

CONVERGENCE OF UNADJUSTED LANGEVIN IN HIGH DIMENSIONS: DELOCALIZATION OF BIAS

YIFAN CHEN, XIAOOU CHENG, JONATHAN NILES-WEED, AND JONATHAN WEARE

ABSTRACT. The unadjusted Langevin algorithm is commonly used to sample probability distributions in extremely high-dimensional settings. However, existing analyses of the algorithm for strongly log-concave distributions suggest that, as the dimension d of the problem increases, the number of iterations required to ensure convergence within a desired error in the W_2 metric scales in proportion to d or $\sqrt{d}$. In this paper, we argue that, despite this poor scaling of the W_2 error for the full set of variables, the behavior for a *small number* of variables can be significantly better: a number of iterations proportional to K , up to logarithmic terms in d , often suffices for the algorithm to converge to within a desired W_2 error for all K -marginals. We refer to this effect as *delocalization of bias*. We show that the delocalization effect does not hold universally and prove its validity for Gaussian distributions and strongly log-concave distributions with certain sparse interactions. Our analysis relies on a novel W_{2,ℓ^∞} metric to measure convergence. A key technical challenge we address is the lack of a one-step contraction property in this metric. Finally, we use asymptotic arguments to explore potential generalizations of the delocalization effect beyond the Gaussian and sparse interactions setting.

1. INTRODUCTION

Overdamped Langevin dynamics

$$(1.1) \quad dX_t = -\nabla V(X_t)dt + \sqrt{2}dB_t,$$

has been used extensively to sample from high dimensional probability distributions in applications ranging from molecular dynamics to Bayesian inverse problems and data assimilation [33]. Here $V : \mathbb{R}^d \rightarrow \mathbb{R}$ is a function in $\mathbb{R}^d$ referred to as the potential and B_t is the d -dimensional Brownian motion. The target distribution π is proportional to $\exp(-V)$, which is the stationary distribution of (1.1).

The overdamped Langevin Monte Carlo algorithm, also known as the unadjusted Langevin algorithm, is obtained by applying the Euler–Maruyama scheme to (1.1):

$$(1.2) \quad X_{(k+1)h} = X_{kh} - h\nabla V(X_{kh}) + \sqrt{2h}(B_{(k+1)h} - B_{kh}),$$

where h is the step size. The distribution of X_{kh} is denoted by ρ_{kh} .

1.1. Motivation. The unadjusted Langevin algorithm (1.2) is biased, meaning that even when $k \rightarrow \infty$, the distribution ρ_{kh} will not converge to the exact target distribution π for any finite step size h . For strongly log-concave distributions, ρ_{kh} converges in the W_2 metric to the stationary distribution of (1.2), π_h , at a rate that is independent of dimension (see, e.g., [17]). Consequently, the dependence of the error $W_2(\rho_{kh}, \pi)$ on dimension, d , is completely determined by the bias $W_2(\pi_h, \pi)$. This bias depends on the

COURANT INSTITUTE, NEW YORK UNIVERSITY, NY, USA

E-mail addresses: yifan.chen@nyu.edu, xc2108@nyu.edu, jdw453@nyu.edu, jqw200@nyu.edu.

step size h , which, according to state-of-the-art analyses, needs to scale proportionally to $1/d$ or $1/\sqrt{d}$ to achieve a bounded bias for any d . Since the iteration complexity typically scales as $1/h$ (up to logarithmic factors), this small step size implies a computational cost of order $\sqrt{d}$ or d .

Unbiased samplers based on (1.1), such as the Metropolis-adjusted Langevin algorithm [37, 36] and proximal samplers [26, 5], are available, but these algorithms also require a small step size when d is large in order to ensure that the acceptance rate is non-negligible; see the review in Section 1.2. In short, existing theoretical analyses of both the unadjusted Langevin algorithm and unbiased variants require step sizes scaling as d^{-c} for some $c > 0$. In the former case, small step sizes are required to ensure small bias; in the latter, small step sizes are required to maintain a reasonable acceptance rate.

In the case of the unadjusted Langevin algorithm, however, these theoretical predictions seem at odds with abundant empirical evidence that the scheme samples efficiently in extremely high dimensions. As an example, molecular dynamics simulations using integrators closely related to (1.2) with up to billions of atoms are not uncommon [22]. In these simulations, the step size is typically set to several femtoseconds, irrespective of the system size [28].

In this paper, we argue that the mismatch between theory and practice is due, in large part, to the metric used to measure convergence. Standard metrics such as W_2 measure accuracy of the entire distribution, but practitioners are often interested in averages of functions involving a relatively small set of variables. Based on this observation, we aim to investigate the convergence behavior of the algorithm under alternative metrics designed to characterize the accuracy of low-dimensional marginals. We identify a new “delocalization” phenomenon for the bias of the unadjusted Langevin algorithm to show that, even in high-dimensional settings, the step size may not need to be very small if the quantities of interest depend on low-dimensional marginals only. Consequently, the iteration complexity required to achieve a bounded error exhibits a benign dependence on dimension d .

1.2. Literature review. We begin by reviewing existing works on the analysis of Langevin algorithms, with a particular focus on the dependence on dimension d . Analysis of (1.2) dates back to the work of Roberts and Tweedie [36], where asymptotic properties such as ergodicity of the discrete Markov chain are studied. Asymptotic bias of the discrete SDEs can be investigated using Taylor’s expansion and the Poisson equation [39, 40, 1, 31]. Non-asymptotic analysis of the algorithm (1.2) aims to characterize the step size and iteration complexity to achieve bounded error under certain metrics; again, the iteration complexity typically scales inversely with the step size (up to logarithmic factors of d). Here, we focus on the scaling of the step size with dimension. A large body of work considers strongly log-concave distributions:

Assumption 1.1. *Let $\pi \propto \exp(-V)$ and $V \in C^2(\mathbb{R}^d)$. Assume V is α -strongly convex and β -smooth such that $\alpha I \preceq \nabla^2 V(x) \preceq \beta I$ at any $x \in \mathbb{R}^d$. Here $0 < \alpha \leq \beta < \infty$.*

Under Assumption 1.1, existing non-asymptotic bounds for the unadjusted Langevin algorithm (1.2) mainly focus on metrics such as the total variation distance, the W_2 distance, and the Kullback-Leibler (KL) divergence; see [13, 12, 16, 17, 8, 15]. In these

studies, to maintain a bounded error or bias as the dimension d increases, the step size h must scale inversely with d , i.e., $h \sim 1/d$. With additional smoothness assumptions on the Hessian of V , the bound on the step size can be improved to $h \sim 1/\sqrt{d}$; see [29] and the mean squared analysis framework developed in [30]. Other convergence results exist, such as those using the χ^2 and Rényi divergences as metrics [41, 20], under weaker assumptions like the log-Sobolev inequality [41, 9] or other conditions on π [4, 19, 32, 27], and alternative discretization schemes [38, 23]. All these results lead to power law scaling of the step size with d (up to logarithmic terms) to attain a desired error for any d .

Metropolis-adjusted Langevin algorithms (MALA) [36] and proximal samplers [26] are two common unbiased schemes for the Langevin dynamics (1.1). For MALA, the asymptotic analysis in [35] shows that, when $h \sim 1/d^{1/3}$, the algorithm admits a non-degenerate diffusion limit as $d \rightarrow \infty$, provided π is a smooth product measure and the algorithm starts at stationarity; see also further generalizations beyond product measures and stationarity [3, 11, 34, 24] and note that out of stationarity, the results suggest the scaling $h \sim 1/\sqrt{d}$. Non-asymptotic bounds on the mixing time of MALA have also been investigated in the literature [18, 6, 10, 43, 7]. These bounds require a specific choice of step size, scaling as $h \sim 1/\sqrt{d}$ for MALA initialized at a warm start and $h \sim 1/d$ for a feasible start, up to $\log d$ terms; see also related studies on the lower bounds of the mixing time [10, 25, 43]. For proximal samplers [26], a key component is the implementation of the restricted Gaussian oracle (RGO), and the step size h needs to be small to ensure efficient implementation. Existing analyses suggest choosing $h \sim 1/d$ if the RGO is implemented via rejection sampling [5], and $h \sim 1/\sqrt{d}$ if it is implemented via approximate rejection sampling [21].

In summary, for both MALA and proximal samplers, choosing the step size inversely proportional to d or $\sqrt{d}$ is necessary for efficient implementation and to achieve the desired mixing time bounds. For the unadjusted Langevin algorithm, similar scaling is required to attain a bounded error on the bias under the aforementioned metrics.

Under this context, the goal of this paper is to show that the requirement on the step size in the unadjusted Langevin algorithm can be significantly improved if we use an alternative W_{2,ℓ^∞} metric to measure convergence. The convergence bounds we prove in this metric indicate that, in some situations, the bias of the unadjusted Langevin algorithm is *delocalized*, in the sense that the bias in individual coordinates is nearly dimension-free. This is precisely the behavior observed when π is a product measure. Our results show that this phenomenon holds in wider generality.

References [2] and [14] studied target distributions for which larger stepsizes are possible in high-dimensions if the observable of interest has Lipschitz constant that scales as $1/\sqrt{d}$, such as the averaged observable $f(x) = \frac{1}{d} \sum_{i=1}^d \Phi(x^{(i)})$ with ℓ^2 -Lipschitz Φ . In these cases, a constant step size is sufficient to achieve a bounded error regardless of dimensionality. Our results go beyond these studies by establishing bounds that are not limited to averaged observables. Our analysis implies that the algorithm could achieve a bounded error for low-dimensional marginals in any dimension using a step size that is dimension-independent up to logarithmic terms.

1.3. Main results. In this section, we outline our main results, including a novel metric W_{2,ℓ^∞} to measure convergence, two positive and one negative example that illustrate

the delocalization effects of bias, and our main theorem regarding the convergence result for strongly log-concave distributions with sparse interactions.

1.3.1. *A new metric to measure convergence.* We introduce the following W_{p,ℓ^∞} metric:

$$(1.3) \quad W_{p,\ell^\infty}(\mu, \nu) = \left(\min_{\gamma \in \Pi(\mu, \nu)} \int |x - y|_\infty^p \gamma(dx, dy) \right)^{1/p},$$

where $|\cdot|_\infty$ is the ℓ^∞ norm of a vector and $\Pi(\mu, \nu)$ represents the set of measures in the joint space $\mathbb{R}^d \times \mathbb{R}^d$ that have marginals μ, ν .

We note that $|x - y|_\infty \geq |x^{(j)} - y^{(j)}|$ for any $1 \leq j \leq d$, where we use the superscript (j) to denote the j -th component of a vector. This observation implies that $W_{p,\ell^\infty}(\mu, \nu)$ serves as an upper bound for the W_p distance between one-dimensional marginals of μ and ν . Moreover, since $K|x - y|_\infty^p \geq \sum_{t=1}^K |x^{(j_t)} - y^{(j_t)}|^p$ for any $1 \leq j_t \leq d$, we have that $K^{1/p} \cdot W_{p,\ell^\infty}(\mu, \nu)$ serves as the upper bound for the W_p distance between any K -dimensional marginals of μ and ν . In summary, the metric W_{p,ℓ^∞} is capable of measuring the accuracy of low-dimensional marginals. In this paper, we specifically focus on the case $p = 2$, i.e., the W_{2,ℓ^∞} metric.

Another important observation is that

$$W_{2,\ell^\infty}(\rho_{kh}, \pi) \leq W_{2,\ell^\infty}(\rho_{kh}, \pi_h) + W_{2,\ell^\infty}(\pi_h, \pi) \leq W_2(\rho_{kh}, \pi_h) + W_{2,\ell^\infty}(\pi_h, \pi),$$

which, combined with the dimension-independent contraction result in the W_2 metric, implies that the W_{2,ℓ^∞} bias governs the dependence of the W_{2,ℓ^∞} error on d .

1.3.2. *Positive and negative examples: delocalization effect.* How should we expect the W_{2,ℓ^∞} bias of (1.2) behave? To motivate the discussion, we first consider the examples of product measures and Gaussian measures. We begin with the product measure case for which a bound on $W_{2,\ell^\infty}(\pi_h, \pi)$ can be obtained by a contraction argument similar to the W_2 analysis (e.g., [12]). We include a sketch of the argument for later reference in the sketch of the proof of our main result, Theorem 2.2 in Section 2.2.

Example 1.2. Consider $\pi \propto \exp(-V)$ where $V(x) = \sum_{i=1}^d V_i(x^{(i)})$ satisfies $\alpha \leq \nabla^2 V_i \leq \beta$. Then, for $h \leq \alpha/\beta^2$, it holds that

$$W_{2,\ell^\infty}(\pi_h, \pi) = O\left(\frac{\beta}{\alpha} \sqrt{h \log(2d)}\right).$$

Sketch of Proof. Given $k \in \mathbb{N}$, we couple the continuous-time Langevin dynamics at stationarity Y_t , $kh \leq t \leq (k+1)h$ and the discrete-time iterates X_{kh} in (1.2) so that they share the same Brownian motion. Introducing an auxiliary random variable $\bar{Y}_{(k+1)h} = Y_{kh} - h\nabla V(Y_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh})$ and using the triangle inequality then lead to

$$\sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq \underbrace{\sqrt{\mathbb{E}[|X_{(k+1)h} - \bar{Y}_{(k+1)h}|_\infty^2]}}_{(a)} + \underbrace{\sqrt{\mathbb{E}[|\bar{Y}_{(k+1)h} - Y_{(k+1)h}|_\infty^2]}}_{(b)}.$$

Here (b) is the one-step discretization error which is bounded by $O(\beta h^{3/2} \sqrt{\log(2d)})$ given $h \leq \alpha/\beta^2$. And (a) $= \sqrt{\mathbb{E}[|X_{kh} - Y_{kh} - h(\nabla V(X_{kh}) - \nabla V(Y_{kh}))|_\infty^2]} \leq (1 - \alpha h) \sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]} \leq \exp(-\alpha h) \sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]}$ where we apply the strong convexity of V_i to bound each coordinate of the vector yielding the final ℓ^∞ norm bound.

This shows that there is a one-step contraction in the ℓ^∞ norm. We then couple the distribution of X_{kh} and Y_{kh} such that $\sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]} = W_{2,\ell^\infty}(\rho_{kh}, \pi)$. Using the definition of the W_{2,ℓ^∞} norm, we get

$$W_{2,\ell^\infty}(\rho_{(k+1)h}, \pi) \leq \exp(-\alpha h) W_{2,\ell^\infty}(\rho_{kh}, \pi) + O(\beta h^{3/2} \sqrt{\log(2d)}).$$

Iterating this inequality implies $W_{2,\ell^\infty}(\pi_h, \pi) = O(\frac{\beta}{\alpha} \sqrt{h \log(2d)})$. A detailed proof can be found in Appendix A. $\square$

Example 1.2 shows that the W_{2,ℓ^∞} bias scales only as the square root of $\log(2d)$; that is, the bias is nearly independent of the dimension. The key in the proof for Example 1.2 is the one-step contraction property in the ℓ^∞ norm, which relies on the structure of the product measures. This property does not hold for general π . On the other hand, we know that for Gaussian distributions, we have an explicit formula for the law of iterates in the algorithm and thus the biased distribution π_h , so we can investigate the W_{2,ℓ^∞} bias directly without concerning the one-step iteration property; see the following Example 1.3 with proof in Appendix B.

Example 1.3. Consider $\pi \propto \exp(-V)$ and $V(x) = \frac{1}{2}(x-m)^T \Sigma^{-1}(x-m)$ where $m \in \mathbb{R}^d$ and $\alpha I \preceq \Sigma^{-1} \preceq \beta I$. Then, for $h \leq 1/\beta$, it holds that

$$W_{2,\ell^\infty}(\pi_h, \pi) = O(h \sqrt{\beta \log(2d)}) = O(\sqrt{h \log(2d)}).$$

Again, Example 1.3 shows that the W_{2,ℓ^∞} bias is nearly independent of the dimension. Using the property of the W_{2,ℓ^∞} metric, this further implies that the W_2 distance between K -marginals of π_h and π is bounded by $O(\sqrt{Kh \log(2d)})$, nearly independent of d . We can interpret this as the overall bias being nearly *delocalized* across all one-dimensional marginals. A step size of $O(1/K)$, up to logarithmic terms, suffices for a bounded bias and error in K -marginals for any dimension; the iteration complexity scales with K and is also nearly independent of d .

On the other hand, a simple example shows that delocalization of bias does not always hold:

Example 1.4. Consider $\pi = \rho^{\otimes d}$ where ρ is a one-dimensional centered log concave distribution for which the biased distribution ρ_h obtained by the corresponding one dimensional unadjusted Langevin has nonzero mean¹, so that their mean differs by $\delta > 0$. Consider the observable $f(x) = \frac{1}{\sqrt{d}} \sum_{i=1}^d x^{(i)}$ for $x \in \mathbb{R}^d$. It holds that

$$\left| \int f(\pi - \pi_h) \right| = \sqrt{d} \delta.$$

Now, consider the rotation matrix Q which satisfies $(Qx)^{(1)} = \frac{1}{\sqrt{d}} \sum_{i=1}^d x^{(i)}$. Let $\tilde{\pi} = Q\#\pi$. We have $\tilde{\pi}_h = Q\#\pi_h$. Consequently, it holds that

$$\left| \int x^{(1)}(\tilde{\pi} - \tilde{\pi}_h) \right| = \left| \int f(\pi - \pi_h) \right| = \sqrt{d} \delta.$$

¹For example, Proposition 3.1 can be used to show that ρ can be taken to be a univariate mixture of Gaussians $p\mathcal{N}(\mu_1, 1) + (1-p)\mathcal{N}(\mu_2, 1)$ with $p\mu_1 + (1-p)\mu_2 = 0$, $0 < p < 1/2$, and $|\mu_1 - \mu_2|$ sufficiently small.

We then find $W_{2,\ell^\infty}(\tilde{\pi}, \tilde{\pi}_h) \geq W_{1,\ell^\infty}(\tilde{\pi}, \tilde{\pi}_h) \geq \left| \int x^{(1)}(\tilde{\pi} - \tilde{\pi}_h) \right| = \sqrt{d}\delta$.

Example 1.4 shows that for the rotated product measure $\tilde{\pi}$, the W_{2,ℓ^∞} bias is of order $\sqrt{d}$. This indicates that the bias is not *delocalized*, but *concentrated* on one specific dimension. Thus, the *delocalization effect* of the bias over marginals is a delicate phenomenon and does not hold universally. We note that Example 1.4 is characterized by *strong, dense* interactions between coordinates after the rotation.

1.3.3. Strongly log-concave target distributions with sparse interactions. In this paper, our main result is to show that the delocalization effect holds for distributions with sparse interactions; see the following Figure 1 and Theorem 1.5.




FIGURE 1. Illustration of a potential $V(x), x \in \mathbb{R}^d$ with sparse interactions.

Theorem 1.5 (informal). *Let Assumption 1.1 hold. Suppose V further satisfies the sparsity conditions illustrated in Figure 1 and rigorously formulated in Section 2.1 below with the sparsity parameter $s_k \leq C(k+1)^n$, then*

$$W_{2,\ell^\infty}(\pi_h, \pi) \leq \sqrt{h \log(2d)} \min \left\{ \left(O\left(\frac{\beta}{\alpha} \log(2d)\right) \right)^{\frac{n}{2}+1}, O\left(\frac{\beta}{\alpha} \sqrt{d}\right) \right\}.$$

The precise statement of the theorem is in Section 2. The key technical challenge in extending the delocalization of bias result beyond product measures and Gaussian measures is the lack of one-step contraction in the W_{2,ℓ^∞} metric and the lack of explicit formulas of the solutions. We employ a novel multi-step coupling argument to derive a multi-step contraction result and crucially use the sparsity of the potential to control the accumulated errors over these steps to prove the theorem.

The theorem implies that for such a sparse potential, the W_{2,ℓ^∞} bias is nearly independent of d , up to logarithmic terms and assuming the condition number β/α is also nearly independent of d . As a consequence, by taking $h \sim 1/K$, the iteration complexity scales with K to ensure a desired W_2 error for all K -marginals, up to logarithmic terms on d . Thus, for these potentials, the unadjusted Langevin algorithm proves more scalable than its unbiased variants when the quantities of interest are low-dimensional

marginals. We note that these sparse potential illustrated in Figure 1 can arise in physical applications and Bayesian inference; see Section 2.3.

Generalizing the delocalization of bias effect beyond sparse potentials can provide further insights on the behavior of the algorithm in high dimensions. We also approach this through an alternative asymptotic perspective, in the spirit of methods using Taylor's expansion and the Poisson equation [39, 40, 1, 31]. We derive explicit formulas for the bias of observables in first orders and demonstrate how the delocalization of bias, in the context of observables, can hold in a wider generality.

1.4. Organization of this paper. Section 2 presents our main result on sparse potentials. In Section 3, we discuss generalizations of our result through asymptotic arguments. We conclude the paper in Section 4. All proofs are deferred to the appendices.

1.5. Notations. We write $A = O(B)$ or $A \lesssim B$ to mean that there exists a constant C independent of α, β, h, d such that $A \leq CB$. On the other hand, $A = \Omega(B)$ or $A \gtrsim B$ means that there exists a constant C independent of α, β, h, d such that $A \geq CB$. We use $\|\cdot\|_\infty$ and $\|\cdot\|_2$ to represent the ℓ^∞ and ℓ^2 norms for vectors and matrices. When applied to matrices, they stand for the corresponding operator norms. We use $\preceq$ for the Loewner order such that if $M \preceq N$ where M, N are symmetric matrices, then the matrix $N - M$ is positive semi-definite.

2. STRONGLY LOG-CONCAVE DISTRIBUTIONS WITH SPARSE INTERACTIONS

In this section, we begin by mathematically defining the potential with sparse interactions (as illustrated in Figure 1) in subsection 2.1. Subsection 2.2 presents the convergence bound. Finally, in subsection 2.3, we discuss examples of sparse potentials that satisfy the assumptions of our theorem.

2.1. Sparse graphical models. Consider an undirected graph G with d nodes, labeled by $1 \leq i \leq d$. If there is an edge between node i and j , we write $i \sim j$, meaning that the two nodes are connected. We denote by $\mathbf{N}(i)$ the *neighborhood* of the node i , which is a set of nodes connected to i . Here the neighborhood relationship is symmetric, and without loss of generality, we assume $i \sim i$ for all $1 \leq i \leq d$. Furthermore, we define recursively that $\mathbf{N}_k(i) = \{1 \leq j \leq d : \exists \ell \in \mathbf{N}_{k-1}(i), \text{ such that } j \sim \ell\}$, for $k \geq 2$. Here $\mathbf{N}_1(i) := \mathbf{N}(i)$. Denote the cardinality of the set $\mathbf{N}_k(i)$ by $|\mathbf{N}_k(i)|$ and let $\max_{1 \leq i \leq d} |\mathbf{N}_k(i)| = s_k$. We also write $s = s_1$.

Assumption 2.1. *Let the potential V be of the form*

$$(2.1) \quad V(x) = \sum_{i=1}^d V_i(\mathcal{X}^i)$$

where $\mathcal{X}^i = \{x^{(j)} : j \in \mathbf{N}(i)\}$ and V_i is a function depending on variables in the set $\mathcal{X}^i$.

We can also understand G as a *factor graph*, where $F = \{\mathcal{X}^i, 1 \leq i \leq d\}$ represents the *factor vertices*. By definition, the potential V described above exhibits sparse interaction when the graph is sparse.

2.2. Convergence bound for unadjusted Langevin. Below we present the bias and convergence bounds in terms of the W_{2,ℓ^∞} metric. The sparsity parameters s_k will play important roles in our bound.

Theorem 2.2. *Let Assumptions 1.1 and 2.1 hold. Assume $0 \leq h \leq 1/\beta$ and denote $q = \exp(-h\alpha)$. Let $r_i = \lceil e^2 i h \beta + \log \sqrt{d} \rceil$ for $i \in \mathbb{N}$. Then, for any $N \in \mathbb{N}$ such that $2q^N \sqrt{s_{r_N}} < 1$, we have the bound*

$$(2.2) \quad \begin{aligned} W_{2,\ell^\infty}(\pi_h, \pi) &\leq \text{bias}(N, q, \beta, h) \\ &:= \frac{2\beta \sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}}}{1 - 2q^N \sqrt{s_{r_N}}} \left(h^2 \sqrt{\mathbb{E}_\pi[|\nabla V|_\infty^2]} + 3h^{3/2} \sqrt{\log(2d)} \right). \end{aligned}$$

Moreover, the following convergence bound holds for any $k \in \mathbb{N}$:

$$(2.3) \quad W_{2,\ell^\infty}(\rho_{kh}, \pi) \leq q^k W_2(\rho_0, \pi_h) + \text{bias}(N, q, \beta, h).$$

Sketch of proof. Using the same coupling argument and notations in Example 1.2, we have

$$\sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq (a) + O(\beta h^{3/2} \sqrt{\log(2d)}),$$

where $(a) = \sqrt{\mathbb{E}[|X_{kh} - Y_{kh} - h(\nabla V(X_{kh}) - \nabla V(Y_{kh}))|_\infty^2]}$. Now π is not a product measure and we no longer have the contraction $(a) \leq (1 - \alpha h) \sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]}$. In fact, using Taylor's expansion, we can write

$$X_{kh} - Y_{kh} - h(\nabla V(X_{kh}) - \nabla V(Y_{kh})) = H_k(X_{kh} - Y_{kh})$$

where $H_k = I - h \int_0^1 \nabla^2 V(uX_{kh} + (1-u)Y_{kh}) du$. Here $|H_k|_2 \leq 1 - \alpha h$. When π is a product measure, H_k is a diagonal matrix and thus its ℓ^∞ norm equals ℓ^2 norm so we can get the contraction property. However, in general, $|H_k|_\infty$ can be much larger than $|H_k|_2$ so the one-step contraction fails. Specifically $|H_k|_\infty \leq \sqrt{s_1} |H_k|_2$ based on the fact that H_k has only s_1 nonzero entries in each column. As a result, we must employ a multi-step coupling argument to analyze the iterations. We now have

$$\sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq \sqrt{\mathbb{E}[|H_k(X_{kh} - Y_{kh})|_\infty^2]} + \text{error}(1),$$

where $\text{error}(1) = O(\beta h^{3/2} \sqrt{\log(2d)})$. Applying the bound again within the first term on the right hand side, yields

$$\sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq \sqrt{\mathbb{E}[|H_k H_{k-1}(X_{(k-1)h} - Y_{(k-1)h})|_\infty^2]} + \text{error}(2),$$

where $\text{error}(2)$ represents the two-step error. More generally, for any $k, N \in \mathbb{N}$, we get

$$\sqrt{\mathbb{E}[|X_{(k+N)h} - Y_{(k+N)h}|_\infty^2]} \leq \sqrt{\mathbb{E}[|H_{k+N-1} H_{k+N-2} \cdots H_k(X_{kh} - Y_{kh})|_\infty^2]} + \text{error}(N).$$

Now by choosing a large N , we will get $|H_{k+N-1} H_{k+N-2} \cdots H_k|_\infty < 1$. In fact, a simple bound is

$$|H_{k+N-1} H_{k+N-2} \cdots H_k|_\infty \leq \sqrt{d} |H_{k+N-1} H_{k+N-2} \cdots H_k|_2 \leq \sqrt{d} \exp(-N\alpha h).$$

Thus $N \sim \frac{\log d}{h}$ will lead to a contraction. In our proof, we leverage the sparsity of the potential to produce a potentially tighter bound $|H_{k+N-1} H_{k+N-2} \cdots H_k|_\infty \leq 2q^N \sqrt{s_{r_N}}$.

The primary challenge lies in controlling the growth of $\text{error}(N)$. A naïve approach would lead to polynomial dependence on d since the bound may grow exponentially

fast in N without contraction. In contrast, our analysis hinges on a sparsity bound of the propagator of the unadjusted Langevin algorithm over multiple steps; see Appendix C.1. This sparsity enables us to obtain tighter bounds on these ℓ^∞ errors across multiple iterations. In fact, we can see that in (2.2), r_i scales with the physical time ih , which justifies that s_{r_i} can characterize the sparsity of interactions after i steps of the algorithm. Using the sparsity bound, we finally control the accumulated discretization errors in N steps by

$$\text{error}(N) \leq 2\beta \left(\sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}} \right) \left(h^2 \sqrt{\mathbb{E}_\pi[|\nabla V|_\infty^2]} + 3h^{3/2} \sqrt{\log(2d)} \right).$$

With the N -step contraction and the bound on the accumulated discretization errors, we will get

$$W_{2,\ell^\infty}(\rho_{(k+N)h}, \pi) \leq 2q^N \sqrt{s_{r_N}} W_{2,\ell^\infty}(\rho_{kh}, \pi) + \text{error}(N),$$

for any $k, N \in \mathbb{N}$. When $2q^N \sqrt{s_{r_N}} < 1$, we can let $k \rightarrow \infty$ and get the final bound in (2.2). The complete proof for Theorem 2.2 can be found in Appendix C. $\square$

Overall, Theorem 2.2 implies that the bias depends on the sparsity growth, the expectation of $|\nabla V|_\infty^2$, and $\log(2d)$. In the following, we state a result that provides an upper bound on the expected $|\nabla V|_\infty^2$.

Proposition 2.3. *Let Assumptions 1.1 hold. Then, it holds that*

$$(2.4) \quad \sqrt{\mathbb{E}_\pi[|\nabla V|_\infty^2]} \leq 2 \frac{\beta}{\sqrt{\alpha}} \sqrt{\log(2d)}.$$

Sketch of Proof. When V is strongly convex, by the Bakry-Émery criterion, π satisfies the log-Sobolev inequality with constant $1/\alpha$. Using this fact and Herbst's argument leads to the bound (2.4). The complete proof can be found in Appendix C.3. $\square$

Now, we are ready to provide one instantiation of Theorem 2.2 and Proposition 2.3 for the case that the sparsity of the neighborhood of the graphical model grows at most polynomially. We need to carefully calculate an explicit bound on the terms in (2.2); the detailed proof is in Appendix C. Section 2.3 will provide concrete examples illustrating this case.

Theorem 2.4. *Let Assumptions 1.1 and 2.1 hold with $s_k \leq C(k+1)^n$ for any $k \in \mathbb{N}$, where $C > 0$ is a generic constant and $n \in \mathbb{N}$. Then for $h \leq \frac{\alpha}{\beta^2}$, we have*

$$(2.5) \quad W_{2,\ell^\infty}(\pi_h, \pi) \leq \sqrt{h \log(2d)} \min \left\{ \left(O\left(\frac{\beta}{\alpha} \log(2d)\right) \right)^{\frac{n}{2}+1}, O\left(\frac{\beta}{\alpha} \sqrt{d}\right) \right\}.$$

Moreover, it holds that

$$(2.6) \quad W_{2,\ell^\infty}(\rho_{kh}, \pi) \leq q^k W_2(\rho_0, \pi_h) + \sqrt{h \log(2d)} \min \left\{ \left(O\left(\frac{\beta}{\alpha} \log(2d)\right) \right)^{\frac{n}{2}+1}, O\left(\frac{\beta}{\alpha} \sqrt{d}\right) \right\}.$$

In Theorem 2.4, we express $\left(O\left(\frac{\beta}{\alpha} \log(2d)\right) \right)^{\frac{n}{2}+1}$ in this manner to ensure that the hidden constant is also independent of n . In the worst-case scenario, where there is no sparsity present at all, n scales with $\log d$.

Theorem 2.4 implies that the W_{2,ℓ^∞} bias is nearly independent of d , up to logarithmic terms, if we assume n is independent of d and the condition number β/α is also nearly independent of d . In such case, by taking $h \sim 1/K$, the iteration complexity scales with K to ensure a bounded W_2 error for all K -marginals, up to logarithmic terms on d , per the discussion in Section 1.3.

Remark 2.5. The current upper bound in Theorem 2.4 involves $\sqrt{h}$. This is in alignment with the W_2 bound proved in the literature (under Assumption 1.1) based on coupling arguments where such square root dependence manifests. The second part of the bound, $\sqrt{h \log(2d)} O\left(\frac{\beta}{\alpha} \sqrt{d}\right)$, is consistent, up to logarithmic terms in d , with the established result for the W_2 bound. In fact, using the W_2 bound, we have that $W_{2,\ell^\infty}(\pi_h, \pi) \leq W_2(\pi_h, \pi) = O\left(\frac{\beta}{\alpha} \sqrt{dh}\right)$. $\diamond$

2.3. Examples of sparse graphical models. In this subsection, we provide examples that satisfy Assumptions 1.1 and 2.1. For the bound in Theorem 2.4 to be non-trivial, the ratio β/α should not grow too rapidly with the dimension. The simplest case would be when α and β are fixed constants, independent of the dimension.

It is not immediately apparent that given constant values of α and β such that $\alpha \leq \beta$, there exists a potential function V for any dimension $d \in \mathbb{N}$ with $\alpha I \preceq \nabla^2 V(x) \preceq \beta I$ at any point $x \in \mathbb{R}^d$, and where s_k grows at most polynomially in k , as required in Theorem 2.4. As the following example shows, a natural class of examples of such potentials arises from Laplacians of sparse graphs.

Example 2.6. We provide an example of $V : \mathbb{R}^d \rightarrow \mathbb{R}$ for which $\alpha I \preceq \nabla^2 V(x) \preceq \beta I$ at any $x \in \mathbb{R}^d$ with α, β independent of d , and $s_k = \min\{2k + 1, d\}$.

Consider the following matrix

$$\begin{bmatrix} 2 + \lambda(x) & -1 & 0 & 0 & \cdots & 0 \\ -1 & 2 + \lambda(x) & -1 & 0 & \cdots & 0 \\ 0 & -1 & 2 + \lambda(x) & -1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & -1 & 2 + \lambda(x) & -1 \\ 0 & 0 & \cdots & 0 & -1 & 2 + \lambda(x) \end{bmatrix} \in \mathbb{R}^{d \times d}.$$

This matrix can be seen as $-\Delta + \lambda(x)I$, where Δ is the discretization of the 1D Laplace operator with homogeneous Dirichlet boundary conditions. This operator satisfies $0 \preceq -\Delta \preceq 4I$. Therefore, if V is a potential such that $\nabla^2 V(x)$ is equal to the above matrix and λ is a bounded function satisfying $\min_{x \in \mathbb{R}^d} \lambda(x) = \alpha > 0$, then we have $s_k = \max\{2k + 1, d\}$ and $\alpha I \preceq \nabla^2 V(x) \preceq \beta I$ with $\beta = 4 + \max_{x \in \mathbb{R}^d} \lambda(x)$, in any dimension d . Thus this model satisfies the assumption in Theorem 2.4 with $n = 1$.

More generally, consider a sparse, bounded-degree graph G on the variables $x^{(1)}, \dots, x^{(d)}$ and potentials of the form $V(x) = \sum_{i=1}^d V_i(x^{(i)}) + \frac{1}{2} \sum_{i,j:(i,j) \in E(G)} (x^{(i)} - x^{(j)})^2$, where $V_1, \dots, V_d$ are arbitrary smooth and strongly convex functions and $E(G)$ is the edge set of G . The resulting probability measures are log-concave perturbations of a Gaussian free field. Then V will satisfy Assumptions 1.1 and 2.1 with α and β independent of the dimension, analogous to Example 2.6. An important special case is when G is a subset of the lattice $\mathbb{Z}^n$, in which case the sparsity parameter satisfies $s_k \leq C(k + 1)^n$.

These examples can arise in physical systems where there are local interactions and a restoring force for each state which plays the role of $V_i(x^{(i)})$, or in Bayesian inverse problems where such $V_i(x^{(i)})$ terms can come from priors.

3. DISCUSSIONS AND GENERALIZATION WITH ASYMPTOTIC ARGUMENTS

In this section, we discuss generalizations of the results beyond log-concave distributions with sparse interactions. While our analysis requires strict sparsity, where most interactions are zero, we anticipate that the analysis can be extended to cases where all but a small number of interactions are weak. Moreover, our sparsity assumptions treat different coordinates equally, and it is of interest to consider heterogeneous sparse models where some coordinates have more, but potentially weaker, connections with others. The study of these scenarios is left for future endeavors.

Alternatively, we can approach the question from a different perspective, employing asymptotic analysis. More precisely, we adopt the Poisson argument and utilize the generator of the Markov process to derive an asymptotic formula for the bias of certain observables. By doing so, we can gain additional insights into the algorithm's bias.

Recall our definition $V(x) = -\log \pi(x)$. Let $\mathcal{L}$ and $\mathcal{L}_h$ be the generators of the continuous Langevin dynamics and the unadjusted Langevin algorithm, respectively. By definition, $\mathcal{L}u = \nabla \log \pi \cdot \nabla u + \Delta u$ and

$$\mathcal{L}_h u(x) = \frac{\mathbb{E}[u(x + h \nabla \log \pi(x) + \sqrt{2h} \xi)] - u(x)}{h}.$$

Below we calculate the first order formula for the bias of an observable f . We assume that f is sufficiently regular so that the Taylor expansions and integration by parts used in the Proof of Proposition 3.1 below are valid. In this section, we use the notation $A = o(h)$ to mean that the term A is a higher order term than h as $h \rightarrow 0$.

Proposition 3.1. *Without loss of generality, assume $\int f \pi = 0$. Then it holds that*

$$(3.1) \quad \int f \pi - \int f \pi_h = \frac{1}{4} h \int (-2\Delta f + |\nabla \log \pi|_2^2 f) \pi + o(h).$$

Moreover, we also have the following formula:

$$(3.2) \quad \int f \pi - \int f \pi_h = -\frac{1}{4} h \int (\Delta f + f \Delta \log \pi) \pi + o(h).$$

The proof of the formula can be found in Appendix D. Let us make some observations regarding Proposition 3.1. If π is Gaussian, then $\int f(\Delta \log \pi) \pi = 0$ since $\int f \pi = 0$. We have

$$\int f \pi - \int f \pi_h = -\frac{1}{4} h \int (\Delta f) \pi + o(h).$$

This means that if f is a linear observable, then the first-order term of the bias is zero. Moreover, if f depends only on a small number of coordinates of $x \in \mathbb{R}^d$, then the integral $\int (\Delta f) \pi$ will scale with that number, rather than d , because the entire integral will rely only on the marginal distribution of π at these coordinates. Thus, to leading order, the bias of the observable is delocalized across dimensions.

In fact, the above argument can be further generalized to π that is a perturbation of Gaussians; see the following Proposition 3.2, and its proof is in Appendix D.

Proposition 3.2. *Let $\pi \propto \exp(-V(x)) \propto \mathcal{N}(x; m, \Sigma) \exp(-U(x))$ where $\mathcal{N}(x; m, \Sigma)$ is the density of a Gaussian with mean m and covariance Σ . Assume π is centered, and $\int (\Delta U)^2 \pi \leq C_0$ for some constant C_0 independent of d . Suppose f depends only on K coordinates of $x \in \mathbb{R}^d$, then the first-order term of the bias depends on K , not d . In particular, for the observable $f(x) = \sum_{k=1}^K x^{(k)}$, under the additional assumption that π satisfies a Poincaré inequality with a positive dimension-independent constant, we have*

$$\left| \int f \pi - \int f \pi_h \right| \leq C_1 \sqrt{K} h + o(h),$$

where C_1 is a constant independent of d .

Proposition 3.2 shows that the delocalization of bias for observables holds if π is a perturbation of a Gaussian distribution, suggesting that the delocalization effect may hold in significant generality. Note, however, the asymptotic arguments in this section do not provide insight into the behavior of the higher order contributions to the bias. On the other hand, the asymptotic formula is applicable beyond log-concave distributions and can be applied to study the bias of any observable of interest.

Proposition 3.2 also suggests that the Gaussian part may not have a significant effect on the bias of the observable f . Recall that our W_{2,ℓ^∞} bounds apply to Gaussian distributions (which can have dense interactions) and log-concave distributions with sparse interactions. It is natural to inquire whether our W_{2,ℓ^∞} analysis can be extended to the case of a distribution that is the product of a dense Gaussian and a sparse log-concave distribution.

4. CONCLUSIONS

In this article, we studied the convergence of unadjusted Langevin algorithms in high dimensions. For strongly log-concave distributions, existing results showed that the iteration complexity scales proportionally to d or $\sqrt{d}$ in order to achieve a desired error in the W_2 metric for any dimension d . We demonstrate that for Gaussian distributions or distributions with certain sparsity structures, a constant number of iterations, up to some logarithmic terms in d , suffices to achieve a bounded error in the W_{2,ℓ^∞} metric. Consequently, a number of iterations proportional to K (up to some $\log d$ terms) can achieve a bounded W_2 error for all K -marginals. This result implies that even in extremely high-dimensional settings, unadjusted Langevin algorithms can still be highly scalable if the quantities of interest depend only on low-dimensional marginals. We note that this desirable property is not satisfied for unbiased schemes such as the MALA or proximal samplers, where the constraint on the step size necessitates poor scaling in d .

Our discussion hinges on a delocalization effect of the bias: the overall bias of the algorithm is delocalized to a comparable order across each one-dimensional marginal distribution. Thus, to obtain accurate low-dimensional marginals, one does not suffer from a poor dependence on the full dimension d . The delocalization effect is rigorously shown in this paper for Gaussian distributions and distributions with certain sparsity structures. Our error bounds have a favorable dependence on d ; however, the dependence on the condition number β/α may potentially be improved. We also provide counterexamples of a rotated product measure for which the bias is not delocalized.

Extending the delocalization analysis to other broader classes of distributions is of future interest. In fact, our asymptotic arguments based on Poisson equations provided further insights into how the first-order terms of the bias of general observables scale with dimension.

More broadly, it is an important question to understand how the complexity of other unadjusted MCMC algorithms scales with dimension when the quantities of interest depend solely on low-dimensional marginals. Addressing this question can provide valuable insights to help practitioners select appropriate algorithms and understand their computational complexity when dealing with high-dimensional sampling problems.

Acknowledgments We thank Nawaf Bou-Rabee and Aaron Dinner for helpful discussions. Y. Chen is supported in part by the Office of Naval Research project under award N00014-22-1-2728 (PI: Benjamin Peherstorfer) and Vannevar Bush award “Mathematical Foundations and Scientific Applications of Machine Learning” (PI: Eric Vanden-Eijnden). J. Weare and X. Cheng are supported in part by National Science Foundation award DMS-2054306. J. Niles-Weed is supported in part by National Science Foundation awards DMS-2210583 and DMS-2339829.

REFERENCES

- [1] Vlad Bally and Denis Talay. The law of the Euler scheme for stochastic differential equations: I. convergence rate of the distribution function. *Probability theory and related fields*, 104:43–60, 1996.
- [2] Nawaf Bou-Rabee and Katharina Schuh. Convergence of unadjusted Hamiltonian Monte Carlo for mean-field models. *Electronic Journal of Probability*, 28:1–40, 2023.
- [3] Laird Arnault Breyer, Mauro Piccioni, and Sergio Scarlatti. Optimal scaling of MALA for nonlinear regression. *The Annals of Applied Probability*, 14(3):1479, 2004.
- [4] Niladri Chatterji, Jelena Diakonikolas, Michael I Jordan, and Peter Bartlett. Langevin Monte Carlo without smoothness. In *International Conference on Artificial Intelligence and Statistics*, pages 1716–1726. PMLR, 2020.
- [5] Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 2984–3014. PMLR, 02–05 Jul 2022.
- [6] Yuansi Chen, Raaz Dwivedi, Martin J. Wainwright, and Bin Yu. Fast mixing of Metropolized Hamiltonian Monte Carlo: Benefits of multi-step gradients. *J. Mach. Learn. Res.*, 21:92:1–92:72, 2020.
- [7] Yuansi Chen and Khashayar Gatmiry. A simple proof of the mixing of Metropolis-adjusted Langevin algorithm under smoothness and isoperimetry. *arXiv preprint arXiv:2304.04095*, 2023.
- [8] Xiang Cheng and Peter Bartlett. Convergence of Langevin MCMC in KL-divergence. In *Algorithmic Learning Theory*, pages 186–211. PMLR, 2018.
- [9] Sinho Chewi, Murat A Erdogdu, Mufan Li, Ruqi Shen, and Shunshi Zhang. Analysis of Langevin Monte Carlo from Poincare to log-Sobolev. In *Conference on Learning Theory*, pages 1–2. PMLR, 2022.
- [10] Sinho Chewi, Chen Lu, Kwangjun Ahn, Xiang Cheng, Thibaut Le Gouic, and Philippe Rigollet. Optimal dimension dependence of the Metropolis-adjusted Langevin algorithm. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 1260–1300. PMLR, 15–19 Aug 2021.
- [11] Ole F Christensen, Gareth O Roberts, and Jeffrey S Rosenthal. Scaling limits for the transient phase of local Metropolis–Hastings algorithms. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(2):253–268, 2005.

- [12] Arnak Dalalyan. Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 678–689. PMLR, 07–10 Jul 2017.
- [13] Arnak S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 79(3):651–676, 2017.
- [14] Alain Durmus and Andreas Eberle. Asymptotic bias of inexact Markov chain Monte Carlo methods in high dimension. *The Annals of Applied Probability*, 34(4):3435–3468, August 2024.
- [15] Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019.
- [16] Alain Durmus and Éric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017.
- [17] Alain Durmus and Éric Moulines. High-dimensional Bayesian inference via the unadjusted Langevin algorithm. *Bernoulli*, 25(4A):2854 – 2882, 2019.
- [18] Raaz Dwivedi, Yuansi Chen, Martin J Wainwright, and Bin Yu. Log-concave sampling: Metropolis-Hastings algorithms are fast! In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 793–797. PMLR, 06–09 Jul 2018.
- [19] Murat A Erdogdu and Rasa Hosseinzadeh. On the convergence of Langevin Monte Carlo: The interplay between tail growth and smoothness. In *Conference on Learning Theory*, pages 1776–1822. PMLR, 2021.
- [20] Murat A. Erdogdu, Rasa Hosseinzadeh, and Shunshi Zhang. Convergence of Langevin Monte Carlo in Chi-squared and Rényi divergence. In Gustau Camps-Valls, Francisco J. R. Ruiz, and Isabel Valera, editors, *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 8151–8175. PMLR, 28–30 Mar 2022.
- [21] Jiaojiao Fan, Bo Yuan, and Yongxin Chen. Improved dimension dependence of a proximal algorithm for sampling. In Gergely Neu and Lorenzo Rosasco, editors, *The Thirty Sixth Annual Conference on Learning Theory, COLT 2023, 12-15 July 2023, Bangalore, India*, volume 195 of *Proceedings of Machine Learning Research*, pages 1473–1521. PMLR, 2023.
- [22] Vytautas Gapsys, Wojciech Kopec, Dirk Matthes, and Bert L. de Groot. Biomolecular simulations at the exascale: From drug design to organelles and beyond. *Current Opinion in Structural Biology*, 88:102887, 2024.
- [23] Ye He, Krishnakumar Balasubramanian, and Murat A Erdogdu. On the ergodicity, bias and asymptotic normality of randomized midpoint sampling method. *Advances in Neural Information Processing Systems*, 33:7366–7376, 2020.
- [24] Juan Kuntz, Michela Ottobre, and Andrew M Stuart. Non-stationary phase of the MALA algorithm. *Stochastics and Partial Differential Equations: Analysis and Computations*, 6:446–499, 2018.
- [25] Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Lower bounds on Metropolized sampling methods for well-conditioned distributions. *Advances in Neural Information Processing Systems*, 34:18812–18824, 2021.
- [26] Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured logconcave sampling with a restricted Gaussian oracle. In Mikhail Belkin and Samory Kpotufe, editors, *Proceedings of Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pages 2993–3050. PMLR, 15–19 Aug 2021.
- [27] Joseph Lehec. The Langevin Monte Carlo algorithm in the non-smooth log-concave case. *The Annals of Applied Probability*, 33(6A):4858–4874, 2023.
- [28] Ben Leimkuhler and Charles Matthews. *Molecular Dynamics : With Deterministic and Stochastic Numerical Methods*. Interdisciplinary Applied Mathematics, 39. Springer International Publishing, Cham, 1st ed. 2015. edition, 2015.
- [29] Ruilin Li, Hongyuan Zha, and Molei Tao. Sqrt(d) dimension dependence of Langevin Monte Carlo. In *International Conference on Learning Representations*, 2022.

- [30] Xuechen Li, Yi Wu, Lester Mackey, and Murat A Erdogdu. Stochastic Runge-Kutta accelerates Langevin Monte Carlo and beyond. *Advances in neural information processing systems*, 32, 2019.
- [31] Jonathan C Mattingly, Andrew M Stuart, and Michael V Tretyakov. Convergence of numerical time-averaging and stationary measures via Poisson equations. *SIAM Journal on Numerical Analysis*, 48(2):552–577, 2010.
- [32] Wenlong Mou, Nicolas Flammarion, Martin J Wainwright, and Peter L Bartlett. Improved bounds for discretization of Langevin diffusions: Near-optimal rates without convexity. *Bernoulli*, 28(3):1577–1601, 2022.
- [33] Grigorios A Pavliotis. Stochastic processes and applications. *Springer*, 2014.
- [34] Natesh S. Pillai, Andrew M. Stuart, and Alexandre H. Thiéry. Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions. *The Annals of Applied Probability*, 22(6):2320 – 2356, 2012.
- [35] Gareth O. Roberts and Jeffrey S. Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 60(1):255–268, 1998.
- [36] Gareth O Roberts and Richard L Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996.
- [37] Peter J Rossy, Jimmie D Doll, and Harold L Friedman. Brownian dynamics as smart Monte Carlo simulation. *The Journal of Chemical Physics*, 69(10):4628–4633, 1978.
- [38] Ruoyi Shen and Yin Tat Lee. The randomized midpoint method for log-concave sampling. *Advances in Neural Information Processing Systems*, 32, 2019.
- [39] Denis Talay. Second-order discretization schemes of stochastic differential systems for the computation of the invariant law. *Stochastics: An International Journal of Probability and Stochastic Processes*, 29(1):13–36, 1990.
- [40] Denis Talay and Luciano Tubaro. Expansion of the global error for numerical schemes solving stochastic differential equations. *Stochastic analysis and applications*, 8(4):483–509, 1990.
- [41] Santosh Vempala and Andre Wibisono. Rapid convergence of the unadjusted Langevin algorithm: Isoperimetry suffices. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [42] Martin J Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019.
- [43] Keru Wu, Scott Schmidler, and Yuansi Chen. Minimax mixing time of the Metropolis-adjusted Langevin algorithm for log-concave sampling. *Journal of Machine Learning Research*, 23(270):1–63, 2022.

APPENDIX A. PROOFS FOR LOG-CONCAVE PRODUCT MEASURES

Proposition A.1. *Under Assumption 1.1, let $V(x) = \sum_{i=1}^d V_i(x^{(i)})$ such that $\pi \propto \exp(-V)$ is a product measure. Assume $0 \leq h \leq 1/\beta$ and denote $q = \exp(-h\alpha)$. Consider the Langevin dynamics $dY_t = -\nabla V(Y_t)dt + \sqrt{2}dB_t$ with $Y_0 \sim \pi$. Let ρ_{kh} denote the probability distribution of X_{kh} from the iterations of the unadjusted Langevin algorithm with stepsize h . Then, the following estimates hold for any $k \in \mathbb{N}$:*

$$(A.1) \quad W_{2,\ell^\infty}(\rho_{(k+1)h}, \pi) \leq qW_{2,\ell^\infty}(\rho_{kh}, \pi) + \beta \sqrt{\left(\frac{8\beta^2}{3\alpha}h^4 + 8h^3\right) \log(2d)}.$$

This implies that

$$W_{2,\ell^\infty}(\pi_h, \pi) \leq \frac{\beta}{\alpha} \sqrt{\left(\frac{8\beta^2}{3\alpha}h^2 + 8h\right) \log(2d)}.$$

Furthermore, if $h \leq \frac{\alpha}{\beta^2}$, we get

$$W_{2,\ell^\infty}(\pi_h, \pi) \leq \frac{4\beta}{\alpha} \sqrt{h \log(2d)}.$$

Proof. First, for the continuous-time Langevin dynamics Y_t and the discrete-time iterates X_{kh} in the unadjusted Langevin algorithm, we have:

$$(A.2) \quad \begin{aligned} Y_{(k+1)h} &= Y_{kh} - \int_{kh}^{(k+1)h} \nabla V(Y_t) dt + \sqrt{2}(B_{(k+1)h} - B_{kh}) \\ X_{(k+1)h} &= X_{kh} - h\nabla V(X_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh}). \end{aligned}$$

We couple the two processes using the same Brownian motion. We aim to estimate $\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]$. For this purpose we introduce an auxiliary random variable

$$(A.3) \quad \bar{Y}_{(k+1)h} = Y_{kh} - h\nabla V(Y_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh}).$$

Using the triangle inequality then leads to

$$(A.4) \quad \sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq \underbrace{\sqrt{\mathbb{E}[|X_{(k+1)h} - \bar{Y}_{(k+1)h}|_\infty^2]}}_{(a)} + \underbrace{\sqrt{\mathbb{E}[|\bar{Y}_{(k+1)h} - Y_{(k+1)h}|_\infty^2]}}_{(b)}.$$

For (a), we have

$$(A.5) \quad \begin{aligned} (a) &= \sqrt{\mathbb{E}[|X_{kh} - Y_{kh} - h(\nabla V(X_{kh}) - \nabla V(Y_{kh}))|_\infty^2]} \\ &= \sqrt{\mathbb{E}[\max_{1 \leq i \leq d} |X_{kh}^{(i)} - Y_{kh}^{(i)} - h(\nabla V_i(X_{kh}^{(i)}) - \nabla V_i(Y_{kh}^{(i)}))|^2]} \\ &\leq \sqrt{\mathbb{E}[\max_{1 \leq i \leq d} (1 - h\alpha)^2 |X_{kh}^{(i)} - Y_{kh}^{(i)}|^2]} \\ &\leq q\sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]}, \end{aligned}$$

where we used the facts that $V(x) = \sum_{i=1}^d V_i(x^{(i)})$ and $0 \leq 1 - h\beta \leq 1 - h\alpha$. For (b), we have

$$(A.6) \quad \begin{aligned} \mathbb{E}[|\bar{Y}_{(k+1)h} - Y_{(k+1)h}|_\infty^2] &= \mathbb{E}[|\int_{kh}^{(k+1)h} \nabla V(Y_t) - \nabla V(Y_{kh}) dt|_\infty^2] \\ &\leq h \int_{kh}^{(k+1)h} \mathbb{E}[|\nabla V(Y_t) - \nabla V(Y_{kh})|_\infty^2] dt \\ &\leq h \int_{kh}^{(k+1)h} \int_0^1 \mathbb{E}[|\nabla^2 V(uY_t + (1-u)Y_{kh})(Y_t - Y_{kh})|_\infty^2] du dt \\ &\leq h\beta^2 \int_{kh}^{(k+1)h} \mathbb{E}[|Y_t - Y_{kh}|_\infty^2] dt, \end{aligned}$$

where in the last inequality, we used the fact that $\nabla^2 V(uY_t + (1-u)Y_{kh})$ is a diagonal matrix with each diagonal entry bounded by β in magnitude. We further have

$$\begin{aligned}
& \int_{kh}^{(k+1)h} \mathbb{E}[|Y_t - Y_{kh}|_\infty^2] dt \\
&= \int_{kh}^{(k+1)h} \mathbb{E}[|\int_{kh}^t \nabla V(Y_s) ds + \sqrt{2}B_{t-kh}|_\infty^2] dt \\
&\leq \int_{kh}^{(k+1)h} \left(2\mathbb{E}[|\int_{kh}^t \nabla V(Y_s) ds|_\infty^2] + 2\mathbb{E}[\sqrt{2}B_{t-kh}|_\infty^2] \right) dt \\
&\leq \int_{kh}^{(k+1)h} 2(t-kh) \int_{kh}^t \mathbb{E}[|\nabla V(Y_s)|_\infty^2] ds dt + \int_{kh}^{(k+1)h} 2\mathbb{E}[\sqrt{2}B_{t-kh}|_\infty^2] dt \\
&\leq 2 \int_{kh}^{(k+1)h} (t-kh)^2 \mathbb{E}_\pi[|\nabla V(Y)|_\infty^2] dt + \int_{kh}^{(k+1)h} 16(t-kh) \log(2d) dt \\
&= \frac{2}{3} h^3 \mathbb{E}_\pi[|\nabla V(Y)|_\infty^2] + 8h^2 \log(2d),
\end{aligned} \tag{A.7}$$

where we used the fact that all $Y_t \sim \pi$, and the bound $\mathbb{E}|B_u|_\infty^2 \leq 4u \log(2d)$ holds for any $u \geq 0$ due to Lemma B.1.

Then by Proposition 2.3, we have $\sqrt{\mathbb{E}_\pi[|\nabla V|_\infty^2]} \leq 2\frac{\beta}{\sqrt{\alpha}} \sqrt{\log(2d)}$. Combining all the bounds above, we get

$$\sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq q \sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]} + \frac{\beta}{\alpha} \sqrt{\left(\frac{8\beta^2}{3\alpha} h^4 + 8h^3\right) \log(2d)}.$$

We can now couple the distribution of X_{kh} and Y_{kh} such that $\sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]} = W_{2,\ell^\infty}(\rho_{kh}, \pi)$. With this and using the definition of the W_{2,ℓ^∞} norm, we get

$$W_{2,\ell^\infty}(\rho_{(k+1)h}, \pi) \leq q W_{2,\ell^\infty}(\rho_{kh}, \pi) + \frac{\beta}{\alpha} \sqrt{\left(\frac{8\beta^2}{3\alpha} h^2 + 8h\right) \log(2d)}.$$

Iterating this inequality leads to the bound on the $W_{2,\infty}$ bias. Furthermore under the assumption $h \leq \frac{\alpha}{\beta^2}$, we get

$$W_{2,\ell^\infty}(\pi_h, \pi) \leq \frac{\beta}{\alpha} \sqrt{\left(\frac{8}{3} + 8\right) h \log(2d)} \leq \frac{4\beta}{\alpha} \sqrt{h \log(2d)}.$$

The proof is complete. $\square$

APPENDIX B. PROOFS FOR GAUSSIAN DISTRIBUTIONS

We first state a lemma for the expected squared maximal norm of a random vector whose entries are sub-Gaussians. This lemma will also be used in the proof of our main theorem.

Lemma B.1. *Suppose $Y = (Y^{(1)}, Y^{(2)}, \dots, Y^{(d)}) \in \mathbb{R}^d$ and each $Y^{(i)}$ is centered and sub-Gaussian with variance proxy σ^2 , namely*

$$\mathbb{E}[\exp(\lambda Y^{(i)})] \leq \exp\left(\frac{1}{2} \lambda^2 \sigma^2\right). \tag{B.1}$$

Then, it holds that

$$(B.2) \quad \mathbb{E}[|Y|_\infty^2] \leq 4\sigma^2 \log(2d).$$

Proof of Lemma B.1. By the property of sub-Gaussian random variables [42], we have that for $0 \leq \lambda < 1/(2\sigma^2)$,

$$(B.3) \quad \mathbb{E}[\exp(\lambda(Y^{(i)})^2)] \leq \frac{1}{\sqrt{1-2\lambda\sigma^2}}.$$

Using the convexity of the exponential function, we can derive

$$(B.4) \quad \exp(\lambda \mathbb{E}[\max_{1 \leq i \leq d} (Y^{(i)})^2]) \leq \mathbb{E}[\exp(\lambda \max_{1 \leq i \leq d} (Y^{(i)})^2)] = \mathbb{E}[\max_{1 \leq i \leq d} \exp(\lambda(Y^{(i)})^2)].$$

Then, we can bound the right hand side as follows:

$$(B.5) \quad \mathbb{E}[\max_{1 \leq i \leq d} \exp(\lambda(Y^{(i)})^2)] \leq \sum_{i=1}^d \mathbb{E}[\exp(\lambda(Y^{(i)})^2)] \leq \frac{d}{\sqrt{1-2\lambda\sigma^2}}.$$

Thus, combining the above two inequalities and taking logarithms, we find

$$(B.6) \quad \mathbb{E}[|Y|_\infty^2] = \mathbb{E}[\max_{1 \leq i \leq d} (Y^{(i)})^2] \leq \frac{\log d}{\lambda} - \frac{1}{2\lambda} \log(1-2\lambda\sigma^2),$$

for any $0 \leq \lambda < 1/(2\sigma^2)$. Taking $\lambda = \frac{1}{4\sigma^2}$, we arrive at

$$(B.7) \quad \mathbb{E}[|Y|_\infty^2] \leq 4\sigma^2 \log d + 2\sigma^2 \log 2 \leq 4\sigma^2 \log(2d).$$

The proof is complete. $\square$

With Lemma B.1, we present the proof for the statement in Example 1.3 as follows.

Proof for Example 1.3. When π is Gaussian, the potential $V = \frac{1}{2}(x-m)^T \Sigma^{-1}(x-m)$ is quadratic. In this case, the iteration takes the form

$$X_{(k+1)h} - m = (I - h\Sigma^{-1})(X_{kh} - m) + \sqrt{2}(B_{(k+1)h} - B_{kh}).$$

Suppose $X_{kh} \sim \mathcal{N}(m_k, \Sigma_k)$, we have

$$m_{k+1} - m = (I - h\Sigma^{-1})(m_k - m), \Sigma_{k+1} = (I - h\Sigma^{-1})\Sigma_k(I - h\Sigma^{-1}) + 2hI.$$

Let $0 < h \leq 1/|\Sigma^{-1}|_2 = 1/\beta$, then as $k \rightarrow \infty$, we have $m_\infty = m$. For Σ_k , we have the identity:

$$\begin{aligned} \Sigma_{k+1} &= (I - h\Sigma^{-1})\Sigma_k(I - h\Sigma^{-1}) + 2hI \\ &= (I - h\Sigma^{-1})^{k+1}\Sigma_0(I - h\Sigma^{-1})^{k+1} + 2h \sum_{\ell=0}^k (I - h\Sigma^{-1})^{2\ell}. \end{aligned}$$

Letting $k \rightarrow \infty$, we get

$$\Sigma_\infty = 2h(I - (I - h\Sigma^{-1})^2)^{-1} = \Sigma(I - \frac{h}{2}\Sigma^{-1})^{-1}.$$

Thus, $\pi_h = \mathcal{N}(m_\infty, \Sigma_\infty)$ when $\pi = \mathcal{N}(m, \Sigma)$. Consider the coupling $X = \Sigma^{1/2}Z + m$ and $Y = \Sigma_\infty^{1/2}Z + m$ where $Z \sim \mathcal{N}(0, I)$. Then, for this specific coupling, we have the bound

$$(B.8) \quad W_{2,\ell^\infty}^2(\pi, \pi_h) \leq \mathbb{E}[|X - Y|_\infty^2] = \mathbb{E}[|(\Sigma^{1/2} - \Sigma_\infty^{1/2})Z|_\infty^2].$$

Setting $Y = (\Sigma^{1/2} - \Sigma_\infty^{1/2})Z$, we see that $Y \sim \mathcal{N}(0, \sigma_i^2)$, where $\sigma_i^2 = ((\Sigma^{1/2} - \Sigma_\infty^{1/2})^2)_{ii} \leq |\Sigma^{1/2} - \Sigma_\infty^{1/2}|_2^2$. In particular, the entries of Y are $|\Sigma^{1/2} - \Sigma_\infty^{1/2}|_2^2$ -subgaussian.

Consider the eigendecomposition $\Sigma = Q^T \Lambda Q$ where Q is an orthogonal matrix and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_d)$; we note that $\alpha \leq \frac{1}{\lambda_i} \leq \beta, 1 \leq i \leq d$. Then we have the formula

$$\Sigma_\infty = Q^T \Lambda (I - \frac{h}{2} \Lambda^{-1})^{-1} Q = Q^T \Lambda_\infty Q,$$

where

$$\Lambda_\infty = \text{diag} \left(\frac{\lambda_1}{1 - \frac{h}{2\lambda_1}}, \dots, \frac{\lambda_d}{1 - \frac{h}{2\lambda_d}} \right).$$

We obtain

$$(B.9) \quad |\Sigma^{1/2} - \Sigma_\infty^{1/2}|_2 = \max_{1 \leq i \leq d} \left| \sqrt{\lambda_i} - \sqrt{\frac{\lambda_i}{1 - \frac{h}{2\lambda_i}}} \right| = O(\sqrt{\beta}h).$$

Combining this bound with Lemma B.1, we get $W_{2,\ell^\infty}(\pi_h, \pi) = O(\sqrt{\beta}h\sqrt{\log(2d)}) = O(\sqrt{h\log(2d)})$ as $h \leq 1/\beta$. $\square$

APPENDIX C. PROOFS FOR LOG-CONCAVE DISTRIBUTIONS WITH SPARSE INTERACTIONS

C.1. Sparsity of the propagator of unadjusted Langevin. First, we present a proposition concerning the ℓ^∞ norm of matrices that include the propagator of the Langevin Monte Carlo algorithm as a special case. We will use this proposition when analyzing the convergence of the algorithm in the next subsection.

Proposition C.1. *Assume $0 \leq h \leq 1/\beta$ and denote $q = \exp(-h\alpha)$. For the potential function in Assumption 2.1, the following facts hold:*

- (i) *The matrix $(\nabla^2 V(x_1))(\nabla^2 V(x_2)) \dots (\nabla^2 V(x_r))$, which is the multiplication of r Hessian matrices, is s_r sparse, for any $x_1, x_2, \dots, x_r \in \mathbb{R}^d, r \in \mathbb{N}$. Here we say a matrix is s_r -sparse if each row of the matrix contains at most s_r nonzero entries.*
- (ii) *Let $\nu_1, \dots, \nu_N$ be any probability measures in $\mathbb{R}^d$. Define the matrix*

$$(C.1) \quad P_N = \left(I - h \int \nabla^2 V d\nu_1 \right) \left(I - h \int \nabla^2 V d\nu_2 \right) \dots \left(I - h \int \nabla^2 V d\nu_N \right).$$

We have the inequality

$$(C.2) \quad |P_N|_\infty \leq \sqrt{s_r} q^N + \sqrt{d} \exp(-r)$$

for any $r \geq e^2 N h \beta$. In particular, taking $r_N = \lceil e^2 N h \beta + \log \sqrt{d} \rceil$, we get that

$$(C.3) \quad |P_N|_\infty \leq 2\sqrt{s_{r_N}} q^N.$$

- (iii) *In the context of (ii), consider additionally a probability measure ν_0 and the matrix*

$$(C.4) \quad J_N = \left(\int \nabla^2 V d\nu_0 \right) \left(I - h \int \nabla^2 V d\nu_1 \right) \left(I - h \int \nabla^2 V d\nu_2 \right) \dots \left(I - h \int \nabla^2 V d\nu_N \right).$$

We have the inequality

$$(C.5) \quad |J_N|_\infty \leq \beta(\sqrt{s_r}q^{N-1} + \sqrt{d}\exp(-r))$$

for any $r \geq e^2Nh\beta$. In particular, taking $r_N = \lceil e^2Nh\beta + \log \sqrt{d} \rceil$, we get that

$$(C.6) \quad |J_N|_\infty \leq 2\beta\sqrt{s_{r_N}}q^N.$$

Proof of Proposition C.1. We prove the above facts one by one.

Proof for (i). We know that the ij -th entry of the matrix

$$(\nabla^2 V(x_1))(\nabla^2 V(x_2)) \dots (\nabla^2 V(x_r))$$

is nonzero if and only if $j \in \mathbf{N}_r(i)$. Since $\max_{1 \leq i \leq d} |\mathbf{N}_r(i)| = s_r$, we will have at most s_r nonzero entries in each row. This implies that the matrix is s_r -sparse.

Proof for (ii). Let us denote $A_k = \int \nabla^2 V d\nu_k$, then $P_N = (I - hA_1)(I - hA_2) \cdots (I - hA_N)$. Expanding the product, we have

$$(C.7) \quad \begin{aligned} P_N &= \sum_{k=0}^N (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} \\ &= \sum_{k=0}^r (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} + \sum_{k=r+1}^N (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k}, \end{aligned}$$

For $k \geq r+1$, we bound

$$(C.8) \quad \left| (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} \right|_2 \leq h^k \beta^k \binom{N}{k} \leq \frac{(Nh\beta)^k}{k!} \leq \left(\frac{eNh\beta}{k} \right)^k,$$

where in the last inequality, we used the fact that $k! \geq k^k / \exp(k)$. This fact can be seen by taking $x = k$ in the inequality $\exp(x) \geq x^k / k!$.

When $r \geq e^2Nh\beta$, we can bound

$$(C.9) \quad \left| \sum_{k=r+1}^N (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} \right|_2 \leq \sum_{k=r+1}^N \left(\frac{eNh\beta}{k} \right)^k \leq \sum_{k=r+1}^N \exp(-k) \leq \exp(-r).$$

Furthermore, we obtain $\left| \sum_{k=r+1}^N (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} \right|_\infty \leq \sqrt{d} \exp(-r)$ by the Cauchy-Schwarz inequality.

On the other hand, $|P_N|_2 \leq \Pi_{k=1}^N |I - hA_k|_2 \leq (1 - h\alpha)^N \leq q^N$ where we used the fact that $h \leq 1/\beta$. Then, for each row i of P_N ,

$$\begin{aligned}
 \sum_{j=1}^N |(P_N)_{ij}| &= \sum_{j \in \mathbf{N}_r(i)} |(P_N)_{ij}| + \sum_{j \notin \mathbf{N}_r(i)} |(P_N)_{ij}| \\
 &\leq \sqrt{|\mathbf{N}_r(i)|} \sqrt{\sum_{j \in \mathbf{N}_r(i)} |(P_N)_{ij}|^2} + \sqrt{d} \exp(-r) \\
 &\leq \sqrt{s_r} \sqrt{\sum_{j=1}^N |(P_N)_{ij}|^2} + \sqrt{d} \exp(-r) \\
 &\leq \sqrt{s_r} q^N + \sqrt{d} \exp(-r),
 \end{aligned}
 \tag{C.10}$$

where in the second inequality, we used the fact that

$$\sum_{j \notin \mathbf{N}_r(i)} |(P_N)_{ij}| \leq \left| \sum_{k=r+1}^N (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} \right|_\infty,$$

since a nonzero term $(P_N)_{ij}$ for $j \notin \mathbf{N}_r(i)$ can only be produced by product of more than r matrices of the kind of A_{i_k} , as a consequence of the argument in (i). In summary, we get the result that if $r \geq e^2 N h \beta$, then

$$|P_N|_\infty = \max_i \sum_{j=1}^N |(P_N)_{ij}| \leq \sqrt{s_r} q^N + \sqrt{d} \exp(-r).$$

We note that the above inequality will also hold true when $r \geq d$.

Taking $r_N = \lceil e^2 N h \beta + \log \sqrt{d} \rceil$, we get

$$\begin{aligned}
 |P_N|_\infty &\leq \sqrt{s_{r_N}} q^N + \sqrt{d} \exp(-r_N) \\
 &\leq \sqrt{s_{r_N}} \exp(-h\alpha N) + \sqrt{d} \exp(-e^2 N h \beta - \log \sqrt{d}) \\
 &\leq \sqrt{s_{r_N}} \exp(-h\alpha N) + \exp(-e^2 N h \beta) \leq 2\sqrt{s_{r_N}} \exp(-h\alpha N) = 2\sqrt{s_{r_N}} q^N.
 \end{aligned}
 \tag{C.11}$$

Proof for (iii). Note that $J_N = (\int \nabla^2 V d\nu_0) H_N$. Following the proof for (ii), we have that when $r \geq e^2 N h \beta$,

$$\begin{aligned}
 &\left| \left(\int \nabla^2 V d\nu_0 \right) \sum_{k=r+1}^N (-1)^k h^k \sum_{1 \leq i_1 < \dots < i_k \leq N} A_{i_1} \dots A_{i_k} \right|_2 \\
 &\leq \beta \sum_{k=r+1}^N \left(\frac{e N h \beta}{k} \right)^k \leq \beta \sum_{k=r+1}^N \exp(-k) \leq \beta \exp(-r).
 \end{aligned}
 \tag{C.12}$$

Thus the ∞ -norm of the above matrix is bounded by $\sqrt{d} \beta \exp(-r)$. Then, for any $1 \leq i \leq d$,

$$\sum_{j=1}^N |(J_N)_{ij}| \leq \beta \sqrt{s_r} q^N + \sqrt{d} \beta \exp(-r).
 \tag{C.13}$$

Therefore, $|J_N|_\infty \leq \beta(\sqrt{s_r}q^{N-1} + \sqrt{d}\exp(-r))$ for $r \geq e^2Nh\beta$. Taking $r_N = \lceil e^2Nh\beta + \log \sqrt{d} \rceil$, we get $|J_N|_\infty \leq 2\beta\sqrt{s_{r_N}}q^N$. The proof is complete. $\square$

C.2. A multistep coupling argument. To prove Theorem 2.2, we first show the following Proposition C.2, which bounds the W_{2,ℓ^∞} metric through a multistep coupling argument. Note that the one-step coupling argument that is commonly used in bounding W_2 distance is not enough here, as the one-step contraction property is lost in the W_{2,ℓ^∞} metric.

Proposition C.2. *Let Assumptions 1.1 and 2.1 hold. Assume $0 \leq h \leq 1/\beta$ and denote $q = \exp(-h\alpha)$. Consider the Langevin dynamics $dY_t = -\nabla V(Y_t)dt + \sqrt{2}dB_t$ with $Y_0 \sim \pi$. Let ρ_{kh} represent the law of X_{kh} from the iterations of Langevin Monte Carlo with stepsize h . Then, the following estimates hold for any $k, N \in \mathbb{N}$:*

$$(C.14) \quad W_{2,\ell^\infty}(\rho_{(k+N)h}, \pi) \leq 2q^N \sqrt{s_{r_N}} W_{2,\ell^\infty}(\rho_{kh}, \pi) + 2\beta \left(\sum_{i=1}^N q^{i-1} \sqrt{h s_{r_i}} \epsilon_{k+N-i} \right),$$

where $r_i, 1 \leq i \leq N$ is defined in Proposition C.1, and $\epsilon_j^2 = \int_{jh}^{(j+1)h} \mathbb{E}[|Y_t - Y_{jh}|_\infty^2] dt$ for $j \in \mathbb{N}$, which satisfies

$$\epsilon_j \leq h^{3/2} \sqrt{\mathbb{E}_\pi[|\nabla V(Y)|_\infty^2]} + 3h \sqrt{\log(2d)}.$$

Proof of Proposition C.2. Let $dY_t = -\nabla V(Y_t)dt + \sqrt{2}dB_t$. We write down the following identity:

$$(C.15) \quad \begin{aligned} Y_{(k+1)h} &= Y_{kh} - \int_{kh}^{(k+1)h} \nabla V(Y_t)dt + \sqrt{2}(B_{(k+1)h} - B_{kh}) \\ X_{(k+1)h} &= X_{kh} - h\nabla V(X_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh}). \end{aligned}$$

We couple the two processes using the same Brownian motion. The goal is to estimate $\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]$. To do so we introduce an auxiliary random variable

$$(C.16) \quad \bar{Y}_{(k+1)h} = Y_{kh} - h\nabla V(Y_{kh}) + \sqrt{2}(B_{(k+1)h} - B_{kh}).$$

Using the triangle inequality then leads to

$$(C.17) \quad \sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq \underbrace{\sqrt{\mathbb{E}[|X_{(k+1)h} - \bar{Y}_{(k+1)h}|_\infty^2]}}_{(a)} + \underbrace{\sqrt{\mathbb{E}[|\bar{Y}_{(k+1)h} - Y_{(k+1)h}|_\infty^2]}}_{(b)}.$$

For (a), we have

$$(C.18) \quad \begin{aligned} (a) &= \sqrt{\mathbb{E}[|X_{kh} - Y_{kh} - h(\nabla V(X_{kh}) - \nabla V(Y_{kh}))|_\infty^2]} \\ &= \sqrt{\mathbb{E}[|X_{kh} - Y_{kh} - h(\int_0^1 \nabla^2 V(uX_{kh} + (1-u)Y_{kh})du)(X_{kh} - Y_{kh})|_\infty^2]} \\ &= \sqrt{\mathbb{E}[|H_k(X_{kh} - Y_{kh})|_\infty^2]}, \end{aligned}$$

where $H_k(X_{kh} - Y_{kh}) = (I - h \int_0^1 \nabla^2 V(uX_{kh} + (1-u)Y_{kh})du)(X_{kh} - Y_{kh})$. We can view H_k as a random matrix depending on X_{kh} and Y_{kh} . In particular, each realization of this random matrix can be written in the form $(I - h \int \nabla^2 V d\nu)$ where ν is a probability

measure; here ν has its mass on a line. This form allows us to use Proposition C.1 to analyze the iterations.

For (b), we have

$$\begin{aligned}
 (C.19) \quad \mathbb{E}[|\bar{Y}_{(k+1)h} - Y_{(k+1)h}|_\infty^2] &= \mathbb{E}\left[\left|\int_{kh}^{(k+1)h} \nabla V(Y_t) - \nabla V(Y_{kh}) dt\right|_\infty^2\right] \\
 &\leq h \int_{kh}^{(k+1)h} \mathbb{E}[|\nabla V(Y_t) - \nabla V(Y_{kh})|_\infty^2] dt \\
 &\leq h \int_{kh}^{(k+1)h} \int_0^1 \mathbb{E}[|\nabla^2 V(uY_t + (1-u)Y_{kh})(Y_t - Y_{kh})|_\infty^2] du dt \\
 &\leq hs_1\beta^2 \int_{kh}^{(k+1)h} \mathbb{E}[|Y_t - Y_{kh}|_\infty^2] dt = hs_1\beta^2 \epsilon_k^2,
 \end{aligned}$$

which can be understood as the one-step discretization error of the overdamped Langevin dynamics. In the above, we used the fact that $\nabla^2 V(uY_t + (1-u)Y_{kh})$ is s_1 -sparse and $|\nabla^2 V(uY_t + (1-u)Y_{kh})|_2 \leq \beta$; thus $|\nabla^2 V(uY_t + (1-u)Y_{kh})|_\infty \leq \sqrt{s_1}\beta$. We overestimate the right hand side $s_1\beta^2 \leq 4s_{r_1}\beta^2$.

Now, combining the above estimates, we arrive at

$$(C.20) \quad \sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \leq \sqrt{\mathbb{E}[|H_k(X_{kh} - Y_{kh})|_\infty^2]} + 2\sqrt{hs_{r_1}}\beta\epsilon_k.$$

We move one step back further and get

$$(C.21) \quad \sqrt{\mathbb{E}[|H_k(X_{kh} - Y_{kh})|_\infty^2]} \leq \underbrace{\sqrt{\mathbb{E}[|H_k(X_{kh} - \bar{Y}_{kh})|_\infty^2]}}_{(c)} + \underbrace{\sqrt{\mathbb{E}[|H_k(\bar{Y}_{kh} - Y_{kh})|_\infty^2]}}_{(d)},$$

where, similar as before, we define $\bar{Y}_{kh} = Y_{(k-1)h} - h\nabla V(Y_{(k-1)h}) + \sqrt{2}(B_{kh} - B_{(k-1)h})$. For (c), we use the same argument as earlier to get

$$(c) = \sqrt{\mathbb{E}[|H_k H_{k-1}(X_{(k-1)h} - Y_{(k-1)h})|_\infty^2]}.$$

For (d), it holds that

$$\begin{aligned}
 &\mathbb{E}[|H_k(\bar{Y}_{kh} - Y_{kh})|_\infty^2] \\
 &\leq h \int_{(k-1)h}^{kh} \mathbb{E}[|H_k(\nabla V(Y_t) - \nabla V(Y_{(k-1)h}))|_\infty^2] dt \\
 &\leq h \int_{(k-1)h}^{kh} \int_0^1 \mathbb{E}[|H_k(\nabla^2 V(uY_t + (1-u)Y_{(k-1)h}))(Y_t - Y_{(k-1)h})|_\infty^2] du dt \\
 &\leq 4hs_{r_2}\beta^2 q^2 \int_{(k-1)h}^{kh} \mathbb{E}[|Y_t - Y_{(k-1)h}|_\infty^2] dt = 4hs_{r_2}\beta^2 q^2 \epsilon_{k-1}^2,
 \end{aligned}$$

where we applied (iii) of Proposition C.1 to the the matrix $H_k \nabla^2 V(uY_t + (1-u)Y_{(k-1)h})$. As a summary, we get

$$\begin{aligned}
 (C.22) \quad &\sqrt{\mathbb{E}[|X_{(k+1)h} - Y_{(k+1)h}|_\infty^2]} \\
 &\leq \sqrt{\mathbb{E}[|H_k H_{k-1}(X_{(k-1)h} - Y_{(k-1)h})|_\infty^2]} + 2\beta(\sqrt{hs_{r_1}}\epsilon_k + \sqrt{s_{r_2}h}q\epsilon_{k-1}).
 \end{aligned}$$

Iterating the above arguments N times, and for simplicity of notations writing $k + N$ in place of $k + 1$, we get

$$\begin{aligned}
& \sqrt{\mathbb{E}[|X_{(k+N)h} - Y_{(k+N)h}|_\infty^2]} \\
(C.23) \quad & \leq \sqrt{\mathbb{E}[|H_{k+N-1}H_{k+N-2} \cdots H_k(X_{kh} - Y_{kh})|_\infty^2]} + 2\beta \left(\sum_{i=1}^N q^{i-1} \sqrt{hs_{r_i}\epsilon_{k+N-i}} \right), \\
& \leq 2q^N \sqrt{s_{r_N}} \sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]} + 2\beta \left(\sum_{i=1}^N q^{i-1} \sqrt{hs_{r_i}\epsilon_{k+N-i}} \right),
\end{aligned}$$

where in the last step, we applied (ii) of Proposition C.1 to the matrix $H_{k+r-1}H_{k+r-2} \cdots H_k$. We can now couple the distribution of X_{kh} and Y_{kh} such that $\sqrt{\mathbb{E}[|X_{kh} - Y_{kh}|_\infty^2]} = W_{2,\ell^\infty}(\rho_{kh}, \pi)$. With this and using the definition of the W_{2,ℓ^∞} norm, we get

$$W_{2,\ell^\infty}(\rho_{(k+N)h}, \pi) \leq 2q^N \sqrt{s_{r_N}} W_{2,\ell^\infty}(\rho_{kh}, \pi) + 2\beta \left(\sum_{i=1}^N q^{i-1} \sqrt{hs_{r_i}\epsilon_{k+N-i}} \right).$$

With this expression, we can ensure contraction in the W_{2,ℓ^∞} metric by iterating sufficiently large number of steps N , so that the factor q^N is small enough to offset the increasing sparsity parameter s_{r_N} . In particular, in the dense case where $s_{r_N} = d$, such a contraction can be achieved by taking $N \sim \frac{\log d}{h}$.

We can further bound ϵ_j^2 using the same approach as in (A.7), which implies that

$$(C.24) \quad \epsilon_j \leq h^{3/2} \sqrt{\mathbb{E}_\pi[|\nabla V(Y)|_\infty^2]} + 3h \sqrt{\log(2d)}.$$

The proof is complete. $\square$

C.3. Convergence bounds.

Proof of Theorem 2.2. With the bound on ϵ_j in Proposition C.2, we take the limit $k \rightarrow \infty$ in (C.14) to obtain the bias bound.

Once we have the bound on the bias, we can utilize the convergence bound for W_2 to establish the following:

$$\begin{aligned}
(C.25) \quad W_{2,\ell^\infty}(\rho_{kh}, \pi) & \leq W_{2,\ell^\infty}(\rho_{kh}, \pi_h) + W_{2,\ell^\infty}(\pi_h, \pi) \leq W_2(\rho_{kh}, \pi_h) + W_{2,\ell^\infty}(\pi_h, \pi) \\
& \leq q^k W_2(\rho_0, \pi_h) + \text{bias}(N, q, \beta, h),
\end{aligned}$$

where we used the contraction in W_2 to get $W_2(\rho_{kh}, \pi_h) \leq q^k W_2(\rho_0, \pi_h)$. The proof is complete. $\square$

Proof of Proposition 2.3. By the Bakry-Émery criterion, for distribution π that is α -strongly log concave, π satisfies the log-Sobolev inequality with constant $1/\alpha$. Then, by Herbst's argument [42], we have

$$(C.26) \quad \mathbb{E}_\pi[\exp(\lambda(Y^{(i)} - \mathbb{E}_\pi Y^{(i)}))] \leq \exp\left(\frac{\lambda^2}{2\alpha}\right).$$

Furthermore, since ∇V is β -Lipschitz, Herbst's argument leads to

$$(C.27) \quad \mathbb{E}_\pi[\exp(\lambda(\nabla_i V(Y) - \mathbb{E}_\pi \nabla_i V(Y)))] \leq \exp\left(\frac{\lambda^2 \beta^2}{2\alpha}\right),$$

where $\nabla_i V$ is the i -th component of the vector ∇V . This implies that the distribution of $\nabla_i V(Y)$ has a sub-Gaussian tail. Thus, by Lemma B.1, we get

$$(C.28) \quad \sqrt{\mathbb{E}_\pi[|\nabla V(Y) - \mathbb{E}_\pi[\nabla V(Y)]|_\infty^2]} \leq 2\frac{\beta}{\sqrt{\alpha}}\sqrt{\log(2d)}.$$

Note that $\mathbb{E}_\pi[\nabla V(Y)] = -\int \pi \nabla \log \pi = \int \nabla \pi = 0$. The proof is complete. $\square$

Proof of Proposition 2.4. We will use Theorem 2.2. Recall the definition $r_N = \lceil e^2 N h \beta + \log \sqrt{d} \rceil$ and $q = \exp(-h\alpha)$. There is a universal upper bound on s_i , given by $s_i \leq d$. Choose $N = \lceil \frac{\log(4\sqrt{d})}{h\alpha} \rceil$, which leads to $2q^N \sqrt{s_{r_N}} \leq 1/2$. Based on Theorem 2.2, it remains to calculate the bound on $\sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}}$.

Since $q \leq 1$, we have the bound

$$(C.29) \quad \begin{aligned} \sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}} &\leq \sum_{i=1}^N \sqrt{C}(e^2 i h \beta + \log \sqrt{d} + 2)^{\frac{n}{2}} \\ &\leq \sqrt{C} \int_1^{N+1} (e^2 y h \beta + \log \sqrt{d} + 2)^{\frac{n}{2}} dy \\ &\leq \sqrt{C} \frac{(e^2(N+1)h\beta + \log \sqrt{d} + 2)^{\frac{n}{2}+1}}{(n/2+1)(e^2 h \beta)} \\ &\leq \sqrt{C} \frac{(4e^2 \log(4\sqrt{d})\frac{\beta}{\alpha} + \log \sqrt{d} + 2)^{\frac{n}{2}+1}}{(n/2+1)(e^2 h \beta)}. \end{aligned}$$

Therefore, we get

$$(C.30) \quad \begin{aligned} \frac{\beta}{1 - 2q^N \sqrt{s_{r_N}}} (2 \sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}}) &\leq 4\sqrt{C} \frac{(4e^2 \log(4\sqrt{d})\frac{\beta}{\alpha} + \log \sqrt{d} + 2)^{\frac{n}{2}+1}}{(n/2+1)(e^2 h)} \\ &= \frac{1}{h} \left(O\left(\frac{\beta}{\alpha} \log(2d)\right) \right)^{\frac{n}{2}+1}. \end{aligned}$$

On the other hand, the trivial bound $s_i \leq d$ leads to

$$(C.31) \quad \sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}} \leq \frac{\sqrt{d}}{\alpha h}.$$

Thus, we have another bound

$$(C.32) \quad \frac{\beta}{1 - 2q^N \sqrt{s_{r_N}}} (2 \sum_{i=1}^N q^{i-1} \sqrt{s_{r_i}}) \leq \frac{4\beta\sqrt{d}}{\alpha h}.$$

Moreover, by Proposition 2.3, we have

$$(C.33) \quad h^2 \sqrt{\mathbb{E}_\pi[|\nabla V(Y)|_\infty^2]} + 3h^{3/2} \sqrt{\log(2d)} \lesssim \left(\frac{\beta}{\sqrt{\alpha}} h^2 + h^{3/2} \right) \sqrt{\log(2d)}.$$

Using the fact that $h \leq \frac{\alpha}{\beta^2}$ and combining the above two inequalities leads to the final result. $\square$

APPENDIX D. ASYMPTOTIC BIAS FOR GENERAL OBSERVABLES

In the proof, we use the notation $x = (x_1, \dots, x_d) \in \mathbb{R}^d$.

Proof of Proposition 3.1. Let $\mathcal{L}u = f$. Then, we get $\int f\pi - \int f\pi_h = -\int \mathcal{L}u\pi_h = \int (\mathcal{L}_h u - \mathcal{L}u)\pi_h$.

$$(D.1) \quad \mathcal{L}_h u(x) = \frac{\mathbb{E}[u(\bar{x} + \sqrt{2h}\xi)] - u(\bar{x})}{h},$$

where $\bar{x} = x + h\nabla \log \pi(x)$. Then, by Taylor's expansion, we get

$$(D.2) \quad \begin{aligned} & \mathbb{E}[u(\bar{x} + \sqrt{2h}\xi)] - u(\bar{x}) \\ &= \frac{1}{2} \cdot 2h \cdot \mathbb{E}[\xi^T \nabla^2 u(\bar{x}) \xi] + \sum_{|\alpha|=4} \frac{1}{\alpha!} D^\alpha u(\bar{x}) \mathbb{E}[\xi^\alpha] \cdot (4h^2) + o(h^2) \\ &= h\Delta u(\bar{x}) + h^2 \left(\frac{1}{2} \sum_{i=1}^d D_i^4 u(\bar{x}) + \sum_{1 \leq i < j \leq d} D_i^2 D_j^2 u(\bar{x}) \right) + o(h^2) \\ &= h\Delta u(x) + h^2 (\nabla \Delta u(x)) \cdot \nabla \log \pi(x) + h^2 \left(\frac{1}{2} \sum_{i=1}^d D_i^4 u(x) + \sum_{1 \leq i < j \leq d} D_i^2 D_j^2 u(x) \right) + o(h^2) \\ &= h\Delta u(x) + h^2 (\nabla \Delta u(x)) \cdot \nabla \log \pi(x) + \frac{1}{2} h^2 \Delta^2 u + o(h^2), \end{aligned}$$

where in the first identity, we used the notation that $\alpha = (\alpha_1, \dots, \alpha_d)$ is a multi-index with non-negative entries and $|\alpha| = 4$ implies that $\sum_i \alpha_i = 4$. Moreover $\alpha! = \alpha_1! \dots \alpha_d!$ and $\xi^\alpha = \xi_1^{\alpha_1} \dots \xi_d^{\alpha_d}$. In the second identity, we used the notation $D_i^k u = \frac{\partial^k}{\partial x_i^k} u$ and the fact that $\mathbb{E}[\xi_i^2] = 1$ and $\mathbb{E}[\xi_i^4] = 3$ for $\xi_i \sim \mathcal{N}(0, 1)$. In the third identity, we performed the Taylor expansion at x , based on the fact $\bar{x} = x + h\nabla \log \pi(x)$. In the last identity, we noticed the fact that the terms in the big bracket equals $\Delta^2 u$.

Moreover, $u(\bar{x}) - u(x) = h\nabla u(x) \cdot \nabla \log \pi(x) + \frac{1}{2} h^2 (\nabla \log \pi(x))^T \nabla^2 u(x) \nabla \log \pi(x) + o(h^2)$. Therefore, we get

$$(D.3) \quad \begin{aligned} \mathcal{L}_h u(x) - \mathcal{L}u(x) &= h(\nabla \Delta u(x)) \cdot \nabla \log \pi(x) + \frac{1}{2} h(\nabla \log \pi(x))^T \nabla^2 u(x) \nabla \log \pi(x) \\ &\quad + \frac{1}{2} h \Delta^2 u(x) + o(h), \end{aligned}$$

where we have used the definition $\mathcal{L}u = \nabla \log \pi \cdot \nabla u + \Delta u$.

We note that

$$(D.4) \quad \int \pi \nabla \Delta u \cdot \nabla \log \pi = \int \nabla \Delta u \cdot \nabla \pi = - \int \pi \Delta^2 u.$$

Therefore, we get

$$(D.5) \quad \int \pi (\mathcal{L}_h u - \mathcal{L}u) = \frac{1}{2} h \int \pi \left(\nabla \Delta u \cdot \nabla \log \pi + (\nabla \log \pi)^T (\nabla^2 u) \nabla \log \pi \right) + o(h).$$

Now, note that,

$$(D.6) \quad \begin{aligned} \nabla f &= \nabla \mathcal{L}u = \nabla(\nabla \log \pi \cdot \nabla u + \Delta u) \\ &= (\nabla^2 \log \pi) \nabla u + (\nabla^2 u) \nabla \log \pi + \nabla \Delta u, \end{aligned}$$

we get

$$(D.7) \quad \int \pi(\mathcal{L}_h u - \mathcal{L}u) = \frac{1}{2}h \int \pi \left(\nabla f \cdot \nabla \log \pi - (\nabla \log \pi)^T (\nabla^2 \log \pi) \nabla u \right) + o(h).$$

Let $g = \frac{1}{2}|\nabla \log \pi|_2^2$ which satisfies the equation $\nabla g = (\nabla^2 \log \pi) \nabla \log \pi$. We know that the adjoint of the generator satisfies $\mathcal{L}^*(g\pi) = \nabla \cdot (\pi \nabla g)$. Thus

$$(D.8) \quad \begin{aligned} &\int \pi(\nabla \log \pi)^T (\nabla^2 \log \pi) \nabla u \\ &= - \int \nabla \cdot (\pi \nabla g) u \\ &= - \int (\mathcal{L}^*(g\pi)) u \\ &= - \int g f \pi. \end{aligned}$$

Therefore,

$$(D.9) \quad \begin{aligned} \int f \pi - \int f \pi_h &= \int \pi_h(\mathcal{L}_h u - \mathcal{L}u) = \int \pi(\mathcal{L}_h u - \mathcal{L}u) + o(h) \\ &= \frac{1}{2}h \int \pi \left(\nabla f \cdot \nabla \log \pi + \frac{1}{2}|\nabla \log \pi|_2^2 f \right) \\ &= \frac{1}{4}h \int (-2\Delta f + |\nabla \log \pi|_2^2 f) \pi + o(h), \end{aligned}$$

where in the last identity, we used the fact $\pi \nabla \log \pi = \nabla \pi$ and integration by parts.

Moreover, note that $\nabla \pi = \pi \nabla \log \pi$ and thus

$$\Delta \pi = \nabla \pi \cdot \nabla \log \pi + \pi \Delta \log \pi = \pi |\nabla \log \pi|_2^2 + \pi \Delta \log \pi,$$

which implies that

$$\int |\nabla \log \pi|_2^2 f \pi = \int f \Delta \pi - f \pi \Delta \log \pi = \int \pi \Delta f - f \pi \Delta \log \pi.$$

Therefore, we also have another representation of the bias

$$(D.10) \quad \int f \pi - \int f \pi_h = -\frac{1}{4}h \int (\Delta f + f \Delta \log \pi) \pi + o(h).$$

□

Proof of Proposition 3.2. We have

$$(D.11) \quad \left| \int f(\Delta \log \pi) \pi \right| = \left| \int f(\Delta U) \pi \right| \leq \sqrt{\int (f)^2 \pi} \sqrt{\int (\Delta U)^2 \pi} \leq \sqrt{C_0} \sqrt{\int (f)^2 \pi}.$$

When f only depends on K number of coordinates of $x \in \mathbb{R}^d$, then the integral $\int (f)^2 \pi$ will only scale with that number, rather than d , because the whole integral will only rely on the marginal distribution of π at these coordinates. As such argument also applies

to the term $-\frac{1}{4} \int (\Delta f) \pi$, we get that the first order term of the bias depend only on K rather than d .

In particular, for $f(x) = \sum_{k=1}^K x^{(k)}$, we have $\Delta f = 0$. Then, using the assumed Poincaré inequality, we get $\sqrt{\int (f)^2 \pi} \leq \sqrt{C_{\text{PI}} \int |\nabla f|^2 \pi} = O(\sqrt{K})$ where C_{PI} is the Poincaré constant. Thus, in such case,

$$(D.12) \quad \left| \int f(\Delta \log \pi) \pi \right| = O(\sqrt{K}h) + o(h).$$

The proof is complete. □