

Perspectives on Stochastic Localization

Bobby Shi* Kevin Tian† Matthew S. Zhang‡

Abstract

We survey different perspectives on the *stochastic localization* process of [Eld13], a powerful construction that has had many exciting recent applications in high-dimensional probability and algorithm design. Unlike prior surveys on this topic, our focus is on giving a self-contained presentation of all known alternative constructions of Eldan’s stochastic localization, with an emphasis on connections between different constructions. Our hope is that by collecting these perspectives, some of which had primarily arisen within a particular community (e.g., probability theory, theoretical computer science, information theory, or machine learning), we can broaden the accessibility of stochastic localization, and ease its future use.

Contents

1	Introduction	2
1.1	Notation	3
1.2	Preliminaries	4
2	Measure-valued process	5
3	Posterior estimation	6
4	Diffusion models	7
5	Renormalization	9
6	Static Schrödinger bridge	12
7	Dynamic Schrödinger bridge	17
8	Entropic optimal transport	18
A	Anisotropic stochastic localization	22
B	Restricted Gaussian dynamics	23
B.1	The perspective from entropic stability	23
B.1.1	Localization schemes	23
B.1.2	Entropy conservation from entropic stability	24
B.1.3	Log-Sobolev inequality from entropy conservation	27
B.2	The perspective from renormalization	28
B.3	The perspective from time reversal	31

*University of Texas at Austin, bhshi@utexas.edu

†University of Texas at Austin, kjtian@cs.utexas.edu

‡University of Toronto, matthew.zhang@mail.utoronto.ca

1 Introduction

Stochastic localization, an elegant stochastic process introduced by [Eld13], has proven to be a particularly useful tool in the analysis of high-dimensional distributions, and in designing algorithms to sample from them. Use of this process has notably resulted in a series of gradual improvements to the estimate of the KLS constant [Eld13; LV17; Che21; KL22; JLV22; Kla23], i.e., the smallest isoperimetric constant among all isotropic log-concave densities on $\mathbb{R}^d$ [KLS95]. It has also enabled various other breakthroughs in probability theory and theoretical computer science [EL14; Eld18; Kla18; LV18; Eld20; EMZ20; KP21; ES22; EKZ22; AMS22; AHLVXY23; Gua24; KL24b].

Among the many useful properties of stochastic localization, the most pertinent to us is the realization that the process is in fact *equivalent* to various other constructions that have arisen from probability theory, theoretical computer science, information theory, and machine learning. This is not a new observation; many results have connected two or more equivalent forms of stochastic localization to obtain novel characterizations or tighter estimates for at least one of the forms [Leh13; EL15; LST21; CCSW22; EM22; CE22; BBDD24; MS24; KO25].

This survey’s goal is to present a relatively complete collection of different perspectives on stochastic localization, each equating the base process (cf. Perspective 1) with another natural probabilistic object. While the proofs presented here are not new (except for providing some missing details), we believe this exposition benefits the community for at least the following reasons.

- (1) Although prior expositions of stochastic localization [KP21; Eld22; Mon23; KL24a; Che25] have presented several forms of the process, none make it a goal to connect all existing perspectives. Our presentation includes all perspectives contained in these prior works, and our explicit focus is a streamlined presentation of their connections.
- (2) Several perspectives may be more familiar to a subcommunity interacting with stochastic localization in a particular way. Our presentation simplifies translation between different perspectives by clarifying the connections, which may enable further applications.
- (3) We make an effort to review the relevant background and keep our exposition self-contained, without being overly verbose. While some derivations are folklore to subcommunities where such calculations are routine, we believe there is value in providing explicit derivations for unfamiliar readers, particularly in an introductory survey.

At this juncture, it is useful to introduce our first perspective on stochastic localization, based on its definition in [LV17], a small modification of its original definition in [Eld13].¹

Perspective 1. Let $\pi_0 \in \mathcal{P}(\mathbb{R}^d)$ be such that $\mathbf{m}_0 := \mathbb{E}_{\mathbf{x} \sim \pi_0}[\mathbf{x}]$ exists. Define a stochastic process $\{\mathbf{c}_t \in \mathbb{R}^d\}_{t \geq 0}$ as follows, where $\{\mathbf{W}_t \in \mathbb{R}^d\}_{t \geq 0}$ is a Wiener process:

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{0}_d, \quad d\mathbf{c}_t = \mathbf{m}_t dt + d\mathbf{W}_t, \\ \text{where } \mathbf{m}_t &:= \mathbb{E}_{\mathbf{x} \sim \pi_t}[\mathbf{x}], \quad \pi_t(\mathbf{x}) \propto \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2\right) \pi_0(\mathbf{x}), \text{ for all } t \geq 0. \end{aligned} \tag{SL-I}$$

We call the (random) induced measures $\{\pi_t\}_{t \geq 0}$ the stochastic localization of π_0 .

In other words, π_t is an *exponential tilt* of the “Gaussian-regularized” measure $\propto \exp(-\frac{t}{2} \|\mathbf{x}\|_2^2) \pi_0(\mathbf{x})$, by the log-linear function $\exp(\langle \mathbf{c}_t, \mathbf{x} \rangle)$. The dynamics of the random tilt $\mathbf{c}_t$ are governed by (SL-I), which includes a bias towards the mean $\mathbf{m}_t$ of the current measure π_t .

Why is Perspective 1 useful? For one thing, it replaces π_0 with a distribution over measures π_t , each of which regularizes π_0 by a randomly-shifted Gaussian. These Gaussians always have variance $\frac{1}{t}$ at time t , and hence as $t \rightarrow \infty$, the measure π_t obtains strong concentration properties, potentially much stronger than those of π_0 . For this reason, (SL-I) is termed a *localization process* (Definition 3, [CE22]), in that π_t successively “localizes” towards a Dirac delta distribution.

¹More generally, the stochastic localization process of [Eld13] allows specifying *control matrices* in the dynamics, which affect the covariance of the Gaussian regularization in π_t . We recall this more general form in Appendix A, focusing on the isotropic variant throughout the main body for simplicity.

In addition to always inducing a fixed amount of Gaussian regularization by time t , (SL-I) satisfies a second important property: it is a *measure-valued martingale*, i.e., $\mathbb{E}[\pi_t(\mathbf{x})] = \pi_0(\mathbf{x})$ pointwise over $\mathbf{x} \in \mathbb{R}^d$, and for all $t \geq 0$. These two facts let us view the process $\pi_0 \rightarrow \pi_t$ as a random decomposition of π_0 into better-behaved components, whose properties are inherited by π_0 “on average.” Many applications of stochastic localization build on precisely this perspective.

We now outline the remaining perspectives found in this survey.

- (1) **Section 2** formalizes the martingale property of (SL-I), using a dual perspective (observed in [Eld13]) of (SL-I) as a stochastic process on the measure π_t itself, rather than the tilt $\mathbf{c}_t$.
- (2) **Section 3** presents an alternative information-theoretic perspective (due to [EM22]) on the tilt dynamics (SL-I) as *posterior sampling* from a noisy Gaussian channel.
- (3) **Section 4** presents a perspective that equates a time-changed variant of (SL-I) with *denoising diffusion probabilistic models* [HJA20; SSKKEP21], a dominant paradigm in modern diffusion modeling. This connection was observed by [KP21; Mon23].
- (4) **Section 5** presents a perspective that uses a continuous renormalization procedure — Gaussian integration and the Polchinski equation — to decompose a measure in order to derive functional inequalities; [BBD24] connects this explicitly to stochastic localization.
- (5) **Section 6** presents a perspective via the classical (static) Schrödinger bridge problem, which gives another derivation of the Polchinski semigroup in Section 5.
- (6) **Section 7** and **Section 8** present two alternative perspectives on the static Schrödinger bridge problem, respectively connecting it to a dynamic reformulation based on Girsanov’s theorem, and an entropy-regularized optimal transport problem.

Each section begins by introducing a new probabilistic object of self-contained interest, along with the relevant background needed to understand it. It then shows how the newly-introduced object is equivalent to the process in Perspective 1, after appropriate transformations.

Finally, while this survey focuses on the stochastic localization process itself, a related topic of interest is *algorithmic applications* of stochastic localization. There are natural sampling algorithms derived from stochastic localization, several of which have enabled improved runtimes for important applications in theoretical computer science, statistical physics, and machine learning [LV18; LST21; CCSW22; CE22; EKZ22; AMS22; AHLVXY23; BBDD24]. To this end, in Appendix B, we provide an exposition on how applying the *localization schemes* framework of [CE22] to the process (SL-I) yields the *restricted Gaussian dynamics* Markov chain [LST21]. We also show how techniques developed in the main body provide a mixing time estimate for this Markov chain.

Remark 1 (On solutions to (SL-I)). *A sufficient condition for the existence and uniqueness of a solution to the process $\mathbf{c}_t$ is uniform Lipschitz-continuity of $\mathbf{m}_t$. A straightforward calculation (i.e., (44) under time reparameterization) shows $\nabla_{\mathbf{c}_t} \mathbf{m}_t = \nabla_{\mathbf{c}_t} \mathbb{E}_{\mathbf{x} \sim \pi_t}[\mathbf{x}] = \mathbf{Cov}(\pi_t)$; thus, it suffices that $\|\mathbf{Cov}(\pi_t)\|_{\text{op}}$ is uniformly bounded. If π_t is compactly supported then this is clear. More generally, if π_t satisfies a Poincaré inequality (37), i.e.,*

$$\text{Var}_{\pi_t}[f] \leq \frac{1}{\alpha} \mathbb{E}_{\pi_t}[\|\nabla f\|_2^2]$$

for suitable functions f , then plugging in linear functions $f(\mathbf{x}) = \langle \mathbf{v}, \mathbf{x} \rangle$ we deduce that $\|\mathbf{Cov}(\pi_t)\|_{\text{op}} \leq \frac{1}{\alpha}$. If π itself is α -strongly log-concave, then the arguments of Section B.3 show that for all finite t , π_t satisfies a (weakly) improved Poincaré inequality. We caution that for more general usage, solution existence and uniqueness for (SL-I) do not necessarily hold, so the reader should carefully check in the case of their application before applying results in this survey.

1.1 Notation

For $n \in \mathbb{N}$ we denote $[n] := \{i \in \mathbb{N} \mid i \leq n\}$. Vectors are denoted in upright boldface lowercase, and matrices are denoted in upright boldface uppercase, unless specified otherwise. We reserve use of italic boldface uppercase letters, e.g., $\{\mathbf{W}_t\}_{t \geq 0}$, $\{\mathbf{B}_t\}_{t \geq 0}$ for Wiener processes in $\mathbb{R}^d$. All probability measures are given as densities relative to the Lebesgue measure over $\mathbb{R}^d$, and all integrals are over $\mathbb{R}^d$ unless specified otherwise. We use $\mathcal{P}(\mathbb{R}^d)$ to denote the set of probability measures over $\mathbb{R}^d$.

When $f : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is integrable, we use $\pi \propto f$ to mean that π is the measure that equals f up to a normalization constant $(\int f(\mathbf{x})d\mathbf{x})^{-1}$. We use $\mathcal{N}(\mathbf{m}, \mathbf{\Sigma})$ to denote the multivariate normal distribution with specified mean and covariance, $\mathbf{m} \in \mathbb{R}^d$ and $\mathbf{\Sigma} \in \mathbb{R}^{d \times d}$. We use $\mathbf{I}_d$ to denote the identity matrix in $\mathbb{R}^d$, and $\mathbf{0}_d$ to denote the all-zeroes vector in $\mathbb{R}^d$. We use $\nabla, \nabla^2, \nabla \cdot, \Delta$ to denote the gradient, Hessian, divergence, and Laplacian respectively. When f is a function depending on a variable t , we often use $\partial_t f$ as a shorthand for $\frac{\partial f}{\partial t}$.

In Sections 3, 6, and 8, we require tools specialized to *path measures*, i.e., probability measures supported on continuous paths on $\mathbb{R}^d$ indexed by a time in $[0, t]$. We denote this support by $\mathcal{C}([0, t] \times \mathbb{R}^d)$, and the space of associated path measures by $\mathcal{P}([0, t] \times \mathbb{R}^d)$. We always denote path measures with capital letters, e.g., $P \in \mathcal{P}([0, t] \times \mathbb{R}^d)$ is supported on continuous paths $\mathbf{p}_{[0, t]} \in \mathcal{C}([0, t] \times \mathbb{R}^d)$. Hence, $\mathbf{p}_s \in \mathbb{R}^d$ for any $s \in [0, t]$, and P_s denotes the law of $\mathbf{p}_s$.

1.2 Preliminaries

We repeatedly use the following two standard facts in stochastic calculus.

Lemma 1 (Itô's lemma). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be twice-differentiable, and suppose $\{\mathbf{x}_t\}_{t \geq 0}$ follows the SDE $d\mathbf{x}_t = \mathbf{m}_t dt + \mathbf{\Sigma}_t d\mathbf{W}_t$ where $\{\mathbf{m}_t, \mathbf{\Sigma}_t\}_{t \geq 0}$ are adapted to the filtration generated by a Wiener process $\{\mathbf{W}_t\}_{t \geq 0}$. Then $\{f(\mathbf{x}_t)\}_{t \geq 0}$ is also a drift-diffusion process, following the SDE*

$$\begin{aligned} df(\mathbf{x}_t) &= \langle \nabla f(\mathbf{x}_t), d\mathbf{x}_t \rangle + \frac{1}{2} \langle \nabla^2 f(\mathbf{x}_t), \mathbf{\Sigma}_t \mathbf{\Sigma}_t^\top \rangle dt \\ &= \left(\langle \nabla f(\mathbf{x}_t), \mathbf{m}_t \rangle + \frac{1}{2} \langle \nabla^2 f(\mathbf{x}_t), \mathbf{\Sigma}_t \mathbf{\Sigma}_t^\top \rangle \right) dt + \langle \nabla f(\mathbf{x}_t), \mathbf{\Sigma}_t d\mathbf{W}_t \rangle. \end{aligned}$$

Lemma 2 (Fokker-Planck equation). *Suppose $\{\mathbf{x}_t\}_{t \geq 0}$ follows the SDE $d\mathbf{x}_t = \mathbf{m}_t(\mathbf{x}_t) dt + \mathbf{\Sigma}_t(\mathbf{x}_t) d\mathbf{W}_t$ where $\{\mathbf{m}_t, \mathbf{\Sigma}_t\}_{t \geq 0}$ are adapted to the filtration generated by a Wiener process $\{\mathbf{W}_t\}_{t \geq 0}$. Then letting π_t denote the law of $\mathbf{x}_t$, we have for all $\mathbf{x} \in \mathbb{R}^d$ that*

$$\partial_t \pi_t(\mathbf{x}) = -\nabla \cdot (\mathbf{m}_t(\mathbf{x}) \pi_t(\mathbf{x})) + \frac{1}{2} \sum_{(i,j) \in [d] \times [d]} \partial_{\mathbf{x}_i} \partial_{\mathbf{x}_j} (\mathbf{\Sigma}_t(\mathbf{x}) \mathbf{\Sigma}_t(\mathbf{x})^\top \pi_t(\mathbf{x}))_{ij}.$$

In Sections 5 and 6, and Appendix B, we use tools from Markov semigroup theory. A (time-inhomogeneous) Markov semigroup $\{P_{\sigma, \tau}\}_{0 \leq \sigma \leq \tau}$ is induced by a stochastic process $\{\mathbf{x}_\tau\}_{\tau \geq 0}$ via the following definition: for a compactly supported test function $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we let

$$P_{\sigma, \tau} f(\mathbf{x}) := \mathbb{E}[f(\mathbf{x}_\tau) \mid \mathbf{x}_\sigma = \mathbf{x}].$$

Note that by iterating expectations, the semigroup property $P_{\rho, \sigma} P_{\sigma, \tau} = P_{\rho, \tau}$ holds.²

The adjoint operator $P_{\sigma, \tau}^*$ is then interpreted in duality as $(P_{\sigma, \tau}^* \delta_{\mathbf{x}})(f) = P_{\sigma, \tau} f(\mathbf{x})$. This means that applying $P_{\sigma, \tau}^*$ to a density over $\mathbf{x}$ at time σ advances the density to that of $\mathbf{x}$ at time τ .

We denote the *infinitesimal generator* of the semigroup at time τ by L_τ , which operates as:

$$L_\tau f := \lim_{\eta \searrow 0} \frac{P_{\tau, \tau+\eta} f - f}{\eta}. \quad (1)$$

Finally, we define the *carré du champ* operator at time τ as:

$$\begin{aligned} \Gamma_\tau(f, g)(\mathbf{x}) &:= \frac{1}{2} \lim_{\eta \searrow 0} \frac{\mathbb{E}[(f(\mathbf{x}_{\tau+\eta}) - f(\mathbf{x}_\tau))(g(\mathbf{x}_{\tau+\eta}) - g(\mathbf{x}_\tau)) \mid \mathbf{x}_\tau = \mathbf{x}]}{\eta} \\ &= \frac{1}{2} (L_\tau(fg)(\mathbf{x}) - f(\mathbf{x})L_\tau g(\mathbf{x}) - g(\mathbf{x})L_\tau f(\mathbf{x})). \end{aligned} \quad (2)$$

² $\{P_{\sigma, \tau}\}_{0 \leq \sigma \leq \tau \leq 1}$ is not a semigroup in the standard mathematical sense, as elements can only be composed if they share an index. We follow the terminology used by the literature on time-inhomogeneous stochastic processes.

2 Measure-valued process

We next present a dual perspective on the process (SL-I): rather than track how the random tilt $\mathbf{c}_t$ evolves, we track the dynamics of its induced distribution $\pi_t(\mathbf{x}) \propto \exp(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2) \pi_0(\mathbf{x})$.

Perspective 2. Let $\pi_0 \in \mathcal{P}(\mathbb{R}^d)$ be such that $\mathbf{m}_0 := \mathbb{E}_{\mathbf{x} \sim \pi_0}[\mathbf{x}]$ exists. Define a measure-valued stochastic process $\{\pi_t\}_{t \geq 0}$ as follows, where $\{\mathbf{W}_t\}_{t \geq 0}$ is a Wiener process:

$$d\pi_t(\mathbf{x}) = \langle \mathbf{x} - \mathbf{m}_t, d\mathbf{W}_t \rangle \pi_t(\mathbf{x}), \text{ pointwise over } \mathbf{x} \in \mathbb{R}^d, \text{ where } \mathbf{m}_t := \mathbb{E}_{\mathbf{x} \sim \pi_t}[\mathbf{x}]. \quad (\text{SL-II})$$

We first make a simple observation regarding correctness of the dynamics (SL-II).

Lemma 3. In (SL-II), $\int \pi_t(\mathbf{x}) d\mathbf{x} = 1$ for all $t \geq 0$, and $\mathbb{E}[\pi_t(\mathbf{x})] = \pi_0(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^d$ and $t \geq 0$.

Proof. The first part asks to show that

$$d\left(\int \pi_t(\mathbf{x}) d\mathbf{x}\right) = 0,$$

from which we can conclude $\int \pi_t(\mathbf{x}) d\mathbf{x} = 1$ for all $t \geq 0$. To see this, we apply (SL-II):

$$d\left(\int \pi_t(\mathbf{x}) d\mathbf{x}\right) = \int \langle \mathbf{x} - \mathbf{m}_t, d\mathbf{W}_t \rangle \pi_t(\mathbf{x}) d\mathbf{x} = \langle \mathbf{m}_t - \mathbf{m}_t, d\mathbf{W}_t \rangle = 0.$$

The second part is immediate because $\pi_t(\mathbf{x})$ is a martingale in (SL-II) for all $\mathbf{x} \in \mathbb{R}^d$. $\square$

Thus, π_t always remains a valid probability measure, regardless of the realization of $\{\mathbf{W}_t\}_{t \geq 0}$. We now state the main result of this section: the following equivalence, shown by [Eld13].

Theorem 1. The dynamics of π_t given by (SL-I) and (SL-II) are the same.

Proof. Let us start with the dynamics (SL-II). By Itô's lemma (Lemma 1),

$$d \log \pi_t(\mathbf{x}) = \langle \mathbf{x} - \mathbf{m}_t, d\mathbf{W}_t \rangle - \frac{1}{2} \|\mathbf{x} - \mathbf{m}_t\|_2^2 dt = \langle \mathbf{x}, d\mathbf{W}_t + \mathbf{m}_t dt \rangle - \frac{1}{2} \|\mathbf{x}\|_2^2 dt + C_t, \quad (3)$$

where $C_t = -\langle \mathbf{m}_t, d\mathbf{W}_t \rangle - \frac{1}{2} \|\mathbf{m}_t\|_2^2 dt$ is independent of $\mathbf{x}$. We know from Lemma 3 that C_t must be chosen so that π_t stays a density, so by integrating (3), we conclude that (SL-I) holds:

$$\pi_t(\mathbf{x}) \propto \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2\right) \pi_0(\mathbf{x}), \quad d\mathbf{c}_t = d\mathbf{W}_t + \mathbf{m}_t dt.$$

$\square$

Combining Lemma 3 and Theorem 1 shows our earlier claim from Section 1: that (SL-I) induces a measure-valued martingale decomposing π_0 into randomly-tilted measures π_t .

It is an instructive exercise to reverse Theorem 1, i.e., start from the dynamics (SL-I) and derive the equivalence to (SL-II). This follows by rewriting $d \log \pi_t(\mathbf{x}) = \langle d\mathbf{c}_t, \mathbf{x} \rangle - \frac{1}{2} \|\mathbf{x}\|_2^2 - d \log(Z_t)$, where $Z_t := \int \exp(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2) \pi_0(\mathbf{x}) d\mathbf{x}$, and using Itô's lemma to show

$$d \log(Z_t) = \langle d\mathbf{c}_t, \mathbf{m}_t \rangle - \frac{1}{2} \|\mathbf{m}_t\|_2^2 dt.$$

This gives the same form of $d \log \pi_t(\mathbf{x})$ as computed earlier in (3), and indeed, applying Itô's lemma once more completes the derivation of $d\pi_t(\mathbf{x})$ as in Perspective 2. We provide this calculation in Appendix A, for more general anisotropic stochastic localization processes.

3 Posterior estimation

In this section, we present an information-theoretic perspective on stochastic localization due to [EM22]. This perspective recasts the tilt $\mathbf{c}_t$ as a noisy observation from a Gaussian channel.

Perspective 3. Let $\pi_0 \in \mathcal{P}(\mathbb{R}^d)$ be such that $\mathbf{m}_0 := \mathbb{E}_{\mathbf{x} \sim \pi_0}[\mathbf{x}]$ exists, and let $\mathbf{x} \sim \pi_0$. Define noisy observations $\{\mathbf{c}_t := t\mathbf{x} + \mathbf{B}_t\}_{t \geq 0}$, where $\{\mathbf{B}_t\}_{t \geq 0}$ is a Wiener process.

Conditioned on $\mathbf{x} \sim \pi_0$, the output of the noisy channel $\mathbf{c}_t$ follows $\mathbf{c}_t \sim \mathcal{N}(t\mathbf{x}, t\mathbf{I}_d)$. This implies that the information content of $\mathbf{c}_t$ is increasing, in the sense that $\frac{1}{t}\mathbf{c}_t = \mathcal{N}(\mathbf{x}, \frac{1}{t}\mathbf{I}_d)$ becomes a better estimate of $\mathbf{x}$ as $t \rightarrow \infty$, which again captures the “localization” behavior of the process.

Before introducing the main result of this section, we require two helper facts. The first is a reparameterization of a stochastic process in terms of conditional means.

Lemma 4 ([LS01, Theorem 7.12]). Let $d\mathbf{c}_t = \mathbf{v}_t dt + d\mathbf{B}_t$, $\mathbf{c}_0 = \mathbf{0}_d$, where $\{\mathbf{B}_t\}_{t \geq 0}$ is a Wiener process. Let $\mathbf{m}_t := \mathbb{E}[\mathbf{v}_t \mid \mathcal{F}_t]$ for all $t \geq 0$, where $\{\mathcal{F}_t\}_{t \geq 0}$ is the filtration generated by $\{\mathbf{c}_t\}_{t \geq 0}$. Then we also have that $d\mathbf{c}_t = \mathbf{m}_t dt + d\mathbf{W}_t$ where $\{\mathbf{W}_t\}_{t \geq 0}$ is a Wiener process.

We omit the proof of Lemma 4, but mention here that it follows by using Itô’s lemma to compute the characteristic function of $\mathbf{c}_t - \int_0^t \mathbf{m}_s ds$, which is the same as that of a Wiener process.

The second helper fact is the Cameron-Martin theorem (a specialization of the more general Girsanov’s theorem, cf. Lemma 16), derived by similar means as those used to show Lemma 4. It characterizes how adding a drift $\mathbf{v}_s ds$ changes the *path measure* of a Wiener process.

Lemma 5 (Cameron-Martin theorem; [Gal16, Theorem 5.24]). Let $\{\mathbf{h}_s\}_{s \in [0, t]}$ follow the ODE $\mathbf{h}_0 = \mathbf{0}_d$, $d\mathbf{h}_s = \mathbf{v}_s ds$, where $\int_0^t \|\mathbf{v}_s\|_2^2 ds < \infty$, and let $\mathbf{c}_s = \mathbf{h}_s + \mathbf{W}_s$ for all $s \in [0, t]$, where $\{\mathbf{W}_s\}_{s \geq 0}$ is a Wiener process.

Let $\mathbf{p}_{[0, t]} \in \mathcal{C}([0, t])$ index a continuous path $\{\mathbf{p}_s\}_{s \in [0, t]}$ in $\mathbb{R}^d$, let W denote the Wiener measure of $\mathbf{p}_{[0, t]} = \mathbf{W}_{[0, t]}$ (i.e., the density of $\mathbf{W}_s = \mathbf{p}_s$ for all $s \in [0, t]$), and let P denote the path measure of $\mathbf{p}_{[0, t]} = \mathbf{c}_{[0, t]}$. Then for all $\mathbf{p}_{[0, t]} \in \mathcal{C}([0, t])$,

$$\frac{P(\mathbf{p}_{[0, t]})}{W(\mathbf{p}_{[0, t]})} = \exp\left(\int_0^t \langle \mathbf{v}_s, d\mathbf{p}_s \rangle - \frac{1}{2} \int_0^t \|\mathbf{v}_s\|_2^2 ds\right).$$

We are now ready to state the equivalence between Perspectives 1 and 3, due to [EM22].

Theorem 2. The dynamics of $\mathbf{c}_t$ given by (SL-I) and in Perspective 3 are the same. Moreover, the induced measure π_t in (SL-I) is the same as the posterior distribution $\pi_0(\mathbf{x} \mid \mathbf{c}_t)$ in Perspective 3.

Proof. To see the second claim, regardless of the realization of $\mathbf{c}_t$, we have that

$$\pi_0(\mathbf{x} \mid \mathbf{c}_t) \propto \exp\left(-\frac{1}{2t} \|t\mathbf{x} - \mathbf{c}_t\|_2^2\right) \pi_0(\mathbf{x}) \propto \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2\right) \pi_0(\mathbf{x}),$$

where we used Bayes’ theorem in the above derivation. This agrees with the induced π_t in (SL-I).

Next, we claim that the distribution of $\mathbf{x} \mid \mathcal{F}_t$ is the same as $\mathbf{x} \mid \mathbf{c}_t$, i.e., specifying just the endpoint $\mathbf{c}_t$ gives as much information about $\mathbf{x}$ as the entire path $\{\mathbf{c}_s\}_{0 \leq s \leq t}$. Equivalently, $\mathbf{c}_t$ is a sufficient statistic for $\mathbf{x}$ with respect to $\mathcal{F}_t$. This follows from Lemma 5 with $\mathbf{v}_s = \mathbf{x}$ for all $s \in [0, t]$, which shows that the joint measure for $\mathbf{x} \sim \pi_0$ and the induced path $\mathbf{c}_{[0, t]} \in \mathcal{C}([0, t])$ is

$$\propto \exp\left(\int_0^t \langle \mathbf{x}, d\mathbf{c}_s \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2\right) \pi_0(\mathbf{x}) W(\mathbf{c}_{[0, t]}) \propto \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2} \|\mathbf{x}\|_2^2\right) \pi_0(\mathbf{x}) W(\mathbf{c}_{[0, t]}).$$

Note that $\mathbf{x}$ only interacts with $\mathbf{c}_t$ in the above factorization, so sufficiency of $\mathbf{c}_t$ for $\mathbf{x}$ follows from the Fisher–Neyman factorization theorem. We conclude by proving the first claim. Starting from $\{\mathbf{c}_t\}_{t \geq 0}$ in Perspective 3, we apply Lemma 4 with the substitution $\mathbf{v}_t \leftarrow \mathbf{x}$. This shows that

$$d\mathbf{c}_t = \mathbf{m}_t dt + d\mathbf{W}_t \text{ where } \mathbf{m}_t := \mathbb{E}[\mathbf{x} \mid \mathcal{F}_t] = \mathbb{E}[\mathbf{x} \mid \mathbf{c}_t] = \mathbb{E}_{\pi_t}[\mathbf{x}]. \quad (4)$$

□

4 Diffusion models

In this section, we relate stochastic localization to a framework known as *denoising diffusion probabilistic models* (DDPMs), which was popularized by [HJA20; SSKKEP21] and is widely used in the practice of generative modeling. The goal of diffusion models in the context of generative modeling is to sample from a distribution π , given a dataset of i.i.d. draws from it.

The DDPM framework in particular is based on a time-changed backwards Ornstein–Uhlenbeck (OU) process. Recall that the OU process is the Langevin dynamics with a standard Gaussian as its stationary measure, i.e., it follows the following SDE from $\mathbf{x}_0 \in \mathbb{R}^d$:

$$d\mathbf{x}_t = -\mathbf{x}_t dt + \sqrt{2} d\mathbf{B}_t, \text{ where } \{\mathbf{B}_t\}_{t \geq 0} \text{ is a Wiener process.} \quad (5)$$

We make the following simple observation.

Lemma 6. *The distribution of $\mathbf{x}_t \mid \mathbf{x}_0$ in (5) is $\mathcal{N}(\exp(-t)\mathbf{x}_0, (1 - \exp(-2t))\mathbf{I}_d)$.*

Proof. Solving the SDE in (5) as an Itô integral, we have

$$\mathbf{x}_t = \exp(-t)\mathbf{x}_0 + \sqrt{2} \int_0^t \exp(-(t-s)) d\mathbf{B}_s.$$

The mean of $\mathbf{x}_t$ is thus $\exp(-t)\mathbf{x}_0$ as claimed, and the covariance scales proportionally to the quadratic variation, i.e., $2 \int_0^t \exp(-2(t-s)) ds = 1 - \exp(-2t)$. $\square$

Let us define π_t to be the distribution of $\mathbf{x}_t$ according to (5), where we draw $\mathbf{x}_0 \sim \pi$, the target of our sampling algorithm. We apply a one-to-one backwards time change $\tau : [0, \infty] \rightarrow [0, \infty]$, i.e., a τ that satisfies $\tau'(t) < 0$ for all times $t \geq 0$, and sets $\tau(0) = \infty$ and $\tau(\infty) = 0$. This induces an equivalent backwards process, indexed by a backwards time $u \geq 0$:

$$\mathbf{x}_u^\leftarrow := \mathbf{x}_{\tau^{-1}(u)}, \quad \pi_u^\leftarrow := \pi_{\tau^{-1}(u)}. \quad (6)$$

To see how (6) relates to generative models, if we let $\mathbf{x}_0^\leftarrow = \mathbf{x}_\infty$, which we can simulate with a draw from the stationary distribution $\mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, then running the “backwards process” and producing a sample $\mathbf{x}_\infty^\leftarrow$ is equivalent to drawing from the target distribution $\pi_\infty^\leftarrow = \pi_0 = \pi$.

We now give a helpful characterization of the backwards process $\{\mathbf{x}_u^\leftarrow\}_{u \geq 0}$. The following result is standard, and an early derivation of it can be found, e.g., in [And82].

Lemma 7. *Let $d\mathbf{x}_t = \mathbf{x}_t dt + \sqrt{2} d\mathbf{B}_t$, where $\{\mathbf{B}_t\}_{t \geq 0}$ is a Wiener process. Let $\tau : [0, \infty] \rightarrow [0, \infty]$ satisfy $\tau'(t) < 0$ for all $t \geq 0$, $\tau(0) = \infty$, $\tau(\infty) = 0$. Letting π_t be the law of $\mathbf{x}_t$, and following the notation (6), we have for a Wiener process $\{\mathbf{W}_u\}_{u \geq 0}$ that*

$$d\mathbf{x}_u^\leftarrow = |(\tau^{-1})'(u)| (\mathbf{x}_u^\leftarrow + 2\nabla \log \pi_u^\leftarrow(\mathbf{x}_u^\leftarrow)) du + \sqrt{2|(\tau^{-1})'(u)|} d\mathbf{W}_u.$$

Proof. By the Fokker-Planck equation (Lemma 2), the measures $\{\pi_t\}_{t \geq 0}$ follow the PDE:

$$\begin{aligned} \partial_t \pi_t(\mathbf{x}) &= -\nabla \cdot (-\mathbf{x} \pi_t(\mathbf{x})) + \Delta \pi_t(\mathbf{x}) \\ &= \nabla \cdot ((\mathbf{x} + \nabla \log \pi_t(\mathbf{x})) \pi_t(\mathbf{x})) \\ &= \nabla \cdot ((\mathbf{x} + 2\nabla \log \pi_t(\mathbf{x})) \pi_t(\mathbf{x})) - \Delta \pi_t(\mathbf{x}). \end{aligned}$$

Performing a change of variables, we have

$$\partial_u \pi_u^\leftarrow(\mathbf{x}^\leftarrow) = -|(\tau^{-1})'(u)| \nabla \cdot ((\mathbf{x}^\leftarrow + 2\nabla \log \pi_u^\leftarrow(\mathbf{x}^\leftarrow)) \pi_u^\leftarrow(\mathbf{x}^\leftarrow)) + |(\tau^{-1})'(u)| \Delta \pi_u^\leftarrow(\mathbf{x}^\leftarrow).$$

The conclusion follows by inverting the Fokker-Planck equation with respect to the above display. $\square$

We can now introduce our main object of study in this section.

Perspective 4. Let $\pi \in \mathcal{P}(\mathbb{R}^d)$ be such that $\mathbb{E}_{\mathbf{x} \sim \pi}[\mathbf{x}]$ exists. Define a stochastic process $\{\mathbf{x}_u^\leftarrow\}_{u \geq 0}$ as follows, where $\{\mathbf{W}_u\}_{u \geq 0}$ is a Wiener process:

$$\begin{aligned} \mathbf{x}_0^\leftarrow &\sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d), \\ d\mathbf{x}_u^\leftarrow &= \left(\frac{\mathbf{x}_u^\leftarrow}{2u(u+1)} + \frac{1}{u(u+1)} \nabla \log \pi_u^\leftarrow(\mathbf{x}_u^\leftarrow) \right) du + \frac{1}{\sqrt{u(u+1)}} d\mathbf{W}_u. \end{aligned} \quad (\text{rev})$$

Note that (rev) is simply the SDE driving the backwards dynamics (6), according to our derivation in Lemma 7, under the time change

$$\tau(t) = \frac{1}{\exp(2t) - 1}, \quad \tau^{-1}(u) = \frac{1}{2} \log\left(\frac{u+1}{u}\right), \quad (\tau^{-1})'(u) = -\frac{1}{2u(u+1)}. \quad (7)$$

We remark that (rev) is not the original backwards SDE from [SSKKEP21], Eq. (6), which reads

$$d\mathbf{x}_u^\leftarrow = (\mathbf{x}_u^\leftarrow + 2\nabla \log \pi_u^\leftarrow(\mathbf{x}_u^\leftarrow)) du + \sqrt{2} d\mathbf{W}_u. \quad (8)$$

The above SDE is the result of applying Lemma 7 with a time change satisfying $(\tau^{-1})'(u) = -u$. However, because there is no well-defined $\tau : [0, \infty] \rightarrow [0, \infty]$ with this property, applications in finite time require cutting off the forwards SDE at a time $T > 0$ such that $\pi_T \approx \pi_\infty = \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)$, and only defining the backwards SDE up to time T (i.e., $\tau(t) = T - t$). To avoid these boundary issues, Perspective 4 uses (7), which just gives a reparameterization of (8) if we restrict to $[0, T]$.

As a final technical tool, we require Tweedie’s formula (cf. [Rob56]) to rewrite the $\nabla \log \pi^\leftarrow$ term, which is often called the *score* of the backwards process.

Lemma 8 (Tweedie’s formula). Let $\mathbf{y} \sim \mathcal{N}(\mathbf{x}, \sigma^2 \mathbf{I}_d)$ be the output of $\mathbf{x} \sim \pi$ passed through a noisy Gaussian channel, let $\pi(\cdot | \mathbf{y})$ be the posterior distribution, and let $\nu(\mathbf{y})$ be the marginal of $\mathbf{y}$:

$$\pi(\mathbf{x} | \mathbf{y}) \propto \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{y}\|_2^2\right) \pi(\mathbf{x}), \quad \nu(\mathbf{y}) \propto \int \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{y}\|_2^2\right) \pi(\mathbf{x}) d\mathbf{x}.$$

Then, $\nabla \log \nu(\mathbf{y}) = \frac{1}{\sigma^2} (\mathbb{E}_{\pi(\cdot | \mathbf{y})}[\mathbf{x}] - \mathbf{y})$.

Proof. Note that the proportionality constant in ν does not affect $\nabla \log \nu(\mathbf{y})$. Hence,

$$\begin{aligned} \nabla \log \nu(\mathbf{y}) &= \frac{1}{\int \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{y}\|_2^2\right) \pi(\mathbf{x}) d\mathbf{x}} \int \frac{\mathbf{x} - \mathbf{y}}{\sigma^2} \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x} - \mathbf{y}\|_2^2\right) \pi(\mathbf{x}) d\mathbf{x} \\ &= \frac{1}{\sigma^2} (\mathbb{E}_{\pi(\cdot | \mathbf{y})}[\mathbf{x}] - \mathbf{y}). \end{aligned}$$

□

Lemma 8 is remarkably useful in practical applications. Recall from Lemma 6 that intermediate distributions in the forwards and backwards processes are outputs of noisy Gaussian channels, initialized from π . Further, the backwards SDEs (rev), (8) are entirely explicit, except for the score term $\nabla \log \pi_u^\leftarrow$. Lemma 8 says that we can approximate this score term by predicting appropriate posterior means $\mathbb{E}_{\pi(\cdot | \mathbf{y})}$, where $\mathbf{y} = \mathbf{x}_u^\leftarrow$ is an intermediate iterate. In practice, this predictor can be learned by minimizing an empirical risk over samples from π , and modern machine learning models (e.g., deep neural networks) achieve good prediction error. Moreover, fairly strong bounds are available that convert errors arising from score approximation and finite-time discretization to errors in the sampling process, see e.g., a line of work initiated by [CCLLSZ23; LIT23].

We conclude by relating Perspective 4 to Perspective 1, as observed by [KP21; Mon23].

Theorem 3. The processes $\{\mathbf{c}_u\}_{u \geq 0}$ and $\{\mathbf{x}_u^\leftarrow\}_{u \geq 0}$ in (SL-I), (rev), satisfy $\sqrt{u(u+1)} \mathbf{x}_u^\leftarrow = \mathbf{c}_u$.

Proof. First, consider the SDE (SL-I). Performing the change of variables $\mathbf{x}_u^\leftarrow = \frac{1}{\sqrt{u(u+1)}} \mathbf{c}_u$,

$$\begin{aligned} d\mathbf{x}_u^\leftarrow &= \left(\frac{d}{du} \frac{1}{\sqrt{u(u+1)}} \right) \mathbf{c}_u du + \frac{1}{\sqrt{u(u+1)}} d\mathbf{c}_u \\ &= -\frac{2u+1}{2(u(u+1))^{\frac{3}{2}}} \mathbf{c}_u du + \frac{1}{\sqrt{u(u+1)}} \mathbf{m}_u du + \frac{1}{\sqrt{u(u+1)}} d\mathbf{W}_u \\ &= -\frac{2u+1}{2u(u+1)} \mathbf{x}_u^\leftarrow du + \frac{1}{\sqrt{u(u+1)}} \mathbf{m}_u du + \frac{1}{\sqrt{u(u+1)}} d\mathbf{W}_u, \end{aligned} \quad (9)$$

where $\mathbf{m}_u$ is the mean of the measure $\propto \exp(\mathbf{c}_u^\top \mathbf{x} - \frac{u}{2} \|\mathbf{x}\|_2^2) \pi(\mathbf{x})$, as in (SL-I).

On the other hand, observe that $\mathbf{x}_u^\leftarrow$ in (rev) is distributed as $\mathbf{x}_t$ in (5), for $t = \frac{1}{2} \log(\frac{u+1}{u})$. By Lemma 6, we have $\mathbf{x}_t \sim \mathcal{N}(\exp(-t)\mathbf{x}_0, (1 - \exp(-2t))\mathbf{I}_d)$. Now Tweedie's formula (Lemma 8), applied with $\mathbf{y} \leftarrow \mathbf{x}_t$, $\mathbf{x} \leftarrow \exp(-t)\mathbf{x}_0$, and $\sigma^2 \leftarrow 1 - \exp(-2t)$, gives

$$\begin{aligned} \nabla \log \pi_t(\mathbf{x}_t) &= \frac{1}{1 - \exp(-2t)} (\mathbb{E}_{\pi_t}[\exp(-t)\mathbf{x}_0] - \mathbf{x}_t) \\ \implies \nabla \log \pi_u^\leftarrow(\mathbf{x}_u^\leftarrow) &= \sqrt{u(u+1)} \mathbb{E}_{\pi_u^\leftarrow}[\mathbf{x}_0] - (u+1)\mathbf{x}_u^\leftarrow. \end{aligned} \quad (10)$$

Combining (rev) and (10) yields

$$\begin{aligned} d\mathbf{x}_u^\leftarrow &= \left(\frac{\mathbf{x}_u^\leftarrow}{2u(u+1)} + \frac{1}{\sqrt{u(u+1)}} \mathbb{E}_{\pi_u^\leftarrow}[\mathbf{x}_0] - \frac{\mathbf{x}_u^\leftarrow}{u} \right) du + \frac{1}{\sqrt{u(u+1)}} d\mathbf{W}_u \\ &= \left(-\frac{2u+1}{2u(u+1)} \mathbf{x}_u^\leftarrow + \frac{1}{\sqrt{u(u+1)}} \mathbb{E}_{\pi_u^\leftarrow}[\mathbf{x}_0] \right) du + \frac{1}{\sqrt{u(u+1)}} d\mathbf{W}_u. \end{aligned} \quad (11)$$

To connect (9) and (11), recall that $\pi_u^\leftarrow$ is the density of $\mathcal{N}(\exp(-t)\mathbf{x}_0, (1 - \exp(-2t))\mathbf{I}_d) = \mathcal{N}(\sqrt{\frac{u}{u+1}}\mathbf{x}_0, \frac{1}{u+1}\mathbf{I}_d)$, so under $\pi_u^\leftarrow$, the posterior distribution that $\mathbf{x}_0$ follows (cf. Lemma 8) is

$$\propto \exp\left(-\frac{u+1}{2} \left\| \sqrt{\frac{u}{u+1}} \mathbf{x}_0 - \mathbf{x}_u^\leftarrow \right\|_2^2\right) \pi(\mathbf{x}_0) \propto \exp\left(\mathbf{c}_u^\top \mathbf{x}_0 - \frac{u}{2} \|\mathbf{x}_0\|_2^2\right) \pi(\mathbf{x}_0),$$

which agrees with the distribution used to compute $\mathbf{m}_u$ in (9) as desired. $\square$

5 Renormalization

We present another perspective called renormalization [WK74; Pol84]. The idea is to decompose a measure π via infinitesimal convolutions, and describe how so-called “renormalized potentials” evolve as a Markov process. We follow the recent presentation of [BBD24].

In this section, we use σ, τ to denote times in $[0, 1]$, and $\mathcal{F}(\mathbb{R}^d)$ denotes the set of smooth compactly-supported measurable functions on $\mathbb{R}^d$.³ We now introduce our renormalized potentials and measures $\{V_\tau, \nu_\tau\}_{\tau \in [0, 1]}$, and the Polchinski semigroup $\{P_{\sigma, \tau}\}_{0 \leq \sigma \leq \tau \leq 1}$ that they induce.

Perspective 5. Let $\pi \in \mathcal{P}(\mathbb{R}^d)$ be such that $\mathbb{E}_{\mathbf{x} \sim \pi}[\mathbf{x}]$ exists. Define a family of renormalized potentials and measures as follows: let $V_1(\mathbf{x}) := -\log \pi(\mathbf{x}) - \frac{1}{2} \|\mathbf{x}\|_2^2$, and

$$\begin{aligned} V_\tau(\mathbf{x}) &:= -\log(\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)}[\exp(-V_1(\mathbf{x} + \mathbf{z}))]), \\ \nu_\tau(\mathbf{x}) &\propto \exp\left(-V_\tau(\mathbf{x}) - \frac{1}{2\tau} \|\mathbf{x}\|_2^2\right), \text{ for all } \tau \in [0, 1]. \end{aligned} \quad (\text{renorm})$$

Also, define the induced Polchinski semigroup: for all $0 \leq \sigma \leq \tau \leq 1$, let

$$P_{\sigma, \tau} f(\mathbf{x}) := \exp(V_\sigma(\mathbf{x})) \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (\tau-\sigma)\mathbf{I}_d)}[\exp(-V_\tau(\mathbf{x} + \mathbf{z})) f(\mathbf{x} + \mathbf{z})]. \quad (\text{PSG})$$

³An extended discussion on Markov semigroup theory is beyond our scope. The results here extend beyond the class $\mathcal{F}(\mathbb{R}^d)$ defined here via approximation; we refer the reader to [BGL13] for additional background.

It is clear that $\nu_1 = \pi$ in (renorm), and we let $\nu_0 := \delta_{\mathbf{0}_d}$ be a Dirac measure by weak continuity.

We next clarify the relationship between (renorm) and (PSG). Intuitively, $P_{\sigma,\tau}$ advances time from σ to τ (in the sense described in Section 1.2), and the corresponding densities are ν_σ, ν_τ .

Lemma 9. *Following notation in Perspective 5, we have that $P_{\rho,\sigma}P_{\sigma,\tau} = P_{\rho,\tau}$ for all $0 \leq \rho \leq \sigma \leq \tau \leq 1$. Moreover, $P_{\sigma,\tau}^*\nu_\sigma = \nu_\tau$ for all $0 \leq \sigma \leq \tau \leq 1$, in the sense that*

$$\mathbb{E}_{\nu_\tau}[f] = \mathbb{E}_{\nu_\sigma}[P_{\sigma,\tau}f] \text{ for all } f \in \mathcal{F}(\mathbb{R}^d). \quad (12)$$

Proof. By applying (PSG) twice,

$$\begin{aligned} P_{\rho,\sigma}P_{\sigma,\tau}f(\mathbf{x}) &= \exp(V_\rho(\mathbf{x}))\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}_d, (\sigma-\rho)\mathbf{I}_d)}[\exp(-V_\sigma(\mathbf{x} + \mathbf{y}))\exp(V_\sigma(\mathbf{x} + \mathbf{y})) \\ &\quad \cdot \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (\tau-\sigma)\mathbf{I}_d)}[\exp(-V_\tau(\mathbf{x} + \mathbf{y} + \mathbf{z}))f(\mathbf{x} + \mathbf{y} + \mathbf{z})]] \\ &= \exp(V_\rho(\mathbf{x}))\mathbb{E}_{\mathbf{y} \sim \mathcal{N}(\mathbf{0}_d, (\sigma-\rho)\mathbf{I}_d)}[\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (\tau-\sigma)\mathbf{I}_d)}[\exp(-V_\tau(\mathbf{x} + \mathbf{y} + \mathbf{z}))f(\mathbf{x} + \mathbf{y} + \mathbf{z})]] \\ &= \exp(V_\rho(\mathbf{x}))\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (\tau-\rho)\mathbf{I}_d)}[\exp(-V_\tau(\mathbf{x} + \mathbf{z}))f(\mathbf{x} + \mathbf{z})] = P_{\rho,\tau}f(\mathbf{x}), \end{aligned}$$

where we use the fact that if $\mathbf{x}_1 \sim \mathcal{N}(\mathbf{0}_d, \Sigma_1), \mathbf{x}_2 \sim \mathcal{N}(\mathbf{0}_d, \Sigma_2)$ then $\mathbf{x}_1 + \mathbf{x}_2 \sim \mathcal{N}(\mathbf{0}_d, \Sigma_1 + \Sigma_2)$.

Next, note that for all $\tau \in [0, 1]$,

$$\begin{aligned} (2\pi\tau)^{-\frac{d}{2}} \int \exp\left(-V_\tau(\mathbf{x}) - \frac{1}{2\tau}\|\mathbf{x}\|_2^2\right) d\mathbf{x} &= \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{0}_d, \tau\mathbf{I}_d)}[\exp(-V_\tau(\mathbf{x}))] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{N}(\mathbf{0}_d, \tau\mathbf{I}_d)}[\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)}[\exp(-V_1(\mathbf{x} + \mathbf{z}))]] \\ &= \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, \mathbf{I}_d)}[\exp(-V_1(\mathbf{z}))] = \exp(-V_0(\mathbf{0}_d)), \end{aligned}$$

so the normalizing constant of ν_τ is $(2\pi\tau)^{\frac{d}{2}} \exp(V_0(\mathbf{0}_d))$. Thus, for all $f \in \mathcal{F}(\mathbb{R}^d)$,

$$\begin{aligned} P_{0,\tau}f(\mathbf{0}_d) &= \exp(V_0(\mathbf{0}_d)) \cdot \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, \tau\mathbf{I}_d)}[\exp(-V_\tau(\mathbf{z}))f(\mathbf{z})] \\ &= \frac{(2\pi\tau)^{\frac{d}{2}}}{\int \exp\left(-V_\tau(\mathbf{z}) - \frac{1}{2\tau}\|\mathbf{z}\|_2^2\right) d\mathbf{z}} \cdot \frac{\int \exp\left(-V_\tau(\mathbf{z}) - \frac{1}{2\tau}\|\mathbf{z}\|_2^2\right) f(\mathbf{z}) d\mathbf{z}}{(2\pi\tau)^{\frac{d}{2}}} = \mathbb{E}_{\nu_\tau}[f]. \end{aligned} \quad (13)$$

The second claim now follows as

$$\mathbb{E}_{\nu_\tau}[f] = P_{0,\tau}f(\mathbf{0}_d) = P_{0,\sigma}P_{\sigma,\tau}f(\mathbf{0}_d) = \mathbb{E}_{\nu_\sigma}[P_{\sigma,\tau}f].$$

□

We next describe $\{P_{\sigma,\tau}\}_{0 \leq \sigma \leq \tau \leq 1}$ via its infinitesimal generators $\{L_\tau\}_{\tau \in [0,1]}$.

Lemma 10. *Define a family of operators $\{L_\tau\}_{\tau \in [0,1]}$ that act on $f \in \mathcal{F}(\mathbb{R}^d)$ via*

$$L_\tau f = \frac{1}{2}\Delta f - \langle \nabla V_\tau, \nabla f \rangle, \text{ for all } \tau \in [0, 1]. \quad (14)$$

Then, we have for all $0 \leq \sigma \leq \tau \leq 1$ and $f \in \mathcal{F}(\mathbb{R}^d)$ that

$$\partial_\sigma P_{\sigma,\tau}f = -L_\sigma P_{\sigma,\tau}f, \quad \partial_\tau P_{\sigma,\tau}f = P_{\sigma,\tau}L_\tau f, \quad \partial_\tau \mathbb{E}_{\nu_\tau}[f] = \mathbb{E}_{\nu_\tau}[L_\tau f]. \quad (15)$$

Proof. Let $\gamma_\tau(\mathbf{z}) \propto \exp(-\frac{1}{2\tau}\|\mathbf{z}\|_2^2)$ be the density of $\mathcal{N}(\mathbf{0}_d, \tau\mathbf{I}_d)$, and recall the heat equation, $\partial_\tau \gamma_\tau = \frac{1}{2}\Delta \gamma_\tau$. Then, letting $(f * g)(\mathbf{x}) := \int f(\mathbf{x} - \mathbf{z})g(\mathbf{z}) d\mathbf{z}$ denote convolution, we have

$$\partial_\tau (f * \gamma_{1-\tau})(\mathbf{x}) = \partial_\tau \int f(\mathbf{z})\gamma_{1-\tau}(\mathbf{x} - \mathbf{z})d\mathbf{z} = -\frac{1}{2}\Delta (f * \gamma_{1-\tau})(\mathbf{x}). \quad (16)$$

Next, define $Z_\tau(\mathbf{x}) := \exp(-V_\tau(\mathbf{x})) = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)}[\exp(-V_1(\mathbf{x} - \mathbf{z}))]$, where we used symmetry of $\mathbf{z}$ in the last equality. Because $Z_\tau = \exp(-V_1) * \gamma_{1-\tau}$, applying (16) yields

$$\partial_\tau Z_\tau(\mathbf{x}) = -\frac{1}{2}\Delta Z_\tau(\mathbf{x}).$$

Now applying the chain rule to the definition of Z_τ yields

$$\nabla Z_\tau = -\nabla V_\tau \exp(-V_\tau) \implies \Delta Z_\tau = \left(-\Delta V_\tau + \|\nabla V_\tau\|_2^2\right) \exp(-V_\tau).$$

Combining the above two displays, we derive

$$\partial_\tau V_\tau = -\frac{\partial_\tau Z_\tau}{Z_\tau} = \frac{\Delta Z_\tau}{2Z_\tau} = -\frac{1}{2}\Delta V_\tau + \frac{1}{2}\|\nabla V_\tau\|_2^2. \quad (17)$$

Finally,

$$\begin{aligned} \partial_\sigma P_{\sigma,\tau} f &= (\partial_\sigma V_\sigma) P_{\sigma,\tau} f - \frac{1}{2} \exp(V_\sigma) \Delta (\exp(-V_\sigma) P_{\sigma,\tau} f) \\ &= (\partial_\sigma V_\sigma) P_{\sigma,\tau} f - \frac{1}{2} \exp(V_\sigma) \nabla \cdot (-\nabla V_\sigma \exp(-V_\sigma) P_{\sigma,\tau} f + \exp(-V_\sigma) \nabla P_{\sigma,\tau} f) \\ &= (\partial_\sigma V_\sigma) P_{\sigma,\tau} f + \frac{1}{2} (\Delta V_\sigma) P_{\sigma,\tau} f - \frac{1}{2} \|\nabla V_\sigma\|_2^2 P_{\sigma,\tau} f - \frac{1}{2} \Delta P_{\sigma,\tau} f + \langle \nabla V_\sigma, \nabla P_{\sigma,\tau} f \rangle \\ &= -\frac{1}{2} \Delta P_{\sigma,\tau} f + \langle \nabla V_\sigma, \nabla P_{\sigma,\tau} f \rangle = -L_\sigma P_{\sigma,\tau} f, \end{aligned}$$

where we used the chain rule and (16) in the first line, and (17) in the last. The claim about $\partial_\tau P_{\sigma,\tau} f$ follows similarly. For the last claim, applying (12) twice for $\sigma \leq \tau$,

$$\partial_\tau \mathbb{E}_{\nu_\tau} [f] = \partial_\tau \mathbb{E}_{\nu_\sigma} [P_{\sigma,\tau} f] = \mathbb{E}_{\nu_\sigma} [\partial_\tau P_{\sigma,\tau} f] = \mathbb{E}_{\nu_\sigma} [P_{\sigma,\tau} L_\tau f] = \mathbb{E}_{\nu_\tau} [L_\tau f].$$

□

The evolution of the renormalized potential V_τ in (17) is known as the *Polchinski equation*. Further, note that (15) is consistent with our earlier definition of infinitesimal generators in (1), because the following *Kolmogorov equations* hold: using (1) as our definition,

$$\begin{aligned} \partial_\sigma P_{\sigma,\tau} f &= \lim_{\eta \searrow 0} \frac{P_{\sigma+\eta,\tau} f - P_{\sigma,\tau} f}{\eta} = \lim_{\eta \searrow 0} \frac{P_{\sigma,\sigma+\eta}(-P_{\sigma+\eta,\tau} f) - (-P_{\sigma+\eta,\tau} f)}{\eta} = -L_\sigma P_{\sigma,\tau} f, \\ \partial_\tau P_{\sigma,\tau} f &= \lim_{\eta \searrow 0} \frac{P_{\sigma,\tau+\eta} f - P_{\sigma,\tau} f}{\eta} = P_{\sigma,\tau} \lim_{\eta \searrow 0} \frac{P_{\tau,\tau+\eta} f - f}{\eta} = P_{\sigma,\tau} L_\tau f. \end{aligned}$$

To complete our exposition of the Polchinski semigroup, we note that the infinitesimal generators $\{L_\tau\}_{\tau \in [0,1]}$ induce a SDE on particles $\{\mathbf{v}_\tau \in \mathbb{R}^d\}_{\tau \in [0,1]}$, such that $\mathbf{v}_\tau \sim \nu_\tau$.

Lemma 11. *For all $\mathbf{v} \in \mathbb{R}^d$ and $\tau \in [0,1]$, define an induced fluctuation measure*

$$\pi_\tau^{\mathbf{v}}(\mathbf{x}) \propto \exp\left(\frac{1}{1-\tau} \langle \mathbf{v}, \mathbf{x} \rangle - \frac{\tau}{2(1-\tau)} \|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}). \quad (18)$$

Consider the following SDE, where $\{\mathbf{W}_\tau\}_{\tau \in [0,1]}$ is a Wiener process:

$$d\mathbf{v}_\tau = -\frac{1}{1-\tau}(\mathbf{v}_\tau - \mathbf{m}_\tau) d\tau + d\mathbf{W}_\tau, \text{ where } \mathbf{m}_\tau := \mathbb{E}_{\mathbf{x} \sim \pi_\tau^{\mathbf{v}_\tau}}[\mathbf{x}], \mathbf{v}_0 = \mathbf{0}_d. \quad (19)$$

The infinitesimal generator of (19) is L_τ in (14), and for any $\tau \in [0,1]$, we have $\mathbf{v}_\tau \sim \nu_\tau$.

Proof. We first claim that

$$\mathbb{E}_{\pi_\tau^{\mathbf{v}}} [f] = P_{\tau,1} f(\mathbf{v}) \text{ for all } f \in \mathcal{F}(\mathbb{R}^d), \quad (20)$$

which follows from

$$\begin{aligned} P_{\tau,1} f(\mathbf{v}) &= \frac{\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)} [\exp(-V_1(\mathbf{v} + \mathbf{z})) f(\mathbf{v} + \mathbf{z})]}{\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)} [\exp(-V_1(\mathbf{v} + \mathbf{z}))]} \\ &= \frac{\int \exp\left(-V_1(\mathbf{x}) - \frac{1}{2(1-\tau)} \|\mathbf{x} - \mathbf{v}\|_2^2\right) f(\mathbf{x}) d\mathbf{x}}{\int \exp\left(-V_1(\mathbf{x}) - \frac{1}{2(1-\tau)} \|\mathbf{x} - \mathbf{v}\|_2^2\right) d\mathbf{x}}, \end{aligned} \quad (21)$$

where the second line substituted $\mathbf{x} = \mathbf{v} + \mathbf{z}$. Next, Lemma 2 shows that if

$$d\mathbf{v}_\tau = -\nabla V_\tau(\mathbf{v}_\tau) d\tau + d\mathbf{W}_\tau, \quad (22)$$

the generator of (22) is precisely L_τ , and from $P_{\sigma,\tau}^* \nu_\sigma = \nu_\tau$ (Lemma 9) we obtain the final statement, assuming we can show (22) is equivalent to (19). To obtain this equivalence, note that

$$\nabla V_\tau(\mathbf{v}) = \frac{\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)}[\exp(-V_1(\mathbf{v} + \mathbf{z})) \nabla V_1(\mathbf{v} + \mathbf{z})]}{\mathbb{E}_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)}[\exp(-V_1(\mathbf{v} + \mathbf{z}))]} = (P_{\tau,1} \nabla V_1)(\mathbf{v}) = \mathbb{E}_{\pi_\tau^\mathbf{v}}[\nabla V_1],$$

where the last equality uses our earlier derivation (21) coordinatewise. We conclude by recalling that for any measure μ with sufficiently fast decay, $\mathbb{E}_{\mathbf{x} \sim \mu}[\nabla \log \mu(\mathbf{x})] = \int \nabla \mu(\mathbf{x}) d\mathbf{x} = \mathbf{0}_d$, where we integrated by parts coordinatewise. Applying this fact with $\mu = \pi_\tau^\mathbf{v}$ completes the proof:

$$\mathbb{E}_{\mathbf{x} \sim \pi_\tau^\mathbf{v}} \left[-\nabla V_1(\mathbf{x}) + \frac{1}{1-\tau}(\mathbf{v}_\tau - \mathbf{x}) \right] = \mathbf{0}_d \implies \mathbb{E}_{\pi_\tau^\mathbf{v}}[\nabla V_1] = \frac{1}{1-\tau}(\mathbf{v}_\tau - \mathbf{m}_\tau).$$

Alternatively, this can be seen by an application of Tweedie's formula (Lemma 8), noting that $-\nabla V_\tau(\mathbf{x})$ is nothing more than $\nabla \log \mu(\mathbf{x})$, where

$$\mu(\mathbf{x}) \propto \int \exp\left(-\frac{1}{2(1-\tau)}\|\mathbf{x} - \mathbf{z}\|^2\right) \exp(-V_1(\mathbf{z})) d\mathbf{z},$$

employing a change of variables. $\square$

We are finally ready to connect the Polchinski semigroup to stochastic localization.

Theorem 4. *The processes $\{\mathbf{c}_t\}_{t \geq 0}$ and $\{\mathbf{v}_\tau\}_{\tau \in [0,1]}$ in (SL-I), (19) satisfy $\mathbf{c}_t = \frac{1}{1-\tau} \mathbf{v}_\tau$, where $t = \frac{\tau}{1-\tau}$, and the induced π_t in (SL-I) and $\pi_\tau^\mathbf{v}$ in (18) are identical under this reparameterization.*

Proof. We begin by equating the measures in (SL-I) and (18): starting from (SL-I),

$$\pi_t(\mathbf{x}) \propto \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{t}{2}\|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}) = \exp\left(\frac{1}{1-\tau} \langle \mathbf{v}_\tau, \mathbf{x} \rangle - \frac{\tau}{2(1-\tau)}\|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}),$$

which matches the definition of $\pi_\tau^\mathbf{v}$ in (18). Therefore, $\mathbf{m}_t$ in (SL-I) and $\mathbf{m}_\tau$ in (19) have the same definition, under the time change $t = \frac{\tau}{1-\tau}$. Next, letting $\mathbf{u}_\tau := \frac{1}{1-\tau} \mathbf{v}_\tau$ where $\mathbf{v}_\tau$ follows (19),

$$\partial_\tau \mathbf{u}_\tau = \frac{1}{(1-\tau)^2} \mathbf{v}_\tau + \frac{1}{1-\tau} \frac{\partial}{\partial \tau} \mathbf{v}_\tau = \frac{1}{(1-\tau)^2} \mathbf{m}_\tau + \frac{1}{1-\tau} d\mathbf{W}_\tau. \quad (23)$$

At this point, applying the time change formula⁴ for Itô diffusions (cf. Theorem 8.5.1, [Øks00]), with $\partial_\tau \frac{1}{1-\tau} = \frac{1}{(1-\tau)^2}$ shows equivalence of $\mathbf{u}_\tau$ in (23) and $\mathbf{c}_t$ in (SL-I). $\square$

6 Static Schrödinger bridge

In this section, we place stochastic localization in the context of the Schrödinger bridge [Sch31; Sch32]. We again restrict time to $\tau \in [0, 1]$, and follow the presentation of [Léo14] throughout.

We begin by describing a general formulation of the Schrödinger bridge optimization problem.

Perspective 6. *Let $\mu, \pi \in \mathcal{P}(\mathbb{R}^d)$, and let $R \in \mathcal{P}([0, 1] \times \mathbb{R}^d)$ be a reference path measure on $\mathcal{C}([0, 1] \times \mathbb{R}^d)$. The induced static Schrödinger bridge problem is:*

$$\inf_{P \in \mathcal{P}([0, 1] \times \mathbb{R}^d)} \text{KL}(P \parallel R) \text{ such that } P_0 = \mu, P_1 = \pi. \quad (\text{SSB})$$

In other words, we seek the path measure on $\mathcal{C}([0, 1] \times \mathbb{R}^d)$, with prescribed starting measure $P_0 = \mu$ and ending measure $P_1 = \pi$, that is as close as possible in KL divergence to a given path measure

⁴This formula can be arrived at via a similar calculation as in Lemma 7, that changes but does not reverse time.

R . It may be helpful to keep the running example of R being the Wiener path measure in mind, as we specialize calculations at the end of the section to this case.

We now introduce some helpful notation. When $P \in \mathcal{P}([0, 1] \times \mathbb{R}^d)$ is a path measure, P_{01} denotes the marginal density of its endpoints (P_0, P_1) , and similarly we denote the endpoints of $\mathbf{p}_{[0,1]} \in \mathcal{C}([0, 1] \times \mathbb{R}^d)$ by $\mathbf{p}_{01} \in \mathbb{R}^d \times \mathbb{R}^d$. For $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^d$, we let $P_{(0,1)|(\mathbf{x}, \mathbf{y})}$ denote the conditional density of $\mathbf{p}_{(0,1)} \mid \mathbf{p}_{01} = (\mathbf{x}, \mathbf{y})$ for $\mathbf{p}_{[0,1]} \sim P$, i.e., the intermediate path given the endpoints. Finally, we let $\Gamma(\mu, \pi)$ denote the set of *couplings* of μ and π , i.e., measures γ on $\mathbb{R}^d \times \mathbb{R}^d$ whose left and right marginals, denoted γ_0 and γ_1 , respectively equal μ and π .

The first observation regarding (SSB) is that to find its solution (when it is attainable), it suffices to find the coupling of the marginals $P_0 = \mu$ and $P_1 = \pi$ closest to R_{01} in KL divergence. This coupling is then extended to paths on $[0, 1]$ via matching conditional distributions.

Lemma 12. *Let P^* optimally solve (SSB), and let $\gamma^* \in \Gamma(\mu, \pi)$ optimally solve*

$$\inf_{\gamma \in \Gamma(\mu, \pi)} \text{KL}(\gamma \parallel R_{01}). \quad (24)$$

Then P^ and γ^* are related as follows: for all $\mathbf{p}_{[0,1]} \in \mathcal{C}([0, 1] \times \mathbb{R}^d)$ and $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^d$,*

$$P_{01}^* = \gamma^*, \quad P_{(0,1)|(\mathbf{x}, \mathbf{y})}(\mathbf{p}_{(0,1)} \mid \mathbf{p}_{01} = (\mathbf{x}, \mathbf{y})) = R_{(0,1)|(\mathbf{x}, \mathbf{y})}(\mathbf{p}_{(0,1)} \mid \mathbf{p}_{01} = (\mathbf{x}, \mathbf{y})).$$

Proof. The chain rule for the KL divergence gives

$$\text{KL}(P \parallel R) = \text{KL}(P_{01} \parallel R_{01}) + \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P_{01}} [\text{KL}(P_{(0,1)|(\mathbf{x}, \mathbf{y})} \parallel R_{(0,1)|(\mathbf{x}, \mathbf{y})})].$$

We can choose P_{01} and $P_{(0,1)|(\mathbf{x}, \mathbf{y})}$ separately: the first term above is the objective in (24), and the second is optimized (in fact zero) by setting $P_{(0,1)|(\mathbf{x}, \mathbf{y})} = R_{(0,1)|(\mathbf{x}, \mathbf{y})}$ for all $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^d$. $\square$

Next, we characterize the solution to the (convex) problem (24) via taking its dual.

Lemma 13. *Let $\gamma^* \in \Gamma(\mu, \pi)$ optimally solve (24). There exist $f, g : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ such that*

$$\gamma^*(\mathbf{x}, \mathbf{y}) = R_{01}(\mathbf{x}, \mathbf{y})f(\mathbf{x})g(\mathbf{y}) \quad (25)$$

for all $(\mathbf{x}, \mathbf{y}) \in \mathbb{R}^d \times \mathbb{R}^d$, and the following Schrödinger system is satisfied:

$$\frac{\mu(\mathbf{x})}{R_0(\mathbf{x})} = f(\mathbf{x})\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim R_{01}}[g(\mathbf{y}) \mid \mathbf{x}], \quad \frac{\pi(\mathbf{y})}{R_1(\mathbf{y})} = g(\mathbf{y})\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim R_{01}}[f(\mathbf{x}) \mid \mathbf{y}]. \quad (26)$$

Proof. We begin by writing the Lagrangian of (24): denoting $\rho := R_{01}$ for short, it is

$$\text{KL}(\gamma \parallel \rho) + \int \phi(\mathbf{x}) \left(\mu(\mathbf{x}) - \int \gamma(\mathbf{x}, \mathbf{y}) d\mathbf{y} \right) d\mathbf{x} + \int \psi(\mathbf{y}) \left(\pi(\mathbf{y}) - \int \gamma(\mathbf{x}, \mathbf{y}) d\mathbf{x} \right) d\mathbf{y},$$

where $\phi, \psi : \mathbb{R}^d \rightarrow \mathbb{R}$ are absolutely integrable with respect to μ, π respectively. To minimize this over γ ,⁵ we derive using the Donsker–Varadhan variational formula:

$$\begin{aligned} \min_{\gamma \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)} & \left\{ \text{KL}(\gamma \parallel \rho) - \iint (\phi(\mathbf{x}) + \psi(\mathbf{y})) \gamma(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} \right\} \\ & = -\log \left(\iint \exp(\phi(\mathbf{x}) + \psi(\mathbf{y})) \rho(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} \right) \end{aligned}$$

where $\mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ is the set of probability measures over $\mathbb{R}^d \times \mathbb{R}^d$. Thus (24) is equivalent to

$$\max_{\phi, \psi} \int \phi(\mathbf{x}) \mu(\mathbf{x}) d\mathbf{x} + \int \psi(\mathbf{y}) \pi(\mathbf{y}) d\mathbf{y} - \log \left(\iint \exp(\phi(\mathbf{x}) + \psi(\mathbf{y})) \rho(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} \right).$$

⁵Strong duality holds for this problem under mild regularity conditions, justifying the exchanging of min and max; we refer the reader to Appendix A of [Léo14] for technicalities regarding this point.

Now letting ϕ^*, ψ^* optimize the above expression, and writing $f = \exp(\phi^*)$, $g = \exp(\psi^*)$, first-order optimality conditions now show the optimizer of (24) satisfies $\gamma^*(\mathbf{x}, \mathbf{y}) = \rho(\mathbf{x}, \mathbf{y})f(\mathbf{x})g(\mathbf{y})$, as claimed. Finally, to compare $\mathbf{x}$ -marginals between γ^* and ρ , we compute

$$\frac{\mu(\mathbf{x})}{R_0(\mathbf{x})} = \frac{\int \rho(\mathbf{x}, \mathbf{y})f(\mathbf{x})g(\mathbf{y}) d\mathbf{y}}{\int \rho(\mathbf{x}, \mathbf{y}) d\mathbf{y}} = f(\mathbf{x})\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho}[g(\mathbf{y}) \mid \mathbf{x}],$$

which matches the expression in (26); the calculation for $\mathbf{y}$ -marginals follows similarly. $\square$

The factorizations in (25), (26) extend more generally to characterize the marginal measure P_τ for all $\tau \in [0, 1]$, when the reference path measure R is *Markov*. Recall that R is Markov if for all times $\tau \in [0, 1]$, we have the conditional independence property

$$\mathbf{p}_{[0, \tau]} \perp \mathbf{p}_{[\tau, 1]} \mid \mathbf{p}_\tau, \text{ for } \mathbf{p}_{[0, 1]} \sim R. \quad (27)$$

Lemma 14. *Suppose that R is a Markov reference path measure, i.e., (27) holds, and define $f, g : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ as in (25). For all $\tau \in [0, 1]$, define*

$$f_\tau(\mathbf{z}) := \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0) \mid \mathbf{p}_\tau = \mathbf{z}], \quad g_\tau(\mathbf{z}) := \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[g(\mathbf{p}_1) \mid \mathbf{p}_\tau = \mathbf{z}], \quad (28)$$

let P^* optimally solve (SSB) and define $P_{0|\tau}^*, P_{1|\tau}^*$ to be the laws of $\mathbf{p}_0 \mid \mathbf{p}_\tau$, $\mathbf{p}_1 \mid \mathbf{p}_\tau$ for $\mathbf{p}_{[0, 1]} \sim P^*$ (and $R_{0|\tau}$, $R_{1|\tau}$ similarly for $\mathbf{p}_{[0, 1]} \sim R$). Then P^* is Markov, and for all $0 \leq \sigma \leq \tau \leq 1$,

$$\begin{aligned} P^*(\mathbf{p}_{[0, 1]}) &= f(\mathbf{p}_0)g(\mathbf{p}_1)R(\mathbf{p}_{[0, 1]}), \quad P_\tau^*(\mathbf{z}) = f_\tau(\mathbf{z})g_\tau(\mathbf{z})R_\tau(\mathbf{z}), \\ P_{\sigma|\tau}^*(\mathbf{s} \mid \mathbf{t}) &= \frac{f_\sigma(\mathbf{s})}{f_\tau(\mathbf{t})}R_{\sigma|\tau}(\mathbf{s} \mid \mathbf{t}), \quad P_{\tau|\sigma}^*(\mathbf{t} \mid \mathbf{s}) = \frac{g_\tau(\mathbf{t})}{g_\sigma(\mathbf{s})}R_{\tau|\sigma}(\mathbf{t} \mid \mathbf{s}). \end{aligned} \quad (29)$$

Proof. Combining Lemmas 12 and 13, and following notation (28), we have shown that $P_0^* = f_0g_0R_0$, $P_1^* = f_1g_1R_1$, and $P_{0,1}^*(\mathbf{x}, \mathbf{y}) = R_{0,1}(\mathbf{x}, \mathbf{y})f(\mathbf{x})g(\mathbf{y})$. Multiplying both sides of the last equation by $P_{(0,1)|(\mathbf{x}, \mathbf{y})}^*$ from Lemma 12 then gives the first claim in (29).

We next prove that P^* is Markov. Let $0 \leq \tau \leq 1$ and consider bounded functions a, b such that a is measurable with respect to (the sigma algebra generated by) $\mathbf{p}_{[0, \tau]}$, and b is measurable with respect to $\mathbf{p}_{[\tau, 1]}$. Then

$$\begin{aligned} \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim P^*}[a(\mathbf{p}_{[0, \tau]})b(\mathbf{p}_{[\tau, 1]}) \mid \mathbf{p}_\tau] &= \frac{\mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0)a(\mathbf{p}_{[0, \tau]})b(\mathbf{p}_{[\tau, 1]})g(\mathbf{p}_1) \mid \mathbf{p}_\tau]}{\mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0)g(\mathbf{p}_1) \mid \mathbf{p}_\tau]} \\ &= \frac{\mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0)a(\mathbf{p}_{[0, \tau]}) \mid \mathbf{p}_\tau] \cdot \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[b(\mathbf{p}_{[\tau, 1]})g(\mathbf{p}_1) \mid \mathbf{p}_\tau]}{\mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0) \mid \mathbf{p}_\tau] \cdot \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[g(\mathbf{p}_1) \mid \mathbf{p}_\tau]} \\ &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim P^*}[a(\mathbf{p}_{[0, \tau]}) \mid \mathbf{p}_\tau] \cdot \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim P^*}[b(\mathbf{p}_{[\tau, 1]}) \mid \mathbf{p}_\tau], \end{aligned}$$

where the first and last equalities are general results about conditioning, and the second equality uses the Markov property of R . We note that we do not divide by 0, P^* -almost surely.

The second claim in (29) applies the Markov property of R :

$$\begin{aligned} \frac{P_\tau^*(\mathbf{z})}{R_\tau(\mathbf{z})} &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}\left[\frac{P^*(\mathbf{p}_{[0, 1]})}{R(\mathbf{p}_{[0, 1]})} \mid \mathbf{p}_\tau = \mathbf{z}\right] \\ &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0)g(\mathbf{p}_1) \mid \mathbf{p}_\tau = \mathbf{z}] \\ &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0) \mid \mathbf{p}_\tau = \mathbf{z}] \cdot \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[g(\mathbf{p}_1) \mid \mathbf{p}_\tau = \mathbf{z}] = f_\tau(\mathbf{z})g_\tau(\mathbf{z}). \end{aligned}$$

Similarly, for any $0 \leq \sigma \leq \tau \leq 1$,

$$\begin{aligned} \frac{P_{\sigma, \tau}^*(\mathbf{s}, \mathbf{t})}{R_{\sigma, \tau}(\mathbf{s}, \mathbf{t})} &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}\left[\frac{P^*(\mathbf{p}_{[0, 1]})}{R(\mathbf{p}_{[0, 1]})} \mid \mathbf{p}_{\sigma, \tau} = (\mathbf{s}, \mathbf{t})\right] \\ &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0)g(\mathbf{p}_1) \mid \mathbf{p}_{\sigma, \tau} = (\mathbf{s}, \mathbf{t})] \\ &= \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[f(\mathbf{p}_0) \mid \mathbf{p}_\sigma = \mathbf{s}] \cdot \mathbb{E}_{\mathbf{p}_{[0, 1]} \sim R}[g(\mathbf{p}_1) \mid \mathbf{p}_\tau = \mathbf{t}] = f_\sigma(\mathbf{s})g_\tau(\mathbf{t}), \end{aligned}$$

where the third equality applies that R is Markov. The last two parts of (29) now follow:

$$\begin{aligned} P_{\sigma|\tau}^*(\mathbf{s} | \mathbf{t}) &= \frac{P_{\sigma,\tau}^*(\mathbf{s}, \mathbf{t})}{P_{\tau}^*(\mathbf{t})} = \frac{f_{\sigma}(\mathbf{s})}{f_{\tau}(\mathbf{t})} R_{\sigma|\tau}(\mathbf{s} | \mathbf{t}), \\ P_{\tau|\sigma}^*(\mathbf{t} | \mathbf{s}) &= \frac{P_{\sigma,\tau}^*(\mathbf{s}, \mathbf{t})}{P_{\sigma}^*(\mathbf{s})} = \frac{g_{\tau}(\mathbf{t})}{g_{\sigma}(\mathbf{s})} R_{\tau|\sigma}(\mathbf{t} | \mathbf{s}). \end{aligned}$$

□

Our next step is to derive the infinitesimal generator of P^* . Consistent with (1), we define the infinitesimal generator corresponding to a path measure $P \in \mathcal{P}([0, 1] \times \mathbb{R}^d)$ at time $\tau \in [0, 1]$ by:

$$\mathbb{L}_{\tau}^P u(\mathbf{z}) := \lim_{\eta \searrow 0} \frac{\mathbb{E}_{\mathbf{p}_{[0,1]} \sim P}[u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}) | \mathbf{p}_{\tau} = \mathbf{z}]}{\eta}, \quad (30)$$

and we extend (2) slightly to the setting where $\{v_{\tau}\}_{\tau \in [0,1]}$ are time-inhomogeneous functions:

$$\Gamma_{\tau}^P(u, v_{\tau})(\mathbf{z}) := \frac{1}{2} \lim_{\eta \searrow 0} \frac{\mathbb{E}[(u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}))(v_{\tau+\eta}(\mathbf{p}_{\tau+\eta}) - v_{\tau}(\mathbf{p}_{\tau})) | \mathbf{p}_{\tau} = \mathbf{z}]}{\eta}. \quad (31)$$

Lemma 15. *Following notation in Lemma 14, (30), and (31), for all $u \in \mathcal{F}(\mathbb{R}^d)$ and $\mathbf{z} \in \mathbb{R}^d$,*

$$\mathbb{L}_{\tau}^{P^*} u(\mathbf{z}) = \mathbb{L}_{\tau}^R u(\mathbf{z}) + \frac{2\Gamma_{\tau}^R(g_{\tau}, u)(\mathbf{z})}{g_{\tau}(\mathbf{z})}. \quad (32)$$

Proof. First, by using (29) and the Markov property of R ,

$$\begin{aligned} \mathbb{L}_{\tau}^{P^*} u(\mathbf{z}) &= \lim_{\eta \searrow 0} \frac{1}{\eta} \cdot \frac{\mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[f(\mathbf{p}_0)(u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}))g(\mathbf{p}_1) | \mathbf{p}_{\tau} = \mathbf{z}]}{\mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[f(\mathbf{p}_0)g(\mathbf{p}_1) | \mathbf{p}_{\tau} = \mathbf{z}]} \\ &= \lim_{\eta \searrow 0} \frac{f_{\tau}(\mathbf{z}) \mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[(u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}))g(\mathbf{p}_1) | \mathbf{p}_{\tau} = \mathbf{z}]}{\eta f_{\tau}(\mathbf{z}) g_{\tau}(\mathbf{z})} \\ &= \frac{1}{g_{\tau}(\mathbf{z})} \lim_{\eta \searrow 0} \frac{\mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[(u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}))g(\mathbf{p}_1) | \mathbf{p}_{\tau} = \mathbf{z}]}{\eta}. \end{aligned}$$

However, we also may expand

$$\begin{aligned} &\mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[(u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}))g(\mathbf{p}_1) | \mathbf{p}_{\tau} = \mathbf{z}] \\ &= \int (u(\mathbf{w}) - u(\mathbf{z}))g_{\tau+\eta}(\mathbf{w})R_{\tau+\eta|\tau}(\mathbf{w} | \mathbf{z}) d\mathbf{w} \\ &= \int (u(\mathbf{w}) - u(\mathbf{z}))g_{\tau}(\mathbf{z})R_{\tau+\eta|\tau}(\mathbf{w} | \mathbf{z}) d\mathbf{w} \\ &\quad + \int (u(\mathbf{w}) - u(\mathbf{z}))(g_{\tau+\eta}(\mathbf{w}) - g_{\tau}(\mathbf{z}))R_{\tau+\eta|\tau}(\mathbf{w} | \mathbf{z}) d\mathbf{w} \\ &= g_{\tau}(\mathbf{z}) \mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}) | \mathbf{p}_{\tau} = \mathbf{z}] \\ &\quad + \mathbb{E}_{\mathbf{p}_{[0,1]} \sim R}[(u(\mathbf{p}_{\tau+\eta}) - u(\mathbf{p}_{\tau}))(g_{\tau+\eta}(\mathbf{p}_{\tau+\eta}) - g_{\tau}(\mathbf{p}_{\tau})) | \mathbf{p}_{\tau} = \mathbf{z}] \end{aligned}$$

and combining the above two displays with the definition (31) yields the claim. □

In the remainder of the section, we specialize Perspective 6 to the setting where $\mu = \delta_{\mathbf{0}_d}$, and $R = W$ is the Wiener path measure on $[0, 1]$ (so that $\mu = R_0$ agree). We recall the well-known identities for when $\mathbf{p}_{0,1} \sim W$ follows the heat flow $d\mathbf{p}_{\tau} = dW_{\tau}$:

$$\mathbb{L}_{\tau}^W u = \frac{1}{2} \Delta u, \quad \Gamma_{\tau}^W(u, v_{\tau}) = \frac{1}{2} \langle \nabla u, \nabla v_{\tau} \rangle. \quad (33)$$

The second identity can be verified using an Itô–Taylor expansion. We can now relate this special case of the Schrödinger bridge to stochastic localization.

Theorem 5. Let P^* solve (SSB) when $\mu = \delta_{\mathbf{0}_d}$ and $R = W$ is the Wiener path measure on $[0, 1]$. Then if $\mathbf{v}_\tau \sim P_\tau^*$, the processes $\{\mathbf{c}_t\}_{t \geq 0}$ in (SL-I) and $\{\mathbf{v}_\tau\}_{\tau \in [0, 1]}$ satisfy $\mathbf{c}_t = \frac{1}{1-\tau} \mathbf{v}_\tau$, where $t = \frac{\tau}{1-\tau}$, and the induced π_t in (SL-I) and $P_{1|\tau}^*(\cdot | \mathbf{v}_\tau)$ are identical under this reparameterization.

Proof. We prove this by appealing to Theorem 4. In particular, we show that $P_{1|\tau}^*(\cdot | \mathbf{v}_\tau)$ as defined in the theorem statement agrees with the induced fluctuation measure $\pi_\tau^{\mathbf{v}_\tau}$ defined in (18), and that $\mathbf{v}_\tau$ defined in the theorem statement follows the SDE (19).

To see the first claim, because $\mu = W_0 = \delta_{\mathbf{0}_d}$, the optimal (and only) coupling of μ and π in (24) is the product coupling, so that $f \equiv 1$ and $g = \frac{\pi}{W_1}$ in (25). Now, the last part of (29) gives

$$\begin{aligned} P_{1|\tau}^*(\mathbf{y} | \mathbf{z}) &= \frac{g(\mathbf{y})}{g_\tau(\mathbf{z})} W_{1|\tau}(\mathbf{y} | \mathbf{z}) = \frac{\pi(\mathbf{y})}{g_\tau(\mathbf{z})} \frac{W_{1|\tau}(\mathbf{y} | \mathbf{z})}{W_1(\mathbf{y})} \\ &\propto \exp\left(-\frac{1}{2(1-\tau)} \|\mathbf{y} - \mathbf{z}\|_2^2 + \frac{1}{2} \|\mathbf{y}\|_2^2\right) \pi(\mathbf{y}) \\ &\propto \exp\left(\frac{1}{1-\tau} \langle \mathbf{y}, \mathbf{z} \rangle - \frac{\tau}{2(1-\tau)} \|\mathbf{y}\|_2^2\right) \pi(\mathbf{y}). \end{aligned}$$

This agrees with our definition of the fluctuation measure in (18), when $\mathbf{v}_\tau = \mathbf{z}$.

To see the second claim, we instead show that $L_\tau^{P^*}$, the infinitesimal generator of P_τ^* , agrees with (14). If we can show this, Lemma 11 implies the SDE (19) holds for $\mathbf{v}_\tau$. By combining (32) and the specialized formulas for the heat flow (33), we have

$$L_\tau^{P^*} u = \frac{1}{2} \Delta u + \frac{1}{g_\tau} \langle \nabla g_\tau, \nabla u \rangle = \frac{1}{2} \Delta u + \langle \nabla \log g_\tau, \nabla u \rangle.$$

This is exactly the same as (14), as long as we can show $\nabla \log g_\tau = -\nabla V_\tau$. Applying (28),

$$g_\tau(\mathbf{z}) = \mathbb{E}_{\mathbf{p}_{[0,1]} \sim W} [g(\mathbf{p}_1) | \mathbf{p}_\tau = \mathbf{z}] = \mathbb{E}_{\mathbf{p}_{[0,1]} \sim W} \left[\frac{\pi(\mathbf{p})}{W_1(\mathbf{p})} | \mathbf{p}_\tau = \mathbf{z} \right].$$

Further, $\frac{\pi}{W_1} = \frac{1}{Z} \exp(-V_1)$ as defined in Perspective 5 for some normalizing constant Z . Moreover, conditioned on $\mathbf{p}_\tau = \mathbf{z}$, we know that $\mathbf{p}_1 \sim \mathcal{N}(\mathbf{z}, (1-\tau)\mathbf{I}_d)$ under W . Thus,

$$\begin{aligned} g_\tau(\mathbf{z}) &= \frac{1}{Z} \mathbb{E}_{\mathbf{p}_1 \sim \mathcal{N}(\mathbf{z}, (1-\tau)\mathbf{I}_d)} [\exp(-V_1(\mathbf{p}_1))] \\ &= \frac{1}{Z} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(\mathbf{0}_d, (1-\tau)\mathbf{I}_d)} [\exp(-V_1(\mathbf{z} + \mathbf{g}))] = \frac{1}{Z} \exp(-V_\tau(\mathbf{z})), \end{aligned}$$

at which point it is clear that $\nabla \log g_\tau = -\nabla V_\tau$ pointwise. $\square$

The special case of the Schrödinger bridge to a Wiener reference measure in Theorem 5 is extremely well-studied. In particular, the transition kernels in (29) are an instance of Doob's h -transform [Doo84], and the resulting optimal path measure P^* is induced by the Föllmer drift [Föl85; Föl86]; see further discussion in Section 7, as well as [Leh13] for applications of this construction.

We include a remark demonstrating the broader applicability of these techniques.

Remark 2. Consider the path measure R , defined as follows: $R_0 = \delta_{\mathbf{0}_d}$, and $\mathbf{v}_\tau \sim R_\tau$ satisfies the SDE $d\mathbf{v}_\tau = \mathbf{m}_\tau(\mathbf{v}_\tau) d\tau + \mathbf{\Sigma}_\tau(\mathbf{v}_\tau) d\mathbf{W}_\tau$. In general, it is difficult to calculate the transition densities for P^* , the optimizer to (SSB) with the stated R and $\mu = \delta_{\mathbf{0}_d}$. Nevertheless, using the techniques from this section, we can derive an SDE for $\mathbf{v}_{[0,1]} \sim P^*$: defining $g := \frac{P_1^*}{R_1}$ and g_τ as in (28),

$$d\mathbf{v}_\tau = (\mathbf{m}_\tau(\mathbf{v}_\tau) + \mathbf{\Sigma}_\tau(\mathbf{v}_\tau) \mathbf{\Sigma}_\tau(\mathbf{v}_\tau)^\top \nabla \log g_\tau(\mathbf{v}_\tau)) d\tau + \mathbf{\Sigma}_\tau(\mathbf{v}_\tau) d\mathbf{W}_\tau.$$

For a complete derivation, see Section 3.3 of [Che25].

7 Dynamic Schrödinger bridge

We present the first of two alternative perspectives on the static Schrödinger bridge problem: a dynamic reformulation framed as computing an optimal drift with a target end measure.

In this section, let $L^2([0, 1]; \mathbb{R}^d) := \{\mathbf{u} : [0, 1] \rightarrow \mathbb{R}^d, \int_0^1 \|\mathbf{u}_\tau\|_2^2 d\tau < \infty\}$, let $\{\mathbf{W}_\tau\}_{\tau \in [0, 1]}$ be a Wiener process, and let $\{\mathcal{F}_\tau\}_{\tau \in [0, 1]}$ be the filtration generated by $\{\mathbf{W}_\tau\}_{\tau \in [0, 1]}$.

Perspective 7. Let $\mu, \pi \in \mathcal{P}(\mathbb{R}^d)$, and define

$$\mathcal{U} := \{\mathbf{u} \in L^2([0, 1], \mathbb{R}^d) \mid \mathbf{u} \text{ is adapted to } \{\mathcal{F}_\tau\}_{\tau \in [0, 1]}\}.$$

The induced dynamic Schrödinger bridge problem is:

$$\inf_{\mathbf{u} \in \mathcal{U}} \mathbb{E} \left[\frac{1}{2} \int_0^1 \|\mathbf{u}_\tau\|_2^2 d\tau \right], \quad (\text{DSB})$$

such that $d\mathbf{x}_\tau = \mathbf{u}_\tau d\tau + d\mathbf{W}_\tau$, $\text{law}(\mathbf{x}_0) = \mu$, $\text{law}(\mathbf{x}_1) = \pi$.

We make use of Girsanov's theorem, which is a classic change-of-measure principle on the space $L^2([0, 1]; \mathbb{R}^d)$. We note that the specialization of Girsanov's theorem to deterministic drifts was already introduced as the Cameron-Martin theorem (Lemma 5).

Lemma 16 (Girsanov's theorem; adapted from [Che25, Theorem 3.2.8, Theorem 4.4.1]). Suppose $\{\mathbf{x}_\tau\}_{\tau \in [0, 1]}$ is simultaneously driven by the following SDEs, for $\{\mathbf{u}_\tau\}_{\tau \in [0, 1]}$ adapted to $\{\mathcal{F}_\tau\}_{\tau \in [0, 1]}$:

$$d\mathbf{x}_\tau = \mathbf{u}_\tau d\tau + d\mathbf{W}_\tau, \quad d\mathbf{x}_\tau = d\widetilde{\mathbf{W}}_\tau,$$

with initial condition $\mathbf{x}_0 = \mathbf{a}$. Let $P_{\mathbf{a}}$ be the probability measure under which $\{\mathbf{W}_\tau\}_{\tau \in [0, 1]}$ is a standard Brownian motion, and similarly let $W_{\mathbf{a}}$ be the probability measure under which $\{\widetilde{\mathbf{W}}_\tau\}_{\tau \in [0, 1]}$ is Brownian. Then,⁶ $P_{\mathbf{a}}$ is absolutely continuous with respect to $W_{\mathbf{a}} \in \mathcal{P}([0, 1] \times \mathbb{R}^d)$, and

$$\mathbb{E}_{P_{\mathbf{a}}} \left[\log \frac{P_{\mathbf{a}}}{W_{\mathbf{a}}} \right] = \frac{1}{2} \int_0^1 \mathbb{E}_{P_{\mathbf{a}}} [\|\mathbf{u}_\tau\|_2^2] d\tau.$$

As a last helper definition, when μ is a probability measure on $\mathbb{R}^d$, we define the *generalized Wiener measure* $W_\mu \in \mathcal{P}([0, 1] \times \mathbb{R}^d)$ as $W_\mu := \mathbb{E}_{\mathbf{a} \sim \mu} W_{\mathbf{a}}$. Equivalently, $\mathbf{p}_{[0, 1]} \sim W_\mu$ follows

$$\mathbf{p}_0 \sim \mu, \quad \mathbf{p}_{(0, 1] | 0} \sim W_{\mathbf{p}_0}.$$

Theorem 6. The path measure $P = \text{law}(\{\mathbf{x}_\tau\}_{\tau \in [0, 1]})$ induced by the optimal solution to (DSB) is the optimal solution to (SSB) with $R = W_\mu$.

Proof. Starting from (SSB) and using the chain rule for KL, we have

$$\begin{aligned} \text{KL}(P \parallel R_\mu) &= \text{KL}(P_0 \parallel \mu) + \mathbb{E}_{\mathbf{x}_0 \sim \mu} [\text{KL}(P_{(0, 1] | 0}(\cdot \mid \mathbf{x}_0) \parallel W_{\mathbf{x}_0})] \\ &= \mathbb{E}_{\mathbf{x}_0 \sim \mu} [\text{KL}(P_{(0, 1] | 0}(\cdot \mid \mathbf{x}_0) \parallel W_{\mathbf{x}_0})] \leq \frac{1}{2} \mathbb{E}_P \left[\int_0^1 \|\mathbf{u}_\tau\|_2^2 d\tau \right]. \end{aligned}$$

The second line used $P_0 = \mu$. We then applied Lemma 16: $\{\mathbf{x}_\tau\}_{\tau \in [0, 1]}$ is adapted to the underlying Brownian path measures, so the KL divergence is bounded via the data processing inequality. $\square$

When $\mu = \delta_{\mathbf{0}_d}$ in Theorem 6, the measure W_μ is just the standard Wiener measure. In this special case (the setting of Theorem 5), the optimal drift $\{\mathbf{u}_\tau\}_{\tau \in [0, 1]}$ in (DSB) is called the Föllmer drift. The full discussion around the equivalence of these two areas (and generalizations thereof) contains many subtle points and is a fertile area of mathematical research.

⁶Formally, Girsanov's theorem requires that a technical condition called *Novikov's condition*. This condition can be ignored for our application, as for the KL divergence, it suffices to apply a standard localization argument combined with lower semicontinuity. See [CCLLSZ23, Section 5.2] for a similar argument.

8 Entropic optimal transport

We conclude with a second alternative perspective on the static Schrödinger bridge problem: a connection to entropic optimal transport, following [PC19; CNR25].

Perspective 8. Let $\mu, \pi \in \mathcal{P}(\mathbb{R}^d)$, and let $\varepsilon > 0$. The induced entropic optimal transport problem is:

$$\min_{\gamma \in \Gamma(\mu, \pi)} \left\{ \int \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 d\gamma(\mathbf{x}, \mathbf{y}) + \varepsilon \text{KL}(\gamma \parallel \mu \otimes \pi) \right\}. \quad (\text{EOT})$$

Note that the solution to (EOT) exists and is unique, by strict convexity of KL in its first argument, and compactness of the set $\Gamma(\mu, \pi)$ (defined in Section 6), the set of all couplings of μ and π . Here, $\mu \otimes \pi$ refers to the trivial coupling (product measure) between μ and π .

When $\varepsilon \searrow 0$, (EOT) recovers the classical optimal transport problem [CNR25]. One motivation for studying (EOT) is that while solving classical optimal transport can often be computationally intractable, (EOT) can be efficiently solved via methods such as Sinkhorn’s algorithm [Sin67].

We now relate (EOT) to (SSB) in the case of generalized Wiener reference measures.

Theorem 7. The optimization problems in (EOT) (with $\varepsilon = 1$) and (SSB) (with $R = W_\mu$) are the same, up to an additive constant shift (depending only on μ, ν).

Proof. We start from (SSB). As the initial condition $R_0 = \mu$ is fixed,

$$\begin{aligned} \min_{\gamma \in \Gamma(\mu, \pi)} \text{KL}(\gamma \parallel R_{0,1}) &= \min_{\gamma \in \Gamma(\mu, \pi)} \mathbb{E}_{\mathbf{x} \sim \mu} [\text{KL}(\gamma(\mathbf{y} \mid \mathbf{x}) \parallel R_{1|0}(\mathbf{y} \mid \mathbf{x}))] + \text{KL}(\mu(\mathbf{x}) \parallel R_0(\mathbf{x})) \\ &= \min_{\gamma \in \Gamma(\mu, \pi)} \int -\log R_{0,1}(\mathbf{x}, \mathbf{y}) \gamma(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} + \int \gamma(\mathbf{x}, \mathbf{y}) \log \gamma(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} \end{aligned}$$

The first line used Lemma 12. Now, as $R = W_\mu$ is a generalized Wiener measure, using the explicit form of $\log R_{0,1}(\mathbf{x}, \mathbf{y}) = -\frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2 + C$ for a constant C , and

$$\int \gamma(\mathbf{x}, \mathbf{y}) \log \gamma(\mathbf{x}, \mathbf{y}) d\mathbf{x} d\mathbf{y} - \text{KL}(\gamma \parallel \mu \otimes \pi) = \int \mu(\mathbf{x}) \log \mu(\mathbf{x}) d\mathbf{x} + \int \nu(\mathbf{y}) \log \nu(\mathbf{y}) d\mathbf{y},$$

where the right-hand side is independent of γ , reduces the problem to (EOT). $\square$

Acknowledgements

MSZ was supported by an NSERC CGS-D fellowship.

We thank Sinho Chewi and Ronen Eldan for helpful discussions.

References

- [AHLVXY23] Nima Anari, Yizhi Huang, Tianyu Liu, Thuy-Duong Vuong, Brian Xu, and Katherine Yu. “Parallel Discrete Sampling via Continuous Walks”. In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023*. ACM, 2023, pp. 103–116 (pages 2, 3).
- [AMS22] Ahmed El Alaoui, Andrea Montanari, and Mark Sellke. “Sampling from the Sherrington-Kirkpatrick Gibbs measure via algorithmic stochastic localization”. In: *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*. IEEE, 2022, pp. 323–334 (pages 2, 3).
- [And82] Brian DO Anderson. “Reverse-time diffusion equation models”. In: *Stochastic Processes and their Applications* 12.3 (1982), pp. 313–326 (page 7).
- [BBD24] Roland Bauerschmidt, Thierry Bodineau, and Benoit Dagallier. “Stochastic dynamics and the Polchinski equation: an introduction”. In: *Probability Surveys* 21 (2024), pp. 200–290 (pages 3, 9, 28).
- [BBDD24] Joe Benton, Valentin De Bortoli, Arnaud Doucet, and George Deligiannidis. “Nearly d-Linear Convergence Bounds for Diffusion Models via Stochastic Localization”. In: *The Twelfth International Conference on Learning Representations, ICLR 2024*. OpenReview.net, 2024 (pages 2, 3).
- [BCPSV21] Antonio Blanca, Pietro Caputo, Daniel Parisi, Alistair Sinclair, and Eric Vigoda. “Entropy decay in the Swendsen-Wang dynamics on $\mathbb{Z}^d$ ”. In: *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2021, pp. 1551–1564 (page 27).
- [BÉ06] Dominique Bakry and Michel Émery. “Diffusions hypercontractives”. In: *Séminaire de Probabilités XIX 1983/84: Proceedings*. Springer, 2006, pp. 177–206 (page 28).
- [BE19] Sébastien Bubeck and Ronen Eldan. “The Entropic Barrier: Exponential Families, Log-Concave Geometry, and Self-Concordance”. In: *Math. Oper. Res.* 44.1 (2019), pp. 264–276 (page 25).
- [BGL13] Dominique Bakry, Ivan Gentil, and Michel Ledoux. *Analysis and geometry of Markov diffusion operators*. Vol. 348. Springer Science & Business Media, 2013 (page 9).
- [BL76] Herm Jan Brascamp and Elliott H Lieb. “On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation”. In: *Journal of Functional Analysis* 22.4 (1976), pp. 366–389 (page 28).
- [CCLLSZ23] Sitan Chen, Sinho Chewi, Jerry Li, Yuanzhi Li, Adil Salim, and Anru Zhang. “Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions”. In: *The Eleventh International Conference on Learning Representations, ICLR 2023*. OpenReview.net, 2023 (pages 8, 17).
- [CCSW22] Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. “Improved analysis for a proximal algorithm for sampling”. In: *Conference on Learning Theory*. Vol. 178. Proceedings of Machine Learning Research. PMLR, 2022, pp. 2984–3014 (pages 2, 3, 31).
- [CE22] Yuansi Chen and Ronen Eldan. “Localization Schemes: A Framework for Proving Mixing Bounds for Markov Chains”. In: *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022*. IEEE, 2022, pp. 110–122 (pages 2, 3, 23–28, 31).
- [Che21] Yuansi Chen. “An almost constant lower bound of the isoperimetric coefficient in the KLS conjecture”. In: *Geometric and Functional Analysis* 31.1 (2021), pp. 34–61 (page 2).
- [Che25] Sinho Chewi. *Log-Concave Sampling*. 2025 (pages 2, 16, 17, 28–30, 32, 33).
- [CNR25] Sinho Chewi, Jonathan Niles-Weed, and Philippe Rigollet. *Statistical Optimal Transport*. en. Springer Nature Switzerland, 2025. DOI: [10.1007/978-3-031-85160-5](https://doi.org/10.1007/978-3-031-85160-5). URL: <http://dx.doi.org/10.1007/978-3-031-85160-5> (page 18).

- [Doo84] Joseph L. Doob. *Classical potential theory and its probabilistic counterpart*. Classics in mathematics. Berlin: Springer, 2001 - 1984 (page 16).
- [EKZ22] Ronen Eldan, Frederic Koehler, and Ofer Zeitouni. “A spectral condition for spectral gap: fast mixing in high-temperature Ising models”. In: *Probability theory and related fields* 182.3 (2022), pp. 1035–1051 (pages 2, 3).
- [EL14] Ronen Eldan and Joseph Lehec. “Bounding the norm of a log-concave vector via thin-shell estimates”. In: *Geometric Aspects of Functional Analysis: Israel Seminar (GAFA) 2011-2013*. Springer. 2014, pp. 107–122 (page 2).
- [EL15] Ronen Eldan and James R. Lee. “Talagrand’s Convolution Conjecture on Gaussian Space”. In: *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*. IEEE Computer Society, 2015, pp. 1395–1408 (page 2).
- [Eld13] Ronen Eldan. “Thin shell implies spectral gap up to polylog via a stochastic localization scheme”. In: *Geometric and Functional Analysis* 23.2 (2013), pp. 532–569 (pages 1–3, 5).
- [Eld18] Ronen Eldan. “Gaussian-width gradient complexity, reverse log-Sobolev inequalities and nonlinear large deviations”. In: *Geometric and Functional Analysis* 28.6 (2018), pp. 1548–1596 (page 2).
- [Eld20] Ronen Eldan. “Taming correlations through entropy-efficient measure decompositions with applications to mean-field approximation”. In: *Probability Theory and Related Fields* 176.3 (2020), pp. 737–755 (page 2).
- [Eld22] Ronen Eldan. “Analysis of high-dimensional distributions using pathwise methods”. In: *Proc. Int. Cong. Math.* Vol. 6. 2022, pp. 4246–4270 (page 2).
- [EM22] Ahmed El Alaoui and Andrea Montanari. “An information-theoretic view of stochastic localization”. In: *IEEE Transactions on Information Theory* 68.11 (2022), pp. 7423–7426 (pages 2, 3, 6).
- [EMZ20] Ronen Eldan, Dan Mikulincer, and Alex Zhai. “The CLT in high dimensions: quantitative bounds via martingale embedding”. In: *The Annals of Probability* 48.5 (2020), pp. 2494–2524 (pages 2, 22).
- [ES22] Ronen Eldan and Omer Shamir. “Log concavity and concentration of Lipschitz functions on the Boolean hypercube”. In: *Journal of functional analysis* 282.8 (2022), p. 109392 (page 2).
- [Föl85] H. Föllmer. *An entropy approach to the time reversal of diffusion processes*. en. 1985. DOI: [10.1007/bfb0005070](https://doi.org/10.1007/bfb0005070). URL: <http://dx.doi.org/10.1007/bfb0005070> (page 16).
- [Föl86] H. Föllmer. *Time reversal on Wiener space*. 1986. DOI: [10.1007/bfb0080212](https://doi.org/10.1007/bfb0080212). URL: <http://dx.doi.org/10.1007/bfb0080212> (page 16).
- [Gal16] Jean-François Le Gall. *Brownian motion, martingales, and stochastic calculus*. Springer, 2016 (page 6).
- [Gua24] Qingyang Guan. “A note on Bourgain’s slicing problem”. In: *arXiv preprint arXiv:2412.09075* (2024) (page 2).
- [HJA20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. “Denoising Diffusion Probabilistic Models”. In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020*. 2020 (pages 3, 7).
- [JLV22] Arun Jambulapati, Yin Tat Lee, and Santosh S Vempala. “A slightly improved bound for the KLS constant”. In: *arXiv preprint arXiv:2208.11644* (2022) (page 2).
- [KL22] Bo’az Klartag and Joseph Lehec. “Bourgain’s slicing problem and KLS isoperimetry up to polylog”. In: *Geometric and functional analysis* 32.5 (2022), pp. 1134–1159 (page 2).
- [KL24a] Bo’az Klartag and Joseph Lehec. “Isoperimetric inequalities in high-dimensional convex sets”. In: *arXiv preprint arXiv:2406.01324* (2024) (page 2).
- [KL24b] Boaz Klartag and Joseph Lehec. “Affirmative resolution of Bourgain’s slicing problem using Guan’s bound”. In: *arXiv preprint arXiv:2412.15044* (2024) (page 2).

- [Kla18] Bo'az Klartag. "Eldan's stochastic localization and tubular neighborhoods of complex-analytic sets". In: *The Journal of Geometric Analysis* 28.3 (2018), pp. 2008–2027 (page 2).
- [Kla23] Bo'az Klartag. "Logarithmic bounds for isoperimetry and slices of convex sets". In: *Ars Inveniendi Analytica* (2023) (page 2).
- [KLS95] Ravi Kannan, László Lovász, and Miklós Simonovits. "Isoperimetric Problems for Convex Bodies and a Localization Lemama". In: *Discret. Comput. Geom.* 13 (1995), pp. 541–559 (page 2).
- [KO25] Bo'az Klartag and Or Ordentlich. "The strong data processing inequality under the heat flow". In: *IEEE Transactions on Information Theory* (2025) (page 2).
- [KP21] Bo'az Klartag and Eli Putterman. "Spectral monotonicity under Gaussian convolution". In: *arXiv preprint arXiv:2107.09496* (2021) (pages 2, 3, 8, 31).
- [Leh13] Joseph Lehec. "Representation formula for the entropy and functional inequalities". In: *Annales de l'I.H.P. Probabilités et statistiques* 49.3 (2013), pp. 885–899 (pages 2, 16).
- [Léo14] Christian Léonard. "A survey of the Schrödinger problem and some of its connections with optimal transport". In: *Discrete and Continuous Dynamical Systems – Series A* 34.4 (2014), pp. 1533–1574 (pages 12, 13).
- [LLT23] Holden Lee, Jianfeng Lu, and Yixin Tan. "Convergence of score-based generative modeling for general data distributions". In: *International Conference on Algorithmic Learning Theory*. Vol. 201. Proceedings of Machine Learning Research. PMLR, 2023, pp. 946–985 (page 8).
- [LS01] Robert S. Lipster and Albert N. Shiryaev. *Statistics of Random Processes I: General Theory*. Springer Nature, 2001 (page 6).
- [LST21] Yin Tat Lee, Ruoqi Shen, and Kevin Tian. "Structured Logconcave Sampling with a Restricted Gaussian Oracle". In: *Conference on Learning Theory, COLT 2021*. Vol. 134. Proceedings of Machine Learning Research. PMLR, 2021, pp. 2993–3050 (pages 2, 3, 23, 24).
- [LV17] Yin Tat Lee and Santosh Srinivas Vempala. "Eldan's Stochastic Localization and the KLS Hyperplane Conjecture: An Improved Lower Bound for Expansion". In: *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*. IEEE Computer Society, 2017, pp. 998–1007 (page 2).
- [LV18] Yin Tat Lee and Santosh S. Vempala. "Stochastic localization + Stieltjes barrier = tight bound for log-Sobolev". In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018*. ACM, 2018, pp. 1122–1129 (pages 2, 3).
- [Mon23] Andrea Montanari. "Sampling, diffusions, and stochastic localization". In: *arXiv preprint arXiv:2305.10690* (2023) (pages 2, 3, 8).
- [MS24] Dan Mikulincer and Yair Shenfeld. "The Brownian transport map". In: *Probability Theory and Related Fields* 190.1 (2024), pp. 379–444 (page 2).
- [Øks00] Bernt Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. 5th ed. 2000 (page 12).
- [PC19] Gabriel Peyré and Marco Cuturi. "Computational optimal transport: With applications to data science". In: *Foundations and Trends® in Machine Learning* 11.5-6 (2019), pp. 355–607 (page 18).
- [Pol84] Joseph Polchinski. "Renormalization and effective Lagrangians". In: *Nuclear Physics B* 231.2 (1984), pp. 269–295 (page 9).
- [Rob56] Herbert Robbins. "An empirical Bayes approach to statistics". In: *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*. 1956, pp. 157–163 (page 8).
- [Sch31] Erwin Schrödinger. "Über die Umkehrung der Naturgesetze". In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften* 144 (1931), pp. 144–153 (page 12).

- [Sch32] Erwin Schrödinger. “Sur la théorie relativiste de l’électron et l’interprétation de la mécanique quantique”. In: *Annales de l’institut Henri Poincaré*. Vol. 2. 4. 1932, pp. 269–310 (page 12).
- [Sin67] Richard Sinkhorn. “Diagonal equivalence to matrices with prescribed row and column sums”. In: *The American Mathematical Monthly* 74.4 (1967), pp. 402–405 (page 18).
- [SSKKEP21] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. “Score-Based Generative Modeling through Stochastic Differential Equations”. In: *9th International Conference on Learning Representations, ICLR 2021*. OpenReview.net, 2021 (pages 3, 7, 8).
- [WK74] Kenneth G Wilson and John Kogut. “The renormalization group and the ε expansion”. In: *Physics reports* 12.2 (1974), pp. 75–199 (page 9).

A Anisotropic stochastic localization

In this section we give a more general variant of (SL-I) based on its presentation in [EMZ20], that allows for non-identity “control matrices” to drive the process.

Let $\{\mathbf{C}_t\}_{t \geq 0}$ be an $\mathbb{R}^{d \times d}$ -valued process that is adapted to the filtration $\{\mathcal{F}_t\}_{t \geq 0}$ generated by a Wiener process $\{\mathbf{W}_t\}_{t \geq 0}$, and let $\pi_0 \in \mathcal{P}(\mathbb{R}^d)$ be such that $\mathbf{m}_0 := \mathbb{E}_{\mathbf{x} \sim \pi_0}[\mathbf{x}]$ exists. We define the following tilt-valued process, and $\|\mathbf{x}\|_{\Sigma}^2 := \mathbf{x}^\top \Sigma \mathbf{x}$ for positive semidefinite Σ :

$$\begin{aligned} \mathbf{c}_0 &= \mathbf{0}_d, \quad d\mathbf{c}_t = \mathbf{C}_t \mathbf{C}_t^\top \mathbf{m}_t dt + \mathbf{C}_t d\mathbf{W}_t, \\ \text{where } \mathbf{m}_t &:= \mathbb{E}_{\mathbf{x} \sim \pi_t}[\mathbf{x}], \quad \pi_t(\mathbf{x}) \propto \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{1}{2} \|\mathbf{x}\|_{\Sigma_t}^2\right) \pi_0(\mathbf{x}), \quad d\Sigma_t = \mathbf{C}_t \mathbf{C}_t^\top dt. \end{aligned} \quad (\text{ASL-I})$$

Note that when $\mathbf{C}_t = \mathbf{I}_d$ for all $t \geq 0$, (ASL-I) reduces to (SL-I). Similarly to the duality between Perspectives 1 and 2, there is an equivalent anisotropic measure-valued process to (ASL-I):

$$d\pi_t(\mathbf{x}) = \langle \mathbf{x} - \mathbf{m}_t, \mathbf{C}_t d\mathbf{W}_t \rangle \pi_t(\mathbf{x}), \quad \text{where } \mathbf{m}_t := \mathbb{E}_{\mathbf{x} \sim \pi_t}[\mathbf{x}]. \quad (\text{ASL-II})$$

Theorem 8. *The dynamics of π_t given by (ASL-I) and (ASL-II) are the same.*

Proof. The direction that starts from (ASL-II) and derives (ASL-I) is a straightforward generalization of the proofs of Lemma 3 and Theorem 1. For convenience to the reader, in this proof we provide the opposite direction for the more general processes (ASL-I), (ASL-II).

Let us begin with (ASL-I). Define the normalization constant and its renormalizing potential:

$$Z_t := \int \exp\left(\langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{1}{2} \|\mathbf{x}\|_{\Sigma_t}^2\right) \pi_0(\mathbf{x}) d\mathbf{x}, \quad h_t(\mathbf{x}) := \langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{1}{2} \|\mathbf{x}\|_{\Sigma_t}^2.$$

Then by applying Itô’s lemma (Lemma 1) to $Z_t = \int \exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x}$, where we note that the diffusion term in $dh_t(\mathbf{x})$ is $\langle \mathbf{C}_t^\top \mathbf{x}, d\mathbf{W}_t \rangle$, we have

$$\begin{aligned} dZ_t &= \int d\exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x} \\ &= \int \left(-\frac{1}{2} \|\mathbf{x}\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt + \frac{1}{2} \|\mathbf{x}\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt \right) \exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x} \\ &\quad + \int \langle \mathbf{x}, d\mathbf{c}_t \rangle \exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x} = \int \langle \mathbf{x}, d\mathbf{c}_t \rangle \exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x}. \end{aligned}$$

Thus, applying Itô’s lemma again,

$$\begin{aligned} d\log Z_t &= \frac{1}{Z_t} \left(\int \langle \mathbf{x}, d\mathbf{c}_t \rangle \exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x} \right) - \frac{1}{2Z_t^2} \left\| \int \mathbf{x} \exp(h_t(\mathbf{x})) \pi_0(\mathbf{x}) d\mathbf{x} \right\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt \\ &= \langle \mathbf{m}_t, d\mathbf{c}_t \rangle - \frac{1}{2} \|\mathbf{m}_t\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt. \end{aligned}$$

Now because $\log \pi_t(\mathbf{x}) = \langle \mathbf{c}_t, \mathbf{x} \rangle - \frac{1}{2} \|\mathbf{x}\|_{\Sigma_t}^2 + \log(\pi_0(\mathbf{x})) - \log Z_t$,

$$\begin{aligned} d \log \pi_t(\mathbf{x}) &= \langle \mathbf{x}, d\mathbf{c}_t \rangle - \frac{1}{2} \|\mathbf{x}\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt - d \log Z_t \\ &= \langle \mathbf{x} - \mathbf{m}_t, d\mathbf{c}_t \rangle + \left(\frac{1}{2} \|\mathbf{m}_t\|_{\mathbf{C}_t \mathbf{C}_t}^2 - \frac{1}{2} \|\mathbf{x}\|_{\mathbf{C}_t \mathbf{C}_t}^2 \right) dt \\ &= \langle \mathbf{x} - \mathbf{m}_t, d\mathbf{c}_t - \mathbf{C}_t \mathbf{C}_t^\top \mathbf{m}_t dt \rangle - \frac{1}{2} \|\mathbf{x} - \mathbf{m}_t\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt \\ &= \langle \mathbf{x} - \mathbf{m}_t, \mathbf{C}_t d\mathbf{W}_t \rangle - \frac{1}{2} \|\mathbf{x} - \mathbf{m}_t\|_{\mathbf{C}_t \mathbf{C}_t^\top}^2 dt. \end{aligned}$$

Applying Itô's lemma one final time then yields the SDE (ASL-II) as claimed. $\square$

B Restricted Gaussian dynamics

B.1 The perspective from entropic stability

In this appendix, we discuss an algorithmic application of stochastic localization: the restricted Gaussian dynamics [LST21]. Our presentation follows the *localization schemes* analysis framework of [CE22], and in particular, this appendix serves to replicate the results of [CE22] that are specialized to the restricted Gaussian dynamics. We organize this appendix as follows.

- (1) Appendix B.1.1 introduces the concept of a localization process and describes its induced localization dynamics Markov chain. Specializing this framework to the stochastic localization process (SL-I) gives rise to the restricted Gaussian dynamics (RGD).
- (2) Appendix B.1.2 develops tools for proving that localization processes conserve entropy. We describe the notion of entropic stability (Definition 4) from [CE22], give a sufficient condition for this notion (Lemma 19), and show how it implies entropy conservation (Lemma 21).
- (3) Appendix B.1.3 shows how conservation of entropy, in the sense of Lemma 21, implies a modified log-Sobolev inequality. We then apply this machinery to the restricted Gaussian dynamics (RGD) to derive a mixing time bound for strongly log-concave stationary measures.

B.1.1 Localization schemes

The restricted Gaussian dynamics is an instance of the localization schemes framework of [CE22], specialized to the stochastic localization process (SL-I). We first require a general definition.

Definition 1. A localization process is a measure-valued stochastic process $\{\pi_t\}_{t \geq 0}$ on a state space Ω , which satisfies the following conditions.

- (1) π_t is a probability measure over Ω for all t almost surely.
- (2) For all $A \subseteq \Omega$, $t \mapsto \pi_t(A)$ is a martingale.
- (3) For all $A \subseteq \Omega$, $\lim_{t \rightarrow \infty} \pi_t(A) \in \{0, 1\}$ almost surely.

A localization scheme maps π , a measure over Ω , to a localization process $\{\pi_t\}_{t \geq 0}$ with initial condition $\pi_0 = \pi$. Lastly, the localization dynamics associated with the localization process and some fixed time $T > 0$ is the Markov chain with transition kernel defined as

$$P_T^\pi(A \mid \omega) = \mathbb{E} \left[\frac{\pi_T(\omega) \pi_T(A)}{\pi(\omega)} \right], \text{ for all } \omega \in \Omega, A \subseteq \Omega. \quad (34)$$

We recall a basic fact about localization dynamics.

Lemma 17. For any $T > 0$, π is stationary for the transition kernel P_T^π .

Proof. By Fubini's theorem,

$$\begin{aligned} \int_{\Omega} P_T^\pi(A \mid \omega) \pi(\omega) d\omega &= \int_{\Omega} \mathbb{E} \left[\frac{\pi_T(\omega) \pi_T(A)}{\pi(\omega)} \right] \pi(\omega) d\omega = \mathbb{E} \left[\int_{\Omega} \frac{\pi_T(\omega) \pi_T(A)}{\pi(\omega)} \pi(\omega) d\omega \right] \\ &= \mathbb{E} \left[\pi_T(A) \int_{\Omega} \pi_T(\omega) d\omega \right] = \pi(A), \end{aligned}$$

where we used properties (1) and (2) in Definition 1. $\square$

We have seen a localization scheme for measures π on $\Omega := \mathbb{R}^d$: the stochastic localization process (SL-I), (SL-II). Condition (1) of Definition 1 is Lemma 3, condition (2) is immediate from the dynamics (SL-II), and condition (3) follows from the form of π_t in (SL-I).

We next define the restricted Gaussian dynamics [LST21], the focus of the rest of the section. For a measure π on $\mathbb{R}^d$, the restricted Gaussian dynamics have transitions $\mathbf{x} \rightarrow \mathbf{x}'$ given by

$$\mathbf{y} \sim \mathcal{N}(\mathbf{x}, \eta \mathbf{I}_d), \quad \text{Law}(\mathbf{x}') \propto \exp\left(-\frac{1}{2\eta} \|\mathbf{x}' - \mathbf{y}\|_2^2\right) \pi(\mathbf{x}'), \quad (\text{RGD})$$

for some $\eta > 0$. We observe that (RGD) is an instance of localization dynamics, following [CE22].

Theorem 9. *The Markov chain with transition kernel defined by (RGD) is the same as the localization dynamics associated with the localization scheme given by (SL-I), (SL-II), with $T = \frac{1}{\eta}$.*

Proof. Theorem 2 shows that π_T in (SL-I), (SL-II) can be equivalently described as follows: it is the posterior distribution $\pi(\cdot | \mathbf{c})$, where $\mathbf{w} \sim \pi$ and $\mathbf{c} \sim \mathcal{N}(T\mathbf{w}, T\mathbf{I}_d)$. For notational convenience in this proof, let ρ denote the joint distribution of $(\mathbf{w}, \mathbf{c})$, let μ denote the $\mathbf{c}$ -marginal of ρ , and let $\pi(\cdot | \mathbf{c})$, $\mu(\cdot | \mathbf{w})$ denote the conditional laws of one marginal of ρ given the other.

We now compute the transition density from $\mathbf{x}$ to $\mathbf{x}'$ according to (34):

$$\begin{aligned} P_T^\pi(\mathbf{x}' | \mathbf{x}) &= \mathbb{E}_{(\mathbf{w}, \mathbf{c}) \sim \rho} \left[\frac{\pi(\mathbf{x} | \mathbf{c}) \pi(\mathbf{x}' | \mathbf{c})}{\pi(\mathbf{x})} \right] = \mathbb{E}_{\mathbf{c} \sim \mu} \left[\frac{\rho(\mathbf{x}, \mathbf{c}) \pi(\mathbf{x}' | \mathbf{c})}{\mu(\mathbf{c}) \pi(\mathbf{x})} \right] \\ &= \int \frac{\rho(\mathbf{x}, \mathbf{c}) \pi(\mathbf{x}' | \mathbf{c})}{\pi(\mathbf{x})} d\mathbf{c} = \int \mu(\mathbf{c} | \mathbf{x}) \pi(\mathbf{x}' | \mathbf{c}) d\mathbf{c}. \end{aligned}$$

Thus, an equivalent way to perform this transition is to sample $\mathbf{c} \sim \mu(\cdot | \mathbf{x}) = \mathcal{N}(T\mathbf{x}, T\mathbf{I}_d)$ and then $\mathbf{x}' \sim \pi(\cdot | \mathbf{c})$. Under the variable transformation $\mathbf{y} := \frac{1}{T}\mathbf{c}$ and $\eta := \frac{1}{T}$, it is equivalent to sample $\mathbf{y} \sim \mathcal{N}(\mathbf{x}, \eta \mathbf{I}_d)$ and then $\mathbf{x}' \sim \pi(\cdot | T\mathbf{y})$. By Bayes' theorem,

$$\pi(\mathbf{x}' | T\mathbf{y}) \propto \rho(\mathbf{x}', T\mathbf{y}) \propto \exp\left(-\frac{1}{2T} \|T\mathbf{x}' - T\mathbf{y}\|_2^2\right) \pi(\mathbf{x}') = \exp\left(-\frac{1}{2\eta} \|\mathbf{x}' - \mathbf{y}\|_2^2\right) \pi(\mathbf{x}').$$

This equivalent sampling process is exactly as described in (RGD). $\square$

The rest of the section establishes a mixing time estimate on the dynamics (RGD).

B.1.2 Entropy conservation from entropic stability

In this section, we use the notion of *entropic stability* to establish that entropy is conserved, in a precise sense, along the localization process $\{\pi_t\}_{t \geq 0}$ given by (SL-I). Before proceeding, we introduce a suite of definitions, which are adapted from [CE22, Section 3.2.1].

Definition 2 (Barycentres). *Denote the barycentre $\mathbf{b} : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ of a measure as $\mathbf{b}(\pi) = \mathbb{E}_\pi[\mathbf{x}]$.*

Definition 3 (Exponential tilt). *Define the exponential tilt operator $\mathcal{T} : \mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d) \rightarrow \mathcal{P}(\mathbb{R}^d)$ as:*

$$\mathcal{T}_{\mathbf{y}}\pi(\mathbf{x}) \propto \exp(\langle \mathbf{y}, \mathbf{x} \rangle) \pi(\mathbf{x}),$$

whenever $\exp(\langle \mathbf{y}, \cdot \rangle) \pi$ is integrable over $\mathbb{R}^d$.

Definition 4 (Entropic stability). *We say that $\pi \in \mathcal{P}(\mathbb{R}^n)$ is α -entropically stable if*

$$\frac{1}{2} \|\mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi) - \mathbf{b}(\pi)\|_2^2 \leq \alpha \text{KL}(\mathcal{T}_{\mathbf{y}}\pi \| \pi), \text{ for all } \mathbf{y} \in \mathbb{R}^d.$$

We remark that Definition 4 is used in greater generality in [CE22], with alternative distance functions on the left-hand side. For example, a variant where $\frac{1}{2} \|\cdot - \cdot\|_2^2$ is replaced with a modified KL divergence is used to study the Glauber dynamics on the hypercube.

Let us record an auxiliary lemma. In its statement, we denote the covariance of $\pi \in \mathcal{P}(\mathbb{R}^d)$ by

$$\text{Cov}(\pi) := \mathbb{E}_{\mathbf{x} \sim \pi} [(\mathbf{x} - \mathbf{b}(\pi))(\mathbf{x} - \mathbf{b}(\pi))^\top].$$

Lemma 18 ([BE19, Lemma 1]). *Let $\pi \in \mathcal{P}(\mathbb{R}^d)$ be a measure such that $\mathbf{Cov}(\pi)$ exists and is everywhere invertible, and that the cumulant generating function (a.k.a. log-Laplace transform)*

$$\chi(\mathbf{y}) = \log \mathbb{E}_\pi[\exp(\langle \mathbf{y}, \cdot \rangle)]$$

exists. There exists a unique function $\mathbf{y} : \mathcal{K} := \text{int}(\text{conv}(\text{supp}(\pi))) \rightarrow \mathbb{R}^d$ such that for all $\mathbf{x} \in \mathcal{K}$,

$$\mathbf{b}(\mathcal{T}_{\mathbf{y}(\mathbf{x})}\pi) = \mathbf{x},$$

and furthermore if $\Phi(\mathbf{x}) := \text{KL}(\mathcal{T}_{\mathbf{y}(\mathbf{x})}\pi \parallel \pi)$, Φ and χ are convex conjugates, and on $\mathcal{K}$,

$$\nabla \Phi(\mathbf{x}) = \mathbf{y}(\mathbf{x}), \quad \nabla^2 \Phi(\mathbf{x}) = \mathbf{Cov}(\mathcal{T}_{\mathbf{y}(\mathbf{x})}\pi)^{-1}.$$

Proof. We provide a brief sketch here, deferring more details to [BE19]. The key observation is

$$\nabla \chi(\mathbf{y}) = \mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi), \quad \nabla^2 \chi(\mathbf{y}) = \mathbf{Cov}(\mathcal{T}_{\mathbf{y}}\pi).$$

which follows from a direct computation (and because the mean and covariance are the first two cumulants). Now, for any $\mathbf{y} \in \mathbb{R}^d$, first-order optimality shows

$$\mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi) \in \arg \max_{\mathbf{x} \in \mathbb{R}^d} \langle \mathbf{x}, \mathbf{y} \rangle - \chi(\mathbf{y}),$$

and we can check that

$$\begin{aligned} \chi^*(\mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi)) &= \langle \mathbf{y}, \mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi) \rangle - \chi(\mathbf{y}) \\ &= \int \langle \mathbf{y}, \mathbf{x} \rangle \mathcal{T}_{\mathbf{y}}\pi(\mathbf{x}) \, d\mathbf{x} - \chi(\mathbf{y}) \\ &= \int \mathcal{T}_{\mathbf{y}}\pi(\mathbf{x}) \log \left(\frac{\mathcal{T}_{\mathbf{y}}\pi(\mathbf{x})}{\pi(\mathbf{x})} \right) \, d\mathbf{x} = \text{KL}(\mathcal{T}_{\mathbf{y}}\pi \parallel \pi). \end{aligned}$$

This proves the conjugacy of χ and Φ . The fact that $\nabla \Phi$ is the inverse mapping of $\mathbf{y} \rightarrow \mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi) = \nabla \chi(\mathbf{y})$, and that $\nabla^2 \Phi(\mathbf{x})$ is the inverse of $\nabla^2 \chi(\mathbf{y}(\mathbf{x}))$, then follow from standard properties of convex conjugates. We defer the proof of uniqueness of $\mathbf{y}$ to [BE19]. $\square$

Using Lemma 18, we give a sufficient condition for α -entropic stability.

Lemma 19 ([CE22, Lemma 40]). *If $\mathbf{Cov}(\mathcal{T}_{\mathbf{y}}\pi) \preceq \alpha \mathbf{I}_d$ for all $\mathbf{y} \in \mathbb{R}^d$, π is α -entropically stable.*

Proof. Let $\mathbf{y}(\mathbf{x})$ and $\Phi(\mathbf{x}) := \text{KL}(\mathcal{T}_{\mathbf{y}(\mathbf{x})}\pi \parallel \pi)$ be as in Lemma 18. Then,

$$\nabla^2 \Phi(\mathbf{x}) = \mathbf{Cov}(\mathcal{T}_{\mathbf{y}(\mathbf{x})}\pi)^{-1} \succeq \frac{1}{\alpha} \mathbf{I}_d,$$

where the first equality is Lemma 18, whereas the second is by assumption. Now, define $\tilde{\Phi}(\mathbf{x}) = \frac{1}{2\alpha} \|\mathbf{x} - \mathbf{b}(\pi)\|_2^2$. The lemma statement asks to show $\Phi \geq \tilde{\Phi}$ pointwise. Observe that

$$\nabla \Phi(\mathbf{b}(\pi)) = \nabla \tilde{\Phi}(\mathbf{b}(\pi)) = \mathbf{0}_d, \quad \nabla^2 \Phi \succeq \nabla^2 \tilde{\Phi} = \frac{1}{\alpha} \mathbf{I}_d.$$

Thus, as Φ and $\tilde{\Phi}$ agree up to first order around $\mathbf{b}(\pi)$, the second order term in Φ dominates that of $\tilde{\Phi}$, and $\tilde{\Phi}$ is a quadratic, it follows that $\Phi(\mathbf{x}) \geq \tilde{\Phi}(\mathbf{x})$ everywhere, since

$$\Phi(\mathbf{x}) - \tilde{\Phi}(\mathbf{x}) = \int_0^1 \int_0^t (\mathbf{x} - \mathbf{b}(\pi))^\top \nabla^2(\Phi - \tilde{\Phi})(\mathbf{b}(\pi) + s(\mathbf{x} - \mathbf{b}(\pi))) \cdot (\mathbf{x} - \mathbf{b}(\pi)) \, ds \, dt \geq 0.$$

$\square$

We next use the *maximum entropy principle* to piggyback off Definition 4 and show that entropic stability holds not just for exponential tilts, but all absolutely continuous measures.

Lemma 20. *Suppose π is α -entropically stable, and that $\mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi)$ exists for every $\mathbf{y} \in \mathbb{R}^d$. For every measure ν absolutely continuous with respect to π where $\mathbf{b}(\nu)$ exists, we have*

$$\frac{1}{2}\|\mathbf{b}(\nu) - \mathbf{b}(\pi)\|_2^2 \leq \alpha \text{KL}(\nu \parallel \pi).$$

Proof. We recall the maximum entropy principle: among all absolutely continuous measures to π with a given mean, the one closest to π in KL divergence is an exponential tilt of π . First, define

$$\nu_{\mathbf{w}}^* := \arg \min_{\substack{\nu \in \mathcal{P}(\mathbb{R}^d) \\ \mathbf{b}(\nu) = \mathbf{w}}} \text{KL}(\nu \parallel \pi),$$

and note that by standard variational calculus, defining the Lagrangian $\mathcal{L}(\nu, \lambda) = \text{KL}(\nu \parallel \pi) + \langle \lambda_1, \mathbf{b}(\nu) - \mathbf{w} \rangle + \lambda_2(\mathbb{E}_{\nu}[1] - 1)$ for Lagrange multipliers $\lambda_1 \in \mathbb{R}^d, \lambda_2 \in \mathbb{R}$, we have

$$\log \frac{\nu_{\mathbf{w}}^*}{\pi}(\mathbf{x}) + \lambda_2 + \langle \lambda_1, \mathbf{x} \rangle = 0.$$

This tells us that

$$d\nu_{\mathbf{w}}^*(\mathbf{x}) \propto \exp(-\langle \lambda_1, \mathbf{x} \rangle) d\pi(\mathbf{x}), \quad (35)$$

for some vector λ_1 fixing the barycentre, or in other words, that $\nu_{\mathbf{w}}^*$ is a linear tilt of π .

Then, we note that for any function $g : \mathbb{R}^d \rightarrow \mathbb{R}_+$,

$$\inf_{\mathbf{y} \in \mathbb{R}^d} \frac{\text{KL}(\mathcal{T}_{\mathbf{y}}\pi \parallel \pi)}{g(\mathbf{b}(\mathcal{T}_{\mathbf{y}}\pi))} \stackrel{(i)}{=} \inf_{\mathbf{w} \in \mathbb{R}^d} \inf_{\substack{\nu \in \mathcal{P}(\mathbb{R}^d) \\ \mathbf{b}(\nu) = \mathbf{w}}} \frac{\text{KL}(\nu \parallel \pi)}{g(\mathbf{w})} = \inf_{\nu \in \mathcal{P}(\mathbb{R}^d)} \frac{\text{KL}(\nu \parallel \pi)}{g(\mathbf{b}(\nu))},$$

where (i) uses (35).⁷ Applying this to $g(\mathbf{x}) := \frac{1}{2}\|\mathbf{x} - \mathbf{b}(\pi)\|_2^2$ concludes the proof. $\square$

We are now ready to state the main result of this section, which shows that on average over an entropically stable localization scheme, the entropy (defined below) is conserved:

$$\text{Ent}_{\pi}[f] := \mathbb{E}_{\pi}[f \log f] - \mathbb{E}_{\pi}[f] \log \mathbb{E}_{\pi}[f], \text{ for } \pi \in \mathcal{P}(\mathbb{R}^d), f : \mathbb{R}^d \rightarrow \mathbb{R}_+.$$

We will use the following identity: for $\nu \propto \pi f$,

$$\text{KL}(\nu \parallel \pi) = \mathbb{E}_{\nu} \left[\log \frac{f}{\mathbb{E}_{\pi} f} \right] = \mathbb{E}_{\nu} [\log f] - \log \mathbb{E}_{\pi}[f] = \frac{\text{Ent}_{\pi}[f]}{\mathbb{E}_{\pi}[f]}. \quad (36)$$

Lemma 21 ([CE22, Proposition 39]). *Let $\{\pi_t\}_{t \geq 0}$ be a localization process for $\pi \in \mathcal{P}(\mathbb{R}^d)$. Fix $T > 0$, and suppose that π_t is almost surely α_t -entropically stable for all $t \in [0, T]$. Then,*

$$\mathbb{E}[\text{Ent}_{\pi_T}[f]] \geq \exp\left(-\int_0^T \alpha_t dt\right) \text{Ent}_{\pi}[f].$$

Proof. Fix a measurable test function $f : \mathbb{R}^d \rightarrow \mathbb{R}_+$. Define for $t \geq 0$

$$\nu_t(\mathbf{x}) \propto f(\mathbf{x}) \pi_t(\mathbf{x}).$$

If we consider the process $M_t := \mathbb{E}_{\pi_t}[f]$, then (SL-II) gives

$$dM_t = \mathbb{E}_{\pi_t}[f \langle \mathbf{x} - \boldsymbol{\mu}_t, d\mathbf{W}_t \rangle] = M_t \langle \mathbf{b}(\nu_t) - \mathbf{b}(\pi_t), d\mathbf{W}_t \rangle,$$

and so M_t is a martingale. Itô's lemma then tells us that

$$d(M_t \log M_t) = \frac{1}{2} M_t \|\mathbf{b}(\nu_t) - \mathbf{b}(\pi_t)\|_2^2 dt + \text{martingale}.$$

⁷Here, we need to restrict to measures with finite barycentres.

Thus, we can compute the entropy differential,

$$\begin{aligned} d\text{Ent}_{\pi_t}[f] &= d\mathbb{E}_{\pi_t}[f \log f] - d(M_t \log M_t) \\ &= -\frac{1}{2}M_t \|\mathbf{b}(\nu_t) - \mathbf{b}(\pi_t)\|_2^2 dt + \text{martingale}, \end{aligned}$$

where $d\mathbb{E}_{\pi_t}[f \log f]$ is a martingale because π_t is a martingale pointwise. Lemma 20 and (36) imply

$$d\text{Ent}_{\pi_t}[f] \geq -\alpha_t M_t \text{KL}(\nu_t \parallel \pi_t) dt + \text{martingale} = -\alpha_t \text{Ent}_{\pi_t}[f] + \text{martingale}.$$

Taking expectations, and applying Grönwall's inequality, we conclude that

$$\mathbb{E}[\text{Ent}_{\pi_T}[f]] \geq \exp\left(-\int_0^T \alpha_t dt\right) \text{Ent}_{\pi}[f].$$

□

B.1.3 Log-Sobolev inequality from entropy conservation

We now show how to use the entropy conservation bound from Lemma 21 to establish rapid mixing of (RGD). Recall that a transition kernel $\mathbf{P}$ with stationary measure π satisfies a (modified) *log-Sobolev inequality* (LSI) with constant C_{LS} if the following inequality holds:

$$C_{\text{LS}}(\mathbf{P}) := 1 - \sup_{f: \Omega \rightarrow \mathbb{R}_+} \frac{\text{Ent}_{\pi}[\mathbf{P}f]}{\text{Ent}_{\pi} f}.$$

Lower bounding the LSI constant is useful because it implies rapid mixing. Specifically, a Markov chain with transition kernel $\mathbf{P}$, initialized at a distribution μ with $\frac{\mu}{\pi} \leq \beta$ pointwise, mixes to within ε in total variation of π in $O(\frac{1}{C_{\text{LS}}(\mathbf{P})} \log(\frac{\log \beta}{\varepsilon}))$ steps; see e.g., Lemma 2.4 in [BCPSV21]. This is because $\text{KL}(\cdot \parallel \pi)$ improves by a $1 - C_{\text{LS}}(\mathbf{P})$ factor in each step.

We next reproduce a key observation of [CE22]: entropy conservation of a localization process, as in Lemma 21, lower bounds the LSI constant of the associated localization dynamics.

Lemma 22 ([CE22, Proposition 19]). *Assume that $\{\pi_t\}_{t \geq 0}$ is a localization process associated to π , and that $\mathbf{P}_T^\pi$ is the transition kernel (34) associated to its localization dynamics. Then,*

$$C_{\text{LS}}(\mathbf{P}_T^\pi) \geq \inf_{f: \Omega \rightarrow \mathbb{R}_+} \frac{\mathbb{E}[\text{Ent}_{\pi_T}[f]]}{\text{Ent}_{\pi}[f]}.$$

Proof. Throughout this proof, $\mathbb{E}$ with no subscript implies that the expectation is over the randomness used in defining π_T (i.e., the localization process), and we let $\mathbf{P} := \mathbf{P}_T^\pi$ for short.

Let $\mathbf{x} \in \mathbb{R}^d$ be arbitrary. By Jensen's inequality applied to the convex function $c \rightarrow c \log c$, with random variable $\mathbb{E}_{\pi_T}[f]$ and measure $\mathbb{E}[\frac{\pi_T(\mathbf{x})}{\pi(\mathbf{x})} \cdot]$ (as $\mathbb{E}[\frac{\pi_T(\mathbf{x})}{\pi(\mathbf{x})}] = 1$), we have that

$$\mathbb{E}\left[\frac{\pi_T(\mathbf{x})}{\pi(\mathbf{x})} \mathbb{E}_{\pi_T}[f]\right] \log\left(\mathbb{E}\left[\frac{\pi_T(\mathbf{x})}{\pi(\mathbf{x})} \mathbb{E}_{\pi_T}[f]\right]\right) \leq \mathbb{E}\left[\frac{\pi_T(\mathbf{x})}{\pi(\mathbf{x})} \left(\mathbb{E}_{\pi_T}[f] \log \mathbb{E}_{\pi_T}[f]\right)\right].$$

Integrating over $\mathbf{x} \sim \pi$, this implies via Fubini's theorem that

$$\begin{aligned} \mathbb{E}_{\pi}[\mathbf{P}f \log \mathbf{P}f] &= \mathbb{E}_{\pi}\left[\mathbb{E}\left[\frac{\pi_T}{\pi} \mathbb{E}_{\pi_T}[f]\right] \log \mathbb{E}\left[\frac{\pi_T}{\pi} \mathbb{E}_{\pi_T}[f]\right]\right] \\ &\leq \mathbb{E}_{\pi}\left[\mathbb{E}\left[\frac{\pi_T}{\pi} \left(\mathbb{E}_{\pi_T}[f] \log \mathbb{E}_{\pi_T}[f]\right)\right]\right] \\ &= \mathbb{E}\left[\mathbb{E}_{\pi}\left[\frac{\pi_T}{\pi} \mathbb{E}_{\pi_T}[f] \log \mathbb{E}_{\pi_T}[f]\right]\right] = \mathbb{E}\left[\mathbb{E}_{\pi_T}[f] \log \mathbb{E}_{\pi_T}[f]\right]. \end{aligned}$$

Finally,

$$\begin{aligned}
\frac{\mathbb{E}[\text{Ent}_{\pi_T}[f]]}{\text{Ent}_{\pi}[f]} &= \frac{\mathbb{E}[\mathbb{E}_{\pi_T}[f \log f]] - \mathbb{E}[\mathbb{E}_{\pi_T} f \log \mathbb{E}_{\pi_T} f]}{\text{Ent}_{\pi}[f]} \\
&\leq \frac{\mathbb{E}[\mathbb{E}_{\pi_T}[f \log f]] - \mathbb{E}_{\pi}[P f \log P f]}{\text{Ent}_{\pi}[f]} \\
&= \frac{(\mathbb{E}_{\pi}[f \log f] - \mathbb{E}_{\pi}[f] \log \mathbb{E}_{\pi}[f]) - (\mathbb{E}_{\pi}[P f \log P f] - \mathbb{E}_{\pi}[P f] \log \mathbb{E}_{\pi}[P f])}{\text{Ent}_{\pi}[f]} \\
&= 1 - \frac{\text{Ent}_{\pi}[P f]}{\text{Ent}_{\pi}[f]},
\end{aligned}$$

where we used that π is stationary for P , as well as $\mathbb{E}_{\pi_T} = \pi$, in the third line. This concludes the proof upon infimizing over f on both sides. $\square$

Our final result applies the development thus far to the localization dynamics (RGD).

Theorem 10 ([CE22, Theorem 58]). *If π is α -strongly log-concave (i.e., $-\nabla^2 \log \pi \succeq \alpha \mathbf{I}_d$ pointwise on $\mathbb{R}^d$), then the Markov chain with transition kernel P defined by (RGD) satisfies $C_{\text{LS}}(P) \geq \frac{\alpha}{\alpha + \eta - 1}$.*

Proof. Throughout this proof, let $T := \frac{1}{\eta}$. From the definition of the localization process (SL-I), the localized measure π_t is always of the form

$$\pi_t \propto \exp\left(\langle \mathbf{c}_t, \cdot \rangle - \frac{t}{2} \|\cdot\|_2^2\right) \pi,$$

and therefore $\nabla^2 \log \pi_t \preceq -(\alpha + t) \mathbf{I}_d$ pointwise, as linear tilts do not affect the second derivative matrix. The famous Brascamp-Lieb inequality [BL76] then states that for all $\mathbf{y} \in \mathbb{R}^d$,

$$\|\mathbf{Cov}(\mathcal{J}_{\mathbf{y}} \pi_t)\|_{\text{op}} \leq \frac{1}{\alpha + t}.$$

Lemma 19 now implies that π_t is $\frac{1}{\alpha + t}$ -entropically stable. Then, applying Lemma 21 gives

$$\mathbb{E}[\text{Ent}_{\pi_T}[f]] \geq \exp\left(-\int_0^T \frac{1}{\alpha + t} dt\right) \text{Ent}_{\pi}[f] = \frac{\alpha}{\alpha + T} \text{Ent}_{\pi}[f],$$

for all $f : \mathbb{R}^d \rightarrow \mathbb{R}_+$. Finally, Lemma 22 tells us that $C_{\text{LS}}(P) \geq \frac{\alpha}{\alpha + T}$, as desired. $\square$

B.2 The perspective from renormalization

In this appendix, we show how the Polchinski semigroup perspective from Section 5 can be used to derive log-Sobolev inequalities and entropic stability estimates for the localization scheme (Definition 1) induced by the stochastic localization process (Perspective 1).

For consistency with Section 5, in this section we let $\tau \in [0, 1]$. Recall that under the time change $t \leftarrow \frac{\tau}{1-\tau}$, Theorem 4 shows that the renormalized measure $\pi^{\mathbf{v}\tau}$ defined in (18), (19) is identical in distribution to the localized measure π_t used to define the stochastic localization scheme.

We begin with a technical lemma from [BBD24]. The proof follows similarly to the classical Bakry-Émery theorem [BÉ06] (see Theorem 1.2.30 in [Che25] for a recent exposition): the key steps use the Bochner formula and a curvature condition implied by strong log-concavity.

Lemma 23 ([BBD24, Lemma 3.10]). *Following notation in Perspective 5, assume that there exist $\{\lambda_{\tau}\}_{\tau \in [0, 1]}$ such that for all $\tau \in [0, 1]$, $\nabla^2 V_{\tau} \succeq \lambda_{\tau} \mathbf{I}_d$. Then for all $0 \leq \sigma \leq \tau \leq 1$ and $f \in \mathcal{F}(\mathbb{R}^d)$,*

$$\left\| \nabla \sqrt{P_{\sigma, \tau} f} \right\|_2^2 \leq \exp(-2(\Lambda_{\sigma} - \Lambda_{\tau})) P_{\sigma, \tau} \left\| \nabla \sqrt{f} \right\|_2^2, \text{ where } \Lambda_{\tau} := \int_{\tau}^1 \lambda_{\sigma} d\sigma.$$

We next derive functional inequalities for the renormalized measures in Perspective 5. Recall that we say a measure π satisfies a *Poincaré inequality* with constant α if for all $f \in \mathcal{F}(\mathbb{R}^d)$,

$$\text{Var}_{\pi}[f] \leq \frac{1}{\alpha} \mathbb{E}_{\pi} \left[\left\| \nabla f \right\|_2^2 \right], \quad (37)$$

and similarly, π satisfies a *log-Sobolev inequality* with constant α if for all $f \in \mathcal{F}(\mathbb{R}^d)$,

$$\text{Ent}_\pi[f] \leq \frac{2}{\alpha} \mathbb{E}_\pi \left[\left\| \nabla \sqrt{f} \right\|_2^2 \right]. \quad (38)$$

We remark that a standard linearization argument shows that a log-Sobolev inequality implies a Poincaré inequality with the same constant (Exercise 1.7, [Che25]).

Proposition 1. *Following notation in Perspective 5, assume that there exist $\{\lambda_\tau\}_{\tau \in [0,1]}$ such that for all $\tau \in [0,1]$, $\nabla^2 V_\tau \succeq \lambda_\tau \mathbf{I}_d$. Define for all $\tau \in [0,1]$,*

$$\Lambda_\tau := \int_\tau^1 \lambda_\sigma d\sigma, \quad \frac{1}{\gamma_\tau} := \int_0^\tau \exp(-2(\Lambda_\sigma - \Lambda_\tau)) d\sigma.$$

Then for all $\tau \in [0,1]$ and $f \in \mathcal{F}(\mathbb{R}^d)$,

$$\text{Ent}_{\nu_\tau}[f] \leq \frac{2}{\gamma_\tau} \mathbb{E}_{\nu_\tau} \left[\left\| \nabla \sqrt{f} \right\|_2^2 \right], \quad (39)$$

and

$$\left(1 - \exp \left(- \int_\tau^1 \gamma_\sigma d\sigma \right) \right) \text{Ent}_\pi[f] \leq \mathbb{E}_{\nu_\tau} [\text{Ent}_{\pi_\tau}[f]]. \quad (40)$$

Proof of (39). Here we prove the log-Sobolev inequality (39), postponing (40) until later. We will develop a more general claim that subsumes (40) as well as an analog for variance stability.

Write $\Phi(x) = x \log x$. By compactness of f 's support, we have

$$\lim_{\sigma \rightarrow 0} \mathbb{E}_{\nu_\sigma} [\Phi(P_{\sigma,\tau} f)] = \Phi(\mathbb{E}_{\nu_\tau}[f]),$$

recalling that $P_{0,\tau} f = \mathbb{E}_{\nu_\tau}[f]$ (Lemma 9). Then,

$$\text{Ent}_{\nu_\tau}[f] = \mathbb{E}_{\nu_\tau}[\Phi(f)] - \Phi(\mathbb{E}_{\nu_\tau}[f]) = \int_0^\tau \frac{\partial}{\partial \sigma} \mathbb{E}_{\nu_\sigma} [\Phi(P_{\sigma,\tau} f)] d\sigma.$$

Then by first applying the chain rule, and then recalling the action of $\mathbf{L}_\sigma$ given in Lemma 10,

$$\begin{aligned} \partial_\sigma \mathbb{E}_{\nu_\sigma} [\Phi(P_{\sigma,\tau} f)] &= \mathbb{E}_{\nu_\sigma} \left[\mathbf{L}_\sigma(\Phi(P_{\sigma,\tau} f)) + \Phi'(P_{\sigma,\tau} f) \frac{\partial}{\partial \sigma} P_{\sigma,\tau} f \right] \\ &= \mathbb{E}_{\nu_\sigma} \left[\frac{1}{2} \Phi''(P_{\sigma,\tau} f) \|\nabla P_{\sigma,\tau} f\|_2^2 + \Phi'(P_{\sigma,\tau} f) \mathbf{L}_\sigma P_{\sigma,\tau} f + \Phi'(P_{\sigma,\tau} f) \frac{\partial}{\partial \sigma} P_{\sigma,\tau} f \right] \\ &= \frac{1}{2} \mathbb{E}_{\nu_\sigma} \left[\Phi''(P_{\sigma,\tau} f) \|\nabla P_{\sigma,\tau} f\|_2^2 \right] = 2 \mathbb{E}_{\nu_\sigma} \left[\left\| \nabla \sqrt{P_{\sigma,\tau} f} \right\|_2^2 \right], \end{aligned}$$

so that

$$\begin{aligned} \text{Ent}_{\nu_\tau}[f] &= 2 \int_0^\tau \mathbb{E}_{\nu_\sigma} \left[\left\| \nabla \sqrt{P_{\sigma,\tau} f} \right\|_2^2 \right] d\sigma \\ &\leq 2 \int_0^\tau \exp(-2(\Lambda_\sigma - \Lambda_\tau)) \mathbb{E}_{\nu_\sigma} \left[P_{\sigma,\tau} \left\| \nabla \sqrt{f} \right\|_2^2 \right] d\sigma = \frac{2}{\gamma_\tau} \mathbb{E}_{\nu_\tau} \left[\left\| \nabla \sqrt{f} \right\|_2^2 \right]. \end{aligned}$$

The only inequality was Lemma 23, and then we used the definition of γ_τ and Lemma 9. $\square$

We have thus far shown the bound (39), i.e., a log-Sobolev inequality for ν_τ . As discussed previously, this also implies a Poincaré inequality holds with the same constant.

We next describe how to obtain the stability estimate (40) in somewhat greater generality. Let Φ be a convex function, and let Ψ be such that

$$\Psi_\nu[f] := \mathbb{E}_\nu[\Phi(f)] - \Phi(\mathbb{E}_\nu[f]). \quad (41)$$

For example, when $\Phi(x) = x \log x$, then $\Psi = \text{Ent}$, and when $\Phi(x) = x^2$, then $\Psi = \text{Var}$. From this definition and the fact that $\pi = \nu_1$ in Perspective 5, we also derive the decomposition

$$\begin{aligned}\Psi_\pi[f] &= \mathbb{E}_{\nu_1}[\Phi(f)] - \Phi(\mathbb{E}_{\nu_1}[f]) \\ &= \mathbb{E}_{\nu_\tau}[\mathbb{P}_{\tau,1}\Phi(f)] - \mathbb{E}_{\nu_\tau}[\Phi(\mathbb{P}_{\tau,1}f)] + \mathbb{E}_{\nu_\tau}[\Phi(\mathbb{P}_{\tau,1}f)] - \Phi(\mathbb{E}_{\nu_\tau}[\mathbb{P}_{\tau,1}f]) \\ &= \mathbb{E}_{\nu_\tau}[\mathbb{P}_{\tau,1}\Phi(f) - \Phi(\mathbb{P}_{\tau,1}f)] + \Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f] \\ &= \mathbb{E}_{\nu_\tau}[\Psi_{\pi_\tau^\vee}[f]] + \Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f]\end{aligned}\tag{42}$$

for any $\tau \in [0, 1]$, where we used Lemma 9 in the second line, and (20) twice in the last.

Proposition 2. *Let Φ be a convex function, and define Ψ as in (41). Following notation in Perspective 5, assume that for all $\tau \in [0, 1]$, ν_τ satisfies the following functional Ψ -inequality:*

$$\Psi_{\nu_\tau}[f] \leq \frac{1}{2\gamma_\tau} \mathbb{E}_{\nu_\tau}[\Phi''(f) \|\nabla f\|_2^2],\tag{43}$$

for some $\{\gamma_\tau\}_{\tau \in [0,1]}$. Then π satisfies the following Ψ -stability bound:

$$\left(1 - \exp\left(-\int_\tau^1 \gamma_\sigma d\sigma\right)\right) \Psi_\pi[f] \leq \mathbb{E}_{\nu_\tau}[\Psi_{\pi_\tau^\vee}[f]]$$

Proof. First, note that for all $\tau \in [0, 1]$,

$$\begin{aligned}\Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f] &= \mathbb{E}_{\nu_\tau}[\Phi(\mathbb{P}_{\tau,1}f)] - \Phi(\mathbb{E}_{\nu_\tau}[\mathbb{P}_{\tau,1}f]) \\ &= \mathbb{E}_{\nu_\tau}[\Phi(\mathbb{P}_{\tau,1}f)] - \Phi(\mathbb{E}_\pi[f]) \\ \implies \partial_\tau \Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f] &= \partial_\tau \mathbb{E}_{\nu_\tau}[\Phi(\mathbb{P}_{\tau,1}f)].\end{aligned}$$

Our earlier proof of (39) showed

$$\partial_\tau \mathbb{E}_{\nu_\tau}[\Phi(\mathbb{P}_{\tau,1}f)] = \frac{1}{2} \mathbb{E}_{\nu_\tau}[\Phi''(\mathbb{P}_{\tau,1}f) \|\nabla \mathbb{P}_{\tau,1}f\|_2^2].$$

Although we derived this directly for $\Phi(x) = x \log x$, we never used the form of Φ . Indeed, this is also a consequence of the *diffusion chain rule* [Che25, Definition 2.2.13], as the carré du champ of $\mathbb{P}_{\sigma,\tau}$ is $\frac{1}{2}\langle \nabla, \nabla \rangle$. Then, the Ψ -inequality (43) implies

$$\partial_\tau \Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f] = \frac{1}{2} \mathbb{E}_{\nu_\tau}[\Phi''(\mathbb{P}_{\tau,1}f) \|\nabla \mathbb{P}_{\tau,1}f\|_2^2] \geq \gamma_\tau \Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f].$$

Defining $E(\tau) := \Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f]$, we have

$$\log E(1) - \log E(\tau) = \int_\tau^1 \frac{E'(\sigma)}{E(\sigma)} d\sigma \geq \int_\tau^1 \gamma_\sigma d\sigma.$$

Rearranging,

$$\Psi_{\nu_\tau}[\mathbb{P}_{\tau,1}f] \leq \exp\left(\int_1^\tau \gamma_\sigma d\sigma\right) \Psi_{\nu_1}[f].$$

Now using $\nu_1 = \pi$ and (42) we obtain the conclusion, which also recovers (40). $\square$

We now derive the consequences of Propositions 1 and 2 for strongly log-concave π .

Corollary 1. *Let $\pi \propto \exp(-V)$ be α -strongly log-concave. Then π satisfies Poincaré and log-Sobolev inequalities (37), (38) with constant α . Moreover, for all $\tau \in [0, 1]$, we have the Ψ -stability bounds*

$$\begin{aligned}\frac{\alpha}{\alpha + \frac{\tau}{1-\tau}} \text{Ent}_\pi[f] &\leq \mathbb{E}_{\nu_\tau}[\text{Ent}_{\pi_\tau^\vee}[f]], \\ \frac{\alpha}{\alpha + \frac{\tau}{1-\tau}} \text{Var}_\pi[f] &\leq \mathbb{E}_{\nu_\tau}[\text{Var}_{\pi_\tau^\vee}[f]].\end{aligned}$$

Proof. Our first step is to obtain bounds λ_τ for use of Propositions 1 and 2. Recall from Lemma 11 that $\nabla V_\tau(\mathbf{v}) = \frac{1}{1-\tau}(\mathbf{v} - \mathbf{m}_\tau)$ with $\mathbf{m}_\tau = \mathbb{E}_{\mathbf{x} \sim \pi_\tau^\mathbf{v}}[\mathbf{x}]$. Then we can compute $\nabla_{\mathbf{v}} \mathbf{m}_\tau$ as

$$\begin{aligned} \nabla_{\mathbf{v}} \mathbf{m}_\tau &= \nabla_{\mathbf{v}} \left(\frac{\int \mathbf{x} \exp\left(\frac{1}{1-\tau} \langle \mathbf{v}, \mathbf{x} \rangle - \frac{\tau}{2(1-\tau)} \|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}) \, d\mathbf{x}}{\int \exp\left(\frac{1}{1-\tau} \langle \mathbf{v}, \mathbf{x} \rangle - \frac{\tau}{2(1-\tau)} \|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}) \, d\mathbf{x}} \right) \\ &= \frac{1}{1-\tau} \left(\frac{\int \mathbf{x} \mathbf{x}^\top \exp\left(\frac{1}{1-\tau} \langle \mathbf{v}, \mathbf{x} \rangle - \frac{\tau}{2(1-\tau)} \|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}) \, d\mathbf{x}}{\int \exp\left(\frac{1}{1-\tau} \langle \mathbf{v}, \mathbf{x} \rangle - \frac{\tau}{2(1-\tau)} \|\mathbf{x}\|_2^2\right) \pi(\mathbf{x}) \, d\mathbf{x}} - \mathbf{m}_\tau \mathbf{m}_\tau^\top \right) \\ &= \frac{1}{1-\tau} \mathbf{Cov}(\pi_\tau^\mathbf{v}). \end{aligned} \tag{44}$$

Thus we have the lower bound $\lambda_\tau = \frac{\alpha-1}{(1-\tau)\alpha+\tau}$ which follows from

$$\begin{aligned} \nabla^2 V_\tau &= \frac{1}{1-\tau} \mathbf{I}_d - \frac{1}{(1-\tau)^2} \mathbf{Cov}(\pi_\tau^\mathbf{v}) \\ &\succeq \left(\frac{1}{1-\tau} - \frac{1}{(1-\tau)^2} \cdot \frac{1}{\alpha + \frac{\tau}{1-\tau}} \right) \mathbf{I}_d = \frac{\alpha-1}{(1-\tau)\alpha+\tau} \mathbf{I}_d, \end{aligned}$$

where the inequality is due to Brascamp-Lieb and V being α -strongly convex. Then

$$\Lambda_\tau = \int_\tau^1 \lambda_\sigma \, d\sigma = \log((1-\alpha)\tau + \alpha),$$

and

$$\frac{1}{\gamma_\tau} = \int_0^\tau e^{-2(\Lambda_\sigma - \Lambda_\tau)} \, d\sigma = \frac{\tau((1-\alpha)\tau + \alpha)}{\alpha}.$$

In particular, $\gamma_1 = \alpha$, so Proposition 1 demonstrates a log-Sobolev inequality (38) holds with constant α . As discussed, this implies a Poincaré inequality (37) with the same constant. Finally, the Ψ -stability claims follow from our bound on γ_τ and Proposition 2. $\square$

Notably, under the reparameterization in Theorem 4, i.e., the time-change $t \leftarrow \frac{\tau}{1-\tau}$ and equivalence of localized measures, Theorem 10 and Corollary 1 yield the same bound. The generality of Proposition 2 further establishes an analogous variance stability bound, yielding a contraction $\chi^2(\cdot \parallel \pi)$ along the dynamics (RGD), similar to our derivation in Appendix B.1.3 for $\text{KL}(\cdot \parallel \pi)$. We defer a formal proof that variance stability implies χ^2 mixing to Proposition 19, [CE22].

B.3 The perspective from time reversal

In this appendix, we show how our perspective in Section 4 of the stochastic localization process as a backwards heat flow also yields quantitative estimates on how quickly the dynamics (RGD) converge. The computations here are replicated from [KP21; CCSW22].

Formally, Section 4 considers a time reversal of the OU process, $d\mathbf{x}_t = -\mathbf{x}_t dt + \sqrt{2} d\mathbf{B}_t$. Here, we instead reverse the *heat flow*, which advances $d\mathbf{x}_t = d\mathbf{B}_t$, which behaves the same under reparameterization. If $\mathbf{x}_0 \sim \pi \propto \exp(-V)$, we then have that $\mathbf{x}_t = \pi * \gamma_t$, where γ_t is the density of the Gaussian $\mathcal{N}(\mathbf{0}_d, t\mathbf{I}_d)$. We write this convolved measure as $\pi \mathbf{Q}_t$.

The forwards and backwards processes we consider in this appendix are as follows, defined in a time range $t \in [0, \eta]$ for some $\eta \in (0, \infty)$:

$$d\mathbf{x}_t = d\mathbf{B}_t \tag{45}$$

$$d\mathbf{x}_t^\leftarrow = \nabla \log \pi \mathbf{Q}_{\eta-t}(\mathbf{x}_t^\leftarrow) dt + d\mathbf{W}_t. \tag{46}$$

Here, the form of (46) follows directly from an analogous calculation as used in Lemma 7.

The reason we are interested in the dynamics (45), (46) is because they exactly capture the restricted Gaussian dynamics (RGD). Indeed, note that advancing the forward process (45) for time η from $\mathbf{x}_0 \leftarrow \mathbf{x}$ implements the first step of (RGD). Similarly, advancing the backwards process (46) for

time η from $\mathbf{x}_0^\leftarrow \leftarrow \mathbf{x}_\eta = \mathbf{y}$ implements posterior sampling, which gives the second step of (RGD) by using a similar calculation as in Theorem 2.

We summarize some key facts about (45), (46): the proofs are routine applications of Lemma 2 and integration by parts to compute generators as adjoints of time evolutions.

Lemma 24. *The generators and carré du champs of (45) and (46) are as follows.*

(i) *If $\mathbf{x}_0 \sim \mu_0$ and $\mu_t := \text{Law}(\mathbf{x}_t)$ where $\mathbf{x}_t$ follows (45), then*

$$\partial_t \mu_t = \frac{1}{2} \Delta \mu_t.$$

The generator $\mathsf{L}f = \frac{1}{2} \Delta f$ is time-invariant, and the carré du champ is $\Gamma(f, g) = \frac{1}{2} \langle \nabla f, \nabla g \rangle$.

(ii) *If $\mathbf{x}_0^\leftarrow \sim \mu_0^\leftarrow$ and $\mu_t^\leftarrow := \text{Law}(\mathbf{x}_t^\leftarrow)$ where $\mathbf{x}_t^\leftarrow$ follows (46), then*

$$\partial_t \mu_t^\leftarrow = \frac{1}{2} \Delta \mu_t^\leftarrow - \nabla \cdot (\mu_t^\leftarrow \nabla \log \pi \mathbf{Q}_{\eta-t}).$$

The generator $\mathsf{L}_t^\leftarrow f = \frac{1}{2} \Delta f + \langle \nabla \log \pi \mathbf{Q}_{\eta-t}, \nabla f \rangle$ is time-dependent, and the carré du champ $\Gamma_t(f, g) \equiv \Gamma(f, g) = \frac{1}{2} \langle \nabla f, \nabla g \rangle$ is time-invariant.

We also require a technical lemma about the time change in an f -divergence between two densities going simultaneous evolution. We defer a proof to the excellent reference [Che25].

Lemma 25 (Simultaneous heat flow; adapted from [Che25, Theorem 8.3.1]). *Consider two measures co-evolving according the same dynamics,*

$$\partial_t \mu_t = (\mathsf{L}_t)^* \mu_t, \quad \partial_t \nu_t = (\mathsf{L}_t)^* \nu_t.$$

Then, assuming $(\mathsf{L}_t)^$ satisfies a technical condition⁸ and Γ_t is its associated carré du champ operator, we have for an f -divergence $\mathcal{D}_f(\mu \parallel \nu) := \int f(\frac{\mu}{\nu}) d\nu$,*

$$\partial_t \mathcal{D}_f(\mu_t \parallel \nu_t) = - \int f''\left(\frac{\mu_t}{\nu_t}\right) \Gamma_t\left(\frac{\mu_t}{\nu_t}, \frac{\mu_t}{\nu_t}\right) d\nu_t.$$

Finally, we require an ancilliary lemma on convolving strongly log-concave functions.

Lemma 26. *Let $\mu : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be α -strongly log-concave and let $\nu : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be β -strongly log-concave, for some $\alpha, \beta > 0$. Then their convolution $\mu * \nu$ is $\frac{\alpha\beta}{\alpha+\beta}$ -strongly log-concave.*

Proof. By definition, γ -strong log-concavity of a function $f : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is equivalent to $f \exp(\frac{\gamma}{2} \|\cdot\|_2^2)$ being log-concave. Now, define the function

$$f(\mathbf{x}, \mathbf{y}) := \mu(\mathbf{y}) \nu(\mathbf{x} - \mathbf{y}) \exp\left(\frac{\gamma}{2} \|\mathbf{x}\|_2^2\right).$$

We claim that f is log-concave if $\gamma \leq \frac{\alpha\beta}{\alpha+\beta}$. This holds because

$$\begin{aligned} & -\nabla^2 \log \mu(\mathbf{y}) - \nabla^2 \log \nu(\mathbf{x} - \mathbf{y}) - \nabla^2 \left(\frac{\gamma}{2} \|\mathbf{x}\|_2^2\right) \\ & \succeq \begin{pmatrix} \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \alpha \mathbf{I}_d \end{pmatrix} + \begin{pmatrix} \beta \mathbf{I}_d & -\beta \mathbf{I}_d \\ -\beta \mathbf{I}_d & \beta \mathbf{I}_d \end{pmatrix} - \gamma \begin{pmatrix} \mathbf{I}_d & \mathbf{0}_{d \times d} \\ \mathbf{0}_{d \times d} & \mathbf{0}_{d \times d} \end{pmatrix} = \begin{pmatrix} (\beta - \gamma) \mathbf{I}_d & -\beta \mathbf{I}_d \\ -\beta \mathbf{I}_d & (\alpha + \beta) \mathbf{I}_d \end{pmatrix}, \end{aligned}$$

and we can verify that the last matrix is positive semidefinite by checking its determinant. Finally, the conclusion holds because the $\mathbf{x}$ -marginal of f is log-concave by the Prékopa-Leindler inequality, which simplifies to the convolution of μ and ν being γ -strongly log-concave. $\square$

To interpret Lemma 26, recall that classical Bakry-Émery theory shows that strong log-concavity implies a log-Sobolev inequality (38). Combining this fact with Lemma 26 thus gives a log-Sobolev inequality for densities which result from applying a heat kernel to a strongly log-concave measure.

⁸This is the diffusion chain rule [Che25, Definition 2.2.13] discussed in the proof of Proposition 2, which is true for our applications of interest (and carré du champs of the form $\Gamma(f, f) = c \|\nabla f\|_2^2$).

This is the dual of our calculation in Proposition 1, where log-Sobolev inequalities at convolved measures were used to demonstrate a log-Sobolev inequality for the original measure π .

Finally, we are ready to state our main result, a direct bound on the contraction of (RGD) to its stationary distribution in the KL divergence. Note that the estimate established here is in fact slightly stronger than that concluded by Theorem 10.

Theorem 11. *Let π be α -strongly log-concave, and consider following the dynamics (RGD) from $\mathbf{x} \sim \mu$, where we define $\mu' := \text{Law}(\mathbf{x}')$. Then,*

$$\text{KL}(\mu' \parallel \pi) \leq \frac{\text{KL}(\mu \parallel \pi)}{(1 + \alpha\eta)^2}.$$

Proof. As remarked in our discussion following (45), (46), an equivalent way to simulate (RGD) is by running (45) from $\mathbf{x}_0 \leftarrow \mathbf{x}$, and then running (46) from $\mathbf{x}_0^\leftarrow \leftarrow \mathbf{x}_\eta$, where we output $\mathbf{x}' \leftarrow \mathbf{x}_0^\leftarrow$.

In the sequel, let $\pi\mathbf{Q}_t$ denote the law of $\mathbf{x}_t$ if $\mathbf{x}_0 \sim \pi$; likewise, it is the law of $\mathbf{x}_{\eta-t}^\leftarrow$ if $\mathbf{x}_0^\leftarrow \sim \pi\mathbf{Q}_\eta$, which follows from the time reversal property. We also define $\{\mu_t\}_{t \in [0, \eta]}$, $\{\mu_t^\leftarrow\}_{t \in [0, \eta]}$ as in Lemma 24, where $\mu_0 := \mu$. We consider the forward and backward heat flows separately.

Forward heat: Applying Lemma 25 to the KL divergence, which is an f -divergence with $f(x) = x \log x$, and with the carré du champ $\Gamma_t(f, f) = \frac{1}{2} \|\nabla f\|_2^2$ from Lemma 24(i), we have

$$\partial_t \text{KL}(\mu_t \parallel \pi\mathbf{Q}_t) = -\frac{1}{2} \int \frac{\pi\mathbf{Q}_t}{\mu_t} \left\| \nabla \frac{\mu_t}{\pi\mathbf{Q}_t} \right\|_2^2 d\pi\mathbf{Q}_t \geq -\frac{\alpha}{1 + \alpha t} \text{KL}(\mu_t \parallel \pi\mathbf{Q}_t),$$

where the inequality follows from Lemma 26 and the fact that a γ -strongly log-concave measure satisfies a log-Sobolev inequality (38) with constant γ . This implies via Gronwall's inequality that

$$\text{KL}(\mu_\eta \parallel \pi\mathbf{Q}_\eta) \leq \exp\left(-\int_0^\eta \frac{\alpha}{1 + \alpha t} dt\right) \text{KL}(\mu_0 \parallel \pi) = \frac{\text{KL}(\mu_0 \parallel \pi)}{1 + \alpha\eta}. \quad (47)$$

Backward heat: Applying Lemma 25 and Lemma 24(ii), we have completely analogously that

$$\partial_t \text{KL}(\mu_t^\leftarrow \parallel \pi\mathbf{Q}_{\eta-t}) = \frac{1}{2} \int \frac{\pi\mathbf{Q}_{\eta-t}}{\mu_t^\leftarrow} \left\| \nabla \frac{\mu_t^\leftarrow}{\pi\mathbf{Q}_{\eta-t}} \right\|_2^2 d\pi\mathbf{Q}_{\eta-t} \geq -\frac{\alpha}{1 + \alpha(\eta - t)} \text{KL}(\mu_0^\leftarrow \parallel \pi\mathbf{Q}_{\eta-t}),$$

and as a result, Gronwall's inequality yields

$$\partial_t \text{KL}(\mu_\eta^\leftarrow \parallel \pi) \leq \exp\left(-\int_0^\eta \frac{\alpha}{1 + \alpha(\eta - t)} dt\right) \text{KL}(\mu_0^\leftarrow \parallel \pi\mathbf{Q}_\eta) = \frac{\text{KL}(\mu_0^\leftarrow \parallel \pi\mathbf{Q}_\eta)}{1 + \alpha\eta}. \quad (48)$$

Finally, combining (47) and (48) completes the proof. $\square$

Lastly, we note that Theorem 11 can be strengthened in a few ways. A fortiori, Lemma 25 allows us to write a similar result for Rényi divergences, and direct use of functional inequalities allows extensions beyond the strong log-concavity. We defer more discussion to [Che25, Chapter 8].