TIGHT SAMPLING COMPLEXITY WITH STOCHASTIC GRADIENT ORACLES IN FIXED DIMENSIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

We establish the stochastic-gradient query complexity of sampling smooth strongly log-concave distributions in every fixed dimension $d \geq 1$. For μ-strongly convex, L-smooth potentials with unknown minimizers $x_f^\star$ in the ball $\mathbb{B}(0, \mu^{-1/2})$, and unbiased gradient oracles with variance at most σ^2, the minimax worst-case expected query complexity is $\Theta\left(\log(1+\kappa) + \frac{\sigma^2}{\mu\varepsilon}\right)$, jointly optimal for the condition number $\kappa := L/\mu$, variance $\sigma^2 \geq 0$, and TV accuracy $0 < \varepsilon \leq 1/10$.

In the noiseless setting $\sigma = 0$, this tight complexity $\Theta(\log(1+\kappa))$ is independent of ε. Moreover, a gradient-only sampler generates an exact sample with $O(\log(1+\kappa))$ worst-case expected queries. Without previous ball $\mathbb{B}(0, \mu^{-1/2})$, exact sampling from any initial point x_0 can be implemented with expected cost $O\left(\log(1+\kappa) + \log(1+\sqrt{\mu}\|x_0 - x_f^\star\|)\right)$, without knowing the initial distance. We also show that dependence on initial distance is generally unavoidable.

1 INTRODUCTION

Sampling from a density proportional to e^{-f} is central to Bayesian inference and uncertainty quantification. When full gradients are costly, stochastic gradients offer a natural source of information about the target (Welling & Teh, 2011). For μ-strongly convex, L-smooth potentials, a basic question is how many gradient queries are needed to attain total-variation (TV) accuracy ε. Both ε and the condition number $\kappa = L/\mu$ enter nonasymptotic sampling guarantees (Dalalyan, 2017).

Recent work has clarified the optimal accuracy dependence. In *High-accuracy log-concave sampling with stochastic queries*, Chen et al. (2026b) establish polylogarithmic accuracy dependence under sub-exponential gradient noise. Under only a variance ceiling σ^2, their strongly convex upper bound in $\mathbb{R}^d$ has order $\widetilde{O}(\kappa\sqrt{d}\,[1 + \sigma^2/(\mu\varepsilon)])$, with logarithmic and initialization factors suppressed. At fixed positive variance, their lower bound matches the $1/\varepsilon$ exponent but not the upper bound's multiplicative κ dependence. Thus accuracy optimality, up to logarithmic factors, does not give joint optimality in ε and κ.

This distinction matters for ill-conditioned targets. An accuracy-optimal bound can still overestimate the cost by a large factor when κ grows, and it does not determine whether curvature must multiply the statistical cost. The joint ε–κ dependence is left open by these bounds. Resolving it requires upper and lower bounds that match as both parameters vary, not merely the correct power of $1/\varepsilon$.

Noise adaptivity is equally important. Gradient estimators can have very different variances, and averaging reduces variance only by spending additional queries. A useful complexity characterization should improve as the variance ceiling σ^2 decreases, distinguish noise-dominated from nearly noiseless regimes, and recover the exact-gradient benchmark at $\sigma^2 = 0$. Fixing a positive noise level in an accuracy asymptotic can hide this transition. Merely removing a stochastic term from a lower bound can also leave no meaningful noiseless guarantee.

We therefore seek matching expected query bounds under unbiased, variance-bounded gradient access without potential values. Here *noise-adaptive* means dependence on the supplied variance ceiling, including zero; it does not mean that this ceiling is unknown to the algorithm. This leads to our central question:

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.

Can we obtain tight sampling complexity jointly in ε and κ, with a guarantee that also adapts to the noise level?

Our contributions. We resolve this question in the *fixed-dimensional setting*, under the oracle, localization, and expected-cost conventions of Section 2.

1. For every fixed $d \geq 1$, we establish the minimax rate $\Theta(\log(1+\kappa) + \sigma^2/(\mu\varepsilon))$, jointly in ε, κ, and all $\sigma^2 \geq 0$, with constants depending only on d (Theorems 2.2 and 2.3).
2. We construct a gradient-only sampler with fully charged geometric preprocessing and variance-controlled Poisson acceptance, including all failed histories and transcript markers (Algorithms 1–2 and Section 3).
3. We give a zero-noise exact sampler, its radius-free extension, and a fixed-curvature lower bound isolating logarithmic worst-case location cost (Section 5).

1.1 Closest related work

The bounded-variance comparison above uses Chen et al. (2026b), Theorem 3.4(i) and Remark 5, with their lower bound in Proposition 4.1 and Remark 6. Their bounds already track noise strength; neither the $1/\varepsilon$ exponent nor the rare-informative-reply mechanism is new here. Our contribution is joint fixed-dimensional optimality, without multiplicative curvature or logarithmic losses, under the conditional oracle and stopping conventions below.

For exact information, Chewi et al. (2022) obtain a $\Theta(\log\log\kappa)$ scalar benchmark with a known mode and potential values. Chewi et al. (2024) give value-and-gradient query bounds and fixed-dimensional ellipsoidal preprocessing. Their geometry also permits accuracy-independent expected value-based exact rejection, as derived in Appendix F. Our noiseless contribution is therefore gradient-only exact sampling with a coarsely localized unknown mode, not accuracy-independent expected rejection cost alone.

Poisson rejection and randomized gradient integration have precedents in the first-order framework of Chen et al. (2026a) and the exact-diffusion constructions of Beskos et al. (2006). We analyze these principles with noisy, history-dependent replies and explicitly charge geometric preprocessing and random stopping. Appendix A gives the extended comparison, including recent exact-gradient methods and the relevant oracle, initialization, and error-metric distinctions.

2 Model and main minimax results

Fix an integer $d \geq 1$ and public parameters $0 < \mu \leq L < \infty$, $\sigma^2 \geq 0$, and $0 < \varepsilon \leq 1/10$. Norms are Euclidean and logarithms are natural unless a base is displayed. With d fixed, constants hidden in O, Ω, Θ may depend on d but are uniform in the other parameters unless stated otherwise. The notation C_d denotes a finite constant depending only on d and may change between bounds. Write $\mathbb{B}(a,r) = \{x \in \mathbb{R}^d : \|x-a\| \leq r\}$ and $\mathbb{B}_d = \mathbb{B}(0,1)$ for the closed Euclidean balls.

Definition 2.1 (Potential class)**.** The class $\mathcal{C}_d(L,\mu)$ consists of all $f \in C^1(\mathbb{R}^d)$ such that

$$\frac{\mu}{2}\|x-y\|^2 \leq f(x) - f(y) - \langle \nabla f(y), x-y\rangle \leq \frac{L}{2}\|x-y\|^2 \qquad (x,y \in \mathbb{R}^d), \tag{2.1}$$

and whose unique minimizer satisfies $\|x_f^\star\| \leq \mu^{-1/2}$. The target is $\nu_f(dx) = Z_f^{-1} e^{-f(x)}dx$, where $Z_f = \int_{\mathbb{R}^d} e^{-f(x)}dx$.

Strong convexity ensures $0 < Z_f < \infty$. The curvature endpoints are public bounds, not a requirement that every function attain them; no second derivative is assumed. The ball localizes the unknown mode, not the target's support. Additive constants in f are irrelevant. For initialization extensions, $\mathcal{U}_d(L,\mu)$ denotes the same curvature class without localization, and $\mathcal{C}_d(L,\mu;x_0,R)$ restricts its modes to $\mathbb{B}(x_0,R)$.

Oracle and stopping conventions. A physical query at x_t returns a full vector g_t. Let $\mathcal{H}_t^{\mathrm{in}}$ be the complete information σ-field after x_t is selected and immediately before its reply is generated. It

contains x_t, all previous queries and replies, and all internal randomness already generated, but no future random draws. Every admissible conditional kernel satisfies, almost surely,

$$\mathbb{E}[g_t \mid \mathcal{H}_t^{\mathrm{in}}] = \nabla f(x_t), \qquad \mathbb{E}[\|g_t - \nabla f(x_t)\|^2 \mid \mathcal{H}_t^{\mathrm{in}}] \le \sigma^2. \tag{2.2}$$

This is a *total vector* variance bound; successive replies need not be independent or identically distributed. Potential values, Hessians, proximal oracles, and convex-body sampling oracles are unavailable. Exact real arithmetic and internal random generation are free.

The call count is a stopping time for the physical filtration

$$\mathcal{F}_t = \sigma(x_1, g_1, \dots, x_t, g_t), \qquad \mathcal{F}_0 = \{\varnothing, \Omega\}. \tag{2.3}$$

The output may use private randomness. Since that randomness is not itself in $\mathcal{F}_t$, we encode private branch decisions in charged query locations before stopping. This stricter convention is explicit: it accounts for the exact one-call lower bound below, and costs only a constant-factor overhead in our construction.

Remark. Conditioning on $\mathcal{H}_t^{\mathrm{in}}$ fixes the current query as well as the randomness already used to select it. This matters because a fresh random query need not be $\mathcal{F}_{t-1}$-measurable. For example, an exact oracle for $f(x) = x^2/2$, queried first at either 1 or -1 with equal probability, satisfies (2.2), whereas $\mathbb{E}[g_1 \mid \mathcal{F}_0] = 0 \neq \nabla f(x_1)$. Fresh-data oracles satisfy (2.2) when their pointwise moments obey these bounds and their new data are independent of the incoming history. The assumption also permits history-dependent reply laws; it does not require independent replies. The physical filtration $(\mathcal{F}_t)$ serves the separate stopping requirement.

Minimax quantity and normalization. Write $\|P - Q\|_{\mathrm{TV}} = \sup_D |P(D) - Q(D)|$ for any two probability distributions P and Q. The complexity is

$$N^\star_{\mathrm{TV}}(d, L, \mu, \sigma, \varepsilon) = \inf_{\mathcal{A}} \sup_{f \in \mathcal{C}_d(L,\mu),\, \mathcal{O}} \mathbb{E}_{f,\mathcal{O},\mathcal{A}}[T], \tag{2.4}$$

where the infimum is over measurable, almost surely terminating algorithms obeying the information and stopping rules, with TV error at most ε for every f and every oracle $\mathcal{O}$ satisfying (2.2). An empty infimum is $+\infty$. Normalize without changing query count by

$$z = \sqrt{\mu}\, x, \quad F(z) = f(z/\sqrt{\mu}), \quad G = g/\sqrt{\mu}, \quad \kappa = L/\mu, \quad A = \sigma^2/\mu. \tag{2.5}$$

Then F has curvature bounds $1, \kappa$, its mode $x_\star = \sqrt{\mu} x_f^\star$ is in $\mathbb{B}(0,1)$, and normalized noise has total variance at most A. Inverse scaling preserves TV and physical stopping. Throughout, put

$$X = A/\varepsilon, \qquad S = 1 + \log(1+\kappa) + X. \tag{2.6}$$

2.1 Sampling procedure and guarantees

In this section, we describe the algorithms achieving the required upper bound, including Algorithms 1 and 2. A call $\mathsf{Oracle}(y)$ queries $y/\sqrt{\mu}$ and divides the reply by $\sqrt{\mu}$. All calls are sequential and charged, including repeated locations. Fresh random draws are generated only when needed, after preceding replies. The algorithms mainly consist of several ingredients below:

Geometric preprocessing, step 1: $\mathsf{CenterPilot}(\kappa, A, B, \delta_m)$**.** Given the center tolerance B and failure budget δ_m, initialize $E = c_E + Q_E \mathbb{B}_d = 2\mathbb{B}_d$. At each round inspect the symmetric vertices $z_i^\pm = c_E \pm \gamma Q_E e_i$, where $\gamma > 0$ is a fixed geometric parameter depending only on d. A vertex outside $2\mathbb{B}_d$ gives the separating normal $v = z_i^\pm$ without an oracle call. Otherwise form fresh means $\bar{g}_{i,b}^\pm = b^{-1} \sum_{\ell=1}^b \mathsf{Oracle}(z_i^\pm)$ at all $2d$ vertices and the scores

$$\widehat{s}_{i,b}^\pm = \langle \bar{g}_{i,b}^\pm, z_i^\pm - c_E \rangle, \qquad \widehat{M}_b = \max\{1, \max_{i,\pm} \widehat{s}_{i,b}^\pm / m_0\}.$$

Here $m_0 > 0$ is a score floor determined by B and d. Double the fresh batch size $b = 1, 2, 4, \dots$ until $b\widehat{M}_b \ge \Lambda$, for a variance–confidence threshold Λ. If every returned score is at most $2m_0$, return $m = c_E$. Otherwise use the mean at a maximizing vertex as a cut normal v. For either type of cut, replace E by the prescribed enclosing ellipsoid of $E \cap \{y : \langle v, y - c_E \rangle \le 2\gamma \|Q_E^\top v\|\}$ and repeat. Fixed tie-breaking makes the choice unambiguous. Return $m = 0$ if the next full batch would exceed the call cap or the update cap is reached. Symmetry lets small scores certify the center's potential gap, while a large score selects a region to discard. Thus this step seeks only $F(m) - F(x_\star) \le B$ with failure probability at most δ_m, not optimization to accuracy ε.

Geometric preprocessing, step 2: $\mathsf{RoundingPilot}(m, \kappa, A, \delta_q)$**.** Starting from $E = c_E + Q_E\mathbb{B}_d = \mathbb{B}(m, 2)$, inspect the same vertices. A vertex outside $\mathbb{B}(m, 2)$ gives the free cut normal $z_i^{\pm} - m$. For the remaining rounds, query instead the midpoints $y_i^{\pm} = (m + z_i^{\pm})/2$ and form scores $\widehat{s}_{i,b}^{\pm} = \langle \bar{g}_{i,b}^{\pm}, y_i^{\pm} - m\rangle$. Double fresh batches until $b \max\{1, \max_{i,\pm} \widehat{s}_{i,b}^{\pm}\} \geq \lambda$. If all scores meet the prescribed small-score threshold, contract the inscribed ellipsoid $(m + c_E)/2 + [\gamma/(2\sqrt{d})]Q_E\mathbb{B}_d$ toward m by the prescribed factor. Return its center $z \in \mathbb{R}^d$ and invertible shape matrix $\mathsf{T} \in \mathbb{R}^{d\times d}$. Otherwise, use a maximizing midpoint's mean as the shallow-cut normal at its original vertex and repeat. Before a call exceeding the global budget, or at the round cap, return $(z, \mathsf{T}) = (m, I_d)$. The midpoint test is designed to enclose a whole level set, not just find a low-value point. When the center and rounding succeed, the returned ellipsoid satisfies

$$z + \mathsf{T}\mathbb{B}_d \subseteq K_1 := \{x : F(x) - F(m) \leq 1\} \subseteq z + \rho\mathsf{T}\mathbb{B}_d.$$

Thus z locates the ellipsoid and T determines its shape; z need not equal m. The failure budget δ_q controls this sandwich conditional on a valid center; no potential value is tested.

Proposal generation: $\mathsf{RadialProposal}(z, \mathsf{T}, \rho)$**.** From the returned geometry define

$$U(x) = \tfrac{1}{4}\big(\|\mathsf{T}^{-1}(x - z)\|/\rho - 1\big)_+, \qquad q(x) \propto e^{-U(x)}.$$

Draw a uniform unit direction Θ and, independently, a radius R with density proportional to $v^{d-1}e^{-(v-1)_+/4}$ for $v > 0$; return $x = z + \rho\mathsf{T}R\Theta$. A finite mixture of a unit-ball radial law and shifted gamma laws generates R without oracle calls. The flat ellipsoidal region adapts the proposal to the level set, and its tails give full support rather than discarding target mass.

Directional envelopes: $\mathsf{TwoGrids}(m, s, r, \kappa, A, \delta)$**.** For a proposal $x \neq m$, set $r = \|x - m\|$ and $s = (x - m)/r$. Write $w(t) = \langle s, \nabla F(m + st)\rangle$ only to describe the quantities being estimated. Starting at $t_0 = 0$, use $t_j = r(1 - 2^{-j})$ for $1 \leq j \leq J$, where $J \geq 0$ is the smallest integer with $r2^{-J} \leq \kappa^{-1/2}$, and add the terminal endpoint r. The inputs determine batch sizes $\beta_j \geq 1$ and paddings $e_j > 0$; form the signed averages

$$\widehat{w}_{+,j} = \left\langle s, \frac{1}{\beta_j}\sum_{i=1}^{\beta_j} \mathsf{Oracle}(m + st_j)\right\rangle, \qquad \widehat{w}_{-,j} = \left\langle -s, \frac{1}{\beta_j}\sum_{i=1}^{\beta_j} \mathsf{Oracle}(m + s(r - t_j))\right\rangle.$$

Use separate fresh batches: complete the positive grid and its terminal node before the reflected negative grid. On the interval ending at t_j, place height $\max\{0, \widehat{w}_{\pm,j} + e_j\}$, with the terminal batch and padding treated likewise. At a shared boundary use the larger adjacent height, and at an endpoint its one-sided value. Reflect the negative rectangles back to $[0, r]$ and return g_+, g_- and their areas $Z_g = \int_0^r (g_+ + g_-)\,dt$, $Z_- = \int_0^r g_-\,dt$. Padding protects against underestimation; reflection makes both signed parts nondecreasing before rectangles are formed. With the calibrated schedules, the envelopes dominate the corresponding derivative parts except on an event of conditional probability at most δ.

Trial control: $\mathsf{FiniteCategory}(Q, H)$ **and** $\mathsf{Position}(a_0/Q)$**.** Given the grids, a margin $c > 0$, an offset $b_{\mathrm{acc}} > 0$, and an integer cap $H \geq 1$, set $a_0 = g_+ + g_- + 2c/r$, $Q = Z_g + 2c$, and $v = \min\{1, e^{Z_- + c - b_{\mathrm{acc}}}\}$. A failed Bernoulli$(v)$ coin rejects the proposal immediately. Otherwise $\mathsf{FiniteCategory}$ returns N_c by assigning Poisson(Q) masses to $0, \ldots, H$ and the entire tail mass to $H + 1$. Overflow rejects; zero accepts without nodes. For each remaining node, $\mathsf{Position}$ selects the uniform background with weight $2c/Q$ or an envelope rectangle with weight its area divided by Q, then draws uniformly in the selected interval. It returns $t \in [0, r]$ with density a_0/Q; this density remains fixed throughout the trial. The positive background prevents a zero denominator, and the overflow category keeps the number of processed nodes finite.

Given a correction batch size $n \geq 1$, average n new directional gradients at $m + st$ as $\bar{w}$ and draw a rejection bit with probability $\mathrm{clip}_{[0,1]}((\bar{w} + g_-(t) - u'(t) + c/r)/a_0(t))$, where $u(t) = U(m + st)$ and $\mathrm{clip}_{[0,1]}(a) = \min\{1, \max\{0, a\}\}$. Compute u' from U outside its flat region and use zero on that region and its boundary. Compensation by g_- offsets negative directional values, subtraction of u' corrects the proposal, and clipping makes the probability legal on every history. Reject at the first marked node; accept after N_c unmarked nodes. If $x = m$, bypass all segment operations and use the coin $\min\{1, e^{U(m) - b_{\mathrm{acc}}}\}$ instead.

Recording and returning the result. $\mathsf{Marker}(y)$ makes a charged oracle call at y in normalized coordinates, waits, and discards the reply for estimation while retaining it in the history. The markers record the proposal, category, and decisions, so stopping is visible in the physical transcript. Bernoulli and category selections use fresh open-interval uniforms. $\mathsf{OneTrial}$ returns (a, x) with acceptance bit a; Algorithm 1 returns the first accepted $x/\sqrt{\mu}$ within K trials, or $m/\sqrt{\mu}$ after all trials reject.

Algorithm 1 Gradient-only sampler with finite caps

Require: Public $d, L, \mu, \sigma^2, \varepsilon$ and the stochastic gradient oracle
Require: Center tolerance $0 < B \le 1/16$, expansion $\rho \ge 1$, margin $c > 0$, offset $b_{\mathrm{acc}} > 0$
Require: Failure budgets $\delta_m, \delta_q, \delta \in (0,1)$; batch size $n \ge 1$ and caps $K, H \ge 1$ (integers)
Ensure: A sample Y in the original coordinates

$\kappa \leftarrow L/\mu$, $A \leftarrow \sigma^2/\mu$; form Oracle as above.
$m \leftarrow \textsc{CenterPilot}(\kappa, A, B, \delta_m)$
$(z, \mathsf{T}) \leftarrow \textsc{RoundingPilot}(m, \kappa, A, \delta_q)$
Form the displayed U, q; fix $\Pi = (m, z, \mathsf{T}, q, U)$.
for $k = 1, \dots, K$ **do**
  $(a, x) \leftarrow \textsc{OneTrial}(\Pi)$ ▷ Shared parameters remain fixed
  if $a = 1$ **then**
    return $Y = x/\sqrt{\mu}$
  end if
end for
return $Y = m/\sqrt{\mu}$ ▷ After the final marker of trial K

Algorithm 2 One trial, including every transcript marker

Require: Shared $\rho, \kappa, A, \delta, c, b_{\mathrm{acc}}, n, H$ from Algorithm 1

function $\textsc{OneTrial}(\Pi = (m, z, \mathsf{T}, q, U))$
  $x \leftarrow \textsc{RadialProposal}(z, \mathsf{T}, \rho)$; $\textsc{Marker}(x)$ ▷ Reveal the proposal; wait for the reply
  $r \leftarrow \|x - m\|$; $a \leftarrow 0$
  if $r = 0$ **then**
    $a \leftarrow \mathrm{Bernoulli}(\min\{1, e^{U(m) - b_{\mathrm{acc}}}\})$
  else
    $s \leftarrow (x - m)/r$
    $(g_+, g_-, Z_g, Z_-) \leftarrow \textsc{TwoGrids}(m, s, r, \kappa, A, \delta)$
    $a_0(t) \leftarrow g_+(t) + g_-(t) + 2c/r$; $Q \leftarrow Z_g + 2c$
    $v \leftarrow \min\{1, e^{Z_- + c - b_{\mathrm{acc}}}\}$; $b_{\mathrm{pre}} \leftarrow \mathrm{Bernoulli}(v)$
    if $b_{\mathrm{pre}} = 0$ **then**
      $\textsc{Marker}(0e_1)$ ▷ Control marker: failed prefactor
    else
      $N_c \leftarrow \textsc{FiniteCategory}(Q, H)$ ▷ $H + 1$ denotes overflow
      $\textsc{Marker}((N_c + 1)e_1)$ ▷ Control marker: passed prefactor and category
      if $N_c \le H$ **then**
        $a \leftarrow 1$ ▷ Category 0 has an empty node loop
        for $j = 1, \dots, N_c$ **do**
          $t \leftarrow \textsc{Position}(a_0/Q)$
          $\overline{w} \leftarrow \big\langle s, \frac{1}{n}\sum_{i=1}^{n} \mathsf{Oracle}(m + st) \big\rangle$ ▷ Exactly n fresh calls
          $b \leftarrow \mathrm{Bernoulli}\Big(\mathrm{clip}_{[0,1]}\ \frac{\overline{w} + g_-(t) - u'(t) + c/r}{a_0(t)}\Big)$
          $\textsc{Marker}(be_1)$ ▷ Node mark bit; wait for its reply
          if $b = 1$ **then**
            $a \leftarrow 0$; **break**
          end if
        end for
      else
        $a \leftarrow 0$ ▷ Overflow: no correction calls
      end if
    end if
  end if
  $\textsc{Marker}(ae_1)$ ▷ Final acceptance marker; wait for its reply
  return (a, x)
end function

Stage order distinguishes coincident marker and data-query locations. Complete implementation is shown in Algorithms 1 and 2. The detailed setting of parameters in the two algorithms are all stated in the proof of Theorem 2.2.

Theorem 2.2 (Constructive upper bound)**.** *For every fixed $d \geq 1$, a finite C_d exists such that, for every public tuple above, Algorithms 1 and 2 with the public parameter choices specified in the proof return Y with*

$$\|\mathcal{L}(Y) - \nu_f\|_{\mathrm{TV}} \leq 7\varepsilon/16, \qquad \sup_{f,\mathcal{O}} \mathbb{E}[T] \leq C_d S. \tag{2.7}$$

The supremum includes all admissible conditional oracles. The algorithm terminates on every finite-real-reply path with its specified open-interval uniforms, and T is a physical-transcript stopping time. In particular, $N^{\star}_{\mathrm{TV}} \leq C_d S$, including when $\sigma^2 = 0$.

Theorem 2.3 (Minimax lower bound)**.** *For every fixed $d \geq 1$ and every public tuple above,*

$$N^{\star}_{\mathrm{TV}}(d, L, \mu, \sigma, \varepsilon) \geq \max\{1, X/24, S/62\}. \tag{2.8}$$

The bound permits arbitrary adaptive full-vector queries and random stopping under the worst-case expected physical-query convention.

Together these theorems imply

$$N^{\star}_{\mathrm{TV}}(d, L, \mu, \sigma, \varepsilon) = \Theta\left(\log(1+\kappa) + \frac{\sigma^2}{\mu\varepsilon}\right). \tag{2.9}$$

The constants are jointly uniform in κ, $\sigma^2 \geq 0$, and $0 < \varepsilon \leq 1/10$. The ratio of the displayed upper and lower benchmarks is at most $62C_d$. We prove the theorems in Appendices C and D; the next two sections explain their mechanisms. These are information bounds in fixed dimension, not polynomial-time implementations or dimension-uniform optimal rates.

3 Upper bound: geometric amortization and variance-controlled acceptance

We sketch Theorem 2.2 in this section. The above Algorithms 1 and 2 must avoid multiplying geometric search by statistical precision, control acceptance error without potential values, and charge all failed histories. The next three subsections address these requirements.

3.1 Constant-energy geometry at additive cost

Upper bound: idea. The pilots need only constant energy accuracy. A high-score round can therefore use a less accurate gradient than a low-score round. What matters is the sum of reciprocal energy scores, not a uniform accuracy requirement at every geometric cut. Let $\mathcal{M} \geq 1$ denote the normalized true score: $\widetilde{M}$ for centering and M for rounding in Appendix C. With failure budget δ, a successful round uses $O(1+A/(\delta\mathcal{M}))$ calls and incurs failure probability $O(\delta/\mathcal{M})$ (Propositions C.2 and C.7).

Proof sketch. A valid shallow cut contracts volume by a dimension-dependent constant. At a round with score $\mathcal{M}$, convexity bounds the current ellipsoid's volume by $C_d\mathcal{M}^d$ times that of the protected sublevel set. After the first visit to $[2^j, 2^{j+1})$, only $O(j+1)$ further valid cuts can preserve that set, regardless of subsequent score increases. Hence

$$\sum_{\text{valid rounds}} \frac{1}{\mathcal{M}} \leq C_d \sum_{j \geq 0} (j+1) 2^{-j} < \infty. \tag{3.1}$$

The total number of cuts is $O(\log(1+\kappa))$, so summing costs gives $O(\log(1+\kappa) + A/\delta)$. The batching test does not know the true score: an early stop requires a large score error, while inaccurate later stages have summable dyadic probabilities. Martingale orthogonality gives mean-square error at most A/b for a fresh mean of b replies. The early-stop probability bound grows with b, whereas the late-error bound decays as $1/b$. Summing over the respective dyadic ranges avoids a logarithmic union penalty. First-failure events are summed only along good prefixes, never by conditioning on future success. Public caps extend the cost bound to failed histories.

Both pilots succeed with probability at least $1 - \varepsilon/4$. For their proposal, put $I(x) = F(x) - F(m)$. Proposition C.11 gives

$$\alpha_*(x) = e^{-I(x)+U(x)-b_{\mathrm{acc}}} \le 1, \qquad \int q(x)\alpha_*(x)\,dx \ge p_0 > 0, \tag{3.2}$$

where p_0 depends only on d. The inner ellipsoid supplies target mass; convexity outside the outer ellipsoid controls the proposal's tails. Since $q\alpha_*$ is proportional to e^{-F}, this is the desired ideal rejection law. The same construction also bounds proposal moments on failed pilot histories, which will matter when counting work.

3.2 A COMPENSATED POISSON EXPERIMENT INSTEAD OF EXPONENTIATING AN ESTIMATE

Upper bound: idea. Exponentiating an unbiased estimate of I introduces bias. Instead, the absence of Poisson marks represents the exponential exactly; random positions belong to the experiment rather than approximating an integral. The negative derivative is compensated to keep the marking intensity nonnegative.

Proof sketch. On a good pilot, the directional derivative w introduced above satisfies $\int_0^r |w(t)|\,dt \le I(x) + 2B$ and $r^2 \le 4(I(x) + 2B)$. Right-endpoint rectangles on a dyadic grid are charged to neighboring larger-value intervals, with only a constant terminal remainder. The padding adds a constant mass and costs $O(1 + \log(1 + \kappa r^2) + Ar^2/\delta)$ calls on every history (Lemma C.14). This controls the positive and negative envelope masses separately, rather than paying for the largest endpoint slope along the whole segment.

For the computable a_0, Q above, define only for analysis $b_0 = w + g_- - u' + c/r$ and $\ell = \int_0^r b_0(t)\,dt$. The proposal is flat while the target may decrease; after the ray leaves that region its target derivative dominates the proposal derivative (Lemma C.15). On good pilot and envelope histories this yields

$$c/r \le b_0(t) \le a_0(t) - c/r, \qquad \ell \ge c, \quad r^2 \le 8\ell, \quad Q \le 68\ell. \tag{3.3}$$

Ideal marks have mean probability $p = \ell/Q \ge 1/68$. For $N \sim \mathrm{Poisson}(Q)$ independent ideal positions and marks give no-mark probability $\mathbb{E}[(1-p)^N] = e^{-\ell}$. Since $\ell = I + Z_- - U + c$, the prefactor cancels the compensation:

$$ve^{-\ell} = e^{-I+U-b_{\mathrm{acc}}} = \alpha_*(x). \tag{3.4}$$

Here clipping of v is inactive on the good event.

For the actual noisy marks, a two-sided margin bounds clipping bias by variance divided by margin (Lemma C.17). After integrating over the position density, the conditional mark error is at most $\eta_0 = Ar^2/(ncQ)$, with $\eta_0/p \le 32A/n$. Comparing survival recursively with $(1-p)^j$ gives total error at most $\eta_0 \sum_{j\ge 0}(1-p)^j = \eta_0/p$, not a factor proportional to Q. This uses conditional bounds at reached histories, not independent oracle replies. Thus $n = O(1 + A/\varepsilon)$ controls noise-induced bias, and this bias vanishes at $A = 0$. The survival estimate also controls work: even if Q is large, nodes after the first mark are never queried. The expected number actually processed is bounded by a geometric sum, uniformly in the generated Poisson count. No moment bound on Q is required. Finite caps incur separate errors.

3.3 ACCURACY, FAILED HISTORIES, AND PHYSICAL STOPPING

The calibrated tolerances and caps make each reached good-pilot trial's acceptance function differ from α_* by at most τ for proposal-almost every point. Normalization contributes at most $\tau/(p_0 - \tau) \le \varepsilon/8$ in TV. Mixing over the first acceptance preserves this bound; the outer-cap fallback contributes $\varepsilon/16$ and pilot failure $\varepsilon/4$, giving $7\varepsilon/16$.

On good pilots the expected number of trials is $O(1)$, and on good envelopes the expected number of examined nodes is $O(1)$. Failed envelopes cost at most H, offset by their small probability. Failed pilots cost at most $KH(n+1)$ in correction calls and their markers; $\varepsilon[1 + \log(1/\varepsilon)]^2$ is uniformly bounded, so this term also has expected order $1 + A/\varepsilon$. All-history proposal moments control grid costs. Including every marker gives $\mathbb{E}[T] \le C_d S$, not a cost conditioned on success. The marker sequence makes each stopping decision measurable in the physical transcript; ignored replies remain in all later incoming histories (Section C.4).

4 LOWER BOUND: NOISE AND CURVATURE ENCODE DIFFERENT INFORMATION

We sketch Theorem 2.3. Each scalar witness h is embedded as $h(x_1) + \frac{1}{2}\|x_{2:d}\|^2$ and supplied with exact gradients in the added coordinates. Total noise variance is unchanged and projecting the target onto its first coordinate preserves the scalar law. The couplings and decision trees below are applied directly to full-vector algorithms; no projected stopping-filtration assumption is needed.

4.1 RARE INFORMATIVE REPLIES FORCE THE NOISE TERM

Hard instance: idea. A small linear tilt changes the target law while producing the same gradient gap at every query point. Hiding that tilt behind a rare oracle innovation prevents a remote or adaptive query from amplifying the signal. Accuracy then requires enough physical queries to encounter informative replies with appreciable probability.

Proof sketch. Let $\tau_0 = 4\varepsilon$ and $H_{\mathrm{a}}(u) = u^2/2 + (\kappa - 1)(|u| - 8)_+^2/2$. The two potentials $F_\pm = H_{\mathrm{a}} \pm \tau_0 u$ attain the curvature endpoints and have modes $\mp\tau_0 \in [-1, 1]$. Their normalized target laws are separated by more than $8\varepsilon/3$ in TV. For $A > 0$, let $p_{\mathrm{r}} = \tau_0^2/(A + \tau_0^2)$ and return

$$G_\pm(u) = H'_{\mathrm{a}}(u) \pm \frac{\tau_0}{p_{\mathrm{r}}}\xi, \qquad \xi \sim \mathrm{Bernoulli}(p_{\mathrm{r}})$$

with fresh innovations. These replies are unbiased and have variance exactly A. Couple both runs with the same private tape and innovations. Their transcripts and outputs agree until a queried innovation is one. Nonnegative summation at a random stopping time gives $\mathbb{E}\big[\sum_{j\le T}\xi_j\big] = p_{\mathrm{r}}\mathbb{E}[T]$. Hence their output-law separation is at most $p_{\mathrm{r}}\mathbb{E}[T]$ in each marginal. The target separation and two accuracy errors imply $p_{\mathrm{r}}\mathbb{E}[T] > 2\varepsilon/3$, and therefore $\mathbb{E}[T] \ge A/(24\varepsilon)$ (Proposition D.1). This does not impose a fixed query horizon or a positive lower bound on A.

Where the $1/\varepsilon$ scale comes from. The target perturbation is of size $\tau_0 \asymp \varepsilon$, so its gradient gap is also $O(\varepsilon)$. Under a variance ceiling A, that gap can be hidden in innovations of probability on the order of $\varepsilon^2/(A + \varepsilon^2)$. To separate the output laws by order ε, the coupled executions must split with probability of that order. Dividing the required splitting probability by the innovation rate gives A/ε, rather than A/ε^2. The exact calculation retains the additive ε term and works for every positive A, even when $A \ll \varepsilon^2$. The distant curvature anchor ensures that changing κ does not alter this noise argument or let far-away queries amplify the tilt.

4.2 A FINITE-REPLY FAMILY FORCES LOGARITHMIC CURVATURE COST

Hard instance: idea. Exact real replies can normally carry unlimited information, so a packing of modes alone is insufficient. We choose a family for which any fixed query has at most three possible gradient replies. A sample nevertheless identifies one of $\Theta(\sqrt{\kappa})$ separated mass concentrations. Only logarithmically many ternary decisions can reveal that index.

Proof sketch. For $s_0 = \sqrt{\kappa} \ge 640$, take $M = \lfloor s_0/80 \rfloor$ disjoint narrow ramps of width $s_0/(s_0^2 - 1)$ with thresholds $\theta_j = -1/4 + 40(j-1)/s_0$. The derivative of the jth potential equals $u - s_0/2$ to the left, $u + s_0/2$ to the right, and has slope κ on its ramp. Its slope elsewhere is one. Its mode lies in $(-1/4, 1/4)$, and its target gives mass greater than $99/100$ to an interval D_j disjoint from all the other decoding intervals.

At any real query, at most one ramp can be active: every remaining reply is one of the two common outer values. Fixing the private tape therefore produces a ternary transcript tree. Under a uniform family index, TV accuracy gives decoding probability greater than $89/100$. Truncating at $h = \lceil 4\bar{n} \rceil$, where $\bar{n}$ is the prior-average expected query count, loses at most $1/4$ by Markov's inequality. At most 3^h leaves remain, and each fixed output decodes at most one index. Thus $3^h/M > 16/25$, which yields

$$N^\star_{\mathrm{TV}} \ge \tfrac{1}{4}\big[\log_3(16M/25) - 1\big] = \Omega(\log \kappa)$$

in the large-curvature regime (Proposition D.2). Exact gradients are admissible for every advertised variance ceiling.

Two separated Gaussian targets additionally force at least one physical query, since the zero-call event belongs to the trivial $\mathcal{F}_0$ (Proposition D.3). Combining the three inequalities, using the constant bound when $\kappa < 640^2$, gives $N^\star_{\mathrm{TV}} \geq S/62$. This combines worst-case lower bounds, not costs on a single instance; Appendix D.4 records the constants.

5 Exact sampling and the role of initialization

We first examine zero noise and then the quadratic endpoint $L = \mu$. Appendix C.4 gives both arguments.

Corollary 5.1 (Exact gradient-only sampling). *For every fixed $d \geq 1$, if $\sigma^2 = 0$ and $f \in \mathcal{C}_d(L, \mu)$, an accuracy-free gradient-only sampler returns an exact sample from ν_f, terminates almost surely, and satisfies $\sup_f \mathbb{E}[T] \leq C_d \log(1+\kappa)$ with physical-transcript stopping.*

The exact sampler is uncapped; it is not obtained by setting $\varepsilon = 0$ in the finite theorem. The quadratic case admits a sharper bound.

Corollary 5.2 (Quadratic endpoint). *When $\kappa = 1$, with $X = \sigma^2/(\mu\varepsilon)$,*

$$\max\{1, X/24\} \leq N^\star_{\mathrm{TV}} \leq \max\{1, \lceil X/2 \rceil\} \leq 1 + X/2.$$

At $\sigma^2 = 0$, exactly one gradient query suffices for exact sampling.

We next allow an arbitrary supplied ball, then remove the radius input at zero noise; Appendix E gives both reductions.

Corollary 5.3 (An arbitrary supplied ball). *For $f \in \mathcal{C}_d(L, \mu; x_0, R)$ with x_0, R supplied, let $\Gamma_R = \max\{1, \mu R^2\}$. There is an ε-TV sampler with*

$$\sup_{f,\mathcal{O}} \mathbb{E}[T] \leq C_d \left[\log(1 + \kappa\Gamma_R) + \frac{\Gamma_R \sigma^2}{\mu\varepsilon} \right]. \tag{5.1}$$

At zero noise an almost surely terminating exact sampler satisfies

$$\sup_f \mathbb{E}[T] \leq C_d \log(1 + \kappa\Gamma_R) \leq C'_d [\log(1+\kappa) + \log(1 + \sqrt{\mu} R)]. \tag{5.2}$$

Both algorithms use physical-transcript stopping.

No joint radius–noise optimality is claimed. With exact gradients, a supplied radius is unnecessary for an instance-wise guarantee.

Corollary 5.4 (No supplied radius at zero noise). *For fixed x_0 and every $f \in \mathcal{U}_d(L, \mu)$, exact gradients suffice for exact sampling, without a radius or accuracy input, at instance-wise expected cost*

$$\mathbb{E}_f[T] \leq C_d \big[\log(1+\kappa) + \log\big(1 + \sqrt{\mu}\|x_0 - x^\star_f\|\big) \big]. \tag{5.3}$$

The sampler terminates almost surely and its query count is a physical stopping time.

Finite instance-wise cost need not yield a finite uniform bound over locations. A fixed-curvature family isolates this obstruction.

Proposition 5.5 (Fixed-curvature location obstruction). *At $\mu = 1, L = 2$, every $R \geq 64$, fixed $d \geq 1$, and $0 < \varepsilon \leq 1/10$, the noiseless expected minimax cost on $\mathcal{C}_d(2, 1; 0, R)$ is $\Theta(\log(1+R))$. On $\mathcal{U}_d(2, 1)$ the uniformly accurate minimax expected cost is infinite, even at fixed positive accuracy.*

Appendix G proves this worst-case-over-locations obstruction; it is not per-instance and does not apply when $L = \mu$.

6 Conclusion

We characterized fixed-dimensional gradient-only sampling by the joint expected query bound $\Theta(\log(1+\kappa) + \sigma^2/(\mu\varepsilon))$. At zero noise, an uncapped variant gives exact sampling with almost-sure termination. The bounds allow dimension-dependent constants and do not assert an accuracy-independent deterministic cap. Sharper radius–noise dependence and finite-precision implementations remain natural directions.

AI USE STATEMENT

In this work, we use generative AI tools to propose the algorithms of upper bound and complete its proof. Additionally, we guide gererative AI to establish lower bound via more than 10 rounds interaction with AI.

REPRODUCIBILITY STATEMENT

All assumptions and theoretical claims are stated in the paper, with complete proofs provided in the appendix.

REFERENCES

Jason M. Altschuler and Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Journal of the ACM*, 71(3):24:1–24:55, 2024. doi: 10.1145/3653446. URL https://arxiv.org/abs/2302.10249.

Christophe Andrieu and Gareth O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009. doi: 10.1214/07-AOS574. URL https://arxiv.org/abs/0903.5480.

Alexandros Beskos, Omiros Papaspiliopoulos, and Gareth O. Roberts. Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli*, 12(6):1077–1098, 2006. doi: 10.3150/bj/1165269151. URL https://discovery.ucl.ac.uk/10050411/1/Beskos_PapaspiliopoulosRoberts06.pdf.

Niladri Chatterji, Nicolas Flammarion, Yian Ma, Peter Bartlett, and Michael Jordan. On the theory of variance reduction for stochastic gradient Monte Carlo. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 764–773. PMLR, 2018. URL https://proceedings.mlr.press/v80/chatterji18a.html.

Niladri S. Chatterji, Peter L. Bartlett, and Philip M. Long. Oracle lower bounds for stochastic gradient sampling algorithms. *Bernoulli*, 28(2):1074–1092, 2022. doi: 10.3150/21-BEJ1377. URL https://arxiv.org/abs/2002.00291v3. Journal Theorem 4.1 corresponds to Theorem 1 in arXiv v3.

Fan Chen, Sinho Chewi, Constantinos Daskalakis, and Alexander Rakhlin. High-accuracy sampling for diffusion models and log-concave distributions. In *Proceedings of the 43rd International Conference on Machine Learning*, 2026a. URL https://arxiv.org/abs/2602.01338v2. Extended version: arXiv:2602.01338v2, April 27, 2026.

Fan Chen, Sinho Chewi, Constantinos Daskalakis, and Alexander Rakhlin. High-accuracy log-concave sampling with stochastic queries. In *Proceedings of the Thirty Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pp. 1341–1372. PMLR, 2026b. URL https://proceedings.mlr.press/v336/chen26g.html. Extended version: arXiv:2602.14342v2, May 14, 2026.

Fan Chen, Sinho Chewi, Jianfeng Lu, and Matthew S. Zhang. Accelerated high-accuracy sampling from a warm start via the Proximal Bouncy Particle Sampler, 2026c. URL https://arxiv.org/abs/2609.06905. Version 1, September 7, 2026.

Fan Chen, Sinho Chewi, Jianfeng Lu, and Matthew S. Zhang. Smoothed Picard Hamiltonian Monte Carlo, 2026d. URL https://arxiv.org/abs/2609.06906. Version 1, September 7, 2026.

Fan Chen, Sinho Chewi, Alexander Rakhlin, and Matthew S. Zhang. Exact simulation of diffusions and improved algorithms for log-concave sampling. arXiv:2608.05022, 2026e. URL https://arxiv.org/abs/2608.05022. Version 1, August 5, 2026.

Yongxin Chen, Sinho Chewi, Adil Salim, and Andre Wibisono. Improved analysis for a proximal algorithm for sampling. In *Proceedings of the Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 2984–3014. PMLR, 2022. URL https://proceedings.mlr.press/v178/chen22c.html.

Sinho Chewi, Patrik R. Gerber, Chen Lu, Thibaut Le Gouic, and Philippe Rigollet. The query complexity of sampling from strongly log-concave distributions in one dimension. In *Proceedings of the Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pp. 2041–2059. PMLR, 2022. URL https://proceedings.mlr.press/v178/chewi22b.html.

Sinho Chewi, Jaume de Dios Pont, Jerry Li, Chen Lu, and Shyam Narayanan. Query lower bounds for log-concave sampling. *Journal of the ACM*, 71(4):29:1–29:42, 2024. doi: 10.1145/3673651. URL https://arxiv.org/abs/2304.02599. Numbered results cited from arXiv:2304.02599v2 (October 30, 2023). Preliminary version in FOCS 2023, pp. 2139–2148.

Arnak S. Dalalyan. Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(3): 651–676, 2017. doi: 10.1111/rssb.12183. URL https://arxiv.org/abs/1412.7392.

Arnak S. Dalalyan and Avetik Karagulyan. User-friendly guarantees for the Langevin Monte Carlo with inaccurate gradient. *Stochastic Processes and their Applications*, 129(12):5278–5311, 2019. doi: 10.1016/j.spa.2019.02.016. URL https://arxiv.org/abs/1710.00095.

Alain Durmus and Éric Moulines. Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *The Annals of Applied Probability*, 27(3):1551–1587, 2017. doi: 10.1214/16-AAP1238. URL https://arxiv.org/abs/1507.05021.

Alain Durmus, Szymon Majewski, and Błażej Miasojedow. Analysis of Langevin Monte Carlo via convex optimization. *Journal of Machine Learning Research*, 20(73):1–46, 2019. URL https://jmlr.org/papers/v20/18-173.html.

Jiaojiao Fan, Bo Yuan, and Yongxin Chen. Improved dimension dependence of a proximal algorithm for sampling. In *Proceedings of the Thirty Sixth Conference on Learning Theory*, volume 195 of *Proceedings of Machine Learning Research*, pp. 1473–1521. PMLR, 2023. URL https://proceedings.mlr.press/v195/fan23a.html.

W. R. Gilks and P. Wild. Adaptive rejection sampling for Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 41(2):337–348, 1992. doi: 10.2307/2347565. URL https://rss.onlinelibrary.wiley.com/doi/10.2307/2347565.

M. S. Keane and G. L. O'Brien. A Bernoulli factory. *ACM Transactions on Modeling and Computer Simulation*, 4(2):213–219, 1994. doi: 10.1145/175007.175019. URL https://doi.org/10.1145/175007.175019.

Yin Tat Lee, Ruoqi Shen, and Kevin Tian. Structured logconcave sampling with a restricted Gaussian oracle. In *Proceedings of the Thirty Fourth Conference on Learning Theory*, volume 134 of *Proceedings of Machine Learning Research*, pp. 2993–3050. PMLR, 2021. URL https://proceedings.mlr.press/v134/lee21a.html.

László Lovász and Santosh Vempala. The geometry of logconcave functions and sampling algorithms. *Random Structures & Algorithms*, 30(3):307–358, 2007. doi: 10.1002/rsa.20135. URL https://doi.org/10.1002/rsa.20135.

Şerban Nacu and Yuval Peres. Fast simulation of new coins from old. *The Annals of Applied Probability*, 15(1A):93–115, 2005. doi: 10.1214/105051604000000549. URL https://arxiv.org/abs/math/0309222.

Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient Langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning*, pp. 681–688, 2011. URL https://icml.cc/2011/papers/398_icmlpaper.pdf.

A Extended related work

The complexity of sampling depends on the information supplied by the oracle, the approximation metric, the initialization, and the convention for counting queries. These distinctions are essential when dimension is fixed and other parameters vary. A value oracle can evaluate an acceptance ratio directly, a known minimizer removes a localization problem, and a fixed query budget is different from a worst-case expected stopping cost. We organize the comparison around these distinctions. Throughout this discussion, $\kappa = L/\mu$ and $A = \sigma^2/\mu$; logarithmic factors are suppressed only when reporting a bound from the literature using $\widetilde{O}$ notation.

The geometric approach of Lovász & Vempala (2007) analyzes ball-walk and hit-and-run sampling for log-concave densities without local smoothness assumptions. Its information model supplies density values up to a common multiplicative constant, together with geometric bounds, and its analysis includes preprocessing toward isotropic position. This is foundational background for using geometry to construct efficient sampling procedures. Direct density evaluation and geometric initialization differ from the noisy-gradient information available here; their cost guarantees therefore cannot be transferred merely by fixing the ambient dimension.

The closest stochastic-oracle result. Chen et al. (2026b) already establish the inverse-accuracy exponent under bounded variance. Their Theorem 3.4(i) and Remark 5 give the smooth log-Sobolev upper bound, specializing under strong convexity to

$$\widetilde{O}\left(\kappa\sqrt{d}\left(1+\frac{A}{\varepsilon}\right)\right).$$

Proposition 4.1 and Remark 6 supply the scalar Gaussian rare-reply lower bound with a fixed horizon. Their upper bound allows small variance; neither its inverse-accuracy exponent nor that lower-bound mechanism is a contribution here.

Finite initial divergence is already available under our localization promise: the query-free law $q_0 = N(0, L^{-1}I_d)$ satisfies

$$1 + \chi^2(q_0\|\nu_f) \le \kappa^{d/2}e^{\kappa}.$$

Indeed, $Z_f e^{f(x_f^\star)} \le (2\pi/\mu)^{d/2}$ and $f(x) - f(x_f^\star) \le L\|x - x_f^\star\|^2/2$ give this bound by completing the square in $\int q_0^2/\nu_f$ and using $L\|x_f^\star\|^2 \le \kappa$. Thus initialization alone is not a barrier to applying that fixed-kernel benchmark here; the joint rate and the oracle and stopping conventions are the relevant distinctions.

Our comparison concerns the joint minimax rate and the operational model. We prove matching constant-factor bounds for $1 + \log(1+\kappa) + A/\varepsilon$, for every $A \ge 0$, with constants depending only on the fixed dimension. The mode is unknown, with only a prescribed localization ball. The response distribution may change with the entire history, provided that its conditional mean and total variance obey the oracle conditions. The TV algorithm has worst expected physical-call cost of this order and terminates on every real-valued execution path. At $A = 0$, an uncapped variant additionally produces an exact sample with $O(1 + \log(1+\kappa))$ expected calls and almost-sure termination. These are distinct stopping statements. Restricting an existing theorem to fixed dimension does not supply either the joint rate or its implementation in this model.

Noise tails are a separate assumption. The sub-exponential branch of Chen et al. (2026b, Remark 5) has only polylogarithmic accuracy dependence, with a separate tail proxy. A small variance bound alone does not imply such a tail bound: at every positive A, the admissible class still contains rare, large replies. The transition to $A = 0$ therefore has to be analyzed within the variance-only model. Our acceptance construction retains a noise error proportional to A, without treating randomness in the integration locations as additional oracle noise.

Scalar sampling and query lower bounds. Chewi et al. (2022) obtain the known-mode scalar $\Theta(\log\log\kappa)$ rate. Theorem 3 uses values to construct a rejection envelope, followed by constant expected additional cost for an exact sample; Theorem 2 permits values, gradients, and second derivatives in its lower bound. Their Appendix C gives bisection-based localization for a conditional sampling subproblem with an unknown mode. Its upper construction still uses values. Thus the double-logarithmic rate does not contradict the localization cost proved here. Likewise, the logarithmic

accuracy cost of their capped rejection procedure is a maximum-budget statement; it does not replace their constant expected additional cost.

Exact-information fixed-dimensional sampling. For fixed $d \geq 2$, Chewi et al. (2024, arXiv v2, Theorem 49) give a known-mode value-and-gradient algorithm with a query budget $O(\log(1+\kappa) + \log^{O(d)}(1/\varepsilon))$. Their proof truncates the rejection procedure. In contrast, the quantity (2.4) measures expected cost. Appendix F derives a full-support rejection sampler from their ellipsoidal sublevel-set approximation. This gives exact sampling with $O(\log(1+\kappa))$ expected value-and-gradient queries; truncation yields a deterministic budget $O(\log(1+\kappa) + \log(1/\varepsilon))$. These are consequences proved here for comparison, not statements attributed to their Theorem 49.

Our additional noiseless guarantee is gradient-only exact sampling: Corollary 5.1 attains the optimal expected order with a coarsely localized unknown mode, and Corollary 5.4 requires no supplied radius, at logarithmic instance-wise initialization cost. Positive-noise guarantees additionally cover history-dependent, conditionally unbiased variance-bounded replies. The ellipsoid geometry is classical; the gradient-only pilots, their summable stochastic costs, and value-free acceptance are the additional components analyzed here. We neither assume sublevel membership access nor simulate every geometric query at a common noise precision.

The lower-bound oracle comparison has the opposite direction. Chewi et al. (2024, arXiv v2, Theorem 4) establish an $\Omega(\log \kappa)$ lower bound at fixed constant TV accuracy already in dimension two, with known mode and even when values and gradients are available. Our scalar curvature family encodes an unknown mode, and its proof treats expected gradient-query cost. It proves optimality for our stated model, including $d = 1$, rather than a stronger lower bound for their more informative model. Appendix G separately isolates location hardness at fixed curvature. Scalar witnesses are completed with Gaussian coordinates and analyzed directly for full-vector algorithms.

Chatterji et al. (2022) develop a stochastic-gradient lower-bound framework using statistical decision theory. In our total-variance convention, their journal Theorem 4.1 (Theorem 1 in arXiv v3) gives TV error at least $\sigma/(16\sqrt{n})$ when $n \geq \sigma^2/4$ and $\alpha \leq \sigma^2/(256n)$, over α-smooth, $\alpha/2$-strongly convex potentials whose target mean has norm at most $1/\alpha$. Their notation instead writes the total variance as $d\sigma^2_{\text{paper}}$. The inverse-square accuracy dependence must be read with the curvature restriction: it does not give such a lower bound at arbitrarily small accuracy with a fixed positive strong-convexity parameter, and so does not contradict the present inverse-linear rate.

Rejection sampling with limited information. Adaptive rejection sampling has long exploited the geometry of univariate log-concave densities. Gilks & Wild (1992) use tangent bounds on the log-density to construct piecewise exponential proposals, refining the envelope as sampling proceeds. This establishes the classical role of shape-adapted proposals and acceptance bounds in scalar sampling. In that setting, the evaluations used to construct and test the envelope are exact. With the oracle studied here, a reported gradient cannot be treated as a deterministic supporting slope. Our preprocessing therefore controls its own failure probability, and our acceptance analysis separately accounts for clipping noisy marking probabilities.

The first-order rejection sampling framework, FORS, of Chen et al. (2026a) is a direct methodological predecessor. Section 3.1 and Theorem 3.1 use bounded log-ratio estimators and a Poisson product to sample their exponential tilt exactly; randomized gradient integration supplies potential differences. Their Gaussian-tilt instantiation clips a generally unbounded estimator and approximates the original target. The smooth-LSI application in the extended version (arXiv v2, Theorem G.1(i)) assumes exact gradients and proximal access, fixed Rényi order $\lambda \geq 2$, and finite initial R_λ. For $0 < \varepsilon \leq 1/2$, it attains $R_\lambda \leq \varepsilon^2$; its strongly convex specialization costs $\widetilde{O}(\kappa\sqrt{d})$ expected calls to those oracles, retaining accuracy and initialization logarithms. Primitive exactness is therefore different from an exact-target complexity theorem.

The Poisson acceptance identity also has classical precedents. Beskos et al. (2006) simulate an exponential of a nonnegative path integral by testing whether a Poisson process places any points under its graph. Their exact diffusion algorithms reveal only the required path skeleton, under drift regularity and implementable path-envelope assumptions. This is an exact-simulation principle, not a strongly convex gradient-query bound. Our proof applies the same no-point identity on a deterministic segment, with envelopes constructed from gradient replies and with their full query costs included.

The additional acceptance analysis here is quantitative. Compensating the negative directional gradient produces a nonnegative intensity. A fixed slack keeps its normalized marking probabilities inside $[0, 1]$, so clipping an averaged noisy mark introduces bias controlled by the oracle variance. A survival-probability recursion bounds the accumulated error under history-dependent replies. Early rejection controls the expected number of processed nodes, including when the proposal-dependent Poisson intensity is large. At zero variance the uncapped experiment with exact marks recovers exact acceptance. We claim neither Poisson rejection, gradient-path integration, nor exact simulation from incomplete numerical information as new principles.

Bernoulli factories provide another perspective on acceptance decisions from incomplete numerical information. The existence theory of Keane & O'Brien (1994) and the quantitative constructions of Nacu & Peres (2005) study simulation of a coin with probability $h(p)$ from independent coins of unknown probability p. The latter obtain fast simulation for analytic functions on suitable compact intervals. These results distinguish exact simulation of a transformed probability from applying a nonlinear function to an estimator. Their bounded independent input coins are not automatically supplied by a real-valued conditional variance oracle. Our proof uses a direct Poisson marking experiment and specifies its bias and stopping costs, without assuming an additional Bernoulli-factory oracle.

A related distinction appears in the pseudo-marginal approach of Andrieu & Roberts (2009): appropriate nonnegative unbiased estimates of an unnormalized density can be incorporated into an augmented Markov chain while preserving the desired marginal stationary distribution. An unbiased estimate of a potential difference is a different object, since exponentiation does not preserve unbiasedness. Moreover, invariance of a Markov chain and a finite-query total-variation guarantee for its output are separate requirements. These observations explain why the integral estimator alone is insufficient here. The acceptance bias, the probability of reaching each trial, and the number of calls before termination all enter the complexity proof.

Langevin methods and stochastic gradients. The nonasymptotic analysis of Langevin algorithms forms a major part of the broader sampling literature. Dalalyan (2017) derive explicit approximation guarantees for Langevin Monte Carlo under smooth log-concavity, using comparisons with continuous-time diffusion. Durmus & Moulines (2017) analyze the unadjusted Langevin algorithm in total variation, including constant and decreasing step sizes. These results show how smoothness, curvature, dimension, initialization, and discretization interact in concrete sampling schemes. Their algorithm-specific bounds are not lower bounds on all possible gradient-query algorithms. In the present fixed-dimensional problem, rejection from a directly constructed proposal replaces the task of following a discretized diffusion for a prescribed number of steps.

Welling & Teh (2011) introduce stochastic gradient Langevin dynamics for Bayesian learning from minibatches. Their construction combines stochastic optimization updates with injected Gaussian noise and a decreasing step size. It motivates the computational setting in which full gradients are expensive and random gradient estimates are available. The minimax question considered here fixes the dimension and is more abstract in its oracle specification: there need not be a dataset, a finite-sum decomposition, or access to individual likelihood terms. The theorem concerns the marginal distribution of one output after a finite, explicitly accounted sampling procedure, rather than an asymptotic annealing statement.

Dalalyan & Karagulyan (2019) give quantitative Wasserstein guarantees for Langevin Monte Carlo with inaccurate gradients, encompassing stochastic and deterministic errors. Durmus et al. (2019) develop a convex-optimization analysis of Langevin Monte Carlo and extensions involving stochastic gradients. These works make clear that gradient error must be incorporated into the sampling analysis instead of being treated as an implementation detail. The metric remains important in comparisons: a Wasserstein estimate does not, without a further argument, give the total-variation accuracy required here. Likewise, a result for a particular noisy Langevin recursion does not determine the optimum over adaptive algorithms that can choose batches, integration locations, proposals, and stopping decisions.

Variance reduction can substantially change the information available to a sampler. Chatterji et al. (2018) analyze SAGA, SVRG, and control-variate variants of stochastic-gradient Monte Carlo in Wasserstein distance under smoothness, strong convexity, and additional regularity assumptions. Their constructions exploit finite-sum structure and correlated evaluations of component gradients.

Such access is useful, but is not supplied by an arbitrary conditional variance bound. Our oracle may alter its response law after every call and need not permit the same random component to be evaluated at two locations. We therefore cannot assume the cancellation identities underlying a finite-sum control variate. Conversely, our worst-case guarantee does not describe the additional gains possible when this structure is available.

Proximal sampling, warm starts, and recent exact-gradient methods. Lee et al. (2021) introduce the restricted Gaussian oracle as a building block for structured log-concave sampling. It samples from the target after adding a quadratic localization term and supports reductions for composite and finite-sum potentials. Their convention for a first-order oracle includes both a potential value and its gradient, which is stronger than the gradient-only access assumed here. Chen et al. (2022) extend the analysis of the resulting proximal sampler to weak log-concavity and functional-inequality assumptions, and relate it to optimization in Wasserstein space. These reductions explain the importance of accurately sampling localized subproblems, but implementing the subproblem remains part of an oracle-complexity theorem.

Fan et al. (2023) improve the dimension dependence of proximal sampling by refining the implementation of the restricted Gaussian oracle, obtaining high-accuracy guarantees under several regularity regimes. Their strongly log-concave results achieve an $\widetilde{O}(\kappa\sqrt{d})$ dependence, with the corresponding accuracy factors, and analyze initialization as well as approximate rejection. This is a multivariate algorithmic benchmark, not a joint fixed-dimensional stochastic-oracle minimax theorem. Our proof also constructs a region on which rejection is efficient, but does so once for the original target. It does not run an outer proximal chain or assume a pre-existing implementation of a restricted Gaussian oracle.

Initialization can be a substantive computational cost even for exact-gradient methods. Altschuler & Chewi (2024) construct algorithmic warm starts for high-accuracy log-concave sampling, combining Rényi-divergence control for underdamped Langevin discretizations with subsequent high-accuracy sampling. This addresses the gap between a mixing result from a favorable initial law and an implementable algorithm that must obtain such a law. Our initialization requirement is instead a deterministic ball containing the unknown minimizer. The pilot converts this information into a point with a constant potential gap and then into a proposal with constant overlap. Its cost is included in the displayed minimax bound.

More recently, Chen et al. (2026e) use rejection on diffusion path space with Girsanov-based ratio estimators. Their Theorem 3.2(ii) assumes a C^2 convex, smooth potential satisfying an LSI, and controls a fixed-order Rényi divergence from a supplied phase-space initialization. With suitable initial divergence, its exact-gradient complexity is $\widetilde{O}(\kappa^{2/3}d^{1/3}\,\mathrm{polylog}(1/\varepsilon))$; additional derivative assumptions yield further improvements. The target-law guarantee retains a positive accuracy parameter. This multivariate result does not implement an arbitrary conditional variance oracle or supply the zero-noise fixed-dimensional exact-target statement proved here.

The September 2026 preprints of Chen et al. (2026c;d) further improve exact-gradient sampling. The Proximal Bouncy Particle Sampler gives expected $\widetilde{O}(\sqrt{\kappa}\,d^{1/4})$ queries from a supplied bounded order-2 Rényi warm start. The companion Smoothed Picard HMC theorem, including implemented proximal computations, gives $\widetilde{O}(\kappa^{7/6}d^{1/6}+\sqrt{\kappa}\,d^{1/4})$ from a reference point with gradient norm at most $\sqrt{\mu d}$. These statements concern exact gradients and C^2 curvature bounds; they do not implement an arbitrary conditional variance-only oracle.

The result established here is the joint fixed-dimensional rate over the full range $A \geq 0$, together with its fully charged gradient-only construction and exact zero-noise variant. The constant may depend on d, but not on κ, A, ε. Classical ellipsoid geometry, Poisson and gradient-path rejection, and rare-message lower bounds remain their respective precedents. The proofs below supply the energy amortization, marking-error analysis, and stopping arguments needed to combine these ingredients in the specified oracle model.

B SUBROUTINE AND IMPLEMENTATION GUIDE

Algorithms 1 and 2 state the sampling rules with symbolic parameters in Section 2.1. The proof of Theorem 2.2 in Appendix C.4 gives their public calibration in (C.107)–(C.108) and the finite-category weights in (C.109). The named routines abbreviate the finite constructions in Appendix C; they are not additional oracles. This guide locates their definitions and guarantees. The working-ellipsoid symbols c_E, Q_E in Section 2.1 are c, Q in the two pilot proofs; its temporary batch size b is denoted n there, not the fixed correction batch size. The midpoint probes $y_i^{\pm}$ are written $z_{i,\pm}$ in the rounding proof. These local names separate geometric matrices from the scalar correction normalizer Q.

The center routine is the procedure in Proposition C.2, including its update cap, physical-call cap, and fallback. The rounding routine is Proposition C.7, with its own caps and nonsingular fallback geometry. Its terminal score threshold is 4; contracting the displayed midpoint ellipsoid toward m by $1/5$ yields exactly $z = m + (c_E - m)/10$ and $\mathsf{T} = \gamma Q_E/(10\sqrt{d})$ in (C.53). The returned parameters determine U, q through (C.56)–(C.57); Proposition C.11 gives the finite radial sampling rule. Neither pilot queries a potential value or a sublevel-set membership oracle.

For each proposed point, the two directional grids are defined by (C.71)–(C.76). Lemma C.15 supplies the derivative convention used in the mark probability, while the finite rectangle mixture after (C.82) samples a correction position. Every repeated location in a new batch still incurs a new physical call. The overflow branch and all marker calls are explicit in the main-body specification; (C.109) gives the category weights used in its implementation.

Section C.4 proves measurability, physical-transcript stopping, pathwise finite execution of the capped procedure, and its unconditional expected cost. Its proof includes failed pilots, failed envelopes, ignored marker replies, and the zero-radius branch. The zero-noise exact modification is specified separately in the proof of Corollary 5.1; it removes both rejection caps and asserts almost-sure, rather than every-path, termination.

C PROOF OF THE UPPER BOUND

All algorithms below operate in normalized coordinates. Their outputs are divided by $\sqrt{\mu}$ before being returned.

For the global sampler, take $B = 3/64$ and $b_{\mathrm{acc}} = 10$. Local pilot statements that allow a variable B retain their stated ranges. Thus the numerical offset 10 in the acceptance calculations below is the specialization of b_{acc} in Algorithm 2. We collect the full public parameter recipe when completing the proof of Theorem 2.2; the preceding estimates justify that choice.

C.1 CONDITIONAL AVERAGES AND A CONSTANT-GAP CENTER

The statistical input to both geometric pilots is a conditional mean-square estimate. We establish it first because every later batch is entered at a random, history-dependent time.

Lemma C.1 (Vector averages under adaptive kernels)**.** *At a point fixed by an incoming history, the mean $\overline{G}$ of n new vector replies satisfies*

$$\mathbb{E}[\overline{G} \mid \mathcal{H}] = \nabla F(z), \qquad \mathbb{E}[\|\overline{G} - \nabla F(z)\|^2 \mid \mathcal{H}] \le A/n.$$

Consequently $\mathbb{P}(\|\overline{G} - \nabla F(z)\| > t \mid \mathcal{H}) \le A/(nt^2)$. Projection onto any unit vector fixed before the calls has the same scalar variance bound.

Proof. The centered replies form martingale differences. If $i < j$, the ith difference is measurable before reply j, whose conditional mean is zero. Their conditional inner-product expectation is therefore zero, by the tower property. Expanding the squared norm of the sum and using the individual variance bounds proves the assertion. Chebyshev and norm contraction under projection give the consequences. At an adaptively reached batch, condition on its incoming history; the event that it is reached is already determined. □

This estimate will be used to find a single point with a constant potential gap, rather than to optimize to the final sampling accuracy.

We work in the normalized model: F is 1-strongly convex and κ-smooth, its minimizer $x_\star$ satisfies $\|x_\star\| \le 1$, and the total conditional squared noise of a vector reply is at most $A \ge 0$. All conditional statements below refer to the complete incoming history. We write $\mathbb{B}_d = \{x \in \mathbb{R}^d : \|x\| \le 1\}$ and

$$K_v = \{x \in \mathbb{R}^d : F(x) - F(x_\star) \le v\}, \qquad v > 0.$$

These level sets are used in the analysis and are not supplied to the algorithm.

Proposition C.2 (Center construction)**.** *Fix $d \ge 1$, $0 < B \le 1/16$, and $0 < \delta < 1$. There is an algorithm which returns a point m such that*

$$\mathbb{P}\{F(m) - F(x_\star) > B\} \le \delta. \tag{C.1}$$

Every actual gradient query and every possible output belongs to $2\mathbb{B}_d$. The algorithm terminates on every real-valued reply history, and its number of physical calls is a stopping time for the physical transcript. On every such history its call count obeys

$$T_{\text{center}} \le C_{d,B}\left(\log(1+\kappa) + \frac{A}{\delta}\right), \tag{C.2}$$

where $C_{d,B} < \infty$ is independent of κ, A, δ and of the admissible conditional oracle. In particular, fixing $B = 3/64$ gives a constant depending only on d.

The proof gives explicit public budgets. A symmetric collection of vertices certifies the center's potential gap, and also determines the accuracy needed for the next cut. The inverse energy scales of the estimated iterations have a bounded sum. This allows both the sample costs and the conditional failure probabilities to be summed without a factor of $\log \kappa$ multiplying A/δ.

We begin with the deterministic operation that every subsequent geometric update will use. Its explicit formula fixes both the amount of retained slack and the contraction charged in the volume argument.

Set

$$\gamma = \frac{1}{8d}, \qquad \alpha = 2\gamma = \frac{1}{4d}, \qquad a_d = \frac{9}{32(d+1)}. \tag{C.3}$$

An ellipsoid is represented as $E = c + Q\mathbb{B}_d$, with Q invertible; Q need not be symmetric.

Lemma C.3 (Fixed shallow-cut update)**.** *For any nonzero $v \in \mathbb{R}^d$, one can obtain an ellipsoid E' containing*

$$E \cap \{x : \langle v, x - c\rangle \le \alpha \|Q^T v\|\} \tag{C.4}$$

by finitely many arithmetic operations and square roots. The update preserves invertibility of the shape matrix and satisfies

$$\operatorname{vol}(E') \le e^{-a_d} \operatorname{vol}(E) \le e^{-1/(8d)} \operatorname{vol}(E). \tag{C.5}$$

Proof. Put $u = Q^T v/\|Q^T v\|$. For $d \ge 2$, define

$$\begin{aligned} \beta &= \frac{1 - d\alpha}{d+1}, \qquad a = \frac{d(1+\alpha)}{d+1}, \qquad h = \frac{d\sqrt{1-\alpha^2}}{\sqrt{d^2-1}}, \\ c' &= c - \beta Q u, \qquad Q' = Q\big(auu^T + h(I - uu^T)\big). \end{aligned} \tag{C.6}$$

Both eigenvalues a, h in the last factor are positive, so Q' is invertible. In unit coordinates, write $s = \langle u, z\rangle$ for a point $z \in \mathbb{B}_d$ with $s \le \alpha$. Its squared norm in the new ellipsoid is at most

$$\frac{(s+\beta)^2}{a^2} + \frac{1 - s^2}{h^2}, \qquad -1 \le s \le \alpha.$$

Since $\alpha < 1/d$, one has $a < h$, so this expression is a convex quadratic in s. It equals one at $s = -1$ and at $s = \alpha$ and hence is at most one throughout the interval. This proves inclusion of (C.4).

The volume ratio is ah^{d-1}. At $\alpha = 1/(4d)$,

$$a = 1 - \frac{3}{4(d+1)}, \qquad h^2 = 1 + \frac{15}{16(d^2-1)}.$$

Using $\log(1-x) \le -x$ and $\log(1+x) \le x$ gives

$$\log a \le -\frac{3}{4(d+1)}, \qquad (d-1)\log h \le \frac{15}{32(d+1)}.$$

Their sum is at most $-a_d$.

For $d = 1$, use the interval update

$$\beta = \frac{1-\alpha}{2}, \qquad a = \frac{1+\alpha}{2}, \qquad c' = c - \beta Q u, \qquad Q' = aQ. \tag{C.7}$$

In unit coordinates this is precisely the interval $[-1, \alpha]$. Here $\alpha = 1/4$ and $a = 5/8$, and $\log a \le -3/8 \le -a_1$. Finally, $9/[32(d+1)] \ge 1/(8d)$ for every integer $d \ge 1$, which proves the second bound in (C.5). □

The shallow-cut update is deterministic. To use it with noisy normals, we next identify the quantities that a bounded set of test vertices can certify and then estimate those quantities with fresh batches.

Consider an ellipsoid $E = c + Q\mathbb{B}_d$ whose symmetric vertices

$$z_i^+ = c + \gamma Q e_i, \qquad z_i^- = c - \gamma Q e_i, \qquad 1 \le i \le d, \tag{C.8}$$

all belong to $2\mathbb{B}_d$. There are $D = 2d$ vertices; here D counts points, not a diameter. Let $w_i^\pm = z_i^\pm - c$. The midpoint and difference of each opposite pair give the bounds below, for which we set $W = 2$ and $R_Q = 2\sqrt{d}/\gamma$:

$$\|c\| \le 2, \qquad \|w_i^\pm\| \le W = 2, \qquad \|Q\|_{\mathrm{op}} \le R_Q. \tag{C.9}$$

Symmetry will turn a bound on the vertex scores into a center-gap certificate with factor $C_0 = \sqrt{d}/\gamma$. Accordingly, introduce the score floor

$$m_0 = \frac{B}{4C_0}. \tag{C.10}$$

This floor is fixed by the requested center accuracy, not by the unknown potential values. Define the unobserved scores and energy scale

$$s_i^\pm = \langle \nabla F(z_i^\pm), w_i^\pm \rangle, \qquad M = \max\{m_0, s_i^\pm : 1 \le i \le d\}, \qquad \widetilde{M} = \frac{M}{m_0} \ge 1. \tag{C.11}$$

The algorithm estimates this scale without being given M.

For this estimation subroutine, fix a confidence coefficient $0 < \zeta < 1$; its value for the complete pilot will be chosen after the reciprocal-energy bound is proved. To keep noisy cut normals within the geometric tolerance, set

$$t_0 = \frac{1}{100 R_Q}, \qquad V = \frac{A}{m_0^2}. \tag{C.12}$$

The second quantity is the variance ceiling after dividing replies by m_0. The stopping threshold is

$$\Lambda = \max\left\{1, \frac{8DV(W^2 + t_0^{-2})}{\zeta}\right\}. \tag{C.13}$$

The unit floor keeps the batch schedule nonempty even at $A = 0$; only the other term is needed for the probability estimates.

At fresh stages $n = 1, 2, 4, \ldots$, take n new replies at each of the D vertices. Divide each resulting mean by m_0 and call the normalized vector $\widehat{g}_i^\pm$. Form

$$\widehat{M} = \max\{1, \langle \widehat{g}_i^\pm, w_i^\pm \rangle : 1 \le i \le d\}. \tag{C.14}$$

End the subroutine at the first stage with $n\widehat{M} \ge \Lambda$. Return the unnormalized means $G_i^\pm = m_0 \widehat{g}_i^\pm$ and scores $\widehat{s}_i^\pm = \langle G_i^\pm, w_i^\pm \rangle$. Earlier stages are discarded. Since $\widehat{M} \ge 1$, the first dyadic $n \ge \Lambda$ always ends the subroutine on every reply history.

Lemma C.4 (Adaptive batches for the center)**.** *Conditional on the complete incoming history fixing the bounded vertices (C.8), the subroutine has an event $\mathcal{E}_{\mathrm{est}}$ satisfying*

$$\mathbb{P}(\mathcal{E}_{\mathrm{est}}^{c} \mid \mathcal{H}) \le \frac{\zeta}{\widetilde{M}}. \tag{C.15}$$

On $\mathcal{E}_{\mathrm{est}}$ all returned vectors obey

$$\|G_i^{\pm} - \nabla F(z_i^{\pm})\| \le t_0 M, \tag{C.16}$$

and the total number of calls in all its fresh stages is at most

$$2D + \frac{8D\Lambda}{\widetilde{M}}. \tag{C.17}$$

The event can be chosen measurable at the end of the subroutine.

Proof. By Lemma C.1, the normalized mean at any fixed vertex and reached stage has conditional squared error at most V/n. This remains true under all intervening adaptation of the oracle.

First consider a reached stage with $n < \Lambda/(2\widetilde{M})$. Stopping there requires $\widehat{M} \ge \Lambda/n > 2\widetilde{M}$. The floor value one in (C.14) cannot cause this. Some estimated scalar score must therefore have error at least $\Lambda/(2n)$. By (C.9), the corresponding vector error is at least $\Lambda/(2nW)$. Conditional Markov inequality and a union over the vertices give the upper bound

$$\frac{4DVW^2n}{\Lambda^2}$$

for stopping at this early stage. The sum of all early dyadic sizes is at most $\Lambda/\widetilde{M}$, so the probability of any early stop is at most

$$\frac{4DVW^2}{\Lambda\widetilde{M}}. \tag{C.18}$$

At each reached late stage $n \ge \Lambda/(2\widetilde{M})$, the probability that any normalized vector error exceeds $t_0\widetilde{M}$ is at most

$$\frac{DV}{nt_0^2\widetilde{M}^2}.$$

The sum of reciprocal late dyadic sizes is at most $4\widetilde{M}/\Lambda$, including when the first such size is one. The probability of any inaccurate reached late stage is thus at most

$$\frac{4DV}{t_0^2\Lambda\widetilde{M}}. \tag{C.19}$$

Define $\mathcal{E}_{\mathrm{est}}$ by excluding the early-stop event and these inaccurate reached late stages. Only stages actually executed enter this definition, so the event is measurable when the subroutine ends. Equations (C.18)–(C.19) and (C.13) give (C.15). Each union is over fixed dyadic indices; reaching an index is decided before its new replies. These arguments therefore require no independence across stages.

On $\mathcal{E}_{\mathrm{est}}$ the returned stage is late and accurate. At any accurate stage, putting $\theta = Wt_0 < 1/100$, one has

$$\widehat{M} \ge (1-\theta)\widetilde{M} \ge \widetilde{M}/2.$$

If $\widetilde{M} = 1$, the first inequality also follows from the floor $\widehat{M} \ge 1$. Consequently the first dyadic $n \ge 2\Lambda/\widetilde{M}$, if reached, triggers stopping. The final size obeys $n_{\mathrm{stop}} \le \max\{1, 4\Lambda/\widetilde{M}\}$. The sum of the fresh stage sizes is at most $2n_{\mathrm{stop}}$, proving (C.17). Rescaling the normalized vector errors proves (C.16). □

We now specify the local iteration rules. They either make a free geometric cut, estimate the vertex scores, or return a certified center. The global caps are chosen after their preservation and energy bounds have been established.

1. Form the vertices (C.8). If some vertex z lies outside $2\mathbb{B}_d$, choose the first such vertex in a fixed ordering and apply Lemma C.3 with normal $v = z$. Count one update and continue. This step makes no oracle calls.
2. If all vertices lie in $2\mathbb{B}_d$, execute the adaptive subroutine (C.14).
3. If the returned scores satisfy
$$\max_{i,\pm} \widehat{s}_i^{\pm} \le 2m_0, \tag{C.20}$$
return the current c and stop.
4. Otherwise choose the first maximizing vertex z and its returned mean G. Thus $\widehat{s} = \langle G, z - c\rangle > 2m_0$ and $G \neq 0$. Apply Lemma C.3 with normal $v = G$, count one update, and continue.

All choices use a fixed index order to break ties. The next lemma explains why an accurate estimate makes either the return or the cut valid. The set to be preserved is K_b, where we now set $b = m_0$; thus the protected level matches the score floor in the stopping rule.

Lemma C.5 (Accurate center iterations)**.** *As long as all preceding estimation events $\mathcal{E}_{\mathrm{est}}$ hold, the current ellipsoid contains K_b. At an estimated iteration with $K_b \subset E$, its true scale satisfies*
$$F(c) - F(x_\star) \le C_0 M. \tag{C.21}$$
If its estimation event also holds, a return under (C.20) has gap at most $3B/4$, and any update made instead preserves K_b.

Proof. Strong convexity and smoothness at the minimizer give
$$\mathbb{B}\left(x_\star, \sqrt{\frac{2b}{\kappa}}\right) \subset K_b \subset \mathbb{B}(x_\star, \sqrt{2b}) \subset 2\mathbb{B}_d. \tag{C.22}$$
Thus the initial ellipsoid contains K_b. The final containment in (C.22) uses the normalized localization promise $\|x_\star\| \le 1$; it is not implied by the curvature bounds alone. The extensions in Section 5 restore this promise by translation and a conservative public strong-convexity bound before invoking the existing sampler. They do not delete the promise from this pilot or enlarge its initial ellipsoid without changing its analysis.

For a free update, z is outside $2\mathbb{B}_d$. Every $y \in 2\mathbb{B}_d$ satisfies
$$\langle z, y - z\rangle \le 2\|z\| - \|z\|^2 < 0.$$
Since $z - c = \pm\gamma Q e_i$ for some i,
$$\langle z, z - c\rangle \le \gamma\|Q^T z\| \le \alpha\|Q^T z\|.$$
The halfspace used by the update therefore contains $2\mathbb{B}_d$, and the new ellipsoid contains the previously retained K_b.

For an estimated iteration, monotonicity of the gradient on each of the coordinate lines through c gives
$$-s_i^- \le \gamma(Q^T \nabla F(c))_i \le s_i^+.$$
Every coordinate in the middle is therefore bounded in absolute value by M, and
$$\|Q^T \nabla F(c)\| \le \frac{\sqrt{d}}{\gamma} M = C_0 M.$$
The inclusion $x_\star \in E$ supplies a vector $u_\star$ of norm at most one with $x_\star = c + Q u_\star$. Convexity now proves (C.21). This is a certificate at the single point c; no function values or identification of an unknown best iterate are needed.

Suppose the current estimation event holds. With $\theta = Wt_0 < 1/100$,
$$|\widehat{s}_i^{\pm} - s_i^{\pm}| \le \theta M. \tag{C.23}$$
If (C.20) holds, then $M \le 2m_0 + \theta M$, whence $M \le 3m_0$. Equation (C.21) bounds the returned gap by $3C_0 m_0 = 3B/4$.

Otherwise, $\widehat{s} > 2m_0$. The possibility $\max_{i,\pm} s_i^{\pm} \le m_0$ would contradict (C.23); hence $M = \max_{i,\pm} s_i^{\pm}$. Since the chosen estimated score is maximal,

$$M \le \frac{\widehat{s}}{1-\theta}.$$

For $y \in K_b$, convexity at its selected vertex z yields

$$\langle \nabla F(z), y - z \rangle \le F(y) - F(z) \le b.$$

The points $y \in E$ and $z = c \pm \gamma Q e_i$ satisfy $\|y - z\| \le (1+\gamma)\|Q\|_{\mathrm{op}} \le 2R_Q$. Using (C.16) and $2R_Q t_0 = 1/50$ gives

$$\langle G, y - z \rangle \le b + \frac{M}{50} \le b + \frac{\widehat{s}}{25}.$$

Also $\|Q^T G\| \ge \widehat{s}/\gamma$ because $\widehat{s} = \langle G, z - c \rangle$. Re-centering the preceding inequality at c gives

$$\frac{\langle G, y - c \rangle}{\|Q^T G\|} \le \gamma \left(1 + \frac{1}{25} + \frac{b}{\widehat{s}}\right) < 2\gamma = \alpha,$$

where $b = m_0$ and $\widehat{s} > 2m_0$ were used. The actual shallow halfspace therefore contains K_b. Lemma C.3 proves its preservation after this update, completing the induction. □

Preservation alone bounds the number of cuts, but does not yet control their noise cost. To obtain an additive bound, we relate each current energy to the remaining ellipsoid volume. The vertex gap adds one copy of M to the center-gap bound, so put

$$C_1 = C_0 + 1, \qquad H_d = (C_0 C_1)^d. \tag{C.24}$$

The resulting volume comparison also controls the reciprocal energies, even when the algorithm revisits an energy range.

Lemma C.6 (Energy and ellipsoid volume)**.** *At any estimated iteration whose preceding cuts preserve K_b, the current ellipsoid and its true scale satisfy*

$$\mathrm{vol}(E) \le H_d \widetilde{M}^d \,\mathrm{vol}(K_b). \tag{C.25}$$

Along any prefix of the iteration rules above in which all preceding estimation events hold, including one final estimated iteration before its own accuracy is decided,

$$\sum_{\text{estimated iterations in the prefix}} \frac{1}{\widetilde{M}} \le C_{\mathrm{sum}}, \qquad C_{\mathrm{sum}} := 2 + \frac{2\log H_d + 4d\log 2}{a_d}. \tag{C.26}$$

Proof. Convexity at a vertex gives

$$F(z_i^{\pm}) - F(c) \le \langle \nabla F(z_i^{\pm}), z_i^{\pm} - c \rangle \le M.$$

Together with (C.21), this places every vertex in $K_{C_1 M}$. Their convex hull contains $c + (\gamma/\sqrt{d}) Q \mathbb{B}_d$, since the Euclidean unit ball scaled by $1/\sqrt{d}$ is contained in the unit ℓ_1 ball. It follows that

$$\mathrm{vol}(E) \le \left(\frac{\sqrt{d}}{\gamma}\right)^d \mathrm{vol}(K_{C_1 M}).$$

For $v \ge b$, convexity about the minimizer implies

$$K_v \subset x_\star + \frac{v}{b}(K_b - x_\star).$$

Indeed, for $x \in K_v$ the point $x_\star + (b/v)(x - x_\star)$ has gap at most b. Since $M = m_0 \widetilde{M} = b\widetilde{M}$, the preceding volume bounds prove (C.25) with $H_d = (C_0 C_1)^d$.

Fix an integer $k \ge 0$. At the first estimated iteration in the band $2^k \le \widetilde{M} < 2^{k+1}$, the current ellipsoid has volume at most

$$H_d 2^{d(k+1)} \,\mathrm{vol}(K_b).$$

Every intervening update, free or estimated, reduces volume by at least e^{-a_d}. All ellipsoids before the last iteration of the prefix still contain K_b. Therefore the number of estimated iterations in this band is at most

$$1 + \frac{\log H_d + d(k+1)\log 2}{a_d}.$$

This argument does not assume that the energy scales are monotone. Summing the contribution of each band gives

$$\begin{aligned}\sum \frac{1}{\widetilde{M}} &\le \sum_{k=0}^{\infty} 2^{-k}\left(1 + \frac{\log H_d + d(k+1)\log 2}{a_d}\right) \\ &= 2 + \frac{2\log H_d + 4d\log 2}{a_d} = C_{\mathrm{sum}}.\end{aligned}$$

This proves (C.26) also for a prefix ending just before an estimation outcome is declared inaccurate. □

We can now assemble the finite pilot. The reciprocal-energy sum fixes the confidence allocation, rather than splitting confidence equally among all possible iterations:

$$\zeta = \frac{\delta}{2C_{\mathrm{sum}}}. \tag{C.27}$$

Use this value in the batch threshold (C.13). The lower volume of K_b and the shallow-cut contraction suggest the update cap

$$N = \left\lceil \frac{d\log(2\kappa/b)}{2a_d} \right\rceil + 1. \tag{C.28}$$

Summing the per-round batch bound over at most N estimated iterations and the reciprocal-energy budget then gives the physical-call cap

$$Q_{\max} = \lceil 2DN + 8D\Lambda C_{\mathrm{sum}} \rceil + 1. \tag{C.29}$$

All these quantities are public and finite. Initialize $E_0 = 2\mathbb{B}_d$ with $c = 0$ and $Q = 2I$, and apply the local rules above, allowing at most N updates in total, free or estimated. Before every fresh batch stage, check its entire cost Dn; if it would exceed $Q_{\max}$ cumulatively, return 0. Also return 0 immediately after the Nth update unless a center has already been returned. There is no further iteration after that update. The following proof shows that neither cap interrupts an all-good execution, while both caps remain effective on failed histories.

Proof of Proposition C.2. First analyze the same procedure with the N-update cap but without the global physical-call cap. This auxiliary procedure is also finite on every history: each estimation subroutine stops by its first dyadic $n \ge \Lambda$, and at most N updates are performed.

Condition on the incoming history of each reached estimation round. Lemma C.4 bounds its failure probability by $\zeta/\widetilde{M}$. Let $\mathcal{I}$ index reached estimation rounds with no earlier failed estimation event. All preceding cuts preserve K_b, so (C.26) gives

$$\begin{aligned}\mathbb{P}\{\text{some estimation event fails}\} &\le \zeta\,\mathbb{E}\left[\sum_{i\in\mathcal{I}} \frac{1}{\widetilde{M}_i}\right] \\ &\le \zeta C_{\mathrm{sum}} = \frac{\delta}{2}.\end{aligned} \tag{C.30}$$

The summands include at most one failed iteration. The reciprocal sum is bounded pathwise on just this prefix, so no bound conditional on future success has been used.

On the event that every estimation event holds, the N-update fallback is impossible. All updates retain K_b, which contains the ball in (C.22). But after N updates,

$$\mathrm{vol}(E_N) \le e^{-a_d N}\,\mathrm{vol}(2\mathbb{B}_d) < \mathrm{vol}\left(\mathbb{B}\left(x_\star, \sqrt{2b/\kappa}\right)\right)$$

by the definition of N in (C.28). Thus the auxiliary procedure must have returned a center c under (C.20), and its gap is at most $3B/4$ by Lemma C.5.

On the same event, there are at most N estimated iterations. Their individual pathwise costs (C.17) and the sum (C.26) bound the entire physical call count by

$$2DN + 8D\Lambda C_{\mathrm{sum}} < Q_{\max}. \tag{C.31}$$

Consequently the global physical-call cap cannot alter any such run. Formally, use the same conditional oracle kernels to couple the capped and auxiliary procedures up to their first possible separation. On the successful auxiliary event, no stage would exceed the global budget, so separation does not occur. Combining this fact with (C.30) proves (C.1), with the stronger upper bound $\delta/2$.

On every history, including those with estimation failures, the capped algorithm uses at most $Q_{\max}$ physical calls. Its formula (C.29) gives

$$Q_{\max} = O\left(\log(1+\kappa) + \frac{A}{\delta}\right),$$

where the hidden constant may also depend on B: the quantities $D, W, R_Q, C_0, m_0, t_0, H_d, C_{\mathrm{sum}}$ depend only on d, B, while (C.27) gives $\zeta = \delta/(2C_{\mathrm{sum}})$ and (C.12) gives $V = A/m_0^2$. The unit floor in Λ adds only a constant depending on d, B, which is absorbed by $\log(1+\kappa)$ because $\kappa \geq 1$. This proves the deterministic bound (C.2), and therefore the same bound for the expected cost under every admissible oracle. No cost incurred after an erroneous estimate is discarded.

For completeness, the execution is well defined on arbitrary real-valued reply histories. The shape matrix stays invertible by Lemma C.3. A free-cut normal is nonzero because it is an out-of-ball vertex. An estimated-cut normal is nonzero because its score exceeds $2m_0$. All computed quantities at a finite stage are finite, and the stage sizes and the number of updates have finite public caps. Every actual query is made only after all vertices have passed the test for membership in $2\mathbb{B}_d$. A center returned under (C.20) is the midpoint of an opposite in-ball pair and hence is also in $2\mathbb{B}_d$; a fallback returns 0.

Finally, the pilot generates no private random variables. Starting from its known initial state, its free updates, query locations, stage lengths, continuation decisions, and possible budget termination can all be reconstructed from the physical replies already received. All operations between replies are finite. Thus the event that no further call will be made after a given physical reply is measurable from the physical transcript at that reply. The physical call count is a stopping time, and the completed pilot state can be parsed from that transcript for use by subsequent phases. This finishes the proof. □

C.2 Constructing a proposal in fixed dimension

We next construct the global proposal from a center of bounded potential error. Throughout this subsection, $d \geq 1$, the normalized potential F is 1-strongly convex and κ-smooth, and the total conditional gradient-noise variance is at most $A \geq 0$. Set

$$B = \frac{3}{64}, \qquad \alpha = \frac{1}{4d}, \qquad \gamma = \frac{1}{8d}, \qquad D = 2d, \qquad \eta_d = \frac{\gamma}{2\sqrt{d}}, \qquad C_{\mathrm{geom}} = \eta_d^{-d}. \tag{C.32}$$

The symbol D counts test points and is unrelated to the radial normalizing constant introduced below. Define further

$$L_{\mathrm{geom},d} = 4 + 16d \log C_{\mathrm{geom}} + 32d^2 \log 2, \tag{C.33}$$

$$H_\kappa = 2 + \left\lceil 8d^2 \log(8\sqrt{\kappa}) \right\rceil. \tag{C.34}$$

All constants with a subscript d depend only on the fixed dimension.

Proposition C.7 (Proposal construction from a constant-gap center)**.** *Suppose the public inputs and incoming physical transcript supply a point m satisfying $F(m) - \min F \leq B$, and let $0 < \delta < 1$. Put*

$$\zeta = \frac{\delta}{L_{\mathrm{geom},d}}, \qquad \lambda = \max\left\{1, \frac{2^{16}DA}{\zeta}\right\}, \qquad B_{\mathrm{geom}} = \left\lceil 2DH_\kappa + 8D\lambda L_{\mathrm{geom},d} \right\rceil. \tag{C.35}$$

There is a procedure using at most B_{geom} physical calls on every execution which returns a point $z \in \mathbb{R}^d$ and an invertible matrix T. With conditional probability at least $1-\delta$, these parameters satisfy

$$z + \mathsf{T}\mathbb{B}_d \;\subset\; \{x : F(x) - F(m) \leq 1\} \;\subset\; z + \rho\mathsf{T}\mathbb{B}_d, \qquad \rho = \frac{20\sqrt{d}}{\gamma}. \tag{C.36}$$

On every history, including failures and histories for which the supplied center does not satisfy its promise, the procedure is finite and its output obeys

$$\|z - m\| \le 1, \qquad \|\mathsf{T}\|_{\mathrm{op}} \le 1. \tag{C.37}$$

Its rules use only public parameters and physical queries and replies. In particular, its stopping decision and returned parameters are measurable in the physical transcript. Its call bound is $O(1 + \log\kappa + A/\delta)$.

We prove this proposition by first establishing the geometric and statistical components, then specifying their finite implementation.

We first identify the set that the rounding procedure must preserve. Its inner and outer radii determine the initial ellipsoid and the maximum number of valid cuts; neither radius requires a value query.

For this rounding argument, measure levels relative to the supplied center m, rather than to the unknown minimum used in the center pilot:

$$I(x) = F(x) - F(m), \qquad K_t = \{x : I(x) \le t\}, \qquad K = K_1. \tag{C.38}$$

Let $x_\star$ be the minimizer of F. Strong convexity and $B \le 1/8$ give

$$\|m - x_\star\| \le \sqrt{2B} \le \frac{1}{2}.$$

For $x \in K$, the same inequality gives $\|x - x_\star\| \le \sqrt{2(1+B)} \le 3/2$. Therefore

$$K \subset \mathbb{B}(m, 2). \tag{C.39}$$

The descent inequality

$$F\left(m - \frac{\nabla F(m)}{\kappa}\right) \le F(m) - \frac{\|\nabla F(m)\|^2}{2\kappa}$$

implies $\|\nabla F(m)\| \le \sqrt{2\kappa B} \le \sqrt{\kappa}/2$. Consequently, if $r_0 = 1/(4\sqrt{\kappa})$ and $\|y - m\| \le r_0$, smoothness gives

$$I(y) \le \frac{\sqrt{\kappa}}{2}\|y - m\| + \frac{\kappa}{2}\|y - m\|^2 \le \frac{1}{8} + \frac{1}{32} < 1.$$

We have proved

$$\mathbb{B}\left(m, \frac{1}{4\sqrt{\kappa}}\right) \subset K \subset \mathbb{B}(m, 2). \tag{C.40}$$

For $t \ge 1$, convexity and $I(m) = 0$ imply

$$K_t \subset m + t(K - m), \qquad \mathrm{vol}(K_t) \le t^d \,\mathrm{vol}(K). \tag{C.41}$$

Indeed, if $x \in K_t$, then $I(m + (x-m)/t) \le I(x)/t \le 1$. All these sets are compact, with positive finite volume, by strong convexity and (C.40).

The initial enclosure is now available. We next describe how symmetric vertices produce a valid cut or certify an inner ellipsoid. Points outside the known ball can be eliminated without querying the oracle; only the remaining vertices need statistical tests.

Maintain an ellipsoid $E = c + Q\mathbb{B}_d$, with Q invertible, starting from $c = m$, $Q = 2I_d$. Its $D = 2d$ test vertices are

$$v_{i,+} = c + \gamma Q e_i, \qquad v_{i,-} = c - \gamma Q e_i, \qquad 1 \le i \le d. \tag{C.42}$$

Whenever a cut with nonzero normal g at a test vertex v preserves K, its normalized offset is

$$\frac{g \cdot (v - c)}{\|Q^T g\|} = \frac{\pm\gamma\, g \cdot Q e_i}{\|Q^T g\|} \le \gamma < \alpha.$$

We weaken it to the shallow cut of offset α and apply Lemma C.3. The updated ellipsoid contains the retained part and its volume is at most $e^{-1/(8d)}$ times the previous volume. The interval update in that lemma covers $d = 1$.

First inspect the vertices without making oracle calls. If some v lies outside $\mathbb{B}(m, 2)$, use $g = v - m$ and the cut $g \cdot (y - v) \le 0$. This is a valid cut: for every $y \in \mathbb{B}(m, 2)$,

$$g \cdot (y - v) \le 2\|g\| - \|g\|^2 < 0.$$

It has a nonzero normal, preserves K by (C.39), and costs no gradient query. Update the ellipsoid and begin the next round.

Only if all vertices belong to $\mathbb{B}(m, 2)$ do we make a gradient-testing round. For each vertex define

$$z_{i,\pm} = \frac{m + v_{i,\pm}}{2}, \qquad w_{i,\pm} = z_{i,\pm} - m. \tag{C.43}$$

Then $\|w_{i,\pm}\| \le 1$. Relabel these D pairs as $(z_i, w_i)_{i=1}^{D}$, and introduce their true, unknown scores and energy:

$$s_i = \nabla F(z_i) \cdot w_i, \qquad M = \max\{1, s_1, \dots, s_D\}. \tag{C.44}$$

Lemma C.8 (A reliable test either stops or cuts safely)**.** *Suppose estimates $\widehat{g}_i$ satisfy*

$$\|\widehat{g}_i - \nabla F(z_i)\| \le \frac{M}{100} \qquad (1 \le i \le D). \tag{C.45}$$

Set $\widehat{s}_i = \widehat{g}_i \cdot w_i$ and $\widehat{M} = \max\{1, \widehat{s}_1, \dots, \widehat{s}_D\}$. If $\max_i \widehat{s}_i \le 4$, then the ellipsoid

$$\frac{m + c}{2} + \eta_d Q \mathbb{B}_d \tag{C.46}$$

is contained in K_5. Otherwise, choosing a maximizing index i, the cut with normal $\widehat{g}_i$ at $v_i = m + 2w_i$ preserves K.

Proof. The bound on the w_i implies

$$|\widehat{s}_i - s_i| \le M/100, \qquad |\widehat{M} - M| \le M/100. \tag{C.47}$$

If all estimated scores are at most 4, then $99M/100 \le \widehat{M} \le 4$, so $M \le 400/99 < 5$. Convexity relative to m gives $I(z_i) \le s_i \le M$. Thus the convex hull of the points z_i lies in K_5. It contains (C.46), because

$$\frac{1}{\sqrt{d}} \mathbb{B}_d \subset \operatorname{conv}\{e_1, -e_1, \dots, e_d, -e_d\}$$

and the test points have center $(m + c)/2$ and generating vectors $\gamma Q e_i/2$.

For the other case, fix a maximizing index, so $\widehat{s}_i = \widehat{M} > 4$. For every $y \in K$, convexity gives

$$\nabla F(z_i) \cdot (y - z_i) \le I(y) - I(z_i) \le 1 + B.$$

Writing $e_i = \widehat{g}_i - \nabla F(z_i)$ and using $\|y - m\| \le 2$, $\|w_i\| \le 1$, we obtain

$$\widehat{g}_i \cdot (y - m) \le \widehat{s}_i + 1 + B + 3\|e_i\| \le \widehat{s}_i + 1 + B + 3M/100.$$

By (C.47), $M \le 100\widehat{s}_i/99$, and hence

$$1 + B + 3M/100 \le 9/8 + \widehat{s}_i/33 < \widehat{s}_i.$$

The last inequality follows from $\widehat{s}_i > 4$. Therefore $\widehat{g}_i \cdot (y - v_i) < 0$. The normal is nonzero, since its inner product with w_i exceeds 4. □

The preceding lemma establishes correctness of an individual reliable test. To avoid paying its noise cost once for every possible cut, we must also count how often low-energy tests can occur. The following volume argument provides exactly that accounting.

Lemma C.9 (Reciprocal-energy bound)**.** *Consider the geometric execution up to termination or up to and including its first gradient-testing round that fails (C.45). The same conclusion holds if a round is declared bad for violating an additional promised batch-cost bound. Then*

$$\sum_{\text{testing rounds in this prefix}} \frac{1}{M} \le L_{\mathrm{geom},d}. \tag{C.48}$$

If every testing round is good, the execution terminates before H_κ outer rounds.

Proof. At the start of every round in the prefix, E contains K: all preceding cuts are valid by Lemma C.8 or by the free ball-cut calculation. At a testing round, the true inequalities $I(z_i) \le s_i \le M$ imply

$$\frac{m+c}{2} + \eta_d Q \mathbb{B}_d \subset K_M.$$

This conclusion does not require accuracy of the current estimates. Taking volumes and using (C.41) gives

$$\operatorname{vol}(E) \le C_{\mathrm{geom}} \operatorname{vol}(K_M) \le C_{\mathrm{geom}} M^d \operatorname{vol}(K). \tag{C.49}$$

Fix $k \ge 0$. If there is a testing round with $2^k \le M < 2^{k+1}$, consider the first such round in the prefix. At that time the volume ratio to K is at most $C_{\mathrm{geom}} 2^{d(k+1)}$. Each subsequent testing round, except possibly the last one in the prefix, produces a valid cut if the algorithm has not stopped. Each such cut decreases volume by a factor at most $e^{-1/(8d)}$; any intervening free cuts decrease it as well. The maintained ellipsoid cannot have volume below that of K. Thus the number of testing rounds in this energy bin is at most

$$2 + 8d \log C_{\mathrm{geom}} + 8d^2(k+1)\log 2.$$

The extra two safely include the initial round and a possible terminal or first bad round. Summing over bins yields

$$\begin{aligned}\sum \frac{1}{M} &\le \sum_{k\ge 0} 2^{-k}[2 + 8d \log C_{\mathrm{geom}} + 8d^2(k+1)\log 2] \\ &= 4 + 16d \log C_{\mathrm{geom}} + 32d^2 \log 2 = L_{\mathrm{geom},d}.\end{aligned}$$

Finally, on an all-good execution every cut preserves K. The initial-to-minimum possible volume ratio is at most $(8\sqrt{\kappa})^d$, by (C.40). There can therefore be at most $8d^2 \log(8\sqrt{\kappa})$ cuts. Every nonterminal outer round performs one cut, so the cap (C.34) cannot be reached on such an execution. □

The reciprocal-energy bound is expressed in terms of the unknown true scores. We now give the observable batching rule that achieves the corresponding cost and failure probability without knowing those scores. Its early-stop and late-error bounds mirror the center pilot.

Lemma C.10 (Adaptive batches at a finite collection of points)**.** *Condition on the complete incoming history of a testing round, so the points z_i and vectors w_i, $\|w_i\| \le 1$, are fixed. Let M be given by (C.44). For $0 < \zeta < 1$ put $\lambda = \max\{1, 2^{16} DA/\zeta\}$. There is a finite procedure returning estimates $\widehat{g}_1, \dots, \widehat{g}_D$ such that, with conditional probability at least $1 - \zeta/M$, both (C.45) and the call bound*

$$N_{\mathrm{batch}} \le 2D + \frac{8D\lambda}{M} \tag{C.50}$$

hold. On all histories it uses fewer than $4D\lambda$ calls.

Proof. For the successive batch sizes

$$n = 1, 2, 4, \dots, 2^{\lceil \log_2 \lambda \rceil},$$

take a fresh mean $\widehat{g}_{i,n}$ of n replies at each of the D points. All D means are completed before testing this stage. Set

$$\widehat{M}_n = \max\{1, \widehat{g}_{1,n} \cdot w_1, \dots, \widehat{g}_{D,n} \cdot w_D\}.$$

Stop at the first stage with $n\widehat{M}_n \ge \lambda$ and return its means. At the last stage $n \ge \lambda$, stopping is automatic. Since $\lambda \ge 1$, the sum of all possible batch sizes is less than 4λ, proving the all-history call bound.

By Lemma C.1, each vector mean has conditional squared error at most A/n. The query points and their true gradients remain fixed throughout the testing round, even though the reply distributions can change with the complete history.

Call a stop early if $n < \lambda/(2M)$. An early stop requires $\widehat{M}_n \ge \lambda/n > 2M$. Since $M \ge 1$ and all true scores are at most M, some mean then has error of norm at least $\lambda/(2n)$. Conditional Chebyshev

and a union bound over the D means give probability at most $4DAn/\lambda^2$ for this event at stage n. Summing over the early dyadic stages bounds the probability of any early stop by

$$\sum_{n<\lambda/(2M)} \frac{4DAn}{\lambda^2} \le \frac{4DA}{\lambda M}. \tag{C.51}$$

At a stage with $n \ge \lambda/(2M)$, the probability that some mean has error greater than $M/100$ is at most $10^4 DA/(nM^2)$. The sum of reciprocals of dyadic batch sizes starting at their first such stage is at most $4M/\lambda$. This also holds when the first stage is $n = 1$, because then $\lambda \le 2M$. Consequently the probability of any such late-stage error is at most

$$\frac{40000DA}{\lambda M}. \tag{C.52}$$

For adaptive execution these bounds are applied only to stages actually reached, conditioning at each stage entry. The fixed points and M do not change, so the same deterministic sums bound the resulting probabilities. No independence between different stages or different pointwise batches is needed.

Outside the events in (C.51) and (C.52), the return is at a non-early stage and all returned means satisfy (C.45). Moreover, a stage with $n \ge 2\lambda/M$, if reached, has $\widehat{M}_n \ge 99M/100$ and must stop. The deterministic final stage may stop sooner. Thus

$$n_{\mathrm{stop}} \le \max\{1, 4\lambda/M\}.$$

The sum of fresh batch sizes up to this stage is at most $2n_{\mathrm{stop}}$, which proves (C.50). The total exceptional probability is at most

$$\frac{40004DA}{\lambda M} \le \frac{\zeta}{M}.$$

Only $\lambda \ge 2^{16} DA/\zeta$ is needed in this estimate. The additional floor $\lambda \ge 1$ makes the same finite schedule valid at zero noise, where every estimation error is zero almost surely. □

We now combine the test, the reciprocal-energy bound, and the adaptive batches. The remaining point is to impose finite public budgets without losing the success probability: an all-good uncapped execution must fit inside those budgets.

Proof of Proposition C.7. Use the following algorithm, with the constants in (C.35). Start at $E = \mathbb{B}(m, 2)$ and a zero call count. Allow at most H_κ outer rounds. In each round perform the vertex inspection and, if needed, the free ball cut described above. Otherwise apply Lemma C.10 at the D testing points. If every estimated score is at most 4, return the parameters

$$z = m + \frac{c - m}{10}, \qquad \mathsf{T} = \frac{\gamma}{10\sqrt{d}} Q. \tag{C.53}$$

If the test does not stop, choose a maximizing index and perform the gradient cut of Lemma C.8. Before each physical call, check the cap B_{geom}; if the call would exceed it, return $z = m$, $\mathsf{T} = I_d$. Return this same fallback if the outer-round cap is reached. Ties in selecting vertices or indices are resolved by their fixed index order.

For the success analysis, first consider the version with the finite outer-round and per-round batch caps but without the global call cap. Declare a batch good if both its accuracy and its call bound in Lemma C.10 hold. At any testing round whose prior batches were good, the conditional probability that the current batch is bad is at most ζ/M. By Lemma C.9, the sum of $1/M$ along a good prefix, including a possible first bad round, is at most $L_{\mathrm{geom},d}$ on every path. Taking conditional expectations of the indicators of first failure therefore gives

$$\mathbb{P}(\text{some batch is bad before termination}) \le \zeta L_{\mathrm{geom},d} = \delta. \tag{C.54}$$

This is a first-failure union bound under the complete history; it does not condition future replies on simultaneous success events.

On the complementary event, every cut preserves K, and the algorithm terminates by a successful test before the outer-round cap. Its physical call count is bounded by

$$\sum_{\text{testing rounds}} (2D + 8D\lambda/M) \le 2DH_\kappa + 8D\lambda L_{\mathrm{geom},d} \le B_{\mathrm{geom}}.$$

Hence adding the global cap cannot change any all-good execution. For a formal comparison, couple the capped and uncapped versions through their common prefix, continuing the oracle kernels on the counterfactual uncapped history after a cap. Both finite procedures have identical histories before the cap. Since the uncapped version cannot reach this cap on its all-good event, the implemented version has success probability at least $1 - \delta$ as well.

At successful termination, Lemma C.8 supplies (C.46). Contract this ellipsoid toward m by a factor $1/5$. Convexity and $I(m) = 0$ imply $z + \mathsf{T}\mathbb{B}_d \subset K_1$, with the parameters in (C.53). Also $K_1 \subset E$, and $m \in E$ implies $\|Q^{-1}(c - m)\| \le 1$. For $x \in E$,

$$\begin{aligned}\|\mathsf{T}^{-1}(x - z)\| &\le \|\mathsf{T}^{-1}(x - c)\| + \|\mathsf{T}^{-1}(c - z)\| \\ &\le \frac{10\sqrt{d}}{\gamma} + \frac{9\sqrt{d}}{\gamma} < \rho.\end{aligned}$$

This proves (C.36).

For the all-history conclusions, every accepted testing ellipsoid has all vertices in $\mathbb{B}(m, 2)$, whether or not its estimates are accurate. Averaging opposite vertices gives $\|c - m\| \le 2$, and subtracting them gives $\|Qe_i\| \le 2/\gamma$. Therefore

$$\|Q\|_{\mathrm{op}} \le \frac{2\sqrt{d}}{\gamma}, \qquad \|z - m\| \le \frac{1}{5}, \qquad \|\mathsf{T}\|_{\mathrm{op}} \le \frac{1}{5}.$$

The fallback $z = m$, $\mathsf{T} = I_d$ obeys (C.37). The matrices remain invertible by Lemma C.3. Every selected cut has a nonzero normal even on a bad history: a ball cut has $\|v - m\| > 2$, and a gradient cut has $\widehat{g}_i \cdot w_i > 4$. All replies, averages, and geometric operations are finite on the specified real-valued paths, and both loop and call counts are capped. All choices and stage boundaries are deterministic functions of the public inputs and the physical transcript. This completes the proof. □

Rounding supplies only an ellipsoid sandwich, not a sampler for the target. We turn it into a full-support proposal whose normalization and random-generation rule are explicit. The same rule must remain valid when the rounding procedure returns its fallback.

Let $v_d = \mathrm{vol}(\mathbb{B}_d)$ and define

$$D_d = 1 + d\sum_{j=0}^{d-1} \binom{d-1}{j} 4^{j+1} j!. \tag{C.55}$$

For every output (z, T) of the preceding procedure, including a fallback, set

$$U(x) = \frac{1}{4}\left(\frac{\|\mathsf{T}^{-1}(x - z)\|}{\rho} - 1\right)_+, \tag{C.56}$$

$$q(x) = \frac{e^{-U(x)}}{v_d |\det \mathsf{T}| \rho^d D_d}. \tag{C.57}$$

The success-probability benchmark used subsequently is the public constant

$$p_0 = \frac{e^{-11}}{\rho^d D_d} > 0. \tag{C.58}$$

Proposition C.11 (Domination, acceptance, and all-history moments)**.** *The formula* (C.57) *is a probability density on every history and can be sampled with a finite number of internal random draws and no oracle queries. On the good event of Proposition C.7,*

$$I(x) + B \ge 2U(x) \qquad (x \in \mathbb{R}^d). \tag{C.59}$$

In particular, with $h(x) = I(x) - U(x) + 10$,

$$h(x) \geq 10 - B + U(x) > 1, \qquad 0 < 1 + I(x) \leq 2(1 + h(x)), \tag{C.60}$$

and the ideal acceptance probability satisfies

$$\int q(x) e^{-h(x)}\, dx \geq p_0. \tag{C.61}$$

On all histories, for $Y_q \sim q$,

$$\begin{aligned} \mathbb{E}_q[\|Y_q - m\|^2] &\leq 2 + 2\rho^2[1 + 8d + 16d(d+1)], \\ \mathbb{E}_q[\log(1 + \|Y_q - m\|)] &\leq \log\big(2 + \rho(1 + 4d)\big). \end{aligned} \tag{C.62}$$

Proof. First verify normalization and give a sampling rule. In coordinates $x = z + \rho \mathsf{T} y$, the unnormalized density is $\exp[-(\|y\| - 1)_+/4]$. Its integral equals

$$\begin{aligned} v_d + d v_d \int_1^\infty r^{d-1} e^{-(r-1)/4}\, dr &= v_d + d v_d \sum_{j=0}^{d-1} \binom{d-1}{j} \int_0^\infty s^j e^{-s/4}\, ds \\ &= v_d D_d. \end{aligned}$$

The last integral is $4^{j+1} j!$, by substitution and repeated integration by parts. The change of variables supplies the factor $\rho^d |\det \mathsf{T}|$ in (C.57).

Draw a direction Θ uniformly on the Euclidean unit sphere. Independently select a radial component as follows. With probability $1/D_d$, set $R = U_0^{1/d}$ for $U_0 \sim \mathrm{Unif}(0,1)$. For each $j \in \{0, \dots, d-1\}$, with probability

$$\frac{d\binom{d-1}{j} 4^{j+1} j!}{D_d}, \tag{C.63}$$

set $R = 1 + 4\sum_{\ell=1}^{j+1} E_\ell$, where the E_ℓ are independent unit exponentials. Return $Y_q = z + \rho \mathsf{T} R \Theta$. The interior component has the radial law of the uniform unit ball; the other components expand $(1 + s)^{d-1} e^{-s/4}$ exactly as above. Their weights sum to one, proving the claimed sampling rule. For $d = 1$, the spherical direction is an independent uniform sign. For general d, normalize a standard Gaussian vector. Generate its coordinates in independent pairs using $\sqrt{-2\log U_1}(\cos(2\pi U_2), \sin(2\pi U_2))$, with fresh open-interval uniforms U_1, U_2, discarding one coordinate when needed. On the probability-zero zero-vector event define a fixed unit direction. Each exponential is $-\log U$ for another fresh open-interval uniform. The number of radial exponential draws is at most d, and there is no internal rejection loop or convex-body sampling oracle.

Now assume the good sandwich (C.36). Write $E_{\mathrm{out}} = z + \rho \mathsf{T} \mathbb{B}_d$ and $t(x) = \|\mathsf{T}^{-1}(x - z)\|/\rho$. For $t(x) > 1$, let $y = m + \theta(x - m)$ be the intersection of the segment from m to x with the boundary of E_{out}. The point m lies in the interior of K_1 and hence of E_{out}, so $0 < \theta < 1$. Continuity implies $I(y) \geq 1$: if $I(y) < 1$, an open neighborhood of y would lie in $K_1 \subset E_{\mathrm{out}}$, contradicting its location on the boundary.

Put $u_m = (\rho \mathsf{T})^{-1}(m - z)$, so $\|u_m\| \leq 1$. The reverse triangle inequality gives

$$1 = \|(1 - \theta)u_m + \theta(\rho \mathsf{T})^{-1}(x - z)\| \geq \theta(t(x) + 1) - 1.$$

Thus $\theta \leq 2/(t(x) + 1)$. Convexity and $I(m) = 0$ yield $1 \leq I(y) \leq \theta I(x)$, and hence

$$I(x) \geq \frac{t(x) + 1}{2} \geq \frac{t(x) - 1}{2} = 2U(x).$$

Inside E_{out}, $U = 0$ and $I \geq -B$. This proves (C.59). It implies $h \geq 10 - B + U > 1$. Also $I \geq -B$ gives $1 + I > 0$, and $U \leq (I + B)/2$ gives $h \geq I/2 + 10 - B/2$, from which $1 + I \leq 2(1 + h)$ follows. This proves (C.60).

The inner ellipsoid $z + \mathsf{T}\mathbb{B}_d$ lies in K_1, so

$$\int_{\mathbb{R}^d} e^{-I(x)}\, dx \geq e^{-1} v_d |\det \mathsf{T}|.$$

Using the cancellation between q and e^{-h},

$$\int q(x)e^{-h(x)}\,dx = \frac{e^{-10}}{v_d|\det \mathsf{T}|\rho^d D_d}\int e^{-I(x)}\,dx \geq \frac{e^{-11}}{\rho^d D_d} = p_0.$$

Finally, the explicit radial mixture gives, on every history,

$$\mathbb{E}[R] \leq 1 + 4d, \qquad \mathbb{E}[R^2] \leq 1 + 8d + 16d(d+1).$$

For a tail component, these follow from the mean k and second moment $k(k+1)$ of a sum of $k \leq d$ unit exponentials; the interior component has $0 \leq R \leq 1$. By (C.37), $\|Y_q - m\| \leq 1 + \rho R$. The inequality $(a+b)^2 \leq 2a^2 + 2b^2$ gives the first moment bound in (C.62); Jensen's inequality for the concave logarithm gives the second. Neither calculation uses the good event or any property of the oracle's past errors. □

The preceding proposal bounds can now be coupled with the center construction. We record the successful-pilot event using only the implemented histories, so later conditioning never imposes a success requirement on future oracle replies. Run the center procedure with $\delta_m = \varepsilon/8$, then run the rounding procedure with $\delta_q = \varepsilon/8$ at the returned center. Use their specified fallbacks and physical-call caps on every history. The two procedures use at most

$$T_{\text{pilot}} \leq C_d\big(1 + \log(1+\kappa) + A/\varepsilon\big) \tag{C.64}$$

physical calls on every history. The output satisfies $\|m\| \leq 2$; the proposal parameters always satisfy the moment bounds of Proposition C.11.

Fix a potential F and an admissible oracle. We define the successful-pilot event directly from the implemented, capped procedures.

Let $\mathcal{S}_m$ be the event that the center procedure returns through the certificate test (C.20), and let $\mathcal{S}_q$ be the event that the rounding procedure returns through its successful score test with the parameters (C.53). Both events exclude fallback returns caused by the geometric-round or physical-call caps.

Let $\mathcal{I}_m$ and $\mathcal{I}_q$ be the finite random index sets of estimation subroutines actually completed by the center and rounding procedures, respectively. For $i \in \mathcal{I}_m$, let $\mathcal{E}_i^m$ be the completed-subroutine event $\mathcal{E}_{\text{est}}$ defined in the proof of Lemma C.4. For $j \in \mathcal{I}_q$, let $\mathcal{E}_j^q$ be the event that the actual completed subroutine satisfies both (C.45) and (C.50), with its corresponding true energy M. Define

$$\mathcal{G}_m = \mathcal{S}_m \cap \bigcap_{i\in\mathcal{I}_m} \mathcal{E}_i^m, \qquad \mathcal{G}_q = \mathcal{S}_q \cap \bigcap_{j\in\mathcal{I}_q} \mathcal{E}_j^q, \qquad \mathcal{G} = \mathcal{G}_m \cap \mathcal{G}_q.$$

An intersection over an empty index set is understood to be the whole sample space.

All these conditions refer to queries, replies, and decisions that have actually occurred. For fixed F, the true gradients appearing in the estimation events are deterministic continuous functions of the recorded query points. The events need not be computable by the algorithm to be measurable. Since the two pilots use no private random draws, their states, completed subroutines, and return branches can be reconstructed from the physical transcript. Consequently,

$$\mathcal{G}_m \in \mathcal{F}_{T_{\text{center}}}, \qquad \mathcal{G}_q, \mathcal{G} \in \mathcal{F}_{T_{\text{pilot}}}.$$

The auxiliary executions without the global physical-call cap in the proofs of Propositions C.2 and C.7 are used to establish probability bounds for these actual events. Under the common-prefix coupling, an auxiliary execution on which every estimation event is good has total cost within the implemented budget and terminates through the certificate branch. It therefore coincides with the implemented execution and implies the corresponding event $\mathcal{G}_m$ or $\mathcal{G}_q$. Those proofs give

$$\mathbb{P}(\mathcal{G}_m^c) \leq \delta_m, \qquad \mathbf{1}_{\mathcal{G}_m}\mathbb{P}(\mathcal{G}_q^c \mid \mathcal{F}_{T_{\text{center}}}) \leq \delta_q \mathbf{1}_{\mathcal{G}_m} \quad \text{almost surely.}$$

Here the rounding guarantee applies on $\mathcal{G}_m$, because the returned center satisfies its required potential-gap promise. The tower property therefore yields

$$\mathbb{P}(\mathcal{G}^c) \leq \delta_m + \delta_q = \varepsilon/4. \tag{C.65}$$

On $\mathcal{G}$, $F(m) - F(x_\star) \leq B$, and the proposal has ideal acceptance mass at least $p_0 > 0$, where p_0 depends only on d.

For clarity, the analytical quantities used for acceptance are

$$I(x) = F(x) - F(m), \qquad h(x) = I(x) - U(x) + 10, \qquad \alpha_*(x) = e^{-h(x)}. \tag{C.66}$$

The pilot does not evaluate these quantities. Its proved domination $I \geq 2U - B$, with $U \geq 0$, implies

$$h \geq 10 - B > 1, \qquad h \geq I/2 + 10 - B/2, \qquad 1 + I \leq 2(1 + h). \tag{C.67}$$

Also, with $p_* := \int q\alpha_*$,

$$p_* \geq p_0, \qquad \frac{q(x)\alpha_*(x)}{p_*} = \frac{e^{-F(x)}}{\int_{\mathbb{R}^d} e^{-F(y)} dy}. \tag{C.68}$$

These assertions refer to a fixed successful pilot history; proposal normalization and moment bounds remain valid on all histories.

C.3 Directional envelopes and acceptance at every noise level

Fix a proposed point $x \neq m$ and set

$$r = \|x - m\|, \qquad s = \frac{x - m}{r} \in \mathbb{R}^d, \qquad w(t) = \langle s, \nabla F(m + st)\rangle, \quad 0 \leq t \leq r. \tag{C.69}$$

The function w is increasing and κ-Lipschitz, for every such unit direction s. The following bound allows its positive and negative parts to be treated separately even when $\|\nabla F(m)\|$ is large.

Lemma C.12 (Absolute gradient mass along a segment)**.** *On a good pilot history, put $V_I = 1 + I(x)$. Then*

$$\int_0^r |w(t)|\, dt \leq I(x) + 2B \leq V_I, \qquad r^2 \leq 4(I(x) + 2B) \leq 4V_I, \qquad V_I \geq \frac{61}{64}. \tag{C.70}$$

Proof. Since w is increasing, its negative part occupies an initial segment, possibly empty or all of $[0, r]$. Its negative mass is exactly the decrease of the potential from m to the minimum along this segment:

$$\int_0^r \max\{-w(t), 0\}\, dt = F(m) - \min_{0 \leq t \leq r} F(m + st) \leq B.$$

Using $\int_0^r w = I(x)$ gives the first inequality. Strong convexity at $x_\star$ gives

$$\|x - x_\star\|^2 \leq 2(I(x) + B), \qquad \|m - x_\star\|^2 \leq 2B.$$

The triangle inequality and $(u + v)^2 \leq 2u^2 + 2v^2$ yield the stated bound on r^2. Finally, $I \geq -B$, $2B < 1$, and $B = 3/64$ give the remaining assertions. □

This absolute-mass bound suggests separate upper envelopes for the positive derivative and the reflected negative derivative. The next lemma explains why a geometric grid controls their integrals with only a constant terminal remainder.

Lemma C.13 (A deterministic geometric envelope)**.** *Let $v : [0, r] \to [0, \infty)$ be nondecreasing and κ-Lipschitz, with $r > 0$. Define*

$$J = \max\left\{0, \left\lceil \log_2(r\sqrt{\kappa}) \right\rceil\right\}, \qquad d_j = r2^{-j}, \qquad t_j = r - d_j \quad (0 \leq j \leq J). \tag{C.71}$$

On each ordinary interval (t_{j-1}, t_j), $1 \leq j \leq J$, place a rectangle of height $v(t_j)$. On the terminal interval (t_J, r), place a rectangle of height $v(r)$. Their union is an upper envelope g_{exact} satisfying

$$\int_0^r g_{\mathrm{exact}}(t)\, dt \leq 2 \int_0^r v(t)\, dt + \frac{1}{2}. \tag{C.72}$$

Proof. The definition gives $d_J \leq \kappa^{-1/2}$. Monotonicity gives the upper-envelope property. If $J = 0$, there is only a terminal rectangle, and Lipschitz continuity gives

$$rv(r) \leq \int_0^r v(t)\, dt + \frac{\kappa r^2}{2} \leq \int_0^r v(t)\, dt + \frac{1}{2}.$$

Suppose $J \geq 1$. An ordinary interval has length d_j, while its successor has length $d_j/2$. For $1 \leq j \leq J-1$,

$$d_j v(t_j) \leq 2 \int_{t_j}^{t_{j+1}} v(t)\, dt.$$

The last ordinary rectangle satisfies $d_J v(t_J) \leq \int_{t_J}^{r} v(t)\, dt$. The terminal rectangle satisfies

$$d_J v(r) \leq \int_{t_J}^{r} v(t)\, dt + \frac{\kappa d_J^2}{2}.$$

Summing these inequalities charges every integral on the right at most twice and leaves an error at most $1/2$, proving (C.72). □

Partition-point values do not affect these integrals. To fix the algorithm on all real-valued paths, at an interior partition point we use the maximum of the two adjacent rectangle heights, and at 0 or r the corresponding one-sided value. We use the same convention for the estimated rectangle functions below. It preserves an upper-envelope inequality whenever that inequality holds on the adjacent intervals.

Constructing the two noisy envelopes. Let $0 < \delta < 1$ be a public error budget, to be fixed below. We apply the grid (C.71) to the two nonnegative, nondecreasing, κ-Lipschitz functions

$$v_+(t) = \max\{w(t), 0\}, \qquad v_-(u) = \max\{-w(r-u), 0\}. \tag{C.73}$$

For each of the two grids and each ordinary node $j = 1, \ldots, J$, use

$$e_j = \frac{2^{j/2}}{r}, \qquad \delta_j = \frac{\delta}{8}\, 2^{-j/2}, \qquad \beta_j = \max\left\{1, \left\lceil \frac{8Ar^2}{\delta} 2^{-j/2} \right\rceil \right\}. \tag{C.74}$$

For the terminal node, use

$$e_E = \frac{1}{d_J}, \qquad \delta_E = \frac{\delta}{8}, \qquad \beta_E = \max\left\{1, \left\lceil \frac{8Ad_J^2}{\delta} \right\rceil \right\}. \tag{C.75}$$

Execute the positive grid first, in increasing node order, followed by its terminal node. At its ordinary node t_j, query $m + st_j$ exactly β_j times and take the inner product of its vector average with s. At its terminal node, query $m + sr = x$ exactly β_E times and again project the vector average onto s. Denote a resulting signed average by $\widehat{w}$. The height used on the corresponding ordinary or terminal rectangle is $\max\{0, \widehat{w} + e\}$, with its corresponding e_j or e_E.

Next execute the negative grid in the same order. At its node t_j, query $m + s(r - t_j)$ exactly β_j times and project the vector average onto $-s$; at its terminal node, query m exactly β_E times and project its vector average onto $-s$. This estimates the signed quantity $-w(r-u)$ before its positive part is taken. Use the same rectangle-height rule. The grids use separate batches; coincident query points are still queried and charged separately.

Let $\widehat{v}_+(t)$ and $\widehat{v}_-(u)$ be the resulting rectangle functions, with the stated partition-point convention, and put

$$\begin{aligned} g_+(t) &= \widehat{v}_+(t), \qquad g_-(t) = \widehat{v}_-(r-t), \qquad g(t) = g_+(t) + g_-(t), \\ Z_g &= \int_0^r g(t)\, dt, \qquad Z_- = \int_0^r g_-(t)\, dt. \end{aligned} \tag{C.76}$$

The function g is nonnegative, finite, and piecewise constant on every real-valued reply history. The algorithm computes its integral by summing rectangle areas.

Lemma C.14 (Accuracy and cost of the noisy envelopes)**.** *Condition on the complete history before this envelope construction, including the point $x \neq m$. With conditional probability at least $1 - \delta$, all signed averages are within their assigned tolerances. Denote this event by $\mathcal{E}$. On $\mathcal{E}$,*

$$g(t) \geq |w(t)| \quad (0 \leq t \leq r), \qquad Z_g \leq 2 \int_0^r |w(t)|\, dt + 16. \tag{C.77}$$

More precisely, $g_+ \geq w_+$ and $g_- \geq w_-$, and each envelope has integral at most twice its corresponding true mass plus 8. Consequently, on a good pilot history and $\mathcal{E}$,

$$Z_- \leq 2B + 8, \qquad Z_g \leq 2I(x) + 4B + 16. \tag{C.78}$$

On every real-valued history, the number of calls used by both grids is bounded by

$$N_{\text{grid}} \leq 2J + 2 + \frac{56Ar^2}{\delta}. \tag{C.79}$$

The grid call count depends only on x, the pilot values, and public parameters, not on the observed heights.

Proof. The two types of batch size can equivalently be written as

$$\beta_j = \max\{1, \lceil A/(\delta_j e_j^2)\rceil\}, \qquad \beta_E = \max\{1, \lceil A/(\delta_E e_E^2)\rceil\}.$$

The lower bound of one makes all means well defined also when $A = 0$. Lemma C.1 applies to each signed average, even when its oracle kernels depend on replies from earlier nodes. Therefore

$$\mathbb{P}(\mathcal{E}^c \mid \text{incoming history}) \leq 2\left(\sum_{j=1}^{\infty} \delta_j + \delta_E\right) = \frac{\delta}{4}(\sqrt{2}+2) < \delta. \tag{C.80}$$

Here $\sum_{j\geq 1} 2^{-j/2} = \sqrt{2}+1$.

If $|\widehat{w} - a| \leq e$, then

$$\max\{a, 0\} \leq \max\{\widehat{w} + e, 0\} \leq \max\{a, 0\} + 2e.$$

Consequently each estimated rectangle is an upper bound for the corresponding true function on $\mathcal{E}$. In one grid the excess area from ordinary-node estimation is at most

$$2\sum_{j=1}^{J} d_j e_j \leq 2\sum_{j=1}^{\infty} 2^{-j/2} = 2(\sqrt{2}+1) < 5.$$

The terminal-node excess area is $2d_J e_E = 2$. Adding the deterministic error $1/2$ from Lemma C.13 gives an excess less than 8 for each grid beyond twice the corresponding true integral. Reflection preserves the negative-part integral, proving (C.77). The negative true mass is at most B by the proof of Lemma C.12, which also proves (C.78).

Finally, using $\max\{1, \lceil a\rceil\} \leq 1 + a$ for $a \geq 0$ gives

$$N_{\text{grid}} \leq 2J + 2 + \frac{16Ar^2}{\delta}\left(\sum_{j=1}^{J} 2^{-j/2} + 2^{-2J}\right).$$

The coefficient in parentheses is at most $\sqrt{2}+2$, and $16(\sqrt{2}+2) < 55 < 56$. This proves (C.79) on all histories, without using $\mathcal{E}$ or a good pilot. □

The envelope lemma controls the two parts of the target derivative. Acceptance also subtracts the derivative of the proposal potential; we therefore need a pointwise comparison along the same ray, not only an integrated potential bound.

Lemma C.15 (Derivative domination along a proposed ray)**.** *Fix a good pilot and a unit vector s, and set $u(t) = U(m + st)$ and $w(t) = \langle s, \nabla F(m + st)\rangle$ for every $t \geq 0$. Give u' the value zero in the flat part and at the boundary of the outer proposal ellipsoid. Then*

$$u'(t) \geq 0, \qquad w(t) \geq 2u'(t) \quad \textit{whenever } u(t) > 0. \tag{C.81}$$

The stated version of u' agrees with the derivative almost everywhere and is computable from the proposal parameters without potential values.

Proof. Put $a = (\rho\mathsf{T})^{-1}(m - z)$ and $b = (\rho\mathsf{T})^{-1}s$. The good-pilot sandwich is $z + \mathsf{T}\mathbb{B}_d \subset K_1 \subset z + \rho\mathsf{T}\mathbb{B}_d$. Since $I(m) = 0 < 1$, continuity gives a neighborhood of m in K_1; thus $\|a\| < 1$. Let $t_0 > 0$ be the positive exit time determined by $\|a + t_0 b\| = 1$. Since $b \neq 0$, this time is finite and unique. The triangle inequality gives $t_0\|b\| \leq 2$. Before exit, $u' = 0$. After exit, the norm of the affine ray is increasing, and the explicit proposal formula gives

$$0 \leq u'(t) = \frac{\langle b, a + tb\rangle}{4\|a + tb\|} \leq \frac{\|b\|}{4} \leq \frac{1}{2t_0}.$$

At the boundary, $I(m + st_0) \geq 1$. Otherwise continuity would give a neighborhood of that boundary point contained in K_1, contradicting the outer containment. Convexity along the ray implies $t_0 w(t_0) \geq I(m + st_0) \geq 1$. Monotonicity of w then proves (C.81) for $t > t_0$. At the boundary our convention sets $u' = 0$; the single chosen value does not change any integral. The proof also covers $t_0 > r$, when the entire segment used by the algorithm is in the flat part.

For operational purposes evaluate the displayed derivative formula when $\|a + tb\| > 1$ and use zero otherwise. This prescription is finite and measurable on every history, including a failed pilot. Its correctness inequalities are used only on good pilots. □

The derivative comparison lets us compensate the negative part of the target derivative while retaining a nonnegative marking intensity. We introduce the computable position density first, and only then the unknown intensity used to analyze its marks.

For the remainder of the acceptance construction fix $c = 1/4$. After the grids have been completed, define the known density numerator and its normalizer by

$$a_0(t) = g_+(t) + g_-(t) + \frac{2c}{r}, \qquad Q = Z_g + 2c = \int_0^r a_0(t)\, dt. \tag{C.82}$$

These definitions are valid on every history: $Q \geq 2c > 0$ and $a_0(t) > 0$. To sample from a_0/Q, assign weight $2c/Q$ to the uniform distribution on $(0, r)$ and assign each envelope rectangle its area divided by Q. Reflect the negative-grid rectangles into the t coordinate. Omit zero-weight components. The weights sum to one; overlapping rectangles simply contribute additively to the density. There are finitely many components on every finite reply history.

For analysis, introduce the unknown intensity and its integral

$$b_0(t) = w(t) + g_-(t) - u'(t) + \frac{c}{r}, \qquad \ell = I(x) + Z_- - U(x) + c = \int_0^r b_0(t)\, dt. \tag{C.83}$$

The last equality uses $U(m) = 0$, so it is asserted on a good pilot. Neither b_0 nor ℓ is evaluated by the algorithm.

Lemma C.16 (Margins and a uniform mark probability)**.** *On a good pilot and a completed good envelope history,*

$$\frac{c}{r} \leq b_0(t) \leq a_0(t) - \frac{c}{r} \qquad (0 \leq t \leq r). \tag{C.84}$$

Moreover,

$$\ell \geq c, \qquad r^2 \leq 8\ell, \qquad Q \leq 68\ell, \qquad p := \frac{\ell}{Q} \geq \frac{1}{68} =: p_{\min}. \tag{C.85}$$

Proof. Inside the outer ellipsoid, $u' = 0$ and $w + g_- \geq 0$. Outside it, Lemma C.15 gives $w - u' \geq 0$. These facts prove the lower margin. For the upper margin use $a_0 - b_0 = g_+ - w + u' + c/r \geq c/r$. Integrating gives $\ell \geq c$.

The good-pilot inequality $U \leq (I + B)/2$ yields

$$I \leq 2\ell - 2Z_- + B - 2c.$$

Combine this with (C.70) to obtain

$$r^2 \leq 4(I + 2B) \leq 8\ell - 8Z_- + 12B - 8c \leq 8\ell,$$

since $Z_- \geq 0$ and $3B \leq 2c$. By (C.78),

$$Q \leq 2I + 4B + 16 + 2c \leq 4\ell - 4Z_- + 6B + 16 - 2c \leq 4\ell + 6B + 16 - 2c.$$

As $\ell \geq 1/4$, the ratio Q/ℓ is at most $4 + 4(6B + 16 - 2c) = 67.125 < 68$. This proves all the claims. □

On all histories, set the computable prefactor probability to

$$v = \min\{1, \exp(Z_- + c - 10)\}. \tag{C.86}$$

On a good pilot and good envelope, $Z_- + c \le 2B + 8 + 1/4 < 10$, so clipping is inactive. The ideal identity is then

$$ve^{-\ell} = \exp[-I(x) + U(x) - 10] = \alpha_*(x). \tag{C.87}$$

Indeed, if N has the Poisson distribution with mean Q, and a point with density a_0/Q is independently marked with probability $b_0(t)/a_0(t)$, its mean mark probability is $p = \ell/Q$. The probability of no marks at N such independent points is $\mathbb{E}[(1-p)^N] = e^{-Qp} = e^{-\ell}$. This is an analytical identity; the implementable finite construction below does not evaluate the unknown intensity or generate an unbounded number of nodes. The random negative-envelope mass Z_- cancels exactly in (C.87).

The ideal identity reduces sampling to Bernoulli marks, but the oracle provides only unbiased estimates of their probabilities. The two-sided margins above are what make clipping affordable: its bias is controlled by a variance rather than a standard deviation.

For a real number y, write $\mathrm{clip}_{[0,1]}\, y = \min\{1, \max\{0, y\}\}$.

Lemma C.17 (Clipping with a two-sided margin)**.** *Let Y be square integrable, with mean $z \in [\eta, 1-\eta]$, where $\eta > 0$. Then*

$$\left|\mathbb{E}[\mathrm{clip}_{[0,1]}(Y)] - z\right| \le \frac{\mathrm{Var}(Y)}{\eta}. \tag{C.88}$$

Proof. Clipping changes Y only on the disjoint events $Y < 0$ and $Y > 1$. On either event the absolute change is at most $(Y-z)^2/\eta$: for example, if $Y < 0$, then $z - Y \ge \eta$ and $-Y \le z - Y \le (z-Y)^2/\eta$. The upper event is identical after reflection about $1/2$. Taking expectations and using $\mathbb{E}[Y] = z$ proves the assertion. □

For each correction node that is reached, draw a fresh position t with density a_0/Q. Make n fresh vector-gradient calls at $m + st$, average them, and project onto s to obtain $\overline{w}$. After all these replies, use a new uniform to make a rejection mark with probability

$$\mathrm{clip}_{[0,1]}\left(\frac{\overline{w} + g_-(t) - u'(t) + c/r}{a_0(t)}\right). \tag{C.89}$$

This is a legal probability on every history. The density a_0/Q is held fixed throughout this trial's correction stage, even if subsequent oracle kernels adapt to the entire history.

The clipping estimate applies conditionally at each fresh position. Integrating against the fixed position density cancels its denominator, which is the key step in the following relative-error bound.

Lemma C.18 (Conditional mark error)**.** *Fix a good pilot and a completed good envelope history. Let p_j be the actual mark probability conditional on any complete history before a fresh correction position is generated. Such a history may include the already drawn Poisson category, prior marks, and all ignored marker replies. Then, almost surely,*

$$|p_j - p| \le \eta_0 := \frac{Ar^2}{ncQ}, \qquad \frac{\eta_0}{p} \le \frac{32A}{n}. \tag{C.90}$$

Proof. Condition also on the fresh position t. Lemma C.1 gives conditional mean $w(t)$ and variance at most A/n for $\overline{w}$, with no independence assumption on the n oracle replies. The untruncated fraction in (C.89) has mean $b_0(t)/a_0(t)$, variance at most $A/(na_0(t)^2)$, and distance at least $c/(ra_0(t))$ from each endpoint of $[0,1]$, by (C.84). Lemma C.17 therefore bounds its clipping bias by $Ar/(nca_0(t))$. The fresh marking uniform realizes the clipped expectation as a mark probability. Integrating over t with density a_0/Q bounds the unconditional node bias by

$$\int_0^\tau \frac{Ar}{nca_0(t)} \frac{a_0(t)}{Q}\, dt = \frac{Ar^2}{ncQ}.$$

Division by $p = \ell/Q$, followed by $r^2 \le 8\ell$ and $c = 1/4$, gives the second inequality.

The good envelope event concerns only completed calls. Fixing such a history does not condition on future oracle errors. After any subsequent marker reply, the complete-history oracle assumption and the same bounds apply again. In particular, revealing a privately drawn category before the node does not change this proof. □

The relative mark-error bound is uniform over all reached histories. We now select its accuracy budget and cap the correction procedure. The proof will propagate conditional survival probabilities instead of assuming that noisy marks are independent.

Use the public parameters

$$\tau = \frac{p_0\varepsilon}{16}, \qquad \delta = \frac{\tau}{2}, \qquad n = \max\left\{1, \left\lceil \frac{256A}{\tau} \right\rceil\right\}, \\ H = \left\lceil 136 \log \frac{4}{\tau} \right\rceil, \qquad K = \left\lceil \frac{2}{p_0} \log \frac{16}{\varepsilon} \right\rceil. \tag{C.91}$$

Here $p_0 = e^{-11}/(\rho^d D_d)$ is the constant for the ideal acceptance probability with offset 10. Use this δ for each envelope construction and the same fixed n for every correction node. Write

$$\beta = \frac{32A}{n} \le \frac{\tau}{8} < \frac{1}{2}. \tag{C.92}$$

Thus, on a good envelope, every reached node has conditional mark probability at least $(1-\beta)p \ge 1/136$.

After a successful prefactor coin, generate a finite-valued category

$$N_c \overset{\text{law}}{=} \min\{\mathrm{Pois}(Q), H+1\}. \tag{C.93}$$

Specifically, assign mass $e^{-Q}Q^j/j!$ to each $j = 0, \dots, H$ and the remaining mass to $H+1$, and use one fresh uniform to select a category. Only these finitely many weights are computed. If the prefactor fails, reject immediately. If it succeeds, category 0 accepts immediately and category $H+1$ rejects immediately. For $1 \le N_c \le H$, process at most N_c nodes using (C.89); reject on the first mark, and accept if every node is unmarked. The charged marker calls specified below are made before a trial is left, including in the immediate branches.

If $x = m$, omit the grids, category generation, and every formula involving $1/r$. Use a fresh uniform to accept with probability $\min\{1, \exp(U(m) - 10)\}$. On a good pilot this equals $e^{-10} = \alpha_*(m)$.

Lemma C.19 (Conditional acceptance error)**.** *Fix a potential F and an admissible oracle. For $1 \le k \le K$, let R_k be the event that trial k is reached. On R_k, let H_k be the complete history immediately before its fresh proposal Y_k is generated, and let B_k be its acceptance indicator. Thus H_k includes all internal randomness already generated, but no future random draws.*

For an incoming history h, write m_h, q_h, U_h for the center, proposal density, and proposal potential determined by its completed pilot, and define

$$\alpha_{*,h}(x) = \exp[-F(x) + F(m_h) + U_h(x) - 10].$$

When $\mathbb{P}(R_k) > 0$, choose a jointly measurable version of the regular conditional acceptance probability

$$\alpha_k(h, x) = \mathbb{P}(B_k = 1 \mid R_k, H_k = h, Y_k = x).$$

This conditional probability averages over the initial marker reply, the envelopes, the prefactor and category, correction nodes, all remaining marker replies, and future internal random draws. Define the subprobability measure on incoming histories by

$$\lambda_k(D) = \mathbb{P}\big(\mathcal{G} \cap R_k \cap \{H_k \in D\}\big).$$

Then, for λ_k-almost every h and for $q_h(x)\,dx$-almost every x,

$$|\alpha_k(h, x) - \alpha_{*,h}(x)| \le \tau. \tag{C.94}$$

If $\mathbb{P}(R_k) = 0$, the assertion is vacuous. The exceptional history sets may depend on F, the oracle, and k; the exceptional proposal set may also depend on h.

At $x = m_h$, we may choose the version

$$\alpha_k(h, m_h) = \min\{1, \exp(U_h(m_h) - 10)\},$$

consistently with the specified zero-radius branch. Since q_h has a Lebesgue density, this changes the conditional function only on a set of zero joint history–proposal measure.

When the trial index and incoming history are understood, we retain the notation $\alpha_{\mathcal{H}}(x) = \alpha_k(h, x)$ and $\alpha_*(x) = \alpha_{*,h}(x)$, and write $q = q_h$.

Proof. The finite histories of this algorithm can be encoded in a countable disjoint union of finite-dimensional Borel spaces, so regular conditional probabilities may be chosen. All conditional estimates below are interpreted under the law of the reached kth trial and disintegrated with respect to (H_k, Y_k).

On R_k, the event $\mathcal{G}$ is determined by H_k. Restricting to good incoming histories therefore imposes no additional success condition on future oracle errors. The conditional oracle bounds are used only almost surely at actually reached calls. Conditional expectation and disintegration transfer the ensuing estimates to λ_k-almost every incoming history and $q_h(x)\,dx$-almost every proposal. No pointwise oracle guarantee on exceptional histories is required.

For $x \neq m$, condition first on the initial marker reply and on a completed good envelope history. The parameters p, Q, and v are then fixed for analysis. Condition also on successful prefactor selection, a category $N_c = j \leq H$, and the reply to its control marker. Lemma C.18 applies to every subsequent complete history. Let s_l be the probability that the first l nodes are all unmarked, with $s_0 = 1$. On the event of surviving $l - 1$ nodes, the next conditional mark probability differs from p by at most η_0. Therefore

$$|s_l - (1 - p)s_{l-1}| \leq \eta_0 s_{l-1} \leq \eta_0.$$

Iterating this scalar inequality gives

$$|s_j - (1 - p)^j| \leq \eta_0 \sum_{l=0}^{j-1} (1 - p)^l \leq \frac{\eta_0}{p} \leq \beta \leq \frac{\tau}{8}. \tag{C.95}$$

The formula includes $j = 0$, with zero error. It uses neither independent oracle replies nor hypothetical replies after the algorithm has stopped. Averaging over the control reply and the categories $j \leq H$ preserves the error bound $\tau/8$. The ideal Poisson survival mass omitted by the overflow category is at most

$$\sum_{j>H} e^{-Q} \frac{Q^j}{j!} (1 - p)^j \leq (1 - p)^H \leq e^{-H/68} \leq \frac{\tau}{4}. \tag{C.96}$$

The full ideal sum is $e^{-\ell}$. Multiplying by $v \leq 1$ and using (C.87) proves an acceptance error at most $3\tau/8$ on this completed good envelope history.

Conditional on the history before the grids, their failure probability is at most $\delta = \tau/2$ by (C.80). Both actual and ideal acceptance probabilities are in $[0, 1]$ on the entire fixed good pilot, so failed envelopes contribute at most δ. Averaging the completed grid histories and the initial marker reply gives an error at most $3\tau/8 + \tau/2 < \tau$. At $x = m$, the explicit branch has zero error. Final marker replies cannot alter the previously selected acceptance bit or proposal. This proves the assertion. □

The preceding lemma controls one trial after any good incoming pilot history. We now specify how trials are concatenated and how every private decision is recorded in the physical transcript; the next lemma then converts the one-trial bound into a bound on the final law.

Execute the two pilots and fix their m, z, T, q, U. Attempt at most K trials, each with a fresh proposal from q. In normalized coordinates, implement a reached trial as follows. All marker replies are charged and ignored, but belong to the complete history conditioning of every later oracle call.

1. Generate $x \sim q$ by the finite radial mixture of Proposition C.11, and make one initial marker call at x.
2. If $x = m$, draw the zero-radius acceptance bit specified above and proceed directly to the final marker in step 5. Otherwise construct the two grids in their prescribed order, with (C.74) and (C.75).

3. Generate the prefactor coin with probability (C.86). If it fails, set the trial's acceptance bit to zero and make a control marker call at $0e_1$. If it succeeds, generate N_c by (C.93), and make the control marker call at $(N_c+1)e_1$. Category 0 sets the acceptance bit to one; category $H+1$ sets it to zero. These branches then proceed to step 5.
4. For a passed prefactor with $1 \le N_c \le H$, generate each reached node position freshly from a_0/Q, make exactly n calls there, and draw its rejection mark by (C.89). Make one node marker call at be_1, where $b=1$ denotes a mark and $b=0$ its absence. A mark ends the node stage and sets the trial's acceptance bit to zero. After N_c unmarked nodes, set that bit to one.
5. Make one final marker call at ae_1, where $a \in \{0,1\}$ is the trial's acceptance bit. Only after its reply, if $a=1$, stop and return $Y = x/\sqrt{\mu}$. If $a=0$ on trial K, stop and return $Y = m/\sqrt{\mu}$. Otherwise begin the next trial.

Thus an immediate decision in the probabilistic description always includes the control and final markers prescribed by its branch. The algorithm queries no potential values, mode, Hessian, or target sample. Section C.4 proves that its stopping rule is measurable in the smaller physical filtration.

Although trial acceptance probabilities may depend on earlier replies, the accepted law after each reached good-pilot history is close to the same target. The argument below retains both normalizing constants and then mixes over the first accepting trial.

Lemma C.20 (Total variation under adaptive trials)**.** *The complete sampler satisfies*

$$\|\mathcal{L}(Y) - \nu_f\|_{\mathrm{TV}} \le \frac{7\varepsilon}{16}. \tag{C.97}$$

Proof. Throughout this proof, Y denotes the algorithm's output in the original coordinates, and

$$Y_{\mathrm{norm}} := \sqrt{\mu}\, Y$$

denotes its output before inverse scaling. Thus Y_{norm} is the first accepted proposal, or m if all K trials reject. Write

$$\nu_F(dz) = \frac{e^{-F(z)}}{\int_{\mathbb{R}^d} e^{-F(u)}\, du}\, dz$$

for the target distribution in normalized coordinates. We first bound the total variation distance between the law of Y_{norm} and ν_F, and then transfer this bound to the original coordinates. All statements below conditional on a completed pilot history or a reached trial history are understood almost surely under the corresponding disintegrated laws. For the finitely many trials, we discard the union of the exceptional null sets supplied by Lemma C.19. Fix a good completed pilot history. For any subsequent complete history $\mathcal{H}$ before a reached trial's proposal, put $p_{\mathcal{H}} = \int q(x)\alpha_{\mathcal{H}}(x)\, dx$. The ideal acceptance mass (C.68) and Lemma C.19 give

$$|p_{\mathcal{H}} - p_*| \le \tau, \qquad p_{\mathcal{H}} \ge p_0 - \tau \ge p_0/2. \tag{C.98}$$

Conditional on this trial accepting, its proposal density is $q\alpha_{\mathcal{H}}/p_{\mathcal{H}}$. Retaining both normalizers, its total variation distance from ν_F is at most

$$\begin{aligned}
&\frac{1}{2}\int_{\mathbb{R}^d} \left| \frac{q(x)\alpha_{\mathcal{H}}(x)}{p_{\mathcal{H}}} - \frac{q(x)\alpha_*(x)}{p_*} \right| dx \\
&\quad \le \frac{1}{2p_{\mathcal{H}}} \left(\int_{\mathbb{R}^d} q(x)|\alpha_{\mathcal{H}}(x) - \alpha_*(x)|\, dx + |p_* - p_{\mathcal{H}}| \right) \\
&\quad \le \frac{\tau}{p_0 - \tau} = \frac{\varepsilon}{16 - \varepsilon} \le \frac{\varepsilon}{8}.
\end{aligned} \tag{C.99}$$

All preceding rejection decisions and ignored replies belong to $\mathcal{H}$. Thus this bound holds after any such history, even if the oracle kernels changed adaptively.

Disintegrate over the index of the first accepting trial and its incoming history. Earlier rejections are already determined by that history, and conditioning its current trial on acceptance gives exactly the law bounded in (C.99). The final marker reply changes neither the selected bit nor the recorded proposal. Consequently, conditional on the fixed good pilot history and on at least one acceptance,

the law of Y_{norm} is a mixture of laws each within $\varepsilon/8$ of ν_F. Repeated conditional expectation in (C.98) also yields

$$\mathbb{P}(\text{all } K \text{ trials reject} \mid \text{good pilot history}) \le (1 - p_0/2)^K \le e^{-p_0 K/2} \le \varepsilon/16. \tag{C.100}$$

If all trials reject, the normalized output is $Y_{\mathrm{norm}} = m$. This fallback contributes at most the probability above to total variation. Therefore, conditional on almost every good pilot history, the law of Y_{norm} is within $3\varepsilon/16$ of ν_F.

The good-pilot event $\mathcal{G}$ is measurable at the completion of the pilot and imposes no condition on future noise. Mixing over good pilot histories preserves the preceding bound, while failed pilot histories contribute at most $\mathbb{P}(\mathcal{G}^c) \le \varepsilon/4$. Hence

$$\|\mathcal{L}(Y_{\mathrm{norm}}) - \nu_F\|_{\mathrm{TV}} \le \frac{3\varepsilon}{16} + \mathbb{P}(\mathcal{G}^c) \le \frac{7\varepsilon}{16}.$$

Finally, since $F(z) = f(z/\sqrt{\mu})$, the map $z \mapsto z/\sqrt{\mu}$ sends ν_F to ν_f. This map is a measurable bijection with a measurable inverse, so total variation is invariant under it. Since $Y = Y_{\mathrm{norm}}/\sqrt{\mu}$, we obtain

$$\|\mathcal{L}(Y) - \nu_f\|_{\mathrm{TV}} = \|\mathcal{L}(Y_{\mathrm{norm}}) - \nu_F\|_{\mathrm{TV}} \le \frac{7\varepsilon}{16}.$$

This proves the assertion. □

Remark C.21. The failure probability of the envelopes is included in each trial's conditional acceptance error. It is not union-bounded over all K trials. Random integration positions are used directly to generate marks; they create no approximation error when $A = 0$. Only the clipping bias from the oracle noise is bounded in Lemma C.18.

C.4 Operational completion and endpoint cases

Proposition C.22 (Upper bound and operational admissibility)**.** *For every $A \ge 0$, the complete sampler is measurable, has a physical-transcript stopping time, terminates on every specified finite-real-reply path with open-interval uniforms, and obeys*

$$\|\mathcal{L}(Y) - \nu_f\|_{\mathrm{TV}} \le 7\varepsilon/16, \qquad \mathbb{E}[T] \le C_d S.$$

The cost includes failed pilot histories and every ignored reply.

Proof. Lemma C.20 already proves the accuracy assertion. We verify execution, physical stopping, and the uniform expected cost.

Finite execution and measurability. Both pilots have deterministic limits on geometric updates and physical calls. Every matrix update is nonsingular, and every fallback preserves the stated bounded, nonsingular proposal geometry. The pilots' averages, free cuts, and budget decisions depend only on public parameters and physical replies. Proposal generation uses a finite mixture and finitely many open-interval uniforms and Gaussian coordinates. Its specified fixed-direction convention on the zero-Gaussian-vector event completes the definition on every path.

For a finite proposal with $r > 0$, the grid length J and all positive batch sizes are finite. Every envelope height and its integral are finite on a finite-real-reply path. The correction normalizer satisfies $Q \ge 2c > 0$, and the position density a_0/Q is everywhere positive. The category generation in (C.93) uses exactly $H + 2$ categories and a finite list of explicitly given weights. It does not require an unbounded Poisson simulation or the generation of any unused future correction nodes. Every reached node has a fixed finite batch of n calls; there are at most H such nodes and at most K trials. The zero-radius branch contains no division by zero. The clipped probabilities remain defined on all bad histories. Thus every specified path terminates after finitely many calls.

All operations are measurable: matrix arithmetic and norms on nonsingular matrices, finite sums, exponential and logarithmic functions, comparisons with fixed tie-breaking, ceilings, and finite mixtures. Countably many possible finite grid and batch lengths give countably many measurable branches. Free exact real arithmetic is understood in the sense of the model; no numerical floating-point claim is needed.

Parsing the physical transcript. First reconstruct the two pilots from their deterministic rules, including free ellipsoid updates and budget fallbacks. After their last reply, the parser knows m, z, T and all public acceptance parameters. The initial marker of a trial reveals x, and hence whether $x = m$. In the zero-radius branch, the next call is the final acceptance marker. Otherwise the proposal fixes r, J and every grid batch size, so the parser recognizes the end of the two grids without any private uniforms.

The next call is the control marker. Its location is $0e_1$ after a failed prefactor, or $(N_c + 1)e_1$ after a successful prefactor. Thus it announces the category as well as whether the prefactor passed. A failed prefactor, category 0, or category $H + 1$ has no correction nodes and is followed immediately by the final marker. In a category $1 \le N_c \le H$, each node consists of exactly n gradient calls followed by one mark-bit marker. A marked node ends the node stage; otherwise it continues until N_c nodes have been completed. The next call is then the final acceptance-bit marker. The parser can identify all these positions and count trials from the transcript alone.

Every stop occurs after a final marker reply, when its acceptance bit is one or the public trial cap has been reached. Thus $\{T = t\}$ is measurable in the physical transcript after reply t. Spatial coincidences between data points and marker locations cause no ambiguity: the current stage and position specify the role of a call. No proposal, category, or acceptance uniform needs to be adjoined to the physical filtration. Ignored replies cannot change a bit already selected or the recorded output. They may influence later oracle kernels; the conditional proofs explicitly include them.

Trial counts and grid costs. Write N_{tr} for the number of attempted trials and set

$$L_\kappa = 1 + \log(1 + \kappa), \qquad R_\varepsilon = 1 + \log(1/\varepsilon).$$

Since p_0 depends only on d, the public choices imply

$$n + 1 \le C_d(1 + X), \qquad H \le C_d R_\varepsilon, \qquad K \le C_d R_\varepsilon, \qquad A/\delta = C_d X. \tag{C.101}$$

Here the final equality denotes the explicit positive dimension-dependent factor $32/p_0$ multiplying X and remains valid when $A = X = 0$. There are at most three non-node markers per trial: initial, control, and final. Every correction node costs $n + 1$ calls including its marker. Consequently, on every history, a reached trial has the bound

$$C_{\mathrm{trial}} \le N_{\mathrm{grid}} + 3 + (n + 1)H, \tag{C.102}$$

where $N_{\mathrm{grid}} = 0$ if $x = m$. We do not use this crude bound on all good trials; that would retain an unnecessary factor H.

For $r > 0$ the grid lemma gives $N_{\mathrm{grid}} \le 2J + 2 + 56Ar^2/\delta$, and

$$J \le 1 + \frac{\log \kappa}{2 \log 2} + \frac{\log(1 + r)}{\log 2}.$$

Set $J = 0$ when $r = 0$. Conditional on any history before a reached proposal, the new proposal has the fixed pilot's density q. Its second radius moment and logarithmic radius moment are bounded by constants depending only on d, on every pilot history, including fallback histories. Therefore

$$\mathbb{E}[N_{\mathrm{grid}} \mid \text{incoming trial history}] \le C_d(L_\kappa + X) = C_d S. \tag{C.103}$$

On any good completed pilot history, (C.98) gives conditional acceptance probability at least $p_0/2$ in every reached trial. Repeated conditioning implies $\mathbb{E}[N_{\mathrm{tr}} \mid \text{good pilot history}] \le 2/p_0$. On a bad pilot use only $N_{\mathrm{tr}} \le K$. It follows that

$$\mathbb{E}[N_{\mathrm{tr}}] \le \frac{2}{p_0} + \mathbb{P}(\mathcal{G}^c)K \le \frac{2}{p_0} + \frac{\varepsilon}{4}K \le C_d. \tag{C.104}$$

The event of reaching a trial is determined before its new proposal. Multiply (C.103) by this event's indicator, take expectations, and sum over the trial index. The total expected grid cost is at most $C_d S\mathbb{E}[N_{\mathrm{tr}}] \le C_d S$. The total cost of non-node markers is at most $3\mathbb{E}[N_{\mathrm{tr}}] \le C_d$.

Correction nodes after a good pilot. Fix a good completed pilot and any history before a reached trial. On a completed good envelope, condition on a passed prefactor, any category $N_c = j \le H$, and its control marker reply. By (C.90) and (C.92), every reached node has conditional mark probability at least $1/136$. Hence the expected number of processed nodes is at most

$$\sum_{l\ge 0}(1-1/136)^l = 136.$$

This bound holds uniformly in j; category $H+1$ uses no nodes. It also holds when the prefactor is not conditioned to pass, since a failed prefactor has no nodes. On failed envelope histories use the bound H instead. Their conditional probability, before the envelope construction, is at most δ, independently of any estimate of its realized heights. Averaging gives the conditional node-count bound

$$\mathbb{E}[N_{\text{nodes}} \mid \text{incoming good-pilot trial history}] \le 136 + \delta H \le C_d. \tag{C.105}$$

The same bound covers the zero-radius branch. The inequality $\delta H \le C_d$ follows from boundedness of $\tau\log(4/\tau)$ over the stipulated range of τ. Multiply by $n+1$ and sum over reached trials as above. Conditional on any good pilot the total expected node cost is at most $C_d(n+1) \le C_d(1+X)$.

Correction nodes after a bad pilot. On a bad pilot there need not be a positive lower bound on the mark probability. Both public caps must be retained in this part of the argument. If T_{nodes} counts all correction calls and their node markers, then

$$\begin{aligned}\mathbb{E}[\mathbf{1}_{\mathcal{G}^c}T_{\text{nodes}}] &\le \mathbb{P}(\mathcal{G}^c)KH(n+1)\\ &\le C_d\varepsilon R_\varepsilon^2(1+X) \le C_d(1+X).\end{aligned} \tag{C.106}$$

The last step uses the uniform boundedness of $\varepsilon[1+\log(1/\varepsilon)]^2$ on $(0,1/10]$. The event $\mathcal{G}^c$ is already measurable at the end of the pilot; this estimate makes no independence assumption about later replies. In particular, it does not treat bad-pilot trials as if they enjoyed the good-pilot mark bound.

Adding the grid and marker bounds, the two node-cost contributions, and the pathwise pilot bound (C.64) proves $\mathbb{E}[T] \le C_d S$. All grid ceilings, node means, ignored replies, zero-radius branches, and failure histories have been charged. No bound on the expectation of a noisy envelope height was used. Normalization preserves both call count and TV, completing the proof. □

The complete construction now proves the main upper-bound theorem; its two endpoint consequences are treated separately below.

Completion of the proof of Theorem 2.2. We now specify the parameters left symbolic in Algorithms 1 and 2. Use $B = 3/64$, $b_{\text{acc}} = 10$, and the dimension-dependent constants

$$\gamma = \frac{1}{8d}, \qquad \rho = \frac{20\sqrt{d}}{\gamma}, \qquad D_d = 1 + d\sum_{j=0}^{d-1}\binom{d-1}{j}4^{j+1}j!, \qquad p_0 = \frac{e^{-11}}{\rho^d D_d}. \tag{C.107}$$

Choose the failure budgets, correction batch size, and caps as

$$\begin{aligned}\delta_m = \delta_q = \varepsilon/8, \qquad c = 1/4, \qquad \tau = p_0\varepsilon/16, \qquad \delta = \tau/2,\\ n = \max\{1, \lceil 256A/\tau\rceil\}, \qquad H = \lceil 136\log(4/\tau)\rceil,\\ N_{\text{overflow}} = H+1, \qquad K = \left\lceil \frac{2}{p_0}\log\frac{16}{\varepsilon}\right\rceil.\end{aligned} \tag{C.108}$$

These agree with (C.91); τ and p_0 are calibration quantities, not additional unknown inputs. Implement the two pilots by Propositions C.2 and C.7, including their caps, fallbacks, and tie-breaking. The grid sizes, padding, and batches are (C.71)–(C.76). Use the finite radial mixture in Proposition C.11 and the rectangle mixture following (C.82). For each fixed input Q, FiniteCategory uses the following weights

$$\mathbb{P}(N_c = j) = e^{-Q}\frac{Q^j}{j!} \quad (0 \le j \le H), \qquad \mathbb{P}(N_c = N_{\text{overflow}}) = 1 - \sum_{j=0}^{H} e^{-Q}\frac{Q^j}{j!}. \tag{C.109}$$

Thus its output has precisely the law (C.93) without simulating an uncapped Poisson count. All choices depend only on public parameters and completed replies. With $b_{\mathrm{acc}} = 10$, the zero-radius probability, the prefactor, and the correction marks in the main-body pseudocode coincide with the ones analyzed above. The algorithms therefore implement exactly that capped sampler, including every marker call and its timing. Proposition C.22 proves its accuracy, physical stopping, termination, and unconditional expected cost. It is therefore admissible in (2.4), and taking the infimum over admissible algorithms gives the asserted upper bound. □

The finite-accuracy construction is now complete. At zero oracle variance the marking bias disappears, so removing the two rejection caps gives a separate exact algorithm. Its proof must establish almost-sure termination and expected cost anew.

Proof sketch and intuition for Corollary 5.1. At zero variance every queried gradient is exact almost surely. The pilots and directional envelopes then succeed without a statistical failure event; each batch has size one. Remove the outer trial cap and Poisson overflow cap and use the exact marking probability $b_0(t)/a_0(t)$ in (C.89). The Poisson acceptance identity (3.4) then yields exact acceptance. Every accepted sample has law ν_f. Acceptance mass is at least p_0, so the trial count is geometrically bounded, while early rejection keeps the expected number of processed nodes constant even for large Poisson intensity. Expected grid and pilot costs are $O(\log(1+\kappa))$. This is a separate construction, not substitution of $\varepsilon = 0$ into an approximation theorem. Almost-sure termination and bounded expected cost do not assert an accuracy-independent deterministic query budget.

Proof of Corollary 5.1.

Proof. Run the modified center and rounding procedures with fixed budgets $\delta_m = \delta_q = 1/8$. Since the conditional oracle variance is zero, every reply equals the exact gradient almost surely. This statement holds simultaneously for all calls that are made: there are only countably many possible physical call indices, so the union of their null error events is null. Every reached pilot estimation is therefore accurate, its good cost promise holds, and its cap cannot preempt successful completion. Both pilots are correct almost surely, and their combined call bound is $C_d \log(1+\kappa)$. Their thresholds have lower bound one and their noise-dependent parts are zero.

Use the same proposal and both envelopes, now with the fixed budget $\delta = 1/2$. Every grid batch has size one and every signed estimate is exact almost surely; the nonnegative padding remains in the construction. The conclusions of Lemmas C.14, C.15, and C.16 consequently hold almost surely in every reached trial. Use one exact gradient call at each correction node, so the mark probability in (C.89) is exactly $b_0(t)/a_0(t)$. It already belongs to $[0,1]$, by (C.84).

Replace the finite category by a genuine Poisson variable $N \sim \mathrm{Pois}(Q)$, and omit the overflow branch. For definiteness, generate N by inverse CDF from a fresh uniform on $(0,1)$; for finite Q this returns a finite nonnegative integer for every such uniform. Following a successful prefactor, process at most N nodes, with immediate rejection after the first mark. If all N nodes are unmarked, accept. Category $N = 0$ accepts without any correction nodes. Remove the outer trial cap and keep trying until an acceptance occurs. All clipped-probability and zero-radius conventions are retained on exceptional histories.

Conditional on every exact completed grid, the fresh positions and marking uniforms give

$$\mathbb{P}(\text{accept} \mid x, \text{completed grids}) = v\mathbb{E}[(1-p)^N] = ve^{-\ell} = \alpha_*(x).$$

The expression is independent of the realized envelope mass because of (C.87). Conditional on each incoming trial history, the accepted proposal therefore has exactly the target law by (C.68), and the trial acceptance mass is $p_* \geq p_0$. Repeated conditioning gives $\mathbb{P}(N_{\mathrm{tr}} > k) \leq (1-p_0)^k$ and $\mathbb{E}[N_{\mathrm{tr}}] \leq 1/p_0$. Thus an acceptance occurs almost surely, and disintegrating over its index and preceding history proves exactness of the returned sample.

For each fixed N, the expected number of processed nodes is at most $\sum_{l\geq 0}(1-p)^l = 1/p \leq 68$. This is uniform in N, so averaging the Poisson count preserves the bound; no bound on $\mathbb{E}[Q]$ is needed. Each node has one data call and one mark-bit marker. The grid count is at most $2J+2$. The proposal's uniform logarithmic moment gives conditional expected grid cost at most $C_d \log(1+\kappa)$ in every reached trial. The constantly many other markers and the bounded node cost have the same

upper order. Summing over reached trials using their geometric tail, and adding the pilots, gives the claimed uniform expected cost.

The physical implementation is the earlier parser with the genuine integer N announced by the control marker at $(N+1)e_1$ after a passed prefactor, and $0e_1$ after a failed one. The parser processes exactly one known-size node block at a time, ending the node stage after a marked node or after N unmarked nodes. A final acceptance marker is always charged before leaving the trial. The zero-radius branch still has its initial and final markers. Stop only after a successful final marker. Thus the query count remains a stopping time for the physical filtration. Every Poisson count and each individual trial is finite almost surely, and the outer geometric tail proves almost-sure termination of the whole procedure.

An infinite sequence of rejected trials can be specified as an exceptional internal-randomness path. No claim of termination on that path is made for this uncapped exact algorithm. For every positive ε, the capped algorithm of Proposition C.22 instead retains termination on every specified real-valued path and has the same expected asymptotic cost when $A = 0$. □

The general proof is not needed at $\kappa = 1$: the target is a Gaussian with an unknown shift. Directly averaging that shift gives the sharper query bound in Corollary 5.2.

Proof sketch and intuition for Corollary 5.2. In normalized coordinates $F(x) = F(0) + \langle b, x\rangle + \|x\|^2/2$. Average $n = \max\{1, \lceil A/(2\varepsilon)\rceil\}$ replies at zero and return $Z - \widehat{b}$, with an independent $Z \sim N(0, I_d)$. The mean error has zero expectation and second moment at most A/n. In a Taylor expansion of the Gaussian density, the first-order shift cancels after averaging. The remaining TV error is at most $A/(2n)$, not merely proportional to the standard deviation. This explains the inverse-linear accuracy order already for quadratic potentials. The full proof below gives the density calculation. The same construction works without any localization assumption when $L = \mu$.

Proof of Corollary 5.2.

Proof. When $\kappa = 1$, the two Bregman inequalities coincide. Setting their second argument to zero gives

$$F(x) = F(0) + \langle b, x\rangle + \tfrac{1}{2}\|x\|^2, \qquad b = \nabla F(0), \quad \|b\| \le 1.$$

The target is therefore $N(-b, I_d)$. Query zero a deterministic number $n = \max\{1, \lceil A/(2\varepsilon)\rceil\}$ of times and let $\widehat{b}$ be the mean reply. Lemma C.1 gives

$$\mathbb{E}[\widehat{b}] = b, \qquad \mathbb{E}[\|\widehat{b} - b\|^2] \le A/n.$$

Draw an independent $Z \sim N(0, I_d)$ using private randomness and return $Z - \widehat{b}$. The stopping count is the deterministic number n, so it is a physical stopping time.

To bound the error, write $e = \widehat{b} - b$ and let ϕ_d be the standard Gaussian density. For fixed e, Taylor's formula gives

$$\phi_d(y+e) - \phi_d(y) - \langle \nabla\phi_d(y), e\rangle = \int_0^1 (1-t)e^\top \nabla^2\phi_d(y+te)e\, dt.$$

For each t, rotation and translation invariance imply

$$\int_{\mathbb{R}^d} \left|e^\top \nabla^2\phi_d(y+te)e\right| dy = \|e\|^2\mathbb{E}[|Z_1^2 - 1|] \le 2\|e\|^2.$$

The integrated absolute Taylor remainder is at most $\|e\|^2$. Upon averaging, the linear term vanishes because $\mathbb{E}[e] = 0$. Tonelli's theorem and the finite second moment consequently give

$$\|\mathcal{L}(Z - e) - N(0, I_d)\|_{\mathrm{TV}} \le \tfrac{1}{2}\mathbb{E}[\|e\|^2] \le \frac{A}{2n} \le \varepsilon.$$

Translation proves the required target guarantee. This argument uses no independence assumption on successive replies; their conditional mean and variance suffice. The lower bounds follow from Theorem 2.3, proved independently in Section D. Finally, if $A = 0$, the single reply equals b almost surely, so the same construction is exact and the lower bound one is attained. □

D Proof of the Lower Bound

We separate the noise, curvature, and constant obstructions before combining them at the same public parameter tuple. In this appendix, $N^\star$ abbreviates the minimax quantity (2.4); normalization preserves that quantity and gives the parameters κ, A, X, S defined in Section 2.

D.1 Hard-family embedding for vector algorithms

For each scalar potential h, use the completed potential and oracle

$$f_h(x) = h(x_1) + \tfrac{1}{2}\|x_{2:d}\|^2, \qquad \mathcal{G}_h(x) = (G_h(x_1), x_2, \dots, x_d).$$

The Bregman remainder is the scalar remainder plus $\|x_{2:d} - y_{2:d}\|^2/2$, so the curvature bounds and localization hold in $\mathbb{R}^d$. The total vector variance equals the scalar variance. The target is the scalar law times a fixed standard Gaussian, so its first marginal is exactly the scalar law and projection contracts TV.

In the two proofs below, u denotes the first coordinate of an arbitrary vector query; ν_j, Q_j, Y_j denote first-coordinate target laws, output laws, and outputs. The algorithm and its stopping count remain fully vector-valued. All extra reply coordinates are deterministic functions of the current query, identical across instances. Therefore the rare-message coupling preserves the entire vector transcript until its first informative reply. Fixing the complete private tape makes the vector transcript tree have at most the same three replies per node. The proofs consequently establish lower bounds on the full quantity (2.4), including algorithms that query outside the first coordinate axis. No assertion about preserving a projected physical stopping filtration is needed.

D.2 The stochastic-noise lower bound

We work in the normalized model throughout this subsection. Thus the curvature bounds are 1 and κ, the first coordinate of the mode belongs to $[-1, 1]$ (the others are zero), and the advertised conditional variance ceiling is A. Recall that $A \geq 0$, $0 < \varepsilon \leq 1/10$, and $X = A/\varepsilon$. All instances below have the same public parameters. In particular, the algorithm receives no instance label.

Proposition D.1 (Full-variance stochastic-noise lower bound)**.** *For every public tuple in the stated regime,*

$$N^\star \geq \frac{A}{24\varepsilon} = \frac{X}{24}. \tag{D.1}$$

For $A > 0$, the witnesses attain the advertised curvature endpoints and use fresh, conditionally unbiased gradient replies whose conditional variance is exactly A. The bound allows arbitrary adaptive vector-valued queries and random stopping with expected query cost.

Proof. If $A = 0$, the asserted lower bound is zero and is immediate. Assume $A > 0$ for the rest of this proof; no lower bound on its positive value will be used. We first construct the potentials and separate their normalized target laws. We then couple two valid stochastic oracles up to their first informative reply, without imposing a deterministic query horizon.

Potentials and class membership. Set

$$\tau_0 = 4\varepsilon, \qquad H_{\mathrm{a}}(u) = \frac{u^2}{2} + \frac{\kappa - 1}{2}(|u| - 8)_+^2, \qquad F_\pm(u) = H_{\mathrm{a}}(u) \pm \tau_0 u, \tag{D.2}$$

where $v_+ = \max\{v, 0\}$. The function

$$r_0(u) = \operatorname{sgn}(u)(|u| - 8)_+, \qquad \operatorname{sgn}(0) = 0,$$

is continuous, nondecreasing, and one-Lipschitz. Indeed, it equals $u + 8$ on the left tail, zero in the core, and $u - 8$ on the right tail. Consequently

$$H_{\mathrm{a}}'(u) = u + (\kappa - 1) r_0(u),$$

and, for $u > v$,

$$u - v \leq H_{\mathrm{a}}'(u) - H_{\mathrm{a}}'(v) \leq \kappa(u - v). \tag{D.3}$$

Integrating along the segment between two points, in either orientation, gives the strong-convexity and smoothness inequalities with constants $1, \kappa$. The linear tilts in (D.2) leave their Bregman remainders unchanged. Thus $F_\pm$ belong to the required C^1 class. On a nonempty interval in the core the curvature is exactly 1, and on an interval strictly outside the core it is exactly κ. When $\kappa = 1$, the two advertised endpoints coincide.

The unique minimizers of F_- and F_+ are τ_0 and $-\tau_0$, respectively. They lie in the core, solve the derivative equations, and obey $\tau_0 = 4\varepsilon \le 2/5 < 1$. Both instances therefore satisfy the localization promise. Their gradient gap is independent of the query:

$$|F'_+(u) - F'_-(u)| = \Delta := 2\tau_0 = 8\varepsilon \qquad (u \in \mathbb{R}). \tag{D.4}$$

In particular, arbitrarily remote queries cannot amplify this gap.

Separation after normalization. Let $\nu_\pm$ be the probability measures with densities proportional to $e^{-F_\pm}$. Their normalizers are finite and positive because $H_{\mathrm{a}}(u) \ge u^2/2$ and the remaining terms are linear. Evenness of H_{a} makes the two normalizers equal to

$$2\int_0^\infty e^{-H_{\mathrm{a}}(u)} \cosh(\tau_0 u)\, du.$$

It follows, with the normalizers retained, that

$$\nu_+((-\infty, 0]) - \nu_-((-\infty, 0]) = \frac{\displaystyle\int_0^\infty e^{-H_{\mathrm{a}}(u)} \sinh(\tau_0 u)\, du}{\displaystyle\int_0^\infty e^{-H_{\mathrm{a}}(u)} \cosh(\tau_0 u)\, du}. \tag{D.5}$$

The numerator is nonnegative. Restrict it to $[0, 8]$ and use $\sinh(\tau_0 u) \ge \tau_0 u$. Enlarge the denominator by removing the nonnegative additional anchor term. This bounds the ratio below by

$$\frac{\tau_0 \displaystyle\int_0^8 u e^{-u^2/2}\, du}{\displaystyle\int_0^\infty e^{-u^2/2} \cosh(\tau_0 u)\, du} = \tau_0 (1 - e^{-32}) e^{-\tau_0^2/2} \sqrt{\frac{2}{\pi}}. \tag{D.6}$$

For the denominator identity, expand the hyperbolic cosine, complete the squares, and reflect one integral. The Gaussian integral itself follows from Tonelli's theorem and polar coordinates:

$$\left(\int_{\mathbb{R}} e^{-u^2/2}\, du\right)^2 = \int_0^{2\pi} \int_0^\infty e^{-r^2/2} r\, dr\, d\theta = 2\pi.$$

We record elementary numerical bounds to keep the constant explicit. For $0 \le u \le 1$, $(1 - u^2 + u^4)(1 + u^2) = 1 + u^6 \ge 1$, so

$$\pi = 4\int_0^1 \frac{du}{1 + u^2} \le 4\int_0^1 (1 - u^2 + u^4)\, du = \frac{52}{15} < \frac{32}{9}.$$

The first identity follows from $u = \tan\theta$ on $[0, \pi/4]$. Therefore $\sqrt{2/\pi} > 3/4$. Since $\tau_0 \le 2/5$,

$$e^{-\tau_0^2/2} \ge 1 - \tau_0^2/2 \ge \frac{23}{25}, \qquad 1 - e^{-32} > \frac{99}{100}.$$

The last inequality follows from $e > 2$ and $2^{32} > 100$. Combining these estimates with (D.5) and (D.6) yields

$$\|\nu_+ - \nu_-\|_{\mathrm{TV}} > \frac{99}{100}\frac{3}{4}\frac{23}{25}(4\varepsilon) = \frac{6831}{2500}\varepsilon > \frac{8}{3}\varepsilon. \tag{D.7}$$

The last comparison is $20493 > 20000$. None of these estimates deteriorates when $\kappa - 1$ increases.

Fresh kernels using the entire variance budget. Define

$$p_{\mathrm{r}} = \frac{\Delta^2}{4A + \Delta^2} \in (0, 1). \tag{D.8}$$

On every physical query draw a new Bernoulli variable ξ with success probability p_{r}, independently of the complete incoming history. On instance $F_{\pm}$, return only

$$G_{\pm}(u) = H'_{\mathrm{a}}(u) \pm \frac{\tau_0}{p_{\mathrm{r}}}\xi. \tag{D.9}$$

The Bernoulli flag is not a separate observation. These are measurable conditional kernels at every real query. For a query measurable in the complete incoming history,

$$\mathbb{E}[G_{\pm}(u) \mid \mathcal{H}^{\mathrm{in}}] = H'_{\mathrm{a}}(u) \pm \tau_0 = F'_{\pm}(u),$$

and their conditional variance is exactly

$$\tau_0^2 \frac{1 - p_{\mathrm{r}}}{p_{\mathrm{r}}} = \frac{\Delta^2}{4} \frac{4A}{\Delta^2} = A. \tag{D.10}$$

Thus these kernels satisfy the stipulated oracle contract, including at adaptive and unbounded query points.

Coupling with arbitrary expected stopping. Fix any algorithm uniformly ε-accurate over the entire vector class and all valid oracles. If its expected count is infinite on either constructed instance, the claimed lower bound already holds. Otherwise couple the two runs with the same complete private random tape, including randomness used for the output, and the same Bernoulli sequence $\xi_1, \xi_2, \ldots$, independent of that tape.

Until the first success on a queried round, both replies equal $H'_{\mathrm{a}}(u)$. Inductively, their queries, physical transcripts, and stopping decisions agree throughout this common execution. If they stop before any such success, their outputs agree as well. Let E_{split} be the event of a success during the common queried execution. Its first success is a queried call in both marginals.

For marginal $i \in \{-, +\}$ let T_i be its physical call count. The event $\{T_i \geq j\}$ is determined before the fresh jth innovation, also after adjoining the independent complete private tape for this coupling calculation. Nonnegative summation and conditional expectation give

$$\begin{aligned} \mathbb{E}_i\left[\sum_{j=1}^{T_i} \xi_j\right] &= \sum_{j \geq 1} \mathbb{E}_i\left[\mathbf{1}_{\{T_i \geq j\}} \xi_j\right] \\ &= p_{\mathrm{r}} \sum_{j \geq 1} \mathbb{P}_i(T_i \geq j) = p_{\mathrm{r}} \mathbb{E}_i[T_i]. \end{aligned} \tag{D.11}$$

This calculation enlarges the information used in the proof, not the physical stopping information available to the algorithm.

Let $Q_{\pm}$ be the output laws and $Y_{\pm}$ the coupled outputs. On E_{split} the stopped sum in (D.11) is at least one in each marginal. The coupling inequality therefore gives, for either i,

$$\|Q_+ - Q_-\|_{\mathrm{TV}} \leq \mathbb{P}(Y_+ \neq Y_-) \leq \mathbb{P}(E_{\mathrm{split}}) \leq p_{\mathrm{r}} \mathbb{E}_i[T_i].$$

Uniform accuracy and (D.7) imply

$$p_{\mathrm{r}} \mathbb{E}_i[T_i] \geq \|\nu_+ - \nu_-\|_{\mathrm{TV}} - 2\varepsilon > \frac{2}{3}\varepsilon.$$

Consequently the algorithm's worst-case expected cost is at least

$$\frac{2\varepsilon}{3p_{\mathrm{r}}} = \frac{2}{3}\varepsilon + \frac{8A\varepsilon}{3\Delta^2} = \frac{2}{3}\varepsilon + \frac{A}{24\varepsilon} \geq \frac{X}{24}.$$

Taking the infimum over all full-class successful algorithms proves (D.1). The argument includes zero-call branches, random output kernels, and rare arbitrarily long executions. □

D.3 THE SCALAR CURVATURE LOWER BOUND

The second obstruction is present even with exact gradients. Such a kernel is admissible under every advertised variance ceiling, so the construction applies for all $A \geq 0$. Its family encodes an unknown mode inside the prescribed localization ball. The finite-branch argument controls gradient replies, not joint value-and-gradient replies, and is not asserted to survive revelation of the exact mode. The result proves the curvature dependence in our stated localized model. Appendix G gives a different family that varies the location scale while keeping $\mu = 1$ and $L = 2$; that argument isolates why unrestricted initialization cannot have a uniformly finite bound in general.

Proposition D.2 (Scalar curvature lower bound)**.** *For $\kappa \geq 640^2$, put*

$$M = \left\lfloor \frac{\sqrt{\kappa}}{80} \right\rfloor, \qquad B_\kappa = \frac{1}{4}\left[\log_3\left(\frac{16M}{25}\right) - 1\right]. \tag{D.12}$$

Then $N^\star \geq B_\kappa$. The first coordinates of the witness modes belong to $(-1/4, 1/4)$, and all other coordinates are zero. The witnesses attain the curvature endpoints $1, \kappa$ and use exact gradients. The bound allows arbitrary adaptive vector queries and expected stopping.

Proof. We construct the family, prove concentration on disjoint intervals, and then apply a finite-branch decoding argument.

A localized threshold-ramp family. Let

$$s_0 = \sqrt{\kappa}, \qquad w = \frac{s_0}{s_0^2 - 1}, \qquad \theta_j = -\frac{1}{4} + \frac{40(j-1)}{s_0} \quad (1 \leq j \leq M). \tag{D.13}$$

Then $s_0 \geq 640$, $M \geq 8$, $w < 2/s_0$, and

$$-\frac{1}{4} \leq \theta_j \leq \frac{1}{4} - \frac{40}{s_0} < \frac{1}{4}. \tag{D.14}$$

Define the continuous derivative

$$p_j(u) = \begin{cases} u - s_0/2, & u \leq \theta_j, \\ u - s_0/2 + (s_0^2 - 1)(u - \theta_j), & \theta_j \leq u \leq \theta_j + w, \\ u + s_0/2, & u \geq \theta_j + w. \end{cases} \tag{D.15}$$

The formulas agree at both joins, since $(s_0^2 - 1)w = s_0$. Let F_j be any primitive of p_j. The function p_j is absolutely continuous, with derivative 1 off the ramp and s_0^2 in its interior almost everywhere. For $u > v$ this gives

$$u - v \leq p_j(u) - p_j(v) \leq s_0^2(u - v).$$

Integrating once more, with the orientation reversed when necessary, proves

$$\frac{1}{2}(u-v)^2 \leq F_j(u) - F_j(v) - p_j(v)(u - v) \leq \frac{s_0^2}{2}(u-v)^2. \tag{D.16}$$

Hence the functions are C^1, 1-strongly convex and κ-smooth. Pairs of distinct points in an outer interval attain the first curvature endpoint, while pairs in the nonempty ramp interior attain the second.

Let $a_j = s_0/2 - \theta_j$. The left endpoint derivative is negative and the right endpoint derivative is positive. The unique root is consequently inside the ramp and equals

$$z_j = \theta_j + \frac{a_j}{s_0^2}. \tag{D.17}$$

On the ramp, $p_j(u) = s_0^2(u - z_j)$. By (D.14) and $w < 2/s_0$,

$$-\frac{1}{4} < z_j < \frac{1}{4} - \frac{40}{s_0} + \frac{2}{s_0} < \frac{1}{4}.$$

The localization promise holds. Moreover, $F_j(u) \geq F_j(z_j) + (u - z_j)^2/2$, so each target normalizer is finite and positive.

Disjoint events carrying almost all target mass. Define

$$D_j = \left[\theta_j - \frac{16}{s_0},\, \theta_j + w + \frac{16}{s_0}\right]. \tag{D.18}$$

For adjacent indices,

$$\inf D_{j+1} = \theta_j + \frac{24}{s_0}, \qquad \sup D_j = \theta_j + w + \frac{16}{s_0}.$$

Since $w < 2/s_0 < 8/s_0$, these closed intervals are pairwise disjoint with strict gaps.

Write

$$V_j(u) = F_j(u) - F_j(z_j), \qquad b_j = s_0/2 + \theta_j + w, \qquad c = s_0/2 - 1/4.$$

Both a_j and b_j are at least c. On the ramp,

$$V_j(u) = \frac{s_0^2}{2}(u - z_j)^2.$$

The distances from z_j to the left and right ramp endpoints are a_j/s_0^2 and b_j/s_0^2, respectively. Since $c/s_0^2 > 1/(4s_0)$, the interval $[z_j - 1/(4s_0), z_j + 1/(4s_0)]$ lies in the ramp. Therefore the normalizer satisfies

$$Z_j := \int_{\mathbb{R}} e^{-F_j(u)}\, du \geq e^{-F_j(z_j)} \frac{1}{2s_0} e^{-1/32}. \tag{D.19}$$

Direct integration of the outer affine derivatives gives, for $t \geq 0$,

$$V_j(\theta_j - t) = \frac{a_j^2}{2s_0^2} + a_j t + \frac{t^2}{2},$$
$$V_j(\theta_j + w + t) = \frac{b_j^2}{2s_0^2} + b_j t + \frac{t^2}{2}.$$

Dropping the nonnegative square terms and integrating for $t \geq 16/s_0$ bounds the unnormalized mass outside D_j by

$$e^{-F_j(z_j)} \left(\frac{e^{-16a_j/s_0}}{a_j} + \frac{e^{-16b_j/s_0}}{b_j} \right) \leq e^{-F_j(z_j)} \frac{2e^{-16c/s_0}}{c}.$$

Dividing by (D.19), the normalized tail probability is at most

$$\nu_j(D_j^c) \leq \frac{4s_0}{c} \exp\left(\frac{1}{32} - \frac{16c}{s_0} \right), \tag{D.20}$$

where ν_j denotes the target corresponding to F_j. Here

$$\frac{4s_0}{c} = \frac{16s_0}{2s_0 - 1} < 9, \qquad \frac{1}{32} - \frac{16c}{s_0} = -8 + \frac{1}{32} + \frac{4}{s_0} \leq -\frac{637}{80} < -7.$$

Since $e > 1 + 1 + 1/2 + 1/6 = 8/3$ and $(8/3)^7 > 900$, the right side of (D.20) is smaller than $9e^{-7} < 1/100$. We have thus proved the normalized probability statement

$$\nu_j(D_j) > \frac{99}{100} \qquad (1 \leq j \leq M). \tag{D.21}$$

At most three replies per arbitrary query. For any fixed real query u, all indices satisfying $u \leq \theta_j$ return the same derivative $u - s_0/2$. All indices satisfying $u \geq \theta_j + w$ return the same derivative $u + s_0/2$. At most one index has $\theta_j < u < \theta_j + w$, because the ramps are disjoint. Such an index supplies at most one additional value, strictly between the two outer values. At either ramp endpoint its value equals the corresponding outer value by continuity. Thus every real query, including an endpoint or an arbitrarily remote point, has at most three exact replies across this family.

Fixing the complete private random tape of an adaptive algorithm makes its query a deterministic function of prior replies. The query itself does not create additional branches. Its exact-gradient transcript tree therefore has branching factor at most three.

TV decoding with expected query cost. Consider any algorithm uniformly ε-accurate on the full class and all valid oracles. Let Q_j be its output law on the exact-gradient instance F_j. By (D.21),

$$Q_j(D_j) > \frac{99}{100} - \frac{1}{10} = \frac{89}{100}.$$

Define a decoder $\widehat{J}(y) = j$ when $y \in D_j$, and assign a failure symbol outside their union. Let the true index J_0 be uniform on $\{1, \dots, M\}$, independently of the algorithm's private tape. Then

$$\mathbb{P}(\widehat{J}(Y) = J_0) > \frac{89}{100}. \tag{D.22}$$

Write T_j for the physical query count on instance j and set

$$\bar{n} = \frac{1}{M} \sum_{j=1}^{M} \mathbb{E}[T_j].$$

If $\bar{n} = \infty$, the desired lower bound already holds. If $\bar{n} = 0$, there are almost surely no queries on each member of this finite family. The common depth-zero output can correctly decode at most one index for any fixed private tape, giving success probability at most $1/M \le 1/8$, contrary to (D.22). It remains to consider $0 < \bar{n} < \infty$.

Let T be the query count in the uniform experiment and take $H_{\mathrm{tree}} = \lceil 4\bar{n} \rceil$. Markov's inequality gives

$$\mathbb{P}(T > H_{\mathrm{tree}}) \le \frac{\bar{n}}{H_{\mathrm{tree}}} \le \frac{1}{4}. \tag{D.23}$$

Fix the entire private tape, including the randomness used for the output. Truncate the resulting deterministic transcript tree at depth H_{tree}. Pad an earlier stopping leaf with one dummy edge per remaining level, retaining its original output. Each active node has at most three children and each padded node has one, so at most $3^{H_{\mathrm{tree}}}$ leaves occur at the common depth. A leaf corresponding to stopping by the horizon has a single output and hence a single decoder label. Among the indices reaching that leaf, at most one is decoded correctly. Averaging over the private tape gives

$$\mathbb{P}(\widehat{J}(Y) = J_0,\ T \le H_{\mathrm{tree}}) \le \frac{3^{H_{\mathrm{tree}}}}{M}. \tag{D.24}$$

On the other hand, (D.22) and (D.23) make this probability greater than $89/100 - 1/4 = 16/25$. Therefore

$$H_{\mathrm{tree}} > \log_3\left(\frac{16M}{25}\right).$$

Since $H_{\mathrm{tree}} < 4\bar{n} + 1$,

$$\bar{n} > \frac{1}{4}\left[\log_3\left(\frac{16M}{25}\right) - 1\right] = B_\kappa.$$

The worst-case expected cost on the full class is at least this prior average. Infimizing over all admissible uniformly accurate algorithms proves $N^\star \ge B_\kappa$. The exact-gradient kernel used here is valid at the same advertised variance ceiling as every other instance in the minimax definition. □

D.4 The constant lower bound and completion

The ramp family yields useful information only for large curvature, whereas the noise lower bound vanishes at $A = 0$. To cover the remaining parameter range, we use two separated Gaussian targets and the stated physical stopping convention.

Proposition D.3 (Constant lower bound under physical stopping)**.** *For every public tuple in the stated regime, including $A = 0$,*

$$N^\star \ge 1. \tag{D.25}$$

Proof. The event $\{T = 0\}$ belongs to the trivial physical sigma-field $\mathcal{F}_0$ in (2.3), so it has probability zero or one. The decision whether to make a first query uses no instance-dependent information. On a common private random tape, therefore, this zero-call decision and any resulting output are the same for all instances sharing the public parameters.

In normalized coordinates consider

$$F_{\pm}(x) = \tfrac{1}{2}\|x \mp e_1\|^2,$$

where e_1 is the first coordinate vector. Both potentials satisfy the curvature bounds for every $\kappa \geq 1$, and their modes obey the localization bound. Use their exact gradients, which are admissible for every $A \geq 0$. Their targets are $N(e_1, I_d)$ and $N(-e_1, I_d)$. Their probability difference on the half-space $\{x : x_1 \geq 0\}$ equals

$$2\Phi(1) - 1 = \sqrt{\frac{2}{\pi}} \int_0^1 e^{-t^2/2}\, dt > \frac{3}{4} \cdot \frac{1}{2} = \frac{3}{8} > 2\varepsilon,$$

where Φ is the standard normal distribution function. The two elementary bounds used here are $\sqrt{2/\pi} > 3/4$ and $e^{-1/2} > 1/2$.

A common zero-query output law cannot have TV distance at most ε from both targets, by the triangle inequality. Uniform accuracy therefore rules out the probability-one zero-call decision. It follows that $T \geq 1$ almost surely for every admissible uniformly successful algorithm, so its worst-case expected query count is at least one. Taking the infimum proves (D.25). □

All three obstructions are now available. Their witnesses need not coincide: the minimax cost dominates each witness family separately. It remains to compare their maximum with the common benchmark S.

Proof of Theorem 2.3. If $N^\star = \infty$, the lower bound is immediate. Otherwise, combine the three lower bounds at the same public tuple. No assertion that their costs add on a single instance is needed.

By Propositions D.1 and D.3,

$$X \leq 24N^\star, \qquad N^\star \geq 1. \tag{D.26}$$

These statements also hold when $A = X = 0$. If $1 \leq \kappa < 640^2$, then $1 + \kappa < 409601 < 2^{21}$ and $\log 2 < 1$, so

$$S < X + 22 \leq 24N^\star + 22N^\star = 46N^\star.$$

This branch includes $\kappa = 1$ and does not invoke the ramp construction outside its stated range.

If $\kappa \geq 640^2$, the argument of the floor in (D.12) is at least eight. Hence

$$M \geq \frac{\sqrt{\kappa}}{160}.$$

Proposition D.2 gives

$$N^\star \geq \frac{\log \kappa}{8 \log 3} - \frac{\log 250}{4 \log 3} - \frac{1}{4}. \tag{D.27}$$

Since $\log 3 > 0$, this can be rearranged. Using $\log 3 < 2$, $\log 250 < 8$, and $N^\star \geq 0$, we obtain

$$\log \kappa \leq 8 \log 3\, N^\star + 2 \log 250 + 2 \log 3 < 16N^\star + 20.$$

Also $1 + \kappa \leq 2\kappa$. It follows that

$$S < 16N^\star + 22 + X \leq 40N^\star + 22 \leq 62N^\star,$$

where the last step uses $N^\star \geq 1$ from (D.26). Both curvature ranges therefore yield $N^\star \geq S/62$.

Together with the stochastic-noise and constant lower bounds, this proves

$$N^\star \geq \max\left\{1, \frac{X}{24}, \frac{S}{62}\right\}. \tag{D.28}$$

This proves Theorem 2.3. Together with Theorem 2.2, it gives the joint minimax rate in Section 2. The ratio of the displayed upper benchmark $C_d S$ to the lower benchmark $S/62$ is at most $62C_d$, independently of $L, \mu, \sigma, \varepsilon$ for each fixed d. Restoring $A = \sigma^2/\mu$ and $\kappa = L/\mu$ through the query-preserving normalization yields the asserted constant-factor characterization in the original physical coordinates, throughout $\sigma^2 \geq 0$. □

E PROOFS OF THE INITIALIZATION EXTENSIONS

The main theorem uses $\mathcal{C}_d(L, \mu)$ as defined in Section 2. We now derive extensions without changing that theorem or its pilot algorithms. All query bounds below count physical replies, including markers, and retain the physical-transcript stopping convention.

Proof sketch and intuition for Corollary 5.3. Translate by x_0 and advertise the conservative strong-convexity bound $\mu_R = \mu/\Gamma_R$. Then $R \le \mu_R^{-1/2}$, so the translated potential belongs to the class of the main theorem with condition number $\kappa\Gamma_R$. Only the public curvature bound changes; the target potential is not multiplied or tempered. Applying Theorem 2.2 or Corollary 5.1 proves the claims. The factor Γ_R in the noisy term is a sufficient cost of this reduction, not a claimed matching radius-noise lower bound.

Proof of Corollary 5.3.

Proof. Put $\Gamma_R = \max\{1, \mu R^2\}$ and $\mu_R = \mu/\Gamma_R$. Let $h(u) = f(x_0 + u)$. Translation preserves the curvature bounds, and $\mu_R \le \mu$, so h is also μ_R-strongly convex. Its minimizer satisfies

$$\|u_h^\star\| \le R \le \mu_R^{-1/2}.$$

Thus $h \in \mathcal{C}_d(L, \mu_R)$. Simulate a gradient query at u by one physical query at $x_0 + u$, returning the same vector. The conditional variance ceiling remains σ^2. Apply Theorem 2.2 with public parameters $(L, \mu_R, \sigma, \varepsilon)$ and translate its output back by x_0. Its stronger error bound $7\varepsilon/16$ in particular implies ε-accuracy. Its condition number and normalized variance are $\kappa\Gamma_R$ and $\Gamma_R\sigma^2/\mu$, proving (5.1). The additive constant in its bound is absorbed because $\kappa\Gamma_R \ge 1$.

For $\sigma = 0$, instead invoke Corollary 5.1 on h with the same public curvature bounds. This gives exact sampling and the first bound in (5.2). If $D_R = \sqrt{\mu}\, R$, then $\Gamma_R \le (1 + D_R)^2$ and

$$1 + \kappa\Gamma_R \le (1 + \kappa)(1 + D_R)^2,$$

which proves the second bound. Translation and the sampler's internal normalization are invertible transformations known from the inputs, so they preserve the query count, the physical stopping rule, and TV. Only a more conservative lower curvature bound is supplied to the sampler; the potential values and target temperature are not rescaled. □

Scope of the noisy-radius bound. Equation (5.1) is a direct upper-bound consequence, not a claim of optimal joint dependence on R and σ. In particular, no matching lower bound for its factor $\Gamma_R\sigma^2/(\mu\varepsilon)$ is asserted.

Proof sketch and intuition for Corollary 5.4. A single exact reply $g_0 = \nabla f(x_0)$ certifies $\|x_0 - x_f^\star\| \le \|g_0\|/\mu$ by strong monotonicity. Use this radius in Corollary 5.3. Smoothness gives $\|g_0\| \le L\|x_0 - x_f^\star\|$ and converts the resulting cost $1 + C_d \log(1 + \kappa \max\{1, \|g_0\|^2/\mu\})$ into (5.3). The exact first reply thus supplies its own initialization certificate. This certificate is not valid for a single noisy reply. The full proof below includes the initialization cost and physical stopping.

Proof of Corollary 5.4.

Proof. Put $g_0 = \nabla f(x_0)$ and $\widehat{\Gamma} = \max\{1, \|g_0\|^2/\mu\}$. Make one physical gradient query at x_0. Its reply equals g_0 almost surely. Strong monotonicity of the gradient and $\nabla f(x_f^\star) = 0$ imply

$$\mu\|x_0 - x_f^\star\|^2 \le \langle g_0, x_0 - x_f^\star\rangle \le \|g_0\|\,\|x_0 - x_f^\star\|.$$

Hence

$$\|x_0 - x_f^\star\| \le \widehat{R} := \frac{\|g_0\|}{\mu}.$$

This also holds when $g_0 = 0$, in which case $x_0 = x_f^\star$. Compute $\widehat{\mu} = \mu/\widehat{\Gamma}$ and apply the exact sampler of Corollary 5.3 with center x_0 and radius $\widehat{R}$. Its required input promise is certified by the first reply. The additional expected cost is at most $C_d \log(1 + \kappa\widehat{\Gamma})$. Including the initial reply gives

$$\mathbb{E}_f[T] \le 1 + C_d \log(1 + \kappa\widehat{\Gamma}), \qquad \widehat{\Gamma} = \max\{1, \|g_0\|^2/\mu\}, \tag{E.1}$$

and the returned sample is exact.

To express the bound through the actual distance, put $D_f = \sqrt{\mu}\,\|x_0 - x_f^\star\|$. Smoothness gives

$$\|g_0\| \le L\|x_0 - x_f^\star\|, \qquad \widehat{\Gamma} \le \max\{1, \kappa^2 D_f^2\}.$$

Since $\kappa \ge 1$,

$$1 + \kappa\widehat{\Gamma} \le (1+\kappa)^3(1+D_f)^2.$$

Taking logarithms and absorbing the initial one-call cost proves (5.3).

For a fixed f, the first exact reply fixes all subsequent parameters almost surely. The new parameter choices are measurable functions of that physical reply, and the translated, normalized query locations can be reconstructed from it and the later transcript. Thus composition with the first call preserves physical stopping and charges every marker. Almost-sure termination follows from Corollary 5.1. No estimate of D_f is supplied as an input, and no potential-value query is used. A noisy reply would not provide the deterministic certificate above; this corollary is stated only for zero variance. □

The quadratic endpoint without localization. At $\kappa = 1$, even the initialization term can be avoided. For every $f \in \mathcal{U}_d(\mu, \mu)$,

$$f(x) = \frac{\mu}{2}\|x\|^2 + \langle b, x\rangle + c, \qquad \nu_f = N(-b/\mu, \mu^{-1} I_d),$$

where b is unrestricted. The argument in Corollary 5.2 never uses the bound on b in its error estimate. In original coordinates, query zero

$$n = \max\left\{1, \left\lceil \frac{\sigma^2}{2\mu\varepsilon} \right\rceil \right\}$$

times, average the replies as $\widehat{b}$, and return $-\widehat{b}/\mu + Z/\sqrt{\mu}$ for independent $Z \sim N(0, I_d)$. Translation of the Taylor-remainder argument gives TV error at most $\sigma^2/(2\mu n)$. At zero noise one query is sufficient for exact sampling, and the constant lower bound follows from the localized subclass. Thus the location lower bound in Appendix G, proved at $\kappa = 2$, is a general obstruction, not a claim about the quadratic endpoint.

F A VALUE-AND-GRADIENT BENCHMARK FOR THE NOISELESS COMPARISON

This appendix records an elementary comparison argument, not an additional gradient-only algorithm. We work with a 1-strongly convex, κ-smooth potential V, a known mode at the origin, and exact access to both V and ∇V. Normalize $V(0) = \min V = 0$ by one value query and subtraction when necessary. For fixed $d \geq 2$, the ellipsoidal approximation in Chewi et al. (2024, arXiv v2, Corollary 48) supplies, in $O(\log(1+\kappa))$ queries,

$$E = z + M\mathbb{B}_d \subseteq K := \{x : V(x) \leq 1\} \subseteq z + \rho M\mathbb{B}_d,$$

where M is invertible and $1 \leq \rho = O(1)$. The following argument explains why this geometry already permits accuracy-independent expected cost when values are available.

Proposition F.1 (Value-based rejection benchmark)**.** *Given the displayed sandwich, there is an exact sampler for $\pi(dx) = Z^{-1}e^{-V(x)}\,dx$ with at most $e^2 d!(2\rho)^d$ expected additional value queries under the usual randomized stopping convention. Hence preprocessing and sampling together cost $O(\log(1+\kappa))$ expected queries. For $0 < \varepsilon < 1$, truncation gives an ε-TV sampler with a deterministic budget $O(\log(1+\kappa)+\log(1/\varepsilon))$. Imposing physical-transcript stopping changes these bounds by only a constant factor.*

Proof. Because $0 \in K \subseteq z + \rho M\mathbb{B}_d$, we have $\|M^{-1}z\| \leq \rho$. Consequently

$$K \subseteq E_0 := 2\rho M\mathbb{B}_d, \qquad \mathrm{vol}(E_0) = (2\rho)^d \,\mathrm{vol}(E).$$

Define

$$r(x) = \frac{\|M^{-1}x\|}{2\rho}.$$

We claim that $V(x) \geq r(x) - 1$ for all x. If $r(x) \leq 1$, this follows from $V(x) \geq 0$. If $r(x) > 1$, then $x \notin K$, so $V(x) > 1$. Convexity and $V(0) = 0$ give

$$V\left(\frac{x}{V(x)}\right) \leq 1.$$

Thus $x/V(x) \in K \subseteq E_0$, which implies $r(x) \leq V(x)$ and proves the claim.

Use the full-support proposal

$$q(x) = \frac{e^{-r(x)}}{d!\,\mathrm{vol}(E_0)}.$$

Polar integration verifies

$$\int_{\mathbb{R}^d} e^{-r(x)}\,dx = d\,\mathrm{vol}(E_0)\int_0^\infty e^{-s}s^{d-1}\,ds = d!\,\mathrm{vol}(E_0).$$

To sample q, draw a uniform direction Θ on the unit sphere and, independently, $R \sim \mathrm{Gamma}(d,1)$, and return $Y_q = 2\rho MR\Theta$. Here R can be generated as a sum of d independent unit exponentials, using only internal randomness.

Given $Y_q = x$, query $V(x)$ and accept with probability

$$a(x) = \exp\big(r(x) - 1 - V(x)\big).$$

The proved inequality ensures $0 < a(x) \leq 1$, and

$$q(x)a(x) = \frac{e^{-V(x)}}{e\,d!\,\mathrm{vol}(E_0)}.$$

Hence the accepted law is exactly π. Since $V \leq 1$ on E,

$$Z \geq e^{-1}\,\mathrm{vol}(E), \qquad p_{\mathrm{acc}} = \frac{Z}{e\,d!\,\mathrm{vol}(E_0)} \geq p_d := \frac{1}{e^2 d!(2\rho)^d} > 0.$$

Independent repetitions therefore terminate almost surely and use at most p_d^{-1} expected value queries. Adding the geometric preprocessing and the optional normalization query proves the expected-cost bound.

For a deterministic budget, attempt at most

$$m_\varepsilon = \left\lceil p_d^{-1} \log \frac{1}{\varepsilon} \right\rceil$$

trials and return a fixed point if all reject. Couple this algorithm with the uncapped exact sampler. Their outputs can differ only if the first m_ε trials all reject, an event of probability at most

$$(1 - p_{\mathrm{acc}})^{m_\varepsilon} \le e^{-p_d m_\varepsilon} \le \varepsilon.$$

The coupling inequality gives the TV guarantee and the stated deterministic query budget.

Finally, to enforce the physical-transcript stopping rule, follow each trial's privately drawn acceptance bit $b \in \{0, 1\}$ by a charged dummy oracle call at be_1 and ignore its reply. Stop only after the marker announcing acceptance, or after the marker in the last capped trial. The query locations then reveal every stopping decision. This at most doubles the additional query cost and does not change either complexity order. □

The acceptance test above explicitly uses $V(x)$ and therefore does not implement the oracle model of the main paper. Its role is to separate the distinction between expected and deterministic cost from the separate issue of implementing exact acceptance without potential values. Corollary 5.1 addresses the latter issue while also allowing an unknown, coarsely localized mode.

G A FIXED-CURVATURE LOWER BOUND FOR UNRESTRICTED LOCALIZATION

This appendix isolates location hardness while keeping curvature fixed. It uses the unrestricted class $\mathcal{U}_d$, defined in Section 2, and exact gradient queries only.

Proposition G.1 (Logarithmic location cost)**.** *Fix $d \ge 1$ and $0 < \varepsilon \le 1/10$. For every integer $m \ge 1$, put $R_m = 8^{m+1}$. There is a finite family*

$$\{F_s : s \in \{-1, +1\}^m\} \subseteq \mathcal{C}_d(2, 1; 0, R_m)$$

such that every randomized gradient-query sampler with TV error at most ε on each family member satisfies

$$\sup_s \mathbb{E}_s[T] \ge \frac{27m}{80}.$$

The bound permits arbitrary real full-vector queries and adaptive random stopping. Consequently, over $\mathcal{U}_d(2, 1)$ with no restriction on the minimizer,

$$\inf_{\mathcal{A} \text{ uniformly } \varepsilon\text{-accurate}} \sup_{f \in \mathcal{U}_d(2,1)} \mathbb{E}_f[T] = +\infty.$$

Hard instance: idea and proof sketch. In one dimension, construct a depth-m binary family on an interval of radius 8^{m+1}. At each level the active interval has two possible children, each of one-eighth the radius and with disjoint interiors. Affine derivative connectors make the alternatives agree outside the parent while keeping all derivative slopes in $[1, 2]$. Terminal modes are at least 16 apart; each target puts more than $99/100$ of its mass within distance four of its mode. A TV-accurate sample therefore decodes the entire hidden string.

The children of $[c - r, c + r]$ are $[c - r/4, c]$ and $[c, c + r/4]$. The connecting derivative slopes are $21/16$ and $7/4$, exchanged between the two choices. Thus each hidden sign moves the possible zero without changing external replies or the curvature bounds. The final target width stays constant as the enclosing radius grows.

A crucial distinction from the curvature family is that one query *can* reveal more than one sign: it may land inside several nested active intervals. An argument asserting one bit per real query would therefore be invalid. The proof instead charges the expected number of newly exposed signs. Long revelations are possible, but geometrically unlikely under the hidden-sign prior, even when the query location is chosen adaptively from all previous replies.

To control arbitrary real replies, augment the oracle by revealing the hidden prefix needed to evaluate a query. At most one child interior contains that query. Under uniform unrevealed signs, proceeding to

each further level has conditional probability at most one half, so a physical query reveals at most two new signs in expectation. The suffix remains uniform after every adaptive stopping history. Decoding with constant success forces $\Omega(m)$ exposed signs and hence $\Omega(m)$ expected queries. More precisely, if K_T is the exposed-prefix length at stopping, deferred decisions and nonnegative summation yield

$$\mathbb{E}[K_T] \le 2\mathbb{E}[T], \qquad \mathbb{P}(\text{correct decoding}) \le \mathbb{E}[2^{-(m-K_T)}] \le \frac{1}{2} + \frac{\mathbb{E}[K_T]}{2m}.$$

The constant decoding probability therefore forces expected cost linear in m. Unqueried signs cannot be inferred from the stopping event itself: terminality depends only on the exposed prefix and private randomness, not on the still-uniform suffix. This is why the same lower bound covers algorithms that occasionally make very long runs, rather than only deterministic-budget algorithms. For the supplied-radius conclusion in Proposition 5.5, let $R \ge 64$ and take $m = \lfloor \log_8 R \rfloor - 1 \ge 1$; then $8^{m+1} \le R$, so all modes satisfy the supplied-ball promise, including at $R = 64$. Adding independent standard Gaussian coordinates preserves the scalar decoding argument and the curvature bounds in every fixed dimension. The full proof of Proposition G.1 below supplies the construction and random-stopping calculation. Combining this choice of m with Corollary 5.3 yields the matching radius order. This is a worst-case-over-locations statement, not a per-instance lower bound, and it does not apply to the one-query quadratic endpoint $L = \mu$.

Proof. It suffices to construct scalar witnesses and then use the Gaussian completion in Section D.1. We give a nested family whose curvature bounds do not depend on its depth. An augmented oracle will make the information calculation explicit; it only strengthens the algorithm relative to exact-gradient access.

A binary family of uniformly curved derivatives. Write $R = R_m = 8^{m+1}$ and $r_k = R/8^k$ for $0 \le k \le m$. For a sign string $s = (s_1, \dots, s_m)$, put

$$c_0 = 0, \qquad c_k = \sum_{i=1}^{k} s_i r_i, \qquad I_k = [c_k - r_k, c_k + r_k].$$

The two possible children of an interval $[c - r, c + r]$ have radius $r/8$ and centers $c \pm r/8$. Their interiors are disjoint:

$$I_- = [c - r/4, c], \qquad I_+ = [c, c + r/4].$$

Define a continuous piecewise-affine derivative p_s as follows. Outside I_0, set $p_s(u) = 3u/2$. On each of the two components of $I_{k-1} \setminus \operatorname{int} I_k$, interpolate affinely between its prescribed endpoint values

$$p_s(c_{k-1} \pm r_{k-1}) = \pm\tfrac{3}{2} r_{k-1}, \qquad p_s(c_k \pm r_k) = \pm\tfrac{3}{2} r_k.$$

Finally, on I_m set $p_s(u) = 3(u - c_m)/2$. When the chosen child is the right child, the left and right connector slopes are respectively $21/16$ and $7/4$; for the left child these slopes are exchanged. The exterior and terminal slopes are $3/2$. Consequently the definitions agree at every joint and all slopes lie in $[1, 2]$. With $F_s(u) = \int_{c_m}^{u} p_s(v)\, dv$, integration gives

$$\tfrac{1}{2}(u - v)^2 \le F_s(u) - F_s(v) - p_s(v)(u - v) \le (u - v)^2.$$

Thus $F_s \in C^1$, with the advertised curvature bounds $1, 2$. These are bounds, not a requirement that either endpoint be attained. Its unique minimizer is c_m, and $|c_m| \le \sum_{k=1}^{m} r_k < R/7 < R$.

A sample identifies the entire string. Distinct centers are separated by at least 16. Indeed, if k is the first differing sign, their separation is at least $2r_k - 2\sum_{j>k} r_j$; it equals $2r_m = 16$ when $k = m$, and exceeds $12r_k/7 \ge 12 \cdot 64/7 > 16$ when $k < m$. Therefore the intervals $D_s = [c_m - 4, c_m + 4]$ are pairwise disjoint. The curvature bounds and $F_s(c_m) = 0$ imply

$$Z_s \ge \int_{c_m - 1}^{c_m + 1} e^{-F_s(u)}\, du \ge 2e^{-1}, \qquad \int_{D_s^c} e^{-F_s(u)}\, du \le 2\int_4^\infty e^{-t^2/2}\, dt \le \tfrac{1}{2} e^{-8}.$$

For the final inequality use $1 \le t/4$ on $[4, \infty)$. Hence $\nu_s(D_s^c) \le e^{-7}/4 < 1/100$. Any uniformly ε-accurate sampler has $\mathbb{P}_s(Y \in D_s) > 89/100$. Under a uniform prior on the m signs, decoding by the disjoint intervals consequently succeeds with probability greater than $89/100$.

One arbitrary query reveals at most two new bits in expectation. Expose the signs only as needed to answer a query. Suppose a prefix of length k has been revealed. At $k = m$, the entire derivative is known and no more signs are exposed. If $k < m$ and the queried scalar coordinate u is outside $\operatorname{int} I_k$, its derivative is already fixed by that prefix, so reveal no further sign. Otherwise reveal s_{k+1}. If u is outside the interior of the selected child, the connector formula determines its derivative and the query is answered. If it is inside, continue recursively. Stop also at depth m, where the terminal affine formula determines the reply. At a child boundary the derivative is fixed and no deeper sign is needed. Return both the derivative and the newly exposed signs. This is an augmented oracle, and an exact-gradient algorithm can be simulated by ignoring the exposed signs. Previously answered queries impose no restriction on the unexposed suffix: their replies are functions only of the exposed prefix. Thus, conditionally on the augmented past and all private randomness used to choose the next query, the remaining signs are still independent and uniform.

At most one child interior contains a fixed u. Every additional level after the first therefore requires an independent sign of conditional probability at most $1/2$. If Δ_j is the number of newly exposed signs at the jth query, then, on a reached query,

$$\mathbb{P}(\Delta_j \geq a \mid \text{incoming augmented history}) \leq 2^{-(a-1)} \quad (a \geq 1),$$
$$\mathbb{E}[\Delta_j \mid \text{incoming augmented history}] \leq 2.$$

These inequalities hold for arbitrary real query locations, including points outside I_0 and points at any connector boundary. Let K_T be the exposed-prefix length when the algorithm stops. If the prior-average expected call count is infinite, there is nothing to prove. Otherwise nonnegative summation and conditioning before each reached query give

$$\mathbb{E}[K_T] = \sum_{j \geq 1} \mathbb{E}[\mathbf{1}_{\{T \geq j\}} \Delta_j] \leq 2 \sum_{j \geq 1} \mathbb{P}(T \geq j) = 2\mathbb{E}[T]. \tag{G.1}$$

This argument permits arbitrary adaptive random stopping; it uses no fixed-horizon reduction.

Decoding forces a linear number of exposures. Conditionally on the terminal augmented transcript, the unrevealed suffix remains uniform. This deferred-decisions fact also holds at a stopping time: whether that transcript is terminal is determined by its exposed information and private randomness, not by unused signs. Adjoin any remaining private randomness used for the output; it is independent of the unexposed suffix. A fixed output can decode at most one of the 2^{m-K_T} completions. Hence

$$\mathbb{P}(\widehat{s}(Y) = s) \leq \mathbb{E}[2^{-(m-K_T)}] \leq \frac{1}{2} + \frac{\mathbb{E}[K_T]}{2m}.$$

The last inequality follows from convexity: for $0 \leq k \leq m$, 2^{k-m} lies below the chord joining $(0, 2^{-m})$ and $(m, 1)$, which in turn lies below $1/2 + k/(2m)$. Success greater than $89/100$ implies $\mathbb{E}[K_T] > 39m/50$. By (G.1),

$$\sup_s \mathbb{E}_s[T] \geq \mathbb{E}[T] > \frac{39m}{100} \geq \frac{27m}{80}.$$

For $d > 1$, append $\frac{1}{2}\|x_{2:d}\|^2$ and return the exact known gradients in those coordinates. The same argument applies directly to the first coordinate of every full-vector query and to the first coordinate of the output. Finally, m is arbitrary, so a uniformly accurate algorithm on $\mathcal{U}_d(2, 1)$ cannot have a finite worst-case expected query count. Taking the infimum proves the claim. □

Proof of Proposition 5.5. For a supplied radius $R \geq 64$, choose $m = \lfloor \log_8 R \rfloor - 1$ in the construction. Its minimizers lie inside $\mathbb{B}(0, R)$. Writing $t = \log_8 R \geq 2$, we have $m = \lfloor t \rfloor - 1 \geq t/3$: for $2 \leq t < 3$ use $m = 1$; for $t \geq 3$ use $m > t - 2 \geq t/3$. Combined with Corollary 5.3, this gives $\Theta(\log(1 + R))$ noiseless expected minimax cost at $\mu = 1$, $L = 2$, uniformly for $0 < \varepsilon \leq 1/10$. This is a worst-case-over-locations statement, not a per-instance lower bound. It does not assert joint optimality for all curvature ratios or extend the lower bound to a value oracle. In particular, it does not apply to the quadratic endpoint $L = \mu$. □