RIEMANNIAN BANDIT CONVEX OPTIMIZATION WITH SELF-CONCORDANT GEOMETRY

Anonymous authors
Paper under double-blind review

ABSTRACT

We study adversarial bandit convex optimization on Riemannian manifolds without a global affine structure. Our intrinsic barrier-proximal method combines exponential-map Dikin exploration with single-step updates based on scalar logarithmic surrogates. A squared-distance analysis avoids aggregating tangent vectors across base points or assuming surrogate convexity. We bound the curvature-induced estimator-to-smoothing support error by $C_{\mathrm{sm}}\delta^2$ under C^2 loss regularity. Conditional on compact-interior iterates and uniform smooth normal-tube data, we establish $\widetilde{\mathcal{O}}(T^{2/3})$ one-point regret against nonanticipating adversaries and provide a two-point extension. Covariance-tracking experiments on symmetric positive definite manifolds demonstrate effective parameter transfer and query-loss improvements over feedback-matched baselines under both one- and two-point feedback.

1 INTRODUCTION

Random smoothing and self-concordant barriers give strong Euclidean bandit online convex optimization (OCO) guarantees (Zinkevich, 2003; Flaxman et al., 2005; Dani et al., 2008; Abernethy et al., 2012; Agarwal et al., 2010; Saha & Tewari, 2011; Duchi et al., 2015; Dekel et al., 2015; Shamir, 2017; Bubeck et al., 2021; Fokkema et al., 2024; Lattimore, 2024). Our goal is to extend this barrier-based template to curved spaces using geodesics, exponential maps, and local curvature.

Motivating domains include positive definite matrices (Bhatia, 2007; Pennec et al., 2006), Stiefel and Grassmann manifolds (Edelman et al., 1998; Absil et al., 2008; Boumal, 2023), and hyperbolic spaces (Nickel & Kiela, 2017; Ganea et al., 2018). Euclidean embedding can inflate dimension, require expensive projections, and destroy geodesic convexity, motivating intrinsic Riemannian algorithms and regret bounds.

We study adversarial bandit OCO on a Riemannian manifold $\mathcal{M}$ over a geodesically convex domain $\mathcal{K}$. A barrier R supplies interior regularization and local exploration geometry. At round t, a tangent direction at the center x_t is scaled by the inverse square root of the Riemannian Hessian $\nabla^2 R(x_t)$ and mapped through Exp_{x_t} to query $f_t(y_t)$. This intrinsic Dikin sampling yields the one-point estimator for a single-step barrier-proximal center update.

The primary estimator challenge is that the fixed-tangent-space smoothing identity does not by itself establish the required support inequality on a curved manifold. The exponential map and moving Dikin geometry introduce an estimator-to-smoothing mismatch. Under C^2 loss regularity and uniform smooth normal-tube data, we bound the resulting support error explicitly by $C_{\mathrm{sm}}\delta^2$. Its cumulative contribution is $C_{\mathrm{sm}}T\delta^2$, and the resulting one-point regret is $\widetilde{\mathcal{O}}(T^{2/3})$ under the stated local conditions. For interpretation, we also retain the smoothed-comparator mismatch $\varepsilon_{\mathrm{sm}}(\delta)$ and verify that it vanishes in the flat, fixed-ellipsoid case.

A second difficulty is that tangent vectors at different base points do not form a common vector space. The update therefore does not sum gradients living in different tangent spaces. Given an estimator $g_t \in T_{x_t}\mathcal{M}$, the algorithm forms the scalar surrogate

$$h_t(x) = \langle g_t, \mathrm{Log}_{x_t}(x)\rangle_{x_t},$$

a real-valued function on $\mathcal{K}$, and uses it only in the current proximal step, together with squared geodesic distance and a small barrier term. This distinction is not cosmetic. On a curved manifold,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

logarithmic "affine" functions are not automatically globally geodesically convex; in fact, affine-type convexity conditions on Hadamard manifolds can be highly restrictive (Kristály et al., 2016). Our stability proof instead uses the proximal first-order condition and a telescoping squared-distance potential; it does not assume convexity of h_t. The guarantees remain conditional on feasible Dikin balls, compact-interior center iterates, and uniform smooth logarithm-map data for the analyzed pairs. The barrier ensures that each proximal minimizer is interior, but does not by itself establish horizon-uniform compact containment. Suitable Hadamard domains and SPD windows allow the local geometric data to be checked; the same requirements must remain explicit on positively curved manifolds.

Contributions. The paper makes three contributions.

(1) One-point regret bounds with isolated geometric error. For barrier-local exploration via the exponential map, we establish a one-point regret bound that separates optimization errors from the geometric estimator-to-smoothing mismatch. A direct support argument bounds its contribution by $C_{\mathrm{sm}}T\delta^2$ under the stated C^2 and uniform local hypotheses. Together with a uniform martingale argument for nonanticipating adversaries, this yields an $\widetilde{\mathcal{O}}(T^{2/3})$ regret guarantee.

(2) Intrinsic barrier-proximal stability. We prove a pathwise regret bound for a one-step proximal update using the current scalar logarithmic surrogate. Its first-order condition and squared-distance potential give a telescoping argument under the existing compact normal-tube conditions, without summing tangent vectors at different iterates or assuming that logarithmic surrogates are convex.

(3) Two-point regret bounds. We give a two-point extension based on symmetric exponential-map queries. The central-difference estimator removes the one-point radius-dependent variance factor and yields $\widetilde{O}(\sqrt{T})$ regret under local symmetry and bounded moment assumptions.

Empirical evaluation. Section 6 first evaluates one-point breast-cancer covariance tracking. With parameters selected at $T = 200$ and frozen, proximal Dikin achieves a 5.08% mean paired query-loss reduction over projected Euclidean at $T = 1000$ on 100 held-out query seeds; the $T = 500$ interval includes zero. Separate retuning and scalar-baseline checks provide additional held-out comparisons (Section 6.1). We then evaluate two-point digits, chronological stock blocks, and synthetic SPD(3)/SPD(5) streams, comparing against projected Euclidean on real data and intrinsic two-point Riemannian OGD on synthetic tasks (Section 6.2).

Each comparison matches the number of function observations, without gradient access. The results demonstrate effective parameter transfer and lower query losses for the selected proximal configurations. For a common comparator class, these improvements also translate into regret advantages. The covariance tasks provide structured SPD benchmarks with strongly geodesically convex squared-distance losses.

Related work and positioning are discussed in Appendix D.

2 Problem Setup

Let $(\mathcal{M}, \langle\cdot,\cdot\rangle)$ be a complete d-dimensional Riemannian manifold with distance $d_{\mathcal{M}}$, exponential map Exp_x, logarithm map Log_x, and tangent space $T_x\mathcal{M}$. The feasible set $\mathcal{K} \subset \mathcal{M}$ is geodesically convex with nonempty relative interior $\mathcal{K}^\circ$. We write $\gamma_{x,y}$ for the unique minimizing geodesic between x and y whenever it is well defined.

At each round $t = 1, \ldots, T$, a nonanticipating adversary chooses a loss $f_t : \mathcal{K} \to \mathbb{R}$ before seeing the learner's fresh random direction. The learner maintains a center $x_t \in \mathcal{K}^\circ$, queries a feasible point $y_t \in \mathcal{K}$, observes only $f_t(y_t)$, and incurs $f_t(y_t)$. The regret is

$$\mathrm{Reg}_T = \mathbb{E}\left[\sum_{t=1}^{T} f_t(y_t) - \inf_{x\in\mathcal{K}} \sum_{t=1}^{T} f_t(x)\right].$$

The expectation is over the learner's randomization and any adversary randomness. The hindsight infimum is inside the expectation because a nonanticipating adversary may depend on past observations. Throughout, $\mathbb{E}_t$ denotes conditioning on the full pre-query history, including x_t and the chosen loss f_t, but not the fresh direction at round t.

Assumption 1 (Geometric regularity)**.** *The set $\mathcal{K}$ lies in a geodesically convex chart with injectivity radius at least $\rho > 0$. Along all geodesics used by the algorithm and comparator argument, the sectional curvature is bounded as $-\kappa_- \le \sec \le \kappa_+$. The logarithm map $\mathrm{Log}_x(y)$ is single valued for all pairs considered in the analysis.*

Assumption 2 (Loss regularity)**.** *Each f_t is geodesically convex, H-smooth along geodesics, and uniformly bounded by $|f_t(x)| \le C$ on the feasible query region. When a Lipschitz bound is needed for boundary-comparator truncation, we assume $|f_t(x) - f_t(y)| \le L d_{\mathcal{M}}(x, y)$.*

Assumption 3 (Barrier)**.** *There is a twice continuously differentiable barrier $R : \mathcal{K}^\circ \to \mathbb{R}$ whose Riemannian Hessian $\nabla^2 R(x)$ is positive definite on $T_x\mathcal{M}$ and whose local norm is*

$$\|u\|_{R,x} = \sqrt{\nabla^2 R(x)[u, u]}.$$

The barrier is ν-self-concordant in the sense of Section 3.

We also use the dual local norm

$$\|\zeta\|^*_{R,x} = \sqrt{(\nabla^2 R(x))^{-1}[\zeta, \zeta]}, \qquad \zeta \in T^*_x\mathcal{M},$$

identifying covectors and tangent vectors through the Riemannian metric when no ambiguity arises.

Assumption 4 (Interior comparator path)**.** *Let $x_0 \in \mathcal{K}^\circ$ be an analytic center of R. For every $u \in \mathcal{K}$ and every $\epsilon \in (0, 1)$, the point $u_\epsilon = \gamma_{x_0,u}(1 - \epsilon)$ is well defined, belongs to $\mathcal{K}^\circ$, and satisfies*

$$R(u_\epsilon) - R(x_0) \le c_{\mathrm{sh}} \nu \log(1/\epsilon).$$

These conditions identify where curvature, injectivity radius, loss smoothness, and boundary behavior enter the regret argument. Appendix A.1 gives conditional sufficient conditions for Hadamard domains and SPD spectral windows. The numerical comparisons in Section 6 do not by themselves establish the theorem's horizon-uniform interiority and regularity hypotheses.

3 Self-Concordant Geometry

Self-concordant barriers are used here as a local geometry for exploration and regularization, following the interior-point perspective of Nesterov & Nemirovski (1994) and the Riemannian Hessian-barrier tools developed in Hirai (2022). For a smooth curve γ and tangent vectors transported along γ, write $\nabla^3 R_x[u, v, w]$ for the covariant derivative of the Hessian.

Definition 1 (Riemannian self-concordant barrier)**.** *A function $R : \mathcal{K}^\circ \to \mathbb{R}$ is a ν-self-concordant barrier if $\nabla^2 R(x)$ is positive definite for all $x \in \mathcal{K}^\circ$, $R(x) \to \infty$ as x approaches $\partial\mathcal{K}$, and*

$$|\nabla^3 R_x[u, u, u]| \le 2\big(\nabla^2 R(x)[u, u]\big)^{3/2}$$

for all $x \in \mathcal{K}^\circ$ and $u \in T_x\mathcal{M}$, together with the standard barrier-gradient bound

$$|\nabla R(x)[u]| \le \sqrt{\nu}\, \|u\|_{R,x}.$$

The induced Dikin ball at x and radius r is

$$\mathcal{D}_R(x, r) = \{\mathrm{Exp}_x(u) : \|u\|_{R,x} \le r\}.$$

The algorithm uses perturbations inside $\mathcal{D}_R(x_t, \delta)$. Thus the relevant feasibility condition is not a global Euclidean shrinkage but a local statement: for δ below a geometry-dependent threshold, $\mathcal{D}_R(x_t, \delta) \subset \mathcal{K}$ and Exp_{x_t} is injective on the sampled tangent ball.

Lemma 1 (Compact-region Hessian comparison)**.** *Suppose R satisfies Definition 1, and let $\mathcal{X} \subset \mathcal{K}^\circ$ be a fixed compact region. There are $0 < m_R \le M_R < \infty$ such that, for $x \in \mathcal{X}$,*

$$m_R \|u\|_x^2 \le \|u\|^2_{R,x} \le M_R \|u\|_x^2, \qquad \frac{\|\zeta\|_x^2}{M_R} \le \|\zeta\|^{*2}_{R,x} \le \frac{\|\zeta\|_x^2}{m_R}.$$

Proof. The continuous positive-definite Hessian has positive minimum and finite maximum eigenvalues on $\mathcal{X}$. Inverting these eigenvalue bounds gives the dual-norm comparison. □

Lemma 1 supplies the norm comparison used by the proximal regret proof. No transported Euclidean Hessian-stability formula or convexity of logarithmic surrogates is needed.

Algorithm 1 Riemannian proximal updates with self-concordant Dikin exploration

Require: domain $\mathcal{K}$, barrier R, horizon T, parameters δ, η, initial point $x_1 \in \mathcal{K}^\circ$

for $t = 1, \dots, T$ **do**
  $A_t \leftarrow (\nabla^2 R(x_t))^{-1/2}$
  sample ξ_t uniformly from $\{u \in T_{x_t}\mathcal{M} : \|u\|_{x_t} = 1\}$
  $y_t \leftarrow \mathrm{Exp}_{x_t}(\delta A_t \xi_t)$
  observe $f_t(y_t)$
  $g_t \leftarrow \frac{d}{\delta} f_t(y_t) A_t^{-1} \xi_t \in T_{x_t}\mathcal{M}$
  define the scalar surrogate $h_t(x) = \langle g_t, \mathrm{Log}_{x_t}(x) \rangle_{x_t}$
  $x_{t+1} \in \arg\min_{x \in \mathcal{K}^\circ} \left\{ \eta h_t(x) + \frac{1}{2} d_{\mathcal{M}}(x_t, x)^2 + \frac{\eta}{T} R(x) \right\}$
end for

4 Algorithm and One-Point Regret

The one-point algorithm keeps a central point x_t and queries a randomized Dikin perturbation. Let $A_t = (\nabla^2 R(x_t))^{-1/2} : T_{x_t}\mathcal{M} \to T_{x_t}\mathcal{M}$, interpreted with respect to the Riemannian metric at x_t.

The center update uses a single scalar surrogate, squared Riemannian distance, and a barrier term of weight η/T. It never sums tangent vectors from different base points or requires logarithmic surrogates to be geodesically convex. Barrier blow-up and the quadratic distance term give an interior minimizer; remaining in a fixed compact interior region is still a separate local premise of the regret theorem. Exploration is specified by the Dikin query and estimator in Algorithm 1.

For each loss define the Dikin-smoothed function

$$\hat{f}_t(x) = \mathbb{E}_v \left[f_t\big(\mathrm{Exp}_x(\delta (\nabla^2 R(x))^{-1/2} v)\big) \right],$$

where v is uniform on the unit tangent ball at x. The algorithm queries a unit-sphere direction; this is the surface estimator associated with the fixed-center tangent-ball average.

This distinction matters. If $\hat{f}_t$ were defined as a sphere average, the one-point estimator would not be the usual gradient estimator even in Euclidean space. The tangent-space divergence identity remains exact at a fixed center, but it does not identify the estimator with the gradient of the moving-center function $\hat{f}_t$. The next definition isolates the required support comparison instead of assuming that identification.

Definition 2 (Local smoothing error)**.** *Fix the compact interior region reached by the algorithm. For a smoothing radius δ, define $\varepsilon_{\mathrm{sm}}(\delta)$ as the smallest nonnegative number such that for every round t, every feasible center x_t in this region, and every interior comparator z in the same normal neighborhood with feasible smoothing queries,*

$$\mathbb{E}_t \left[\left\langle \frac{d}{\delta} f_t(y_t) A_t^{-1} \xi_t, -\mathrm{Log}_{x_t}(z) \right\rangle_{x_t} \right] \geq \hat{f}_t(x_t) - \hat{f}_t(z) - \varepsilon_{\mathrm{sm}}(\delta).$$

Proposition 1 (Local estimator targets the smoothed surrogate)**.** *Under assumption 1, assumption 2, and assumption 3, suppose x_t lies in the compact interior region and $0 < \delta \leq \bar{\delta}$. For every interior comparator z in the same normal neighborhood with feasible smoothing queries,*

$$\mathbb{E}_t[h_t(x_t) - h_t(z)] \geq \hat{f}_t(x_t) - \hat{f}_t(z) - \varepsilon_{\mathrm{sm}}(\delta),$$

where $\mathbb{E}_t$ conditions on the pre-query history defined in Section 2. Moreover

$$\mathbb{E}_t \|g_t\|_{R,x_t}^{*2} \leq c_g \left(\frac{dC}{\delta} \right)^2.$$

Lemma 2 (Flat ellipsoidal smoothing has zero local error)**.** *If $\mathcal{M} = \mathbb{R}^d$, $\mathcal{K}$ is a Euclidean convex body, $\mathrm{Exp}_x(v) = x + v$, and A_t is fixed with respect to x over the smoothing region, then $\varepsilon_{\mathrm{sm}}(\delta) = 0$ for the ellipsoidal ball-smoothing estimator.*

Proposition 2 (Scalar-surrogate proximal stability)**.** *Let $h_t(x) = \langle g_t, \mathrm{Log}_{x_t}(x) \rangle_{x_t}$ be the scalar surrogate used in Algorithm 1. Suppose the centers, including update endpoints, remain in a fixed*

compact interior region, and the uniform smooth normal-tube bounds hold along update and comparator geodesics. Then, pathwise for every admissible comparator z,

$$\sum_{t=1}^{T}\big(h_t(x_t) - h_t(z)\big) \le c_f\eta \sum_{t=1}^{T} \|g_t\|_{R,x_t}^{*2} + \frac{d_{\mathcal{M}}(x_1, z)^2}{2\eta} + c_0(1 + \eta/T).$$

Here c_f, c_0 depend only on the fixed local geometric and barrier data, not on t, T, η, or δ. The proof telescopes squared distance; it does not invoke convexity of h_t or cumulative-surrogate FTRL stability.

Theorem 1 (One-point regret)**.** *Under assumption 1, assumption 2, assumption 3, and assumption 4, suppose the proximal centers remain in a fixed compact interior sublevel set of R. Assume the uniform local hypotheses of Lemma 6 and Proposition 2, including the comparator shrinkage paths, hold with constants independent of T. Let $0 < \delta \le \bar{\delta}$ ensure that the Dikin queries are feasible and inside the normal tube, and choose $\eta > 0$ so that the update segments lie in that tube. For $T \ge 2$, Algorithm 1 satisfies*

$$\begin{aligned}\mathrm{Reg}_T \le (c_1 H + C_{\mathrm{sm}})T\delta^2 + c_2\eta T\left(\frac{dC}{\delta}\right)^2 + \frac{c_3}{\eta} + c_4 L\sqrt{T} \\ + c_5\frac{dC}{\delta}\sqrt{T\log(2T)} + c_6(1 + \eta/T),\end{aligned}$$

where C_{sm} is given in Equation (1) *and the c_i depend only on the fixed local geometric and barrier data and the dimension. The c_5 term accounts for an adaptive nonanticipating adversary; it is absent for a fixed comparator, or for an oblivious loss sequence.*

Corollary 1 (Second-order local smoothing rate)**.** *In the setting of Theorem 1, define*

$$C_{\mathrm{sm}} := \frac{d}{(d+2)m_R}\left(H\beta_1 + \frac{L\beta_2}{2}\right), \tag{1}$$

where m_R, β_1, β_2 are the uniform constants in Lemma 6. Then $\varepsilon_{\mathrm{sm}}(\delta) \le C_{\mathrm{sm}}\delta^2$. If the admissible Dikin radius and update-tube condition of Theorem 1 allow

$$\delta \asymp T^{-1/6}, \qquad \eta \asymp \frac{\delta}{dC\sqrt{T}},$$

then Algorithm 1 satisfies $\mathrm{Reg}_T = \widetilde{O}(T^{2/3})$ for fixed dimension, loss bounds, and local geometric constants.

Proof. Lemma 6 gives the stated C^2 support coefficient. With the displayed learning rate, the variance and initial distance terms in Theorem 1 are $O(dC\sqrt{T}/\delta)$. The adaptive-comparator term is $O(dC\sqrt{T\log(2T)}/\delta)$. Taking $\delta \asymp T^{-1/6}$ makes these terms $\widetilde{O}(T^{2/3})$ and the smoothing term $O(T^{2/3})$; the remaining terms are $O(\sqrt{T} + 1)$. No loss third derivative is used. □

Proof sketch. Shrink each boundary comparator toward the analytic center, at cumulative cost $O(L\sqrt{T})$. The proof of Lemma 6 also gives the direct support bound with $f_t(z)$ in place of $\hat{f}_t(z)$. Together with the query-center bias, it yields

$$\mathbb{E}_t f_t(y_t) - f_t(z) \le \langle \mathbb{E}_t g_t, -\mathrm{Log}_{x_t} z\rangle + (c_1 H + C_{\mathrm{sm}})\delta^2.$$

The proximal first-order condition and a squared-distance Taylor bound give Proposition 2. For adaptive losses, a uniform martingale bound over a comparator net permits comparison with the random hindsight optimum, adding $O((dC/\delta)\sqrt{T\log(2T)})$. This proves the result without smoothing at a boundary-approaching comparator or assuming convexity of the moving-kernel smoothed loss. The complete argument is in Appendix A. □

Table 1 records the separate roles of exploration, support, and center-update stability. The smoothed-comparator error $\varepsilon_{\mathrm{sm}}$ is zero in the flat fixed-ellipsoid case by Lemma 2; the sufficient raw-comparator coefficient C_{sm} used by the theorem need not vanish there. The final row is unnecessary for a fixed comparator or oblivious losses.

Table 1: Proof obligations in the one-point theorem. Each term in Theorem 1 corresponds to a separately checkable step.

Term in the bound	Source of the term	Main proof device
$HT\delta^2$	Replacing raw losses by local Dikin smoothing	Smoothness in normal coordinates
$C_{\mathrm{sm}}T\delta^2$	Estimator-to-loss support error	Symmetric geometric cancellation in Lemma 6
$\eta T(dC/\delta)^2$	Bandit estimator variance in barrier dual norm	Moment bound for g_t
$1/\eta + 1 + \eta/T$	Initial distance and barrier correction	Proximal first-order condition and distance telescoping
$L\sqrt{T}$	Moving a boundary comparator to the interior	Geodesic shrinkage from the analytic center
$(dC/\delta)\sqrt{T\log(2T)}$	Adaptive hindsight comparator	Uniform martingale bound on a finite net

Remark 1. *The theorem is stated with explicit locality requirements because global logarithm maps and global geodesic convexity cannot be assumed on arbitrary manifolds. Interior existence of each proximal update does not prove horizon-uniform compactness of its trajectory. Likewise, comparator paths approaching the boundary require uniform geometric normal-tube bounds; no uniform bound on the barrier at those comparators is asserted. These requirements must be verified for the application under consideration.*

5 Two-Point Extension

The two-point feedback model allows the learner to query two nearby points before receiving their loss values. The natural Riemannian construction samples a tangent direction at x_t and queries the symmetric pair

$$y_t^+ = \mathrm{Exp}_{x_t}(\delta A_t \xi_t), \qquad y_t^- = \mathrm{Exp}_{x_t}(-\delta A_t \xi_t).$$

This construction requires both points to remain feasible and inside the injectivity region. On symmetric manifolds, the second point can equivalently be viewed through the geodesic reflection map $y_t^- = -_{x_t} y_t^+$.

The estimator becomes

$$g_t^{(2)} = \frac{d}{2\delta}\left(f_t(y_t^+) - f_t(y_t^-)\right) A_t^{-1}\xi_t,$$

and the proximal center update in Algorithm 1 uses $h_t^{(2)}(x) = \langle g_t^{(2)}, \mathrm{Log}_{x_t}(x)\rangle_{x_t}$ in place of the one-point surrogate. We charge the average of the two queried losses; charging either member has the same conditional expectation.

Theorem 2 (Two-point regret)**.** *Assume assumption 1, assumption 2, and assumption 3, the fixed compact-interior and normal-tube conditions of Theorem 1, and feasible symmetric query pairs. Suppose the uniform local hypotheses of Lemma 6 hold for the centers, queries, and admissible comparators. Let $\|g_t^{(2)}\|_{R,x_t}^* \le G_2$ almost surely, with G_2 independent of δ, as verified in Lemma 12. With a constant step size η, the two-point proximal variant satisfies*

$$\mathrm{Reg}_T^{(2)} \le (c_1' H + C_{\mathrm{sm}})T\delta^2 + c_2'\eta T G_2^2 + \frac{c_3'}{\eta} + c_4' L\sqrt{T} + c_5' G_2\sqrt{T\log(2T)} + c_6'(1+\eta/T),$$

where the constants depend on the fixed local data, not on T, η, δ. The boundary-comparator term uses assumption 4 and the stated Lipschitz bound; it is absent for comparators already in the admissible interior region. The additional square-root term allows a nonanticipating adaptive loss sequence and a hindsight comparator; it is absent for a fixed comparator. The support-error coefficient is unchanged because antipodal sampling gives $\mathbb{E}_t g_t^{(2)} = G_\delta f_t(x_t)$, as detailed in Appendix A.5. Balancing the two stability terms with $\eta \asymp 1/(G_2\sqrt{T})$ and taking $\delta = O(T^{-1/4})$ within the feasibility threshold yields $\widetilde{O}(\sqrt{T})$ regret up to the stated smoothing and geometry terms.

The statement separates a feedback improvement from a geometric assumption. The $\sqrt{T}$ scaling comes from central differencing and lower estimator variance; the manifold assumptions ensure that the symmetric pair is well defined and that the local Taylor comparison is legitimate. This avoids presenting the two-point rate as a consequence of self-concordance alone. The proof is included in Appendix A.5. In compact locally Lipschitz regions, Lemma 12 verifies the finite G_2 condition, so the estimator-moment requirement is a checkable local regularity statement rather than a separate algorithmic assumption.

6 NUMERICAL EXPERIMENTS

We evaluate the one- and two-point proximal methods on covariance streams over the affine-invariant SPD manifold, with $f_t(X) = \frac{1}{2} d_{\mathrm{AI}}(X, \Sigma_t)^2$. These are strongly geodesically convex benchmark losses. Each comparison shares the target stream and feedback budget across methods. Let $L^a_{s,T}$ be method a's mean charged query loss on paired seed s: one queried value per round for one-point feedback, and the average of two for two-point feedback. We report $r_{s,T} = 100(L^{\mathrm{base}}_{s,T} - L^{\mathrm{prox}}_{s,T})/L^{\mathrm{base}}_{s,T}$ averaged over seeds, with paired 95% confidence intervals (CIs); positive values favor proximal. This differs from the percentage reduction of aggregate loss means. Protocols and complete aggregate results appear in Appendices B and C.

6.1 ONE-POINT COVARIANCE-STREAM TRANSFER

Setup and frozen transfer. On the breast-cancer covariance stream, proximal Dikin, isotropic proximal, and projected Euclidean each receive one exact function value per round, without gradient access. Each method is selected from a 64-candidate schedule grid at $T = 200$, using disjoint screening and validation seeds; its parameters are then frozen for evaluation at larger horizons. The selected Dikin schedule is $\eta_t = 2 \times 10^{-7}/\sqrt{t}$, with radius scale 1 and barrier factors $(\lambda, \gamma) = (20, 1000)$. Isotropic proximal and projection use separately selected step scales 2×10^{-5} and 2×10^{-4}, with radius scales 0.5 and 0.125, respectively. This tests short-horizon parameter transfer without using the longer-horizon evaluation losses for selection.

Table 2 reports absolute losses, and Figure 1 displays Dikin's paired effects. The frozen configuration gives a 2.93% reduction at $T = 500$ ($[-0.72, 6.59]\%$; 24/40 wins), whose CI includes zero. At $T = 1000$, the reduction is 5.08% ($[3.08, 7.07]\%$; 74/100 wins). The latter pools 100 distinct seeds held out from selection; the 40 $T = 500$ seeds are a subset, so the two horizons are not independent tests.

Table 2: One-point breast-cancer covariance tracking: mean queried loss per round (lower is better). Each row is a paired evaluation under the indicated selection protocol; retuning rows are separate tests, not pooled with frozen transfer. n is the number of paired seeds.

Selection protocol	T	n	Dikin	Isotropic	Projected
Frozen transfer	500	40	7.930	7.967	8.302
Frozen transfer	1000	100	7.966	8.037	8.506
Barrier-only retuning	1000	60	7.851	8.013	8.260
Step/radius retuning	1000	60	7.851	8.103	8.461

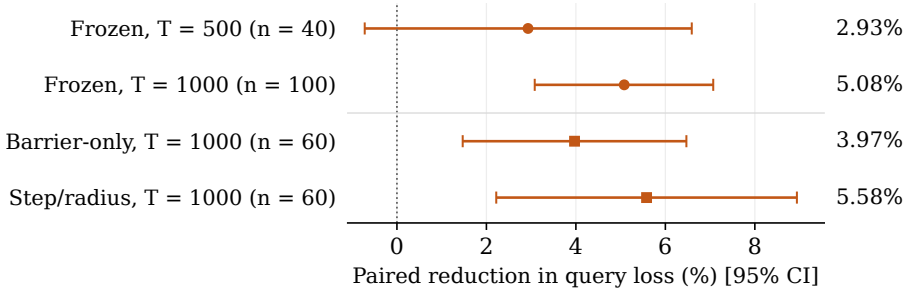


Figure 1: One-point Dikin versus projected Euclidean: mean paired percentage loss reductions and reported 95% CIs. The vertical line is zero improvement. Rows distinguish frozen transfer from two separate $T = 1000$ retuning protocols.

Separate long-horizon retuning. Selecting only Dikin's barrier scale at $T = 1000$ gives a 3.97% reduction versus projection ($[1.47, 6.47]\%$; 44/60 wins). A subsequent step/radius selection, evalu-

ated on another 60 held-out seeds, gives 5.58% ($[2.22, 8.94]\%$; 44/60 wins) versus projection and 2.70% ($[1.08, 4.32]\%$; 41/60 wins) versus isotropic proximal. Both controls retain their earlier settings. Thus the latter is an advantage of the *retuned Dikin configuration*, not a query-map-only ablation: Dikin and isotropic use different step, radius, and barrier scales.

Predictable-baseline check. A separate 60-seed $T = 1000$ evaluation compares the step/radius-retuned Dikin settings, frozen for this check, with no scalar baseline and with an EMA-0.5 baseline. Their reductions versus projection are 5.02% and 5.01%, respectively, with CIs excluding zero. The direct EMA-versus-no-baseline effect is 0.004% ($[-0.047, 0.054]\%$), with no detectable difference in this test. This protocol uses a different projected step scale and is reported separately in Appendix B.4. Overall, proximal Dikin performs well on this covariance stream, with lower mean query loss at $T = 1000$ across the frozen-transfer, retuning, and scalar-baseline comparisons. For scale, the frozen $T = 1000$ means, 7.966 versus 8.506, correspond to approximately 540 less cumulative charged loss. With a common static comparator this is also the mean regret advantage, by (2).

6.2 Two-Point Covariance Tracking

Tasks and protocol. We next give each method two exact function values per round. Real-data tasks use PCA-6 digits covariances and two chronological six-stock blocks; the comparator is gradient-free projected Euclidean. Digits selection and evaluation use disjoint query seeds, and both methods' settings transfer unchanged to stocks. Controlled SPD(3) cycles and SPD(5) tasks instead use intrinsic two-point Riemannian OGD (ROGD), with independent calibration for each method. Central and boundary-stress SPD(3) comparisons use equal 48-candidate budgets and 30 fresh test seeds. SPD(5) is calibrated separately through $T = 300$; longer horizons transfer both methods' frozen $T = 300$ settings.

Table 3: Two-point real-data charged loss (mean $\pm$ across-seed SD). Digits uses 20 paired seeds per horizon; each stock block uses five. Stock settings transfer unchanged from digits. Complete paired effects are in Table 11.

Stream	T	Proximal Dikin	Isotropic proximal	Projected Euclidean
Digits	200	12.242 ± 0.360	12.243 ± 0.358	15.206 ± 3.010
Digits	500	9.316 ± 0.368	9.317 ± 0.368	13.716 ± 3.096
Digits	1000	7.174 ± 0.303	7.176 ± 0.303	12.830 ± 3.258
Stock 1	200	5.641 ± 0.362	5.641 ± 0.362	10.707 ± 1.148
Stock 2	200	8.275 ± 0.550	8.275 ± 0.551	14.125 ± 0.995

Results. Dikin has lower mean loss than projection in all five real-data cells (Table 3). Digits paired reductions increase from 16.37% at $T = 200$ to 40.55% at $T = 1000$ (CIs $[8.44, 24.31]\%$ and $[33.46, 47.63]\%$; 17/20 and 20/20 wins). Stock reductions are 46.91% and 41.31%, with 5/5 wins in each block. The matched isotropic-query arm is nearly identical to Dikin. On central SPD(3), Dikin improves over ROGD by 12.48% ($[5.32, 19.63]\%$; 21/30 wins); boundary stress selects *isotropic* proximal, giving 15.36% ($[9.87, 20.86]\%$; 25/30 wins). SPD(5) reductions range from 5.12% at $T = 100$ to 35.97% at $T = 300$, then 44.98% and 44.43% under frozen transfer to $T = 500, 1000$ (20/20 wins each). Selected absolute synthetic losses are in Table 4; complete effects and the four-panel overview are in Appendix C. These improvements are configuration-dependent: isotropic results do not isolate a Dikin-specific gain, and the separate inverse-time configuration check favors projection (Appendix C.5).

Table 4: Selected two-point synthetic tests at $T = 300$, each with 30 fresh paired seeds. Losses are charged per-round means; reductions and 95% CIs are paired versus ROGD. The stress arm is isotropic, not Dikin.

Regime	Proximal query	Proximal	ROGD	Reduction [95% CI]
SPD(3), central	Dikin	0.00549	0.00644	12.48% $[5.32, 19.63]$
SPD(3), stress	Isotropic	0.02064	0.02503	15.36% $[9.87, 20.86]$
SPD(5)	Dikin	0.01691	0.02669	35.97% $[33.00, 38.95]$

Relation to regret. For either feedback budget, the common best-fixed-comparator loss cancels:

$$R_{s,T}^{\mathrm{base}} - R_{s,T}^{\mathrm{prox}} = T\big(L_{s,T}^{\mathrm{base}} - L_{s,T}^{\mathrm{prox}}\big). \tag{2}$$

Thus lower charged loss means lower regret on a shared sequence and comparator domain: the approximate mean two-point gaps are 5656 on digits at $T = 1000$ and 0.285 on central SPD(3). Percentage effects and CIs concern loss, not regret percentages.

7 Limitations and Conclusion

Our guarantees require compact-interior centers, feasible Dikin query balls, and uniform smooth normal-tube data for center–comparator pairs. The barrier ensures interior updates, not horizon-uniform compact containment. Geometric constants may deteriorate near singularities; uniformity must also hold as shrunken comparators approach the boundary.

The one- and two-point SPD experiments demonstrate effective parameter transfer and query-loss improvements over feedback-matched baselines, with corresponding regret advantages for a common comparator class. Isotropic comparisons do not isolate a Dikin-specific gain, and the two-point inverse-time check favors projection. All benchmark losses are strongly geodesically convex; merely-convex tasks remain untested. Numerical runs certify neither all theorem premises nor exact global solution of the proximal subproblems.

We combine barrier-local Dikin exploration with intrinsic single-step proximal updates. A direct C^2 support bound and squared-distance telescoping yield conditional $\widetilde{O}(T^{2/3})$ one-point regret without assuming convexity of logarithmic surrogates. The two-point extension and SPD experiments complement this result.

References

Jacob D. Abernethy, Elad Hazan, and Alexander Rakhlin. Interior-point methods for full-information and bandit online learning. *IEEE Transactions on Information Theory*, 58(7):4164–4175, 2012. doi: 10.1109/TIT.2012.2192096.

P.-A. Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 2008.

Alekh Agarwal, Ofer Dekel, and Lin Xiao. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *Proceedings of the 23rd Annual Conference on Learning Theory*, pp. 28–40, 2010.

Rajendra Bhatia. *Positive Definite Matrices*. Princeton University Press, 2007.

Nicolas Boumal. *An Introduction to Optimization on Smooth Manifolds*. Cambridge University Press, 2023.

Sébastien Bubeck, Ronen Eldan, and Yin Tat Lee. Kernel-based methods for bandit convex optimization. *Journal of the ACM*, 68(4):25:1–25:35, 2021. doi: 10.1145/3453721.

Varsha Dani, Thomas P. Hayes, and Sham M. Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory*, pp. 355–366, 2008.

Ofer Dekel, Ronen Eldan, and Tomer Koren. Bandit smooth convex optimization: Improving the bias-variance tradeoff. In *Advances in Neural Information Processing Systems*, volume 28, 2015.

John C. Duchi, Michael I. Jordan, Martin J. Wainwright, and Andre Wibisono. Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806, 2015. doi: 10.1109/TIT.2015.2409256.

Alan Edelman, Tomás A. Arias, and Steven T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998. doi: 10.1137/S0895479895290954.

Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 385–394. SIAM, 2005.

Hidde Fokkema, Dirk van der Hoeven, Tor Lattimore, and Jack J. Mayo. Online newton method for bandit convex optimisation, 2024. arXiv:2406.06506.

Octavian-Eugen Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. In *Advances in Neural Information Processing Systems*, volume 31, 2018.

Hiroshi Hirai. On a manifold formulation of self-concordant functions, 2022. arXiv:2212.10981.

Hiroshi Hirai, Harold Nieuwboer, and Michael Walter. Interior-point methods on manifolds: theory and applications, 2023. arXiv:2303.04771.

Zihao Hu, Guanghui Wang, and Jacob D. Abernethy. Minimizing dynamic regret on geodesic metric spaces. In *Proceedings of Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 4336–4383. PMLR, 2023a.

Zihao Hu, Guanghui Wang, and Jacob D. Abernethy. Riemannian projection-free online learning. In *Advances in Neural Information Processing Systems*, volume 36, 2023b.

Alexandru Kristály, Chong Li, Genaro López-Acedo, and Adriana Nicolae. What do 'convexities' imply on hadamard manifolds? *Journal of Optimization Theory and Applications*, 170(3):1068–1074, 2016. doi: 10.1007/s10957-015-0780-2.

Tor Lattimore. Bandit convex optimisation, 2024. arXiv:2402.06535.

Jiaxiang Li, Krishnakumar Balasubramanian, and Shiqian Ma. Stochastic zeroth-order riemannian derivative estimation and optimization. *Mathematics of Operations Research*, 48(2):1183–1211, 2023. doi: 10.1287/moor.2022.1302.

Yurii Nesterov and Arkadii Nemirovski. *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM, 1994.

Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006. doi: 10.1007/s11263-005-3222-z.

Ankan Saha and Ambuj Tewari. Improved regret guarantees for online smooth convex optimization with bandit feedback. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 636–642. PMLR, 2011.

Ohad Shamir. An optimal algorithm for bandit and zero-order convex optimization with two-point feedback. *Journal of Machine Learning Research*, 18(52):1–11, 2017.

Xi Wang, Zhipeng Tu, Yiguang Hong, Yingyi Wu, and Guodong Shi. Online optimization over riemannian manifolds. *Journal of Machine Learning Research*, 24(84):1–67, 2023.

Hongyi Zhang and Suvrit Sra. First-order methods for geodesically convex optimization. In *Proceedings of the 29th Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pp. 1617–1638. PMLR, 2016.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning*, pp. 928–936, 2003.

A PROOFS

We give the formal proof at the level of the constants stated in the main text. All constants below are uniform over the compact interior region reached by the algorithm. They may depend on the injectivity radius, curvature bounds, the chosen compact sublevel set of the barrier, and the dimension, but not on T.

A.1 VERIFYING ASSUMPTIONS

This subsection records sufficient conditions that turn the assumptions in Section 2 into checkable geometric statements. The point is not to remove locality from the theorem. Rather, the propositions below separate primitive geometry from the proof obligations used later in the regret argument.

Proposition 3 (Hadamard compact-region verification)**.** *Suppose $\mathcal{M}$ is Hadamard and $\mathcal{X} \subset \mathcal{K}^\circ$ is a compact geodesically convex region containing all iterates, feasible queries, and interior comparators. If the losses are geodesically convex, twice differentiable, and have uniformly bounded value, gradient, and Hessian on $\mathcal{X}$, and if R is a ν-self-concordant barrier on $\mathcal{K}^\circ$, then the geometric, loss, and barrier conditions used by Theorem 1 hold on $\mathcal{X}$ with finite constants.*

Proof. On a Hadamard manifold, the exponential map is globally nonsingular and the logarithm map is single valued. Compactness of $\mathcal{X}$ gives finite curvature bounds on the region, a positive distance from $\mathcal{X}$ to $\partial\mathcal{K}$, and finite upper and lower bounds for the barrier Hessian eigenvalues. The assumed bounds on the losses give the constants C, L, H after enlarging them to cover the compact query tube. The self-concordant-barrier condition is exactly assumption 3. These observations verify the local hypotheses invoked in Lemmas 3 to 5. □

Proposition 4 (SPD spectral windows verify the application geometry)**.** *Let $\mathcal{M} = \mathbb{S}^m_{++}$ carry the affine-invariant metric and fix $0 < \alpha < \beta < \infty$. On the spectral window*

$$\mathcal{K}_{\alpha,\beta} = \{X : \alpha I \preceq X \preceq \beta I\},$$

the logarithm map is single valued, sectional curvature and affine-invariant norm comparisons are uniformly bounded, and the squared-distance losses

$$X \mapsto \frac{1}{2} d_{\mathrm{AI}}(X, \Sigma)^2, \qquad \Sigma \in \mathcal{K}_{\alpha,\beta},$$

are geodesically convex and locally smooth with constants depending only on (α, β, m). On every strict subwindow $\mathcal{K}_{\alpha',\beta'}$ with $\alpha < \alpha' < \beta' < \beta$, any smooth spectral-window barrier whose Riemannian Hessian is positive definite has finite local Hessian-comparison and derivative constants.

Proof. The affine-invariant SPD manifold is a Hadamard manifold (Bhatia, 2007), so geodesics between positive definite matrices are unique and squared distance to a fixed point is geodesically convex. The spectral window bounds eigenvalues away from zero and infinity; consequently the affine-invariant metric, tangent norms, and curvature-dependent constants are uniformly controlled on the window. Smoothness of the squared-distance loss on a compact subset follows from smooth dependence of the affine-invariant logarithm on its arguments away from the boundary. A strict spectral subwindow is compactly contained in the domain of any spectral-window barrier; positive definiteness and smoothness of its Riemannian Hessian then give finite upper and lower Hessian-comparison constants and finite derivative bounds by compactness. These are the smoothness and local-norm comparison constants used in the regret analysis. □

Proposition 5 (Concrete local SPD spectral-window barrier). *Fix $0 < \alpha < \alpha' < \beta' < \beta < \infty$ and let the reached iterates and queries stay in the strict spectral window $\mathcal{K}_{\alpha',\beta'}$. Define*

$$W_{\alpha,\beta}(X) = -\log\det(X - \alpha I) - \log\det(\beta I - X), \qquad X_c = \frac{\alpha+\beta}{2} I.$$

There exist finite constants $\lambda > 0$ and $\gamma > 0$ such that

$$R_{\lambda,\gamma}(X) = \gamma\left\{W_{\alpha,\beta}(X) + \frac{\lambda}{2} d_{\mathrm{AI}}(X, X_c)^2\right\}$$

is a local ν-self-concordant barrier on $\mathcal{K}_{\alpha',\beta'}$ in the sense needed by the Dikin analysis. In particular, on $\mathcal{K}_{\alpha',\beta'}$ its Riemannian Hessian is positive definite and uniformly comparable to the affine-invariant metric:

$$m_R \|U\|_{\mathrm{AI,X}}^2 \le \nabla^2 R_{\lambda,\gamma}(X)[U,U] \le M_R \|U\|_{\mathrm{AI,X}}^2$$

for finite constants $0 < m_R \le M_R < \infty$ depending only on $(\alpha, \alpha', \beta', \beta, m, \lambda, \gamma)$.

Proof. The potential $W_{\alpha,\beta}$ is smooth on the open window $\{X : \alpha I \prec X \prec \beta I\}$ and diverges when any eigenvalue approaches either boundary α or β. On the compact strict window $\mathcal{K}_{\alpha',\beta'}$, all covariant derivatives of $W_{\alpha,\beta}$ are bounded and its Riemannian Hessian has a finite lower bound: there is $B < \infty$ such that

$$\nabla^2 W_{\alpha,\beta}(X)[U,U] \ge -B\|U\|_{\mathrm{AI,X}}^2 \qquad \text{for all } X \in \mathcal{K}_{\alpha',\beta'}.$$

Because the affine-invariant SPD manifold is Hadamard, the squared-distance function $X \mapsto \frac{1}{2} d_{\mathrm{AI}}(X, X_c)^2$ is strongly geodesically convex. Hence its Hessian dominates the affine-invariant metric on the strict window. Choosing $\lambda > B + 1$ gives a smooth potential

$$\Phi_\lambda(X) = W_{\alpha,\beta}(X) + \frac{\lambda}{2} d_{\mathrm{AI}}(X, X_c)^2$$

whose Riemannian Hessian is uniformly positive definite. Compactness also gives a finite upper Hessian bound, so $\nabla^2\Phi_\lambda$ is comparable to the affine-invariant metric.

It remains to check the self-concordance constants. On $\mathcal{K}_{\alpha',\beta'}$, compactness gives constants $m_\lambda > 0$, $M_1 < \infty$, and $M_3 < \infty$ such that

$$\nabla^2\Phi_\lambda(X)[U,U] \ge m_\lambda \|U\|_{\mathrm{AI,X}}^2, \qquad |\nabla\Phi_\lambda(X)[U]| \le M_1\|U\|_{\mathrm{AI,X}},$$

and

$$|\nabla^3\Phi_\lambda(X)[U,U,U]| \le M_3\|U\|_{\mathrm{AI,X}}^3.$$

For $R_{\lambda,\gamma} = \gamma\Phi_\lambda$, the Hessian and third derivative both scale by γ. Taking

$$\gamma \ge \left(\frac{M_3}{2m_\lambda^{3/2}}\right)^2$$

implies

$$|\nabla^3 R_{\lambda,\gamma}(X)[U,U,U]| \le 2\big(\nabla^2 R_{\lambda,\gamma}(X)[U,U]\big)^{3/2}.$$

The barrier-gradient inequality holds with any finite

$$\nu \ge \gamma M_1^2/m_\lambda.$$

Finally, the Hessian-comparison constants for $R_{\lambda,\gamma}$ are $m_R = \gamma m_\lambda$ and a corresponding compact upper bound M_R. Thus this concrete spectral-window regularizer supplies the local barrier constants used in the regret analysis. The statement is conditional on a reached strict subwindow and on choosing the barrier parameters large enough for the displayed bounds. The finite-horizon experimental results do not independently certify these uniform conditions. □

A.2 Local Geometric Lemmas

Fix a compact interior region $\mathcal{X} \subset \mathcal{K}^\circ$ containing all centers, update endpoints, and feasible query points, uniformly in T. Comparator paths and the geodesics used in the analysis lie in a fixed compact geometric normal tube on which the relevant logarithms are smooth. This tube may meet $\partial\mathcal{K}$: a boundary-approaching comparator need not stay in a fixed compact subset of $\mathcal{K}^\circ$. All geometric bounds involving those comparators are required to remain uniform. Define

$$\rho_{\mathcal{X}} = \inf_{x\in\mathcal{X}} \mathrm{inj}(x), \qquad b_A = \sup_{x\in\mathcal{X}} \|(\nabla^2 R(x))^{-1/2}\|_{\mathrm{op}}, \qquad m_A = \inf_{x\in\mathcal{X}} \lambda_{\min}(\nabla^2 R(x)).$$

Let r_∂ be a positive lower bound on the distance from $\mathcal{X}$ to $\partial\mathcal{K}$ measured along geodesics in the chart used by the analysis. These constants are finite and positive by compactness and positive definiteness of the barrier Hessian. Keeping them visible is useful because the proof is local: every hidden constant below can be traced to one of $\rho_{\mathcal{X}}$, r_∂, b_A, m_A, curvature bounds, smoothness bounds, or the dimension. Write $m_R = m_A$ and let M_R be the upper Hessian bound from Lemma 1. For $V_z(x) = d_{\mathcal{M}}(x,z)^2/2$, fix uniform constants $D, L_V \geq 1, K_R, L_{\log} < \infty$ such that

$$d_{\mathcal{M}}(x,z) \leq D, \qquad |\nabla^2 V_z(w)[a,a]| \leq L_V \|a\|_w^2,$$
$$\left\|(d\,\mathrm{Exp}_x)^*_{\mathrm{Log}_x y} \nabla R(y)\right\|_x \leq K_R, \qquad \|(d_z\,\mathrm{Log}_x)_z\|_{\mathrm{op}} \leq L_{\log}.$$

Here x, y are center/update pairs, w lies on the relevant geodesics, and z ranges over the comparator tube; enlarge D also to bound $d_{\mathcal{M}}(x_0, u)$ in the shrinkage argument. Smoothness on the stated compact sets gives these constants. In particular, K_R requires a barrier bound only at centers, not at boundary-approaching comparators.

Lemma 3 (Dikin feasibility)**.** *Under assumption 1 and assumption 3, there is a radius $\bar{\delta} > 0$ such that for every iterate x_t in the compact interior region and every tangent vector u with $\|u\|_{x_t} \leq 1$, the point*

$$\mathrm{Exp}_{x_t}\Big(\delta(\nabla^2 R(x_t))^{-1/2} u\Big)$$

belongs to $\mathcal{K}^\circ$ whenever $0 < \delta \leq \bar{\delta}$.

Proof. Because the iterates remain in a compact subset of $\mathcal{K}^\circ$, the Riemannian injectivity radius and the distance to $\partial\mathcal{K}$ are bounded below by positive constants on that set. The operator norm of $(\nabla^2 R(x))^{-1/2}$ is also bounded above on the same compact set. In the notation above, it is enough to take

$$\bar{\delta} \leq \frac{1}{2}\min\{\rho_{\mathcal{X}}, r_\partial\}/b_A.$$

Then $\|\delta(\nabla^2 R(x_t))^{-1/2} u\|_{x_t} \leq \min\{\rho_{\mathcal{X}}, r_\partial\}/2$. The exponential image is inside the injectivity region and remains in the interior tube around $\mathcal{X}$. Geodesic convexity then keeps the local query segment inside $\mathcal{K}^\circ$. □

Lemma 4 (Smoothing bias)**.** *Let $\hat{f}_t$ be the Dikin-smoothed loss defined in Section 4. Under assumption 1, assumption 2, and assumption 3, at any base point x whose symmetric Dikin ball is feasible,*

$$0 \leq \hat{f}_t(x) - f_t(x) \leq \frac{H\delta^2 \,\mathrm{tr}(A_x^2)}{2(d+2)}.$$

If $\nabla^2 R(x) \succeq m_R I$, this is at most $c_s H\delta^2$, where $c_s = d/(2(d+2)m_R)$. The same pointwise statement holds at a comparator whenever its smoothing queries are feasible.

Proof. Apply the geodesic Hessian bound along $s \mapsto \mathrm{Exp}_x(sv)$:

$$f_t(\mathrm{Exp}_x(v)) \leq f_t(x) + \langle \nabla f_t(x), v\rangle_x + \frac{H}{2}\|v\|_x^2.$$

No third-order remainder is needed along a geodesic. For the tangent-ball perturbation $v = \delta A_x u$, $\mathbb{E}v = 0$ and $\mathbb{E}\|v\|^2 = \delta^2 \,\mathrm{tr}(A_x^2)/(d+2)$. Taking expectation gives the upper bound; geodesic convexity and antipodal symmetry give the lower bound. □

Lemma 5 (Query-center displacement)**.** *Under the same conditions as Lemma 4, the random one-point query satisfies*

$$\mathbb{E}_t f_t(y_t) - \hat{f}_t(x_t) \le \frac{H\delta^2 \operatorname{tr}(A_t^2)}{2d} \le c_q H \delta^2, \qquad c_q = \frac{1}{2m_R}.$$

Proof. Write $v_t = \delta A_t \xi_t$. Geodesic smoothness gives

$$f_t(\operatorname{Exp}_{x_t}(v_t)) - f_t(x_t) \le \langle \nabla f_t(x_t), v_t \rangle_{x_t} + \frac{H}{2}\|v_t\|_{x_t}^2.$$

Conditional symmetry removes the linear term, and the uniform sphere moment gives $\mathbb{E}_t\|v_t\|^2 = \delta^2 \operatorname{tr}(A_t^2)/d$. Since $\hat{f}_t(x_t) \ge f_t(x_t)$, the result follows. This argument does not assume smoothness of the moving-kernel function $\hat{f}_t$. □

Lemma 6 (Sufficient condition for second-order local smoothing error)**.** *Work under assumption 1, assumption 2, and assumption 3 on the fixed compact normal tube fixed above. All bounds below are uniform over the admissible centers, comparators, and radial query segments in this tube, which lies in the smooth domain of the logarithm map. Suppose the geometric maps have uniformly bounded derivatives through order two there, and the losses are C^2 with $0 \preceq \nabla^2 f_t \preceq H\langle\cdot,\cdot\rangle$. Write $A_x = (\nabla^2 R(x))^{-1/2}$ and let L bound $\|\nabla f_t(x)\|_x$ at the centers. This local gradient bound can be obtained from the existing bounded-loss and interior-margin data, as shown below; no third derivative of a loss is required.*

There are finite geometry-only constants β_1, β_2 on this tube such that, for every admissible x, z and every sufficiently small $\delta > 0$, the direct support inequality is

$$\hat{f}_t(x) - f_t(z) - \big\langle G_\delta f_t(x), -\operatorname{Log}_x(z)\big\rangle_x \le \frac{\delta^2 \operatorname{tr}(A_x^2)}{d+2}\left(H\beta_1 + \frac{L\beta_2}{2}\right). \tag{3}$$

No smoothing at z is required for this conclusion. If its symmetric Dikin ball is also feasible, then additionally

$$\hat{f}_t(x) - \hat{f}_t(z) - \big\langle G_\delta f_t(x), -\operatorname{Log}_x(z)\big\rangle_x \le \frac{\delta^2 \operatorname{tr}(A_x^2)}{d+2}\left(H\beta_1 + \frac{L\beta_2}{2}\right), \tag{4}$$

where

$$G_\delta f_t(x) = \mathbb{E}_{\xi\sim\operatorname{Unif}(S_x)}\left[\frac{d}{\delta} f_t(\operatorname{Exp}_x(\delta A_x \xi)) A_x^{-1}\xi\right]$$

and S_x is the unit sphere in $T_x\mathcal{M}$. In particular, if $\nabla^2 R(x) \succeq m_R I$ on the compact region, then

$$\varepsilon_{\mathrm{sm}}(\delta) \le \frac{d}{(d+2)m_R}\left(H\beta_1 + \frac{L\beta_2}{2}\right)\delta^2 = C_{\mathrm{sm}}\delta^2 = O(\delta^2), \tag{5}$$

where the supremum defining $\varepsilon_{\mathrm{sm}}$ is taken over comparators with feasible smoothing. The same $C_{\mathrm{sm}}\delta^2$ bounds the direct support error in Equation (3) *without that feasibility restriction. Here C_{sm} is the coefficient in Equation* (1)*. It includes both the Hessian and gradient contributions and is not asserted to depend on geometry alone. The constants are independent of t, T, and δ, for fixed local data. No division by H is needed; if $H = 0$, the sharper local smoothing error is zero.*

Proof. Fix a round, condition on the history before its fresh direction is sampled, and abbreviate f_t to f. We prove the support inequality directly, without assuming convexity of the moving-kernel function $\hat{f}$ or a second-order approximation of $G_\delta f(x)$ by $\nabla f(x)$.

Step 1: the geometric defect. Fix admissible x, z, and set

$$q = \operatorname{Log}_x(z), \qquad y_u = \operatorname{Exp}_x(u), \qquad E_u = (d\operatorname{Exp}_x)_u : T_x\mathcal{M} \longrightarrow T_{y_u}\mathcal{M}.$$

Let $\Pi_u : T_x\mathcal{M} \to T_{y_u}\mathcal{M}$ be parallel transport along $s \mapsto \operatorname{Exp}_x(su)$, and define

$$W_{x,z}(u) = \Pi_u^{-1}\big(\operatorname{Log}_{y_u}(z) - E_u q\big). \tag{6}$$

We first justify uniform constants satisfying

$$\|W_{x,z}(u)\|_x \le \beta_1 \|u\|_x, \qquad \|W_{x,z}(u) + W_{x,z}(-u)\|_x \le \beta_2 \|u\|_x^2. \tag{7}$$

All norms and derivatives in this step are taken in a Riemannian orthonormal frame. Put

$$V_{x,z}(u) = \Pi_u^{-1} \operatorname{Log}_{\operatorname{Exp}_x(u)}(z), \qquad \widetilde{E}_u = \Pi_u^{-1} E_u.$$

On the fixed compact normal tube, smooth radial parallel transport and the stated second-order geometric bounds give finite uniform constants M_1, M_2 bounding $\|D_u V_{x,z}\|_{\mathrm{op}}$ and $\|D_u^2 V_{x,z}\|_{\mathrm{op}}$, respectively. They are uniform over z as well as x. Let $Q = \sup_{x,z} \|\operatorname{Log}_x(z)\|_x < \infty$.

For completeness, $\widetilde{E}_u - I$ is quadratic in u, rather than merely linear. Let $\kappa = \max\{\kappa_-, \kappa_+\}$ and restrict to $\|u\|_x \le r$ with $\kappa r^2 \le 1$. The Jacobi field $J_a(s) = (d\operatorname{Exp}_x)_{su}(sa)$ has initial data $J_a(0) = 0$, $D_s J_a(0) = a$. After parallel transport back to x, its Jacobi equation reads

$$j_a(s) = sa - \int_0^s (s-\tau)\mathcal{R}_u(\tau) j_a(\tau)\, d\tau, \qquad \|\mathcal{R}_u(\tau)\|_{\mathrm{op}} \le \kappa \|u\|_x^2.$$

Consequently $\sup_{0\le s\le 1} \|j_a(s)\|_x \le 2\|a\|_x$, and substitution at $s = 1$ gives

$$\|\widetilde{E}_u - I\|_{\mathrm{op}} \le c_E \|u\|_x^2, \qquad c_E = \kappa.$$

Since $V_{x,z}(0) = q$, the first-order Taylor bound and the symmetric second-order Taylor bound for V yield Equation (7) with

$$\beta_1 = M_1 + c_E Q r, \qquad \beta_2 = M_2 + 2c_E Q.$$

Indeed, $W(u) = V(u) - q - (\widetilde{E}_u - I)q$. This argument does not differentiate E_u twice, and introduces no third-order loss data. The normal tube and its derivative bounds, not just the sectional curvature interval, enter these constants. Shrinking $\bar{\delta}$ so that $b_A \bar{\delta} \le r$ ensures the bounds apply to all perturbations.

Step 2: an exact identity at a fixed center. Let B_x be the unit ball in $T_x\mathcal{M}$, with its Euclidean volume induced by $\langle\cdot,\cdot\rangle_x$, and set $F_x(u) = f(\operatorname{Exp}_x(u))$. For any $a \in T_x\mathcal{M}$, apply the divergence theorem on B_x to the constant direction $A_x^{-1}a$ and the function $v \mapsto F_x(\delta A_x v)$. The chain rule and the sphere-to-ball volume ratio d give

$$\begin{aligned} \big\langle G_\delta f(x), a \big\rangle_x &= \frac{d}{\delta} \mathbb{E}_{\xi \sim \mathrm{Unif}(S_x)} F_x(\delta A_x \xi) \langle \xi, A_x^{-1} a \rangle_x \\ &= \mathbb{E}_{v \sim \mathrm{Unif}(B_x)} DF_x(\delta A_x v)[a]. \end{aligned} \tag{8}$$

Here A_x is self-adjoint and is held fixed. This identity concerns a derivative in the fixed tangent space, not differentiation of $\hat{f}$ with respect to its center. In particular, it requires neither a derivative of A_x nor a Riemannian volume Jacobian: the defining smoothing measure is the tangent-ball measure.

Step 3: geodesic convexity and symmetric cancellation. Take $a = -q$ and $u = \delta A_x v$ in Equation (8). Geodesic convexity on the segment from y_u to z implies

$$df_{y_u}[-\operatorname{Log}_{y_u}(z)] \ge f(y_u) - f(z).$$

Writing $p(u) = \Pi_u^{-1} \nabla f(y_u)$ and using Equation (6), we obtain the pointwise inequality

$$df_{y_u}[-E_u q] \ge f(y_u) - f(z) + \langle p(u), W_{x,z}(u) \rangle_x. \tag{9}$$

Integrating the Hessian along the radial geodesic gives

$$p(u) - p(0) = \int_0^1 \Pi_{su}^{-1} \nabla^2 f(y_{su})[\Pi_{su} u]\, ds, \qquad \|p(u) - p(0)\|_x \le H \|u\|_x,$$

where the Hessian is identified with a self-adjoint linear operator. The distribution of u is symmetric, so

$$\|\mathbb{E} W_{x,z}(u)\|_x \le \frac{1}{2} \mathbb{E} \|W_{x,z}(u) + W_{x,z}(-u)\|_x \le \frac{\beta_2}{2} \mathbb{E} \|u\|_x^2.$$

Together with $\|p(0)\|_x \le L$ and Equation (7), this yields

$$\begin{aligned}\mathbb{E}\langle p(u), W_{x,z}(u)\rangle_x &= \mathbb{E}\langle p(u)-p(0), W_{x,z}(u)\rangle_x + \langle p(0), \mathbb{E}W_{x,z}(u)\rangle_x \\ &\ge -\left(H\beta_1 + \frac{L\beta_2}{2}\right)\mathbb{E}\|u\|_x^2. \end{aligned} \tag{10}$$

For a uniform unit-ball vector, rotational invariance and $\mathbb{E}\|v\|_x^2 = d/(d+2)$ give

$$\mathbb{E}[vv^\top] = \frac{I}{d+2}, \qquad \mathbb{E}\|u\|_x^2 = \frac{\delta^2 \operatorname{tr}(A_x^2)}{d+2}.$$

Averaging Equation (9) and combining these facts proves

$$\langle G_\delta f(x), -q\rangle_x \ge \hat{f}(x) - f(z) - \frac{\delta^2 \operatorname{tr}(A_x^2)}{d+2}\left(H\beta_1 + \frac{L\beta_2}{2}\right).$$

This is Equation (3), the conclusion used by the main regret theorem, including for boundary-approaching comparators.

Step 4: optional comparator smoothing and uniformity. If the two-sided radial segments at z are feasible for the chosen radius, their midpoint is z, so geodesic convexity gives

$$f(z) \le \frac{1}{2}\left\{f(\operatorname{Exp}_z(\delta A_z v)) + f(\operatorname{Exp}_z(-\delta A_z v))\right\}.$$

Averaging gives $f(z) \le \hat{f}(z)$, proving Equation (4). Since $\operatorname{tr}(A_x^2) \le d/m_R$, taking the supremum over all admissible t, x, z proves Equation (5).

To obtain L from existing local data, choose a fixed $r_0 > 0$ such that every center's radius-r_0 geodesic ball is feasible and lies in the bounded-loss query tube. Compact interiority and positive lower eigenvalue bounds for A_x allow such a choice within the family of admissible Dikin balls. If $\nabla f(x) \ne 0$, put $e = \nabla f(x)/\|\nabla f(x)\|_x$. Convexity gives

$$r_0\|\nabla f(x)\|_x \le f(\operatorname{Exp}_x(r_0 e)) - f(x) \le 2C.$$

Thus one may use $L = 2C/r_0$, or the Lipschitz constant already present in assumption 2. This uses no derivative beyond the Hessian. Finally, when $H = 0$, f is affine along every relevant geodesic. Then $\hat{f} = f$, $G_\delta f(x) = \nabla f(x)$, and $f(z) - f(x) = \langle \nabla f(x), q\rangle_x$, so the support error is zero. □

Lemma 7 (Estimator moment bound)**.** *For the one-point estimator in Algorithm 1,*

$$\mathbb{E}_t\|g_t\|_{R,x_t}^{*2} \le \left(\frac{dC}{\delta}\right)^2.$$

Proof. The loss is bounded by C, and ξ_t has unit norm. The estimator is $g_t = (d/\delta) f_t(y_t) A_t^{-1}\xi_t$, where $A_t = (\nabla^2 R(x_t))^{-1/2}$. The dual barrier norm satisfies

$$\|A_t^{-1}\xi_t\|_{R,x_t}^{*2} = \xi_t^\top A_t^{-1}(\nabla^2 R(x_t))^{-1} A_t^{-1}\xi_t = \|\xi_t\|_{x_t}^2,$$

exactly when operators and norms are taken in the Riemannian metric, as in the algorithm. In fact $\|g_t\|_{R,x_t}^* \le dC/\delta$ holds pathwise, so one may take $c_g = 1$. □

Proof of Proposition 1. Definition 2 isolates the smoothed-comparator support inequality. In flat fixed-ellipsoid smoothing, the surface-gradient identity and convexity establish this inequality with zero error. Since

$$h_t(x_t) - h_t(z) = \left\langle g_t, -\operatorname{Log}_{x_t}(z)\right\rangle_{x_t},$$

substituting the definition of g_t into Definition 2 gives the displayed lower bound for comparators with feasible smoothing. The main regret proof instead uses the stronger raw-comparator bound Equation (3). The moment claim follows from Lemma 7. □

Proof of Lemma 2. In the flat case the exponential and logarithm maps are ordinary addition and subtraction, so the smoothed loss is

$$\hat{f}_t(x) = \mathbb{E}_{v\sim\mathrm{Unif}(B)} f_t(x + \delta A v),$$

with A independent of the center over the smoothing region. Differentiating the ball average and applying the divergence theorem gives the standard surface identity

$$\nabla \hat{f}_t(x) = \mathbb{E}_{\xi\sim\mathrm{Unif}(S)} \left[\frac{d}{\delta} f_t(x + \delta A\xi) A^{-1}\xi \right],$$

where S is the unit sphere and A^{-1} denotes the self-adjoint inverse under the Euclidean metric. Convexity of $\hat{f}_t$ therefore implies

$$\hat{f}_t(x_t) - \hat{f}_t(z) \le \langle \nabla \hat{f}_t(x_t), x_t - z\rangle = \mathbb{E}_t \left[\left\langle \frac{d}{\delta} f_t(y_t) A^{-1}\xi_t, x_t - z \right\rangle \right].$$

Because $x_t - z = -\operatorname{Log}_{x_t}(z)$ in Euclidean space, this is precisely Definition 2 with $\varepsilon_{\mathrm{sm}}(\delta) = 0$. □

Lemma 8 (Proximal existence and one-step stability)**.** *Under the barrier and geometric hypotheses, each center update in Algorithm 1 admits an interior global minimizer. Retain the fixed compact-center and uniform normal-tube premises. With any positive step size η_t and barrier weight η_t/T, every such minimizer satisfies, pathwise for each admissible comparator z,*

$$\begin{aligned} \langle g_t, -\operatorname{Log}_{x_t} z\rangle \le \frac{V_z(x_t) - V_z(x_{t+1})}{\eta_t} + L_V M_R \eta_t \|g_t\|^{*2}_{R,x_t} \\ + \frac{K_R D}{T} + \frac{L_V K_R^2 \eta_t}{T^2}. \end{aligned} \tag{11}$$

Proof. Fix $x = x_t$ and write $r = d_{\mathcal{M}}(x, y)$. Geodesic convexity gives $R(y) \ge R(x) - \|\nabla R(x)\|_x r$. The update objective is bounded below by

$$\tfrac{1}{2} r^2 - \eta_t \big(\|g_t\|_x + \|\nabla R(x)\|_x/T\big) r + (\eta_t/T) R(x).$$

Hence a finite objective sublevel has bounded distance from x. Completeness makes closed bounded balls compact, and barrier blow-up excludes a limit point on $\partial\mathcal{K}$. Continuity in $\mathcal{K}^\circ$ therefore gives an interior minimizer. This proves existence, not horizon-uniform compactness of the resulting sequence.

For this minimizer write $y = x_{t+1}$, $v = \operatorname{Log}_x y$ and $q = \operatorname{Log}_x z$. Pulling the objective back by Exp_x gives the exact first-order condition

$$v + \eta_t g_t + \frac{\eta_t}{T} p_t = 0, \qquad p_t = (d\operatorname{Exp}_x)^*_v \nabla R(y), \quad \|p_t\|_x \le K_R. \tag{12}$$

Consequently $\|v\|_x \le \eta_t(\|g_t\|_x + K_R/T)$. Along the update geodesic, $\nabla V_z(x) = -q$ and the Hessian bound give

$$\langle v, q\rangle_x \le V_z(x) - V_z(y) + \frac{L_V}{2}\|v\|_x^2.$$

Taking the inner product of Equation (12) with q yields

$$\begin{aligned} \langle g_t, -q\rangle_x &\le \frac{V_z(x) - V_z(y)}{\eta_t} + \frac{L_V \eta_t}{2}(\|g_t\|_x + K_R/T)^2 + \frac{K_R D}{T} \\ &\le \frac{V_z(x) - V_z(y)}{\eta_t} + L_V M_R \eta_t \|g_t\|^{*2}_{R,x} + \frac{L_V K_R^2 \eta_t}{T^2} + \frac{K_R D}{T}, \end{aligned}$$

where the last step uses Lemma 1. Neither convexity of h_t nor an accumulated-surrogate Hessian bound is used. □

Proof of Proposition 2. Set $\eta_t = \eta$ in Equation (11), sum over t, and discard $-V_z(x_{T+1})/\eta \le 0$. Since $h_t(x_t) - h_t(z) = \langle g_t, -\operatorname{Log}_{x_t} z\rangle$, the result follows with $c_f = L_V M_R$ and, for example, $c_0 = K_R D + L_V K_R^2$. □

Lemma 9 (Concrete update-tube sufficient condition)**.** *Let $B_R = \sup_{x\in\mathcal{X}} \|\nabla R(x)\|_x$, so $B_R^2 \le \nu M_R$. Every proximal step satisfies*

$$d_{\mathcal{M}}(x_t, x_{t+1}) \le 2\eta_t(\|g_t\|_{x_t} + B_R/T).$$

Thus, for a normal update radius r_u covered by the local data, the one-point condition $2\eta(\sqrt{M_R}dC/\delta + B_R/T) \le r_u$ suffices.

Proof. Compare the update objective at x_{t+1} and x_t. With $v = \mathrm{Log}_{x_t} x_{t+1}$, geodesic convexity of R gives

$$\tfrac{1}{2}\|v\|^2 \le -\eta_t\langle g_t, v\rangle - (\eta_t/T)\langle \nabla R(x_t), v\rangle \le \eta_t(\|g_t\| + B_R/T)\|v\|.$$

Divide by $\|v\|$ if it is nonzero. The norm comparison and Lemma 7 give the stated sufficient condition. It controls step length, not the full trajectory's uniform interior margin. □

Lemma 10 (Comparator shrinkage)**.** *Under assumption 4 and assumption 2, for every comparator $x^\star \in \mathcal{K}$ and horizon $T \ge 2$, there is an interior comparator $\tilde{x} \in \mathcal{K}^\circ$ such that*

$$R(\tilde{x}) - R(x_0) \le c_{\mathrm{sh}}\nu \log T, \qquad \sum_{t=1}^{T} f_t(\tilde{x}) - \sum_{t=1}^{T} f_t(x^\star) \le c_L L\sqrt{T}.$$

Proof. Choose $\epsilon = T^{-1/2}$ and set $\tilde{x} = \gamma_{x_0,x^\star}(1-\epsilon)$. assumption 4 gives the barrier bound. Geodesic convexity gives

$$f_t(\tilde{x}) \le (1-\epsilon)f_t(x^\star) + \epsilon f_t(x_0),$$

and the Lipschitz bound gives a per-round increase at most $L\epsilon d_{\mathcal{M}}(x_0, x^\star) \le LD\epsilon$. The assertion holds pathwise simultaneously for every comparator. The geometric constants along these paths must remain uniform as specified above; their individual interiority does not establish this. □

Lemma 11 (Uniform adaptive-comparator bound)**.** *Suppose $\|g_t\|_{x_t} \le G$ almost surely and put $\bar{g}_t = \mathbb{E}_t g_t$. Let $\mathcal{Z}_T$ be a deterministic admissible comparator family in the fixed compact normal tube. It admits a $1/T$-net of size $N_T \le C_{\mathrm{cov}}(1+T)^d$, with C_{cov} independent of T. Then*

$$\mathbb{E}\sup_{z\in\mathcal{Z}_T}\sum_{t=1}^{T}\langle g_t - \bar{g}_t, \mathrm{Log}_{x_t} z\rangle \le 2GD\sqrt{2T\log(\max\{N_T, 2\})} + 2GL_{\log}. \tag{13}$$

For a deterministic fixed comparator, the corresponding expected sum is zero.

Proof. A fixed compact smooth d-dimensional neighborhood has the stated covering bound, by a finite coordinate cover. For each fixed net point z_i, the increments of $M_T(z_i) = \sum_t \langle g_t - \bar{g}_t, \mathrm{Log}_{x_t} z_i\rangle$ are conditionally centered and have absolute value at most $2GD$. Conditional Hoeffding's inequality gives $\mathbb{E}e^{\lambda M_T(z_i)} \le e^{2\lambda^2 G^2 D^2 T}$. Taking the maximum over the net and optimizing $\lambda > 0$ yields $\mathbb{E}\max_i M_T(z_i) \le 2GD\sqrt{2T\log(\max\{N_T, 2\})}$. If $d_{\mathcal{M}}(z, z_i) \le 1/T$, the uniform logarithm derivative bound gives $|M_T(z) - M_T(z_i)| \le 2GL_{\log}$, proving the result. The logarithm bound is applied along the comparator geodesics in the same tube. □

A.3 Proof of the One-Point Regret Theorem

Proof of Theorem 1. Let $\mathcal{Z}_T = \{\gamma_{x_0,u}(1-T^{-1/2}) : u \in \mathcal{K}\}$ be the deterministic family of shrunken comparators. By Equation (3) and Lemma 5, on every pre-query history and simultaneously for all $z \in \mathcal{Z}_T$,

$$\mathbb{E}_t f_t(y_t) - f_t(z) \le \langle \bar{g}_t, -\mathrm{Log}_{x_t} z\rangle + (c_q H + C_{\mathrm{sm}})\delta^2, \qquad \bar{g}_t = \mathbb{E}_t g_t. \tag{14}$$

This uses neither smoothing at z nor smoothness of $\hat{f}_t$ as a moving-center function. Set

$$M_T(z) = \sum_t \langle g_t - \bar{g}_t, \mathrm{Log}_{x_t} z\rangle, \qquad e_t = f_t(y_t) - \mathbb{E}_t f_t(y_t).$$

Summing Equation (14) gives the pathwise inequality

$$\begin{aligned}\sum_t [f_t(y_t) - f_t(z)] &\le \sum_t \langle g_t, -\operatorname{Log}_{x_t} z\rangle + M_T(z) + \sum_t e_t + (c_q H + C_{\mathrm{sm}})T\delta^2 \\ &\le \frac{D^2}{2\eta} + L_V M_R \eta T (dC/\delta)^2 + K_R D + \frac{L_V K_R^2 \eta}{T} \\ &\quad + M_T(z) + \sum_t e_t + (c_q H + C_{\mathrm{sm}})T\delta^2.\end{aligned}$$

The second line uses the pathwise proximal bound and Lemma 7. Take the supremum over $z \in \mathcal{Z}_T$ before taking expectation. The scalar noise $\sum_t e_t$ is common to all comparators and has zero mean. With $G = \sqrt{M_R} dC/\delta$, Lemma 11 bounds the remaining supremum by $O((dC/\delta)\sqrt{T\log(2T)})$. Finally Lemma 10 applies pathwise to every $u \in \mathcal{K}$ and adds $LD\sqrt{T}$. Taking the hindsight infimum then proves the theorem. For an oblivious loss sequence one may instead fix an optimal comparator (or a deterministic approximating sequence), and the martingale term has zero expectation. □

A.4 Rate Calculation

Under the uniform hypotheses of Theorem 1, the coefficient C_{sm} from Equation (1) controls the direct support remainder. Balance the variance and initial-distance terms by $\eta \asymp \delta/(dC\sqrt{T})$, subject to the update-tube restriction. The bound becomes

$$(c_1 H + C_{\mathrm{sm}})T\delta^2 + O\left(\frac{dC\sqrt{T\log(2T)}}{\delta}\right) + c_4 L\sqrt{T} + O(1 + \eta/T).$$

Ignoring logarithms and fixed data, the balance remains $T\delta^2 \asymp \sqrt{T}/\delta$. Thus

$$\delta \asymp T^{-1/6}, \qquad \mathrm{Reg}_T = \widetilde{O}(T^{2/3}).$$

Moreover $\eta\|g_t\| = O(T^{-1/2})$, so the sufficient step-radius condition holds for all sufficiently large T; constant caps handle the remaining horizons. For an oblivious loss sequence the logarithm from the comparator net is absent. The two-point theorem differs because its estimator bound is independent of $1/\delta$.

A.5 Proof of the Two-Point Results

Lemma 12 (Two-point estimator moment from local Lipschitzness)**.** *Suppose the symmetric Dikin queries y_t^+ and y_t^- are feasible and lie in the fixed compact normal tube. If the losses are L_q-Lipschitz there and $\|A_t\|_{\mathrm{op}} \le b_A$, then, independently of δ,*

$$\|g_t^{(2)}\|_{R,x_t}^* \le G_2 := dL_q b_A \quad \textit{almost surely}, \qquad \mathbb{E}_t\|g_t^{(2)}\|_{R,x_t}^{*2} \le G_2^2.$$

The local gradient and Hessian bounds in Lemma 6 provide such an L_q on the radial query tube.

Proof. The triangle inequality along the two radial geodesics gives

$$d_{\mathcal{M}}(y_t^+, y_t^-) \le 2\delta\|A_t\xi_t\| \le 2b_A\delta.$$

Consequently $|f_t(y_t^+) - f_t(y_t^-)| \le 2L_q b_A \delta$. The local-norm identity is exact:

$$\|A_t^{-1}\xi_t\|_{R,x_t}^* = \|A_t A_t^{-1}\xi_t\| = 1.$$

Substitution into the estimator formula proves the almost-sure bound and hence the conditional second-moment bound. Finally, if the gradient at a center is bounded by L_0 and the radial Hessian bound is H, integration along a query geodesic gives $L_q \le L_0 + Hb_A\bar{\delta}$ on the radial tube. The bounded-loss and interior-margin argument in Lemma 6 supplies L_0 when needed. No third derivative of a loss is used. □

Proof of Theorem 2. Condition on the history before the fresh direction, including the nonanticipating loss f_t. Antipodal symmetry gives the exact identity

$$\begin{aligned}\mathbb{E}_t g_t^{(2)} &= \frac{d}{2\delta} A_t^{-1} \left(\mathbb{E}_t[f_t(y_t^+)\xi_t] - \mathbb{E}_t[f_t(y_t^-)\xi_t] \right) \\ &= \frac{d}{\delta} A_t^{-1} \mathbb{E}_t[f_t(y_t^+)\xi_t] = G_\delta f_t(x_t).\end{aligned}$$

Let $\ell_t = (f_t(y_t^+) + f_t(y_t^-))/2$. The direct support inequality in Lemma 6, together with the query-smoothing bound, gives for every admissible comparator z

$$\mathbb{E}_t \ell_t - f_t(z) \le \langle \mathbb{E}_t g_t^{(2)}, -\operatorname{Log}_{x_t} z \rangle + (c_1' H + C_{\mathrm{sm}})\delta^2.$$

Apply Equation (11) and sum with constant η. Its distance potential telescopes, the estimator term is at most $L_V M_R \eta T G_2^2$, and the barrier correction contributes $K_R D + L_V K_R^2 \eta / T$. For a hindsight comparator under adaptive losses, Lemma 11 applies with the ordinary-norm bound $G = \sqrt{M_R} G_2$, yielding the additional $c_5' G_2 \sqrt{T \log(2T)}$ term and a fixed-data remainder. For a fixed comparator this martingale term has zero expectation. The centered query noise $\ell_t - \mathbb{E}_t \ell_t$ is common to all comparators and also has zero expectation. Boundary truncation, when needed, is exactly the Lipschitz comparison used in Theorem 1. These bounds give the displayed theorem. Choosing $\eta \asymp 1/(G_2\sqrt{T})$ and $\delta = O(T^{-1/4})$ proves the stated rate. □

B ONE-POINT EXPERIMENTAL DETAILS

This experiment uses the breast-cancer covariance targets and the loss $f_t(X) = \frac{1}{2} d_{\mathrm{AI}}(X, \Sigma_t)^2$. Each method observes one exact function value per round, and the metric is mean queried loss per round. The comparisons below use separate held-out protocols: frozen-parameter horizon transfer, two forms of retuning, and a scalar-baseline check.

B.1 FROZEN-PARAMETER HORIZON TRANSFER

The screen uses seeds 116026–116028 and the disjoint validation uses seeds 117026–117030, both at $T = 200$. Each method receives the same 64-candidate schedule grid. We pass the eight best admissible screen candidates to validation (four for Dikin, since only four candidates pass its solver and feasibility checks) and freeze the configuration with the lowest validation loss. Dikin uses $\eta_t = 2 \times 10^{-7}/\sqrt{t}$, radius scale 1.0, and $(\lambda, \gamma) = (20, 1000)$; isotropic proximal uses $\eta_t = 2 \times 10^{-5}/\sqrt{t}$ and radius scale 0.5; projected Euclidean uses $\eta_t = 2 \times 10^{-4}/\sqrt{t}$ and radius scale 0.125.

The original held-out evaluation used seeds 118026–118045 at $T \in \{200, 500, 1000\}$. We retained these frozen parameters for an independent 20-seed replication at $T = 500, 1000$ (seeds 120100–120119) and a 60-seed confirmation at $T = 1000$ (seeds 121100–121159). The pooled $T = 1000$ result therefore uses 100 distinct paired seeds, none used for screening or validation. The $T = 500$ evaluation uses the first two blocks, so its 40 seeds are a subset of the $T = 1000$ seed set; the two horizons are not independent replications of each other. The solver uses 20 continuation stages and at most 60 iterations per stage.

Table 5: One-point breast-cancer covariance results at $T = 1000$. Loss is mean queried loss per round. Reductions are paired relative to projected Euclidean; intervals use 100 independent test seeds.

Method	Mean loss	Reduction (95% CI)	Wins
Proximal Dikin	7.966	5.08% [3.08, 7.07]	74/100
Proximal isotropic	8.037	4.39% [2.56, 6.22]	69/100
Projected Euclidean	8.506	–	–

The same frozen $T = 200$ selection was evaluated at $T = 500$ on two disjoint 20-seed blocks: the original seeds 118026–118045 and the independent replication seeds 120100–120119. Pooling all 40 paired runs gives positive point estimates for both proximal methods, although the paired intervals include zero at this horizon.

Table 6: One-point results at $T = 500$ under the parameters selected at $T = 200$. Each row pools 40 fresh paired seeds from two disjoint batches; reductions are paired relative to projected Euclidean.

Method	Mean loss	Reduction (95% CI)	Wins
Proximal Dikin	7.930	2.93% [−0.72, 6.59]	24/40
Proximal isotropic	7.967	2.60% [−0.83, 6.02]	24/40
Projected Euclidean	8.302	–	–

The $T = 1000$ comparisons favor both proximal configurations, while the $T = 500$ intervals include zero. The isotropic result is similar to Dikin, and the two arms have different step and radius settings. These results demonstrate successful horizon transfer for the evaluated proximal configurations.

B.2 BARRIER-SCALE RETUNING

We also selected only the Dikin barrier scale at $T = 1000$, holding its step and radius at $\eta_t = 2 \times 10^{-7}/\sqrt{t}$ and radius scale 1.0. Screening tested $\gamma \in \{30, 100, 300, 1000, 3000, 10000\}$ on seeds 126026–126028; validation on seeds 126031–126035 selected $\gamma = 30$. The isotropic and projected settings were held at their previously selected values. On 60 fresh paired seeds 126100–126159, Dikin lowers mean query loss by 3.97% relative to projected Euclidean (95% paired CI

$[1.47, 6.47]\%$; 44/60 wins) and by 1.75% relative to the isotropic proximal arm (95% paired CI $[0.50, 3.00]\%$; 36/60 wins; $p = 0.0067$).

Table 7: One-point $T = 1000$ results from the independent barrier-scale-only retuning confirmation. Reductions are paired relative to projected Euclidean.

Method	Mean loss	Reduction (95% CI)	Wins
Proximal Dikin ($\gamma = 30$)	7.851	3.97% $[1.47, 6.47]$	44/60
Proximal isotropic	8.013	2.34% $[0.41, 4.26]$	40/60
Projected Euclidean	8.260	–	–

B.3 Step/Radius Retuning

This protocol uses the same barrier-scale selection at $T = 1000$: we screened $\gamma \in \{30, 100, 300, 1000, 3000, 10000\}$ on seeds 126026–126028 and selected the lowest-loss scale among those passing all five validation seeds 126031–126035; this selected $\gamma = 30$. Holding $\gamma = 30$ fixed, a 12-candidate step/radius screen followed by validation on the same calibration and validation blocks selected $\eta_t = 6 \times 10^{-8}/\sqrt{t}$ and radius scale 2.0. The final evaluation uses 60 fresh paired seeds 126200–126259. The isotropic and projected baselines retain their previously frozen settings, $\eta_t = 2 \times 10^{-5}/\sqrt{t}$ with $(\lambda, \gamma) = (20, 1000)$ and $\eta_t = 2 \times 10^{-4}/\sqrt{t}$, respectively; their radius scales remain 0.5 and 0.125. The two retuning protocols share calibration and validation seeds, but use disjoint test blocks and are reported separately.

Table 8: One-point breast-cancer covariance results for the separate $T = 1000$ step/radius retuning confirmation. Reductions are paired relative to projected Euclidean; intervals use 60 fresh test seeds.

Method	Mean loss	Reduction (95% CI)	Wins
Proximal Dikin ($\gamma = 30$)	7.851	5.58% $[2.22, 8.94]$	44/60
Proximal isotropic	8.103	3.19% $[0.81, 5.57]$	40/60
Projected Euclidean	8.461	–	–

On these paired seeds, Dikin also reduces mean query loss by 2.70% relative to the separately configured isotropic proximal arm (95% paired CI $[1.08, 4.32]\%$; 41/60 wins; $p = 0.0015$). The retuned Dikin configuration therefore improves over both controls in this comparison. Its step, radius, and barrier scale differ from the isotropic arm, so the comparison does not isolate the effect of the Dikin query map.

B.4 Predictable-Baseline Check

We evaluated the frozen Dikin configuration from the predictable-baseline screen on 60 new paired seeds 290100–290159 at $T = 1000$. The Dikin step and radius scales were 6×10^{-8} and 2.0, respectively, with $\gamma = 30$ and $\eta_t \propto t^{-1/2}$. We compared the same Dikin update using no scalar baseline and an EMA-0.5 baseline, together with the projected baseline fixed at step scale 10^{-4} and radius scale 0.125. The projected step scale differs from the 2×10^{-4} setting in the preceding protocols; this is a separate comparison, not a pooled extension of those controls.

Table 9: One-point $T = 1000$ evaluation on 60 fresh paired seeds with frozen settings. Reductions are paired relative to projected Euclidean.

Method	Mean loss	Reduction (95% CI)	Wins
Proximal Dikin, no baseline	7.850	5.02% $[2.14, 7.91]$	41/60
Proximal Dikin, EMA-0.5	7.850	5.01% $[2.09, 7.93]$	41/60
Projected Euclidean	8.376	–	–

The paired EMA-versus-no-baseline reduction is 0.004% (95% CI $[-0.047, 0.054]\%$; 29/60 wins). There is no detectable difference between the two variants in this evaluation, and the no-baseline arm already improves over projection. This is not a formal equivalence test or a claim that predictable baselines have no effect in other settings.

B.5 Reporting and Reproducibility

Reductions are averages of per-seed percentage loss reductions against the stated comparator, not percentage differences between the displayed mean losses. Confidence intervals and the two reported p-values are retained from the experiment summaries, without reconstructing them from rounded aggregates. The summaries report that all evaluated runs passed the solver, KKT, and query/center-feasibility checks; the present manuscript package does not contain the underlying solver logs for independent verification.

The available one-point records are aggregate results, not per-seed loss records or trajectories. In the `analysis` directory, `one_point_summary_data.json` transcribes these aggregates and `one_point_paper_artifacts.py` generates the main-text loss table and effect plot; these reproduce the aggregates, not the experimental runs. The retained records do not specify every schedule-grid entry, the confidence-interval construction, or the predictable-baseline screen and EMA initialization, recursion, and update timing. Those details are not inferred from the summaries.

C TWO-POINT EXPERIMENTAL DETAILS

Figure 2 summarizes the complete two-point real-data and synthetic evaluations. We provide their metrics, protocols, and numerical results below.

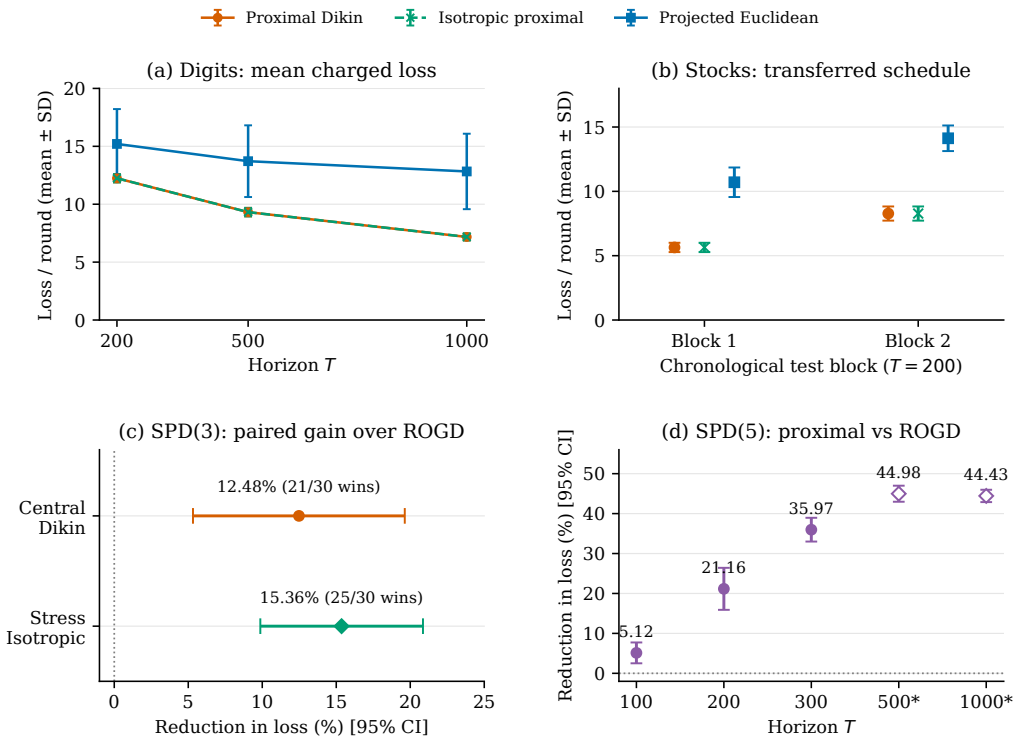


Figure 2: Two-point covariance tracking. (a,b) Mean charged query loss per round with across-seed SD: digits uses 20 seeds per horizon, and each stock block uses five. (c,d) Mean paired percentage reduction relative to intrinsic ROGD, with 95% confidence intervals; positive values favor the proximal arm. The SPD(3) stress result uses isotropic proximal, not Dikin. In (d), open diamonds and starred horizons transfer the frozen $T = 300$ configurations without retuning; shorter horizons are calibrated separately.

C.1 METRICS AND STATISTICAL REPORTING

All two-point comparisons use the affine-invariant squared-distance loss $f_t(X) = \frac{1}{2} d_{\mathrm{AI}}(X, \Sigma_t)^2$ and two exact function observations per round. The per-round charge is the arithmetic average of the two queried losses, not a center loss or a best-fixed-comparator regret. For each paired seed we first compute mean charged loss over the horizon, then compute the percentage reduction relative to that seed's baseline. The reported effect is the average of these per-seed percentages. Win counts use the same paired comparison. Real-data loss spreads are SDs, whereas effect bars are 95% confidence intervals; they must not be read interchangeably.

The real-data effect intervals are paired Student-t intervals. We retain the supplied synthetic intervals without refitting them from rounded means. Query seeds on real-data tasks quantify algorithmic randomness conditional on the fixed target stream, not uncertainty over a population of datasets. For SPD(3), seeds also determine the generated target paths. All intervals are per-cell summaries, without a simultaneous multiple-comparison claim.

The common static-comparator term cancels exactly in a paired regret difference, as shown in (2). Consequently, the saved means recover the mean regret advantage whenever both absolute losses are available. Table 10 reports all nine such cells, using the common target sequence and comparator class for each pair. Positive entries mean lower proximal regret. Rounding of the archived loss

means makes these recovered values approximate. Loss-reduction percentages alone do not recover the three remaining SPD(5) regret gaps. Moreover, the reported intervals for mean percentage loss reduction cannot be reused as intervals for mean absolute regret differences without the paired seed-level data.

Table 10: Mean cumulative regret advantage recovered from paired-task loss means: $\overline{R}_T^{\mathrm{base}} - \overline{R}_T^{\mathrm{prox}} = T(\overline{L}_T^{\mathrm{base}} - \overline{L}_T^{\mathrm{prox}})$. All entries are approximate because the input means are rounded. No comparator loss is needed for this difference, but absolute regret values are not available. The SPD(3) stress arm is isotropic proximal.

Task	T	Baseline	Mean regret advantage
Digits	200	Projected Euclidean	592.8
Digits	500	Projected Euclidean	2200.0
Digits	1000	Projected Euclidean	5656.0
Stock 1	200	Projected Euclidean	1013.2
Stock 2	200	Projected Euclidean	1170.0
SPD(3), central	300	Intrinsic ROGD	0.285
SPD(3), stress	300	Intrinsic ROGD	1.317
SPD(5)	100	Intrinsic ROGD	0.161
SPD(5)	300	Intrinsic ROGD	2.934

C.2 Digits Calibration and Chronological Stock Transfer

Digits uses PCA dimension six and a fixed class-covariance target sequence. Calibration uses seeds 14026–14028 and 15026–15030; evaluation uses disjoint seeds 24026–24030 and 25026–25040. Thus each of $T = 200, 500, 1000$ has 20 paired evaluation seeds, which vary query randomness rather than data samples. The proximal schedule is $\eta_t = 0.02/\sqrt{t}$, with barrier factors $(\lambda, \gamma) = (20, 10)$; the projected Euclidean baseline uses its separately calibrated $\eta_t = 8/t$ schedule. The query-radius scales are 0.125 for the Dikin and isotropic proximal arms and 0.25 for projection. These are the recorded method-specific scales, not an assertion of matched physical query radii. The isotropic ablation changes only the query map and preserves the proximal subproblem and all its other parameters. Proximal subproblems use a 20-stage BFGS continuation solve.

The stock stream contains AAPL, MSFT, IBM, SBUX, JPM, and BAC in that order. The first 300 rolling covariance targets determine the spectral window. The following 400 targets form two consecutive, non-overlapping 200-round test blocks; each is evaluated using five paired query seeds. Both proximal and projected configurations are transferred unchanged from digits, without cycling the stream or tuning on stock losses. These are two chronological blocks of the same rolling series, not independent stock datasets. Table 11 supplies the complete paired statistics accompanying Table 3.

Table 11: Paired two-point real-data effects of proximal Dikin relative to projected Euclidean. Reduction is the mean per-seed percentage decrease with a 95% paired-t interval. n counts paired query seeds and wins count strictly lower Dikin loss.

Stream	T	n	Reduction [95% CI]	Wins
Digits	200	20	16.37% $[8.44, 24.31]$	17/20
Digits	500	20	28.76% $[21.26, 36.27]$	19/20
Digits	1000	20	40.55% $[33.46, 47.63]$	20/20
Stock 1	200	5	46.91% $[39.92, 53.90]$	5/5
Stock 2	200	5	41.31% $[36.75, 45.86]$	5/5

C.3 CONTROLLED SPD(3) CYCLES

For the central cycle, we generate a horizon-300 target path by exponentiating smooth cyclic combinations of three seed-specific symmetric matrices in the tangent space at $1.1I$. The directions are generally noncommuting, and the central targets remain in the spectral interval $[0.2I, 2I]$. The intrinsic ROGD baseline uses symmetric exponential-map queries with feasible dyadic radius backtracking, followed by an affine-invariant exponential center step with feasible-step backtracking. It does not use the barrier-proximal subproblem.

Dikin, isotropic proximal, and ROGD are independently calibrated, each with 48 candidates, four screening seeds, and eight disjoint validation seeds. The frozen configurations are then evaluated on 30 fresh paired test seeds. The selected central Dikin configuration has step scale 0.16, query-radius scale $0.0009765625 = 1/1024$, recorded decay power 0.1, and barrier factors $(\lambda, \gamma) = (50, 1)$. ROGD selects step scale 0.32, the same radius scale, and the same recorded decay power. Dikin and isotropic proximal both have mean charged loss 0.00549, to the displayed precision, compared with 0.00644 for ROGD. The supplied paired reduction is 12.48% ($[5.32, 19.63]\%$; 21/30 wins).

The separate boundary-stress regime stretches affine-log target paths by a factor of two. This transformation is not assumed to preserve the central target interval. Each method again receives 48 candidates, with disjoint screening, validation, and test seeds. Validation selects the isotropic proximal query map, with step scale 0.16, radius scale $1/1024$, recorded decay power 0.1, and $(\lambda, \gamma) = (50, 1)$; ROGD selects step scale 0.32 with the same radius scale and decay power. A fresh 30-seed test gives mean losses 0.02064 and 0.02503, respectively, and a paired reduction of 15.36% ($[9.87, 20.86]\%$; 25/30 wins). This result concerns isotropic proximal; no Dikin stress-test effect is inferred from it.

C.4 SPD(5) HORIZON EVALUATIONS AND FROZEN TRANSFER

SPD(5) uses independent horizon-specific calibration for proximal and ROGD at $T = 100, 200, 300$. The $T = 100$ result is an independent 100-seed replication of its equal-budget selection; $T = 200$ uses 20 fresh paired seeds and $T = 300$ uses 30. The evaluations at $T = 500, 1000$ each use 20 fresh paired seeds and transfer both methods' frozen $T = 300$ settings without retuning. Hence the last two rows test horizon transfer, whereas the first three are separately calibrated comparisons, not a single fixed-policy scaling curve.

All five paired effects and their win counts appear in Table 12. The supplied summaries contain absolute loss means for $T = 100, 300$, but not for $T = 200, 500, 1000$. Table 13 displays this distinction explicitly; no missing absolute values or per-round trajectories are reconstructed from percentage reductions. The reported isotropic arm has essentially the same performance as Dikin on SPD(5), but its precise means and SDs are not available in these summaries.

Table 12: Complete synthetic two-point effects versus intrinsic ROGD. Reduction is the mean paired percentage decrease in charged loss with its 95% confidence interval; wins count paired seeds with lower proximal loss. The query column identifies the reported proximal arm; "Unspecified" means the source does not identify its query map for that cell. †: both methods transfer frozen $T = 300$ settings without retuning. Available absolute losses are in Table 13.

Regime	T	Query	Reduction [95% CI]	Wins
SPD(3), central	300	Dikin	12.48% $[5.32, 19.63]$	21/30
SPD(3), stress	300	Isotropic	15.36% $[9.87, 20.86]$	25/30
SPD(5)	100	Dikin	5.12% $[2.51, 7.73]$	67/100
SPD(5)	200	Unspecified	21.16% $[15.90, 26.42]$	19/20
SPD(5)	300	Dikin	35.97% $[33.00, 38.95]$	30/30
SPD(5)	500†	Unspecified	44.98% $[42.97, 46.99]$	20/20
SPD(5)	1000†	Unspecified	44.43% $[42.88, 45.97]$	20/20

Table 13: Available absolute two-point synthetic losses (mean charged loss per round). The proximal query type and paired test size are shown for each cell. A dash means the absolute mean is absent from the supplied aggregate records; the paired effect for that cell is nevertheless available in Table 12. Synthetic per-arm SDs are not supplied.

Regime	T	n	Query	Proximal	ROGD
SPD(3), central	300	30	Dikin	0.00549	0.00644
SPD(3), stress	300	30	Isotropic	0.02064	0.02503
SPD(5)	100	100	Dikin	0.02237	0.02398
SPD(5)	200	20	Unspecified	–	–
SPD(5)	300	30	Dikin	0.01691	0.02669
SPD(5)	500	20	Unspecified	–	–
SPD(5)	1000	20	Unspecified	–	–

C.5 Inverse-Time Schedule Sensitivity

A separate check uses the inverse-time proximal schedule $\eta_t = 1/t$ on breast-cancer and digits covariance streams, with fresh seeds 21026–21030, 20-stage continuation, and $\gamma = 1000$ and 300, respectively. Table 14 reports all five complete cells. Projection has lower mean charged loss in each, showing that the main real-data advantage is not uniform over schedules and barrier scales. This is a separate configuration check, not an ablation that isolates the step-size alone. The archived $T = 1000$ digits cell is incomplete, with an unspecified aggregation convention; it is excluded from the comparative table and retained with that qualification in the summary data. No confidence intervals are inferred for this diagnostic from aggregate means.

Table 14: Complete cells of the separate inverse-time schedule check. Entries are mean charged query losses across five seeds; lower is better. The table has no SD or confidence interval because neither is provided in the source summary.

Dataset	T	Proximal	Projected Euclidean
Breast cancer	200	7.736	5.302
Breast cancer	500	7.551	4.587
Breast cancer	1000	7.274	4.126
Digits	200	17.489	15.456
Digits	500	16.285	14.202

C.6 Two-Point Artifact Provenance

The two-point source package contains rounded aggregate summaries rather than raw per-seed records or the full experiment implementation. We preserve these summaries, their source-file hashes, and missing-value flags in `analysis/two_point_summary_data.json`. The following command generates the two-point tables and PDF/PNG figure from those records only:

```
python analysis/two_point_paper_artifacts.py
```

Artifact hashes are written to `analysis/two_point_artifact_manifest.json`. This reproduces the displayed aggregates, not the underlying runs or an independent calculation of their confidence intervals. In particular, the archived synthetic description records a decay power but does not identify its precise schedule term; we retain the reported parameter without inventing a missing schedule formula. Unavailable synthetic SDs, loss means, and seed-level observations are left unreported.

D Related Work and Positioning

Euclidean BCO provides the algorithmic starting point. The one-point smoothing method of Flaxman et al. (2005) and the multi-point framework of Agarwal et al. (2010) show how function values can replace gradients in online convex learning. Self-concordant barriers enter bandit online learning through the SCRiBLe/interior-point line of work (Abernethy et al., 2012), and Saha & Tewari (2011) use this geometry for smooth Euclidean one-point BCO. Later Euclidean work improves the dimension and horizon dependence through more sophisticated smoothing and second-order ideas (Dekel et al., 2015; Shamir, 2017; Bubeck et al., 2021; Fokkema et al., 2024; Lattimore, 2024). These papers are not direct competitors on curved domains, but they are essential for calibration: our $\widetilde{O}(T^{2/3})$ one-point result should be read as an intrinsic barrier-local smooth-bandit guarantee, not as a new best Euclidean rate.

Riemannian optimization and Riemannian OCO provide the geometric background. General-purpose algorithms on matrix manifolds and smooth manifolds are treated by Edelman et al. (1998); Absil et al. (2008); Boumal (2023). Geodesically convex optimization on Hadamard manifolds was studied by Zhang & Sra (2016), and online optimization over Riemannian manifolds was developed by Wang et al. (2023). The latter is especially relevant because it analyzes bandit feedback on manifolds and shows why estimator validity is not automatic; homogeneous and Hadamard-type structure are used in important parts of that analysis. Other recent online manifold work studies dynamic regret and projection-free oracle models (Hu et al., 2023a;b). Stochastic zeroth-order Riemannian optimization is also related (Li et al., 2023), but it usually studies offline stochastic oracle complexity rather than adversarial regret.

Self-concordance on manifolds is not new. Euclidean self-concordant barriers come from the interior-point literature, while Riemannian versions and manifold interior-point methods have been developed by Hirai (2022); Hirai et al. (2023). Our contribution is to use this geometry in an online bandit setting: the barrier defines both feasible local exploration and interior regularization of the proximal step, while the regret proof tracks exponential-map smoothing, scalar surrogates, and local curvature through a squared-distance potential.

Table 15: Positioning relative to the closest algorithmic and geometric lines of work.

Line of work	Feedback	Geometry	What is missing for our setting
Euclidean BCO with smoothing	Function values	Euclidean convex bodies	No tangent-space or injectivity issue
Self-concordant bandits	Function values	Euclidean barrier geometry	Dikin steps are not exponential-map queries
Riemannian OCO	Gradients or oracles	Geodesically convex sets	Usually not bandit and not barrier-local
Manifold self-concordance	Deterministic local geometry	Manifold barriers	No online bandit estimator or regret proof
This paper	Function values	Riemannian domain with barrier	Intrinsic proximal step plus local Dikin smoothing

Table 15 summarizes the gap. The issue is not only to replace Euclidean lines by geodesics. A bandit estimator is anchored at the current center, the proximal step uses its scalar logarithmic surrogate, and the barrier changes the exploration ellipsoid as the iterate moves. The analysis must therefore track three maps simultaneously: the exponential map used for feasible queries, the logarithm map used to form scalar surrogates, and the barrier Hessian used to shape local perturbations. Our proofs are devoted to making these maps explicit; suppressing them leads either to invalid sums of tangent vectors or to an estimator that targets the wrong smoothed function.