm{ref}})$, and $\mathcal{C} \coloneqq \{\mathsf{P} : \mathsf{P}_0 = \mu_0\}$.

In other words, (IPF) is equivalent to an MD iteration with step-size 1.

§ γ-IPF iterations. Upon recognizing that IPF can be interpreted as MD iterations with a step-size of 1, we can proceed to investigate the MD iteration (26) with an arbitrary step-size γ_n:

$$\mathsf{P}^{n+1} = \arg\min_{\mathsf{P} \in \mathcal{C}} \left\{ \langle \delta F(\mathsf{P}^n), \mathsf{P} - \mathsf{P}^n \rangle + \frac{D_\varphi(\mathsf{P} \,\|\, \mathsf{P}^n)}{\gamma_n} \right\}. \tag{27}$$

A similar calculation to that of Lemma 1 reveals that (27) can be equivalently expressed as:

$$\mathsf{P}^{n+1} = \arg\min_{\mathsf{P} \in \mathcal{C}} \Big\{ \gamma_n H(\mathsf{P} \,\|\, \mathsf{P}^{n+1/2}) + (1 - \gamma_n) H(\mathsf{P} \,\|\, \mathsf{P}^n) \Big\}. \tag{γ-IPF}$$

In analogy to the γ-Sinkhorn iterations, we call the update rule in (γ-IPF) the *γ-IPF scheme.*

4.3 The SDE Representation of γ-IPF

So far, we have shown that the results in Section 3 can be straightforwardly extended to SB setting. However, the true significance of the SB formulation becomes apparent in its representation as *stochastic differential equations* (SDEs), enabling the utilization of powerful diffusion models (Ho et al., 2020; Song and Ermon, 2019). Thus, the primary objective of this section is to establish that the minimizer of (γ-IPF) can indeed be expressed as an SDE with a readily available drift formula.

Before delving into the derivation, we first recall the important fact that the IPF iterates can be expressed in terms of the *time-reversal* of SDEs, which can be solved in practice via score matching techniques (Chen et al., 2022; De Bortoli et al., 2021). We provide a proof for this result for completeness.

Theorem 4.1. *Suppose that P^n is an SDE given by*

$$\mathrm{d}X_t^n = v_t^n(X_t^n)\,\mathrm{d}t + \sigma \mathrm{d}W_t,\ X_0 \sim \mu_0, \tag{28}$$

and that the time-reversal of $\mathsf{P}^{n+1/2}$ is given by

$$\mathrm{d}Y_t^{n+1/2} = w_{T-t}^{n+1/2}(Y_t^{n+1/2})\,\mathrm{d}t + \sigma \mathrm{d}W_t,\ Y_0^{n+1/2} \sim \mu_T. \tag{29}$$

Then the drift vector field $w_t^{n+1/2}$ satisfies:

$$-v_t^n(x) + \sigma^2 \nabla \log p_t^n(x) = w_t^{n+1/2}(x) \tag{30}$$

where p_t^n is the density of P_t^n.

It is well-known that the condition (30) expresses precisely the fact that $\mathsf{P}^{n+1/2}$ is given by the time-reversal of (28) (Föllmer, 1985; Haussmann and Pardoux, 1986; Song et al., 2020).

In what follows, define the likelihood ratio to be $\ell_t^n \coloneqq {p_t^{n+1/2}}/{p_t^n}$. We are now ready to present the main result of this section.

Theorem 4.2. *Let P^n be given by the scheme (γ-IPF), and let $v_t^n(\cdot)$ be the (forward) vector field corresponding to the SDE representation of P^n in (28). Then $v_t^n(\cdot)$ satisfies the following recursive formula:*

$$v_t^{n+1} = v_t^n + \gamma\sigma^2 \nabla \log \ell_t^n - \sigma^2 \nabla V_t, \tag{SDE$_\gamma$}$$

where

$$V_t(x) = -\log \mathbb{E}\left[e^{-\frac{\sigma^2\gamma(1-\gamma)}{2}\int_t^T |\nabla \log \ell_s^n(Y_s)|^2\,\mathrm{d}s}\right], \tag{31}$$

and the expectation is with respect to the law of the SDE $(Y_s)_{s\geq t}$ starting at $Y_t = x$ and following

$$\mathrm{d}Y_s = \left\{v_s^n(Y_s) + \gamma\sigma^2\nabla \log \ell_s^n(Y_s)\right\}\mathrm{d}s + \sigma \mathrm{d}W_s. \tag{32}$$

When $\gamma = 1$, the ∇V_t term in (SDE$_\gamma$) disappears so that Theorem 4.2 is Theorem 4.1 applied twice, and we recover the iterative formula for the SDE representation of IPF (Pariset et al., 2023, **Proposition 4.2**).

Remark. Comparing (SDE$_\gamma$) with the classical IPF iterates, we see that γ-IPF *can be efficiently implemented* provided we can calculate the additional ∇V_t component. In Proposition C.2, we make this computational step possible by a establishing novel link to stochastic optimal control. As our paper primarily focuses on the theoretical understanding of the Sinkhorn and IPF iterations, we defer the details to Appendix C.

4.4 The Flow of Schrödinger SDEs

In this section, we show how the results in Section 3 naturally lead to *a flow of SDEs*, i.e., an evolution of path measures $(\mathsf{P}^s)_{s\geq 0}$ where each P^s is the law of an SDE with certain drift $v_t^s(\cdot)$:

$$\mathrm{d}X_t^s = v_t^s(X_t^s)\mathrm{d}t + \sigma \mathrm{d}W_t. \tag{33}$$

To streamline the exposition, we make the simplifying assumption that $T = \sigma = 1$ and the reference measure $\mathsf{P}^{\mathrm{ref}}$ is given by the law of the reversible Brownian motion $(W_t^{\mathrm{R}})_{t\in[0,1]}$ (Léonard, 2014). Our conclusions remain applicable in the general case, but the notation becomes more cumbersome in that scenario.

Consider the static SB problem in (25), which is nothing but (OT$_\varepsilon$) with cost function $c(x,y) = \frac{1}{2}|x-y|^2$ and $\varepsilon = 1$. Recall its associated Schrödinger$_\varepsilon$ flow (14) defined via the Schrödinger potentials f^s, g^s. Consider the path measures P^s defined by

$$\frac{\mathrm{d}\mathsf{P}^s}{\mathrm{d}\mathsf{P}^{\mathrm{ref}}} = \exp\{(f^s \oplus g^s)(W_0^{\mathrm{R}}, W_1^{\mathrm{R}})\}. \tag{34}$$

Similar to the static case, these path measures are known to solve the SB problem for their corresponding marginals μ_0^s, μ_1^s (Léonard, 2014) and, by construction, $\mu_0^s = \mu_0$ for all s.

We can now formally define an evolution of the path measures P^s, where at each time s, P^s admits an SDE representation which can be described using the Schrödinger potentials f^s, g^s: For each s, define the function on $[0,1]\times\mathbb{R}^d$ by

$$g_t^s(z) \coloneqq \log \mathbb{E}\left[e^{g^s(W_1^{\mathrm{R}})} \mid W_t^{\mathrm{R}} = z\right] \tag{35}$$

so that $g_1^s \equiv g^s$. Then Léonard (2014, **Prop. 6**) implies that P^s is the law of the SDE:

$$\mathrm{d}X_t^s = \nabla g_t^s(X_t^s)\mathrm{d}t + \mathrm{d}W_t, \quad X_0^s \sim \mu_0. \tag{36}$$

As a result, the mapping $s \mapsto (g_t^s)_{t\in[0,1]}$ can be regarded as the *dynamic* Schrödinger$_\varepsilon$ flow associated with $(g^s)_{s\geq 0}$, while (36) can be considered as the continuous-time limit of the SDE representation of (γ-IPF), as $\gamma \to 0$.

5 CONCLUSIONS AND FUTURE WORK

In summary, our work introduces the continuous-time Sinkhorn algorithm as a novel approach to design schemes that maintain convergence in the presence of noise and bias. It also unifies previously isolated dynamics through the mirror descent perspective. We extend these insights to Schrödinger bridges and the IPF procedure. These findings open doors to exciting future research directions, including exploring connections with existing dynamics and the potential for achieving acceleration through momentum terms.

Acknowledgements

This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement No 815943. YPH acknowledges funding through an ETH Foundations of Data Science (ETH-FDS) postdoctoral fellowship.

References

Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.

Pierre-Cyril Aubin-Frankowski, Anna Korba, and Flavien Léger. Mirror Descent with Relative Smoothness in Measure Spaces, with application to Sinkhorn and EM, October 2022.

Pierre-Cyril Aubin-Frankowski, Anna Korba, and Flavien Léger. Mirror descent with relative smoothness in measure spaces, with application to sinkhorn and em. *Advances in Neural Information Processing Systems*, 35:17263–17275, 2022.

Marin Ballu and Quentin Berthet. Mirror sinkhorn: Fast online optimization on transport polytopes. In *International Conference on Machine Learning*, pages 1595–1613. PMLR, 2023.

Heinz H Bauschke, Jonathan M Borwein, and Patrick L Combettes. Bregman monotone optimization algorithms. *SIAM Journal on control and optimization*, 42(2):596–636, 2003.

Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3): 167–175, 2003.

Michel Benaïm. Dynamics of stochastic approximation algorithms. In Jacques Azéma, Michel Émery, Michel Ledoux, and Marc Yor, editors, *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pages 1–68. Springer Berlin Heidelberg, 1999.

Jean-David Benamou, Colin Cotter, and Hugo Malamut. Entropic Optimal Transport Solutions of the Semigeostrophic Equations. 2023.

Espen Bernton, Jeremy Heng, Arnaud Doucet, and Pierre E Jacob. Schrödinger Bridge Samplers. In *arXiv preprint arXiv:1912.13170*, 2019.

Charlotte Bunne, Ya-Ping Hsieh, Marco Cuturi, and Andreas Krause. The schrödinger bridge between gaussian measures has a closed form. In *International Conference on Artificial Intelligence and Statistics*, pages 5802–5833. PMLR, 2023.

Tianrong Chen, Guan-Horng Liu, and Evangelos A Theodorou. Likelihood Training of Schrödinger Bridge using Forward-Backward SDEs Theory. In *International Conference on Learning Representations (ICLR)*, 2022.

Yongxin Chen, Tryphon Georgiou, and Michele Pavon. Entropic and displacement interpolation: a computational approach using the hilbert metric. *SIAM Journal on Applied Mathematics*, 76(6):2375–2396, 2016.

Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Optimal Transport in Systems and Control. *Annual Review of Control, Robotics, and Autonomous Systems*, 4, 2021.

Alberto Chiarini, Giovanni Conforti, Giacomo Greco, and Luca Tamanini. Gradient estimates for the schrödinger potentials: convergence to the brenier map and quantitative stability. *Communications in Partial Differential Equations*, pages 1–49, 2023.

Julien Claisse, Giovanni Conforti, Zhenjie Ren, and Songbo Wang. Mean field optimization problem regularized by fisher information. *arXiv preprint arXiv:2302.05938*, 2023.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.

Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion Schrödinger Bridge with Applications to Score-Based Generative Modeling. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 35, 2021.

Nabarun Deb, Young-Heon Kim, Soumik Pal, and Geoffrey Schiebinger. Wasserstein mirror gradient flow as the limit of the sinkhorn algorithm. *arXiv preprint arXiv:2307.16421*, 2023.

Radu Alexandru Dragomir, Mathieu Even, and Hadrien Hendrikx. Fast stochastic bregman gradient methods: Sharp analysis and variance reduction. In *International Conference on Machine Learning*, pages 2815–2825. PMLR, 2021.

Wendell H. Fleming. Exit probabilities and optimal stochastic control. *Applied Mathematics and Optimization*, 4(1):329–346, March 1977. ISSN 1432-0606. doi: 10.1007/BF01442148.

H Föllmer. An entropy approach to the time reversal of diffusion processes, stochastic differential systems filtering and control (m. metivier, e. pardoux, ed.), 1985.

Robert Fortet. Résolution d'un systeme d'équations de M. Schrödinger. *J. Math. Pure Appl. IX*, 1, 1940.

Promit Ghosal and Marcel Nutz. On the convergence rate of sinkhorn's algorithm. *arXiv preprint arXiv:2212.06000*, 2022.

Paul R Halmos. *Measure theory*, volume 18. Springer, 2013.

Filip Hanzely and Peter Richtárik. Fastest rates for stochastic mirror descent methods. *Computational Optimization and Applications*, 79:717–766, 2021.

U. G. Haussmann and E. Pardoux. Time Reversal of Diffusions. *The Annals of Probability*, 14(4):1188–1205, October 1986. ISSN 0091-1798, 2168-894X. doi: 10.1214/aop/1176992362.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Lars Holdijk, Yuanqi Du, Ferry Hooft, Priyank Jaini, Bernd Ensing, and Max Welling. Path Integral Stochastic Optimal Control for Sampling Transition Paths. *arXiv preprint arXiv:2207.02149*, 2022.

Ya-Ping Hsieh, Chen Liu, and Volkan Cevher. Finding mixed nash equilibria of generative adversarial networks. In *International Conference on Machine Learning*, pages 2810–2819, 2019.

Jian Huang, Yuling Jiao, Lican Kang, Xu Liao, Jin Liu, and Yanyan Liu. Schrödinger-Föllmer Sampler: Sampling without Ergodicity. *arXiv preprint arXiv:2106.10880*, 2021.

Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the fokker–planck equation. *SIAM journal on mathematical analysis*, 29(1):1–17, 1998.

Mohammad Reza Karimi, Ya-Ping Hsieh, and Andreas Krause. A dynamical system view of langevin-based non-convex sampling. *arXiv preprint arXiv:2210.13867*, 2022a.

Mohammad Reza Karimi, Ya-Ping Hsieh, Panayotis Mertikopoulos, and Andreas Krause. The dynamics of Riemannian Robbins-Monro algorithms. In *COLT 2022-35th Annual Conference on Learning Theory*, pages 1–31, 2022b.

Walid Krichene, Alexandre Bayen, and Peter L Bartlett. Accelerated mirror descent in continuous and discrete time. In *Advances in neural information processing systems*, pages 2845–2853, 2015.

Solomon Kullback. Probability densities with given marginals. *The Annals of Mathematical Statistics*, 39(4):1236–1243, 1968.

Flavien Léger. A gradient descent perspective on sinkhorn. *Applied Mathematics & Optimization*, 84 (2):1843–1855, 2021.

Christian Léonard. Girsanov theory under a finite entropy condition, January 2011.

Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete and Continuous Dynamical Systems - Series A*, 34(4):1533–1574, 2014.

Guan-Horng Liu, Tianrong Chen, Oswin So, and Evangelos A Theodorou. Deep Generalized Schrödinger Bridge. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Arthur Mensch and Gabriel Peyré. Online sinkhorn: Optimal transport distances from sample streams. *Advances in Neural Information Processing Systems*, 33:1657–1667, 2020.

Panayotis Mertikopoulos, Ya-Ping Hsieh, and Volkan Cevher. A unified stochastic approximation framework for learning in games. *Mathematical Programming*, pages 1–51, 2023.

Konstantin Mishchenko. Sinkhorn algorithm as a special case of stochastic mirror descent. *arXiv preprint arXiv:1909.06918*, 2019.

AS Nemirovsky and DB Yudin. Problem complexity and method efficiency in optimization. 1983.

Marcel Nutz and Johannes Wiesel. Entropic optimal transport: Convergence of potentials. *Probability Theory and Related Fields*, 184(1-2):401–424, October 2022. ISSN 0178-8051, 1432-2064. doi: 10.1007/s00440-021-01096-8.

Felix Otto. The geometry of dissipative evolution equations: the porous medium equation. 2001.

Matteo Pariset, Ya-Ping Hsieh, Charlotte Bunne, Andreas Krause, and Valentin De Bortoli. Unbalanced diffusion schr\" odinger bridge. *arXiv preprint arXiv:2306.09099*, 2023.

Gabriel Peyré and Marco Cuturi. Computational optimal transport: With applications to data science.

Foundations and Trends® in Machine Learning, 11 (5-6):355–607, 2019.

Aram-Alexandre Pooladian and Jonathan Niles-Weed. Entropic estimation of optimal transport maps. *arXiv preprint arXiv:2109.12004*, 2021.

Paolo Dai Pra and Michele Pavon. On the Markov processes of Schrödinger, the Feynman-Kac formula and stochastic control. In M. A. Kaashoek, J. H. Van Schuppen, and A. C. M. Ran, editors, *Realization and Modelling in System Theory*, pages 497–504. Birkhäuser Boston, Boston, MA, 1990. ISBN 978-1-4612-8033-0 978-1-4612-3462-3. doi: 10.1007/978-1-4612-3462-3_55.

R Tyrrell Rockafellar. *Convex analysis*, volume 11. Princeton university press, 1997.

Vivien Seguy, Bharath Bhushan Damodaran, Rémi Flamary, Nicolas Courty, Antoine Rolet, and Mathieu Blondel. Large-scale optimal transport and mapping estimation. *arXiv preprint arXiv:1711.02283*, 2017.

Richard Sinkhorn and Paul Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2):343–348, 1967.

Vignesh Ram Somnath, Matteo Pariset, Ya-Ping Hsieh, Maria Rodriguez Martinez, Andreas Krause, and Charlotte Bunne. Aligned diffusion schr\" odinger bridges. *arXiv preprint arXiv:2302.11419*, 2023.

Yang Song and Stefano Ermon. Generative Modeling by Estimating Gradients of the Data Distribution. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.

Belinda Tzen, Anant Raj, Maxim Raginsky, and Francis Bach. Variational principles for mirror descent and mirror langevin dynamics. *IEEE Control Systems Letters*, 2023.

Francisco Vargas, Pierre Thodoroff, Neil D Lawrence, and Austen Lamacraft. Solving Schrödinger Bridges via Maximum Likelihood. *Entropy*, 23(9), 2021.

Cédric Villani. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.

Gefei Wang, Yuling Jiao, Qian Xu, Yang Wang, and Can Yang. Deep Generative Learning via Schrödinger Bridge. In *International Conference on Machine Learning (ICML)*, 2021.

Qinsheng Zhang and Yongxin Chen. Path Integral Sampler: A Stochastic Control Approach For Sampling. In *International Conference on Learning Representations (ICLR)*, 2022.

Checklist

1. For all models and algorithms presented, check if you include:
   (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
   (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
   (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
   (a) Statements of the full set of assumptions of all theoretical results. [Yes]
   (b) Complete proofs of all theoretical results. [Yes]
   (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
   (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Not Applicable]
   (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
   (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Not Applicable]
   (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
   (a) Citations of the creator If your work uses existing assets. [Not Applicable]
   (b) The license information of the assets, if applicable. [Not Applicable]
   (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
   (d) Information about consent from data providers/curators. [Not Applicable]
   (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
   (a) The full text of instructions given to participants and screenshots. [Not Applicable]
   (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
   (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A Background

A.1 Notions of Derivative

Throughout this paper, we use the notation δ to denote the *First Variation* operator. In this section, we bring the necessary background for defining and using this operator. Our discussion mostly follows (Aubin-Frankowski et al., 2022) with a slight change of notation.

Let $\mathcal{M}$ be a vector space of (signed) finite measures on the space $\mathcal{X}$.

Definition 3 (Gâteaux and Fréchet Differentiability)**.** A functional F is called Gâteaux differentiable at $\mu \in \mathcal{M}$, if there exists a linear operator $\nabla_{\text{Gât}} F(\mu)$ such that for any direction $\nu \in \mathcal{M}$, one has

$$\nabla_{\text{Gât}} F(\mu)[\nu] = \lim_{h \to 0} \frac{F(\mu + h\nu) - F(\mu)}{h}.$$

If the limit above holds uniformly in the unit ball in $\mathcal{M}$, the function F is called Fréchet differentiable, and we denote the resulting Fréchet derivative as $\nabla_{\text{Fré}} F(\mu)$.

The issue with the aforementioned definitions is that the limit must exist in all directions, which means that the points of differentiability must be within the functional's domain, as stated. Nevertheless, in the case of functionals like the relative entropy, the domain of F has an empty interior, as discussed by Aubin-Frankowski et al. (2022).

For this, Aubin-Frankowski et al. (2022) propose to use the weaker notion of *directional derivative*, defined as follows:

Definition 4 (Directional Derivative)**.** For a functional F and $\mu \in \mathcal{M}$ define the directional derivative of F at μ in the direction of $\nu \in \mathcal{M}$ as

$$F'(\mu; \nu) = \lim_{h \downarrow 0} \frac{F(\mu + h\nu) - F(\mu)}{h}.$$

See (Aubin-Frankowski et al., 2022, Remark 1) for a discussion on when this notion of derivative exists. Specifically, for convex and proper functions (such as relative entropy) this notion of derivative exists.

We are now ready to recall the notion of first variation, see (Aubin-Frankowski et al., 2022, Definition 2) and the discussion afterward for a more in-depth exposition.

Definition 5 (First Variation)**.** Let $\mathcal{C}$ be a subset of $\mathcal{M}$. For a functional F and $\mu \in \mathcal{C} \cap \operatorname{dom}(F)$ define the first variation of F at μ to be the element $\delta_{\mathcal{C}} F(\mu) \in \mathcal{M}^*$, where $\mathcal{M}^*$ is the topological dual of $\mathcal{M}$, such that it holds for all $\nu \in \mathcal{C} \cap \operatorname{dom}(F)$ and $\xi = \nu - \mu \in \mathcal{M}$:

$$\langle \delta_{\mathcal{C}} F(\mu), \xi \rangle = F'(\mu; \xi).$$

Here, $\langle \cdot, \cdot \rangle$ is the duality pairing of $\mathcal{M}$ and $\mathcal{M}^*$.

A.2 Selection of Results regarding SDEs

The first important result is the time-reversal formula for diffusions. The result, along with the necessary regularity conditions, can be found in, e.g., (Haussmann and Pardoux, 1986).

Theorem A.1 (Time-Reversal of Diffusions)**.** *Let $(X_t)_{t \in [0,1]}$ be the (strong) solution of $dX_t = v_t(X_t)\,dt + \sigma dW_t$, and assume X_t has density p_t. Define the vector field*

$$w_t(x) = -v_{1-t}(x) + \sigma^2 \nabla \log p_{1-t}(x)$$

and the operator L_t which, when evaluated on $f \in C_c^\infty(\mathbb{R}^d)$, gives

$$(L_t f)(x) = \frac{\sigma^2}{2} \Delta f(x) + \langle w_t(x), \nabla f(x) \rangle.$$

Then, under some assumptions (see, e.g., (A) in (Haussmann and Pardoux, 1986)), the process $(\hat{X}_t := X_{1-t})_{0 \le t < 1}$ is a Markov diffusion process with generator L_t.

The following two theorems are special cases of a general Girsanov formula; see e.g., (Léonard, 2011, Theorem 2.3).

Theorem A.2 (Girsanov)**.** *Let* P *be the law of the semi-martingale*

$$X_t = X_0 + \int_0^t b_s \, ds + \sigma W_t,$$

and suppose we are given a probability measure Q *with* $H(\mathsf{Q} \,\|\, \mathsf{P}) < \infty$*. Then, there exists an* $\mathbb{R}^d$*-valued adapted process* β_t *with* $\mathbb{E}_{\mathsf{Q}}\left[\int_0^1 |\beta_t|^2 \, dt\right] < \infty$*, such that* X *has the semi-martingale decomposition*

$$X_t = X_0 + \int_0^t (b_s + \beta_s) \, ds + \sigma W_t^{\mathsf{Q}}, \quad \mathsf{Q}\text{-}a.s.,$$

where W^{Q} *is a* Q*-Brownian motion. Moreover,*

$$\frac{d\mathsf{Q}}{d\mathsf{P}} = \frac{d\mathsf{Q}_0}{d\mathsf{P}_0}(X_0) \cdot \exp\left\{ \int_0^1 \frac{1}{\sigma} \langle \beta_t, dW_t \rangle - \frac{1}{2} \int_0^1 \frac{|\beta_t|^2}{\sigma^2} \, dt \right\}.$$

Corollary A.1 (Relative Entropy of Diffusions)**.** *Let* P *and* Q *be two path measures, with* $H(\mathsf{Q} \,\|\, \mathsf{P}) < \infty$*. Moreover, assume that under* P*, the canonical process* X *has the semi-martingale decomposition*

$$X_t = X_0 + \int_0^t b_s \, ds + W_t^{\mathsf{P}},$$

where W^{P} *is a* P*-Brownian motion, and under* Q*,*

$$X_t = X_0 + \int_0^t c_s \, ds + W_t^{\mathsf{Q}},$$

where W^{Q} *is a* Q*-Brownian motion. Then,*

$$H(\mathsf{Q} \,\|\, \mathsf{P}) = \mathbb{E}_{\mathsf{Q}}\left[\log \frac{d\mathsf{Q}}{d\mathsf{P}}\right] = H(\mathsf{Q}_0 \,\|\, \mathsf{P}_0) + \mathbb{E}_{\mathsf{Q}}\left[\frac{1}{2} \int_0^T |c_t - b_t|^2 \, dt\right].$$

B Proofs of Section 3

B.1 Results about the γ-Sinkhorn Iterates

Recall from Definition 1 that the γ-Sinkhorn iterates are defined via the recursion

$$\pi_\gamma^{n+1} = \operatorname*{arg\,min}_{\pi \in \mathcal{C}} \left\{ \langle \delta F(\pi_\gamma^n), \pi - \pi_\gamma^n \rangle + \frac{D_\varphi(\pi \,\|\, \pi_\gamma^n)}{\gamma_n} \right\}, \qquad \pi_\gamma^0 = \pi_\varepsilon^{\mathrm{ref}}.$$

Using the values of δF and $D_\varphi(\cdot \,\|\, \cdot)$ from (Aubin-Frankowski et al., 2022), this is equivalent to

$$\pi_\gamma^{n+1} = \operatorname*{arg\,min}_{\pi \in \mathcal{C}} \left\{ \int \mathrm{d}(\pi - \pi_\gamma^n) \log \frac{\mathrm{d}(\pi_\gamma^n)_{\mathcal{Y}}}{\mathrm{d}\nu} + \frac{H(\pi \,\|\, \pi_\gamma^n)}{\gamma_n} \right\}, \qquad \pi_\gamma^0 = \pi_\varepsilon^{\mathrm{ref}}. \tag{B.1}$$

Lemma 1. *The* γ*-Sinkhorn iterates* π_γ^n *defined as in Definition 1 correspond to the update rule:*

$$\begin{aligned} \pi_\gamma^{n+1/2} &:= \arg\min\{H(\pi \,\|\, \pi_\gamma^n) : \pi_{\mathcal{Y}} = \nu\}, \\ \pi_\gamma^{n+1} &:= \arg\min\{\gamma_n H(\pi \,\|\, \pi_\gamma^{n+1/2}) \\ &\qquad + (1-\gamma_n) H(\pi \,\|\, \pi_\gamma^n) : \pi_{\mathcal{X}} = \mu\}. \end{aligned} \tag{Sink$_\gamma$}$$

Proof. For brevity, we drop the γ in π^n_γ. First, observe that by the chain rule of the relative entropy,

$$H(\pi \,\|\, \pi^n) = H(\pi_{\mathcal{Y}} \,\|\, \pi^n_{\mathcal{Y}}) + \int \mathrm{d}\pi^n_{\mathcal{Y}}(y)\, H(\pi^n(\cdot \mid y) \,\|\, \pi(\cdot \mid y))$$

As $\pi^{n+1/2}$ is the minimizer of the KL above among all couplings with $\mathcal{Y}$-marginal ν, this means that for $\mathrm{d}\pi^n_{\mathcal{Y}}(y)$-almost sure y, we have $\pi^n(\cdot \mid y) = \pi(\cdot \mid y)$. This implies that

$$\frac{\mathrm{d}\pi^{n+1/2}}{\mathrm{d}\pi^n} = \frac{\mathrm{d}\pi^{n+1/2}_{\mathcal{Y}}}{\mathrm{d}\pi^n_{\mathcal{Y}}} = \frac{\mathrm{d}\nu}{\mathrm{d}\pi^n_{\mathcal{Y}}}. \tag{B.2}$$

From the computation above and (B.1), we have

$$\begin{aligned}
F(\pi^n) + \langle \delta F(\pi^n), \pi - \pi^n \rangle + \frac{D_\varphi(\pi \,\|\, \pi^n)}{\gamma_n} &= \int \mathrm{d}\pi^n \log \frac{\mathrm{d}\pi^n_{\mathcal{Y}}}{\mathrm{d}\nu} + \int \mathrm{d}(\pi - \pi^n) \log \frac{\mathrm{d}\pi^n_{\mathcal{Y}}}{\mathrm{d}\nu} + \frac{H(\pi \,\|\, \pi^n)}{\gamma_n} \\
&= \int \mathrm{d}\pi \log \frac{\mathrm{d}\pi^n_{\mathcal{Y}}}{\mathrm{d}\nu} + \frac{H(\pi \,\|\, \pi^n)}{\gamma_n} \\
&= \int \mathrm{d}\pi \log \frac{\mathrm{d}\pi^n}{\mathrm{d}\pi^{n+1/2}} + \frac{H(\pi \,\|\, \pi^n)}{\gamma_n} && \text{by (B.2)} \\
&= \int \mathrm{d}\pi \log \left\{ \frac{\mathrm{d}\pi^n}{\mathrm{d}\pi^{n+1/2}} \cdot \left(\frac{\mathrm{d}\pi}{\mathrm{d}\pi^n} \right)^{1/\gamma_n} \right\} \\
&= \frac{1}{\gamma_n} \int \mathrm{d}\pi \log \left\{ \left(\frac{\mathrm{d}\pi}{\mathrm{d}\pi^{n+1/2}} \right)^{\gamma_n} \cdot \left(\frac{\mathrm{d}\pi}{\mathrm{d}\pi^n} \right)^{1-\gamma_n} \right\} \\
&= \frac{1}{\gamma_n} \Big(\gamma_n H(\pi \,\|\, \pi^{n+1/2}) + (1-\gamma_n) H(\pi \,\|\, \pi^n) \Big).
\end{aligned}$$

■

As a corollary to Lemma 1, we have

Corollary B.1 (Closed-form of the γ-Sinkhorn Step)**.** *The γ-Sinkhorn iterates have the density*

$$\pi^{n+1}_\gamma(\mathrm{d}x, \mathrm{d}y) \propto \mu(\mathrm{d}x) \left(\pi^{n+1/2}(\mathrm{d}y \mid x) \right)^{\gamma_n} (\pi^n(\mathrm{d}y \mid x))^{1-\gamma_n}.$$

Proof. Let us drop the γ for brevity. From Lemma 1 we know that

$$\pi^{n+1} = \arg\min \Big\{ \gamma_n H(\pi \,\|\, \pi^{n+1/2}) + (1-\gamma_n) H(\pi \,\|\, \pi^n) : \pi_{\mathcal{X}} = \mu \Big\}$$

By the chain rule of relative entropy, and considering couplings π with $\pi_{\mathcal{X}} = \mu$, we have

$$\begin{aligned}
&\gamma_n H(\pi \,\|\, \pi^{n+1/2}) + (1-\gamma_n) H(\pi \,\|\, \pi^n) \\
&= \text{constant} + \int \mu(\mathrm{d}x) \int \pi(\mathrm{d}y \mid x) \log \left(\left(\frac{\pi(y \mid x)}{\pi^{n+1/2}(y \mid x)} \right)^{\gamma_n} \left(\frac{\pi(y \mid x)}{\pi^n(y \mid x)} \right)^{1-\gamma_n} \right).
\end{aligned}$$

Thus, for each x, we have to set $\pi(\cdot \mid x)$ to the minimizer of

$$H(\pi(\cdot \mid x)) - \int \pi(\mathrm{d}y \mid x)\, \log \Big\{ \pi^{n+1/2}(y \mid x)^{1-\gamma_n} \pi^n(y \mid x)^{1-\gamma_n} \Big\},$$

which by standard optimality conditions in calculus of variations, we see that

$$\pi^{n+1}(\cdot \mid x) \propto \pi^{n+1/2}(y \mid x)^{1-\gamma_n} \pi^n(y \mid x)^{1-\gamma_n},$$

and the claim of the corollary follows. ■

Lemma 2. *The γ-Sinkhorn iterates π_γ^n in (Sink$_\gamma$) admit the representation*

$$\mathrm{d}\pi_\gamma^n = \exp\big(f_\gamma^n \oplus g_\gamma^n\big)\,\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}. \tag{11}$$

Moreover, the potentials g_γ^n satisfy the recursion

$$g_\gamma^{n+1} = g_\gamma^n - \gamma_n \log \frac{\mathrm{d}(\pi_\gamma^n)_{\mathcal{Y}}}{\mathrm{d}\nu},$$

and f_γ^{n+1} is computed from g_γ^{n+1} as in (9).

Proof. For easier readability, we drop the γ subscript of f_γ^n and g_γ^n. We prove this lemma by induction. For $n=0$, (11) holds because of the initialization of the iterations. Now suppose that $\frac{d\pi^n}{d\pi_\varepsilon^{\mathrm{ref}}} = \exp(f^n \oplus g^n)$. Then,

$$\pi^{n+1/2}(x,y) = \nu(y)\,\pi^n(x \mid y) = \frac{\mu(x)\exp\Big(f^n(x) - \frac{c(x,y)}{\varepsilon}\Big)\nu(y)}{\int \mu(x')\exp\Big(f^n(x') - \frac{c(x',y)}{\varepsilon}\Big)\,dx'} = \pi_\varepsilon^{\mathrm{ref}}(x,y)\exp\big(f_{n+1/2} \oplus g_{n+1/2}\big),$$

where $f^{n+1/2} = f^n$, and

$$g^{n+1/2}(y) = -\log \int \mu(x)\exp\left(f^n(x) - \frac{c(x,y)}{\varepsilon}\right)dx = -\log \frac{\pi_{\mathcal{Y}}^n(y)}{\nu(y)\exp(g^n(y))} = -\log\frac{\pi_{\mathcal{Y}}^n(y)}{\nu(y)} + g^n(y)$$

Now compute

$$\pi^n(y \mid x) = \frac{\exp\Big(g^n(y) - \frac{c(x,y)}{\varepsilon}\Big)\nu(y)}{\int \nu(y')\exp\Big(g^n(y') - \frac{c(x,y')}{\varepsilon}\Big)\,dy'} =: \frac{1}{A^n(x)}\exp\left(g^n(y) - \frac{c(x,y)}{\varepsilon}\right)\nu(y)$$

and

$$\pi^{n+1/2}(y \mid x) = \frac{\exp\Big(g^{n+1/2}(y) - \frac{c(x,y)}{\varepsilon}\Big)\nu(y)}{\int \nu(y')\exp\Big(g^{n+1/2}(y') - \frac{c(x,y')}{\varepsilon}\Big)\,dy'} =: \frac{1}{A^{n+1/2}(x)}\exp\left(g^{n+1/2}(y) - \frac{c(x,y)}{\varepsilon}\right)\nu(y).$$

Thus,

$$\pi^{n+1/2}(y \mid x)^\gamma \pi^n(y \mid x)^{1-\gamma} = \frac{1}{A^n(x)^{1-\gamma}A^{n+1/2}(x)^\gamma}\exp\left(\gamma g_{n+1/2}(y) + (1-\gamma)g_n(y) - \frac{c(x,y)}{\varepsilon}\right)\nu(y).$$

Recall that $\pi^{n+1}(dx,dy) = \mu(dx)\frac{1}{Z^n(x)}\pi^{n+1/2}(y \mid x)^\gamma \pi^n(y \mid x)^{1-\gamma}$. Putting the values computed above, and gathering all terms that only depend on x into f^{n+1} shows that

$$\pi^{n+1} = \exp\left(f^{n+1}(x) + g^{n+1}(y) - \frac{c(x,y)}{\varepsilon}\right)\mu(x)\nu(y),$$

where

$$g^{n+1} = \gamma g^{n+1/2} + (1-\gamma)g^n = -\gamma\log\frac{\pi_{\mathcal{Y}}^n}{\nu} + \gamma g^n + (1-\gamma)g^n = g^n - \gamma\log\frac{\pi_{\mathcal{Y}}^n}{\nu}. \qquad \blacksquare$$

B.2 Results about the Sinkhorn$_\varepsilon$ and Schrödinger$_\varepsilon$ Flows

Proposition 1. *Fix a coupling $\pi \in \{\pi : \pi_{\mathcal{X}} = \mu\}$. For any $\gamma > 0$, let $\pi^\gamma = \mathrm{S}_\gamma[\pi]$. Then,*

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0} \log \pi^\gamma(x,y) = -\log\frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}(y) + \mathbb{E}_{\pi(\cdot|x)}\left[\log\frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}\right].$$

Moreover, if $\mathrm{d}\pi = \exp(f \oplus g)\,\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}$, then for all $\gamma > 0$, $\mathrm{d}\pi^\gamma = \exp(f^\gamma \oplus g^\gamma)\,\mathrm{d}\pi_\varepsilon^{\mathrm{ref}}$, and f^γ and g^γ satisfy

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0} g^\gamma(y) = -\log\frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}(y),$$
$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0} f^\gamma(x) = \mathbb{E}_{\pi(\cdot|x)}\left[\log\frac{\mathrm{d}\pi_{\mathcal{Y}}}{\mathrm{d}\nu}\right].$$

Proof. Define $\pi^{1/2}(x,y) = \nu(y)\,\pi(x \mid y)$ and for $\gamma \geq 0$, by Corollary B.1

$$\pi^\gamma(x,y) = \mu(x)\frac{1}{Z^\gamma(x)}\pi^{1/2}(y \mid x)^\gamma \pi(y \mid x)^{1-\gamma},$$

where $Z^\gamma(x) = \int \pi^{1/2}(y \mid x)^\gamma \pi(y \mid x)^{1-\gamma}\,dy$ is the normalization factor. Note that $\pi^\gamma(y \mid x) = \pi^\gamma(x,y)/\mu(x)$. Our goal is to characterize the derivative $\frac{\mathrm{d}}{\mathrm{d}\gamma}\big|_{\gamma=0}\log \pi^\gamma(y \mid x)$ for all fixed x. For that, we compute

$$\frac{1}{\gamma}(\log \pi^\gamma(y \mid x) - \log \pi(y \mid x)) = -\frac{1}{\gamma}\log Z^\gamma(x) + \log\frac{\pi^{1/2}(y \mid x)}{\pi(y \mid x)}.$$

Thus, we only have to compute the limit

$$\lim_{\gamma\downarrow 0}\frac{1}{\gamma}\log Z^\gamma(x) = \frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0}\log Z^\gamma(x) = \frac{1}{Z^0(x)}\int \pi(y \mid x)\log\frac{\pi^{1/2}(y \mid x)}{\pi(y \mid x)}\,dy = -H(\pi(\cdot \mid x)\,\|\,\pi^{1/2}(\cdot \mid x)).$$

Thus,

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0}\log \pi^\gamma(y \mid x) = -\log\frac{\pi(y \mid x)}{\pi^{1/2}(y \mid x)} + H(\pi(\cdot \mid x)\,\|\,\pi^{1/2}(\cdot \mid x)).$$

Notice that since the first marginal of π is assumed to be fixed (to μ), we have

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0}\log \pi^\gamma(y \mid x) = \frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0}\log \pi^\gamma(x,y) = -\log\frac{\pi(y \mid x)}{\pi^{1/2}(y \mid x)} + H(\pi(\cdot \mid x)\,\|\,\pi^{1/2}(\cdot \mid x)).$$

We have the identity:

$$\log\frac{\pi(y \mid x)}{\pi^{1/2}(y \mid x)} = \log\frac{\pi_{\mathcal{Y}}(y)}{\nu(y)} + \log\frac{\pi^{1/2}_{\mathcal{X}}(x)}{\pi_{\mathcal{X}}(x)}.$$

With this, we can rewrite the evolution above as

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\gamma}\Big|_{\gamma=0}\log \pi^\gamma(x,y) &= -\log\frac{\pi_{\mathcal{Y}}(y)}{\nu(y)} - \log\frac{\pi^{1/2}_{\mathcal{X}}(x)}{\pi_{\mathcal{X}}(x)} + \int \pi(z \mid x)\left(\log\frac{\pi_{\mathcal{Y}}(z)}{\nu(z)} + \log\frac{\pi^{1/2}_{\mathcal{X}}(x)}{\pi_{\mathcal{X}}(x)}\right)dz \\ &= -\log\frac{\pi_{\mathcal{Y}}(y)}{\nu(y)} + \int \pi(z \mid x)\log\frac{\pi_{\mathcal{Y}}(z)}{\nu(z)}\,dz.\end{aligned}$$

This proves the first part of the proposition.

For the evolution of g^γ, recall the recursion in Lemma 2, and see that when $\gamma \to 0$, the limit of g^γ satisfies the evolution in the proposition. For the evolution of f^γ, observe that

$$\log \pi^\gamma = (f^\gamma \oplus g^\gamma) + \log \pi^{\mathrm{ref}}_\varepsilon$$

Thus,

$$\frac{\mathrm{d}}{\mathrm{d}\gamma}\log \pi^\gamma(x,y) = \frac{\mathrm{d}}{\mathrm{d}\gamma}f^\gamma(x) + \frac{\mathrm{d}}{\mathrm{d}\gamma}g^\gamma.$$

Replacing the formula for $\frac{\mathrm{d}}{\mathrm{d}\gamma}\log \pi^\gamma(x,y)$ and $\frac{\mathrm{d}}{\mathrm{d}\gamma}g^\gamma$ above, gives the formula for $\frac{\mathrm{d}}{\mathrm{d}\gamma}f^\gamma$. ■

Lemma 3. *The Fenchel conjugate φ^* of $\varphi + \mathrm{I}_{\mathcal{C}}$ evaluated at $h \in \mathcal{D}$ is given by $\varphi^*(h) = \langle\hat\pi, h\rangle - H(\hat\pi\,\|\,\pi^{\mathrm{ref}}_\varepsilon)$, where*

$$\hat\pi(x,y) \coloneqq \frac{\pi^{\mathrm{ref}}_\varepsilon(x,y)\,e^{h(x,y)}}{\int \pi^{\mathrm{ref}}_\varepsilon(x,y')\,e^{h(x,y')}\,dy'}\,\mu(x) \in \mathcal{C}. \tag{16}$$

Moreover, one has $\delta\varphi^*(h) = \hat\pi$ where $\hat\pi$ is defined in (16).

Proof. For $h \in \mathcal{D}$, we know that

$$\varphi^*(h) = \sup_{\pi\in\mathcal{C}}\langle\pi, h\rangle - \varphi(\pi).$$

For brevity, define $k(x, y) \coloneqq -h(x, y) - \log \pi_\varepsilon^{\mathrm{ref}}(x, y)$. Then,

$$\varphi^*(h) = -\inf_{\pi \in \mathcal{C}} \langle \pi, k \rangle + H(\pi),$$

where $H(\pi)$ is the entropy of π. Now defined the Lagrangian

$$L(\pi, \lambda, \psi) \coloneqq \langle \pi, k \rangle + H(\pi) + \lambda(\langle \pi, 1 \rangle - 1) + (\langle \pi, \psi \rangle - \langle \mu, \psi \rangle), \quad \text{with} \quad \lambda \in \mathbb{R}, \quad \psi : \mathbb{R}^d \to \mathbb{R}.$$

For a fixed λ and ψ, we have that the minimizer of the Lagrangian is

$$\pi(x, y) = e^{-k(x,y) - \psi(x) - \lambda - 1}.$$

Moreover, we have that

$$\begin{aligned}
\psi &= \arg\max_{\psi} \iint dx\, dy\, e^{-k(x,y)-\psi(x)-\lambda-1}(k(x, y) - k(x, y) - \psi(x) - \lambda - 1 + \lambda + \psi(x)) - \int \psi(x)\mu(x)\, dx \\
&= \arg\max_{\psi} -\iint dx\, dy\, e^{-k(x,y)-\psi(x)-\lambda-1} - \int \psi(x)\mu(x)\, dx \\
&= \arg\min_{\psi} \iint dx \left(\int dy\, e^{-k(x,y)} \right) e^{-\psi(x)-\lambda-1} + \psi(x)\mu(x),
\end{aligned}$$

which gives

$$\psi(x) = -\log \mu(x) + \log \int e^{-k(x,y)}\, dy - \lambda - 1,$$

implying that

$$\pi(x, y) = \frac{e^{-k(x,y)}}{\int e^{-k(x,y')}\, dy'} \mu(x) = \frac{\pi_\varepsilon^{\mathrm{ref}}(x, y)\, e^{h(x,y)}}{\int \pi_\varepsilon^{\mathrm{ref}}(x, y')\, e^{h(x,y')}\, dy'} \mu(x).$$

∎

Remark. An easy consequence of Lemma 3 is the following generalization of the classical dual isometry of Bregman divergences present in MD for Euclidean spaces to our setting.

Lemma B.1 (Dual Isometry of Bregman Divergences). *Let $\pi, \pi' \in \mathcal{C} = \{\pi : \pi_{\mathcal{X}} = \mu\}$ having densities with respect to $\pi_\varepsilon^{\mathrm{ref}}$. Then we have*

$$H(\pi \,\|\, \pi') = D_\varphi(\pi \,\|\, \pi') = D_{\varphi^*}(\delta\varphi(\pi') \,\|\, \delta\varphi(\pi)). \tag{B.3}$$

Proof. Write $\pi = e^h \pi_\varepsilon^{\mathrm{ref}}$ and $\pi' = e^{h'} \pi_\varepsilon^{\mathrm{ref}}$. We first note that, by (16) and the fact that $\delta\varphi(\pi) = h$, we have

$$\begin{aligned}
\delta\varphi^*(\delta\varphi(\pi))(x, y) &= \frac{\pi_\varepsilon^{\mathrm{ref}}(x, y)\, e^{h(x,y)}}{\int \pi_\varepsilon^{\mathrm{ref}}(x, y')\, e^{h(x,y')}\, dy'} \mu(x) \\
&= \frac{\pi(x, y)}{\int \pi(x, y')\, dy'} \mu(x) \\
&= \pi(x, y)
\end{aligned} \tag{B.4}$$

where the last equality follows from $\pi \in \mathcal{C}$. Similarly, we have $\delta\varphi^*(\delta\varphi(\pi')) = \pi'$. Using the formula for φ^* in Lemma 3, we have

$$\begin{aligned}
\varphi^*(h) &= \langle \pi, h \rangle - H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}}) \\
&= \int \pi \log \frac{e^h \pi_\varepsilon^{\mathrm{ref}}}{\pi_\varepsilon^{\mathrm{ref}}} - H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}}) \\
&= \int \pi \log \frac{\pi}{\pi_\varepsilon^{\mathrm{ref}}} - H(\pi \,\|\, \pi_\varepsilon^{\mathrm{ref}}) \\
&= 0,
\end{aligned} \tag{B.5}$$

and the exact same computation shows that $\varphi^*(h') = 0$. We therefore have

$$\begin{aligned} D_{\varphi^*}(\delta\varphi(\pi') \,\|\, \delta\varphi(\pi)) &= D_{\varphi^*}(h' \,\|\, h) \\ &= \varphi^*(h') - \varphi^*(h) - \langle \delta\varphi^*(h), h' - h \rangle \\ &= \langle \pi, h - h' \rangle \\ &= H(\pi \,\|\, \pi') \end{aligned} \tag{B.6}$$

where the third equality follows from (B.4) and (B.5). ■

Theorem 3.1. *The dynamics (17) or (18) coincide with the Schrödinger$_\varepsilon$ flow (14), and the corresponding $(\hat{\pi}^t_\varepsilon)_{t\geq 0}$ solves the Sinkhorn$_\varepsilon$ flow (13) starting at $\hat{\pi}^0_\varepsilon$.*

Proof. It is clear that (18) is equivalent to (17). Before stating the proof of the theorem, we state some useful facts about the Bregman potential and the first variation of its convex dual.

By the formula for $\delta\varphi^*$ in Lemma 3, we see that for any $h \in \mathcal{D}$ and any $f : \mathcal{X} \to \mathbb{R}$, we have $\delta\varphi^*(h+f) = \delta\varphi^*(h)$:

$$\delta\varphi^*(h+f)(x,y) = \mu(x)\frac{\pi^{\mathrm{ref}}_\varepsilon(x,y)e^{h(x,y)+f(x)}}{\int \pi^{\mathrm{ref}}_\varepsilon(x,y')e^{h(x,y')+f(x)}\,\mathrm{d}y'} = \mu(x)\frac{e^{f(x)}\pi^{\mathrm{ref}}_\varepsilon(x,y)e^{h(x,y)}}{e^{f(x)}\int \pi^{\mathrm{ref}}_\varepsilon(x,y')e^{h(x,y')}\,\mathrm{d}y'} = \delta\varphi^*(h)(x,y),$$

that is, if two functions in dual space $\mathcal{D}$ only differ by a function of x, they correspond to the same primal point.

Consider a coupling $\pi \in \Pi_{c,\varepsilon} \cap \mathcal{C}$ written as $\frac{\mathrm{d}\pi}{\mathrm{d}\pi^{\mathrm{ref}}_\varepsilon} = \exp(f \oplus g)$. By the definition of the Bregman potential $\varphi(\pi) = H(\pi \,\|\, \pi^{\mathrm{ref}}_\varepsilon)$, we see that

$$\delta\varphi(\pi) = \log\frac{\mathrm{d}\pi}{\mathrm{d}\pi^{\mathrm{ref}}_\varepsilon} = f \oplus g.$$

From the discussion above, $\pi = \delta\varphi^*(\delta\varphi(\pi)) = \delta\varphi^*(f \oplus g) = \delta\varphi^*(g)$.

Now, consider the flow (17) of h^t_ε, started at $h^0_\varepsilon = \delta\varphi(\hat{\pi}^0_\varepsilon) =: f^0_\varepsilon \oplus g^0_\varepsilon$. Note that since $\delta F(\cdot)$ is only a function of y, we have

$$h^t_\varepsilon = f^0_\varepsilon \oplus g^t_\varepsilon.$$

Noticing that (by construction) $\hat{\pi}^t_\varepsilon = \delta\varphi^*(h^t_\varepsilon) \in \Pi_{c,\varepsilon} \cap \mathcal{C}$ for all t, our previous discussion implies

$$\hat{\pi}^t_\varepsilon = \delta\varphi^*(h^t_\varepsilon) = \delta\varphi^*(f^0_\varepsilon \oplus g^t_\varepsilon) = \delta\varphi^*(g^t_\varepsilon).$$

Thus, in the evolution (17), if one only looks at the y variable, one gets

$$\frac{\mathrm{d}}{\mathrm{d}t} g^t_\varepsilon = -\delta F(\hat{\pi}^t_\varepsilon),$$

which is exactly the Schrödinger$_\varepsilon$ flow. ■

Theorem 3.2. *Starting from $\pi^0_\varepsilon \in \Pi_{c,\varepsilon} \cap \{\pi : \pi_{\mathcal{X}} = \mu\}$, consider the Sinkhorn$_\varepsilon$ flow π^t_ε and the corresponding Schrödinger$_\varepsilon$ flow g^t_ε. Then,*

$$F(\pi^t_\varepsilon) \leq \frac{D_{\varphi^*}(g^0_\varepsilon \,\|\, g^{\mathrm{opt}}_\varepsilon)}{t} = \mathcal{O}(t^{-1}),$$

where $g^{\mathrm{opt}}_\varepsilon$ is the Schrödinger potential of the optimal coupling for (OT$_\varepsilon$). That is, the $\mathcal{Y}$-marginal of π^t_ε converges (in relative entropy) to ν with the rate $1/t$.

Proof. For brevity, we drop the ε. First, we show that the objective function F is decreasing along the flow:

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t} F(\pi^t) \\ &= \left\langle \delta F(\pi^t), \pi^t \frac{\mathrm{d}}{\mathrm{d}t} \log \pi^t \right\rangle \\ &= -\iint \pi^t(x,y)\left(\log\frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y)\right)^2 \mathrm{d}x\,\mathrm{d}y + \iint \pi^t(x,y)\log\frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y)\int \pi^t(z \mid x)\log\frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(z)\,\mathrm{d}z\,\mathrm{d}x\,\mathrm{d}y \end{aligned}$$

Defining $k(x) = \int \pi^t(y \mid x) \log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y)\,\mathrm{d}y$, we see that the second term above writes

$$\begin{aligned}\iint \pi^t(x,y) \log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y) \int \pi^t(z \mid x) \log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(z)\,\mathrm{d}z\,\mathrm{d}x\,\mathrm{d}y &= \iint \pi^t(x,y) \log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y) k(x)\,\mathrm{d}x\,\mathrm{d}y \\ &= \iint \mu(x)\pi^t(y \mid x) \log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y) k(x)\,\mathrm{d}x\,\mathrm{d}y \\ &= \int \mu(x) k(x)^2\,\mathrm{d}x\end{aligned}$$

Now, by Jensen inequality, we have

$$\begin{aligned}\int \mu(x)k(x)^2\,\mathrm{d}x &= \int \mu(x)\left(\int \pi^t(y \mid x) \log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y)\,\mathrm{d}y\right)^2 \mathrm{d}x \\ &\leq \int \mu(x) \int \pi^t(y \mid x)\left(\log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y)\right)^2 \mathrm{d}y\,\mathrm{d}x \\ &= \iint \pi^t(x,y)\left(\log \frac{\mathrm{d}\pi^t_{\mathcal{Y}}}{\mathrm{d}\nu}(y)\right)^2 \mathrm{d}y\,\mathrm{d}x.\end{aligned}$$

We thus have shown that

$$\frac{\mathrm{d}}{\mathrm{d}t} F(\pi^t) \leq 0.$$

For a Schrödinger potential g and its corresponding coupling $\pi = \delta\varphi^*(g)$, define

$$L(g) = D_{\varphi^*}(g \,\|\, g^{\mathrm{opt}}) = \varphi^*(g) - \varphi^*(g^{\mathrm{opt}}) - \langle \delta\varphi^*(g^{\mathrm{opt}}), g - g^{\mathrm{opt}}\rangle = \varphi^*(g) - \varphi^*(g^{\mathrm{opt}}) - \langle \pi^{\mathrm{opt}}, g - g^{\mathrm{opt}}\rangle$$

and observe that $\delta L(g) = \delta\varphi^*(g) - \delta\varphi^*(g^{\mathrm{opt}}) = \pi - \pi^{\mathrm{opt}}$. We treat L as a Lyapunov function of the Schrödinger$_\varepsilon$ flow. For that, we compute

$$\frac{\mathrm{d}}{\mathrm{d}t} L(g^t) = \langle \delta L(g^t), \frac{\mathrm{d}}{\mathrm{d}t} g^t\rangle = -\langle \pi^t - \pi^{\mathrm{opt}}, \delta F(\pi^t)\rangle \leq F(\pi^{\mathrm{opt}}) - F(\pi^t), \tag{B.7}$$

where the inequality is due to convexity of F. Thus,

$$L(\pi^t) - L(\pi^0) = \int_0^t \frac{\mathrm{d}}{\mathrm{d}s} L(\pi^s)\,\mathrm{d}s \leq \int_0^t F(\pi^{\mathrm{opt}}) - F(\pi^s)\,\mathrm{d}s \leq t(F(\pi^{\mathrm{opt}}) - F(\pi^t)),$$

where the last inequality is due to the fact that $F(\pi^t)$ is non-increasing. Using the fact that $L \geq 0$, we obtain the result. ■

B.3 Guarantees on Noisy γ-Sinkhorn

Theorem 3.3. *Suppose that we have a stochastic estimate $\tilde{\delta}F$ of δF such that $\mathbb{E}[\tilde{\delta}F(\pi)] = \delta F(\pi)$ and $\mathbb{E}[\|\tilde{\delta}F(\pi)\|_\infty^2] \leq \sigma^2 < \infty$ for all π. Consider the iterations π^n_γ generated by (Sink$_\gamma$) using $\tilde{\delta}F$ and a fixed step-size γ. Then we have, with $\bar{\pi}^n_\gamma \coloneqq \frac{1}{n}\sum_{k=0}^n \pi^k_\gamma$,*

$$\mathbb{E}\big[H((\bar{\pi}^n_\gamma)_{\mathcal{Y}} \,\|\, \nu)\big] \leq \frac{H(\pi^{\mathrm{opt}} \,\|\, \pi^{\mathrm{ref}}_\varepsilon)}{\gamma n} + \gamma\sigma^2. \tag{19}$$

Proof. Since F is convex and 1-smooth relative to φ (Aubin-Frankowski et al., 2022), Hanzely and Richtárik (2021, **Lemma 5.2**) with $\mu = 0$, $L_t \leftarrow \frac{1}{\gamma}$, and $x \leftarrow \pi^{\mathrm{opt}}_\varepsilon$ gives

$$\mathbb{E}\big[F(\pi^{n+1}_\gamma)|\mathcal{F}_n\big] \leq \frac{1}{\gamma}\big(H(\pi^{\mathrm{opt}}_\varepsilon \,\|\, \pi^n_\gamma) - \mathbb{E}\big[H(\pi^{\mathrm{opt}}_\varepsilon \,\|\, \pi^{n+1}_\gamma)|\mathcal{F}_n\big]\big) + \gamma\sigma^2 \tag{B.8}$$

where $\mathcal{F}_n$ denotes the filtration generated by the stochastic algorithm up to step n. Taking expectation on both sides and summing over n, we get

$$\frac{1}{n}\sum_{k=0}^{n-1} \mathbb{E}\big[F(\pi^k_\gamma)\big] \leq \frac{H(\pi^{\mathrm{opt}}_\varepsilon \,\|\, \pi^{\mathrm{ref}}_\varepsilon)}{\gamma n} + \gamma\sigma^2. \tag{B.9}$$

The proof follows by using the convexity of $F(\cdot) \coloneqq H(\cdot \,\|\, \nu)$. ■

§ A formal statment for Theorem 3.4. We will now present a formal theorem addressing the convergence of the method described in (Sink_γ), taking into account a noisy and biased oracle denoted as δF. To accomplish this, we will rely on the framework of *stochastic approximation* (Benaïm, 1999; Karimi et al., 2022a; Mertikopoulos et al., 2023). Although this approach is well-established, it involves some technical complexities, so we have chosen to defer the details to this appendix.

Let $(\pi^n)_{n\in\mathbb{N}}$ be the sequence of measures generated by (Sink_γ) with a noisy oracle $\tilde{\delta}F$ and step-sizes $(\gamma_n)_{n\in\mathbb{N}}$, and let $(g^n)_{n\in\mathbb{N}}$ be its corresponding Schrödinger potentials. We first define the "effective time" τ^n to be $\tau^n := \sum_{k=1}^n \gamma_k$, which is the time that has elapsed at the n-th iteration of the discrete-time process g^n. Using τ^n, we consider the *continuous-time interpolation* $g(t)$ of g^n:

$$g(t) := g^n + \frac{t - \tau^n}{\tau^{n+1} - \tau^n}(g^{n+1} - g^n). \tag{B.10}$$

Note that each $g(t)$ is a function on $\mathcal{Y}$. The following assumption is standard in the stochastic approximation literature:

Assumption B.1. *Let π^n and g^n be given as above. We assume that (i) δF is Lipschitz and bounded on a neighborhood of $(\pi^n)_{n\in\mathbb{N}}$, and (ii) $(g(t))_{t\geq 0}$ is a* ***precompact*** *set in the topology of L^∞.*

It is worth highlighting that Assumption B.1 is a relatively mild technical condition that finds applicability in a wide range of practical scenarios. For example, it remains satisfied when employing bounded and Hölder continuous neural networks to parameterize distributions with compact support; see, e.g., (Seguy et al., 2017).

We are now ready to state the formal version of Theorem 3.4.

Theorem B.1. *Let π^n and g^n be given as above such that Assumption B.1 holds. Suppose that the step-size rule γ_n is such that $\sum \gamma_n = \infty$ and $\sum \gamma_n^2 < \infty$. Denote by $\mathcal{F}_n$ the filtration of the stochastic algorithm up to iteration n, and its martingale noise and bias by*

$$\begin{aligned}
\lambda^n &:= \mathbb{E}[\tilde{\delta}F(\pi^n) \,|\, \mathcal{F}_n] - \delta F(\pi^n), \\
\omega^n &:= \tilde{\delta}F(\pi^n) - \mathbb{E}[\tilde{\delta}F(\pi^n) \,|\, \mathcal{F}_n].
\end{aligned}$$

Then $\mathbb{P}(\lim_{n\to\infty} \pi^n = \pi^{\mathrm{opt}}) = 1$ if the following holds almost surely:

$$\lim_{n\to\infty} \|\lambda^n\|_\infty = 0, \quad \textit{and} \quad \sup_n \mathbb{E}\left[\|\omega^n\|_\infty^2\right] \leq \sigma^2 < \infty. \tag{B.11}$$

Proof. The proof of the theorem is established through a combination of well-established results in stochastic approximation theory with our continuous-time framework.

Assumption B.1 and (B.11) ensure that $g(\cdot)$ is a precompact *asymptotic pseudo-trajectory* of the associated continuous-time $\text{Schrödinger}_\varepsilon$ flow given in (14); see e.g., (Benaïm, 1999, **Proposition 4.1**). It follows from this association that the iterates $(g^n)_{n\geq 0}$ converge almost surely to an *internally chain-transitive* (ICT) set of the $\text{Schrödinger}_\varepsilon$ flow. On the other hand, within the course of our proof for Theorem 3.2, we have established the existence of a Lyapunov function for the $\text{Schrödinger}_\varepsilon$ flow; see (B.7). Consequently, the only possible ICT set is identified as the singleton set $\{g_\varepsilon^{\mathrm{opt}}\}$ (Benaïm, 1999, **Proposition 6.4**). This, in turn, implies that the following event happens almost surely:

$$\begin{aligned}
\lim_{n\to\infty} \pi^n &= \lim_{n\to\infty} \delta F^*(g^n) \\
&= \delta F^*(g_\varepsilon^{\mathrm{opt}}) \\
&= \pi_\varepsilon^{\mathrm{opt}}.
\end{aligned}$$

■

C On Schrödinger Bridges

C.1 Proof of Proposition 2

First, we need some lemmas.

Lemma C.1. *Suppose* $\mathsf{P} \in \mathcal{P}(\Omega)$ *and* $\mathsf{Q} \in \mathcal{M}(\Omega)$ *be a finite measure. For* $F(\mathsf{P}) = H(\mathsf{P}_T \,\|\, \mu_T)$ *we have*

$$F'(\mathsf{P}; \mathsf{Q}) = \int_{\mathbb{R}^d} \mathrm{d}\mathsf{Q}_T \log \frac{\mathrm{d}\mathsf{P}_T}{\mathrm{d}\mu_T}.$$

Proof. As the function F only depends on the marginals at time 1, the claim follows from a similar calculation as in (Aubin-Frankowski et al., 2022, Prop. 5). ∎

Lemma C.2. *For* $\varphi(\mathsf{P}) = H(\mathsf{P} \,\|\, \mathsf{P}^{\mathrm{ref}})$*, we have* $D_\varphi(\mathsf{P} \,\|\, \mathsf{Q}) = H(\mathsf{P} \,\|\, \mathsf{Q})$.

Proof. Similar to (Aubin-Frankowski et al., 2022, Example 2). ∎

Lemma C.3. *For the iteration* (IPF)*, it holds that*

$$\frac{\mathrm{d}\mathsf{P}^{n+1/2}}{\mathrm{d}\mathsf{P}^n} = \frac{\mathrm{d}\mu_T}{\mathrm{d}\mathsf{P}^n_T}.$$

Proof. By the chain rule of KL divergence, we know that

$$H(\mathsf{P} \,\|\, \mathsf{P}^n) = H(\mathsf{P}_T \,\|\, \mathsf{P}^n_T) + \int H(\mathsf{P}(\cdot \mid X_T = x) \,\|\, \mathsf{P}^n(\cdot \mid X_T = x))\, \mathrm{d}\mathsf{P}_T(x)$$

Notice that in the first part of the iteration (IPF), the last marginal is fixed, thus, the first term above is fixed, and the minimizer shall be

$$\mathsf{P}^{n+1/2}(\cdot) = \int \mathrm{d}\mu_T(x)\, \mathsf{P}^n(\cdot \mid X_T = x).$$

From this representation, the claim of the lemma is clear. ∎

We can now go ahead and prove Proposition 2, stated below for convenience.

Proposition 2. *The iterations* P^n *of* (IPF) *satisfy*

$$\mathsf{P}^{n+1} = \arg\min_{\mathsf{P} \in \mathcal{C}} \{\langle \delta F(\mathsf{P}^n), \mathsf{P} - \mathsf{P}^n \rangle + D_\varphi(\mathsf{P} \,\|\, \mathsf{P}^n)\}, \tag{26}$$

with $F(\mathsf{P}) \coloneqq H(\mathsf{P}_T \,\|\, \mu_T)$*,* $\varphi(\mathsf{P}) \coloneqq H(\mathsf{P} \,\|\, \mathsf{P}^{\mathrm{ref}})$*, and* $\mathcal{C} \coloneqq \{\mathsf{P} : \mathsf{P}_0 = \mu_0\}$.

Proof. For a path measure P, compute the following:

$$\begin{aligned}
&F(\mathsf{P}^n) + F'(\mathsf{P}^n; \mathsf{P} - \mathsf{P}^n) + D_\varphi(\mathsf{P} \,\|\, \mathsf{P}^n) \\
&= H(\mathsf{P}^n_T \,\|\, \mu_T) + \int_{\mathbb{R}^d} d(\mathsf{P} - \mathsf{P}^n)_T \log \frac{d\mathsf{P}^n_T}{d\mu_T} + H(\mathsf{P} \,\|\, \mathsf{P}^n) && \text{by Lemmas C.1 and C.2} \\
&= \int_{\mathbb{R}^d} d\mathsf{P}_T \log \frac{d\mathsf{P}^n_T}{d\mu_T} + H(\mathsf{P} \,\|\, \mathsf{P}^n) \\
&= \int_\Omega d\mathsf{P} \log \frac{d\mathsf{P}^n_T}{d\mu_T} + H(\mathsf{P} \,\|\, \mathsf{P}^n) \\
&= \int_\Omega d\mathsf{P} \log \frac{d\mathsf{P}^n}{d\mathsf{P}^{n+1/2}} + H(\mathsf{P} \,\|\, \mathsf{P}^n) && \text{by Lemma C.3} \\
&= \int_\Omega d\mathsf{P} \log \left\{ \frac{d\mathsf{P}^n}{d\mathsf{P}^{n+1/2}} \cdot \frac{d\mathsf{P}}{d\mathsf{P}^n} \right\} \\
&= H(\mathsf{P} \,\|\, \mathsf{P}^{n+1/2}).
\end{aligned}$$

Now it is clear that the minimizer of the above in the set $\mathcal{C}$ is exactly P^{n+1}. ∎

C.2 SDE Representation and the Drift Formula

Theorem 4.1. *Suppose that* P^n *is an SDE given by*

$$\mathrm{d}X_t^n = v_t^n(X_t^n)\,\mathrm{d}t + \sigma \mathrm{d}W_t,\ X_0 \sim \mu_0, \tag{28}$$

and that the time-reversal of $\mathsf{P}^{n+1/2}$ *is given by*

$$\mathrm{d}Y_t^{n+1/2} = w_{T-t}^{n+1/2}(Y_t^{n+1/2})\,\mathrm{d}t + \sigma \mathrm{d}W_t,\ Y_0^{n+1/2} \sim \mu_T. \tag{29}$$

Then the drift vector field $w_t^{n+1/2}$ *satisfies:*

$$-v_t^n(x) + \sigma^2 \nabla \log p_t^n(x) = w_t^{n+1/2}(x) \tag{30}$$

where p_t^n *is the density of* P_t^n.

Proof. First, letting $\hat{\mathsf{P}}^n$ to denote the law of the time-reversal of P^n, observe that since the time reversal of $\mathsf{P}^{n+1/2}$ solves $\arg\min\{H(\mathsf{P} \,\|\, \hat{\mathsf{P}}^n) : \mathsf{P}_0 = \mu_T\}$, its SDE representation is the same as the one for $\hat{\mathsf{P}}^n$, and only its initial datum is set to be μ_T. By the time reversal formula (Theorem A.1), $\hat{\mathsf{P}}^n$ corresponds to

$$dY_t^n = \left\{-v_{T-t}^n(Y_t^n) + \sigma^2 \nabla \log p_{T-t}^n(Y_t^n)\right\} dt + \sigma dW_t, \qquad Y_T^n \sim \mu_0,$$

where p_t^n is the density of P_t^n. This means that this should coincide with the SDE for time reversal of $\mathsf{P}^{n+1/2}$, that is

$$-v_t^n(x) + \sigma^2 \nabla \log p_t^n(x) = w_t^{n+1/2}(x). \qquad \blacksquare$$

Theorem 4.2. *Let* P^n *be given by the scheme (*γ*-IPF), and let* $v_t^n(\cdot)$ *be the (forward) vector field corresponding to the SDE representation of* P^n *in (28). Then* $v_t^n(\cdot)$ *satisfies the following recursive formula:*

$$v_t^{n+1} = v_t^n + \gamma\sigma^2 \nabla \log \ell_t^n - \sigma^2 \nabla V_t, \tag{SDE$_\gamma$}$$

where

$$V_t(x) = -\log \mathbb{E}\left[e^{-\frac{\sigma^2\gamma(1-\gamma)}{2}\int_t^T |\nabla \log \ell_s^n(Y_s)|^2\,\mathrm{d}s}\right], \tag{31}$$

and the expectation is with respect to the law of the SDE $(Y_s)_{s\geq t}$ *starting at* $Y_t = x$ *and following*

$$\mathrm{d}Y_s = \left\{v_s^n(Y_s) + \gamma\sigma^2 \nabla \log \ell_s^n(Y_s)\right\} \mathrm{d}s + \sigma \mathrm{d}W_s. \tag{32}$$

Proof. Note that the path measure $\mathsf{P}^{n+1/2}$ corresponds to the reversal of (29), which is a process with drift $v^{n+1/2} := -w_t^{n+1/2} + \sigma^2 \nabla \log p_t^{n+1/2}$, with $p_t^{n+1/2}$ being the density of $\mathsf{P}_t^{n+1/2}$. Recall that P^{n+1} is the solution to the minimization

$$\mathsf{P}^{n+1} = \underset{\mathsf{P}\in\mathcal{C}}{\arg\min}\left\{\gamma_n\, H(\mathsf{P} \,\|\, \mathsf{P}^{n+1/2}) + (1-\gamma_n)\, H(\mathsf{P} \,\|\, \mathsf{P}^n)\right\},$$

and suppose that it corresponds to the SDE

$$\mathrm{d}X_t^u = (b_t^\gamma(X_t^u) + u_t)\,\mathrm{d}t + \sigma \mathrm{d}W_t, \tag{C.1}$$

with $X_0 \sim \mu_0$, where we define the drift b_t^γ as

$$b_t^\gamma := \gamma v_t^{n+1/2} + (1-\gamma) v_t^n = \gamma \cdot (-w_t^{n+1/2} + \sigma^2 \nabla \log p_t^{n+1/2}) + (1-\gamma)\cdot v_t^n = v_t^n + \gamma\sigma^2 \nabla \log \ell_t^n$$

by Theorem 4.1. The reason that we take (C.1) as an SDE representation of P^{n+1} is that, firstly, it should be a diffusion with the same diffusion coefficient, and its drift shall be the "weighted average" of the drifts of $\mathsf{P}^{n+1/2}$ and P^n, with some correction u_t.

It turns out that characterizing u_t corresponds to solving a stochastic optimal control problem. Concretely, by the Girsanov theorem (see Corollary A.1), we obtain

$$\begin{aligned}
&\gamma H(\mathsf{P} \,\|\, \mathsf{P}^{n+1/2}) + (1-\gamma) H(\mathsf{P} \,\|\, \mathsf{P}^{n}) \\
&\quad = \text{constant} + \mathbb{E}_{\mathsf{P}}\left[\frac{1}{2\sigma^2}\int_0^T \left\{\gamma |u_t + b_t^\gamma(X_t) - v_t^{n+1/2}(X_t)|^2 + (1-\gamma)|u_t + b_t^\gamma(X_t) - v_t^n(X_t)|^2\right\} \mathrm{d}t\right] \\
&\quad = \text{constant} + \frac{1}{\sigma^2}\mathbb{E}_{\mathsf{P}}\left[\frac{1}{2}\int_0^T |u_t|^2 \,\mathrm{d}t + \frac{\gamma(1-\gamma)}{2}\int_0^T |v_t^{n+1/2}(X_t) - v_t^n(X_t)|^2 \,\mathrm{d}t\right] \\
&\quad = \text{constant} + \frac{1}{\sigma^2}\mathbb{E}_{\mathsf{P}}\left[\int_0^T \frac{1}{2}|u_t|^2 \,\mathrm{d}t + \frac{\sigma^4\gamma(1-\gamma)}{2}\int_0^T |\nabla \log \ell_t^n(X_t)|^2 \,\mathrm{d}t\right]
\end{aligned}$$

where the constant is the weighted sum of KL divergences for time marginals at 0, and is some fixed number (as we fixed the initial distribution of P). Now we recognize that the minimization problem (γ-IPF) is a stochastic optimal control problem with running cost

$$r(t, x, \alpha) = \frac{1}{2}|\alpha|^2 + \frac{\sigma^4\gamma(1-\gamma)}{2}|\nabla \log \ell_t^n(x)|^2,$$

the cost functional $J[u] = \mathbb{E}_{\mathsf{P}}\left[\int_0^t r(t, X_t^u, u_t)\, dt\right]$, and zero terminal cost. By Proposition C.1 below (setting $c_t = \frac{1}{2}\sigma^4\gamma(1-\gamma)|\nabla \log \ell_t^n|^2$), the value function $V_t(x)$ of this control problem is

$$\begin{aligned}
V_t(x) &= -\sigma^2 \log \mathbb{E}^{t,x}\left[\exp\left(-\frac{1}{\sigma^2}\int_t^T \frac{1}{2}\sigma^4\gamma(1-\gamma)|\nabla \log \ell_s^n(Y_s)|^2 \,\mathrm{d}s\right)\right] \\
&= -\sigma^2 \log \mathbb{E}^{t,x}\left[\exp\left(-\frac{\sigma^2\gamma(1-\gamma)}{2}\int_t^T |\nabla \log \ell_s^n(Y_s)|^2 \,\mathrm{d}s\right)\right]
\end{aligned}$$

where $\mathbb{E}^{t,x}$ is expectation with respect to the law of the process Y, given that it starts at x at time t; $Y_t = x$, and the optimal control $u_t = -\nabla V_t$. Thus, denoting the drift of P^{n+1} as v^{n+1}, we see that P^{n+1} is the law of the SDE

$$\mathrm{d}X_t = v_t^{n+1}(X_t)\,\mathrm{d}t + \sigma \mathrm{d}W_t, \quad X_0 \sim \mu_0,$$

with

$$\begin{aligned}
v_t^{n+1}(x) &= b_t^\gamma(x) + u_t(x) \\
&= v_t^n(x) + \gamma\sigma^2 \nabla \log \ell^n(x) + u_t(x) \\
&= v_t^n(x) + \gamma\sigma^2 \nabla \log \ell^n(x) + \sigma^2 \nabla \log \mathbb{E}^{t,x}\left[\exp\left(-\frac{\sigma^2\gamma(1-\gamma)}{2}\int_t^T |\nabla \log \ell_s^n(Y_s)|^2 \,\mathrm{d}s\right)\right],
\end{aligned}$$

where $(Y_s)_{s\geq t}$ follows the SDE

$$\mathrm{d}Y_s = \left\{v_s^n(Y_s) + \gamma\sigma^2 \nabla \log \ell_s^n(Y_s)\right\} \mathrm{d}s + \sigma \mathrm{d}W_s, \quad Y_t = x. \qquad \blacksquare$$

We used the following result regarding computation of the value function of a specific stochastic optimal control problem in Theorem 4.2.

Proposition C.1. *Let b_t be a given drift, and consider the controlled SDE*

$$dX_t^u = (b_t(X_t^u) + u_t)\, dt + \sigma dW_t$$

along with the following stochastic optimal control problem:

$$\min_u J[u] \coloneqq \mathbb{E}\left[\int_0^T \tfrac{1}{2}|u_t|^2 + c_t(X_t^u)\, dt\right].$$

Then, the value function is given by

$$V_s(x) = -\sigma^2 \log \mathbb{E}^{s,x}\left[\exp\left(-\frac{1}{\sigma^2}\int_s^T c_t(Y_t)\,dt\right)\right],$$

where Y_t is the solution of the uncontrolled SDE $dY_t = b_t(Y_t)\,dt + \sigma dW_t$. Moreover, the optimal control is of feedback type $u_t(x) = -\nabla V_t(x)$.

Proof. The value function $V_t(x)$ of this control problem shall satisfy the HJB equation, which writes

$$\partial_t V_t(x) + \min_{\alpha\in\mathbb{R}^d}\left\{\langle b_t(x)+\alpha, \nabla V_t(x)\rangle + \frac{\sigma^2}{2}\Delta V_t(x) + \tfrac{1}{2}|\alpha|^2 + c_t(x)\right\} = 0, \quad V_T(x) = 0,$$

and evaluates to the optimal control $\alpha^*(t,x) = -\nabla V_t(x)$. Replacing α with its optimal value gives

$$\partial_t V_t(x) - \frac{1}{2}|\nabla V_t(x)|^2 + \frac{\sigma^2}{2}\Delta V_t(x) + \langle b_t(x), \nabla V_t(x)\rangle + c_t(x) = 0, \quad V_T(x) = 0.$$

Inspired by the Fleming log transform (Fleming, 1977), make the change of variables $V_t(x) = -\sigma^2 \log E_t(x)$, and observe that the equation above becomes

$$\partial_t E_t(x) + \frac{\sigma^2}{2}\Delta E_t(x) + \langle b_t(x), \nabla E_t(x)\rangle = \frac{1}{\sigma^2}E_t(x)\,c_t(x), \quad E_T(x) = 1.$$

Following a similar argument as in (Pra and Pavon, 1990), consider the uncontrolled diffusion process Y_t with generator $\frac{\sigma^2}{2}\Delta + b_t\cdot\nabla$. By the Feynman-Kac formula,

$$E_s(x) = \mathbb{E}^{s,x}\left[\exp\left(-\frac{1}{\sigma^2}\int_s^T c_t(Y_t)\,dt\right)\right],$$

where $\mathbb{E}^{s,x}$ means expectation with respect to the law of the process Y, given that it starts at x at time s. ■

C.3 On the Implementation of the γ-IPF iteration

Although our paper focuses on the theoretical understanding of the Sinkhorn and IPF iterates, we briefly remark that the formula in Theorem 4.2 admits a practically efficient implementation. To see this, notice that the $\nabla \log \ell_t^n$ term in (SDE$_\gamma$) is the standard *Stein score* ratio that can be estimated by various diffusion models and is present in most practical training procedures of SB. On the other hand, the computation of the additional term V_t in (SDE$_\gamma$) is facilitated by the following connection to stochastic optimal control, whose proof is already present in the proof of Theorem 4.2.

Proposition C.2. *The minimization (γ-IPF) is equivalent to solving the following stochastic optimal control problem: Consider the following controlled SDE with drift $b_t = v_t^n + \gamma\nabla\log\ell_t^n$ and control u_t:*

$$\mathrm{d}X_t^u = (b_t(X_t^u) + u_t)\,\mathrm{d}t + \sigma \mathrm{d}W_t, \tag{C.2}$$

and the cost functional

$$J[u] \coloneqq \mathbb{E}\left[\int_0^T \frac{1}{2}|u_t|^2 + c_t(X_t^u)\,dt\right] \tag{C.3}$$

with $c_t = \frac{1}{2}\sigma^2\gamma(1-\gamma)|\nabla\log\ell_t^n|^2$. Then, the value function of the stochastic optimal control problem $\min_u J[u]$ is $\sigma^2 V_t$, where V_t is defined in (31) and the optimal control is $u^\star = -\sigma^2\nabla V_t$.

It turns out that the value function V_t can also be evaluated as an expectation that involves only the law of the standard Brownian motion:

Lemma C.4. *The function V_t in (31) also satisfies*

$$V_t(x) = -\log\mathbb{E}\left[\exp\left(\frac{1}{\sigma}\int_t^T \langle b_s(x+\sigma W_{s-t}), \mathrm{d}W_{s-t}\rangle - \frac{1}{2\sigma^2}\int_t^T |b_s(x+\sigma W_{s-t})|^2 + c_s(x+\sigma W_{s-t})\,\mathrm{d}s\right)\right],$$

where the expectation is with respect to the standard Brownian motion $(W_t)_{t\geq 0}$.

Proof. Proof is a simple application of the Girsanov theorem (Theorem A.2), applied to the SDEs (C.2) and σW_t, where W_t is a standard Brownian motion. ■

Proposition C.2 and Lemma C.4 furnishes at least two different ways of computing V_t in (SDE$_\gamma$). First, given $\nabla \log \ell_t^n$ which is given by the usual score matching procedure in SB training, one can compute the value function using standard approximation techniques in control theory for integration with respect to standard Brownian motion (Zhang and Chen, 2022). Alternatively, another common practice is to connect the value function via the *Feynman-Kac formula* to SDEs with *killing*. Concretely, since the cost c_t above is non-negative, one can simulate the uncontrolled SDE (32), and kill it at a rate $\frac{1}{\sigma^2} c_t = \frac{\gamma(1-\gamma)}{2} |\nabla \log \ell_t^n|^2$, that is,

$$\mathbb{P}[Y_{t+h} \text{ is killed} \mid Y_t] = \frac{\gamma(1-\gamma)}{2} |\nabla \log \ell_t^n(Y_t)|^2 + o(h).$$

These procedures are already employed in the SB community in other contexts (Liu et al., 2022; Pariset et al., 2023).