Rao–Blackwellized Score Matching on Manifolds

Divit Rawal [1]

Abstract

We study the tangent channel of denoising score matching (DSM) when the latent law is supported on a smooth embedded submanifold $M \subset \mathbb{R}^D$. For ambient Gaussian corruption, this channel is the projected denoising residual $T_\sigma \doteq P_{T_{\pi(X)}M}(Z - X)/\sigma^2$, whose conditional variance diverges at rate d/σ^2 as $\sigma \to 0^+$. Within the class of fiber-collapsing summaries $S(X)$, we identify the nearest-point projection $\pi(X)$ as the canonical finest such summary, so that $r_\sigma(z) \doteq \mathbb{E}[T_\sigma \mid \pi(X) = z]$ is the unique L^2-optimal Rao-Blackwellized predictor of T_σ in this class; the singular d/σ^2 term is an irreducible Bayes-risk floor for any coarser summary. Expanding this canonical target in σ, we show

$$r_\sigma(z) = \nabla_M \log q(z) + \sigma^2 \left[b_q(z) + g_M^{\text{ext}}(z) \right] + \mathrm{o}(\sigma^2),$$

where $g_M^{\text{ext}}(z) = \left(\frac{1}{2} W_{H(z)} - \mathrm{Ric}_z^\sharp\right) \nabla_M \log q(z)$, uniformly on M. Here b_q is the intrinsic Tweedie term and the extrinsic term combines the Weingarten operator in the mean-curvature direction with the Ricci endomorphism. This extrinsic correction is absent from intrinsic-noising analyses and is the σ^2 obstruction to recovering the Riemannian score by ambient DSM. On S^d it collapses to the scalar $(1 - d/2)\,\mathrm{Id}$, which vanishes at $d = 2$, a common test case for manifold DSM. These results separate the removable singular variance of the raw ambient target from the intrinsic and extrinsic second-order biases of the canonical target.

1. Introduction

The manifold hypothesis — that high-dimensional data concentrates on or near a lower-dimensional submanifold $M \subset \mathbb{R}^D$ — underpins much of modern generative modeling. For score-based generative models in particular, it creates a tension: the latent law $Z \sim q d\,\mathrm{Vol}_M$ is singular with respect to the ambient Lebesgue measure, so the ambient score $\nabla \log q$ is not defined, and denoising score matching (DSM) is only well-posed at strictly positive noise level $\sigma > 0$. As a result, the DSM regression learns only a σ-dependent surrogate of the intrinsic Riemannian score instead of the true score.

Two primary lines of work have emerged in response. Intrinsic methods (Bortoli et al., 2022; Huang et al., 2022) replace the ambient noising process by Riemannian Brownian motion on M and parameterize scores as tangent vector fields, gaining manifold tools (exponential maps, logarithm maps, heat kernels) at the cost of sampler complexity. Ambient methods (Levy-Jurgenson et al., 2026) keep the standard Euclidean DSM process and work under the assumption that as $\sigma \to 0^+$, the learned ambient field recovers the intrinsic score after projection. Both lines have yielded impressive empirical results; both also recognize that the manifold hypothesis causes ambient DSM targets to carry singular normal-fiber nuisance and existing generalization bounds degrade as $\sigma \to 0^+$ (Yakovlev & Puchkin, 2025).

However, a basic question remains open about the ambient approach: ambient DSM on manifold-supported data does not directly regress the intrinsic Riemannian score. Instead, at finite noise level it regresses a σ-dependent tangent target containing both intrinsic score signal and singular normal-fiber noise. This raises three related questions: *what is the exact finite-σ target, how can the singular variance be removed in a statistically canonical way, and what curvature-dependent bias remains relative to the intrinsic score?*

In this work, we provide such an answer. We study the Rao-Blackwellized denoising target

$$r_\sigma(z) \doteq \mathbb{E}\left[\frac{P_{T_z M}(Z - X)}{\sigma^2} \mid \pi(X) = z \right],$$

where $\pi : \mathrm{Tub}(M) \to M$ is the nearest-point projection and $P_{T_z M}$ is the orthogonal projection onto the tangent space at z. We show that this object arises naturally by conditioning the tangent component of the raw DSM denoising residual on the canonical manifold-valued summary $\pi(X)$. We show that r_σ is the unique L^2-optimal Rao-Blackwellized predictor of the tangent DSM target T_σ

[1] Department of Statistics, University of California, Berkeley, Berkeley CA, U.S.A.. Correspondence to: Divit Rawal <divit.rawal@berkeley.edu>.

Preprint. May 5, 2026.

among fiber-collapsing summaries of X, and that it agrees with the intrinsic Riemannian score up to order $O(\sigma^2)$. In particular, conditioning on $\pi(X)$ removes the singular fiber-noise channel, yielding a bounded-variance target.

We then compute $r_\sigma(z)$ to second order on an arbitrary smooth embedded submanifold. The leading term is the intrinsic Riemannian score; the σ^2 correction decomposes cleanly into a Tweedie-style intrinsic piece and an extrinsic piece. Concretely,

$$\begin{aligned} r_\sigma(z) = \nabla_M \log q(z) &\\ + \sigma^2 \Big[\tfrac{1}{2} \nabla_M \big(\Delta_M \log q + \|\nabla_M \log q\|^2 \big)(z) &\\ + \big(\tfrac{1}{2} W_{H(z)} - \mathrm{Ric}^\sharp_z \big) \nabla_M \log q(z) \Big] &\\ + \mathrm{o}(\sigma^2), & \end{aligned} \tag{1}$$

uniformly on M, where $W_{H(z)}$ is the Weingarten operator in the mean-curvature direction and $\mathrm{Ric}^\sharp_z$ is the Ricci endomorphism on $T_z M$ (a refresher of relevant concepts from differential geometry is given in Appendix A). The extrinsic operator is a curvature tensor in the Riemannian sense: using the Gauss equation $\mathrm{Ric}^\sharp = W_H - \sum_\alpha W_{n_\alpha}^2$, one may equivalently write it as $\big(\sum_\alpha W_{n_\alpha}^2 - \frac{1}{2} W_H\big)$. On the unit sphere S^d this collapses to the scalar $(1 - d/2)\,\mathrm{Id}$. On S^2, then, the extrinsic σ^2 obstruction vanishes, so ambient DSM matches the intrinsic score up to the intrinsic flat-Tweedie bias $\sigma^2 b_q$. This cancellation helps explain why ambient approaches can perform well on spherical scientific data, such as Earth and climate data.

Equation (1) allows one, given any embedding of M and any σ, to compute the finite σ bias of ambient DSM to second order from two elementary curvature tensors. Practitioners are then free to (i) ignore it, (ii) subtract it off, or (iii) choose σ to dominate it.

Contributions. In this work, we first identify the canonical object to study, then compute it. Concretely, our contributions are as follows:

1. *Canonicality of $\pi(X)$ and a singular Bayes-risk floor.* Among all fiber-collapsing summaries $S(X)$ of the noisy observation (those satisfying $\sigma(S) \subseteq \sigma(\pi(X))$), the nearest-point projection $\pi(X)$ is the canonical finest such summary; consequently $r_\sigma(\pi(X)) = \mathbb{E}[T_\sigma \mid \pi(X)]$ is the unique (up to null sets) L^2-optimal Rao-Blackwellized predictor of the tangent DSM target T_σ in this class. We also show that $\mathrm{Var}(T_\sigma) = d/\sigma^2 + O(1)$ while $\mathrm{Var}(r_\sigma(\pi(X))) = O(1)$ and the singular d/σ^2 is an irreducible Bayes-risk floor for regression against T_σ through any coarser summary.

2. *Exact reduction in the flat case.* When M is an affine subspace, the Rao-Blackwellized target reduces exactly to ordinary lower-dimensional Gaussian DSM on M. In this case there is no curvature correction: the tangent component of ambient DSM, the intrinsic Gaussian-smoothed score, and the Rao-Blackwellized target all coincide exactly.

3. *Extrinsic σ^2 correction to the canonical target.* Expanding the canonical target r_σ in σ on an arbitrary smooth embedded submanifold, we obtain the expansion in (1) with explicit intrinsic and extrinsic pieces. We isolate two geometrically distinct sources of this correction: a curvature deformation of the volume measure ($-\mathrm{Ric}^\sharp$) and a Laplacian mismatch between ambient and intrinsic differentials ($+\frac{1}{2} W_H$). On S^d, the operator collapses to the scalar $(1 - d/2)\,\mathrm{Id}$.

2. Related Work

Score matching and denoising score matching. Score matching was introduced by Hyvärinen (2005) as a way of fitting non-normalized statistical models without evaluating normalizing constants. Vincent (2011) established the equivalence between denoising score matching (DSM) at a fixed noise level σ and regression of the denoising residual $(Z - X)/\sigma^2$, via Tweedie's formula (Efron, 2011; Robbins & Neyman, 1956). Both derivations assume that the latent law q admits a density with respect to ambient Lebesgue measure. When q is supported on a lower-dimensional submanifold, the ambient score $\nabla \log q$ is not defined and DSM is only meaningful at positive σ. Our results clarify what DSM is actually estimating in this regime: the raw tangent denoising target T_σ contains a nuisance normal-fiber component whose conditional variance diverges at rate d/σ^2, and its projection onto $\pi(X)$ isolates the signal-bearing part.

Score-based generative models and the manifold hypothesis. The diffusion-model literature (Song & Ermon, 2020; Ho et al., 2020; Song et al., 2021) formulates generation as reversing a Gaussian noising process, which implicitly regularizes any singular latent law by Gaussian convolution. Theoretical work on this regime has repeatedly observed that the ambient score blows up near the support of a low-dimensional latent. Pidstrigach (2022) showed that score-based generative models trained on data concentrated near a manifold recover drift fields that align with the normal direction, effectively detecting the manifold. Bortoli (2023) gave quantitative convergence guarantees for denoising diffusion models under a manifold hypothesis, and Bortoli et al. (2022) introduced Riemannian score-based generative modelling (RSGM), which replaces the ambient Gaussian noising process by an intrinsic heat-kernel diffusion on a known manifold. RSGM estimates an intrinsic Riemannian score directly, but requires the ability to simulate the heat kernel on M. Our work is complementary: we ask what the ambient DSM target actually identifies when the latent law is singular with respect to Lebesgue measure, and show that a single Rao-Blackwellization step against the nearest-point projection $\pi(X)$ converts the ambient target into an $O(\sigma^2)$-accurate estimator of the intrinsic Riemannian score, without simulating any heat kernel.

Sample complexity and rates. Chen et al. (2023) and Oko et al. (2023) derive sample-complexity and minimax-rate guarantees for diffusion models, working in the ambient formulation. Our results are population-level identification theorems rather than rate theorems, but the variance collapse in Theorem 4.2 and the finite-sample bound in Proposition G.1 quantify the price of ignoring the manifold structure. They suggest that Rao-Blackwellization can yield an asymptotically stronger signal-to-noise improvement as $\sigma \to 0^+$, rather than merely a constant-factor gain.

Manifolds and Geometry. The tubular-neighborhood and positive-reach calculus we use was developed by Federer (1959); we use it in the form developed by Niyogi et al. (2008) for manifold learning. The local-linear regression proposition in Appendix G is based on the classical rate theory of Fan (1992).

Existing DSM theory assumes an absolutely continuous latent law (Hyvärinen, 2005; Vincent, 2011); the manifold-hypothesis literature has observed blowup of the ambient score (Pidstrigach, 2022; Bortoli, 2023); and the intrinsic heat-kernel route (Bortoli et al., 2022) bypasses the ambient score entirely. In this work, we provide three pieces absent from prior analyses: (i) a canonicality statement identifying $\pi(X)$ as the finest fiber-collapsing summary; (ii) the exact constant d/σ^2 in the raw-target variance, together with a matching Bayes-risk floor; and (iii) an exact equality reduction of ambient DSM to lower-dimensional DSM in the flat case, which pins down the baseline against which curvature effects must be measured.

3. Setup

Let $M \subset \mathbb{R}^D$ be a compact embedded C^5 submanifold of dimension d and positive reach. Let $q \in C^5(M)$ be strictly positive with respect to the Riemannian volume measure $d\,\mathrm{Vol}_M$.

We consider the Gaussian corruption model

$$Z \sim q d\,\mathrm{Vol}_M, \quad X = Z + \sigma\xi, \quad \xi \sim \mathcal{N}(0, I_D). \tag{2}$$

Let $r_0 < \mathrm{reach}(M)$ and write $\mathrm{Tub}_{r_0}(M)$ for the corresponding tubular neighborhood. For sufficiently small σ, the event

$$\mathcal{E}_\sigma \doteq \{X \in \mathrm{Tub}_{r_0}(M)\}$$

has probability $1 - \exp(-c/\sigma^2)$ for some $c > 0$. On $\mathcal{E}_\sigma$, the nearest-point projection

$$\pi : \mathrm{Tub}_{r_0}(M) \to M$$

is well-defined and smooth. Throughout we work on $\mathcal{E}_\sigma$; all omitted tails are exponentially small in σ^{-2} and do not affect any polynomial-order expansions below.

For $z \in M$, let $P_T(z)$ and $P_\perp(z)$ denote the orthogonal projections onto T_zM and N_zM.

The standard denoising target is

$$Y_\sigma \doteq \frac{Z - X}{\sigma^2}. \tag{3}$$

We define the raw tangent denoising target

$$T_\sigma \doteq P_T(\pi(X))Y_\sigma. \tag{4}$$

Since $X - \pi(X) \in N_{\pi(X)}M$, this simplifies to

$$T_\sigma = \frac{1}{\sigma^2} P_T(\pi(X))(Z - \pi(X)). \tag{5}$$

Definition 3.1 (Rao-Blackwellized tangent target)**.** For $z \in M$, define

$$r_\sigma(z) \doteq \mathbb{E}[T_\sigma \mid \pi(X) = z] \in T_zM. \tag{6}$$

For any measurable tangent field $h : M \to TM$, define the projected denoising risk

$$\mathcal{R}_\sigma(h) \doteq \mathbb{E}\,\|T_\sigma - h(\pi(X))\|^2. \tag{7}$$

4. Canonicality of $\pi(X)$

In this section, we prove that within the family of fiber-measurable tangent fields, the Rao-Blackwellized target r_σ is the unique L^2-risk minimizer — a straightforward Pythagorean identity that also gives canonicality of $\pi(X)$ among fiber-collapsing summaries. We also show that r_σ agrees with the intrinsic Riemannian score $\nabla_M \log q$ to leading order in σ and that the residual risk of the raw target T_σ against any fiber-collapsing predictor diverges at the exact rate d/σ^2; an irreducible Bayes-risk floor. Full proofs of the claims are in Appendix C.

4.1. L^2 Projection Identity

The risk $\mathcal{R}_\sigma(h)$ in (7) penalizes a tangent field h for its L^2 distance to T_σ. Because $r_\sigma(\pi(X))$ is a conditional expectation, an L^2-projection decomposition applies.

Theorem 4.1 (Projected denoising risk)**.** *For every measurable tangent field $h : M \to TM$,*

$$\mathcal{R}_\sigma(h) = \mathcal{R}_\sigma(r_\sigma) + \mathbb{E}\,\|r_\sigma(\pi(X)) - h(\pi(X))\|^2. \tag{8}$$

In particular, r_σ is the unique (up to null sets) minimizer of $\mathcal{R}_\sigma$.

Proof. Since $r_\sigma(\pi(X)) = \mathbb{E}[T_\sigma \mid \pi(X)]$, we may write

$$T_\sigma - h(\pi(X)) = \big(T_\sigma - r_\sigma(\pi(X))\big) + \big(r_\sigma(\pi(X)) - h(\pi(X))\big).$$

The second term is $\sigma(\pi(X))$-measurable, whereas the first is orthogonal in L^2 to every $\sigma(\pi(X))$-measurable square-integrable random variable. Hence the cross term in $\mathbb{E}\,\|T_\sigma - h(\pi(X))\|^2$ vanishes, which is (8). Uniqueness is immediate: the second summand is nonnegative and vanishes iff $h(\pi(X)) = r_\sigma(\pi(X))$ a.s. ☐

The same identity, applied with h replaced by an estimator measurable with respect to a coarser statistic, extends the Pythagorean decomposition to the full family of fiber-collapsing summaries. These are statistics $S = S(X)$ satisfying $\sigma(S) \subseteq \sigma(\pi(X))$; equivalently, there exists a measurable function $\bar{S}$ such that $S = \bar{S}(\pi(X))$ a.s. For any such S, define $\eta_S \doteq \mathbb{E}[T_\sigma \mid S]$. The tower property gives

$$\eta_S = \mathbb{E}[T_\sigma \mid S] = \mathbb{E}\big[r_\sigma(\pi(X)) \mid S\big], \tag{9}$$

so the optimal S-measurable predictor of T_σ equals the optimal S-measurable predictor of $r_\sigma(\pi(X))$, and the excess risk of any S-measurable estimator η over $r_\sigma(\pi(X))$ decomposes exactly as

$$\begin{aligned}\mathbb{E}\,\|T_\sigma - \eta\|^2 = \mathbb{E}\,\|T_\sigma - r_\sigma(\pi(X))\|^2 \\ + \mathbb{E}\,\|r_\sigma(\pi(X)) - \eta_S\|^2 \\ + \mathbb{E}\,\|\eta_S - \eta\|^2 .\end{aligned} \tag{10}$$

Among estimators measurable with respect to fiber-collapsing summaries, $r_\sigma(\pi(X))$ is therefore the unique minimum-risk choice, and (10) identifies the exact cost $\mathbb{E}\,\|r_\sigma(\pi(X)) - \eta_S\|^2$ of coarsening $\pi(X)$ to S. This is the Rao-Blackwell statement for our setup: among fiber-collapsing statistics, $\pi(X)$ is the canonical finest summary and $r_\sigma(\pi(X))$ is the corresponding L^2-optimal predictor of T_σ.

4.2. Leading-Order Behavior

Having identified the correct object to target, we now ask how close r_σ is to the intrinsic Riemannian score. A tubular-coordinate Bayes calculation combined with a manifold Stein identity (Appendix C) yields, uniformly in $z \in M$,

$$r_\sigma(z) = \nabla_M \log q(z) + O(\sigma^2), \qquad \sigma \to 0^+. \tag{11}$$

This is a small-noise *population* statement; a finite-sample consistency rate in the number of observations used to form the conditional expectation is given in Appendix G. The σ^2 residual is nonzero on a curved support and is precisely the extrinsic correction studied in Section 5.

4.3. Variance Collapse and Singular Floor

We now quantify how much of the raw target's risk is supervision noise rather than signal. Throughout, for a vector-valued random variable V, the symbol $\mathrm{Var}(V)$ denotes the trace covariance, equivalently $\mathbb{E}\,\|V - \mathbb{E}V\|^2$.

Theorem 4.2 (Variance collapse under Rao-Blackwellization)**.** *Uniformly in $z \in M$,*

$$\mathrm{Var}(T_\sigma \mid \pi(X) = z) = \frac{d}{\sigma^2} + O(1). \tag{12}$$

Consequently,

$$\mathrm{Var}(T_\sigma) = \mathrm{Var}(r_\sigma(\pi(X))) + \frac{d}{\sigma^2} + O(1), \tag{13}$$

so the raw tangent denoising target has variance diverging like d/σ^2 while the Rao-Blackwellized target has $O(1)$ variance.

The constant d/σ^2 is not an artifact of the specific estimator r_σ; it is the Bayes-risk floor for regression against T_σ within every fiber-collapsing summary. Indeed, for any S with $\sigma(S) \subseteq \sigma(\pi(X))$, (10) bounds

$$\begin{aligned}\inf_{\eta \in L^2(\sigma(S))} \mathbb{E}\,\|T_\sigma - \eta\|^2 &\geq \mathbb{E}\,\|T_\sigma - r_\sigma(\pi(X))\|^2 \\ &= \mathbb{E}[\mathrm{Var}(T_\sigma \mid \pi(X))] \\ &= \frac{d}{\sigma^2} + O(1),\end{aligned}$$

where the last equality is the tower-rule average of (12). Equality at leading order holds only when $\sigma(S) = \sigma(\pi(X))$ modulo null sets. The roles of π and r_σ are distinct: $\pi(X)$ is the canonical finest fiber-collapsing summary, and r_σ is the L^2-optimal predictor of T_σ given that summary. Every fiber-collapsing S inherits the d/σ^2 floor; equality at leading order holds only when $\sigma(S) = \sigma(\pi(X))$. Figure 1 demonstrates the variance-collapse rate of Theorem 4.2 on S^2 with the von-Mises Fisher density.

5. Extrinsic σ^2 Correction

We now provide a closed-form expansion of the canonical target r_σ in σ, with an extrinsic curvature correction that is absent from intrinsic-noising analyses. First, we exactly solve the flat case, then state the curved case expansion, and finally specialize to S^d.

5.1. Affine Case

We first analyze the affine case, where all identities are exact.

Let $V \subset \mathbb{R}^D$ be an affine d-plane, P the orthogonal projection onto V, and $Q \doteq I - P$. Assume $Z \in V$ with

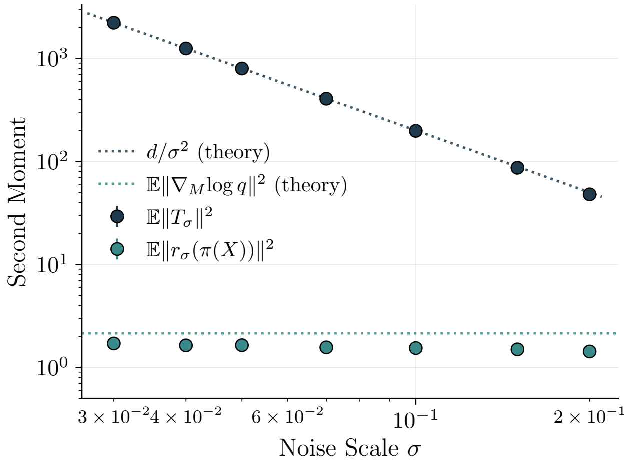

Figure 1. Variance collapse on S^2 under vMF(μ, κ=2). Second moment of the raw target T_σ (black, slope -2 in $\log \sigma$, matching d/σ^2 with $d = 2$) versus the Rao-Blackwellized target $r_\sigma(\pi(X))$ (blue, flat at the theoretical $\mathbb{E}\,\|\nabla_M \log q\|^2$). The gap is the irreducible d/σ^2 Bayes-risk floor of Section 4.3.

$Z \sim q$ on V, $X = Z + \sigma\xi$ with $\xi \sim \mathcal{N}(0, I_D)$, and write $T \doteq PX$, $N \doteq QX$. Then $T = Z + \sigma P\xi$ and $N = \sigma Q\xi$ are independent. Let

$$p_T \doteq q * \phi_\sigma^{(d)}$$

be the d-dimensional Gaussian convolution of q on V.

Proposition 5.1 (Flat Case: Tweedie Identity and Exact DSM Reduction)**.** *In the setting above, the ambient score factorizes as*

$$\nabla \log p_\sigma(x) = P\nabla_V \log p_T(Px) - \frac{Qx}{\sigma^2}, \tag{14}$$

the Rao-Blackwellized target coincides with the flat Tweedie score along V,

$$r_\sigma(t) = \nabla_V \log p_T(t), \qquad p_T = q * \phi_\sigma^{(d)}, \tag{15}$$

and ambient DSM restricted to the tangent channel reduces exactly to d-dimensional Gaussian DSM on V: for every measurable $h : V \to V$, with $g_h(x) \doteq Ph(Px) - Qx/\sigma^2$,

$$\mathbb{E}\,\|Y_\sigma - g_h(X)\|^2 = \mathbb{E}\left\|\tfrac{Z-T}{\sigma^2} - h(T)\right\|^2. \tag{16}$$

Proof. For (16), write

$$Y_\sigma = \frac{Z - X}{\sigma^2} = \frac{Z - T}{\sigma^2} - \frac{N}{\sigma^2}, \qquad g_h(X) = h(T) - \frac{N}{\sigma^2}.$$

Therefore

$$Y_\sigma - g_h(X) = \frac{Z - T}{\sigma^2} - h(T),$$

which lies in V. Taking squared norms and expectations gives (16). □

Thus, on a flat support, the intrinsic score, the V-component of the ambient score, and the Rao-Blackwellized tangent target all coincide; no σ^2 correction is needed. The curved case, which we examine next, breaks all three coincidences by an amount that is second order in σ and carried entirely by the second fundamental form of $M \hookrightarrow \mathbb{R}^D$.

5.2. Curved Case

We now state the main expansion of r_σ on a curved support. The coefficient decomposes into an intrinsic flat-Tweedie term and an extrinsic curvature term.

Theorem 5.2 (Extrinsic σ^2 correction)**.** *Under the assumptions of Section 3, uniformly in $z \in M$,*

$$r_\sigma(z) = \nabla_M \log q(z) + \sigma^2\big[b_q(z) + g_M^{\text{ext}}(z)\big] + \mathrm{o}(\sigma^2), \tag{17}$$

as $\sigma \to 0^+$, where the intrinsic flat-Tweedie term is

$$b_q(z) = \tfrac{1}{2}\nabla_M\left[\Delta_M \log q + \|\nabla_M \log q\|^2\right](z), \tag{18}$$

and the extrinsic term is

$$g_M^{\text{ext}}(z) = \left(\tfrac{1}{2}W_{H(z)} - \mathrm{Ric}_z^\sharp\right)\big(\nabla_M \log q(z)\big). \tag{19}$$

Here $W_u : T_zM \to T_zM$ is the Weingarten operator in normal direction u, $H(z) = \sum_{i=1}^d \mathrm{II}_z(e_i, e_i) \in N_zM$ is the mean curvature vector of $M \hookrightarrow \mathbb{R}^D$ at z, and $\mathrm{Ric}_z^\sharp$ is the Ricci endomorphism of (M, g_M) at z. In the flat case $M = V$ both W_H and $\mathrm{Ric}^\sharp$ vanish, so $g_M^{\text{ext}} \equiv 0$ and (17) reduces to the flat Tweedie identity (15).

The takeaway is that ambient Gaussian corruption does not merely smooth the intrinsic density on M; it also changes the target score through a curvature-dependent embedding term. Thus, even after projecting onto the tangent bundle and Rao–Blackwellizing away the singular fiber noise, ambient DSM is biased relative to intrinsic noising by the explicit $\sigma^2 g_M^{\text{ext}}$ term. The proof is a graph-coordinate Bayes calculation that isolates two independent curvature contributions: the induced volume-form correction and the tube-Jacobian mean-curvature correction, which combine without cancellation into the operator $\frac{1}{2}W_{H(z)} - \mathrm{Ric}_z^\sharp$. Full details are provided in Appendix F.

Remark 5.3 (Comments on (19))**.** We note that (i) the extrinsic term is invisible to any analysis that corrupts Z by intrinsic manifold noise (e.g. Brownian motion on M or geodesic noising), where only $b_q(z)$ appears. The term $g_M^{\text{ext}}(z)$ is a result of ambient Gaussian corruption and vanishes identically in intrinsic-noising analyses. (ii) Using the Gauss equation $\mathrm{Ric}_z^\sharp = W_{H(z)} - \mathcal{S}_z$ with $\mathcal{S}_z \doteq \sum_\alpha W_{n_\alpha}^2$ for any orthonormal normal frame, one can equivalently write $g_M^{\text{ext}}(z) = (\mathcal{S}_z - \frac{1}{2}W_{H(z)})\nabla_M \log q(z)$.

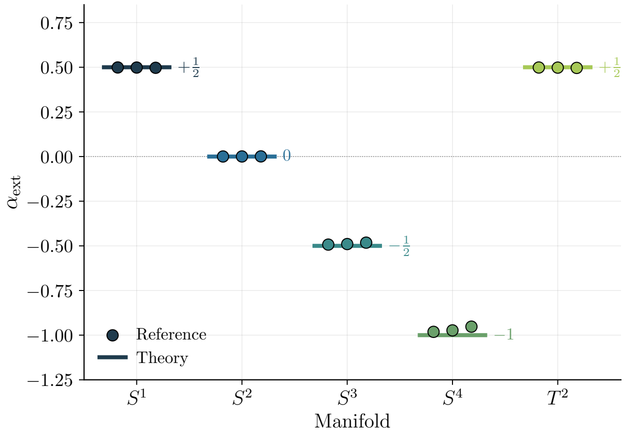

Figure 2. Extrinsic coefficient α_{ext} across manifolds. Quadrature estimates (one dot per $\sigma \in \{0.05, 0.06, 0.08\}$) are computed by Gauss–Hermite quadrature of $r_\sigma(z)$ and compared to the predicted coefficients $\alpha_d = 1 - d/2$ on S^d and $\alpha = +\frac{1}{2}$ on T^2. Latent density q is von Mises–Fisher with $\kappa = 2$ on each S^d and a wrapped Gaussian with $\kappa = 1.5$ on T^2.

Figure 2 computes $r_\sigma(z)$ numerically, subtracts the intrinsic score and the Tweedie term, and divides by $\sigma^2 \nabla_M \log q(z)$ to isolate the dimensionless extrinsic coefficient α_{ext} we predict. The T^2 entry is a non-sphere control: T^2 is intrinsically flat and extrinsically curved; its nonzero $\alpha = +\frac{1}{2}$ confirms the correction is an embedding effect, not an intrinsic-curvature artifact.

5.3. Specialization to S^d

We record the specialization of (17) to the unit sphere $S^d \subset \mathbb{R}^{d+1}$, a common testbed for manifold DSM.

Corollary 5.4 (Extrinsic coefficient on S^d). *For $M = S^d \subset \mathbb{R}^{d+1}$ with the round metric and outward normal $\nu = z$, $W_\nu = -\operatorname{Id}_{T_zM}$, $H(z) = -dz$, $W_{H(z)} = d \operatorname{Id}_{T_zM}$, and $\operatorname{Ric}^\sharp_z = (d-1)\operatorname{Id}_{T_zM}$, so*

$$g^{\text{ext}}_{S^d}(z) = \big(1 - \tfrac{d}{2}\big)\nabla_M \log q(z). \tag{20}$$

In particular, the scalar multiplier $\alpha_d \doteq 1 - d/2$ attached to $\nabla_M \log q(z)$ in the σ^2 correction is $\alpha_1 = +\frac{1}{2}$, $\alpha_2 = 0$, $\alpha_3 = -\frac{1}{2}$, and $\alpha_4 = -1$.

The S^2 zero is structural — on S^2 the Einstein identity $\frac{1}{2}W_H = \operatorname{Ric}^\sharp = \operatorname{Id}$ forces the two curvature contributions to cancel. Proof is a direct substitution; see Appendix F.4.

Figure 3 reads Corollary 5.4 at the level of generated samples. Running Langevin with the population-level ambient-DSM drift converges to $\operatorname{vMF}\big(\kappa(1 + \sigma^2\alpha_d)\big)$, an under-concentration that is 0% on S^2 (the structural α_2=0), $\approx$ 17% on S^6, and $\approx$ 33% on S^{10} at σ=0.3, κ=5. The post-hoc multiplication $(1 - \sigma^2\alpha_d)$ recovers the truth in every panel, with a residual $O(\sigma^4)$ floor from the intrinsic Tweedie bias b_q that is shared between ambient and intrinsic DSM and is therefore separately removable.

6. Discussion

Intrinsic generation and inference. (11) identifies the canonical target as an intrinsic score surrogate, $r_\sigma(z) = \nabla_M \log q(z) + O(\sigma^2)$, so the intrinsic Langevin dynamics

$$dZ_t = r_\sigma(Z_t)dt + \sqrt{2}dB_t^M \tag{21}$$

are an $O(\sigma^2)$-accurate drift surrogate for Langevin sampling targeting q on M. Theorem 4.2 makes this surrogate statistically usable: the raw tangent denoising target has variance diverging at the d/σ^2 rate, whereas $r_\sigma(\pi(X))$ has bounded variance, so Rao-Blackwellization is not a constant-factor improvement; it removes the singular fiber-noise channel and isolates the intrinsic score signal.

Ambient vs. intrinsic DSM. The comparison between ambient and intrinsic methods is often framed qualitatively: intrinsic methods require manifold infrastructure while ambient methods do not. Equation (17) quantifies the exact bias resulting from using ambient instead of intrinsic methods. Given any embedded M and any σ, the finite-σ extrinsic bias of ambient DSM is $\sigma^2 g^{\text{ext}}_M(z) = \sigma^2\left(\frac{1}{2}W_H(z) - \operatorname{Ric}^\sharp_z\right)\nabla_M \log q(z)$, computable from two elementary curvature tensors and the intrinsic score. Practitioners may (i) accept it (if σ is small enough that $\sigma^2\,\|g^{\text{ext}}\|$ is below the target accuracy), (ii) subtract it off, leaving only the intrinsic flat-Tweedie bias $\sigma^2 b_q$, or (iii) avoid regimes where it dominates.

S^2 bias. Corollary 5.4 shows that S^2 — the most commonly used test for manifold-aware score matching — is a coincidentally benign case: the extrinsic coefficient $1 - d/2$ vanishes exactly at $d = 2$, so on S^2 ambient DSM matches the intrinsic score up to the intrinsic flat-Tweedie bias $\sigma^2 b_q$, without any extrinsic debiasing. This structural coincidence arises from the Einstein identity $\frac{1}{2}W_H = \operatorname{Ric}^\sharp$ on unit S^2, and the formula predicts measurable second-order bias on S^1, S^3, or S^d for $d \geq 3$ (Figure 3). The torus control (T^2 with $\alpha = +\frac{1}{2}$ despite being intrinsically flat) rules out the interpretation that the effect is a pure intrinsic-curvature artifact; it is an embedding effect.

Limitations and open directions. Our results are small-noise population statements: we characterize r_σ in terms of the true joint law of (Z, X) and its expansion as $\sigma \to 0^+$. A finite-sample local-averaging rate is given in Appendix G, but the interplay between σ and sample size in the regime where the extrinsic bias and the statistical error are of comparable magnitude (i.e. how to choose σ given N) is a natural next question. Extending the expansion to higher order, to

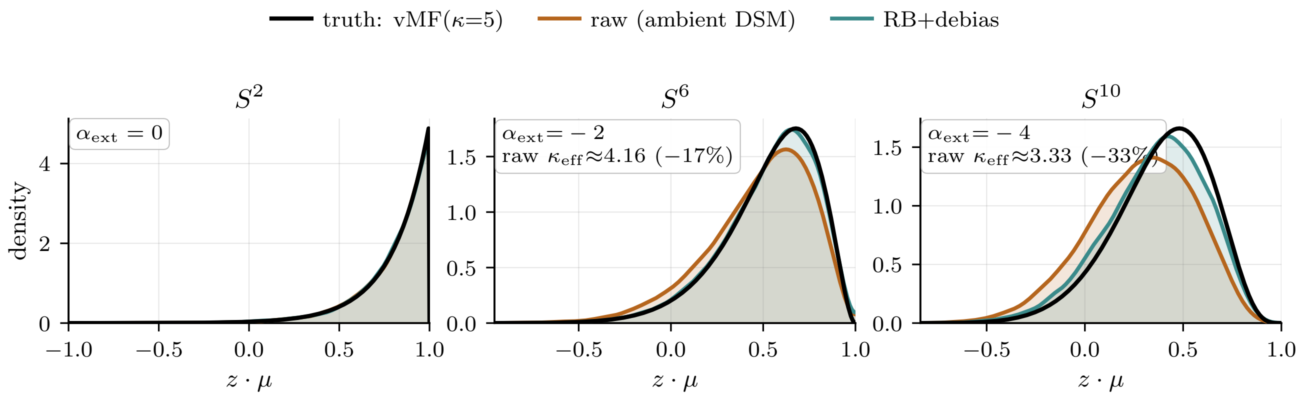

Figure 3. Effective concentration of ambient DSM samples on S^d under vMF(μ, κ=5), σ=0.3. Densities of $z \cdot \mu$ at the equilibrium of three closed-form Langevin drifts (no network trained): the intrinsic score $\nabla_M \log q$ (black, analytic vMF marginal), the ambient-DSM target $(1 + \sigma^2\alpha_d)\nabla_M \log q$ with α_d=$1 - d/2$ (orange), and its Corollary 5.4 debias $(1 - \sigma^2\alpha_d)(1 + \sigma^2\alpha_d)\nabla_M \log q$ (blue). The Einstein cancellation α_2=0 makes ambient DSM exact on S^2; on S^6 and S^{10} it under-concentrates the equilibrium by $\approx$ 17% and $\approx$ 33% in effective κ, both removed by the scalar correction up to $O(\sigma^4)$.

manifolds with boundary, or to anisotropic noise are also open.

Scientific data are often constrained to embedded submanifolds — e.g. orientation and rotation groups in robotics and protein structure, directional and spherical data in geosciences and astronomy, and products of spaces. The question of whether, and to what order, ambient score models learn the correct intrinsic object is both practically and conceptually central: it informs whether a generic score-matching pipeline may be deployed on manifold-valued data without heavy geometric tools. Our results answer that question quantitatively at small σ and identify a canonical Rao-Blackwellized target that is implementation-agnostic, informing the development of score-matching-based pipelines for manifold-constrained data.

References

Bortoli, V. D. Convergence of denoising diffusion models under the manifold hypothesis, 2023. URL https://arxiv.org/abs/2208.05314.

Bortoli, V. D., Mathieu, E., Hutchinson, M., Thornton, J., Teh, Y. W., and Doucet, A. Riemannian score-based generative modelling, 2022. URL https://arxiv.org/abs/2202.02763.

Chen, S., Chewi, S., Li, J., Li, Y., Salim, A., and Zhang, A. R. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions, 2023. URL https://arxiv.org/abs/2209.11215.

do Carmo, M. P. *Riemannian Geometry*. Birkhäuser, Boston, 1992.

Efron, B. Tweedie's formula and selection bias. *Journal of the American Statistical Association*, 106:1602 – 1614, 2011. URL https://api.semanticscholar.org/CorpusID:23284154.

Fan, J. Design-adaptive nonparametric regression. *Journal of the American Statistical Association*, 87:998–1004, 1992. URL https://api.semanticscholar.org/CorpusID:53587425.

Federer, H. Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491, 1959. doi: 10.1090/S0002-9947-1959-0110078-1.

Ho, J., Jain, A., and Abbeel, P. Denoising diffusion probabilistic models, 2020. URL https://arxiv.org/abs/2006.11239.

Huang, C.-W., Aghajohari, M., Bose, A. J., Panangaden, P., and Courville, A. Riemannian diffusion models, 2022. URL https://arxiv.org/abs/2208.07949.

Hyvärinen, A. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(24):695–709, 2005. URL http://jmlr.org/papers/v6/hyvarinen05a.html.

Levy-Jurgenson, A., Prat, A., Cuin, J., and Teh, Y. W. Manifold aware denoising score matching (mad), 2026. URL https://arxiv.org/abs/2603.02452.

Niyogi, P., Smale, S., and Weinberger, S. Finding the homology of submanifolds with high confidence from random samples. *Discrete Comput. Geom.*, 39(1–3):419–441, March 2008. ISSN 0179-5376.

Oko, K., Akiyama, S., and Suzuki, T. Diffusion models are minimax optimal distribution estimators, 2023. URL https://arxiv.org/abs/2303.01861.

Pidstrigach, J. Score-based generative models detect manifolds, 2022. URL https://arxiv.org/abs/2206.01018.

Robbins, H. and Neyman, J. *An Empirical Bayes Approach to Statistics*. University of California Press, 1956. URL https://books.google.com/books?id=KANOwwEACAAJ.

Song, Y. and Ermon, S. Generative modeling by estimating gradients of the data distribution, 2020. URL https://arxiv.org/abs/1907.05600.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations, 2021. URL https://arxiv.org/abs/2011.13456.

Vincent, P. A connection between score matching and denoising autoencoders. *Neural Comput.*, 23(7):1661–1674, July 2011. ISSN 0899-7667. doi: 10.1162/NECO_a_00142. URL https://doi.org/10.1162/NECO_a_00142.

Yakovlev, K. and Puchkin, N. Generalization error bound for denoising score matching under relaxed manifold assumption, 2025. URL https://arxiv.org/abs/2502.13662.

A. Notation and Geometric Background

We define notation and record some differential-geometric identities used in the paper. We follow do Carmo (1992) and all statements are given without proof here.

Submanifold setup. Throughout the paper $M \subset \mathbb{R}^D$ is a C^4 embedded submanifold without boundary, of dimension $d < D$, with positive reach $\mathrm{reach}(M) > 0$. For $z \in M$ we write T_zM and N_zM for its tangent and normal spaces in $\mathbb{R}^D$, and $P_T(z) : \mathbb{R}^D \to T_zM$, $P_N(z) \doteq I - P_T(z) : \mathbb{R}^D \to N_zM$ for the orthogonal projections. The tangent bundle is $TM = \bigsqcup_{z\in M} T_zM$. The induced Riemannian metric g_M on M is the restriction of the Euclidean inner product on $\mathbb{R}^D$ to each T_zM. Gradients on M are denoted ∇_M and are defined by $\langle \nabla_M f(z), v\rangle = df(z)[v]$ for $v \in T_zM$; equivalently, $\nabla_M f(z) = P_T(z)\nabla_{\mathbb{R}^D} \tilde{f}(z)$ for any smooth extension $\tilde{f}$.

Second fundamental form. The second fundamental form of $M \subset \mathbb{R}^D$ at z is the symmetric bilinear map

$$\mathrm{II}_z : T_zM \times T_zM \to N_zM, \qquad \mathrm{II}_z(u, v) \doteq P_N(z)\overline{\nabla}_u \bar{v},$$

where $\overline{\nabla}$ is the ambient Euclidean connection and $\bar{v}$ is any tangent extension of v. Informally, $\mathrm{II}_z(u, v)$ measures how M bends away from its tangent plane in the (u, v) direction.

Weingarten operator. For a normal vector $\nu \in N_zM$, the Weingarten operator $W_\nu : T_zM \to T_zM$ is the symmetric endomorphism characterized by

$$\langle W_\nu u, v\rangle_{\mathbb{R}^D} = \langle \mathrm{II}_z(u, v), \nu\rangle_{\mathbb{R}^D}\,, \qquad u, v \in T_zM. \tag{22}$$

Equivalently, $W_\nu u = -P_T(z)\overline{\nabla}_u \tilde{\nu}$ for any normal extension $\tilde{\nu}$ of ν. On the unit sphere $S^d \subset \mathbb{R}^{d+1}$ with outward normal $\nu = z$, $W_\nu = -\mathrm{Id}_{T_zM}$.

Mean curvature vector. The mean curvature vector at z is

$$H(z) \doteq \sum_{i=1}^{d} \mathrm{II}_z(e_i, e_i) \in N_zM,$$

for any orthonormal basis $(e_i)_{i=1}^d$ of T_zM (the sum is basis-independent). The operator $W_{H(z)}$, obtained by first forming the mean curvature vector and then taking its Weingarten operator, appears in the extrinsic correction. On S^d with outward normal $\nu = z$, $H(z) = -dz$ and hence $W_{H(z)} = d\mathrm{Id}_{T_zM}$.

Ricci endomorphism. The intrinsic Ricci tensor Ric of (M, g_M) is raised to an endomorphism $\mathrm{Ric}^\sharp_z : T_zM \to T_zM$ via the metric: $\left\langle \mathrm{Ric}^\sharp_z\, u, v\right\rangle = \mathrm{Ric}_z(u, v)$. On S^d, $\mathrm{Ric}^\sharp_z = (d-1)\mathrm{Id}_{T_zM}$.

Gauss equation. The identity linking intrinsic and extrinsic curvature used in the paper is the *Gauss equation* for a submanifold $M^d \subset \mathbb{R}^D$ with orthonormal normal frame $(n_\alpha)_{\alpha=1}^{D-d}$:

$$\mathrm{Ric}^\sharp_z = W_{H(z)} - \mathcal{S}_z, \qquad \mathcal{S}_z \doteq \sum_{\alpha=1}^{D-d} W^2_{n_\alpha}. \tag{23}$$

This identity lets us rewrite the extrinsic operator $\frac{1}{2}W_{H(z)} - \mathrm{Ric}^\sharp_z$ equivalently as $\mathcal{S}_z - \frac{1}{2}W_{H(z)}$; Remark 5.3 uses this alternative form.

Reach and tubular coordinates. Positive reach, $\mathrm{reach}(M) > 0$, is the largest r_0 such that every point $x \in \mathbb{R}^D$ with $\mathrm{dist}(x, M) < r_0$ has a unique nearest point $\pi(x) \in M$. For any $r_0 < \mathrm{reach}(M)$, the map

$$\Psi : \{(y, u) : y \in M,\ u \in N_yM,\ \|u\| < r_0\} \to \mathrm{Tub}_{r_0}(M), \qquad \Psi(y, u) \doteq y + u,$$

is a C^3 diffeomorphism onto the open r_0-tube around M. Its inverse is $(\pi(x), x - \pi(x))$.

Exponential map. The intrinsic exponential map at $z \in M$ is $\mathrm{Exp}_z : T_zM \supseteq B(0,r) \to M$, defined by $\mathrm{Exp}_z(v) = \gamma_v(1)$ where γ_v is the geodesic with $\gamma_v(0) = z$ and $\dot\gamma_v(0) = v$. For r smaller than the injectivity radius at z, Exp_z is a C^3 diffeomorphism onto its image; normal coordinates on M near z are its inverse.

B. Derivation of the Fiber Posterior Normal Form

This appendix establishes Proposition C.2 by a sequence of exact changes of variables.

Roadmap. The Federer–Gray tube map provides a global C^3 diffeomorphism between a neighborhood of M in $\mathbb{R}^D$ and the normal bundle, with an explicit Jacobian controlled by the Weingarten operator. Normal coordinates on M centered at z decompose the chord $F_z(v) - z$ into its tangential and normal components, with the tangential part agreeing with v up to a cubic odd remainder and the normal part given to leading order by the second fundamental form. Using these two changes of variables, one obtains an explicit expression for the joint density of (Z, X) in tubular–normal coordinates. Then, integrating out the ambient normal coordinate of X yields the conditional law of $V_z = \mathrm{Exp}_z^{-1}(Z)$ given $\pi(X) = z$, whose heart is an $(D-d)$-dimensional Gaussian integral $I_{\sigma,z}(v)$ against the tube Jacobian. The substitution trick that defines $\Lambda_{\sigma,z}(v)$ absorbs the tangential distortion of the chord into a single multiplicative correction factor whose deviation from one is quartic in v, not cubic; this is where the normal-form claim (38) becomes explicit. Finally, polynomial moment bounds on $\mu_{\sigma,z}$ follow from compactness of M, strict positivity of q, and Gaussian tail estimates, and the same tail estimate shows that the event $X \notin \mathrm{Tub}_{r_0}(M)$ contributes only $O(\exp(-c/\sigma^2))$ and hence does not affect any of the polynomial-in-σ expansions used elsewhere.

Throughout this appendix, fix $z \in M$, and recall $F_z(v) = \mathrm{Exp}_z(v)$. Constants C_k, C'_k depend only on M, $\mathrm{reach}(M)$, $\|q\|_{C^4(M)}$, $\inf_M q$, and ambient bounds on the second fundamental form; in particular not on z or σ.

Since $r_0 < \mathrm{reach}(M)$, the Federer tube map

$$\Psi : \{(y,u) : y \in M,\ u \in N_yM,\ \|u\| < r_0\} \to \mathrm{Tub}_{r_0}(M), \qquad \Psi(y,u) \doteq y + u$$

is a C^3 diffeomorphism onto its image, with inverse $(\pi(x), x - \pi(x))$. Choose an orthonormal frame $(e_1, \ldots, e_d)$ of TM on a neighborhood $U \ni z$ and complete it to an ambient orthonormal frame $(e_1, \ldots, e_d, n_{d+1}, \ldots, n_D)$ smoothly in $y \in U$, with $n_\alpha(y) \in N_yM$. Writing $u = \sum_\alpha u^\alpha n_\alpha(y)$, the Federer–Gray tube formula on positive reach yields

$$dx = \det\big(I_d - W_u(y)\big)\, d\,\mathrm{Vol}_M(y)du, \tag{24}$$

where $W_u(y) : T_yM \to T_yM$ is the Weingarten operator in the normal direction u, so $W_u(y)$ is linear in u and $\mathrm{tr}(W_u(y)) = \langle H(y), u\rangle$ with $H(y) \in N_yM$ the mean curvature vector. In particular,

$$J_z^{\mathrm{tub}}(u) \doteq \det\big(I_d - W_u(z)\big) = 1 - \langle H(z), u\rangle + \mathcal{Q}_z(u), \tag{25}$$

where $\mathcal{Q}_z(u) = O(\|u\|^2)$ uniformly in z, and $J_z^{\mathrm{tub}}(u)$ is bounded above and below by positive constants on $\|u\| < r_0$.

Passing to normal coordinates on M at z, write $y = F_z(v)$. The Riemannian volume satisfies

$$d\,\mathrm{Vol}_M(F_z(v)) = J_z(v)dv, \qquad J_z(v) = 1 - \tfrac{1}{6}\langle \mathrm{Ric}_z\, v, v\rangle + O(\|v\|^3), \tag{26}$$

so in particular $J_z(0) = 1$, $\nabla_v J_z(0) = 0$, and $J_z(v) = 1 + O(\|v\|^2)$ uniformly in z. Decomposing the chord along the tangent and normal spaces at z,

$$F_z(v) - z = G_z(v) + N_z(v), \quad G_z(v) \doteq P_T(z)(F_z(v) - z), \quad N_z(v) \doteq P_\perp(z)(F_z(v) - z), \tag{27}$$

Lemma C.1 gives $G_z(v) = v + R_{3,z}(v) + R_{\geq 4,z}(v)$ with $R_{3,z}(v) = O(\|v\|^3)$, $R_{\geq 4,z}(v) = O(\|v\|^4)$, and the cubic part odd: $R_{3,z}(-v) = -R_{3,z}(v)$. The Gauss lemma and the standard expansion of the exponential map yield

$$N_z(v) = \tfrac{1}{2}\mathrm{II}_z(v,v) + O(\|v\|^3), \tag{28}$$

where $\mathrm{II}_z(v,v) \in N_zM$ is the second fundamental form, so $N_z(v)$ is even in v to leading order and purely quadratic.

The joint law of (Z, X) has density, with respect to $d\,\mathrm{Vol}_M(Z) \otimes dX$,

$$f_{Z,X}(y,x) = q(y) \cdot (2\pi\sigma^2)^{-D/2} \exp\left(-\frac{\|x-y\|^2}{2\sigma^2}\right).$$

Changing variables in x via Ψ and using (24),

$$f_{Z,\pi(X),\perp}(y, z^*, u) = q(y) \cdot (2\pi\sigma^2)^{-D/2} \exp\left(-\frac{\|z^* + u - y\|^2}{2\sigma^2}\right) \det\!\big(I_d - W_u(z^*)\big) \tag{29}$$

with respect to $d\,\mathrm{Vol}_M(y) \otimes d\,\mathrm{Vol}_M(z^*) \otimes du$, where $(z^*, u) \in \{(z', u') : z' \in M,\ u' \in N_{z'}M,\ \|u'\| < r_0\}$.

To obtain the conditional law $\mu_{\sigma,z}(dv)$ of $V_z = \mathrm{Exp}_z^{-1}(Z)$ given $\pi(X) = z$, parameterize $y = F_z(v)$ and integrate out the normal coordinate u of X in (29) evaluated at $z^* = z$:

$$\mathcal{Z}_{\sigma,z}\mu_{\sigma,z}(dv) = q(F_z(v))J_z(v) \cdot k_{\sigma,z}(v)dv, \tag{30}$$

where

$$k_{\sigma,z}(v) \doteq (2\pi\sigma^2)^{-D/2} \int_{N_zM} \exp\left(-\frac{\|F_z(v) - z - u\|^2}{2\sigma^2}\right) J_z^{\mathrm{tub}}(u)du,$$

and $\mathcal{Z}_{\sigma,z}$ is the marginal density of $\pi(X)$ at z. By (27) and the orthogonality $T_zM \perp N_zM$,

$$\|F_z(v) - z - u\|^2 = \|G_z(v)\|^2 + \|N_z(v) - u\|^2,$$

so the tangential and normal components of the quadratic form decouple and

$$k_{\sigma,z}(v) = (2\pi\sigma^2)^{-d/2} \exp\left(-\frac{\|G_z(v)\|^2}{2\sigma^2}\right) \cdot I_{\sigma,z}(v), \tag{31}$$

where

$$I_{\sigma,z}(v) \doteq \int_{N_zM} \phi_\sigma^{(D-d)}(u - N_z(v)) J_z^{\mathrm{tub}}(u)du, \tag{32}$$

and $\phi_\sigma^{(D-d)}$ is the isotropic Gaussian density on $N_zM \simeq \mathbb{R}^{D-d}$.

The tube integral (32) is an ambient Gaussian expectation with mean $N_z(v)$ and isotropic variance $\sigma^2 I_{D-d}$. By (25) and Gaussian-moment calculations,

$$\begin{aligned} I_{\sigma,z}(v) &= 1 - \langle H(z), \mathbb{E}[u]\rangle + \mathbb{E}[\mathcal{Q}_z(u)] \\ &= 1 - \langle H(z), N_z(v)\rangle + \mathcal{Q}_z(N_z(v)) + \sigma^2 \cdot \tfrac{1}{2}\operatorname{tr}(D^2\mathcal{Q}_z(0)) + O(\sigma^4 + \sigma^2\|N_z(v)\|^2). \end{aligned}$$

By (28), $\langle H(z), N_z(v)\rangle = O(\|v\|^2)$, and every term above is $1 + O(\sigma^2 + \|v\|^2)$ with vanishing first-order part in v. The tangential part of the chord is not exactly v, however, and the mismatch must be absorbed to expose the Gaussian density γ_σ. Define

$$\Lambda_{\sigma,z}(v) \doteq \exp\left(\frac{\|v\|^2 - \|G_z(v)\|^2}{2\sigma^2}\right) \cdot I_{\sigma,z}(v). \tag{33}$$

Writing $R_z(v) \doteq R_{3,z}(v) + R_{\geq 4,z}(v) = O(\|v\|^3)$ so that $G_z(v) = v + R_z(v)$, we have $\|G_z(v)\|^2 = \|v\|^2 + 2\langle v, R_z(v)\rangle + \|R_z(v)\|^2$, whence the exponent in (33) is $O(\|v\|^4/\sigma^2)$. Since we will apply this only on the range $\|v\| \lesssim \sigma\sqrt{\log(1/\sigma)}$ (the Gaussian typical set), this exponent is $O(\sigma^2\log^2(1/\sigma))$, and in particular

$$\Lambda_{\sigma,z}(v) = 1 + O(\sigma^2 + \|v\|^2), \qquad \nabla_v\Lambda_{\sigma,z}(0) = 0,$$

uniformly in z, as claimed in (41). Substituting (31) and (33) into (30),

$$\mu_{\sigma,z}(dv) = \frac{1}{\mathcal{Z}'_{\sigma,z}}\gamma_\sigma(v)q(F_z(v))J_z(v)\Lambda_{\sigma,z}(v)dv, \tag{34}$$

with $\mathcal{Z}'_{\sigma,z}$ the normalizing constant of (34). This is exactly (38), proving the normal-form claim.

For the moment bounds (42), write $a_{\sigma,z}(v) = q(F_z(v))J_z(v)\Lambda_{\sigma,z}(v)$. By compactness of M, strict positivity of q, and the expansions above, there exist constants $c_0, C_0 > 0$ and $\sigma_0 > 0$ such that $c_0 \le a_{\sigma,z}(v) \le C_0$ for all $z \in M$, $\sigma \in (0, \sigma_0]$, and all $\|v\| < r_0/2$. Outside $\{\|v\| < r_0/2\}$, the Gaussian factor $\gamma_\sigma(v)$ decays faster than any polynomial in σ, contributing only $O(\exp(-c/\sigma^2))$ to any moment. Hence, for every integer $k \ge 1$,

$$\int \|v\|^k \mu_{\sigma,z}(dv) \le \frac{C_0}{c_0} \int \|v\|^k \gamma_\sigma(v)dv + O(e^{-c/\sigma^2}) \le C_k \sigma^k,$$

proving (42) and completing the proof of Proposition C.2. □

Remark B.1 (Restriction to $\mathcal{E}_\sigma$)**.** The identities above are derived on the event $\mathcal{E}_\sigma \doteq \{X \in \mathrm{Tub}_{r_0}(M)\}$. Under the Gaussian corruption model with $Z \in M$ compact, the complement $\mathcal{E}_\sigma^c$ has probability $O(\exp(-c/\sigma^2))$ for some $c > 0$ depending only on r_0. Since every moment of T_σ is polynomial in σ^{-1}, the contribution of $\mathcal{E}_\sigma^c$ is negligible in every expansion of the form $O(\sigma^k)$ used in Sections 3 to 4.3.

C. Proof of the Leading-Order Identification and Variance Collapse

This appendix proves (11) and Theorem 4.2 of Section 4. The argument reduces both statements to a tubular-coordinate Bayes calculation combined with a manifold Stein identity. It is self-contained given Proposition C.2 below, whose tubular-coordinate proof is in Appendix B.

C.1. Local Coordinates on a Curved Manifold

Fix $z \in M$. Let

$$F_z(v) \doteq \mathrm{Exp}_z(v), \qquad v \in T_zM,$$

be geodesic normal coordinates on a sufficiently small neighborhood of $0 \in T_zM$. Define the tangent chord map

$$G_z(v) \doteq P_T(z)\big(F_z(v) - z\big) \in T_zM. \tag{35}$$

Lemma C.1 (Chord expansion)**.** *Uniformly in $z \in M$, for v sufficiently small,*

$$G_z(v) = v + R_{3,z}(v) + R_{\ge 4,z}(v), \qquad R_{3,z}(v) = O(\|v\|^3), \qquad R_{\ge 4,z}(v) = O(\|v\|^4), \tag{36}$$

where the cubic part $R_{3,z}$ is odd: $R_{3,z}(-v) = -R_{3,z}(v)$.

Proof. The ambient Taylor expansion of the exponential map at z is

$$F_z(v) = z + v + \tfrac{1}{2}\mathrm{II}_z(v,v) + C_z(v,v,v) + O(\|v\|^4),$$

where II_z is the (symmetric, quadratic) second fundamental form and C_z is a symmetric trilinear form valued in $\mathbb{R}^D$ coming from the third Taylor coefficient of F_z. Applying $P_T(z)$ annihilates the normal quadratic term $\frac{1}{2}\mathrm{II}_z(v,v) \in N_zM$. The tangent component of the cubic term, $R_{3,z}(v) \doteq P_T(z)C_z(v,v,v)$, is odd in v because C_z is symmetric-trilinear and $v \mapsto (v,v,v)$ changes sign under $v \mapsto -v$. Higher-order corrections are collected in $R_{\ge 4,z}(v) = O(\|v\|^4)$. □

On the event $\{\pi(X) = z\}$, define

$$V_z \doteq \mathrm{Exp}_z^{-1}(Z) \in T_zM.$$

Then (5) becomes

$$T_\sigma = \frac{1}{\sigma^2} G_z(V_z) \qquad \text{on } \{\pi(X) = z\}. \tag{37}$$

C.2. Fiber Posterior Normal Form

Proposition C.2 (Fiber posterior normal form)**.** *There exists $\sigma_0 > 0$ such that for each $\sigma \in (0, \sigma_0]$ and each $z \in M$, the conditional law of V_z given $\pi(X) = z$ admits a density of the form*

$$\mu_{\sigma,z}(dv) = \frac{1}{\mathcal{Z}_{\sigma,z}} \gamma_\sigma(v) a_{\sigma,z}(v) dv, \tag{38}$$

where

$$\gamma_\sigma(v) \doteq (2\pi\sigma^2)^{-d/2} \exp\left(-\frac{\|v\|^2}{2\sigma^2}\right)$$

and

$$a_{\sigma,z}(v) = q(F_z(v))J_z(v)\Lambda_{\sigma,z}(v). \tag{39}$$

Here J_z is the geodesic-coordinate volume Jacobian and $\Lambda_{\sigma,z}$ is a smooth positive correction factor satisfying, uniformly in z,

$$J_z(v) = 1 + O(\|v\|^2), \qquad \nabla_v J_z(0) = 0, \tag{40}$$

$$\Lambda_{\sigma,z}(v) = 1 + O(\sigma^2 + \|v\|^2), \qquad \nabla_v \Lambda_{\sigma,z}(0) = 0. \tag{41}$$

Moreover, for every integer $k \geq 1$,

$$\sup_{z\in M} \int \|v\|^k \mu_{\sigma,z}(dv) \leq C_k \sigma^k. \tag{42}$$

Proof. See Appendix B. The proof is a tubular-coordinate Bayes calculation. The main point is that conditioning on $\pi(X) = z$ produces a posterior over latent tangent displacements v whose leading factor is the centered Gaussian $\exp(-\|v\|^2/(2\sigma^2))$, while all geometric corrections are smooth and even to first order. □

C.3. A Manifold Stein Identity

Lemma C.3 (Posterior Stein identity)**.** *For every $z \in M$,*

$$\frac{1}{\sigma^2}\mathbb{E}_{\mu_{\sigma,z}}[v] = \mathbb{E}_{\mu_{\sigma,z}}[\nabla_v \log a_{\sigma,z}(v)]. \tag{43}$$

Proof. Fix an orthonormal basis in $T_zM \cong \mathbb{R}^d$. For coordinate i,

$$0 = \int \partial_i\big(\gamma_\sigma(v)a_{\sigma,z}(v)\big)dv$$

because $\gamma_\sigma(v)a_{\sigma,z}(v)$ decays rapidly. Using

$$\partial_i\gamma_\sigma(v) = -\frac{v_i}{\sigma^2}\gamma_\sigma(v),$$

we obtain

$$0 = \int \left(\partial_i a_{\sigma,z}(v) - \frac{v_i}{\sigma^2}a_{\sigma,z}(v)\right)\gamma_\sigma(v)dv.$$

Divide by the normalizing constant $\mathcal{Z}_{\sigma,z}$ to get

$$\frac{1}{\sigma^2}\mathbb{E}[v_i] = \mathbb{E}[\partial_i \log a_{\sigma,z}(v)].$$

Collecting coordinates gives (43). □

C.4. Proof of the Leading-Order Expansion (11)

Proof of (11). Fix $z \in M$. From (37),

$$r_\sigma(z) = \frac{1}{\sigma^2}\mathbb{E}_{\mu_{\sigma,z}}[G_z(v)].$$

Using Lemma C.1, $G_z(v) = v + R_z(v)$ with $R_z(v) \doteq R_{3,z}(v) + R_{\geq 4,z}(v)$, where $R_{3,z}(v) = O(\|v\|^3)$ is odd in v and $R_{\geq 4,z}(v) = O(\|v\|^4)$ is a higher-order remainder, we obtain

$$r_\sigma(z) = \frac{1}{\sigma^2}\mathbb{E}[v] + \frac{1}{\sigma^2}\mathbb{E}[R_z(v)]. \tag{44}$$

By Lemma C.3,

$$\frac{1}{\sigma^2}\mathbb{E}[v] = \mathbb{E}[\nabla_v \log a_{\sigma,z}(v)].$$

Taylor-expand $\nabla_v \log a_{\sigma,z}(v)$ around 0:

$$\nabla_v \log a_{\sigma,z}(v) = \nabla_v \log a_{\sigma,z}(0) + B_{\sigma,z} v + O(\|v\|^2),$$

where $B_{\sigma,z}$ is uniformly bounded in z and σ for sufficiently small σ. Taking expectation and using (42),

$$\frac{1}{\sigma^2}\mathbb{E}[v] = \nabla_v \log a_{\sigma,z}(0) + B_{\sigma,z}\mathbb{E}[v] + O(\mathbb{E}\,\|v\|^2).$$

Since $\mathbb{E}\,\|v\|^2 = O(\sigma^2)$, the left-hand side is bounded, so $\mathbb{E}[v] = O(\sigma^2)$ and therefore

$$\frac{1}{\sigma^2}\mathbb{E}[v] = \nabla_v \log a_{\sigma,z}(0) + O(\sigma^2). \tag{45}$$

Now use (39):

$$\nabla_v \log a_{\sigma,z}(0) = \nabla_v \log q(F_z(v))\big|_{v=0} + \nabla_v \log J_z(0) + \nabla_v \log \Lambda_{\sigma,z}(0).$$

By construction,

$$\nabla_v \log q(F_z(v))\big|_{v=0} = \nabla_M \log q(z).$$

By (40) and (41),

$$\nabla_v \log J_z(0) = 0, \qquad \nabla_v \log \Lambda_{\sigma,z}(0) = 0.$$

Hence

$$\frac{1}{\sigma^2}\mathbb{E}[v] = \nabla_M \log q(z) + O(\sigma^2). \tag{46}$$

The cubic part $R_{3,z}$ has zero Gaussian expectation by oddness; under the posterior density (38), the first nonzero contribution from $R_{3,z}$ comes from the odd linear perturbation of $a_{\sigma,z}(v)$, which is order $\|v\|$. The quartic remainder $R_{\geq 4,z}(v) = O(\|v\|^4)$ contributes $O(\sigma^4)$ directly by moment bounds. Combining with $R_{3,z}(v) = O(\|v\|^3)$ and Gaussian moments yields

$$\mathbb{E}[R_z(v)] = O(\sigma^4),$$

uniformly in z (a more explicit estimate is given in Appendix D). Therefore

$$\frac{1}{\sigma^2}\mathbb{E}[R_z(v)] = O(\sigma^2). \tag{47}$$

Substituting (46) and (47) into (44) gives

$$r_\sigma(z) = \nabla_M \log q(z) + O(\sigma^2),$$

uniformly in z. □

C.5. Proof of Theorem 4.2

Proof of Theorem 4.2. Fix $z \in M$. By (37),

$$T_\sigma = \frac{1}{\sigma^2} G_z(V_z) \qquad \text{conditionally on } \pi(X) = z.$$

Hence

$$\mathrm{Var}(T_\sigma \mid \pi(X) = z) = \frac{1}{\sigma^4}\,\mathrm{Var}_{\mu_{\sigma,z}}(G_z(v)).$$

By Lemma C.1,

$$\|G_z(v)\|^2 = \|v\|^2 + O(\|v\|^4).$$

Using (42),

$$\mathbb{E}\,\|v\|^2 = d\sigma^2 + O(\sigma^4), \qquad \mathbb{E}\,\|v\|^4 = O(\sigma^4),$$

so

$$\mathbb{E}\,\|G_z(v)\|^2 = d\sigma^2 + O(\sigma^4). \tag{48}$$

On the other hand,

$$\mathbb{E}[G_z(v)] = \sigma^2 r_\sigma(z) = \sigma^2 \nabla_M \log q(z) + O(\sigma^4)$$

by (11), hence

$$\|\mathbb{E}[G_z(v)]\|^2 = O(\sigma^4). \tag{49}$$

Subtracting (49) from (48),

$$\mathrm{Var}_{\mu_{\sigma,z}}(G_z(v)) = d\sigma^2 + O(\sigma^4).$$

Dividing by σ^4 yields (12).

Finally, the law of total variance applied to T_σ with conditioning on $\pi(X)$ gives

$$\mathrm{Var}(T_\sigma) = \mathrm{Var}(r_\sigma(\pi(X))) + \mathbb{E}[\mathrm{Var}(T_\sigma \mid \pi(X))].$$

Averaging (12) gives (13). Since (11) implies

$$r_\sigma(\pi(X)) = \nabla_M \log q(\pi(X)) + O_{L^2}(\sigma^2),$$

the variance of the Rao-Blackwellized target remains $O(1)$. □

D. A More Explicit Estimate for the Chord Correction

Lemma D.1 (Chord Remainder Moment)**.** *Let $R_z(v) \doteq R_{3,z}(v) + R_{\geq 4,z}(v) = G_z(v) - v$ be the chord remainder of Lemma C.1, with $R_{3,z}$ the odd cubic part and $R_{\geq 4,z}(v) = O(\|v\|^4)$, and let $\mu_{\sigma,z}$ be the fiber posterior of Proposition C.2. Then, uniformly in $z \in M$,*

$$\mathbb{E}_{\mu_{\sigma,z}}[R_z(v)] = O(\sigma^4).$$

Proof. Expand the posterior density as

$$a_{\sigma,z}(v) = a_{\sigma,z}(0)\big(1 + \ell_z(v) + m_{\sigma,z}(v)\big), \qquad \ell_z(v) = \langle b_z, v\rangle\,, \quad m_{\sigma,z}(v) = O(\|v\|^2 + \sigma^2),$$

and split the expectation into cubic and higher-order pieces.

Cubic part. Since $R_{3,z}(v) = O(\|v\|^3)$ is odd in v,

$$\int R_{3,z}(v)\gamma_\sigma(v)dv = 0 \quad \text{and} \quad \int R_{3,z}(v)\,(\text{even part of } m_{\sigma,z})\,\gamma_\sigma(v)dv = 0.$$

The first nonzero contribution comes from the linear piece, $\int R_{3,z}(v)\ell_z(v)\gamma_\sigma(v)dv$, which is $O(\sigma^4)$ by Gaussian moment bounds; all remaining cubic-part contributions are of order at least σ^5.

Higher-order part. Since $R_{\geq 4,z}(v) = O(\|v\|^4)$, a direct moment bound gives $\mathbb{E}_{\mu_{\sigma,z}}[R_{\geq 4,z}(v)] = O(\sigma^4)$ without needing any cancellation.

Adding the two pieces yields the stated $O(\sigma^4)$ bound. □

E. Alternative Intrinsic Target Based on the Logarithmic Map

An intrinsically cleaner target is

$$\widetilde{T}_\sigma \doteq \frac{1}{\sigma^2}\,\mathrm{Exp}^{-1}_{\pi(X)}(Z) \in T_{\pi(X)}M.$$

In local coordinates, $\widetilde{T}_\sigma = V_z/\sigma^2$ on $\{\pi(X) = z\}$. The following L^2 estimate shows that T_σ and $\widetilde{T}_\sigma$ agree at leading order.

Proposition E.1 (Equivalence of Ambient and Logmap Targets)**.** *Under the assumptions of Section 3,*

$$T_\sigma - \widetilde{T}_\sigma = O_{L^2}(\sigma).$$

Proof. By Lemma C.1, $T_\sigma - \widetilde{T}_\sigma = R_z(V_z)/\sigma^2$. Since $R_z(V_z) = O(\|V_z\|^3)$ and $\mathbb{E}\|V_z\|^2 = O(\sigma^2)$,

$$\mathbb{E}\left\|T_\sigma - \widetilde{T}_\sigma\right\|^2 = \frac{1}{\sigma^4}\mathbb{E}O(\|V_z\|^6) = O(\sigma^2).$$

□

We chose T_σ in the main text because it arises directly from ambient denoising score matching.

F. Second-Order Refinement

The proof of (11) shows that every first-order geometric correction vanishes by symmetry. In this appendix we compute the σ^2 coefficient exactly in the flat case and state the structural form of the curved-case expansion. The flat computation is a one-line consequence of the exact Tweedie identity (15); the curved statement follows the same strategy as (11) with one additional order of Taylor expansion.

F.1. Exact Flat-Case Coefficient

Proposition F.1 (Flat-case σ^2 expansion)**.** *Assume the flat setting of Section 5.1. Assume further that $q \in C^5(V)$ is strictly positive with $\left\|\nabla_V^k \log q\right\|_\infty < \infty$ for $k \leq 5$. Then, uniformly on compact subsets of V,*

$$r_\sigma(t) = \nabla_V \log q(t) + \frac{\sigma^2}{2}\nabla_V\left[\Delta_V \log q + \|\nabla_V \log q\|^2\right](t) + O(\sigma^4), \qquad \sigma \to 0^+. \tag{50}$$

Proof. By (15), $r_\sigma(t) = \nabla_V \log p_T(t)$ where $p_T = q * \phi_\sigma^{(d)}$. A fourth-order Taylor expansion of q inside the convolution gives

$$p_T(t) = \mathbb{E}_{w\sim\mathcal{N}(0,\sigma^2 I_d)}[q(t-w)] = q(t) + \frac{\sigma^2}{2}\Delta_V q(t) + \frac{\sigma^4}{8}\Delta_V^2 q(t) + O(\sigma^6).$$

Dividing by $q(t) > 0$,

$$\frac{p_T(t)}{q(t)} = 1 + \frac{\sigma^2}{2}\frac{\Delta_V q(t)}{q(t)} + O(\sigma^4).$$

Using the identity $\Delta_V q/q = \Delta_V \log q + \|\nabla_V \log q\|^2$,

$$\log p_T(t) - \log q(t) = \frac{\sigma^2}{2}\left[\Delta_V \log q(t) + \|\nabla_V \log q(t)\|^2\right] + O(\sigma^4).$$

Taking ∇_V of both sides yields (50). The fifth-derivative bound on $\log q$ controls the gradient of the $O(\sigma^4)$ remainder, giving uniform $O(\sigma^4)$ control of the score remainder on compact sets. □

The Tweedie identity $r_\sigma(t) = \nabla_V \log p_T(t)$ is exact; (50) is the corresponding asymptotic expansion of its $\sigma \to 0^+$ behavior, with the Gaussian-smoothing bias of q entering as the explicit σ^2 correction of the score.

F.2. Structural Form in the Curved Case

For curved M, the expansion (50) picks up two additional geometric contributions: an intrinsic Riemannian smoothing bias depending on the Ricci tensor of (M, g_M), and an extrinsic curvature correction depending on the second fundamental form of the embedding.

Proposition F.2 (Curved-case σ^2 expansion)**.** *Under the assumptions of Section 3, uniformly in $z \in M$,*

$$r_\sigma(z) = \nabla_M \log q(z) + \sigma^2\left[b_q(z) + g_M^{\text{ext}}(z)\right] + \text{o}(\sigma^2), \qquad \sigma \to 0^+, \tag{51}$$

where the intrinsic (flat-Tweedie) term is

$$b_q(z) = \tfrac{1}{2}\nabla_M \Big[\Delta_M \log q + \|\nabla_M \log q\|^2\Big](z), \tag{52}$$

and the extrinsic term is

$$g_M^{\mathrm{ext}}(z) = \Big(\tfrac{1}{2}W_{H(z)} - \mathrm{Ric}^\sharp_z\Big)\big(\nabla_M \log q(z)\big) = \Big(\mathcal{S}_z - \tfrac{1}{2}W_{H(z)}\Big)\big(\nabla_M \log q(z)\big), \tag{53}$$

with W_u the Weingarten operator in normal direction u, $H(z) = \sum_{i=1}^d \mathrm{II}_z(e_i, e_i) \in N_zM$ the mean curvature vector, $\mathrm{Ric}^\sharp_z : T_zM \to T_zM$ the Ricci endomorphism, and $\mathcal{S}_z \doteq \sum_\alpha W_{n_\alpha}^2$ for any orthonormal normal frame $\{n_\alpha\}$. The two expressions in (53) are equivalent via the Gauss equation $\mathrm{Ric}^\sharp_z = W_{H(z)} - \mathcal{S}_z$. In the flat case $M = V$ one has $\mathrm{II} \equiv 0$, so $g_M^{\mathrm{ext}} \equiv 0$ and (51) reduces to Proposition F.1.

Proof outline; the full coordinate derivation and uniformity bounds are given in Appendix F.3. We work in graph coordinates at z: choose orthonormal frames $\{e_i\}$ of T_zM and $\{n_\alpha\}$ of N_zM, and parametrize a neighborhood of z in M by $s \mapsto y(s) \doteq z + s + h(s) \in \mathbb{R}^D$, where $s \in T_zM$ is tangential and $h(s) \in N_zM$ satisfies $h(0) = 0$, $dh(0) = 0$, $\partial_i\partial_j h^\alpha(0) = \mathrm{II}^\alpha_{ij}(z)$. A direct computation from $g_{ij}(s) = \delta_{ij} + \sum_\alpha \partial_i h^\alpha(s)\partial_j h^\alpha(s)$ and the Taylor expansion $\partial_i h^\alpha(s) = \mathrm{II}^\alpha_{ik}(z)s^k + O(\|s\|^2)$ gives the induced graph-Jacobian expansion

$$J_{\mathrm{gr}}(s) \doteq \sqrt{\det g(s)}, \qquad J_{\mathrm{gr}}(s) = 1 + \tfrac{1}{2}\langle \mathcal{S}_z s, s\rangle + O(\|s\|^3), \qquad \mathcal{S}_z = \sum_\alpha W_{n_\alpha}^2. \tag{54}$$

This graph-coord expansion differs from the classical $\log\sqrt{\det g} = -\frac{1}{6}\mathrm{Ric}(s,s) + O$ of geodesic normal coordinates; graph coordinates use the ambient tangent offset s, not arclength, and the quadratic term is the Gram form $\sum_\alpha \mathrm{II}^\alpha(s,\cdot)\mathrm{II}^\alpha(s,\cdot)$, i.e. $\langle \mathcal{S}s, s\rangle$, rather than the Ricci form. The chord projection is trivial: $P_{T_zM}(y(s) - z) = s$ by construction, so the tangent component of $r_\sigma(z)$ equals the posterior mean of s/σ^2 exactly, without any separate chord correction.

The second ingredient is the fiber factor. Given $\pi(X) = z$, write $X = z + \sum_\alpha u_\alpha n_\alpha$ with $u \in N_zM \cong \mathbb{R}^{D-d}$; the ambient volume element in tubular coordinates contributes the tube Jacobian $\det(I - \sum_\alpha u_\alpha W_{n_\alpha})$, and the ambient Gaussian factors as $e^{-\|s\|^2/(2\sigma^2)}e^{-\|u-h(s)\|^2/(2\sigma^2)}$. Integrating u out gives the fiber factor

$$F_\sigma(s) \doteq \int_{N_zM} \phi_{\sigma^2 I}(u - h(s)) \det\big(I - \textstyle\sum_\alpha u_\alpha W_{n_\alpha}\big)du.$$

Substitute $u = h(s) + \eta$ and expand $\det(I - W_{h(s)+\eta})$ as a polynomial in u. Gaussian averaging in η kills all odd moments, so no term of the form $\sigma^2 \cdot$ (linear in s) survives; the remaining s-independent moments absorb into a constant $C_\sigma = 1 + O(\sigma^2)$. The leading s-dependent piece is $-\langle H(z), h(s)\rangle$, and since $h(s) = \frac{1}{2}\sum_\alpha \langle \mathrm{II}^\alpha s, s\rangle n_\alpha + O(\|s\|^3)$ with $\langle H(z), \mathrm{II}(s,s)\rangle = \langle W_{H(z)}s, s\rangle$, we obtain

$$F_\sigma(s) = C_\sigma\big[1 - \tfrac{1}{2}\langle W_{H(z)}s, s\rangle + O(\|s\|^3) + O(\sigma^2\|s\|^2) + O(\sigma^4)\big]. \tag{55}$$

Combining (54) and (55) and using the Gauss equation $\mathrm{Ric}^\sharp_z = W_{H(z)} - \mathcal{S}_z$, the combined geometric weight $M_\sigma(s) \doteq J_{\mathrm{gr}}(s)F_\sigma(s)$ satisfies

$$\log M_\sigma(s) = \mathrm{const} - \tfrac{1}{2}\Big\langle \mathrm{Ric}^\sharp_z\, s, s\Big\rangle + O(\|s\|^3) + O(\sigma^2\|s\|^2) + O(\sigma^4). \tag{56}$$

Writing $a_\sigma(s) \doteq q(y(s))M_\sigma(s) = e^{f(s)+m_\sigma(s)}$ with $f(s) = \lambda(y(s))$ and $m_\sigma(s) = \log M_\sigma(s)$, the Euclidean posterior mean of s under $\gamma_\sigma \cdot a_\sigma$ is given by the standard Gaussian-moment expansion

$$\tfrac{1}{\sigma^2}\mathbb{E}[s] = \nabla f(0) + \tfrac{\sigma^2}{2}\nabla\big(\Delta f + \|\nabla f\|^2\big)(0) - \sigma^2\,\mathrm{Ric}^\sharp_z\,\nabla f(0) + O(\sigma^4), \tag{57}$$

where ∇, Δ are Euclidean in the graph variable s and we used $\nabla m_\sigma(0) = 0$, $\nabla^2 m_\sigma(0) = -\mathrm{Ric}^\sharp_z$. The final step converts the Euclidean graph-coord derivatives of f to intrinsic derivatives of λ on M. Because $Dy(0) = \mathrm{Id}_{T_zM}$ and the Christoffel symbols of the induced metric vanish at $s = 0$ in graph coordinates, $\nabla f(0) = \nabla_M\lambda(z)$ and $\nabla^2 f(0) = \nabla^2_M\lambda(z)$, so

$\nabla \|\nabla f\|^2 (0) = \nabla_M \|\nabla_M \lambda\|^2 (z)$. For the Laplacian, in graph coordinates the Laplace–Beltrami operator is $\Delta_M f = (1/\sqrt{g})\partial_i(\sqrt{g} g^{ij} \partial_j f)$, and a direct computation from (54) with $g^{ij}(s) = \delta_{ij} + O(\|s\|^2)$ gives

$$\Delta_M f = \Delta_s f - \langle W_{H(z)} s, \nabla_s f\rangle + O(\|s\|^2 |\nabla_s f|) + O(\|s\| |\nabla_s^2 f|), \tag{58}$$

so at the base point $\nabla_s \Delta_s f(0) = \nabla_M \Delta_M \lambda(z) + W_{H(z)} \nabla_M \lambda(z)$. Substituting into (57) and collecting the $W_{H(z)}$ shift yields

$$r_\sigma(z) = \ell + \sigma^2 \Big[\tfrac{1}{2} \nabla_M(\Delta_M \log q + \|\nabla_M \log q\|^2) + \big(\tfrac{1}{2} W_{H(z)} - \mathrm{Ric}^\sharp_z\big)\ell\Big] + O(\sigma^4), \tag{59}$$

which is (53), and equivalently $(\mathcal{S}_z - \frac{1}{2} W_{H(z)})\ell$ via the Gauss equation. In the flat case $\mathrm{II} \equiv 0$, both operators vanish and only the flat-Tweedie term survives. Full uniformity bounds (requiring compactness of M, $\mathrm{reach}(M) > 0$, and $\log q \in C^4$) are carried out in Appendix F.3. □

Proposition F.2 shows that the $O(\sigma^2)$ error in (11) is of order σ^2, not smaller; the leading bias has an intrinsic component (present even on a flat support) and a tensorial extrinsic component mixing the mean-curvature Weingarten operator W_H with either the Ricci endomorphism $\mathrm{Ric}^\sharp$ or equivalently the normal-sum operator $\mathcal{S}$. Two a priori independent curvature sources (tube-Jacobian mean curvature and volume-form Ricci) combine without cancellation into the operator $\frac{1}{2} W_H - \mathrm{Ric}^\sharp$. The extrinsic correction vanishes on any totally geodesic submanifold (where $\mathrm{II} \equiv 0$, hence $W_H = \mathrm{Ric}^\sharp = 0$) and whenever $\nabla_M \log q(z) = 0$.

Remark F.3 (Frame formula)**.** Fix orthonormal frames $\{e_i\}_{i=1}^d$ of T_zM and $\{n_\alpha\}_{\alpha=1}^{D-d}$ of N_zM, and write $h^\alpha_{ij} = \langle \mathrm{II}_z(e_i, e_j), n_\alpha\rangle$, $H^\alpha = \sum_i h^\alpha_{ii}$, $\ell_j = e_j(\log q)$. Using the equivalent form $\mathcal{S}_z - \frac{1}{2} W_{H(z)}$ of (53),

$$(g_M^{\mathrm{ext}})_k = \sum_j \Big[\sum_{\alpha,i} h^\alpha_{ki} h^\alpha_{ij} - \tfrac{1}{2} \sum_\alpha H^\alpha h^\alpha_{kj}\Big]\ell_j, \qquad (b_q)_k = \tfrac{1}{2} e_k\big(\Delta_M \log q + \|\nabla_M \log q\|^2\big). \tag{60}$$

The extrinsic term is a tensorial combination of $\mathcal{S}$ and W_H: on a hypersurface with single normal n it reduces to $(W_n^2 - \frac{1}{2} H W_n)\ell$, which is not proportional to $W_H \ell$ in general (e.g. on the asymmetric product $S^1(R_1) \times S^1(R_2)$ in $\mathbb{R}^4$ the operator is $\frac{1}{2}\mathrm{diag}(1/R_1^2, 1/R_2^2)$, which has distinct eigenvalues when $R_1 \neq R_2$).

F.3. General Extrinsic Correction

We now derive the explicit formula (53) for the extrinsic σ^2-coefficient on an arbitrary smooth embedded submanifold $M \subset \mathbb{R}^D$. Fix $z \in M$, write $\lambda = \log q$, $\ell = \nabla_M \lambda(z) \in T_zM$, and let $\{e_i\}_{i=1}^d$ be an orthonormal basis of T_zM and $\{n_\alpha\}_{\alpha=1}^{D-d}$ an orthonormal basis of N_zM. For $u \in N_zM$ let $W_u : T_zM \to T_zM$ be the Weingarten operator, $\langle W_u v, w\rangle = \langle \mathrm{II}_z(v, w), u\rangle$, write $\mathrm{II}^\alpha_{ij} \doteq \langle \mathrm{II}_z(e_i, e_j), n_\alpha\rangle$, so that $W_{n_\alpha} = \mathrm{II}^\alpha$ in the basis $\{e_i\}$, and let $H(z) = \sum_i \mathrm{II}_z(e_i, e_i) \in N_zM$ be the mean curvature vector. The two symmetric operators on T_zM that appear in the final answer are

$$W_{H(z)} = \sum_\alpha H^\alpha \mathrm{II}^\alpha, \qquad \mathcal{S}_z \doteq \sum_\alpha W_{n_\alpha}^2 = \sum_\alpha (\mathrm{II}^\alpha)^2, \qquad H^\alpha \doteq \mathrm{tr}(\mathrm{II}^\alpha), \tag{61}$$

connected to the intrinsic Ricci operator by the Gauss equation $\mathrm{Ric}^\sharp_z = W_{H(z)} - \mathcal{S}_z$.

The derivation proceeds in four steps: We first set up tangent-graph coordinates where the chord map is linear, then expand the two geometric factors (graph Jacobian and fiber factor) to quadratic order, apply a Euclidean Tweedie expansion to the resulting posterior, and finally convert the Euclidean graph-coord derivatives to intrinsic derivatives using a Laplacian mismatch identity.

Parameterize a neighborhood of z in M by

$$y(s) \doteq z + \sum_i s^i e_i + \sum_\alpha h^\alpha(s) n_\alpha, \qquad s = (s^i) \in T_zM, \tag{62}$$

with $h^\alpha(0) = 0$, $\partial_i h^\alpha(0) = 0$, and $\partial_i \partial_j h^\alpha(0) = \mathrm{II}^\alpha_{ij}$. In these coordinates $P_{T_zM}(y(s) - z) = s$ exactly, so the Rao-Blackwell target is the posterior mean of s/σ^2. Let $J_{\mathrm{gr}}(s) \doteq \sqrt{\det g(s)}$ be the graph Jacobian with $g_{ij}(s) = \delta_{ij} + \sum_\alpha \partial_i h^\alpha \partial_j h^\alpha$, and define the fiber factor

$$F_\sigma(s) \doteq \int_{N_zM} \phi_{\sigma^2 I}(u - h(s)) \det\big(I - \textstyle\sum_\alpha u_\alpha W_{n_\alpha}\big) du, \tag{63}$$

where $\phi_{\sigma^2 I}$ is the centered Gaussian density on $N_z M$ with covariance $\sigma^2 I$. The ambient Gaussian factorizes under the orthogonal decomposition $X - Z = -\sum s^i e_i + \sum (u_\alpha - h^\alpha(s)) n_\alpha$ as $\phi_{\sigma^2 I}(X - Z) \propto e^{-\|s\|^2/(2\sigma^2)} e^{-\|u-h(s)\|^2/(2\sigma^2)}$, so the conditional law of s given $\pi(X) = z$ reads

$$r_\sigma(z) = \frac{1}{\sigma^2} \frac{\int s q(y(s)) J_{\mathrm{gr}}(s) F_\sigma(s) e^{-\|s\|^2/(2\sigma^2)} ds}{\int q(y(s)) J_{\mathrm{gr}}(s) F_\sigma(s) e^{-\|s\|^2/(2\sigma^2)} ds}. \tag{64}$$

Lemma F.4 (Quadratic graph Jacobian)**.** *In the graph coordinates (62),*

$$J_{\mathrm{gr}}(s) = 1 + \tfrac{1}{2} \langle \mathcal{S}_z s, s \rangle + O(\|s\|^3).$$

Proof. From $\partial_i h^\alpha(s) = \mathrm{II}^\alpha_{im} s^m + O(\|s\|^2)$ we have $g_{ij}(s) = \delta_{ij} + \sum_{\alpha,m,n} \mathrm{II}^\alpha_{im} \mathrm{II}^\alpha_{jn} s^m s^n + O(\|s\|^3)$. Taking determinants, $J_{\mathrm{gr}}(s) = 1 + \frac{1}{2} \sum_{\alpha,i,m,n} \mathrm{II}^\alpha_{im} \mathrm{II}^\alpha_{in} s^m s^n + O(\|s\|^3)$, and the quadratic form is precisely $\sum_\alpha W^2_{n_\alpha} = \mathcal{S}_z$. □

Lemma F.5 (Quadratic fiber factor)**.** *As $s \to 0$ and $\sigma \to 0$,*

$$F_\sigma(s) = C_\sigma \big[1 - \tfrac{1}{2} \langle W_{H(z)} s, s \rangle + O(\|s\|^3) + O(\sigma^2 \|s\|^2) + O(\sigma^4) \big],$$

where $C_\sigma = 1 + O(\sigma^2)$ is independent of s. In particular, the remainder contains no term linear in s at order σ^2.

Proof. Substitute $u = h(s) + \eta$ in (63) so that $F_\sigma(s) = \int \phi_{\sigma^2 I}(\eta) \det(I - W_{h(s)+\eta}) d\eta$. Expand the determinant about $u = 0$ as a polynomial in the components of u: $\det(I - W_u) = 1 - \langle H(z), u \rangle + Q(u) + C(u) + O(\|u\|^4)$, where Q is homogeneous quadratic in u and C is homogeneous cubic. Taking Gaussian expectation in η (with $\mathbb{E}[\eta] = 0$, $\mathbb{E}[\eta_\alpha \eta_\beta] = \sigma^2 \delta_{\alpha\beta}$, and all odd moments vanishing) yields

$$F_\sigma(s) = 1 - \langle H(z), h(s) \rangle + Q(h(s)) + \sigma^2 \ \mathrm{tr}\big(Q^{\mathrm{quad}}\big) + C(h(s)) + 3\sigma^2 \ \langle \nabla C(0), h(s) \rangle + O\big((\|h(s)\| + \sigma)^4\big),$$

where Q^{quad} is the symmetric matrix representing Q and $\nabla C(0)$ is the vector picking out the quadratic-in-η, linear-in-$h(s)$ part of C. Because $h(s) = \frac{1}{2} \sum_\alpha \langle \mathrm{II}^\alpha s, s \rangle n_\alpha + O(\|s\|^3)$ vanishes to *second* order in s, every occurrence of $h(s)$ in the display above is $O(\|s\|^2)$. Therefore: the $-\langle H(z), h(s) \rangle$ term gives the stated $-\frac{1}{2} \langle W_{H(z)} s, s \rangle + O(\|s\|^3)$ via $\langle H(z), \mathrm{II}(s,s) \rangle = \langle W_{H(z)} s, s \rangle$; $Q(h(s)) = O(\|s\|^4)$ and $C(h(s)) = O(\|s\|^6)$; the s-independent $\sigma^2 \ \mathrm{tr}(Q^{\mathrm{quad}})$ piece is absorbed into C_σ; the $\sigma^2 \langle \nabla C(0), h(s) \rangle$ term is $O(\sigma^2 \|s\|^2)$; and all remaining contributions are $O(\sigma^4)$ or of higher combined order. Crucially, no term of the form $\sigma^2 \cdot$ (linear in s) appears: such a term would require an odd moment of η, which vanishes. Collecting these into $C_\sigma = 1 + O(\sigma^2)$ and the remainder $O(\|s\|^3) + O(\sigma^2 \|s\|^2) + O(\sigma^4)$ gives the claim. □

Proposition F.6 (Combined geometric weight)**.** *Let $M_\sigma(s) \doteq J_{\mathrm{gr}}(s) F_\sigma(s)$. Then*

$$\log M_\sigma(s) = c_\sigma - \tfrac{1}{2} \left\langle \mathrm{Ric}^\sharp_z \, s, s \right\rangle + O(\|s\|^3) + O(\sigma^2 \|s\|^2) + O(\sigma^4),$$

for some scalar c_σ. In particular, $\nabla m_\sigma(0) = 0$ and the remainder contains no term linear in s at order σ^2.

Proof. Multiply the expansions of Lemma F.4 and Lemma F.5 to get $M_\sigma(s) = C_\sigma[1 + \frac{1}{2} \langle (\mathcal{S}_z - W_{H(z)}) s, s \rangle + R(s,\sigma)]$ where $R(s,\sigma) = O(\|s\|^3) + O(\sigma^2 \|s\|^2) + O(\sigma^4)$. Taking the logarithm preserves the quadratic coefficient and the form of the remainder, and the Gauss equation $\mathrm{Ric}^\sharp_z = W_{H(z)} - \mathcal{S}_z$ gives $\mathcal{S}_z - W_{H(z)} = -\mathrm{Ric}^\sharp_z$. □

Lemma F.7 (Euclidean score expansion)**.** *Let $a_\sigma : \mathbb{R}^d \to (0, \infty)$ be smooth in a neighborhood of 0, with uniformly bounded derivatives up to order 4. Then*

$$\frac{1}{\sigma^2} \frac{\int s a_\sigma(s) e^{-\|s\|^2/(2\sigma^2)} ds}{\int a_\sigma(s) e^{-\|s\|^2/(2\sigma^2)} ds} = \nabla \log a_\sigma(0) + \tfrac{\sigma^2}{2} \nabla \big(\Delta \log a_\sigma + \|\nabla \log a_\sigma\|^2 \big)(0) + O(\sigma^4),$$

where ∇, Δ are the Euclidean gradient and Laplacian in the s variables.

Proof. Write $a_\sigma = e^{b_\sigma}$ and Taylor expand b_σ to order 3 at the origin; then expand $e^{b_\sigma(s)-b_\sigma(0)}$ to cubic order. Under the centered isotropic Gaussian, only linear and cubic odd moments survive in the numerator, and the denominator contributes the normalization correction. A direct Gaussian-moment calculation (Proposition F.1) yields the stated coefficient. $\square$

Apply Lemma F.7 to $a_\sigma(s) = q(y(s))M_\sigma(s) = e^{f(s)+m_\sigma(s)}$ with $f(s) \doteq \lambda(y(s))$ and $m_\sigma(s) \doteq \log M_\sigma(s)$. By Proposition F.6, $\nabla m_\sigma(0) = 0$, $\nabla^2 m_\sigma(0) = -\operatorname{Ric}_z^\sharp$, and $\nabla\Delta m_\sigma(0) = 0$ (the leading s-dependent term is quadratic). Therefore

$$r_\sigma(z) = \nabla f(0) + \tfrac{\sigma^2}{2}\nabla\big(\Delta f + \|\nabla f\|^2\big)(0) - \sigma^2 \operatorname{Ric}_z^\sharp \nabla f(0) + O(\sigma^4), \tag{65}$$

where ∇, Δ are Euclidean in the graph variable s.

Lemma F.8 (First and second derivatives)**.** *At $s = 0$,*

$$\nabla f(0) = \nabla_M \lambda(z), \qquad \nabla^2 f(0) = \nabla_M^2 \lambda(z),$$

and consequently $\nabla \|\nabla f\|^2 (0) = \nabla_M \|\nabla_M \lambda\|^2 (z)$.

Proof. Because $Dy(0) = \mathrm{Id}_{T_zM}$ and the Christoffels of g vanish at 0 in graph coordinates (the mixed partials $\partial_i\partial_j y(0) = \sum_\alpha \mathrm{II}_{ij}^\alpha n_\alpha$ are purely normal to T_zM), the first and second coordinate derivatives at the base point agree with the intrinsic gradient and covariant Hessian. $\square$

Lemma F.9 (Laplacian mismatch in graph coordinates)**.** *For any smooth scalar f on the graph patch,*

$$\Delta_M f = \Delta_s f - \big\langle W_{H(z)} s, \nabla_s f\big\rangle + O(\|s\|^2 |\nabla_s f|) + O(\|s\| |\nabla_s^2 f|).$$

In particular, $\nabla_s \Delta_s f(0) = \nabla_M \Delta_M f(0) + W_{H(z)} \nabla_s f(0)$.

Proof. The Laplace–Beltrami operator in local coordinates is $\Delta_M f = (1/\sqrt{g})\partial_i(\sqrt{g} g^{ij} \partial_j f)$. By the proof of Lemma F.4, $g^{ij}(s) = \delta_{ij} + O(\|s\|^2)$ and $\sqrt{g}(s) = 1 + \frac{1}{2}\langle \mathcal{S}_z s, s\rangle + O(\|s\|^3)$. Differentiating, $\partial_i \log\sqrt{g}(s) = (\mathcal{S}_z s)^i + O(\|s\|^2)$ and $\partial_i g^{ij}(s) = O(\|s\|)$, so $(1/\sqrt{g})\partial_i(\sqrt{g} g^{ij}) = (\mathcal{S}_z s)^j - (\mathcal{S}_z s)^j + O(\|s\|^2)$ from the g^{ij} part, leaving the $g^{ij}\partial_i \log\sqrt{g}$ contribution; using the symmetric form of the second fundamental form gives exactly $-(W_{H(z)} s)^j + O(\|s\|^2)$. Substituting into the divergence formula yields the stated expansion; taking the gradient at 0 gives the derivative identity. $\square$

Main theorem. Combining Lemmas F.8 and F.9 with (65),

$$r_\sigma(z) = \ell + \tfrac{\sigma^2}{2}\nabla_M\big(\Delta_M \lambda + \|\nabla_M \lambda\|^2\big)(z) + \sigma^2\big(\tfrac{1}{2} W_{H(z)} - \operatorname{Ric}_z^\sharp\big)\ell + O(\sigma^4). \tag{66}$$

Using the Gauss equation (61) this is equivalently $\sigma^2(\mathcal{S}_z - \frac{1}{2} W_{H(z)})\ell$, establishing (53). Because the chord map in graph coordinates is the identity, no further chord correction is needed: the tangent component of $r_\sigma(z)$ equals $\mathbb{E}[s]/\sigma^2$ exactly.

F.4. Specialization: The Unit Sphere S^d

For the unit sphere $S^d \subset \mathbb{R}^{d+1}$, the second fundamental form in the outward-normal $\nu = z$ convention is $\mathrm{II}_z(v, w) = -\langle v, w\rangle z$, so the single Weingarten $W_\nu = -\mathrm{Id}_{T_z S^d}$, $\mathcal{S}_z = W_\nu^2 = \mathrm{Id}_{T_z S^d}$, and $H(z) = -dz$ gives $W_{H(z)} = d\,\mathrm{Id}_{T_z S^d}$. Substituting into (53) in the form $(\mathcal{S}_z - \frac{1}{2} W_{H(z)})\ell$,

$$g_{S^d}^{\mathrm{ext}}(z) = \big(1 - \tfrac{d}{2}\big)\nabla_{S^d} \log q(z). \tag{67}$$

The full σ^2-coefficient on S^d is therefore

$$b_q(z) + g_{S^d}^{\mathrm{ext}}(z) = \tfrac{1}{2}\nabla_{S^d}\big(\Delta_{S^d} \log q + \|\nabla_{S^d} \log q\|^2\big)(z) + \big(1 - \tfrac{d}{2}\big)\nabla_{S^d} \log q(z). \tag{68}$$

The dimensional pattern $(1 - d/2) = +\frac{1}{2}, 0, -\frac{1}{2}, -1, \ldots$ for $d = 1, 2, 3, 4, \ldots$ is:

- $d = 1$: coefficient $+\frac{1}{2}$;

- $d = 2$: coefficient 0, so $g^{\text{ext}}(z) \equiv 0$ on S^2 and the σ^2-bias is *purely* the flat-Tweedie term $b_q(z)$ even though S^2 is positively curved. This exact vanishing is a consequence of the Einstein property $W_H = 2\,\mathrm{Id}$, $\mathrm{Ric}^\sharp = \mathrm{Id}$ on unit S^2, which balances the volume-form and tube contributions.

- $d = 3$: coefficient $-\frac{1}{2}$, the first nontrivial case with a negative extrinsic correction.

- $d \geq 4$: increasingly negative extrinsic corrections, scaling linearly in d.

Equivalently, using $\mathrm{Ric}^\sharp_z = (d-1)\,\mathrm{Id}$ on unit S^d, the $(\frac{1}{2}W_H - \mathrm{Ric}^\sharp)$ form of (53) gives $(\frac{d}{2} - (d-1))\,\mathrm{Id} = (1 - \frac{d}{2})\,\mathrm{Id}$.

G. A Finite-Sample Rate for Local-Averaging Estimation of r_σ

Equation (11) and Theorem 4.2 are population statements: they describe r_σ and its variance under the true joint law of (Z, X). This appendix gives a corresponding finite-sample statement that connects our variance-collapse identity to an explicit estimation rate.

Let $(Z_1, X_1), \dots, (Z_N, X_N)$ be i.i.d. copies of (Z, X) under (2). For each i, define the observed Rao-Blackwell sample

$$\pi_i \doteq \pi(X_i) \in M, \qquad T_{\sigma,i} \doteq \frac{1}{\sigma^2} P_T(\pi_i)(Z_i - \pi_i) \in T_{\pi_i} M.$$

Fix a bandwidth $h > 0$ and a bounded, nonnegative kernel $K : [0, \infty) \to [0, \infty)$ supported on $[0, 1]$ and bounded away from zero on $[0, 1/2]$. For $z \in M$, define the local-averaging estimator

$$\widehat{r}_\sigma^{(h)}(z) \doteq \frac{\sum_{i=1}^N P_T(z) \mathcal{T}_{\pi_i \to z}(T_{\sigma,i}) K(d_M(\pi_i, z)/h)}{\sum_{i=1}^N K(d_M(\pi_i, z)/h)}, \tag{69}$$

where d_M is the intrinsic distance on M and $\mathcal{T}_{\pi_i \to z} : T_{\pi_i} M \to T_z M$ is parallel transport along the minimizing geodesic. Let f_π denote the density of $\pi(X)$ with respect to $d\,\mathrm{Vol}_M$; by continuity of π and strict positivity of q, f_π is bounded above and below by positive constants on M.

Proposition G.1 (Finite-sample MSE bound). *Under the assumptions of Section 3, there exist constants $C_1, C_2, h_0, \sigma_0 > 0$ depending only on (M, q) and the kernel K such that, for every $z \in M$, every $\sigma \in (0, \sigma_0]$, every $h \in (0, h_0]$, and every N with $Nh^d \geq 1/(C_2 \sigma^2)$,*

$$\mathbb{E}\big\|\widehat{r}_\sigma^{(h)}(z) - r_\sigma(z)\big\|^2 \leq C_1 h^2 + \frac{C_2 d}{\sigma^2 N h^d f_\pi(z)}. \tag{70}$$

Then, the minimax-optimal bandwidth

$$h^\star = \left(\frac{C_2 d}{C_1 \sigma^2 N f_\pi(z)}\right)^{1/(d+2)}$$

yields

$$\mathbb{E}\big\|\widehat{r}_\sigma^{(h^\star)}(z) - r_\sigma(z)\big\|^2 \leq C' \left(\frac{1}{\sigma^2 N}\right)^{2/(d+2)}. \tag{71}$$

In particular, the natural sample-size scaling $N \asymp \sigma^{-(d+2)}$ is both sufficient and necessary for $\widehat{r}_\sigma^{(\sigma)}(z)$ to have bounded variance at bandwidth $h = \sigma$.

Proof. Write $\widehat{r}_\sigma^{(h)}(z)$ as a sum of a bias term and a variance term.

Bias. By (11), r_σ is C^1 on M uniformly in $\sigma \in (0, \sigma_0]$, with Lipschitz constant L bounded independently of σ. Parallel transport along geodesics preserves norms. Hence

$$\big\|\mathbb{E}[\widehat{r}_\sigma^{(h)}(z)] - r_\sigma(z)\big\| \leq Lh,$$

and the squared bias is at most $L^2 h^2$.

Variance. By Theorem 4.2, $\mathrm{Var}(T_{\sigma,i} \mid \pi_i = y) = d/\sigma^2 + O(1)$ uniformly in $y \in M$. The effective sample size inside the bandwidth ball is $N_{\text{eff}} = Nh^d f_\pi(z) \cdot (1 + o_N(1))$, and the standard local-average variance bound gives

$$\mathrm{Var}\big(\widehat{r}_\sigma^{(h)}(z)\big) \le \frac{C_2'(d/\sigma^2 + 1)}{Nh^d f_\pi(z)} \le \frac{C_2 d}{\sigma^2 N h^d f_\pi(z)},$$

for $\sigma \le \sigma_0$. Summing bias^2 and variance yields (70).

Minimizing (70) over h gives $h^\star$ as displayed, and substituting back yields (71). Finally, setting $h = \sigma$ in (70) gives MSE bounded by $C_1\sigma^2 + C_2 d/(N\sigma^{d+2} f_\pi(z))$, whose variance term is $O(1)$ iff $N \gtrsim \sigma^{-(d+2)}$. □

The finite-sample rate in (71) matches, up to constants and logarithmic factors, the standard d-dimensional nonparametric regression rate (Fan, 1992), with the noise level replaced by the conditional-variance constant d/σ^2 from Theorem 4.2. Thus the variance collapse under Rao-Blackwellization directly translates into a better finite-sample rate than an analogous estimator of the raw target T_σ would achieve without the Rao-Blackwell step.