

Schrödinger bridge for generative AI: Soft-constrained formulation and convergence analysis*

In celebration of Prof. Xun Yu Zhou's 60th birthday

Jin Ma[†], Ying Tan[‡] and Renyuan Xu[§]

October 29, 2025

Abstract

Generative AI can be framed as the problem of learning a model that maps simple reference measures into complex data distributions, and it has recently found a strong connection to the classical theory of the Schrödinger bridge problems (SBPs) due partly to their common nature of interpolating between prescribed marginals via entropy-regularized stochastic dynamics. However, the classical SBP enforces hard terminal constraints, which often leads to instability in practical implementations, especially in high-dimensional or data-scarce regimes. To address this challenge, we follow the idea of the so-called *soft-constrained Schrödinger bridge problem* (SCSBP), in which the terminal constraint is replaced by a general penalty function. This relaxation leads to a more flexible stochastic control formulation of McKean–Vlasov type.

We establish the existence of optimal solutions for all penalty levels and prove that, as the penalty grows, both the controls and value functions converge to those of the classical SBP at a linear rate. Our analysis builds on Doob's h -transform representations, the stability results of Schrödinger potentials, Γ -convergence, and a novel fixed-point argument that couples an optimization problem over the space of measures with an auxiliary entropic optimal transport problem. These results not only provide the first quantitative convergence guarantees for soft-constrained bridges but also shed light on how penalty regularization enables robust generative modeling, fine-tuning, and transfer learning.

Keywords. Schrödinger bridge, soft-constrained Schrödinger bridge, entropic optimal transport stability, Schrödinger potentials, generative AI, h -transform, converse of Scheffé's Theorem, Γ -convergence, Schauder's fixed-point.

2020 AMS Mathematics subject classification: 60H10, 60J60, 49J21, 68T01, 93E20.

*We thank helpful comments and discussions from Beatrice Acciaio, Rama Cont, Huyền Pham, Xin Zhang, and Xun Yu Zhou.

[†]Department of Mathematics, University of Southern California, Los Angeles, CA, 90089. Email: jinma@usc.edu. This author is supported in part by US NSF grants DMS#2205972 and #2510403.

[‡]Department of Statistics and Applied Probability, University of California, Santa Barbara, CA 93106. Email: yingtang@ucsb.edu.

[§]Department of Management of Science & Engineering, Stanford University, Stanford, CA 94305. Email: renyuanxu@stanford.edu. This author is supported in part by the NSF CAREER Award DMS-2524465.

1 Introduction

Generative modeling provides a powerful framework for synthesizing data that preserves the statistical structure of real-world samples while introducing controlled variability. Among the most prominent approaches, diffusion models — such as those introduced by [36, 55, 57, 58] — have achieved remarkable success, underpinning state-of-the-art systems like DALL·E 2 and 3 [6, 51], Stable Diffusion [53], and Sora [47]. These models learn to reverse a diffusion process that gradually adds noise to data, enabling the generation of realistic samples from pure noise. Such a structure, namely, transforming a noise distribution into a data distribution, closely mirrors the Schrödinger bridge problem (or dynamic optimal transport), which has recently gained renewed attention as a principled framework for generative modeling due to its structural parallels with diffusion models and its ability to interpolate between distributions in a statistically grounded manner.

The Schrödinger bridge problem (SBP for short), originally proposed as an entropy-regularized variant of optimal transport, seeks the most likely evolution of a process, subject to a reference diffusion process, that matches prescribed marginal distributions $\mu_{\text{ini}}, \mu_{\text{tar}}$ at two endpoints. Under suitable regularity conditions, the optimally controlled process remains a diffusion but with an additional drift term added to the reference process. This result has been established through various approaches and levels of generality, with seminal contributions by [7, 18, 25, 26, 37].

The recent generative modeling literature has seen a surge in the use of Schrödinger bridges. In these applications, one typically starts or chooses some distribution μ_{ini} that is easy to sample from, and tries to “learn” an unknown distribution μ_{tar} of a given data set. By numerically approximating the solution to the Schrödinger bridge problem, one can generate unlimited samples (i.e., synthetic data points that resemble the original data set). One such algorithm is presented by De Bortoli et al. [20] and Vargas et al. [64], who approximate the iterative proportional fitting procedure (Deming-Stephan [21]), using score matching with neural networks and maximum likelihood, respectively. Concurrently, Wang et al. [66] proposed a two-stage method with an auxiliary bridge handling possible non-smooth μ_{tar} . Some more recent developments include [2, 13, 32, 49, 52, 54, 56, 62, 63, 67] as well as developing optimal transport techniques for generative AI tasks [1, 5, 42, 44, 68].

However, the classical SBP imposes hard terminal constraints on the marginal distributions, which can result in computational difficulties, instability in high-dimensional settings, and limited adaptability when aligning with empirical data in generative tasks. In practice, most numerical schemes for solving the SBP rely on iterative procedures that alternately relax the initial and terminal constraints. These algorithms can exhibit instability, particularly when the two constraints differ significantly, and their convergence guarantees for general target distributions remain an open problem in the literature.

In this paper, in light of Garg et al. [29], we study a soft-constrained Schrödinger bridge problem (SCSBP). Mathematically, we consider a penalty function $G : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}_+$, satisfying $G(\mu; \mu_{\text{tar}}) = 0$ if and only if $\mu = \mu_{\text{tar}}$, where $\mathcal{P}_2(\mathbb{R}^d)$ denotes all the probability measures on $\mathbb{R}^d$ with finite second moment equipped with 2-Wasserstein metric, and $\mu_{\text{tar}} \in \mathcal{P}_2(\mathbb{R}^d)$ is some given “target” distribution. For each $k \in \mathbb{N}$ and a given initial distribution $\mu_{\text{ini}} \in \mathcal{P}_2(\mathbb{R}^d)$, we consider the following stochastic control problem with dynamics on a given probability space $(\Omega, \mathcal{F}, \mathbb{P})$:

$$dX_t^\alpha = (b(t, X_t^\alpha) + \sigma(t)\alpha_t)dt + \sigma(t)dW_t, \quad \mathbb{P} \circ (X_0^\alpha)^{-1} = \mu_{\text{ini}},$$

and cost functional

$$J^k(\alpha) = \mathbb{E} \left[\frac{1}{2} \int_0^T |\alpha_s|^2 ds + kG(\mathbb{P}_{X_T^\alpha}) \right],$$

where $\mathbb{P}_{X_T^\alpha}$ is the law of X_T^α under probability $\mathbb{P}$ and the control α is chosen from a square-integrable progressively measurable *admissible control* set $\mathcal{A}$. The goal is to find $V^k := \inf_{\alpha \in \mathcal{A}} J^k(\alpha)$ and

optimal control α^k , for each $k \in \mathbb{N}$, and study the limiting behavior of $\{\alpha^k\}$ and $\{V^k\}$. Clearly, the dependence of the terminal cost on $\mathbb{P}_{X_T^\alpha}$ renders this relaxed formulation a *non-standard* stochastic control problem, leading to a McKean-Vlasov type of stochastic control. In contrast to Garg et al. [29], which focuses on the case where the penalty G is given by the KL divergence and μ_{ini} is a delta measure, we investigate the problem under more general cost functions and initial distributions.

Compared to existing methods using SBP under hard constraints, there are several advantages to considering SCSBP. First, when the KL divergence between μ_{tar} and $\mathbb{P}_{X_T^\alpha}$ is infinite, the Schrödinger bridge does not admit a solution, whereas SCSBP always does (cf. e.g., [29]). More importantly, the penalty parameter k acts as a regularization factor, preventing the algorithm from overfitting to μ_{tar} , which is crucial for certain generative modeling tasks with limited data [30, 45]. In addition, SCSBP provides a more general framework in generative AI, with applications beyond pre-training in data generation, and can be applied to fine-tuning and transfer learning (see Examples 2.6 and 2.7).

1.1 Outline of the Main Results and Contributions

The soft-constrained Schrödinger bridge problem (SCSBP) studied in this paper replaces the terminal distribution constraint by a general penalty function, which leads to a McKean–Vlasov type stochastic control problem. The main results include the existence of the optimal solution to the SCSBP at each penalty level k , and the convergence of the solutions to the SCSBP to that of the SBP, in terms of both the optimal policy and the corresponding value function, as $k \rightarrow \infty$.

More precisely, we begin with the special case where the initial measure is a delta measure. In this setting, we derive the explicit form of the optimal control policy for SCSBP via Doob’s h -transform (Proposition 3.6) and more importantly, we establish a linear convergence rate for the optimal control (Theorem 4.1). To the best of our knowledge, such a rate of convergence is novel in the literature. Moreover, by applying the so-called *early stopping*, we are able to obtain the linear convergence results for the corresponding value functions (Proposition 4.4) as well as the Wasserstein distance between the target distribution and the output distribution of the SCSBP (Proposition 4.5).

The similar results in the case of a general initial distribution is much more involved. Among other things, we establish and/or extend some recently observed stability results of the solutions to the SBP, as the foundation for a fixed-point argument. A key element in our argument is the continuous dependence (or stability) of a mapping that is well known in the (static) optimal transport literature. Specifically, fix $\mu_{\text{ini}} \in \mathcal{P}_2(\mathbb{R}^d)$ and consider any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, it is known (cf. e.g., [7]) that there exists a unique pair of σ -finite measures $(\nu_0, \nu_T) := \mathcal{T}(\mu)$ such that the measure π on $\mathbb{R}^d \times \mathbb{R}^d$ defined by

$$\pi(E) = \int_E p(T, y; 0, x) \nu_0(dx) \nu_T(dy), \quad E \in \mathcal{B}(\mathbb{R}^d \times \mathbb{R}^d) \quad (1.1)$$

has the marginals μ_{ini} and μ , where $p(\cdot, \cdot; \cdot, \cdot)$ is the transition density of a given diffusion process. Then it turns out that the solution to the SCSBP is the fixed-point of the mapping $\Gamma : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^d)$, where

$$\Gamma(\mu) := \arg \min_{\bar{\mu} \in \mathcal{P}_2(\mathbb{R}^d)} \{kG(\bar{\mu}) + \mathbb{E}_{X \sim \bar{\mu}}[\log \rho^\mu(X)]\}, \quad (1.2)$$

where ρ^μ is the density of the σ -finite measure ν in (1.1) (which may not be a probability measure). The successful application of Schauder’s fixed-point theorem on the Wasserstein space relies on several key elements, in particular a crucial continuous dependence result of the mapping $\mu \mapsto \rho^\mu$, for

which we introduce an auxiliary *entropic optimal transport problem*, and identify its solution to the measure π in (1.1). By utilizing some important stability results of the corresponding Schrödinger potentials (Proposition 5.9), together with some arguments in the spirit of the converse of Scheffé’s theorem (i.e., the weak convergence in Prohorov metric vs. the convergence in densities) as well as the so-called Γ -convergence of the minimizers of the optimization problem (1.2), we are able to verify the required properties so the mapping Γ has a fixed-point (Theorem 6.1). As a direct consequence, we establish that the optimal control of the SCSBP converges linearly with respect to the penalty parameter k (Theorem 6.4). We believe that such a fixed-point perspective is novel in the literature, as it not only offers a constructive framework for characterizing solutions in the general case but also yields insights into how the penalty parameter affects the convergence rate.

1.2 Closely Related Literature

Our general formulation is largely inspired by Garg et al. [29], which investigates the SCSBP with the KL divergence as the penalty function G . Within that framework, the authors established the asymptotic convergence of the optimal policies as the penalty parameter k tends to infinity, under the assumption that the initial measure is a delta measure. However, the use of KL divergence presents practical difficulties: if the model distribution μ assigns zero probability to any region where the data distribution μ_{tar} has positive mass, then $D_{\text{KL}}(\mu_{\text{tar}}\|\mu) = +\infty$, rendering the divergence ill-posed under support mismatch [15, 38]. Moreover, a delta initial measure is rarely employed in generative tasks, as it lacks the diversity and randomness required for effective training. Finally, no convergence rate is quantified therein.

On a technical level, our formulation is closely related to Hernández-Tangpi [35], in which the authors use a probabilistic approach to recast a mean-field Schrödinger bridge into a stochastic optimization problem with McKean-Vlasov dynamics, and connect the optimal control to a solution to a forward backward SDE of McKean-Vlasov type (MKV FBSDE). However, given the generality of the drift, diffusion, and running cost functions, the associated MKV FBSDE is derived without a discussion of uniqueness. In fact, it is not completely obvious that a McKean-Vlasov-type SBP can be converted to a McKean-Vlasov stochastic control problem via the usual Girsanov theorem argument, as we show in Remark 2.1 below. Moreover, the conditions imposed on the penalty function G are abstract and can be difficult to verify in common examples. By contrast, our framework leverages the PDE formulation and Doob’s h -transform representation, requiring only mild growth conditions on G and control of density gaps (see Assumption (3.2)). Several concrete examples of admissible G are provided in Example 3.4 and 3.5.

The rest of the paper is organized as follows. Section 2 introduces the necessary concepts and notations. In particular, we present the connection between the underlying SBP and its stochastic control formulation, and introduce the notion of the SCSBP together with some potential applications. Section 3 is devoted to the existence of optimal policies for the SCSBP at each penalty parameter k , while Section 4.1 establishes that the penalized optimal policies converge to those of the original SBP as $k \rightarrow \infty$, with a linear rate of convergence. In addition, we prove convergence of the corresponding value functions and quantify the distance between the terminal distribution and the target distribution in terms of the Wasserstein distance. Sections 5 and 6 are devoted to the case with a general initial distribution. A crucial stability result is established in Section 5, via the stability of Schrödinger potentials of an auxiliary entropy optimal transport problem, and in Section 6 we complete the fixed-point argument.

2 Preliminary

Throughout this paper, we consider a generic Euclidean space $\mathbb{X}$, and regardless of its dimension, we denote $(\cdot, \cdot)$ and $|\cdot|$ be its inner product and norm, respectively. We denote $\mathbb{C}([0, T]; \mathbb{X})$ to be the space of $\mathbb{X}$ -valued continuous functions defined on $[0, T]$ with the usual sup-norm, and in particular, we denote $\mathbb{C}_T^d := \mathbb{C}([0, T]; \mathbb{R}^d)$, and let $\mathcal{B}(\mathbb{C}_T^d)$ be its topological Borel field. We shall consider the following *canonical probabilistic space*: $(\Omega, \mathcal{F}, \mathbb{P})$, where $(\Omega, \mathcal{F}) := (\mathbb{C}_T^d, \mathcal{B}(\mathbb{C}_T^d))$ and $\mathbb{P} \in \mathcal{P}(\Omega)$, the space of all probability measures defined on $(\Omega, \mathcal{F})$. Finally, we let $\mathbb{P}^0 \in \mathcal{P}(\Omega)$ be the Wiener measure on $(\Omega, \mathcal{F})$; $W_t(\omega) := \omega(t)$, $\omega \in \Omega$, the canonical process; and $\mathbb{F}^0 := \{\mathcal{F}_t^0\}_{t \in [0, T]}$, where $\mathcal{F}_t^0 := \mathcal{B}_t(\mathbb{C}_T^d) := \sigma\{\omega(\cdot \wedge t) : \omega \in \mathbb{C}_T^d\}$, $t \in [0, T]$. We assume that $\mathbb{F}^0$ has the usual $\mathbb{P}^0$ -augmentation so that it satisfies the *usual hypotheses* (cf. e.g., [50]), and for $p \geq 1$, we denote $\mathbb{L}_{\mathbb{F}^0}^p([0, T]; \mathbb{X})$ to be all $\mathbb{X}$ -valued, p -integrable, $\mathbb{F}^0$ -adapted processes. Finally, we denote $\mathcal{M}(\mathbb{R}^d)$ to be all σ -finite measures on $\mathbb{R}^d$ and $\mathcal{P}_p(\mathbb{R}^d)$ to be all probability measures with finite p -th moment on $\mathbb{R}^d$ equipped with p -Wasserstein metric, denoted by $W_p(\cdot, \cdot)$.

We recall that a classical Schrödinger bridge problem (SBP) amounts to finding, for $\mathbb{P} \in \mathcal{P}(\Omega)$,

$$V(\mathbb{P}) := \inf_{\mathbb{Q} \in \mathcal{P}(\mu_{\text{ini}}, \mu_{\text{tar}})} D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}), \quad (2.1)$$

where $D_{\text{KL}}(\cdot \parallel \cdot)$ is the *Kullback-Leibler Divergence* or the *Relative Entropy* (cf. [39])¹, defined by

$$D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}) := \begin{cases} \mathbb{E}^{\mathbb{Q}}[\log(\frac{d\mathbb{Q}}{d\mathbb{P}})], & \text{if } d\mathbb{Q} \ll d\mathbb{P}; \\ \infty, & \text{otherwise.} \end{cases}$$

We note that the admissible set $\mathcal{P}(\mu_{\text{ini}}, \mu_{\text{tar}}) \subseteq \mathcal{P}(\Omega)$ in (2.1) is defined as the collection of all probability measures $\mathbb{Q} \in \mathcal{P}(\Omega)$, such that $\mathbb{Q}$ is the law of some continuous $\mathbb{R}^d$ -valued continuous process X defined on the canonical space with probability measure $\bar{\mathbb{P}} \in \mathcal{P}(\Omega)$ (i.e., $\mathbb{Q} = \bar{\mathbb{P}} \circ X^{-1}$) satisfying $\mathbb{Q} \circ X_0^{-1} = \mu_{\text{ini}}$ and $\mathbb{Q} \circ X_T^{-1} = \mu_{\text{tar}}$. In other words, $\mathcal{P}(\mu_{\text{ini}}, \mu_{\text{tar}})$ contains all “path measures” with initial and terminal distributions being μ_{ini} and μ_{tar} , respectively.

We note that if $\mathbb{P} = \mathbb{P}^0$ is Wiener measure and $\mathbb{Q} \in \mathcal{P}(\Omega)$ is absolutely continuous with respect to $\mathbb{P}^0$, then by the Cameron-Martin-Girsanov theorem, there exists $\alpha \in \mathbb{L}_{\mathbb{F}^0}^2([0, T]; \mathbb{R}^d)$, such that

$$\frac{d\mathbb{Q}}{d\mathbb{P}^0} \Big|_{\mathcal{F}_t} := \mathcal{E}_t(\alpha) = \exp \left\{ \int_0^t \alpha_s dW_s - \frac{1}{2} \int_0^t |\alpha_s|^2 ds \right\}, \quad t \in [0, T],$$

is a $\mathbb{P}^0$ -martingale, and $\widetilde{W}_t := W_t - \int_0^t \alpha_s ds$, $t \in [0, T]$, is a $\mathbb{Q}$ -Brownian motion. In this case, one can then easily check that

$$D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}^0) = \frac{1}{2} \mathbb{E}^{\mathbb{Q}} \left[\int_0^T |\alpha_t|^2 dt \right]. \quad (2.2)$$

Schrödinger Bridge and Related Control Problem. In light of (2.2), one can easily associate a SBP to a stochastic control problem (see, e.g., [14, §4.4], [17, §1] and [40, p36]). Consider a standard SDE on a canonical space $(\Omega, \mathcal{F}, \mathbb{P})$, with $\mathbb{P} \in \mathcal{P}(\Omega)$:

$$dX_t = b(t, X_t)dt + dW_t, \quad \mathbb{P} \circ (X_0)^{-1} = \mu_{\text{ini}}. \quad (2.3)$$

¹We remark here that the KL divergence $D_{\text{KL}}(\cdot \parallel \cdot)$ can be easily extended to any σ -finite measures. In this case $D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}) := \int \log(\frac{\mathbb{Q}(dx)}{\mathbb{P}(dx)}) \mathbb{Q}(dx)$, if $d\mathbb{Q} \ll d\mathbb{P}$.

In what follows we shall denote, for $t \in [0, T]$, $\mathbb{P}_{X_t} := \mathbb{P} \circ (X_t^{-1}) \in \mathcal{P}_2(\mathbb{R}^d)$.

Then we consider the following SBP:

$$V(\mu_{\text{ini}}, \mu_{\text{tar}}) := \inf_{\mathbb{Q} \in \mathcal{P}(\mu_{\text{ini}}, \mu_{\text{tar}})} D_{\text{KL}}(\mathbb{Q} \parallel \mathbb{P}). \quad (2.4)$$

Similar to (2.2), we now recast (2.4) as the following stochastic control problem:

$$V(\mu_{\text{ini}}, \mu_{\text{tar}}) = \inf_{\alpha \in \mathcal{A}} J(\alpha) = \inf_{\alpha \in \mathcal{A}} \mathbb{E}^{\mathbb{Q}} \left[\frac{1}{2} \int_0^T |\alpha_s|^2 ds \right], \quad (2.5)$$

where $\mathbb{Q} \in \mathcal{P}(\Omega)$ is such that $\frac{d\mathbb{Q}}{d\mathbb{P}} = \mathcal{E}(\alpha)$ for some $\alpha \in \mathcal{A} \subseteq \mathbb{L}_{\mathbb{F}_0}^2([0, T]; \mathbb{R}^d)$, under which the underlying controlled dynamics takes the form:

$$dX_t^\alpha = [b(t, X_t^\alpha) + \alpha_t]dt + d\widetilde{W}_t, \quad \mathbb{Q}_{X_0^\alpha} = \mu_{\text{ini}}, \quad \mathbb{Q}_{X_T^\alpha} = \mu_{\text{tar}}, \quad (2.6)$$

where $\widetilde{W}$ is a $\mathbb{Q}$ -Brownian motion.

Remark 2.1 (Subtlety in formulating the McKean-Vlasov version of the problem). It is rather tempting to apply the idea above to the so-called McKean-Vlasov SBP (MVSBP). Consider, for example, the following McKean-Vlasov SDE on $(\Omega, \mathcal{F}, \mathbb{P})$:

$$dX_t = b(t, X_t, \mathbb{P}_{X_t})dt + dW_t, \quad \mathbb{P}_{X_0} = \mu_{\text{ini}}.$$

Similar to (2.4) we can define an SBP, which we shall refer to as an MVSBP. Again, by (2.2), we recast such MVSBP as the following (weak form) stochastic control problem:

$$V(\mu_{\text{ini}}, \mu_{\text{tar}}) := \inf_{\alpha} \mathbb{E}^{\mathbb{Q}} \left[\int_0^T \frac{1}{2} |\alpha_t|^2 dt \right]. \quad (2.7)$$

Then, $\mathbb{Q} \in \mathcal{P}(\Omega)$ must be such that the underlying controlled process X^α becomes, under $\mathbb{Q}$:

$$dX_t^\alpha = [\alpha_t + b(t, X_t^\alpha, \mathbb{P}_{X_t^\alpha})]dt + d\widetilde{W}_t, \quad \mathbb{Q}_{X_0^\alpha} = \mu_{\text{ini}}, \quad \mathbb{Q}_{X_T^\alpha} = \mu_{\text{tar}}, \quad (2.8)$$

where $\widetilde{W}$ is a $\mathbb{Q}$ -Brownian motion, and $\alpha \in \mathbb{L}_{\mathbb{F}_0}^2([0, T]; \mathbb{R}^d)$. However, by looking at (2.8) more closely we see that the problem (2.7) and (2.8) *do not* form a McKean-Vlasov control problem, since $\mathbb{P}_{X_t} \neq \mathbb{Q}_{X_t}(!)$. Therefore, an MVSBP should be formulated more carefully so as to connect to an McKean-Vlasov stochastic control problem. $\square$

Ideally, the optimal solution to the Schrödinger bridge problem (2.5)-(2.6) provides a transport map from the initial distribution μ_{ini} to the target distribution μ_{tar} . This transport is interpolated by a diffusion process that most closely resembles the canonical Brownian motion in the space of path measures. However, designing training algorithms to (approximately) learn the optimal solution to (2.5)–(2.6) typically involves an iterative scheme that alternately relaxes the initial and terminal constraints [20, 54, 64, 66], and whose convergence rate and computational complexity in high-dimensional settings remain unclear.

While there could be techniques in stochastic control theory to deal with such a constraint, we shall follow the idea of [35], by approximating the original control problem by a family of *unconstrained* McKean-Vlasov stochastic control problems with terminal penalties. More precisely, we shall allow the law of X_T^α in (2.6) to be different from μ_{tar} , but add a corresponding penalty function to it in the cost functional $J(\cdot)$.

To this end, let us still denote $\mathcal{A} \subseteq \mathcal{P}(\Omega)$ to be the set of all $\mathbb{Q} \in \mathcal{P}(\Omega)$ such that

- (i) $\frac{d\mathbb{Q}}{d\mathbb{P}}\Big|_{\mathcal{F}_T} = \mathcal{E}(\alpha) = \exp\left\{\int_0^T \alpha_s dW_s - \frac{1}{2} \int_0^T |\alpha_s|^2 ds\right\}$, $\alpha \in \mathbb{L}_{\mathbb{F}^0}^2([0, T]; \mathbb{R}^d)$;
- (ii) Under $\mathbb{Q}$, the underlying state process X follows the dynamics:

$$dX_t^\alpha = [b(t, X_t^\alpha) + \sigma(t)\alpha_t]dt + \sigma(t)d\widetilde{W}_t, \quad \mathbb{Q}_{X_0^\alpha} = \mu_{\text{ini}}, \quad (2.9)$$

where $\widetilde{W}$ is a $\mathbb{Q}$ -Brownian motion.

Throughout this paper, we shall make the following *Standing Assumption* on the coefficients b and σ .

Assumption 2.2. *The coefficients $b : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : [0, T] \rightarrow \mathbb{R}^{d \times d}$ are given deterministic continuous functions, such that there exists $L > 0$, it holds that*

$$|b(t, x) - b(t, y)| \leq L|x - y|, \quad t \in [0, T], \quad x, y \in \mathbb{R}^d.$$

Clearly, under Assumption 2.2, the SDE (2.9) has a unique strong solution X^α on $(\Omega, \mathcal{F}, \mathbb{P}^0)$ for any given $\alpha \in \mathbb{L}_{\mathbb{F}^0}^1([0, T]; \mathbb{R}^{d \times d})$ (see [50, 69]). We shall often identify $\mathbb{Q} \in \mathcal{A}$ with its associated process α , and denote $\mathbb{Q} \sim \alpha$ and $\alpha \in \mathcal{A}$ when the context is clear. The key element of the soft-constrained Schrödinger bridge problem is the following penalty function.

Definition 2.3. *A function $G(\cdot) = G(\cdot; \mu_{\text{tar}}) : \mathcal{P}_2(\mathbb{R}^d) \rightarrow [0, \infty)$ is called a continuous penalty function associated to $\mu_{\text{tar}} \in \mathcal{P}_2(\mathbb{R}^d)$ if: (1) $G(\cdot)$ is continuous on $\mathcal{P}_2(\mathbb{R}^d)$; and (2) $G(\mu; \mu_{\text{tar}}) = 0$ if and only if $\mu = \mu_{\text{tar}}$. $\square$*

Now let us introduce the following family of McKean-Vlasov-type stochastic control problems:

Problem 2.4 (Soft-constrained Schrödinger bridge problem (SCSBP)). *For $k \in \mathbb{N}$, find $\hat{\alpha}^k \in \mathcal{A}$ such that $V^k := J^k(\hat{\alpha}^k) = \inf_{\alpha \in \mathcal{A}} J^k(\alpha)$, where*

$$J^k(\alpha) = \mathbb{E}^\mathbb{Q} \left[\frac{1}{2} \int_0^T |\alpha_s|^2 ds + kG(\mathbb{Q}_{X_T^\alpha}) \right], \quad (2.10)$$

and $G(\cdot) = G(\cdot; \mu_{\text{tar}})$ is the given penalty function satisfying Definition 2.3 and $\mathbb{Q} \sim \alpha$. $\square$

Applications in Generative AI. We remark that the SCSBP Problem 2.4 offers a general framework that can be applied to address multiple problems in generative AI. We briefly mention a few motivational examples.

Example 2.5 (Data generation). The goal of generative AI is to train a data generation procedure using a finite number of iid. data samples $\{x_1, \dots, x_N\}$ under a (unknown) target distribution μ_{tar} , in order to simulate unlimited number of data samples whose underlying distribution is close to μ_{tar} [33, 36, 58].

To cast this problem into our framework, we can take, for example, $\mu_{\text{ini}} = \mathcal{N}(0, I)$ and $\mu_{\text{tar}} = p_{\text{data}}$ in the theoretical framework (or $\mu_{\text{tar}} = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$ in the practical implementation). Then the optimal control $\hat{\alpha}$ of SCSBP leads to a controlled process $(X_t^{\hat{\alpha}})_{0 \leq t \leq T}$ that simulates the data output $X_T^{\hat{\alpha}}$. Our key results (see Theorem 4.1 and Theorem 6.4 below) show that the terminal measure $\mathbb{Q}_{X_T^{\hat{\alpha}}}$ is close to μ_{tar} , when k is sufficiently large. $\square$

Example 2.6 (Fine-tuning under a reward signal). Fine-tuning a diffusion model means taking a pre-trained model and training it further on a smaller, task-specific dataset so it learns to generate

outputs more suited to that new context [34, 60, 62, 70]. For example, a diffusion model trained on general images can be fine-tuned to generate a specific style (evaluated via a reward function). This process updates the model’s parameters just enough to adapt to the new data, without starting training from scratch.

In terms of our framework, we can consider (2.3) as a pre-trained model with the drift $b(t, x) := s_{\hat{\theta}}(t, x)$ being a well-trained score function, and $\hat{\theta}$ is the trained parameter. Note that, as the result of pre-training, the output measure $\mathbb{Q}_{X_T}$ is sufficiently close to some data distribution μ_{tar} . We then introduce a fine-tuning procedure through a reference measure p_{ref} with density $\frac{\exp(r(x))}{\int_{\mathbb{R}^d} \exp(r(x))dx}$, where $r : \mathbb{R}^d \rightarrow \mathbb{R}$ is a given reward function satisfying $\int_{\mathbb{R}^d} \exp(r(x))dx < \infty$. Now replacing $\mu_{\text{tar}} = p_{\text{ref}}$, the optimal control $\hat{\alpha}$ of the corresponding SCSBP can then serve as the fine-tuning score function; and consequently, the new drift term $b(t, X_t^{\hat{\alpha}}) + \sigma(t)\hat{\alpha}$ acts as a combined score function.

Clearly, in this application the penalty parameter k should not be chosen too large; otherwise, the effect of the preference function may dominate the fidelity to the original data distribution. With an appropriately selected k , the resulting measure $\mathbb{Q}_{X_T^{\hat{\alpha}}}$ not only reflects p_{data} but also integrates the reward function r . In contrast, we remark that the classic SBP (2.1) is not capable of handling this application as it has a *pre-fixed* target distribution. $\square$

Example 2.7 (Transfer learning). Transfer learning is a machine learning approach where knowledge gained from a “source task” is reused to improve learning in a related but different “target task” [9, 48, 61]. In what follows we shall consider transfer learning in the context of data generation.

Let us consider a source task (Y_{sou}) , characterized by a distribution p_{sou} , and a target task (Y_{tar}) with distribution p_{tar} . Typically, p_{sou} and p_{tar} are assumed to be close under a certain divergence or distance function $G(p_{\text{tar}}; p_{\text{sou}})$ (assuming $G \geq 0$), such as the Wasserstein distance [48].

To fit the transfer learning into our framework, we can take $\mu_{\text{ini}} = p_{\text{sou}}$, $\mu_{\text{tar}} = p_{\text{tar}}$, and set $b \equiv 0$ for simplicity. In this case, if we choose $\alpha = 0$, and $X_0 \sim p_{\text{sou}}$, then $X_T = X_0 + W_T \sim p_{\text{sou}} * \mathcal{N}(0, T\mathbb{I}_d)$, where $\mathcal{N}(0, T\mathbb{I}_d) = \mathbb{P}_{W_T}^0$ and $\mathbb{I}_d$ denotes the $d \times d$ identity matrix. Thus, denoting the optimal control by $\hat{\alpha}$ and noting that $\alpha \equiv 0$ is sub-optimal, we must have

$$\mathbb{E}^{\mathbb{Q}} \left[\frac{1}{2} \int_0^T |\hat{\alpha}_s|^2 ds \right] \leq \mathbb{E}^{\mathbb{Q}} \left[\frac{1}{2} \int_0^T |\hat{\alpha}_s|^2 ds + kG(\mathbb{Q}_{X_T^{\hat{\alpha}}}) \right] \leq kG(p_{\text{sou}} * \mathcal{N}(0, T\mathbb{I}_d); p_{\text{tar}}).$$

This implies that the optimal control $\hat{\alpha}$ has a small L^2 -norm, indicating only minor adjustments are required during sampling—provided k is not too large. $\square$

3 Existence of Optimal Policies for SCSBP’s

In this section we study the stochastic control problem (2.9)-(2.10) and the associated soft-constrained SBP. In particular, we shall prove that the optimal control for each $k \in \mathbb{N}$ exists and in next section we will show that these optimal policies will converge to the solution of the original SBP, with a linear rate of convergence. We shall assume that the target distribution for the SCSBP has density $f_{\text{tar}} \in \mathbb{L}^1(\mathbb{R}^d)$. Also, we assume $\sigma(\cdot) = I_d$, that is, we consider the following SDE:

$$dX_t = b(t, X_t)dt + dW_t, \quad \mathbb{P}_{X_0} = \mu_{\text{ini}}. \quad (3.1)$$

In what follows we shall often consider the SDE with initial data $(t, x) \in [0, T] \times \mathbb{R}^d$:

$$X_s = x + \int_t^s b(r, X_r)dr + W_s - W_t, \quad s \in [t, T]. \quad (3.2)$$

We denote the solution of (3.2) by $X^{t,x}$. Clearly, if we let $p(\cdot, \cdot; \cdot, \cdot)$ be the transition density of the solution X to (3.1), that is, $\mathbb{P}\{X_s \in dz | X_t = x\} = p(s, z; t, x)dz$, $0 \leq t < s \leq T$, $z, x \in \mathbb{R}^d$, then $p(s, \cdot; t, x)$ is the density of $X_s^{t,x}$. It is well known that $p(\cdot, \cdot; \cdot, \cdot)$ is the fundamental solution to Kolmogorov backward (parabolic) PDE, and under mild conditions (see, e.g., [3]), there exist $c_1, c_2, \lambda, \Lambda > 0$, it holds that

$$c_1(s-t)^{-\frac{d}{2}} e^{-\frac{\lambda|z-x|^2}{s-t}} < p(s, z; t, x) < c_2(s-t)^{-\frac{d}{2}} e^{-\frac{\Lambda|z-x|^2}{4(s-t)}}. \quad (3.3)$$

Keeping the original SBP (2.5) associated with (3.1) in mind, let us now recall the Problem 2.4 and the cost functional $J^k(\alpha)$ defined by (2.10). For notational simplicity, we let $\mathbb{P}_{X^\alpha}$ (resp. $\mathbb{P}_X$) denote the probability measure on $\mathbb{C}_T^d$ induced by X^α (resp. X). Clearly, when $\alpha \equiv 0$, we have $\mathbb{P}_{X^0} = \mathbb{P}_X$, where X solves (3.1). Furthermore, we shall denote $\mathbb{E}[\cdot] = \mathbb{E}^\mathbb{P}[\cdot]$ when context is clear, and for each $k \in \mathbb{N}$ we can easily check that

$$J^k(\alpha) = \mathbb{E}\left[\frac{1}{2} \int_0^T |\alpha_s|^2 ds + kG(\mathbb{P}_{X_T^\alpha})\right] = D_{\text{KL}}(\mathbb{P}_{X^\alpha} \|\mathbb{P}_X) + kG(\mathbb{P}_{X_T^\alpha}). \quad (3.4)$$

Now let us define, for each $k \in \mathbb{N}$, a mapping $D_k(\cdot) : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ by

$$D_k(\mu) = D_{\text{KL}}(\mu \|\mathbb{P}_{X_T}) + kG(\mu),$$

and note that $D_{\text{KL}}(\mathbb{P}_{X^\alpha} \|\mathbb{P}_X) \geq D_{\text{KL}}(\mathbb{P}_{X_T^\alpha} \|\mathbb{P}_{X_T})$, we deduce from (3.4) that

$$J^k(\alpha) \geq D_{\text{KL}}(\mathbb{P}_{X_T^\alpha} \|\mathbb{P}_{X_T}) + kG(\mathbb{P}_{X_T^\alpha}) = D_k(\mathbb{P}_{X_T^\alpha}). \quad (3.5)$$

If $\hat{\alpha}$ is the optimal control corresponding to the original SBP, that is, $\mathbb{P}_{X_T^{\hat{\alpha}}} = \mu_{\text{tar}}$, then by definition of the penalty function $G(\cdot)$, we should have $G(\mu_{\text{tar}}) = 0$, and therefore,

$$D_k(\mathbb{P}_{X_T^{\hat{\alpha}}}) = D_k(\mu_{\text{tar}}) = D_{\text{KL}}(\mu_{\text{tar}} \|\mathbb{P}_{X_T}) + kG(\mu_{\text{tar}}) = D_{\text{KL}}(\mu_{\text{tar}} \|\mathbb{P}_{X_T}).$$

Throughout the rest of this section, we shall focus on the special case: $\mu_{\text{ini}} = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^d$. The case with general initial distribution μ_{ini} will be studied in Sections 5 and 6. We begin with the following well-known result from [18], which will play an important role in our discussion.

Lemma 3.1 ([18, Theorem 3.1]). *Let X be a weak solution to (3.1) with $X_0 = x_0 \in \mathbb{R}^d$ (i.e., $\mu_{\text{ini}} = \delta_{x_0}$). Assume that $D_{\text{KL}}(\mu_{\text{tar}} \|\mathbb{P}_{X_T}) < \infty$. Then, the optimal solution to the SBP (2.4)-(2.5) is given by $\hat{\alpha}_t = \nabla \log h(t, X_t^{\hat{\alpha}})$, where*

$$h(t, x) = \int_{\mathbb{R}^d} p(T, z; t, x) \frac{f_{\text{tar}}(z)}{p(T, z; 0, x_0)} dz = \mathbb{E}\left[\frac{f_{\text{tar}}(X_T)}{p(T, X_T; 0, x_0)} \middle| X_t = x\right], \quad (3.6)$$

for $(t, x) \in [0, T] \times \mathbb{R}^d$. □

Next, we make the following assumptions on the penalty function G :

Assumption 3.2. (i) *There exists some small constant $\varepsilon > 0$ such that*

$$G(\mu) \rightarrow +\infty, \quad \text{as } \|\mu\|^{2+\varepsilon} \rightarrow +\infty. \quad (3.7)$$

where $\|\mu\|^p := \int_{\mathbb{R}^d} |x|^p \mu(dx)$ for any $p > 0$.

(ii) *There exist $C, \lambda > 0$, and a function $\phi : \mathbb{R}^d \rightarrow (0, 1]$ satisfying $\phi(x)e^{\lambda|x-x_0|^2} \leq C$, such that for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ with density function f_μ , it holds that*

$$|f_\mu(x) - f_{\text{tar}}(x)| \leq C\phi(x)G(\mu), \quad x \in \mathbb{R}^d. \quad (3.8)$$

Remark 3.3. (i) The function $G(\mu)$ on the right hand side of (3.8) should read $|G(\mu) - G(\mu_{\text{tar}})|$, as $G(\mu_{\text{tar}}) = 0$, which essentially states that if μ is close to μ_{tar} in terms of G , then f_μ is close to f_{tar} in $\mathbb{L}^1$.

(ii) A slightly stronger consequence of (3.8) is the following. Recall the (generalized) Kantorovich and Rubinstein dual representation (cf. e.g., [23]): denoting $\text{Lip}(1)$ to be all Lipschitz functions $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ with Lipschitz constant $\text{Lip}_\varphi \leq 1$ (hence $|\varphi(x)| \leq C(1 + |x|)$), then it holds that

$$W_1(\mu, \mu_{\text{tar}}) = \sup_{\varphi \in \text{Lip}(1)} \left\{ \int_{\mathbb{R}^d} \varphi(x)(f_\mu(x) - f_{\text{tar}}(x))dx \right\} \leq CG(\mu) \int_{\mathbb{R}^d} (1 + |x|)\phi(x)dx.$$

This suggests that $G(\mu) \sim 0$ implies that μ is close to μ_{tar} in the sense of W_1 . $\square$

Before we proceed, let us give two examples that justify Assumption 3.2.

Example 3.4. We consider the class of μ and μ_{tar} such that

$$G(\mu) := \int_{\mathbb{R}} |x|^p |f_\mu(x) - f_{\text{tar}}(x)|dx,$$

is well-defined for a given p . Clearly, Definition 2.3-(i) and Assumption 3.2-(i) are satisfied when $p > 2$. Assumption 3.2-(i) holds when

$$\phi(x) \leq \frac{|f_\mu(x) - f_{\text{tar}}(x)|}{c \int_{\mathbb{R}} |x|^p |f_\mu(x) - f_{\text{tar}}(x)|dx},$$

for all μ in the collection one may consider. $\square$

Another natural example of G satisfying Assumption 3.2 would be the Wasserstein distance or the KL divergence, augmented with a small “guardrail” term that enforces the uniform (weighted) pointwise control in (3.8). This guardrail can be taken as a weighted $\mathbb{L}_\infty$ norm, a Hölder $\mathbb{C}^\alpha$ seminorm, or an RKHS norm (e.g., with kernel $k(x, y) = \phi(x)\phi(y)\kappa(x - y)$). In next example we illustrate such a choice with the Wasserstein distance plus an $\mathbb{L}_\infty$ guardrail.

Example 3.5. Consider the case that $\mu_{\text{tar}} \in \mathcal{P}_p(\mathbb{R}^d)$ with $p > 2$. We define, for $c > 0$ and $\phi(x) = \exp(-\lambda|x - x_0|^2)$ with some $x_0 \in \mathbb{R}^d$,

$$G(\mu) := W_2(\mu, \mu_{\text{tar}}) + c \left\| \frac{f_\mu - f_{\text{tar}}}{\phi} \right\|_{\mathbb{L}^\infty}.$$

Then, it is easy to check that

$$|f_\mu(x) - f_{\text{tar}}(x)| \leq \left\| \frac{f_\mu - f_{\text{tar}}}{\phi} \right\|_{\mathbb{L}^\infty} \phi(x) \leq \frac{1}{c} \phi(x) G(\mu).$$

Thus (3.8) holds and $\phi(x)e^{\lambda|x-x_0|^2} \leq C$ holds with $C = \max\{1, \frac{1}{c}\}$.

Let $\{\mu_n\}_{n \geq 1} \subset \mathcal{P}_2(\mathbb{R}^d)$ with $\|\mu_n\|^{2+\varepsilon} \rightarrow \infty$. We claim that $G(\mu_n)$ must be unbounded. Indeed, suppose not. Then there exists $M, M' > 0$ such that $W_2^2(\mu_n, \mu_{\text{tar}}) \leq M$ and $\|\phi^{-1}(f_{\mu_n} - f_{\text{tar}})\|_{\mathbb{L}^\infty} < M'$ for all $n \in \mathbb{N}_+$. Hence $f_{\mu_n}(x) \leq f_{\text{tar}}(x) + M'\phi(x)$, $x \in \mathbb{R}^d$. Integrating against $|x|^{2+\varepsilon}$ and using the facts that $\mu_{\text{tar}} \in \mathcal{P}_{2+\varepsilon}(\mathbb{R}^d)$ with $\varepsilon = p - 2 > 0$ and $\int |x|^{2+\varepsilon} \phi(x)dx < \infty$, we have

$$\|\mu_n\|^{2+\varepsilon} \leq \|\mu_{\text{tar}}\|^{2+\varepsilon} + M' \int |x|^{2+\varepsilon} \phi(x)dx < \infty, \quad n \in \mathbb{N}.$$

This contradicts the fact that $\|\mu_n\|^{2+\varepsilon} \rightarrow \infty$, proving the claim. Hence (3.7) holds. $\square$

We are now ready to investigate the existence of optimal control of Problem 2.4 for each $k \in \mathbb{N}$, which would be essential for our approximation scheme. Recall that in the rest of the section we assume that $\mu_{\text{ini}} = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^d$. To begin with, we first claim that for each $k \in \mathbb{N}$, there exists $\hat{\mu}_k \in \mathcal{P}_2(\mathbb{R}^d)$ such that the static optimization problem on the measure space has a solution

$$D_k(\hat{\mu}_k) = \inf_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} D_k(\mu). \quad (3.9)$$

Indeed, let X be the solution to uncontrolled SDE (3.1), and μ_{tar} be given such that $D_{\text{KL}}(\mu_{\text{tar}} \|\mathbb{P}_{X_T}) < \infty$. Since $G(\mu_{\text{tar}}) = 0$, we have

$$m := D_k(\mu_{\text{tar}}) = D_{\text{KL}}(\mu_{\text{tar}} \|\mathbb{P}_{X_T}) + kG(\mu_{\text{tar}}) = D_{\text{KL}}(\mu_{\text{tar}} \|\mathbb{P}_{X_T}) < \infty.$$

Next, let us define, for fixed $k \in \mathbb{N}$, a set

$$\mathcal{S}_k := \left\{ \mu \in \mathcal{P}_2(\mathbb{R}^d) : D_k(\mu) \leq m \right\}.$$

Clearly, $\mathcal{S}_k \neq \emptyset$ since $\mu_{\text{tar}} \in \mathcal{S}_k$, and by (3.7), there exists $M_k > 0$ such that $\|\mu\|^{2+\varepsilon} \leq M_k$, for all $\mu \in \mathcal{S}_k$. Thus $\mathcal{S}_k$ is uniformly integrable in $\mathbb{L}^2$. Now let $\{\mu_k^{(i)}\}_{i=1}^\infty \subset \mathcal{P}_2(\mathbb{R}^d)$ be a minimizing sequence, namely,

$$\lim_{i \rightarrow \infty} D_k(\mu_k^{(i)}) = \inf_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} D_k(\mu).$$

Since $\inf_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} D_k(\mu) \leq D_k(\mu_{\text{tar}}) = m$, we may assume without loss of generality that $\{\mu_k^{(i)}\}_{i=1}^\infty \subset \mathcal{S}_k$. Since $\mathcal{S}_k$ is uniformly integrable and is tight in $\mathcal{P}_2(\mathbb{R}^d)$, there exists subsequence $\{\mu_k^{(i_l)}\}_{l=1}^\infty$ such that $\hat{\mu}_k := \lim_{l \rightarrow \infty} \mu_k^{(i_l)} \in \mathcal{P}_2(\mathbb{R}^d)$.² Since the mapping $\mu \mapsto D_k(\mu)$ is continuous, we have

$$D_k(\hat{\mu}_k) = D_{\text{KL}}(\hat{\mu}_k \|\mathbb{P}_{X_T}) + kG(\hat{\mu}_k) = \inf_{\mu \in \mathcal{P}_2(\mathbb{R}^d)} D_k(\mu),$$

proving the claim. Furthermore, if we denote the density of $\mathbb{P}_{X_T}$ by f_{X_T} , and note that $D_{\text{KL}}(\hat{\mu}_k \|\mathbb{P}_{X_T}) \leq D_k(\hat{\mu}_k) \leq m < \infty$, we know that $\frac{d\hat{\mu}_k}{d\mathbb{P}_{X_T}}$ exists and

$$\frac{d\hat{\mu}_k}{dx}(x) = \frac{d\hat{\mu}_k}{d\mathbb{P}_{X_T}} \cdot f_{X_T}(x) =: f_k(x).$$

Keeping the above discussion in mind, we are now ready to prove the following theorem.

Proposition 3.6. *Assume that Assumption 3.2 is in force, and that $\mu_{\text{ini}} = \delta_{x_0}$, $x_0 \in \mathbb{R}^d$. Then, for each $k \in \mathbb{N}$, the optimal control for Problem 2.4, denoted by $\hat{\alpha}^k$, exists. Furthermore, $\hat{\alpha}^k$ has the following explicit feedback form: $\hat{\alpha}_t^k := \nabla \log h^k(t, X_t^{\hat{\alpha}^k})$, where*

$$h^k(t, x) = \int_{\mathbb{R}^d} \frac{f_k(z)}{p(T, z; 0, x_0)} p(T, z; t, x) dz = \mathbb{E} \left[\frac{f_k(X_T)}{p(T, X_T; 0, x_0)} \middle| X_t = x \right]. \quad (3.10)$$

Proof. Let $k \in \mathbb{N}$ be fixed, and let $\hat{\mu}_k$ be the minimizer of $D_k(\cdot)$ defined by (3.9). Then, by (3.5), for any $\alpha \in \mathbb{L}_{\mathbb{P}^0}^2([0, T]; \mathbb{R}^d)$, we have

$$J^k(\alpha) \geq D_k(\mathbb{P}_{X_T}^\alpha) \geq D_k(\hat{\mu}_k),$$

²This follows from the result on Wasserstein distance vs. weak convergence (see, e.g., [65, Theorem 7.12]), which states that $W_p(\mu_k, \mu) \rightarrow 0$ if and only if $\mu_k \Rightarrow \mu$ weakly, and $\lim_{R \rightarrow \infty} \overline{\lim}_{k \rightarrow \infty} \int_{\{|x - x_0| \geq R\}} |x - x_0|^p \mu_k(dx) = 0$.

Therefore, in order to find the optimal control for Problem 2.4, it suffices to find $\hat{\alpha}^k$ such that (i) $X_T^{\hat{\alpha}^k} \sim \hat{\mu}_k$; and (ii) $J^k(\hat{\alpha}^k) = D_k(\mathbb{P}_{X_T^{\hat{\alpha}^k}})$.

To this end, we first recall that $\hat{\mu}_k$ is the minimizer of the function $D_k(\cdot)$ with density f_k . Next, we apply Lemma 3.1 with μ_{tar} being replaced by $\hat{\mu}_k$ to get the optimal control $\hat{\alpha}^k \in \mathcal{P}(\mu_{\text{ini}}, \hat{\mu}_k)$ for the original SBP (2.5)-(2.6), which satisfies $\hat{\alpha}_t^k = \nabla \log h^k(t, X_t^{\hat{\alpha}^k})$, where h^k is defined by (3.10), and $X_T^{\hat{\alpha}^k} \sim \hat{\mu}_k$. Now, note that for this SBP we have

$$V(\mu_{\text{ini}}, \hat{\mu}_k) = \frac{1}{2} \mathbb{E} \left[\int_0^T |\hat{\alpha}_s^k|^2 ds \right] = D_{\text{KL}}(\hat{\mu}_k \| \mathbb{P}_{X_T}),$$

we conclude that

$$J^k(\hat{\alpha}^k) = D_{\text{KL}}(\hat{\mu}_k \| \mathbb{P}_{X_T}) + kG(\hat{\mu}_k) = D_k(\hat{\mu}_k) = D_k(\mathbb{P}_{X_T^{\hat{\alpha}^k}}).$$

In other words, $\hat{\alpha}^k$ is indeed the optimal control for the Problem 2.4, proving the proposition. $\square$

4 Convergence Results under Delta Initial Distribution

To show the convergence of both policies and value functions, we first make the following observations. First, if we denote $g(x) := \frac{f_{\text{tar}}(x)}{p(T, x; 0, x_0)}$ and $\mathbb{E}_{t,x}[\cdot] := \mathbb{E}[\cdot | X_t = x]$, then by (3.6) we can write $h(t, x) = \mathbb{E}_{t,x}[g(X_T)]$, where X is the solution to (3.1) with $X_0 = x_0$. By Feynman-Kac formula, we see that h satisfy the PDE:

$$\begin{cases} \partial_t h(t, x) + \mathcal{L}_t h(t, x) = 0; \\ h(T, x) = g(x) = \frac{f_{\text{tar}}(x)}{p(T, x; 0, x_0)}, \end{cases} \quad (4.1)$$

where the infinitesimal generator $\mathcal{L}_t$ is defined by $\mathcal{L}_t := b(t, x) \cdot \nabla + \frac{1}{2} \Delta$. Similarly, we define $g_k(x) := \frac{f_k(x)}{p(T, x; 0, x_0)}$, then the function $h^k(t, x)$ can also be represented as the solution of the PDE:

$$\begin{cases} \partial_t h^k(t, x) + \mathcal{L}_t h^k(t, x) = 0; \\ h^k(T, x) = g_k(x) = \frac{f_k(x)}{p(T, x; 0, x_0)}. \end{cases} \quad (4.2)$$

Another useful observation is that for $\hat{\mu}_k \in \mathcal{S}_k$, we have $kG(\hat{\mu}_k) \leq D_k(\hat{\mu}_k) \leq m$, or

$$G(\hat{\mu}_k) \leq m/k. \quad (4.3)$$

With the above observations, the key is to establish the convergence rate of $|h(t, x) - h^k(t, x)|$ and $|\nabla h(t, x) - \nabla h^k(t, x)|$ as k tends to infinity.

4.1 The Convergence of Optimal Policies

We shall now argue that the optimal controls for Problem 2.4, $\{\hat{\alpha}^k(\cdot, \cdot)\}$, given by Proposition 3.6, actually converges to the solution of the original SBP $\hat{\alpha}(\cdot, \cdot)$ given by Lemma 3.1, and also establish its rate of convergence. More precisely, we have the following theorem.

Theorem 4.1. *Assume that the Assumption 3.2 is in force, and that $\mu_{\text{ini}} = \delta_{x_0}$ for some $x_0 \in \mathbb{R}^d$. Furthermore, assume that there exists constants $C, \delta > 0$, such that $\delta \leq g(x), g_k(x) \leq C$, $x \in \mathbb{R}^d$,*

$k \in \mathbb{N}$. Let $\hat{\alpha}(t, x)$ and $\hat{\alpha}^k(t, x)$, $(t, x) \in [0, T] \times \mathbb{R}^d$ be the optimal controls given in Lemma 3.1 and Proposition 3.6, respectively. Then, it holds that

$$\int_0^T |\hat{\alpha}^k(t, x) - \hat{\alpha}(t, x)| dt \leq \frac{C}{k}, \quad x \in \mathbb{R}^d,$$

where $C > 0$ is some constant independent of k .

Remark 4.2. We note that the assumption $\delta \leq g(x) = \frac{f_{\text{tar}}(x)}{P(T, x; 0, x_0)} \leq C$ (resp. $\delta \leq g_k(x) \leq C$) amounts to saying that $f_{\text{tar}}(x)$ (resp. $f_k(x)$) $\propto P(T, x; 0, x_0)$ as $x \rightarrow \infty$, which is not particularly a stringent condition in light of the general estimate (3.3), and the arbitrariness of the sample data selection for the data generation procedure. $\square$

Proof. First, by definition $\hat{\alpha}(t, x) = \nabla \log h(t, x)$ and $\hat{\alpha}^k(t, x) = \nabla \log h^k(t, x)$, where h^k and h are the solution to (4.1) and (4.2), respectively, and $\nabla = \partial_x$. We can easily deduce that

$$\begin{aligned} |\hat{\alpha}^k(t, x) - \hat{\alpha}(t, x)| &= |\nabla \log h^k(t, x) - \nabla \log h(t, x)| = \left| \frac{\nabla h^k(t, x)}{h^k(t, x)} - \frac{\nabla h(t, x)}{h(t, x)} \right| \\ &= \left| \frac{\nabla h^k(t, x)h(t, x) - \nabla h(t, x)h(t, x) + \nabla h(t, x)h(t, x) - \nabla h(t, x)h^k(t, x)}{h^k(t, x)h(t, x)} \right| \\ &\leq \left| \frac{\nabla h^k(t, x) - \nabla h(t, x)}{h^k(t, x)} \right| + |\nabla h(t, x)| \left| \frac{h(t, x) - h^k(t, x)}{h^k(t, x)h(t, x)} \right| =: I_1 + I_2. \end{aligned} \quad (4.4)$$

We now estimate I_1 and I_2 , respectively. To this end we first apply the well-known Bismut-Elworthy-Li formula [8, 24] (see also the representation formula in [27, 43]) to get

$$\begin{cases} \nabla h(t, x) = \partial_x \mathbb{E}_{t, x}[g(X_T)] = \mathbb{E}_{t, x}[g(X_T)N_T], \\ \nabla h^k(t, x) = \partial_x \mathbb{E}_{t, x}[g_k(X_T)] = \mathbb{E}_{t, x}[g_k(X_T)N_T], \end{cases} \quad (4.5)$$

where,

$$N_s = N_s^{t, x} := \frac{1}{s - t} \int_t^s (\nabla X_r^{t, x})^\top dW_r, \quad s \in [t, T],$$

and $\nabla X = \nabla X^{t, x}$ is a $\mathbb{R}^{d \times d}$ -valued variational process satisfying the (random) ODE:

$$\partial_{x^j} X_s^i = \delta_{ij} + \int_t^s \sum_{\ell=1}^d \partial_{x^\ell} b^i(r, X_r) \partial_{x^j} X_r^\ell dr, \quad 1 \leq i, j \leq d, \quad s \in [t, T].$$

Furthermore, one can easily check that

$$\mathbb{E}[|\nabla X_s^{t, x}|^2] \leq C e^{C(s-t)}, \quad \mathbb{E}[|N_s^{t, x}|^2] \leq \frac{C}{s-t} e^{C(s-t)}, \quad 0 \leq t \leq s \leq T.$$

Therefore, denoting $C > 0$ to be a generic constant that is allowed to vary from line to line, and applying Assumption 3.8 and estimate (3.3) we have

$$\begin{aligned} |\nabla h^k(t, x) - \nabla h(t, x)| &\leq \mathbb{E}[|g(X_T^{t, x}) - g_k(X_T^{t, x})| |N_T^{t, x}|] = \mathbb{E}\left[\left| \frac{f_k(X_T^{t, x}) - f_{\text{tar}}(X_T^{t, x})}{p(T, X_T^{t, x}; 0, x_0)} \right| |N_T^{t, x}| \right] \\ &\leq \left(\mathbb{E}[|N_T^{t, x}|^2] \right)^{\frac{1}{2}} \left[\mathbb{E}\left[\left| \frac{f_k(X_T^{t, x}) - f_{\text{tar}}(X_T^{t, x})}{p(T, X_T^{t, x}; 0, x_0)} \right|^2 \right] \right]^{\frac{1}{2}} \\ &\leq \frac{C e^{C(T-t)}}{\sqrt{T-t}} G(\hat{\mu}_k) \left[\mathbb{E}\left[\left| \frac{\phi(X_T^{t, x})}{p(T, X_T^{t, x}; 0, x_0)} \right|^2 \right] \right]^{\frac{1}{2}} \leq \frac{C}{k \sqrt{T-t}}. \end{aligned} \quad (4.6)$$

Next, we note that by assumption $\delta \leq g(x), g_k(x) \leq C$ for all $x \in \mathbb{R}^d$ and $k \in \mathbb{N}$, by the weak maximum principle we conclude that as the solutions to the PDEs (4.1) and (4.2), respectively, it holds that $\delta \leq h(t, x), h^k(t, x) \leq C$, for all $(t, x) \in \mathbb{R}^d \times [0, T)$. Consequently, we have

$$I_1 \leq \frac{C}{\delta k \sqrt{T-t}} \leq \frac{C}{k \sqrt{T-t}}. \quad (4.7)$$

Similarly, we can argue that $|\nabla h(t, x)| \leq \frac{C}{\sqrt{T-t}}$, and that

$$|h^k(t, x) - h(t, x)| \leq \mathbb{E} \left[\left| \frac{f_k(X_T^{t,x}) - f_{\text{tar}}(X_T^{t,x})}{p(T, X_T^{t,x}; 0, x_0)} \right| \right] \leq CG(\hat{\mu}_k) \mathbb{E} \left[\left| \frac{\phi(X_T^{t,x})}{p(T, X_T^{t,x}; 0, x_0)} \right| \right] \leq \frac{C}{k}, \quad (4.8)$$

where the last inequality is due to (4.3) and hence $I_2 \leq \frac{C}{k \sqrt{T-t}}$. This, together with (4.7) and (4.4), we obtain

$$|\hat{\alpha}^k(t, x) - \hat{\alpha}(t, x)| \leq \frac{C}{k \sqrt{T-t}} \quad (4.9)$$

and hence convergence result :

$$\int_0^T |\hat{\alpha}^k(t, x) - \hat{\alpha}(t, x)| dt \leq \int_0^T \frac{C}{k \sqrt{T-t}} dt \leq \frac{C \sqrt{T}}{k},$$

proving the theorem. $\square$

Remark 4.3. A particular example is when we take the penalty function $G(\mu) = D_{\text{KL}}(\mu \| \mu_{\text{tar}})$. In this case, it is known (see, e.g., [29, Theorem 2]) that the optimal control for (2.9)-(2.10) is given by $\hat{\alpha}_t^k = \nabla \log h^k(X_t^k, t)$, where

$$h^k(t, x) = d_k^{-1} \int p(T, z; t, x) \left(\frac{f_{\text{tar}}(z)}{p(T, z; 0, x_0)} \right)^{\frac{k}{k+1}} dz,$$

with $d_k = \int f_{\text{tar}}(x)^{\frac{k}{1+k}} p(x, T; x_0, 0)^{\frac{1}{1+k}} dx$. Consequently, Assumption 3.2-(ii) can be reduced to that $\mathbb{E} \left[\left| \frac{f_{\text{tar}}(X_T)}{p(T, X_T; 0, x_0)} \right|^2 \right]$ is bounded (see Assumption 5.4 for similar conditions); and the linear rate of convergence can be proved with the same arguments. $\square$

4.2 The Convergence of the Value Function

Having worked out the convergence analysis for the optimal controls, it is natural to extend the results to the convergence of value functions. However, the singularity at the terminal time T in (4.9) requires some technical care. It turns out that the popular notion of *early stopping* in diffusion models as well as the flow-based method literature [4, 33, 41] is exactly the remedy to this issue.

To be more precise, for any $\varepsilon > 0$, we introduce the following ε -value function.

$$J_\varepsilon(\alpha) := \mathbb{E} \left[\int_0^{T-\varepsilon} \frac{1}{2} |\alpha|^2 dt \right].$$

There are many practical reasons, mainly for computational purposes, to invoke the notion of early stopping, as elaborated in [4, 33, 41]. But on the other hand, it is clear that the ε -value function effectively excludes the singularity at the terminal time T . This leads to the following straightforward result.

Proposition 4.4. *Assume that all the assumptions of Theorem 4.1 are in force. Then, for any $\varepsilon > 0$, there exists a generic constant $C := C(\varepsilon) = \mathcal{O}(\frac{1}{\sqrt{\varepsilon}}) > 0$, independent of k , such that*

$$|J_\varepsilon(\hat{\alpha}^k) - J_\varepsilon(\hat{\alpha})| \leq \frac{C}{k}, \quad k \in \mathbb{N}. \quad (4.10)$$

where $\hat{\alpha}^k$ and $\hat{\alpha}$ are the optimal controls in Theorem 4.1, respectively.

Proof. The proof is straightforward. For any $k \in \mathbb{N}$, let $\hat{\alpha}^k$ and $\hat{\alpha}$ be the optimal controls in Theorem 4.1, respectively. Then, for any $\varepsilon > 0$, applying (4.9) we have

$$\begin{aligned} |J_\varepsilon(\hat{\alpha}^k) - J_\varepsilon(\hat{\alpha})| &\leq \mathbb{E} \left[\frac{1}{2} \int_0^{T-\varepsilon} ||\hat{\alpha}_s^k|^2 - |\hat{\alpha}_s|^2| ds \right] \leq \mathbb{E} \left[\frac{1}{2} \int_0^{T-\varepsilon} |\hat{\alpha}_s^k - \hat{\alpha}_s| (|\hat{\alpha}_s^k| + |\hat{\alpha}_s|) ds \right] \\ &\leq \frac{C}{k} \mathbb{E} \left[\frac{1}{2} \int_0^{T-\varepsilon} \frac{1}{\sqrt{T-s}} (|\hat{\alpha}_s^k| + |\hat{\alpha}_s|) ds \right] \\ &\leq \frac{C}{k\sqrt{\varepsilon}} \mathbb{E} \left[\frac{1}{2} \int_0^T (|\hat{\alpha}_s^k| + |\hat{\alpha}_s|) ds \right], \end{aligned} \quad (4.11)$$

where the last inequality is due to the fact that $\frac{1}{\sqrt{T-s}} \leq \frac{1}{\sqrt{\varepsilon}}$ for $s \in [0, T-\varepsilon]$. To further bound (4.11), we recall that the definitions of $J(\cdot)$ (2.5) and $J^k(\cdot)$ (2.10), $k \in \mathbb{N}$, and define

$$V^* = J(\hat{\alpha}) = \inf_{\alpha \in \mathcal{A}} J(\alpha); \quad V^{k,*} = J^k(\hat{\alpha}^k) = \inf_{\alpha \in \mathcal{A}} J^k(\alpha).$$

We should note that $X^{\hat{\alpha}}$ follows the constrained dynamics (2.6), whereas $X^{\hat{\alpha}^k}$ follows the soft-constrained dynamics (2.9). Clearly, by definition (2.10) we have

$$\sup_{k \geq 1} J^k(\alpha) = \begin{cases} \mathbb{E} \left[\int_0^T \frac{1}{2} |\alpha_t|^2 dt \right] & \text{if } \mathbb{P}_{X_T^\alpha} = \mu_{\text{tar}} \\ \infty & \text{otherwise.} \end{cases}$$

Thus, since $\hat{\alpha}$ satisfies the constrained dynamics (2.6), we have

$$V^* = J(\hat{\alpha}) = \inf_{\alpha \in \mathcal{A}} \sup_{k \geq 1} J^k(\alpha) \geq \inf_{\alpha \in \mathcal{A}} J^k(\alpha) = J^k(\hat{\alpha}^k) = V^{k,*}, \quad k \in \mathbb{N}. \quad (4.12)$$

Consequently, we have, for each $k \in \mathbb{N}$, a simple application of Cauchy–Schwarz inequality and the fact (4.12) yields

$$\mathbb{E} \left[\frac{1}{2} \int_0^T (|\hat{\alpha}_s^k| + |\hat{\alpha}_s|) ds \right] \leq \frac{\sqrt{T}}{2} \left(\mathbb{E} \left[\int_0^T |\hat{\alpha}_s^k|^2 ds \right] \right)^{1/2} + \left(\mathbb{E} \left[\int_0^T |\hat{\alpha}_s|^2 ds \right] \right)^{1/2} \leq \sqrt{T} \sqrt{V^*}. \quad (4.13)$$

Combining (4.11) and (4.13), we obtain (4.10). $\square$

Besides the convergence of the value functions, another important convergence, that relies crucially on the convergence of the optimal controls, is the convergence of the terminal law $\mathbb{P}_{X_T^{\hat{\alpha}^k}}$ (with respect to the target distribution μ_{tar}), measured, for instance, in the Wasserstein distance. Again, to avoid the technicalities that the singularity at terminal time T might cause, we shall focus on the early stopped state $X_{T-\varepsilon}^{\hat{\alpha}^k}$, which is a commonly used criterion in statistical estimation results for generative diffusion models (see, e.g., [12, 28, 33]). More precisely, we have the following result.

Proposition 4.5. *Let all assumptions in Theorem 4.1 be in force. Assume further that the optimal policy $\hat{\alpha}$ of the original SBP is Lipschitz in x : there exists $\kappa > 0$, such that*

$$|\hat{\alpha}(t, x) - \hat{\alpha}(t, y)| \leq \kappa|x - y|, \quad t \in [0, T]. \quad (4.14)$$

Then there exists a constant $C > 0$, depending on the Lipschitz constants L in Assumption 2.2 and κ in (4.14), but independent of $k \in \mathbb{N}$, such that for any $\varepsilon > 0$, it holds that

$$W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}^k}}, \mu_{\text{tar}}) \leq \frac{C\sqrt{\ln T - \ln \varepsilon}}{k} + C\varepsilon. \quad (4.15)$$

In particular, if we choose $\varepsilon = \frac{1}{k}$, then it holds that

$$W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}^k}}, \mu_{\text{tar}}) \leq \frac{C}{k}(\sqrt{\ln k} + \sqrt{\ln T}) + 1 = \mathcal{O}\left(\frac{\sqrt{\ln k}}{k}\right). \quad (4.16)$$

Remark 4.6. (i) The linear (i.e., $\sim \frac{1}{k}$) “closeness” between the law of the optimal state and μ_{tar} has appeared several times so far. For example, (4.3) implies that $G(\mathbb{P}_{X_T^{\hat{\alpha}^k}}) = G(\mathbb{P}_{X_T^{\hat{\alpha}^k}}; \mu_{\text{tar}}) \leq \frac{c}{k}$, and by Remark 3.3-(ii), this implies that $W_1(\mathbb{P}_{X_T^{\hat{\alpha}^k}}, \mu_{\text{tar}}) \sim \frac{1}{k}$. The result in (4.16) is in the same spirit, but under the stronger W_2 -distance, but compensated by an early stopping.

(ii) The Lipschitz condition (4.14) for the optimal control $\hat{\alpha}$ is not unusual in the diffusion model literature (see, e.g., [11, 12, 60]). In fact, this can be argued via regularity of the solution to the PDE (4.1) combined with the speed of decay of the density f_{tar} , which can be assumed and analyzed rigorously (see Assumption 5.4 below). We therefore consider such an assumption non-stringent. $\square$

[Proof of Proposition 4.5.] First note that $X^{\hat{\alpha}}$ and $X^{\hat{\alpha}^k}$ satisfy the following SDEs, respectively:

$$\begin{cases} dX_t^{\hat{\alpha}} = [b(t, X_t^{\hat{\alpha}}) + \hat{\alpha}_t(X_t^{\hat{\alpha}})]dt + dW_t & X_0^{\hat{\alpha}} = x_0; \\ dX_t^{\hat{\alpha}^k} = [b(t, X_t^{\hat{\alpha}^k}) + \hat{\alpha}_t^k(X_t^{\hat{\alpha}^k})]dt + dW_t, & X_0^{\hat{\alpha}^k} = x_0. \end{cases} \quad (4.17)$$

Let us now denote $\hat{\alpha}_t(x) = \hat{\alpha}(t, x)$, $\hat{\alpha}_t^k(x) = \hat{\alpha}^k(t, x)$, and define

$$b^{\hat{\alpha}}(t, x) = b(t, x) + \hat{\alpha}_t(x), \quad \Delta \hat{\alpha}_t^k(x) = \hat{\alpha}_t^k(x) - \hat{\alpha}_t(x), \quad (t, x) \in [0, T] \times \mathbb{R}^d.$$

Then we see that SDE (4.17) can be written as

$$\begin{cases} dX_t^{\hat{\alpha}} = b^{\hat{\alpha}}(t, X_t^{\hat{\alpha}})dt + dW_t & X_0^{\hat{\alpha}} = x_0; \\ dX_t^{\hat{\alpha}^k} = [b^{\hat{\alpha}}(t, X_t^{\hat{\alpha}^k}) + \Delta \hat{\alpha}_t^k(X_t^{\hat{\alpha}^k})]dt + dW_t, & X_0^{\hat{\alpha}^k} = x_0. \end{cases}$$

That is,

$$X_t^{\hat{\alpha}} - X_t^{\hat{\alpha}^k} = \int_0^t [b^{\hat{\alpha}}(s, X_s^{\hat{\alpha}}) - b^{\hat{\alpha}}(s, X_s^{\hat{\alpha}^k}) + \Delta \hat{\alpha}_s^k(X_s^{\hat{\alpha}^k})]ds, \quad t \in [0, T]$$

Note that by Assumption 2.2 and (4.14), $b^{\hat{\alpha}}$ is uniform Lipschitz in x (with Lipschitz constant $L + \kappa$), and applying the estimate (4.9), we deduce easily that

$$\mathbb{E}[|X_t^{\hat{\alpha}} - X_t^{\hat{\alpha}^k}|^2] \leq 2T \int_0^t [(L + \kappa)^2 \mathbb{E}[|X_s^{\hat{\alpha}} - X_s^{\hat{\alpha}^k}|^2]ds + \frac{2c^2}{k^2} \ln \left[\frac{T}{T-t} \right]. \quad (4.18)$$

In what follows let us denote $C > 0$ to be a generic constant depending only on L, κ, c , but independent of k , and we allow it to vary from line to line. Then, by a simple calculation using Gronwall's inequality, we see that (4.18) lead to that

$$\mathbb{E}[|X_t^{\hat{\alpha}} - X_t^{\hat{\alpha}^k}|^2] \leq \frac{C}{k^2}(\ln T - \ln(T-t))e^{Ct}, \quad t \in [0, T]. \quad (4.19)$$

Furthermore, for any $\varepsilon > 0$, using the monotonicity of the log function we deduce from (4.19) that

$$\mathbb{E}[|X_{T-\varepsilon}^{\hat{\alpha}} - X_{T-\varepsilon}^{\hat{\alpha}^k}|^2] \leq \frac{C}{k^2}(\ln T - \ln \varepsilon)$$

It then follows that

$$W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}}}, \mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}^k}}) \leq \mathbb{E}[|X_{T-\varepsilon}^{\hat{\alpha}} - X_{T-\varepsilon}^{\hat{\alpha}^k}|^2]^{1/2} \leq \frac{\sqrt{C}}{k} \sqrt{\ln T - \ln \varepsilon}.$$

Finally, since the function $b^{\hat{\alpha}} = b + \hat{\alpha}$ is Lipschitz, by standard $\mathbb{L}^2$ -continuity result of SDE, we have

$$W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}}}, \mathbb{P}_{X_T^{\hat{\alpha}}}) \leq C\varepsilon,$$

and consequently, noting that $\mathbb{P}_{X_T^{\hat{\alpha}}} = \mu_{\text{tar}}$, we obtain

$$W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}}}, \mu_{\text{tar}}) \leq W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}}}, \mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}^k}}) + W_2(\mathbb{P}_{X_{T-\varepsilon}^{\hat{\alpha}^k}}, \mathbb{P}_{X_T^{\hat{\alpha}}}) \leq \frac{\sqrt{C(\ln T - \ln \varepsilon)}}{k} + C\varepsilon,$$

proving (4.15), whence the proposition. $\square$

5 Stability of the Solutions to the SBP

We note that all the results in the previous section are based on an important assumption: $\mu_{\text{ini}} = \delta_{x_0}$, for some $x_0 \in \mathbb{R}^d$. In this and the next section, we shall extend the results to more general initial condition $\mu_{\text{ini}} \in \mathcal{P}_2(\mathbb{R}^d)$ with density f_{ini} , and establish a similar rate of convergence.

We shall begin an important aspect in probability theory, which is the basis for the so-called stability issues of the solutions to the classic Schrödinger bridge problem. For notational convenience, we still denote $p(\cdot, \cdot; \cdot, \cdot)$ to be the transition density of a standard $\mathbb{R}^d$ -valued diffusion (3.1). We begin with the following well-known result in diffusion theory (cf. e.g., [7]).

Proposition 5.1 ([7]). *For any $\mu_{\text{ini}}, \mu \in \mathcal{P}(\mathbb{R}^d)$, there exists a unique pair of σ -finite measures $\nu_0, \nu_T \in \mathcal{M}(\mathbb{R}^d)$ such that the measure π on $\mathbb{R}^d \times \mathbb{R}^d$ defined by*

$$\pi(E) = \int_E p(T, y; 0, x) \nu_0(dx) \nu_T(dy), \quad E \in \mathcal{B}(\mathbb{R}^d \times \mathbb{R}^d) \quad (5.1)$$

has marginals μ_{ini} and μ . Furthermore, ν_T and μ (resp. ν_0 and μ_{ini}) are mutually absolutely continuous, denoted by $\nu_T \simeq \mu$ (resp. $\nu_0 \simeq \mu_{\text{ini}}$). $\square$

Fix $\mu_{\text{ini}} \in \mathcal{P}(\mathbb{R}^d)$ apriori. Let us denote a (well-defined) mapping $\mathcal{T} : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{M}(\mathbb{R}^d) \times \mathcal{M}(\mathbb{R}^d)$ by $\mathcal{T}(\mu) = (\nu_0, \nu_T)$. Note that in Proposition 5.1 the measures (ν_0, ν_T) are only σ -finite in general, to facilitate our discussion, we shall consider, the following set:

$$\mathcal{D}_{\mu_{\text{ini}}} := \{\mu \in \mathcal{P}_2(\mathbb{R}^d) : \mathcal{T}(\mu) \ll \text{Leb}(\cdot); \mathcal{T}(\mu)(\mathbb{R}^d \times \mathbb{R}^d) < \infty\}.$$

Here $Leb(\cdot)$ denotes the Lebesgue measure on $\mathbb{R}^d \times \mathbb{R}^d$.

We note that if $\mu \in \mathcal{D}_{\mu_{\text{ini}}}$ and $(\nu_0, \nu_T) = \mathcal{T}(\mu)$, then ν_T must have a density function, which we shall denote by $\rho^\mu \in \mathbb{L}^1(\mathbb{R}^d)$. Moreover, we define an operator $S : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ by

$$S[\mu](dy) = \int_{\mathbb{R}^d} p(T, y; 0, x) \mu(dx) dy, \quad \mu \in \mathcal{P}_2(\mathbb{R}^d). \quad (5.2)$$

Clearly, if $\mu \in \mathcal{P}(\mathbb{R}^d)$, then $S[\mu](dy) = f_{X_T^{0,\mu}}(y) dy$, where $X^{0,\mu} = \{X_t^{0,\mu}\}_{t \in [0,T]}$ denotes the solution to (3.1) with $X_0^{0,\mu} \sim \mu$. But the operator S can be naturally extended to any $\mu \in \mathcal{M}(\mathbb{R}^d)$, provided the right-hand side of (5.2) is well-defined.

Let us now recall a well-known analogue of Lemma 3.1 in the case of general initial condition $\mu_{\text{ini}} \in \mathcal{P}_2(\mathbb{R}^d)$.

Proposition 5.2 ([18, Theorem 3.2]). *Let $\mu_{\text{ini}}, \mu_{\text{tar}} \in \mathcal{P}_2(\mathbb{R}^d)$ and $(\nu_0, \nu_T) = \mathcal{T}(\mu_{\text{tar}})$. Assume that $D_{\text{KL}}(\mu_{\text{ini}} \parallel \nu_0) < \infty$ and $D_{\text{KL}}(\mu_{\text{tar}} \parallel S[\nu_0]) < \infty$. Then, the optimal control for the (original) SBP (2.5)-(2.6) is given by $\hat{\alpha}_t = \nabla \log h(t, X_t^{\hat{\alpha}})$ where, denoting $\rho^{\mu_{\text{tar}}}(\cdot)$ to be the density function of ν_T ,*

$$h(t, x) := \int_{\mathbb{R}^d} p(T, z; t, x) \rho^{\mu_{\text{tar}}}(z) dz. \quad (5.3)$$

Moreover, it holds that

$$J(\hat{\alpha}) = \int_{\mathbb{R}^d} \log \rho^{\mu_{\text{tar}}}(y) \mu_{\text{tar}}(dy) - D_{\text{KL}}(\mu_{\text{ini}} \parallel \nu_0). \quad (5.4)$$

We note that in the above $D_{\text{KL}}(\mu_{\text{ini}} \parallel \nu_0) = \int \log \frac{\mu_{\text{ini}}(dx)}{\nu_0(dx)} \mu_{\text{ini}}(dx)$ (see footnote 1), and (5.1) implies that $\frac{\mu_{\text{ini}}(dx)}{\nu_0(dx)} = \int p(T, y; 0, x) \rho^{\mu_{\text{tar}}}(y) dy$. Therefore (5.4) can be rewritten as

$$\begin{aligned} J(\hat{\alpha}) &= \int_{\mathbb{R}^d} \log \rho^{\mu_{\text{tar}}}(y) \mu_{\text{tar}}(dy) - \int_{\mathbb{R}^d} \log \left(\int_{\mathbb{R}^d} p(T, y; 0, x) \rho^{\mu_{\text{tar}}}(y) dy \right) \mu_{\text{ini}}(dx) \\ &= \mathbb{E}[\log \rho^{\mu_{\text{tar}}}(X_T^{\hat{\alpha}})] - \int_{\mathbb{R}^d} \log h(0, x) \mu_{\text{ini}}(dx) = \mathbb{E}[\log \rho^{\mu_{\text{tar}}}(X_T^{\hat{\alpha}})] - \mathbb{E}[\log h(0, X_0^{\hat{\alpha}})]. \end{aligned}$$

Moreover, for a fixed $\mu_{\text{ini}} \in \mathcal{P}_2(\mathbb{R}^d)$, we define $h^\mu(t, x) = \int_{\mathbb{R}^d} p(T, z; t, x) \rho^\mu(z) dz$. Then, we have the following result.

Lemma 5.3 ([29, Lemma 3.1]). *Let $\mu \in \mathcal{D}_{\mu_{\text{ini}}}$. Then, for any $\{\alpha_t\} \subset \mathbb{L}_{\mathbb{F}_0}^2([0, T]; \mathbb{R}^d)$, it holds that*

$$J(\alpha) \geq \mathbb{E}[\log \rho^\mu(X_T^\alpha)] - \mathbb{E}[\log h^\mu(0, X_0^\alpha)].$$

The equality holds when $\alpha_t = \alpha_t^\mu = \nabla \log h^\mu(t, X_t^{\alpha_t^\mu})$, $t \in [0, T]$ and $X_T^{\alpha_t^\mu} \sim \mu$. $\square$

From Proposition 5.2 and Lemma 5.3 we see that the density function ρ^μ plays an important role in the structure of the solution of SBP. We shall be particularly interested in the continuous dependence of $\rho^\mu := \Gamma_1(\mu)$ on $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, which we shall refer to as the *Stability of the SBP*, borrowing the well-known concept of the SBP theory (cf. e.g., [10, 22, 46]).

To continue our discussion, we shall identify a set $\mathcal{E} \subset \mathcal{P}_2(\mathbb{R}^d)$ on which an argument based on Schauder's fixed-point theorem can be carried out. We begin by denoting

$$\mathcal{K} := \left\{ \mu \in \mathcal{P}_2(\mathbb{R}^d) : \mu \text{ has density } f_\mu \in \mathbb{L}^1(\mathbb{R}^d) \right\}.$$

Furthermore, we shall make use of the following assumption.

Assumption 5.4. (i) There exists a constant $C > 0$, independent of $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, such that $|\log f_\mu(x)| \leq C|x|^2$, $x \in \mathbb{R}^d$;

(ii) There exists a function $g \in \mathbb{L}^2(\mathbb{R}^d; (0, 1])$, with $\int_{\mathbb{R}^d} |x|^2 g(x) dx < \infty$, and a constant $K > 0$, independent of $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, such that $\left\| \frac{f_\mu}{g^2} \right\|_\infty \leq K$.

We shall consider the following two sets that will play a crucial role in our discussion.

$$\mathcal{E} := \{\mu \in \mathcal{K} : \text{Assumption 5.4 holds}\} \subset \mathcal{P}_2(\mathbb{R}^d); \quad \mathcal{S}_{\mathcal{E}} := \{f_\mu : \mu \in \mathcal{E}\} \subset \mathbb{L}^1(\mathbb{R}^d). \quad (5.5)$$

Remark 5.5. We note that Assumption 5.4-(i) provides a uniform lower bound for the densities $f_\mu \in \mathcal{K}$, whereas Assumption 5.4-(ii) controls the decay of $f_\mu \in \mathcal{K}$ as $x \sim \infty$, hence mutually non-inclusive. A typical example of the function g is $e^{-c|x|^2}$, $x \in \mathbb{R}^d$, $c > 0$. In fact, in light of the estimate (3.3), such a property holds essentially for all transition probabilities of diffusion processes. $\square$

The following lemma lists some basic properties of the set $\mathcal{E}$ (or Assumption 5.4).

Lemma 5.6. Assume that Assumption 5.4 is in force. Then it holds that

- (i) The set $\{f_\mu\}_{\mu \in \mathcal{E}}$ is uniformly bounded in $\mathbb{L}^2(\mathbb{R}^d)$.
- (ii) The set $\{f_\mu\}_{\mu \in \mathcal{E}}$ is uniformly integrable in $\mathcal{P}_2(\mathbb{R}^d)$, in the sense that

$$\lim_{R \rightarrow \infty} \sup_{\mu \in \mathcal{E}} \int_{\{|x| \geq R\}} |x|^2 f_\mu(x) dx = 0. \quad (5.6)$$

- (iii) If $\{\mu_n\}_{n \geq 1} \subset \mathcal{E}$ such that $\mu_n \Rightarrow \mu$, as $n \rightarrow \infty$, then $\|f_{\mu_n} - f_\mu\|_{\mathbb{L}^1} \rightarrow 0$.

Proof. For any $\mu \in \mathcal{E}$, we note that $0 < g(x) \leq 1$, and by assumption,

$$\int_{\mathbb{R}^d} |f_\mu(x)|^2 dx \leq K^2 \int_{\mathbb{R}^d} |g(x)|^4 dx \leq K^2 \|g\|_{\mathbb{L}^2}^2,$$

That is $\{f_\mu\}_{\mu \in \mathcal{E}}$ is uniformly bounded (by $K\|g\|_{\mathbb{L}^2}$) in $\mathbb{L}^2(\mathbb{R}^d)$, proving (i).

Similarly, for any $\mu \in \mathcal{E}$, by the absolute continuity of the integral we have

$$\sup_{\mu \in \mathcal{E}} \int_{\{|x| \geq R\}} |x|^2 f_\mu(x) dx \leq K \int_{\{|x| \geq R\}} |x|^2 g(x) dx \rightarrow 0, \quad \text{as } R \rightarrow \infty,$$

This proves (5.6), whence (ii).

The proof of part (iii) is slightly more involved, which is in the spirit of the so-called Scheffé's theorem (cf. [59]). We note that $\mu_n \Rightarrow \mu$ amounts to saying that $f_{\mu_n} \xrightarrow{w} f_\mu$, as $n \rightarrow \infty$, in $\mathbb{L}^2(\mathbb{R}^d)$. To show $f_{\mu_n} \rightarrow f_\mu$ in $\mathbb{L}^1(\mathbb{R}^d)$, we first consider, for each $m > 0$, the smooth mollifiers $\varphi^m \in \mathbb{C}^\infty(\mathbb{R}^d; \mathbb{R}_+)$ such that $\int_{\mathbb{R}^d} \varphi^m(z) dz = 1$, $m \geq 1$, and denote

$$f_{\mu_n}^m(x) = [\varphi^m * f_{\mu_n}](x) = \int_{\mathbb{R}^d} \varphi^m(x-z) f_{\mu_n}(z) dz; \quad f_\mu^m(x) = [\varphi^m * f_\mu](x), \quad x \in \mathbb{R}^d.$$

Then it is clear that for each $n \in \mathbb{N}$, $\lim_{m \rightarrow \infty} f_{\mu_n}^m(x) = f_{\mu_n}(x)$ and $\lim_{m \rightarrow \infty} f_\mu^m(x) = f_\mu(x)$, for a.e. $x \in \mathbb{R}^d$. We should remark that the convergence is uniform in n . Indeed, by Assumption 5.4 and Dominated Convergence Theorem we have, as $m \rightarrow \infty$, for all $n \geq 0$, and $x \in \mathbb{R}^d$,

$$|f_{\mu_n}^m(x) - f_{\mu_n}(x)| \leq \int_{\mathbb{R}^d} |\varphi^m(x-z) - \delta_x(z)| f_{\mu_n}(z) dz \leq K \int_{\mathbb{R}^d} |\varphi^m(x-z) - \delta_x(z)| g^2(z) dz \rightarrow 0.$$

Furthermore, since $\sup_{\mu \in \mathcal{E}} |f_\mu| \leq Kg^2 \in \mathbb{L}^1(\mathbb{R}^d)$, by Dominated Convergence Theorem we have $\lim_{m \rightarrow \infty} f_{\mu_n}^m = f_{\mu_n}$ in $\mathbb{L}^1(\mathbb{R}^d)$, uniformly for $n \geq 0$. That is, for any $\varepsilon > 0$, there exists $M(\varepsilon) > 0$, such that for all $n \geq 1$, it holds that

$$\|f_{\mu_n}^m - f_{\mu_n}\|_{\mathbb{L}^1} < \frac{\varepsilon}{3}; \quad \|f_\mu^m - f_\mu\|_{\mathbb{L}^1} < \frac{\varepsilon}{3}, \quad \text{whenever } m > M. \quad (5.7)$$

In the sequel we fix $m > M(\varepsilon)$, and take a closer look at the sequence $\{f_{\mu_n}^m\}_{n \geq 1}$. Clearly, each $f_{\mu_n}^m$ is still a density function, and it holds that

$$\sup_n |f_{\mu_n}^m(y)| \leq \sup_n [\varphi^m * |f_{\mu_n}|](y) \leq K. \quad (5.8)$$

Moreover, since φ^m is continuous, thus for any $x, y \in \mathbb{R}^d$, applying the Dominated Convergence Theorem we have

$$\begin{aligned} |f_{\mu_n}^m(x+y) - f_{\mu_n}^m(x)| &\leq \int_{\mathbb{R}^d} |\varphi^m(x+y-z) - \varphi^m(x-z)| |f_{\mu_n}(z)| dz \\ &\leq K \int_{\mathbb{R}^d} |\varphi^m(z'+y) - \varphi^m(z')| dz' \rightarrow 0, \quad \text{as } y \rightarrow 0. \end{aligned}$$

Clearly, the convergence above is uniform in n . That is, the sequence $\{f_{\mu_n}^m\}_{n \geq 1}$ is so-called *asymptotically equi-continuous* in the sense of Sweeting [59]. This, together with (5.8), implies that $\lim_{n \rightarrow \infty} f_{\mu_n}^m = f_\mu^m$, uniformly on compacts in $\mathbb{R}^d$ (cf. [59, Theorem 1]). Applying the Dominated Convergence Theorem again we have $\lim_{n \rightarrow \infty} \|f_{\mu_n}^m - f_\mu^m\|_{\mathbb{L}^1} = 0$. That is, for the given $\varepsilon > 0$ in (5.7), there exists $N > 0$ such that $\|f_{\mu_n}^m - f_\mu^m\|_{\mathbb{L}^1} < \frac{\varepsilon}{3}$, whenever $n > N$. This, together with (5.7), yields

$$\|f_{\mu_n} - f_\mu\|_{\mathbb{L}^1} \leq \|f_{\mu_n} - f_{\mu_n}^m\|_{\mathbb{L}^1} + \|f_{\mu_n}^m - f_\mu^m\|_{\mathbb{L}^1} + \|f_\mu^m - f_\mu\|_{\mathbb{L}^1} < \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon, \quad n > N,$$

proving (iii), whence the Lemma. $\square$

We are now ready to study our main stability result. More precisely, we shall argue that the mapping $\Gamma_1 : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{L}^1(\mathbb{R}^d)$ is continuous. That is, that $\mu_n \in \mathcal{P}_2(\mathbb{R}^d)$ weakly converges to some $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ in Prohorov metric would imply that ρ^{μ_n} converges to ρ^μ in $\mathbb{L}^1(\mathbb{R}^d)$. Such a result, to the best of our knowledge, is novel in the literature.

To simplify our discussion, in what follows, we assume $T = 1$. Recall that π in (5.1) has marginals μ_{ini} and μ , and in what follows, we shall assume that μ_{ini} is fixed and $\mu \in \mathcal{E}$. Let us now consider the following *entropic optimal transport problem*:

$$I(\mu) := \inf_{\pi \in \Pi(\mu_{\text{ini}}, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \mathbf{c}(x, y) \pi(dxdy) + D_{\text{KL}}(\pi \| \mu_{\text{ini}} \otimes \mu), \quad (5.9)$$

where $\Pi(\mu_{\text{ini}}, \mu)$ is the set of all coupling probability measures π on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals μ_{ini} and μ ; and $\mathbf{c}(\cdot, \cdot)$ is a continuous *cost function*. It is well-known (see, e.g., [16, 22, 31, 46]) that the minimization (5.9) admits a unique solution $\hat{\pi}$, whose density takes the form:

$$\hat{\pi}(dxdy) = \exp(-\mathbf{c}(x, y) + \phi^\mu(x) + \psi^\mu(y)) \mu_{\text{ini}}(dx) \mu(dy), \quad (5.10)$$

where $\phi^\mu, \psi^\mu: \mathbb{R}^d \rightarrow \mathbb{R}$ are two measurable functions, often referred to as the *Schrödinger potentials*. It is clear that the pair (ϕ^μ, ψ^μ) is unique up to an additive constant. That is, if (ϕ^μ, ψ^μ) is a pair of

Schrödinger potentials, then so is $(\phi^\mu + c, \psi^\mu - c)$. Furthermore, since both μ_{ini} and μ are probability measures, we can easily choose a constant c so that the following *symmetric normalization* holds:

$$\int \phi^\mu(x) \mu_{\text{ini}}(dx) = \int \psi^\mu(y) \mu(dy). \quad (5.11)$$

(Otherwise we take $c = \frac{1}{2}[-\int \phi^\mu(x) \mu_{\text{ini}}(dx) + \int \psi^\mu(y) \mu(dy)]$.) Note that under the symmetric normalization, the Schrödinger potentials is unique. The following stability result for the mappings $\mu \mapsto (\phi^\mu, \psi^\mu)$ is crucial for our discussion.

Lemma 5.7 ([10, Theorem 1.1]). *Assume that the cost function $\mathbf{c}(\cdot, \cdot) \in \mathbb{C}^{k+1}(\mathbb{R}^d \times \mathbb{R}^d)$ for some $k \in \mathbb{N}$. Then there exists $C > 0$ depending only on $\|\mathbf{c}\|_{\mathbb{C}^{k+1}}$, such that for all $\mu_1, \mu_2 \in \mathcal{P}_2(\mathbb{R}^d)$, it holds that*

$$\|(\phi^{\mu_1} - \phi^{\mu_2}, \psi^{\mu_1} - \psi^{\mu_2})\|_* \leq CW_2(\mu_1, \mu_2),$$

where $\|(\phi, \psi)\|_* := \inf_{c \in \mathbb{R}} \left\{ \|\phi - c\|_{\mathbb{C}^k(\mathbb{R}^d)} + \|\psi + c\|_{\mathbb{C}^k(\mathbb{R}^d)} \right\}$. $\square$

We now proceed to prove the main result of this section. To begin with, let us consider the entropic optimal transport problem (5.9) with $\mathbf{c}(x, y) := -\log p(1, y; 0, x)$, $x, y \in \mathbb{R}^d$, where $p(s, y; t, x)$, $0 \leq t < s \leq 1$ and $x, y \in \mathbb{R}^d$, is the transition density of the diffusion (3.1). By (5.10), for fixed $\mu_{\text{ini}}, \mu \in \mathcal{P}_2(\mathbb{R}^d)$, the unique solution for this entropic optimal transport problem is given by (see also [46])

$$\hat{\pi}(dxdy) = p(1, y; 0, x) e^{\phi^\mu(x) + \psi^\mu(y)} f_{\text{ini}}(x) f_\mu(y) dxdy,$$

where ϕ^μ and ψ^μ are the Schrödinger potentials, and we shall enforce the symmetric normalization so that they satisfy (5.11). Since $\hat{\pi}$ has the marginals μ_{ini} and μ , by the uniqueness of $(\nu_0, \nu_T) = \mathcal{T}(\mu)$, in Proposition 5.1, we can conclude that

$$p(1, y; 0, x) \rho_0(x) \rho^\mu(y) = p(1, y; 0, x) e^{\phi^\mu(x) + \psi^\mu(y)} f_{\text{ini}}(x) f_\mu(y), \quad x, y \in \mathbb{R}^d,$$

where ρ_0 is the density function of ν_0 . An easy argument of separation of variables then yields that

$$\rho^\mu(y) = e^{\psi^\mu(y)} f_\mu(y); \quad \rho_0(x) = e^{\phi^\mu(x)} f_{\text{ini}}(x), \quad x, y \in \mathbb{R}^d. \quad (5.12)$$

Now note that the transition density $p(\cdot, \cdot; \cdot, \cdot)$ is a classical solution to the Kolmogorov PDE. Thanks to Assumption 2.2, we can assume without loss of generality that $\mathbf{c}(\cdot, \cdot) = -\log p(1, \cdot; 0, \cdot) \in \mathbb{C}^2(\mathbb{R}^d \times \mathbb{R}^d)$. Thus according to Lemma 5.7 and noting the definition of $\|\cdot\|_*$, we see that, modulo some constant normalization, we have that the Schrödinger potential $(\phi^{\mu_n}, \psi^{\mu_n})$ itself satisfies the estimate:

$$\|\phi^{\mu_n} - \phi^\mu\|_{\mathbb{L}^\infty} + \|\psi^{\mu_n} - \psi^\mu\|_{\mathbb{L}^\infty} \leq CW_2(\mu_n, \mu). \quad (5.13)$$

Here in the above the constant $C > 0$ depending only on $\|\mathbf{c}\|_{\mathbb{C}^2}$, but independent of n .

Furthermore, we note that $\mathbf{c} \in \mathbb{C}^2$ also lead to the following *a priori* estimate of the Schrödinger potential (see, e.g., [46, Lemma 2.1]):

$$\psi^\mu(y) \leq \int_{\mathbb{R}^d} \mathbf{c}(x, y) \mu_{\text{ini}}(dx) =: \xi(y), \quad y \in \mathbb{R}^d. \quad (5.14)$$

Recall the fundamental estimate (3.3) and the definition of $\mathbf{c}(\cdot, \cdot)$, it is readily seen that $|\xi(y)| \sim \lambda'|y|^2$, as $y \rightarrow \infty$, for some constant $\lambda' > 0$ depending only on the coefficient $b(\cdot, \cdot)$ in SDE(3.1). In light of Remark 5.5 and [46, Lemma 2.1], in what follows we shall impose some explicit assumptions on the Schrödinger potentials $\{\psi^\mu(\cdot)\}_{\mu \in \mathcal{E}}$ and their upper bound $\xi(\cdot)$ in (5.14), which are in line with Assumption 5.4.

Assumption 5.8. (i) There exists a constant $C > 0$, such that for any $\mu \in \mathcal{E}$, it holds that

$$|\psi^\mu(y)| \leq C|y|^2, \quad y \in \mathbb{R}^d. \quad (5.15)$$

(ii) In Assumption 5.4, the control function g satisfies

$$\eta(\cdot) := e^{\xi(\cdot)} g^2(\cdot) \in \mathbb{L}^1(\mathbb{R}^d). \quad (5.16)$$

Now for any $f_\mu \in \mathcal{S}_\mathcal{E}$, by Assumption 5.4 and (5.14) we have

$$0 \leq \rho^\mu(y) = e^{\psi^\mu(y)} f_\mu(y) \leq e^{\xi(y)} f_\mu(y) \leq e^{\xi(y)} K g^2(y) = K \eta(y), \quad y \in \mathbb{R}^d.$$

Consequently, we conclude that $\rho^\mu \in \mathbb{L}^1(\mathbb{R}^d)$ for any $\mu \in \mathcal{E}$, thanks to (5.16).

Bearing the above discussion in mind, we are now ready to present the main result of this section.

Proposition 5.9. Assume that Assumptions 2.2 and 5.4 are in force. Assume further that $\{\mu_n\}_{n \geq 1} \subset \mathcal{E}$ and $\mu_n \Rightarrow \mu$ in Prohorov metric. Then $\|\rho^{\mu_n} - \rho^\mu\|_{\mathbb{L}^1} = \|\Gamma_1(\mu_n) - \Gamma_1(\mu)\|_{\mathbb{L}^1} \rightarrow 0$, as $n \rightarrow \infty$.

Proof. Assume $\{\mu_n\}_{n \geq 1} \subset \mathcal{E}$, and $\mu_n \Rightarrow \mu$, in Prohorov metric. By Lemma 5.6-(ii), $\{\mu_n\}$ is uniformly integrable in $\mathbb{L}^2$, thanks to Assumption 5.4, and thus by the relationship between Wasserstein distance and Prohorov metric (see, [65, Theorem 7.12]), we have $W_2(\mu_n, \mu) \rightarrow 0$, as $n \rightarrow \infty$. Thus, it follows from (5.13) that $\|\psi^{\mu_n} - \psi^\mu\|_{\mathbb{L}^\infty} \rightarrow 0$, as $n \rightarrow \infty$.

Next, for each $\mu_n \in \mathcal{E}$, $n \geq 1$, and μ , we apply (5.12) and write

$$\rho^{\mu_n}(y) = e^{\psi^{\mu_n}(y)} f_{\mu_n}(y), \quad \rho^\mu(y) = e^{\psi^\mu(y)} f_\mu(y), \quad x, y \in \mathbb{R}^d.$$

Therefore, for $y \in \mathbb{R}^d$, we have

$$\begin{aligned} |\rho^\mu(y) - \rho^{\mu_n}(y)| &= |e^{\psi^\mu(y)} f_\mu(y) - e^{\psi^{\mu_n}(y)} f_{\mu_n}(y)| \\ &\leq |e^{\psi^\mu(y)} - e^{\psi^{\mu_n}(y)}| f_{\mu_n}(y) + e^{\psi^\mu(y)} |f_{\mu_n}(y) - f_\mu(y)| =: I_n^1(y) + I_n^2(y), \end{aligned}$$

where I_n^i , $i = 1, 2$ are defined in an obvious way. It then suffices to show that both I_n^1 and $I_n^2 \rightarrow 0$ in $\mathbb{L}^1$, as $n \rightarrow \infty$.

To this end, we first recall that $\|\psi^{\mu_n} - \psi^\mu\|_{\mathbb{L}^\infty} \rightarrow 0$, as $n \rightarrow \infty$. Hence there exists $N > 0$, such that $\psi^{\mu_n}(y) \leq \psi^\mu(y) + 1$, for all $y \in \mathbb{R}^d$, whenever $n \geq N$. Thus, for $n \geq N$, we have

$$0 \leq I_n^1(y) \leq (|e^{\psi^\mu(y)}| + |e^{\psi^{\mu_n}(y)}|) f_{\mu_n}(y) \leq 2e^{\psi^\mu(y)+1} f_{\mu_n}(y) \leq 2e \cdot e^{\xi(y)} g^2(y) = 2e\eta(y).$$

Here in the above, the last inequality holds due to Assumption 5.4 and (5.16). Since $\eta \in \mathbb{L}^1$ by (5.16), the Dominated Convergence Theorem implies that $I_n^1(\cdot)$ converges to 0 in $\mathbb{L}^1(\mathbb{R}^d)$ as $n \rightarrow \infty$, because ψ^{μ_n} converges uniformly to ψ^μ on $\mathbb{R}^d$.

Finally, since $I_n^2(y) \leq 2K\eta(y)$, and $f_{\mu_n} \rightarrow f_\mu$ in $\mathbb{L}^1(\mathbb{R}^d)$, thanks to Lemma 5.6-(iii), we can apply Dominated Convergence again to get $I_n^2(\cdot)$ converges to 0 in $\mathbb{L}^1(\mathbb{R}^d)$, as $n \rightarrow \infty$, proving the proposition. $\square$

6 Existence of optimal control and convergence for general μ_{ini}

In this section, we shall extend the results of Section 3 and show that the Problem 2.4 has solution for each $k \in \mathbb{N}$ when μ_{ini} is an arbitrary distribution with density f_{ini} in $\mathcal{P}_2(\mathbb{R}^d)$. To be more

precise, for fixed $k \in \mathbb{N}$, let $J^k(\alpha)$ be the cost functional in Problem 2.4. Applying Lemma 5.3, for any $\mu \in \mathcal{D}_{\mu_{\text{ini}}}$, we have

$$\begin{aligned} J^k(\alpha) &= \mathbb{E} \left[\frac{1}{2} \int_0^T |\alpha_s|^2 ds + kG(\mathbb{P}_{X_T^\alpha}; \mu_{\text{tar}}) \right] \\ &\geq kG(\mathbb{P}_{X_T^\alpha}; \mu_{\text{tar}}) + \mathbb{E}[\log \rho^\mu(X_T^\alpha)] - \mathbb{E}[\log h^\mu(0, X_0^\alpha)] \end{aligned} \quad (6.1)$$

and the equality holds when $\alpha_t = \alpha_t^\mu = \nabla \log h^\mu(t, X_t^{\alpha^\mu})$ and $X_T^{\alpha^\mu} \sim \mu$. Our main goal of this section is to determine the probability measure $\hat{\mu}$, such that $\hat{\alpha}_t = \nabla \log h^{\hat{\mu}}(t, X_t^{\hat{\alpha}})$ is the optimal control to Problem 2.4, where $h^{\hat{\mu}}(t, x) = \mathbb{E}_{t,x}[\rho^{\hat{\mu}}(X_T)]$.

We now give the heuristic idea of the construction of "solution mapping" Γ . Let μ_{ini} be given. For any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, first apply Lemma 5.3 to get the feedback control $\alpha_t^\mu = \nabla \log h^\mu(t, X_t^{\alpha^\mu})$ so that $\mathbb{P}_{X_T^{\alpha^\mu}} = \mu$ and

$$J(\alpha^\mu) = \mathbb{E}[\log \rho^\mu(X_T^{\alpha^\mu})] - \mathbb{E}[\log h^\mu(0, X_0^{\alpha^\mu})]. \quad (6.2)$$

In what follows we fix $k \in \mathbb{N}$. To find the $\hat{\mu}^k$ such that $J^k(\alpha^{\hat{\mu}^k}) = \inf J^k(\alpha)$, we consider a mapping: $\Gamma_2 : \mathbb{L}^1(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^d)$ by $\Gamma_2(\rho^\mu) = \mu'$ where

$$\mu' = \arg \min_{\bar{\mu} \in \mathcal{P}_2(\mathbb{R}^d)} \left\{ kG(\bar{\mu}) + \int_{\mathbb{R}^d} \log \rho^\mu(y) \bar{\mu}(dy) \right\}. \quad (6.3)$$

Finally, we define $\Gamma = \Gamma_2 \circ \Gamma_1 : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathcal{P}_2(\mathbb{R}^d)$, and we shall argue that the mapping Γ has a fixed point $\hat{\mu} \in \mathcal{E}$, where $\mathcal{E}$ is defined by (5.5). Clearly, if $\Gamma(\hat{\mu}) = \hat{\mu}$, then we can still define $\hat{\alpha} = \alpha^{\hat{\mu}}$, and by Lemma 5.3 we have $\mathbb{P}_{X_T^{\hat{\alpha}}} = \hat{\mu}$. Thus by (6.2), for any $\alpha \in \mathbb{L}_{\mathbb{F}^0}^2([0, T]; \mathbb{R}^d)$ we have

$$\begin{aligned} J^k(\hat{\alpha}) &= J(\hat{\alpha}) + kG(\hat{\mu}) = kG(\hat{\mu}) + \mathbb{E}[\log \rho^{\hat{\mu}}(X_T^{\hat{\alpha}})] - \mathbb{E}[\log h^{\hat{\mu}}(0, X_0^{\hat{\alpha}})] \\ &\leq kG(\mathbb{P}_{X_T^\alpha}) + \mathbb{E}[\log \rho^{\hat{\mu}}(X_T^\alpha)] - \mathbb{E}[\log h^{\hat{\mu}}(0, X_0^\alpha)] \leq J^k(\alpha). \end{aligned}$$

Here in the above the first inequality is due to (6.3), and the last inequality is due to (6.1). This shows that $\hat{\alpha}$ is the minimizer of $J^k(\cdot)$.

We now show that the set $\mathcal{E} \subset \mathcal{D}_{\mu_{\text{ini}}}$ defined by (5.5) satisfies all the necessary properties, thanks to the Lemma 5.6 and Proposition 5.9 that we established in the last section, so that the mapping Γ possesses a fixed point on $\mathcal{E}$ by Schauder's fixed-point theorem. Our main result is as follows.

Theorem 6.1. *Assume that Assumptions 5.4 is in force. Consider the set $\mathcal{E}$ defined by (5.5). Then the following hold:*

- (i) $\mathcal{E}$ is convex and closed under Prohorov metric, and $\mathcal{S}_{\mathcal{E}}$ is convex and closed in $\mathbb{L}^1(\mathbb{R}^d)$;
- (ii) $\Gamma(\mathcal{E}) \subseteq \mathcal{E}$, and is precompact in $\mathcal{P}_2(\mathbb{R}^d)$, under both Prohorov and Wasserstein metric;
- (iii) Γ is continuous on $\mathcal{E}$, under Prohorov metric.

Consequently, the mapping Γ has a fixed point in $\mathcal{E}$.

Proof. Since the last statement is a direct consequence of Schauder's fixed point theorem, applying to the space $\mathcal{P}(\mathbb{R}^d)$ with Prohorov metric, we need only prove the properties (i)-(iii).

(i) is obvious.

(ii) By definition of $\mathcal{S}_{\mathcal{E}}$ we rewrite (6.3) as

$$f_{\mu'} = \arg \min_{f_{\bar{\mu}} \in \mathcal{S}_{\mathcal{E}}} \left\{ kG(\bar{\mu}) + \int_{\mathbb{R}^d} \log \rho^{\mu}(y) f_{\bar{\mu}}(y) dy \right\}.$$

Since $\mathcal{S}_{\mathcal{E}}$ is convex and closed in $\mathbb{L}^1(\mathbb{R}^d)$, it follows that $f_{\mu'} \in \mathcal{S}_{\mathcal{E}}$, and thus $\Gamma(\mathcal{E}) \subseteq \mathcal{E}$. We are to show that $\Gamma(\mathcal{E})$ is precompact in $\mathcal{P}_2(\mathbb{R}^d)$.

To this end, let $\{\mu_n\} \subseteq \Gamma(\mathcal{E})$ be any sequence, we shall find a subsequence $\{\mu_{n_k}\}_{k \geq 1}$ such that $\lim_{k \rightarrow \infty} \mu_{n_k} = \mu \in \mathcal{P}_2(\mathbb{R}^d)$, under both Prohorov metric and W_2 -metric. Since $\Gamma(\mathcal{E}) \subseteq \mathcal{E}$, by Lemma 5.6-(i), $\{f_{\mu_n}\}$ is bounded in $\mathbb{L}^2(\mathbb{R}^d)$. Thus by Banach-Alaoglu Theorem and noting that $\mathbb{L}^2$ is reflexive, $\{f_{\mu_n}\}$ is weakly compact, that is, there exists a subsequence $\{f_{\mu_{n_k}}\}$ such that $f_{\mu_{n_k}} \xrightarrow{w} f_{\mu}$ in $\mathbb{L}^2(\mathbb{R}^d)$, as $k \rightarrow \infty$. But this amounts to saying the $\mu_{n_k} \Rightarrow \mu$ in Prohorov metric. This, together with Lemma 5.6-(ii) and the relationship between Wasserstein distance and weak convergence (see, e.g., [65, Theorem 7.12]), leads to that $\lim_{k \rightarrow \infty} \mu_{n_k} = \mu$ in $\mathcal{P}_2(\mathbb{R}^d)$, proving (ii).

(iii) Let us assume that $\{\mu_n\} \subset \mathcal{E}$ such that $\mu_n \Rightarrow \mu$ in Prohorov metric. The stability result in Proposition 5.9 shows that $\rho^{\mu_n} = \Gamma_1(\mu_n) \rightarrow \Gamma_1(\mu) = \rho^{\mu} \in \mathcal{S}$ in $\mathbb{L}^1(\mathbb{R}^d)$. Next, we show that $\Gamma_2(\rho^{\mu_n}) \Rightarrow \Gamma_2(\rho^{\mu})$ in Prohorov metric. Recall the definition of Γ_2 , we define a family of functionals on $\mathcal{E}$: for each $k, n \in \mathbb{N}$ and $\bar{\mu} \in \mathcal{E}$,

$$\begin{cases} F_n^k(\bar{\mu}) := kG(\bar{\mu}) + \int_{\mathbb{R}^d} \log \rho^{\mu_n}(y) \bar{\mu}(dy); \\ F^k(\bar{\mu}) := kG(\bar{\mu}) + \int_{\mathbb{R}^d} \log \rho^{\mu}(y) \bar{\mu}(dy). \end{cases} \quad (6.4)$$

Then $\mu'_n = \Gamma_2(\rho^{\mu_n}) := \arg \min_{\bar{\mu} \in \mathcal{E}} F_n^k(\bar{\mu})$ and $\mu' = \Gamma_2(\rho^{\mu}) := \arg \min_{\bar{\mu} \in \mathcal{E}} F^k(\bar{\mu})$.

To show that the minimizers $\mu'_n \Rightarrow \mu'$, we shall invoke the notion of Γ -convergence (cf. [19]). To be more precise, to prove the sequence $\{F_n^k\}_{n \geq 1}$ Γ -converge to F^k as $n \rightarrow \infty$, it is sufficient to show

$$\begin{cases} \text{For every sequence } \bar{\mu}_n \Rightarrow \bar{\mu}, \text{ it holds that } F^k(\bar{\mu}) \leq \varliminf_n F_n^k(\bar{\mu}_n); \\ \text{There exists a sequence } \bar{\mu}_n \Rightarrow \bar{\mu}, \text{ such that } F^k(\bar{\mu}) \geq \varlimsup_n F_n^k(\bar{\mu}_n). \end{cases} \quad (6.5)$$

Now, note that G is convex and $\mathcal{E}$ is compact under Prohorov metric, we see that both $\{F_n^k\}$ and F^k are *coercive* (in the sense that there exists minimizing sequence in $\mathcal{E} \subset \mathcal{P}_2(\mathbb{R}^d)$). Thus, in light of the Γ -convergence result (see [19, Theorem 7.1]), in order to show $\mu'_n \Rightarrow \mu'$, it suffices to check that $\{F_n^k\}$ Γ -converges to F^k , $k \in \mathbb{N}$. To see this, let $\{\bar{\mu}_n\} \subset \mathcal{E}$ such that $\bar{\mu}_n \Rightarrow \bar{\mu}$, we have

$$\begin{aligned} & \left| \int_{\mathbb{R}^d} (\log \rho^{\mu_n}(y) \bar{\mu}_n(dy) - \int_{\mathbb{R}^d} \log \rho^{\mu}(y) \bar{\mu}(dy) \right| \\ & \leq \int_{\mathbb{R}^d} |\log \rho^{\mu_n}(y) - \log \rho^{\mu}(y)| f_{\bar{\mu}_n}(y) dy + \int_{\mathbb{R}^d} |\log \rho^{\mu}(y)| |f_{\bar{\mu}_n}(y) - f_{\bar{\mu}}(y)| dy =: I_1 + I_2, \end{aligned}$$

where I_1 and I_2 are the two integrals on the right side above. We now prove that both I_1 and I_2 converge to 0, as $n \rightarrow \infty$. To see this, first note that by (5.12) we have

$$\begin{aligned} I_1 & \leq \int_{\mathbb{R}^d} |\psi^{\mu_n}(y) - \psi^{\mu}(y)| f_{\bar{\mu}_n}(y) dy + \int_{\mathbb{R}^d} |\log f_{\mu_n}(y) - \log f_{\mu}(y)| f_{\bar{\mu}_n}(y) dy \\ & \leq W_2(\mu_n, \mu) \int_{\mathbb{R}^d} f_{\bar{\mu}_n}(y) dy + K \int_{\mathbb{R}^d} |\log f_{\mu_n}(y) - \log f_{\mu}(y)| g(y) dy =: I_{11} + I_{12}, \end{aligned}$$

where the last inequality is due to (5.13) and Assumption 5.4. Since $\mu_n \Rightarrow \mu$, again by Lemma 5.6-(ii) and [65, Theorem 7.12], we have $W_2(\mu_n, \mu) \rightarrow 0$. Hence $I_{11} \rightarrow 0$ as $n \rightarrow \infty$. Furthermore, by Assumption 5.4-(i) we see that, for $\mu \in \mathcal{E}$, one has $|\log f_\mu(y)| \leq C|y|^2$ for some constant $C > 0$, that is, $|\log f_{\mu_n}(y) - \log f_\mu(y)|g(y) \leq 2C|y|^2g(y)$ for $y \in \mathbb{R}^d$. On the other hand, applying Lemma 5.6-(iii), we have $f_{\mu_n} \rightarrow f_\mu$ in $\mathbb{L}^1$. Thus, possibly along a subsequence (may assume itself, by virtue of Heine's Lemma), it holds $\log f_{\mu_n}(y) \rightarrow \log f_\mu(y)$, a.e. $y \in \mathbb{R}^d$, thanks to the continuity of $\log$ function. Therefore, by Assumption 5.4 and Dominated Convergence Theorem, we see that $I_{12} \rightarrow 0$ as $n \rightarrow \infty$, whence $I_1 \rightarrow 0$ as $n \rightarrow \infty$. Similarly, by (5.12), (5.15), and Assumption 5.4-(i), we have

$$\begin{aligned} I_2 &\leq \int_{\mathbb{R}^d} |\psi^\mu(y)| |f_{\bar{\mu}_n}(y) - f_{\bar{\mu}}(y)| dy + \int_{\mathbb{R}^d} |\log f_\mu(y)| |f_{\bar{\mu}_n}(y) - f_{\bar{\mu}}(y)| dy \\ &\leq 2C \int_{\mathbb{R}^d} |y|^2 |f_{\bar{\mu}_n}(y) - f_{\bar{\mu}}(y)| dy \rightarrow 0, \quad \text{as } n \rightarrow \infty, \end{aligned}$$

again thanks to Assumption 5.4-(ii) and the Dominated Convergence Theorem. Finally, since $G(\cdot)$ is continuous on $\mathcal{P}_2(\mathbb{R}^d)$, by definition (6.4) we see that $F_n^k(\bar{\mu}_n) \rightarrow F^k(\bar{\mu})$ whenever $\bar{\mu}_n \Rightarrow \bar{\mu}$. Thus, by (6.5), $\{F_n^k\}$ Γ -converges to F^k , as $n \rightarrow \infty$. The proof is now complete. $\square$

Finally, we shall establish an analogue of Theorem 4.1 in the case of general $\mu_{\text{ini}} \in \mathcal{P}_2(\mathbb{R}^d)$. For technical convenience, in what follows we shall make use of the following extra assumptions to facilitate our discussion. Recall the function ξ and η defined by (5.14) and (5.16), respectively.

Assumption 6.2. (i) The penalty function G satisfies $G(\mu; \mu_{\text{tar}}) \geq W_2(\mu, \mu_{\text{tar}})$;

(ii) In Assumption 3.2-(ii), the function ϕ satisfies $\|\phi(\cdot)e^{\xi(\cdot)}\|_\infty < \infty$;

(iii) For any $R > 0$, there exists $M_R > 0$, such that

$$\mathbb{E}[\eta^2(X_T^{t,x})] = \int_{\mathbb{R}^d} \eta^2(y)p(T, y; t, x)dy \leq M_R, \quad (t, x) \in [0, T] \times B_R, \quad (6.6)$$

where $B_R := \{x \in \mathbb{R}^d : |x| \leq R\}$. $\square$

Remark 6.3. (1) Assumption 6.2-(i) is not overly restrictive, and can be justified by Example 3.5.

(2) Note that the function ϕ in Assumption 3.2-(ii) satisfies $\phi(y)e^{\lambda|y|^2} \leq C$, for some $c > 0$ and $\xi(y) \sim c|y|^2$, as $|y| \rightarrow \infty$, Assumption 6.2-(ii) amounts to saying that ϕ and ξ are compatible.

(3) While Assumption 6.2-(iii) is slightly stronger than the requirement (5.16), it would be trivial if the mapping $(t, x) \mapsto \mathbb{E}[\eta^2(X_T^{t,x})]$ is continuous, which is by no means stringent. $\square$

Our last result of this section is the following.

Theorem 6.4. Assume that the Assumptions 3.2, 5.4 and 6.2 are in force. Let $\mu_{\text{ini}} \in \mathcal{E}$ be given, and let $\hat{\alpha}(t, x)$ and $\hat{\alpha}^k(t, x)$, $(t, x) \in [0, T] \times \mathbb{R}^d$ are optimal controls given in Proposition 3.6 and Lemma 5.3, respectively. Then, for any $R > 0$, there exists $C_R > 0$, such that for any $k \in \mathbb{N}$, it holds that

$$\int_0^T |\hat{\alpha}^k(t, x) - \hat{\alpha}(t, x)| dt \leq \frac{C_R}{k}, \quad (t, x) \in [0, T] \times B_R. \quad (6.7)$$

Proof. We begin by denoting $\mu_k := \mathbb{P}_{X_T^{\hat{\alpha}^k}}$, $\mu_{\text{tar}} := \mathbb{P}_{X_T^{\hat{\alpha}}}$, and let $\rho^{\mu_k} = \Gamma_1(\mu_k)$, $k \in \mathbb{N}$, $\rho^{\mu_{\text{tar}}} = \Gamma_1(\mu_{\text{tar}})$, respectively, as we defined before. Next, applying (5.12), we have

$$\begin{aligned} |\rho^{\mu_k}(y) - \rho^{\mu_{\text{tar}}}(y)| &= |e^{\psi^{\mu_k}(y)} f_{\mu_k}(y) - e^{\psi^{\mu_{\text{tar}}}(y)} f_{\text{tar}}(y)| \\ &\leq e^{\psi^{\mu_{\text{tar}}}(y)} |f_{\mu_k}(y) - f_{\text{tar}}(y)| + f_{\mu_k}(y) |e^{\psi^{\mu_k}(y)} - e^{\psi^{\mu_{\text{tar}}}(y)}|. \end{aligned} \quad (6.8)$$

Let us now recall some facts from Section 4.1. First, note that the optimality of μ_k implies that $G(\mu_k) \leq \frac{C}{k}$ (cf. (4.3)), for some generic constant $C > 0$ independent of k , which we shall allow to vary from line to line below. Thus, by virtue of Assumption 3.2-(ii), we can write

$$|f_{\mu_k}(y) - f_{\text{tar}}(y)| \leq \frac{C}{k} \phi(y), \quad y \in \mathbb{R}^d,$$

where $\phi(y)e^{\xi(y)} \leq C$, $y \in \mathbb{R}^d$, thanks to Assumption 6.2-(ii). Furthermore, under Assumption 6.2-(i), we can assume without loss of generality that the Schrödinger potentials ψ^{μ_k} and $\psi^{\mu_{\text{tar}}}$ all satisfy estimates (5.13) and (5.14). Consequently, by Assumption 6.2-(i) and an easy application of Lemma 5.7 and Newton-Leibniz formula we have

$$\begin{aligned} |e^{\psi^{\mu_{\text{tar}}}(y)} - e^{\psi^{\mu_k}(y)}| &= |\psi^{\mu_{\text{tar}}}(y) - \psi^{\mu_k}(y)| \int_0^1 \exp\{\psi^{\mu_{\text{tar}}}(y) + \theta(\psi^{\mu_k}(y) - \psi^{\mu_{\text{tar}}}(y))\} d\theta \\ &\leq C e^{\xi(y)} W_2(\mu_k, \mu_{\text{tar}}) \leq C e^{\xi(y)} G(\mu_k) \leq \frac{C}{k} e^{\xi(y)}, \quad y \in \mathbb{R}^d. \end{aligned}$$

Summarizing above and recalling Assumption 5.4 and (5.16), we derive from (6.8) that

$$|\rho^{\mu_k}(y) - \rho^{\mu_{\text{tar}}}(y)| \leq \frac{C}{k} \left(e^{\psi^{\mu_{\text{tar}}}(y)} \phi(y) + f_{\mu_k}(y) e^{\xi(y)} \right) \leq \frac{C}{k} \left(e^{\xi(y)} \phi(y) + g^2(y) e^{\xi(y)} \right) \leq \frac{C}{k} (1 + \eta(y)),$$

and therefore, given $R > 0$, and $(t, x) \in [0, T] \times B_R$, we apply Assumption 6.2-(iii) to get

$$\mathbb{E}[|\rho^{\mu_k}(X_T^{t,x}) - \rho^{\mu_{\text{tar}}}(X_T^{t,x})|] \leq \mathbb{E}[|\rho^{\mu_k}(X_T^{t,x}) - \rho^{\mu_{\text{tar}}}(X_T^{t,x})|^2]^{\frac{1}{2}} \leq \frac{C_R}{k}, \quad (6.9)$$

where $C_R > 0$ is some constant depending on the generic constant C above and M_R in (6.6).

To complete the proof, let us recall that optimal strategies are of the form $\hat{\alpha}^k(t, x) = \nabla \log h^{\mu_k}(t, x)$, $k \in \mathbb{N}$, and $\hat{\alpha}(t, x) = \nabla \log h(t, x)$, and $h^{\mu_k}(t, x)$ and $h(t, x)$ are the solutions to the respective PDEs:

$$\begin{cases} \partial_t h^{\mu_k}(t, x) + \mathcal{L}_t h^{\mu_k}(t, x) = 0; \\ h^{\mu_k}(T, x) = \rho^{\mu_k}(x). \end{cases} \quad \begin{cases} \partial_t h(t, x) + \mathcal{L}_t h(t, x) = 0; \\ h(T, x) = \rho^{\mu_{\text{tar}}}(x), \end{cases} \quad (6.10)$$

Furthermore, noting that $h^{\mu_k}(t, x) = \mathbb{E}[\rho^{\mu_k}(X_T^{t,x})]$, $h(t, x) = \mathbb{E}[\rho^{\mu_{\text{tar}}}(X_T^{t,x})]$, and by the Bismut-Elworthy-Li formula we have

$$\nabla h^{\mu_k}(t, x) = \mathbb{E}[\rho^{\mu_k}(X_T^{t,x}) N_T^{t,x}]; \quad \nabla h(t, x) = \mathbb{E}[\rho^{\mu_{\text{tar}}}(X_T^{t,x}) N_T^{t,x}].$$

Thus, we have whenever $(t, x) \in [0, T] \times B_R$,

$$|\nabla h(t, x)| \leq K \mathbb{E}[\eta^2(X_T^{t,x})]^{\frac{1}{2}} \mathbb{E}[|N_T^{t,x}|]^{\frac{1}{2}} \leq \frac{C_R}{\sqrt{T-t}},$$

and a similar argument as in (4.6) and (4.8), together with the estimate (6.9), leads to that

$$|\nabla h^{\mu_k}(t, x) - \nabla h(t, x)| \leq \frac{C_R}{k\sqrt{T-t}}; \quad |h^{\mu_k}(t, x) - h(t, x)| \leq \frac{C_R}{k}, \quad (t, x) \in [0, T] \times B_R. \quad (6.11)$$

Finally, we note that by definition the function h is positive everywhere, and being the classical solution to the PDE (6.10) it is continuous. Thus, given $R > 0$, there exists $\delta_R > 0$, such that $h(t, x) \geq \delta_R$, for all $(t, x) \in [0, T] \times B_R$. Since (6.11) implies that h^{μ_k} converges to h uniformly on compacts, thus it must hold that $h^{\mu_k}(t, x) \geq \delta_R/2$, for $(t, x) \in [0, T] \times B_R$, and k large enough. We therefore conclude, similar to (4.4), that

$$|\hat{\alpha}^k(t, x) - \hat{\alpha}(t, x)| \leq \left| \frac{\nabla h^{\mu_k}(t, x) - \nabla h(t, x)}{h^{\mu_k}(t, x)} \right| + |\nabla h(t, x)| \left| \frac{h(t, x) - h^{\mu_k}(t, x)}{h^{\mu_k}(t, x)h(t, x)} \right| \leq \frac{C_R}{k\sqrt{T-t}},$$

as $k \rightarrow \infty$, for $(t, x) \in [0, T] \times B_R$, where C_R depends on M_R and δ_R above, but independent of k . Integrating in t we obtain (6.7). $\square$

7 Conclusion

We study the soft-constrained Schrödinger bridge problem (SCSBP) as a flexible alternative to the classical formulation for generative modeling. By replacing hard terminal constraints with a general penalty function, the SCSBP potentially offers greater flexibility and stability for generative AI tasks. Moreover, we establish linear convergence of both the value functions and the optimal controls as the penalty parameter tends to infinity, thereby providing a theoretical guarantee for the framework.

In future work, we will develop efficient algorithms for learning the SCSBP solutions and test the performance on benchmark generative AI tasks. This will allow us to translate the theoretical framework into practical tools, further demonstrating the potential of regularized stochastic control formulations for modern generative modeling.

References

- [1] B. Acciaio, S. Eckstein, and S. Hou. Time-causal vae: Robust financial time series generator. *arXiv preprint arXiv:2411.02947*, 2024.
- [2] A. Alouadi, B. Barreau, L. Carlier, and H. Pham. Robust time series generation via schrödinger bridge: a comprehensive evaluation. *arXiv preprint arXiv:2503.02943*, 2025.
- [3] D. G. Aronson. Bounds for the fundamental solution of a parabolic equation. 1967.
- [4] Y. Bai, E. Yang, B. Han, Y. Yang, J. Li, Y. Mao, G. Niu, and T. Liu. Understanding and improving early stopping for learning with noisy labels. *Advances in Neural Information Processing Systems*, 34:24392–24403, 2021.
- [5] J.-D. Benamou, G. Chazareix, and G. Loeper. From entropic transport to martingale transport, and applications to model calibration. *arXiv preprint arXiv:2406.11537*, 2024.
- [6] J. Betker, G. Goh, L. Jing, T. Brooks, J. Wang, L. Li, L. Ouyang, J. Zhuang, J. Lee, Y. Guo, et al. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3):8, 2023.
- [7] A. Beurling. An automorphism of product measures. *Annals of Mathematics*, 72(1):189–200, 1960.
- [8] J.-M. Bismut. *Large Deviation and Malliavin Calculus*, volume 45 of *Progress in Mathematics*. Birkhäuser, 1984.
- [9] H. Cao, H. Gu, X. Guo, and M. Rosenbaum. Risk of transfer learning and its applications in finance. *arXiv preprint arXiv:2311.03283*, 2023.
- [10] G. Carlier, L. Chizat, and M. Laborde. Displacement smoothness of entropic optimal transport. *ESAIM: Control, Optimisation and Calculus of Variations*, 30:25, 2024.
- [11] M. Chen, K. Huang, T. Zhao, and M. Wang. Score approximation, estimation and distribution recovery of diffusion models on low-dimensional data. In *International Conference on Machine Learning*, pages 4672–4712. PMLR, 2023.
- [12] M. Chen, R. Xu, Y. Xu, and R. Zhang. Diffusion factor models: Generating high-dimensional returns with factor structure. *arXiv preprint arXiv:2504.06566*, 2025.

- [13] T. Chen, G.-H. Liu, and E. A. Theodorou. Likelihood training of schrödinger bridge using forward-backward sdes theory. *arXiv preprint arXiv:2110.11291*, 2021.
- [14] Y. Chen, T. T. Georgiou, and M. Pavon. Stochastic control liaisons: Richard sinkhorn meets gaspard monge on a schrödinger bridge. *Siam Review*, 63(2):249–313, 2021.
- [15] Z. Chen, A. Mustafi, P. Glaser, A. Korba, A. Gretton, and B. K. Sriperumbudur. (de)-regularized maximum mean discrepancy gradient flow. *arXiv preprint arXiv:2409.14980*, 2024.
- [16] A. Chiarini, G. Conforti, G. Greco, and L. Tamanini. Gradient estimates for the schrödinger potentials: convergence to the brenier map and quantitative stability. *Communications in Partial Differential Equations*, 48(6):895–943, 2023.
- [17] S.-N. Chow, W. Li, C. Mou, and H. Zhou. Dynamical schrödinger bridge problems on graphs. *Journal of Dynamics and Differential Equations*, pages 1–20, 2022.
- [18] P. Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied mathematics and Optimization*, 23(1):313–329, 1991.
- [19] G. Dal Maso. *Introduction to Gamma-convergence*. Springer Science + Business Media, LLC, 1993.
- [20] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- [21] W. E. Deming and F. F. Stephan. On a least squares adjustment of a sampled frequency table when the expected marginal totals are known. *The Annals of Mathematical Statistics*, 11(4):427–444, 1940.
- [22] V. Divol, J. Niles-Weed, and A.-A. Pooladian. Tight stability bounds for entropic brenier maps. *International Mathematics Research Notices*, 2025(7):rnaf078, 2025.
- [23] D. A. Edwards. On the kantorovich-rubinstein theorem. *Expositiones Mathematicae*, 29:387–398, 2011.
- [24] K. D. Elworthy and X. M. Li. Formulae for the derivatives of heat semigroups. *Journal of Functional Analysis*, 125:252–286, 1994.
- [25] H. Föllmer. Random fields and diffusion processes. In *École d’Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, pages 101–203. Springer, 2006.
- [26] R. Fortet. Résolution d’un système d’équations de m. schrödinger. *Journal de mathématiques pures et appliquées*, 19(1-4):83–105, 1940.
- [27] E. Fournié, J.-M. Lasry, J. Lebuchoux, P.-L. Lions, and N. Touzi. Applications of malliavin calculus to monte carlo methods in finance. *Finance and Stochastics*, 3:391–412, 1999.
- [28] H. Fu, Z. Yang, M. Wang, and M. Chen. Unveil conditional diffusion models with classifier-free guidance: A sharp statistical theory. *arXiv preprint arXiv:2403.11968*, 2024.
- [29] J. Garg, X. Zhang, and Q. Zhou. Soft-constrained schrödinger bridge: a stochastic control approach. In *International Conference on Artificial Intelligence and Statistics*, pages 4429–4437. PMLR, 2024.

- [30] X. Gu, C. Du, T. Pang, C. Li, M. Lin, and Y. Wang. On memorization in diffusion models. *arXiv preprint arXiv:2310.02664*, 2023.
- [31] F. F. Gunsilius. On the convergence rate of potentials of brenier maps. *Econometric Theory*, 38(2):381–417, 2022.
- [32] M. Hamdouche, P. Henry-Labordere, and H. Pham. Generative modeling for time series via schrödinger bridge. *arXiv preprint arXiv:2304.05093*, 2023.
- [33] Y. Han, M. Razaviyayn, and R. Xu. Neural network-based score estimation in diffusion models: Optimization and generalization. *arXiv preprint arXiv:2401.15604*, 2024.
- [34] Y. Han, M. Razaviyayn, and R. Xu. Stochastic control for fine-tuning diffusion models: Optimality, regularity, and convergence. *arXiv preprint arXiv:2412.18164*, 2024.
- [35] C. Hernández and L. Tangpi. Propagation of chaos for mean field schrödinger problems. *arXiv preprint arXiv:2304.09340*, 2023.
- [36] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. In *Neurips*, volume 33, pages 6840–6851, 2020.
- [37] B. Jamison. The markov processes of schrödinger. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 32(4):323–331, 1975.
- [38] L. Kong, H. Wang, T. Wang, G. Xiong, and M. Tambe. Composite flow matching for reinforcement learning with shifted-dynamics data. *arXiv preprint arXiv:2505.23062*, 2025.
- [39] S. Kullback and R. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, 22(1):79–86, 1951.
- [40] C. Léonard. A survey of the schrödinger problem and some of its connections with optimal transport. *arXiv preprint arXiv:1308.0215*, 2013.
- [41] P. Li, Z. Li, H. Zhang, and J. Bian. On the generalization properties of diffusion models. *Advances in Neural Information Processing Systems*, 36:2097–2127, 2023.
- [42] G. Loeper, J. Obloj, and B. Joseph. Calibration of local volatility models with stochastic interest rates using optimal transport. *arXiv preprint arXiv:2305.00200*, 2023.
- [43] J. Ma and J. Zhang. Representation theorems for backward stochastic differential equations. *Annals of Applied Probability*, 12(4):1390–1418, 2002.
- [44] E. F. Montesuma, F. M. N. Mboula, and A. Souloumiac. Recent advances in optimal transport for machine learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [45] T. Moon, M. Choi, G. Lee, J.-W. Ha, and J. Lee. Fine-tuning diffusion models with limited data. In *NeurIPS 2022 Workshop on Score-Based Methods*, 2022.
- [46] M. Nutz and J. Wiesel. Entropic optimal transport: Convergence of potentials. *Probability Theory and Related Fields*, 184(1):401–424, 2022.
- [47] OpenAI. Sora: Creating video from text. <https://openai.com/sora>, 2024.
- [48] Y. Ouyang, L. Xie, H. Zha, and G. Cheng. Transfer learning for diffusion models. *Advances in Neural Information Processing Systems*, 37:136962–136989, 2024.

- [49] S. Peluchetti. Diffusion bridge mixture transports, schrödinger bridge problems and generative modeling. *Journal of Machine Learning Research*, 24(374):1–51, 2023.
- [50] P. E. Protter. *Stochastic Integration and Differential Equations*. Springer-Verlag, Heidelberg, second edition, 2005.
- [51] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [52] L. Richter and J. Berner. Improved sampling via learned diffusions. *arXiv preprint arXiv:2307.01198*, 2023.
- [53] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [54] Y. Shi, V. De Bortoli, A. Campbell, and A. Doucet. Diffusion schrödinger bridge matching. *Advances in Neural Information Processing Systems*, 36, 2024.
- [55] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [56] K.-U. Song. Applying regularized schrödinger-bridge-based stochastic process in generative modeling. *arXiv preprint arXiv:2208.07131*, 2022.
- [57] Y. Song and S. Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [58] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020.
- [59] T. Sweeting. On a converse to scheffé’s theorem. *The Annals of Statistics*, 14(3):1252–1256, 1986.
- [60] W. Tang. Fine-tuning of diffusion models via stochastic control: entropy regularization and beyond. *arXiv preprint arXiv:2403.06279*, 2024.
- [61] L. Torrey and J. Shavlik. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI Global Scientific Publishing, 2010.
- [62] M. Uehara, Y. Zhao, K. Black, E. Hajiramezanali, G. Scalia, N. L. Diamant, A. M. Tseng, T. Biancalani, and S. Levine. Fine-tuning of continuous-time diffusion models as entropy-regularized control. *arXiv preprint arXiv:2402.15194*, 2024.
- [63] F. Vargas, S. Padhy, D. Blessing, and N. Nüsken. Transport meets variational inference: Controlled monte carlo diffusions. *arXiv preprint arXiv:2307.01050*, 2023.
- [64] F. Vargas, P. Thodoroff, A. Lamacraft, and N. Lawrence. Solving schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.

- [65] C. Villani. *Topics in Optimal Transportation*. Graduate Studies in Mathematics, 58. AMS, Providence, RI, 2003.
- [66] G. Wang, Y. Jiao, Q. Xu, Y. Wang, and C. Yang. Deep generative learning via schrödinger bridge. In *International conference on machine learning*, pages 10794–10804. PMLR, 2021.
- [67] L. Winkler, C. Ojeda, and M. Opper. A score-based approach for training schrödinger bridges for data modelling. *Entropy*, 25(2):316, 2023.
- [68] T. Xu, L. K. Wenliang, M. Munn, and B. Acciaio. Cot-gan: Generating sequential data via causal optimal transport. *Advances in neural information processing systems*, 33:8798–8809, 2020.
- [69] J. Zhang. *Backward stochastic differential equations*. Springer, 2017.
- [70] H. Zhao, H. Chen, Y. Guo, G. I. Winata, T. Ou, Z. Huang, D. D. Yao, and W. Tang. Fine-tuning diffusion generative models via rich preference optimization. *arXiv preprint arXiv:2503.11720*, 2025.