On Schrödinger Bridge Matching and Expectation Maximization

Rob Brekelmans
Vector Institute

Kirill Neklyudov
Vector Institute

Abstract

In this work, we analyze methods for solving the Schrödinger Bridge problem from the perspective of alternating KL divergence minimization. While existing methods such as Iterative Proportional or Iterative Markovian Fitting require exact updates due to the fact that each step optimizes a different objective, we propose a joint optimization of a single KL divergence objective which is motivated using tools from information geometry. As in the variational EM algorithm, this allows for inexact or stochastic gradient updates to decrease a unified objective. We highlight connections with related bridge matching, flow matching, and few-step generative modeling approaches, where various parameterizations of the coupling distributions are contextualized from the perspective of marginal-preserving inference.

1 Introduction

Optimal mass transport problems have a rich history [15, 31] and find wide-ranging applications throughout machine learning, ranging from generative modeling [33, 23] to predicting the evolution of biological cells in single-cell RNA sequencing [30]. Fundamentally, we are interested in learning a transport map or dynamical process to transform samples from an initial distribution to samples from a final, target distribution. Entropic regularization of the transport problem introduces stochasticity into these mappings, and also leads to computational benefits using the famous Sinkhorn algorithm [8]. While scaling these transport methods to high dimensional, continuous spaces remains a challenge, recent successful approaches in score-based generative modeling [33, 12] have been shown to be related to the dynamical entropy-regularized problem known as Schrödinger Bridge (SB) [9, 19, 5].

In this work, we interpret the Sinkhorn, or Iterative Proportional Fitting (IMF), algorithm for static entropy-regularized OT [8, 3, 17] and the recent Iterative Markovian Fitting (IMF) algorithm for the dynamical Schrödinger Bridge problem [32, 27] from the perspective of information geometry. While it is clear that each of these approaches perform alternating KL divergence projections onto sets satisfying desirable properties [3, 32], the asymmetry of the KL divergence raises the question of which order of arguments should be used for each projection [6, 7]. Furthermore, both IPF and IMF involve minimization over the same argument of the divergence in each iteration, which means that a different objective is optimized at each step and exact iterations are required to ensure convergence.

We address these questions by characterizing the solution to the Schrödinger Bridge problem as the intersection of four sets of measures (Prop. 2.2), and highlighting the properties of these sets which are relevant for KL divergence projection (Sec. 3). Motivated by this analysis, we propose a novel alternating projection algorithm in Sec. 4 which, in contrast to IPF and IMF, minimizes a unified KL divergence objective across iterations. This provides justification for the inexact, variational updates which are inevitably necessary in practice [26]. Our analysis sheds new light on recent bridge- and flow-matching approaches, which update couplings of the initial and final distributions and learn a vector field parameterizing the dynamical transport (Sec. 5), and clarifies the need for expressive, marginal-preserving parameterizations of the couplings which are amenable to joint optimization.

Accepted to the Optimal Transport and Machine Learning Workshop at NeurIPS 2023.

2 Background

Entropic OT and Schrödinger Bridge We begin by introducing the entropy-regularized optimal transport (OT) problem between two given measures with finite second moment $\mu_0, \nu_T \in \mathcal{P}_2(\mathbb{R}^d)$. The OT problem minimizes the cost over couplings $\mathbb{Q}_{0,T} \in \Pi(\mu_0, \nu_T) := \{\mathbb{Q}_{0,T} | \mathbb{Q}_0 = \mu_0, \mathbb{Q}_T = \nu_T\}$, or joint measures with the desired endpoint marginals,

$$W^{\mathrm{KL}}_{c,\epsilon}(\mu_0, \nu_T) = \inf_{\mathbb{Q}_{0,T} \in \Pi(\mu_0,\nu_T)} \int c(x_0, x_T)\, d\mathbb{Q}_{0,T} - \epsilon\, \mathcal{H}[\mathbb{Q}_{0,T}] = \inf_{\mathbb{Q}_{0,T} \in \Pi_{(\mu_0,\nu_T)}} \epsilon\, D_{\mathrm{KL}}[\mathbb{Q}_{0,T} : e^{-\frac{1}{\epsilon} c(x_0,x_T)}] \tag{eOT}$$

where $\mathcal{H}[\mathbb{Q}_{0,T}]$ is the Shannon entropy of the coupling.

Closely related is the Schrödinger Bridge (SB) problem, which seeks a stochastic process $\mathbb{Q}_{0:T}$ that matches the given endpoint marginals (μ_0, ν_T) while minimizing the KL divergence with a reference Brownian diffusion process $\mathbb{Q}^{\text{ref}}_{0:T}$ [4, 5, 19],

$$\mathrm{SB}(\mu_0, \nu_T) = \inf_{\mathbb{Q}_{0:T} \in \Pi(\mu_0,\nu_T)} D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{Q}^{\text{ref}}_{0:T}] = \inf_{\mathbb{Q}_{0:T} \in \Pi(\mu_0,\nu_T)} \int \log \frac{d\mathbb{Q}_{0:T}}{d\mathbb{Q}^{\text{ref}}_{0:T}}\, d\mathbb{Q}_{0:T} \tag{SB}$$

where $\mathbb{Q}_{0:T}$ and $\mathbb{Q}^{\text{ref}}_{0:T}$ are measures on the space of continuous paths $\mathcal{C} : [0, T] \to \mathbb{R}^d$ and our naming of $\mathbb{Q}_{0:T}, \mathbb{Q}^{\text{ref}}_{0:T}$ is inspired by later connections with variational inference. In what follows, we consider $x_T \sim \nu_T$ as the data distribution and $x_0 \sim \mu_0$ as, for example, a noise distribution.

Equivalence of eOT and SB Assume a reference $\mathbb{Q}^{\text{ref}}_{0:T}$ which induces an endpoint coupling $\mathbb{Q}^{\text{ref}}_{0,T} := e^{-\frac{1}{\epsilon} c(x_0,x_T)}$ ([18] Sec. 3), as is natural for the Euclidean cost and $\mathbb{Q}^{\text{ref}}_{0:T}$ as pure Brownian motion (see Ex. 2.3 below). In this case, we can relate the solutions to the dynamical Eq. (SB) and static Eq. (eOT) problems by decomposing the path space KL divergence using disintegration,

$$\mathrm{SB}_{c,\epsilon}(\mu_0, \nu_T) = \inf_{\mathbb{Q}_{0:T} \in \Pi(\mu_0,\nu_T)} D_{\mathrm{KL}}\Big[\mathbb{Q}_{0,T} : e^{-\frac{1}{\epsilon} c(x_0,x_T)}\Big] + \mathbb{E}_{\mathbb{Q}_{0,T}}\Big[D_{\mathrm{KL}}[\mathbb{Q}_{\circ|0,T} : \mathbb{Q}^{\text{ref}}_{\circ|0,T}]\Big] \tag{1}$$

where $\mathbb{Q}^{\text{ref}}_{\circ|0,T}(\cdot|x_0, x_T)$ denotes the conditional path measure on $t \in (0, T)$ given the endpoints. Note that the constraint $\mathbb{Q}_{0,T} \in \Pi(\mu_0, \nu_T)$ depends only on the endpoint coupling $\mathbb{Q}_{0,T}$, which means we can bring the second term to zero in the optimal solution if $\mathbb{Q}^*_{\circ|0,T} = \mathbb{Q}^{\text{ref}}_{\circ|0,T}$. Since only the optimization over $\mathbb{Q}_{0,T}$ remains in Eq. (1), we conclude that the optimal couplings in Eq. (eOT) and Eq. (SB) coincide, with $W^{\mathrm{KL}}_{c,\epsilon}(\mu_0, \nu_T) = \epsilon\, \mathrm{SB}_{c,\epsilon}(\mu_0, \nu_T)$.

Characterizing the Solution of Schrödinger Bridge Problem To refer to the optimality condition $\mathbb{Q}^*_{\circ|0,T} = \mathbb{Q}^{\text{ref}}_{\circ|0,T}$, we recall the definition of the *reciprocal class* of a reference path measure [20, 32]. A member of the reciprocal class may also be described as a 'mixture of bridges' as in [24, 32, 27].

Definition 2.1 (Reciprocal Class). *The reciprocal class of a reference process $\mathbb{Q}^{ref}_{0:T}$ is the set of measures $\mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T}) = \{\Pi_{0:T} |\, \Pi_{0:T} = \Pi_{0,T}\mathbb{Q}^{ref}_{\circ|0,T}\}$, where $\mathbb{Q}^{ref}_{\circ|0,T}$ is the 'bridge' process obtained by conditioning $\mathbb{Q}^{ref}_{0:T}$ on its endpoint values $X_0 = x_0, X_T = x_T$.*

Along with the reciprocal class, we will consider the set of Markov path measures $\mathcal{M}$. We defer detailed definitions of reciprocal and Markov path measures to App. A.

The focal point of the current work is the following proposition, which states that, under mild conditions [19], the unique solution $\mathbb{Q}^*_{0:T}$ to the SB problem with a Markov reference process $\mathbb{Q}^{\text{ref}}_{0:T} \in \mathcal{M}$ can be characterized as a path measure in the intersection of four sets.

Proposition 2.2 ([20] Thm. 3.2, [19] Thm. 2.12)**.** *Under suitable conditions, if there exists*

$$\mathbb{Q}_{0:T} \;\in\; \mathcal{M} \;\cap\; \mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T}) \;\cap\; \Pi(\mu_0, \cdot) \;\cap \Pi(\cdot, \nu_T), \tag{2}$$

which is a Markov path measure in the reciprocal class of $\mathbb{Q}^{ref}_{0:T}$, with endpoint marginals (μ_0, ν_T), then $\mathbb{Q}_{0:T}$ uniquely solves the Schrödinger Bridge problem Eq. (SB) *with reference $\mathbb{Q}^{ref}_{0:T}$.*

SB Solution with Brownian Diffusion as Reference $\mathbb{Q}^{\text{ref}}_{0:T}$ We proceed to consider each property in Prop. 2.2 for the class of reference path measures given as the law of a Brownian diffusion with initial $\mathbb{Q}^{\text{ref}}_0 = \mu_0$, which will be our focus for the remainder of this work.

$$\mathbb{Q}^{\text{ref}}_{0:T}: \qquad dx_t = b(x_t, t)dt + \sigma_t dB_t, \qquad x_0 \sim \mu_0. \tag{3}$$

Any reference process of this form satisfies the Markov property $\mathbb{Q}^{\text{ref}}_{0:T} \in \mathcal{M}$. For given (x_0, x_T), the corresponding bridge process $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$ is obtained via Doob's h-transform [10, 14] as the law of

$$\mathbb{Q}^{\text{ref}}_{\circ|0,T}: \qquad dx_{t|0,T} = \Big(b_t(x_{t|0,T}, t) + \sigma_t^2 \nabla_{x_t} \log \mathbb{Q}^{\text{ref}}_{T|t}(x_T|x_{t|0,T})\Big)dt + \sigma_t dB_t \tag{4}$$

We will denote the h-transform term as $\dot{x}_{t|T}(x_{t|0,T}, t) = \dot{x}_{t|T} = \sigma_t^2 \nabla_{x_t} \log \mathbb{Q}^{\text{ref}}_{T|t}(x_T|x_{t|0,T})$.

Example 2.3 ([32] Eq. 3-4)**.** *For the $\mathbb{Q}^{ref}_{0:T}$ as a σ-Brownian motion (with no drift $b_t = 0$ and $\sigma_t = \sigma$ in Eq.* (3)*), the bridge $\mathbb{Q}^{ref}_{\circ|0,T}$ is the law of a simple linear interpolation of x_0, x_T plus noise*

$$\mathbb{Q}^{ref}_{\circ|0,T}: \qquad dx_{t|0,T} = \frac{x_T - x_{t|0,T}}{T-t}dt + \sigma dB_t \tag{5}$$

where $\dot{x}_{t|T} = \frac{1}{T-t}(x_T - x_{t|0,T}) = \frac{1}{T}(x_T - x_0)$ is the time-derivative of the linear interpolation. Note that $\mathbb{Q}^{ref}_{0,T} = e^{-\frac{1}{2T\sigma^2}\|x_T - x_0\|^2}$, which corresponds to Eq. (eOT) *with Euclidean cost and $\epsilon = 2T\sigma^2$.*

Using Prop. 2.2 (or the reasoning below Eq. (1)), the optimal solution to the dynamical SB problem can be constructed as a member of the reciprocal class $\mathbb{Q}^*_{0:T} = \mathbb{Q}^*_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ with the reference bridge in Eq. (4) and the optimal coupling solving Eq. (eOT). Finally, to illustrate the Markov property, it can be shown that $\mathbb{Q}^*_{0:T} \in \mathcal{M}$ may also be expressed as the law of a Brownian diffusion. In the forward direction, we have [14]

$$\begin{aligned} \mathbb{Q}^*_{0:T}: \qquad & dx_t = \Big(b(x_t, t) + v^*(x_t, t)\Big)dt + \sigma_t dB_t, \qquad x_0 \sim \mu_0, \\ \text{where} \quad & v^*(x_t, t) = \mathbb{E}_{\mathbb{Q}^*_{T|t}}[\dot{x}_{t|T}] = \sigma_t^2\, \mathbb{E}_{\mathbb{Q}^*_{T|t}}\big[\nabla_{x_t} \log \mathbb{Q}^*_{T|t}(x_T|x_t)\big]. \end{aligned} \tag{6}$$

Inspecting Eq. (6), note that the optimal $v^*(x_t, t)$ depends explicitly on the optimal path measure $\mathbb{Q}^*_{0:T}$ via the conditional $\mathbb{Q}^*_{T|t}$. For the computational methods described later, this will motivate either alternating optimizations (Sec. 3.2, [32, 27]) or separate parameterizations (Sec. 4, ours) of (i) the reciprocal path measure $\mathbb{Q}_{0:T} = \mathbb{Q}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$ induced by a coupling $\mathbb{Q}_{0,T}$, and (ii) the Markov path measure induced by a learned $v(x_t, t)$.

3 Alternating Projection Algorithms for Solving Schrödinger Bridge

In this section, we view methods for solving the SB problem from the perspective of alternating KL divergence projection [3, 32] onto sets of measures satisfying the optimality properties in Prop. 2.2. Our eventual goal is to propose and analyse a new alternating projection scheme in Sec. 4.

We first recall two notions of KL divergence projection from information geometry ([7, 2] 1.6,2.8).

Definition 3.1 (e**-Projection**)**.** *The* e-projection *of a reference $\mathbb{P}^{(i)}$ onto a set $\mathcal{S}$ is defined*

$$\mathbb{Q}^e = \text{proj}^e_{\mathcal{S}}(\mathbb{P}^{(i)}) = \arg\min_{\mathbb{Q} \in \mathcal{S}} D_{KL}[\mathbb{Q} : \mathbb{P}^{(i)}].$$

Definition 3.2 (m**-Projection**)**.** *The* m-projection *of a reference $\mathbb{Q}^{(i)}$ onto a set $\mathcal{S}$ is defined*

$$\mathbb{P}^m = \text{proj}^m_{\mathcal{S}}(\mathbb{Q}^{(i)}) = \arg\min_{\mathbb{P} \in \mathcal{S}} D_{KL}[\mathbb{Q}^{(i)} : \mathbb{P}].$$

To distinguish the projections, note that e-projection optimizes over the first argument (under which expectations are taken), while the m-projection optimizes over the second argument (as in maximum likelihood). Due to the asymmetry of the KL divergence, these projections have fundamentally different properties. Thm. 1, 3, and 4 of [7] establish conditions for the existence and uniqueness of each projection.[1]

Theorem 3.3 ([7])**.** *If $\mathcal{S}_C$ is* convex*, in other words if $\mathbb{Q}_a, \mathbb{Q}_b \in \mathcal{S}_C$ implies that $(1-\alpha)\mathbb{Q}_a + \alpha\mathbb{Q}_b \in \mathcal{S}_C$, then the e-projection of $\mathbb{P}^{(i)}$ onto $\mathcal{S}_C$ is unique and satisfies a Pythagorean relation for any $\mathbb{Q} \in \mathcal{S}_C$,*

$$D_{KL}[\mathbb{Q} : \mathbb{P}^{(i)}] = D_{KL}[\mathbb{Q} : \mathbb{Q}^e] + D_{KL}[\mathbb{Q}^e : \mathbb{P}^{(i)}].$$

Theorem 3.4 ([7])**.** *If $\mathcal{S}_{LC}$ is* log-convex*, in other words if $\mathbb{P}_a, \mathbb{P}_b \in \mathcal{S}_{LC}$ implies that $\exp\{(1-\alpha)\log \mathbb{P}_a + \alpha \log \mathbb{P}_b - \log Z_\alpha\} \in \mathcal{S}_{LC}$, then the m-projection of $\mathbb{Q}^{(i)}$ onto $\mathcal{S}_{LC}$ is unique and satisfies a Pythagorean relation for any $\mathbb{P} \in \mathcal{S}_{LC}$,*

$$D_{KL}[\mathbb{Q}^{(i)} : \mathbb{P}] = D_{KL}[\mathbb{Q}^{(i)} : \mathbb{P}^m] + D_{KL}[\mathbb{P}^m : \mathbb{P}].$$

[1] As written, the Pythagorean relations in Thm. 3.3-3.4 also require $\mathcal{S}_{\text{C}}$ or $\mathcal{S}_{\text{LC}}$ to satisfy a notion of closure with respect to a KL divergence in the appropriate direction. Uniqueness holds without these conditions [7].

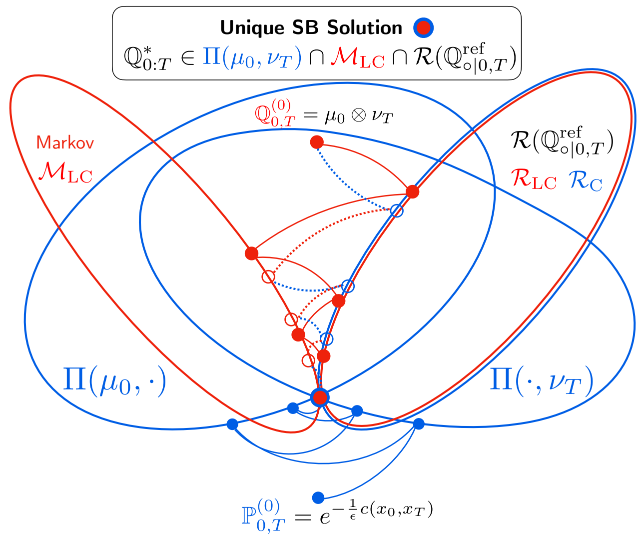

Figure 1: Alternating Projection Algorithms. Colors indicate exact e- or m- projections underlying IPF (both e, $\bullet$), IMF (both m, $\bullet$), and ours (e-m, $\circ$,$\circ$, dashed) with example initial $\mathbb{Q}^{(0)}_{0,T} \in \Pi(\mu_0, \nu_T)$ and $\mathbb{P}^{(0)}_{0,T} = e^{-\frac{1}{\epsilon}c(x_0,x_T)}$. Note that the reciprocal class $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ is both log-convex and convex.

Note that we use subscripts $\mathcal{S}_{\text{C}}$ or $\mathcal{S}_{\text{LC}}$ to emphasize the convex or log-convex properties of a set. We next highlight the properties of the set of Markov, reciprocal, and marginal-constrained path measures which are relevant to the KL projections above. See App. A for proofs.

Proposition 3.5. *The set of measures with a given marginal $\Pi(\mu_0, \cdot)$ or $\Pi(\cdot, \nu_T)$ is convex.*

Proposition 3.6. *The set of Markov path measures $\mathcal{M}_{LC}$ is log-convex.*

Proposition 3.7. *The reciprocal class $\mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T})$ is both convex ($\mathcal{R}_C$) and log-convex ($\mathcal{R}_{LC}$).*

While Shi et al. [32] note that the set of Markov path measures is not convex, we emphasize the log-convex property of $\mathcal{M}_{\text{LC}}$ and $\mathcal{R}_{\text{LC}}$ and interpret the IMF algorithm of [32, 27] using alternating m-projections. Our proposed approach in Sec. 4 will leverage the convexity of $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ to instead perform the e-projection onto the reciprocal class. Finally, we interpret IPF as performing alternating e-projections onto the sets $\Pi(\mu_0, \cdot)$, $\Pi(\cdot, \nu_T)$. We illustrate these projections geometrically in Fig. 1.

In what follows, our notation is chosen carefully such that we always denote the first argument of the KL divergence using $\mathbb{Q}$ and second argument using $\mathbb{P}$.

3.1 Iterative Proportional Fitting

The classical Iterative Proportional Fitting (IPF) or Sinkhorn algorithm [11, 6, 29, 8, 3, 17] performs iterative 'half-bridge' updates $\mathbb{P}^{(n+1)}_{0,T} \in \Pi(\mu_0, \cdot)$ and $\mathbb{P}^{(n+2)} \in \Pi(\cdot, \nu_T)$ to satisfy the marginal constraints, where we reset $n \leftarrow n+2$ after each even iteration

$$\mathbb{P}^{(n+1)}_{0,T} \leftarrow \underset{\mathbb{Q}_{0,T} \in \Pi(\cdot, \nu_T)}{\arg\min}\; D_{\text{KL}}[\mathbb{Q}_{0,T} : \mathbb{P}^{(n)}_{0,T}] \qquad \mathbb{P}^{(n+2)}_{0,T} \leftarrow \underset{\mathbb{Q}_{0,T} \in \Pi(\mu_0, \cdot)}{\arg\min}\; D_{\text{KL}}[\mathbb{Q}_{0,T} : \mathbb{P}^{(n+1)}_{0,T}] \qquad \text{(IPF)}$$

Using Def. 3.1 and Thm. 3.3, we interpret IPF as performing alternating e-projections onto the convex sets defined by marginal constraints $\Pi(\mu_0, \cdot)$ or $\Pi(\cdot, \nu_T)$ (Prop. 3.5). The Pythagorean relation in Thm. 3.3 can be used to establish monotonicity of the KL divergence to the optimum $D_{\text{KL}}[\mathbb{Q}^*_{0,T} : \mathbb{P}^{(n)}_{0,T}]$ (see App. C.2, [6, 29]), while [29] prove convergence of IPF to the solution of Eq. (eOT). However, note that exact, alternating iterates are required for IPF, since $\mathbb{P}^{(n+1)}_{0:T}$ is found using an optimization in the first argument, and then used in the second argument to find $\mathbb{P}^{(n+2)}_{0:T}$.

While we write the projections in Eq. (IPF) using couplings for the static problem, recent work [9, 37] proposes to solve SB problems using *path-space* versions of IPF. For example, the forward $(\mathbb{P}^{(n)_f}_{0:T})$ and backward $(\mathbb{P}^{(n+1)_b}_{0:T})$ iterates may be SDEs parameterized by drifts $v_{\theta_f}(x_t, t)$ and $v_{\theta_b}(x_t, t)$. From a dynamical perspective, path-space IPF converges to the optimal SB solution, a Markov path measure in the reciprocal class $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ with endpoint marginals μ_0, ν_T [9, Prop 5].

3.2 Iterative Markovian Fitting

Iterative Markovian Fitting (IMF) [24, 32, 27] (see Alg. 1) iteratively enforces the Markov $\mathbb{Q}^{(n+1)}_{0:T} \in \mathcal{M}_{\text{LC}}$ and reciprocal $\mathbb{Q}^{(n+2)}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \coloneqq \mathcal{R}_{\text{LC}}$ properties using KL divergence projections of path measures

$$\mathbb{Q}^{(n+1)}_{0:T} \leftarrow \arg\min_{\mathbb{P}\in\mathcal{M}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^{(n)}_{0:T} : \mathbb{P}_{0:T}] \qquad \mathbb{Q}^{(n+2)}_{0:T} \leftarrow \arg\min_{\mathbb{P}\in\mathcal{R}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^{(n+1)}_{0:T} : \mathbb{P}_{0:T}]. \qquad \text{(IMF)}$$

Using Def. 3.2, Thm. 3.4, and the log-convexity of $\mathcal{M}_{\text{LC}}$ and $\mathcal{R}_{\text{LC}}$ from Prop. 3.6-3.7, we interpret IMF as performing alternating m-projections. Again, the Pythagorean relation in Thm. 3.4 can be used to establish monotonicity in KL divergence to the optimum $D_{\text{KL}}[\mathbb{Q}^{(n)}_{0,T} : \mathbb{Q}^{*}_{0,T}]$ (App. C.2, [32, 27]). We again note the need for exact iterates due to optimization over the same argument of the KL divergence in each step.

We next review the Markov and reciprocal m-projections of IMF in detail. In particular, when initializing with $\mathbb{Q}^{(0)}_{0:T} \in \Pi(\mu_0, \nu_T)$, we will see that each exact projection in Eq. (IMF) preserves the endpoint marginals. This suggests that exact IMF iterates will converge to the optimal $\mathbb{Q}^{*}_{0:T} \in \mathcal{M} \cap \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$ with the desired properties in Prop. 2.2 (see [27] Thm. 2 for proof).

Markov m-Projection Inspired by the form of the SB solution in Eq. (6), the KL divergence minimization over Markov processes in Eq. (IMF) can be parameterized using a (learned) vector field $v_\theta(x_t, t)$. Using the Girsanov theorem on Eq. (IMF) and assuming $\mathbb{Q}^{(n+1)}_{0} = \mu_0$, the Markov m-projection can be implemented via the solution to the following optimization problem,

$$v^{*(n+1)}(x_t, t) = \mathbb{E}_{\mathbb{Q}^{(n)}_{T|t}}\left[\dot{x}_{t|T}\right] = \arg\min_{v_\theta} \int_0^T \mathbb{E}_{\mathbb{Q}^{(n)}_{t}}\left[\frac{1}{2\sigma_t^2}\mathbb{E}_{\mathbb{Q}^{(n)}_{T|t}}\left[\left\|\dot{x}_{t|T} - v_\theta(x_{t|0,T}, t)\right\|^2\right]\right] dt \quad (7)$$

which matches Eq. (6) for a possibly suboptimal $\mathbb{Q}^{(n)}_{0:T}$.

The Markov projection $\mathbb{Q}^{(n+1)}_{0:T} = \text{proj}^m_{\mathcal{M}_{\text{LC}}}(\mathbb{Q}^{(n)}_{0:T})$ is now given as the law of

$$\mathbb{Q}^{(n+1)}_{0:T}: \qquad dx_t = \Big(b(x_t, t) + v^{*(n+1)}(x_t, t)\Big)\, dt + \sigma_t dB_t, \qquad x_0 \sim \mu_0, \quad (8)$$

which can be shown to preserve *all* marginals of $\mathbb{Q}^{(n)}_{0:T}$ [27, Thm. 1], in particular $\mathbb{Q}^{(n+1)}_{0:T} \in \Pi(\mu_0, \nu_T)$ if $\mathbb{Q}^{(n)}_{0:T} \in \Pi(\mu_0, \nu_T)$. However, $\mathbb{Q}^{(n+1)}_{0:T} \notin \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ may not have the correct bridges and we thus require further projection.

Reciprocal m-Projection via SDE Simulation The reciprocal m-projection is given by $\mathbb{Q}^{(n+2)}_{0:T} = \mathbb{Q}^{(n+1)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T} = \text{proj}^m_{\mathcal{R}_{\text{LC}}}(\mathbb{Q}^{(n+1)}_{0:T})$, which can be seen by noting that the only degree of freedom for optimization over the reciprocal class $\mathbb{P}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ is the coupling $\mathbb{P}_{0,T}$ and that the expectation in the KL divergence is under fixed $\mathbb{Q}^{(n+1)}_{0:T}$ [32, 27].

The exact reciprocal m-projection can thus be performed by simulating the Markov process in Eq. (11) to obtain couplings $\mathbb{Q}^{(n+2)}_{0,T} \leftarrow \mathbb{Q}^{(n+1)}_{0,T}$. In the deterministic limit $\sigma \to 0$, the corresponding ODE simulation recovers to the *rectification* step for updating the couplings in Rectified Flow [23, 22]. While the exact projection clearly preserves the endpoint marginals, $\mathbb{Q}^{(n+2)}_{0,T} \in \Pi(\mu_0, \nu_T)$ if $\mathbb{Q}^{(n+1)}_{0,T} \in \Pi(\mu_0, \nu_T)$, the interpolating bridge process changes the intermediate marginals [23, 22].

4 Bridge Matching via Expectation Maximization

In this section, we propose an alternating projection approach to the SB problem which minimizes a *single* KL divergence objective $D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}_{0:T}]$ by updating $\mathbb{Q}^{(n)}_{0:T}$ and $\mathbb{P}^{(n+1)}_{0:T}$ in the same argument across iterations. We analyze the exact iterates and demonstrate their convergence in Sec. 4.1. In notable contrast to the need for exact iterates in IMF and IPF, the benefit of our approach is that minimization of a unified objective provides principled justification for inexact updates or partial descent steps, as in the variational Expectation Maximization (EM) algorithm [26] (Sec. 4.2).

Algorithm 1 Iterative Markovian Fitting [32]

input: reference bridge process $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$

input: initial coupling $\mathbb{Q}^{(0)}_{0,T} \in \Pi(\mu,\nu)$

initialize: $\mathbb{Q}^{(0)}_{0:T} = \mathbb{Q}^{(0)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$

while not converged **do**

Markov m-Projection (Flow Matching):

SAMPLE($\mathbb{Q}^{(n)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$) (see Alg. 3)

$$v_t^{*(n+1)} \leftarrow \arg\min_{v_t} \mathbb{E}_{\mathbb{Q}^{(n)}_{t,T}}\left[\|\dot{x}_{t|T} - v_t\|^2\right] = \mathbb{E}_{\mathbb{Q}^{(n)}_{T|t}}[\dot{x}_{t|T}]$$

$$\mathbb{Q}^{(n+1)}_{0:T} \leftarrow \text{Law}[\text{SDE}(b_t, v_t^{*(n+1)}, \sigma_t)] = \arg\min_{\mathbb{P}\in\mathcal{M}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^{(n)}_{0:T} : \mathbb{P}_{0:T}]$$

Reciprocal m-Projection (Rectification):

$$\mathbb{Q}^{(n+2)}_{0:T} \leftarrow \mathbb{Q}^{(n+1)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T} = \arg\min_{\mathbb{P}\in\mathcal{R}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^{(n+1)}_{0:T} : \mathbb{P}_{0:T}]$$

$n \leftarrow n+2$

end while

return: $\mathbb{Q}^{(n)}_{0:T} \in \mathcal{M}_{\text{LC}} \cap \mathcal{R}_{\text{LC}} \cap \Pi(\mu_0,\nu_T)$

Algorithm 2 Exact Bridge Matching em

input: reference bridge process $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$

input: initial coupling $\mathbb{Q}^{(0)}_{0,T} \in \Pi(\mu,\nu)$

initialize: $\mathbb{Q}^{(0)}_{0:T} = \mathbb{Q}^{(0)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}, \;\; n=0$

while not converged **do**

Markov m-Projection (Flow Matching):

$$\mathbb{P}^{(n+1)}_{0:T} \leftarrow \arg\min_{\mathbb{P}\in\mathcal{M}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^{(n)}_{0:T} : \mathbb{P}_{0:T}]$$

Reciprocal e-Projection:

$$\mathbb{Q}^{(n+2)}_{0:T} \leftarrow \arg\min_{\mathbb{Q}\in\mathcal{R}_{\text{C}}} D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}]$$

initialize $i=0$, $\mathbb{P}^{(0)_{n+2}}_{0:T} = \mathbb{Q}^{(n+2)}_{0:T}$

while not converged **do**

e-Projection onto $\Pi(\cdot,\nu_T)$:

$$\mathbb{P}^{(i+1)_{n+2}}_{0:T} \leftarrow \arg\min_{\mathbb{P}\in\Pi(\cdot,\nu_T)} D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(i)_{n+2}}_{0:T}]$$

e-Projection onto $\Pi(\mu_0,\cdot)$:

$$\mathbb{P}^{(i+2)_{n+2}}_{0:T} \leftarrow \arg\min_{\mathbb{P}\in\Pi(\mu_0,\cdot)} D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(i+1)_{n+2}}_{0:T}]$$

$i \leftarrow i+2$

end while

$n \leftarrow n+2, \;\; \mathbb{Q}^{(n)}_{0:T} \leftarrow \mathbb{P}^{(i)_{n+2}}_{0:T}$

end while

return: $\mathbb{Q}^{(n)}_{0:T} \in \mathcal{M}_{\text{LC}} \cap \mathcal{R}_{\text{C}} \cap \Pi(\mu_0,\nu_T)$

Our approach is motivated by the properties of the sets of Markov, reciprocal, and marginal-constrained path measures in Prop. 3.5-3.7. While the Markov m-projection is natural due to Prop. 3.6 and the success of previous work minimizing the regression loss in Eq. (7) [21, 23], Prop. 3.7 suggests that *either* order of the arguments might be used for the reciprocal projection onto $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$, since it is both convex and log-convex. We treat $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) = \mathcal{R}_{\text{C}}$ as a convex set and perform the e-projection onto $\mathcal{R}_{\text{C}} \cap \Pi(\mu_0,\nu_T)$, which is convex as the intersection of the convex sets.

We thus propose to perform alternating e- and m-projections as follows,

$$\mathbb{P}^{(n+1)}_{0:T} \leftarrow \arg\min_{\mathbb{P}\in\mathcal{M}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^{(n)}_{0:T} : \mathbb{P}_{0:T}] \qquad \mathbb{Q}^{(n+2)}_{0:T} \leftarrow \arg\min_{\mathbb{Q}\in\, \mathcal{R}_{\text{C}}\,\cap\,\Pi(\mu_0,\nu_T)} D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}]. \quad \text{(em)}$$

Crucially, each step of Eq. (em) optimizes the unified objective $D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}_{0:T}]$, since $\mathbb{Q}^{(n)}_{0:T}$ and $\mathbb{P}^{(n+1)}_{0:T}$ always appear in first and second argument, respectively. Furthermore, our projections consider all four of the sets characterizing the optimal SB solution in Prop. 2.2.

4.1 Exact em Procedure

We next analyze the exact projections in Eq. (em), with a particular focus on the e-projection onto $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0,\nu_T)$. We describe an algorithm for computing exact iterates in Alg. 2, before presenting a variational approach in Sec. 4.2.

Reciprocal e-Projection with Marginal Constraints Ignoring the $\Pi(\mu_0,\nu_T)$ constraint for the moment, note that the KL divergencpe in the reciprocal e-projection in Eq. (em) involves an expectation under $\mathbb{Q}_{0:T} = \mathbb{Q}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$. Compared to the reciprocal m-projection in Eq. (IMF), this changes the exact projection such that $\text{proj}^e_{\mathcal{R}_\mathbb{C}}(\mathbb{P}^{(n+1)}_{0:T})$ no longer preserves the endpoint marginals of $\mathbb{P}^{(n+1)}_{0,T}$ (see App. B.1). Thus, we must impose the marginal constraints in Eq. (em) in order to solve SB with fixed $\Pi(\mu_0, \nu_T)$.

This leads to the following form for the exact e-projection $\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0:T} = \text{proj}^e_{\mathcal{R}_\mathbb{C}\cap\Pi(\mu_0,\nu_T)}(\mathbb{P}^{(n+1)}_{0:T})$.

Proposition 4.1 (Marginal-Constrained Reciprocal e-Projection)**.** *The e-projection of a path measure $\mathbb{P}^{(n+1)}_{0:T}$ onto the convex set $\mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$ is of the form $\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0:T} = \mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T}\mathbb{Q}^{ref}_{\circ|0,T}$, where*

$$\frac{d\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T}}{d\mathbb{P}^{(n+1)}_{0,T}} = \exp\left\{-D_{KL}\left[\mathbb{Q}^{ref}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}\right] - \phi_0 - \phi_T\right\}. \tag{9}$$

The projection in Eq. (9) matches the solution of the static regularized OT problem [19], with a cost $c(x_0, x_T) = D_{\text{KL}}[\mathbb{Q}^{\text{ref}}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}]$ that depends on $\mathbb{P}^{(n+1)}_{0:T} \in \mathcal{M}_{\text{LC}}$ and regularization with $D_{\text{KL}}[\mathbb{Q}_{0,T} : \mathbb{P}^{(n+1)}_{0,T}]$ instead of the entropy in Eq. (eOT) (see App. C.1). For given endpoints (x_0, x_T), the cost measures the mismatch between the target bridge $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$ and the bridge $\mathbb{P}^{(n+1)}_{\circ|0,T}$ induced by the current Markov path measure (as in Eq. (4)). Our updates reach a fixed point if the support of $\mathbb{P}^{(n+1)}_{0,T}$ concentrates on (x_0, x_T) with the correct bridges $\mathbb{P}^{(n+1)}_{\circ|0,T} = \mathbb{Q}^{\text{ref}}_{\circ|0,T}$.

Exact em Algorithm In Alg. 2, we describe an algorithm to calculate the projections in Eq. (em). While the Markov m-projection is the same as in IMF (Sec. 3.2), the e-projection is more involved due to the set intersection $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$. We propose to first perform the reciprocal e-projection $\mathbb{Q}^{e(\mathcal{R})}_{0:T} = \text{proj}^e_{\mathcal{R}_\mathbb{C}}(\mathbb{P}^{(n+1)}_{0:T})$ onto $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$ only (see Prop. B.1). Motivated by the interpretation of Prop. 4.1 as the solution to a static regularized OT problem, we next perform alternating IPF e-projections starting from an initial reference measure $\mathbb{P}^{(0)_{n+2}}_{0,T} = \mathbb{Q}^{e(\mathcal{R})}_{0:T} \propto \mathbb{P}^{(n+1)}_{0,T} \exp\{-D_{\text{KL}}[\mathbb{Q}^{\text{ref}}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}]\}$. These IPF iterations converge to the correct potentials ϕ_0, ϕ_T in Eq. (9) [29], and remain within the reciprocal class since only the couplings are updated.

We show the convergence of this procedure in the following proposition (see App. C for proof).

Proposition 4.2. *The exact alternating projection algorithm in Eq.* (em) *(in particular, Alg. 2) converges to $\mathbb{Q}^{(\infty)}_{0:T} \in \mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$ and $\mathbb{P}^{(\infty)}_{0:T} \in \mathcal{M}_{LC}$. If $\mathbb{Q}^{(\infty)}_{0:T}$ or $\mathbb{P}^{(\infty)}_{0:T}$ is a fixed point of the procedure, then it is equal to the optimal SB solution, $\mathbb{Q}^{(\infty)}_{0:T} = \mathbb{P}^{(\infty)}_{0:T} = \mathbb{Q}^*_{0:T}$.*

Nevertheless, we note two issues which make it impractical to perform the exact e-projection in Prop. 4.1. First, even if we could calculate the cost $D_{\text{KL}}[\mathbb{Q}^{\text{ref}}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}]$, the IPF iterations to compute the e-projection in each step are as difficult as the original problem Eq. (eOT). Further, the Doob h-transform with a nonlinear drift $v^{(n+1)}(x_t, t)$ is intractable in general, so we do not expect to be able to calculate the bridge $\mathbb{P}^{(n+1)}_{\circ|0,T}$ or KL divergence cost.

4.2 Variational Bridge Matching

Despite the intractability of the exact updates, the fact that our alternating projections in Eq. (em) minimize a unified KL divergence objective suggests maintaining separate (parametric) representations of a reciprocal path measure $\mathbb{Q}^{\phi(n)}_{0:T} = \mathbb{Q}^{\phi(n)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$ and Markov process $\mathbb{P}^{\theta(n+1)}_{0:T}$ (via $v^{(n+1)}_\theta(x_t, t)$). As in the variational EM algorithm, we may perform partial descent steps which serve to decrease $D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}_{0:T}]$. In particular, the m-step performs *learning* of a Markovian diffusion model, while the e-step performs *inference* of a (marginal-constrained) coupling and reciprocal path measure.

Using the Girsanov theorem as in [24, 32], the KL divergence minimization in Eq. (em) becomes

$$\min_{\mathbb{Q}^\phi_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{0:T})} \min_{\mathbb{P}^\theta_{0:T} \in \mathcal{M}_{\text{LC}}} D_{\text{KL}}[\mathbb{Q}^\phi_{0:T} : \mathbb{P}^\theta_{0:T}] \tag{10}$$
$$= \min_{\mathbb{Q}^\phi_{0,T}} \min_{v^\theta_t, \mathbb{P}^\theta_0} \int_0^T \mathbb{E}_{\mathbb{Q}^\phi_t}\left[\frac{1}{2\sigma_t^2}\mathbb{E}_{\mathbb{Q}^\phi_{T|t}}\left[\left\|\dot{x}_{t|T} - v_\theta(x_{t|0,T}, t)\right\|^2\right]\right] dt + \mathbb{E}_{\nu_T}\left[D_{\text{KL}}[\mathbb{Q}^\phi_{0|T} : \mathbb{P}^\theta_0]\right]$$

In particular, we assume $\mathbb{Q}^{\phi}_{T} = \nu_T$ and apply the Girsanov theorem between the forward bridge process in Eq. (4) (with initial $x_0 \sim \mathbb{Q}^{\phi}_{0|T}$) and the Markov process $\mathbb{P}^{\theta}_{0:T}$ given by

$$\mathbb{P}^{\theta}_{0:T}: \qquad dx_t = \Big(b(x_t,t) + v_\theta(x_t,t)\Big)dt + \sigma_t dB_t, \qquad x_0 \sim \mathbb{P}^{\theta}_0. \tag{11}$$

where the optimal $v_\theta(x_t,t)$ for a given $\mathbb{Q}^{\phi}_{0:T}$ is given by Eq. (7). Note, the KL divergence term at $t=0$ in Eq. (10) is between the initial measures of the two forward processes, whereas the final endpoint term in the limit as $t \to T$ is ignored, as in [24] Sec. 4, [32].

Remark 4.3 (Marginal Preservation). While we write a possible optimization over $\mathbb{P}^{\theta}_0$ in Eq. (10) (as in [24]), note that its optimum occurs at $\mathbb{P}^{\theta}_0 = \mathbb{Q}^{\phi}_0$ for a given $\mathbb{Q}^{\phi}_{0:T}$. Fixing $\mathbb{Q}^{\phi}_T = \nu_T$, we thus require $\mathbb{Q}^{\phi}_0 = \int \mathbb{Q}^{\phi}_{0|T} d\nu_T = \mathbb{P}^{\theta}_0 = \mu_0$ in order for Eq. (10) to solve the SB problem for (μ_0, ν_T). This suggests the need to design clever parameterizations of couplings $\mathbb{Q}^{\phi}_{0:T} \in \Pi(\mu_0, \nu_T)$ which preserve the desired marginals. We discuss approaches from previous work in Sec. 5

Finally, as in Neal and Hinton [26], we consider joint optimization of the KL divergence objective in Eq. (11) using, for example, gradient descent with learning rates η_θ, η_ϕ

$$\theta \leftarrow \theta - \eta_\theta \nabla_\theta D_{\mathrm{KL}}[\mathbb{Q}^{\phi}_{0:T} : \mathbb{P}^{\theta}_{0:T}] \qquad\qquad \phi \leftarrow \phi - \eta_\phi \nabla_\phi D_{\mathrm{KL}}[\mathbb{Q}^{\phi}_{0:T} : \mathbb{P}^{\theta}_{0:T}].$$

where we might also choose to alternate between K_θ gradient steps of θ for fixed ϕ, and K_ϕ gradient steps of ϕ for fixed θ. Again, our unified objective provides justification for these inexact, partial e and m steps, in contrast to IPF and IMF whose convergence relies on exact iterates.

5 Discussion

While the Markov m-step is well understood in the literature, our em perspective sheds light on how to perform the reciprocal e-step such that inexact, variational updates are justified. Our exact projection in Prop. 4.1 differs from the SDE simulation used in IMF, and thus suggests searching for expressive, marginal-preserving parameterizations of the reciprocal couplings to approximate Eq. (9). We discuss existing parameterizations [35, 28, 16] and other related work

Markov m-Step and Path Straightness The family of flow matching [21, 1, 28, 35, 34] and rectified flow [23, 22, 25] methods learn a vector field v_θ using a regression loss similar to Eq. (10),

$$v^{*(n+1)}(x_t,t) = \mathbb{E}_{\mathbb{Q}^{(n)}_{T|t}}\left[\dot{x}_{t|T}\right] = \operatorname*{arg\,min}_{v_\theta} \int_0^T \mathbb{E}_{\mathbb{Q}^{(n)}_{t,T}}\left[\left\|\dot{x}_{t|T} - v_\theta(x_{t|0,T},t)\right\|^2\right]dt, \tag{12}$$

which matches the KL divergence minimization in Eq. (7) up to the $1/\sigma_t^2$ weighting factor.

We can view the Markov m-projection or flow-matching objective as associating the optimal value in Eq. (12) to a given coupling $\mathbb{Q}^{(n)}_{0,T}$. In particular, the optimal value at $v^*(x_t,t) = \mathbb{E}_{\mathbb{Q}^{(n)}_{T|t}}[\dot{x}_{t|T}]$ corresponds to the *conditional variance* of $\dot{x}_{t|T}$ under the reciprocal path measure $\mathbb{Q}^{(n)}_{0,T} = \mathbb{Q}^{(n)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$,

$$\mathrm{Var}_{\mathbb{Q}^{(n)}_{0:T}}\left[\dot{x}_{t|T}\right] = \int_0^T \mathbb{E}_{\mathbb{Q}^{(n)}_{t,T}}\left[\left\|\dot{x}_{t|T} - \mathbb{E}_{\mathbb{Q}^{(n)}_{T|t}}[\dot{x}_{t|T}]\right\|^2\right]dt. \tag{13}$$

This quantity has been interpreted to measure the 'straightness' of bridge paths induced by a coupling $\mathbb{Q}^{(n)}_{0,T}$ (lower is straighter) [23, 22, 28, 16], and obtaining straighter paths has been the motivation for various related work described below.

Relation to Rectified Flow Rectified flow [23, 22] may be viewed as the deterministic limit of the IMF algorithm as $\sigma_t \to 0$, although reasoning using the KL divergence and Girsanov theorem does not appear to translate directly to the deterministic case. As in IMF, couplings are updated by simulating an ODE from either endpoint marginal, and $v_\theta(x_t,t)$ is updated using the regression loss to the bridge vector fields in Eq. (12). Exact iterations of this procedure have been shown to reduce the conditional variance in Eq. (13) and yield straighter paths [23]. Liu et al. [23, 25] demonstrate impressive results for one-step generative modeling and distillation.

Relation to Flow Matching Initial work on flow matching [21, 1] uses fixed, independent couplings $\mathbb{Q}^{\phi}_{0,T} = \mu_0 \otimes \nu_T$ to construct a reciprocal process $\mathbb{Q}_{0:T}$, and learns a vector field v_θ using Eq. (12). However, to straighten the paths in Eq. (13) and approach the solution to a dynamical OT problem, [28, 35, 34] use (regularized) OT solvers to obtain couplings on mini-batches of data. From our perspective, this defines a particular inference procedure for the reciprocal coupling $\mathbb{Q}_{0,T}$. However, marginal preservation and convergence to the OT or SB solution can only be guaranteed in the limit as the minibatch size $n \to \infty$ [28].

Relation to Lee et al. [16] Lee et al. [16] propose a similar joint optimization of the KL divergence in Eq. (10), with fixed $\mathbb{P}^{\theta}_0 = \mu_0 = \mathcal{N}(0, \mathbb{I})$ and couplings $\mathbb{Q}^{\phi}_{0,T} = \nu_T \mathbb{Q}^{\phi}_{0|T}$ parameterized by an encoder $\mathbb{Q}^{\phi}_{0|T}$ which maps from empirical data $x_T \sim \nu_T$ to noise. However, this coupling distribution may not preserve the correct initial marginal $\mathbb{Q}^{\phi}_0 \neq \mu_0$ (see Remark 4.3). Lee et al. [16] thus add additional regularization by $D_{\mathrm{KL}}[\mathbb{Q}^{\phi(n)}_0 : \mu_0]$ with weight $\beta \gg 1$, which translates to reweighting the term $\beta\, \mathbb{E}_{\nu_T} D_{\mathrm{KL}}[\mathbb{Q}^{\phi}_{0|T} : \mathbb{P}^{\theta}_0]$ in Eq. (10). The method in [16] is motivated by optimizing the forward process in diffusion models to obtain straighter paths as in Eq. (13), and does not make explicit connections with the SB problem. While the joint optimization is similar to our proposed variational bridge matching, exact marginal preservation remains a challenge using this approach.

Previous EM Approaches Similarly to our proposed interpretation in the stochastic case, Liu [22] Sec. 5.4 views (deterministic) rectified flow as a majorization-minimization algorithm, of which the EM algorithm is the most famous example [13]. Liu et al. [24] discuss an analogy with EM, but ignore inference in the e-step and only optimize the model parameters.

Vargas and Nüsken [36] also propose an EM-style optimization for solving a general class of divergence minimization problems involving forward and backward SDEs. Compared to their approach, we restrict attention to reciprocal projections for reference processes with tractable bridges [24, 32, 27], which simplifies the learning process in the first argument since we only optimize over the coupling $\mathbb{Q}^{\phi}_{0,T}$. In particular, we avoid backpropagation through sampling dynamics and the need for Hamilton-Jacobi regularizers, which [36] argue play a similar role to the reciprocal projection.

6 Conclusion

In this work, we have understood alternating projection methods for solving Schrödinger Bridge problems from the perspective of information geometry. Motivated the properties of the sets of Markov and reciprocal path measures, we proposed a new projection approach which yields a single KL divergence objective and allows for inexact updates in the style of variational EM. The perspective sheds light on methods from previous work, and suggests searching for expressive parameterizations of marginal-preserving coupling distributions to solve Schrödinger Bridge problems.

References

[1] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023.

[2] Shun-ichi Amari. *Information geometry and its applications*, volume 194. Springer, 2016.

[3] Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.

[4] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169:671–691, 2016.

[5] Yongxin Chen, Tryphon T Georgiou, and Michele Pavon. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrodinger bridge. *Siam Review*, 63(2):249–313, 2021.

[6] Imre Csiszár. I-divergence geometry of probability distributions and minimization problems. *The annals of probability*, pages 146–158, 1975.

[7] Imre Csiszár and Frantisek Matus. Information projections revisited. *IEEE Transactions on Information Theory*, 49(6):1474–1490, 2003.

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

[9] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.

[10] Joseph L Doob. Conditional brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société Mathématique de France*, 85:431–458, 1957.

[11] Robert Fortet. Résolution d'un systeme d'équations de m. schrödinger. *J. Math. Pure Appl. IX*, 1:83–105, 1940.

[12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.

[13] David R Hunter and Kenneth Lange. A tutorial on mm algorithms. *The American Statistician*, 58(1):30–37, 2004.

[14] Benton Jamison. The markov processes of schrödinger. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 32(4):323–331, 1975.

[15] Leonid V Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pages 199–201, 1942.

[16] Sangyun Lee, Beomsu Kim, and Jong Chul Ye. Minimizing trajectory curvature of ode-based generative models. *International Conference on Machine Learning*, 2023.

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

[18] Christian Léonard. From the schrödinger problem to the monge–kantorovich problem. *Journal of Functional Analysis*, 262(4):1879–1920, 2012.

[19] Christian Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *Discrete & Continuous Dynamical Systems*, 34(4):1533, 2014.

[20] Christian Léonard, Sylvie Rœlly, and Jean-Claude Zambrini. Reciprocal processes. a measure-theoretical point of view. *Probability Surveys*, 11:237–269, 2014.

[21] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *International Conference on Learning Representations*, 2022.

[22] Qiang Liu. Rectified flow: A marginal preserving approach to optimal transport. *arXiv preprint arXiv:2209.14577*, 2022.

[23] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *International Conference on Learning Representations*, 2022.

[24] Xingchao Liu, Lemeng Wu, Mao Ye, and Qiang Liu. Let us build bridges: Understanding and extending diffusion generative models. *arXiv preprint arXiv:2208.14699*, 2022.

[25] Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, and Qiang Liu. Instaflow: One step is enough for high-quality diffusion-based text-to-image generation. *arXiv preprint arXiv:2309.06380*, 2023.

[26] Radford M Neal and Geoffrey E Hinton. A view of the em algorithm that justifies incremental, sparse, and other variants. *Learning in graphical models*, pages 355–368, 1998.

[27] Stefano Peluchetti. Diffusion bridge mixture transports, schrodinger bridge problems and generative modeling. *arXiv preprint arXiv:2304.00917*, 2023.

[28] Aram-Alexandre Pooladian, Heli Ben-Hamu, Carles Domingo-Enrich, Brandon Amos, Yaron Lipman, and Ricky Chen. Multisample flow matching: Straightening flows with minibatch couplings. *arXiv preprint arXiv:2304.14772*, 2023.

[29] Ludger Ruschendorf. Convergence of the iterative proportional fitting procedure. *The Annals of Statistics*, pages 1160–1174, 1995.

[30] Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176 (4):928–943, 2019. URL `https://www.cell.com/cms/10.1016/j.cell.2019.01.006/attachment/d251b88d-a356-436a-a9a7-b7e04b56b41f/mmc8.pdf`.

[31] Erwin Schrödinger. Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique. In *Annales de l'institut Henri Poincaré*, volume 2, pages 269–310, 1932.

[32] Yuyang Shi, Valentin De Bortoli, Andrew Campbell, and Arnaud Doucet. Diffusion schr\" odinger bridge matching. *arXiv preprint arXiv:2303.16852*, 2023.

[33] 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.

[34] Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguet, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023.

[35] Alexander Tong, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Kilian Fatras, Guy Wolf, and Yoshua Bengio. Conditional flow matching: Simulation-free dynamic optimal transport. *arXiv preprint arXiv:2302.00482*, 2023.

[36] Francisco Vargas and Nikolas Nüsken. Transport, vi, and diffusions. In *ICML Workshop on New Frontiers in Learning, Control, and Dynamical Systems*, 2023.

[37] Francisco Vargas, Pierre Thodoroff, Austen Lamacraft, and Neil Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.

A Markov and Reciprocal Path Measures

In this section, we review definitions of Markov and reciprocal path measures, and demonstrate their properties of convexity or log-convexity.

Definition A.1 (Markov Path Measure (Léonard et al. [20] Thm. 1.2))**.** *For times $0 \le s < t \le 1$ and events $A_{(0,s)}$ in σ-algebra generated by $\{X_r : 0 \le r < s\}$ and $B_{(t,1)}$ in the σ-algebra generated by $\{X_r : t < r \le 1\}$, then*

$$Pr[A_{(0,s)} \cap B_{(t,1)} | X_s, X_t] = Pr[A_{(0,s)} | X_s] \; Pr[B_{(t,1)} | X_t] \tag{14}$$

Proposition 3.6. *The set of Markov path measures $\mathcal{M}_{LC}$ is log-convex.*

Proof. For $\mathbb{Q}_a, \mathbb{Q}_b \in \mathcal{E}_M$, consider $\mathbb{Q}^{e_\alpha} := \exp\{(1-\alpha)\log \mathbb{Q}_a + \alpha \log \mathbb{Q}_b - \log Z_\alpha\}$. We would like to show that this is Markov.

$$\mathbb{Q}^{e_\alpha}[A_{(0,s)} \cap B_{(t,1)} | X_s, X_t] \tag{15}$$

$$= \exp\{(1-\alpha)\log \mathbb{Q}_a[A_{(0,s)} \cap B_{(t,1)} | X_s, X_t] + \alpha \log \mathbb{Q}_b[A_{(0,s)} \cap B_{(t,1)} | X_s, X_t] - \log Z_\alpha\} \tag{16}$$

$$= \exp\{(1-\alpha)\log \mathbb{Q}_a[A_{(0,s)} | X_s] + \alpha \log \mathbb{Q}_b[A_{(0,s)} | X_s] + (1-\alpha)\log \mathbb{Q}_a[B_{(t,1)} | X_t] + \alpha \log \mathbb{Q}_b[B_{(t,1)} | X_t] - \log Z_\alpha\} \tag{17}$$

$$= \frac{1}{Z_\alpha} \mathbb{Q}^{e_\alpha}[A_{(0,s)} | X_s] \mathbb{Q}^{e_\alpha}[B_{(t,1)} | X_t] \tag{18}$$

which satisfies Eq. (14) as desired. □

Definition A.2 (Reciprocal Path Measure)**.** *For times $0 \le s < t \le 1$ and events $A_{(s,t)^c}$ in σ-algebra generated by $\{X_r : 0 \le r < s\} \cup \{X_r : t < r \le 1\}$ and $B_{(s,t)}$ in the σ-algebra generated by $\{X_r : s < r < t\}$, then*

$$Pr[A_{(s,t)^c} \cap B_{(s,t)} | X_s, X_t] = Pr[A_{(s,t)^c} | X_s, X_t] \; Pr[B_{(s,t)} | X_s, X_t]. \tag{19}$$

Note that reciprocal path measures are not Markov in general, but Markov path measures are reciprocal [20]. Nevertheless, a Markov $\mathbb{P}_{0:T} \in \mathcal{M}$ will not be in the reciprocal class for a given reference $\mathbb{Q}^{\text{ref}}_{0:T}$ if it does not have the desired bridges $\mathbb{P}_{\circ|0,T} \ne \mathbb{Q}^{\text{ref}}_{\circ|0,T}$.

Proposition 3.7. *The reciprocal class $\mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T})$ is both convex ($\mathcal{R}_C$) and log-convex ($\mathcal{R}_{LC}$).*

Proof. Clearly, the set of reciprocal measures is m-affine, since for $\mathbb{Q}^a_{0:T}, \mathbb{Q}^b_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{0:T})$, then

$$\mathbb{Q}^{m(\alpha)}_{0:T} = (1-\alpha)\mathbb{Q}^a_{0:T} + \alpha \mathbb{Q}^b_{0:T} = (1-\alpha)\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}^a_{0,T} + \alpha \mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}^b_{0,T} \tag{20}$$

$$= \mathbb{Q}^{\text{ref}}_{\circ|0,T}\left((1-\alpha)\mathbb{Q}^a_{0,T} + \alpha \mathbb{Q}^b_{0,T}\right) \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{0:T}) \tag{21}$$

Consider the e-affine property. For $\mathbb{Q}^a_{0:T}, \mathbb{Q}^b_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{0:T})$, then

$$\mathbb{Q}^{e(\alpha)}_{0:T} = \frac{1}{Z_\alpha}(\mathbb{Q}^a_{0:T})^{1-\alpha}(\mathbb{Q}^b_{0:T})^{\alpha} = \frac{1}{Z_\alpha}\left(\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}^a_{0,T}\right)^{1-\alpha}\left(\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{R}^b_{0,T}\right)^{\alpha} \tag{22}$$

$$= \mathbb{Q}^{\text{ref}}_{\circ|0,T}\frac{1}{Z_\alpha}(\mathbb{Q}^a_{0,T})^{1-\alpha}(\mathbb{Q}_{0,T})^{\alpha} \tag{23}$$

$$= \mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}^{e(\alpha)}_{0,T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}),$$

as desired. □

Sampling from Reciprocal Measure For completeness, we provide Alg. 3 describing how to sample from a reciprocal measure $\mathbb{Q}_{0:T} = \mathbb{Q}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{0:T})$ via ancestral sampling.

B Reciprocal e- and m-Projections

In this section, we calculate both the e- and m-projections on the reciprocal class $\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$. **??** is novel, while Prop. B.2 recovers the projection from Shi et al. [32].

Algorithm 3 SAMPLE($\mathbb{Q}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$)

Sample $\{(x_0^{(k)}, x_T^{(k)})\}_{k=1}^K \sim \mathbb{Q}_{0,T}$

Simulate $\{x_{t|0,T}^{(k)}\}_{k=1}^K$ using $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$

Compute $\{\dot{x}_{t|T}(x_{t|0,T}^{(k)})^{(k)}\}_{k=1}^K$ using $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$, $(x_0^{(k)}, x_T^{(k)}, x_{t|0,T}^{(k)})$

return: $\{(x_0^{(k)}, x_T^{(k)}, x_{t|0,T}^{(k)}, \dot{x}_{t|T}^{(k)})\}_{k=1}^K$

B.1 Reciprocal e-Projection (Ours)

We first state the result for the reciprocal e-projection $\mathbb{Q}_{0:T}^{e(\mathcal{R}_\text{C})} = \text{proj}^e_{\mathcal{R}_\text{C}}(\mathbb{P}_{0:T}^{(n+1)})$, without the endpoint marginal constraints.

Proposition B.1 (Reciprocal e-Projection). *The e-projection $\mathbb{Q}_{0:T}^{e(\mathcal{R})} = \text{proj}^e_{\mathcal{R}(\mathbb{Q}^{ref})}(\mathbb{P}_{0:T}^{(n+1)})$ of a path measure $\mathbb{P}_{0:T}^{(n+1)}$ onto the reciprocal class $\mathcal{R}(\mathbb{Q}^{ref})$ satisfies*

$$\frac{d\mathbb{Q}_{0,T}^{e(\mathcal{R})}}{d\mathbb{P}_{0,T}^{\theta(n+1)}} = \frac{1}{Z}\exp\Big\{ - D_{KL}\Big[\mathbb{Q}^{ref}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}\Big]\Big\} \tag{24}$$

where the proof follows similar derivations as in the proof of Prop. 4.1 below. Notably, this projection does not preserve the endpoint marginals of $\mathbb{P}_{0,T}^{\theta(n+1)}$, since the KL divergence in the exponential is a function of x_0, x_T. Thus, as in the main text, we must consider further constraining the projection to match the marginals $\mathbb{Q}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$ in order to solve the SB problem.

Proposition 4.1 (Marginal-Constrained Reciprocal e-Projection). *The e-projection of a path measure $\mathbb{P}_{0:T}^{(n+1)}$ onto the convex set $\mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$ is of the form $\mathbb{Q}_{0:T}^{e(\mathcal{R}\cap\Pi)} = \mathbb{Q}_{0,T}^{e(\mathcal{R}\cap\Pi)}\mathbb{Q}^{ref}_{\circ|0,T}$, where*

$$\frac{d\mathbb{Q}_{0,T}^{e(\mathcal{R}\cap\Pi)}}{d\mathbb{P}_{0,T}^{(n+1)}} = \exp\Big\{-D_{KL}\Big[\mathbb{Q}^{ref}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}\Big] - \phi_0 - \phi_T\Big\}. \tag{9}$$

Proof. Consider the KL divergence minimization defining the e-projection of a path measure $\mathbb{P}_{0:T}^{(n+1)}$ onto the convex set $\mathcal{R}_\text{C} \cap \Pi(\mu_0, \nu_T)$, where $\mathcal{R}_\text{C} := \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$.

$$\mathbb{Q}_{0:T}^{e(\mathcal{R}\cap\Pi)} := \text{proj}^e_{\mathcal{R}_\text{C}\cap\Pi(\mu_0,\nu_T)}(\mathbb{P}_{0:T}^{(n+1)}) = \underset{\mathbb{Q}\in\mathcal{R}_\text{C}\cap\Pi(\mu_0,\nu_T)}{\arg\min} D_\text{KL}[\mathbb{Q}_{0:T} : \mathbb{P}_{0:T}^{(n+1)}]$$

Introducing Lagrange multipliers $\phi_0(x_0)$ and $\phi_T(x_T)$ to enforce the endpoint constraints and taking the variation with respect to $\mathbb{Q}_{0,T}$, we have

$$0 = \int d\mathbb{Q}^{\text{ref}}_{\circ|0,T} \log \frac{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}}{d\mathbb{P}_{0:T}^{(n+1)}} + \int \frac{d\mathbb{Q}_{0,T}}{d\mathbb{Q}_{0,T}} d\mathbb{Q}^{\text{ref}}_{\circ|0,T} + \phi_0 + \phi_T$$

which, ignoring the constant, implies

$$d\mathbb{Q}^e_{0,T} = d\mathbb{P}_{0,T}^{(n+1)} \exp\left\{ - \int d\mathbb{Q}^{\text{ref}}_{\circ|0,T} \log \frac{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}}{d\mathbb{P}^{(n+1)}_{\circ|0,T}} - \phi_0 - \phi_T \right\} \tag{25}$$

$$= d\mathbb{P}_{0,T}^{(n+1)} \exp\Big\{ - D_\text{KL}[\mathbb{Q}^{\text{ref}}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}] - \phi_0 - \phi_T \Big\} \tag{26}$$

Note that normalization is automatically enforced if μ_0, ν_T are normalized.

Pythagorean Relation: We now confirm that the Pythagorean relation holds for some (other) $\mathbb{Q}_{0:T} \in \mathcal{R}_\text{C} \cap \Pi(\mu_0, \nu_T)$ with $\mathbb{Q}_{0:T} = \mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}$. Writing $\mathbb{Q}_{0:T}^{e(\mathcal{R}\cap\Pi)} = \text{proj}^e_{\mathcal{R}_\text{C}:=\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})}(\mathbb{P}_{0:T}^{(n+1)}) =$

$\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$ (with $\mathbb{Q}^{e}_{0,T}$ as in Eq. (26)), we have

$$
\begin{aligned}
&D_{\text{KL}}[\mathbb{Q}_{0:T} : \text{proj}^{e}_{\mathcal{R}_{\text{C}}\cap\Pi(\mu_0,\nu_T)}(\mathbb{P}^{(n+1)}_{0:T})] + D_{\text{KL}}[\text{proj}^{e}_{\mathcal{R}_{\text{C}}\cap\Pi(\mu_0,\nu_T)}(\mathbb{P}^{(n+1)}_{0:T}) : \mathbb{P}^{(n+1)}_{0:T}] \\
&= \int d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}\left(\log \cancel{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}} + \log \mathbb{Q}_{0,T} - \log \cancel{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}} - \log \mathbb{P}^{(n+1)}_{0,T} + D_{\text{KL}}[\mathbb{Q}^{\text{ref}}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}] + \log\phi_0 + \log\phi_T\right) \\
&\quad + \int d\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0:T}\left(\log \mathbb{Q}^{\text{ref}}_{\circ|0,T} + \log \cancel{\mathbb{P}^{(n+1)}_{0,T}} - D_{\text{KL}}[\mathbb{Q}^{\text{ref}}_{\circ|0,T} : \mathbb{P}^{(n+1)}_{\circ|0,T}] - \log\phi_0 - \log\phi_T - \log \cancel{d\mathbb{P}^{(n+1)}_{0,T}}\, d\mathbb{P}^{(n+1)}_{\circ|0,T}\right) \\
&= \int d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}\left(\log \frac{d\mathbb{Q}_{0,T}}{d\mathbb{P}^{(n+1)}_{0,T}}\right) + \int d\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T}\int d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\left(\log \frac{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}}{d\mathbb{P}^{(n+1)}_{\circ|0,T}}\right) \\
&= \int d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}\left(\log \frac{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}}{d\mathbb{P}^{(n+1)}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}}\right) = D_{\text{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}]
\end{aligned}
$$

as desired. □

B.2 Reciprocal m-Projection (IMF)

We repeat the derivation of the m-projection on the reciprocal class and its Pythagorean relation from [32, 27, 24].

Proposition B.2 (m-Projection onto Reciprocal Class)**.** *The m-projection of a (Markov) $\mathbb{P}^{(n+1)}_{0:T}$ on the reciprocal class $\mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T})$ is given by $\text{proj}^{m}_{\mathcal{R}_{LC}}(\mathbb{P}^{(n+1)}_{0:T}) = \mathbb{Q}^{ref}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}$.*

Proof. Consider the projection onto $\mathcal{R}_{\text{LC}} := \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T})$, the set of $\mathbb{Q}_{0:T} \in \mathcal{R}_{\text{LC}}$ taking the form $\mathbb{Q}_{0:T} = \mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}$

$$
\begin{aligned}
\mathbb{Q}^{m}_{0:T} := \text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}) &= \underset{\mathbb{Q}\in\mathcal{R}_{\text{LC}}}{\arg\min}\, D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \mathbb{Q}_{0:T}] && (27)\\
&= \underset{\mathbb{Q}\in\mathcal{R}_{\text{LC}}}{\arg\min}\, D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0,T} : \mathbb{Q}_{0,T}] + \mathbb{E}_{\mathbb{P}^{(n+1)}_{0,T}}\left[D_{\text{KL}}[\mathbb{P}^{(n+1)}_{\circ|0,T} : \mathbb{Q}^{\text{ref}}_{\circ|0,T}]\right] && (28)
\end{aligned}
$$

Since $\mathbb{Q}^{\text{ref}}_{\circ|0,T}$ and $\mathbb{P}^{(n+1)}_{\circ|0,T}$ are fixed, we ignore the second term and conclude that the m-projection is

$$
\mathbb{Q}^{m}_{0,T} = \mathbb{P}^{(n+1)}_{0,T} \qquad\qquad \mathbb{Q}^{m}_{0:T} = \mathbb{P}^{(n+1)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T} = \text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}). \tag{29}
$$

Pythagorean Relation: To confirm the Pythagorean relation holds, we would like to show

$$
D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \mathbb{Q}_{0:T}] = D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}_{0:T})] + D_{\text{KL}}[\text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}_{0:T}) : \mathbb{Q}_{0:T}]. \tag{30}
$$

For any (other) $\mathbb{Q}_{0:T} \in \mathcal{R}_{\text{LC}}$, we can write $\mathbb{Q}_{0:T} = \mathbb{Q}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$. Using $\mathbb{Q}^{m}_{0:T} = \mathbb{P}^{(n+1)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$,

$$
\begin{aligned}
&D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}_{0:T})] + D_{\text{KL}}[\text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}_{0:T}) : \mathbb{Q}_{0:T}] && (31)\\
&\quad = \int d\mathbb{P}^{(n+1)}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}\left(\log \frac{d\mathbb{P}^{(n+1)}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}}{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}}\right) + \int d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}\left(\log \frac{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}}{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}}\right)\\
&\quad = \mathbb{E}_{\mathbb{P}^{(n)}_{0,T}}\left[D_{\text{KL}}[\mathbb{P}^{(n+1)}_{\circ|0,T} : \mathbb{Q}^{\text{ref}}_{\circ|0,T}]\right] + D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0,T} : \mathbb{Q}_{0,T}] && (32)
\end{aligned}
$$

which we confirm is equal to

$$
\begin{aligned}
D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \mathbb{Q}_{0:T}] &= \int d\mathbb{P}^{(n+1)}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}\left(\log \frac{d\mathbb{P}^{(n+1)}_{\circ|0,T}\mathbb{P}^{(n+1)}_{0,T}}{d\mathbb{Q}^{\text{ref}}_{\circ|0,T}\mathbb{Q}_{0,T}}\right) && (33)\\
&= \mathbb{E}_{\mathbb{P}^{(n)}_{0,T}}\left[D_{\text{KL}}[\mathbb{P}^{(n+1)}_{\circ|0,T} : \mathbb{Q}^{\text{ref}}_{\circ|0,T}]\right] + D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0,T} : \mathbb{Q}_{0,T}] && (34)
\end{aligned}
$$

Comparing Eq. (32) and Eq. (34), we have $D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \mathbb{Q}_{0:T}] = D_{\text{KL}}[\mathbb{P}^{(n+1)}_{0:T} : \text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}_{0:T})] + D_{\text{KL}}[\text{proj}^{m}_{\mathcal{R}_{\text{LC}}}(\mathbb{P}^{(n+1)}_{0:T}) : \mathbb{Q}_{0:T}]$, as desired in Eq. (30). □

C Convergence Analysis

In App. C.1, we prove that the exact em iterates from Alg. 2 converge to the optimal SB solution. We also recall the proofs of marginal convergence for IPF and IMF from [6] and [32] respectively, which rely on the Pythagorean relations in Thm. 3.3-3.4.

C.1 Convergence of Exact em Iterations

We first recall the form of the optimal solution to the 'static SB' problem for general reference measure $\mathbb{R}_{0,T}$ [19], which recovers the entropic- or KL -regularized OT examples for $\mathbb{R}_{0,T} = e^{-\frac{1}{\epsilon}c(x_0,x_T)}$ and $\mathbb{R}_{0,T} = \mu_0 \otimes \nu_T \; e^{-\frac{1}{\epsilon}c(x_0,x_T)}$ (see Eq. (eOT)). Under suitable conditions detailed in Léonard [19] Thm. 2.8 and 2.12, the form of the static SB solution has the form

$$\mathbb{Q}^*_{0,T} = \mathop{\arg\min}_{\mathbb{Q}_{0,T}\in\Pi(\mu_0,\nu_T)} D_{\mathrm{KL}}\left[\mathbb{Q}_{0,T} : \mathbb{R}_{0,T}\right] \qquad \Longrightarrow \qquad \frac{d\mathbb{Q}^*_{0,T}}{d\mathbb{R}_{0,T}} = e^{-\phi_0(x_0)-\phi_T(x_T)}. \tag{35}$$

Static IPF iterates may be parameterized explicitly in terms of ϕ_0, ϕ_T (see, e.g. Léger [17]). We will make use of the flexibility to choose the reference measure $\mathbb{R}_{0,T}$ in the proof of Prop. 4.2 below.

Proposition 4.2. *The exact alternating projection algorithm in Eq.* (em) *(in particular, Alg. 2) converges to* $\mathbb{Q}^{(\infty)}_{0:T} \in \mathcal{R}(\mathbb{Q}^{ref}_{\circ|0,T}) \cap \Pi(\mu_0,\nu_T)$ *and* $\mathbb{P}^{(\infty)}_{0:T} \in \mathcal{M}_{LC}$*. If* $\mathbb{Q}^{(\infty)}_{0:T}$ *or* $\mathbb{P}^{(\infty)}_{0:T}$ *is a fixed point of the procedure, then it is equal to the optimal SB solution,* $\mathbb{Q}^{(\infty)}_{0:T} = \mathbb{P}^{(\infty)}_{0:T} = \mathbb{Q}^*_{0:T}$.

Proof. We begin by showing that the e-projection onto $\mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T}) \cap \Pi(\mu_0,\nu_T)$,

$$\mathbb{Q}^{(n+2)}_{0:T} \leftarrow \mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0:T} \coloneqq \mathop{\arg\min}_{\mathbb{Q}\in\, \mathcal{R}_{\mathrm{C}}\,\cap\,\Pi(\mu_0,\nu_T)} D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}] \tag{36}$$

can be obtained by using steps from Alg. 2. First, we perform the e-projection onto $\mathcal{R}_{\mathrm{C}}$, $\mathbb{Q}^{e(\mathcal{R})}_{0:T} = \mathrm{proj}_{\mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})}(\mathbb{P}^{(n+1)}_{0:T})$, and then perform iterative IPF updates of couplings $\mathbb{Q}^{(n+i)}_{0,T}$ associated with reciprocal class measures $\mathbb{Q}^{(n+i)}_{0:T} = \mathbb{Q}^{(n+i)}_{0,T}\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T} \in \mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})$.

Consider the e-projection $\mathbb{Q}^{e(\mathcal{R})}_{0:T} = \mathrm{proj}^e_{\mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})}(\mathbb{P}^{(n+1)}_{0:T})$,

$$\mathbb{Q}^{e(\mathcal{R})}_{0:T} \coloneqq \mathop{\arg\min}_{\mathbb{Q}\in\mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})} D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}] \tag{37}$$

whose explicit form is given in Prop. 4.1. Consider the Pythagorean relation in Thm. 3.3, for arbitrary $\mathbb{Q}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})$

$$D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}] = D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{Q}^{e(\mathcal{R})}_{0:T}] + D_{\mathrm{KL}}[\mathbb{Q}^{e(\mathcal{R})}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}]. \tag{38}$$

We are eventually interested in the e-projection $\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0:T}$ of $\mathbb{P}^{(n+1)}_{0:T}$ onto the intersection $\mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T}) \cap \Pi(\mu_0,\nu_T)$ as in Eq. (36). Note that the KL divergence $D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}]$ to be minimized appears on the left-hand side of Eq. (38), with $\mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T}) \cap \Pi(\mu_0,\nu_T) \subset \mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})$.

On the right side of Eq. (38), we further note that $\mathbb{Q}^{e(\mathcal{R})}_{0:T}$ and $\mathbb{P}^{(n+1)}_{0:T}$ are fixed, so that minimizing $D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}]$ with respect to $\mathbb{Q}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})$ reduces to minimizing $D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{Q}^{e(\mathcal{R})}_{0:T}]$. Finally, the latter divergence $D_{\mathrm{KL}}[\mathbb{Q}_{0:T} : \mathbb{Q}^{e(\mathcal{R})}_{0:T}] = D_{\mathrm{KL}}[\mathbb{Q}_{0,T} : \mathbb{Q}^{e(\mathcal{R})}_{0,T}]$ reduces to a KL divergence over couplings since $\mathbb{Q}_{0,T}, \mathbb{Q}^{e(\mathcal{R})}_{0,T} \in \mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T})$.

Together, this reasoning suggests that solving for $\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\mathrm{ref}}_{\circ|0,T}) \cap \Pi(\mu_0,\nu_T)$ in Eq. (36) simply corresponds to finding

$$\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T} = \mathop{\arg\min}_{\mathbb{Q}_{0,T}\in\Pi(\mu_0,\nu_T)} D_{\mathrm{KL}}[\mathbb{Q}_{0,T} : \mathbb{Q}^{e(\mathcal{R})}_{0,T}]. \tag{39}$$

where we have dropped the reciprocal condition since we may construct the path measure as $\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T} = \mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T}\mathbb{Q}^{\text{ref}}_{\circ|0,T}$.

Eq. (39) corresponds exactly to the static problem in Eq. (35) (c.f. Eq. (eOT)), where the reference coupling measure in the second argument is now given by $\mathbb{Q}^{e(\mathcal{R})}_{0,T}$. Thus, we may use IPF iteration with $\mathbb{P}^{(0)} = \mathbb{Q}^{e(\mathcal{R})}_{0,T}$ as an inner loop to finally solve for $\mathbb{Q}^{e(\mathcal{R}\cap\Pi)}_{0,T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$. Since IPF converges to the e-projection in Eq. (39) [6, 29], we conclude that Alg. 2 converges to the appropriate projection in Eq. (em) or Eq. (36).

Convergence: Finally, to show convergence, note that $D_{\text{KL}}[\mathbb{Q}^{(n)}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}] \geq D_{\text{KL}}[\mathbb{Q}^{(n+2)}_{0:T} : \mathbb{P}^{(n+1)}_{0:T}] \geq D_{\text{KL}}[\mathbb{Q}^{(n+2)}_{0:T} : \mathbb{P}^{(n+3)}_{0:T}]$ and so on, since each exact iterate minimizes the same KL divergence. Along with the nonnegativity of the KL divergence, this implies that the procedure converges to $\mathbb{Q}^{(\infty)}_{0:T} \in \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$ and $\mathbb{P}^{(\infty)}_{0:T} \in \mathcal{M}_{\text{LC}}$.

If $\mathbb{Q}^{(\infty)}_{0:T}$ is a fixed point, then $\mathbb{Q}^{(\infty)}_{0:T} = \mathbb{P}^{(\infty)}_{0:T}$ and $\mathbb{Q}^{(\infty)}_{0:T} \in \mathcal{M}_{\text{LC}} \cap \mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)$. By Prop. 2.2, $\mathbb{Q}^{(\infty)}_{0:T} = \mathbb{Q}^{*}_{0:T}$ uniquely solves the SB problem with reference $\mathbb{Q}^{\text{ref}}_{0:T}$ and endpoints (μ_0, ν_T).

We suspect it is possible to provide a rigorous proof that if $\mathcal{M}_{\text{LC}} \cap [\mathcal{R}(\mathbb{Q}^{\text{ref}}_{\circ|0,T}) \cap \Pi(\mu_0, \nu_T)] \neq \emptyset$, i.e. if the unique SB solution exists, then $\lim_{n\to\infty} \mathbb{Q}^{(n)}_{0:T} \to \mathbb{Q}^{*}_{0:T}$ and $\lim_{n\to\infty} \mathbb{P}^{(n+1)}_{0:T} \to \mathbb{Q}^{*}_{0:T}$. □

C.2 Pythagorean Relations for Analysis of IPF and IMF

We recall the following results, which are derived from direct application of the Pythagorean relations. However, additional reasoning is required to show convergence to the optimal solution, for example $\lim_{n\to\infty} D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(n)}_{0,T}] = 0$ for IPF ([29] Sec. 3) and $\lim_{n\to\infty} D_{KL}[\mathbb{Q}^{(n)}_{0:T} : \mathbb{Q}^{*}_{0:T}] = 0$ for IMF ([27] Thm. 2).

Proposition C.1. [Convergence of IPF Marginals] *Consider initializing IPF with* $\mathbb{P}^{(0)}_{0,T}$ *such that* $D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(0)}_{0,T}] < \infty$*. Then, IPF iterates satisfy* $\lim_{n\to\infty} D_{KL}[\mathbb{P}^{(n+1)}_{0,T} : \mathbb{P}^{(n)}_{0,T}] = 0$ *and converge to the correct marginals,* $\lim_{n\to\infty} D_{KL}[\mathbb{P}^{(n)}_{0} : \mu_0] = 0$ *and* $\lim_{n\to\infty} D_{KL}[\mathbb{P}^{(n+1)}_{T} : \nu_T] = 0$.

Proof. Following [6] Thm. 3.2, [29] Prop 2.1, consider projecting $\mathbb{P}^{(0)}_{0,T}$ onto the convex set $\Pi(\cdot, \nu_T)$ using $\mathbb{P}^{(1)}_{0,T} = \underset{\mathbb{Q}_{0,T}\in\Pi(\cdot,\nu_T)}{\arg\min}\ D_{KL}[\mathbb{Q}_{0,T} : \mathbb{P}^{(0)}_{0,T}]$. We will use the Pythagorean relation for the e-projection for any $\mathbb{Q}_{0,T} \in \Pi(\cdot, \nu_T)$. In particular, consider the optimal solution $\mathbb{Q}_{0,T} = \mathbb{Q}^{*}_{0,T}$ to the EOT problem, with $\mathbb{Q}^{*}_{0,T} \in \Pi(\mu_0, \nu_T)$ and

$$D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(0)}_{0,T}] = D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(1)}_{0,T}] + D_{KL}[\mathbb{P}^{(1)}_{0,T} : \mathbb{P}^{(0)}_{0,T}]. \tag{40}$$

Next, project $\mathbb{P}^{(1)}_{0,T} \in \Pi(\cdot, \nu_T)$ onto $\Pi(\mu_0, \cdot)$ using $\mathbb{P}^{(2)}_{0,T} = \underset{\mathbb{Q}_{0,T}\in\Pi(\mu_0,\cdot)}{\arg\min}\ D_{KL}[\mathbb{Q}_{0,T} : \mathbb{P}^{(1)}_{0,T}]$. Using the Pythagorean relation,

$$D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(1)}_{0,T}] = D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(2)}_{0,T}] + D_{KL}[\mathbb{P}^{(2)}_{0,T} : \mathbb{P}^{(1)}_{0,T}] \tag{41}$$

which we can plug into Eq. (40). Continuing to apply these ite rations, we have

$$D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(0)}_{0,T}] = D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(n)}_{0,T}] + \sum_{i=0}^{n-1} D_{KL}[\mathbb{P}^{(i+1)}_{0,T} : \mathbb{P}^{(i)}_{0,T}] \tag{42}$$

which implies $\sum_{i=0}^{n-1} D_{KL}[\mathbb{P}^{(i+1)}_{0,T} : \mathbb{P}^{(i)}_{0,T}] \leq D_{KL}[\mathbb{Q}^{*}_{0,T} : \mathbb{P}^{(0)}_{0,T}] < \infty$ by the nonnegativity of KL divergence and the assumption on the initial $\mathbb{P}^{(0)}_{0,T}$. Since the sum can not grow to infinity, we have $\lim_{n\to\infty} D_{KL}[\mathbb{P}^{(n+1)}_{0,T} : \mathbb{P}^{(n)}_{0,T}] = 0$ as desired.

To show that $\lim_{n\to\infty} D_{KL}[\mathbb{P}_T^{(n+1)} : \nu_T] = 0$, note that $\mathbb{P}_{0,T}^{(n)} \in \Pi(\cdot, \nu_T)$ for even iterations. We use the fact that marginalization can only reduce the KL divergence and the limiting behavior of $\lim_{n\to\infty} D_{KL}[\mathbb{P}_{0,T}^{(n+1)} : \mathbb{P}_{0,T}^{(n)}] = 0$ to conclude that

$$D_{KL}[\mathbb{P}_T^{(n+1)} : \nu_T] \leq D_{KL}[\mathbb{P}_{0,T}^{(n+1)} : \mathbb{P}_{0,T}^{(n)}] \implies \lim_{n\to\infty} D_{KL}[\mathbb{P}_T^{(n+1)} : \nu_T] \leq \lim_{n\to\infty} D_{KL}[\mathbb{P}_{0,T}^{(n+1)} : \mathbb{P}_{0,T}^{(n)}] = 0$$

which shows $\lim_{n\to\infty} D_{KL}[\mathbb{P}_T^{(n+1)} : \nu_T] = 0$. Similar reasoning applies for convergence to μ_0. □

Proposition C.2. [Convergence of IMF Marginals] *Consider initializing IMF with $\mathbb{Q}_{0:T}^{(0)} \in \Pi(\mu_0, \nu_T)$ such that $D_{KL}[\mathbb{Q}_{0:T}^{(0)} : \mathbb{Q}_{0:T}^*] < \infty$. Then, IMF iterates satisfy $\lim_{n\to\infty} D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^{(n+1)}] = 0$. If $D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^*] = D_{KL}[\mathbb{Q}_{0:T}^{(n+1)} : \mathbb{Q}_{0:T}^*]$ or $D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^{(n+1)}] = 0$, then $\mathbb{Q}_{0:T}^{(n)} = \mathbb{Q}_{0:T}^{(n+1)} = \mathbb{Q}_{0:T}^*$ is the optimal SB solution (under the conditions of Prop. 2.2).*

Proof. Following [27] Thm. 2, [32] Prop. 7, first consider projecting $\mathbb{Q}_{0:T}^{(0)}$ onto the set of Markov measures $\mathcal{M}_{\text{LC}}$ using $\mathbb{Q}_{0:T}^{(1)} = \underset{\mathbb{P}_{0:T} \in \mathcal{M}_{\text{LC}}}{\arg\min} D_{KL}[\mathbb{Q}_{0:T}^{(0)} : \mathbb{P}_{0:T}]$. For $\mathbb{Q}_{0:T}^* \in \mathcal{M}_{\text{LC}}$, we use the Pythagorean relation to write

$$D_{KL}[\mathbb{Q}_{0:T}^{(0)} : \mathbb{Q}_{0:T}^*] = D_{KL}[\mathbb{Q}_{0:T}^{(0)} : \mathbb{Q}_{0:T}^{(1)}] + D_{KL}[\mathbb{Q}_{0:T}^{(1)} : \mathbb{Q}_{0:T}^*]. \tag{43}$$

Next, project $\mathbb{Q}_{0:T}^{(1)}$ onto the reciprocal class $\mathcal{R}_{\text{LC}}$ using $\mathbb{Q}_{0:T}^{(2)} = \underset{\mathbb{P}_{0:T} \in \mathcal{R}_{\text{LC}}}{\arg\min} D_{KL}[\mathbb{Q}_{0:T}^{(1)} : \mathbb{P}_{0:T}]$. For $\mathbb{Q}_{0:T}^* \in \mathcal{R}_{\text{LC}}$, the Pythagorean relation implies

$$D_{KL}[\mathbb{Q}_{0:T}^{(1)} : \mathbb{Q}_{0:T}^*] = D_{KL}[\mathbb{Q}_{0:T}^{(1)} : \mathbb{Q}_{0:T}^{(2)}] + D_{KL}[\mathbb{Q}_{0:T}^{(2)} : \mathbb{Q}_{0:T}^*] \tag{44}$$

Plugging back into Eq. (43) and iterating the above decomposition steps, we obtain

$$D_{KL}[\mathbb{Q}_{0:T}^{(0)} : \mathbb{Q}_{0:T}^*] = D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^*] + \sum_{i=0}^{n-1} D_{KL}[\mathbb{Q}_{0:T}^{(i)} : \mathbb{Q}_{0:T}^{(i+1)}] \tag{45}$$

which implies $\sum_{i=0}^{n-1} D_{KL}[\mathbb{Q}_{0:T}^{(i)} : \mathbb{Q}_{0:T}^{(i+1)}] \leq D_{KL}[\mathbb{Q}_{0:T}^{(0)} : \mathbb{Q}_{0:T}^*] < \infty$ by the nonnegativity of KL divergence and the assumption on the initial $\mathbb{Q}_{0:T}^{(0)}$. Since the sum can not grow to infinity, we have $\lim_{n\to\infty} D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^{(n+1)}] = 0$ as desired.

Using the Pythagorean relation such as in Eq. (44), we can see that $D_{KL}[\mathbb{Q}_{0:T}^{(n+1)} : \mathbb{Q}_{0:T}^*] \leq D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^*]$ with equality iff $D_{KL}[\mathbb{Q}_{0:T}^{(n)} : \mathbb{Q}_{0:T}^{(n+1)}] = 0$. This implies that $\mathbb{Q}_{0:T}^{(n)} = \mathbb{Q}_{0:T}^{(n+1)} \in \mathcal{M}_{\text{LC}} \cap \mathcal{R}_{\text{LC}}$ since $\mathbb{Q}_{0:T}^{(n)} \in \mathcal{M}_{\text{LC}}$ and $\mathbb{Q}_{0:T}^{(n+1)} \in \mathcal{R}_{\text{LC}}$. Since exact reciprocal and Markov projections preserve the condition $\mathbb{Q}_{0:T}^{(0)} \in \Pi(\mu_0, \nu_T)$ in later iterations ([27] Thm. 1), we have $\mathbb{Q}_{0:T}^{(n)} \in \mathcal{M}_{\text{LC}} \cap \mathcal{R}_{\text{LC}} \cap \Pi(\mu_0, \nu_T)$, which is the optimal SB solution by Prop. 2.2. □