

Mixing Time of the Proximal Sampler in Relative Fisher Information via Strong Data Processing Inequality

Andre Wibisono 

Abstract—We study the mixing time guarantee for sampling in relative Fisher information via the Proximal Sampler algorithm, which is an approximate proximal discretization of the Langevin dynamics. We show that when the target probability distribution is strongly log-concave, the relative Fisher information converges exponentially fast along the Proximal Sampler; this matches the exponential convergence rate of the relative Fisher information along the continuous-time Langevin dynamics for strongly log-concave target. When combined with a standard implementation of the Proximal Sampler via rejection sampling, this exponential convergence rate provides a high-accuracy iteration complexity guarantee for the Proximal Sampler in relative Fisher information when the target distribution is strongly log-concave and log-smooth. Our proof proceeds by establishing a strong data processing inequality for relative Fisher information along the Gaussian channel under strong log-concavity, and a data processing inequality along the reverse Gaussian channel for a special distribution. The forward and reverse Gaussian channels compose to form the Proximal Sampler, and these data processing inequalities imply the exponential convergence rate of the relative Fisher information along the Proximal Sampler.

Index Terms—Mixing time, Fisher information, Langevin dynamics, data processing inequality.

I. INTRODUCTION

IN THIS paper we establish the mixing time guarantee of the Proximal Sampler algorithm for sampling from a strongly log-concave target distribution with error guarantee in relative Fisher information. We describe the background from sampling and the analogy with optimization in Section I-A. Our technical contribution is to establish a strong data processing inequality in relative Fisher information under strong log-concavity, which we describe in Section I-B.

A. Mixing Time of the Proximal Sampler in Relative Fisher Information

Sampling from a probability distribution is a fundamental algorithmic task in many applications, including in computer

Received 15 June 2025; revised 21 December 2025; accepted 15 February 2026. Date of publication 9 March 2026; date of current version 23 April 2026. This work was supported by NSF under Award CCF-2403391 and Award CCF-2443097. An earlier version of this paper was presented in part at the 38th Annual Conference on Learning Theory (COLT 2025).

The author is with the Department of Computer Science, Yale University, New Haven, CT 06520 USA (e-mail: andre.wibisono@yale.edu).

Communicated by S.-H. Lee, Associate Editor for Shannon Theory and Information Measures.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TIT.2026.3670723>.

Digital Object Identifier 10.1109/TIT.2026.3670723

science, machine learning, and Bayesian statistics. Sampling algorithms typically take the form of random walks or Markov chains which have the target probability distribution as their stationary distribution. The *mixing time* of the Markov chain, i.e., how fast the Markov chain converges to its stationary distribution, determines the iteration complexity of the sampling algorithm. For many applications, we need mixing time guarantees where the error is measured in a certain metric or statistical divergence, such as in the total variation (TV) distance, Wasserstein distance, Kullback-Leibler (KL) divergence, chi-square (χ^2) divergence, Rényi divergence, or even in the general Φ -divergence [1], which broadly generalizes the TV distance, KL divergence, and χ^2 -divergence. Thus, there have been great interests in developing Markov chains with fast mixing time guarantees in various divergences which can be implemented efficiently in practice.

For target distributions supported on a continuous domain such as the Euclidean space, many sampling algorithms can be derived as discretizations of continuous-time stochastic processes such as the Langevin dynamics; these include the Unadjusted Langevin Algorithm (ULA) [2] and the Proximal Sampler algorithm [3], which we also study in this paper. The Langevin dynamics for sampling has a natural optimization interpretation as the gradient flow for minimizing KL divergence in the space of probability distributions [4]. This perspective of sampling as optimization is useful to translate the tools and techniques from optimization to sampling, to derive and analyze sampling algorithms in discrete time, and to suggest sampling guarantees one may hope to obtain by drawing parallels from optimization; see [5], and Appendix B for a review.

When the target distribution is well-behaved, e.g., when it is strongly log-concave or satisfies isoperimetric conditions such as the log-Sobolev inequality (LSI), existing results in the literature have established mixing time guarantees in various divergences. In particular, for strongly log-concave target distributions, the Langevin dynamics in continuous time is known to converge exponentially fast in Wasserstein distance and KL divergence [6], in Rényi divergence [7], and in all Φ -divergence [1]. In discrete time, the convergence guarantees are more subtle due to a potential bias issue, since some algorithms may have stationary distributions which are different from the target distribution.

A simple discretization of the Langevin dynamics is the *Unadjusted Langevin Algorithm (ULA)*, which is an explicit

algorithm that is easy to implement in practice. However, ULA is *biased*, which means for each fixed step size, it converges to a biased limiting distribution [2]; this bias can be seen as arising from using a mismatched discretization method for discretizing the Langevin dynamics when viewed as a gradient flow for optimizing a composite objective function [5]. When the target distribution is strongly log-concave and log-smooth, one can show the biased convergence guarantees of ULA to the target distribution in Wasserstein distance [8], [9], [10], in KL divergence [11], [12], or in Rényi divergence [13]; one can also show the exponential convergence guarantees of ULA to its biased stationary distribution in KL and Rényi divergence [14], [15], and in Φ -divergence [16]. However, the presence of the bias only yields a “low-accuracy” iteration complexity guarantee for ULA, where the number of iterations to reach an error ε in the specified divergence scales polynomially in ε^{-1} , rather than logarithmically as suggested by the exponential convergence in continuous time. We note one way to remove the bias of ULA is to add a Metropolis-Hastings filter or the accept-reject step, resulting in the Metropolis-Adjusted Langevin Algorithm (MALA) [2]. Since MALA is unbiased, its iteration complexity has the desired logarithmic dependency on the error parameter; however, the convergence guarantees of MALA are only known in weaker metrics such as in the total variation distance [17], [18], [19], [20], [21].

Another algorithm that implements the Langevin dynamics is the *Proximal Sampler*; it was derived as a Gibbs sampling algorithm for sampling from a joint distribution defined on an extended phase space [3], [22], and it can be viewed as an approximate proximal discretization of the Langevin dynamics [3], [23]. The Proximal Sampler is unbiased, so it converges to the correct target distribution. When the target distribution is strongly log-concave, one can show an exponential convergence guarantee of the Proximal Sampler in Wasserstein distance [3], in KL and Rényi divergence [23], and in all Φ -divergence [16]. The Proximal Sampler is an implicit algorithm, since in each iteration it requires sampling from a regularized distribution, known as the *Restricted Gaussian Oracle (RGO)* [3]; when the target distribution is log-smooth, the RGO can be implemented efficiently via rejection sampling [23]. Thus, for strongly log-concave and log-smooth target distributions, the Proximal Sampler has a “high-accuracy” iteration complexity guarantee that scales logarithmically with the inverse error parameter, as expected from the exponential convergence of the continuous-time Langevin dynamics.

1) *Our Work*: In this paper, we study the mixing time of the Proximal Sampler algorithm for sampling with error guarantee in *relative Fisher information*, which measures the expected square distance between the score functions (gradient of the log-density) between two distributions:

$$\text{FI}(\rho \parallel \nu) = \mathbb{E}_\rho \left[\left\| \nabla \log \rho - \nabla \log \nu \right\|^2 \right].$$

Accordingly, it is a natural convergence criterion in settings where gradients of log-densities matter, including Stein-type discrepancies, score matching, and diffusion/score-based generative modeling [24], [25]. In such applications, $\text{KL}(\rho \parallel \nu) = \mathbb{E}_\rho[\log \rho - \log \nu]$ can be small while the score error $\text{FI}(\rho \parallel \nu)$ remains large (see Lemma 5 in Appendix D), so a bound in FI provides strictly stronger information. The (counter)examples

we describe below show that one should not expect monotonicity of FI (i.e., weak data processing inequality) to hold without additional structure; this motivates the strong data processing viewpoint we develop here. Our results present a notable departure from prior work since relative Fisher information is not a Φ -divergence, so previous results for Φ -divergence do not apply. We develop new tools for understanding FI along the Fokker-Planck channel (Lemma 2), which enables us to characterize situations under which weak or strong data processing inequality holds (Theorem 1 for the heat flow and Theorem 2 for the Ornstein-Uhlenbeck flow).

2) *Analogy With Optimization*: Another motivation for our work is to strengthen the links between sampling and optimization, in particular via the perspective of sampling as optimization. We can view sampling from a target distribution ν over \mathbb{R}^d as the problem of optimizing the KL divergence $\text{KL}(\cdot \parallel \nu)$ on the space of probability distributions $\mathcal{P}(\mathbb{R}^d)$, and the gradient flow under the Wasserstein metric for optimizing KL divergence is the Langevin dynamics [4]. Relative Fisher information has the interpretation as the squared Wasserstein gradient norm of the KL divergence. In the Wasserstein geometry, conditions on the distribution ν translate to properties of the KL divergence that enable fast optimization; e.g., the strong log-concavity of ν corresponds to the geodesic strong convexity of the KL divergence, while the LSI on ν corresponds to the Polyak-Łojaciewicz or gradient domination condition of KL divergence [26]. See Appendix B for further review.

In strongly convex optimization, we can obtain exponential convergence guarantees in both the function value and in the squared gradient norm, in continuous time via the gradient flow dynamics and in discrete time via the proximal gradient algorithm (see Appendix C for a review). Then for sampling a strongly log-concave target distribution, it is natural to expect exponential convergence guarantees in both the KL divergence (the objective function), as well as in the relative Fisher information (the squared gradient norm). The exponential convergence guarantees in KL divergence are known both for the Langevin dynamics in continuous time (see Section II-E.1) and for the Proximal Sampler in discrete time (see Section IV-B). When the target distribution is strongly log-concave, the Langevin dynamics is also known to converge exponentially fast in relative Fisher information (see Section II-E). However, in discrete time, the corresponding exponential convergence guarantee in relative Fisher information seems to be unknown; this is what we fill out in this work. Indeed, we show that the Proximal Sampler has an exponential convergence guarantee in relative Fisher information when the target distribution is strongly log-concave (Theorem 3). This implies a high-accuracy iteration complexity guarantee of the Proximal Sampler in relative Fisher information for log-smooth and strongly log-concave target distribution (Corollary 1).

3) *Related Work*: A few works have studied mixing time in relative Fisher information as a sampling analogue of the gradient norm convergence in optimization. Reference [27] study mixing time in relative Fisher information for sampling from non-log-concave and log-smooth target distributions, analogous to finding approximate stationary points in non-convex optimization; they study the biased convergence of ULA

in relative Fisher information, and provide a low-accuracy iteration complexity guarantee on the average iterate of ULA. Reference [28] study the iteration complexity lower bounds for relative Fisher information in non-log-concave sampling, analogous to lower bounds for finding stationary points in non-convex optimization. We note that [28, Appendix A.2] provide an argument for obtaining a high-accuracy mixing time in relative Fisher information for log-concave target distributions, involving a reduction from the high-accuracy mixing time of the Proximal Sampler in χ^2 -square divergence from [23] and a post-processing via the heat flow. In this work, we give a direct proof of the exponential convergence of the Proximal Sampler in relative Fisher information for strongly log-concave target distributions. We provide further discussion on related works in Appendix A. Our high-level proof technique follows the template of information-dissipation to contraction as in prior analyses of the Proximal Sampler [23]. Our work is technically distinct in an essential way: relative Fisher information is not a Φ -divergence, so standard data processing or convexity arguments do not apply directly. Our main new ingredient is a time derivative identity for relative Fisher information along simultaneous Fokker–Planck channels (Lemma 2) and its SDPI consequences (Theorem 1).

B. Strong Data Processing Inequality in Relative Fisher Information

Data processing inequality (DPI) is a fundamental tenet in information theory which states that information is never increasing along any noisy channel [29]. Here “information” is measured via statistical divergences such as the KL divergence or any Φ -divergence, and their monotone transformations including Rényi divergence. The *strong data processing inequality* (SDPI) is a strengthening of DPI that quantifies the rate at which information is contracting along the noisy channel, where the contraction rate may depend on properties of the channel and the input distributions. For a review of SDPIs and their applications, see e.g. [30], [31], [32], and [33].

SDPIs are useful because many sampling algorithms can be viewed as composition of noisy channels. In continuous time, the Langevin dynamics can be viewed as a channel driven by a diffusion process. In discrete time, ULA can be viewed as the composition of a deterministic gradient map and a Gaussian channel, while the Proximal Sampler can be viewed as the composition of the forward and reverse Gaussian channels. If we have an SDPI for each channel in the algorithm, then we immediately obtain a mixing time guarantee for the sampling algorithm by iterating the contraction rate across multiple steps of the algorithm. This strategy has been successful in establishing the convergence guarantees of ULA and the Proximal Sampler in KL and Rényi divergence, e.g., by [14], [23], [34], and [35], in Φ -divergence by [16], and in mutual information by [36].

For our goal of proving mixing time guarantees in relative Fisher information, ideally we want to establish SDPIs in relative Fisher information. Relative Fisher information is not a Φ -divergence, nor a monotone transformation of it. It seems unknown whether DPI holds for relative Fisher information.

We note we work with the *non-parametric* relative Fisher information, which has a natural geometric interpretation as the squared gradient norm of KL divergence in Wasserstein metric; for the classical parametric Fisher information, some form of DPI is known, see e.g. [37] and also Appendix A-A for further discussion. One way to prove DPI for Φ -divergence is to use the joint convexity of the Φ -divergence and then apply Jensen’s inequality; however, this fails in our case since relative Fisher information is convex in its first argument, but not convex in its second argument [38]. In fact, we show that DPI in relative Fisher information does not always hold in general, by constructing a counterexample for the Gaussian channel (see Example 3).

To proceed, we focus on the *Fokker-Planck (FP) channels*, a family of noisy channels generated by continuous-time diffusion processes (see Section II-D); this includes the Langevin dynamics, the Gaussian channel, and the reverse Gaussian channel, which compose the Proximal Sampler. One way to prove SDPIs for Φ -divergence along the FP channels is via time-differentiation of the Φ -divergence along the simultaneous evolution along the FP channel (see Lemma 1 for a review in the case of KL divergence). We extend this technique to analyze the time derivative of the relative Fisher information along a simultaneous evolution along the FP channel. Under log-concavity assumptions on one of the input distributions, we can extract DPI and SDPI results for relative Fisher information.

C. Contributions

Our main algorithmic result is the exponential convergence of the Proximal Sampler in relative Fisher information when the target distribution is strongly log-concave (Theorem 3). We combine this with a standard rejection sampling implementation of the Proximal Sampler to conclude a high-accuracy iteration complexity guarantee of the Proximal Sampler in relative Fisher information for log-smooth and strongly log-concave target distribution (Corollary 1).

Our main technical contribution is to establish a formula for the time derivative of relative Fisher information along a simultaneous application of the Fokker-Planck channel (Lemma 2). We use this to show that along the Gaussian channel, (S)DPI in relative Fisher information holds when the second distribution is (strongly) log-concave. We also show that along the Gaussian channel, SDPI in relative Fisher information eventually holds when the second distribution is a log-Lipschitz perturbation of a strongly log-concave distribution (Theorem 1). However, in general, DPI in relative Fisher information may not hold. We show an example when DPI initially does not hold along the Gaussian channel even when the first distribution is a Gaussian (Example 3).

We also show that along the Ornstein-Uhlenbeck (OU) channel, SDPI in relative Fisher information eventually holds when the second distribution is strongly log-concave (Theorem 2). However, we also show an example when DPI in relative Fisher information initially does not hold along the OU channel, even when both input distributions are Gaussians (Example 4).

II. PRELIMINARIES AND REVIEWS

A. Notations

In this paper, we are working on the Euclidean space \mathbb{R}^d for some $d \geq 1$. We denote the ℓ_2 -inner product between vectors $u = (u_1, \dots, u_d), v = (v_1, \dots, v_d) \in \mathbb{R}^d$ as $u^\top v = \langle u, v \rangle = \sum_{i=1}^d u_i v_i$. For a symmetric matrix $A \in \mathbb{R}^{d \times d}$, the notation $A \geq 0$ means A is positive semidefinite, i.e., $u^\top A u \geq 0$ for all $u \in \mathbb{R}^d$. For symmetric matrices $A, B \in \mathbb{R}^{d \times d}$, the notation $A \geq B$ means $A - B \geq 0$ is positive semidefinite. Throughout, let $I \in \mathbb{R}^{d \times d}$ denote the identity matrix.

For a matrix $A \in \mathbb{R}^{d \times d}$, let $\text{Tr}(A) = \sum_{i=1}^d A_{ii}$ be the trace of A . Let $\|A\|_{\text{HS}}^2 = \text{Tr}(A^\top A)$ denote the squared Hilbert-Schmidt (Frobenius) norm of A . For $u \in \mathbb{R}^d$ and $A \in \mathbb{R}^{d \times d}$, we denote

$$\|u\|_A^2 := u^\top A u$$

to be the ‘‘squared norm’’ of a vector $u \in \mathbb{R}^d$ with respect to the matrix A . If $A \geq 0$, then $\|u\|_A^2 \geq 0$ for all $u \in \mathbb{R}^d$. However, if A is not positive semidefinite, then $\|u\|_A^2$ may be negative for some u . Note $\|u\|_A^2$ only depends on the symmetric part $A_{\text{sym}} = \frac{1}{2}(A + A^\top)$ of A , i.e., $\|u\|_A^2 = \|u\|_{A_{\text{sym}}}^2$.

For a differentiable function $f: \mathbb{R}^d \rightarrow \mathbb{R}$, let $\nabla f(x) \in \mathbb{R}^d$ denote the gradient vector at $x \in \mathbb{R}^d$ of the partial derivatives: $(\nabla f(x))_i = \frac{\partial f(x)}{\partial x_i}$. Let $\nabla^2 f(x) \in \mathbb{R}^{d \times d}$ be the Hessian matrix of second partial derivatives: $(\nabla^2 f(x))_{i,j} = \frac{\partial^2 f(x)}{\partial x_i \partial x_j}$. Let $\Delta f(x) = \text{Tr}(\nabla^2 f(x))$ be the Laplacian.

For a vector field $v: \mathbb{R}^d \rightarrow \mathbb{R}^d$ with $v(x) = (v_1(x), \dots, v_d(x)) \in \mathbb{R}^d$, let $\nabla v: \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ be the Jacobian matrix of mixed partial derivatives: $(\nabla v(x))_{i,j} = \frac{\partial v_j(x)}{\partial x_i}$. Let $\nabla \cdot v: \mathbb{R}^d \rightarrow \mathbb{R}$ be the divergence of v , defined by $(\nabla \cdot v)(x) = \sum_{i=1}^d \frac{\partial v_i(x)}{\partial x_i} = \text{Tr}(\nabla v(x))$.

B. Relative Entropy and Relative Fisher Information

Let $\mathcal{P}(\mathbb{R}^d)$ denote the space of probability distributions ρ over \mathbb{R}^d which are absolutely continuous with respect to the Lebesgue measure and have a finite second moment $\mathbb{E}_\rho[\|X\|^2] < \infty$. We identify a probability distribution $\rho \in \mathcal{P}(\mathbb{R}^d)$ with its probability density function with respect to the Lebesgue measure, which we also denote by $\rho: \mathbb{R}^d \rightarrow \mathbb{R}$, so $\rho(x) > 0$ and $\int_{\mathbb{R}^d} \rho(x) dx = 1$. We say ρ is absolutely continuous with respect to another distribution ν , denoted by $\rho \ll \nu$, if $\nu(A) = 0$ implies $\rho(A) = 0$ for any $A \subseteq \mathbb{R}^d$; if ρ and ν both have density functions, then $\rho \ll \nu$ means $\nu(x) = 0$ implies $\rho(x) = 0$ for all $x \in \mathbb{R}^d$. Throughout, we assume ρ and ν have full support \mathbb{R}^d and their density functions have the required regularity properties for the various operations below to be well-defined.

Let $H: \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ be the (differential) entropy functional:

$$H(\rho) = -\mathbb{E}_\rho[\log \rho] = -\int_{\mathbb{R}^d} \rho(x) \log \rho(x) dx.$$

Let $J: \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ be the Fisher information:

$$J(\rho) = \mathbb{E}_\rho \left[\|\nabla \log \rho\|^2 \right] = -\mathbb{E}_\rho [\Delta \log \rho],$$

and we define $J(\rho) = +\infty$ if ρ does not have a differentiable density. The second equality in the definition of $J(\rho)$ above follows by integration by parts assuming the boundary terms are zero.

For probability distributions $\rho \ll \nu$ on \mathbb{R}^d , the *Kullback-Leibler (KL) divergence* or the *relative entropy* of ρ with respect to ν is defined by:

$$\text{KL}(\rho \parallel \nu) = \mathbb{E}_\rho \left[\log \frac{\rho}{\nu} \right] = \int_{\mathbb{R}^d} \rho(x) \log \frac{\rho(x)}{\nu(x)} dx.$$

If ρ and ν have differentiable density functions, then the *relative Fisher information* of ρ with respect to ν is defined by:

$$\text{FI}(\rho \parallel \nu) = \mathbb{E}_\rho \left[\left\| \nabla \log \frac{\rho}{\nu} \right\|^2 \right] = \int_{\mathbb{R}^d} \rho(x) \left\| \nabla \log \frac{\rho(x)}{\nu(x)} \right\|^2 dx.$$

We note that the KL divergence KL is the Bregman divergence of the negative entropy $-H$, and similarly, the relative Fisher information FI is the Bregman divergence of the Fisher information J .

C. Strong Log-Concavity, Log-Sobolev Inequality, and Poincaré Inequality

We say a probability distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$ is α -strongly log-concave (α -SLC) for some $\alpha > 0$ if $-\log \nu: \mathbb{R}^d \rightarrow \mathbb{R}$ is an α -strongly convex function; if $\log \nu$ is twice differentiable, then this is equivalent to $-\nabla^2 \log \nu(x) \geq \alpha I$ for all $x \in \mathbb{R}^d$. We say ν is (weakly) log-concave if $-\log \nu$ is a convex function, or $-\nabla^2 \log \nu(x) \geq 0$ for all $x \in \mathbb{R}^d$.

We say ν is L -log-smooth for some $0 < L < \infty$ if $-LI \leq -\nabla^2 \log \nu(x) \leq LI$ for all $x \in \mathbb{R}^d$.

We say ν satisfies the α -log-Sobolev inequality (α -LSI) for some $\alpha > 0$ if for any probability distribution ρ , the following inequality holds:

$$\text{FI}(\rho \parallel \nu) \geq 2\alpha \text{KL}(\rho \parallel \nu).$$

We recall by the Bakry-Émery criterion that if ν is α -SLC, then ν satisfies α -LSI [6], [39]. We note that even under SLC, the gap between KL divergence and relative Fisher information can be arbitrarily large; see Lemma 5 in Appendix D for an example.

We say ν satisfies the α -Poincaré inequality (α -PI) with $\alpha > 0$ if the following inequality holds for any smooth function $\phi: \mathbb{R}^d \rightarrow \mathbb{R}$:

$$\text{Var}_\nu(\phi) \leq \frac{1}{\alpha} \mathbb{E}_\nu[\|\nabla \phi\|^2].$$

Equivalently, for any smooth vector field $\psi: \mathbb{R}^d \rightarrow \mathbb{R}^d$, we have $\text{Var}_\nu(\psi) \leq \frac{1}{\alpha} \mathbb{E}_\nu[\|\nabla \psi\|_{\text{HS}}^2]$. We recall that if ν satisfies α -LSI, then ν also satisfies α -PI [39].

We say a probability distribution ν on \mathbb{R}^d is symmetric if $\nu(x) = \nu(-x)$ for all $x \in \mathbb{R}^d$.

D. Fokker-Planck Channels

We consider the continuous-time stochastic process $(X_t)_{t \geq 0}$ on \mathbb{R}^d which evolves following the stochastic differential equation (SDE):

$$dX_t = b_t(X_t) dt + \sqrt{c} dW_t \quad (1)$$

where $b_t: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a Lipschitz and twice continuously-differentiable vector field, $c \geq 0$ is a constant, and $(W_t)_{t \geq 0}$ is the standard Brownian motion in \mathbb{R}^d that starts from $W_0 = 0$. Recall that if $X_t \sim \rho_t$ evolves following the SDE (1), then

its probability density function $\rho_t: \mathbb{R}^d \rightarrow \mathbb{R}$ evolves via the *Fokker-Planck equation*, which is the following partial differential equation (PDE):

$$\partial_t \rho_t = -\nabla \cdot (\rho_t b_t) + \frac{c}{2} \Delta \rho_t. \quad (2)$$

We refer to either the SDE (1) or the PDE (2) as the *Fokker-Planck (FP) channel*, which maps any initial random variable $X_0 \sim \rho_0$ to the random variable $X_t \sim \rho_t$ which is the output of the SDE (1) at time $t > 0$; or equivalently, it maps any initial distribution $\rho_0 \in \mathcal{P}(\mathbb{R}^d)$ to the probability distribution $\rho_t \in \mathcal{P}(\mathbb{R}^d)$ which is the solution to the PDE (2) at time $t > 0$. Fokker-Planck channels have been studied, e.g., in [40], [41], [42], and [43].

Some examples of the particular FP channels we consider in this paper are the following.

- 1) **Gaussian channel:** Consider $b_t \equiv 0$, so the SDE is simply a Brownian motion: $dX_t = dW_t$. The exact solution is:

$$X_t = X_0 + W_t \stackrel{d}{=} X_0 + \sqrt{t}Z \quad (3)$$

where $Z \sim \mathcal{N}(0, I)$ is an independent Gaussian random variable in \mathbb{R}^d . This is the classical *Gaussian channel* in information theory. The Fokker-Planck equation is the *heat equation*:

$$\partial_t \rho_t = \frac{1}{2} \Delta \rho_t. \quad (4)$$

- 2) **Reverse Gaussian channel:** Let $\pi_t = \pi_0 * \mathcal{N}(0, tI)$ where $\mathcal{N}(0, tI)$ is the Gaussian distribution with mean 0 and covariance tI , so π_t satisfies the heat equation: $\partial_t \pi_t = \frac{1}{2} \Delta \pi_t$. Let $T > 0$, and define $\nu_t := \pi_{T-t}$, for $0 \leq t \leq T$. Then ν_t satisfies the *backward heat equation*:

$$\begin{aligned} \partial_t \nu_t &= -\frac{1}{2} \Delta \nu_t = -\nabla \cdot (\nu_t \nabla \log \nu_t) + \frac{1}{2} \Delta \nu_t \\ &= -\nabla \cdot (\nu_t \nabla \log \pi_{T-t}) + \frac{1}{2} \Delta \nu_t. \end{aligned} \quad (5)$$

This is the Fokker-Planck equation for the SDE:

$$dX_t = \nabla \log \pi_{T-t}(X_t) dt + dW_t. \quad (6)$$

This is a Fokker-Planck channel with time-dependent vector field $b_t = \nabla \log \nu_t = \nabla \log \pi_{T-t}$, which we call the *reverse Gaussian channel*. It has the property that if we run the channel (6) from $X_0 \sim \rho_0 = \pi_T$, then $X_T \sim \rho_T = \pi_0$, i.e., we “reverse” the heat equation. (However, note that if $\rho_0 \neq \pi_T$, then $\rho_T \neq \pi_0$, so the channel is only suitably constructed for π_0 .)

- 3) **Langevin dynamics:** Let $b_t(x) = -\nabla g(x)$ for some differentiable function $g: \mathbb{R}^d \rightarrow \mathbb{R}$ with $\int_{\mathbb{R}^d} e^{-g(x)} dx < \infty$, and set $c = 2$. Then the Fokker-Planck channel becomes:

$$dX_t = -\nabla g(X_t) dt + \sqrt{2} dW_t. \quad (7)$$

This is the *Langevin dynamics* for sampling from the probability distribution $\nu(x) \propto e^{-g(x)}$. Indeed, the Fokker-Planck equation becomes:

$$\partial_t \rho_t = \nabla \cdot (\rho_t \nabla g) + \Delta \rho_t = \nabla \cdot \left(\rho_t \nabla \log \frac{\rho_t}{\nu} \right) \quad (8)$$

and we see that ν is a stationary distribution.

- 4) **Ornstein-Uhlenbeck channel:** Let $b_t(x) = -\gamma x$ for some $\gamma > 0$, and set $c = 2$. Then the Fokker-Planck channel becomes the *Ornstein-Uhlenbeck (OU) process*:

$$dX_t = -\gamma X_t dt + \sqrt{2} dW_t. \quad (9)$$

This is the Langevin dynamics for the Gaussian target distribution $\nu = \mathcal{N}\left(0, \frac{1}{\gamma}I\right)$; we also call (9) as the *OU channel*. The SDE (9) has an explicit solution at each $t \geq 0$:

$$X_t \stackrel{d}{=} e^{-\gamma t} X_0 + \sqrt{\frac{1 - e^{-2\gamma t}}{\gamma}} Z \quad (10)$$

where $Z \sim \mathcal{N}(0, I)$ is an independent Gaussian random variable in \mathbb{R}^d . Note as $\gamma \rightarrow 0$, the OU channel (9) recovers the Gaussian channel (3) at twice the time speed ($2t$ in place of t), and the solution (10) recovers the Gaussian channel solution $X_t \stackrel{d}{=} X_0 + \sqrt{2t}Z$.

E. Review of the Convergence Rates Along the Langevin Dynamics

We review the convergence rates of the KL divergence and relative Fisher information along the Langevin dynamics. This provides a useful context for the formulas and the convergence rates that we find along the more general Fokker-Planck channel.

1) *Convergence Rate of the KL Divergence:* We recall if $(\rho_t)_{t \geq 0}$ evolves following the Fokker-Planck equation (8) for the Langevin dynamics (7) targeting $\nu \propto e^{-g}$, then we have the following *de Bruijn’s identity* on the time derivative of the KL divergence in terms of the relative Fisher information:

$$\frac{d}{dt} \text{KL}(\rho_t \| \nu) = -\text{FI}(\rho_t \| \nu). \quad (11)$$

This follows by direct time differentiation and integration by parts. If ν satisfies α -LSI, then from de Bruijn’s identity above, along the Langevin dynamics we have $\frac{d}{dt} \text{KL}(\rho_t \| \nu) \leq -2\alpha \text{KL}(\rho_t \| \nu)$. Therefore, we conclude an exponential convergence rate of the KL divergence along the Langevin dynamics when the target distribution ν satisfies LSI:

$$\text{KL}(\rho_t \| \nu) \leq e^{-2\alpha t} \text{KL}(\rho_0 \| \nu). \quad (12)$$

2) *Convergence Rate of the Relative Fisher Information:* We recall if $(\rho_t)_{t \geq 0}$ evolves following the Fokker-Planck equation (8) for the Langevin dynamics (7) targeting $\nu \propto e^{-g}$, then we have the following identity on the time derivative of the relative Fisher information, in terms of the second-order relative Fisher information:

$$\begin{aligned} \frac{d}{dt} \text{FI}(\rho_t \| \nu) &= -2\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{\nu} \right\|_{\text{HS}}^2 \right] \\ &\quad - 2\mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu} \right\|_{\nabla^2 g}^2 \right]. \end{aligned} \quad (13)$$

This identity follows by direct computation via integration by parts, or via the Otto calculus formula for the Hessian of KL divergence; see [39, Ch. 15]. If $\nu \propto e^{-g}$ is α -SLC, so $\nabla^2 g(x) \geq \alpha I$ for all $x \in \mathbb{R}^d$, then from the identity (13), along the Langevin dynamics we have $\frac{d}{dt} \text{FI}(\rho_t \| \nu) \leq -2\alpha \text{FI}(\rho_t \| \nu)$. Therefore, we conclude an exponential convergence rate of

the relative Fisher information along the Langevin dynamics when the target distribution ν is α -strongly log-concave:

$$\text{FI}(\rho_t \parallel \nu) \leq e^{-2\alpha t} \text{FI}(\rho_0 \parallel \nu). \quad (14)$$

We can verify that this convergence rate is tight in the Gaussian case.

Example 1 (Convergence rate of the relative Fisher information along the OU process): Consider the OU process (9), which is the Langevin dynamics (7) for the target Gaussian distribution $\nu = \mathcal{N}(0, \gamma^{-1}I)$, which is γ -strongly log-concave. Suppose we start the OU process from $\rho_0 = \mathcal{N}(m, s^{-1}I)$ for some $m \in \mathbb{R}^d$ and $s > 0$. For all $t \geq 0$, the solution (10) to the OU channel is $\rho_t = \mathcal{N}(e^{-\gamma t}m, s_t^{-1}I)$ where $s_t^{-1} = e^{-2\gamma t}s^{-1} + (1 - e^{-2\gamma t})\gamma^{-1}$. Note that $\min\{s, \gamma\} \leq s_t \leq \max\{s, \gamma\}$, so $s_t = \Theta(1)$ for all $t \geq 0$. Furthermore, note that $(s_t - \gamma)^2 = s_t^2\gamma^2(s_t^{-1} - \gamma^{-1})^2 = e^{-4\gamma t}s_t^2\gamma^2(s^{-1} - \gamma^{-1})^2$. Then we can compute:

$$\begin{aligned} \text{FI}(\rho_t \parallel \nu) &= e^{-2\gamma t} \gamma^2 \|m\|^2 + \frac{d}{s_t} (s_t - \gamma)^2 \\ &= e^{-2\gamma t} \gamma^2 \|m\|^2 + e^{-4\gamma t} d\gamma^2 s_t (s^{-1} - \gamma^{-1})^2. \end{aligned}$$

We observe that if $m \neq 0$, then $\text{FI}(\rho_t \parallel \nu) = O(e^{-2\gamma t})$, as predicted in the general convergence rate (14). (Note that if $m = 0$, then the rate improves to $\text{FI}(\rho_t \parallel \nu) = O(e^{-4\gamma t})$; this is because in this case ρ_t and ν are symmetric, and ρ_t satisfies a Poincaré inequality, so we can use the second-order relative Fisher information in (13) to improve the convergence rate, see Theorem 2(ii).)

F. Review of the SDPI in KL Divergence Along the Fokker-Planck Channel

We recall the following formula for the time derivative of the KL divergence along simultaneous applications of the Fokker-Planck channel; this implies the SDPI in KL divergence along the Fokker-Planck channel when the second distribution satisfies LSI. We note the similarity with de Bruijn's identity (11), which is a special case of the identity (15) below for the Langevin dynamics when the second distribution is stationary. For completeness, we provide the proof of Lemma 1 in Appendix E.

Lemma 1: Suppose $(\rho_t)_{t \geq 0}$ and $(\nu_t)_{t \geq 0}$ evolve following the Fokker-Planck channel (2):

$$\begin{aligned} \partial_t \rho_t &= -\nabla \cdot (\rho_t b_t) + \frac{c}{2} \Delta \rho_t, \\ \partial_t \nu_t &= -\nabla \cdot (\nu_t b_t) + \frac{c}{2} \Delta \nu_t \end{aligned}$$

for some vector fields $(b_t)_{t \geq 0}$ and constant $c \geq 0$ so Assumption 1 is satisfied. Then for any $t \geq 0$:

$$\frac{d}{dt} \text{KL}(\rho_t \parallel \nu_t) = -\frac{c}{2} \text{FI}(\rho_t \parallel \nu_t). \quad (15)$$

In particular, if for all $t \geq 0$ we know that ν_t satisfies α_t -LSI, then:

$$\text{KL}(\rho_t \parallel \nu_t) \leq \exp\left(-c \int_0^t \alpha_s ds\right) \text{KL}(\rho_0 \parallel \nu_0). \quad (16)$$

A similar formula to (15) holds for all Φ -divergence and Rényi divergence [16], [23], [35]; we only present it for KL divergence to draw the comparison with the relative Fisher

information that we study in this paper. In the case of the heat flow (4), the formula (15) can be used to show an SDPI along the Gaussian channel under LSI [38]. Generalizations of the formula above have been used to prove the mixing time guarantees of ULA and the Proximal Sampler in [14], [16], [23], [34], [35], and [36]. See [44] for a rigorous derivation of the identity (15) in the case of the heat flow (the case $b_t = 0$).

III. BEHAVIOR OF THE RELATIVE FISHER INFORMATION ALONG THE FOKKER-PLANCK CHANNEL

We present results on the contraction of relative Fisher information along Fokker-Planck channels.

A. Time Derivative of the Relative Fisher Information Along the Fokker-Planck Channel

Our key tool is the following formula for the time derivative of the relative Fisher information along simultaneous Fokker-Planck channel (2). This formula is the Fisher information analogue of the identity (15) for KL divergence. We provide the proof of Lemma 2 in Appendix F-A, assuming sufficient regularity and decay of the densities such that we can differentiate under the integral sign, and the boundary terms in integration by parts vanish; see Assumption 1.

Assumption 1: Assume ρ_t and ν_t below have strictly positive C^2 (twice continuously-differentiable) densities on \mathbb{R}^d , with sufficient decay so the boundary terms in integrations by parts vanish (e.g., ρ_t, ν_t and their first derivatives are integrable and $\lim_{\|x\| \rightarrow \infty} \rho_t(x) \|b_t(x)\| = 0$); the drift b_t is C^1 in x with locally-Lipschitz derivatives, and the map $t \mapsto (\rho_t, \nu_t, b_t)$ is sufficiently regular to justify differentiations under the integral sign in the computations below.

We note [44, Conjecture 4.2] also provides a time derivative formula for the relative Φ -Fisher information (here we study the case $\Phi = \text{KL}$ divergence) along the heat flow (which is the case $b_t = 0$, whereas here we study Fokker-Planck channel with general b_t).

Lemma 2: Suppose $(\rho_t)_{t \geq 0}$ and $(\nu_t)_{t \geq 0}$ evolve following the Fokker-Planck equation (2):

$$\begin{aligned} \partial_t \rho_t &= -\nabla \cdot (\rho_t b_t) + \frac{c}{2} \Delta \rho_t, \\ \partial_t \nu_t &= -\nabla \cdot (\nu_t b_t) + \frac{c}{2} \Delta \nu_t \end{aligned}$$

for some vector fields $(b_t)_{t \geq 0}$ and constant $c \geq 0$ so Assumption 1 is satisfied. Then for any $t \geq 0$:

$$\begin{aligned} \frac{d}{dt} \text{FI}(\rho_t \parallel \nu_t) &= -c \mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{\nu_t} \right\|_{\text{HS}}^2 \right] \\ &\quad - 2 \mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|_{(-c \nabla^2 \log \nu_t + (\nabla b_t)_{\text{sym}})}^2 \right]. \quad (17) \end{aligned}$$

We consider a few special cases of interest:

- 1) **Langevin dynamics:** Consider when $b_t = -\nabla g$ and $c = 2$, so the channel is the Langevin dynamics (7) for target distribution $\nu \propto e^{-g}$. Suppose that $\nu_0 = \nu$, so $\nu_t = \nu$ for all $t \geq 0$. In this case, $-c \nabla^2 \log \nu_t + (\nabla b_t)_{\text{sym}} = 2 \nabla^2 g - \nabla^2 g = \nabla^2 g$, so the identity (17) recovers the time derivative formula (13) along the Langevin dynamics from Section II-E.

- 2) **Gaussian channel:** When $b_t = 0$ and $c = 1$, the identity (17) implies (SD)PIs in relative Fisher information for the Gaussian channel under (strong) log-concavity assumptions on ν_0 ; see Section III-B.
- 3) **Reverse Gaussian channel:** Consider when $b_t = \nabla \log \nu_t$ and $c = 1$, so the Fokker-Planck channel is the reverse Gaussian channel (6). In this case, $-c\nabla^2 \log \nu_t + (\nabla b_t)_{\text{sym}} = 0$. Thus, DPI always holds along the reverse Gaussian channel. We use this property to prove the mixing time guarantee of the Proximal Sampler in relative Fisher information; see Section IV-C.
- 4) **OU channel:** When $b_t(x) = -\gamma x$ and $c = 2$, the identity (17) implies SDPI in relative Fisher information eventually holds for the OU channel (9) under strong log-concavity assumption on ν_0 ; see Section III-C.

B. SDPI in Relative Fisher Information Along the Gaussian Channel

In Theorem 1 below we show for relative Fisher information along the Gaussian channel: (i) DPI holds when the second distribution is log-concave; (ii) SDPI holds when the second distribution is SLC; (iii) SDPI with an improved rate holds when the first distribution additionally satisfies Poincaré inequality and symmetry; and (iv) SDPI eventually holds when the second distribution is a log-Lipschitz perturbation of an SLC distribution. The proofs use the time derivative formula from Lemma 2; for part (iv), we also invoke the result of [45] on the evolution of log-concavity along the heat flow. We provide the proof of Theorem 1 in Appendix F-B. We present an application of Theorem 1 for analyzing the Proximal Sampler algorithm in Section IV.

Theorem 1: Let ρ_0, ν_0 be probability distributions on \mathbb{R}^d with $\text{FI}(\rho_0 \| \nu_0) < \infty$. Let ρ_t, ν_t be the output of the Gaussian channel (3) from ρ_0, ν_0 , respectively, at time $t > 0$. Then we have:

- (i) Assume ν_0 is log-concave. Then for all $t \geq 0$:

$$\text{FI}(\rho_t \| \nu_t) \leq \text{FI}(\rho_0 \| \nu_0).$$

- (ii) Assume ν_0 is α -SLC for some $\alpha > 0$. Then for all $t \geq 0$:

$$\text{FI}(\rho_t \| \nu_t) \leq \frac{\text{FI}(\rho_0 \| \nu_0)}{(1 + \alpha t)^2}.$$

- (iii) Assume ν_0 is α -SLC and ρ_0 satisfies β -PI for some $\alpha, \beta > 0$. Assume ρ_0, ν_0 are symmetric. Then for all $t \geq 0$:

$$\text{FI}(\rho_t \| \nu_t) \leq \frac{\text{FI}(\rho_0 \| \nu_0)}{(1 + \beta t)(1 + \alpha t)^2}.$$

- (iv) Assume $\nu_0 \propto e^{-g-\psi}$ where $g: \mathbb{R}^d \rightarrow \mathbb{R}$ is α -strongly convex and $\psi: \mathbb{R}^d \rightarrow \mathbb{R}$ is L -Lipschitz for some $\alpha > 0$, $0 \leq L < \infty$. Then for all $t \geq 0$:

$$\text{FI}(\rho_t \| \nu_t) \leq \frac{\text{FI}(\rho_0 \| \nu_0)}{(1 + \alpha t)^2} \exp\left(\frac{2tL^2}{\alpha t + 1} + \frac{8L\sqrt{t}}{\sqrt{\alpha t + 1}}\right). \quad (18)$$

We can verify the rate in part (ii) is tight by a calculation in the Gaussian case, and that we get an improved rate under Poincaré and symmetry as in part (iii), see Example 2 in Section III-B.1.

In Theorem 1 part (iv) above, we note that the upper bound (18) is initially increasing for small $t > 0$, and it is decreasing for large t , so it shows we eventually have SDPI. Furthermore, for all $t \geq 0$, we can further upper bound (18) by: $\text{FI}(\rho_t \| \nu_t) \leq C \cdot \frac{\text{FI}(\rho_0 \| \nu_0)}{(1 + \alpha t)^2}$ where $C = \exp\left(\frac{2L^2}{\alpha} + \frac{8L}{\sqrt{\alpha}}\right)$. We can construct an explicit example where the relative Fisher information along the heat flow indeed exhibits this behavior of initial increase and eventual decrease, see Example 3 below.

1) *Example of SDPI in Fisher Information Along Gaussian Channel:*

Example 2: Let $\rho_0 = \mathcal{N}(m, sI)$ and $\nu_0 = \mathcal{N}(0, I)$ for some $m \in \mathbb{R}^d$ and $s > 0$, $s \neq 1$. Note that ν_0 is 1-SLC, ν_0 is symmetric, and ρ_0 satisfies $(1/s)$ -PI. Along the Gaussian channel (3), we know $\rho_t = \mathcal{N}(m, (s + t)I)$ and $\nu_t = \mathcal{N}(0, (1 + t)I)$. Then we can compute:

$$\text{FI}(\rho_t \| \nu_t) = \frac{(s - 1)^2 d}{(1 + t)^2 (s + t)^2} + \frac{\|m\|^2}{(1 + t)^2}.$$

We note: (1) If $m \neq 0$, then $\text{FI}(\rho_t \| \nu_t) = O(t^{-2})$ as predicted in Theorem 1(ii). (2) If $m = 0$, then ρ_0 is symmetric, so Theorem 1(iii) predicts an improved rate of $O(t^{-3})$; in this example, we actually have a faster rate of $\text{FI}(\rho_t \| \nu_t) = O(t^{-4})$.

We remark that the faster rate $\text{FI}(\rho_t \| \nu_t) = O(t^{-4})$ in the case $m = 0$ above is likely because ρ_0 satisfies an improved Poincaré inequality for symmetric test functions, which follows from the improved LSI result by [46].

2) *Counterexample of DPI in Fisher Information Along Gaussian Channel:* We show DPI for relative Fisher information does *not* always hold even for the Gaussian channel. Note by Theorem 1(i), the counterexample needs the second distribution ν_0 to be non-log-concave.

Example 3: Let $\rho_0 = \mathcal{N}(0, 1)$ on \mathbb{R} ($d = 1$). Let $\nu_0 \propto e^{-g}$ where $g: \mathbb{R} \rightarrow \mathbb{R}$ is defined by:

$$g(x) = \begin{cases} -\frac{M}{2}x^2 & \text{if } |x| \leq L, \\ \frac{1}{2}(x - L)^2 - ML(x - L) - \frac{ML^2}{2} & \text{if } x > L, \\ \frac{1}{2}(x + L)^2 + ML(x + L) - \frac{ML^2}{2} & \text{if } x < -L, \end{cases}$$

where $M, L \geq 2$ are arbitrary. Then along the Gaussian channel (3), for small $t > 0$, the relative Fisher information is increasing: $\text{FI}(\rho_t \| \nu_t) > \text{FI}(\rho_0 \| \nu_0)$. We prove this in Proposition 1 in Appendix F-C. See Figure 1 for an illustration.

Figure 1 shows an illustration of Example 3 above. Figure 1(a) shows the density of ν_0 as defined in Proposition 1 when $M = L = 2$, with $\rho_0 = \mathcal{N}(0, 1)$. Figure 1(b) shows the plot of $t \mapsto \text{FI}(\rho_t \| \nu_t)$, which is initially increasing (hence DPI does not hold), then eventually decreasing. Recall $\frac{d}{dt} \text{KL}(\rho_t \| \nu_t) = -\frac{1}{2} \text{FI}(\rho_t \| \nu_t)$ along the Gaussian channel. This means the KL divergence (which is always decreasing along the Gaussian channel by classical DPI) is initially decreasing in a concave way, then eventually in a convex way, as shown in Figure 1(c).

C. Eventual SDPI in Relative Fisher Information Along the OU Channel

We derive the following results on the eventual SDPI for the relative Fisher information along OU channel (9) under strong log-concavity. In Theorem 2 below we show along the OU channel: (i) SDPI for relative Fisher information eventually




Fig. 1. An example where DPI in relative Fisher information does not hold along the Gaussian channel; see Example 3 and Proposition 1 for details.




Fig. 2. An example where DPI in relative Fisher information does not hold along the OU channel. Here $\rho_0 = \mathcal{N}(0, 0.01)$, $\nu_0 = \mathcal{N}(0, 10)$, and the OU channel is targeting $\mathcal{N}(0, 1)$. See Example 4 for details.

holds with rate $O(e^{-2\gamma t})$ when the second distribution is strongly log-concave; and (ii) SDPI eventually holds with an improved rate $O(e^{-4\gamma t})$ when the first distribution additionally satisfies Poincaré inequality and both distributions are symmetric. We provide the proof of Theorem 2 in Appendix F-D. However, even when both distributions are Gaussian, the relative Fisher information can be initially increasing (see Figure 2 and Example 4).

Theorem 2: Let ρ_0, ν_0 be probability distributions on \mathbb{R}^d with $\text{FI}(\rho_0 \parallel \nu_0) < \infty$. Along the OU channel (9) for target distribution $\mathcal{N}(0, \gamma^{-1}I)$, let ρ_t, ν_t be the output distributions at time $t > 0$ from input distributions ρ_0, ν_0 , respectively. Then we have:

- (i) Assume ν_0 is α -SLC for some $\alpha > 0$. Then for all $t \geq 0$:

$$\text{FI}(\rho_t \parallel \nu_t) \leq \frac{\gamma^2 e^{-2\gamma t}}{(\alpha + e^{-2\gamma t}(\gamma - \alpha))^2} \text{FI}(\rho_0 \parallel \nu_0).$$

- (ii) Assume ν_0 is α -SLC and ρ_0 satisfies β -PI for some $\alpha, \beta > 0$. Assume ρ_0, ν_0 are symmetric. Then for all $t \geq 0$, we have:

$$\text{FI}(\rho_t \parallel \nu_t) \leq \frac{\gamma^3 e^{-4\gamma t} \text{FI}(\rho_0 \parallel \nu_0)}{(\beta + e^{-2\gamma t}(\gamma - \beta))(\alpha + e^{-2\gamma t}(\gamma - \alpha))^2}.$$

We make a few remarks: Recall that as $\gamma \rightarrow 0$, the OU channel recovers the Gaussian channel (3) at twice the time speed. We note that as $\gamma \rightarrow 0$, the upper bound in

Theorem 2(i) recovers the upper bound for the Gaussian channel in Theorem 1(ii) at time $2t$, and the upper bound in Theorem 2(ii) recovers the upper bound for the Gaussian channel in Theorem 1(iii) at time $2t$.

We also recall from Theorem 1(i) that for the Gaussian channel, DPI for relative Fisher information holds when the second distribution is weakly log-concave. In contrast, for the OU channel, we are unable to show DPI when the second distribution is only weakly log-concave. By taking $\alpha \rightarrow 0$ in Theorem 2(i), the upper bound becomes $e^{2\gamma t}$ which is increasing in t since $\gamma > 0$, so we cannot conclude DPI (but as we further take $\gamma \rightarrow 0$, this upper bound becomes 1, which agrees with the DPI in Theorem 1(i) as the OU channel recovers the Gaussian channel).

The SDPI of relative Fisher information along the OU channel has an application in the stability of the Kim–Milman (probability flow) map between distributions [47].

1) Example of Eventual SDPI Along the OU Channel:

We can verify that the contraction rates in Theorem 2 are tight for large t by explicit calculation in the Gaussian case; see Example 4. However, we also note that even when both distributions are strongly log-concave, SDPI for relative Fisher information may not hold for small t ; see Figure 2.

Example 4 (Eventual SDPI along the OU channel): Let $\rho_0 = \mathcal{N}(m, \beta^{-1}I)$ and $\nu_0 = \mathcal{N}(0, \alpha^{-1}I)$ for some $m \in \mathbb{R}^d$ and $\alpha, \beta > 0$. Note that ν_0 is α -SLC and symmetric, and ρ_0 is β -SLC so it satisfies β -PI. Along the OU channel (9) to target distribution $\mathcal{N}(0, \gamma^{-1}I)$, the solutions are:

$$\rho_t = \mathcal{N}(e^{-\gamma t}m, \beta_t^{-1}I), \quad \nu_t = \mathcal{N}(0, \alpha_t^{-1}I)$$

where $\beta_t^{-1} := e^{-2\gamma t}\beta^{-1} + (1 - e^{-2\gamma t})\gamma^{-1}$, and $\alpha_t^{-1} := e^{-2\gamma t}\alpha^{-1} + (1 - e^{-2\gamma t})\gamma^{-1}$. Note $\min\{\beta, \gamma\} \leq \beta_t \leq \max\{\beta, \gamma\}$ and $\min\{\alpha, \gamma\} \leq \alpha_t \leq \max\{\alpha, \gamma\}$, so $\beta_t = \Theta(1)$ and $\alpha_t = \Theta(1)$ for all $t \geq 0$. Furthermore, $\frac{1}{\beta_t}(\alpha_t - \beta_t)^2 = \alpha_t^2\beta_t(\beta_t^{-1} - \alpha_t^{-1})^2 = e^{-4\gamma t}\alpha_t^2\beta_t(\beta^{-1} - \alpha^{-1})^2$. Then we can compute:

$$\begin{aligned} \text{FI}(\rho_t \parallel \nu_t) &= e^{-2\gamma t}\beta_t^2\|m\|^2 + \frac{d}{\beta_t}(\alpha_t - \beta_t)^2 \\ &= e^{-2\gamma t}\beta_t^2\|m\|^2 + e^{-4\gamma t}d\alpha_t^2\beta_t(\beta^{-1} - \alpha^{-1})^2. \end{aligned} \quad (19)$$

We observe for large t : (1) If $m \neq 0$, then $\text{FI}(\rho_t \parallel \nu_t) = O(e^{-2\gamma t})$ as predicted in Theorem 2(i); and (2) If $m = 0$, then ρ_0 is symmetric, and $\text{FI}(\rho_t \parallel \nu_t) = O(e^{-4\gamma t})$ as predicted in Theorem 2(ii).

However, note that for small $t > 0$, the value (19) can be increasing. Consider in $d = 1$ dimension when $m = 0$, $\gamma = 1$, $\beta = 100$, and $\alpha = 0.1$. See the plot in Figure 2, which shows the relative Fisher information is initially increasing.

We also note that since the initial behaviors of the OU channel and the Gaussian channel are similar, the same example from Example 3 (for the counterexample for DPI along the Gaussian channel) also shows that the relative Fisher information is initially increasing along the OU channel; this can be verified by the same calculation as in the proof of Proposition 1, which we omit.

IV. APPLICATION: ANALYSIS OF THE PROXIMAL SAMPLER

As an algorithmic application of our result on the SDPI along the heat flow, we prove a mixing time result of the Proximal Sampler algorithm for sampling with error guarantee in relative Fisher information. We review the Proximal Sampler in Section IV-A and its convergence guarantee in KL divergence in Lemma 3 in Section IV-B. We present our new result on the convergence rate of the Proximal Sampler in relative Fisher information in Theorem 3 in Section IV-C; this is the relative Fisher information analogue of the KL divergence guarantee in Lemma 3.

A. Review of the Proximal Sampler

Suppose we want to sample from a target probability distribution $\nu^X(x) \propto e^{-g(x)}$ on \mathbb{R}^d . We define a target joint distribution ν^{XY} on the extended state space $\mathbb{R}^{2d} = \mathbb{R}^d \times \mathbb{R}^d$ with density function:

$$\nu^{XY}(x, y) \propto \exp\left(-g(x) - \frac{1}{2\eta}\|y - x\|^2\right) \quad (20)$$

where $\eta > 0$ is the *step size*. Observe that the X -marginal of ν^{XY} is equal to ν^X . Therefore, if we can sample $(X, Y) \sim \nu^{XY}$, then by ignoring the Y component, we obtain a sample $X \sim \nu^X$.

The *Proximal Sampler* [3] for sampling from ν^X performs Gibbs sampling to sample from the joint distribution ν^{XY} . In each iteration $k \geq 0$, from the current iterate $x_k \sim \rho_k^X$, the Proximal Sampler performs the following two steps:

- 1) (*Forward step*): Sample

$$y_k | x_k \sim \nu^{Y|X}(\cdot | x_k) = \mathcal{N}(x_k, \eta I).$$

This results in a new iterate $y_k \sim \rho_k^Y$, where $\rho_k^Y = \rho_k^X * \mathcal{N}(0, \eta I)$.

- 2) (*Backward step*): Sample

$$x_{k+1} | y_k \sim \nu^{X|Y}(\cdot | y_k).$$

This gives the next iterate $x_{k+1} \sim \rho_{k+1}^X$.

Since Gibbs sampling is a reversible Markov chain, ν^{XY} is a stationary distribution of the forward and backward steps above. Thus, the Proximal Sampler is *unbiased*: If $\rho_k^X = \nu^X$, then $\rho_{k+1}^X = \nu^X$.

We note the forward step is easy to implement by adding an independent Gaussian noise. The backward step requires sampling from the conditional distribution $\nu^{X|Y}$ with density:

$$\nu^{X|Y}(x | y_k) \propto_x \exp\left(-g(x) - \frac{1}{2\eta}\|x - y_k\|^2\right). \quad (21)$$

We assume we can sample from $\nu^{X|Y}$ (21); this is called the *Restricted Gaussian Oracle (RGO)* in [3]. When ν^X is log-smooth (i.e., when g has bounded second derivatives), for sufficiently small step size η , we can implement the RGO via rejection sampling with an $O(1)$ expected number of queries to g ; see Appendix G-A for a review.

B. Review of the Mixing Time of Proximal Sampler in KL Divergence

We can interpret each step of the Proximal Sampler as the evolution along a Fokker-Planck channel:

- 1) The forward step is the result of passing ρ_k^X via the Gaussian channel (3) for time η to obtain $\rho_k^Y = \rho_k^X * \mathcal{N}(0, \eta I)$. We also apply the forward step to ν^X to obtain $\nu^Y = \nu^X * \mathcal{N}(0, \eta I)$.
- 2) The backward step is the result of passing ρ_k^Y via the reverse Gaussian channel (6) for time η to obtain ρ_{k+1}^X . When we apply the backward step to ν^Y , we obtain ν^X back.

By using the interpretations above and invoking the SDPI for the forward and reverse Gaussian channels, [23] show the following mixing time for Proximal Sampler in KL divergence under LSI, where each forward and backward step induces a contraction in KL divergence.

Lemma 3 ([23, Theorem 3]): Assume ν^X satisfies α -LSI for some $\alpha > 0$. Along each iteration of the Proximal Sampler with any $\eta > 0$, we have:

$$\text{KL}(\rho_k^Y \| \nu^Y) \leq \frac{\text{KL}(\rho_k^X \| \nu^X)}{(1 + \alpha\eta)} \quad (22a)$$

$$\text{KL}(\rho_{k+1}^X \| \nu^X) \leq \frac{\text{KL}(\rho_k^Y \| \nu^Y)}{(1 + \alpha\eta)}. \quad (22b)$$

Therefore, for any $k \geq 0$, the KL divergence converges exponentially fast:

$$\text{KL}(\rho_k^X \| \nu^X) \leq \frac{\text{KL}(\rho_0^X \| \nu^X)}{(1 + \alpha\eta)^{2k}}. \quad (23)$$

Note the exponential convergence rate in (23) matches the exponential convergence rate of the KL divergence along the Langevin dynamics (12), since for small $\eta > 0$, $(1 + \alpha\eta)^{-2k} \approx e^{-2\alpha\eta k}$.

C. Convergence Rate of the Proximal Sampler in Relative Fisher Information

We show that when the target distribution is strongly log-concave, relative Fisher information converges exponentially fast along the Proximal Sampler. We use the same notation and interpretation of the forward and backward steps in Section IV-B. We provide the proof of Theorem 3 in Appendix G-B.

Theorem 3: Assume ν^X is α -strongly log-concave for some $\alpha > 0$. Along each iteration of the Proximal Sampler with any $\eta > 0$, we have:

$$\text{FI}(\rho_k^Y \| \nu^Y) \leq \frac{\text{FI}(\rho_k^X \| \nu^X)}{(1 + \alpha\eta)^2} \quad (24a)$$

$$\text{FI}(\rho_{k+1}^X \| \nu^X) \leq \text{FI}(\rho_k^Y \| \nu^Y). \quad (24b)$$

Therefore, for any $k \geq 0$, the relative Fisher information converges exponentially fast:

$$\text{FI}(\rho_k^X \parallel \nu^X) \leq \frac{\text{FI}(\rho_0^X \parallel \nu^X)}{(1 + \alpha\eta)^{2k}}. \quad (25)$$

Note the exponential convergence rate in (25) matches the exponential convergence rate of the relative Fisher information along the Langevin dynamics (14). The convergence rate in (25) also matches the exponential convergence rate of the squared gradient norm along the proximal gradient method in optimization for a strongly convex objective function (see Lemma 4). We can verify the convergence rate in (25) is tight in the Gaussian case; see Example 5 in Section IV-C.1.

We note the convergence rate (25) in relative Fisher information is similar to the convergence rate (23) in KL divergence. However, we also note the difference: For KL divergence, there is a contraction in both the forward and backward steps (22), since we have SDPI for both the forward and reverse Gaussian channels. For relative Fisher information (24), there is a contraction only in the forward step since we have SDPI along the Gaussian channel (Theorem 1), while the backward step does not provide a strict contraction since we only have DPI along the reverse Gaussian channel.

1) *Convergence of Relative Fisher Information Along the Proximal Sampler:*

Example 5: Consider the Proximal Sampler with step size $\eta > 0$ for the target distribution $\nu^X = \mathcal{N}(0, \alpha^{-1}I)$ on \mathbb{R}^d , which is α -SLC, from the initial distribution $\rho_0^X = \mathcal{N}(m_0, \sigma_0^2 I)$ for some $m_0 \in \mathbb{R}^d$, $m_0 \neq 0$, and $\sigma_0^2 > 0$. In this case, the RGO conditional distribution is, for each $y \in \mathbb{R}^d$:

$$\nu^{X|Y}(\cdot | y) = \mathcal{N}\left(\frac{1}{1 + \alpha\eta}y, \frac{\eta}{1 + \alpha\eta}I\right).$$

Therefore, the iterates remain Gaussian $\rho_k^X = \mathcal{N}(m_k, \sigma_k^2 I)$ for all $k \geq 0$, where:

$$m_k = \frac{m_0}{(1 + \alpha\eta)^k}, \quad \sigma_k^2 = \frac{1}{(1 + \alpha\eta)^{2k}} \left(\sigma_0^2 - \frac{1}{\alpha} \right) + \frac{1}{\alpha}.$$

Then we can compute:

$$\begin{aligned} \text{FI}(\rho_k^X \parallel \nu^X) &= \alpha^2 \|m_k\|^2 + \left(\frac{1}{\sigma_k^2} - \alpha \right)^2 \sigma_k^2 d \\ &= \frac{\alpha^2 \|m_0\|^2}{(1 + \alpha\eta)^{2k}} + \frac{d}{\sigma_k^2} \frac{(1 - \alpha\sigma_0^2)^2}{(1 + \alpha\eta)^{4k}}. \end{aligned}$$

Since $m_0 \neq 0$, we have $\text{FI}(\rho_k^X \parallel \nu^X) = O((1 + \alpha\eta)^{-2k})$, as claimed in Theorem 3. (Note if $m_0 = 0$, then we have an improved rate $\text{FI}(\rho_k^X \parallel \nu^X) = O((1 + \alpha\eta)^{-4k})$; this is because in this case ρ_k^X is symmetric and satisfies a Poincaré inequality, so the improved SDPI from Theorem 1(iii) applies.)

D. Iteration Complexity of the Proximal Sampler in Relative Fisher Information

With the rejection sampling implementation of RGO (see Appendix G-A), the convergence rate in Theorem 3 implies the following high-accuracy iteration complexity guarantee for the Proximal Sampler in relative Fisher information. We provide the proof of Corollary 1 in Appendix G-C.

Corollary 1: Assume $\nu^X \propto e^{-g}$ is α -strongly log-concave and L -log-smooth for some $0 < \alpha \leq L < \infty$. Let $x^* = \arg \min_{x \in \mathbb{R}^d} g(x)$ be the minimizer of g . Consider the Proximal Sampler with step size $\eta = \frac{1}{Ld}$ with the rejection sampling implementation of the RGO. Suppose we start the algorithm from $x_0 \sim \rho_0^X = \mathcal{N}(x^*, \frac{1}{L}I)$. Then for any $\varepsilon > 0$, along the iteration of the Proximal Sampler $x_k \sim \rho_k^X$, we have $\text{FI}(\rho_k^X \parallel \nu^X) \leq \varepsilon$ whenever the number of iterations k satisfies

$$k \geq \frac{dL}{\alpha} \log \frac{dL}{\varepsilon}$$

and the expected number of queries to g in each iteration for the rejection sampling is $O(1)$.

In the above, we assume we can find the minimizer x^* of g . We also assume we can find the minimizer x_k^* for the regularized objective in each implementation of the RGO; see discussion in Appendix G-A for how this optimization error does not affect the overall iteration complexity.

We note the logarithmic dependence of the iteration complexity above on the error parameter ε^{-1} , which comes from the exponential convergence rate of the Proximal Sampler in relative Fisher information. This is in contrast to the polynomial dependence on ε^{-1} in the result of [27], which comes from the bias of ULA. However, the result of [27] only assumes the target is log-smooth, while our result assumes the target is SLC and log-smooth. Furthermore, our guarantee holds for the last iterate of the Proximal Sampler, while the guarantee in [27] holds for the average iterate of ULA.

V. DISCUSSION

We have shown that the Proximal Sampler has a high-accuracy iteration complexity guarantee for sampling in relative Fisher information when the target distribution is strongly log-concave and log-smooth. This complements the result of [27], who show a low-accuracy iteration complexity guarantee of ULA for non-log-concave, log-smooth target distribution.

Our analysis technique is to establish SDPI in relative Fisher information along several channels, including the Gaussian channel under strong log-concavity, and along the reverse Gaussian channel, which compose to form the Proximal Sampler algorithm. We show that DPI in relative Fisher information may not hold initially along the Gaussian channel and the OU process. We also show that SDPI in relative Fisher information will eventually hold along the Gaussian channel when the second input distribution is a log-Lipschitz perturbation of a strongly log-concave distribution. We have also shown SDPI in relative Fisher information along the OU process when the second input distribution is strongly log-concave; this has found an application in the stability analysis of the Kim–Milman (probability flow) map between distributions [47].

Going forward, it would be interesting to study if we can weaken the SLC assumption to isoperimetry such as LSI. It would be interesting to show mixing time in relative Fisher information for other sampling algorithms beyond what we study in this paper, including for discretization of the underdamped Langevin dynamics [48] and for the Hamiltonian Monte Carlo algorithm [49].

APPENDIX A
ADDITIONAL RELATED WORKS

Mixing time guarantees in relative Fisher information have been studied in [27] and [28], both focusing on non-log-concave sampling as we discussed in Section I. The work of [50] use convergence in relative Fisher information to characterize conditional mixing for non-log-concave target distributions.

For other sampling algorithms, mixing time analyses based on hypocoercivity [51] involve the construction of a Lyapunov function which combines the KL divergence and a rescaled relative Fisher information; see [48] for the analysis of the underdamped Langevin algorithm, and [49] for the analysis of the ideal Hamiltonian Monte Carlo algorithm, where both work only assume the target distribution satisfy LSI and smoothness. From such Lyapunov analyses, one can in principle extract convergence guarantees in relative Fisher information for the corresponding algorithms. However, the result will be a biased convergence guarantee for the underdamped Langevin algorithm in [48] since the discretization introduces bias similar to ULA. On the other hand, the result will be an exponential convergence guarantee for the ideal Hamiltonian Monte Carlo algorithm in [49] since the algorithm is unbiased, but the algorithm is still idealized and requires a further discretization step to implement in practice, and it is unclear whether the convergence guarantee in relative Fisher information will be preserved.

Convergence guarantees in relative Fisher information have been useful as an intermediate step in showing the uniqueness of the equilibrium distribution and analyzing the convergence behaviors of the mean-field Langevin dynamics for zero-sum games in the space of probability distributions under strong convexity, see e.g. [52, Lemma 7.2] and [53, Theorem 4].

The work of [54] study the decay of Fisher information along a general class of stochastic processes which allow a time and position-dependent covariance term multiplying the (possibly degenerate) Brownian motion; this includes Fokker-Planck channels as a special case. Reference [54] derive a formula for the time derivative of a generalized relative Fisher information along such processes, using the machinery of information gamma calculus, and extract convergence guarantees in relative Fisher information under certain assumptions, see [54, Theorem 2]; they provide various applications including analysis of annealed Langevin dynamics. In this work, we provide a self-contained proof of the time derivative formula along the Fokker-Planck channel (see Lemma 2); furthermore, our SDPI perspective and our algorithmic application for deriving the mixing time of the Proximal Sampler in relative Fisher information seem novel.

A. Parametric Vs. Non-Parametric Fisher Information

In this paper, we are working with the *non-parametric* relative Fisher information $\text{Fl}(\rho \parallel \nu)$, where the derivative of the log-density is taken with respect to the space variable; this has a natural geometric interpretation as the squared gradient norm of the KL divergence in the Wasserstein metric. The relative Fisher information is a relative form of the non-parametric Fisher information $J(\rho) = \mathbb{E}_\rho[\|\nabla \log \rho\|^2]$. More precisely, the

relative Fisher information $\text{Fl}(\rho \parallel \nu)$ is the Bregman divergence (i.e., excess in the linear approximation) of the absolute Fisher information; see e.g. [40, Lemma 2].

In the literature, another popular notion is the *parametric* Fisher information, where the derivative of the log-density is taken with respect to a parameter that indexes the distribution; for example, this is the notion of Fisher information that appears in the Cramer-Rao bound. When the parameter is a location parameter, i.e., the distribution is $\rho_\theta(x) = \rho_0(x - \theta)$ for some base distribution ρ_0 and parameter θ , then the parametric Fisher information coincides with the non-parametric Fisher information; however, in general they are different. In particular, the non-parametric Fisher information requires working with a continuous state space such as \mathbb{R}^d , while the parametric Fisher information is well-defined even for a discrete state space as long as the parameter is a continuous variable.

The data processing inequality for the parametric Fisher information is known under post-processing, see e.g. [37, Lemma 3]. For the non-parametric Fisher information $J(\rho)$, we have the *Fisher information inequality* or the Blachman-Stam inequality, which states that for any probability distributions $\rho, \nu \in \mathcal{P}(\mathbb{R}^d)$:

$$J(\rho * \nu)^{-1} \geq J(\rho)^{-1} + J(\nu)^{-1}$$

see e.g. [55] and [37, Theorem 1]. In particular, for any distribution ρ , we have $J(\rho * \nu) \leq J(\rho)$, which can be seen as a data processing inequality in non-parametric Fisher information for a noisy channel which performs convolution with ν (i.e., adding an independent noise drawn from ν). However, a version of the inequality above for relative Fisher information is not known, and does not appear to be true in general; in Example 3, we show an example where convolution with a Gaussian (i.e., along the Gaussian channel) increases relative Fisher information.

APPENDIX B
REVIEW OF SAMPLING AS OPTIMIZATION IN THE
SPACE OF DISTRIBUTIONS

In the perspective of sampling as optimization in the space of distributions [4], [5], the problem of sampling from a target probability distribution $\nu \in \mathcal{P}(\mathbb{R}^d)$ can be formulated as minimizing some objective functions $\mathcal{F}: \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}$ on the space of probability distributions which is minimized at the target distribution $\rho^* = \nu$; for example, a good objective function is the KL divergence $\mathcal{F}(\rho) = \text{KL}(\rho \parallel \nu)$.

We recall that the Langevin dynamics (7) to the target distribution $\nu \propto e^{-g}$, or more precisely its associated Fokker-Planck equation (8), has an optimization interpretation as the gradient flow dynamics:

$$\dot{\rho}_t = -\text{grad KL}(\rho_t \parallel \nu)$$

for minimizing the KL divergence objective function $\mathcal{F}(\rho) = \text{KL}(\rho \parallel \nu)$ on the space of probability distributions $\mathcal{P}(\mathbb{R}^d)$ endowed with the Wasserstein W_2 metric [4]. Here grad is the gradient operator with respect to the W_2 metric, following the computation rule of Otto calculus, see e.g. [39, Ch. 15].

We also recall the optimization interpretation of the definitions above [26]: if ν is α -SLC, then the KL divergence

$\rho \mapsto \text{KL}(\rho \parallel \nu)$ is an α -geodesically strongly convex function on the space of probability distributions on \mathbb{R}^d with respect to the W_2 metric. Furthermore, recall the relative Fisher information is the squared gradient norm of the KL divergence in the W_2 metric:

$$\text{FI}(\rho \parallel \nu) = \|\text{grad KL}(\rho \parallel \nu)\|_{\rho}^2.$$

The property that ν satisfies LSI is equivalent to the Polyak-Łojaciewicz (PL) or gradient domination property of the KL divergence objective function, which means for all probability distributions ρ :

$$\|\text{grad KL}(\rho \parallel \nu)\|_{\rho}^2 \geq 2\alpha \text{KL}(\rho \parallel \nu).$$

We recall that α -strong convexity implies α -gradient domination, and gradient domination is a sufficient condition for the exponential convergence rate in function value along gradient flow.

Via the gradient flow interpretation of the Langevin dynamics, de Bruijn's identity (11) can be seen as the energy identity along gradient flow:

$$\frac{d}{dt} \text{KL}(\rho_t \parallel \nu) = -\|\text{grad KL}(\rho_t \parallel \nu)\|_{\rho_t}^2.$$

Thus, the exponential convergence rate of KL divergence along the Langevin dynamics under LSI (12) is a manifestation of the fact that along the gradient flow for an objective function which is gradient dominated, the function value converges exponentially fast to its minimum.

Via the gradient flow interpretation of the Langevin dynamics ($\dot{\rho}_t = -\text{grad KL}(\rho_t \parallel \nu)$), the identity (13) also follows from a general time derivative computation along the gradient flow:

$$\begin{aligned} \frac{d}{dt} \|\text{grad KL}(\rho_t \parallel \nu)\|_{\rho_t}^2 \\ = -2(\text{Hess KL}(\rho_t \parallel \nu)) \left[(\text{grad KL}(\rho_t \parallel \nu))^{\otimes 2} \right] \end{aligned}$$

where Hess is the Hessian operator in the W_2 metric, which when evaluated yields the right-hand side of the identity (13) above; see [39, Ch. 15]. Thus, the exponential convergence rate (14) of the relative Fisher information along the Langevin dynamics under SLC is a manifestation of the fact that along the gradient flow for a strongly convex objective function, the squared gradient norm converges exponentially fast; see Lemma 4 in Appendix C for a review.

APPENDIX C REVIEW OF THE GRADIENT NORM CONVERGENCE IN OPTIMIZATION

We review the behavior of the squared gradient norm in finite-dimensional optimization.

Suppose we want to minimize an objective function $f: \mathbb{R}^d \rightarrow \mathbb{R}$. We assume f is twice differentiable and α -strongly convex, so the Hessian matrix satisfies $\nabla^2 f(x) \geq \alpha I$ for all $x \in \mathbb{R}^d$.

In continuous time, we consider the *gradient flow* dynamics for $(X_t)_{t \geq 0}$ in \mathbb{R}^d :

$$\dot{X}_t = -\nabla f(X_t). \quad (26)$$

In discrete time, we consider the *proximal gradient algorithm* for the iterates $x_k \in \mathbb{R}^d$:

$$x_{k+1} = \arg \min_{x \in \mathbb{R}^d} \left\{ f(x) + \frac{1}{2\eta} \|x - x_k\|^2 \right\} \quad (27)$$

where $\eta > 0$ is step size. Recall that the proximal gradient method algorithm can be equivalently written as the following implicit update:

$$x_{k+1} = x_k - \eta \nabla f(x_{k+1}).$$

We have the following guarantees for the convergence of the squared gradient norm along the gradient flow and the proximal gradient algorithm above. Note the rate for the proximal gradient algorithm below matches the convergence rate for the gradient flow: $(1 + \alpha\eta)^{2k} \rightarrow \exp(-2\alpha t)$, in the limit of vanishing step size $\eta \rightarrow 0$ with $\eta k \rightarrow t$.

Lemma 4: Assume f is α -strongly convex and twice continuously differentiable.

- 1) Along the gradient flow dynamics (26), we have for all $t \geq 0$:

$$\|\nabla f(X_t)\|^2 \leq e^{-2\alpha t} \|\nabla f(X_0)\|^2.$$

- 2) Along the proximal gradient algorithm (27) with any step size $\eta > 0$, we have for all $k \geq 0$:

$$\|\nabla f(x_k)\|^2 \leq \frac{\|\nabla f(x_0)\|^2}{(1 + \alpha\eta)^{2k}}.$$

Proof:

- (1) **Analysis for gradient flow:** Using the gradient flow dynamics (26), we can compute:

$$\begin{aligned} \frac{d}{dt} \|\nabla f(X_t)\|^2 &= 2\langle \nabla f(X_t), \nabla^2 f(X_t) \dot{X}_t \rangle \\ &= -2\|\nabla f(X_t)\|_{\nabla^2 f(X_t)}^2 \\ &\leq -2\alpha \|\nabla f(X_t)\|^2, \end{aligned}$$

where in the last step above we have used the property that $\nabla^2 f(X_t) \geq \alpha I$. We can write the differential inequality above equivalently as: $\frac{d}{dt} (e^{2\alpha t} \|\nabla f(X_t)\|^2) \leq 0$. Integrating from 0 to t and rearranging, we get the desired convergence bound along gradient flow:

$$\|\nabla f(X_t)\|^2 \leq e^{-2\alpha t} \|\nabla f(X_0)\|^2.$$

- (2) **Analysis for proximal gradient algorithm:** We use a continuous-time interpolation of each step of the proximal gradient algorithm as follows. From any $Y_0 \in \mathbb{R}^d$, consider the continuous-time curve $(Y_t)_{t \geq 0}$ defined by:

$$Y_t = Y_0 - t \nabla f(Y_t).$$

Observe that if $Y_0 = x_k$, then the next iterate x_{k+1} of the proximal gradient algorithm is equal to the value Y_η at time $t = \eta$, since indeed at $t = \eta$, we have $Y_\eta = Y_0 - \eta \nabla f(Y_\eta) = x_k - \eta \nabla f(Y_\eta)$, so $Y_\eta = (I + \eta \nabla f)^{-1}(x_k) = x_{k+1}$. Along the curve Y_t above, we can compute:

$$\dot{Y}_t = -\nabla f(Y_t) - t \nabla^2 f(Y_t) \dot{Y}_t$$

and therefore,

$$\dot{Y}_t = -(I + t \nabla^2 f(Y_t))^{-1} \nabla f(Y_t).$$

Therefore, we can compute along the curve Y_t above:

$$\begin{aligned} \frac{d}{dt} \|\nabla f(Y_t)\|^2 &= 2\langle \nabla f(Y_t), \nabla^2 f(Y_t) \dot{Y}_t \rangle \\ &= -2\|\nabla f(Y_t)\|_{\nabla^2 f(Y_t)(I+t\nabla^2 f(Y_t))^{-1}}^2. \end{aligned}$$

Note since $\nabla^2 f(Y_t) \geq \alpha I$, we have

$$\nabla^2 f(Y_t)(I+t\nabla^2 f(Y_t))^{-1} \geq \frac{\alpha}{1+\alpha t} I$$

for all $t \geq 0$. Therefore, continuing the calculation above,

$$\frac{d}{dt} \|\nabla f(Y_t)\|^2 \leq -\frac{2\alpha}{1+\alpha t} \|\nabla f(Y_t)\|^2 = -A_t \|\nabla f(Y_t)\|^2$$

where in the last step we have defined $A_t = 2 \log(1 + \alpha t)$, so $\dot{A}_t = \frac{2\alpha}{1+\alpha t}$. The above differential inequality is equivalent to: $\frac{d}{dt} (e^{A_t} \|\nabla f(Y_t)\|^2) \leq 0$. Integrating from 0 to t and rearranging yields:

$$\|\nabla f(Y_t)\|^2 \leq e^{-A_t} \|\nabla f(Y_0)\|^2 = \frac{\|\nabla f(Y_0)\|^2}{(1+\alpha t)^2}.$$

Plugging in $t = \eta$, and substituting $Y_0 = x_k$, $Y_\eta = x_{k+1}$, we obtain the one-step contraction bound for the proximal gradient algorithm:

$$\|\nabla f(x_{k+1})\|^2 \leq \frac{\|\nabla f(x_k)\|^2}{(1+\alpha\eta)^2}.$$

Unrolling this recursion gives the desired convergence bound for the proximal gradient algorithm. \square

APPENDIX D

LARGE GAP BETWEEN RELATIVE FISHER INFORMATION AND KL DIVERGENCE

We exhibit an example of two distributions where their relative Fisher information is large, but the KL divergence is small. In fact, we can make the Rényi divergence of order ∞ be small.

Recall the *Rényi divergence* of order $1 < q < \infty$ between probability distributions ρ and ν is:

$$\mathbf{R}_q(\rho \parallel \nu) = \frac{1}{q-1} \log \mathbb{E}_\nu \left[\left(\frac{\rho}{\nu} \right)^q \right].$$

Recall that $q \mapsto \mathbf{R}_q(\rho \parallel \nu)$ is increasing, $\lim_{q \rightarrow 1} \mathbf{R}_q(\rho \parallel \nu) = \text{KL}(\rho \parallel \nu)$, and the Rényi divergence of order ∞ is the logarithm of the *warmness* parameter:

$$\mathbf{R}_\infty(\rho \parallel \nu) = \lim_{q \rightarrow \infty} \mathbf{R}_q(\rho \parallel \nu) = \text{ess sup}_\nu \log \frac{\rho}{\nu}$$

where ess sup_ν denotes the supremum over the support of ν except on sets of measure zero with respect to ν . Therefore, if we can show $\mathbf{R}_\infty(\rho \parallel \nu)$ is small, then it implies $\mathbf{R}_q(\rho \parallel \nu)$ is small for all $1 < q < \infty$, and in particular, $\text{KL}(\rho \parallel \nu)$ is also small.

We have the following construction in $d = 1$ dimension, where the second distribution ν is a standard Gaussian (which is SLC and thus satisfies LSI), and the first distribution ρ is a small perturbation of ν , such that the ∞ -Rényi divergence can be made arbitrarily small, but the relative Fisher information can be made arbitrarily large.

Lemma 5: Let $\nu = \mathcal{N}(0, 1)$ on \mathbb{R} . Let $0 < \epsilon < 1 < L$ be arbitrary. There exists a probability distribution ρ on \mathbb{R} (constructed below) such that $\mathbf{R}_\infty(\rho \parallel \nu) \leq \epsilon$ and $\text{FI}(\rho \parallel \nu) \geq L$.

Proof: Define $a \equiv a(\epsilon) > 0$ such that $\nu([-a, a]) = \epsilon$. Define $M = \max\{1/a, \sqrt{(eL)/\epsilon}\}$. Let $K \geq 0$ be the smallest integer such that $(2K+1)/M \geq a$. Define $\eta := a/(2K+1)$, so $\eta \leq \frac{1}{M} \leq a$.

We define $g: \mathbb{R} \rightarrow \mathbb{R}$ to be a periodic spike function with height 1 and width 2η supported on the interval $[-a, a]$, that is:

$$g(x) = \begin{cases} 1 - |x - 2k\eta|/\eta & \text{if } |x| \leq a \\ & \text{and } x \in [(2k-1)\eta, (2k+1)\eta] \\ & \text{for some } k \in \{-K, \dots, K\}, \\ 0 & \text{if } |x| > a. \end{cases}$$

Define the probability distribution $\rho \propto \nu \cdot e^{-g}$, or explicitly:

$$\rho(x) = \frac{\nu(x) e^{-g(x)}}{\mathbb{E}_\nu[e^{-g}]}.$$

Since $0 \leq g(x) \leq 1$ for $|x| \leq a$ and $g(x) = 0$ for $|x| > a$, we have $\mathbb{E}_\nu[g] \leq \nu([-a, a]) = \epsilon$. Then by Jensen's inequality, $\log \mathbb{E}_\nu[e^{-g}] \geq -\mathbb{E}_\nu[g] \geq -\epsilon$. Then for all $x \in \mathbb{R}^d$, we have:

$$\log \frac{\rho(x)}{\nu(x)} = -g(x) - \log \mathbb{E}_\nu[e^{-g}] \leq 0 + \epsilon = \epsilon.$$

Therefore, $\mathbf{R}_\infty(\rho \parallel \nu) \leq \epsilon$, as claimed.

Next, since $g(x) \geq 0$, we have $\mathbb{E}_\nu[e^{-g}] \leq 1$. Since $|g'(x)| = 1/\eta$ almost everywhere for $|x| \leq a$ (except for x which are integer multiples of η) and $g'(x) = 0$ for $|x| > a$, and since $g(x) \leq 1$, we have

$$\begin{aligned} \mathbb{E}_\nu[(g')^2 e^{-g}] &= \frac{1}{\eta^2} \int_{-a}^a e^{-g(x)} \nu(x) dx \\ &\geq \frac{e^{-1}}{\eta^2} \nu([-a, a]) = \frac{e^{-1}\epsilon}{\eta^2} \geq e^{-1}\epsilon M^2 \geq L. \end{aligned}$$

Therefore, we can bound the relative Fisher information as

$$\begin{aligned} \text{FI}(\rho \parallel \nu) &= \mathbb{E}_\rho \left[\left(\log \frac{\rho}{\nu} \right)^2 \right] \\ &= \mathbb{E}_\rho[(g')^2] = \frac{\mathbb{E}_\nu[(g')^2 e^{-g}]}{\mathbb{E}_\nu[e^{-g}]} \geq L \end{aligned}$$

as claimed. \square

APPENDIX E

PROOF OF LEMMA 1 (SDPI FOR KL DIVERGENCE ALONG THE FP CHANNEL)

We provide the proof assuming the distributions have the regularity in Assumption 1, which ensures the validity of the operations in the proof below, including differentiation under the integral sign and the vanishing of the boundary terms in integration by parts. See [44] for a rigorous derivation of the identity in the case of the heat flow (the case $b_t = 0$).

Proof of Lemma 1: We can compute by differentiating under the integral sign and using chain rule:

$$\begin{aligned} \frac{d}{dt} \text{KL}(\rho_t \parallel \nu_t) &= \frac{d}{dt} \int_{\mathbb{R}^d} \rho_t \log \frac{\rho_t}{\nu_t} dx \\ &= \int_{\mathbb{R}^d} (\partial_t \rho_t) \log \frac{\rho_t}{\nu_t} dx + \int_{\mathbb{R}^d} \rho_t \frac{1}{\rho_t} \partial_t \rho_t dx - \int_{\mathbb{R}^d} \rho_t \frac{1}{\nu_t} \partial_t \nu_t dx. \end{aligned} \quad (28)$$

We compute each term above. For the first term in (28), by the Fokker-Planck equation and using integration by parts, where all the boundary terms vanish:

$$\begin{aligned} & \int_{\mathbb{R}^d} (\partial_t \rho_t) \log \frac{\rho_t}{v_t} dx \\ &= \int_{\mathbb{R}^d} \left(-\nabla \cdot (\rho_t b_t) + \frac{c}{2} \Delta \rho_t \right) \log \frac{\rho_t}{v_t} dx \\ &= \int_{\mathbb{R}^d} \rho_t \left\langle b_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle dx \\ & \quad - \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \left\langle \nabla \log \rho_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle dx \end{aligned}$$

where for the second term above we have used the identity that $\Delta \rho_t = \nabla \cdot \nabla \rho_t = \nabla \cdot (\rho_t \nabla \log \rho_t)$.

For the second term in (28), we can show it is equal to 0:

$$\begin{aligned} \int_{\mathbb{R}^d} \rho_t \frac{1}{\rho_t} \partial_t \rho_t dx &= \int_{\mathbb{R}^d} \partial_t \rho_t dx \\ &= \partial_t \left(\int_{\mathbb{R}^d} \rho_t dx \right) = \partial_t(1) = 0. \end{aligned}$$

For the third term in (28), by the Fokker-Planck equation and using integration by parts:

$$\begin{aligned} & - \int_{\mathbb{R}^d} \rho_t \frac{1}{v_t} \partial_t v_t dx \\ &= - \int_{\mathbb{R}^d} \left(-\nabla \cdot (v_t b_t) + \frac{c}{2} \Delta v_t \right) \frac{\rho_t}{v_t} dx \\ &= - \int_{\mathbb{R}^d} v_t \left\langle b_t, \nabla \frac{\rho_t}{v_t} \right\rangle dx + \frac{c}{2} \int_{\mathbb{R}^d} v_t \left\langle \nabla \log v_t, \nabla \frac{\rho_t}{v_t} \right\rangle dx \\ &= - \int_{\mathbb{R}^d} \rho_t \left\langle b_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle dx \\ & \quad + \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \left\langle \nabla \log v_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle dx \end{aligned}$$

where in the above we have used the identity that $\Delta v_t = \nabla \cdot (v_t \nabla \log v_t)$, and $v_t \nabla \frac{\rho_t}{v_t} = \rho_t \nabla \log \frac{\rho_t}{v_t}$.

Combining the calculations above in (28), we see that the terms involving $\left\langle b_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle$ cancel, and we obtain:

$$\begin{aligned} \frac{d}{dt} \text{KL}(\rho_t \| v_t) &= -\frac{c}{2} \int_{\mathbb{R}^d} \rho_t \left\langle \nabla \log \rho_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle dx \\ & \quad + \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \left\langle \nabla \log v_t, \nabla \log \frac{\rho_t}{v_t} \right\rangle dx \\ &= -\frac{c}{2} \int_{\mathbb{R}^d} \rho_t \left\| \nabla \log \frac{\rho_t}{v_t} \right\|^2 dx \\ &= -\frac{c}{2} \text{FI}(\rho_t \| v_t) \end{aligned}$$

as claimed in (15).

Now suppose for all $t \geq 0$ we know that v_t satisfies α_t -LSI. Then from the identity above, we can bound:

$$\frac{d}{dt} \text{KL}(\rho_t \| v_t) = -\frac{c}{2} \text{FI}(\rho_t \| v_t) \leq -c\alpha_t \text{KL}(\rho_t \| v_t).$$

Integrating this differential inequality yields

$$\text{KL}(\rho_t \| v_t) \leq \exp\left(-c \int_0^t \alpha_s ds\right) \text{KL}(\rho_0 \| v_0)$$

as claimed in (16). \square

APPENDIX F

PROOFS FOR SECTION III (CONTRACTION OF RELATIVE FISHER INFORMATION)

A. Proof of Lemma 2 (Time Derivative of Fisher Information Along FP Channel)

Lemma 2: Suppose $(\rho_t)_{t \geq 0}$ and $(v_t)_{t \geq 0}$ evolve following the Fokker-Planck equation (2):

$$\partial_t \rho_t = -\nabla \cdot (\rho_t b_t) + \frac{c}{2} \Delta \rho_t, \quad (29a)$$

$$\partial_t v_t = -\nabla \cdot (v_t b_t) + \frac{c}{2} \Delta v_t \quad (29b)$$

for some vector fields $(b_t)_{t \geq 0}$ and constant $c \geq 0$ so Assumption 1 is satisfied. Then for any $t \geq 0$:

$$\begin{aligned} \frac{d}{dt} \text{FI}(\rho_t \| v_t) &= -c \mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{v_t} \right\|_{\text{HS}}^2 \right] \\ & \quad - 2 \mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{v_t} \right\|_{(-c \nabla^2 \log v_t + (\nabla b_t)_{\text{sym}})}^2 \right]. \end{aligned}$$

Proof: Let $f_t = -\log \rho_t$ be the negative log-density of ρ_t , and note the relation:

$$\Delta \rho_t = -\rho_t \Delta f_t + \rho_t \|\nabla f_t\|^2.$$

Since ρ_t evolves following the Fokker-Planck equation (29a), we can compute the evolution for f_t :

$$\begin{aligned} \partial_t f_t &= -\frac{1}{\rho_t} \partial_t \rho_t \\ &= -\frac{1}{\rho_t} \left(-\langle \nabla \rho_t, b_t \rangle - \rho_t \nabla \cdot b_t + \frac{c}{2} \Delta \rho_t \right) \\ &= -\langle \nabla f_t, b_t \rangle + \nabla \cdot b_t + \frac{c}{2} \Delta f_t - \frac{c}{2} \|\nabla f_t\|^2. \end{aligned} \quad (30)$$

Similarly, let $g_t = -\log v_t$. Since v_t evolves following (29b), we can compute the evolution for g_t :

$$\partial_t g_t = -\langle \nabla g_t, b_t \rangle + \nabla \cdot b_t + \frac{c}{2} \Delta g_t - \frac{c}{2} \|\nabla g_t\|^2. \quad (31)$$

In the computation below, we will use the following relations:

$$\rho_t \nabla f_t = -\nabla \rho_t \quad (32a)$$

$$\rho_t \nabla^2 f_t = -\nabla^2 \rho_t + \rho_t (\nabla f_t) (\nabla f_t)^\top. \quad (32b)$$

We also use the *Bochner's formula*, which states that for smooth $u: \mathbb{R}^d \rightarrow \mathbb{R}$,

$$\langle \nabla u, \nabla \Delta u \rangle = \frac{1}{2} \Delta \|\nabla u\|^2 - \|\nabla^2 u\|_{\text{HS}}^2. \quad (33)$$

Using relation (32a) and integration by parts, we can write the relative Fisher information as:

$$\begin{aligned} \text{FI}(\rho_t \| v_t) &= \mathbb{E}_{\rho_t} \left[\|\nabla f_t - \nabla g_t\|^2 \right] \\ &= \mathbb{E}_{\rho_t} \left[\|\nabla f_t\|^2 + \|\nabla g_t\|^2 - 2 \langle \nabla f_t, \nabla g_t \rangle \right] \\ &= \mathbb{E}_{\rho_t} \left[\|\nabla f_t\|^2 \right] + \mathbb{E}_{\rho_t} \left[\|\nabla g_t\|^2 \right] - 2 \mathbb{E}_{\rho_t} \left[\Delta g_t \right]. \end{aligned} \quad (34)$$

We compute the time derivative of each term above, spelling out the computations below explicitly, and color coding them for clarity. In the computations below, we assume the density functions ρ_t, v_t are smooth and decay sufficiently fast at infinity. Hence, we can differentiate under the integral sign; furthermore, when we perform integration by parts, all the boundary terms vanish. \square

1) *First Term:* The first term in (34) is the Fisher information of ρ_t . We compute its time derivative using the Fokker-Planck equation (29a) and the formula (30) to get:

$$\begin{aligned} & \frac{d}{dt} \mathbb{E}_{\rho_t} [\|\nabla f_t\|^2] \\ &= \int_{\mathbb{R}^d} (\partial_t \rho_t) \|\nabla f_t\|^2 dx + 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla \partial_t f_t \rangle dx \\ &= - \int_{\mathbb{R}^d} \nabla \cdot (\rho_t b_t) \|\nabla f_t\|^2 dx \end{aligned} \quad (35a)$$

$$+ \frac{c}{2} \int_{\mathbb{R}^d} (\Delta \rho_t) \|\nabla f_t\|^2 dx \quad (35b)$$

$$+ 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla(-\langle \nabla f_t, b_t \rangle + \nabla \cdot b_t) \rangle dx \quad (35c)$$

$$+ c \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla(\Delta f_t) \rangle dx \quad (35d)$$

$$- c \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla \|\nabla f_t\|^2 \rangle dx. \quad (35e)$$

We calculate the terms above one by one.

1) The term in (35a) is, by integration by parts:

$$\begin{aligned} & - \int_{\mathbb{R}^d} \nabla \cdot (\rho_t b_t) \|\nabla f_t\|^2 dx \\ &= \int_{\mathbb{R}^d} \rho_t \langle b_t, \nabla(\|\nabla f_t\|^2) \rangle dx \\ &= 2 \int_{\mathbb{R}^d} \rho_t \langle b_t, (\nabla^2 f_t) \nabla f_t \rangle dx. \end{aligned}$$

2) The term in (35b) is, by integration by parts:

$$\frac{c}{2} \int_{\mathbb{R}^d} (\Delta \rho_t) \|\nabla f_t\|^2 dx = \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \Delta \|\nabla f_t\|^2 dx.$$

3) The term in (35c) is, by distributing the gradient,

$$\begin{aligned} & 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla(-\langle \nabla f_t, b_t \rangle + \nabla \cdot b_t) \rangle dx \\ &= -2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla^2 f_t) b_t \rangle dx \\ &\quad - 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla b_t) \nabla f_t \rangle dx \\ &\quad + 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \text{Tr}(\nabla^2 b_t) \rangle dx. \end{aligned}$$

4) The term in (35d) is, by Bochner's formula (33):

$$\begin{aligned} & c \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla(\Delta f_t) \rangle dx \\ &= \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \Delta \|\nabla f_t\|^2 dx - c \int_{\mathbb{R}^d} \rho_t \|\nabla^2 f_t\|_{\text{HS}}^2 dx. \end{aligned}$$

5) The term in (35e) is, using relation (32b) and integration by parts:

$$\begin{aligned} & -c \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla \|\nabla f_t\|^2 \rangle dx \\ &= c \int_{\mathbb{R}^d} \langle \nabla \rho_t, \nabla \|\nabla f_t\|^2 \rangle dx \\ &= -c \int_{\mathbb{R}^d} \rho_t \Delta \|\nabla f_t\|^2 dx. \end{aligned}$$

Combining the above, we see that the terms involving $\langle b_t, (\nabla^2 f_t) \nabla f_t \rangle$ and $\Delta \|\nabla f_t\|^2$ vanish, and we are left with:

$$\begin{aligned} & \frac{d}{dt} \mathbb{E}_{\rho_t} [\|\nabla f_t\|^2] \\ &= -c \mathbb{E}_{\rho_t} \left[\|\nabla^2 f_t\|_{\text{HS}}^2 \right] - 2 \mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla b_t) \nabla f_t \rangle] \\ &\quad + 2 \mathbb{E}_{\rho_t} [\langle \nabla f_t, \text{Tr}(\nabla^2 b_t) \rangle]. \end{aligned} \quad (36)$$

2) *Second Term:* We compute the time derivative on the second term in (34), using the Fokker-Planck equation (29a) and the formula (31):

$$\begin{aligned} & \frac{d}{dt} \mathbb{E}_{\rho_t} [\|\nabla g_t\|^2] \\ &= \int_{\mathbb{R}^d} (\partial_t \rho_t) \|\nabla g_t\|^2 dx + 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla(\partial_t g_t) \rangle dx \\ &= - \int_{\mathbb{R}^d} \nabla \cdot (\rho_t b_t) \|\nabla g_t\|^2 dx \end{aligned} \quad (37a)$$

$$+ \frac{c}{2} \int_{\mathbb{R}^d} \Delta \rho_t \|\nabla g_t\|^2 dx \quad (37b)$$

$$+ 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla(-\langle \nabla g_t, b_t \rangle) \rangle dx \quad (37c)$$

$$+ 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla(\nabla \cdot b_t) \rangle dx \quad (37d)$$

$$+ c \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla \Delta g_t \rangle dx \quad (37e)$$

$$- c \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla \|\nabla g_t\|^2 \rangle dx. \quad (37f)$$

We calculate the terms above one by one.

1) The term in (37a) is, by integration by parts:

$$\begin{aligned} & - \int_{\mathbb{R}^d} \nabla \cdot (\rho_t b_t) \|\nabla g_t\|^2 dx \\ &= \int_{\mathbb{R}^d} \rho_t \langle b_t, \nabla(\|\nabla g_t\|^2) \rangle dx \\ &= 2 \int_{\mathbb{R}^d} \rho_t \langle b_t, (\nabla^2 g_t) \nabla g_t \rangle dx. \end{aligned}$$

2) The term in (37b) is, by integration by parts and relation (32a):

$$\begin{aligned} & \frac{c}{2} \int_{\mathbb{R}^d} (\Delta \rho_t) \|\nabla g_t\|^2 dx \\ &= \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla \|\nabla g_t\|^2 \rangle dx \\ &= c \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla^2 g_t) \nabla g_t \rangle dx. \end{aligned} \quad (38)$$

3) The term in (37c) is, by chain rule,

$$\begin{aligned} & -2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla(\langle \nabla g_t, b_t \rangle) \rangle dx \\ &= -2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, (\nabla^2 g_t) b_t \rangle dx \\ &\quad - 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, (\nabla b_t) \nabla g_t \rangle dx. \end{aligned}$$

4) The term in (37d) is:

$$\begin{aligned} & 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla(\nabla \cdot b_t) \rangle dx \\ &= 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \text{Tr}(\nabla^2 b_t) \rangle dx \\ &= 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle dx \\ &\quad - 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla b_t, \nabla^2 g_t \rangle_{\text{HS}} dx \end{aligned}$$

where the last equality above follows from the relation (32b) and integration by parts:

$$\begin{aligned} & \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle dx \\ &= - \int_{\mathbb{R}^d} \langle \nabla \rho_t, (\nabla b_t) \nabla g_t \rangle dx \\ &= \int_{\mathbb{R}^d} \rho_t \nabla \cdot ((\nabla b_t) \nabla g_t) dx \\ &= \int_{\mathbb{R}^d} \rho_t (\langle \text{Tr}(\nabla^2 b_t), \nabla g_t \rangle + \langle \nabla b_t, \nabla^2 g_t \rangle_{\text{HS}}) dx. \end{aligned}$$

5) The term in (37e) is, by Bochner's formula (33) and the same calculation as in term (38):

$$\begin{aligned} & c \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla \Delta g_t \rangle dx \\ &= \frac{c}{2} \int_{\mathbb{R}^d} \rho_t \Delta \|\nabla g_t\|^2 dx - c \int_{\mathbb{R}^d} \rho_t \|\nabla^2 g_t\|_{\text{HS}}^2 dx \\ &= c \int_{\mathbb{R}^d} \rho_t (\langle \nabla f_t, (\nabla^2 g_t) \nabla g_t \rangle - \|\nabla^2 g_t\|_{\text{HS}}^2) dx. \end{aligned}$$

6) The term in (37f) is:

$$\begin{aligned} & -c \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, \nabla \|\nabla g_t\|^2 \rangle dx \\ &= -2c \int_{\mathbb{R}^d} \rho_t \langle \nabla g_t, (\nabla^2 g_t) \nabla g_t \rangle dx. \end{aligned}$$

Combining, we see the terms involving $\langle b_t, (\nabla^2 g_t) \nabla g_t \rangle$ vanish, and we obtain:

$$\begin{aligned} & \frac{d}{dt} \mathbb{E}_{\rho_t} [\|\nabla g_t\|^2] \\ &= 2c \mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla^2 g_t) \nabla g_t \rangle] \\ &\quad - 2c \mathbb{E}_{\rho_t} [\langle \nabla g_t, (\nabla^2 g_t) \nabla g_t \rangle] - c \mathbb{E}_{\rho_t} [\|\nabla^2 g_t\|_{\text{HS}}^2] \\ &\quad - 2\mathbb{E}_{\rho_t} [\langle \nabla g_t, (\nabla b_t) \nabla g_t \rangle] + 2\mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle] \\ &\quad - 2\mathbb{E}_{\rho_t} [\langle \nabla b_t, \nabla^2 g_t \rangle_{\text{HS}}]. \end{aligned} \quad (39)$$

3) *Third Term:* We compute the time derivative on the third term in (34), using the Fokker-Planck equation (29a) and the formula (31), and integration by parts, to get:

$$\begin{aligned} & \frac{d}{dt} (-2\mathbb{E}_{\rho_t} [\Delta g_t]) \\ &= -2 \int_{\mathbb{R}^d} (\partial_t \rho_t) \Delta g_t dx - 2 \int_{\mathbb{R}^d} \rho_t \Delta (\partial_t g_t) dx \\ &= 2 \int_{\mathbb{R}^d} \nabla \cdot (\rho_t b_t) (\Delta g_t) dx - c \int_{\mathbb{R}^d} (\Delta \rho_t) (\Delta g_t) dx \end{aligned} \quad (40a)$$

$$+ 2 \int_{\mathbb{R}^d} \rho_t \Delta (\langle \nabla g_t, b_t \rangle - \nabla \cdot b_t) dx \quad (40b)$$

$$-c \int_{\mathbb{R}^d} \rho_t \Delta \Delta g_t dx + c \int_{\mathbb{R}^d} \rho_t \Delta \|\nabla g_t\|^2 dx. \quad (40c)$$

We calculate the terms above one by one.

1) The first term in (40a) is, by integration by parts:

$$2 \int_{\mathbb{R}^d} \nabla \cdot (\rho_t b_t) (\Delta g_t) dx = -2 \int_{\mathbb{R}^d} \rho_t \langle b_t, \nabla \Delta g_t \rangle dx.$$

2) For the second term in (40a), we can write it using integration by parts as follows:

$$\begin{aligned} & -c \int_{\mathbb{R}^d} (\Delta \rho_t) (\Delta g_t) dx \\ &= -c \int_{\mathbb{R}^d} \left(\sum_{i=1}^d \frac{\partial^2}{\partial x_i^2} \rho_t(x) \right) \left(\sum_{j=1}^d \frac{\partial^2}{\partial x_j^2} g_t(x) \right) dx \\ &= -c \int_{\mathbb{R}^d} \sum_{i,j=1}^d \left(\frac{\partial^2}{\partial x_i^2} \rho_t(x) \right) \left(\frac{\partial^2}{\partial x_j^2} g_t(x) \right) dx \\ &= -c \int_{\mathbb{R}^d} \sum_{i,j=1}^d \rho_t(x) \left(\frac{\partial^4}{\partial x_i^2 \partial x_j^2} g_t(x) \right) dx \\ &= -c \int_{\mathbb{R}^d} \sum_{i,j=1}^d \left(\frac{\partial^2}{\partial x_i \partial x_j} \rho_t(x) \right) \left(\frac{\partial^2}{\partial x_i \partial x_j} g_t(x) \right) dx \\ &= -c \int_{\mathbb{R}^d} \langle \nabla^2 \rho_t, \nabla^2 g_t \rangle_{\text{HS}} dx \\ &= c \int_{\mathbb{R}^d} \rho_t (\langle \nabla^2 f_t, \nabla^2 g_t \rangle_{\text{HS}} - \langle \nabla f_t, (\nabla^2 g_t) \nabla f_t \rangle) dx \end{aligned} \quad (41)$$

where in the last step we have used relation (32b).

3) The term in (40b) is, by the relation (32a), integration by parts, and distributing the gradient:

$$\begin{aligned} & 2 \int_{\mathbb{R}^d} \rho_t \Delta (\langle \nabla g_t, b_t \rangle - \nabla \cdot b_t) dx \\ &= 2 \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, \nabla (\langle \nabla g_t, b_t \rangle - \nabla \cdot b_t) \rangle dx \\ &= 2 \int_{\mathbb{R}^d} \rho_t (\langle \nabla f_t, (\nabla^2 g_t) b_t \rangle + \langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle \\ &\quad - \langle \nabla f_t, \text{Tr}(\nabla^2 b_t) \rangle) dx \\ &= \int_{\mathbb{R}^d} \rho_t (2\langle \nabla \Delta g_t, b_t \rangle + 2\langle \nabla^2 g_t, \nabla b_t \rangle_{\text{HS}} \\ &\quad + 2\langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle - 2\langle \nabla f_t, \text{Tr}(\nabla^2 b_t) \rangle) dx \end{aligned}$$

where in the last equality we have used the integration by parts to write:

$$\begin{aligned} & \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla^2 g_t) b_t \rangle dx \\ &= - \int_{\mathbb{R}^d} \langle \nabla \rho_t, (\nabla^2 g_t) b_t \rangle dx \\ &= \int_{\mathbb{R}^d} \rho_t \nabla \cdot ((\nabla^2 g_t) b_t) dx \\ &= \int_{\mathbb{R}^d} \rho_t (\langle \nabla \Delta g_t, b_t \rangle + \langle \nabla^2 g_t, \nabla b_t \rangle_{\text{HS}}) dx \end{aligned}$$

and we have also used the identity $\nabla \cdot (\nabla^2 g_t) = \nabla \Delta g_t$, which can be verified from the definition.

4) The first term in (40c) is, by integration by parts, equal to the second term in (40a), which is also equal to (41):

$$\begin{aligned} & -c \int_{\mathbb{R}^d} \rho_t \Delta \Delta g_t dx \\ & = -c \int_{\mathbb{R}^d} (\Delta \rho_t) (\Delta g_t) dx \\ & = c \int_{\mathbb{R}^d} \rho_t (\langle \nabla^2 f_t, \nabla^2 g_t \rangle_{\text{HS}} - \langle \nabla f_t, (\nabla^2 g_t) \nabla f_t \rangle) dx. \end{aligned}$$

5) The second term in (40c) is, by integration by parts, equal to the term in (37b), which is equal to twice the term in (38):

$$\begin{aligned} c \int_{\mathbb{R}^d} \rho_t \Delta \|\nabla g_t\|^2 dx & = c \int_{\mathbb{R}^d} \Delta \rho_t \|\nabla g_t\|^2 dx \\ & = 2c \int_{\mathbb{R}^d} \rho_t \langle \nabla f_t, (\nabla^2 g_t) \nabla g_t \rangle dx. \end{aligned}$$

Combining, we see the terms involving $\langle b_t, \nabla \Delta g_t \rangle$ vanish, so we get:

$$\begin{aligned} & \frac{d}{dt} (-2\mathbb{E}_{\rho_t} [\Delta g_t]) \\ & = 2c \mathbb{E}_{\rho_t} [\langle \nabla^2 f_t, \nabla^2 g_t \rangle_{\text{HS}}] - 2c \mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla^2 g_t) \nabla f_t \rangle] \\ & \quad + 2\mathbb{E}_{\rho_t} [\langle \nabla^2 g_t, \nabla b_t \rangle_{\text{HS}}] + 2\mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle] \\ & \quad - 2\mathbb{E}_{\rho_t} [\langle \nabla f_t, \text{Tr}(\nabla^2 b_t) \rangle] + 2c \mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla^2 g_t) \nabla g_t \rangle]. \end{aligned} \quad (42)$$

4) *Combining the Terms:* Combining the calculations in (36), (39), and (42), we see the terms involving $\mathbb{E}_{\rho_t} [\langle \nabla f_t, \text{Tr}(\nabla^2 b_t) \rangle]$ and $\mathbb{E}_{\rho_t} [\langle \nabla b_t, \nabla^2 g_t \rangle_{\text{HS}}]$ vanish. After rearranging the terms, we find the time derivative of the relative Fisher information (34) is equal to:

$$\begin{aligned} & \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) \\ & = \frac{d}{dt} \mathbb{E}_{\rho_t} [\|\nabla f_t\|^2] + \frac{d}{dt} \mathbb{E}_{\rho_t} [\|\nabla g_t\|^2] + \frac{d}{dt} (-2\mathbb{E}_{\rho_t} [\Delta g_t]) \\ & = -c \mathbb{E}_{\rho_t} [\|\nabla^2 f_t\|_{\text{HS}}^2] - c \mathbb{E}_{\rho_t} [\|\nabla^2 g_t\|_{\text{HS}}^2] \\ & \quad + 2c \mathbb{E}_{\rho_t} [\langle \nabla^2 f_t, \nabla^2 g_t \rangle_{\text{HS}}] \\ & \quad - 2c \mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla^2 g_t) \nabla f_t \rangle] - 2c \mathbb{E}_{\rho_t} [\langle \nabla g_t, (\nabla^2 g_t) \nabla g_t \rangle] \\ & \quad + 2c \mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla^2 g_t) \nabla g_t \rangle] + 2c \mathbb{E}_{\rho_t} [\langle \nabla g_t, (\nabla^2 g_t) \nabla g_t \rangle] \\ & \quad - 2\mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla b_t) \nabla f_t \rangle] - 2\mathbb{E}_{\rho_t} [\langle \nabla g_t, (\nabla b_t) \nabla g_t \rangle] \\ & \quad + 2\mathbb{E}_{\rho_t} [\langle \nabla f_t, (\nabla b_t) \nabla g_t \rangle] + 2\mathbb{E}_{\rho_t} [\langle \nabla g_t, (\nabla b_t) \nabla g_t \rangle] \\ & = -c \mathbb{E}_{\rho_t} [\|\nabla^2 f_t - \nabla^2 g_t\|_{\text{HS}}^2] - 2c \mathbb{E}_{\rho_t} [\|\nabla f_t - \nabla g_t\|_{\nabla^2 g_t}^2] \\ & \quad - 2\mathbb{E}_{\rho_t} [\|\nabla f_t - \nabla g_t\|_{(\nabla b_t)_{\text{sym}}}^2] \\ & = -c \mathbb{E}_{\rho_t} [\|\nabla^2 f_t - \nabla^2 g_t\|_{\text{HS}}^2] \\ & \quad - 2\mathbb{E}_{\rho_t} [\|\nabla f_t - \nabla g_t\|_{(c \nabla^2 g_t + (\nabla b_t)_{\text{sym}})}^2] \end{aligned}$$

as desired. In the third equality above, we use the property that for all $u, v \in \mathbb{R}^d$ and $A \in \mathbb{R}^{d \times d}$: \square

$$\langle u, Av \rangle + \langle v, Au \rangle = \langle u, (A + A^\top)v \rangle = 2\langle u, A_{\text{sym}}v \rangle.$$

B. *Proof of Theorem 1 (SDPI Along Gaussian Channel)*

Proof of Theorem 1: For the Gaussian channel (3), which is the case $b_t = 0$ and $c = 1$, the time derivative formula from Lemma 2 simplifies to:

$$\begin{aligned} \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) & = -\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{\nu_t} \right\|_{\text{HS}}^2 \right] \\ & \quad - 2\mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|_{(-\nabla^2 \log \nu_t)}^2 \right]. \end{aligned} \quad (43)$$

The first term in the right-hand side above is the second-order Fisher information, which is non-negative: $\mathbb{E}_{\rho_t} [\|\nabla^2 \log \frac{\rho_t}{\nu_t}\|_{\text{HS}}^2] \geq 0$. In parts (i), (ii), and (iv) below we drop this first term, and only consider the second term which is a weighted relative Fisher information. In part (iii), we use the Poincaré inequality to exploit this second-order Fisher information to improve the convergence rate.

Recall along the Gaussian channel (3), we have $\rho_t = \rho_0 * \mathcal{N}(0, tI)$ and $\nu_t = \nu_0 * \mathcal{N}(0, tI)$.

1) *Part (i):* Note the Gaussian distribution $\mathcal{N}(0, tI)$ is log-concave, and recall convolution with log-concave distributions preserves log-concavity [56, Proposition 3.5]. Since ν_0 is log-concave, $\nu_t = \nu_0 * \mathcal{N}(0, tI)$ is also log-concave, so $-\nabla^2 \log \nu_t \geq 0$. Then from the formula (43) by dropping the first term:

$$\frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) \leq -2\mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|_{(-\nabla^2 \log \nu_t)}^2 \right] \leq 0.$$

This shows that relative Fisher information is non-increasing over time: $\text{Fl}(\rho_t \| \nu_t) \leq \text{Fl}(\rho_0 \| \nu_0)$.

2) *Part (ii):* Now assume ν_0 is α -SLC. Since $\mathcal{N}(0, tI)$ is $(1/t)$ -SLC, we recall $\nu_t = \nu_0 * \mathcal{N}(0, tI)$ is $\alpha_t = \frac{\alpha}{1+\alpha t}$ -SLC [56, Theorem 3.7]. Then from the formula (43) by dropping the first term:

$$\begin{aligned} \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) & \leq -2\alpha_t \mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|^2 \right] \\ & = -2\alpha_t \text{Fl}(\rho_t \| \nu_t). \end{aligned} \quad (44)$$

Integrating and plugging in the definition of $\alpha_t = \frac{\alpha}{1+\alpha t} = \frac{d}{dt} \log(1 + \alpha t)$ give the desired result:

$$\begin{aligned} \text{Fl}(\rho_t \| \nu_t) & \leq \exp \left(-2 \int_0^t \alpha_s ds \right) \text{Fl}(\rho_0 \| \nu_0) \\ & = \frac{\text{Fl}(\rho_0 \| \nu_0)}{(1 + \alpha t)^2}. \end{aligned}$$

3) *Part (iii):* As in the proof of part (ii) above, $\nu_t = \nu_0 * \mathcal{N}(0, tI)$ is $\alpha_t = \frac{\alpha}{1+\alpha t}$ -SLC. Furthermore, since $\mathcal{N}(0, tI)$ is $(1/t)$ -SLC and thus satisfies $(1/t)$ -Poincaré inequality, and ρ_0 satisfies β -Poincaré inequality by assumption, we know $\rho_t = \rho_0 * \mathcal{N}(0, tI)$ satisfies $\beta_t = \frac{\beta}{1+\beta t}$ -Poincaré inequality [1, Corollary 3.1].

Since ρ_0 and ν_0 are symmetric, $\rho_t = \rho_0 * \mathcal{N}(0, tI)$ and $\nu_t = \nu_0 * \mathcal{N}(0, tI)$ are also symmetric. In particular, $\log \nu_t(-x) = \log \nu_t(x)$, so $\nabla \log \nu_t$ is an odd function: $\nabla \log \nu_t(-x) = -\nabla \log \nu_t(x)$. Furthermore, since ρ_t is symmetric

($\rho_t(-x) = \rho_t(x)$), this implies $\mathbb{E}_{\rho_t}[\nabla \log v_t] = 0$, since indeed:

$$\begin{aligned} & \mathbb{E}_{\rho_t}[\nabla \log v_t] \\ &= \int_{\mathbb{R}^d} \rho_t(x) \nabla \log v_t(x) dx \\ &= \frac{1}{2} \int_{\mathbb{R}^d} \rho_t(x) \nabla \log v_t(x) dx + \frac{1}{2} \int_{\mathbb{R}^d} \rho_t(-x) \nabla \log v_t(-x) dx \\ &= \frac{1}{2} \int_{\mathbb{R}^d} \rho_t(x) (\nabla \log v_t(x) + \nabla \log v_t(-x)) dx \\ &= 0. \end{aligned}$$

Since $\rho_t = \rho_0 * \mathcal{N}(0, tI)$ has full support on \mathbb{R}^d , by integration by parts, $\mathbb{E}_{\rho_t}[\nabla \log \rho_t] = 0$. Combining, this shows that $\mathbb{E}_{\rho_t}[\nabla \log \frac{\rho_t}{v_t}] = 0$. Therefore, by the Poincaré inequality for ρ_t applied to the function $\phi = \nabla \log \frac{\rho_t}{v_t}$,

$$\begin{aligned} \mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{v_t} \right\|_{\text{HS}}^2 \right] &\geq \beta_t \text{Var}_{\rho_t} \left(\nabla \log \frac{\rho_t}{v_t} \right) \\ &= \beta_t \mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{v_t} \right\|^2 \right] \\ &= \beta_t \text{FI}(\rho_t \| v_t). \end{aligned}$$

Therefore, by the formula (43):

$$\begin{aligned} & \frac{d}{dt} \text{FI}(\rho_t \| v_t) \\ &= -\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{v_t} \right\|_{\text{HS}}^2 \right] - 2\mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{v_t} \right\|_{(-\nabla^2 \log v_t)}^2 \right] \\ &\leq -(\beta_t + 2\alpha_t) \text{FI}(\rho_t \| v_t). \end{aligned}$$

Integrating and plugging in the definitions of $\alpha_t = \frac{\alpha}{1+\alpha t}$ and $\beta_t = \frac{\beta}{1+\beta t}$ give the desired result:

$$\begin{aligned} \text{FI}(\rho_t \| v_t) &\leq \exp \left(-\int_0^t (\beta_s + 2\alpha_s) ds \right) \text{FI}(\rho_0 \| v_0) \\ &= \frac{\text{FI}(\rho_0 \| v_0)}{(1 + \beta t)(1 + \alpha t)^2}. \end{aligned}$$

4) Part (iv): By the result of [45, Theorem 1.3], under the assumption on v_0 , we have that $v_t = v_0 * \mathcal{N}(0, tI)$ satisfies $-\nabla^2 \log v_t(x) \geq \alpha_t I$ for all $x \in \mathbb{R}^d$, where

$$\alpha_t := \frac{1}{t} \left(1 - \frac{1}{t} \left(\frac{L}{\alpha + 1/t} + \frac{1}{\sqrt{\alpha + 1/t}} \right)^2 \right).$$

Note α_t may be negative for small $t > 0$, but we have $\alpha_t > 0$ for instance for all $t \geq \left(\frac{L}{\alpha} + \frac{1}{\sqrt{\alpha}} \right)^2$. By direct computation, we may observe that:

$$\alpha_t = \frac{d}{dt} \left(\log(\alpha t + 1) - \frac{L^2}{\alpha + 1/t} - \frac{4L}{\sqrt{\alpha + 1/t}} \right).$$

By the same computation (44) as in part (ii) above, which holds for $\alpha_t \in \mathbb{R}$, we have for all $t \geq 0$:

$$\begin{aligned} \text{FI}(\rho_t \| v_t) &\leq \exp \left(-2 \int_0^t \alpha_s ds \right) \text{FI}(\rho_0 \| v_0) \\ &= \frac{\text{FI}(\rho_0 \| v_0)}{(1 + \alpha t)^2} \exp \left(\frac{2tL^2}{\alpha t + 1} + \frac{8L\sqrt{t}}{\sqrt{\alpha t + 1}} \right). \end{aligned}$$

□

C. Detail for Counterexample of DPI Along the Gaussian Channel

The following is a detailed computation for Example 3. See Figure 1 for an illustration.

Proposition 1: Let $\rho_0 = \mathcal{N}(0, 1)$ on \mathbb{R} . Let $M \geq 2$ and $L \geq 2$ be arbitrary, and define $v_0 \propto e^{-g}$ where $g: \mathbb{R} \rightarrow \mathbb{R}$ is the function:

$$g(x) = \begin{cases} -\frac{M}{2}x^2 & \text{if } |x| \leq L, \\ \frac{1}{2}(x-L)^2 - ML(x-L) - \frac{ML^2}{2} & \text{if } x > L, \\ \frac{1}{2}(x+L)^2 + ML(x+L) - \frac{ML^2}{2} & \text{if } x < -L. \end{cases}$$

For $t \geq 0$, define $\rho_t = \rho_0 * \mathcal{N}(0, t)$ and $v_t = v_0 * \mathcal{N}(0, t)$, which are the output of the Gaussian channel (3) at time t . Then for small $t > 0$, the relative Fisher information $t \mapsto \text{FI}(\rho_t \| v_t)$ is increasing, and in particular,

$$\text{FI}(\rho_t \| v_t) > \text{FI}(\rho_0 \| v_0).$$

Proof: We show that the time derivative of $t \mapsto \text{FI}(\rho_t \| v_t)$ is positive at $t = 0$; then by continuity, the time derivative will still be positive for sufficiently small $t > 0$, which means $t \mapsto \text{FI}(\rho_t \| v_t)$ is increasing. Since we are in 1-dimension, we denote derivative by $'$ and second derivative by $''$.

Recall by the time derivative formula (43) along the Gaussian channel, we have:

$$\begin{aligned} & \frac{d}{dt} \text{FI}(\rho_t \| v_t) \Big|_{t=0} \\ &= -\mathbb{E}_{\rho_0} \left[\left(\left(\log \frac{\rho_0}{v_0} \right)'' \right)^2 \right] \\ &\quad - 2\mathbb{E}_{\rho_0} \left[(-\log v_0)'' \left(\left(\log \frac{\rho_0}{v_0} \right)' \right)^2 \right] \\ &= -\mathbb{E}_{\rho_0} \left[(-1 + g''(X))^2 \right] \\ &\quad - 2\mathbb{E}_{\rho_0} \left[g''(X) (-X + g'(X))^2 \right]. \end{aligned} \quad (45)$$

By construction, we have:

$$g'(x) = \begin{cases} -Mx & \text{if } |x| \leq L, \\ x - (M+1)L & \text{if } x > L, \\ x + (M+1)L & \text{if } x < -L \end{cases} \quad (46)$$

and

$$g''(x) = \begin{cases} -M & \text{if } |x| \leq L, \\ 1 & \text{if } |x| > L. \end{cases} \quad (47)$$

Then we can bound the first term in (45) by:

$$\begin{aligned} -\mathbb{E}_{\rho_0} \left[(-1 + g''(X))^2 \right] &= -(M+1)^2 \rho_0([-L, L]) \\ &\geq -(M+1)^2. \end{aligned}$$

Furthermore, with $\rho_0 = \mathcal{N}(0, 1)$, we can compute the second term in (45):

$$\begin{aligned} & -2\mathbb{E}_{\rho_0} \left[g''(X) (-X + g'(X))^2 \right] \\ &= 2M(M+1)^2 \mathbb{E}_{\rho_0} [X^2 \mathbf{1}_{\{|X| \leq L\}}] \\ &\quad - 4(M+1)^2 L^2 \rho_0((L, \infty)). \end{aligned}$$

We choose $L > 0$ such that $\mathbb{E}_{\rho_0}[X^2 \mathbf{1}\{|X| \leq L\}] > \frac{1}{2}\mathbb{E}_{\rho_0}[X^2] = \frac{1}{2}$ and $L^2\rho_0((L, \infty)) < \frac{1}{4}$; note that any $L \geq 2$ works. Then combining the bounds above in (45), we get:

$$\begin{aligned} & \left. \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) \right|_{t=0} \\ &= -\mathbb{E}_{\rho_0} \left[(-1 + g''(X))^2 \right] - 2\mathbb{E}_{\rho_0} \left[g''(X)(-X + g'(X))^2 \right] \\ &> -(M+1)^2 + M(M+1)^2 - (M+1)^2 \\ &= (M-2)(M+1)^2 \\ &\geq 0 \end{aligned}$$

where the last inequality holds since $M \geq 2$. \square

1) *Remark:* In Proposition 1, the potential g is C^1 and piecewise quadratic; in particular g'' is the piecewise-constant function in (47) with jump discontinuities at $\pm L$. To apply Lemma 2 (and hence (43)) under Assumption 1, we can replace g by a smooth approximation g_δ that only modifies g in a δ -neighborhood of $\pm L$. We sketch the argument here. Fix $\delta > 0$, and let $\psi \in C_c^\infty(\mathbb{R})$ be a smooth mollifier with $\int_{\mathbb{R}} \psi(x) dx = 1$, supported on $[-1, 1]$. Write $h(x) := g''(x)$ (as defined in (47)), and define $h_\delta := h * \psi_\delta$ where $\psi_\delta(x) = \delta^{-1}\psi(x/\delta)$. Then $h_\delta \in C^\infty(\mathbb{R})$, $h_\delta(x) = h(x)$ whenever $|x| - L > \delta$, and h_δ is bounded between $-M$ and 1. Define g_δ by:

$$\begin{aligned} g_\delta(0) &= g(0), \\ g'_\delta(0) &= 0, \\ g''_\delta(x) &= h_\delta(x) \quad \text{for all } x, \end{aligned}$$

so $g_\delta \in C^\infty(\mathbb{R})$ and is even. Moreover, g_δ has the same quadratic growth as g for $|x| \geq L + \delta$ (since $h_\delta(x) = 1$ there), hence the density $\nu_0^\delta \propto e^{-g_\delta}$ has a Gaussian tail and satisfies Assumption 1; therefore, the time-derivative identity (43) applies to (ρ_t, ν_t^δ) .

Finally, the derivative at $t = 0$ in Proposition 10 is given by the Gaussian expectation in (45), which depends on g only through g' and g'' evaluated at $X \sim \mathcal{N}(0, 1)$. For each fixed $x \neq \pm L$ we have $g'_\delta(x) \rightarrow g'(x)$ and $g''_\delta(x) \rightarrow g''(x)$ as $\delta \downarrow 0$, and the integrand in (45) is dominated by a polynomial in $|x|$ (since g'_δ grows at most linearly and g''_δ is bounded). Thus, by the dominated convergence theorem, the derivative functional is continuous under this smoothing: $D(g_\delta) \rightarrow D(g)$ as $\delta \downarrow 0$, where $D(g)$ denotes the derivative $\left. \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) \right|_{t=0}$ when ν_0 is defined by g , and similarly for $D(g_\delta)$. Since Proposition 1 shows $D(g) > 0$ for the chosen parameters, the same strict positivity holds for all sufficiently small δ . Therefore, the counterexample holds with a C^∞ potential, while satisfying the regularity assumptions needed for Lemma 2.

D. Proof of Theorem 2 (Eventual SDPI Along OU Channel)

Proof of Theorem 2: For the OU channel (9), we have $b_t(x) = \gamma x$ and $c = 2$, so the time derivative formula from Lemma 2 becomes:

$$\begin{aligned} \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) &= -2\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{\nu_t} \right\|_{\text{HS}}^2 \right] \\ &\quad - 2\mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|_{(-2\nabla^2 \log \nu_t - \gamma I)}^2 \right]. \end{aligned} \quad (48)$$

The first term in the right-hand side above is the second-order Fisher information, which is non-negative: $\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{\nu_t} \right\|_{\text{HS}}^2 \right] \geq 0$. In parts (i) we drop this first term, and only consider the second term which is a weighted relative Fisher information. In part (ii), we apply the Poincaré inequality to exploit this second-order Fisher information to improve the convergence rate.

Part (i): We recall the explicit solution (9) of OU at time $t \geq 0$ is $X_t \stackrel{d}{=} e^{-\gamma t} X_0 + \sqrt{\frac{1-e^{-2\gamma t}}{\gamma}} Z$ where $X_0 \sim \nu_0$, $Z \sim \mathcal{N}(0, I)$ is independent, and $X_t \sim \nu_t$. Since ν_0 is α -SLC by assumption, the law of $e^{-\gamma t} X_0$ is $(e^{2\gamma t} \alpha)$ -SLC. Note the law of $\sqrt{\frac{1-e^{-2\gamma t}}{\gamma}} Z$ is $(\frac{\gamma}{1-e^{-2\gamma t}})$ -SLC. Then from the preservation of strong log-concavity under convolution [56, Theorem 3.7], we know ν_t is α_t -SLC, where

$$\alpha_t = \left(\frac{e^{-2\gamma t}}{\alpha} + \frac{1-e^{-2\gamma t}}{\gamma} \right)^{-1} = \frac{\gamma \alpha}{\alpha + e^{-2\gamma t}(\gamma - \alpha)}.$$

Note that

$$\begin{aligned} \lambda_t &:= 2\alpha_t - \gamma = \gamma - \frac{2\gamma(\gamma - \alpha)e^{-2\gamma t}}{\alpha + e^{-2\gamma t}(\gamma - \alpha)} \\ &= \frac{d}{dt} \left(\gamma t + \log(\alpha + e^{-2\gamma t}(\gamma - \alpha)) \right). \end{aligned} \quad (49)$$

Then in this case, $-2\nabla^2 \nu_t - \gamma I \geq (2\alpha_t - \gamma)I = \lambda_t I$, and thus from the formula (48), we obtain

$$\begin{aligned} \frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) &\leq -2\mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|_{(-2\nabla^2 \log \nu_t - \gamma I)}^2 \right] \\ &\leq -2\lambda_t \text{Fl}(\rho_t \| \nu_t). \end{aligned}$$

Integrating the differential inequality above yields the desired bound:

$$\begin{aligned} \text{Fl}(\rho_t \| \nu_t) &\leq \exp \left(-2 \int_0^t \lambda_s ds \right) \text{Fl}(\rho_0 \| \nu_0) \\ &= \frac{\gamma^2 e^{-2\gamma t}}{(\alpha + e^{-2\gamma t}(\gamma - \alpha))^2} \text{Fl}(\rho_0 \| \nu_0). \end{aligned}$$

Part (ii): Since ρ_0 and ν_0 are symmetric, along the OU channel (9) to $\mathcal{N}(0, \gamma^{-1}I)$, the solutions ρ_t and ν_t are also symmetric, as can be seen from the explicit solution (10). Then as in the proof of Theorem 1(iii), we have $\mathbb{E}_{\rho_t} \left[\nabla \log \frac{\rho_t}{\nu_t} \right] = 0$. Furthermore, similar to the calculation in part (i) and using the composition rule for Poincaré inequality under convolution [1], since ρ_0 satisfies β -Poincaré inequality, along the OU channel (9), ρ_t satisfies β_t -Poincaré inequality where:

$$\begin{aligned} \beta_t &= \frac{\gamma \beta}{\beta + e^{-2\gamma t}(\gamma - \beta)} \\ &= \frac{d}{dt} \left(\gamma t + \frac{1}{2} \log(\beta + e^{-2\gamma t}(\gamma - \beta)) \right). \end{aligned}$$

Then by the Poincaré inequality for ρ_t , for the test function $\nabla \log \frac{\rho_t}{\nu_t}$, we have:

$$\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{\nu_t} \right\|_{\text{HS}}^2 \right] \geq \beta_t \text{Var}_{\rho_t} \left(\nabla \log \frac{\rho_t}{\nu_t} \right)$$

$$\begin{aligned}
&= \beta_t \mathbb{E}_{\rho_t} \left[\left\| \nabla \log \frac{\rho_t}{\nu_t} \right\|^2 \right] \\
&= \beta_t \text{Fl}(\rho_t \| \nu_t).
\end{aligned}$$

We also recall our estimate from part (i) that $-2\nabla^2 \nu_t - \gamma I \geq \lambda_t I$ where λ_t is given in (49). Then by the identity (48) along the OU channel (9), we obtain:

$$\frac{d}{dt} \text{Fl}(\rho_t \| \nu_t) \leq -2(\beta_t + \lambda_t) \text{Fl}(\rho_t \| \nu_t).$$

Integrating the differential inequality above gives the desired result:

$$\begin{aligned}
\text{Fl}(\rho_t \| \nu_t) &\leq \exp\left(-2 \int_0^t (\beta_s + \lambda_s) ds\right) \text{Fl}(\rho_0 \| \nu_0) \\
&= \frac{\gamma^3 e^{-4\gamma t} \text{Fl}(\rho_0 \| \nu_0)}{(\beta + e^{-2\gamma t}(\gamma - \beta))(\alpha + e^{-2\gamma t}(\gamma - \alpha))^2}.
\end{aligned}$$

□

APPENDIX G

PROOFS FOR SECTION IV (ANALYSIS OF THE PROXIMAL SAMPLER)

A. Review of the Rejection Sampling Implementation of the RGO

We review the algorithmic implementation of the restricted Gaussian oracle (RGO), which is required in each step of the Proximal Sampler (see Section IV-A). Here we follow the simple rejection sampling implementation of the RGO from [23, Section 4.2] for the setting when the original target distribution $\nu^X \propto e^{-g}$ is log-smooth. We note that one can also implement the RGO using inexact rejection sampling to obtain an improved dimension dependence [57], and other works have shown how to implement the RGO for broader classes of distributions, see e.g. [58], [59], [60], and [61].

Recall that the RGO (21) needs to sample from the conditional distribution ν^{XY} with density, for each fixed $y \in \mathbb{R}^d$:

$$\nu^{XY}(x | y) \propto_x \exp\left(-g(x) - \frac{1}{2\eta} \|x - y\|^2\right).$$

Define $f_y: \mathbb{R}^d \rightarrow \mathbb{R}$ by $f_y(x) := g(x) + \frac{1}{2\eta} \|x - y\|^2$, so that $\nu^{XY}(x | y) \propto \exp(-f_y(x))$. Suppose that the potential function g for the original target distribution $\nu^X \propto e^{-g}$ is L -smooth, which means $-LI \leq \nabla^2 g(x) \leq LI$ for all $x \in \mathbb{R}^d$. If $0 < \eta < \frac{1}{L}$, then we observe that f_y is an $\left(\frac{1}{\eta} - L\right)$ -strongly convex and $\left(\frac{1}{\eta} + L\right)$ -smooth function, with condition number $\kappa := \frac{1+\eta L}{1-\eta L}$.

We can perform rejection sampling on $\nu^{XY}(\cdot | y) \propto \exp(-f_y)$ via the following procedure:

- 1) Compute the minimizer $x_y^* = \arg \min_{x \in \mathbb{R}^d} f_y(x)$; we can do this, e.g., via gradient descent. Since f_y is $(\alpha + \eta^{-1})$ -strongly convex and $(L + \eta^{-1})$ -smooth, gradient descent with step size $1/(L + \eta^{-1})$ reaches an δ_{opt} -accurate minimizer in $O\left(\kappa \log \frac{1}{\delta_{\text{opt}}}\right)$ iterations. Moreover, replacing x_y^* by an approximate minimizer only changes the rejection envelope by a multiplicative factor $\exp(O(\delta_{\text{opt}}))$, so the expected number of trials remains $\tilde{O}(1)$ when δ_{opt} is chosen as an inverse polynomial in the target accuracy.

For ease of analysis, here we assume we can compute the minimizer x_y^* explicitly. See also [57] and [62] for complexity analysis of the Proximal Sampler under approximate RGO.

- 2) Repeat until acceptance: draw a sample $Z \sim \mathcal{N}\left(x_y^*, \frac{\eta}{1-\eta L} I\right)$ and accept it with probability $\exp\left(-f_y(Z) + f_y(x_y^*) + \frac{\eta}{2(1-\eta L)} \|Z - x_y^*\|^2\right)$.

By standard properties of rejection sampling, the output of the procedure above is an exact sample $Z \sim \nu^{XY}(\cdot | y)$, and the expected number of iterations (draws of Z) until acceptance is at most $\kappa^{d/2} = \left(\frac{1+\eta L}{1-\eta L}\right)^{d/2}$; see [63, Theorem 7]. Note that if we choose $\eta = \frac{1}{dL}$, then the expected number of iterations until acceptance is at most $\kappa^{d/2} = \left(\frac{1+1/d}{1-1/d}\right)^{d/2} = O(1)$. This shows that we can implement the RGO via rejection sampling with $O(1)$ expected iteration complexity.

B. Proof of Theorem 3 (Convergence of the Proximal Sampler in Relative Fisher Information)

Proof of Theorem 3: As we reviewed in Section IV-B, we use stochastic process interpretations of the forward and backward steps of the Proximal Sampler.

- (1) **Forward step:** The first step of the Proximal Sampler algorithm is a convolution with the Gaussian distribution: $\rho_k^Y = \rho_k^X * \mathcal{N}(0, \eta I)$, and $\nu^Y = \nu^X * \mathcal{N}(0, \eta I)$. These are the outputs of the Gaussian channel (3) for time η . Since ν^X is α -SLC by assumption, by the SDPI along the Gaussian channel from Theorem 1(ii), we have:

$$\text{Fl}(\rho_k^Y \| \nu^Y) \leq \frac{\text{Fl}(\rho_k^X \| \nu^X)}{(1 + \alpha\eta)^2}.$$

- (2) **Backward step:** As explained in [23, Appendix A.1.4], the second step of the Proximal Sampler can be seen as the output of the reverse Gaussian channel. Concretely, define the distributions $\mu_0 = \nu^X$ and $\mu_t = \mu_0 * \mathcal{N}(0, tI)$ for $0 \leq t \leq \eta$, so $\mu_\eta = \nu^Y$, and μ_t evolves following the heat equation: $\partial_t \mu_t = \frac{1}{2} \Delta \mu_t$. Let $\nu_t = \mu_{\eta-t}$, so $\nu_0 = \nu^X$, $\nu_\eta = \nu^Y$, and ν_t evolves following the backward heat equation (5):

$$\partial_t \nu_t = -\frac{1}{2} \Delta \nu_t = -\nabla \cdot (\nu_t \nabla \log \nu_t) + \frac{1}{2} \Delta \nu_t.$$

Now define the reverse Gaussian channel (6):

$$dX_t = \nabla \log \nu_t(X_t) dt + dW_t$$

so that if X_t evolves along this channel, then its distribution $X_t \sim \rho_t$ evolves following the Fokker-Planck equation:

$$\partial_t \rho_t = -\nabla \cdot (\rho_t \nabla \log \nu_t) + \frac{1}{2} \Delta \rho_t.$$

This channel has the property that if we initialize it at a point mass $\rho_0 = \delta_y$ for any $y \in \mathbb{R}^d$, the output at time $t = \eta$ is exactly the conditional distribution $\rho_\eta = \nu^{XY}(\cdot | y)$ that the RGO is supposed to sample from; see the explanation in [23, Appendix A.1.2], see also the exposition in [64, Chapter 8.3]. In particular, if we initialize the reverse Gaussian channel from $\rho_0 = \rho_k^Y$,

then by construction, the output at time $t = \eta$ is the distribution of the next iterate:

$$\rho_\eta = \int_{\mathbb{R}^d} \rho_0(y) v^{X|Y}(\cdot | y) dy = \rho_{k+1}^X.$$

Thus, we have interpolated from $v_0 = v^Y$ to $v_\eta = v^X$, and from $\rho_0 = \rho_k^Y$ to $\rho_\eta = \rho_{k+1}^X$, as the evolution along the backward heat equation (5) or the reverse Gaussian channel (6), which is an instance of the Fokker-Planck channel (2) with drift function $b_t = \nabla \log v_t$ and $c = 1$. In this case $-c\nabla^2 \log v_t + (\nabla b_t)_{\text{sym}} = 0$, so the second term in the time derivative formula (17) vanishes. Therefore, along the backward heat equation, using the formula (17) from Lemma 2:

$$\frac{d}{dt} \text{Fl}(\rho_t \| v_t) = -\mathbb{E}_{\rho_t} \left[\left\| \nabla^2 \log \frac{\rho_t}{v_t} \right\|_{\text{HS}}^2 \right] \leq 0.$$

Integrating from $t = 0$ to $t = \eta$ yields:

$$\text{Fl}(\rho_{k+1}^X \| v^X) = \text{Fl}(\rho_\eta \| v_\eta) \leq \text{Fl}(\rho_0 \| v_0) = \text{Fl}(\rho_k^Y \| v^Y).$$

Iterating the recursions above yields the desired convergence rate in (25). \square

C. Proof of Corollary 1 (Iteration Complexity of Proximal Sampler in Relative Fisher Information)

Proof of Corollary 1: Since we assume $v^X \propto e^{-g}$ is α -SLC and is L -log-smooth, we have $\alpha I \leq \nabla^2 g(x) \leq LI$ for all $x \in \mathbb{R}^d$. Since x^* is the minimizer of g , we have $\nabla g(x^*) = 0$. Then from $\rho_0^X = \mathcal{N}(x^*, \frac{1}{L}I)$, we can bound the initial relative Fisher information:

$$\begin{aligned} \text{Fl}(\rho_0^X \| v) &= \mathbb{E}_{\rho_0^X} \left[\|(Lx - \nabla g(x)) - (Lx^* - \nabla g(x^*))\|^2 \right] \\ &\leq (L - \alpha)^2 \mathbb{E}_{\rho_0^X} \left[\|x - x^*\|^2 \right] \\ &= \frac{d(L - \alpha)^2}{L} \\ &\leq dL, \end{aligned}$$

where the first inequality above follows from the fact that $x \mapsto Lx - \nabla g(x)$ is $(L - \alpha)$ -Lipschitz, because its Jacobian is $0 \leq LI - \nabla^2 g(x) \leq (L - \alpha)I$.

With step size $\eta = \frac{1}{dL}$, we can implement the RGO via rejection sampling with an $O(1)$ expected number of queries, as explained in Appendix G-A. Then from the exponential convergence rate (25) of the Proximal Sampler from Theorem 3, we have

$$\text{Fl}(\rho_k^X \| v) \leq \frac{\text{Fl}(\rho_0^X \| v^X)}{(1 + \alpha\eta)^{2k}} \leq \frac{dL}{(1 + \frac{\alpha}{dL})^{2k}} \leq dL \exp\left(-\frac{\alpha k}{dL}\right)$$

where in the last inequality we use the bound $1 + c \geq e^{c/2}$ which holds for $0 \leq c = \frac{\alpha}{dL} \leq 1$.

Then to reach $\text{Fl}(\rho_k^X \| v) \leq \varepsilon$, it suffices to set $dL \exp(-\frac{\alpha k}{dL}) \leq \varepsilon$, or equivalently, we run the Proximal Sampler for $k \geq \frac{dL}{\alpha} \log \frac{dL}{\varepsilon}$ iterations, as claimed. \square

ACKNOWLEDGMENT

The author would like to thank Yihong Wu, Santosh Vempala, Sinho Chewi, and Varun Jog for valuable discussions; and Tinsel Wibisono for guidance.

REFERENCES

- [1] D. Chafaï, "Entropies, convexity, and functional inequalities, on Φ -entropies and Φ -sobolev inequalities," *J. Math. Kyoto Univ.*, vol. 44, no. 2, pp. 325–363, Jan. 2004.
- [2] G. O. Roberts and R. L. Tweedie, "Exponential convergence of Langevin distributions and their discrete approximations," *Bernoulli*, vol. 2, no. 4, pp. 341–363, Dec. 1996.
- [3] Y. T. Lee, R. Shen, and K. Tian, "Structured logconcave sampling with a restricted Gaussian Oracle," in *Proc. 34th Conf. Learn. Theory*, 2020, pp. 2993–3050.
- [4] R. Jordan, D. Kinderlehrer, and F. Otto, "The variational formulation of the Fokker-Planck equation," *SIAM J. Math. Anal.*, vol. 29, no. 1, pp. 1–17, Jan. 1998.
- [5] A. Wibisono, "Sampling as optimization in the space of measures: The Langevin dynamics as a composite optimization problem," in *Proc. Conf. Learn. Theory*, 2018, pp. 2093–3027.
- [6] D. Bakry and M. Émery, "Diffusions hypercontractives," in *Séminaire De Probabilités XIX 1983/84*. Cham, Switzerland: Springer, 1985, pp. 177–206.
- [7] Y. Cao, J. Lu, and Y. Lu, "Exponential decay of Rényi divergence under Fokker-Planck equations," *J. Stat. Phys.*, vol. 176, no. 5, pp. 1172–1184, Sep. 2019.
- [8] A. S. Dalalyan, "Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent," in *Proc. Conf. Learn. Theory*, Jul. 2017, pp. 678–689.
- [9] A. S. Dalalyan, "Theoretical guarantees for approximate sampling from smooth and log-concave densities," *J. Roy. Stat. Soc. Ser. B, Stat. Methodol.*, pp. 651–676, 2014.
- [10] A. Durmus and É. Moulines, "Nonasymptotic convergence analysis for the unadjusted Langevin algorithm," *Ann. Appl. Probab.*, vol. 27, no. 3, p. 1551, 2017.
- [11] X. Cheng and P. L. Bartlett, "Convergence of Langevin MCMC in KL-divergence," in *Proc. Algorithmic Learn. Theory*, 2018, pp. 186–211.
- [12] A. Durmus, S. Majewski, and B. Miasojedow, "Analysis of Langevin Monte Carlo via convex optimization," *J. Mach. Learn. Res.*, vol. 20, no. 73, pp. 1–46, 2018.
- [13] S. Chewi, M. A. Erdogdu, M. Li, R. Shen, and M. S. Zhang, "Analysis of Langevin Monte Carlo from Poincare to log-Sobolev," *Found. Comput. Math.*, vol. 25, pp. 1345–1395, Jul. 2024, doi: [10.1007/s10208-024-09667-6](https://doi.org/10.1007/s10208-024-09667-6).
- [14] S. S. Vempala and A. Wibisono, "Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices," in *Geometric Aspects of Functional Analysis: Israel Seminar (GAFA) 2020–2022*, R. Eldan, B. Klartag, A. Litvak, and E. Milman, Eds., Cham, Switzerland: Springer, 2023, pp. 381–438.
- [15] J. M. Altschuler and K. Talwar, "Resolving the mixing time of the Langevin algorithm to its stationary distribution for log-concave sampling," in *Proc. 36th Conf. Learn. Theory (PMLR)*, vol. 195, 2023, pp. 2509–2510.
- [16] S. Mitra and A. Wibisono, "Fast convergence of Φ -divergence along the unadjusted Langevin algorithm and proximal sampler," in *Proc. 36th Int. Conf. Algorithmic Learn. Theory (PMLR)*, vol. 272, 2024, pp. 846–869.
- [17] R. Dwivedi, Y. Chen, M. J. Wainwright, and B. Yu, "Log-concave sampling: Metropolis-hastings algorithms are fast," *J. Mach. Learn. Res.*, pp. 1–42, 2018.
- [18] Y. Chen, R. Dwivedi, M. J. Wainwright, and B. Yu, "Fast mixing of metropolized Hamiltonian Monte Carlo: Benefits of multi-step gradients," *J. Mach. Learn. Res.*, vol. 21, no. 92, pp. 1–72, 2019.
- [19] S. Chewi, C. Lü, K. Ahn, X. Cheng, T. L. Gouic, and P. Rigollet, "Optimal dimension dependence of the metropolis-adjusted Langevin algorithm," in *Proc. Conf. Learn. Theory*, 2020, pp. 1260–1300.
- [20] K. Wu, S. C. Schmidler, and Y. Chen, "Minimax mixing time of the metropolis-adjusted Langevin algorithm for log-concave sampling," *J. Mach. Learn. Res.*, pp. 1–63, 2021.
- [21] Y. Chen and K. Gatmiry, "A simple proof of the mixing of metropolis-adjusted Langevin algorithm under smoothness and isoperimetry," 2023, *arXiv:2304.04095*.
- [22] M. K. Titsias and O. Papaspiliopoulos, "Auxiliary gradient-based sampling algorithms," *J. Roy. Stat. Soc. Ser. B, Stat. Methodol.*, vol. 80, no. 4, pp. 749–767, 2018.
- [23] Y. Chen, S. Chewi, A. Salim, and A. Wibisono, "Improved analysis for a proximal algorithm for sampling," in *Proc. 35th Conf. Learn. Theory*, Jul. 2022, pp. 2984–3014.
- [24] S. Chen, S. Chewi, J. Li, Y. Li, A. Salim, and A. R. Zhang, "Sampling is as easy as learning the score: Theory for diffusion models with minimal data assumptions," in *Proc. Int. Conf. Learn. Represent.*, 2022.

- [25] A. Wibisono, Y. Wu, and K. Y. Yang, "Optimal score estimation via empirical Bayes smoothing," in *Proc. 37th Annu. Conf. Learn. Theory*, 2024, pp. 4958–4991.
- [26] F. Otto and C. Villani, "Generalization of an inequality by Talagrand and links with the logarithmic sobolev inequality," *J. Funct. Anal.*, vol. 173, no. 2, pp. 361–400, Jun. 2000.
- [27] K. Balasubramanian, S. Chewi, M. A. Erdogdu, A. Salim, and S. Zhang, "Towards a theory of non-log-concave sampling: First-order stationarity guarantees for Langevin Monte Carlo," in *Proc. 35th Conf. Learn. Theory*, vol. 178, 2022, pp. 2896–2923.
- [28] S. Chewi, P. Gerber, H. Lee, and L. Chen, "Fisher information lower bounds for sampling," in *Proc. Int. Conf. Algorithmic Learn. Theory*, 2022, pp. 375–410.
- [29] T. Cover and J. Thomas, *Elements of Information Theory*, 2nd ed., Hoboken, NJ, USA: Wiley, 2006.
- [30] M. Raginsky, "Strong data processing inequalities and Φ -Sobolev inequalities for discrete channels," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3355–3389, Jun. 2016.
- [31] Y. Polyanskiy and Y. Wu, "Strong data-processing inequalities for channels and Bayesian networks," in *Convexity and Concentration*. Cham, Switzerland: Springer, 2017, pp. 211–249.
- [32] P. Caputo, Z. Chen, Y. Gu, and Y. Polyanskiy, "Entropy contractions in Markov chains: half-step, full-step and continuous-time," *Electron. J. Probab.*, vol. 30, no. none, pp. 1–29, Jan. 2025.
- [33] Y. Polyanskiy and Y. Wu, *Information Theory: From Coding to Learning*. Cambridge, U.K.: Cambridge Univ. Press, 2025.
- [34] B. Yuan, J. Fan, J. Liang, A. Wibisono, and Y. Chen, "On a class of Gibbs sampling over networks," in *Proc. 36th Conf. Learn. Theory (PMLR)*, vol. 195, 2023, pp. 5754–5780.
- [35] Y. Kook, S. Vempala, and M. Zhang, "In-and-out: Algorithmic diffusion for sampling convex bodies," in *Proc. Adv. Neural Inf. Process. Syst.*, 2024, pp. 108354–108388.
- [36] J. Liang, S. Mitra, and A. Wibisono, "Characterizing dependence of samples along the Langevin dynamics and algorithms via contraction of Φ -mutual information," *IEEE Trans. Inf. Theory*, vol. 72, no. 2, pp. 921–945, Feb. 2026.
- [37] R. Zamir, "A proof of the Fisher information inequality via a data processing argument," *IEEE Trans. Inf. Theory*, vol. 44, no. 3, pp. 1246–1250, May 1998.
- [38] Y. Wu, "A simple proof of (a slightly improved) Gaussian HWI inequality," *Preprint*, 2011.
- [39] C. Villani, *Optimal Transport: Old and New*, vol. 338. Cham, Switzerland: Springer, 2009.
- [40] A. Wibisono, V. Jog, and P.-L. Loh, "Information and estimation in Fokker-Planck channels," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2017, pp. 2673–2677.
- [41] A. Wibisono and V. Jog, "Convexity of mutual information along the heat flow," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2018, pp. 1615–1619.
- [42] A. Wibisono and V. Jog, "Convexity of mutual information along the Ornstein-Uhlenbeck flow," in *Proc. Int. Symp. Inf. Theory Its Appl. (ISITA)*, Oct. 2018, pp. 55–59.
- [43] J. Zou, L. Fan, J. Gao, and J. Wang, "Convexity of mutual information along the Fokker-Planck flow," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2025, pp. 1–6.
- [44] B. Klartag and O. Ordentlich, "The strong data processing inequality under the heat flow," *IEEE Trans. Inf. Theory*, vol. 71, no. 5, pp. 3317–3333, May 2025.
- [45] G. Brigati and F. Pedrotti, "Heat flow, log-concavity, and Lipschitz transport maps," *Electron. Commun. Probab.*, vol. 30, no. none, pp. 1–12, Jan. 2025.
- [46] M. Fathi, E. Indrei, and M. Ledoux, "Quantitative logarithmic sobolev inequalities and stability estimates," *Discrete Continuous Dyn. Syst.*, vol. 36, no. 12, pp. 6835–6853, 2016.
- [47] S. Chewi, A.-A. Pooladian, and M. S. Zhang, "Stability of the kim-milman flow map," 2025, *arXiv:2511.01154*.
- [48] Y.-A. Ma, N. S. Chatterji, X. Cheng, N. Flammarion, P. L. Bartlett, and M. I. Jordan, "Is there an analog of Nesterov acceleration for gradient-based MCMC?" *Bernoulli*, vol. 27, no. 3, pp. 1942–1992, May 2021.
- [49] P. Monmarché, "An entropic approach for Hamiltonian Monte Carlo: The idealized case," *Ann. Appl. Probab.*, pp. 2243–2293, 2022.
- [50] X. Cheng, B. Wang, J. Zhang, and Y. Zhu, "Fast conditional mixing of MCMC algorithms for non-log-concave distributions," in *Proc. Adv. Neural Inf. Process. Syst.* 36, 2023, pp. 13374–13394.
- [51] C. Villani, *Hypocoercivity*. Providence, RI, USA: American Mathematical Soc., 2009.
- [52] L. Conger, F. Hoffmann, E. Mazumdar, and L. J. Ratliff, "Coupled Wasserstein gradient flows for min-max and cooperative games," 2024, *arXiv:2411.07403*.
- [53] Y. Cai, S. Mitra, X. Wang, and A. Wibisono, "On the convergence of min-max Langevin dynamics and algorithm," in *Proc. Thirty 8th Conf. Learn. Theory*, Jul. 2024, pp. 677–754.
- [54] Q. Feng, X. Zuo, and W. Li, "Fisher information dissipation for time-inhomogeneous stochastic differential equations," *J. Mach. Learn. Res.*, vol. 25, no. 392, pp. 1–45, 2024.
- [55] N. Blachman, "The convolution inequality for entropy powers," *IEEE Trans. Inf. Theory*, vol. IT-11, no. 2, pp. 267–271, Apr. 1965.
- [56] A. Saumard and J. A. Wellner, "Log-concavity and strong log-concavity: A review," *Statist. Surveys*, vol. 8, no. none, p. 45, Jan. 2014.
- [57] J. Fan, B. Yuan, and Y. Chen, "Improved dimension dependence of a proximal algorithm for sampling," in *Proc. Annu. Conf. Learn. Theory*, 2023, pp. 1473–1521.
- [58] S. Gopi, Y. T. Lee, and D. Liu, "Private convex optimization via exponential mechanism," in *Proc. Conf. Learn. Theory*, vol. 14, 2024, pp. 1948–1989.
- [59] J. Liang and Y. Chen, "A proximal algorithm for sampling from non-smooth potentials," *Winter Simul. Conf. (WSC)*, pp. 3229–3240, Dec. 2021, doi: [10.1109/WSC57314.2022.10015293](https://doi.org/10.1109/WSC57314.2022.10015293).
- [60] J. Liang and Y. Chen, "A proximal algorithm for sampling," *Trans. Mach. Learn. Res.*, Jul. 2022. [Online]. Available: <https://jmlr.org/tmlr/papers/>
- [61] J. Liang and Y. Chen, "Proximal Oracles for optimization and sampling," 2024, *arXiv:2404.02239*.
- [62] J. M. Altschuler and S. Chewi, "Faster high-accuracy log-concave sampling via algorithmic warm starts," *J. ACM*, vol. 71, no. 3, pp. 1–55, Jun. 2024.
- [63] S. Chewi, P. Gerber, C. Lü, T. L. Gouic, and P. Rigollet, "The query complexity of sampling from strongly log-concave distributions in one dimension," in *Proc. Conf. Learn. Theory*, 2021, pp. 2041–2059.
- [64] S. Chewi, "Log-concave sampling," Yale Univ., New Haven, CT, USA, Tech. Rep., 2025. [Online]. Available: <https://chewisinho.github.io/main.pdf>

Andre Wibisono received the dual B.S. degrees in mathematics and in computer science and the master's degree in computer science from MIT, the master's degree in statistics from UC Berkeley, and the Ph.D. degree in computer science from UC Berkeley. He is an Assistant Professor with the Department of Computer Science, Yale University, with a secondary appointment with the Department of Statistics and Data Science. Before joining Yale University, he has done post-doctoral research with UW Madison and Georgia Tech. His research interests are in the design and analysis of algorithms for machine learning, in particular for problems in optimization, sampling, and game theory.