

MEASURING SAMPLE QUALITY WITH STEIN'S METHOD

By

Jackson Gorham
Lester Mackey

Technical Report No. 2015-19
September 2015

Department of Statistics
STANFORD UNIVERSITY
Stanford, California 94305-4065



MEASURING SAMPLE QUALITY WITH STEIN'S METHOD

By

Jackson Gorham
Lester Mackey
Stanford University

Technical Report No. 2015-19
September 2015

**This research was supported by the Frederick E. Terman Fellowship
and a National Science Foundation Graduate Research Fellowship
under grant DGE-114747.**

Department of Statistics
STANFORD UNIVERSITY
Stanford, California 94305-4065

<http://statistics.stanford.edu>

Measuring Sample Quality with Stein’s Method

Jackson Gorham
Department of Statistics
Stanford University

Lester Mackey
Department of Statistics
Stanford University

Abstract

To improve the efficiency of Monte Carlo estimation, practitioners are turning to biased Markov chain Monte Carlo procedures that trade off asymptotic exactness for computational speed. The reasoning is sound: a reduction in variance due to more rapid sampling can outweigh the bias introduced. However, the inexactness creates new challenges for sampler and parameter selection, since standard measures of sample quality like effective sample size do not account for asymptotic bias. To address these challenges, we introduce a new computable quality measure based on Stein’s method that bounds the discrepancy between sample and target expectations over a large class of test functions. We use our tool to compare exact, biased, and deterministic sample sequences and illustrate applications to hyperparameter selection, convergence rate assessment, and quantifying bias-variance tradeoffs in posterior inference.

1 Introduction

When faced with a complex target distribution, one often turns to Markov chain Monte Carlo (MCMC) [1] to approximate intractable expectations $\mathbb{E}_P[h(Z)] = \int_{\mathcal{X}} p(x)h(x)dx$ with asymptotically exact sample estimates $\mathbb{E}_Q[h(X)] = \sum_{i=1}^n q(x_i)h(x_i)$. These complex targets commonly arise as posterior distributions in Bayesian inference and as candidate distributions in maximum likelihood estimation [2]. In recent years, researchers [e.g., 3, 4, 5] have introduced asymptotic bias into MCMC procedures to trade off asymptotic correctness for improved sampling speed. The rationale is that more rapid sampling can reduce the variance of a Monte Carlo estimate and hence outweigh the bias introduced. However, the added flexibility introduces new challenges for sampler and parameter selection, since standard sample quality measures, like effective sample size, asymptotic variance, trace and mean plots, and pooled and within-chain variance diagnostics, presume eventual convergence to the target [1] and hence do not account for asymptotic bias.

To address this shortcoming, we develop a new measure of sample quality suitable for comparing asymptotically exact, asymptotically biased, and even deterministic sample sequences. The quality measure is based on Stein’s method and is attainable by solving a linear program. After outlining our design criteria in Section 2, we relate the convergence of the quality measure to that of standard probability metrics in Section 3, develop a streamlined implementation based on geometric spanners in Section 4, and illustrate applications to hyperparameter selection, convergence rate assessment, and the quantification of bias-variance tradeoffs in posterior inference in Section 5. We discuss related work in Section 6 and defer all proofs to the appendix.

Notation We denote the ℓ_2 , ℓ_1 , and ℓ_∞ norms on \mathbb{R}^d by $\|\cdot\|_2$, $\|\cdot\|_1$, and $\|\cdot\|_\infty$ respectively. We will often refer to a generic norm $\|\cdot\|$ on \mathbb{R}^d with associated dual norms $\|w\|^* \triangleq \sup_{v \in \mathbb{R}^d: \|v\|=1} \langle w, v \rangle$ for vectors $w \in \mathbb{R}^d$, $\|M\|^* \triangleq \sup_{v \in \mathbb{R}^d: \|v\|=1} \|Mv\|^*$ for matrices $M \in \mathbb{R}^{d \times d}$, and $\|T\|^* \triangleq \sup_{v \in \mathbb{R}^d: \|v\|=1} \|T[v]\|^*$ for tensors $T \in \mathbb{R}^{d \times d \times d}$. We denote the j -th standard basis vector by e_j , the partial derivative $\frac{\partial}{\partial x_k}$ by ∇_k , and the gradient of any \mathbb{R}^d -valued function g by ∇g with components $(\nabla g(x))_{jk} \triangleq \nabla_k g_j(x)$.

2 Quality Measures for Samples

Consider a target distribution P with open convex support $\mathcal{X} \subseteq \mathbb{R}^d$ and continuously differentiable density p . We assume that p is known up to its normalizing constant and that exact integration under P is intractable for most functions of interest. We will approximate expectations under P with the aid of a *weighted sample*, a collection of distinct sample points $x_1, \dots, x_n \in \mathcal{X}$ with weights $q(x_i)$ encoded in a probability mass function q . The probability mass function q induces a discrete distribution Q and an approximation $\mathbb{E}_Q[h(X)] = \sum_{i=1}^n q(x_i)h(x_i)$ for any target expectation $\mathbb{E}_P[h(Z)]$. We make no assumption about the provenance of the sample points; they may arise as random draws from a Markov chain or even be deterministically selected.

Our goal is to compare the fidelity of different samples approximating a common target distribution. That is, we seek to quantify the discrepancy between \mathbb{E}_Q and \mathbb{E}_P in a manner that (i) detects when a sequence of samples is converging to the target, (ii) detects when a sequence of samples is not converging to the target, and (iii) is computationally feasible. A natural starting point is to consider the maximum deviation between sample and target expectations over a class of real-valued test functions \mathcal{H} ,

$$d_{\mathcal{H}}(Q, P) = \sup_{h \in \mathcal{H}} |\mathbb{E}_Q[h(X)] - \mathbb{E}_P[h(Z)]|. \quad (1)$$

When the class of test functions is sufficiently large, the convergence of $d_{\mathcal{H}}(Q_m, P)$ to zero implies that the sequence of sample measures $(Q_m)_{m \geq 1}$ converges weakly to P . In this case, the expression (1) is termed an *integral probability metric* (IPM) [6]. By varying the class of test functions \mathcal{H} , we can recover many well-known probability metrics as IPMs, including the *total variation distance*, generated by $\mathcal{H} = \{h : \mathcal{X} \rightarrow \mathbb{R} \mid \sup_{x \in \mathcal{X}} |h(x)| \leq 1\}$, and the *Wasserstein distance* (also known as the Kantorovich-Rubenstein or earth mover's distance), $d_{\mathcal{W}_{\|\cdot\|}}$, generated by

$$\mathcal{H} = \mathcal{W}_{\|\cdot\|} \triangleq \{h : \mathcal{X} \rightarrow \mathbb{R} \mid \sup_{x \neq y \in \mathcal{X}} \frac{|h(x) - h(y)|}{\|x - y\|} \leq 1\}.$$

The primary impediment to adopting an IPM as a sample quality measure is that exact computation is typically infeasible when generic integration under P is intractable. However, we could skirt this intractability by focusing on classes of test functions with known expectation under P . For example, if we consider only test functions h for which $\mathbb{E}_P[h(Z)] = 0$, then the IPM value $d_{\mathcal{H}}(Q, P)$ is the solution of an optimization problem depending on Q alone. This, at a high level, is our strategy, but many questions remain. How do we select the class of test functions h ? How do we know that the resulting IPM will track convergence and non-convergence of a sample sequence (Desiderata (i) and (ii))? How do we solve the resulting optimization problem in practice (Desideratum (iii))? To address the first two of these questions, we draw upon tools from Charles Stein's method of characterizing distributional convergence. We return to the third question in Section 4.

3 Stein's Method

Stein's method [7] for characterizing convergence in distribution classically proceeds in three steps:

1. Identify a real-valued operator \mathcal{T} acting on a set \mathcal{G} of \mathbb{R}^d -valued¹ functions of \mathcal{X} for which

$$\mathbb{E}_P[(\mathcal{T}g)(Z)] = 0 \quad \text{for all } g \in \mathcal{G}. \quad (2)$$

Together, \mathcal{T} and \mathcal{G} define the *Stein discrepancy*,

$$\mathcal{S}(Q, \mathcal{T}, \mathcal{G}) \triangleq \sup_{g \in \mathcal{G}} |\mathbb{E}_Q[(\mathcal{T}g)(X)]| = \sup_{g \in \mathcal{G}} |\mathbb{E}_Q[(\mathcal{T}g)(X)] - \mathbb{E}_P[(\mathcal{T}g)(Z)]| = d_{\mathcal{T}\mathcal{G}}(Q, P),$$

an IPM-type quality measure with no explicit integration under P .

2. Lower bound the Stein discrepancy by a familiar convergence-determining IPM $d_{\mathcal{H}}$. This step can be performed once, in advance, for large classes of target distributions and ensures that, for any sequence of probability measures $(\mu_m)_{m \geq 1}$, $\mathcal{S}(\mu_m, \mathcal{T}, \mathcal{G})$ converges to zero only if $d_{\mathcal{H}}(\mu_m, P)$ does (Desideratum (ii)).

¹One commonly considers real-valued functions g when applying Stein's method, but we will find it more convenient to work with vector-valued g .

3. Upper bound the Stein discrepancy by any means necessary to demonstrate convergence to zero under suitable conditions (Desideratum (i)). In our case, the universal bound established in Section 3.3 will suffice.

While Stein’s method is typically employed as an analytical tool, we view the Stein discrepancy as a promising candidate for a practical sample quality measure. Indeed, in Section 4, we will adopt an optimization perspective and develop efficient procedures to compute the Stein discrepancy for any sample measure Q and appropriate choices of \mathcal{T} and \mathcal{G} . First, we assess the convergence properties of an equivalent Stein discrepancy in the subsections to follow.

3.1 Identifying a Stein Operator

The *generator method* of Barbour [8] provides a convenient and general means of constructing operators \mathcal{T} which produce mean-zero functions under P (2). Let $(Z_t)_{t \geq 0}$ represent a Markov process with unique stationary distribution P . Then the *infinitesimal generator* \mathcal{A} of $(Z_t)_{t \geq 0}$, defined by

$$(\mathcal{A}u)(x) = \lim_{t \rightarrow 0} (\mathbb{E}[u(Z_t) \mid Z_0 = x] - u(x))/t \quad \text{for } u : \mathbb{R}^d \rightarrow \mathbb{R},$$

satisfies $\mathbb{E}_P[(\mathcal{A}u)(Z)] = 0$ under mild conditions on \mathcal{A} and u . Hence, a candidate operator \mathcal{T} can be constructed from any infinitesimal generator.

For example, the *overdamped Langevin diffusion*, defined by the stochastic differential equation $dZ_t = \frac{1}{2} \nabla \log p(Z_t) dt + dW_t$ for $(W_t)_{t \geq 0}$ a Wiener process, gives rise to the generator

$$(\mathcal{A}_P u)(x) = \frac{1}{2} \langle \nabla u(x), \nabla \log p(x) \rangle + \frac{1}{2} \langle \nabla, \nabla u(x) \rangle. \quad (3)$$

After substituting g for $\frac{1}{2} \nabla u$, we obtain the associated *Stein operator*

$$(\mathcal{T}_P g)(x) \triangleq \langle g(x), \nabla \log p(x) \rangle + \langle \nabla, g(x) \rangle. \quad (4)$$

The Stein operator \mathcal{T}_P is particularly well-suited to our setting as it depends on P only through the derivative of its log density and hence is computable even when the normalizing constant of p is not.

If we let $\partial \mathcal{X}$ denote the boundary of \mathcal{X} (an empty set when $\mathcal{X} = \mathbb{R}^d$) and $n(x)$ represent the outward unit normal vector to the boundary at x , then we may define the *classical Stein set*

$$\mathcal{G}_{\|\cdot\|} \triangleq \left\{ g : \mathcal{X} \rightarrow \mathbb{R}^d \mid \sup_{x \neq y \in \mathcal{X}} \max \left(\|g(x)\|^*, \|\nabla g(x)\|^*, \frac{\|\nabla g(x) - \nabla g(y)\|^*}{\|x - y\|} \right) \leq 1 \quad \text{and} \right. \\ \left. \langle g(x), n(x) \rangle = 0, \forall x \in \partial \mathcal{X} \text{ with } n(x) \text{ defined} \right\}$$

of sufficiently smooth functions satisfying a Neumann-type boundary condition. The following proposition – a consequence of integration by parts – shows that $\mathcal{G}_{\|\cdot\|}$ is a suitable domain for \mathcal{T}_P .

Proposition 1. *If $\mathbb{E}_P[\|\nabla \log p(Z)\|] < \infty$, then $\mathbb{E}_P[(\mathcal{T}_P g)(Z)] = 0$ for all $g \in \mathcal{G}_{\|\cdot\|}$.*

Together, \mathcal{T}_P and $\mathcal{G}_{\|\cdot\|}$ form the *classical Stein discrepancy* $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|})$, our chief object of study.

3.2 Lower Bounding the Classical Stein Discrepancy

In the univariate setting ($d = 1$), it is known for a wide variety of targets P that the classical Stein discrepancy $\mathcal{S}(\mu_m, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|})$ converges to zero only if the Wasserstein distance $d_{\mathcal{W}_{\|\cdot\|}}(\mu_m, P)$ does [9, 10]. In the multivariate setting, analogous statements are available for multivariate Gaussian targets [11, 12, 13], but few other target distributions have been analyzed. To extend the reach of the multivariate literature, we show in Theorem 2 that the classical Stein discrepancy also determines Wasserstein convergence for a large class of strongly log-concave densities, including the Bayesian logistic regression posterior under Gaussian priors.

Theorem 2 (Stein Discrepancy Lower Bound for Strongly Log-concave Densities). *If $\mathcal{X} = \mathbb{R}^d$, and $\log p$ is strongly concave with third and fourth derivatives bounded and continuous, then, for any probability measures $(\mu_m)_{m \geq 1}$, $\mathcal{S}(\mu_m, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}) \rightarrow 0$ only if $d_{\mathcal{W}_{\|\cdot\|}}(\mu_m, P) \rightarrow 0$.*

We emphasize that the sufficient conditions in Theorem 2 are certainly not necessary for lower bounding the classical Stein discrepancy. We hope that the theorem and its proof will provide a template for lower bounding $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|})$ for other large classes of multivariate target distributions.

3.3 Upper Bounding the Classical Stein Discrepancy

We next establish sufficient conditions for the convergence of the classical Stein discrepancy to zero.

Proposition 3 (Stein Discrepancy Upper Bound). *If $X \sim Q$ and $Z \sim P$ with $\nabla \log p(Z)$ integrable,*

$$\begin{aligned} \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}) &\leq \mathbb{E}[\|X - Z\|] + \mathbb{E}[\|\nabla \log p(X) - \nabla \log p(Z)\|] + \mathbb{E}[\|\nabla \log p(Z)(X - Z)^\top\|] \\ &\leq \mathbb{E}[\|X - Z\|] + \mathbb{E}[\|\nabla \log p(X) - \nabla \log p(Z)\|] + \sqrt{\mathbb{E}[\|\nabla \log p(Z)\|^2] \mathbb{E}[\|X - Z\|^2]}. \end{aligned}$$

One implication of Proposition 3 is that $\mathcal{S}(Q_m, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|})$ converges to zero whenever $X_m \sim Q_m$ converges in mean-square to $Z \sim P$ and $\nabla \log p(X_m)$ converges in mean to $\nabla \log p(Z)$.

3.4 Extension to Non-uniform Stein Sets

The analyses and algorithms in this paper readily accommodate non-uniform Stein sets of the form

$$\mathcal{G}_{\|\cdot\|}^{c_1, c_2, c_3} \triangleq \left\{ g : \mathcal{X} \rightarrow \mathbb{R}^d \mid \sup_{x \neq y \in \mathcal{X}} \max \left(\frac{\|g(x)\|^*}{c_1}, \frac{\|\nabla g(x)\|^*}{c_2}, \frac{\|\nabla g(x) - \nabla g(y)\|^*}{c_3 \|x - y\|} \right) \leq 1 \text{ and } \langle g(x), n(x) \rangle = 0, \forall x \in \partial \mathcal{X} \text{ with } n(x) \text{ defined} \right\} \quad (5)$$

for constants $c_1, c_2, c_3 > 0$ known as *Stein factors* in the literature. We will exploit this additional flexibility in Section 5.2 to establish tight lower-bounding relations between the Stein discrepancy and Wasserstein distance for well-studied target distributions. For general use, however, we advocate the parameter-free classical Stein set and graph Stein sets to be introduced in the sequel. Indeed, any non-uniform Stein discrepancy is equivalent to the classical Stein discrepancy in a strong sense:

Proposition 4 (Equivalence of Non-uniform Stein Discrepancies). *For any $c_1, c_2, c_3 > 0$,*

$$\min(c_1, c_2, c_3) \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}) \leq \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}^{c_1, c_2, c_3}) \leq \max(c_1, c_2, c_3) \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}).$$

4 Computing Stein Discrepancies

In this section, we introduce an efficiently computable Stein discrepancy with convergence properties equivalent to those of the classical discrepancy. We restrict attention to the unconstrained domain $\mathcal{X} = \mathbb{R}^d$ in Sections 4.1-4.3 and present extensions for constrained domains in Section 4.4.

4.1 Graph Stein Discrepancies

Evaluating a Stein discrepancy $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G})$ for a fixed (Q, P) pair reduces to solving an optimization program over functions $g \in \mathcal{G}$. For example, the classical Stein discrepancy is the optimum

$$\begin{aligned} \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}) &= \sup_g \sum_{i=1}^n q(x_i) (\langle g(x_i), \nabla \log p(x_i) \rangle + \langle \nabla, g(x_i) \rangle) \\ &\text{s.t. } \|g(x)\|^* \leq 1, \|\nabla g(x)\|^* \leq 1, \|\nabla g(x) - \nabla g(y)\|^* \leq \|x - y\|, \forall x, y \in \mathcal{X}. \end{aligned} \quad (6)$$

Note that the objective associated with any Stein discrepancy $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G})$ is linear in g and, since Q is discrete, only depends on g and ∇g through their values at each of the n sample points x_i . The primary difficulty in solving the classical Stein program (6) stems from the infinitude of constraints imposed by the classical Stein set $\mathcal{G}_{\|\cdot\|}$. One way to avoid this difficulty is to impose the classical smoothness constraints at only a finite collection of points. To this end, for each finite graph $G = (V, E)$ with vertices $V \subset \mathcal{X}$ and edges $E \subset V^2$, we define the *graph Stein set*,

$$\begin{aligned} \mathcal{G}_{\|\cdot\|, Q, G} &\triangleq \left\{ g : \mathcal{X} \rightarrow \mathbb{R}^d \mid \forall x \in V, \max(\|g(x)\|^*, \|\nabla g(x)\|^*) \leq 1 \text{ and, } \forall (x, y) \in E, \right. \\ &\left. \max \left(\frac{\|g(x) - g(y)\|^*}{\|x - y\|}, \frac{\|\nabla g(x) - \nabla g(y)\|^*}{\|x - y\|}, \frac{\|g(x) - g(y) - \nabla g(x)(x - y)\|^*}{\frac{1}{2}\|x - y\|^2}, \frac{\|g(x) - g(y) - \nabla g(y)(x - y)\|^*}{\frac{1}{2}\|x - y\|^2} \right) \leq 1 \right\}, \end{aligned}$$

the family of functions which satisfy the classical constraints and certain implied Taylor compatibility constraints at pairs of points in E . Remarkably, if the graph G_1 consists of edges between all distinct sample points x_i , then the associated *complete graph Stein discrepancy* $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1})$ is equivalent to the classical Stein discrepancy in the following strong sense.

Proposition 5 (Equivalence of Classical and Complete Graph Stein Discrepancies). *If $\mathcal{X} = \mathbb{R}^d$, and $G_1 = (\text{supp}(Q), E_1)$ with $E_1 = \{(x_i, x_l) \in \text{supp}(Q)^2 : x_i \neq x_l\}$, then*

$$\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}) \leq \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1}) \leq \kappa_d \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}),$$

where κ_d is a constant, independent of (Q, P) , depending only on the dimension d and norm $\|\cdot\|$.

Proposition 5 follows from the Whitney-Glaeser extension theorem for smooth functions [14, 15] and implies that the complete graph Stein discrepancy inherits all of the desirable convergence properties of the classical discrepancy. However, the complete graph also introduces order n^2 constraints, rendering computation infeasible for large samples. To achieve the same form of equivalence while enforcing only $O(n)$ constraints, we will make use of sparse *geometric spanner* subgraphs.

4.2 Geometric Spanners

For a given dilation factor $t \geq 1$, a t -spanner [16, 17] is a graph $G = (V, E)$ with weight $\|x - y\|$ on each edge $(x, y) \in E$ and a path between each pair $x' \neq y' \in V$ with total weight no larger than $t\|x' - y'\|$. The next proposition shows that *spanner Stein discrepancies* enjoy the same convergence properties as the complete graph Stein discrepancy.

Proposition 6 (Equivalence of Spanner and Complete Graph Stein Discrepancies). *If $\mathcal{X} = \mathbb{R}^d$, $G_t = (\text{supp}(Q), E)$ is a t -spanner, and $G_1 = (\text{supp}(Q), \{(x_i, x_l) \in \text{supp}(Q)^2 : x_i \neq x_l\})$, then*

$$\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_t}) \leq \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1}) \leq 2t^2 \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1}).$$

Moreover, for any ℓ_p norm, a 2-spanner with $O(\kappa_d n)$ edges can be computed in $O(\kappa_d n \log(n))$ expected time for κ_d a constant depending only on d and $\|\cdot\|$ [18]. As a result, we will adopt a 2-spanner Stein discrepancy, $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_2})$, as our standard quality measure.

4.3 Decoupled Linear Programs

The final unspecified component of our Stein discrepancy is the choice of norm $\|\cdot\|$. We recommend the ℓ_1 norm, as the resulting optimization problem decouples into d independent finite-dimensional linear programs (LPs) that can be solved in parallel. More precisely, $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, (V, E)})$ equals

$$\begin{aligned} & \sum_{j=1}^d \sup_{\gamma_j \in \mathbb{R}^{|V|}, \Gamma_j \in \mathbb{R}^{d \times |V|}} \sum_{i=1}^{|V|} q(v_i) (\gamma_{ji} \nabla_j \log p(v_i) + \Gamma_{jji}) \\ & \text{s.t. } \|\gamma_j\|_\infty \leq 1, \|\Gamma_j\|_\infty \leq 1, \text{ and } \forall i \neq l : (v_i, v_l) \in E, \\ & \max \left(\frac{|\gamma_{ji} - \gamma_{jl}|}{\|v_i - v_l\|_1}, \frac{\|\Gamma_j(e_i - e_l)\|_\infty}{\|v_i - v_l\|_1}, \frac{|\gamma_{ji} - \gamma_{jl} - \langle \Gamma_j e_i, v_i - v_l \rangle|}{\frac{1}{2} \|v_i - v_l\|_1^2}, \frac{|\gamma_{ji} - \gamma_{jl} - \langle \Gamma_j e_l, v_i - v_l \rangle|}{\frac{1}{2} \|v_i - v_l\|_1^2} \right) \leq 1. \end{aligned} \quad (7)$$

We have arbitrarily numbered the elements v_i of the vertex set V so that γ_{ji} represents the function value $g_j(v_i)$, and Γ_{jki} represents the gradient value $\nabla_k g_j(v_i)$.

4.4 Constrained Domains

A small modification to the unconstrained formulation (7) extends our tractable Stein discrepancy computation to any domain defined by coordinate boundary constraints, that is, to $\mathcal{X} = (\alpha_1, \beta_1) \times \dots \times (\alpha_d, \beta_d)$ with $-\infty \leq \alpha_j < \beta_j \leq \infty$ for all j . Specifically, for each dimension j , we augment the j -th coordinate linear program of (7) with the boundary compatibility constraints

$$\max \left(\frac{|\gamma_{ji}|}{|v_{ij} - b_j|}, \frac{|\Gamma_{jki}|}{|v_{ij} - b_j|}, \frac{|\gamma_{ji} - \Gamma_{jji}(v_{ij} - b_j)|}{\frac{1}{2}(v_{ij} - b_j)^2} \right) \leq 1, \text{ for each } i, b_j \in \{\alpha_j, \beta_j\} \cap \mathbb{R}, \text{ and } k \neq j. \quad (8)$$

These additional constraints ensure that our candidate function and gradient values can be extended to a smooth function satisfying the boundary conditions $\langle g(z), n(z) \rangle = 0$ on $\partial \mathcal{X}$. Proposition 15 in the appendix shows that the spanner Stein discrepancy so computed is strongly equivalent to the classical Stein discrepancy on \mathcal{X} .

Algorithm 1 summarizes the complete solution for computing our recommended, parameter-free spanner Stein discrepancy in the multivariate setting. Notably, the spanner step is unnecessary in the univariate setting, as the complete graph Stein discrepancy $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1})$ can be computed directly by sorting the sample and boundary points and only enforcing constraints between consecutive points in this ordering. Thus, the complete graph Stein discrepancy is our recommended quality measure when $d = 1$, and a recipe for its computation is given in Algorithm 2.

Algorithm 1 Multivariate Spanner Stein Discrepancy

input: Q , coordinate bounds $(\alpha_1, \beta_1), \dots, (\alpha_d, \beta_d)$ with $-\infty \leq \alpha_j < \beta_j \leq \infty$ for all j
 $G_2 \leftarrow$ Compute sparse 2-spanner of $\text{supp}(Q)$
for $j = 1$ **to** d **do (in parallel)**
 $r_j \leftarrow$ Solve j -th coordinate linear program (7) with graph G_2 and boundary constraints (8)
return $\sum_{j=1}^d r_j$

Algorithm 2 Univariate Complete Graph Stein Discrepancy

input: Q , bounds (α, β) with $-\infty \leq \alpha < \beta \leq \infty$
 $(x_{(1)}, \dots, x_{(n')}) \leftarrow \text{SORT}(\{x_1, \dots, x_n, \alpha, \beta\} \cap \mathbb{R})$
return $\sup_{\gamma \in \mathbb{R}^{n'}, \Gamma \in \mathbb{R}^{n'}} \sum_{i=1}^{n'} q(x_{(i)}) (\gamma_i \frac{d}{dx} \log p(x_{(i)}) + \Gamma_i)$
s.t. $\|\Gamma\|_\infty \leq 1, \forall i \leq n', |\gamma_i| \leq \mathbb{I}[\alpha < x_{(i)} < \beta]$, and, $\forall i < n'$,
 $\max\left(\frac{|\gamma_i - \gamma_{i+1}|}{x_{(i+1)} - x_{(i)}}, \frac{|\Gamma_i - \Gamma_{i+1}|}{x_{(i+1)} - x_{(i)}}, \frac{|\gamma_i - \gamma_{i+1} - \Gamma_i(x_{(i)} - x_{(i+1)})|}{\frac{1}{2}(x_{(i+1)} - x_{(i)})^2}, \frac{|\gamma_i - \gamma_{i+1} - \Gamma_{i+1}(x_{(i)} - x_{(i+1)})|}{\frac{1}{2}(x_{(i+1)} - x_{(i)})^2}\right) \leq 1$

5 Experiments

We now turn to an empirical evaluation of our proposed quality measures. We compute all spanners using the efficient C++ greedy spanner implementation of Bouts et al. [19] and solve all optimization programs using Julia for Mathematical Programming [20] with the default Gurobi 6.0.4 solver [21]. All reported timings are obtained using a single core of an Intel Xeon CPU E5-2650 v2 @ 2.60GHz.

5.1 A Simple Example

We begin with a simple example to illuminate a few properties of the Stein diagnostic. For the target $P = \mathcal{N}(0, 1)$, we generate a sequence of sample points i.i.d. from the target and a second sequence i.i.d. from a scaled Student's t distribution with matching variance and 10 degrees of freedom. The left panel of Figure 1 shows that the complete graph Stein discrepancy applied to the first n Gaussian sample points decays to zero at an $n^{-0.52}$ rate, while the discrepancy applied to the scaled Student's t sample remains bounded away from zero. The middle panel displays optimal Stein functions g recovered by the Stein program for different sample sizes. Each g yields a test function $h \triangleq \mathcal{T}_P g$, featured in the right panel, that best discriminates the sample Q from the target P . Notably, the Student's t test functions exhibit relatively large magnitude values in the tails of the support.

5.2 Comparing Discrepancies

We show in Theorem 14 in the appendix that, when $d = 1$, the classical Stein discrepancy is the optimum of a convex quadratically constrained quadratic program with a linear objective, $O(n)$ variables, and $O(n)$ constraints. This offers the opportunity to directly compare the behavior of the graph and classical Stein discrepancies. We will also compare to the Wasserstein distance $d_{\mathcal{W}_{\|\cdot\|}}$,

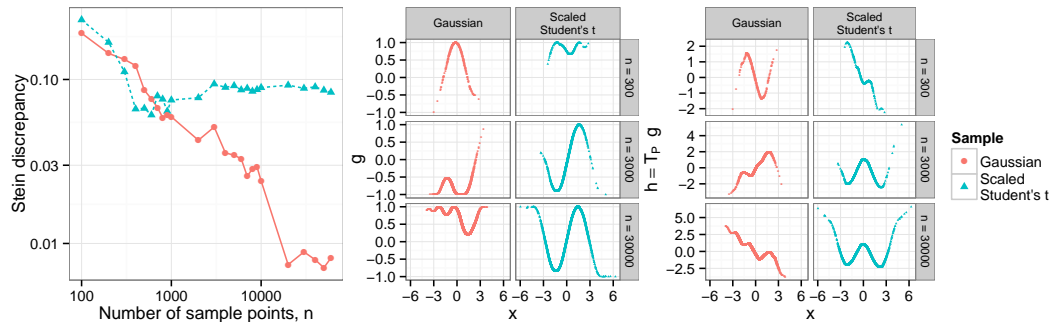


Figure 1: Left: Complete graph Stein discrepancy for a $\mathcal{N}(0, 1)$ target. Middle / right: Optimal Stein functions g and discriminating test functions $h = \mathcal{T}_P g$ recovered by the Stein program.




Figure 2: Comparison of discrepancy measures for sample sequences drawn i.i.d. from their targets.

which is computable for simple univariate target distributions [22] and provably lower bounds the non-uniform Stein discrepancies (5) with $c_{1:3} = (0.5, 0.5, 1)$ for $P = \text{Unif}(0, 1)$ and $c_{1:3} = (1, 4, 2)$ for $P = \mathcal{N}(0, 1)$ [9, 23]. For $\mathcal{N}(0, 1)$ and $\text{Unif}(0, 1)$ targets and several random number generator seeds, we generate a sequence of sample points i.i.d. from the target distribution and plot the non-uniform classical and complete graph Stein discrepancies and the Wasserstein distance as functions of the first n sample points in Figure 2. Two apparent trends are that the graph Stein discrepancy very closely approximates the classical and that both Stein discrepancies track the fluctuations in Wasserstein distance even when a magnitude separation exists. In the $\text{Unif}(0, 1)$ case, the Wasserstein distance in fact equals the classical Stein discrepancy because $\mathcal{T}_P g = g'$ is a Lipschitz function.


5.3 Selecting Sampler Hyperparameters

Stochastic Gradient Langevin Dynamics (SGLD) [3] with constant step size ϵ is a biased MCMC procedure designed for scalable inference. It approximates the overdamped Langevin diffusion, but, because no Metropolis-Hastings (MH) correction is used, the stationary distribution of SGLD deviates increasingly from its target as ϵ grows. If ϵ is too small, however, SGLD explores the sample space too slowly. Hence, an appropriate choice of ϵ is critical for accurate posterior inference. To illustrate the value of the Stein diagnostic for this task, we adopt the bimodal Gaussian mixture model (GMM) posterior of [3] as our target. For a range of step sizes ϵ , we use SGLD with minibatch size 5 to draw 50 independent sequences of length $n = 1000$, and we select the value of ϵ with the highest median quality – either the maximum effective sample size (ESS, a standard diagnostic based on autocorrelation [1]) or the minimum spanner Stein discrepancy – across these sequences. The average discrepancy computation consumes 0.4s for spanner construction and 1.4s per coordinate linear program. As seen in Figure 3a, ESS, which does not detect distributional bias, selects the largest step size presented to it, while the Stein discrepancy prefers an intermediate value. The rightmost plot of Figure 3b shows that a representative SGLD sample of size n using the ϵ selected by ESS is greatly overdispersed; the leftmost is greatly underdispersed due to slow mixing. The middle sample, with ϵ selected by the Stein diagnostic, most closely resembles the true posterior.

5.4 Quantifying a Bias-Variance Trade-off

The approximate random walk MH (ARWMH) sampler [5] is a second biased MCMC procedure designed for scalable posterior inference. Its tolerance parameter ϵ controls the number of datapoint likelihood evaluations used to approximate the standard MH correction step. Qualitatively, a larger ϵ implies fewer likelihood computations, more rapid sampling, and a more rapid reduction of variance. A smaller ϵ yields a closer approximation to the MH correction and less bias in the sampler stationary distribution. We will use the Stein discrepancy to explicitly quantify this bias-variance trade-off.

We analyze a dataset of 53 prostate cancer patients with six binary predictors and a binary outcome indicating whether cancer has spread to surrounding lymph nodes [24]. Our target is the Bayesian logistic regression posterior [1] under a $\mathcal{N}(0, I)$ prior on the parameters. We run RWMH ($\epsilon = 0$) and ARWMH ($\epsilon = 0.1$ and batch size = 2) for 10^5 likelihood evaluations, discard the points from the first 10^3 evaluations, and thin the remaining points to sequences of length 1000. The discrepancy computation time for 1000 points averages 1.3s for the spanner and 12s for a coordinate LP. Figure 4 displays the spanner Stein discrepancy applied to the first n points in each sequence as a function of the likelihood evaluation count. We see that the approximate sample is of higher Stein quality for smaller computational budgets but is eventually overtaken by the asymptotically exact sequence.



(a) Step size selection criteria (b) 1000 SGLD sample points with equidensity contours of p overlaid

Figure 3: (a) ESS maximized at $\epsilon = 5 \times 10^{-2}$; Stein discrepancy minimized at $\epsilon = 5 \times 10^{-3}$.




Figure 4: Bias-variance trade-off curves for Bayesian logistic regression with approximate RWMH.

To corroborate our result, we use a Metropolis-adjusted Langevin chain [25] of length 10^7 as a surrogate Q^* for the target and compute several error measures for each sample Q : normalized probability error $\max_l |\mathbb{E}[\sigma(\langle X, w_l \rangle) - \sigma(\langle Z, w_l \rangle)]| / \|w_l\|_\infty$, mean error $\frac{\max_j |\mathbb{E}[X_j - Z_j]|}{\max_j |\mathbb{E}_{Q^*}[Z_j]|}$, and second moment error $\frac{\max_{j,k} |\mathbb{E}[X_j X_k - Z_j Z_k]|}{\max_{j,k} |\mathbb{E}_{Q^*}[Z_j Z_k]|}$ for $X \sim Q$, $Z \sim Q^*$, $\sigma(t) \triangleq \frac{1}{1+e^{-t}}$, and w_l the l -th datapoint covariate vector. The measures, also found in Figure 4, accord with the Stein discrepancy quantification.

5.5 Assessing Convergence Rates

The Stein discrepancy can also be used to assess the quality of deterministic sample sequences. In Figure 5 in the appendix, for $P = \text{Unif}(0, 1)$, we plot the complete graph Stein discrepancies of the first n points of an i.i.d. $\text{Unif}(0, 1)$ sample, a deterministic Sobol sequence [26], and a deterministic kernel herding sequence [27] defined by the norm $\|h\|_{\mathcal{H}} = \int_0^1 (h'(x))^2 dx$. We use the median value over 50 sequences in the i.i.d. case and estimate the convergence rate for each sampler using the slope of the best least squares affine fit to each log-log plot. The discrepancy computation time averages 0.08s for $n = 200$ points, and the recovered rates of $n^{-0.49}$ and n^{-1} for the i.i.d. and Sobol sequences accord with expected $O(1/\sqrt{n})$ and $O(\log(n)/n)$ bounds from the literature [28, 26]. As witnessed also in other metrics [29], the herding rate of $n^{-0.96}$ outpaces its best known bound of $d_{\mathcal{H}}(Q_n, P) = O(1/\sqrt{n})$, suggesting an opportunity for sharper analysis.

6 Discussion of Related Work

We have developed a quality measure suitable for comparing biased, exact, and deterministic sample sequences by exploiting an infinite class of known target functionals. The diagnostics of [30, 31] also account for asymptotic bias but lose discriminating power by considering only a finite collection of functionals. For example, for a $\mathcal{N}(0, 1)$ target, the score statistic of [31] cannot distinguish two samples with equal first and second moments. Maximum mean discrepancy (MMD) on a characteristic Hilbert space [32] takes full distributional bias into account but is only viable when the expected kernel evaluations are easily computed under the target. One can approximate MMD, but this requires access to a separate trustworthy ground-truth sample from the target.




Figure 5: Comparison of complete graph Stein discrepancy convergence for $P = \text{Unif}(0, 1)$.

A Proof of Proposition 1

Our integrability assumption together with the boundedness of g and ∇g imply that $\mathbb{E}_P[\langle \nabla, g(Z) \rangle]$ and $\mathbb{E}_P[\langle g(Z), \nabla \log p(Z) \rangle]$ exist. Define the ℓ_∞ ball of radius r , $\mathbb{B}_r = \{x \in \mathbb{R}^d : \|x\|_\infty \leq r\}$. Since \mathcal{X} is convex, the intersection $\mathcal{X} \cap \mathbb{B}_r$ is compact and convex with Lipschitz boundary $\partial(\mathcal{X} \cap \mathbb{B}_r)$. Thus, the divergence theorem (integration by parts) implies that

$$\begin{aligned} \mathbb{E}_P[(\mathcal{T}_P g)(Z)] &= \mathbb{E}_P[\langle \nabla, g(Z) \rangle + \langle g(Z), \nabla \log p(Z) \rangle] = \int_{\mathcal{X}} \langle \nabla, p(z)g(z) \rangle dz \\ &= \lim_{r \rightarrow \infty} \int_{\mathcal{X} \cap \mathbb{B}_r} \langle \nabla, p(z)g(z) \rangle dz = \lim_{r \rightarrow \infty} \int_{\partial(\mathcal{X} \cap \mathbb{B}_r)} \langle g(z), n_r(z) \rangle p(z) dz \end{aligned}$$

for n_r the outward unit normal vector to $\partial(\mathcal{X} \cap \mathbb{B}_r)$. The final quantity in this expression equates to zero, as $\langle g(x), n(x) \rangle = 0$ for all x on the boundary $\partial\mathcal{X}$, g is bounded, and $\lim_{m \rightarrow \infty} p(x_m) = 0$ for any $(x_m)_{m=1}^\infty$ with $x_m \in \mathcal{X}$ for all m and $\|x_m\|_\infty \rightarrow \infty$.

B Proof of Theorem 2: Stein Discrepancy Lower Bound for Strongly Log-concave Densities

We let $C^k(\mathcal{X})$ denote the set of real-valued functions on \mathcal{X} with k continuous derivatives and $d_{\mathcal{M}_{\|\cdot\|}}$ denote the *smooth function distance*, the IPM generated by

$$\mathcal{M}_{\|\cdot\|} \triangleq \left\{ h \in C^3(\mathcal{X}) \mid \sup_{x \in \mathcal{X}} \max \left(\|\nabla h(x)\|^*, \|\nabla^2 h(x)\|^*, \|\nabla^3 h(x)\|^* \right) \leq 1 \right\}.$$

We additionally define the operator norms $\|v\|_{op} \triangleq \|v\|_2$ for vectors $v \in \mathbb{R}^d$, $\|M\|_{op} \triangleq \sup_{v \in \mathbb{R}^d, \|v\|_2=1} \|Mv\|_2$ for matrices $M \in \mathbb{R}^{d \times d}$, and $\|T\|_{op} \triangleq \sup_{v \in \mathbb{R}^d, \|v\|_2=1} \|T[v]\|_{op}$ for tensors $T \in \mathbb{R}^{d \times d \times d}$.

The following result, proved in Appendix G, establishes the existence of explicit constants (*Stein factors*) $c_1, c_2, c_3 > 0$, such that, for any test function $h \in \mathcal{M}_{\|\cdot\|}$, the *Stein equation*

$$h(x) - \mathbb{E}_P[h(Z)] = (\mathcal{T}_P g_h)(x)$$

has a solution $g_h = \frac{1}{2} \nabla u_h$ belonging to the non-uniform Stein set $\mathcal{G}_{\|\cdot\|}^{c_1, c_3}$.

Theorem 7 (Stein Factors for Strongly Log-concave Densities). *Suppose that $\mathcal{X} = \mathbb{R}^d$ and that $\log p \in C^4(\mathcal{X})$ is k -strongly concave with*

$$\sup_{z \in \mathcal{X}} \|\nabla^3 \log p(z)\|_{op} \leq L_3 \quad \text{and} \quad \sup_{z \in \mathcal{X}} \|\nabla^4 \log p(z)\|_{op} \leq L_4.$$

For each $x \in \mathcal{X}$, let $(Z_{t,x})_{t \geq 0}$ represent the overdamped Langevin diffusion with infinitesimal generator

$$(\mathcal{A}u)(x) = \frac{1}{2} \langle \nabla u(x), \nabla \log p(x) \rangle + \frac{1}{2} \langle \nabla, \nabla u(x) \rangle \quad (9)$$

and initial state $Z_{0,x} = x$. Then, for each $h \in C^3(\mathcal{X})$ with bounded first, second, and third derivatives, the function

$$u_h(x) \triangleq \int_0^\infty \mathbb{E}_P[h(Z)] - \mathbb{E}[h(Z_{t,x})] dt$$

solves the Stein equation

$$h(x) - \mathbb{E}_P[h(Z)] = (\mathcal{A}u_h)(x) \quad (10)$$

and satisfies

$$\begin{aligned} \sup_{z \in \mathcal{X}} \|\nabla u_h(z)\|_2 &\leq \frac{2}{k} \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2, \\ \sup_{z \in \mathcal{X}} \|\nabla^2 u_h(z)\|_{op} &\leq \frac{2L_3}{k^2} \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 + \frac{1}{k} \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op}, \text{ and} \\ \sup_{z, y \in \mathcal{X}, z \neq y} \frac{\|\nabla^2 u_h(z) - \nabla^2 u_h(y)\|_{op}}{\|z - y\|_2} &\leq \frac{6L_3^2}{k^3} \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 + \frac{L_4}{k^2} \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \\ &\quad + \frac{3L_3}{k^2} \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} + \frac{2}{3k} \sup_{z \in \mathcal{X}} \|\nabla^3 h(z)\|_{op}. \end{aligned}$$

Hence, by the equivalence of non-uniform Stein discrepancies (Proposition 4), $d_{\mathcal{M}_{\|\cdot\|}}(\mu, P) \leq \mathcal{S}(\mu, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}^{c_1, c_3}) \leq \max(c_1, c_2, c_3) \mathcal{S}(\mu, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|})$ for any probability measure μ .

The desired result now follows from Lemma 8, which implies that the Wasserstein distance $d_{\mathcal{W}_{\|\cdot\|}}(\mu_m, P) \rightarrow 0$ whenever $d_{\mathcal{M}_{\|\cdot\|}}(\mu_m, P) \rightarrow 0$ for a sequence of probability measures $(\mu_m)_{m \geq 1}$.

Lemma 8 (Smooth-Wasserstein Inequality). *If μ and ν are probability measures on \mathbb{R}^d , and $\|v\| \geq \|v\|_2$ for all $v \in \mathbb{R}^d$, then*

$$d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \leq d_{\mathcal{W}_{\|\cdot\|}}(\mu, \nu) \leq 3 \max\left(d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu), \sqrt[3]{d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \sqrt{2} \mathbb{E}[\|G\|^2]}\right).$$

for G a standard normal random vector in \mathbb{R}^d .

Proof The first inequality follows directly from the inclusion $\mathcal{M}_{\|\cdot\|} \subset \mathcal{W}_{\|\cdot\|}$.

To establish the second inequality, we fix an $h \in \mathcal{W}_{\|\cdot\|}$ and $t > 0$ and define the smoothed function

$$h_t(x) = \int_{\mathbb{R}^d} h(x + tz) \phi(z) dz \quad \text{for each } x \in \mathbb{R}^d,$$

where ϕ is the density of a vector of d independent standard normal variables. We first show that h_t is a close approximation to h when t is small. Specifically, if $X \in \mathbb{R}^d$ is an integrable random vector, independent of G , then

$$|\mathbb{E}[h(X) - h_t(X)]| = |\mathbb{E}[h(X) - h(X + tG)]| \leq t \mathbb{E}[\|G\|]$$

by the Lipschitz assumption on h .

We next show that the derivatives of h_t are bounded. Fix any $x \in \mathbb{R}^d$. Since h is Lipschitz, it admits a weak gradient, ∇h , bounded uniformly by 1 in $\|\cdot\|^*$. We alternate differentiation and integration by parts to develop the representations

$$\begin{aligned} \nabla h_t(x) &= \int_{\mathbb{R}^d} \nabla h(x + tz) \phi(z) dz = \frac{1}{t} \int_{\mathbb{R}^d} z h(x + tz) \phi(z) dz, \\ \nabla^2 h_t(x) &= \frac{1}{t} \int_{\mathbb{R}^d} \nabla h(x + tz) z^\top \phi(z) dz = \frac{1}{t^2} \int_{\mathbb{R}^d} (zz^\top - I) h(x + tz) \phi(z) dz, \quad \text{and} \\ \nabla^3 h_t(x)[v] &= \frac{1}{t^2} \int_{\mathbb{R}^d} \nabla h(x + tz) v^\top (zz^\top - I) \phi(z) dz \end{aligned}$$

for each $v \in \mathbb{R}^d$. The uniform bound on ∇h and the relation between $\|\cdot\|$ and $\|\cdot\|_2$ now yield

$$\begin{aligned} \|\nabla h_t(x)\|^* &\leq 1, \\ \|\nabla^2 h_t(x)\|^* &\leq \frac{1}{t} \sup_{v \in \mathbb{R}^d: \|v\|=1} \int_{\mathbb{R}^d} |\langle z, v \rangle| \phi(z) dz = \frac{1}{t} \sqrt{\frac{2}{\pi}} \sup_{v \in \mathbb{R}^d: \|v\|=1} \|v\|_2 \leq \frac{1}{t} \sqrt{\frac{2}{\pi}}, \quad \text{and} \\ \|\nabla^3 h_t(x)\|^* &\leq \frac{1}{t^2} \sup_{v, w \in \mathbb{R}^d: \|v\|=\|w\|=1} \int_{\mathbb{R}^d} |v^\top (zz^\top - I)w| \phi(z) dz \\ &\leq \frac{1}{t^2} \sup_{v, w \in \mathbb{R}^d: \|v\|=\|w\|=1} \sqrt{\int_{\mathbb{R}^d} |v^\top (zz^\top - I)w|^2 \phi(z) dz} \\ &= \frac{1}{t^2} \sup_{v, w \in \mathbb{R}^d: \|v\|=\|w\|=1} \sqrt{\langle v, w \rangle^2 + \|v\|_2^2 \|w\|_2^2} \leq \frac{\sqrt{2}}{t^2}. \end{aligned}$$

In final equality we have used the fact that $\langle v, Z \rangle$ and $\langle w, Z \rangle$ are jointly normal with zero mean and covariance $\Sigma = \begin{bmatrix} \|v\|_2^2 & \langle v, w \rangle \\ \langle v, w \rangle & \|w\|_2^2 \end{bmatrix}$, so that the product $\langle v, Z \rangle \langle w, Z \rangle$ has the distribution of the off-diagonal element of the Wishart distribution with scale Σ and 1 degree of freedom.

We can now develop a bound for $d_{\mathcal{W}_{\|\cdot\|}}$ using our smoothed functions. Introduce the shorthand

$$b_t \triangleq \max\left(1, \frac{1}{t} \sqrt{\frac{2}{\pi}}, \frac{\sqrt{2}}{t^2}\right) = \max\left(1, \frac{\sqrt{2}}{t^2}\right)$$

for the maximum derivative bound of h_t , and select $X \sim \mu$ and $Z \sim \nu$ to satisfy $d_{\mathcal{W}_{\|\cdot\|}}(\mu, \nu) = \mathbb{E}[\|X - Z\|]$. We then have

$$\begin{aligned} d_{\mathcal{W}_{\|\cdot\|}}(\mu, \nu) &\leq \inf_{t>0} \sup_{h \in \mathcal{W}_{\|\cdot\|}} |\mathbb{E}_\mu[h(X) - h_t(X)]| + |\mathbb{E}_\nu[h(Z) - h_t(Z)]| + |\mathbb{E}_\mu[h_t(X)] - \mathbb{E}_\nu[h_t(Z)]| \\ &\leq \inf_{t>0} 2t \mathbb{E}[\|G\|] + b_t d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \\ &\leq 2 \sqrt[3]{d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \sqrt{2} \mathbb{E}[\|G\|]^2} + \max\left(d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu), \sqrt[3]{d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \sqrt{2} \mathbb{E}[\|G\|]^2}\right) \\ &\leq 3 \max\left(d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu), \sqrt[3]{d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \sqrt{2} \mathbb{E}[\|G\|]^2}\right), \end{aligned}$$

where we have chosen $t = \sqrt[3]{d_{\mathcal{M}_{\|\cdot\|}}(\mu, \nu) \sqrt{2} / \mathbb{E}[\|G\|]}$ to achieve the penultimate inequality. \square

C Proof of Proposition 3: Stein Discrepancy Upper Bound

Fix any g in $\mathcal{G}_{\|\cdot\|}$. By Proposition 1, $\mathbb{E}[(\mathcal{T}_P g)(Z)] = 0$. The Lipschitz and boundedness constraints on g and ∇g now yield

$$\begin{aligned} \mathbb{E}_Q[(\mathcal{T}_P g)(X)] &= \mathbb{E}[(\mathcal{T}_P g)(X) - (\mathcal{T}_P g)(Z)] \\ &= \mathbb{E}[\langle g(X), \nabla \log p(X) \rangle - \langle g(Z), \nabla \log p(Z) \rangle + \langle \nabla, g(X) - g(Z) \rangle] \\ &= \mathbb{E}[\langle g(X), \nabla \log p(X) - \nabla \log p(Z) \rangle + \langle g(X) - g(Z), \nabla \log p(Z) \rangle] \\ &\quad + \mathbb{E}[\langle \nabla, g(X) - g(Z) \rangle] \\ &\leq \mathbb{E}[\|\nabla \log p(X) - \nabla \log p(Z)\|] + \mathbb{E}[\|\nabla \log p(Z)(X - Z)^\top\|] + \mathbb{E}[\|X - Z\|]. \end{aligned}$$

To derive the second advertised inequality, we use the definition of the matrix norm, the Fenchel-Young inequality for dual norms, the definition of the matrix dual norm, and the Cauchy-Schwarz

inequality in turn:

$$\begin{aligned}
\mathbb{E}[\|\nabla \log p(Z)(X - Z)^\top\|] &= \mathbb{E}\left[\sup_{M:\|M\|^*=1} \langle \nabla \log p(Z), M(X - Z) \rangle\right] \\
&\leq \mathbb{E}\left[\sup_{M:\|M\|^*=1} \|\nabla \log p(Z)\| \|M(X - Z)\|^*\right] \\
&\leq \mathbb{E}[\|\nabla \log p(Z)\| \|X - Z\|] \leq \sqrt{\mathbb{E}[\|\nabla \log p(Z)\|^2] \mathbb{E}[\|X - Z\|^2]}.
\end{aligned}$$

Since our bounds hold uniformly for all g in $\mathcal{G}_{\|\cdot\|}$, the proof is complete.

D Proof of Proposition 4: Equivalence of Non-uniform Stein Discrepancies

Fix any $c_1, c_2, c_3 > 0$, and let $c_{\max} = \max(c_1, c_2, c_3)$ and $c_{\min} = \min(c_1, c_2, c_3)$. Since the Stein discrepancy objective is linear in g , we have $a\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}) = \mathcal{S}(Q, \mathcal{T}_P, a\mathcal{G}_{\|\cdot\|})$ for any $a > 0$. The result now follows from the observation that $c_{\min}\mathcal{G}_{\|\cdot\|} \subseteq \mathcal{G}_{\|\cdot\|}^{c_1, c_3} \subseteq c_{\max}\mathcal{G}_{\|\cdot\|}$.

E Proof of Proposition 5: Equivalence of Classical and Complete Graph Stein Discrepancies

The first inequality follows from the fact that $\mathcal{G}_{\|\cdot\|} \subseteq \mathcal{G}_{\|\cdot\|, Q, G_1}$. By the Whitney-Glaeser extension theorem [15, Thm. 1.4] of Glaeser [14], for every function $g \in \mathcal{G}_{\|\cdot\|, Q, G_1}$, there exists a function $\tilde{g} \in \kappa_d \mathcal{G}_{\|\cdot\|}^*$ with $g(x_i) = \tilde{g}(x_i)$ and $\nabla g(x_i) = \nabla \tilde{g}(x_i)$ for all x_i in the support of Q . Here κ_d is a constant, independent of (Q, P) , depending only on the dimension d and norm $\|\cdot\|$. Since the Stein discrepancy objective is linear in g and depends on g only through the values $g(x_i)$ and $\nabla g(x_i)$, we have $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1}) \leq \mathcal{S}(Q, \mathcal{T}_P, \kappa_d \mathcal{G}_{\|\cdot\|}^*) = \kappa_d \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}^*)$.

F Proof of Proposition 6: Equivalence of Spanner and Complete Graph Stein Discrepancies

The first inequality follows from the fact that $\mathcal{G}_{\|\cdot\|, Q, G_1} \subseteq \mathcal{G}_{\|\cdot\|, Q, G_t}$. Fix any $g \in \mathcal{G}_{\|\cdot\|, Q, G_t}$ and any pair of points $z, z' \in \text{supp}(Q)$. By the definition of $\mathcal{G}_{\|\cdot\|, Q, G_t}$, we have $\max(\|g(z)\|^*, \|\nabla g(z)\|^*) \leq 1$. By the t -spanner property, there exists a sequence of points $z_0, z_1, z_2, \dots, z_{L-1}, z_L \in \text{supp}(Q)$ with $z_0 = z$ and $z_L = z'$ for which $(z_{l-1}, z_l) \in E$ for all $1 \leq l \leq L$ and $\sum_{l=1}^L \|z_{l-1} - z_l\| \leq t\|z_0 - z_L\|$. Since $\max\left(\frac{\|g(z_{l-1}) - g(z_l)\|^*}{\|z_{l-1} - z_l\|}, \frac{\|\nabla g(z_{l-1}) - \nabla g(z_l)\|^*}{\|z_{l-1} - z_l\|}\right) \leq 1$ for each l , the triangle inequality implies that

$$\|\nabla g(z_0) - \nabla g(z_L)\|^* \leq \sum_{l=1}^L \|\nabla g(z_{l-1}) - \nabla g(z_l)\|^* \leq \sum_{l=1}^L \|z_{l-1} - z_l\| \leq t\|z_0 - z_L\|.$$

Identical reasoning establishes that $\|g(z_0) - g(z_L)\|^* \leq t\|z_0 - z_L\|$.

Furthermore, since $\|g(z_{l-1}) - g(z_l) - \nabla g(z_l)(z_{l-1} - z_l)\|^* \leq \frac{1}{2}\|z_{l-1} - z_l\|^2$ for each l , the triangle inequality and the definition of the tensor norm $\|\cdot\|^*$ imply that

$$\begin{aligned}
& \|g(z_0) - g(z_L) - \nabla g(z_L)(z_0 - z_L)\|^* \\
& \leq \sum_{l=1}^L \|g(z_{l-1}) - g(z_l) - \nabla g(z_l)(z_{l-1} - z_l)\|^* + \|(\nabla g(z_l) - \nabla g(z_L))(z_{l-1} - z_l)\|^* \\
& \leq \sum_{l=1}^L \frac{1}{2}\|z_{l-1} - z_l\|^2 + \|\nabla g(z_l) - \nabla g(z_L)\|^* \|z_{l-1} - z_l\| \\
& \leq \sum_{l=1}^L \frac{1}{2}\|z_{l-1} - z_l\|^2 + \sum_{l'=l}^{L-1} \|\nabla g(z_{l'}) - \nabla g(z_{l'+1})\|^* \|z_{l-1} - z_l\| \\
& \leq \sum_{l=1}^L \|z_{l-1} - z_l\| \left(\frac{1}{2}\|z_{l-1} - z_l\| + \sum_{l'=l}^{L-1} \|z_{l'} - z_{l'+1}\| \right) \leq \left(\sum_{l=1}^L \|z_{l-1} - z_l\| \right)^2 \leq t^2 \|z_0 - z_L\|^2.
\end{aligned}$$

Since z, z' were arbitrary, and the Stein discrepancy objective is linear in g , we conclude that $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_t}) \leq \mathcal{S}(Q, \mathcal{T}_P, 2t^2 \mathcal{G}_{\|\cdot\|, Q, G_1}) = 2t^2 \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|, Q, G_1})$.

G Proof of Theorem 7: Stein Factors for Strongly Log-concave Densities

Before tackling the main proof, we will establish a series of useful lemmas. We will make regular use of the following well-known Lipschitz properties throughout:

$$\sup_{x \in \mathcal{X}} \|\nabla h(x)\|_2 = \sup_{x, y \in \mathcal{X}, x \neq y} \frac{|h(x) - h(y)|}{\|x - y\|_2} \quad \text{for all } h \in C^1(\mathcal{X}) \quad \text{and} \quad (11)$$

$$\sup_{x \in \mathcal{X}} \|\nabla^k h(x)\|_{op} = \sup_{x, y \in \mathcal{X}, x \neq y} \frac{\|\nabla^{k-1} h(x) - \nabla^{k-1} h(y)\|_{op}}{\|x - y\|_2} \quad \text{for all } h \in C^k(\mathcal{X}), \quad (12)$$

for each integer $k > 1$.

G.1 Properties of Overdamped Langevin Diffusions

Our first lemma enumerates several well-known properties of the overdamped Langevin diffusion that will prove useful in the proofs to follow.

Lemma 9 (Overdamped Langevin Properties). *If $\mathcal{X} = \mathbb{R}^d$, and $\log p \in C^1(\mathcal{X})$ is strongly concave, then the overdamped Langevin diffusion $(Z_{t,x})_{t \geq 0}$ with infinitesimal generator (9) and $Z_{0,x} = x$ is well-defined for all times $t \in [0, \infty)$, has stationary measure P , and satisfies the strong Feller property.*

Proof Consider the candidate Lyapunov function $V(x) = \|x\|_2^2 + 1$. The strong log-concavity of p , the Cauchy-Schwarz inequality, and the arithmetic-geometric mean inequality together imply that

$$\begin{aligned}
(\mathcal{A}V)(x) &= \langle x, \nabla \log p(x) \rangle + d = \langle x, \nabla \log p(x) - \nabla \log p(0) \rangle + \langle x, \nabla \log p(0) \rangle + d \\
&\leq -k\|x\|_2^2 + \|x\|_2 \|\nabla \log p(0)\|_2 + d \leq \left(\frac{1}{2} - k \right) \|x\|_2^2 + \frac{1}{2} \|\nabla \log p(0)\|_2^2 + d \leq k'V(x)
\end{aligned}$$

for some constants $k, k' \in \mathbb{R}$. Since $\log p$ is continuously differentiable, Theorem 2.1 of Roberts and Tweedie [25] implies the result (see also [33, Thm. 3.5]). \square

G.2 High-order Weighted Difference Bounds

A second, technical lemma bounds the growth of weighted smooth function differences in terms of the proximity of function arguments. The result will be used to characterize the smoothness of $Z_{t,x}$ as a function of the starting point x (Lemma 11) and, ultimately, to establish the smoothness of u_h (Theorem 7).

Lemma 10 (High-order Weighted Difference Bounds). *Fix any open convex set $\mathcal{X} \subseteq \mathbb{R}^d$, any vectors $x, y, z, w, x', y', z', w' \in \mathcal{X}$, and any weights $\lambda, \lambda' > 0$. If $h \in C^2(\mathcal{X})$, then*

$$\begin{aligned} & |\lambda(h(x) - h(y)) - \lambda'(h(x') - h(y')) - \langle \nabla h(y), \lambda(x - y) - \lambda'(x' - y') \rangle| \\ & \leq \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} (2\lambda' \|y - y'\|_2 \|x' - y'\|_2 + \lambda \|x - y\|_2^2 + \lambda' \|x' - y'\|_2^2). \end{aligned} \quad (13)$$

Moreover, if $h \in C^3(\mathcal{X})$, then

$$\begin{aligned} & |\lambda(h(x) - h(y) - (h(z) - h(w))) - \lambda'(h(x') - h(y') - (h(z') - h(w')))| \\ & - \langle \nabla h(z), \lambda(x - y - (z - w)) - \lambda'(x' - y' - (z' - w')) \rangle \\ & \leq \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} \|y' - x'\|_2 \|\lambda(z - x) - \lambda'(z' - x')\|_2 \\ & + \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} (\lambda' \|z - z'\|_2 \|x' - y' - (z' - w')\|_2 + \lambda \|z - x\|_2 \|(y - x) - (y' - x')\|_2) \\ & + \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} \lambda \|x - y - (z - w)\|_2 \|x - y + z - w\|_2 \\ & + \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} \lambda' \|x' - y' - (z' - w')\|_2 \|x' - y' + z' - w'\|_2 \\ & + \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} \|y' - x'\|_2 (2\lambda' \|x - x'\|_2 \|z' - x'\|_2 + \lambda \|z - x\|_2^2 + \lambda' \|z' - x'\|_2^2) \\ & + \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} (\lambda \|z - x\|_2 \|y - x\|_2^2 + \lambda' \|z' - x'\|_2 \|y' - x'\|_2^2) \\ & + \frac{1}{6} \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} (\lambda \|w - z\|_2^3 + \lambda \|y - x\|_2^3 + \lambda' \|w' - z'\|_2^3 + \lambda' \|y' - x'\|_2^3). \end{aligned} \quad (14)$$

Proof To establish the second-order difference bound (13), we first apply Taylor's theorem with mean-value remainder to $h(x) - h(y)$ and $h(x') - h(y')$ to obtain

$$\begin{aligned} & \lambda(h(x) - h(y)) - \lambda'(h(x') - h(y')) - \langle \nabla h(y), \lambda(x - y) - \lambda'(x' - y') \rangle \\ & = \lambda' \langle \nabla h(y) - \nabla h(y'), x' - y' \rangle + \lambda \langle \nabla^2 h(\zeta)(x - y), x - y \rangle / 2 - \lambda' \langle \nabla^2 h(\zeta')(x' - y'), x' - y' \rangle / 2 \end{aligned}$$

for some $\zeta, \zeta' \in \mathcal{X}$. Cauchy-Schwarz, the definition of the operator norm, and the Lipschitz gradient relation (12) now yield

$$\begin{aligned} & |h(x) - h(y) - (h(x') - h(y')) - \langle \nabla h(y), x - y - (x' - y') \rangle| \\ & \leq \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} (2\lambda' \|y - y'\|_2 \|x' - y'\|_2 + \lambda \|x - y\|_2^2 + \lambda' \|x' - y'\|_2^2). \end{aligned}$$

To derive the third-order difference bound (14), we apply Taylor's theorem with mean-value remainder to $h(w) - h(z)$, $h(y) - h(x)$, $h(w') - h(z')$, and $h(y') - h(x')$ to write

$$\begin{aligned} & |\lambda(h(x) - h(y) - (h(z) - h(w))) - \lambda'(h(x') - h(y') - (h(z') - h(w')))| \\ & - \langle \nabla h(z), \lambda(x - y - (z - w)) - \lambda'(x' - y' - (z' - w')) \rangle \\ & = |\lambda' \langle \nabla h(z) - \nabla h(z'), x' - y' - (z' - w') \rangle + \lambda \langle \nabla h(z) - \nabla h(x), (y - x) - (y' - x') \rangle| \\ & + \langle \lambda(\nabla h(z) - \nabla h(x)) - \lambda'(\nabla h(z') - \nabla h(x')), y' - x' \rangle \\ & + \lambda \langle \nabla^2 h(z)(w - z), w - z \rangle / 2 - \lambda \langle \nabla^2 h(x)(y - x), y - x \rangle / 2 \\ & - \lambda' \langle \nabla^2 h(z')(w' - z'), w' - z' \rangle / 2 + \lambda' \langle \nabla^2 h(x')(y' - x'), y' - x' \rangle / 2 \\ & + \lambda \nabla^3 h(\zeta'')[w - z, w - z, w - z] / 6 - \lambda \nabla^3 h(\zeta''')[y - x, y - x, y - x] / 6 \\ & - \lambda' \nabla^3 h(\zeta''')[w' - z', w' - z', w' - z'] / 6 + \lambda' \nabla^3 h(\zeta''''')[y' - x', y' - x', y' - x'] / 6 \end{aligned} \quad (15)$$

for some $\zeta'', \zeta''', \zeta'''' \in \mathcal{X}$. We will bound each line in this expression in turn. First we see, by Cauchy-Schwarz and the Lipschitz property (12), that

$$\begin{aligned} & |\lambda' \langle \nabla h(z) - \nabla h(z'), x' - y' - (z' - w') \rangle + \lambda \langle \nabla h(z) - \nabla h(x), (y - x) - (y' - x') \rangle| \\ & \leq \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} (\lambda' \|z - z'\|_2 \|x' - y' - (z' - w')\|_2 + \lambda \|z - x\|_2 \|(y - x) - (y' - x')\|_2). \end{aligned}$$

Next, we invoke our second-order difference bound (13) on the $C^2(\mathcal{X})$ function $x \mapsto \langle \nabla h(x), y' - x' \rangle$, apply the Cauchy-Schwarz inequality, and use the definition of the operator norm to conclude that

$$\begin{aligned} & |(\lambda(\nabla h(z) - \nabla h(x)) - \lambda'(\nabla h(z') - \nabla h(x')), y' - x')| \\ & \leq \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} \|y' - x'\|_2 \|\lambda(z - x) - \lambda'(z' - x')\|_2 \\ & \quad + \frac{1}{2} \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} \|y' - x'\|_2 (2\lambda' \|x - x'\|_2 \|z' - x'\|_2 + \lambda \|z - x\|_2^2 + \lambda' \|z' - x'\|_2^2). \end{aligned}$$

To bound the subsequent line, we note that Cauchy-Schwarz, the definition of the operator norm, and the Lipschitz property (12) imply that

$$\begin{aligned} & |\langle \nabla^2 h(z)(w - z), w - z \rangle - \langle \nabla^2 h(x)(y - x), y - x \rangle| \\ & = |\langle \nabla^2 h(z)(w - z + y - x), x - y - (z - w) \rangle + \langle (\nabla^2 h(z) - \nabla^2 h(x))(y - x), y - x \rangle| \\ & \leq \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} \|x - y - (z - w)\|_2 \|x - y + z - w\|_2 + \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} \|z - x\|_2 \|y - x\|_2^2. \end{aligned}$$

Similarly,

$$\begin{aligned} & |\langle \nabla^2 h(z')(w' - z'), w' - z' \rangle - \langle \nabla^2 h(x')(y' - x'), y' - x' \rangle| \\ & \leq \sup_{a \in \mathcal{X}} \|\nabla^2 h(a)\|_{op} \|x' - y' - (z' - w')\|_2 \|x' - y' + z' - w'\|_2 + \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} \|z' - x'\|_2 \|y' - x'\|_2^2. \end{aligned}$$

Finally, Cauchy-Schwarz and the definition of the operator norm give

$$\begin{aligned} & |\lambda \nabla^3 h(\zeta'') [w - z, w - z, w - z] - \lambda \nabla^3 h(\zeta''') [y - x, y - x, y - x] \\ & \quad - \lambda' \nabla^3 h(\zeta''') [w' - z', w' - z', w' - z'] + \lambda' \nabla^3 h(\zeta''''') [y' - x', y' - x', y' - x']| \\ & \leq \sup_{a \in \mathcal{X}} \|\nabla^3 h(a)\|_{op} (\lambda \|w - z\|_2^3 + \lambda \|y - x\|_2^3 + \lambda' \|w' - z'\|_2^3 + \lambda' \|y' - x'\|_2^3). \end{aligned}$$

Bounding the third-order difference (15) in terms of these four estimates yields the advertised inequality (14). \square

G.3 Synchronous Coupling Lemma

Our proof of Theorem 7 additionally rests upon a series of coupling inequalities which serve to characterize the smoothness of $Z_{t,x}$ as a function of x . The couplings espoused in the lemma to follow are termed *synchronous*, because the same Brownian motion is used to drive each process.

Lemma 11 (Synchronous Coupling Inequalities). *Suppose that $\mathcal{X} = \mathbb{R}^d$ and that $\log p \in C^4(\mathcal{X})$ is k -strongly concave with*

$$\sup_{z \in \mathcal{X}} \|\nabla^3 \log p(z)\|_{op} \leq L_3 \quad \text{and} \quad \sup_{z \in \mathcal{X}} \|\nabla^4 \log p(z)\|_{op} \leq L_4.$$

Select any vectors $x, x', v, v' \in \mathcal{X}$ with $\|v\|_2 = \|v'\|_2 = 1$ and any weights $\epsilon, \epsilon', \epsilon'' > 0$, and let $(W_t)_{t \geq 0}$ represent a fixed d -dimensional Wiener process.

For each starting point of the form $z + b'v' + bv$ with $z \in \{x, x'\}$, $b' \in \{0, \epsilon', \epsilon''\}$, and $b \in \{0, \epsilon\}$, consider an overdamped Langevin diffusion $(Z_{t, z+b'v'+bv})_{t \geq 0}$ solving the stochastic differential equation

$$dZ_{t, z+b'v'+bv} = \frac{1}{2} \nabla \log p(Z_{t, z+b'v'+bv}) dt + dW_t \quad \text{with} \quad Z_{0, z+b'v'+bv} = z + b'v' + bv, \quad (16)$$

and define the differenced processes

$$\begin{aligned} V_t & \triangleq (Z_{t, x'+\epsilon''v'} - Z_{t, x'})/\epsilon'' - (Z_{t, x+\epsilon'v'} - Z_{t, x})/\epsilon' \quad \text{and} \\ U_t & \triangleq (Z_{t, x'+\epsilon''v'+\epsilon v} - Z_{t, x'+\epsilon''v'} - (Z_{t, x'+\epsilon v} - Z_{t, x'}))/(\epsilon\epsilon'') \\ & \quad - (Z_{t, x+\epsilon'v'+\epsilon v} - Z_{t, x+\epsilon'v'} - (Z_{t, x+\epsilon v} - Z_{t, x}))/(\epsilon\epsilon'). \end{aligned}$$

These coupled processes almost surely satisfy the synchronous coupling bounds,

$$e^{kt/2} \|Z_{t,x+\epsilon v} - Z_{t,x}\|_2 \leq \epsilon, \quad (17)$$

$$e^{kt/2} \|V_t\|_2 \leq \frac{L_3}{k} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2), \quad \text{and} \quad (18)$$

$$\begin{aligned} e^{kt/2} \|U_t\|_2 &\leq \frac{3L_3^2}{k^2} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\ &\quad + \frac{L_4}{2k} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3), \end{aligned} \quad (19)$$

the second-order differenced function bound,

$$\begin{aligned} &(h_2(Z_{t,x'+\epsilon''v'}) - h_2(Z_{t,x'}))/\epsilon'' - (h_2(Z_{t,x+\epsilon'v'}) - h_2(Z_{t,x}))/\epsilon' \\ &\leq \langle \nabla h_2(Z_{t,x'}), V_t \rangle + \sup_{a \in \mathcal{X}} \|\nabla^2 h_2(a)\|_{op} e^{-kt} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2), \end{aligned} \quad (20)$$

and the third-order differenced function bound,

$$\begin{aligned} &(h_3(Z_{t,x'+\epsilon''v'+\epsilon v}) - h_3(Z_{t,x'+\epsilon''v'}) - (h_3(Z_{t,x'+\epsilon v}) - h_3(Z_{t,x'})))/(\epsilon\epsilon'') \\ &\quad - (h_3(Z_{t,x+\epsilon'v'+\epsilon v}) - h_3(Z_{t,x+\epsilon'v'}) - (h_3(Z_{t,x+\epsilon v}) - h_3(Z_{t,x}))) / (\epsilon\epsilon') \\ &\leq \langle \nabla h_3(Z_{t,x'+\epsilon''v'}), U_t \rangle \\ &\quad + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{k} e^{-kt} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\ &\quad + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} e^{-3kt/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3) \end{aligned} \quad (21)$$

for each $t \geq 0$, $h_2 \in C^2(\mathcal{X})$, and $h_3 \in C^3(\mathcal{X})$.

Proof By Lemma 9, each process $(Z_{t,z+b'v'+bv})_{t \geq 0}$ with $z \in \{x, x'\}$, $b' \in \{0, \epsilon', \epsilon''\}$, and $b \in \{0, \epsilon\}$ is well-defined for all times $t \in [0, \infty)$.

The first-order bound The first-order bound (17) is well known, and we include a short proof due to [34] for completeness. Since the differences,

$$Z_{t,x+\epsilon v} - Z_{t,x} = \epsilon v + \int_0^t \frac{1}{2} \nabla \log p(Z_{s,x+\epsilon v}) - \frac{1}{2} \nabla \log p(Z_{s,x}) ds$$

for $t \geq 0$ constitute an Itô process, we first apply Itô's lemma to the function $(t, w) \mapsto e^{kt} \|w\|_2^2$ and then invoke the k -strong log-concavity of p to conclude

$$\begin{aligned} e^{kt} \|Z_{t,x+\epsilon v} - Z_{t,x}\|_2^2 &= \epsilon^2 + \int_0^t k e^{ks} \|Z_{s,x+\epsilon v} - Z_{s,x}\|_2^2 + e^{ks} \frac{d}{ds} \|Z_{s,x+\epsilon v} - Z_{s,x}\|_2^2 ds \\ &= \epsilon^2 + \int_0^t e^{ks} (k \|Z_{s,x+\epsilon v} - Z_{s,x}\|_2^2 + \langle Z_{s,x+\epsilon v} - Z_{s,x}, \nabla \log p(Z_{s,x+\epsilon v}) - \nabla \log p(Z_{s,x}) \rangle) ds \\ &\leq \epsilon^2 + \int_0^t e^{ks} 0 ds = \epsilon^2 \quad \text{almost surely.} \end{aligned}$$

Second-order bounds To establish the second conclusion (18), we consider the Itô process of second-order differences

$$V_t = \frac{1}{2} \int_0^t (\nabla \log p(Z_{s,x'+\epsilon''v'}) - \nabla \log p(Z_{s,x'}))/\epsilon'' - (\nabla \log p(Z_{s,x+\epsilon'v'}) - \nabla \log p(Z_{s,x}))/\epsilon' ds$$

and apply Itô's lemma to the mapping $(t, w) \mapsto e^{kt/2} \|w\|_2$. This yields

$$\begin{aligned} e^{kt/2} \|V_t\|_2 &= e^0 \|V_0\|_2 + \int_0^t k e^{ks} \|V_s\|_2 + e^{ks} \frac{d}{ds} \|V_s\|_2 ds \\ &= \int_0^t \frac{e^{ks/2}}{2 \|V_s\|_2} \left(k \|V_s\|_2^2 \right. \\ &\quad \left. + \langle V_s, (\nabla \log p(Z_{s,x'+\epsilon''v'}) - \nabla \log p(Z_{s,x'}))/\epsilon'' - (\nabla \log p(Z_{s,x+\epsilon'v'}) - \nabla \log p(Z_{s,x}))/\epsilon' \rangle \right) ds. \end{aligned}$$

Fix a value $s \in [0, t]$. For any $h_2 \in C^2(\mathcal{X})$, the Lemma 10 second-order difference inequality (13) and the first order coupling bound (17) together imply the function coupling bound (20) as

$$\begin{aligned} & (h_2(Z_{s,x'+\epsilon''v'}) - h_2(Z_{s,x'}))/\epsilon'' - (h_2(Z_{s,x+\epsilon'v'}) - h_2(Z_{s,x}))/\epsilon' \\ & \leq \langle \nabla h_2(Z_{s,x'}), V_s \rangle + \frac{1}{2} \sup_{z \in \mathcal{X}} \|\nabla^2 h_2(z)\|_{op} (2\|Z_{s,x'} - Z_{s,x}\|_2 \|Z_{s,x+\epsilon'v'} - Z_{s,x}\|_2 / \epsilon' \\ & \quad + \|Z_{s,x'+\epsilon''v'} - Z_{s,x'}\|_2^2 / \epsilon'' + \|Z_{s,x+\epsilon'v'} - Z_{s,x}\|_2^2 / \epsilon') \\ & \leq \langle \nabla h_2(Z_{s,x'}), V_s \rangle + \sup_{z \in \mathcal{X}} \|\nabla^2 h_2(z)\|_{op} e^{-ks} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2). \end{aligned}$$

Applying this bound to the thrice continuously differentiable function $h_2(z) = \langle V_s, \nabla \log p(z) \rangle$ with

$$\sup_{z \in \mathcal{X}} \|\nabla^2 h_2(z)\|_{op} = \sup_{z \in \mathcal{X}} \|\nabla^3 \log p(z)[V_s]\|_{op} \leq L_3 \|V_s\|_2,$$

yields

$$\begin{aligned} & \langle V_s, (\nabla \log p(Z_{s,x'+\epsilon''v'}) - \nabla \log p(Z_{s,x'}))/\epsilon'' - (\nabla \log p(Z_{s,x+\epsilon'v'}) - \nabla \log p(Z_{s,x}))/\epsilon' \rangle \\ & \leq \langle V_s, \nabla^2 \log p(Z_{s,x'}) V_s \rangle + L_3 \|V_s\|_2 e^{-ks} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2) \\ & \leq -k \|V_s\|_2^2 + L_3 \|V_s\|_2 e^{-ks} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2). \end{aligned}$$

To achieve the second inequality, we used the k -strong log-concavity of p . Now we may derive the desired conclusion,

$$e^{kt/2} \|V_t\|_2 \leq \frac{L_3}{2} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2) \int_0^t e^{-ks/2} ds = \frac{L_3}{k} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2).$$

Third-order bounds To establish the third conclusion (19), we consider the Itô process of third-order differences

$$\begin{aligned} U_t &= \frac{1}{2} \int_0^t (\nabla \log p(Z_{s,x'+\epsilon''v'+\epsilon v}) - \nabla \log p(Z_{s,x'+\epsilon''v'}) - (\nabla \log p(Z_{s,x'+\epsilon v}) - \nabla \log p(Z_{s,x'})))/(\epsilon \epsilon'') \\ & \quad - (\nabla \log p(Z_{s,x+\epsilon'v'+\epsilon v}) - \nabla \log p(Z_{s,x+\epsilon'v'}) - (\nabla \log p(Z_{s,x+\epsilon v}) - \nabla \log p(Z_{s,x}))) / (\epsilon \epsilon') ds \end{aligned}$$

and invoke Itô's lemma once more for the mapping $(t, w) \mapsto e^{kt/2} \|w\|_2$. This produces

$$\begin{aligned} e^{kt/2} \|U_t\|_2 &= e^0 \|U_0\|_2 + \int_0^t k e^{ks} \|U_s\|_2 + e^{ks} \frac{d}{ds} \|U_s\|_2 ds \\ &= \int_0^t \frac{e^{ks/2}}{2 \|U_s\|_2} \left(k \|U_s\|_2^2 \right. \\ & \quad + \langle U_s, \nabla \log p(Z_{s,x'+\epsilon''v'+\epsilon v}) - \nabla \log p(Z_{s,x'+\epsilon''v'}) - (\nabla \log p(Z_{s,x'+\epsilon v}) - \nabla \log p(Z_{s,x'})) \rangle / (\epsilon \epsilon'') \\ & \quad \left. - \langle U_s, \nabla \log p(Z_{s,x+\epsilon'v'+\epsilon v}) - \nabla \log p(Z_{s,x+\epsilon'v'}) - (\nabla \log p(Z_{s,x+\epsilon v}) - \nabla \log p(Z_{s,x})) \rangle / (\epsilon \epsilon') \right) ds. \end{aligned}$$

Fix a value $s \in [0, t]$. For any $h_3 \in C^3(\mathcal{X})$, the Lemma 10 third-order difference inequality (14) and the coupling bounds (17) and (18) together imply the third-order function coupling bound (21),

$$\begin{aligned} & (h_3(Z_{s,x'+\epsilon''v'+\epsilon v}) - h_3(Z_{s,x'+\epsilon''v'}) - (h_3(Z_{s,x'+\epsilon v}) - h_3(Z_{s,x'})))/(\epsilon \epsilon'') \\ & \quad - (h_3(Z_{s,x+\epsilon'v'+\epsilon v}) - h_3(Z_{s,x+\epsilon'v'}) - (h_3(Z_{s,x+\epsilon v}) - h_3(Z_{s,x}))) / (\epsilon \epsilon') \\ & \leq \langle \nabla h_3(Z_{s,x'+\epsilon''v'}), U_s \rangle + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{L_3}{k} e^{-ks} (2\|x - x'\|_2 + \|x - x' + (\epsilon' - \epsilon'')v'\|_2) \\ & \quad + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{L_3}{k} e^{-ks} ((\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')) \\ & \quad + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} e^{-3ks/2} (\|x - x' + (\epsilon' - \epsilon'')v'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3). \\ & \leq \langle \nabla h_3(Z_{s,x'+\epsilon''v'}), U_s \rangle \\ & \quad + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{k} e^{-ks} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\ & \quad + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} e^{-3ks/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3), \end{aligned}$$

where we have applied the triangle inequality to achieve the final presentation. Applying this bound to the thrice continuously differentiable function $h_3(z) = \langle U_s, \nabla \log p(z) \rangle$ with

$$\|\nabla^2 h_3(z)\|_{op} = \|\nabla^3 \log p(z)[U_s]\|_{op} \leq L_3 \|U_s\|_2 \quad \text{and} \quad \|\nabla^3 h_3(z)\|_{op} \leq L_4 \|U_s\|_2$$

gives

$$\begin{aligned} & (h_3(Z_{s,x'+\epsilon''v'+\epsilon v}) - h_3(Z_{s,x'+\epsilon''v'}) - (h_3(Z_{s,x'+\epsilon v}) - h_3(Z_{s,x'})))/(\epsilon\epsilon'') \\ & - (h_3(Z_{s,x+\epsilon'v'+\epsilon v}) - h_3(Z_{s,x+\epsilon'v'}) - (h_3(Z_{s,x+\epsilon v}) - h_3(Z_{s,x}))) / (\epsilon\epsilon') \\ & \leq \langle U_s, \nabla^2 \log p(Z_{s,x'+\epsilon''v'}) U_s \rangle \\ & + \|U_s\|_2 \frac{3L_3^2}{k} e^{-ks} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\ & + \|U_s\|_2 L_4 e^{-3ks/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3). \\ & \leq -k \|U_s\|_2^2 + \|U_s\|_2 \frac{3L_3^2}{k} e^{-ks} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\ & + \|U_s\|_2 L_4 e^{-3ks/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3). \end{aligned}$$

In the final line, we used the k -strong log-concavity of p . We can now reproduce the target conclusion, since

$$\begin{aligned} e^{kt/2} \|U_t\|_2 & \leq \int_0^t \frac{3L_3^2}{2k} e^{-ks/2} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) ds \\ & + \int_0^t \frac{L_4}{2} e^{-ks} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3) ds \\ & \leq \frac{3L_3^2}{k^2} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\ & + \frac{L_4}{2k} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3). \end{aligned}$$

□

G.4 Proof of Theorem 7

By Lemma 9, for each $x \in \mathcal{X}$, the overdamped Langevin diffusion $(Z_{t,x})_{t \geq 0}$ is well-defined with stationary distribution P . Moreover, for each $x \in \mathcal{X}$, the diffusion $(Z_{t,x})_{t \geq 0}$, by definition, satisfies

$$dZ_{t,x} = \frac{1}{2} \nabla \log p(Z_{t,x}) dt + dW_t \quad \text{with} \quad Z_{0,x} = x,$$

for $(W_t)_{t \geq 0}$ a d -dimensional Wiener process. In what follows, when considering the joint distribution of a finite collection of overdamped Langevin diffusions, we will assume that the diffusions are coupled in the manner of Lemma 11, so that each diffusion is driven by a shared d -dimensional Wiener process $(W_t)_{t \geq 0}$.

Fix any $x \in \mathcal{X}$ and any $h \in C^3(\mathcal{X})$ with bounded first, second, and third derivatives. We divide the remainder of our proof into five components, establishing that u_h exists, u_h is Lipschitz, u_h has a Lipschitz gradient, u_h has a Lipschitz Hessian, and u_h solves the Stein equation (10).

Existence of u_h To see that the integral representation of $u_h(x)$ is well-defined, note that

$$\begin{aligned} \int_0^\infty |\mathbb{E}_P[h(Z)] - \mathbb{E}[h(Z_{t,x})]| dt & = \int_0^\infty \left| \int_{\mathcal{X}} \mathbb{E}[h(Z_{t,y})] - \mathbb{E}[h(Z_{t,x})] p(y) dy \right| dt \\ & \leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \int_0^\infty \int_{\mathcal{X}} \mathbb{E}[\|Z_{t,y} - Z_{t,x}\|_2] p(y) dy dt \\ & \leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \mathbb{E}_P[\|Z - x\|_2] \int_0^\infty e^{-kt/2} dt < \infty. \end{aligned}$$

The first relation uses the stationarity of P , the second uses the Lipschitz relation (11), the third uses the first-order coupling inequality (17) of Lemma 11, and the last uses the fact that log-concave distributions have subexponential tails and therefore finite moments of all orders [35, Lem. 1].

Lipschitz continuity of u_h We next show that u_h is Lipschitz. Fix any vector $v \in \mathcal{X}$, and consider the difference

$$\begin{aligned} |u_h(x+v) - u_h(x)| &= \left| \int_0^\infty \mathbb{E}[h(Z_{t,x}) - h(Z_{t,x+v})] dt \right| \\ &\leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \int_0^\infty \mathbb{E}[\|Z_{t,x} - Z_{t,x+v}\|_2] dt \\ &\leq \|v\|_2 \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \int_0^\infty e^{-kt/2} dt = \frac{2}{k} \|v\|_2 \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2. \end{aligned} \quad (22)$$

The second relation is an application of the Lipschitz relation (11), and the third applies the first-order coupling inequality (17) of Lemma 11.

Lipschitz continuity of ∇u_h To demonstrate that u_h is differentiable with Lipschitz gradient, we first establish a weighted second-order difference inequality for u_h .

Lemma 12. *For any vectors $x, x', v' \in \mathcal{X}$ with $\|v'\|_2 = 1$ and weights $\epsilon', \epsilon'' > 0$,*

$$\begin{aligned} &|(u_h(x' + \epsilon''v') - u_h(x'))/\epsilon'' - (u_h(x + \epsilon'v') - u_h(x))/\epsilon'| \\ &\leq (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2) \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{2L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} \frac{1}{k} \right). \end{aligned} \quad (23)$$

Proof Introduce the shorthand

$$V_t \triangleq (Z_{t,x'+\epsilon''v'} - Z_{t,x'})/\epsilon'' - (Z_{t,x+\epsilon'v'} - Z_{t,x})/\epsilon'.$$

We apply the Lemma 11 second-order function coupling inequality (20) (to the thrice continuously differentiable function h), the Cauchy-Schwarz inequality, and the second-order process bound (18) in turn to obtain

$$\begin{aligned} &|(u_h(x' + \epsilon''v') - u_h(x'))/\epsilon'' - (u_h(x + \epsilon'v') - u_h(x))/\epsilon'| \\ &= \left| \int_0^\infty \mathbb{E}[h(Z_{t,x'+\epsilon''v'}) - h(Z_{t,x'})]/\epsilon'' - \mathbb{E}[h(Z_{t,x+\epsilon'v'}) - h(Z_{t,x})]/\epsilon' dt \right| \\ &\leq \int_0^\infty \max(\mathbb{E}\langle \nabla h(Z_{t,x'}), V_t \rangle, \mathbb{E}\langle \nabla h(Z_{t,x}), V_t \rangle) + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} e^{-kt} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2) dt \\ &\leq \int_0^\infty \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \mathbb{E}[\|V_t\|_2] + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} e^{-kt} (\|x - x'\|_2 + (\epsilon' + \epsilon)/2) dt \\ &\leq (\|x - x'\|_2 + (\epsilon' + \epsilon)/2) \left| \int_0^\infty \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{L_3}{k} e^{-kt/2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} e^{-kt} dt \right| \\ &= (\|x - x'\|_2 + (\epsilon' + \epsilon)/2) \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{2L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} \frac{1}{k} \right). \end{aligned}$$

□

Now, fix any $x, v \in \mathcal{X}$ with $\|v\|_2 = 1$. As a first application of the Lemma 12 second-order difference inequality (23), we will demonstrate the existence of the directional derivative

$$\nabla_v u_h(x) \triangleq \lim_{\epsilon \rightarrow 0} \frac{u_h(x + \epsilon v) - u_h(x)}{\epsilon}. \quad (24)$$

Indeed, Lemma 12 implies that, for any integers $m, m' > 0$,

$$\begin{aligned} &|m'(u_h(x + v/m') - u_h(x)) - m(u_h(x + v/m) - u_h(x))| \\ &\leq \left(\frac{1}{2m} + \frac{1}{2m'} \right) \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{2L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} \frac{1}{k} \right). \end{aligned}$$

Hence, the sequence $\left(\frac{u_h(x+v/m)-u_h(x)}{1/m}\right)_{m=1}^{\infty}$ is Cauchy, and the directional derivative (24) exists.

To see that the directional derivative (24) is also Lipschitz, fix any $v' \in \mathcal{X}$, and consider the bound

$$\begin{aligned}
& |\nabla_v u_h(x+v') - \nabla_v u_h(x)| \leq \lim_{\epsilon \rightarrow 0} \left| \frac{u_h(x+\epsilon v+v') - u_h(x+v')}{\epsilon} - \frac{u_h(x+\epsilon v) - u_h(x)}{\epsilon} \right| \\
& \leq \lim_{\epsilon \rightarrow 0} (\|v'\|_2 + \epsilon) \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{2L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} \frac{1}{k} \right) \\
& = \|v'\|_2 \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{2L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} \frac{1}{k} \right), \tag{25}
\end{aligned}$$

where the second inequality follows from Lemma 12. Since each directional derivative is Lipschitz continuous, we may conclude that u_h is continuously differentiable with Lipschitz continuous gradient ∇u_h . Our Lipschitz function deduction (22) and the Lipschitz relation (11) additionally supply the uniform bound

$$\sup_{z \in \mathcal{X}} \|\nabla u_h(z)\|_2 \leq \frac{2}{k} \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2.$$

Lipschitz continuity of $\nabla^2 u_h$ To demonstrate that ∇u_h is differentiable with Lipschitz gradient, we begin by establishing a weighted third-order difference inequality for u_h .

Lemma 13. *For any vectors $x, x', v, v' \in \mathcal{X}$ with $\|v\|_2 = \|v'\|_2 = 1$ and weights $\epsilon, \epsilon', \epsilon'' > 0$,*

$$\begin{aligned}
& |(u_h(x'+\epsilon''v'+\epsilon v) - u_h(x'+\epsilon''v')) - (u_h(x'+\epsilon v) - u_h(x'))|/\epsilon\epsilon' \\
& - |(u_h(x+\epsilon'v'+\epsilon v) - u_h(x+\epsilon'v')) - (u_h(x+\epsilon v) - u_h(x))|/\epsilon\epsilon' \tag{26} \\
& \leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{6L_3^2}{k^3} (\|x-x'\|_2 + (\epsilon''+\epsilon')/2 + \epsilon(3+\epsilon/\epsilon''+\epsilon/\epsilon'+\|x-x'\|_2/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{L_4}{k^2} (\|x-x'\|_2 + 3(\epsilon''+\epsilon')/2 + \epsilon(3+\epsilon/\epsilon''+\epsilon/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{k^2} (\|x-x'\|_2 + (\epsilon''+\epsilon')/2 + \epsilon(3+\epsilon/\epsilon''+\epsilon/\epsilon'+\|x-x'\|_2/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} \frac{2}{3k} (\|x-x'\|_2 + 3(\epsilon''+\epsilon')/2 + \epsilon(3+\epsilon/\epsilon''+\epsilon/\epsilon')/3).
\end{aligned}$$

Proof Introduce the shorthand

$$\begin{aligned}
U_t \triangleq & (Z_{t,x'+\epsilon''v'+\epsilon v} - Z_{t,x'+\epsilon''v'} - (Z_{t,x'+\epsilon v} - Z_{t,x'}))/(\epsilon\epsilon'') \\
& - (Z_{t,x+\epsilon'v'+\epsilon v} - Z_{t,x+\epsilon'v'} - (Z_{t,x+\epsilon v} - Z_{t,x}))/(\epsilon\epsilon')
\end{aligned}$$

We apply the Lemma 11 third-order function coupling inequality (21) (to the thrice continuously differentiable function h), the Cauchy-Schwarz inequality, and the third-order process bound (19) in

turn to obtain

$$\begin{aligned}
& |(u_h(x' + \epsilon''v' + \epsilon v) - u_h(x' + \epsilon''v') - (u_h(x' + \epsilon v) - u_h(x')))/\epsilon\epsilon'| \\
& - (u_h(x + \epsilon'v' + \epsilon v) - u_h(x + \epsilon'v') - (u_h(x + \epsilon v) - u_h(x)))/\epsilon\epsilon'| \\
= & \left| \int_0^\infty \mathbb{E}[(h_3(Z_{t,x'+\epsilon''v'+\epsilon v}) - h_3(Z_{t,x'+\epsilon''v'}) - (h_3(Z_{t,x'+\epsilon v}) - h_3(Z_{t,x'})))]/(\epsilon\epsilon'') \right. \\
& \left. - \mathbb{E}[(h_3(Z_{t,x+\epsilon'v'+\epsilon v}) - h_3(Z_{t,x+\epsilon'v'}) - (h_3(Z_{t,x+\epsilon v}) - h_3(Z_{t,x})))]/(\epsilon\epsilon') dt \right| \\
\leq & \int_0^\infty \max(\mathbb{E}\langle \nabla h_3(Z_{t,x'+\epsilon''v'}), U_t \rangle, \mathbb{E}\langle \nabla h_3(Z_{t,x+\epsilon'v'}), U_t \rangle) \\
& + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{k} e^{-kt} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} e^{-3kt/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3) dt \\
\leq & \int_0^\infty \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{3L_3^2}{k^2} e^{-kt/2} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{L_4}{2k} e^{-kt/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{k} e^{-kt} (\|x - x'\|_2 + (\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon' + \|x - x'\|_2/\epsilon')/3) \\
& + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} e^{-3kt/2} (\|x - x'\|_2 + 3(\epsilon'' + \epsilon')/2 + \epsilon(3 + \epsilon/\epsilon'' + \epsilon/\epsilon')/3) dt.
\end{aligned}$$

Integrating this final expression yields the advertised bound. \square

Now, fix any $x, v, v' \in \mathcal{X}$ with $\|v\|_2 = \|v'\|_2 = 1$. As a first application of the Lemma 13 third-order difference inequality (26), we will demonstrate the existence of the second-order directional derivative

$$\begin{aligned}
\nabla_{v'} \nabla_v u_h(x) & \triangleq \lim_{\epsilon' \rightarrow 0} \frac{\nabla_{v'} u_h(x + \epsilon'v') - \nabla_{v'} u_h(x)}{\epsilon'} \\
& = \lim_{\epsilon' \rightarrow 0} \lim_{\epsilon \rightarrow 0} \frac{u_h(x + \epsilon'v' + \epsilon v) - u_h(x + \epsilon v) - (u_h(x + \epsilon'v') - u_h(x))}{\epsilon\epsilon'}.
\end{aligned} \tag{27}$$

Lemma 13 guarantees that, for any integers $m, m' > 0$,

$$\begin{aligned}
& |m'(\nabla_{v'} u_h(x + v'/m') - \nabla_{v'} u_h(x)) - m(\nabla_v u_h(x + v'/m) - \nabla_v u_h(x))| \\
\leq & \lim_{\epsilon \rightarrow 0} |m'(u_h(x + v'/m' + v\epsilon) - u_h(x + v'/m') - (u_h(x + v\epsilon) - u_h(x)))/\epsilon \\
& - m(u_h(x + v'/m + v\epsilon) - u_h(x + v'/m) - (u_h(x + v\epsilon) - u_h(x)))/\epsilon| \\
\leq & \left(\frac{1}{m} + \frac{1}{m'} \right) \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \left(\frac{3L_3^2}{k^3} + \frac{3L_4}{2k^2} \right) + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{2k^2} + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} \frac{1}{k} \right).
\end{aligned}$$

Hence, the sequence $\left(\frac{\nabla_v u_h(x + v'/m) - \nabla_v u_h(x)}{1/m} \right)_{m=1}^\infty$ is Cauchy, and the directional derivative (27) exists.

To see that the directional derivative (27) is also Lipschitz, fix any $v'' \in \mathcal{X}$, and consider the bound

$$\begin{aligned}
& |\nabla_{v'} \nabla_v u_h(x + v'') - \nabla_{v'} \nabla_v u_h(x)| \\
& \leq \lim_{\epsilon' \rightarrow 0} \left| \frac{\nabla_v u_h(x + v'' + \epsilon' v') - \nabla_v u_h(x + v'')}{\epsilon'} - \frac{\nabla_v u_h(x + \epsilon' v') - \nabla_v u_h(x)}{\epsilon'} \right| \\
& \leq \lim_{\epsilon' \rightarrow 0} \lim_{\epsilon \rightarrow 0} \left| \frac{u_h(x + v'' + \epsilon' v' + \epsilon v) - u_h(x + v'' + \epsilon v) - (u_h(x + v'' + \epsilon' v') - u_h(x + v''))}{\epsilon \epsilon'} \right. \\
& \quad \left. - \frac{u_h(x + \epsilon' v' + \epsilon v) - u_h(x + \epsilon v) - (u_h(x + \epsilon' v') - u_h(x))}{\epsilon \epsilon'} \right| \\
& \leq \|v''\|_2 \left(\sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \left(\frac{6L_3^2}{k^3} + \frac{L_4}{k^2} \right) + \sup_{z \in \mathcal{X}} \|\nabla^2 h_3(z)\|_{op} \frac{3L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^3 h_3(z)\|_{op} \frac{2}{3k} \right),
\end{aligned}$$

where the final inequality follows from Lemma 13. Since each second-order directional derivative is Lipschitz continuous, we conclude that $u_h \in C^2(\mathcal{X})$ with Lipschitz continuous Hessian $\nabla^2 u_h$. Our Lipschitz gradient result (25) and the Lipschitz relation (12) further furnish the uniform bound

$$\sup_{z \in \mathcal{X}} \|\nabla^2 u_h(z)\|_{op} \leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \frac{2L_3}{k^2} + \sup_{z \in \mathcal{X}} \|\nabla^2 h(z)\|_{op} \frac{1}{k}.$$

Solving the Stein equation Finally, we show that u_h solves the Stein equation (10). Introduce the notation $(P_t h)(x) \triangleq \mathbb{E}[h(Z_{t,x})]$. Since $(Z_{t,x})_{t \geq 0}$ is strong Feller, its generator \mathcal{A} , defined in (9), satisfies

$$h - P_t h = \mathcal{A} \int_0^t \mathbb{E}_P[h(Z)] - P_s h \, ds$$

for all t by [36, Prop. 1.5]. The left-hand side limits (pointwise) to $h - \mathbb{E}_P[h(Z)]$ as $t \rightarrow \infty$, as

$$\begin{aligned}
|h(x) - \mathbb{E}_P[h(Z)] - (h(x) - (P_t h)(x))| &= \left| \int_{\mathcal{X}} \mathbb{E}[h(Z_{t,y})] - \mathbb{E}[h(Z_{t,x})] p(y) dy \right| \\
&\leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \int_{\mathcal{X}} \mathbb{E}[\|Z_{t,y} - Z_{t,x}\|_2] p(y) dy \\
&\leq \sup_{z \in \mathcal{X}} \|\nabla h(z)\|_2 \mathbb{E}_P[\|Z - x\|_2] e^{-kt/2}
\end{aligned}$$

for each $x \in \mathcal{X}$ and $t \geq 0$. Here we have used the stationarity of P , the Lipschitz relation (11), the first-order coupling inequality (17) of Lemma 11, and the integrability of Z [35, Lem. 1] in turn. Meanwhile, the right-hand side limits to $\mathcal{A}u_h$, since \mathcal{A} is closed [36, Cor. 1.6]. Therefore, u_h solves the Stein equation (10).

H Finite-dimensional Classical Stein Program

Theorem 14 (Finite-dimensional Classical Stein Program). *If $\mathcal{X} = (\alpha, \beta)$ for $-\infty \leq \alpha < \beta \leq \infty$, and $x_{(1)} < \dots < x_{(n')}$ represent the sorted values of $\{x_1, \dots, x_n, \alpha, \beta\} \cap \mathbb{R}$, then the non-uniform classical Stein discrepancy $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}^{c_1;3})$ is the optimal value of the convex program*

$$\max_g \sum_{i=1}^{n'} q(x_{(i)}) \frac{d}{dx} \log p(x_{(i)}) g(x_{(i)}) + q(x_{(i)}) g'(x_{(i)}) \tag{28a}$$

$$s.t. \quad \forall i \in \{1, \dots, n' - 1\}, |g'(x_{(i)})| \leq c_2, |g(x_{(i+1)}) - g(x_{(i)})| \leq c_2(x_{(i+1)} - x_{(i)}), \tag{28b}$$

$$\begin{aligned}
g(x_{(i)}) - g(x_{(i+1)}) + \frac{1}{4c_3} (g'(x_{(i)}) - g'(x_{(i+1)}))^2 + \frac{x_{(i+1)} - x_{(i)}}{2} (g'(x_{(i)}) + g'(x_{(i+1)})) \\
+ \frac{1}{c_3} (L_b)_+^2 \leq \frac{c_3}{4} (x_{(i+1)} - x_{(i)})^2, \tag{28c}
\end{aligned}$$

$$\begin{aligned}
g(x_{(i+1)}) - g(x_{(i)}) + \frac{1}{4c_3} (g'(x_{(i)}) - g'(x_{(i+1)}))^2 - \frac{x_{(i+1)} - x_{(i)}}{2} (g'(x_{(i)}) + g'(x_{(i+1)})) \\
+ \frac{1}{c_3} (L_u)_+^2 \leq \frac{c_3}{4} (x_{(i+1)} - x_{(i)})^2, \quad \text{and} \tag{28d}
\end{aligned}$$

$$\forall i \in \{1, \dots, n'\}, |g(x_{(i)})| \leq \mathbb{I}[\alpha < x_{(i)} < \beta] (c_1 - \frac{1}{2c_3} g'(x_{(i)})^2) \tag{28e}$$

where $(r)_+ \triangleq \max(r, 0)$,

$$\begin{aligned} L_b &\triangleq \frac{c_3}{2}(x_{(i+1)} - x_{(i)}) - \frac{1}{2}(g'(x_{(i)}) + g'(x_{(i+1)})) - c_2, \quad \text{and} \\ L_u &\triangleq \frac{c_3}{2}(x_{(i+1)} - x_{(i)}) + \frac{1}{2}(g'(x_{(i)}) + g'(x_{(i+1)})) - c_2. \end{aligned}$$

We say the program (28) is finite-dimensional, because it suffices to optimize over vectors $\gamma, \Gamma \in \mathbb{R}^{n'}$ representing the function values ($\gamma_i = g(x_{(i)})$) and derivative values ($\Gamma_i = g'(x_{(i)})$) at each sample or boundary point $x_{(i)}$. Indeed, by introducing slack variables, this program is representable as a convex quadratically constrained quadratic program with $O(n)$ constraints, $O(n)$ variables, and a linear objective. Moreover, the pairwise constraints in this program are only enforced between neighboring points in the sequence of ordered locations $x_{(i)}$. Hence the resulting constraint matrix is sparse and banded, making the problem particularly amenable to efficient optimization.

Proof Throughout, we say that \tilde{g} is an extension of g if $\tilde{g}(x_{(i)}) = g(x_{(i)})$ and $\tilde{g}'(x_{(i)}) = g'(x_{(i)})$ for each $x_{(i)} \in \text{supp}(Q)$. Since the Stein objective only depends on g and g' through their values at sample points, g and any extension \tilde{g} have identical objective values.

We will establish our result by showing that every $g \in \mathcal{G}_{\|\cdot\|}^{c_1:3}$ is feasible for the program (28), so that $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}^{c_1:3})$ lower bounds the optimum of (28), and that every feasible g for (28) has an extension in $\tilde{g} \in \mathcal{G}_{\|\cdot\|}^{c_1:3}$, so that $\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|}^{c_1:3})$ also upper bounds the optimum of (28).

H.1 Feasibility of $\mathcal{G}_{\|\cdot\|}^{c_1:3}$

Fix any $g \in \mathcal{G}_{\|\cdot\|}^{c_1:3}$. Also, since g' is c_2 -bounded and c_3 -Lipschitz, the constraints (28b) must be satisfied. Consider now the c_2 -bounded and c_3 -Lipschitz extensions of g'

$$\begin{aligned} B(t) &\triangleq \max(-c_2, \max_{1 \leq i \leq n'} [g'(x_{(i)}) - c_3|t - x_{(i)}|]) \quad \text{and} \\ U(t) &\triangleq \min(c_2, \min_{1 \leq i \leq n'} [g'(x_{(i)}) + c_3|t - x_{(i)}|]). \end{aligned}$$

We know that $B(t) \leq g'(t) \leq U(t)$ for all t , for, if not, there would be a point t_0 and a point $x_{(i)}$ such that $|g'(x_{(i)}) - g'(t_0)| > c_3|x_{(i)} - t_0|$, which combined with the c_3 -Lipschitz property would be a contradiction. Thus, for each sample $x_{(i)}$, the fundamental theorem of calculus gives

$$g(x_{(i+1)}) - g(x_{(i)}) = \int_{x_{(i)}}^{x_{(i+1)}} g'(t) dt \geq \int_{x_{(i)}}^{x_{(i+1)}} B(t) dt.$$

The right-hand side of this inequality evaluates precisely to the right-hand side of the constraint (28c). An analogous upper bound using $U(t)$ yields (28d).

Finally, consider any point $x_{(i)}$. If $x_{(i)} \in \{\alpha, \beta\}$, then (28e) is satisfied as $g(z) = 0$ for any point z on the boundary. Suppose instead that $\alpha < x_{(i)} < \beta$. Without loss of generality, we may assume that $g'(x_{(i)}) \geq 0$. Since g' is c_3 -Lipschitz, we have $g'(t) \geq g'(x_{(i)}) - c_3|t - x_{(i)}|$ for all t . Integrating both sides of this inequality from $x_{(i)}$ to $x_u = x_{(i)} + g'(x_{(i)})/c_3$, we obtain

$$g(x_u) - g(x_{(i)}) = \int_{x_{(i)}}^{x_u} g'(t) dt \geq \int_{x_{(i)}}^{x_u} g'(x_{(i)}) - c_3(t - x_{(i)}) dt = g'(x_{(i)})^2 / (2c_3)$$

Since $g(x_u) \leq c_1$, we have $\frac{1}{2c_3}g'(x_{(i)})^2 + g(x_{(i)}) \leq c_1$. Similarly, by integrating the inequality from $x_b = x_{(i)} - g'(x_{(i)})/c_3$ to $x_{(i)}$, we have $g(x_b) - g(x_{(i)}) \geq g'(x_{(i)})^2 / (2c_3)$, which combined with $g(x_b) \leq c_1$ yields (28e).

H.2 Extending Feasible Solutions

Suppose now that g is any function feasible for the program (28). We will construct an extension $\tilde{g} \in \mathcal{G}_{\|\cdot\|}^{c_1:3}$ by first working independently over each interval $(x_{(i)}, x_{(i+1)})$. Fix an index $i < n'$. Our strategy is to identify a pair of c_2 -bounded, c_3 -Lipschitz functions m_i and M_i defined on the interval $[x_{(i)}, x_{(i+1)}]$ which satisfy $m_i(x) \leq M_i(x)$ for all $x \in [x_{(i)}, x_{(i+1)}]$, $m_i(x) = M_i(x) = g'(x)$ for

$x \in \{x_{(i)}, x_{(i+1)}\}$, and $\int_{x_{(i)}}^{x_{(i+1)}} m_i(t) dt \leq g(x_{(i+1)}) - g(x_{(i)}) \leq \int_{x_{(i)}}^{x_{(i+1)}} M_i(t) dt$. For any such (m_i, M_i) pair, there exists $\zeta_i \in [0, 1]$ satisfying

$$g(x_{(i+1)}) - g(x_{(i)}) = \int_{x_{(i)}}^{x_{(i+1)}} \zeta_i m_i(t) + (1 - \zeta_i) M_i(t) dt,$$

and hence we will define the extension

$$\tilde{g}(x) = g(x_{(i)}) + \int_{x_{(i)}}^x \zeta_i m_i(t) + (1 - \zeta_i) M_i(t) dt.$$

By convexity, the extension derivative \tilde{g}' is c_2 -bounded and c_3 -Lipschitz, so we will only need to check that $\sup_{x \in \mathcal{X}} |\tilde{g}(x)| \leq c_1$. The maximum magnitude values of \tilde{g} occur either at the interval endpoints, which are c_1 -bounded by (28e), or at critical points x satisfying $\tilde{g}'(x) = 0$, so it suffices to ensure that \tilde{g} is c_1 -bounded at all critical points.

We will use the c_2 -bounded, c_3 -Lipschitz functions B and U as building blocks for our extension, since they satisfy $B(t) = U(t) = g'(t)$ for $t \in \{x_{(i)}, x_{(i+1)}\}$ and $B(t) \leq g'(t) \leq U(t)$,

$$\begin{aligned} B(t) &= \max(-c_2, g'(x_{(i)}) - c_3(t - x_{(i)}), g'(x_{(i+1)}) - c_3(x_{(i+1)} - t)), \quad \text{and} \\ U(t) &= \min(c_2, g'(x_{(i)}) + c_3(t - x_{(i)}), g'(x_{(i+1)}) + c_3(x_{(i+1)} - t)), \end{aligned}$$

for $t \in [x_{(i)}, x_{(i+1)}]$. We need only consider three cases.

Case 1: B and U are never negative or never positive on $[x_{(i)}, x_{(i+1)}]$. For this case, we will choose $m_i = B$ and $M_i = U$. By (28c) and (28d) we know $\int_{x_{(i)}}^{x_{(i+1)}} m_i(t) dt \leq g(x_{(i+1)}) - g(x_{(i)}) \leq \int_{x_{(i)}}^{x_{(i+1)}} M_i(t) dt$. Since B and U never change signs, \tilde{g} will be monotonic and hence c_1 -bounded for any choice of ζ_i .

Case 2: Exactly one of B and U changes sign on $[x_{(i)}, x_{(i+1)}]$. Without loss of generality, we may assume that $g'(x_{(i)}), g'(x_{(i+1)}) \geq 0$ and that B changes sign. Consider the quantity $\phi \triangleq \int_{x_{(i)}}^{x_{(i+1)}} \max\{B(t), 0\} dt$. If $g(x_{(i+1)}) - g(x_{(i)}) \leq \phi$, we let $m_i = B$ and $M_i = \max\{B, 0\}$.

Since, on the interval $[x_{(i)}, x_{(i+1)}]$, B is piecewise linear with at most two pieces that can take on the value 0, B has at most two roots within this interval. However, since $B(x)$ is continuous, negative for some value of x , and nonnegative at $x \in \{x_{(i)}, x_{(i+1)}\}$, we know B has at least two roots. Thus let $r_1 < r_2$ be the roots of $B(x)$. For any choice of ζ_i , the convex combination $\zeta_i m_i + (1 - \zeta_i) M_i$ will be exactly B outside (r_1, r_2) . Moreover, if $\zeta_i \neq 0$, then this combination will be less than 0 on (r_1, r_2) , and if $\zeta_i = 0$, the combination will be 0 on the whole interval. Hence it suffices to only check the critical points r_1 and r_2 . By (28e), $m_i(r) = M_i(r) = B(r) \in [-c_1, c_1]$ for $r \in \{r_1, r_2\}$, and so \tilde{g} will be c_1 -bounded.

If instead $g(x_{(i+1)}) - g(x_{(i)}) > \phi$, we can recycle the argument from Case 1 with $m_i = \max\{B, 0\}$ and $M_i = U$ and conclude that \tilde{g} is c_1 -bounded.

Case 3: Both B and U change sign on $[x_{(i)}, x_{(i+1)}]$. Without loss of generality, we may assume that $g'(x_{(i)}) \geq 0, g'(x_{(i+1)}) < 0$. Since B continuously interpolates between $g'(x_{(i)})$ and $g'(x_{(i+1)})$ on $[x_{(i)}, x_{(i+1)}]$, it must have a root r . Let $w_i \in [x_{(i)}, x_{(i+1)}]$ be the point where B changes from one linear portion to another. Then because B is monotonic on each linear portion, the fact that $B(w_i) \leq B(x_{(i+1)}) < 0$ means that B cannot have a root between $[w_i, x_{(i+1)}]$ and hence has at most one root on $[x_{(i)}, x_{(i+1)}]$. Hence r is the unique root of B .

In a similar fashion, let us define s as the root of U , and since $B(x) \leq U(x)$ for all x , we have $s \geq r$. Define

$$W(x) \triangleq \begin{cases} B(x) & x \in [x_{(i)}, r) \\ 0 & x \in [r, s] \\ U(x) & t \in (s, y], \end{cases}$$

and $\psi \triangleq \int_{x_{(i)}}^{x_{(i+1)}} W(t)dt$. As in Case 2, we will consider two subcases. If $g(x_{(i+1)}) - g(x_{(i)}) \leq \psi$, we will let $m_i = B$ and $M_i = W$. By (28e), $m_i(r) = M_i(r) = B(r) \in [-c_1, c_1]$, and since this is the only critical point, \tilde{g} will be c_1 -bounded.

For the other case, in which $g(x_{(i+1)}) - g(x_{(i)}) > \psi$, we choose $m_i = W$ and $M_i = U$. Then (28e) imply that $m_i(s) = M_i(s) = U(s) \in [-c_1, c_1]$, and, since this is the only critical point, the extension is well-defined on $(x_{(i)}, x_{(i+1)})$.

Defining \tilde{g} outside of the interval $[x_1, x_{n'}]$ It only remains to define our extension \tilde{g} outside of the interval $[x_1, x_{n'}]$ when either α or β is infinite. Suppose $\alpha = -\infty$. We extend \tilde{g} to each $x \in (-\infty, x_1)$ using the construction

$$\tilde{g}(x) \triangleq \int_{-\infty}^x \mathbb{I}[t \in (x_1 - |g'(x_1)|/c_3, x_1)](g'(x_1) - c_3 \text{sign}(g'(x_1))t) dt.$$

This extension ensures that \tilde{g}' is c_2 -bounded and c_3 -Lipschitz. Moreover, the constraint (28e) guarantees that $|\tilde{g}(x)| \leq c_1$. Analogous reasoning establishes an extension to $(x_{n'}, \infty)$. \square

I Equivalence of Constrained Classical and Spanner Stein Discrepancies

For P with support $\mathcal{X} = (\alpha_1, \beta_1) \times \dots \times (\alpha_d, \beta_d)$ for $-\infty \leq \alpha_j < \beta_j \leq \infty$, Algorithm 1 computes a Stein discrepancy based on the graph Stein set

$$\begin{aligned} \mathcal{G}_{\|\cdot\|_1, Q, (V, E)} \triangleq & \left\{ g : \mathcal{X} \rightarrow \mathbb{R}^d \mid \forall x \in V, j, k \in \{1, \dots, d\} \text{ with } k \neq j, \text{ and } b_j \in \{\alpha_j, \beta_j\} \cap \mathbb{R}, \right. \\ & \max \left(\|g(x)\|_\infty, \|\nabla g(x)\|_\infty, \frac{|g_j(x)|}{|x_j - b_j|}, \frac{|\nabla_k g_j(x)|}{|x_j - b_j|}, \frac{|g_j(x) - \nabla_j g_j(x)(x_j - b_j)|}{\frac{1}{2}(x_j - b_j)^2} \right) \leq 1, \text{ and, } \forall (x, y) \in E, \\ & \left. \max \left(\frac{\|g(x) - g(y)\|_\infty}{\|x - y\|_1}, \frac{\|\nabla g(x) - \nabla g(y)\|_\infty}{\|x - y\|_1}, \frac{\|g(x) - g(y) - \nabla g(x)(x - y)\|_\infty}{\frac{1}{2}\|x - y\|_1^2}, \frac{\|g(x) - g(y) - \nabla g(y)(x - y)\|_\infty}{\frac{1}{2}\|x - y\|_1^2} \right) \leq 1 \right\}, \end{aligned}$$

Our next result shows that the graph Stein discrepancy based on a t -spanner is strongly equivalent to the classical Stein discrepancy.

Proposition 15 (Equivalence of Constrained Classical and Spanner Stein Discrepancies). *If $\mathcal{X} = (\alpha_1, \beta_1) \times \dots \times (\alpha_d, \beta_d)$, and $G_t = (\text{supp}(Q), E)$ is a t -spanner, then*

$$\mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|_1}) \leq \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|_1, Q, G_t}) \leq t^2 \kappa_d \mathcal{S}(Q, \mathcal{T}_P, \mathcal{G}_{\|\cdot\|_1}),$$

where κ_d is a constant, independent of (Q, P, G_t, t) , depending only on the dimension d .

Proof

Establishing the first inequality Fix any $g \in \mathcal{G}_{\|\cdot\|_1}$, $z \in \text{supp}(Q)$, and $j, k \in \{1, \dots, d\}$ with $k \neq j$, and consider any j -th coordinate boundary projection point

$$b \in \{z + e_j(\alpha_j - z_j), z + e_j(\beta_j - z_j)\} \cap \mathbb{R}^d.$$

Since $b \in \partial\mathcal{X}$, we must have $\langle g(b), n(b) \rangle = \langle g(b), e_j \rangle = g_j(b) = 0$. Moreover, for each dimension $k \neq j$, we have $\nabla_k g_j(x) = 0$, since otherwise, $\langle g(b + \delta e_k), n(b + \delta e_k) \rangle = g_j(b + \delta e_k) \neq 0$ for some $\delta \in \mathbb{R}$ and $b + \delta e_k \in \partial\mathcal{X}$ by the continuity of ∇g_j .

The smoothness constraints of the classical Stein set $\mathcal{G}_{\|\cdot\|_1}$ now imply that

$$|g_j(z)| = |g_j(z) - g_j(b)| \leq |z_j - b_j|, \quad |\nabla_k g_j(x)| = |\nabla_k g_j(z) - \nabla_k g_j(b)| \leq |z_j - b_j|,$$

and

$$|g_j(z) - \nabla_j g_j(x)(z_j - b_j)| = |g_j(b) - g_j(z) - \langle \nabla g_j(z), b - z \rangle| \leq \frac{1}{2}(z_j - b_j)^2$$

so that all graph Stein set boundary compatibility constraints are satisfied. Hence, we have the containment $\mathcal{G}_{\|\cdot\|_1} \subseteq \mathcal{G}_{\|\cdot\|_1, Q, G_t}$, which implies the first advertised inequality.

Establishing the second inequality To establish the second inequality, it suffices to show that for any $\tilde{g} \in \mathcal{G}_{\|\cdot\|_1, Q, G_t}$, each $j \in \{1, \dots, d\}$, and $\zeta \triangleq t$, there exists a function g_j satisfying

$$g_j(z) = \tilde{g}_j(z), \quad \nabla g_j(z) = \nabla \tilde{g}_j(z), \quad g_j(b) = 0, \quad \nabla_k g_j(b) = 0, \quad \forall k \neq j, \quad (29)$$

$$|g_j(b) - g_j(z)| \leq \|b - z\|_1, \quad (30)$$

$$\|\nabla g_j(b) - \nabla g_j(z)\|_\infty \leq \zeta \|b - z\|_1, \quad \|\nabla g_j(b) - \nabla g_j(b')\|_\infty \leq \zeta \|b - b'\|_1, \quad (31)$$

$$|g_j(b) - g_j(z) - \langle \nabla g_j(z), b - z \rangle| \leq \frac{\zeta}{2} \|b - z\|_1^2, \quad (32)$$

$$|g_j(z) - g_j(b) - \langle \nabla g_j(b), z - b \rangle| \leq \frac{3\zeta}{2} \|b - z\|_1^2, \quad \text{and} \quad (33)$$

$$|g_j(b) - g_j(b') - \langle \nabla g_j(b'), b - b' \rangle| \leq \frac{\zeta}{2} \|b - b'\|_1^2 \quad (34)$$

for all $z \in \text{supp}(Q)$ and all b, b' in the j -th coordinate boundary set

$$B_j \triangleq \{b \in \mathbb{R}^d : b = z + e_j(\alpha_j - z_j) \text{ or } b = z + e_j(\beta_j - z_j) \text{ for some } z \in \mathcal{X}\}.$$

Indeed, since such g_j will satisfy $\max(|g_j(z)|, \|\nabla g_j(z)\|_\infty) \leq 1$ for all $z \in \text{supp}(Q) \cup B_j$ and

$$\max\left(\frac{|g_j(x) - g_j(y)|}{\|x - y\|_1}, \frac{\|\nabla g_j(x) - \nabla g_j(y)\|_\infty}{\|x - y\|_1}, \frac{|g_j(x) - g_j(y) - \nabla g_j(x)(x - y)|}{\frac{1}{2}\|x - y\|_1^2}, \frac{|g_j(x) - g_j(y) - \nabla g_j(y)(x - y)|}{\frac{1}{2}\|x - y\|_1^2}\right) \leq 2t^2$$

for all $x, y \in \text{supp}(Q)$ by the argument of Appendix F, the Whitney-Glaeser extension theorem [15, Thm. 1.4] of Glaeser [14] will then imply that there exists $g^* \in t^2 \kappa_d \mathcal{G}_{\|\cdot\|_1}$, for a constant κ_d independent of \tilde{g} depending only on d , with $g^*(z) = g(z)$ and $\nabla g^*(z) = \nabla g(z)$ for all $z \in \text{supp}(Q)$. Since \tilde{g} and g^* will have matching Stein discrepancy objective values, and each objective is linear in g , the second advertised inequality will then follow.

Fix $\tilde{g} \in \mathcal{G}_{\|\cdot\|_1, Q, G_t}$ and $j \in \{1, \dots, d\}$. We will now construct a function g_j satisfying the desired properties. Since g_j and ∇g_j are determined on $\text{supp}(Q)$, and g_j and $\nabla_k g_j$ are determined on B_j for $k \neq j$ by the constraints (29), it remains to define $\nabla_j g_j$ on B_j . We choose the extension

$$\nabla_j g_j(b) \triangleq \min_{z \in \text{supp}(Q)} \{\nabla_j g_j(z) + \zeta \|z - b\|_1\} \quad \text{for all } b \in B_j.$$

Fix any $z \in \text{supp}(Q)$ and $b \in B_j$, and let $b^* = z + e_j(b_j - z_j)$. The argument of Appendix F implies that $\nabla_j g_j$ is ζ -Lipschitz on $\text{supp}(Q)$, and hence it is also ζ -Lipschitz on $\text{supp}(Q) \cup B_j$. Since

$$|\nabla_k g_j(z) - \nabla_k g_j(b)| = |\nabla_k g_j(z)| \leq |z_j - b_j| \leq \|z - b\|_1$$

for all $k \neq j$, we have (31). Moreover, the boundary compatibility constraints of $\mathcal{G}_{\|\cdot\|_1, Q, G_t}$ imply

$$|g_j(b) - g_j(z)| = |g_j(z)| \leq \|b^* - z\|_1 \leq \|b - z\|_1,$$

establishing (30). We next invoke the triangle inequality, the boundary compatibility conditions of $\mathcal{G}_{\|\cdot\|_1, Q, G_t}$, Hölder's inequality, the Lipschitz derivative property (31), and the fact $\|z - b\|_1 = \|b^* - z\|_1 + \|b^* - b\|_1$ in turn to establish (32):

$$\begin{aligned} |g_j(b) - g_j(z) - \langle \nabla g_j(z), b - z \rangle| &= |g_j(z) - \nabla_j g_j(z)(z_j - b_j) - \langle \nabla g_j(z), b^* - b \rangle| \\ &\leq |g_j(z) - \nabla_j g_j(z)(z_j - b_j)| + |\langle \nabla g_j(b^*) - \nabla g_j(z), b^* - b \rangle| \\ &\leq \frac{1}{2} \|b^* - z\|_1^2 + \|\nabla g_j(b^*) - \nabla g_j(z)\|_\infty \|b^* - b\|_1 \\ &\leq \frac{1}{2} \|b^* - z\|_1^2 + \zeta \|b^* - z\|_1 \|b^* - b\|_1 \\ &\leq \frac{\zeta}{2} (\|b^* - z\|_1 + \|b^* - b\|_1)^2 = \frac{\zeta}{2} \|b - z\|_1^2. \end{aligned}$$

A parallel argument yields (34). Finally, we may deduce (33), as

$$\begin{aligned} |g_j(z) - g_j(b) - \langle \nabla g_j(b), z - b \rangle| &\leq |g_j(z) - \nabla_j g_j(z)(z_j - b_j)| + |\nabla_j g_j(b) - \nabla_j g_j(z)| |z_j - b_j| \\ &\leq \frac{1}{2} (z_j - b_j)^2 + \zeta \|b - z\|_1 |z_j - b_j| \leq \frac{3\zeta}{2} \|b - z\|_1^2 \end{aligned}$$

by the triangle inequality, the definition of $\mathcal{G}_{\|\cdot\|_1, Q, G_t}$, and the Lipschitz property (31). \square

Acknowledgments

The authors would like to thank Madeleine Udell for her generous advice concerning optimization in Julia, Quirijn Bouts and Kevin Buchin for sharing their wise counsel and greedy spanner implementation, Francis Bach for sharing his pseudosampling code, Andreas Eberle for his triple coupling pointers, and Jessica Hwang for her feedback on various versions of this manuscript.

This work was supported by the Frederick E. Terman Fellowship and the National Science Foundation Graduate Research Fellowship under Grant No. DGE-114747.

References

- [1] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng. *Handbook of Markov chain monte carlo*. CRC press, 2011.
- [2] C. J. Geyer. Markov chain monte carlo maximum likelihood. *Computer Science and Statistics: Proc. 23rd Symp. Interface*, pages 156–163, 1991.
- [3] M. Welling and Y.-W. Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning*, pages 681–688, 2011.
- [4] S. Ahn, A. Korattikara, and M. Welling. Bayesian posterior sampling via stochastic gradient fisher scoring. In *Proceeding of 29th International Conference on Machine Learning (ICML’12)*, 2012.
- [5] A. Korattikara, Y. Chen, and M. Welling. Austerity in MCMC land: Cutting the Metropolis-Hastings budget. In *Proceeding of 31th International Conference on Machine Learning (ICML’14)*, 2014.
- [6] A. Müller. Integral probability metrics and their generating classes of functions. *Advances in Applied Probability*, 29(2):pp. 429–443, 1997.
- [7] C. Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 2: Probability Theory*, pages 583–602, Berkeley, CA, 1972. University of California Press.
- [8] A. D. Barbour. Stein’s method and Poisson process convergence. *J. Appl. Probab.*, (Special Vol. 25A): 175–184, 1988. A celebration of applied probability.
- [9] L. HY. Chen, L. Goldstein, and Q.-M. Shao. *Normal approximation by Steins method*. Springer Science & Business Media, 2010.
- [10] S. Chatterjee and Q.-M. Shao. Nonnormal approximation by Steins method of exchangeable pairs with application to the Curie–Weiss model. *Annals of Applied Probability*, 21(2):464–483, 2011.
- [11] G. Reinert and A. Röllin. Multivariate normal approximation with Steins method of exchangeable pairs under a general linearity condition. *Annals of Probability*, 37(6):2150–2173, 2009.
- [12] S. Chatterjee and E. Meckes. Multivariate normal approximation using exchange-able pairs. *Alea*, 4: 257–283, 2008.
- [13] E. Meckes. On Steins method for multivariate normal approximation. In *High dimensional probability V: The Luminy volume*, pages 153–178. Institute of Mathematical Statistics, 2009.
- [14] G. Glaeser. Étude de quelques algèbres tayloriennes. *J. Analyse Math.*, 6:1–124; erratum, insert to 6 (1958), no. 2, 1958.
- [15] P. Shvartsman. The Whitney extension problem and Lipschitz selections of set-valued mappings in jet-spaces. *Transactions of the American Mathematical Society*, 360(10):5529–5550, 2008.
- [16] P. Chew. There is a planar graph almost as good as the complete graph. In *Proceedings of the Second Annual Symposium on Computational Geometry*, SCG ’86, pages 169–177, New York, NY, 1986. ACM.
- [17] D. Peleg and A. A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
- [18] S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.
- [19] Q. W. Bouts, A. P. ten Brink, and K. Buchin. A framework for computing the greedy spanner. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, SOCG’14, pages 11:11–11:19, New York, NY, 2014. ACM.
- [20] M. Lubin and I. Dunning. Computing in operations research using Julia. *INFORMS Journal on Computing*, 27(2):238–248, 2015.
- [21] Gurobi Optimization. Gurobi optimizer reference manual, 2015. URL <http://www.gurobi.com>.
- [22] S. S. Vallender. Calculation of the Wasserstein distance between probability distributions on the line. *Theory of Probability & Its Applications*, 18(4):784–786, 1974.
- [23] C. Döbler. Stein’s method of exchangeable pairs for the Beta distribution and generalizations. *arXiv:1411.4477*, 2014.
- [24] A. Canty and B. D. Ripley. *boot: Bootstrap R (S-Plus) Functions*, 2015. R package version 1.3-15.
- [25] G. O. Roberts and R. L. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, pages 341–363, 1996.
- [26] R. E. Caflisch. Monte carlo and quasi-monte carlo methods. *Acta numerica*, 7:1–49, 1998.
- [27] Y. Chen, M. Welling, and A. Smola. Super-samples from kernel herding. In *Proceeding of 26th Uncertainty in Artificial Intelligence (UAI’10)*, 2010.

- [28] E. del Barrio, E. Gin, and C. Matrn. Central limit theorems for the Wasserstein distance between the empirical and the true distributions. *Ann. Probab.*, 27(2):1009–1071, 04 1999.
- [29] F. Bach, S. Lacoste-Julien, and G. Obozinski. On the equivalence between herding and conditional gradient algorithms. In *Proceeding of 29th International Conference on Machine Learning (ICML'12)*, 2012.
- [30] A. Zellner and C.-K. Min. Gibbs sampler convergence criteria. *Journal of the American Statistical Association*, 90(431):921–927, 1995.
- [31] Y. Fan, S. P. Brooks, and A. Gelman. Output assessment for monte carlo simulations via the score statistic. *Journal of Computational and Graphical Statistics*, 15(1), 2006.
- [32] A. Gretton, K. M Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola. A kernel method for the two-sample-problem. In *Advances in Neural Information Processing Systems*, pages 513–520, 2006.
- [33] R. Khasminskii. *Stochastic stability of differential equations*, volume 66. Springer Science & Business Media, 2011.
- [34] F. Bolley, I. Gentil, and A. Guillin. Convergence to equilibrium in Wasserstein distance for Fokker–Planck equations. *Journal of Functional Analysis*, 263(8):2430–2457, 2012.
- [35] M. Cule and R. Samworth. Theoretical properties of the log-concave maximum likelihood estimator of a multidimensional density. *Electron. J. Stat.*, 4:254–270, 2010.
- [36] S. N. Ethier and T. G. Kurtz. *Markov processes: characterization and convergence*, volume 282. John Wiley & Sons, 2009.