How Loud Is a Gradient? An Amplitude Law for Noisy Comparisons

Anonymous authors
Paper under double-blind review

Abstract

A comparison oracle reports which of two points is better. Can it certify that an optimizer may stop, that is, that $\|\nabla f(\bar{x})\| \le \varepsilon$? Exact comparisons cannot: they preserve only the order of values. Under Bradley–Terry noise a query returns 1 with probability $\sigma(\beta(f(x) - f(y)))$, so repeated answers measure value gaps. We show that every question posed to this oracle has a natural size, its *amplitude*: the least uniform change of the logit potential βf, modulo constants, that flips the answer. A question of amplitude a needs at least H/a^2 comparisons, with H the entropy of what must be distinguished, and it is unanswerable when opposite answers can share a potential. Computing amplitudes organizes the theory. A gradient is a small value step, of amplitude exactly $\beta\varepsilon^2/(4L_1)$ for ε-vs-2ε tests over L_1-smooth objectives; for p-th order smoothness, the exponent of Kolmogorov's inequality yields search lower bounds $\Omega(nL^{2/p}/(\beta^2\varepsilon^{2+2/p}))$. When β is known only up to a ratio ρ, the amplitude is proportional to $(c' - \rho)_+^2$: ε-vs-$c'\varepsilon$ verification is impossible exactly when $\rho \ge c'$, and its cost grows as $(c' - \rho)^{-4}$ at the edge, a rate an explicit tester attains. Value-margin tests have amplitude βm_0 and need neither smoothness nor dimension-dependent sampling. The unresolved factor is dimension: at fixed confidence, known-scale verification costs between $\max\{1, a^{-2}\}$ and $\max\{1, n^2 a^{-2}\}$ comparisons up to constants, and nonadaptive verifiers need order $(n-1)a^{-2}$.

> *"When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind."*
> —Lord Kelvin, *Electrical Units of Measurement*, lecture of 1883 (Thomson, 1889)

1 Introduction

Learning from pairwise preferences has become a standard way to align large models with human intent. In reinforcement learning from human feedback (RLHF), annotators compare two outputs, a reward model is fitted to the comparisons under the Bradley–Terry model, and a policy is optimized against the fitted reward (Christiano et al., 2017; Ouyang et al., 2022). Direct preference optimization (DPO) removes the explicit reward model but keeps the same likelihood: the probability that one response is preferred is a logistic function of β times a difference of policy log-ratios, where β sets the strength of the regularization (Rafailov et al., 2023). Public leaderboards rank the models themselves by a Bradley–Terry fit to pairwise votes (Chiang et al., 2024), and the same feedback drives dueling bandits and preferential Bayesian optimization (Yue et al., 2012; González et al., 2017). In each case the learner never observes the objective. It sees only which of two candidates a noisy judge prefers.

Preferences can steer an optimizer. Can they also tell it when to stop? A practitioner who compares checkpoints, prompts or designs by preference would like to certify that no small change is still preferred: that the latent utility $f\colon \mathbb{R}^n \to \mathbb{R}$ is nearly stationary at the current point, $\|\nabla f(\bar{x})\| \le \varepsilon$. Two features of preference data make this delicate. The certificate concerns the *size* of a derivative, whereas a comparison reports an order: f, $2f$ and f^3 order every pair of points alike while their gradients differ, so no number of exact comparisons certifies stationarity (Appendix G). And the utility is not calibrated: under the Bradley–Terry model the temperature of the judge and the scale of the utility enter only through their product, as β and the log-ratios do in the DPO likelihood.

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.

Noise changes the question. Under the Bradley–Terry (BT) model (Bradley & Terry, 1952), a query (x, y) returns a bit B with

$$\mathbb{P}(B = 1 \mid x, y) = \sigma\big(\beta(f(x) - f(y))\big), \qquad \sigma(z) = \frac{1}{1 + e^{-z}}, \quad \beta > 0, \tag{1.1}$$

so the frequency of answers measures the value gap. Stationarity becomes certifiable in principle, and the question becomes quantitative. *How loud is a gradient?* How far must the answer probabilities move before a gradient of size ε can be told from one of size 2ε, and what does it cost to hear the difference? When the scale β is itself unknown, when does the gradient fall silent?

One number. Everything a comparison transcript reveals about f is contained in the logit potential βf modulo additive constants. Two potentials are statistically close exactly when they are uniformly close modulo constants: each comparison contributes at most $\frac{1}{2}d^2$ to the KL divergence, where d is half the oscillation of their difference, and $\frac{1}{2}$ is sharp (Proposition 2.5). This suggests measuring a question by how far the potential must move to change its answer. We call this distance the *amplitude* $a(\mathcal{P})$ of the question $\mathcal{P}$ (Definition 2.6), and it prices the question (Theorem 2.7). If two instances with opposite answers share a potential, no budget suffices. Otherwise every algorithm needs at least $H/a(\mathcal{P})^2$ comparisons, up to a constant, where H is the entropy of the hypotheses that a correct answer must separate. For verification over L_1-smooth objectives and for value margins, explicit algorithms match this law up to factors polynomial in n, paying $\max\{1, a^{-2}\}$ per reading: comparisons can be zoomed but not amplified. The rest of the paper computes amplitudes, and each computation is a theorem.

- *A gradient is a small value step* (Section 3). A gradient change γ under curvature at most M forces a value step of height at least γ^2/M, and this is attained. At a known scale, the amplitude of ε-vs-$c'\varepsilon$ verification over L_1-smooth objectives is exactly $\beta(c' - 1)^2\varepsilon^2/(4L_1)$ (Proposition 3.1). Under an L-Lipschitz p-th derivative the step has height of order $\varepsilon^{1+1/p}L^{-1/p}$, the exponent of Kolmogorov's inequality, and hiding it in $e^{\Omega(n)}$ places makes search cost $\Omega(nL^{2/p}/(\beta^2\varepsilon^{2+2/p}))$ (Theorem 3.3). The ε^{-4} and ε^{-3} rates of first- and second-order smoothness are two samples of this law.

- *Scale uncertainty muffles* (Section 5). If only $\beta \in [\beta_{\min}, \beta_{\max}]$ is known, the amplitude becomes $\beta_{\min}\varepsilon^2(c' - \rho)_+^2/(2(1 + \rho)L_1)$ with $\rho = \beta_{\max}/\beta_{\min}$. Its zero set $\rho \geq c'$ is exactly where verification is impossible, and near the edge the cost grows as $(c' - \rho)^{-4}$, a rate our tester attains (Theorems 5.1 and 5.2).

- *Values need no calculus* (Section 6). A value-margin promise on a probe sphere has amplitude exactly βm_0. No smoothness step intervenes, and its cost is $\Theta((\beta m_0)^{-2})$ up to logarithms, without smoothness or dimension factors (Theorem 6.1).

- *Dimension is what remains* (Section 4). Known-scale verification costs between $\max\{1, a^{-2}\}$ and $\max\{1, n^2a^{-2}\}$ up to constants and $\log(1/\delta)$, and falls to $O(\log(1/\delta))$ once $\beta \geq 2nL_1/\varepsilon^2$ (Theorem 4.1). Nonadaptive verifiers need order $(n - 1)a^{-2}$, because a fixed query hears only a $1/(n - 1)$ share of a hidden direction's variation (Theorem 4.2).

Relation to prior work. Wang et al. (2026) study noiseless comparison oracles for stationarity and ask how their results extend to stochastic oracles, BT included. Nonconvex algorithms exist for ranking feedback (Tang et al., 2023), noisy pairwise feedback (Wang et al., 2025), and nonsmooth objectives with a known link (El Bakkali et al., 2026b). The $O(k^2/\varepsilon^4)$ search bound of El Bakkali et al. (2026a) uses a confidence-score oracle with a gap-dependent second-moment condition, which a raw BT bit, of second moment one, violates near ties; it shares our accuracy exponent but is not an upper bound for our oracle. Comparison lower bounds for strongly convex optimization (Jamieson et al., 2012) and algorithms for dueling convex optimization (Saha et al., 2021; 2022; 2024) address a different objective. Our lower bounds refine change-of-measure arguments (Kaufmann et al., 2016; Lattimore & Szepesvári, 2020): the amplitude turns the choice of hard instances into an extremal problem, which Hadamard's and Kolmogorov's inequalities solve (Hadamard, 1914; Kolmogorov, 1939). The unknown-scale results turn the scale non-identifiability of pairwise-comparison models (Shah et al., 2016) into an exact and quantitative criterion for a stationarity certificate.

Table 1: Amplitudes and comparison complexities at fixed confidence. Scale is known unless stated; $R = L_1^2/(\beta^2\varepsilon^4)$ and $\beta^* = 2nL_1/\varepsilon^2$. Lower bounds hold for adaptive algorithms except in the nonadaptive row. Unknown-scale rows use (A1), with $\beta \in [\beta_{\min}, \beta_{\max}]$, $\rho = \beta_{\max}/\beta_{\min}$ and $\beta^\dagger$ from Proposition F.3; the value-margin row uses $\beta_{\min}$ when the scale is unknown.

Question	Amplitude a	Lower bound	Upper bound
Search, $\mathcal{F}(L_1)$	$\asymp \beta\varepsilon^2/L_1$	$\Omega(nR)$	—
Search, $\mathcal{F}_p(L)$	$\asymp \beta\varepsilon^{1+1/p}L^{-1/p}$	$\Omega\big(\frac{nL^{2/p}}{\beta^2\varepsilon^{2+2/p}}\big)$	—
Verify, $c'=2$, $\beta \le \beta^*$	$\beta\varepsilon^2/(4L_1)$	$\Omega(R)$	$O(n^2R)$
Verify, $c'=2$, $\beta \ge \beta^*$	$\beta\varepsilon^2/(4L_1)$	$\Omega(1)$	$O(1)$
Nonadaptive verify, $\beta\varepsilon^2/L_1 \le \frac{1}{64}$	$\beta\varepsilon^2/(4L_1)$	$\Omega(nR)$	$O(n^2R)$
Verify, unknown β, $\rho < c'$, $\beta_{\min} \le \beta^\dagger$	$\frac{\beta_{\min}\varepsilon^2(c'-\rho)^2}{2(1+\rho)L_1}$	$\Omega\big(\frac{L_1^2}{\beta_{\min}^2\varepsilon^4(c'-\rho)^4}\big)$	$\tilde{O}\big(\frac{n^3\rho^2L_1^2}{\beta_{\min}^2\varepsilon^4(c'-\rho)^4}\big)$
Verify, unknown β, $\rho \ge c'$	0, attained	impossible	impossible
Value margin m_0, no smoothness	βm_0	$\Omega((\beta m_0)^{-2})$	$\tilde{O}(\max\{1, (\beta_{\min}m_0)^{-2}\})$

2 THE COMPARISON GEOMETRY

Notation. Norms are Euclidean, $\mathbb{S}^{n-1}$ is the unit sphere with uniform measure μ, $\mathrm{gap}_f(x,y) = f(x) - f(y)$, and $\mathrm{osc}(h) = \sup h - \inf h \in [0,\infty]$. $\mathcal{F}(L_1)$ is the class of differentiable f with L_1-Lipschitz gradient, $\mathcal{F}(L_1, L_2)$ its subclass with L_2-Lipschitz Hessian, and $\mathcal{F}_p(L)$ the class of $f \in C^p$ whose p-th derivative is L-Lipschitz, all in operator norm. An *instance* is a pair (β, f), and βf is its *logit potential*. The tolerance is ε, the gap ratio is $c' > 1$ ($c' = 2$ unless stated), T counts queries and δ is a failure probability. Appendix A.2 lists all symbols.

Definition 2.1 (Bradley–Terry comparison oracle)**.** Fix $f\colon \mathbb{R}^n \to \mathbb{R}$ and a scale $\beta > 0$. A *query* is an ordered pair $(x, y) \in \mathbb{R}^n \times \mathbb{R}^n$; the oracle answers with a single bit $B \in \{0, 1\}$ distributed as

$$B \sim \mathrm{Bern}\big(\sigma(\beta\ \mathrm{gap}_f(x,y))\big), \qquad \mathbb{P}\big[B=1\big] \;=\; \sigma\big(\beta\,(f(x)-f(y))\big), \tag{2.1}$$

conditionally independent of all previous queries and answers given the pair queried. The bit $B = 1$ is read as the noisy assertion "$f(x) \ge f(y)$".

An algorithm has internal randomness independent of the oracle, chooses each pair measurably from it and the previous answers, and outputs a measurable function of both; it is *nonadaptive* if it chooses all pairs before seeing any answer. Quantitative lower bounds use a fixed worst-case budget and allow arbitrary pairs in $\mathbb{R}^n \times \mathbb{R}^n$, however far from the point of interest. Search must output x with $\|\nabla f(x)\| \le \varepsilon$, starting from x_0; verification is posed at a supplied point $\bar{x}$. By translation, $x_0 = \bar{x} = 0$ when convenient.

Definition 2.2 (ε-vs-$c'\varepsilon$ gradient-norm verification)**.** Fix a class $\mathcal{F}$ of objectives, a point $\bar{x}$, a tolerance $\varepsilon > 0$, and a gap ratio $c' > 1$. A *verifier* is an algorithm querying the BT oracle for $f \in \mathcal{F}$ that outputs accept or reject, under the promise

$$\text{(accept)}\quad \|\nabla f(\bar{x})\| \le \varepsilon, \qquad \text{or} \qquad \text{(reject)}\quad \|\nabla f(\bar{x})\| \ge c'\varepsilon.$$

It is *correct* if on every instance of the promise it answers correctly with probability at least $2/3$ (completeness and soundness $2/3$); the *δ-confidence variant* requires probability at least $1-\delta$ on every promise instance. On f satisfying neither branch the output may be arbitrary.

In the known-scale model the algorithm receives β. In the unknown-scale model it receives only $0 < \beta_{\min} \le \beta \le \beta_{\max} < \infty$, with $\rho = \beta_{\max}/\beta_{\min}$, and the class carries the budget $f(\bar{x}) - \inf f \le \Delta$ together with the following condition.

Assumption 2.3 ((A1): feasibility)**.** The parameters $c' > 1$, $\varepsilon > 0$, $L_1 > 0$, $\Delta > 0$ satisfy $c'\varepsilon \le \sqrt{2L_1\Delta}$.

Since $\|\nabla f(\bar{x})\| \le \sqrt{2L_1\Delta}$ by the descent lemma, the reject class is nonempty exactly under (A1).

2.1 What a transcript sees

Definition 2.4 (Logit distance). For functions $\phi, \psi\colon \mathbb{R}^n \to \mathbb{R}$ put

$$d(\phi,\psi) \;:=\; \tfrac{1}{2}\,\operatorname{osc}(\phi-\psi) \;=\; \tfrac{1}{2}\sup_{x,y\in\mathbb{R}^n}\big|(\phi-\psi)(x)-(\phi-\psi)(y)\big| \;\in\; [0,\infty],$$

the uniform distance between ϕ and ψ modulo additive constants: $d(\phi,\psi)=\inf_{c\in\mathbb{R}}\sup_x|\phi(x)-\psi(x)-c|$. For oracle instances $I=(\beta,f)$ and $I'=(\beta',f')$ write $d(I,I') := d(\beta f,\beta' f')$.

Proposition 2.5 (What a comparison transcript sees). *Let $P_I, P_{I'}$ be the transcript laws of a randomized adaptive algorithm run against instances $I=(\beta,f)$ and $I'=(\beta',f')$.*

(a) (Identity.) *If $d(I,I')=0$, then $P_I=P_{I'}$ and the output laws coincide, for every algorithm and every, possibly data-dependent, budget. If $d(I,I')>0$, some single query has different answer laws under I and I'.*

(b) (Contraction.) *If the algorithm makes T queries, then*

$$\mathrm{KL}(P_I\,\|\,P_{I'}) \;\le\; \frac{T}{2}\,d(I,I')^2.$$

The constant $1/2$ cannot be improved, even for $T=1$.

Part (a) is the scale non-identifiability of pairwise-comparison models (Shah et al., 2016) in transcript form. Part (b) combines the conditional chain rule for relative entropy (Polyanskiy & Wu, 2025, Theorem 2.16(c)), in which the query-selection terms cancel for adaptive algorithms, with the bound $\mathrm{KL}(\mathrm{Bern}(\sigma(u))\|\mathrm{Bern}(\sigma(v))) \le (u-v)^2/8$ (Lemmas B.1, B.4 and B.5). The distance d is the geometry of the oracle: like the oracle, it cannot see additive constants, and it compares potentials over the whole space, since queries may be arbitrarily far apart.

2.2 The amplitude of a question

Definition 2.6 (Amplitude of a promise problem). A *promise problem* $\mathcal{P}=(\mathcal{A},\mathcal{R})$ is a pair of nonempty sets of instances. An algorithm *solves $\mathcal{P}$ with confidence* $1-\delta$ if it outputs accept with probability at least $1-\delta$ on every instance in $\mathcal{A}$ and reject with probability at least $1-\delta$ on every instance in $\mathcal{R}$. The *amplitude* of $\mathcal{P}$ is

$$a(\mathcal{P}) \;:=\; \inf\big\{d(I,I') :\; I\in\mathcal{A},\; I'\in\mathcal{R}\big\},$$

and $\mathcal{P}$ is *silent* if $d(I,I')=0$ for some $I\in\mathcal{A}$ and $I'\in\mathcal{R}$.

Theorem 2.7 (The amplitude law). *Let $\mathcal{P}$ be a promise problem.*

(a) (Silence.) *If $\mathcal{P}$ is silent, no algorithm, at any budget, even unbounded and data-dependent, answers correctly with probability greater than $1/2$ on every instance.*

(b) (Two points.) *Every algorithm that solves $\mathcal{P}$ with confidence $2/3$ makes $T \ge 2\ln(9/8)/a(\mathcal{P})^2$ queries; with confidence $1-\delta$ and $\delta\le 1/16$, it makes $T\ge \ln(1/\delta)/a(\mathcal{P})^2$ queries.*

(c) (Many points.) *Let $I_1,\dots,I_M$, $M\ge 3$, be instances with $d(I_j,I_k)\le a$ for all j,k, and let Ψ map outputs to $\{1,\dots,M\}$ so that every output that is correct for I_j is mapped to j. Every algorithm whose output is correct with probability at least $2/3$ on each I_j makes $T\ge\big(\frac{4}{3}\ln M-2\ln 2\big)/a^2$ queries.*

In (b) and (c), T is a fixed worst-case budget, and $1/0=\infty$.

The proof takes a few lines from Proposition 2.5, via the Bretagnolle–Huber and Fano inequalities (Appendix C.1); the content lies in the choice of metric. All but two of the lower bounds in this paper are instances. Part (a) gives the unknown-scale impossibility, part (b) the two-point bounds on cost, and part (c) the search bounds. The exceptions are the universal confidence floor $\log_2(1/(2\delta))$, which counts answers rather than distances, and Theorem 4.2, which compares a mixture of instances with a single instance. The amplitude belongs to the question alone; what remains is to compute it and to measure algorithms against the resulting bounds.

3 How loud is a gradient

A gradient is a value step. Let two L_1-smooth objectives have gradients of norms ε and $c'\varepsilon$ at $\bar{x}$. Their difference h has a gradient of norm at least $\gamma = (c'-1)\varepsilon$ at $\bar{x}$ and a $2L_1$-Lipschitz gradient. Along the line through $\bar{x}$ in the direction of $\nabla h(\bar{x})$, the slope of h stays above $\gamma - 2L_1|s|$, so h rises by at least $\gamma^2/(2L_1)$ across the window $|s| \le \gamma/(2L_1)$. This is the whole-line inequality $\|q'\|_\infty^2 \le 2\|q\|_\infty\|q''\|_\infty$ of Landau and Hadamard (Landau, 1914; Hadamard, 1914), localized at a point. It is attained: bend one objective up and the other down on the window, so that their difference is a single smoothed step. The same computation with scales in a range gives the general formula.

Proposition 3.1 (The loudness of a gradient). *Let $c' > 1$, $\varepsilon, L_1 > 0$ and $0 < \beta_{\min} \le \beta_{\max} < \infty$, with $\rho = \beta_{\max}/\beta_{\min}$. Let $\mathcal{P}_\rho$ be ε-vs-$c'\varepsilon$ gradient-norm verification at $\bar{x}$ over the instances (β, f) with $\beta \in [\beta_{\min}, \beta_{\max}]$ and $f \in \mathcal{F}(L_1)$, and let $\mathcal{P}_\rho^\Delta$ be the same problem over $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \le \Delta\}$, where Assumption 2.3 holds. Then*

$$a(\mathcal{P}_\rho) \;=\; a(\mathcal{P}_\rho^\Delta) \;=\; \frac{\beta_{\min}\,\varepsilon^2\,(c'-\rho)_+^2}{2\,(1+\rho)\,L_1}.$$

Both problems are silent when $\rho \ge c'$; when $\rho < c'$, a pair of ridge functions attains the infimum. In particular, at a known scale β the amplitude of ε-vs-$c'\varepsilon$ verification over $\mathcal{F}(L_1)$ is $\beta(c'-1)^2\varepsilon^2/(4L_1)$.

The known-scale amplitude and Theorem 2.7(b) give the lower bounds of Theorem 4.1(iii), with constant $32\ln(9/8) \approx 3.769$ at $c' = 2$. Because the amplitude is exact, no other pair of instances improves this bound; the law in this form cannot give more.

Hiding the step: search. Search is harder than verification because the step may be anywhere. For a unit vector v in a packing of the sphere, place a smoothed cone of slope 2ε with apex at $3r_0v$: $f_v(x) = 2\varepsilon\chi(\|x - 3r_0v\|)$, where $\chi(r) = r$ outside a cap of radius r_0. Every ε-stationary point lies in the cap of the selected apex, so a successful search identifies v. The decisive feature is global. Moving the apex changes the cone by at most 2ε times the displacement at every point, so the members of the family are within logit distance $12\beta\varepsilon r_0$ of each other, however far apart the queried points are; a local Taylor argument would control only nearby probes. The cap radius is the width of the value step that carries slope ε: $r_0 = 2\varepsilon/L_1$ for a quadratic cap and $r_0 = 2\sqrt{\varepsilon/L_2}$ for a cubic one. Theorem 2.7(c) with $e^{\Omega(n)}$ apexes gives the following bounds; a two-point argument covers small n.

Theorem 3.2 (Search lower bounds). *Let $n \ge 1$, $\beta > 0$, $\varepsilon > 0$. Consider the known-scale BT oracle (Definition 2.1) and algorithms anchored at $x_0 = 0$.*

(a) (Gradient-Lipschitz.) *Set $\Delta_{\text{inst}} := 10\,\varepsilon^2/L_1$. Every randomized adaptive algorithm that, for every $f \in \mathcal{F}(L_1)$ with $f(x_0) - \inf f \le \Delta_{\text{inst}}$, outputs an ε-stationary point of f with probability at least $2/3$ must make*

$$T \;\ge\; c\,\frac{nL_1^2}{\beta^2\varepsilon^4} \quad \textit{queries.}$$

(b) (Hessian-Lipschitz.) *Assume $2\sqrt{\varepsilon L_2} \le L_1$ and set $\Delta_{\text{inst}} := (32/3)\,\varepsilon^{3/2}/\sqrt{L_2}$. Every randomized adaptive algorithm that, for every $f \in \mathcal{F}(L_1, L_2)$ with $f(x_0) - \inf f \le \Delta_{\text{inst}}$, outputs an ε-stationary point of f with probability at least $2/3$ must make*

$$T \;\ge\; c\,\frac{nL_2}{\beta^2\varepsilon^3} \quad \textit{queries.}$$

Both parts hold with $c = 1/71280$ for all $n \ge 1$, and with $c = 1/13824$ for $n \ge 56$.

The bounds hold for every budget above Δ_{inst} but do not establish a dependence on Δ (Appendices D and H).

The exponent law. The two rates of Theorem 3.2 are samples of one law. Under an L-Lipschitz p-th derivative, a value step that carries slope ε has width of order $(\varepsilon/L)^{1/p}$ and height of order $\varepsilon^{1+1/p}L^{-1/p}$, and by Kolmogorov's inequality no lower step exists (Kolmogorov, 1939). The search construction extends to every p with the cap $\Phi_p(y) = Q_p(\|y\|^2)$ inside the unit ball and $\Phi_p(y) = \|y\|$ outside, where Q_p is the degree-p Taylor polynomial of $\sqrt{s}$ at $s = 1$, so that the cap meets the cone to order p.

Theorem 3.3 (The exponent law). *Let $p \geq 1$ be an integer and $L > 0$, and let $\mathcal{F}_p(L)$ be the class of $f \in C^p(\mathbb{R}^n)$ whose p-th derivative is L-Lipschitz in operator norm. There are constants $K_p, H_p, \varsigma_p, k_p \in (0, \infty)$ depending only on p, with $K_1 = H_1 = 1$, such that the following hold at a known scale β. Let $r_p = (2\varepsilon K_p/L)^{1/p}$.*

(a) (Search.) *Every randomized adaptive algorithm anchored at $x_0 = 0$ that outputs an ε-stationary point with probability at least $2/3$ for every $f \in \mathcal{F}_p(L)$ with $f(x_0) - \inf f \leq \Delta_p := 2\big(3 - 4^{-p}\binom{2p}{p}\big)\varepsilon r_p$ makes*

$$T \;\geq\; \frac{n\, L^{2/p}}{17820\,(2K_p)^{2/p}\,\beta^2\,\varepsilon^{2+2/p}} \quad \textit{queries.}$$

The same holds over $\mathcal{F}(L_1) \cap \mathcal{F}_p(L)$ whenever $2\varepsilon H_p \leq L_1 r_p$.

(b) (Verification.) *If $c'\varepsilon\varsigma_1 \leq L_1(c'\varepsilon\varsigma_p/L)^{1/p}$, every algorithm that solves ε-vs-$c'\varepsilon$ verification over $\mathcal{F}(L_1) \cap \mathcal{F}_p(L)$ with confidence $2/3$ makes $T \geq 2\ln(9/8)\, L^{2/p}/\big(\varsigma_p^{2/p}\beta^2(c'\varepsilon)^{2+2/p}\big)$ queries.*

(c) (The exponent is Kolmogorov's.) *Every accept instance and every reject instance of ε-vs-$c'\varepsilon$ verification over $\mathcal{F}_p(L)$ are at logit distance at least $k_p\,\beta\,((c'-1)\varepsilon)^{1+1/p}L^{-1/p}$; hence no lower bound obtained from Theorem 2.7(b) yields a larger power of $1/\varepsilon$ than $2 + 2/p$.*

The exponent $2 + 2/p$ falls from 4 at $p = 1$ through 3 at $p = 2$ towards the parametric 2: for small ε, smoother objectives spread a gradient over a wider and taller step. Part (b) keeps the first-order constraint for a reason. Without it, steep odd terms saturate the channel, and for $p \geq 3$ verification is impossible at every budget although the amplitude is positive (Remark C.6): the amplitude law bounds cost from below, and saturation is a second obstruction. No algorithm for raw BT bits matches these search bounds yet.

4 Reading the potential: verification at a known scale

A verifier cannot read gradients, only logits. Every upper bound in this paper rests on one principle: comparisons can be *zoomed but not amplified*. A logit in $[-3, 3]$ can be estimated to precision η from $O(\eta^{-2})$ reads, since the inverse link has bounded derivative there, and moving probes closer scales logits down but never up. A decision that must resolve logit differences of size a thus costs $\max\{1, a^{-2}\}$ reads, up to confidence factors, if those differences can be brought to order one or below.

Theorem 4.1 (Known-scale verification: rate interval and saturation). *Fix $c' = 2$, and let $m(\beta, \varepsilon, n, \delta)$ be the query complexity of δ-confidence ε-vs-2ε gradient-norm verification over $\mathcal{F}(L_1)$ at a given $\bar{x}$ under the known-scale BT oracle. There are absolute constants C, C_0 such that, with $c_0 := 32\ln(9/8) \approx 3.77$:*

(i) (Upper bound, subsaturated scales.) *Set $\beta^* := 2nL_1/\varepsilon^2$. For all $0 < \beta \leq \beta^*$,*

$$m(\beta, \varepsilon, n, \delta) \;\leq\; C\,\frac{n^2 L_1^2}{\beta^2\varepsilon^4}\log(1/\delta);$$

for $\beta \geq \beta^*$ part (ii) takes over, so every scale is covered.

(ii) (Saturation.) *If $\beta \geq \beta^*$, then $m(\beta, \varepsilon, n, \delta) \leq C_0 \log(1/\delta)$: the cost collapses to a floor free of n, β, ε, and L_1.*

(iii) (Lower bounds.) *For every $\beta > 0$: $m \geq c_0\, L_1^2/(\beta^2\varepsilon^4)$ whenever $\delta \leq 1/3$; the stronger form*

$$m \;\geq\; 16\,\frac{L_1^2}{\beta^2\varepsilon^4}\ln(1/\delta)$$

holds for $\delta \leq 1/16$; and $m \geq \log_2\big(1/(2\delta)\big)$ unconditionally.

The verifier. For a uniform direction $u \in \mathbb{S}^{n-1}$, compare the antithetic pair $(\bar{x} + su, \bar{x} - su)$, whose logit is $2\beta s\langle \nabla f(\bar{x}), u\rangle + \beta R_u$ with $|R_u| \leq L_1 s^2$. The pair cancels the unknown offset $f(\bar{x})$. A clipped inverse-link estimate of this logit, squared and truncated, is a statistic whose mean separates the two hypotheses: with $g = \nabla f(\bar{x})$, $\mathbb{E}_u\langle g, u\rangle^2 = \|g\|^2/n$ and $\mathbb{E}_u\langle g, u\rangle^4 = 3\|g\|^4/(n(n+2))$,

so its relative variance is free of n and a median of block means over $O(\log(1/\delta))$ directions decides (Appendix E.1). The factor n^2 is the price of reading one direction. Smoothness forces $s \lesssim \varepsilon/(\sqrt{n}L_1)$, and a projection at the decision boundary has size $\varepsilon/\sqrt{n}$, so the logit to be resolved has order $\beta\varepsilon^2/(nL_1)$, the amplitude divided by n. Its inverse square is the number of reads per direction. All pairs are fixed in advance, so this verifier is nonadaptive.

Saturation is zooming. Once $\beta \geq \beta^* = 2nL_1/\varepsilon^2$, the per-direction amplitude exceeds a constant, and the verifier shrinks its radius to $s = \sqrt{n}/(6\beta\varepsilon)$. The logit to be read is then a constant, the curvature bias stays below the decision margin, and each direction needs $O(1)$ reads. The upper bound falls to a dimension-free $O(\log(1/\delta))$ (Theorem 4.1(ii)), which the floor $\log_2(1/(2\delta))$ of part (iii) matches. This is the $\max\{1, \cdot\}$ of the amplitude law: a comparison of amplitude above one logit is no harder to read than one of amplitude one. We claim an achievable envelope, not monotonicity of the minimax complexity in β. The noiseless oracle offers no such rescue, since at $\beta = \infty$ magnitudes cannot be read at any radius (Appendix G).

What a fixed query hears. The two-point bound and the upper bound differ by a factor n^2. A mixture over hidden directions recovers a linear factor when all queries are chosen in advance.

Theorem 4.2 (Nonadaptive verification needs $\Omega(n)$)**.** *Let $n \geq 2$ and suppose $\beta\varepsilon^2/L_1 \leq 1/64$. Every nonadaptive algorithm that, under the known-scale BT oracle, tests $\|\nabla f(0)\| \leq \varepsilon$ against $\|\nabla f(0)\| \geq 2\varepsilon$ over $f \in \mathcal{F}(L_1)$ with success probability at least $2/3$ on every promise instance must make*

$$T \;\geq\; \frac{n-1}{54} \cdot \frac{L_1^2}{\beta^2\varepsilon^4} \quad \textit{queries.}$$

The hard alternative randomizes the direction of a shifted cone. For a fixed query pair, the logit difference is a function of the hidden direction $v \in \mathbb{S}^{n-1}$ with Lipschitz constant of order $\beta\varepsilon^2/L_1$, so by the Poincaré inequality of the sphere its variance is at most a $1/(n-1)$ share of the squared amplitude: a fixed query hears a $1/(n-1)$ share of the variation. A χ^2 comparison of the mixture with the averaged null turns this into the factor $n-1$; a moment-generating-function bound for spherical chaos controls the off-diagonal terms (Lemma B.7), and the exact logistic score is kept throughout (Appendix E.2). With part (iii) of Theorem 4.1, a nonadaptive verifier needs $T \geq \max\{c_0, (n-1)/54\}L_1^2/(\beta^2\varepsilon^4)$ in this regime, with $c_0 = 32\ln(9/8)$ (Corollary E.5). For adaptive verifiers the dimension factor still lies in $[1, n^2]$. An adaptive verifier can steer its queries toward the hidden direction, and the average over directions no longer passes inside the per-query bound (Remark E.4). In amplitude terms, verifiers hear a gradient at an effective amplitude between a/n and a, nonadaptive ones at most at $a/\sqrt{n-1}$; where adaptive ones fall is open.

5 When the scale is uncertain

Suppose only $\beta \in [\beta_{\min}, \beta_{\max}]$ is known, with $\rho = \beta_{\max}/\beta_{\min}$, and consider ε-vs-$c'\varepsilon$ verification over $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \leq \Delta\}$. Now instances with opposite answers can share a potential: the reject instance $(\beta_{\min}, f_0)$ and the accept instance $(c'\beta_{\min}, f_0/c')$ have the same potential $\beta_{\min}f_0$, and both are admissible once $c'\beta_{\min} \leq \beta_{\max}$. Proposition 3.1 gives the complete picture: the amplitude is

$$a(\rho) = \frac{\beta_{\min}\,\varepsilon^2\,(c'-\rho)_+^2}{2(1+\rho)L_1}. \tag{5.1}$$

The worst pair gives the accept side the largest scale and the reject side the smallest. The observable gradient gap $\beta_{\min}c'\varepsilon - \beta_{\max}\varepsilon$ vanishes at $\rho = c'$, and the joint curvature budget $(\beta_{\max} + \beta_{\min})L_1$ turns it into a step of height quadratic in the gap. More consistent answers cannot separate instances that share a comparison law: what matters is the ratio of admissible scales, not their size.

Theorem 5.1 (Unknown-scale dichotomy)**.** *Let $c' > 1$, let $0 < \beta_{\min} \leq \beta_{\max} < \infty$ with $\rho = \beta_{\max}/\beta_{\min}$, and let $\varepsilon, L_1, \Delta > 0$ satisfy Assumption 2.3. Consider ε-vs-$c'\varepsilon$ gradient-norm verification at a given $\bar{x}$ over $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \leq \Delta\}$ under the BT oracle with unknown scale $\beta \in [\beta_{\min}, \beta_{\max}]$. Then the task is possible if and only if $\rho < c'$. More precisely:*

- *(i)* (Impossibility, $\rho \geq c'$.) *If $\rho \geq c'$, no algorithm — at any budget, even unbounded and data-dependent — answers correctly with probability greater than $1/2$ on both hypotheses. In particular the boundary $\rho = c'$ is on the impossible side.*

(ii) (Possibility, $\rho < c'$.) *If $\rho < c'$, for every $\delta \in (0,1)$ there is a verifier with failure probability at most δ, uniform over the class and over $\beta \in [\beta_{\min}, \beta_{\max}]$, of finite query cost (Proposition F.3 gives an explicit bound).*

Theorem 5.2 (The edge of identifiability)**.** *In the setting of Theorem 5.1, let $1 \le \rho < c'$. Every randomized adaptive algorithm that is correct with probability at least $2/3$ on every instance makes*

$$T \;\ge\; 8\ln(9/8)\,\frac{(1+\rho)^2\,L_1^2}{\beta_{\min}^2\,(c'-\rho)^4\,\varepsilon^4}$$

queries, and at confidence $1-\delta$ with $\delta \le 1/16$ it makes $T \ge 4(1+\rho)^2 L_1^2 \ln(1/\delta)/(\beta_{\min}^2(c'-\rho)^4\varepsilon^4)$ queries. Consequently, let $N^\star(\rho)$ be the least fixed worst-case budget that achieves confidence $1-\delta$, $\delta \le 1/16$, in the unknown-scale problem, and let $\beta^\dagger$ be as in Proposition F.3. Whenever $\beta_{\min} \le \beta^\dagger$,

$$\frac{4\,(1+\rho)^2 L_1^2 \ln(1/\delta)}{\beta_{\min}^2\,\varepsilon^4\,(c'-\rho)^4} \;\le\; N^\star(\rho) \;\le\; \frac{3744\,n^3\rho^2 L_1^2 \log(2n/\delta)}{\beta_{\min}^2\,\varepsilon^4\,(c'-\rho)^4} + n.$$

At fixed $(n, L_1, \Delta, \varepsilon, c', \delta, \beta_{\min})$ the complexity therefore grows exactly as $(c'-\rho)^{-4}$ as $\rho \uparrow c'$, and at fixed $\rho < c'$ exactly as $\beta_{\min}^{-2}$ as $\beta_{\min} \to 0$.

The dichotomy is the zero set of the amplitude, and the edge rate is its quadratic vanishing: Theorem 2.7(b) turns $a(\rho)^{-2} \propto (c'-\rho)^{-4}$ into the lower bound. The matching upper bound comes from the tester of Proposition F.3. It estimates the observable gradient norm $G = \|\nabla(\beta f)(\bar{x})\|$ by coordinate central differences and compares it with the midpoint $\varepsilon(\beta_{\max} + c'\beta_{\min})/2$ of the two ranges $G \le \beta_{\max}\varepsilon$ and $G \ge \beta_{\min}c'\varepsilon$; it never estimates β. Its probe radius is the smaller of two radii: one keeps the finite-difference bias below the margin, the other keeps every read logit in $[-9/8, 9/8]$, using the a priori bound $\|\nabla f(\bar{x})\| \le \sqrt{2L_1\Delta}$.

Two consequences deserve plain statement. Reliable comparisons can still be uncalibrated: the impossibility at $\rho \ge c'$ rests on equal observation laws, so no number of repetitions, no distant probe and no adaptivity helps. Reliable comparisons do help a calibrated verifier: at fixed $\rho < c'$ the tester's cost stays bounded as $\beta_{\min} \to \infty$, because the readable radius shrinks in proportion (Proposition F.3). How the saturated cost depends on $c' - \rho$ is open (Remark C.3). The inequality form of the budget matters for identification. An exact budget $f(\bar{x}) - \inf f = \Delta$ would identify β from the infimum $-\beta\Delta$ of the potential $\beta(f - f(\bar{x}))$, so the impossibility witness fails for that class (Remark H.1).

6 Values without derivatives

A question need not ask about a derivative. Fix a radius $r > 0$, a margin $m_0 > 0$ and $p_0 = 3/10$. With $\mathrm{gap}(u) = f(\bar{x} + ru) - f(\bar{x})$ for $u \in \mathbb{S}^{n-1}$, consider

$$\mathcal{M}^+(m_0):\ \mathrm{gap}(u) \ge m_0 \text{ for every } u \in \mathbb{S}^{n-1}, \qquad \mathcal{M}^-(m_0):\ \mu\{u : \mathrm{gap}(u) \le -m_0\} \ge p_0. \tag{6.1}$$

The tester draws independent uniform directions, repeats each comparison $(\bar{x} + ru, \bar{x})$, takes the majority sign, and rejects if any majority is negative.

Theorem 6.1 (Smoothness-free ordinal value-margin verification)**.** *Let $f\colon \mathbb{R}^n \to \mathbb{R}$ be everywhere finite, with $u \mapsto f(\bar{x} + ru)$ μ-measurable;* no *smoothness, boundedness, or gap-budget condition is imposed. Consider the promise problem (6.1) under the BT oracle with scale β known, or unknown in $[\beta_{\min}, \beta_{\max}]$ with $\beta_{\min} > 0$ (for known β read $\beta_{\min} = \beta$ below).*

(i) (Upper bound.) *The majority-vote tester above with $m = \lceil 8\ln(2/\delta) \rceil$ directions and $k = \lceil\, 2\ln(2m/\delta)/\tanh^2(\beta_{\min}m_0/2)\,\rceil$ reads per direction accepts every $f \in \mathcal{M}^+(m_0)$ and rejects every $f \in \mathcal{M}^-(m_0)$ with probability at least $1-\delta$, using*

$$T \;=\; mk \;=\; O\Big(\max\big\{1,\,(\beta_{\min}m_0)^{-2}\big\}\cdot\log(1/\delta)\cdot\log\big(\log(1/\delta)/\delta\big)\Big)$$

comparisons. Its decision rule is scale-free; $\beta_{\min}$ enters only the read count k.

(ii) (Lower bound, no regime cap.) *For every $m_0 > 0$, $r > 0$, $n \geq 1$ and every true scale β: any randomized adaptive tester correct with probability at least $2/3$ on both classes needs*

$$T \;\geq\; \frac{2\ln(9/8)}{(\beta m_0)^2} \;\approx\; \frac{0.2356}{(\beta m_0)^2},$$

and at confidence $\delta \leq 1/16$, $T \geq \ln(1/\delta)/(\beta m_0)^2$.

The amplitude of this question is exactly βm_0, or $\beta_{\min} m_0$ when the scale is unknown: an accept and a reject instance differ by at least $2m_0$ across some probe pair, and a bowl minus a one-signed hemispherical bump of height $2m_0$ attains this (Remark F.5). No smoothness step stands between the promise and the potential. This is why the theorem needs no smoothness and no dimension factor, and why a lower scale bound suffices: a larger β only sharpens every relevant sign; it certifies a value margin, not stationarity. For $f \in \mathcal{F}(L_1)$, a gradient of norm at least ε creates descent directions of mass at least $3/10$ at radius $r = \varepsilon/(2\sqrt{n}L_1)$, so the same tester also rejects such points (Corollary F.7). The compatible margin is then at most $\varepsilon^2/(8nL_1)$, and substituting it restores the dimension and ε dependence of gradient verification: the test does not evade the cost of gradients.

The ladder of oracles. Each oracle determines f up to transformations it cannot see: additive constants at a known scale, maps $f \mapsto \lambda f + c$ with λ in a range of ratio ρ at an unknown scale, and all strictly increasing reparametrizations for exact comparisons. Certifiable properties are the invariants of this ambiguity, and the amplitude measures how far a property is from invariant. Under the largest ambiguity, gradient norms, criticality and even the choice among candidate points are silent (Theorems G.3, G.5 and G.6); value margins survive on every rung.

7 Discussion and conclusion

Can preferences tell an optimizer when to stop? Under Bradley–Terry feedback one quantity decides: the amplitude, how far the logit potential must move to change the certificate. A stopping rule based on $\|\nabla f(\bar{x})\| \leq \varepsilon$ is possible when the preference scale is pinned to within a factor smaller than the slack c' of the certificate. At a known scale it costs between $L_1^2/(\beta^2\varepsilon^4)$ and $n^2L_1^2/(\beta^2\varepsilon^4)$ comparisons up to $\log(1/\delta)$, and only $O(\log(1/\delta))$ when comparisons are very consistent. It is impossible when the scale is uncertain by a factor c' or more, and its cost diverges as $(c' - \rho)^{-4}$ on the way there. Value-margin rules need neither smoothness nor a scale. The questions change from theorem to theorem; the method does not: compute a distance, read its inverse square.

The open problems are about entropy and dimension, not noise. Does adaptivity let a verifier hear a hidden gradient direction at the full amplitude, or must every verifier pay $\Omega(nL_1^2/(\beta^2\varepsilon^4))$? Is there a search algorithm for raw BT bits that attains n/a^2, and how does search cost scale with a larger gap budget Δ, where charging each of several wells separately fails because one comparison can span them all? Does the saturated unknown-scale cost grow as $\rho \to c'$ (Appendix H)?

The amplitude law is one half of a characterization. It bounds cost from below by distance and can be far from tight (Remark C.6): when potentials may be arbitrarily steep, a comparison separates two instances only where their logits are not saturated, and the information vanishes although the amplitude does not. The divergence of one comparison is the squared logit difference weighted by the slope σ' of the link at an intermediate logit, and our upper bounds control that weight by zooming. A two-sided theory would measure questions in this weighted geometry.

Our guarantees count oracle bits and allow queries at arbitrary points, so shrinking a probe radius is free; a minimum query resolution would change the zooming principle and require a different analysis. The bounds describe information under BT access, not computation or the empirical behaviour of a stopping rule. The distance-based lower bounds use the logistic link only through the global bound $\mathrm{KL} \leq (u - v)^2/8$ on logits u, v, and extend to any link with a bound of this form; the nonadaptive mixture bound uses the logistic score itself, and the upper bounds invert the logistic link and require knowing the response law. For preference learning the message is at the level of the oracle: more consistent raters help a calibrated verifier, while uncertainty about the response temperature can make an absolute stationarity claim unidentifiable and leaves value margins as the certificates that survive. Calibrating the temperature, for instance with anchor pairs of known value gap, shrinks the ratio ρ below c' and makes the certificate possible again; value-margin certificates need only a lower bound on the scale.

AI USE STATEMENT

The original research ideas and initial manuscript draft were developed by the authors. AI tools were used to assist with checking mathematical details and the correctness of arguments, filling in routine proof steps, coding assistance and language polishing. The authors take full responsibility for the content and correctness of the final manuscript.

REPRODUCIBILITY STATEMENT

All results are mathematical, stated with their assumptions, and proved in the appendices: Appendices B and C develop the channel lemmas and the amplitude law, Appendices D and E prove the search and verification bounds, Appendix F the unknown-scale and value-margin results, and Appendix G the noiseless baseline. The supplementary scripts recompute every numerical constant quoted in the paper and check the extremal constructions numerically; these checks illustrate the proofs and are not part of them.

REFERENCES

Yossi Arjevani, Yair Carmon, John C. Duchi, Dylan J. Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *Mathematical Programming*, 2023.

Keith Ball. An elementary introduction to modern convex geometry. In *Flavors of Geometry*, volume 31 of *MSRI Publications*. Cambridge University Press, 1997.

Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952. doi: 10.1093/biomet/39.3-4.324.

Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points ii: First-order methods. *Mathematical Programming*, 2019.

Yair Carmon, John C. Duchi, Oliver Hinder, and Aaron Sidford. Lower bounds for finding stationary points i. *Mathematical Programming*, 184(1):71–120, 2020.

Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios N. Angelopoulos, Tianle Li, Dacheng Li, Banghua Zhu, Hao Zhang, Michael I. Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating LLMs by human preference. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, volume 235 of *Proceedings of Machine Learning Research*, pp. 8359–8388, 2024.

Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems 30 (NeurIPS)*, pp. 4299–4307, 2017.

Taha El Bakkali, Rayane Bouftini, Qiuyi Zhang, and Omar Saadi. Noisy pairwise-comparison random search for smooth nonconvex optimization, 2026a. URL https://arxiv.org/abs/2601.21166.

Taha El Bakkali, El Mahdi Chayti, and Omar Saadi. Nonsmooth optimization with zeroth order comparison feedback, 2026b. URL https://arxiv.org/abs/2602.05622.

Javier González, Zhenwen Dai, Andreas Damianou, and Neil D. Lawrence. Preferential Bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, volume 70 of *Proceedings of Machine Learning Research*, pp. 1282–1291, 2017.

Jacques Hadamard. Sur le module maximum d'une fonction et de ses dérivées. *Bulletin de la Société Mathématique de France*, 42:68–72, 1914.

Kevin G. Jamieson, Robert Nowak, and Ben Recht. Query complexity of derivative-free optimization. In *Advances in Neural Information Processing Systems (NeurIPS) 25*, 2012.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1–42, 2016.

Andrei N. Kolmogorov. On inequalities between the upper bounds of the successive derivatives of an arbitrary function on an infinite interval. *Uchenye Zapiski Moskovskogo Gosudarstvennogo Universiteta, Matematika*, 30:3–16, 1939. English translation: *American Mathematical Society Translations, Series 1*, 2:233–243, 1962.

Edmund Landau. Einige Ungleichungen für zweimal differentiierbare Funktionen. *Proceedings of the London Mathematical Society (2)*, 13:43–49, 1914.

Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.

Michel Ledoux. Concentration of measure and logarithmic sobolev inequalities. In *Séminaire de Probabilités XXXIII*, volume 1709 of *Lecture Notes in Mathematics*, pp. 120–216. Springer, 1999.

Aleksandr Lobanov, Alexander Gasnikov, and Andrei Krasnov. Acceleration exists! optimization problems when oracle can only compare objective function values. In *Advances in Neural Information Processing Systems (NeurIPS) 37*, 2024.

Carl E. Mueller and Fred B. Weissler. Hypercontractivity for the heat semigroup for ultraspherical polynomials and on the n-sphere. *Journal of Functional Analysis*, 48(2):252–283, 1982. doi: 10.1016/0022-1236(82)90069-6.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems 35 (NeurIPS)*, pp. 27730–27744, 2022.

Yury Polyanskiy and Yihong Wu. *Information Theory: From Coding to Learning*. Cambridge University Press, 2025. ISBN 9781108832908.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems 36 (NeurIPS)*, pp. 53728–53741, 2023.

O. S. Rothaus. Logarithmic Sobolev inequalities and the spectrum of Sturm–Liouville operators. *Journal of Functional Analysis*, 39:42–56, 1980. doi: 10.1016/0022-1236(80)90018-x.

Aadirupa Saha, Tomer Koren, and Yishay Mansour. Dueling convex optimization. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pp. 9245–9254, 2021.

Aadirupa Saha, Tomer Koren, and Yishay Mansour. Dueling convex optimization with general preferences, 2022. URL https://arxiv.org/abs/2210.02562.

Aadirupa Saha, Vitaly Feldman, Yishay Mansour, and Tomer Koren. Faster convergence with multiway preferences. In *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 238, pp. 433–441, 2024.

Nihar B. Shah, Sivaraman Balakrishnan, Joseph Bradley, Abhay Parekh, Kannan Ramchandran, and Martin J. Wainwright. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *Journal of Machine Learning Research*, 17(58):1–47, 2016.

V. N. Smirnov, K. M. Kazistova, I. A. Sudakov, Valentin Leplat, Alexander Gasnikov, and Aleksandr Lobanov. Ruppert–polyak averaging for stochastic order oracle, 2024. URL https://arxiv.org/abs/2411.15866.

Zhiwei Tang, Dmitry Rybin, and Tsung-Hui Chang. Zeroth-order optimization meets human feedback: Provable learning via ranking oracles, 2023. URL https://arxiv.org/abs/2303.03751.

William (Lord Kelvin) Thomson. *Popular Lectures and Addresses*, volume 1. Macmillan, London, 1889. Lecture "Electrical Units of Measurement", delivered 3 May 1883.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, 2009. doi: 10.1007/b13794.

Helin Wang, Chenyi Zhang, Xiwen Tao, Yexin Zhang, and Tongyang Li. Finding stationary points by comparisons, 2026. URL https://arxiv.org/abs/2606.27082.

Siyi Wang, Zifan Wang, and Karl Henrik Johansson. Preference-based optimization from noisy pairwise comparisons, 2025. URL https://arxiv.org/abs/2512.18511.

Fred B. Weissler. Logarithmic Sobolev inequalities and hypercontractive estimates on the circle. *Journal of Functional Analysis*, 37:218–234, 1980. doi: 10.1016/0022-1236(80)90042-7.

Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The K-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.

ROADMAP TO THE APPENDIX

The appendix contains complete proofs. Appendix A shows how the results depend on one another and collects the notation. Appendix B proves the channel lemmas behind Proposition 2.5, and Appendix C proves the amplitude law and the three results built on it: the loudness of a gradient, the edge of identifiability and the exponent law. The remaining appendices prove the search, verification, unknown-scale and value-margin theorems, establish the noiseless baseline, and discuss open problems.

Topic	Appendix	Main text
Dependency of results and notation	A	whole paper
Channel lemmas and spherical chaos	B	Section 2.1, Theorem 4.2
Amplitude law, loudness, edge, exponent law	C	Sections 2, 3 and 5
Search lower bounds	D	Theorem 3.2
Known-scale verification	E	Section 4
Unknown scale and value margins	F	Sections 5 and 6
Noiseless comparisons	G	Sections 1 and 6
Further discussion and open problems	H	Section 7

A DEPENDENCY OF RESULTS, AND NOTATION

A.1 DEPENDENCY OF RESULTS

Every lower bound descends from Proposition 2.5, through the amplitude law; the upper bounds are separate constructions, measured against it. Table 2 lists, for each main result, the results it uses and where it is proved.

Table 2: Dependency of the main results. "Uses" lists the results a proof relies on beyond standard facts.

Result	Uses	Proof
Lemmas B.1, B.4 and B.5 (channel lemmas)	—	Appendix B
Proposition 2.5 (what a transcript sees)	Lemmas B.1, B.4 and B.5	Appendix C.1
Theorem 2.7 (amplitude law)	Proposition 2.5	Appendix C.1
Proposition 3.1 (loudness of a gradient)	Hadamard-type integration	Appendix C.2
Theorem 3.2 (search)	Lemma B.4, Fano (an instance of Theorem 2.7(c)), shifted-cone packing	Appendix D
Theorem 3.3 (exponent law)	Theorem 2.7, Lemma C.4, Steps 0 and 3–5 of Appendix D, Kolmogorov's inequality	Appendix C.4
Theorem 4.1 (known-scale verification)	(i)–(ii) direct; (iii) Proposition 3.1 and Theorem 2.7	Appendix E.1
Theorem 4.2 (nonadaptive)	Fact E.2 (shifted cones of Appendix D), Lemma B.7, sphere Poincaré	Appendix E.2
Theorem 5.1 (unknown-scale dichotomy)	Lemma B.5; possible side direct	Appendix F.1
Proposition F.3 (bounded cost)	Theorem 5.1 setting	Appendix F.2
Theorem 5.2 (edge of identifiability)	Propositions 3.1 and F.3 and Theorem 2.7	Appendix C.3
Theorem 6.1 (value margins)	Lemma B.4	Appendix F.3
Corollary F.7 (gradient soundness)	Theorem 6.1, Lemma F.8	Appendix F.4
Theorems G.3, G.5 and G.6 (noiseless)	Lemma G.2	Appendix G

A.2 NOTATION

Symbol	Meaning	Defined in
σ, β	logistic link; comparison scale	(1.1)
(β, f), βf	instance; its logit potential	Section 2
$\mathrm{gap}_f(x, y)$	$f(x) - f(y)$	Section 2
$\mathrm{osc}(h)$	$\sup h - \inf h$	Section 2
$d(\phi, \psi)$, $d(I, I')$	logit distance $\frac{1}{2}\,\mathrm{osc}(\phi - \psi)$	Definition 2.4
$\mathcal{P} = (\mathcal{A}, \mathcal{R})$, $a(\mathcal{P})$	promise problem; its amplitude	Definition 2.6
$\mathcal{F}(L_1)$, $\mathcal{F}(L_1, L_2)$, $\mathcal{F}_p(L)$	smoothness classes	Section 2
ε, c'	tolerance; gap ratio ($c' = 2$ unless stated)	Definition 2.2
T, δ	number of queries; failure probability	Section 2
$\bar{x}$, x_0	point to verify; search anchor	Section 2
$\beta_{\min}, \beta_{\max}, \rho$	scale range; $\rho = \beta_{\max}/\beta_{\min}$	Section 2
Δ, (A1)	gap budget; $c'\varepsilon \le \sqrt{2L_1\Delta}$	Assumption 2.3
β^*	saturation scale $2nL_1/\varepsilon^2$	Theorem 4.1
$\beta^\dagger$	readability crossover scale	Proposition F.3
Δ_{inst}	gap budget of the search instances	Theorem 3.2
r_0, r_p, K_p	cap radii; cap constant	Section 3, Lemma C.4
r, m_0, p_0, $\mathcal{M}^+$, $\mathcal{M}^-$	value-margin promise	(6.1)
m_{s}	margin ceiling $\varepsilon^2/(8nL_1)$	Corollary F.7
μ, $\mathbb{S}^{n-1}$	uniform measure; unit sphere	Section 2

B PROOFS OF THE TOOLKIT LEMMAS

This appendix states and proves the three lemmas behind Proposition 2.5: the sharp channel bound (Lemma B.1), the adaptive transcript decomposition (Lemma B.4), and the scale lemma (Lemma B.5). It then proves a moment-generating-function bound for spherical chaos (Lemma B.7), used in the proof of Theorem 4.2. Appendix C assembles the first three into Proposition 2.5 and the amplitude law.

B.1 THE BRADLEY–TERRY CHANNEL: A SHARP KL BOUND

Lemma B.1 (Sharp KL bound for the BT channel). *For $\theta \in \mathbb{R}$ let $B_\theta := \mathrm{Bern}(\sigma(\beta\theta))$. Then for all $\theta_1, \theta_2 \in \mathbb{R}$,*

$$\mathrm{KL}\big(B_{\theta_1} \,\|\, B_{\theta_2}\big) \;\le\; \frac{\beta^2}{8}\,(\theta_1 - \theta_2)^2.$$

The bound is global — no small-gap restriction — and the constant $1/8$ is sharp: the ratio of the two sides tends to 1 as $\theta_1, \theta_2 \to 0$.

Proof. Fix $\beta > 0$.

Setup (the logistic–Bernoulli exponential family). Write the Bernoulli law on $\{0, 1\}$ in canonical form with natural parameter η:

$$p(x; \eta) = \exp\big(\eta x - A(\eta)\big), \qquad x \in \{0, 1\}, \qquad A(\eta) := \log(1 + e^\eta).$$

Then $p(1; \eta) = e^\eta/(1 + e^\eta) = \sigma(\eta)$ and $p(0; \eta) = 1 - \sigma(\eta)$, so $B_\theta = \mathrm{Bern}(\sigma(\beta\theta))$ is exactly the member with natural parameter $\eta = \beta\theta$. The log-partition A is C^∞ with

$$A'(\eta) = \sigma(\eta) \quad \text{(the mean parameter)}, \qquad A''(\eta) = \sigma(\eta)\big(1 - \sigma(\eta)\big) = \sigma'(\eta). \tag{B.1}$$

Step 1 (KL is the Bregman divergence of A; the direction fixed explicitly). For members with natural parameters η_1, η_2,

$$\begin{aligned}\mathrm{KL}(p_{\eta_1} \,\|\, p_{\eta_2}) &= \mathbb{E}_{\eta_1}\big[\log p(X; \eta_1) - \log p(X; \eta_2)\big] = \mathbb{E}_{\eta_1}\big[(\eta_1 - \eta_2)X - A(\eta_1) + A(\eta_2)\big] \\ &= (\eta_1 - \eta_2)\,A'(\eta_1) + A(\eta_2) - A(\eta_1) = A(\eta_2) - A(\eta_1) - A'(\eta_1)(\eta_2 - \eta_1),\end{aligned} \tag{B.2}$$

using $\mathbb{E}_{\eta_1}[X] = A'(\eta_1)$ from (B.1). The gradient is anchored at the *first* KL argument η_1 — the one place a sign could hide. Setting $a := \beta\theta_1 = \eta_1$ and $b := \beta\theta_2 = \eta_2$, (B.2) reads

$$\mathrm{KL}\big(B_{\theta_1} \,\|\, B_{\theta_2}\big) = A(b) - A(a) - A'(a)(b-a).$$

Step 2 (exact second-order remainder). A is C^2, so by Taylor's theorem with the Lagrange remainder, expanding A about a and evaluating at b, there is ξ between a and b ($\xi = a$ if $a = b$) with $A(b) = A(a) + A'(a)(b-a) + \frac{1}{2}A''(\xi)(b-a)^2$. Substituting into the Bregman form of Step 1, the terms $A(a) + A'(a)(b-a)$ cancel and

$$\mathrm{KL}\big(B_{\theta_1} \,\|\, B_{\theta_2}\big) = \tfrac{1}{2}\, A''(\xi)\,(b-a)^2. \tag{B.3}$$

This is an *equality*, valid for all a, b: the Bregman remainder is exactly one curvature evaluation, with no regime split.

Step 3 (global curvature bound $A'' \le 1/4$). By (B.1), $A''(\xi) = \sigma(\xi)(1-\sigma(\xi)) = s(1-s)$ with $s := \sigma(\xi) \in (0,1)$, and $s(1-s) \le 1/4$ for every s, with the maximum at $s = 1/2$, i.e. $\xi = 0$. Hence $A''(\xi) \le 1/4$ for *every* $\xi \in \mathbb{R}$. This is what makes the bound global: near saturation ($|\xi|$ large), $A''(\xi) \to 0$ and the KL is automatically small; the worst case is $\xi = 0$, where $A''(0) = 1/4$.

Step 4 (combine and substitute). From (B.3) and Step 3, $\mathrm{KL} \le \frac{1}{2} \cdot \frac{1}{4}\,(b-a)^2 = \frac{1}{8}(b-a)^2$, and $(b-a)^2 = \beta^2(\theta_1-\theta_2)^2$, which is the claimed bound.

Sharpness. As $\theta_1, \theta_2 \to 0$ with $\theta_1 \neq \theta_2$, the intermediate point satisfies $\xi \to 0$, so $A''(\xi) \to A''(0) = 1/4$ by continuity, and $\mathrm{KL} \,/\, \big(\frac{\beta^2}{8}(\theta_1-\theta_2)^2\big) \to 1$: the constant $1/8$ is attained in the limit and cannot be lowered.

□

Remark B.2 (Global validity; behaviour under saturation). No second, large-gap form of the bound is needed: the ratio $\mathrm{KL} \,/\, \big(\beta^2(\theta_1-\theta_2)^2\big) \le 1/8$ holds globally. When both logits saturate on the *same* side ($\beta\theta_1$ and $\beta\theta_2$ both large positive, or both large negative), the two channels merge and the KL tends to 0; when they saturate on *opposite* sides the KL grows, but $(\theta_1-\theta_2)^2$ grows faster. The single quadratic bound therefore suffices for every application in this paper.

B.2 The adaptive transcript-KL decomposition

The proof instantiates the chain rule for relative entropy in its general conditional form (Polyanskiy & Wu, 2025, Theorem 2.16(c), eq. (2.26)); we state the form we use.

Fact B.3 (Chain rule for relative entropy on standard Borel spaces). *Let $Z = (Z_1, \dots, Z_m)$ be a finite random sequence in which each Z_k takes values in a standard Borel space, and let P and Q be two laws of Z. Then, as an identity in $[0, \infty]$,*

$$\mathrm{KL}(P \,\|\, Q) = \sum_{k=1}^{m} \mathbb{E}_P\Big[\, \mathrm{KL}\big(P(Z_k \in \cdot \mid Z_{1:k-1}) \,\big\|\, Q(Z_k \in \cdot \mid Z_{1:k-1})\big)\Big],$$

where the conditional laws are regular conditional probabilities and the outer expectations are under P, the first argument.

In the bandit literature this fact is usually invoked in a specialization that further resolves the sum into per-arm pull counts over *finitely many* arms (Lattimore & Szepesvári, 2020, Lemma 15.1; Kaufmann et al., 2016, Lemma 1). Our action space is the continuum $\mathbb{R}^n \times \mathbb{R}^n$, so we cite the general conditional form above and prove only its instantiation.

Lemma B.4 (Adaptive transcript KL decomposition). *Let A be any randomized adaptive algorithm making T queries to the BT oracle, run against two instances (β, f) and (β', f'), and let P, P' be the two laws of its transcript. Then*

$$\mathrm{KL}\big(P \,\|\, P'\big) \;\le\; \frac{1}{8}\sum_{t=1}^{T} \mathbb{E}_P\Big[\big(\beta\,\mathrm{gap}_f(X_t, Y_t) - \beta'\,\mathrm{gap}_{f'}(X_t, Y_t)\big)^2\Big] \;\le\; \frac{T}{2}\, d(\beta f, \beta' f')^2,$$

where (X_t, Y_t) is the t-th queried pair and the expectation is under P. For a common scale $\beta = \beta'$, writing $P_f, P_{f'}$ for the two laws, the middle term is $\frac{\beta^2}{8}\sum_t \mathbb{E}_f[(\mathrm{gap}_f(X_t,Y_t) - \mathrm{gap}_{f'}(X_t,Y_t))^2]$; in particular, if $|f - f'| \le a$ pointwise, then $\mathrm{KL}(P_f \| P_{f'}) \le (\beta^2/2)\, a^2 T$.

Proof. **Step 0 (measurability boilerplate: Polish spaces and regular conditional probabilities).** The action (query) space $\mathbb{R}^n \times \mathbb{R}^n$ and the response space $\{0,1\}$ are Polish, hence the transcript space $\big((\mathbb{R}^n \times \mathbb{R}^n) \times \{0,1\}\big)^T$ is standard Borel and regular conditional probabilities exist for every conditioning used below; each conditional law and conditional KL is therefore well defined. The chain rule of Fact B.3 holds in $[0,\infty]$, so no finiteness assumption is needed. Throughout, T is a fixed (deterministic) query budget; data-dependent stopping times are not treated.

Setup (the interaction process). Fix the algorithm A. In round t it (i) selects a query pair (X_t, Y_t) from the observed history $H_{t-1} = \big((X_1,Y_1), B_1, \ldots, (X_{t-1},Y_{t-1}), B_{t-1}\big)$ and its internal randomness, then (ii) receives a Bernoulli response B_t. A is a *comparison* algorithm: it accesses the environment (the instance (β, f) or (β', f')) only through oracle responses — it never reads the objective directly. P, P' are the laws of the full transcript $\tau = \big((X_1,Y_1), B_1, \ldots, (X_T,Y_T), B_T\big)$ under the two environments; for a common scale we also write $P_f, P_{f'}$. We include the algorithm's internal randomness U as an initial entry of τ: its law is the same under both environments, so its chain-rule term vanishes, and given (U, H_{t-1}) the query is determined. Any output or decision is then a measurable function of τ, so by the data-processing inequality its KL is at most $\mathrm{KL}(P\|P')$; bounding the transcript KL suffices.

Step 1 (chain rule over the transcript). τ is a length-$2T$ sequence on standard Borel spaces, so Fact B.3 gives the *equality*

$$\mathrm{KL}(P \,\|\, P') = \sum_{t=1}^{T} \Big\{ \mathbb{E}_P\Big[\mathrm{KL}\big(P\big((X_t,Y_t) \in \cdot \mid H_{t-1}\big) \,\big\|\, P'\big((X_t,Y_t) \in \cdot \mid H_{t-1}\big)\big)\Big] + \mathbb{E}_P\Big[\mathrm{KL}\big(P\big(B_t \in \cdot \mid H_{t-1}, X_t, Y_t\big) \,\big\|\, P'\big(B_t \in \cdot \mid H_{t-1}, X_t, Y_t\big)\big)\Big] \Big\}, \quad \text{(B.4)}$$

the first braces being the *action terms* and the second the *response terms*; each outer expectation is under the *first* measure P, and $P(\cdot \mid \cdot)$, $P'(\cdot \mid \cdot)$ denote regular conditional laws under the two environments.

Step 2 (action terms vanish: zero conditional KL). Given H_{t-1}, the conditional law of the query pair (X_t, Y_t) is the algorithm's policy kernel $\pi_t(\cdot \mid H_{t-1})$ — the *same* stochastic kernel under both environments, because A chooses its queries from the observable history and its own coins, with no dependence on the environment (the comparison property of the Setup). Hence $P\big((X_t,Y_t) \in \cdot \mid H_{t-1}\big) = P'\big((X_t,Y_t) \in \cdot \mid H_{t-1}\big) = \pi_t(\cdot \mid H_{t-1})$ for every value of H_{t-1}, the conditional KL is 0 pointwise, and every action term in (B.4) is 0. (This is where the shared internal randomness and the policy factor out.)

Step 3 (response terms are channel divergences, bounded by Lemma B.1). The oracle is memoryless: given the current query $(X_t, Y_t) = (x, y)$, the response is Bernoulli, independent of the rest of the history, with

$$P\big(B_t \in \cdot \mid H_{t-1}, (X_t,Y_t) = (x,y)\big) = \mathrm{Bern}\big(\sigma(\beta\ \mathrm{gap}_f(x,y))\big),$$

and likewise under the second instance with logit $\beta'\,\mathrm{gap}_{f'}(x,y)$. By Lemma B.1 with unit scale, $\theta_1 = \beta\,\mathrm{gap}_f(X_t,Y_t)$ and $\theta_2 = \beta'\,\mathrm{gap}_{f'}(X_t,Y_t)$,

$$\mathrm{KL}\big(P(B_t \in \cdot \mid H_{t-1}, X_t, Y_t) \,\big\|\, P'(B_t \in \cdot \mid H_{t-1}, X_t, Y_t)\big) \le \frac{1}{8}\big(\beta\,\mathrm{gap}_f(X_t,Y_t) - \beta'\,\mathrm{gap}_{f'}(X_t,Y_t)\big)^2$$

pointwise in (H_{t-1}, X_t, Y_t). Taking $\mathbb{E}_P$ and summing over t (the action terms being 0 by Step 2) yields the first claimed inequality; for $\beta = \beta'$ it is the displayed common-scale form.

Step 4 (uniform-closeness corollary). For every pair (x, y),

$$\begin{aligned}\left|\beta \operatorname{gap}_f(x,y) - \beta' \operatorname{gap}_{f'}(x,y)\right| &= \left|(\beta f - \beta' f')(x) - (\beta f - \beta' f')(y)\right| \\ &\le \operatorname{osc}(\beta f - \beta' f') = 2\, d(\beta f, \beta' f'),\end{aligned}$$

so each summand is at most $4d(\beta f, \beta' f')^2$ and $\mathrm{KL}(P\|P') \le \frac{1}{8} \cdot 4d^2 \cdot T = \frac{T}{2} d(\beta f, \beta' f')^2$. For a common scale and $|f - f'| \le a$ pointwise, $\operatorname{osc}(\beta f - \beta f') \le 2\beta a$, which gives $\mathrm{KL}(P_f\|P_{f'}) \le (\beta^2/2)a^2 T$. □

B.3 THE SCALE LEMMA

Lemma B.5 (Scale lemma: distributional identity)**.** *Let (β_1, f_1) and (β_2, f_2) be two oracle instances whose answer laws coincide on every pair: $\sigma(\beta_1 \operatorname{gap}_{f_1}(x,y)) = \sigma(\beta_2 \operatorname{gap}_{f_2}(x,y))$ for all (x, y). Then every randomized adaptive algorithm, at any — possibly data-dependent — budget, has identical transcript and output laws under the two instances. Moreover, the answer laws coincide on every pair if and only if $\beta_1 f_1 - \beta_2 f_2$ is constant; in particular, (β, f) and $(c\beta, f/c)$ are indistinguishable by every algorithm, for every $c > 0$.*

Proof. Write $q_e(x, y) := \sigma\big(\beta_e \operatorname{gap}_{f_e}(x, y)\big)$ for the success probability of the answer bit on the pair (x, y) under instance $e \in \{1, 2\}$, and assume first the hypothesis of the main implication: $q_1(x, y) = q_2(x, y)$ for every pair. Fix any randomized adaptive algorithm A of any (possibly data-dependent) budget.

Step 1 (coupling). Realize A's internal randomness by one environment-independent seed $U \sim P_U$ (a randomized adaptive algorithm is a measurable deterministic map of U and the observed history). Run A against the two instances on the *same* U, and couple the responses: in round t, at the (identical, by Step 2) query (x, y), draw one shared $W_t \sim \mathrm{Unif}[0, 1]$ and set $B_t = \mathbf{1}\{W_t \le q_1(x, y)\}$ in run 1 and $B'_t = \mathbf{1}\{W_t \le q_2(x, y)\}$ in run 2. Since $q_1(x, y) = q_2(x, y)$ exactly, $B_t = B'_t$ surely; and marginally each response has the correct per-instance law $\mathrm{Bern}(q_e(x, y))$. (The W_t are i.i.d. and independent of U — a valid coupling realizing both instances' true dynamics.)

Step 2 (induction: identical histories). We claim $H_t^{(1)} = H_t^{(2)}$ surely for all t, where $H_t^{(e)}$ is the round-t history in run e. Base $t = 0$: both histories are empty. Induction step: assume $H_{t-1}^{(1)} = H_{t-1}^{(2)} =: h$. The algorithm computes its round-t query as $\phi_t(U, h)$ with the same policy map, the same U and the same h — identical in both runs. Given that common query (x, y), Step 1 gives $B_t = B'_t$ surely, so $H_t^{(1)} = (h, (x, y), B_t) = (h, (x, y), B'_t) = H_t^{(2)}$. By induction the entire transcript satisfies $\tau^{(1)} = \tau^{(2)}$ surely (for an unbounded horizon, by the countable intersection of the sure events $\{H_t^{(1)} = H_t^{(2)}\}$), and any adaptive stopping or budget rule — a function of the common history — fires at the same round in both runs.

Step 3 (equal laws). $\tau^{(1)} = \tau^{(2)}$ surely implies that the transcript laws under the two instances coincide. The output is A's terminal decision, a measurable function of (U, τ), hence identical in both runs, so its law coincides as well. As A and the budget were arbitrary, this proves the main implication.

Step 4 (the characterization). σ is a strictly increasing bijection $\mathbb{R} \to (0, 1)$, so

$$\begin{aligned}q_1 = q_2 \text{ on every pair} &\iff \beta_1 \operatorname{gap}_{f_1}(x,y) = \beta_2 \operatorname{gap}_{f_2}(x,y) \text{ for all } (x,y) \\ &\iff (\beta_1 f_1 - \beta_2 f_2)(x) = (\beta_1 f_1 - \beta_2 f_2)(y) \text{ for all } x, y \\ &\iff \beta_1 f_1 - \beta_2 f_2 \text{ is constant.}\end{aligned}$$

Sufficiency: when $\beta_1 f_1 - \beta_2 f_2$ is constant the answer laws coincide on every pair, and Steps 1–3 give identity of transcript and output laws for every algorithm. *Necessity:* if it is not constant, some pair (x_0, y_0) has $q_1(x_0, y_0) \ne q_2(x_0, y_0)$ (σ is injective); the one-query algorithm "query (x_0, y_0) once and output the bit" then has output laws $\mathrm{Bern}(q_1(x_0, y_0)) \ne \mathrm{Bern}(q_2(x_0, y_0))$. Hence identity of laws for all adaptive algorithms holds iff the per-pair channels coincide iff $\beta_1 f_1 - \beta_2 f_2$ is constant.

The scaling orbit. Finally, with $(\beta_1, f_1) = (\beta, f)$ and $(\beta_2, f_2) = (c\beta, f/c)$ for any $c > 0$, $\beta_1 f_1 - \beta_2 f_2 = \beta f - c\beta\,(f/c) = 0$, a constant; so the channels coincide and (β, f) and $(c\beta, f/c)$ are

indistinguishable by every algorithm, at any budget. This is the mechanism of the impossible direction of Theorem 5.1. □

B.4 A SPHERICAL CHAOS MGF BOUND

The last tool is more specialized: a bound on the moment generating function of the "chaos" $\langle a(v), a(v')\rangle$ of a mean-zero score map on the sphere, at independent uniform arguments. It powers the χ^2-mixture argument behind Theorem 4.2 and is used there at statement level.

Setup. Let v, v' be i.i.d. uniform on $\mathbb{S}^{n-1}$ with $n \geq 2$, and let $\beta > 0$. Let $T \in \mathbb{N}$ and let $a\colon \mathbb{S}^{n-1} \to \mathbb{R}^T$ be measurable with $\mathbb{E}_v[a(v)] = 0$; write $\Sigma := \mathbb{E}_v[a(v)\, a(v)^\top] \in \mathbb{R}^{T\times T}$ (positive semidefinite) with eigenvalues $\sigma_1 \geq \cdots \geq \sigma_T \geq 0$, so that $\|\Sigma\|_{\mathrm{op}} = \sigma_1$ and $\mathrm{tr}(\Sigma^2) = \sum_i \sigma_i^2$. (The range dimension is called T deliberately: in the application, $a(v)$ collects one score per query of a T-query algorithm. The scalar β^2 enters only as a positive multiplier, so applications may substitute any $p > 0$ for it.) For $g \sim \mathcal{N}(0, I_n)$ set

$$\nu_n := \big(\mathbb{E}\|g\|\big)^2 = \frac{2\,\Gamma\big(\frac{n+1}{2}\big)^2}{\Gamma\big(\frac{n}{2}\big)^2}, \qquad \varrho_n := \frac{n^2}{\nu_n^2}.$$

Then $n - \frac{1}{2} \leq \nu_n \leq n$, and ϱ_n is decreasing in n with $\varrho_n \in [1,\ 16/\pi^2]$; in particular $\varrho_2 = 16/\pi^2 \approx 1.6211$ (since $\nu_2 = \pi/2$) and $\varrho_n \to 1$ (Part A, Step 5 of the proof below).

We use two sharp facts about the uniform measure on $\mathbb{S}^{n-1}$, $n \geq 2$; $d(\cdot,\cdot)$ denotes the geodesic distance and $\mathrm{Lip}(\cdot)$ the geodesic Lipschitz constant.

Fact B.6 (Sphere Poincaré inequality and log-Sobolev sub-Gaussianity)**.** *Let $F\colon \mathbb{S}^{n-1} \to \mathbb{R}$ be geodesically Lipschitz, $n \geq 2$.*

(P) If $\mathbb{E}F = 0$, then $\mathrm{Var}(F) \leq \frac{1}{n-1}\,\mathbb{E}\|\nabla F\|^2 \leq \mathrm{Lip}(F)^2/(n-1)$.

(L) For every $u \in \mathbb{R}$, $\mathbb{E}\big[\exp\big(u(F - \mathbb{E}F)\big)\big] \leq \exp\big(u^2\,\mathrm{Lip}(F)^2/(2(n-1))\big)$.

Both rest on the spectral-gap / log-Sobolev constant $n-1$ of the unit sphere (Ledoux, 1999, Section 7.3); (L) follows from the log-Sobolev inequality (Mueller & Weissler, 1982) by the Herbst argument (Ledoux, 1999, Section 2.3), and on $\mathbb{S}^1$ ($n = 2$) the sharp log-Sobolev constant is classical (Rothaus, 1980; Weissler, 1980).

Lemma B.7 (Spherical chaos MGF bound)**.** *In the setup above:*

(A) (Bilinear scores; tight constant.) *Suppose a is linear, $a(v) = Mv$ with $M \in \mathbb{R}^{T\times n}$ — equivalently $\langle a(v), a(v')\rangle = v^\top B v'$ with $B := M^\top M$ positive semidefinite — so that $\Sigma = MM^\top/n$. If $\varrho_n\beta^4\|\Sigma\|_{\mathrm{op}}^2 < 1$, then*

$$\mathbb{E}_{v,v'}\big[\exp\big(\beta^2\langle a(v), a(v')\rangle\big)\big] \;\leq\; \prod_i \big(1 - \varrho_n\beta^4\sigma_i^2\big)^{-1/2} \;\leq\; \exp\left(\frac{(\varrho_n/2)\,\beta^4\,\mathrm{tr}(\Sigma^2)}{1 - \varrho_n\beta^4\|\Sigma\|_{\mathrm{op}}^2}\right),$$

and the leading constant $\varrho_n/2$ lies in $[1/2,\ 8/\pi^2]$, decreasing to the Gaussian value $1/2$.

(B) (General mean-zero Lipschitz scores.) *Suppose each coordinate a_t is L_t-Lipschitz for the geodesic distance; set the Lipschitz energy $L^2 := \sum_t L_t^2$ and*

$$s_1 := \beta^2\,\mathrm{tr}(\Sigma), \qquad s_2 := \frac{\beta^2 L^2}{n-1}.$$

Then $s_1 \leq s_2$, and if $s_2 < 1$,

$$\begin{aligned}\mathbb{E}_{v,v'}\big[\exp\big(\beta^2\langle a(v), a(v')\rangle\big)\big] &\leq (1 - s_2^2)^{-1/2}\exp\left(\frac{s_1 s_2/2}{1 - s_2^2}\right)\\ &\leq (1 - s_2^2)^{-1/2}\exp\left(\frac{s_2^2}{2(1 - s_2^2)}\right).\end{aligned}$$

Proof. We prove Part B first, then Part A.

Part B, Step 1 (trace and Poincaré). Coordinatewise, $\mathbb{E}[a_t] = 0$ and $\mathrm{Var}(a_t) = \Sigma_{tt}$, so $\mathrm{tr}(\Sigma) = \sum_t \mathrm{Var}(a_t) \le \sum_t L_t^2/(n-1) = L^2/(n-1)$ by Fact B.6(P). Multiplying by β^2 gives $s_1 \le s_2$.

Part B, Step 2 (condition on v'). By Tonelli (the integrand is positive),

$$\mathbb{E}_{v,v'}\big[\exp\big(\beta^2\langle a(v), a(v')\rangle\big)\big] = \mathbb{E}_{v'}\Big[\mathbb{E}_v\big[\exp\big(\langle c, a(v)\rangle\big)\big]\Big], \qquad c := \beta^2 a(v') \quad \text{(fixed given } v').$$

The map $v \mapsto \langle c, a(v)\rangle$ has mean $\langle c, \mathbb{E}_v a\rangle = 0$ and is $(\|c\|\, L)$-Lipschitz: coordinatewise $|a_t(v) - a_t(v'')| \le L_t\, d(v, v'')$ gives $\|a(v) - a(v'')\| \le L\, d(v, v'')$, whence $|\langle c, a(v)\rangle - \langle c, a(v'')\rangle| \le \|c\|\, L\, d(v, v'')$. Applying Fact B.6(L) with $u = 1$,

$$\mathbb{E}_v\big[\exp\big(\langle c, a(v)\rangle\big)\big] \le \exp\Big(\frac{\|c\|^2 L^2}{2(n-1)}\Big) = \exp\Big(\frac{\beta^4 \|a(v')\|^2 L^2}{2(n-1)}\Big).$$

Writing $h(v') := \|a(v')\|$ and $s := \beta^4 L^2/(2(n-1))$,

$$\mathbb{E}_{v,v'}\big[\exp\big(\beta^2\langle a(v), a(v')\rangle\big)\big] \le \mathbb{E}_{v'}\big[\exp\big(s\, h(v')^2\big)\big].$$

Part B, Step 3 (h is Lipschitz). $|h(v') - h(v'')| \le \|a(v') - a(v'')\| \le L\, d(v', v'')$, so h is L-Lipschitz. Let $m := \mathbb{E}_{v'}[h]$; then $m^2 \le \mathbb{E}[h^2] = \mathbb{E}\|a\|^2 = \mathrm{tr}(\Sigma)$.

Part B, Step 4 (Hubbard–Stratonovich linearization of h^2). For $Z \sim \mathcal{N}(0,1)$ and any $y \in \mathbb{R}$, $\mathbb{E}_Z[\exp(\sqrt{2s}\, yZ)] = \exp(sy^2)$; hence $\exp(sh^2) = \mathbb{E}_Z[\exp(\sqrt{2s}\, hZ)]$ and, by Tonelli again,

$$\mathbb{E}_{v'}\big[\exp(sh^2)\big] = \mathbb{E}_Z\Big[\mathbb{E}_{v'}\big[\exp\big((\sqrt{2s}\, Z)\, h(v')\big)\big]\Big].$$

For each fixed Z put $u := \sqrt{2s}\, Z$ (any real number); applying Fact B.6(L) to h — keeping the mean term, since h is not centered —

$$\mathbb{E}_{v'}\big[\exp(uh)\big] \le \exp\Big(um + \frac{u^2 L^2}{2(n-1)}\Big) = \exp\Big(\sqrt{2s}\, Zm + \frac{sL^2 Z^2}{n-1}\Big).$$

Part B, Step 5 (Gaussian integral). Set $a_0 := \sqrt{2s}\, m$ and $b_0 := sL^2/(n-1) = \frac{1}{2}s_2^2$; note $b_0 < 1/2$ iff $s_2 < 1$. The Gaussian-quadratic identity $\mathbb{E}_Z[\exp(a_0 Z + b_0 Z^2)] = (1 - 2b_0)^{-1/2} \exp\big(a_0^2/(2(1 - 2b_0))\big)$, valid for $b_0 < 1/2$, yields

$$\mathbb{E}_{v'}\big[\exp(sh^2)\big] \le (1 - s_2^2)^{-1/2} \exp\Big(\frac{sm^2}{1 - s_2^2}\Big).$$

Using $m^2 \le \mathrm{tr}(\Sigma)$ (Step 3) and $s\,\mathrm{tr}(\Sigma) = \beta^4 L^2\, \mathrm{tr}(\Sigma)/(2(n-1)) = s_1 s_2/2$,

$$\mathbb{E}_{v,v'}\big[\exp\big(\beta^2\langle a(v), a(v')\rangle\big)\big] \le (1 - s_2^2)^{-1/2} \exp\Big(\frac{s_1 s_2/2}{1 - s_2^2}\Big),$$

and $s_1 \le s_2$ (Step 1) gives the final form. This proves Part B.

Part A, Step 1 (exact Gaussian bilinear chaos). Let g, g' be i.i.d. $\mathcal{N}(0, I_n)$ and let B be positive semidefinite with eigenvalues $\mu_1 \ge \cdots \ge \mu_n \ge 0$. Diagonalizing B, $g^\top B g' = \sum_i \mu_i G_i G_i'$ with G_i, G_i' i.i.d. $\mathcal{N}(0,1)$. Two-stage Gaussian integration gives, for $\lambda|\mu_i| < 1$,

$$\mathbb{E}\big[\exp(\lambda \mu_i G_i G_i')\big] = \mathbb{E}_{G_i}\big[\exp\big(\lambda^2 \mu_i^2 G_i^2/2\big)\big] = \big(1 - \lambda^2 \mu_i^2\big)^{-1/2},$$

hence, by independence across i,

$$\mathbb{E}_{g,g'}\big[\exp\big(\lambda\, g^\top B g'\big)\big] = \prod_i \big(1 - \lambda^2 \mu_i^2\big)^{-1/2}, \qquad \text{valid for } \lambda \mu_{\max} < 1.$$

Part A, Step 2 (radius factorization). Polar-decompose $g = R\theta$ with $\theta := g/\|g\|$ uniform on $\mathbb{S}^{n-1}$ and $R := \|g\|$ *independent* of θ. Then $g^\top B g' = RR'\, \theta^\top B \theta'$, so with $I(p) := \mathbb{E}_{\theta,\theta'}\big[\exp\big(p\, \theta^\top B \theta'\big)\big]$,

$$\prod_i \big(1 - \lambda^2 \mu_i^2\big)^{-1/2} = \mathbb{E}_{R,R'}\big[I(\lambda RR')\big].$$

Part A, Step 3 (Jensen de-conditioning — the sphere-from-Gaussian step). The map $p \mapsto I(p)$ is convex: $I''(p) = \mathbb{E}\big[(\theta^\top B\theta')^2\, e^{p\,\theta^\top B\theta'}\big] \geq 0$, differentiation under the integral being licensed by $|\theta^\top B\theta'| \leq \mu_{\max}$ on the compact sphere (which also makes I finite for every real p). R, R' are i.i.d. with $\mathbb{E}[RR'] = (\mathbb{E}R)^2 = \nu_n$. Jensen's inequality applied to the random variable $\lambda RR'$ (the right-hand side is finite for $\lambda\mu_{\max} < 1$) gives

$$I(\lambda\nu_n) \leq \mathbb{E}_{R,R'}\big[I(\lambda RR')\big] = \prod_i \big(1 - \lambda^2\mu_i^2\big)^{-1/2}.$$

Choose $\lambda := \beta^2/\nu_n$, so that $\lambda\nu_n = \beta^2$. The validity condition $\lambda\mu_{\max} < 1$ reads $\beta^2\|\Sigma\|_{\mathrm{op}} < \nu_n/n$, i.e. $\varrho_n\beta^4\|\Sigma\|_{\mathrm{op}}^2 < 1$: indeed $\mu_{\max} = n\|\Sigma\|_{\mathrm{op}}$, since the nonzero spectra of $B = M^\top M$ and of $n\Sigma = MM^\top$ coincide, so $\mu_i = n\sigma_i$ up to zero eigenvalues. Hence

$$\mathbb{E}_{v,v'}\big[\exp\big(\beta^2\, v^\top B v'\big)\big] = I(\beta^2) \leq \prod_i \big(1 - \beta^4\mu_i^2/\nu_n^2\big)^{-1/2}.$$

Part A, Step 4 (rewrite via Σ and exponentiate). $\Sigma = B/n$ has eigenvalues $\sigma_i = \mu_i/n$, so $\beta^4\mu_i^2/\nu_n^2 = (n^2/\nu_n^2)\,\beta^4\sigma_i^2 = \varrho_n\beta^4\sigma_i^2$, which proves the first inequality of Part A. For the second, apply $-\log(1-x) \leq x/(1-x_{\max})$, valid for $0 \leq x \leq x_{\max} < 1$, with $x_i := \varrho_n\beta^4\sigma_i^2$ and $x_{\max} := \varrho_n\beta^4\|\Sigma\|_{\mathrm{op}}^2$:

$$\prod_i \big(1 - \varrho_n\beta^4\sigma_i^2\big)^{-1/2} = \exp\Big(\tfrac{1}{2}\textstyle\sum_i - \log(1 - x_i)\Big) \leq \exp\left(\frac{(\varrho_n\beta^4/2)\ \mathrm{tr}(\Sigma^2)}{1 - \varrho_n\beta^4\|\Sigma\|_{\mathrm{op}}^2}\right).$$

Part A, Step 5 (the constant ϱ_n). By Jensen, $\nu_n = (\mathbb{E}R)^2 \leq \mathbb{E}R^2 = n$; moreover $\mathrm{Var}(R) = n - \nu_n \leq 1/2$, so $\nu_n \geq n - 1/2$; and ν_n/n is increasing in n, so $\varrho_n = (n/\nu_n)^2$ is decreasing, with maximum $\varrho_2 = 16/\pi^2$ (from $\nu_2 = \pi/2$) and $\varrho_n \to 1$. Hence $\varrho_n \in [1,\, 16/\pi^2]$ and $\varrho_n/2 \in [1/2,\, 8/\pi^2]$. These Γ-function facts are classical, via the log-convexity of Γ (Kershaw-type bounds); they are used only to package ϱ_n into $[1, 16/\pi^2]$ — the bound itself carries ϱ_n exactly. This proves Part A.

□

Remark B.8 (The naive $\mathrm{tr}(\Sigma^2)$ form fails for general scores). For bilinear scores, Part A of Lemma B.7 is an exact spherical analogue of the Gaussian chaos bound, with the whole radial correction carried by the single explicit factor $\varrho_n \in [1, 16/\pi^2]$. For a *general* mean-zero Lipschitz score map, the clean form $\mathbb{E}[\exp(\beta^2\langle a, a'\rangle)] \leq \exp\big(C\beta^4\,\mathrm{tr}(\Sigma^2)\big)$ is false: a single spike coordinate can drive $\mathrm{tr}(\Sigma^2) \to 0$ while the moment generating function tends to $+\infty$ — it is the Lipschitz energy L^2, not the covariance, that controls the upper tail. (Such witnesses have $s_2 \geq 1$, i.e. they sit outside Part B's validity window, consistently with the lemma.) Part B is the sharp general substitute: its exponent replaces $\mathrm{tr}(\Sigma^2)$ by $L^2\,\mathrm{tr}(\Sigma)/(n-1) \geq \mathrm{tr}(\Sigma)^2 \geq \mathrm{tr}(\Sigma^2)$, i.e. $s_1 s_2/2 = \beta^4 L^2\,\mathrm{tr}(\Sigma)/(2(n-1)) \geq \beta^4\,\mathrm{tr}(\Sigma^2)/2$.

Remark B.9 (Operating-point arithmetic). For orientation, here is the arithmetic of Part B at the operating point of the nonadaptive verification lower bound, with *linearized* scores. On that ensemble the scores satisfy $\mathrm{tr}(\Sigma) \leq 4\varepsilon^2 t^2 T/(n-1)$ and $L^2 \leq 4\varepsilon^2 t^2 T$, so with $\kappa := 4\beta^2\varepsilon^2 t^2 T/(n-1)$ one has $s_1, s_2 \leq \kappa$; both Part-B bounds are increasing in (s_1, s_2) on the validity window, and for $\kappa < 1$,

$$(1-\kappa^2)^{-1/2}\exp\Big(\frac{\kappa^2}{2(1-\kappa^2)}\Big) - 1 \;<\; \frac{4}{9} \qquad \text{for } \kappa \leq 0.5342,$$

which corresponds to the constant $0.5342/16 \geq 1/30$ in the $T \geq c\,(n-1)L_1^2/(\beta^2\varepsilon^4)$ format. The rigorous constants for Theorem 4.2 are established by the exact-sigmoid (non-linearized) argument accompanying its proof, whose $\varrho \to 0$ limit recovers exactly this $1/30$ arithmetic; this remark supports no step of that proof.

C AMPLITUDES: THE CHANNEL, THE LAW, AND THE NEW BOUNDS

This appendix proves the results that organize the paper: Proposition 2.5 (what a transcript sees), Theorem 2.7 (the amplitude law), Proposition 3.1 (the loudness of a gradient), Theorem 5.2 (the edge

of identifiability) and Theorem 3.3 (the exponent law). The first two package the toolkit lemmas of Appendix B; the last three are new. Throughout, $\operatorname{osc}(h) = \sup h - \inf h \in [0, \infty]$, the logit distance is that of Definition 2.4, and $\tau(x) \coloneqq \langle e_1, x - \bar{x} \rangle$ denotes a ridge coordinate. For a scalar profile u, the ridge function $u \circ \tau$ has gradient $u'(\tau(x))e_1$; its gradient is Lipschitz with constant $\sup |u''|$, and more generally its p-th derivative is Lipschitz with constant $\sup |u^{(p+1)}|$ (Equation (G.2)).

C.1 The channel and the amplitude law

Proposition 2.5 (What a comparison transcript sees)**.** *Let $P_I, P_{I'}$ be the transcript laws of a randomized adaptive algorithm run against instances $I = (\beta, f)$ and $I' = (\beta', f')$.*

(a) (Identity.) *If $d(I, I') = 0$, then $P_I = P_{I'}$ and the output laws coincide, for every algorithm and every, possibly data-dependent, budget. If $d(I, I') > 0$, some single query has different answer laws under I and I'.*

(b) (Contraction.) *If the algorithm makes T queries, then*

$$\mathrm{KL}(P_I \,\|\, P_{I'}) \;\le\; \frac{T}{2}\, d(I, I')^2.$$

The constant $1/2$ cannot be improved, even for $T = 1$.

Proof. *(a)* If $d(I, I') = 0$, then $\beta f - \beta' f'$ is constant, and Lemma B.5, applied with $(\beta_1, f_1) = (\beta, f)$ and $(\beta_2, f_2) = (\beta', f')$, gives identical answer laws on every pair and hence identical transcript and output laws for every algorithm at every budget. If $d(I, I') > 0$, then $\beta f - \beta' f'$ is not constant, and the necessity part of Lemma B.5 exhibits a pair whose two answer laws differ.

(b) For every pair (x, y), $|\beta \operatorname{gap}_f(x, y) - \beta' \operatorname{gap}_{f'}(x, y)| = |(\beta f - \beta' f')(x) - (\beta f - \beta' f')(y)| \le \operatorname{osc}(\beta f - \beta' f') = 2d(I, I')$. The first inequality of Lemma B.4 therefore gives $\mathrm{KL}(P_I \| P_{I'}) \le \frac{1}{8}\, T\, (2d(I, I'))^2 = \frac{T}{2}\, d(I, I')^2$.

For the constant, fix $d > 0$ and let $h(x) \coloneqq d \max\{-1, \min\{1, x_1\}\}$, so that $\operatorname{osc}(h) = 2d$ and $h(\pm e_1) = \pm d$. The instances $I = (1, 0)$ and $I' = (1, h)$ have $d(I, I') = d$. The one-query algorithm that queries $(e_1, -e_1)$ has transcript laws $\mathrm{Bern}(\sigma(0))$ and $\mathrm{Bern}(\sigma(2d))$. By the sharpness clause of Lemma B.1 (with $\beta = 1$, $\theta_1 = 0$ and $\theta_2 = 2d$), the ratio of $\mathrm{KL}(\mathrm{Bern}(\sigma(0)) \| \mathrm{Bern}(\sigma(2d)))$ to $\frac{1}{8}(2d)^2 = \frac{1}{2}d^2$ tends to 1 as $d \to 0$. □

Theorem 2.7 (The amplitude law)**.** *Let $\mathcal{P}$ be a promise problem.*

(a) (Silence.) *If $\mathcal{P}$ is silent, no algorithm, at any budget, even unbounded and data-dependent, answers correctly with probability greater than $1/2$ on every instance.*

(b) (Two points.) *Every algorithm that solves $\mathcal{P}$ with confidence $2/3$ makes $T \ge 2\ln(9/8)/a(\mathcal{P})^2$ queries; with confidence $1-\delta$ and $\delta \le 1/16$, it makes $T \ge \ln(1/\delta)/a(\mathcal{P})^2$ queries.*

(c) (Many points.) *Let $I_1, \dots, I_M$, $M \ge 3$, be instances with $d(I_j, I_k) \le a$ for all j, k, and let Ψ map outputs to $\{1, \dots, M\}$ so that every output that is correct for I_j is mapped to j. Every algorithm whose output is correct with probability at least $2/3$ on each I_j makes $T \ge \left(\frac{4}{3}\ln M - 2\ln 2\right)/a^2$ queries.*

In (b) and (c), T is a fixed worst-case budget, and $1/0 = \infty$.

Proof. We include the algorithm's internal randomness as an initial entry of the transcript. It has the same law under every instance, so its term in the chain rule of Lemma B.4 vanishes and the bound of Proposition 2.5(b) is unchanged, while every output becomes a measurable function of the extended transcript.

(a) Let $I \in \mathcal{A}$ and $I' \in \mathcal{R}$ with $d(I, I') = 0$. By Proposition 2.5(a), every algorithm outputs accept with the same probability q under I and I', at any budget. Correctness with probability greater than $1/2$ on both would require $q > 1/2$ and $1 - q > 1/2$.

(b) Let an algorithm solve $\mathcal{P}$ with confidence $1 - \delta$, where $\delta < 1/2$, using T queries, and let $I \in \mathcal{A}$ and $I' \in \mathcal{R}$. The accept event has probability at least $1 - \delta$ under I and at most δ under I', so

$\mathrm{TV}(P_I, P_{I'}) \geq 1 - 2\delta$. The Bretagnolle–Huber inequality $\mathrm{TV} \leq \sqrt{1 - e^{-\mathrm{KL}}}$ (Tsybakov, 2009) gives $e^{-\mathrm{KL}(P_I \| P_{I'})} \leq 1 - (1-2\delta)^2 = 4\delta(1-\delta)$. For $\delta = 1/3$ this is $\mathrm{KL}(P_I \| P_{I'}) \geq \ln(9/8)$. For $\delta \leq 1/16$ it gives $\mathrm{KL}(P_I \| P_{I'}) \geq \ln(1/(4\delta)) \geq \frac{1}{2}\ln(1/\delta)$, since $1/(4\delta) \geq \delta^{-1/2}$ exactly when $\delta \leq 1/16$. With Proposition 2.5(b),

$$T \;\geq\; \frac{2\ln(9/8)}{d(I,I')^2} \qquad \text{respectively} \qquad T \;\geq\; \frac{\ln(1/\delta)}{d(I,I')^2}.$$

Both bounds hold for every such pair; letting $d(I, I')$ decrease to $a(\mathcal{P})$ proves (b). If $a(\mathcal{P}) = 0$, no finite budget suffices.

(c) Let V be uniform on $\{1, \dots, M\}$, let Z be the extended transcript, and let P_j be its law under I_j. By convexity of KL in its second argument,

$$I(V;Z) = \frac{1}{M}\sum_j \mathrm{KL}\Big(P_j \,\Big\|\, \frac{1}{M}\sum_k P_k\Big) \leq \frac{1}{M^2}\sum_{j,k}\mathrm{KL}(P_j \| P_k) \leq \frac{Ta^2}{2},$$

the last step by Proposition 2.5(b). The index $\Psi(\text{output})$ is a function of Z and equals V with probability at least $2/3$. Fano's inequality (Tsybakov, 2009, Chapter 2) gives $1/3 \geq \mathbb{P}(\Psi \neq V) \geq 1 - (I(V;Z) + \ln 2)/\ln M$, so $I(V;Z) \geq \frac{2}{3}\ln M - \ln 2$ and $T \geq 2(\frac{2}{3}\ln M - \ln 2)/a^2 = (\frac{4}{3}\ln M - 2\ln 2)/a^2$. □

Remark C.1 (How the lower bounds of the paper instantiate the law). The search lower bounds are part (c) with the shifted-cone packing, whose members are uniformly within $4\beta\varepsilon t$ of each other in logit distance (Appendix D). The two-point verification bounds are part (b) with the exact amplitude of Proposition 3.1. The unknown-scale impossibility is part (a), and the edge rate of Theorem 5.2 is part (b). The value-margin lower bound is part (b) with amplitude exactly βm_0 (Appendix F.3). Two bounds go beyond pairwise distances: the universal floor of Theorem 4.1(iii), which counts the 2^T possible answer strings, and the nonadaptive bound of Theorem 4.2, which compares a mixture of instances with a single instance through a χ^2 calculation.

C.2 The loudness of a gradient

Proposition 3.1 (The loudness of a gradient). *Let $c' > 1$, $\varepsilon, L_1 > 0$ and $0 < \beta_{\min} \leq \beta_{\max} < \infty$, with $\rho = \beta_{\max}/\beta_{\min}$. Let $\mathcal{P}_\rho$ be ε-vs-$c'\varepsilon$ gradient-norm verification at $\bar{x}$ over the instances (β, f) with $\beta \in [\beta_{\min}, \beta_{\max}]$ and $f \in \mathcal{F}(L_1)$, and let $\mathcal{P}_\rho^\Delta$ be the same problem over $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \leq \Delta\}$, where Assumption 2.3 holds. Then*

$$a(\mathcal{P}_\rho) \;=\; a(\mathcal{P}_\rho^\Delta) \;=\; \frac{\beta_{\min}\,\varepsilon^2\,(c'-\rho)_+^2}{2\,(1+\rho)\,L_1}.$$

Both problems are silent when $\rho \geq c'$; when $\rho < c'$, a pair of ridge functions attains the infimum. In particular, at a known scale β the amplitude of ε-vs-$c'\varepsilon$ verification over $\mathcal{F}(L_1)$ is $\beta(c'-1)^2\varepsilon^2/(4L_1)$.

Proof. Write $\mathcal{A}_\rho, \mathcal{R}_\rho$ for the accept and reject instances of $\mathcal{P}_\rho$, and set

$$\gamma := \beta_{\min}(c'-\rho)\varepsilon, \qquad M_0 := (\beta_{\max} + \beta_{\min})L_1, \qquad w := \gamma/M_0.$$

Since $\mathcal{P}_\rho^\Delta$ has fewer instances, $a(\mathcal{P}_\rho) \leq a(\mathcal{P}_\rho^\Delta)$. We show that every pair in $\mathcal{A}_\rho \times \mathcal{R}_\rho$ is at least as far apart as the claimed value (Step 1), and that a pair in the smaller problem $\mathcal{P}_\rho^\Delta$ attains it (Steps 2 and 3).

Step 1 (every pair is at least this far apart). For $\rho \geq c'$ there is nothing to prove. Let $\rho < c'$, let $(\beta, f) \in \mathcal{A}_\rho$ and $(\beta', f') \in \mathcal{R}_\rho$, and put $h := \beta' f' - \beta f$. Then

$$G := \|\nabla h(\bar{x})\| \geq \beta'\|\nabla f'(\bar{x})\| - \beta\|\nabla f(\bar{x})\| \geq (\beta' c' - \beta)\varepsilon \geq (\beta_{\min} c' - \beta_{\max})\varepsilon = \gamma > 0,$$

and ∇h is Lipschitz with constant $M := (\beta + \beta')L_1$. Let $e := \nabla h(\bar{x})/G$ and $q(s) := h(\bar{x} + se)$ for $s \in \mathbb{R}$. Then $q'(0) = G$ and q' is M-Lipschitz, because $|q'(s) - q'(s')| = |\langle \nabla h(\bar{x} + se) - \nabla h(\bar{x} + s'e), e\rangle| \leq M|s - s'|$. With $s_1 := G/M$ we have $q'(s) \geq G - M|s| \geq 0$ for $|s| \leq s_1$, so

$$\mathrm{osc}(h) \;\geq\; q(s_1) - q(-s_1) \;=\; \int_{-s_1}^{s_1} q'(s)\,ds \;\geq\; \int_{-s_1}^{s_1} (G - M|s|)\,ds \;=\; \frac{G^2}{M} \;\geq\; \frac{(\beta' c' - \beta)^2\varepsilon^2}{(\beta + \beta')L_1}. \tag{C.1}$$

The function $F(\beta, \beta') \coloneqq (\beta' c' - \beta)^2/(\beta + \beta')$ is positive on the box $[\beta_{\min}, \beta_{\max}]^2$ because $\beta' c' \geq \beta_{\min} c' > \beta_{\max} \geq \beta$. On the box,

$$\partial_\beta F = -\frac{(\beta' c' - \beta)\big(2(\beta + \beta') + (\beta' c' - \beta)\big)}{(\beta + \beta')^2} < 0,$$
$$\partial_{\beta'} F = \frac{(\beta' c' - \beta)\big(2c'\beta + c'\beta' + \beta\big)}{(\beta + \beta')^2} > 0,$$

so $F \geq F(\beta_{\max}, \beta_{\min}) = \beta_{\min}(c' - \rho)^2/(1 + \rho)$. Hence $d((\beta, f), (\beta', f')) = \frac{1}{2}\operatorname{osc}(h) \geq \beta_{\min}\varepsilon^2(c' - \rho)^2/(2(1 + \rho)L_1)$.

Step 2 (silence for $\rho \geq c'$). Let $f_0 \coloneqq u_0 \circ \tau$ with $u_0(s) = c'\varepsilon s + \frac{L_1}{2}s^2$. As in Steps I.1–I.5 of Appendix F.1, $f_0 \in \mathcal{F}(L_1)$, $\|\nabla f_0(\bar{x})\| = c'\varepsilon$ and $f_0(\bar{x}) - \inf f_0 = (c'\varepsilon)^2/(2L_1) \leq \Delta$ by Assumption 2.3; the function f_0/c' has the same properties with gradient norm ε. The scales $\beta_{\min}$ and $c'\beta_{\min}$ both lie in $[\beta_{\min}, \beta_{\max}]$ because $c'\beta_{\min} \leq \rho\beta_{\min} = \beta_{\max}$. The reject instance $(\beta_{\min}, f_0)$ and the accept instance $(c'\beta_{\min}, f_0/c')$ have the same logit potential, so their logit distance is 0, and both problems are silent.

Step 3 (an attaining pair for $\rho < c'$). Define profiles u and v on $\mathbb{R}$ by $u(0) = v(0) = 0$, $u'(0) = \varepsilon$, $v'(0) = c'\varepsilon$, and the piecewise constant second derivatives

$$u''(s) = \begin{cases} L_1 \operatorname{sign}(s), & |s| < w, \\ L_1/\rho, & s \leq -w, \\ 0, & s \geq w, \end{cases} \qquad v''(s) = \begin{cases} -L_1 \operatorname{sign}(s), & |s| < w, \\ L_1, & s \leq -w, \\ 0, & s \geq w, \end{cases}$$

and consider the accept instance $I = (\beta_{\max}, u \circ \tau)$ and the reject instance $I' = (\beta_{\min}, v \circ \tau)$.

Class and promise. Both profiles are C^1 with $\sup|u''| \leq L_1$ and $\sup|v''| \leq L_1$, so both ridge functions lie in $\mathcal{F}(L_1)$. Their gradient norms at $\bar{x}$ are $|u'(0)| = \varepsilon$ and $|v'(0)| = c'\varepsilon$, and both scales lie in $[\beta_{\min}, \beta_{\max}]$.

The logit potentials differ by a single step. Let $q \coloneqq \beta_{\min} v - \beta_{\max} u$. Then $q(0) = 0$, $q'(0) = \beta_{\min} c'\varepsilon - \beta_{\max}\varepsilon = \gamma$, and $q'' = \beta_{\min} v'' - \beta_{\max} u''$ equals $-M_0 \operatorname{sign}(s)$ on $(-w, w)$, equals $\beta_{\min} L_1 - \beta_{\max} L_1/\rho = 0$ on $(-\infty, -w]$, and equals 0 on $[w, \infty)$. Hence $q'(s) = \gamma - M_0|s| \geq 0$ on $[-w, w]$, $q'(\pm w) = 0$ and $q' = 0$ outside $[-w, w]$. Thus q is nondecreasing and constant outside $[-w, w]$, and

$$\operatorname{osc}(q) = \int_{-w}^{w} (\gamma - M_0|s|)\, ds = \frac{\gamma^2}{M_0} = \frac{\beta_{\min}(c' - \rho)^2\varepsilon^2}{(1 + \rho)L_1}.$$

Since the logit potentials differ by $q \circ \tau$, $d(I, I') = \frac{1}{2}\operatorname{osc}(q)$ equals the bound of Step 1.

The gap budget. For $s \leq 0$, $v'' = L_1$ on $(-w, 0)$ and on $(-\infty, -w]$, so $v(s) = c'\varepsilon s + \frac{L_1}{2}s^2$. For $s \geq 0$, $v'(s) = c'\varepsilon - L_1 \min\{s, w\} \geq c'\varepsilon - L_1 w > 0$, because $L_1 w = (c' - \rho)\varepsilon/(1 + \rho) < c'\varepsilon$; so $v \geq v(0) = 0$ there. Hence $\inf v = \min_{s \leq 0}(c'\varepsilon s + \frac{L_1}{2}s^2) = -(c'\varepsilon)^2/(2L_1)$ and $v(0) - \inf v = (c'\varepsilon)^2/(2L_1) \leq \Delta$ by Assumption 2.3. For u, $u' \geq \varepsilon > 0$ on $[0, \infty)$, so $u \geq 0$ there. For $s \leq 0$, $q(s) \leq q(0) = 0$ gives $u(s) = (\beta_{\min} v(s) - q(s))/\beta_{\max} \geq \beta_{\min} v(s)/\beta_{\max} \geq -(c'\varepsilon)^2/(2\rho L_1)$. Therefore $u(0) - \inf u \leq (c'\varepsilon)^2/(2\rho L_1) \leq \Delta$.

So I and I' are instances of $\mathcal{P}_\rho^\Delta$ at logit distance equal to the lower bound of Step 1, and $a(\mathcal{P}_\rho) = a(\mathcal{P}_\rho^\Delta) = \beta_{\min}\varepsilon^2(c' - \rho)^2/(2(1 + \rho)L_1)$. The known-scale statement is the case $\beta_{\min} = \beta_{\max} = \beta$, where $\rho = 1$; there the construction does not use the budget. □

Remark C.2 (Hadamard's inequality and the extremal step). Inequality (C.1) is the sharp inequality $\|q'\|_\infty^2 \leq 2\|q\|_\infty\|q''\|_\infty$ of Landau and Hadamard on the whole line (Landau, 1914; Hadamard, 1914), applied to $q - \frac{1}{2}(\sup q + \inf q)$ and localized at one point. The extremal step $q' = (\gamma - M_0|s|)_+$ is a smoothed ramp of height γ^2/M_0. It is the smallest value change that carries a gradient change γ when the difference of two L_1-smooth potentials has curvature at most M_0. The two-point constant it yields is $32\ln(9/8)/(c' - 1)^4$ in units of $L_1^2/(\beta^2\varepsilon^4)$.

C.3 THE EDGE OF IDENTIFIABILITY

Theorem 5.2 (The edge of identifiability). *In the setting of Theorem 5.1, let $1 \le \rho < c'$. Every randomized adaptive algorithm that is correct with probability at least $2/3$ on every instance makes*

$$T \;\ge\; 8\ln(9/8)\,\frac{(1+\rho)^2\,L_1^2}{\beta_{\min}^2\,(c'-\rho)^4\,\varepsilon^4}$$

queries, and at confidence $1-\delta$ with $\delta \le 1/16$ it makes $T \ge 4(1+\rho)^2 L_1^2 \ln(1/\delta)/(\beta_{\min}^2(c'-\rho)^4\varepsilon^4)$ queries. Consequently, let $N^\star(\rho)$ be the least fixed worst-case budget that achieves confidence $1-\delta$, $\delta \le 1/16$, in the unknown-scale problem, and let $\beta^\dagger$ be as in Proposition F.3. Whenever $\beta_{\min} \le \beta^\dagger$,

$$\frac{4\,(1+\rho)^2 L_1^2 \ln(1/\delta)}{\beta_{\min}^2\,\varepsilon^4\,(c'-\rho)^4} \;\le\; N^\star(\rho) \;\le\; \frac{3744\,n^3\rho^2 L_1^2 \log(2n/\delta)}{\beta_{\min}^2\,\varepsilon^4\,(c'-\rho)^4} + n.$$

At fixed $(n, L_1, \Delta, \varepsilon, c', \delta, \beta_{\min})$ the complexity therefore grows exactly as $(c'-\rho)^{-4}$ as $\rho \uparrow c'$, and at fixed $\rho < c'$ exactly as $\beta_{\min}^{-2}$ as $\beta_{\min} \to 0$.

Proof. By Proposition 3.1, the problem of Theorem 5.1 has amplitude

$$a_\rho = \frac{\beta_{\min}\varepsilon^2(c'-\rho)^2}{2(1+\rho)L_1} > 0, \qquad \frac{1}{a_\rho^2} = \frac{4(1+\rho)^2 L_1^2}{\beta_{\min}^2\varepsilon^4(c'-\rho)^4}.$$

Theorem 2.7(b) gives $T \ge 2\ln(9/8)/a_\rho^2$ and, for $\delta \le 1/16$, $T \ge \ln(1/\delta)/a_\rho^2$, which are the two displayed lower bounds. For the two-sided estimate, the left inequality is the second lower bound. For the right one, Proposition F.3 provides a tester with failure probability at most δ and cost at most $1872\,n^2\rho^2 L_1\Delta\log(2n/\delta)\,\varepsilon^{-2}(c'-\rho)^{-2}\max\{1, (\beta^\dagger/\beta_{\min})^2\} + n$. When $\beta_{\min} \le \beta^\dagger$ the maximum equals $(\beta^\dagger/\beta_{\min})^2 = 2nL_1/(\Delta\varepsilon^2(c'-\rho)^2\beta_{\min}^2)$, and the product is $3744\,n^3\rho^2L_1^2\log(2n/\delta)/(\beta_{\min}^2\varepsilon^4(c'-\rho)^4)$.

For the asymptotic statements, fix $(n, L_1, \Delta, \varepsilon, c', \delta, \beta_{\min})$. Since $\beta^\dagger = \sqrt{2nL_1/\Delta}/(\varepsilon(c'-\rho)) \to \infty$ as $\rho \uparrow c'$, the condition $\beta_{\min} \le \beta^\dagger$ holds for all ρ close to c'. There the two sides of the display are $(c'-\rho)^{-4}$ times factors that stay between positive constants, because $(1+\rho)^2 \in [4, (1+c')^2]$ and $\rho^2 \in [1, c'^2]$, while the additive n is of lower order. At fixed $\rho < c'$, the condition holds for all small $\beta_{\min}$, and both sides are $\beta_{\min}^{-2}$ times constants. □

Remark C.3 (What remains open at the edge). The ratio of the two sides of the display is at most $936\,n^3\log(2n/\delta)/\ln(1/\delta)$ plus n divided by the left side, so the edge rate is settled up to a factor polynomial in n in the regime $\beta_{\min} \le \beta^\dagger$. In the saturated regime $\beta_{\min} \ge \beta^\dagger$ the tester of Proposition F.3 costs at most $N_{\mathrm{sat}} \propto n^2\rho^2 L_1\Delta\log(2n/\delta)/(\varepsilon^2(c'-\rho)^2)$, while the lower bound decays like $\beta_{\min}^{-2}$. At $\beta_{\min} = \beta^\dagger$ the lower bound equals $2(1+\rho)^2 L_1\Delta\ln(1/\delta)/(n\varepsilon^2(c'-\rho)^2)$, so there both bounds are proportional to $L_1\Delta/(\varepsilon^2(c'-\rho)^2)$, up to a factor $O(n^3\log(n/\delta))$. Whether the cost stays of order $(c'-\rho)^{-2}$ as $\beta_{\min} \to \infty$ is open.

C.4 THE EXPONENT LAW

The cap of the search construction generalizes to every smoothness order. For $p \ge 1$ let

$$Q_p(s) := \sum_{k=0}^{p} \binom{1/2}{k} (s-1)^k$$

be the Taylor polynomial of degree p of $\sqrt{s}$ at $s = 1$, and define $\chi_p(r) := Q_p(r^2)$ for $r \in [0,1]$ and

$$\Phi_p(y) := \chi_p(\|y\|)\ \ (\|y\| \le 1), \qquad \Phi_p(y) := \|y\|\ \ (\|y\| \ge 1).$$

For $p = 1$ this is $\chi_1(r) = (1+r^2)/2$, the quadratic cap of Appendix D in the units $r_0 = 1$; for $p = 2$ it is $\chi_2(r) = \frac{3}{8} + \frac{3}{4}r^2 - \frac{1}{8}r^4$.

Lemma C.4 (The cap of order p). *Let $n \ge 1$ and $p \ge 1$.*

(i) $\Phi_p \in C^p(\mathbb{R}^n)$*, and* $D^p\Phi_p$ *is Lipschitz in operator norm with constant at most* $K_p := \sup_{m\ge1}\sup_{y\in\mathbb{R}^m,\,\|y\|\neq1}\|D^{p+1}\Phi_p(y)\|$*, where* Φ_p *is read on* $\mathbb{R}^m$*. This supremum is finite, depends only on* p*, and satisfies* $K_p \ge (2p-1)!!$*; moreover* $K_1 = 1$*.*

(ii) $0 \le \chi_p' \le 1$ *on* $[0,1]$*. Consequently* Φ_p *is* 1*-Lipschitz, and* $\|\nabla\Phi_p(y)\| = 1$ *whenever* $\|y\| \ge 1$*.*

(iii) $\min\Phi_p = \chi_p(0) = 4^{-p}\binom{2p}{p} \in (0,1)$*.*

(iv) For $p \ge 2$*,* $H_p := \sup_{m\ge1}\sup_{y\in\mathbb{R}^m}\|D^2\Phi_p(y)\|$ *is finite and depends only on* p*; set* $H_1 := 1$*.*

Proof. (i) Smoothness and matching. Inside the ball, $\Phi_p(y) = Q_p(\|y\|^2)$ is a polynomial in y. Outside, $\|y\|$ is C^∞ on $\mathbb{R}^n\setminus\{0\}$. Since $Q_p - \sqrt{\cdot}$ vanishes to order $p+1$ at $s=1$, the composition $\chi_p(r) - r = (Q_p - \sqrt{\cdot})(r^2)$ vanishes to order $p+1$ at $r=1$, so $\chi_p^{(j)}(1)$ equals the j-th derivative of $r\mapsto r$ at $r=1$ for every $j\le p$. For $y\neq0$, the chain rule expresses D^j of a radial function $g(\|y\|)$ through $g, g', \dots, g^{(j)}$ at $\|y\|$ and derivatives of the norm at y. Hence the two formulas have the same derivatives of order $\le p$ on the unit sphere, and $\Phi_p \in C^p(\mathbb{R}^n)$.

Dimension-free bounds. For a symmetric k-linear form S on $\mathbb{R}^n$, polarization gives $\|S\| \le (k^k/k!)\sup_{\|h\|=1}|S[h,\dots,h]|$. For $\|y\| < 1$ and a unit vector h, $D^{p+1}\Phi_p(y)[h,\dots,h] = \frac{d^{p+1}}{dt^{p+1}}Q_p(\|y\|^2 + 2t\langle y,h\rangle + t^2)\big|_{t=0}$, a polynomial in $(\|y\|^2, \langle y,h\rangle) \in [0,1]\times[-1,1]$ whose coefficients depend only on p. For $\|y\| > 1$, homogeneity gives $\|y+th\| = \|y\|\,(1 + 2(t/\|y\|)c + (t/\|y\|)^2)^{1/2}$ with $c = \langle y,h\rangle/\|y\| \in [-1,1]$, so $D^{p+1}\Phi_p(y)[h,\dots,h] = \|y\|^{-p}\,\frac{d^{p+1}}{ds^{p+1}}(1+2sc+s^2)^{1/2}\big|_{s=0}$, a polynomial in c multiplied by $\|y\|^{-p} \le 1$. These bounds hold in every dimension, so the supremum K_p over all dimensions is finite and depends only on p. For the lower bound, take $h = y/\|y\|$ and let $\|y\| \uparrow 1$: the diagonal value tends to $\chi_p^{(p+1)}(1) = (Q_p(r^2))^{(p+1)}|_{r=1}$. Since the full series $\sum_k \binom{1/2}{k}(r^2-1)^k = r$ has no terms of order ≥ 2 in $r-1$, and $(r^2-1)^k = (2(r-1)+(r-1)^2)^k$ has lowest order k, the truncation at $k=p$ changes the coefficient of $(r-1)^{p+1}$ only through $k = p+1$; hence $|\chi_p^{(p+1)}(1)| = (p+1)!\,2^{p+1}\,|\binom{1/2}{p+1}| = (2p-1)!!$, and $K_p \ge (2p-1)!! \ge 1$.

Lipschitz constant. A segment $[x,x']$ meets the unit sphere in at most two points. Splitting it there, each closed piece lies in the closed ball or in the closed exterior, where $D^p\Phi_p$ is given by one of the two smooth formulas; the bound $\|D^{p+1}\Phi_p\| \le K_p$ extends to the boundary by continuity. Integrating $D^{p+1}\Phi_p$ along each piece and using the continuity of $D^p\Phi_p$ at the junctions, $\|D^p\Phi_p(x) - D^p\Phi_p(x')\| \le K_p\|x - x'\|$. For $p=1$, $D^2\Phi_1 = I$ inside the ball and $D^2\|y\| = (I - yy^\top/\|y\|^2)/\|y\|$ has norm $1/\|y\| < 1$ outside, so $K_1 = 1$.

(iv) Inside the ball, $D^2\Phi_p(y)[h,h] = 2Q_p'(\|y\|^2)\|h\|^2 + 4Q_p''(\|y\|^2)\langle y,h\rangle^2$, which is bounded on the closed ball by a constant depending only on p; outside, $\|D^2\|y\|\| = 1/\|y\| \le 1$. For a symmetric matrix the operator norm is the supremum of $|h^\top Ah|$ over unit h, so $H_p < \infty$ depends only on p.

(ii) We have $\chi_p'(r) = 2rQ_p'(r^2)$, and Q_p' is the Taylor polynomial of degree $p-1$ of $\frac{d}{ds}\sqrt{s} = \frac12 s^{-1/2}$ at $s=1$. For $s\in(0,1]$, write $s^{-1/2} = (1-(1-s))^{-1/2} = \sum_{j\ge0} b_j(1-s)^j$ with $b_j = \binom{2j}{j}4^{-j} > 0$. Then $Q_p'(s) = \frac12\sum_{j=0}^{p-1} b_j(1-s)^j$ is a partial sum of a series of nonnegative terms with sum $\frac12 s^{-1/2}$, so $0 \le Q_p'(s) \le \frac12 s^{-1/2}$ and $0 \le \chi_p'(r) \le 2r\cdot\frac{1}{2r} = 1$ for $r\in(0,1]$; also $\chi_p'(0) = 0$. The radial profile of Φ_p is therefore 1-Lipschitz on $[0,\infty)$, so $|\Phi_p(x) - \Phi_p(x')| \le |\|x\| - \|x'\|| \le \|x - x'\|$. Outside the ball, $\nabla\Phi_p(y) = y/\|y\|$.

(iii) By (ii), χ_p is nondecreasing on $[0,1]$, and $\Phi_p \ge 1 = \chi_p(1)$ outside the ball, so $\min\Phi_p = \chi_p(0) = Q_p(0)$. The binomial series of $\sqrt{1-x}$ converges at $x=1$, and its coefficients of order $k\ge1$ are negative, so $0 = 1 - \sum_{k\ge1}|\binom{1/2}{k}|$ and $Q_p(0) = 1 - \sum_{k=1}^{p}|\binom{1/2}{k}| = \sum_{k>p}|\binom{1/2}{k}| \in (0,1)$. Finally, $|\binom{1/2}{k}| = 4^{-k}\binom{2k}{k}/(2k-1)$ for $k\ge1$, and induction on p gives $1 - \sum_{k=1}^{p}|\binom{1/2}{k}| = 4^{-p}\binom{2p}{p}$: both sides equal $1/2$ at $p=1$, and $4^{-p}\binom{2p}{p} - 4^{-p-1}\binom{2p+2}{p+1} = 4^{-p}\binom{2p}{p}/(2p+2) = |\binom{1/2}{p+1}|$. □

Theorem 3.3 (The exponent law). *Let $p \geq 1$ be an integer and $L > 0$, and let $\mathcal{F}_p(L)$ be the class of $f \in C^p(\mathbb{R}^n)$ whose p-th derivative is L-Lipschitz in operator norm. There are constants $K_p, H_p, \varsigma_p, k_p \in (0,\infty)$ depending only on p, with $K_1 = H_1 = 1$, such that the following hold at a known scale β. Let $r_p = (2\varepsilon K_p/L)^{1/p}$.*

(a) (Search.) *Every randomized adaptive algorithm anchored at $x_0 = 0$ that outputs an ε-stationary point with probability at least $2/3$ for every $f \in \mathcal{F}_p(L)$ with $f(x_0) - \inf f \leq \Delta_p := 2\big(3 - 4^{-p}\binom{2p}{p}\big)\varepsilon r_p$ makes*

$$T \;\geq\; \frac{n\,L^{2/p}}{17820\,(2K_p)^{2/p}\,\beta^2\,\varepsilon^{2+2/p}} \quad \textit{queries.}$$

The same holds over $\mathcal{F}(L_1) \cap \mathcal{F}_p(L)$ whenever $2\varepsilon H_p \leq L_1 r_p$.

(b) (Verification.) *If $c'\varepsilon\varsigma_1 \leq L_1(c'\varepsilon\varsigma_p/L)^{1/p}$, every algorithm that solves ε-vs-$c'\varepsilon$ verification over $\mathcal{F}(L_1) \cap \mathcal{F}_p(L)$ with confidence $2/3$ makes $T \geq 2\ln(9/8)\,L^{2/p}/\big(\varsigma_p^{2/p}\beta^2(c'\varepsilon)^{2+2/p}\big)$ queries.*

(c) (The exponent is Kolmogorov's.) *Every accept instance and every reject instance of ε-vs-$c'\varepsilon$ verification over $\mathcal{F}_p(L)$ are at logit distance at least $k_p\,\beta\,((c'-1)\varepsilon)^{1+1/p}L^{-1/p}$; hence no lower bound obtained from Theorem 2.7(b) yields a larger power of $1/\varepsilon$ than $2 + 2/p$.*

Proof. *(a)* Let K_p be as in Lemma C.4, $r_p = (2\varepsilon K_p/L)^{1/p}$, $t := 3r_p$, let V be the packing of Step 0 of the proof of Theorem 3.2, and for $v \in V$ put $m_v := tv$ and

$$f_v(x) := 2\varepsilon r_p\,\Phi_p\big((x - m_v)/r_p\big).$$

Class. $D^p f_v(x) = 2\varepsilon r_p^{1-p}(D^p\Phi_p)((x-m_v)/r_p)$, which is Lipschitz with constant at most $2\varepsilon r_p^{1-p}K_p/r_p = 2\varepsilon K_p r_p^{-p} = L$. So $f_v \in \mathcal{F}_p(L)$. Moreover $\|D^2 f_v\| \leq 2\varepsilon H_p/r_p$ (for $p = 1$, the gradient of f_v is L-Lipschitz), so $f_v \in \mathcal{F}(L_1)$ whenever $2\varepsilon H_p \leq L_1 r_p$. *Landscape.* $\nabla f_v(x) = 2\varepsilon\nabla\Phi_p((x-m_v)/r_p)$ has norm 2ε whenever $\|x - m_v\| \geq r_p$, so every ε-stationary point of f_v lies in the open ball of radius r_p about m_v. Distinct centers are at least $t = 3r_p$ apart, so the nearest-center map sends every ε-stationary point of f_v to v. *Closeness.* By Lemma C.4(ii), $|f_v(x) - f_{v'}(x)| \leq 2\varepsilon\|m_v - m_{v'}\| \leq 4\varepsilon t$ for every x. *Budget.* $f_v(0) - \inf f_v = 2\varepsilon r_p(\Phi_p(-tv/r_p) - \chi_p(0)) = 2\varepsilon r_p(3 - \chi_p(0)) = \Delta_p$ by Lemma C.4(iii). These are exactly the properties used in Steps 2–5 of the proof of Theorem 3.2, with the 1-Lipschitz profile Φ_p in place of χ. Those steps give $T \geq n/(1980\,\beta^2\varepsilon^2 t^2)$ for every $n \geq 1$: the Fano branch, which is Theorem 2.7(c) with $a = 4\beta\varepsilon t$, gives $T \geq n/(384\beta^2\varepsilon^2t^2)$ for $n \geq 56$, and the two-point branch gives $T \geq 1/(36\beta^2\varepsilon^2t^2) \geq n/(1980\beta^2\varepsilon^2t^2)$ for $n \leq 55$. Since $\varepsilon^2 t^2 = 9\varepsilon^2 r_p^2 = 9(2K_p)^{2/p}\varepsilon^{2+2/p}L^{-2/p}$, this is the bound of (a). For $p = 1$, $K_1 = 1$, $r_1 = 2\varepsilon/L$ and $\Delta_1 = 10\varepsilon^2/L$ reproduce Theorem 3.2(a).

(b) Let $\varsigma_p := \sup_{\mathbb{R}}|\tanh^{(p+1)}|$, which is finite because $\tanh^{(k)}$ is a polynomial in $\tanh$. Let $w_p := (c'\varepsilon\varsigma_p/L)^{1/p}$ and $f_1 := u_1 \circ \tau$ with $u_1(s) = c'\varepsilon w_p \tanh(s/w_p)$, and let $f_0 \equiv 0$. Then $\sup|u_1^{(p+1)}| = c'\varepsilon w_p^{-p}\varsigma_p = L$, so $f_1 \in \mathcal{F}_p(L)$; $\|\nabla f_1(\bar{x})\| = |u_1'(0)| = c'\varepsilon$ and $\nabla f_0 = 0$. Also $\sup|u_1''| = c'\varepsilon\varsigma_1/w_p \leq L_1$ by the hypothesis of (b), so both instances lie in $\mathcal{F}(L_1) \cap \mathcal{F}_p(L)$. The two instances at scale β have logit distance $\frac{1}{2}\beta\,\mathrm{osc}(u_1) = \beta c'\varepsilon w_p$. Theorem 2.7(b) gives $T \geq 2\ln(9/8)/(\beta c'\varepsilon w_p)^2 = 2\ln(9/8)L^{2/p}/(\varsigma_p^{2/p}\beta^2(c'\varepsilon)^{2+2/p})$.

(c) Let (β, f) be an accept and (β, f') a reject instance, put $h := f' - f$ and $\gamma := (c'-1)\varepsilon$. Then $\|\nabla h(\bar{x})\| \geq \gamma$ and $D^p h$ is $2L$-Lipschitz. With $e = \nabla h(\bar{x})/\|\nabla h(\bar{x})\|$, the function $q(s) := h(\bar{x} + se)$ has $q'(0) \geq \gamma$, and $q^{(p)}$ is $2L$-Lipschitz, hence absolutely continuous with $|q^{(p+1)}| \leq 2L$ almost everywhere. If $\mathrm{osc}(h) = \infty$ there is nothing to prove. Otherwise q is bounded, and Kolmogorov's inequality on the line (Kolmogorov, 1939), applied to $q - \frac{1}{2}(\sup q + \inf q)$, gives $\gamma \leq \|q'\|_\infty \leq C_{p+1}(\mathrm{osc}(q)/2)^{p/(p+1)}(2L)^{1/(p+1)}$ with a constant C_{p+1} depending only on p. Rearranging, $\mathrm{osc}(h) \geq \mathrm{osc}(q) \geq 2(\gamma/C_{p+1})^{1+1/p}(2L)^{-1/p}$, and the logit distance is $\frac{1}{2}\beta\,\mathrm{osc}(h) \geq k_p\beta\gamma^{1+1/p}L^{-1/p}$ with $k_p := C_{p+1}^{-1-1/p}2^{-1/p}$. Hence the amplitude of the problem is at least $k_p\beta((c'-1)\varepsilon)^{1+1/p}L^{-1/p}$, and the bound of Theorem 2.7(b) is at most a constant times $L^{2/p}/(\beta^2\varepsilon^{2+2/p})$. □

Remark C.5 (Constants). The proof uses only that K_p and H_p are finite. A grid evaluation of the diagonal values $D^{p+1}\Phi_p(y)[h,\dots,h]$ for $p \le 6$ finds no value above the lower bound $(2p-1)!!$ of Lemma C.4(i) (supplementary script `check_exponent_law.py`). For $p = 2$ the cubic cap of Appendix D has a smaller cap radius, which is why Theorem 3.2(b) has a larger constant.

Remark C.6 (Saturation: why (b) keeps the first-order constraint). The amplitude law bounds information by distance, but information can also vanish because the channel saturates. For $p \ge 3$, verification over $\mathcal{F}_p(L)$ alone is impossible at every fixed budget. Let $f_M(x) = M\tau(x)^3$ and $f'_M = f_M + c'\varepsilon\tau$: both lie in $\mathcal{F}_p(L)$, since their third derivatives are constant, and $\|\nabla f_M(\bar{x})\| = 0$ while $\|\nabla f'_M(\bar{x})\| = c'\varepsilon$. A query (x, y) with $\tau(x) - \tau(y) = \delta$ has logits $\theta = \beta M(\tau(x)^3 - \tau(y)^3)$ and $\theta + \beta c'\varepsilon\delta$, with $|\theta| \ge \beta M|\delta|^3/4$. By the Bregman form of Lemma B.1, its divergence is at most $\frac{1}{2}(\beta c'\varepsilon\delta)^2 \min\{\frac{1}{4}, e^{-(|\theta|-\beta c'\varepsilon|\delta|)_+}\}$, and splitting at $|\delta| = M^{-1/3}$ shows that the supremum over queries is $O(\beta^{4/3}(c'\varepsilon)^2 M^{-2/3})$ for fixed β, ε. So for every budget T some M makes the two transcript laws nearly equal, although the logit distance of the pair is infinite. Zooming in does not help: at any radius the odd cubic term either dominates or leaves a linear signal too small to read. The first-order constraint $\mathcal{F}(L_1)$ restores the zooming principle, and over $\mathcal{F}(L_1) \cap \mathcal{F}_p(L)$ verification is feasible by Theorem 4.1.

D Proof of the search lower bounds

This appendix proves Theorem 3.2. The argument packs exponentially many *shifted smoothed cones* around the anchor, shows that solving the search problem identifies the planted well, bounds the information carried by the transcript through Lemma B.4, and closes with Fano's inequality (a two-point Le Cam argument covering small n). Throughout, T is a fixed (deterministic) query budget, as in Lemma B.4; $B(m, r)$ denotes the closed Euclidean ball of radius r about m; and Δ_{inst} is the construction's own optimality-gap budget — no scaling of the bounds in a larger budget Δ is claimed (Section 3).

Theorem 3.2 (Search lower bounds)**.** *Let $n \ge 1$, $\beta > 0$, $\varepsilon > 0$. Consider the known-scale BT oracle (Definition 2.1) and algorithms anchored at $x_0 = 0$.*

(a) (Gradient-Lipschitz.) *Set $\Delta_{\mathrm{inst}} := 10\,\varepsilon^2/L_1$. Every randomized adaptive algorithm that, for every $f \in \mathcal{F}(L_1)$ with $f(x_0) - \inf f \le \Delta_{\mathrm{inst}}$, outputs an ε-stationary point of f with probability at least $2/3$ must make*

$$T \;\ge\; c\,\frac{nL_1^2}{\beta^2\varepsilon^4} \quad \textit{queries.}$$

(b) (Hessian-Lipschitz.) *Assume $2\sqrt{\varepsilon L_2} \le L_1$ and set $\Delta_{\mathrm{inst}} := (32/3)\,\varepsilon^{3/2}/\sqrt{L_2}$. Every randomized adaptive algorithm that, for every $f \in \mathcal{F}(L_1, L_2)$ with $f(x_0) - \inf f \le \Delta_{\mathrm{inst}}$, outputs an ε-stationary point of f with probability at least $2/3$ must make*

$$T \;\ge\; c\,\frac{nL_2}{\beta^2\varepsilon^3} \quad \textit{queries.}$$

Both parts hold with $c = 1/71280$ for all $n \ge 1$, and with $c = 1/13824$ for $n \ge 56$.

Proof. **The construction.** Fix $\varepsilon > 0$ and the class parameters, and let $\chi\colon [0,\infty) \to \mathbb{R}$ be a radial profile, one per class, specified below, with cap radius $r_0 > 0$. Let $t := 3r_0$ be the shift, and for a unit vector $v \in \mathbb{S}^{n-1}$ put the apex $m_v := tv$ and

$$f_v(x) := 2\varepsilon\,\chi\big(\|x - m_v\|\big), \qquad x \in \mathbb{R}^n.$$

The algorithm is anchored at $x_0 = \bar{x} = 0$. For $x \ne m_v$ write $r := \|x - m_v\|$ and $u := (x - m_v)/r$, and let $P_{\mathrm{rad}} := uu^\top$ and $P_{\mathrm{tan}} := I - uu^\top$. $V \subseteq \mathbb{S}^{n-1}$ is the packing constructed in Step 0. The two profiles:

- *For $\mathcal{F}(L_1)$ (a $C^{1,1}$ quadratic cap):* with $r_0 := 2\varepsilon/L_1$,

$$\chi(r) = r \quad (r \ge r_0), \qquad \chi(r) = \frac{r^2}{2r_0} + \frac{r_0}{2} \quad (0 \le r \le r_0).$$

Then $\chi(r_0) = r_0$ and $\chi'(r_0) = 1$ (C^1 match), and $\chi'' = 1/r_0$ inside the cap, 0 outside.

- *For $\mathcal{F}(L_1, L_2)$ (a $C^{2,1}$ cubic cap):* with $r_0 \coloneqq 2\sqrt{\varepsilon/L_2}$,

$$\chi(r) = r \quad (r \geq r_0), \qquad \chi(r) = \frac{r_0}{3} + \frac{r^2}{r_0} - \frac{r^3}{3r_0^2} \quad (0 \leq r \leq r_0),$$

i.e. $\chi'(r) = 2(r/r_0) - (r/r_0)^2 = \phi(r/r_0)$ with $\phi(s) \coloneqq 2s - s^2$. Then $\chi(r_0) = r_0$, $\chi'(r_0) = 1$ and $\chi''(r_0) = 0$ (C^2 match), and the third radial derivative is $-2/r_0^2$ inside the cap, 0 outside: the profile is $C^{2,1}$, with bounded third radial derivative. This branch is where the regime hypothesis $2\sqrt{\varepsilon L_2} \leq L_1$ (equivalently $4\varepsilon L_2 \leq L_1^2$) enters: it makes the cap's Hessian norm fit inside $\mathcal{F}(L_1)$ (Step 1). (The coefficient 2 in r_0 saturates the Hessian-Lipschitz constraint; taking $r_0 = 2\sqrt{2}\,\sqrt{\varepsilon/L_2}$ instead would weaken the regime hypothesis to $\sqrt{2\varepsilon L_2} \leq L_1$ at the price of a factor 2 in the constant c. We do not pursue this.)

In both cases χ is increasing with $0 \leq \chi' \leq 1$, so χ is 1-Lipschitz and f_v is a smoothed cone of slope 2ε with a rounded well of radius r_0 at the apex m_v.

Step 0 (packing). For $n \geq 15$ there exist $|V| \geq e^{n/16}/2$ unit vectors with pairwise Euclidean distance $\|v - v'\| \geq 1$ (equivalently $\langle v, v' \rangle \leq 1/2$). Indeed, for u, v i.i.d. uniform on $\mathbb{S}^{n-1}$, $\langle u, v \rangle$ is distributed as the first coordinate of a uniform point, and the standard spherical-cap estimate gives $\mathbb{P}\big(\langle u, v \rangle \geq 1/2\big) \leq e^{-n/8}$ (Ball, 1997, Lemma 2.2).

Draw $M = \lceil e^{n/16} \rceil$ i.i.d. uniform points; the expected number of pairs with $\langle v, v' \rangle > 1/2$ is at most $\binom{M}{2} e^{-n/8} \leq (M^2/2)\, e^{-n/8} \leq 1$ for $n \geq 15$ (there $M^2 \leq 2e^{n/8}$). Some realization has at most one bad pair; deleting one of its endpoints leaves $|V| \geq M - 1 \geq e^{n/16}/2$. Hence

$$\ln |V| \;\geq\; \frac{n}{16} - \ln 2.$$

Step 1 (class membership and instance gap). For a radial function $F(x) = g(r)$ with $g \coloneqq 2\varepsilon\,\chi$,

$$\nabla^2 F(x) = g''(r)\, P_{\mathrm{rad}} + \frac{g'(r)}{r}\, P_{\tan}, \tag{D.1}$$

with eigenvalues $g''(r)$ (once) and $g'(r)/r$ (multiplicity $n - 1$).

(i) $\mathcal{F}(L_1)$, quadratic cap. Inside the cap, $g'' = 2\varepsilon/r_0$ and $g'/r = 2\varepsilon/r_0$; outside, $g'' = 0$ and $g'/r = 2\varepsilon/r \leq 2\varepsilon/r_0$. So $\|\nabla^2 f_v\|_{\mathrm{op}} \leq 2\varepsilon/r_0 = L_1$ everywhere, attained on the cap; f_v is $C^{1,1}$ with gradient-Lipschitz constant L_1, i.e. $f_v \in \mathcal{F}(L_1)$.

(ii) $\mathcal{F}(L_1, L_2)$, cubic cap. Inside the cap, $g''(r) = 2\varepsilon\,(2/r_0 - 2r/r_0^2)$, so $\max|g''| = 4\varepsilon/r_0$ at $r = 0$; and $g'(r)/r = 2\varepsilon\,(2/r_0 - r/r_0^2)$, maximal $4\varepsilon/r_0$ at $r = 0$; outside, $g'' = 0$ and $g'/r = 2\varepsilon/r \leq 2\varepsilon/r_0$. Hence

$$\|\nabla^2 f_v\|_{\mathrm{op}} \;\leq\; \frac{4\varepsilon}{r_0} \;=\; 2\sqrt{\varepsilon L_2} \;\leq\; L_1 \quad \text{(the regime hypothesis)},$$

attained near the apex.

For the Hessian-Lipschitz constant we use the following exact description of the directional derivative of the Hessian. Note that its operator norm is *not*, in general, attained at a pure radial or tangential direction, so a genuine maximization over directions is required.

Lemma (radial third-derivative block). For $F(x) = g(\|x\|)$ three times differentiable at $x \neq 0$, set $P \coloneqq g'''(r)$ and $Q \coloneqq \big(g''(r) - g'(r)/r\big)/r$, with $r = \|x\|$ and $u = x/r$. Differentiating (D.1) once more, the fully symmetric third-derivative tensor is

$$T_{ijk} = (P - 3Q)\, u_i u_j u_k + Q\,\big(u_i \delta_{jk} + u_j \delta_{ik} + u_k \delta_{ij}\big).$$

Contracting the last slot with a unit direction $e = au + bw$ (w a unit tangent vector, $a^2 + b^2 = 1$) yields the directional derivative $D_e \nabla^2 F$ of the Hessian, which in an orthonormal frame $\{u, w, n - 2 \text{ further tangent vectors}\}$ is block-diagonal:

$$D_e \nabla^2 F = \begin{pmatrix} aP & bQ \\ bQ & aQ \end{pmatrix} \oplus\; aQ \cdot I_{n-2}.$$

The scalar aQ lies between the two eigenvalues of the 2×2 block (its distance to their midpoint is $|a||P-Q|/2$, at most the half-spread), hence is dominated, and

$$\|D_e \nabla^2 F\|_{\mathrm{op}} = \Lambda(a,b) \coloneqq \frac{|a|\,|P+Q| + \sqrt{a^2(P-Q)^2 + 4b^2Q^2}}{2},$$

so that $\mathrm{Lip}(\nabla^2 F) = \sup_r M(P,Q)$ with $M(P,Q) \coloneqq \max_{a^2+b^2=1} \Lambda(a,b)$. A pure radial direction ($a=1$) gives $\Lambda = \max(|P|,|Q|)$, but a mixed direction can strictly exceed this when $|P| < |Q|$, by up to the factor $c^* = 2/\sqrt{3} \approx 1.1547$: e.g. $M(0,Q) = (2/\sqrt{3})\,|Q| > |Q|$, attained at $a^2 = 1/3$. So M is not in general a pure-direction value and must be maximized over e.

Evaluation for the cubic cap (two regimes).

- *Inside* ($r < r_0$): $P = g''' = -4\varepsilon/r_0^2$ and $Q = (g'' - g'/r)/r = -2\varepsilon/r_0^2$ are both constant in r (indeed $g'' - g'/r = -2\varepsilon\, r/r_0^2$), so $P = 2Q$ — a parameter-free ratio, independent of ε, r_0, L_2. Then $\Lambda(a,b) = \frac{|Q|}{2}\big(3|a| + \sqrt{4-3a^2}\big)$, which is increasing in $|a|$ (indeed $\sqrt{4-3a^2} \geq |a|$ on $[0,1]$), so the maximum is at $a=1$, a pure radial direction: $M = |P| = 4\varepsilon/r_0^2$.
- *Outside* ($r \geq r_0$): $\chi(r) = r$ gives $P = g''' = 0$ and $Q = -2\varepsilon/r^2$, so $M(0,Q) = (2/\sqrt{3})\,|Q| = (2/\sqrt{3})\,(2\varepsilon/r^2) \leq (4/\sqrt{3})\,\varepsilon/r_0^2 \approx 2.309\,\varepsilon/r_0^2 < 4\varepsilon/r_0^2$.
- *Junction* ($r = r_0$): the Hessian is continuous across r_0 — both one-sided limits equal $(2\varepsilon/r_0)(I - uu^\top)$ — and Q is continuous there; only P jumps ($-4\varepsilon/r_0^2 \to 0$). Hence $\nabla^2 f_v$ is continuous everywhere and, at every point of differentiability (all x with $\|x - m_v\| \notin \{0, r_0\}$), $\|D_e \nabla^2 f_v\|_{\mathrm{op}} \leq 4\varepsilon/r_0^2$; integrating along segments of the convex domain $\mathbb{R}^n$, the Hessian is globally Lipschitz with constant the essential supremum over points and directions:
$$\mathrm{Lip}(\nabla^2 f_v) = \max\Big(\frac{4\varepsilon}{r_0^2},\ \frac{4}{\sqrt{3}}\frac{\varepsilon}{r_0^2}\Big) = \frac{4\varepsilon}{r_0^2} = L_2,$$
tight — the supremum is attained inside the cap at a pure radial direction.

Both class constraints hold, so $f_v \in \mathcal{F}(L_1, L_2)$. (At the apex f_v is C^2 but not C^3 — the $-\|x-m_v\|^3$ term — i.e. exactly $C^{2,1}$, which is all that $\mathcal{F}(L_1,L_2)$ requires.)

Instance gap. $\inf f_v = f_v(m_v) = 2\varepsilon\,\chi(0)$, and $f_v(0) = 2\varepsilon\,\chi(t) = 2\varepsilon \cdot 3r_0 = 6\varepsilon r_0$ since $t = 3r_0 \geq r_0$. With $\chi(0) = r_0/2$ (quadratic cap) resp. $\chi(0) = r_0/3$ (cubic cap),

$$f_v(0) - \inf f_v = 5\varepsilon r_0 = \frac{10\,\varepsilon^2}{L_1} \quad [\mathcal{F}(L_1)], \qquad f_v(0) - \inf f_v = \frac{16}{3}\,\varepsilon r_0 = \frac{32}{3}\,\frac{\varepsilon^{3/2}}{\sqrt{L_2}} \quad [\mathcal{F}(L_1,L_2)].$$

These are exactly the budgets Δ_{inst} in the theorem statement. Moreover 0 is *not* ε-stationary for any f_v: $\|\nabla f_v(0)\| = 2\varepsilon\,\chi'(t) = 2\varepsilon > \varepsilon$ (as $t = 3r_0 > r_0$). The anchor is a genuine non-solution, and the algorithm must actually search.

Step 2 (landscape). $\|\nabla f_v(x)\| = 2\varepsilon\,\chi'(r)$, which equals 2ε for all $r \geq r_0$, and for $r \leq r_0$ equals $2\varepsilon\,(r/r_0)$ (quadratic cap) resp. $2\varepsilon\,(2s - s^2)$ with $s = r/r_0$ (cubic cap). Hence $\|\nabla f_v(x)\| > \varepsilon$ for every x with $r \geq r_0$, and the ε-stationary set $\{x : \|\nabla f_v(x)\| \leq \varepsilon\}$ equals $B(m_v, r_0/2)$ for the quadratic cap (the condition is exactly $r/r_0 \leq 1/2$), and for the cubic cap is contained in $B\big(m_v, (1-1/\sqrt{2})\,r_0\big) \subseteq B(m_v, r_0/2)$ (the condition $2s - s^2 \leq 1/2$ gives $s \leq 1 - 1/\sqrt{2} \approx 0.293$). χ' is monotone, so there is no other stationary structure and no leveling far from the apex. In both cases: *every ε-stationary point of f_v lies in $B(m_v, r_0/2)$.*

Step 3 (disjointness; search implies identification). For distinct $v, v' \in V$, $\|m_v - m_{v'}\| = t\,\|v - v'\| \geq t = 3r_0$, so the balls $B(m_v, r_0/2)$ are pairwise disjoint with margin at least $3r_0 - r_0 = 2r_0$; and every $x \in B(m_v, r_0/2)$ has $\|x - m_{v'}\| \geq 3r_0 - r_0/2 = 5r_0/2 > r_0/2$ for every $v' \neq v$. The region map $\Psi(x) \coloneqq \arg\min_{v \in V} \|x - m_v\|$ is therefore single-valued on the union of the balls and equals v on $B(m_v, r_0/2)$. Put $\hat{v} \coloneqq \Psi(\text{output})$. If the algorithm is ε-correct for f_v with probability at least $2/3$, then with probability at least $2/3$ its output lies in $B(m_v, r_0/2)$ (Step 2), whence $\hat{v} = v$. So

$P_v[\hat{v} = v] \geq 2/3$ for *every* $v \in V$, where P_v is the law of the interaction under f_v; the identification error, averaged over a uniform prior on V, is at most $1/3$.

Step 4 (information). χ is 1-Lipschitz, so for every x, by the reverse triangle inequality and $\|v - v'\| \leq 2$,

$$\begin{aligned} |f_v(x) - f_{v'}(x)| &= 2\varepsilon \left|\chi(\|x - m_v\|) - \chi(\|x - m_{v'}\|)\right| \\ &\leq 2\varepsilon \left|\|x - m_v\| - \|x - m_{v'}\|\right| \\ &\leq 2\varepsilon \|m_v - m_{v'}\| \leq 4\varepsilon t. \end{aligned}$$

Hence for every query pair (x, y),

$$\left|\mathrm{gap}_{f_v}(x, y) - \mathrm{gap}_{f_{v'}}(x, y)\right| = \left|(f_v - f_{v'})(x) - (f_v - f_{v'})(y)\right| \leq 8\varepsilon t.$$

Now import Lemma B.4: with P_v the law of the T-query transcript under f_v,

$$\begin{aligned} \mathrm{KL}(P_v \,\|\, P_{v'}) &\leq \frac{\beta^2}{8} \sum_{s=1}^{T} \mathbb{E}_{P_v}\left[\left(\mathrm{gap}_{f_v}(X_s, Y_s) - \mathrm{gap}_{f_{v'}}(X_s, Y_s)\right)^2\right] \\ &\leq \frac{\beta^2}{8} \cdot T \cdot (8\varepsilon t)^2 \;=\; 8\,\beta^2\varepsilon^2 t^2\, T, \end{aligned} \tag{D.2}$$

uniformly over $v, v' \in V$. (Equivalently: the uniform-closeness form of Lemma B.4 with $a = 4\varepsilon t$ gives $(\beta^2/2)(4\varepsilon t)^2 T = 8\beta^2\varepsilon^2 t^2 T$ — the two routes give the coefficient 8 exactly.)

Step 5 (Fano; the two-point branch; conclusion). Let $V \sim \mathrm{Unif}(V)$ and note that $V \to Z \to \hat{v}$ is a Markov chain, Z the transcript ($\hat{v}$ is a function of Z and environment-independent coins, as in the Setup of Lemma B.4). Bounding the mutual information by the average pairwise KL (convexity of KL toward the mixture centroid),

$$I(V; Z) \;\leq\; \frac{1}{|V|^2} \sum_{v, v' \in V} \mathrm{KL}(P_v \| P_{v'}) \;\leq\; \max_{v, v'} \mathrm{KL}(P_v \| P_{v'}) \;\leq\; 8\,\beta^2\varepsilon^2 t^2\, T.$$

By Fano's inequality with data processing (natural logarithms; see, e.g., Tsybakov, 2009, Chapter 2),

$$\mathbb{P}(\hat{v} \neq V) \;\geq\; 1 - \frac{I(V; Z) + \ln 2}{\ln |V|}.$$

The left-hand side is at most $1/3$ (Step 3) and $\ln |V| \geq n/16 - \ln 2$ (Step 0), so

$$I(V; Z) \;\geq\; \tfrac{2}{3} \ln |V| - \ln 2 \;\geq\; \tfrac{2}{3}\left(\tfrac{n}{16} - \ln 2\right) - \ln 2 \;=\; \frac{n}{24} - \frac{5}{3} \ln 2,$$

and for $n \geq n_0 := 56$ the right-hand side is at least $n/48$. Combining with (D.2),

$$8\,\beta^2\varepsilon^2 t^2\, T \;\geq\; \frac{n}{48}, \qquad \text{i.e.} \qquad T \;\geq\; \frac{n}{384\,\beta^2\varepsilon^2 t^2}.$$

Substituting the shift $t = 3r_0$:

$$\begin{aligned} \mathcal{F}(L_1): \quad t = \frac{6\varepsilon}{L_1}, \quad t^2 = \frac{36\varepsilon^2}{L_1^2} \quad &\Longrightarrow \quad T \geq \frac{nL_1^2}{13824\,\beta^2\varepsilon^4}; \\ \mathcal{F}(L_1, L_2): \quad t = 6\sqrt{\varepsilon/L_2}, \quad t^2 = \frac{36\varepsilon}{L_2} \quad &\Longrightarrow \quad T \geq \frac{nL_2}{13824\,\beta^2\varepsilon^3}. \end{aligned}$$

Small n ($1 \leq n \leq 55$): the two-point branch. Use the two-point packing $V = \{v, -v\}$ ($\|v - (-v)\| = 2 \geq 1$, so Steps 1–4 apply verbatim). From Step 3, $P_v[\hat{v} = v] \geq 2/3$ and $P_{-v}[\hat{v} = -v] \geq 2/3$; since $\hat{v}$ is obtained from the transcript through an environment-independent kernel, data processing for total variation gives

$$\mathrm{TV}(P_v, P_{-v}) \;\geq\; P_v[\hat{v} = v] - P_{-v}[\hat{v} = v] \;\geq\; \tfrac{2}{3} - \tfrac{1}{3} = \tfrac{1}{3}$$

(Le Cam), and Pinsker's inequality yields $\mathrm{KL}(P_v \| P_{-v}) \geq 2\,\mathrm{TV}^2 \geq 2/9$. With (D.2), $8\beta^2\varepsilon^2 t^2 T \geq 2/9$, i.e. $T \geq 1/(36\,\beta^2\varepsilon^2 t^2)$, and since $n \leq 55$,

$$T \;\geq\; \frac{1}{36\,\beta^2\varepsilon^2 t^2} \;\geq\; \frac{n}{1980\,\beta^2\varepsilon^2 t^2};$$

substituting t as above reproduces both rates with the absolute constant $1/71280 = 1/(55 \cdot 36^2)$ in place of $1/13824$.

Conclusion. For $n \geq 56$ the Fano branch gives both bounds with $c = 1/13824$; for $1 \leq n \leq 55$ the two-point branch gives them with $c = 1/71280$; and since $1/71280 \leq 1/13824$, the single constant $c = 1/71280$ serves every $n \geq 1$:

$$T \;\geq\; c\,\frac{nL_1^2}{\beta^2\varepsilon^4} \quad [\mathcal{F}(L_1)], \qquad T \;\geq\; c\,\frac{nL_2}{\beta^2\varepsilon^3} \quad [\mathcal{F}(L_1, L_2), \text{ in the regime } 2\sqrt{\varepsilon L_2} \leq L_1]. \qquad \square$$

Remark D.1 (From verification to search; positioning). Any algorithm that *verifies-then-outputs* — producing a candidate $\hat{x}$ together with a certificate, correct with probability at least $2/3$, that $\hat{x}$ is ε-stationary, and outputting $\hat{x}$ only when the certificate accepts — is in particular a $2/3$-correct search algorithm on the packing $\{f_v\}$, so the lower bounds of Theorem 3.2 apply to it verbatim (Steps 0–5 are unchanged). The confidence-score upper bound of El Bakkali et al. (2026a) uses a different oracle, as explained in Section 1; it does not close the upper-bound question for raw BT bits.

E Proofs of the verification theorems

This appendix proves Theorem 4.1 (the known-scale verification interval and its saturation floor) and Theorem 4.2 (the nonadaptive $\Omega(n)$ verification lower bound), and derives the combined nonadaptive lower endpoint (Corollary E.5). Throughout, $\sigma(z) = 1/(1+e^{-z})$ is the logistic link, $\sigma'(z) = \sigma(z)(1-\sigma(z)) \in (0, 1/4]$, and a query on a pair (x, y) returns an independent bit distributed as $\mathrm{Bern}\big(\sigma(\beta(f(x) - f(y)))\big)$ with β known (Definition 2.1). We write $\mathrm{gap}_h(x, y) := h(x) - h(y)$. All logarithms are natural unless subscripted. All constants below are absolute and deliberately unoptimized: the content of both theorems is the scaling in $(n, L_1, \beta, \varepsilon, \delta)$, and we track every constant so that the reader can re-derive the bounds.

E.1 Proof of Theorem 4.1: the interval and the saturation floor

Theorem 4.1 (Known-scale verification: rate interval and saturation)**.** *Fix $c' = 2$, and let $m(\beta, \varepsilon, n, \delta)$ be the query complexity of δ-confidence ε-vs-2ε gradient-norm verification over $\mathcal{F}(L_1)$ at a given $\bar{x}$ under the known-scale BT oracle. There are absolute constants C, C_0 such that, with $c_0 := 32\ln(9/8) \approx 3.77$:*

(i) (Upper bound, subsaturated scales.) *Set $\beta^* := 2nL_1/\varepsilon^2$. For all $0 < \beta \leq \beta^*$,*

$$m(\beta, \varepsilon, n, \delta) \;\leq\; C\,\frac{n^2L_1^2}{\beta^2\varepsilon^4}\log(1/\delta);$$

for $\beta \geq \beta^*$ part (ii) takes over, so every scale is covered.

(ii) (Saturation.) *If $\beta \geq \beta^*$, then $m(\beta, \varepsilon, n, \delta) \leq C_0 \log(1/\delta)$: the cost collapses to a floor free of n, β, ε, and L_1.*

(iii) (Lower bounds.) *For every $\beta > 0$: $m \geq c_0\, L_1^2/(\beta^2\varepsilon^4)$ whenever $\delta \leq 1/3$; the stronger form*

$$m \;\geq\; 16\,\frac{L_1^2}{\beta^2\varepsilon^4}\ln(1/\delta)$$

holds for $\delta \leq 1/16$; and $m \geq \log_2\big(1/(2\delta)\big)$ unconditionally.

Fix $f \in \mathcal{F}(L_1)$ and write $g := \nabla f(\bar{x})$; we test $H_0 : \|g\| \leq \varepsilon$ against $H_1 : \|g\| \geq 2\varepsilon$, i.e. the gap ratio is $c' = 2$. The generalization to an arbitrary $c' > 1$ replaces the factor-4 mean-square gap of Step V5 ($4\varepsilon^2/n$ versus ε^2/n) by c'^2 versus 1 and changes only absolute constants. The proof has four parts: Part A (Steps V1–V6) is the upper bound for $\beta \leq \beta^* := 2nL_1/\varepsilon^2$; Part B (Steps V7–V8) is the saturation analysis for $\beta \geq \beta^*$; Part C (Steps V9–V11) contains the three lower bounds (the constant-confidence two-point bound, its δ-tracked version, and the universal confidence floor); Part D assembles the interval.

Proof. **Part A — the upper bound $T \leq C\, n^2L_1^2/(\beta^2\varepsilon^4)\, \log(1/\delta)$ for $\beta \leq \beta^*$.**

Step V1 (central-difference identity and bias bound over $\mathcal{F}(L_1)$). For a unit vector u and a step $s > 0$, probe the pair $(\bar{x} + su, \bar{x} - su)$. Let $\phi(r) := \langle \nabla f(\bar{x} + ru), u\rangle$; then ϕ is L_1-Lipschitz, since $|\phi(r) - \phi(r')| = |\langle \nabla f(\bar{x} + ru) - \nabla f(\bar{x} + r'u), u\rangle| \le L_1|r - r'|$, and $\phi(0) = \langle g, u\rangle$. By the fundamental theorem of calculus along the segment,

$$f(\bar{x} + su) - f(\bar{x} - su) \;=\; \int_{-s}^{s} \phi(r)\, dr \;=\; 2s\langle g, u\rangle + R, \tag{E.1}$$

where $R := \int_{-s}^{s} \big(\phi(r) - \phi(0)\big)\, dr$ satisfies $|R| \le \int_{-s}^{s} L_1|r|\, dr = L_1 s^2$. Hence the per-direction logit is

$$\mathrm{logit}(u) := \beta\big[f(\bar{x} + su) - f(\bar{x} - su)\big] = 2\beta s\langle g, u\rangle + \beta R, \qquad |\beta R| \le \beta L_1 s^2. \tag{E.2}$$

The bound $|R| \le L_1 s^2$ is tight over $\mathcal{F}(L_1)$: the gradient kink $\phi(r) = L_1|r|$ attains it exactly. The regularity distinction matters: over $\mathcal{F}(L_1)$ the symmetric pair does *not* asymptotically beat a one-sided difference — both have bias-to-signal ratio $L_1 s/(2|\langle g, u\rangle|)$; the strict improvement from Hessian cancellation (bias $O(L_2 s^3)$) would require a Lipschitz Hessian, which Theorem 4.1 does not assume. We keep the symmetric pair because it eliminates the unknown offset $f(\bar{x})$ and matches (E.1); the only fact used below is $|\beta R| \le \beta L_1 s^2$.

Step V2 (probe radius; deterministic bias budget). Fix the probe radius

$$s_* := \frac{\varepsilon}{8\sqrt{n}\, L_1}.$$

The deterministic per-direction bias, measured in the $\langle g, u\rangle$-scale, is then

$$b_0 := \frac{|R|}{2s_*} \le \frac{L_1 s_*}{2} = \frac{\varepsilon}{16\sqrt{n}}. \tag{E.3}$$

Step V3 (precision transfer: reads to logit). Fix a direction u; the bit for this pair is $\mathrm{Bern}(p)$ with $p = \sigma(\mathrm{logit}(u))$. From m i.i.d. reads form the empirical frequency $\hat{p}$; by Hoeffding's inequality (two-sided), with probability at least $1 - \delta'$,

$$|\hat{p} - p| \;\le\; \sqrt{\frac{\log(2/\delta')}{2m}} \;=:\; \mathrm{rad},$$

where we fix the per-read failure probability $\delta' := 1/100$ once and for all (it enters only through $\log(2/\delta') = \log 200$). Introduce the dimensionless scale parameter

$$\nu := \frac{\beta\varepsilon^2}{nL_1} \qquad (\text{so } \beta \le \beta^* \iff \nu \le 2)$$

and the clipped inverse-link estimate $\widehat{\mathrm{logit}} := \mathrm{clip}_{[-3,3]}\big(\sigma^{-1}(\hat{p})\big)$, where $\sigma^{-1}(y) = \log\big(y/(1-y)\big)$. The bulk read logit — a typical direction at the hardest instances, where $|\langle g, u\rangle|$ is of order $\|g\|/\sqrt{n} \le 2\varepsilon/\sqrt{n}$ — is at most $z_{\mathrm{ref}} := 2\beta s_* \cdot (3\varepsilon/\sqrt{n}) = 0.75\,\nu \le 1.5$ for $\beta \le \beta^*$, so the working range $|z| \le 3$ covers with slack every read that needs precision; the rare aligned-tail reads (logit up to order $\sqrt{n}\, z_{\mathrm{ref}}$) are clipped and are negligible to the mean (Steps V4–V5). On the event $\{|\hat{p} - p| \le \mathrm{rad}\}$, wherever $|\mathrm{logit}| \le 3$,

$$\Big|\widehat{\mathrm{logit}} - \mathrm{logit}\Big| \le L_{\mathrm{work}} \cdot \mathrm{rad}, \qquad L_{\mathrm{work}} := \sup_{|z| \le 3} \frac{1}{\sigma'(z)} = \frac{1}{\sigma'(3)} = 22.14\ldots \le 23, \tag{E.4}$$

an absolute constant: indeed $(\sigma^{-1})'(y) = 1/(y(1-y)) = 1/\sigma'(\sigma^{-1}(y))$ and $\sigma'(z) \ge e^{-|z|}/4$, and the mean value theorem applies along the segment from p to $\hat{p}$ (if the raw value $\sigma^{-1}(\hat{p})$ exits $[-3, 3]$, the clip only helps: σ^{-1} is increasing, so the error to a target inside $[-3, 3]$ is realized inside the working range). Choose the target logit precision

$$\tau_* := \frac{\beta\varepsilon^2}{64\, nL_1} \;=\; \frac{\nu}{64},$$

and take

$$m := \left\lceil \frac{L_{\mathrm{work}}^2 \log(2/\delta')}{2\tau_*^2} \right\rceil = C_0\, L_{\mathrm{work}}^2 \Big(\frac{nL_1}{\beta\varepsilon^2}\Big)^2 \log(2/\delta'), \qquad C_0 = \frac{64^2}{2} = 2048 \tag{E.5}$$

(up to the integer rounding, absorbed into the final constant C). Then $|\widehat{\mathrm{logit}} - \mathrm{logit}| \le \tau_*$ with probability at least $1-\delta'$, so the estimation bias in the $\langle g, u\rangle$-scale is

$$b_1 := \frac{\tau_*}{2\beta s_*} = \frac{\varepsilon}{16\sqrt{n}} \qquad \text{with probability} \ge 1-\delta'. \tag{E.6}$$

For $\beta \le \beta^*$ (i.e. $\nu \le 2$) the bulk logit satisfies $z_{\mathrm{ref}} \le 1.5$, so $L_{\mathrm{work}} \le 23$ is a single absolute constant and reads are cheap. The saturated branch $\beta \ge \beta^*$ is Part B: there s is retuned *down* so that the read logit stays $O(1)$ and L_{work} stays $O(1)$ — the magnitude is still read, not just the sign.

Step V4 (per-direction estimator; sphere moments). Draw $u \sim \mathrm{Unif}(\mathbb{S}^{n-1})$ and set

$$W := \frac{\widehat{\mathrm{logit}}}{2\beta s_*}, \qquad Y := \min\big(W^2,\ \Lambda_{\mathrm{cap}}\big), \qquad \Lambda_{\mathrm{cap}} := \frac{25\,\varepsilon^2}{n}.$$

The cap Λ_{cap} (subscripted to avoid the statement-level $\Lambda = \beta_{\max} L_1$ of Proposition F.3) amounts to capping $|W|$ at $5\varepsilon/\sqrt{n}$, i.e. $|\widehat{\mathrm{logit}}| \le (5\varepsilon/\sqrt{n}) \cdot 2\beta s_* = 1.25\,\nu \le 2.5 < 3$, inside the working range of Step V3. On a *good read* (probability $\ge 1-\delta'$), by (E.3) and (E.6),

$$|\,W - \langle g, u\rangle\,| \ \le\ b_0 + b_1 \ =:\ b \ =\ \frac{\varepsilon}{8\sqrt{n}}. \tag{E.7}$$

(For the rare aligned reads with $|\mathrm{logit}| > 3$ one has $|\langle g, u\rangle| \ge 12\varepsilon/(\nu\sqrt{n}) - b_0 > 5\varepsilon/\sqrt{n} = \sqrt{\Lambda_{\mathrm{cap}}} \ge \sqrt{Y}$, so the one-sided bounds used in Step V5 hold trivially through the cap.) For u uniform on the sphere, writing $\langle g, u\rangle = \|g\| u_1$ in a frame whose first axis is $g/\|g\|$, the moments $\mathbb{E}[u_1^2] = 1/n$ and $\mathbb{E}[u_1^4] = 3/(n(n+2))$ give

$$\mathbb{E}\langle g, u\rangle^2 = \frac{\|g\|^2}{n}, \qquad \mathbb{E}\langle g, u\rangle^4 = \frac{3\|g\|^4}{n(n+2)}, \qquad \mathrm{Var}\big(\langle g, u\rangle^2\big) = \frac{2(n-1)\,\|g\|^4}{n^2(n+2)}, \tag{E.8}$$

so the relative variance is $2(n-1)/(n+2) \to 2$: an *n-free* constant — this is why $O(\log(1/\delta))$ directions will suffice regardless of n.

Step V5 (mean separation for all norms, via monotonicity of the clipped mean). Let $\mu := \mathbb{E}[Y]$. Since $\langle g, u\rangle^2 = \|g\|^2 u_1^2$, the map $\|g\| \mapsto \min(\langle g, u\rangle^2, \Lambda_{\mathrm{cap}})$ is pointwise non-decreasing in $\|g\|$, hence so is its mean, and the good-read envelopes below are monotone in $|\langle g, u\rangle|$ as well; it therefore suffices to separate the two *boundary* norms $\|g\| = \varepsilon$ (hardest H_0) and $\|g\| = 2\varepsilon$ (hardest H_1) — every larger H_1 norm only increases the clipped mean, and $\Lambda_{\mathrm{cap}} = 25\varepsilon^2/n$ bites only the rare aligned tail ($\mathbb{P}(nu_1^2 > 25/4)$ is small), removing at most $0.2\,\varepsilon^2/n$ at $\|g\| = 2\varepsilon$.

- Under H_0 ($\|g\| \le \varepsilon$): on a good read $Y \le (|\langle g, u\rangle| + b)^2$, and with $\mathbb{E}|\langle g, u\rangle| \le \sqrt{\mathbb{E}\langle g, u\rangle^2} = \varepsilon/\sqrt{n}$,

  $$\mathbb{E}[Y \mid \text{good}] \le \mathbb{E}\langle g, u\rangle^2 + 2b\,\mathbb{E}|\langle g, u\rangle| + b^2 \le \frac{\varepsilon^2}{n}\Big(1 + \frac{1}{4} + \frac{1}{64}\Big) = 1.266\,\frac{\varepsilon^2}{n}.$$

  A bad read (probability $\le \delta' = 1/100$) contributes at most $\Lambda_{\mathrm{cap}} = 25\varepsilon^2/n$, so $\mu(H_0) \le 1.266\,\varepsilon^2/n + \delta'\Lambda_{\mathrm{cap}} \le 1.52\,\varepsilon^2/n$.

- Under H_1 ($\|g\| \ge 2\varepsilon$): by monotonicity take $\|g\| = 2\varepsilon$. On a good read $Y \ge \big((|\langle g, u\rangle| - b)_+\big)^2$ before the cap, and $\mathbb{E}\big[(|\langle g, u\rangle| - b)_+^2\big] \ge \mathbb{E}\langle g, u\rangle^2 - 2b\,\mathbb{E}|\langle g, u\rangle| = 4\varepsilon^2/n - \varepsilon^2/(2n) = 3.5\,\varepsilon^2/n$; subtracting the cap tail ($\le 0.2\,\varepsilon^2/n$) and the bad-read loss ($\le \delta'\Lambda_{\mathrm{cap}} \le 0.3\,\varepsilon^2/n$) gives $\mu(H_1) \ge 3.0\,\varepsilon^2/n$.

Thus, with the threshold $\theta := 2.25\,\varepsilon^2/n$,

$$\mu(H_0) \le 1.52\,\frac{\varepsilon^2}{n} \ <\ \theta \ <\ 3.0\,\frac{\varepsilon^2}{n} \le \mu(H_1)\,: \tag{E.9}$$

a fixed constant gap on both sides ($\theta - 1.52 = 0.73$ and $3.0 - \theta = 0.75$, in ε^2/n units; the binding gap is 0.73).

Step V6 (median-of-means aggregation; total budget). Estimate μ by a median of means over $K = Bq$ fresh directions (fresh u_i *and* fresh reads per direction), grouped into $B = \lceil 20 \log(1/\delta) \rceil$ blocks of $q = 200$ directions each; bad reads (probability $\le \delta' = 1/100$ each) merely make a Y_i an outlier, absorbed by the median. Each $Y_i \in [0, \Lambda_{\mathrm{cap}}]$ has standard deviation $\sigma_Y \le 6\,\varepsilon^2/n$: the dominant contribution is $\mathrm{Var}(\langle g, u\rangle^2) \le 32\,(\varepsilon^2/n)^2$ at the boundary $\|g\| = 2\varepsilon$ by (E.8) (i.e. standard deviation $\le 5.66\,\varepsilon^2/n$; an n-free relative variance), and the cap only decreases it while the read errors are absorbed in the stated $6\,\varepsilon^2/n$. By Chebyshev's inequality, a block mean of $q = 200$ lands on the correct side of θ with probability

$$p_{\mathrm{blk}} \;\ge\; 1 - \frac{\sigma_Y^2}{q\,(0.73\,\varepsilon^2/n)^2} \;\ge\; 1 - \frac{36}{200 \times 0.5329} \;=\; 0.662 \;=\; \frac{1}{2} + 0.162.$$

By Hoeffding's inequality applied to the B independent block indicators, the median of the B block means errs with probability at most

$$\exp\!\big(-2B\,(0.162)^2\big) = \exp(-0.0525\,B) \le \exp\!\big(-1.05\,\log(1/\delta)\big) \le \delta. \tag{E.10}$$

The decision rule is: *reject H_0 iff the median of block means exceeds θ*. The total query count is

$$T = K\,m = (Bq)\,m = C\,\frac{n^2 L_1^2}{\beta^2 \varepsilon^4}\,\log(1/\delta), \qquad C \coloneqq 20 \cdot 200 \cdot C_0\,L_{\mathrm{work}}^2\,\log 200 \tag{E.11}$$

(integer roundings absorbed into C; $\delta' = 1/100$). Since $z_{\mathrm{ref}} \le 1.5$ for all $\beta \le \beta^*$, $L_{\mathrm{work}} \le 23$ is one absolute constant, so C is absolute (n-, β-, ε-, δ-free). Numerically $C \approx 2 \times 10^{10}$ as written — loose by many orders of magnitude and reducible; the theorem is the scaling $n^2 L_1^2/(\beta^2\varepsilon^4)\log(1/\delta)$ with an absolute coefficient. Note for later use that this tester is *nonadaptive*: all pairs are fixed before any bit is observed.

Part B — saturation: for $\beta \ge \beta^* = 2nL_1/\varepsilon^2$, an n-free $O(\log(1/\delta))$ cost.

The difficulty is that when a read logit z is large the sigmoid saturates, and reading a *magnitude* at logit precision τ costs $m \sim e^{|z|}/\tau^2$ reads — only the *sign* is cheap when saturated. One must not credit free magnitude precision to saturated reads; the resolution is to keep the working interval used by the clipped statistic readable.

Step V7 (tune the probe radius down; magnitudes stay readable). For $\beta \ge \beta^*$, pick the tuned probe radius

$$s_{\mathrm{sat}} \coloneqq \frac{\sqrt{n}}{6\,\beta\varepsilon} \qquad \text{(so the boundary bulk logit is } 2\beta s_{\mathrm{sat}} \cdot 3\varepsilon/\sqrt{n} = 1\text{)}.$$

Then: (i) the working logit range is $O(1)$, so $L_{\mathrm{work}} = 1/\sigma'(O(1)) = O(1)$ — magnitudes *are* readable at $O(1)$ cost per read, not sign-only; and (ii) the bias budget still holds: $b_0 = L_1 s_{\mathrm{sat}}/2 = L_1\sqrt{n}/(12\beta\varepsilon) \le \varepsilon/(16\sqrt{n})$ iff $\beta \ge \frac{4}{3}\,nL_1/\varepsilon^2$, which holds for $\beta \ge \beta^* = 2nL_1/\varepsilon^2$. Shrinking s costs nothing: the large β compensates the smaller value gap. The required logit precision (from $b_1 \le \varepsilon/(16\sqrt{n})$ as in (E.6)) is $\tau = 2\beta s_{\mathrm{sat}} \cdot \varepsilon/(16\sqrt{n}) = 1/48$, an absolute constant; hence the per-direction read count is

$$m_{\mathrm{sat}} = \left\lceil \frac{L_{\mathrm{work}}^2 \log(2/\delta')}{2\tau^2} \right\rceil = O\big(\log(1/\delta')\big) = O(1) \qquad (n\text{-free and } \beta\text{-free}).$$

Aggregation is unchanged (Step V6): $K = O(\log(1/\delta))$ directions suffice, because the relative variance $\to 2$ of Step V4 is n-free. The total in saturation is

$$T_{\mathrm{sat}} = K\,m_{\mathrm{sat}} = O\big(\log(1/\delta)\big), \tag{E.12}$$

an *n-free* polylogarithmic cost. The "sign-only when saturated" fear is avoided because shrinking s keeps the working logit scale bounded. Projections outside the working interval are handled by the clipping from Part A; the argument does not assert a uniform bound on all logits.

Step V8 (a non-increasing upper bound and its threshold). Part A uses the bias-controlled radius s_* for $\beta \le \beta^*$, with a per-direction upper bound of order $(nL_1/(\beta\varepsilon^2))^2$. Part B uses s_{sat} for $\beta \ge \beta^*$, with an absolute per-direction bound. Both choices satisfy the bias and working-interval requirements

in their respective regimes. Increasing the absolute constant if necessary combines them into the non-increasing upper envelope

$$C \max\left\{1, \frac{n^2 L_1^2}{\beta^2 \varepsilon^4}\right\} \log(1/\delta).$$

Thus increasing β does not force the upper bound to grow. This argument establishes an achievable envelope, rather than monotonicity of the exact minimax complexity or a sharp transition at β^*.

Part C — lower bounds (the exact two-point amplitude; the universal floor).

Step V9 (the amplitude). By Proposition 3.1 with $\beta_{\min} = \beta_{\max} = \beta$ and $c' = 2$, the amplitude of the verification problem is

$$a \;=\; \frac{\beta \varepsilon^2}{4L_1},$$

attained by the ridge pair of Step 3 of its proof: an accept instance with gradient norm ε and a reject instance with gradient norm 2ε, both in $\mathcal{F}(L_1)$, whose difference is a single smoothed step of height $\varepsilon^2/(2L_1)$. No pair of promise instances is closer in logit distance.

Step V10 (constant and high confidence). Theorem 2.7(b) gives, for every algorithm, adaptive or not,

$$T \;\geq\; \frac{2\ln(9/8)}{a^2} \;=\; c_0 \frac{L_1^2}{\beta^2 \varepsilon^4}, \qquad c_0 = 32\ln(9/8) = 3.7690\ldots, \tag{E.13}$$

at confidence $2/3$, and, for $\delta \leq 1/16$,

$$T \;\geq\; \frac{\ln(1/\delta)}{a^2} \;=\; 16\, \frac{L_1^2 \ln(1/\delta)}{\beta^2 \varepsilon^4}. \tag{E.14}$$

The qualifier $\delta \leq 1/16$ is exact for the step $\ln(1/(4\delta)) \geq \frac{1}{2}\ln(1/\delta)$ in the proof of Theorem 2.7(b), which fails already at $\delta = 1/8$. For a general gap ratio $c' > 1$ the amplitude is $\beta(c'-1)^2\varepsilon^2/(4L_1)$ and the same argument gives $T \geq 32\ln(9/8)L_1^2/((c'-1)^4\beta^2\varepsilon^4)$.

Step V11 (the universal confidence floor $T \geq \log_2(1/(2\delta))$). Let $f_0 \equiv 0$, an accept instance, and let f_1 be any reject instance, for example $f_1(x) = 2\varepsilon\langle e_1, x - \bar{x}\rangle$, which lies in $\mathcal{F}(L_1)$ for every L_1. Under f_0, every query (x, y) returns $\mathrm{Bern}(\sigma(0)) = \mathrm{Bern}(1/2)$ independently of the (adaptively chosen) pair, so the T-bit transcript law is $P_0 = \mathrm{Unif}\{0,1\}^T$, and the bits are independent of the algorithm's internal coins. Including the coins ω in the transcript, the ω-marginals of P_0 and P_1 agree (the coin law does not depend on the environment), so for any event A, $P_1(A) - P_0(A) = \mathbb{E}_\omega\big[P_1(A_\omega \mid \omega) - 2^{-T}|A_\omega|\big] \leq 1 - 2^{-T}$, i.e. $\mathrm{TV}(P_0, P_1) \leq 1 - 2^{-T}$. A test with both errors $\leq \delta$ forces $\mathrm{TV}(P_0, P_1) \geq 1 - 2\delta$, so $2^{-T} \leq 2\delta$, i.e.

$$T \;\geq\; \log_2 \frac{1}{2\delta} \;=\; \log_2 \frac{1}{\delta} - 1 \;\geq\; \frac{1}{2}\log_2 \frac{1}{\delta} \qquad \text{for } \delta \leq \tfrac{1}{4}. \tag{E.15}$$

This floor is n-, β-, L_1-, and ε-free, and it matches the Part-B upper bound (E.12) up to an absolute constant: in the saturated regime $\beta \geq \beta^*$ the floor is tight.

Part D — assembly of the interval. Part A gives, for every $0 < \beta \leq \beta^*$, the upper bound (E.11); Part B gives, for $\beta \geq \beta^*$, the n-free $O(\log(1/\delta))$ cost (E.12), so verification is possible at every finite $\beta > 0$. Part C gives the lower endpoint: log-free $c_0\, L_1^2/(\beta^2\varepsilon^4)$ at success $2/3$ (E.13), the same endpoint with the factor $16\log(1/\delta)$ in place of c_0 for $\delta \leq 1/16$ (E.14), and the universal floor (E.15), which is operative in the saturated regime where $L_1^2/(\beta^2\varepsilon^4)$ is small. The $(L_1, \beta, \varepsilon)$-dependence $L_1^2/(\beta^2\varepsilon^4)$ is thus pinned at both endpoints. The n-power is *not* resolved by these arguments and lies in the interval $[n^0, n^2]$: the two-point method of Part C is n-oblivious, and since the amplitude of Step V9 is exact, no choice of pair improves the order of (E.13), and we do not claim the $\sim n$ improvement of the upper bound occasionally suggested by a learn-the-direction heuristic (that route was examined and not substantiated — learning the direction costs the same n^2). For *nonadaptive* testers, Theorem 4.2 lifts the lower endpoint's n-power from n^0 to n^1; see Corollary E.5. □

Remark E.1 (the $\beta \to \infty$ discontinuity). For every finite $\beta > 0$, known-scale verification is possible at finite cost, and as $\beta \to \infty$ the achievable upper bound of Part B decreases to the n-free $O(\log(1/\delta))$ floor. Yet *at* $\beta = \infty$ the oracle degenerates to the noiseless comparison $\mathrm{sign}(f(x) - f(y))$, where verification is impossible at every query budget (Theorem G.3, proved in Appendix G). The limit is therefore singular: the supremum over finite β of the cost is the polylogarithmic floor, while the value at $\beta = \infty$ is $+\infty$. The blow-up of both endpoints is as $\beta \to 0$ (more noise), where $T = \Theta(1/\beta^2)$.

E.2 Proof of Theorem 4.2: the χ^2-mixture over the shifted-cone ensemble

Theorem 4.2 (Nonadaptive verification needs $\Omega(n)$)**.** *Let $n \geq 2$ and suppose $\beta\varepsilon^2/L_1 \leq 1/64$. Every nonadaptive algorithm that, under the known-scale BT oracle, tests $\|\nabla f(0)\| \leq \varepsilon$ against $\|\nabla f(0)\| \geq 2\varepsilon$ over $f \in \mathcal{F}(L_1)$ with success probability at least $2/3$ on every promise instance must make*

$$T \;\geq\; \frac{n-1}{54} \cdot \frac{L_1^2}{\beta^2\varepsilon^4} \quad \textit{queries.}$$

The proof mixes over the shifted cones of the search lower bound, with a smaller shift: the search construction uses $t = 3r_0$ to separate the wells of a packing, whereas the mixture below needs only $t \geq r_0$. We use $t = r_0$ and the following facts.

Fact E.2 (the shifted-cone ensemble; as in the proof of Theorem 3.2)**.** *Let $r_0 = 2\varepsilon/L_1$ and $t = r_0 = 2\varepsilon/L_1$, and let χ be the capped radial profile of the proof of Theorem 3.2: χ is 1-Lipschitz with $0 \leq \chi' \leq 1$ and $\chi(r) = r$ for $r \geq r_0$. For a unit vector $v \in \mathbb{S}^{n-1}$ set $m_v := t\,v$ and $f_v(x) := 2\varepsilon\,\chi\big(\|x - m_v\|\big)$. Then, for all unit v, v' and all x: (i) $f_v \in \mathcal{F}(L_1)$; (ii) $\|\nabla f_v(0)\| = 2\varepsilon\,\chi'(t) = 2\varepsilon$ (since $t \geq r_0$ and $\chi'(r_0) = 1$); (iii) $|f_v(x) - f_{v'}(x)| \leq 2\varepsilon\,\big|\|x - m_v\| - \|x - m_{v'}\|\big| \leq 2\varepsilon\,\|m_v - m_{v'}\| \leq 4\varepsilon t$.*

Unlike the search lower bound, no sphere packing is used: we mix over the *full* sphere, which dissolves both the diagonal atom and the packing large-deviation bookkeeping (Step 3). The argument is a χ^2-mixture (second-moment) computation: Steps 1–3 reduce testing to a χ^2 bound; Step 4a handles algorithmic randomization; Step 4 computes the exact one-bit correlation factor of the BT channel; Step 5 caps the per-query information at $O(1/n)$; Step 6 converts the exact — not linearized — per-query score into a recentered chaos functional; Step 7 closes via Lemma B.7 and a master threshold equation. Throughout, $\varrho := 8\beta\varepsilon t = 16\beta\varepsilon^2/L_1$ denotes the uniform logit-deviation bound; the displayed regime $\beta\varepsilon^2/L_1 \leq 1/64$ is exactly $\varrho \leq 1/4$ (equivalently $\beta\varepsilon t \leq 1/32$). Steps 1–7 use the shift only through Fact E.2 and through ϱ, so they hold for every $t \geq r_0$.

Proof. **Step 1 (the two hypotheses; continuous-sphere mixture).** The alternative is the mixture of $\{f_v\}$ with $v \sim \mathrm{Unif}(\mathbb{S}^{n-1})$: by Fact E.2(i)–(ii), every f_v lies in $\mathcal{F}(L_1)$ with $\|\nabla f_v(0)\| = 2\varepsilon$, a valid H_1 instance. The null is the single function $\bar{f} := \mathbb{E}_v[f_v]$:

- $\nabla\bar{f}(0) = 0$: $\bar{f}$ is rotation-invariant (the average of a rotation orbit of radial cones), hence radially symmetric about 0, hence its gradient at the center vanishes. (Alternatively, pair $v \leftrightarrow -v$: $\nabla f_v(0) + \nabla f_{-v}(0) = -2\varepsilon v + 2\varepsilon v = 0$.) So $\|\nabla\bar{f}(0)\| = 0 \leq \varepsilon$: a valid H_0 instance.
- $\bar{f} \in \mathcal{F}(L_1)$: $\nabla\bar{f} = \mathbb{E}_v[\nabla f_v]$ (differentiation under the expectation is justified by the uniform L_1-Lipschitz bound) is an average of L_1-Lipschitz maps and is therefore L_1-Lipschitz.

Step 2 (reduction: test $\Rightarrow \chi^2$ lower bound). Let $Q := P_{\bar{f}}$ and $\bar{P} := \mathbb{E}_v[P_{f_v}]$ be the T-query transcript laws. A $2/3$-correct verifier accepts under $\bar{f}$ with probability $\geq 2/3$ and rejects under each f_v with probability $\geq 2/3$, hence rejects under $\bar{P}$ with probability $\geq 2/3$; so $\mathrm{TV}(Q, \bar{P}) \geq 2/3 - 1/3 = 1/3$. Since $\mathrm{TV}(Q, \bar{P}) \leq \frac{1}{2}\sqrt{\chi^2(\bar{P}\,\|\,Q)}$ (Cauchy–Schwarz; see Tsybakov, 2009), a successful test forces

$$\chi^2(\bar{P}\,\|\,Q) \;\geq\; 4 \cdot (1/3)^2 \;=\; \frac{4}{9}.$$

It therefore suffices to show $\chi^2(\bar{P}\|Q) < 4/9$ whenever $T < c\,(n-1)\,L_1^2/(\beta^2\varepsilon^4)$.

Step 3 (Ingster second moment; the reference is $\bar{f}$). For $v, v' \sim \mathrm{Unif}(\mathbb{S}^{n-1})$ i.i.d.,

$$\chi^2(\bar{P} \,\|\, Q) + 1 = \mathbb{E}_Q\Big[\Big(\tfrac{d\bar{P}}{dQ}\Big)^2\Big] = \mathbb{E}_{v,v'}\big[\langle P_{f_v}, P_{f_{v'}}\rangle_Q\big],$$
$$\text{where} \quad \langle P_v, P_{v'}\rangle_Q := \mathbb{E}_Q\Big[\tfrac{dP_{f_v}}{dQ}\,\tfrac{dP_{f_{v'}}}{dQ}\Big],$$

by expanding $d\bar{P}/dQ = \mathbb{E}_v[dP_{f_v}/dQ]$ and using the independence of v and v'. There is no diagonal atom: $v = v'$ has probability zero under the continuous mixture — the term that would force a packing-cardinality bookkeeping is simply absent, which is why the full sphere is used.

Step 4a (randomized nonadaptive algorithms: coin conditioning). If the query sequence is drawn from a distribution over fixed sequences (independent of the answers — nonadaptive), include the coins ω in the transcript. The ω-marginals of Q and $\bar{P}$ agree (the coin law does not depend on the environment), so the joint divergence decomposes and controls the joint total variation:

$$\chi^2(\bar{P}_{\text{joint}} \,\|\, Q_{\text{joint}}) = \mathbb{E}_\omega\big[\chi^2(\bar{P}_\omega \,\|\, Q_\omega)\big], \qquad \mathrm{TV}(Q_{\text{joint}}, \bar{P}_{\text{joint}}) \le \tfrac{1}{2}\sqrt{\mathbb{E}_\omega\, \chi^2(\bar{P}_\omega \,\|\, Q_\omega)}.$$

The per-sequence bound of Steps 4–7 is uniform over fixed sequences, so the joint bound holds, and Step 2 applies verbatim to the joint transcript. In the remaining steps we therefore fix a deterministic sequence of query pairs $A_1, \dots, A_T$.

Step 4 (the exact per-query factor of the BT channel). Fix a query pair (x, y) and write $a_v = \sigma(\beta\,\mathrm{gap}_{f_v}(x,y))$, $a_{v'} = \sigma(\beta\,\mathrm{gap}_{f_{v'}}(x,y))$, and $\bar{a} = \sigma(\beta\,\mathrm{gap}_{\bar{f}}(x,y))$ for the reference. The one-bit χ^2 correlation factor is

$$\begin{aligned} G_{v,v'}(x,y) &:= \sum_{b\in\{0,1\}} \frac{P_v(b)\,P_{v'}(b)}{Q(b)} = \frac{a_v a_{v'}}{\bar{a}} + \frac{(1-a_v)(1-a_{v'})}{1-\bar{a}} \\ &= 1 + \frac{(a_v - \bar{a})(a_{v'} - \bar{a})}{\bar{a}(1-\bar{a})}. \end{aligned} \tag{E.16}$$

Note $\bar{a}(1-\bar{a}) = \sigma'(\beta\,\mathrm{gap}_{\bar{f}}) > 0$, and $\mathrm{sign}(a_v - \bar{a}) = \mathrm{sign}(\mathrm{gap}_{f_v} - \mathrm{gap}_{\bar{f}})$ since σ is increasing. For nonadaptive (fixed) queries the transcript laws are product measures over the independent bits, so

$$\langle P_v, P_{v'}\rangle_Q = \prod_{t=1}^{T} G_{v,v'}(A_t).$$

Step 5 (the per-query information cap — the load-bearing $1/n$). Define the per-query gap deviation $\Delta_v(x,y) := \mathrm{gap}_{f_v}(x,y) - \mathrm{gap}_{\bar{f}}(x,y)$. Two facts.

(i) *Uniformly small deviation.* $|\Delta_v(x,y)| = |(f_v - \bar{f})(x) - (f_v - \bar{f})(y)|$, and $|f_v - \bar{f}| = |\mathbb{E}_{v'}(f_v - f_{v'})| \le 4\varepsilon t$ by Fact E.2(iii); hence $|\Delta_v(x,y)| \le 8\varepsilon t$, and the reference-relative logit deviation obeys $|\beta\Delta_v| \le 8\beta\varepsilon t = \varrho \le 1/4$, uniformly over queries and v, in the displayed regime. (The Remark's wider regime is $\varrho \le 1/2$, i.e. $\beta\varepsilon^2/L_1 \le 1/32$.) All three logits $\beta\,\mathrm{gap}_{f_v}$, $\beta\,\mathrm{gap}_{f_{v'}}$, $\beta\,\mathrm{gap}_{\bar{f}}$ at a given query lie within ϱ of each other.

(ii) *Variance cap.* For fixed (x,y), the map $F(v) := \chi(\|x - tv\|) - \chi(\|y - tv\|)$ has spherical gradient $\|\nabla_{\mathbb{S}} F\| \le t\big(|\chi'(r_x)| + |\chi'(r_y)|\big) \le 2t$ (as $0 \le \chi' \le 1$), so the Poincaré inequality on $\mathbb{S}^{n-1}$ (spectral gap $n-1$; Ledoux, 1999, Section 7.3) gives $\mathrm{Var}_v[F] \le (2t)^2/(n-1)$, hence, as $\Delta_v = 2\varepsilon\, F$,

$$\mathrm{Var}_v\big[\Delta_v(x,y)\big] \;\le\; \frac{16\,\varepsilon^2 t^2}{n-1}, \qquad \mathbb{E}_v\big[\Delta_v(x,y)\big] = 0 \tag{E.17}$$

(the mean vanishes by the definition of $\bar{f}$). The same gradient bound gives the geodesic Lipschitz constant $\mathrm{Lip}(v \mapsto \Delta_v(x,y)) \le 4\varepsilon t$, used in Step 6.2.

Step 6 (the exact per-query score, recentered). Throughout, $\phi_t := \beta \operatorname{gap}_{\bar{f}}(A_t)$ and $w_t := \sigma'(\phi_t) = \bar{a}_t(1-\bar{a}_t) \in (0, 1/4]$, and by Step 5(i) $|\beta\Delta_v(A_t)| \le \varrho$ uniformly. Since $|\frac{d}{dz}\log\sigma'(z)| = |1-2\sigma(z)| \le 1$,

$$\max\Big(\sigma'(\phi_t+z),\ |\sigma''(\phi_t+z)|\Big) \ \le\ e^{\varrho}\, w_t \qquad \text{for all } |z| \le \varrho, \tag{E.18}$$

since $|\sigma''| = \sigma'\,|1-2\sigma| \le \sigma'$ pointwise.

Step 6.1 (exact score vector; no linearization). Define $\psi_t(v) := \sigma\big(\phi_t + \beta\Delta_v(A_t)\big) - \sigma(\phi_t)$ (which equals $a_v - \bar{a}$ at the query A_t) and $b_t(v) := \psi_t(v)/\sqrt{w_t}$. Then Step 4's factor is exactly $G_{v,v'}(A_t) = 1 + b_t(v)\,b_t(v')$ — write $\xi_t(v,v') := G_{v,v'}(A_t) - 1 = b_t(v)b_t(v')$ for the per-query correlation (Remark E.4 uses this symbol) — so by $1+z \le e^z$ (valid for all real z),

$$\langle P_v, P_{v'}\rangle_Q = \prod_{t=1}^{T}\big(1 + b_t(v)\,b_t(v')\big) \ \le\ \exp\big(\langle b(v), b(v')\rangle\big).$$

Step 6.2 (coordinate Lipschitz constants; the inflation factor is exactly e^{ϱ}). We have $b_t = \Psi_t \circ (v \mapsto \Delta_v(A_t))$ with $\Psi_t(z) := \big(\sigma(\phi_t+\beta z) - \sigma(\phi_t)\big)/\sqrt{w_t}$; on the range $|z| \le 8\varepsilon t$ of Δ_v (Step 5(i)) we have $|\beta z| \le \varrho$, so by (E.18) $\sup|\Psi_t'| \le e^{\varrho}\beta\sqrt{w_t}$. With $\mathrm{Lip}(v \mapsto \Delta_v(A_t)) \le 4\varepsilon t$ (Step 5(ii)), the composition bound gives

$$L_{b,t} := \mathrm{Lip}(b_t) \le 4\, e^{\varrho}\beta\sqrt{w_t}\,\varepsilon t,$$
$$L_b^2 := \sum_{t=1}^{T} L_{b,t}^2 \le 16\, e^{2\varrho}\beta^2\varepsilon^2 t^2 \sum_t w_t \le 4\, e^{2\varrho}\beta^2\varepsilon^2 t^2\, T,$$

using $w_t \le 1/4$. Hence, with

$$\kappa := \frac{4\beta^2\varepsilon^2 t^2\, T}{n-1}, \qquad K := e^{2\varrho}\kappa, \qquad \text{we get} \qquad \frac{L_b^2}{n-1} \le K.$$

(The composition route pins the inflation over the linearized score at exactly e^{ϱ} per coordinate, $e^{2\varrho}$ on L_b^2 and on the trace; no decomposition of the $(1+\xi)$ factor is needed.)

Step 6.3 (second moments; the trace). By (E.18) and the mean value theorem, $|\psi_t| \le e^{\varrho} w_t\,\beta|\Delta_v|$, so $\mathbb{E}_v[b_t^2] \le e^{2\varrho} w_t\,\beta^2\,\mathbb{E}_v[\Delta_v(A_t)^2] \le e^{2\varrho} w_t\,\beta^2 \cdot 16\varepsilon^2 t^2/(n-1)$ by (E.17) (recall $\mathbb{E}_v\Delta_v = 0$). Hence, with $\Sigma_g := \mathrm{Cov}_v(b)$,

$$\mathrm{tr}(\Sigma_g) \le \sum_t \mathbb{E}_v[b_t^2] \le \frac{4\, e^{2\varrho}\beta^2\varepsilon^2 t^2\, T}{n-1} = K.$$

Step 6.4 (the mean term is second-order small, because $\mathbb{E}_v\Delta_v = 0$). Let $m_t := \mathbb{E}_v[b_t(v)]$. Taylor's formula with integral remainder about ϕ_t gives

$$\psi_t = w_t\,\beta\Delta_v + \int_0^{\beta\Delta_v} \big(\beta\Delta_v - u\big)\,\sigma''(\phi_t + u)\,du;$$

the first-order term has $\mathbb{E}_v = 0$ by (E.17), and by (E.18) the remainder is at most $e^{\varrho} w_t(\beta\Delta_v)^2/2$ pointwise, so

$$|m_t| \le \frac{e^{\varrho}\sqrt{w_t}\,\beta^2\,\mathbb{E}_v[\Delta_v^2]}{2} \le \frac{8\, e^{\varrho}\sqrt{w_t}\,\beta^2\varepsilon^2 t^2}{n-1},$$
$$\|m\|^2 \le \frac{64\, e^{2\varrho}\beta^4\varepsilon^4 t^4\,(T/4)}{(n-1)^2} = \frac{\varrho^2 K}{16(n-1)},$$

using $\sum_t w_t \le T/4$, then $\beta\varepsilon t = \varrho/8$ and the definition of κ.

Step 6.5 (recentering and the Hölder split). Write $b(v) = m + g(v)$: the recentered g is mean-zero by construction (the recentering restores exactly what the sigmoid curvature destroyed), $\mathrm{Lip}(g_t) = L_{b,t}$ (constant shifts do not change Lipschitz constants), and $\mathrm{Cov}(g) = \Sigma_g$. Then

$$\langle b(v), b(v')\rangle = \langle g(v), g(v')\rangle + \langle m, g(v)\rangle + \langle m, g(v')\rangle + \|m\|^2,$$

and for any $p > 1$ with $q := p/(p-1)$, Hölder's inequality (exponents p, q) plus the independence of v, v' give

$$\mathbb{E}_{v,v'}\, e^{\langle b,b'\rangle} \;\le\; e^{\|m\|^2}\,\big(\mathbb{E}_{v,v'}\, e^{p\langle g,g'\rangle}\big)^{1/p}\,\big(\mathbb{E}_v\, e^{q\langle m,g(v)\rangle}\big)^{2/q}. \tag{E.19}$$

The linear functional $v \mapsto \langle m, g(v)\rangle$ is mean-zero with

$$\mathrm{Lip}\big(\langle m,g\rangle\big) \le \sum_t |m_t|\, L_{b,t} \le \frac{32\, e^{2\varrho}\beta^3\varepsilon^3 t^3\,(T/4)}{n-1} = \frac{\varrho K}{4}$$

(Step 6.2 × 6.4 summed with $\sum_t w_t \le T/4$; the Cauchy–Schwarz route $\|m\|\, L_b$ gives the same value).

Step 7 (the master threshold: Lemma B.7 on the centered part, sphere sub-Gaussianity on the mean part). Two ingredients, then arithmetic.

- *Centered chaos, via Lemma B.7 (Part B), used at statement level.* Its hypotheses hold for g by Steps 6.2–6.5: g is measurable and mean-zero, coordinatewise geodesically $L_{b,t}$-Lipschitz, and $n \ge 2$. Instantiating the lemma's scalar parameter as p and taking the $1/p$-th power,

$$\big(\mathbb{E}\, e^{p\langle g,g'\rangle}\big)^{1/p} \le (1-s_2^2)^{-1/(2p)}\,\exp\Big(\frac{s_1 s_2/2}{p\,(1-s_2^2)}\Big),$$

$$s_1 := p\ \mathrm{tr}(\Sigma_g) \le pK, \qquad s_2 := \frac{p\,L_b^2}{n-1} \le pK,$$

  valid for $s_2 < 1$. The right side is increasing in (s_1, s_2) (termwise by inspection), so both may be replaced by pK whenever $pK < 1$:

$$\big(\mathbb{E}\, e^{p\langle g,g'\rangle}\big)^{1/p} \le (1-p^2K^2)^{-1/(2p)}\ \exp\Big(\frac{pK^2}{2(1-p^2K^2)}\Big). \tag{E.20}$$

- *The mean part, via fact (L): sub-Gaussianity of Lipschitz functions on the sphere.* For a geodesically λ-Lipschitz $F\colon \mathbb{S}^{n-1} \to \mathbb{R}$ with $\mathbb{E}F = 0$,

$$\mathbb{E}\, e^{uF} \le \exp\Big(\frac{u^2\lambda^2}{2(n-1)}\Big) \qquad \text{for all } u \in \mathbb{R}. \tag{E.21}$$

  This is the classical consequence of the sharp logarithmic Sobolev constant $n-1$ on $\mathbb{S}^{n-1}$ (Mueller & Weissler, 1982) together with the Herbst argument (Ledoux, 1999, Section 2.3, inequality (2.19)) — the same spectral-gap family as Step 5's Poincaré inequality; we import it without reproof. Applied to $F = \langle m, g\rangle$ with $u = q$ and $\lambda \le \varrho K/4$ (Step 6.5),

$$\big(\mathbb{E}_v\, e^{q\langle m,g\rangle}\big)^{2/q} \le \exp\Big(\frac{2}{q}\cdot\frac{q^2(\varrho K/4)^2}{2(n-1)}\Big) = \exp\Big(\frac{q\,\varrho^2K^2}{16(n-1)}\Big).$$

With $e^{\|m\|^2} \le \exp\big(\varrho^2 K/(16(n-1))\big)$ (Step 6.4), the split (E.19) yields the *master bound*: for every $p \in (1,\, 1/K)$,

$$\chi^2(\bar P\,\|\,Q) + 1 \;\le\; \mathbb{E}_{v,v'}\, e^{\langle b,b'\rangle} \;\le\; B_{\varrho,p,n}(K) := (1-p^2K^2)^{-1/(2p)} \times \exp\Big(\frac{pK^2}{2(1-p^2K^2)} + \frac{\varrho^2 K}{16(n-1)}\big(1+qK\big)\Big). \tag{E.22}$$

B is strictly increasing in K on $(0, 1/p)$ (termwise by inspection), $B(0) = 1$, $B \to \infty$ as $K \to 1/p$, and B is decreasing in n. Fix $p = 11/10$ (so $q = 11$; the resulting root is within 2% of the best root over p at both displayed values of ϱ) and worst-case the mean terms at $n = 2$. Let $K^*(\varrho)$ be the unique root of $B_{\varrho,\,11/10,\,2}(K) = 13/9$:

$$K^*(1/4) = 0.497001, \qquad K^*(1/2) = 0.478108 \tag{E.23}$$

Now suppose

$$T \;<\; \frac{\kappa^*(\varrho)\,(n-1)}{4\beta^2\varepsilon^2 t^2}, \qquad \kappa^*(\varrho) := e^{-2\varrho}K^*(\varrho),$$

$$\kappa^*(1/4) = 0.301446, \qquad \kappa^*(1/2) = 0.175886.$$

Then $\kappa < \kappa^*(\varrho)$, so $K = e^{2\varrho}\kappa < K^*(\varrho)$ and $pK < pK^* \le 0.547 < 1$: the master bound applies with every hypothesis strict — in particular Lemma B.7's validity window $s_2 < 1$ *never* binds — and by monotonicity $B(K) < B(K^*) = 13/9$, i.e. $\chi^2(\bar{P}\|Q) < 4/9$, contradicting Step 2. Hence every nonadaptive $2/3$-correct verifier makes

$$T \;\ge\; \frac{\kappa^*(\varrho)\,(n-1)}{4\beta^2\varepsilon^2 t^2} \;=\; \kappa^*(\varrho)\,\frac{(n-1)\,L_1^2}{16\,\beta^2\varepsilon^4} \tag{E.24}$$

queries, using $t = 2\varepsilon/L_1$, i.e. $4\beta^2\varepsilon^2 t^2 = 16\,\beta^2\varepsilon^4/L_1^2$. At the displayed regime $\varrho \le 1/4$ this is the claimed bound with

$$c \;=\; \frac{\kappa^*(1/4)}{16} \;=\; 0.018840\ldots \;\ge\; \frac{1}{54}.$$

The binding condition is $\kappa \le \kappa^*(\varrho)$ is the T-side constraint; the parameter-regime hypothesis $\varrho \le 1/4$ (used in Steps 5–7: the uniform logit window (E.18), the e^{ϱ} inflation, the master threshold) is the only hypothesis on $(\beta, \varepsilon, L_1)$. The distinction between the two parts of the chaos lemma matters: the simpler form $\mathbb{E}\exp(\beta^2\langle a, a'\rangle) \le \exp(C\beta^4\,\mathrm{tr}(\Sigma^2))$ that an earlier version of this argument displayed is *false* for general Lipschitz maps (a bump witness drives $\mathrm{tr}(\Sigma^2) \to 0$ with the moment generating function $\to \infty$; such witnesses live outside $s_2 < 1$); Part B of Lemma B.7 is the applicable bound and absorbs into the same κ bookkeeping. □

Remark E.3 (wider regime, smaller constant; limits of the master equation). In the wider regime $\beta\varepsilon^2/L_1 \le 1/32$ (i.e. $\varrho \le 1/2$) the same proof runs verbatim with $K^*(1/2) = 0.478108$ and $\kappa^*(1/2) = 0.175886$, giving $T \ge c\,(n-1)L_1^2/(\beta^2\varepsilon^4)$ with $c = \kappa^*(1/2)/16 = 0.010993 \ge 1/91$. Two consistency limits of the master equation (E.22): as $\varrho \to 0$ (with $p \to 1$) the root tends to the linearized-chain value 0.534238 and the constant improves to $0.534238/16 \ge 1/30$; as $n \to \infty$ (the mean terms vanish, $p \to 1$) it improves to $\ge 1/50$ at $\varrho = 1/4$, resp. $\ge 1/82$ at $\varrho = 1/2$ — the latter is exactly the corollary estimate of Lemma B.7, identified here as the large-n limit of this chain.

Remark E.4 (why adaptivity resists; scope of the theorem). Theorem 4.2 is a statement about *nonadaptive* verifiers only, and the obstruction to an adaptive extension is localized precisely. For an adaptive algorithm the query $A_t = A_t(B_1, \ldots, B_{t-1})$ is a function of the past bits; in $\langle P_v, P_{v'}\rangle_Q$ the transcript is generated under Q, so A_t is predictable and the tower rule yields only the worst-case recursion $\langle P_v, P_{v'}\rangle_Q \le \big(1 + \sup_{x,y}|\xi(v, v'; x, y)|\big)^T$, with ξ the per-query correlation factor of Step 6.1. Taking the per-query supremum *before* the (v, v')-average discards the sign of $\Delta_v(A_t)\Delta_{v'}(A_t)$: nonadaptively, the query direction d_t is independent of (v, v'), so $\mathbb{E}_{v,v'}[\langle d_t, v\rangle\langle d_t, v'\rangle] = \langle d_t, \mathbb{E}v\rangle\langle d_t, \mathbb{E}v'\rangle = 0$ — the cancellation that powers Step 7 — while adaptively d_t is correlated, through the bits, with the likelihood tilt toward the true (v, v'), the cross term is no longer mean-zero, and the exponential-of-supremum bound is dominated by the aligned-chase pairs: exactly the two-point regime, which is n-oblivious. An adaptive $\Omega(n)$ bound would need a likelihood-ratio supermartingale that keeps the (v, v')-average *inside* the per-step bound (for instance an information-per-query bound on $I(V; B_t \mid \text{past})$); we register this as an open problem. The proven dimension factors remain $[1, n^2]$ for adaptive verification and $[n, n^2]$ for nonadaptive verification in the stated small-logit regime.

E.3 The combined nonadaptive lower endpoint

Corollary E.5 (combined nonadaptive lower endpoint)**.** *Let $n \ge 1$, and consider the promise problem of Theorem 4.1 in the regime $\beta\varepsilon^2/L_1 \le 1/64$, restricted to* nonadaptive *(possibly randomized) verifiers. With $c_0 = 32\ln(9/8)$ and $c = \kappa^*(1/4)/16 \ge 1/54$:*

1. *every nonadaptive verifier correct with probability $\ge 2/3$ on every instance satisfies*

$$T \;\ge\; \max\big(c_0,\; c\,(n-1)\big)\,\frac{L_1^2}{\beta^2\varepsilon^4}\,;$$

2. *every nonadaptive verifier with failure probability at most $\delta \le 1/16$ on every instance satisfies*

$$T \;\ge\; \max\Big(16\log\tfrac{1}{\delta},\; c\,(n-1)\Big)\,\frac{L_1^2}{\beta^2\varepsilon^4}\,.$$

Proof. The first terms are Step V10 of the proof of Theorem 4.1, (E.13) and (E.14): they hold for *all* algorithms, adaptive included, and at every $\beta > 0$; the regime hypothesis is used only by the second term. The second term is Theorem 4.2, proved at success probability $2/3$ (for $n = 1$ it is void); a verifier with failure probability $\delta \le 1/3$ is in particular $2/3$-correct, so the term transfers to every $\delta \le 1/3$ — and in particular to $\delta \le 1/16$ — by monotonicity of the query complexity in δ. Taking the maximum of the valid bounds gives both displays. The two δ-scopings are inherited from Theorem 4.1 and must not be blurred: the $\log(1/\delta)$ factor on the first term is proved for $\delta \le 1/16$ (E.14); at larger δ — for example at the $2/3$-success point — the operative first term is the log-free c_0 (E.13), which is the form the crossover below uses. □

Remark E.6 (crossover; the max form). With $c_0 = 3.76906\ldots$ and $c = 0.0188404\ldots$, the $(n-1)$ term dominates the constant-confidence term precisely for $n \ge 202$; at $n = 201$ it misses by about 0.03 percent ($200c = 3.76808 < c_0$). At confidence $1-\delta$, $\delta \le 1/16$, it dominates once $n - 1 \ge 850 \ln(1/\delta)$. At $n = 1$ the maximum is c_0: the two-point floor, consistent with the $(n-1)$ factor vanishing. Equivalently, the constant-confidence bound of Corollary E.5 may be written $T \ge c \max(n-1,\, 1)\, L_1^2/(\beta^2\varepsilon^4)$, at the price of the smaller constant c in place of c_0 for small n. For *adaptive* verifiers no $(n-1)$ term is claimed (Remark E.4): the adaptive lower endpoint remains (E.13)–(E.15).

F Unknown scale and the ordinal task: proofs

This part of the appendix proves the unknown-scale results (Theorem 5.1 and Proposition F.3) and the ordinal results (Theorem 6.1 and Corollary F.7). Throughout, log denotes the natural logarithm. A convention used across this section: when the slope $\beta \in [\beta_{\min}, \beta_{\max}]$ is unknown, we write

$$\phi := \beta f$$

for the *observable potential*. The logit of the oracle bit on a pair (x, y) is $\beta\, \mathrm{gap}_f(x, y) = \mathrm{gap}_\phi(x, y)$, so the oracle exposes ϕ but never the factorization into (β, f) separately; both proofs below turn on this single fact, in opposite directions.

F.1 Proof of the unknown-scale dichotomy

Theorem 5.1 (Unknown-scale dichotomy). *Let $c' > 1$, let $0 < \beta_{\min} \le \beta_{\max} < \infty$ with $\rho = \beta_{\max}/\beta_{\min}$, and let $\varepsilon, L_1, \Delta > 0$ satisfy Assumption 2.3. Consider ε-vs-$c'\varepsilon$ gradient-norm verification at a given $\bar{x}$ over $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \le \Delta\}$ under the BT oracle with unknown scale $\beta \in [\beta_{\min}, \beta_{\max}]$. Then the task is possible if and only if $\rho < c'$. More precisely:*

- *(i)* (Impossibility, $\rho \ge c'$.) *If $\rho \ge c'$, no algorithm — at any budget, even unbounded and data-dependent — answers correctly with probability greater than $1/2$ on both hypotheses. In particular the boundary $\rho = c'$ is on the impossible side.*
- *(ii)* (Possibility, $\rho < c'$.) *If $\rho < c'$, for every $\delta \in (0, 1)$ there is a verifier with failure probability at most δ, uniform over the class and over $\beta \in [\beta_{\min}, \beta_{\max}]$, of finite query cost (Proposition F.3 gives an explicit bound).*

Proof. Write $g := \nabla f(\bar{x})$, and $G := \|\nabla\phi(\bar{x})\| = \beta\|g\|$ for the gradient norm of the observable potential. Recall $\rho = \beta_{\max}/\beta_{\min}$ and $c' > 1$, and that (A1) reads $c'\varepsilon \le \sqrt{2L_1\Delta}$. The two directions are proved separately and symmetrically: impossibility for $\rho \ge c'$ by exhibiting two admissible instances with the *same* observable potential up to the oracle's channel, and possibility for $\rho < c'$ by an explicit finite-cost estimator of G.

Part I (impossibility, $\rho \ge c'$).

Step I.1 (the witness pair). Let e_1 be a unit vector and set $t(x) := \langle e_1, x - \bar{x}\rangle$, a 1-Lipschitz map with $t(\bar{x}) = 0$ and range all of $\mathbb{R}$. Define the quadratic ridge

$$f_0(x) := h\big(t(x)\big), \qquad h(\tau) := c'\varepsilon\,\tau + \frac{L_1}{2}\,\tau^2,$$

and the two instances

$$\text{H1-instance: } (\beta_a, f_0) \text{ with } \beta_a \coloneqq \beta_{\min}, \qquad \text{H0-instance: } (\beta_b, f_0/c') \text{ with } \beta_b \coloneqq c'\beta_{\min} = c'\beta_a.$$

Step I.2 (f_0 is admissible and sits on the reject boundary). First, $\nabla f_0(x) = h'(t(x))\, e_1 = (c'\varepsilon + L_1 t(x))\, e_1$, so $\nabla f_0(\bar{x}) = c'\varepsilon\, e_1$ and $\|\nabla f_0(\bar{x})\| = c'\varepsilon$: an H1 (reject) boundary point. Second, for any x, y, $\nabla f_0(x) - \nabla f_0(y) = L_1\langle e_1, x - y\rangle\, e_1$, whose norm is at most $L_1\|x - y\|$ with equality along e_1; hence $f_0 \in \mathcal{F}(L_1)$ with the Lipschitz constant attained. Third, h is a convex parabola minimized at $\tau^* = -c'\varepsilon/L_1$ with $h(\tau^*) = -(c'\varepsilon)^2/(2L_1)$, and t is onto $\mathbb{R}$, so

$$f_0(\bar{x}) - \inf f_0 \;=\; 0 - h(\tau^*) \;=\; \frac{(c'\varepsilon)^2}{2L_1} \;\le\; \Delta,$$

where the last inequality is exactly (A1). Assumption (A1) is the sole securer of this budget membership; it is load-bearing.

Step I.3 (f_0/c' is admissible and sits on the accept boundary; $c' > 1$ helps). $\nabla(f_0/c')(\bar{x}) = \varepsilon\, e_1$, so $\|\nabla(f_0/c')(\bar{x})\| = \varepsilon$: an H0 (accept) boundary point. Its gradient is (L_1/c')-Lipschitz with $L_1/c' \le L_1$, so $f_0/c' \in \mathcal{F}(L_1)$; and its gap is $(1/c')\big(f_0(\bar{x}) - \inf f_0\big) \le \Delta$. Both memberships use only $c' > 1$.

Step I.4 (both slopes lie in the range iff $\rho \ge c'$; the boundary case). $\beta_a = \beta_{\min} \in [\beta_{\min}, \beta_{\max}]$ always. For $\beta_b = c'\beta_{\min}$: the lower endpoint holds since $c' > 1$, and $c'\beta_{\min} \le \beta_{\max} \iff c' \le \rho$. So for $\rho \ge c'$ *both* instances lie in the model class $\{(\beta, u) : \beta \in [\beta_{\min}, \beta_{\max}],\ u \in \mathcal{F}(L_1),\ u(\bar{x}) - \inf u \le \Delta\}$. At the boundary $\rho = c'$ we get $\beta_b = \beta_{\max}$, admissible in the closed range: the boundary is on the impossible side.

Step I.5 (identical answer laws on every pair, via Lemma B.5). For every query pair (x, y),

$$\beta_b\ \mathrm{gap}_{f_0/c'}(x, y) = c'\beta_a \cdot \frac{1}{c'}\big(f_0(x) - f_0(y)\big) = \beta_a\ \mathrm{gap}_{f_0}(x, y)\ :$$

the two instances share the same logit potential ($\beta_a f_0 = \beta_b(f_0/c')$ identically), hence $\sigma(\beta_a\, \mathrm{gap}_{f_0}(x, y)) = \sigma(\beta_b\, \mathrm{gap}_{f_0/c'}(x, y))$ on every pair — identical answer channels. By Lemma B.5, every randomized adaptive algorithm, at *any* query budget, has identical transcript and output laws under the H1-instance and the H0-instance.

Step I.6 (no valid test). The H0-instance is an accept instance ($\|\nabla(f_0/c')(\bar{x})\| = \varepsilon \le \varepsilon$) and the H1-instance a reject instance ($\|\nabla f_0(\bar{x})\| = c'\varepsilon \ge c'\varepsilon$). Let p be the probability that a given algorithm outputs "accept"; by Step I.5, p is the same under both instances. Correctness with probability $\ge 2/3$ on both would force $p \ge 2/3$ and $1 - p \ge 2/3$, a contradiction; indeed even success probability $> 1/2$ on both is impossible. Hence for $\rho \ge c'$ no algorithm solves the ε-vs-$c'\varepsilon$ verification problem (Definition 2.2) at any budget. *(End of Part I.)*

Part II (possibility, $\rho < c'$). All tester parameters below are computable from the known quantities $(n, L_1, \Delta, \varepsilon, c', \delta, \beta_{\min}, \beta_{\max})$; the true β enters only the analysis.

Step II.1 (a separating observable). The quantity $G = \beta\|g\| = \|\nabla\phi(\bar{x})\|$ is a functional of the observable potential alone. Under H0 ($\|g\| \le \varepsilon$), $G \le \beta_{\max}\varepsilon$; under H1 ($\|g\| \ge c'\varepsilon$), $G \ge \beta_{\min}c'\varepsilon$. The intervals $[0, \beta_{\max}\varepsilon]$ and $[\beta_{\min}c'\varepsilon, \infty)$ are disjoint iff $\beta_{\max}\varepsilon < \beta_{\min}c'\varepsilon$, i.e. iff $\rho < c'$. On the possible side set the threshold and half-margin

$$\tau_G \coloneqq \frac{\beta_{\max}\varepsilon + \beta_{\min}c'\varepsilon}{2}, \qquad M \coloneqq \frac{\beta_{\min}c'\varepsilon - \beta_{\max}\varepsilon}{2} = \frac{\beta_{\min}\varepsilon\,(c' - \rho)}{2} > 0,$$

so that $\tau_G - M = \beta_{\max}\varepsilon$ and $\tau_G + M = \beta_{\min}c'\varepsilon$ exactly: H0 forces $G \le \tau_G - M$ and H1 forces $G \ge \tau_G + M$. It remains to estimate G to additive accuracy $< M$ with confidence $1 - \delta$ at finite query cost.

Step II.2 (a priori magnitude bounds). Since f is L_1-smooth, bounded below, and $f(\bar{x}) - \inf f \le \Delta$, the descent lemma (take $y = x - \nabla f(x)/L_1$ in $f(y) \le f(x) + \langle\nabla f(x), y - x\rangle + \frac{L_1}{2}\|y - x\|^2$) gives $\|\nabla f(x)\|^2 \le 2L_1(f(x) - \inf f)$; at $x = \bar{x}$, $\|g\| \le \sqrt{2L_1\Delta}$ and hence

$$G \;\le\; G_{\max} \coloneqq \beta_{\max}\sqrt{2L_1\Delta},$$

a known bound. Likewise $\nabla\phi = \beta\,\nabla f$ is (βL_1)-Lipschitz with $\beta L_1 \le \Lambda := \beta_{\max} L_1$, a known effective smoothness.

Step II.3 (coordinate central-difference probes; an exact second-moment identity). Fix a probe radius $t > 0$ (chosen in Step II.5). For each coordinate $i \in \{1, \dots, n\}$ query the pair $(\bar{x} + te_i,\ \bar{x} - te_i)$; its logit is

$$L_i := \phi(\bar{x} + te_i) - \phi(\bar{x} - te_i) = 2t\, s_i + R_i, \qquad |R_i| \le (\beta L_1)\, t^2 \le \Lambda t^2, \tag{F.1}$$

where $s_i := (\nabla\phi(\bar{x}))_i$; this is second-order Taylor expansion of ϕ at $\bar{x} \pm te_i$ with the standard integral remainder $|\phi(\bar{x} \pm te_i) - \phi(\bar{x}) \mp t s_i| \le (\beta L_1/2)t^2$ for a (βL_1)-Lipschitz gradient, subtracted. Aggregating over *all* n coordinate directions, the trivial Parseval identity $G^2 = \|\nabla\phi(\bar{x})\|^2 = \sum_{i=1}^n s_i^2$ is exact: this deterministic sweep carries zero direction-sampling error.

Step II.4 (per-coordinate logit estimation, with clipping). Query pair i for m i.i.d. repetitions and let $\hat{p}_i$ be the empirical bit frequency, unbiased for $p_i := \sigma(L_i)$. Since $|s_i| \le G \le G_{\max}$, Equation (F.1) gives $|L_i| \le B := 2t\,G_{\max} + \Lambda t^2$, a known bound, so $p_i \in [\sigma(-B), \sigma(B)]$. Clip $\hat{p}_i$ to this interval (projection onto an interval containing p_i cannot increase the error) and invert: $\hat{L}_i := \mathrm{logit}(\mathrm{clip}(\hat{p}_i))$, where $\mathrm{logit}(q) = \log \frac{q}{1-q}$, and $\hat{s}_i := \hat{L}_i/(2t)$. On $[\sigma(-B), \sigma(B)]$ the logit has derivative $1/(q(1-q)) \le 1/\kappa_B$ with

$$\kappa_B := \sigma(B)\big(1 - \sigma(B)\big) \ \ge\ \frac{e^{-B}}{4},$$

so $|\hat{L}_i - L_i| \le |\hat{p}_i - p_i|/\kappa_B$. By Hoeffding's inequality, $\mathbb{P}\big[|\hat{p}_i - p_i| > \sqrt{\log(2/\delta_1)/(2m)}\big] \le \delta_1$, and on the complementary event, using Equation (F.1),

$$|\hat{s}_i - s_i| \ \le\ \frac{|\hat{L}_i - L_i|}{2t} + \frac{|R_i|}{2t} \ \le\ \frac{1}{2t\kappa_B}\sqrt{\frac{\log(2/\delta_1)}{2m}} \ +\ \frac{\Lambda t}{2}, \tag{F.2}$$

the second term being the deterministic curvature bias.

Step II.5 (choice of t and m; aggregation). Set $\hat{s} := (\hat{s}_1, \dots, \hat{s}_n)$ and $\hat{G} := \|\hat{s}\|_2$. By the reverse triangle inequality, $|\hat{G} - G| \le \|\hat{s} - \nabla\phi(\bar{x})\|_2$, so it suffices to force $|\hat{s}_i - s_i| \le M/(2\sqrt{n})$ for every i; then $\|\hat{s} - \nabla\phi(\bar{x})\| \le M/2$ and $|\hat{G} - G| \le M/2 < M$. Split Equation (F.2) into two halves of $M/(4\sqrt{n})$ each:

$$\text{(bias)} \quad \frac{\Lambda t}{2} \le \frac{M}{4\sqrt{n}} \iff t := \frac{M}{2\Lambda\sqrt{n}};$$

$$\text{(noise)} \quad \frac{1}{2t\kappa_B}\sqrt{\frac{\log(2/\delta_1)}{2m}} \le \frac{M}{4\sqrt{n}} \iff m \ge \frac{8\log(2/\delta_1)\,\Lambda^2 n^2}{\kappa_B^2 M^4}.$$

Take $\delta_1 := \delta/n$ and union-bound over the n coordinates: with probability $\ge 1 - \delta$ all coordinates obey their bound, so $|\hat{G} - G| \le M/2$, and by Step II.1 the test "accept iff $\hat{G} < \tau_G$" is correct on both hypotheses with probability $\ge 1 - \delta$.

Step II.6 (the explicit query cost, displayed). With $m = \big\lceil 8\log(2n/\delta)\,\Lambda^2 n^2/(\kappa_B^2 M^4)\big\rceil$ reads per coordinate, the total count $N = nm$ obeys

$$N \ \le\ \frac{8\,n^3\Lambda^2\log(2n/\delta)}{\kappa_B^2\,M^4} + n \ =\ \frac{128\,n^3\,\beta_{\max}^2 L_1^2\,\log(2n/\delta)}{\kappa_B^2\,\beta_{\min}^4\,\varepsilon^4\,(c' - \rho)^4} + n, \tag{F.3}$$

with $\kappa_B = \sigma(B)(1 - \sigma(B)) \ge e^{-B}/4$, $B = 2tG_{\max} + \Lambda t^2$, and $t = M/(2\Lambda\sqrt{n})$; note $B \to 0$ and $\kappa_B \to 1/4$ as n grows or the margin shrinks, and $1/\kappa_B^2 \le 16\,e^{2B}$ always. In the regime $B \le 1$ — equivalently, up to the constant chain above, $\beta_{\min}(c' - \rho)\,\varepsilon\sqrt{2\Delta/L_1} \le \sqrt{n}$, since the leading term of B is $2tG_{\max} = \beta_{\min}(c' - \rho)\varepsilon\sqrt{2\Delta/L_1}/(2\sqrt{n})$ — the factor $1/\kappa_B^2 \le 16e^2$ is a constant and Equation (F.3) reads

$$N = O\left(\frac{n^3\,\beta_{\max}^2 L_1^2\,\log(n/\delta)}{\beta_{\min}^4\,\varepsilon^4\,(c' - \rho)^4}\right).$$

Outside that regime the exponential factor is not negligible and must be kept: the bound actually proved is $N = O\big(e^{2B}\,n^3\beta_{\max}^2 L_1^2\log(n/\delta)\,/\,(\beta_{\min}^4\varepsilon^4(c' - \rho)^4)\big)$ with B as above. In every case N

is finite for each $\rho < c'$ and each finite range $[\beta_{\min}, \beta_{\max}]$ — which proves the possible side — and diverges through $(c' - \rho)^{-4}$ continuously as $\rho \to c'^{-}$, matching the impossible boundary $\rho = c'$. The bound is deliberately crude; the dichotomy only needs finite versus infinite cost. Proposition F.3 removes the exponential factor, and Theorem 5.2 shows that the exponent 4 of $(c' - \rho)^{-1}$ is necessary at fixed $\beta_{\min}$.

Conclusion. Combining Parts I and II: under (A1), ε-vs-$c'\varepsilon$ gradient-norm verification at $\bar{x}$ over $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \le \Delta\}$ with unknown $\beta \in [\beta_{\min}, \beta_{\max}]$ is impossible at any budget if $\rho \ge c'$, and possible at the finite explicit cost Equation (F.3) if $\rho < c'$. □

Remark F.1 (Exact-budget asymmetry: the inequality class is load-bearing). The impossibility argument of Part I would *break* if the class fixed $f(\bar{x}) - \inf f = \Delta$ exactly rather than $\le \Delta$. The two witnesses have different gap values, $f_0(\bar{x}) - \inf f_0 = (c'\varepsilon)^2/(2L_1)$ versus $(c'\varepsilon)^2/(2c'L_1)$ — a factor c' — and under Bradley–Terry noise the logit $\beta \cdot \text{gap}$ depends on the gap *magnitude*: at unbounded budget the recoverable potential $\psi := \beta(f - f(\bar{x}))$ has $\inf \psi = -\beta\Delta$, pinning $\beta = -\inf(\psi)/\Delta$ and making the scale identifiable. (A complete β-identifying strategy for the exact-budget class is a separate statement, which we record but do not carry out.) This is the opposite of the noiseless situation, where the sign-only oracle is scale-free and cannot use magnitudes even under an exact budget (the scale-freeness is Lemma G.2(c), the impossibility Theorem G.3, both proved in Appendix G); the inequality form of the budget class is load-bearing here specifically.

Remark F.2 (Scale uncertainty and the noiseless obstruction). For every fixed feasible gap ratio $c' > 1$, sufficiently large ρ makes verification impossible by Theorem 5.1. The noiseless oracle has a related scale-invariance obstruction, proved separately in Theorem G.3. This is a comparison of identifiability mechanisms, not a convergence theorem for the two statistical experiments. The BT model identifies βf up to an additive constant, while the noiseless oracle preserves every strictly increasing reparametrization.

F.2 Proof of the bounded-cost proposition

Proposition F.3 (Bounded unknown-scale cost at saturation)**.** *In the setting of Theorem 5.1 with $\rho < c'$, let $\delta \in (0,1)$ and set*

$$M := \frac{\beta_{\min}\varepsilon(c' - \rho)}{2}, \quad G_{\max} := \beta_{\max}\sqrt{2L_1\Delta}, \quad \Lambda := \beta_{\max}L_1, \quad t := \min\Big\{\frac{M}{2\Lambda\sqrt{n}}, \frac{1}{2G_{\max}}\Big\},$$

and $\beta^\dagger := \sqrt{2nL_1/\Delta}\,/\,(\varepsilon(c' - \rho))$. There is an explicit tester — the $\beta_{\max}$-tuned coordinate central-difference threshold test sketched below, all of whose parameters are computable from $(n, L_1, \Delta, \varepsilon, c', \delta, \beta_{\min}, \beta_{\max})$ — that is correct with probability at least $1 - \delta$ on both hypotheses, uniformly over the class and over the true $\beta \in [\beta_{\min}, \beta_{\max}]$, at total query cost

$$N \;\le\; 1872\,\frac{n^2\rho^2\,L_1\Delta\,\log(2n/\delta)}{\varepsilon^2\,(c' - \rho)^2} \cdot \max\Big\{1,\ \big(\beta^\dagger/\beta_{\min}\big)^2\Big\} \;+\; n.$$

In particular, in the saturated regime $\beta_{\min} \ge \beta^\dagger$,

$$N \;\le\; N_{\text{sat}} := 1872\,\frac{n^2\rho^2\,L_1\Delta\,\log(2n/\delta)}{\varepsilon^2\,(c' - \rho)^2} \;+\; n,$$

which depends on the range $[\beta_{\min}, \beta_{\max}]$ only through ρ. Consequently $\sup_{\beta_{\min} \ge \beta_0} N < \infty$ for every $\beta_0 > 0$: at fixed $\rho < c'$, the unknown-scale verification cost is uniformly bounded as $\beta \to \infty$.

Proof. Notation as in Appendix F.1: $g = \nabla f(\bar{x})$, true slope $\beta \in [\beta_{\min}, \beta_{\max}]$ unknown (the *range* is known), $\phi = \beta f$, $s := \nabla\phi(\bar{x}) = \beta g$ with coordinates s_i, and $G = \|s\| = \beta\|g\|$. All tester parameters $(M, \tau_G, G_{\max}, \Lambda, t, m, \kappa_*)$ are computable from $(n, L_1, \Delta, \varepsilon, c', \delta, \beta_{\min}, \beta_{\max})$; the true β enters only the analysis.

Overview: why tuning the radius removes the exponential. The fixed-radius estimator of Theorem 5.1, Part II, pays $1/\kappa_B^2 = \Theta(e^{2B})$ because its read logits grow linearly in $\beta_{\min}$ at fixed ρ. Three candidate obstructions to removing this are resolved by one design choice. (i) The margin M grows like $\beta_{\min}$, so one might fear the test needs absolute logit precision that saturates; but the per-read logit

accuracy needed is $tM/(2\sqrt{n})$, which *shrinks proportionally with the radius* t, while the ratio $M/G_{\max} \le (c'-\rho)/(2\rho c')$ is a β-free constant (Step 4 below): both the signal and the required accuracy scale with t, and the channel stays non-degenerate, so the queries per unit of normalized precision are constant in β. (ii) The tester commits to a radius built from $\beta_{\max}$ while the truth may be as small as $\beta_{\min}$; no separate mistuning analysis is needed, because every step below holds *pointwise* for each (f,β) with β in the range — the worst mistuning is precisely the endpoint configuration of Step 1's separation — and the price is the factor $\beta/\beta_{\max} \in [1/\rho, 1]$ on every read logit, i.e. at most ρ^2 in queries, visible as the ρ^2 of the final display. (iii) At small β the readability radius would violate the curvature-bias budget; this forces t to be the *minimum* of two radii (both arms load-bearing) and produces exactly the factor $\max\{1, (\beta^\dagger/\beta_{\min})^2\}$: genuine divergence as $\beta \to 0$ (the noise blows up), boundedness as $\beta \to \infty$.

Step 1 (separation and threshold). Under H0, $G \le \beta_{\max}\varepsilon$; under H1, $G \ge \beta_{\min} c'\varepsilon$. Since $\rho < c'$, with $M = \beta_{\min}\varepsilon(c'-\rho)/2 > 0$ and τ_G the midpoint as in the statement, $\tau_G - M = \beta_{\max}\varepsilon$ and $\tau_G + M = \beta_{\min}c'\varepsilon$ exactly: H0 forces $G \le \tau_G - M$ and H1 forces $G \ge \tau_G + M$ (the separating observable of Theorem 5.1, Part II). It suffices to estimate G to additive error $< M$ with probability $\ge 1-\delta$.

Step 2 (a priori magnitude bound). As in Step II.2 of Appendix F.1, the descent lemma gives $\|g\| \le \sqrt{2L_1\Delta}$, hence

$$G = \beta\|g\| \;\le\; \beta_{\max}\sqrt{2L_1\Delta} = G_{\max}, \qquad |s_i| \le G \le G_{\max}. \tag{F.4}$$

Step 3 (probe logits; curvature bias). For coordinate i let $d_i(\tau) := \langle \nabla f(\bar{x} + \tau e_i), e_i\rangle$; d_i is L_1-Lipschitz with $d_i(0) = g_i$ by the definition of $\mathcal{F}(L_1)$. By the fundamental theorem of calculus,

$$f(\bar{x}+te_i) - f(\bar{x}-te_i) = \int_{-t}^{t} d_i(\tau)\,d\tau = 2tg_i + R_i^f, \qquad |R_i^f| \le \int_{-t}^{t} L_1|\tau|\,d\tau = L_1t^2.$$

Multiplying by the true β, the logit of the queried pair is

$$L_i = 2ts_i + R_i, \qquad |R_i| = \beta|R_i^f| \le \beta L_1 t^2 \le \Lambda t^2. \tag{F.5}$$

Step 4 (the readability invariant $B \le 9/8$ — the heart of the proposition). Set $t_{\text{bias}} := M/(2\Lambda\sqrt{n})$, $t_{\text{read}} := 1/(2G_{\max})$, and $t = \min\{t_{\text{bias}}, t_{\text{read}}\}$ as in the statement. Define the logit budget $B := 2tG_{\max} + \Lambda t^2$; by Equations (F.4) and (F.5), every read logit satisfies $|L_i| \le 2tG + |R_i| \le B$. Then:

(i) $2tG_{\max} \le 2t_{\text{read}}G_{\max} = 1$, using $t \le t_{\text{read}}$;

(ii) $\Lambda t^2 = (\Lambda t)\cdot t \le (\Lambda t_{\text{bias}})\cdot t_{\text{read}} = \frac{M}{2\sqrt{n}}\cdot\frac{1}{2G_{\max}} = \frac{M}{4\sqrt{n}\,G_{\max}}$, using $t \le t_{\text{bias}}$ and $t \le t_{\text{read}}$;

(iii) $\dfrac{M}{G_{\max}} = \dfrac{\beta_{\min}\varepsilon(c'-\rho)/2}{\beta_{\max}\sqrt{2L_1\Delta}} \le \dfrac{\beta_{\min}\varepsilon(c'-\rho)/2}{\beta_{\max}\,c'\varepsilon} = \dfrac{c'-\rho}{2\rho c'} < \dfrac{1}{2}$, where the middle inequality is (A1) ($\sqrt{2L_1\Delta} \ge c'\varepsilon$ — the only use of (A1) in this proof) and the last uses $\rho \ge 1$ and $c'-\rho < c'$;

(iv) hence $\Lambda t^2 < \frac{1}{8\sqrt{n}} \le \frac{1}{8}$ and $B \le 1 + \frac{1}{8} = \frac{9}{8}$.

So every read logit lies in the band $[-9/8, 9/8]$: $p_i := \sigma(L_i) \in [\sigma(-9/8), \sigma(9/8)] \subset [0.245, 0.755]$, and since $z \mapsto \sigma(z)(1-\sigma(z))$ is even and decreasing in $|z|$,

$$\min_{|z|\le 9/8} \sigma(z)(1-\sigma(z)) = \sigma(9/8)\big(1-\sigma(9/8)\big) = \kappa_* \ge 0.185. \tag{F.6}$$

This is the invariant the fixed-radius estimator lacks: there $t = t_{\text{bias}}$ always, B grows linearly in $\beta_{\min}$ at fixed ρ, and the cost carries e^{2B}; here $B \le 9/8$ unconditionally. Moreover, in the saturated regime ($t = t_{\text{read}}$), the aggregate signal under H1 is bounded away from degeneracy: $2tG \ge 2t_{\text{read}}\,\beta_{\min}c'\varepsilon = \beta_{\min}c'\varepsilon/G_{\max} = c'\varepsilon/(\rho\sqrt{2L_1\Delta})$, a β-free scale — while in the small-β regime the signal shrinks with $\beta_{\min}$ and the $\max\{1, (\beta^\dagger/\beta_{\min})^2\}$ factor of Step 7 pays for it. [The cap is asymptotically tight: at $n = 1$ with (A1) an equality and t at the regime boundary $t_{\text{bias}} = t_{\text{read}}$,

one computes $B = 1 + \frac{1}{8\rho} - \frac{1}{8c'} \to 9/8$ as $\rho \to 1$, $c' \to \infty$; neither the clip band nor κ_* can be tightened.]

Step 5 (per-coordinate estimation). Query pair i for m i.i.d. reads; let $\hat{p}_i$ be the empirical bit frequency, with $\mathbb{E}[\hat{p}_i] = p_i$. Clip: $\hat{q}_i := \mathrm{clip}(\hat{p}_i, \sigma(-9/8), \sigma(9/8))$; since p_i lies in the clip interval (Step 4), projection onto it is a contraction toward p_i, so $|\hat{q}_i - p_i| \le |\hat{p}_i - p_i|$. Invert: $\hat{L}_i := \mathrm{logit}(\hat{q}_i)$. On the clip interval the logit has derivative $1/(q(1-q)) \le 1/\kappa_*$ by Equation (F.6), so by the mean value theorem $|\hat{L}_i - L_i| \le |\hat{q}_i - p_i|/\kappa_* \le |\hat{p}_i - p_i|/\kappa_*$. By Hoeffding's inequality with $\delta_1 := \delta/n$, $\mathbb{P}\big[|\hat{p}_i - p_i| > \sqrt{\log(2n/\delta)/(2m)}\big] \le \delta/n$. Set $\hat{s}_i := \hat{L}_i/(2t)$. On the good event, using Equation (F.5),

$$|\hat{s}_i - s_i| \le \frac{|\hat{L}_i - L_i|}{2t} + \frac{|R_i|}{2t} \le \frac{1}{2t\kappa_*}\sqrt{\frac{\log(2n/\delta)}{2m}} + \frac{\Lambda t}{2}. \tag{F.7}$$

Bias half: $\Lambda t/2 \le \Lambda t_{\mathrm{bias}}/2 = M/(4\sqrt{n})$. Statistical half: the read count $m := \lceil 2n\log(2n/\delta)/(t^2\kappa_*^2 M^2)\rceil$ forces $\frac{1}{2t\kappa_*}\sqrt{\log(2n/\delta)/(2m)} \le M/(4\sqrt{n})$. So on the good event $|\hat{s}_i - s_i| \le M/(2\sqrt{n})$.

Step 6 (aggregation; correctness, uniformly over f and β). Union bound over the n coordinates: with probability $\ge 1-\delta$ all n bounds of Step 5 hold; then $\|\hat{s} - s\|_2 \le \sqrt{n}\cdot M/(2\sqrt{n}) = M/2$, and by the reverse triangle inequality $|\hat{G} - G| \le M/2$. Combined with Step 1: under H0, $\hat{G} \le \tau_G - M + M/2 < \tau_G$ (accept); under H1, $\hat{G} \ge \tau_G + M - M/2 > \tau_G$ (reject). Every estimate above holds pointwise for the true (f, β), so the tester is correct with probability $\ge 1-\delta$ on both hypotheses, uniformly over f in the class and over $\beta \in [\beta_{\min}, \beta_{\max}]$.

Step 7 (cost: closed forms and regimes). $N = nm \le 2n^2\log(2n/\delta)/(t^2\kappa_*^2 M^2) + n$, and since $t = \min\{t_{\mathrm{bias}}, t_{\mathrm{read}}\}$,

$$\frac{1}{t^2M^2} = \max\left\{\frac{1}{(t_{\mathrm{bias}}M)^2}, \frac{1}{(t_{\mathrm{read}}M)^2}\right\} = \max\left\{\frac{4\Lambda^2 n}{M^4}, \frac{4G_{\max}^2}{M^2}\right\}.$$

The closed forms are exact:

$$\frac{4G_{\max}^2}{M^2} = \frac{32\,\rho^2 L_1 \Delta}{\varepsilon^2(c'-\rho)^2} \quad (\beta\text{-free}), \qquad \frac{4n\Lambda^2}{M^4} = \frac{64\,n\,\rho^2 L_1^2}{\beta_{\min}^2\varepsilon^4(c'-\rho)^4},$$

the regime boundary $t_{\mathrm{bias}} = t_{\mathrm{read}}$ falls exactly at $\beta_{\min} = \beta^\dagger = \sqrt{2nL_1/\Delta}/(\varepsilon(c'-\rho))$, and the two forms merge through the identity

$$\frac{4n\Lambda^2}{M^4} = \frac{4G_{\max}^2}{M^2}\cdot\left(\frac{\beta^\dagger}{\beta_{\min}}\right)^2.$$

With $2/\kappa_*^2 \le 58.5$ (from $\kappa_* \ge 0.18501$) and $58.5 \times 32 \le 1872$, the master display of the statement follows:

$$N \;\le\; 1872\; n^2\rho^2\,\frac{L_1\Delta}{\varepsilon^2}\,\frac{\log(2n/\delta)}{(c'-\rho)^2}\,\max\left\{1, \left(\frac{\beta^\dagger}{\beta_{\min}}\right)^2\right\} \;+\; n. \tag{F.8}$$

In the saturated regime $\beta_{\min} \ge \beta^\dagger$ the maximum is 1 and $N \le N_{\mathrm{sat}}$, which depends on the range only through ρ. In the small-β regime $\beta_{\min} < \beta^\dagger$,

$$N \;\le\; \frac{3744\; n^3\rho^2 L_1^2\log(2n/\delta)}{\beta_{\min}^2\varepsilon^4(c'-\rho)^4} \;+\; n:$$

the fixed-radius scaling, including $(c'-\rho)^{-4}$, reappears — but with no e^{2B} factor, since $B \le 9/8$ always.

Step 8 (uniform boundedness; what diverges and what does not). N_{sat} contains no $\beta_{\min}$: at fixed $(n, L_1, \Delta, \varepsilon, c', \delta)$ and fixed $\rho < c'$, $\sup_{\beta_{\min}\ge\beta_0} N \le N_{\mathrm{sat}}\max\{1, (\beta^\dagger/\beta_0)^2\} < \infty$ for every $\beta_0 > 0$, and $N \le N_{\mathrm{sat}}$ for all $\beta_{\min} \ge \beta^\dagger$, so $\limsup_{\beta_{\min}\to\infty} N \le N_{\mathrm{sat}} < \infty$. Since N is an achievable cost, it upper-bounds the unknown-scale verification complexity, which is therefore uniformly bounded as $\beta \to \infty$ at fixed $\rho < c'$: the divergence exhibited by the fixed-radius estimator of Theorem 5.1, Step II.6, was an artifact of that estimator, not of the model. The two directions in

which *our bound* does diverge are: $\rho \to c'^-$, at rate $(c' - \rho)^{-2}$ in the saturated regime — which itself requires $\beta_{\min} \gtrsim (c' - \rho)^{-1}$, since $\beta^\dagger(c' - \rho)$ is free of $(c' - \rho)$ — and $(c' - \rho)^{-4}$ at any fixed $\beta_{\min}$; and $\beta_{\min} \to 0$, at rate $\beta_{\min}^{-2}$ through the max factor. At $\rho \geq c'$ testing is impossible outright (Theorem 5.1, Part I). In the regime $\beta_{\min} \leq \beta^\dagger$, both divergences are necessary: Theorem 5.2 gives a problem-level lower bound of order $L_1^2/(\beta_{\min}^2 \varepsilon^4 (c' - \rho)^4)$, which matches Equation (F.8) up to a factor $O(n^3 \log(n/\delta))$. In the saturated regime the lower bound decays like $\beta_{\min}^{-2}$, and whether the saturated cost must grow as $\rho \to c'$ is open (Remark C.3). The noise-magnitude limit $\beta \to \infty$, by contrast, is not a divergence — that is the proposition. (No claim is made about the $\beta = \infty$ model itself, which is a different, noiseless oracle; see the remark below.) □

Remark F.4 (What is not claimed; the noiseless limit stays singular). (i) *No optimality.* At $\rho = 1$ (known scale) the display Equation (F.8) reads $1872\, n^2 (L_1 \Delta/\varepsilon^2) \log(2n/\delta)/(c' - 1)^2 + n$, which does not recover the n-free saturated floor of Theorem 4.1: the n^2 (a deterministic coordinate sweep with a per-coordinate $M/(2\sqrt{n})$ error split) and the linear Δ (from worst-case normalization by $G_{\max}$) come from this construction and are not established as necessary. Whether bracketing G can improve the budget dependence is open. The optimal saturated unknown-scale cost — its n-power, its Δ-dependence, and its $(c' - \rho)$-exponent between our 2 and the 0 that Theorem 5.2 yields as $\beta_{\min} \to \infty$ — is open. Below saturation, Theorem 5.2 settles the $(c' - \rho)$- and $\beta_{\min}$-exponents. (ii) *Mechanism.* Tuning $t = \min\{t_{\text{bias}}, t_{\text{read}}\}$ caps the logit budget at $9/8$ unconditionally, so the entire β-dependence sits in the geometry (the radius), not in the channel; this transports the saturation trick of Theorem 4.1 — probe radius tuned down so logits stay $O(1)$ — to the unknown-scale setting, tuning by $\beta_{\max}$ (the range is the model) at a mistuning price of at most ρ^2. (iii) *The noiseless limit stays singular.* As $\beta_{\min} \to \infty$ at fixed ρ, the radius $t = t_{\text{read}} \sim 1/\beta_{\min}$ probes an ever smaller neighborhood, where potential differences are $\Theta(1/\beta)$ and the channel stays informative. The noiseless sign-only oracle of Appendix G does not expose magnitude information at these shrinking scales, so bounded cost at every fixed $\rho < c'$ does not contradict the noiseless impossibility.

F.3 Proof of the ordinal value-margin theorem

Theorem 6.1 (Smoothness-free ordinal value-margin verification). *Let $f : \mathbb{R}^n \to \mathbb{R}$ be everywhere finite, with $u \mapsto f(\bar{x} + ru)$ μ-measurable;* no *smoothness, boundedness, or gap-budget condition is imposed. Consider the promise problem (6.1) under the BT oracle with scale β known, or unknown in $[\beta_{\min}, \beta_{\max}]$ with $\beta_{\min} > 0$ (for known β read $\beta_{\min} = \beta$ below).*

(i) (Upper bound.) *The majority-vote tester above with $m = \lceil 8 \ln(2/\delta) \rceil$ directions and $k = \lceil 2 \ln(2m/\delta) / \tanh^2(\beta_{\min} m_0/2) \rceil$ reads per direction accepts every $f \in \mathcal{M}^+(m_0)$ and rejects every $f \in \mathcal{M}^-(m_0)$ with probability at least $1 - \delta$, using*

$$T = mk = O\Big(\max\big\{1, (\beta_{\min} m_0)^{-2}\big\} \cdot \log(1/\delta) \cdot \log\big(\log(1/\delta)/\delta\big)\Big)$$

comparisons. Its decision rule is scale-free; $\beta_{\min}$ enters only the read count k.

(ii) (Lower bound, no regime cap.) *For every $m_0 > 0$, $r > 0$, $n \geq 1$ and every true scale β: any randomized adaptive tester correct with probability at least $2/3$ on both classes needs*

$$T \geq \frac{2 \ln(9/8)}{(\beta m_0)^2} \approx \frac{0.2356}{(\beta m_0)^2},$$

and at confidence $\delta \leq 1/16$, $T \geq \ln(1/\delta)/(\beta m_0)^2$.

Proof. Recall the objects of the statement: μ is the uniform (rotation-invariant) probability measure on $\mathbb{S}^{n-1}$, $\text{gap}(u) := f(\bar{x} + ru) - f(\bar{x})$, the accept class $M^+(m_0)$ demands $\text{gap}(u) \geq m_0$ for all $u \in \mathbb{S}^{n-1}$, the reject class $M^-(m_0)$ demands $\mu\{u : \text{gap}(u) \leq -m_0\} \geq p_0 = 3/10$, and the only hypotheses on f are that it is real-valued (everywhere finite) with the probe map $u \mapsto f(\bar{x} + ru)$ μ-measurable. For known β, read $\beta_{\min} = \beta_{\max} = \beta$ throughout.

Step 0 (conventions and well-posedness). Finite gaps. f is real-valued, so every $\text{gap}(u)$ is a finite real, every Bernoulli parameter $\sigma(\beta\, \text{gap}(u))$ lies in the open interval $(0, 1)$, and each oracle response is a well-defined non-degenerate Bernoulli bit; $\sigma(\pm\infty)$ never arises. Gaps may be arbitrarily large in magnitude — the class is unbounded — but largeness only sharpens votes (Step U1).

Measurability. Equip $\mathbb{S}^{n-1}$ with the μ-completion of its Borel σ-algebra; the probe map is μ-measurable by hypothesis, hence so is gap, so the descent set $D := \{u : \mathrm{gap}(u) \le -m_0\}$ is measurable and $\mu(D)$ is well-defined (M^+ is a pointwise condition and needs no measurability). Fix a Borel μ-version of the gap map: a completion-measurable function equals a Borel function μ-almost everywhere, response kernels that agree μ-a.e. induce the same transcript law when the probe directions are drawn i.i.d. from μ, and D changes by a μ-null set, so $\mu(D)$ and every event below are unaffected; we therefore take the gap map to be Borel. The tester below is non-adaptive, so its transcript law is the explicit semidirect product $\mu^{\otimes m} \ltimes K^{\otimes m}$ on the standard Borel space $(\mathbb{S}^{n-1})^m \times \{-1,+1\}^{mk}$, where $K : u \mapsto \mathrm{Bern}(\sigma(\beta\,\mathrm{gap}(u)))^{\otimes k}$ is a Markov kernel; every event used below is measurable there, and Tonelli's theorem justifies each conditioning on u_i and each union bound. For the lower bound, adaptive testers appear; their transcript spaces are handled once and for all inside Lemma B.4's framework, and the two lower-bound environments are C^2, hence Borel, so nothing further is needed. (Scope note: a tester other than ours could query points off the probe sphere; if every class member is to induce a well-defined transcript law against every such tester, intersect the class with globally Borel f — a cosmetic restriction that changes neither bound, since our tester queries only sphere–center pairs and the lower-bound witnesses are C^2. For $n = 1$, $S^0 = \{\pm 1\}$ with μ uniform on two points; nothing below uses $n \ge 2$.)

The tester. Draw $m = \lceil 8\log(2/\delta)\rceil$ i.i.d. directions $u_1, \dots, u_m \sim \mu$; for each i take $k = \lceil 2\log(2m/\delta)/\tanh^2(\beta_{\min} m_0/2)\rceil$ Bradley–Terry reads of the pair $(\bar{x} + ru_i,\ \bar{x})$ — each read is $+1$ with probability $\sigma(\beta\,\mathrm{gap}(u_i))$, else -1 — and let $b_i := +1$ iff strictly more than $k/2$ reads are $+1$ (a tie counts as -1). REJECT iff some $b_i = -1$, else ACCEPT. The decision rule mentions no function class and no β: it is a scale-free majority. This is the natural noiseless sphere-sampling verifier — probe m random directions at radius r and reject on any observed descent — with each noiseless comparison replaced by a k-read majority vote.

Phase 0 (the garbage-bit audit). A majority vote reliably recovers $\mathrm{sign}(\mathrm{gap}(u_i))$ only when $|\mathrm{gap}(u_i)|$ is bounded away from 0; on a direction with gap ≈ 0 the vote is a near-fair coin ("garbage"). Classify every consumed bit: (a) on $f \in M^+(m_0)$, $\mathrm{gap}(u_i) \ge m_0 > 0$ for *every* probe — the margin hypothesis constrains all points of the radius-r sphere, exactly where all probes land — so accept instances carry no garbage; (b) on $f \in M^-(m_0)$, a probe landing in D is guaranteed ($|\mathrm{gap}| \ge m_0$, correct answer -1), while a probe outside D has $\mathrm{gap}(u_i) \in (-m_0, \infty)$ with no lower guarantee — its bit may be garbage (for arbitrary f the non-descent gaps can be exactly 0 or astronomically large; neither matters below); (c) the asymmetric rule "reject iff some $b_i = -1$" copes: garbage occurs only on reject instances, where a garbage bit is either -1 (triggering the correct decision) or $+1$ (harmless, since a false accept needs *all* $b_i = +1$ while the mass-$\ge 3/10$ descent set fires a -1 with high probability, Step U3). Garbage can never manufacture a false accept. No step of this audit consulted smoothness, boundedness, or a gap budget.

Step U1 (one direction: the vote lemma, in two regimes). Fix a direction u with $|\mathrm{gap}(u)| \ge \gamma > 0$. The k reads are i.i.d. with $\mathbb{P}[\text{read} = \mathrm{sign}(\mathrm{gap}(u))] = \sigma(\beta|\mathrm{gap}(u)|) \ge \sigma(\beta\gamma) = \frac{1}{2} + \eta$, where $\eta = \frac{1}{2}\tanh(\beta\gamma/2) > 0$ (the identity $\sigma(z) - \frac{1}{2} = \frac{1}{2}\tanh(z/2)$, and σ increasing — a larger $|\mathrm{gap}|$ only strengthens the bound below, which is the only way gaps enter). The majority is wrong only if the empirical fraction of correct reads is $\le 1/2$, a deviation of at least η below its mean; Hoeffding's inequality on k i.i.d. bounded indicators — boundedness of the *reads*, never of f — gives

$$\mathbb{P}[\text{majority wrong}] \le \exp(-2k\eta^2) = \exp\!\big(-\tfrac{k}{2}\tanh^2(\beta\gamma/2)\big), \tag{F.9}$$

so $k = \lceil 2\log(2m/\delta)/\tanh^2(\beta\gamma/2)\rceil$ yields $\mathbb{P}[\text{majority wrong}] \le \delta/(2m)$. Two regimes: since $\tanh(z)/z$ is decreasing, $\tanh(z) \ge z\tanh(1)$ on $[0,1]$, so for $\beta\gamma \le 2$, $\tanh^2(\beta\gamma/2) \ge 0.145\,(\beta\gamma)^2$ and $k = O((\beta\gamma)^{-2}\log(m/\delta))$; for $\beta\gamma > 2$ the tanh factor saturates and k floors at $O(\log(m/\delta))$. Jointly, $k = O\big(\max\{1, (\beta\gamma)^{-2}\}\,\log(m/\delta)\big)$. The tester sets the resolution to the promise margin, $\gamma := m_0$. The *rule* is β-independent (a majority threshold at $k/2$; the per-read bias has the sign of the gap for every $\beta > 0$); only the *count* uses β, monotonically, so computing k at $\beta_{\min}$ over-counts safely for every $\beta \in [\beta_{\min}, \beta_{\max}]$, and $\beta_{\max}$ never enters.

Step U2 (completeness). On $f \in M^+(m_0)$, every $\mathrm{gap}(u_i) \ge m_0$ (Phase 0(a)), so by Equation (F.9) with $\gamma = m_0$, $\mathbb{P}[b_i = -1] \le \delta/(2m)$ for each i, conditionally on u_i and hence unconditionally (Tonelli, Step 0). A union bound over the m probes gives $\mathbb{P}[\text{ACCEPT fails}] = \mathbb{P}[\text{some } b_i = -1] \le m \cdot \delta/(2m) = \delta/2 \le \delta$.

Step U3 (soundness). On $f \in M^-(m_0)$, D is measurable with $\mu(D) \geq 3/10$. *Geometry:* the u_i are i.i.d. μ, so $\mathbb{P}[\text{no probe lands in } D] \leq (7/10)^m$, and with $m \geq 8\log(2/\delta)$, $(7/10)^m \leq (\delta/2)^{8\log(10/7)} = (\delta/2)^{2.853} \leq \delta/2$ (the exponent exceeds 1 and $\delta/2 < 1$). Call E the event that some probe lands in D. *Votes on descent probes:* on $\{u_i \in D\}$, $|\mathrm{gap}(u_i)| \geq m_0$ with correct answer -1, so by Equation (F.9), $\mathbb{P}[u_i \in D \text{ and } b_i \neq -1] \leq \delta/(2m)$; a union over i gives $\mathbb{P}[\text{some descent probe votes} \neq -1] \leq \delta/2$. *Combine:* on E intersected with "every descent probe votes -1", some $b_i = -1$ and the tester rejects; hence $\mathbb{P}[\textsc{reject} \text{ fails}] \leq \mathbb{P}[E^c] + \delta/2 \leq \delta$. Crucially, nothing was assumed about the non-descent probes' bits (the garbage of Phase 0(b)): they enter neither term. (Any constant $p_0 \in (0,1)$ would do here, rescaling only the constant in m.)

Step U4 (total count). $T = mk$ with $m = \lceil 8\log(2/\delta) \rceil = O(\log(1/\delta))$ and, by Step U1's two-regime form, $k = O\big(\max\{1, (\beta_{\min} m_0)^{-2}\}\, \log(\log(1/\delta)/\delta)\big)$; hence

$$T = O\Big(\max\{1, (\beta_{\min} m_0)^{-2}\} \cdot \log(1/\delta) \cdot \log\big(\log(1/\delta)/\delta\big)\Big). \tag{F.10}$$

The count is governed by the promise margin m_0 alone (with $\beta_{\min}$); m_0 is a free parameter of the promise, and no step of Phase 0 or U1–U4 consulted smoothness, boundedness, or a gap budget.

Step L1 (the lower-bound pair: a bowl and a hemisphere bump). Take $\bar{x} = 0$ without loss of generality, fix a unit vector w, and reserve T for the query budget, writing Ψ for the perturbation. Let $S(z) \coloneqq 6z^5 - 15z^4 + 10z^3$ (the quintic smoothstep: $S(0) = S'(0) = S''(0) = 0$, $S(1) = 1$, $S'(1) = S''(1) = 0$), and define the radial and angular profiles

$$q(\tau) \coloneqq \begin{cases} S(\tau), & \tau \in [0,1], \\ S(2-\tau), & \tau \in [1,2], \\ 0, & \text{else}, \end{cases} \qquad A(z) \coloneqq \begin{cases} 1, & z \geq 0, \\ S(1+z), & z \in [-1,0], \\ 0, & z \leq -1. \end{cases}$$

Set

$$f_0(x) \coloneqq \frac{m_0}{r^2}\|x\|^2, \qquad f_1 \coloneqq f_0 - \Psi, \qquad \Psi(x) \coloneqq 2m_0\, q\big(\|x\|/r\big)\, A\big(\langle w, x\rangle / r\big).$$

Ψ is C^2 (the smoothstep vanishes to second order at its knots), supported in $\{\|x\| \leq 2r\}$, one-signed with range $[0, 2m_0]$, and $\nabla\Psi(0) = 0$. *Membership is pure algebra, valid for every $m_0 > 0$, $r > 0$, $n \geq 1$:* $\mathrm{gap}_{f_0}(u) = (m_0/r^2)r^2 = m_0$ for all $u \in \mathbb{S}^{n-1}$, so $f_0 \in M^+(m_0)$ (the tightest accept instance); and on the hemisphere $H = \{u : \langle w, u\rangle \geq 0\}$ we have $q(1) = 1$ and $A(\langle w, u\rangle) = 1$, so $\Psi(ru) = 2m_0$ and $\mathrm{gap}_{f_1}(u) = m_0 - 2m_0 = -m_0$, with $\mu(H) = 1/2 \geq 3/10 = p_0$, so $f_1 \in M^-(m_0)$. *Closeness:* $|f_0 - f_1| = \Psi \leq 2m_0$ everywhere.

Step L2 (admissibility in the widened class; the regime cap is void). A lower bound proved on a two-point subfamily $\{f_0, f_1\}$ applies verbatim to every superclass containing the pair: a tester correct with probability $\geq 2/3$ on the promise classes of the statement is in particular correct on (f_0, f_1), and the transcript argument of Step L3 depends only on the pair and the budget T. Both witnesses are C^2, hence Borel, hence real-valued sphere-measurable functions, so both lie in the theorem's class for *every* $m_0 > 0$, $r > 0$, $n \geq 1$: no smoothness membership is required, and no constraint ties m_0 to (r, n). For a restriction of the promise to $\mathcal{F}(L_1)$, one would additionally need to verify a Hessian-norm bound for the perturbed objective. That restriction is not used in the theorem or the argument below.

Step L3 (transcript divergence and Bretagnolle–Huber). Let P_0, P_1 be the T-query transcript laws of an arbitrary (possibly adaptive, randomized) tester under f_0 and f_1. The per-query gap difference is $\mathrm{gap}_{f_0}(x,y) - \mathrm{gap}_{f_1}(x,y) = \Psi(x) - \Psi(y)$, and since Ψ is *one-signed* with range $[0, 2m_0]$, $|\Psi(x) - \Psi(y)| \leq 2m_0$ pointwise — not the crude $4m_0$ that the two-sided envelope $|f_0 - f_1| \leq 2m_0$ alone would give; the one-signed refinement is free and saves a factor 4 in the final constant. By the general decomposition of Lemma B.4,

$$\mathrm{KL}(P_0 \,\|\, P_1) \;\leq\; \frac{\beta^2}{8} \sum_{t=1}^{T} \mathbb{E}_{f_0}\big[(\Psi(X_t) - \Psi(Y_t))^2\big] \;\leq\; \frac{\beta^2}{8}\,(2m_0)^2\, T \;=\; \frac{\beta^2 m_0^2}{2}\, T. \tag{F.11}$$

Constant confidence. The accept event is transcript-measurable, so correctness with probability $\geq 2/3$ on both instances forces $\mathrm{TV}(P_0, P_1) \geq 1/3$; the Bretagnolle–Huber inequality $\mathrm{TV} \leq \sqrt{1 - e^{-\mathrm{KL}}}$

(Tsybakov, 2009) gives $e^{-\mathrm{KL}} \le 8/9$, i.e. $\mathrm{KL} \ge \log(9/8)$, and with Equation (F.11),

$$T \;\ge\; \frac{2\log(9/8)}{(\beta m_0)^2} \qquad \big(2\log(9/8) \approx 0.2356\big). \tag{F.12}$$

High confidence. If both error probabilities are $\le \delta$, then $\mathrm{TV}(P_0, P_1) \ge 1 - 2\delta$, and Bretagnolle–Huber gives $e^{-\mathrm{KL}} \le 1 - (1-2\delta)^2 = 4\delta(1-\delta) \le 4\delta$, i.e. $\mathrm{KL} \ge \log \frac{1}{4\delta} \ge \frac{1}{2}\log\frac{1}{\delta}$ for $\delta \le 1/16$ (the boundary is exact); with Equation (F.11),

$$T \;\ge\; \frac{\log(1/\delta)}{(\beta m_0)^2} \qquad \text{for } \delta \le \tfrac{1}{16}. \tag{F.13}$$

Both bounds hold for every true slope β and every $m_0 > 0$, $r > 0$, $n \ge 1$. Under unknown $\beta \in [\beta_{\min}, \beta_{\max}]$, the two environments may be presented at the least informative in-range slope $\beta = \beta_{\min}$; a tester correct on the unknown-scale problem is in particular correct on that pair, so $T \ge 2\log(9/8)/(\beta_{\min} m_0)^2$ — matching the $\beta_{\min}$ in the upper bound's count.

Matching, and its scope. In the noise-limited regime $\beta m_0 = O(1)$, Equation (F.10) and Equations (F.12) and (F.13) match in the (βm_0)-exponent — both are $\Theta((\beta m_0)^{-2})$ up to logarithms — with one $\log(1/\delta)$ gap in the δ-powers: the upper bound carries $\log(1/\delta)\log(\log(1/\delta)/\delta)$, the lower bound $\log(1/\delta)$. For $\beta m_0 \gg 1$ the upper bound floors at its logarithmic terms (the k-floor of Step U1) while the lower bound tends to 0 — the shape of the noiseless limit; at $\beta = \infty$ the count recovers a noiseless $O(\log(1/\delta))$-type cost up to the same logarithmic factor. The lower-bound constant $2\log(9/8)$ does not match the upper bound's union-bound constants: the match is in the exponent of βm_0, not in constants. □

Remark F.5 (The amplitude of the value-margin problem). Let (β, f) be an accept and (β', f') a reject instance of (6.1), with $\beta, \beta' \in [\beta_{\min}, \beta_{\max}]$. The descent set of f' has measure at least $p_0 > 0$, so it contains some u, and then $\beta\,\mathrm{gap}_f(u) \ge \beta m_0 \ge \beta_{\min} m_0$ while $\beta'\,\mathrm{gap}_{f'}(u) \le -\beta' m_0 \le -\beta_{\min} m_0$. Hence $\mathrm{osc}(\beta f - \beta' f') \ge 2\beta_{\min} m_0$, and the logit distance is at least $\beta_{\min} m_0$. The pair of Step L1 at the scale $\beta_{\min}$ differs by the one-signed bump $\beta_{\min}\Psi$ with range $[0, 2\beta_{\min} m_0]$, so it attains $\beta_{\min} m_0$. The amplitude is therefore exactly $\beta_{\min} m_0$ (exactly βm_0 at a known scale), and Theorem 2.7(b) reproduces the lower bounds of Step L3.

Remark F.6 (The $\mathcal{F}(L_1)$ ceiling, and why M^+ is broader than local minimality). (i) *Sphere averaging.* If additionally $f \in \mathcal{F}(L_1)$, then $|\mathrm{gap}(u) - r\langle \nabla f(\bar{x}), u\rangle| \le (L_1/2)r^2$ pointwise (the integral Taylor remainder under an L_1-Lipschitz gradient), and $\mathbb{E}_{u\sim\mu}[u] = 0$ by symmetry, so

$$\mathbb{E}_u[\mathrm{gap}(u)] \;=\; r\,\langle \nabla f(\bar{x}), \mathbb{E}_u[u]\rangle + \mathbb{E}_u[\text{remainder}] \;\le\; \frac{L_1}{2}\,r^2 \;=\; m_{\mathrm{s}}.$$

Hence $\mathrm{gap}(u) \ge m_0$ for all u forces $m_0 \le \mathbb{E}_u[\mathrm{gap}] \le m_{\mathrm{s}}$: the class $M^+(m_0) \cap \mathcal{F}(L_1)$ is empty unless $0 < m_0 \le m_{\mathrm{s}}$. At the radius of Corollary F.7, $r = \varepsilon/(2\sqrt{n}\,L_1)$, the ceiling reads $m_{\mathrm{s}} = \varepsilon^2/(8nL_1)$. Over the widened class of Theorem 6.1 no ceiling exists: the bowl f_0 of Step L1 lies in $M^+(m_0)$ for every $m_0 > 0$ and merely exits $\mathcal{F}(L_1)$ once $m_0 > m_{\mathrm{s}}$ (its Hessian norm is $2m_0/r^2$). The cap was never a feature of the ordinal problem; it is the shadow of $\mathcal{F}(L_1)$-membership. (ii) *Scope of the margin promise.* Membership in $M^+(m_0)$ constrains f on the radius-r sphere only; it implies neither that $\bar{x}$ is a local minimum nor that $\nabla f(\bar{x}) = 0$. A counterexample inside $\mathcal{F}(L_1)$: $f(x) = \langle c, x - \bar{x}\rangle + (L_1/2)\|x - \bar{x}\|^2$ with $\|c\| = L_1 r/4$ has every sphere gap at least $-r\|c\| + (L_1/2)r^2 = L_1 r^2/4 = m_{\mathrm{s}}/2$, so $f \in M^+(m_{\mathrm{s}}/2)$, yet $\nabla f(\bar{x}) = c \ne 0$. Thus $M^+(m_0)$ is a *value-margin* class at scale r, strictly broader than the local minima; "local-minimum verification" is the special case, and Theorem 6.1 is a theorem about value margins.

F.4 Proof of the gradient-soundness corollary

Corollary F.7 (Gradient soundness of the ordinal tester). *Let $f \in \mathcal{F}(L_1)$, $r = \varepsilon/(2\sqrt{n}L_1)$, and $0 < m_0 \le m_{\mathrm{s}} = \varepsilon^2/(8nL_1)$. Then every $\bar{x}$ with $\|\nabla f(\bar{x})\| \ge \varepsilon$ lies in $\mathcal{M}^-(m_{\mathrm{s}}) \subseteq \mathcal{M}^-(m_0)$, so the tester of Theorem 6.1(i) also rejects, with probability at least $1 - \delta$, every point whose gradient norm is at least ε — at the ordinal cost $O\big(\max\{1, (\beta_{\min} m_0)^{-2}\} \cdot \log(1/\delta)\log(\log(1/\delta)/\delta)\big)$.*

Recall Corollary F.7: for $f \in \mathcal{F}(L_1)$, probe radius $r = \varepsilon/(2\sqrt{n}\,L_1)$ — so that $m_{\mathrm{s}} = (L_1/2)r^2 = \varepsilon^2/(8nL_1)$ — and any margin $0 < m_0 \le m_{\mathrm{s}}$, every $\bar{x}$ with $\|\nabla f(\bar{x})\| \ge \varepsilon$ lies in $M^-(m_{\mathrm{s}}) \subseteq$

$M^-(m_0)$, so the tester of Theorem 6.1 also rejects every ε-gradient point, with probability $\geq 1-\delta$, at the ordinal cost Equation (F.10). The proof uses the following elementary spherical-cap mass bound.

Lemma F.8 (Cap mass). *For every $n \geq 1$ and every unit vector $v \in \mathbb{R}^n$, $\mu\{u \in \mathbb{S}^{n-1} : \langle v, u\rangle \leq -\frac{1}{2\sqrt{n}}\} \geq \frac{3}{10}$.*

Proof. Let $u \sim \mu$, write $s := \langle v, u\rangle$ and $t := 1/(2\sqrt{n})$. For $n \geq 2$ the marginal s has the even density $p_n(s) = c_n(1-s^2)^{(n-3)/2}$ on $[-1,1]$, with $c_n = \Gamma(n/2)/\big(\sqrt{\pi}\,\Gamma((n-1)/2)\big)$ (the standard normalization, $\int_{-1}^{1}(1-s^2)^{(n-3)/2}\,ds = B\big(\frac{1}{2}, \frac{n-1}{2}\big) = c_n^{-1}$). Since p_n is even, $\int_{-1}^{0} p_n = \frac{1}{2}$, so

$$\mu[s \leq -t] \;=\; \tfrac{1}{2} - \mu[-t < s \leq 0] \;\geq\; \tfrac{1}{2} - t \sup_{s\in[-t,0]} p_n(s). \tag{F.14}$$

A Cauchy–Schwarz bound on c_n. For $x > 0$, splitting the integrand symmetrically,

$$\begin{aligned}\Gamma\big(x+\tfrac{1}{2}\big) &= \int_0^\infty r^{(x-1)/2}e^{-r/2}\cdot r^{x/2}e^{-r/2}\,dr \\ &\leq \Big(\int_0^\infty r^{x-1}e^{-r}\,dr\Big)^{1/2}\Big(\int_0^\infty r^{x}e^{-r}\,dr\Big)^{1/2} \\ &= \Gamma(x)\sqrt{x}\,.\end{aligned}$$

With $x = (n-1)/2$ (so $x+\frac{1}{2} = n/2$), $c_n = \Gamma(x+\frac{1}{2})/(\sqrt{\pi}\,\Gamma(x)) \leq \sqrt{x/\pi} = \sqrt{(n-1)/(2\pi)} \leq \sqrt{n/(2\pi)}$.

Case $n \geq 3$. The exponent $(n-3)/2 \geq 0$, so $(1-s^2)^{(n-3)/2}$ is largest on $[-t,0]$ at $s=0$; thus $\sup_{[-t,0]} p_n = p_n(0) = c_n \leq \sqrt{n/(2\pi)}$, and by (F.14), $\mu[s \leq -t] \geq \frac{1}{2} - \frac{1}{2\sqrt{n}}\sqrt{n/(2\pi)} = \frac{1}{2} - \frac{1}{2\sqrt{2\pi}}$. Since $2\pi > 6.25 = (5/2)^2$ we have $\sqrt{2\pi} > 5/2$, hence $1/(2\sqrt{2\pi}) < 1/5$ and $\mu[s \leq -t] > \frac{1}{2} - \frac{1}{5} = \frac{3}{10}$ (numerically ≥ 0.3005).

Case $n = 2$. Here $c_2 = 1/\pi$ and $p_2(s) = \frac{1}{\pi}(1-s^2)^{-1/2}$ (the arcsine density), increasing in $|s|$; with $t = 1/(2\sqrt{2})$, $t^2 = \frac{1}{8}$, the supremum on $[-t,0]$ is at $s=-t$, where $p_2(-t) = \frac{1}{\pi}\sqrt{8/7}$. Then $t\,p_2(-t) = \frac{1}{2\sqrt{2}}\cdot\frac{1}{\pi}\sqrt{8/7} = \frac{1}{\pi\sqrt{7}}$, so by (F.14), $\mu[s \leq -t] \geq \frac{1}{2} - \frac{1}{\pi\sqrt{7}}$. As $\pi\sqrt{7} = 8.3118\ldots > 5$, the second term is $< 1/5$ and $\mu[s \leq -t] > \frac{3}{10}$ (the bound equals 0.3797; the exact arcsine value is 0.3850).

Case $n = 1$. Then $\mathbb{S}^{n-1} = \{-1,+1\}$ with $u = \pm 1$ each of measure $\frac{1}{2}$, $\langle v, u\rangle \in \{-1,+1\}$, and $t = \frac{1}{2}$; hence $\mu[\langle v,u\rangle \leq -\frac{1}{2}] = \frac{1}{2} \geq \frac{3}{10}$. □

Proof of Corollary F.7. Suppose $\|\nabla f(\bar{x})\| \geq \varepsilon$ and let $v = \nabla f(\bar{x})/\|\nabla f(\bar{x})\|$. By Lemma F.8, $\mu\{u : \langle v,u\rangle \leq -1/(2\sqrt{n})\} \geq 3/10$. For any such u, the descent lemma for $\mathcal{F}(L_1)$ gives

$$\mathrm{gap}(u) = f(\bar{x}+ru) - f(\bar{x}) \;\leq\; r\,\langle \nabla f(\bar{x}), u\rangle + \frac{L_1}{2}\,r^2 \;\leq\; -\frac{r\varepsilon}{2\sqrt{n}} + \frac{L_1}{2}\,r^2 \;=\; -\frac{\varepsilon^2}{8nL_1} \;=\; -m_{\mathrm{s}},$$

using $\langle \nabla f(\bar{x}), u\rangle = \|\nabla f(\bar{x})\|\langle v,u\rangle \leq -\varepsilon/(2\sqrt{n})$ and then the exact arithmetic of $r = \varepsilon/(2\sqrt{n}L_1)$. Hence $\mu\{u : \mathrm{gap}(u) \leq -m_{\mathrm{s}}\} \geq 3/10$, i.e. $f \in M^-(m_{\mathrm{s}})$; and since $m_0 \leq m_{\mathrm{s}}$, $\{\mathrm{gap} \leq -m_{\mathrm{s}}\} \subseteq \{\mathrm{gap} \leq -m_0\}$, so $f \in M^-(m_0)$ as well. Step U3 of Theorem 6.1's proof therefore rejects $\bar{x}$ with probability $\geq 1-\delta$ — the votes, calibrated to resolution $\gamma = m_0$, resolve the larger descent $|\mathrm{gap}| \geq m_{\mathrm{s}}$ a fortiori — at the cost Equation (F.10). This is exactly where smoothness earns its keep: the descent lemma converts a cardinal gradient hypothesis into an ordinal descent mass, whereas Theorem 6.1 itself needs no smoothness at all. □

Remark F.9 (The cardinal corner; a non-claim). At the corner $m_0 \to m_{\mathrm{s}} = \varepsilon^2/(8nL_1)$ the ordinal count's leading factor degrades to

$$\Big(\frac{1}{\beta m_{\mathrm{s}}}\Big)^2 = \frac{64\,n^2L_1^2}{\beta^2\varepsilon^4},$$

recovering the cardinal $L_1^2/(\beta^2\varepsilon^4)$ scaling exactly as the ordinal margin shrinks to the smoothness-imposed ceiling; (ε, n, L_1) enter the ordinal count only through this specialization — through the probe radius and the ceiling — never through the ordinal rate itself, which is purely (βm_0). The cardinal task proper is Theorem 4.1's subject. We explicitly do *not* claim tightness for the mixed task "accept $M^+(m_0)$ versus reject $\|\nabla f(\bar{x})\| \geq \varepsilon$": our tester's count at margin m_0 scales as $(\beta_{\min} m_0)^{-2}$, while a tester thresholding at resolution m_s instead attains $O\big((\beta_{\min} m_s)^{-2} \cdot \mathrm{polylog}\big) = O\big(n^2 L_1^2/(\beta_{\min}^2 \varepsilon^4) \cdot \mathrm{polylog}\big)$ and is not beaten by ours; the tight two-sided object is the ordinal problem of Theorem 6.1.

G THE NOISELESS LIMIT: VERIFICATION IS IMPOSSIBLE

This appendix establishes the paper's noiseless baseline: under the exact, sign-only comparison oracle — the pointwise $\beta \to \infty$ limit of Definition 2.1 — verifying ε-stationarity at a given point is impossible at *every* query budget, no completeness is possible for critical points, and selection among candidates is impossible as well, formalizing a remark of Wang et al. (2026). These results complement Theorem 5.1 by showing how invariance obstructs certification when only exact orderings are observed. As discussed in Remarks F.2 and H.2, this is a comparison of identifiability mechanisms, not a convergence claim for the unknown-scale experiment as $\rho \to \infty$. The oracle below is a different, noiseless channel from Definition 2.1, and no proof in the body depends on this appendix. One identity deserves emphasis at the outset: the feasibility threshold of Theorem G.3 is $c'\varepsilon \leq \sqrt{2L_1\Delta}$ — *literally* the expression of Assumption 2.3. The noiseless zero-window boundary and the noisy model's feasibility assumption are the same inequality: (A1) marks where the reject class becomes nonempty, and in the noiseless model the task is impossible on all of it, at every query budget.

G.1 THE NOISELESS ORACLE AND THE INVARIANCE LEMMA

Definition G.1 (Noiseless comparison oracle)**.** Fix $f\colon \mathbb{R}^n \to \mathbb{R}$. On a query pair $(x, y) \in \mathbb{R}^n \times \mathbb{R}^n$ the *noiseless comparison oracle* returns a single sign,

$$\text{answer}(x, y) \;=\; \mathrm{sign}\big(f(x) - f(y)\big) \;\in\; \{-1, +1\} \qquad \text{whenever } f(x) \neq f(y), \tag{G.1}$$

while on a tie $f(x) = f(y)$ either sign is admissible and the choice is made adversarially by a *tie rule*, possibly at random and inconsistently across repeats. A tie rule is *blind* if it depends only on the transcript so far and the queried pair, never on the values of f; this is the standard adversarial-tie model, and every impossibility below holds under every blind rule.

On any pair with $f(x) \neq f(y)$, the BT oracle of Definition 2.1 has $\mathbb{P}[B = 1] = \sigma(\beta\,\mathrm{gap}_f(x, y)) \to \mathbf{1}\{f(x) > f(y)\}$ as $\beta \to \infty$, so — recoding the bit $B \in \{0, 1\}$ as a sign — the noiseless oracle is the pointwise $\beta \to \infty$ limit of the BT oracle; on a tie, $\sigma(0) = 1/2$ for every β, and the fair coin is one admissible blind tie rule. Algorithms are exactly as in Section 2 with the Bernoulli bit replaced by the sign answer: randomized and adaptive, with internal randomness U, the t-th pair measurable in U and the previous answers, an arbitrary — possibly unbounded and data-dependent — budget, and output measurable in $(U, \text{transcript})$. Throughout this appendix we write, extending the class notation of Section 2 by the gap budget at the anchor,

$$\mathcal{F}(L_1, L_2, \Delta) \;:=\; \big\{f \in \mathcal{F}(L_1) \cap C^2(\mathbb{R}^n) \,:\, \nabla^2 f \text{ is } L_2\text{-Lipschitz},\; f(\bar{x}) - \inf f \leq \Delta\big\};$$

for $L_2 > 0$ this is $\{f \in \mathcal{F}(L_1, L_2) : f(\bar{x}) - \inf f \leq \Delta\}$, and $L_2 = 0$ is allowed where stated (the witnesses of Theorem G.3 have constant Hessians).

The engine of the whole family is one fact: a comparison oracle sees only the *order* that f induces on $\mathbb{R}^n$, and that order is untouched by any strictly increasing relabelling of the values of f.

Lemma G.2 (Monotone invariance of comparisons)**.** *Let $U \subseteq \mathbb{R}$ be any set with $\mathrm{range}(f_0) \subseteq U$ and let $\phi\colon U \to \mathbb{R}$ be strictly increasing on $\mathrm{range}(f_0)$. Put $f_0\colon \mathbb{R}^n \to \mathbb{R}$ and $f_1 := \phi \circ f_0$. Then:*

(a) *for every pair (x, y), $\mathrm{sign}\big(f_1(x) - f_1(y)\big) = \mathrm{sign}\big(f_0(x) - f_0(y)\big)$, ties included: $f_1(x) = f_1(y) \iff f_0(x) = f_0(y)$;*

(b) *consequently, for every* blind *tie rule $\mathcal{T}$ — a (possibly randomized) map from (transcript so far, queried pair) to $\{-1, +1\}$ that does not read the values of f — there is a coupling of*

the algorithm's and the rule's randomness under which, for each realization of the shared randomness, the realized transcripts of the f_0-oracle and the f_1-oracle (both with rule $\mathcal{T}$) coincide pointwise; hence every randomized, adaptive, arbitrary-length comparison algorithm has identical transcript and output distributions on f_0 and f_1;

(c) *in particular, for $\phi(s) = as$ with $a > 0$, $f_1 = af_0$ is comparison-indistinguishable from f_0.*

Proof. (a) For $u, v \in \mathrm{range}(f_0)$, strict monotonicity of ϕ gives $u > v \iff \phi(u) > \phi(v)$, and, by injectivity of a strictly increasing map, $u = v \iff \phi(u) = \phi(v)$. Apply this with $u = f_0(x)$, $v = f_0(y)$: since $f_1 = \phi \circ f_0$,

$$\mathrm{sign}\big(f_1(x) - f_1(y)\big) = \mathrm{sign}\big(\phi(f_0(x)) - \phi(f_0(y))\big) = \mathrm{sign}\big(f_0(x) - f_0(y)\big),$$

and the tie sets $\{(x, y) : f_0(x) = f_0(y)\}$ and $\{(x, y) : f_1(x) = f_1(y)\}$ coincide.

(b) Fix a blind rule $\mathcal{T}$ and couple the two runs to share the algorithm's internal randomness and the randomness of $\mathcal{T}$. We show by induction on the query index t that the two realized transcripts agree. At $t = 1$ the first pair is a function of the algorithm's coins only, hence identical in both runs. Assume the transcripts agree through step $t - 1$, so the same pair (x_t, y_t) is issued in both runs. If $f_0(x_t) \neq f_0(y_t)$, then by (a) $\mathrm{sign}(f_1(x_t) - f_1(y_t)) = \mathrm{sign}(f_0(x_t) - f_0(y_t))$ and both oracles return this common non-tie sign. If $f_0(x_t) = f_0(y_t)$, then by (a) also $f_1(x_t) = f_1(y_t)$; both oracles consult $\mathcal{T}$, which — being blind — sees the same history (inductive hypothesis) and the same coupled coins, and so returns the same answer. Either way the step-t answers agree, and the algorithm (same coins, same history) issues the same next pair or halts with the same output. By induction the whole realized transcript and the output coincide pointwise; taking expectations over the shared randomness, the transcript and output distributions on f_0 and f_1 are equal.

(c) $\phi(s) = as$ with $a > 0$ is strictly increasing on all of $\mathbb{R} \supseteq \mathrm{range}(f_0)$, so (a) and (b) apply. □

Lemma G.2 is the noiseless counterpart of Lemma B.5: there the confounding orbit is the two-parameter rescaling $(\beta, f) \sim (c\beta, f/c)$ of the observable potential; here, with no scale in the model at all, the orbit is the full group of increasing reparametrizations $f \mapsto \phi \circ f$ — vastly larger, which is why the noiseless conclusions are absolute rather than range-conditioned.

G.2 Impossibility of gradient-norm testing

All three constructions of this appendix are ridges: $f(x) = h(t(x))$ for a scalar profile h and the linear coordinate $t(x) := \langle e_1, x - \bar{x} \rangle$, where e_1 is the first coordinate vector (the same device as Step I.1 of Appendix F.1). Here $t(\bar{x}) = 0$ and $|t(x) - t(y)| \leq \|x - y\|$, with equality when $x - y$ is parallel to e_1. For any smooth $h \colon \mathbb{R} \to \mathbb{R}$ the chain rule gives the *ridge identities*

$$\nabla f(x) = h'(t(x))\, e_1, \qquad \nabla^2 f(x) = h''(t(x))\, e_1 e_1^\top, \qquad D^3 f(x) = h'''(t(x))\, e_1^{\otimes 3}. \tag{G.2}$$

Consequently $\|\nabla f(\bar{x})\| = |h'(0)|$; the gradient map $x \mapsto h'(t(x))e_1$ is Lipschitz with constant exactly $\sup_\tau |h''(\tau)|$ (attained along e_1); the Hessian map is Lipschitz with constant $\sup_\tau |h'''(\tau)|$; and $\inf_x f = \inf_{\mathbb{R}} h$ since t is onto $\mathbb{R}$. We use Equation (G.2) throughout without comment; the constant extension along $e_1^\perp$ costs nothing.

Testing whether the gradient at a point is small asks a sign oracle to read a magnitude it has already discarded; the first theorem makes this exact.

Theorem G.3 (Impossibility of gradient-norm testing, noiseless oracle)**.** *Fix $n \geq 1$, constants $L_1, \Delta, \varepsilon > 0$, a gap ratio $c' > 1$, a point $\bar{x} \in \mathbb{R}^n$, and any $L_2 \geq 0$. Consider testing $\|\nabla f(\bar{x})\| \leq \varepsilon$ (ACCEPT) against $\|\nabla f(\bar{x})\| \geq c'\varepsilon$ (REJECT) over $f \in \mathcal{F}(L_1, L_2, \Delta)$ under the noiseless comparison oracle (Definition G.1).*

(a) (Impossibility.) *If $c'\varepsilon \leq \sqrt{2L_1\Delta}$ (i.e. exactly under Assumption 2.3), there exist $f_0, f_1 \in \mathcal{F}(L_1, L_2, \Delta)$ with $\|\nabla f_0(\bar{x})\| \leq \varepsilon$, $\|\nabla f_1(\bar{x})\| \geq c'\varepsilon$, and $f_1 = c' f_0$. By Lemma G.2 they are comparison-indistinguishable, so no comparison verifier — at any query budget, under any blind tie rule — accepts f_0 and rejects f_1 with success probability $> 1/2$ on both.*

(b) (Vacuity.) *If $c'\varepsilon > \sqrt{2L_1\Delta}$, then no $f \in \mathcal{F}(L_1, L_2, \Delta)$ has $\|\nabla f(\bar{x})\| \geq c'\varepsilon$ at all, since smoothness and the gap force $\|\nabla f(\bar{x})\|^2 \leq 2L_1\big(f(\bar{x}) - \inf f\big) \leq 2L_1\Delta$; the* REJECT *class is empty and the always-accept rule is vacuously correct.*

Hence testing is impossible if and only if the REJECT *class is nonempty, i.e. iff $c'\varepsilon \leq \sqrt{2L_1\Delta}$: no parameter window is both non-vacuous and solvable. The witnesses have vanishing third derivative, so the statement holds for every $L_2 \geq 0$, and verbatim in $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \leq \Delta\}$ with no Hessian condition; with unbounded gap the linear ridge makes it impossible for all (ε, c').*

The class in the last sentence is exactly the class of Theorem 5.1: the theorem above is the $\rho = \infty$ wall behind Remark H.2.

Proof. Witness. Take $f_0(x) = h_0(t(x))$ with $h_0(\tau) = \varepsilon\tau + \frac{L_1}{2c'}\tau^2$ and $f_1 := c' f_0$; the leading coefficient $L_1/(2c') > 0$, so h_0 is a strictly convex parabola.

(i) Gradients at $\bar{x}$. $h_0'(\tau) = \varepsilon + (L_1/c')\tau$, so by Equation (G.2) $\nabla f_0(\bar{x}) = \varepsilon\, e_1$ (norm ε, the accept boundary) and $\nabla f_1(\bar{x}) = c'\varepsilon\, e_1$ (norm $c'\varepsilon$, the reject boundary); (i) holds with equality.

(ii) Smoothness. $h_0'' = L_1/c'$ and $h_0''' = 0$. By Equation (G.2), f_0 has gradient-Lipschitz constant $L_1/c' \leq L_1$ and f_1 has gradient-Lipschitz constant L_1, both exactly; the Hessians are constant, so both third derivatives vanish and every $L_2 \geq 0$ is met. Thus L_2 never enters, and the construction lies in $\mathcal{F}(L_1, L_2, \Delta)$ for every $L_2 \geq 0$ (and verbatim in $\{f \in \mathcal{F}(L_1) : f(\bar{x}) - \inf f \leq \Delta\}$).

(iii)–(iv) Gaps and membership. h_0 is minimized at $\tau^\star = -c'\varepsilon/L_1$ with value $-c'\varepsilon^2/(2L_1)$, attained since t is onto; hence

$$f_0(\bar{x}) - \inf f_0 = \frac{c'\varepsilon^2}{2L_1}, \qquad f_1(\bar{x}) - \inf f_1 = c' \cdot \frac{c'\varepsilon^2}{2L_1} = \frac{(c'\varepsilon)^2}{2L_1}.$$

Because $c' > 1$, the gap of f_0 is the smaller, and the gap of f_1 is $\leq \Delta$ iff $(c'\varepsilon)^2 \leq 2L_1\Delta$, i.e. iff the feasibility condition of Assumption 2.3, $c'\varepsilon \leq \sqrt{2L_1\Delta}$, holds. Under it both gaps are $\leq \Delta$, so $f_0, f_1 \in \mathcal{F}(L_1, L_2, \Delta)$ for every $L_2 \geq 0$.

(v) Indistinguishability and the contradiction. Since $f_1 = c' f_0$ with $c' > 0$, Lemma G.2(c) makes every randomized adaptive comparison algorithm — any budget, any blind tie rule — produce identical transcript and output distributions on f_0 and f_1. Let $p := \mathbb{P}[A \text{ accepts } f_0] = \mathbb{P}[A \text{ accepts } f_1]$. But f_0 is an accept instance ($\|\nabla f_0(\bar{x})\| = \varepsilon$) and f_1 a reject instance ($\|\nabla f_1(\bar{x})\| = c'\varepsilon$). In the $2/3$ formulation, correctness needs $p \geq 2/3$ and $1 - p \geq 2/3$, impossible; in the $(\frac{1}{2} + \gamma)$ formulation, $p \geq \frac{1}{2} + \gamma$ and $1 - p \geq \frac{1}{2} + \gamma$ sum to $1 \geq 1 + 2\gamma$, forcing $\gamma \leq 0$. No comparison algorithm beats a coin flip on both instances, at any budget. This is part (a).

(vi) Vacuity. For L_1-smooth f the descent lemma gives, for all x, y,

$$f(y) - f(x) - \langle \nabla f(x), y - x\rangle = \int_0^1 \langle \nabla f(x + s(y - x)) - \nabla f(x), y - x\rangle\, ds \leq \frac{L_1}{2}\|y - x\|^2.$$

Taking $y = x - \nabla f(x)/L_1$ and using $f \geq \inf f$,

$$\|\nabla f(x)\|^2 \leq 2L_1\big(f(x) - \inf f\big). \tag{G.3}$$

At $x = \bar{x}$ with $f(\bar{x}) - \inf f \leq \Delta$ this reads $\|\nabla f(\bar{x})\| \leq \sqrt{2L_1\Delta}$. So if $c'\varepsilon > \sqrt{2L_1\Delta}$ no f in the class has $\|\nabla f(\bar{x})\| \geq c'\varepsilon$: the reject class is empty and the always-accept tester is vacuously correct. This is part (b).

(vii) Normalization. Throughout, $x_0 = \bar{x}$: the class normalizes the gap at the very point to be verified, matching the anchor convention of Section 2.

Dichotomy. By (iii)–(iv) the reject class is nonempty exactly when Assumption 2.3 holds — witnessed by f_1, including the boundary $(c'\varepsilon)^2 = 2L_1\Delta$ where its gap equals Δ — and by Equation (G.3) it is empty when it fails. Testing is impossible when the reject class is nonempty and vacuous when it is empty; no parameter window is both non-vacuous and solvable.

The $\Delta = \infty$ corollary. In $\mathcal{F}(L_1)$ (no gap constraint) take the linear ridge $f_0(x) = \varepsilon\, t(x)$, of gradient-Lipschitz constant 0, and $f_1 = c' f_0$: the feasibility condition is vacuous and Lemma G.2(c) yields impossibility for every pair (ε, c'). □

Remark G.4 (General evaluation points). For a general evaluation point $x_0 \neq \bar{x}$ the vacuity direction and the exact threshold are specific to the convention $x_0 = \bar{x}$ (Section 2, convention (i)); the impossibility direction is expected to persist (recenter the witnesses so that both x_0 and $\bar{x}$ carry the required gaps), but we do not prove this here.

G.3 No completeness for critical points

Rescaling changes a gradient's size; an increasing reparametrization can annihilate it entirely.

Theorem G.5 (No completeness for critical points, noiseless oracle). *Fix $n \geq 1$, $L_1, L_2, \Delta > 0$, and $\varepsilon > 0$ with $\varepsilon \leq 6^{-1/3} \min\big(\sqrt{L_1\Delta},\, (L_2\Delta^2)^{1/3}\big)$. There exist $f_0, f_1 \in \mathcal{F}(L_1, L_2, \Delta)$ with identical comparison behavior ($f_0 = \psi \circ f_1$ for a strictly increasing ψ) such that*

$$\nabla f_0(\bar{x}) = 0 \quad \textit{while} \quad \|\nabla f_1(\bar{x})\| \geq \varepsilon;$$

moreover $\nabla^2 f_0(\bar{x}) = 0$ and $\bar{x}$ is a degenerate, non-extremal critical point of f_0 (monkey-saddle type). Consequently every verifier that is complete for exact critical points *— accepting every f with $\nabla f(\bar{x}) = 0$ with probability $\geq 2/3$ — must also accept some f_1 with $\|\nabla f_1(\bar{x})\| \geq \varepsilon$, and so is unsound. The completeness class of any sound stationarity verifier must exclude some exact critical points.*

Proof. Use the ridge identities Equation (G.2) with $t(x) = \langle e_1, x - \bar{x}\rangle$.

Witness. Put $f_1(x) = g(t(x))$ with $g(\tau) = \varepsilon\ell \tanh(\tau/\ell)$, and $f_0 = \psi \circ f_1$ with $\psi(s) = \alpha s^3$, $\alpha > 0$; equivalently $f_0(x) = h_0(t(x))$, $h_0(\tau) = \alpha(\varepsilon\ell)^3 \tanh^3(\tau/\ell)$. Both are C^∞. The scale $\ell = \Delta/\varepsilon$ and

$$\alpha = \min\Big(\frac{L_1}{12\varepsilon^2\Delta},\ \frac{L_2}{60\varepsilon^3},\ \Delta^{-2}\Big) > 0$$

are fixed below.

ψ is strictly increasing (despite $\psi'(0) = 0$): for $s_1 < s_2$,

$$\psi(s_2) - \psi(s_1) = \alpha(s_2 - s_1)\big(s_2^2 + s_1 s_2 + s_1^2\big), \qquad s_2^2 + s_1 s_2 + s_1^2 = \big(s_1 + \tfrac{s_2}{2}\big)^2 + \tfrac{3}{4}s_2^2 \geq 0,$$

with equality only at $s_1 = s_2 = 0$ (excluded), so $\psi(s_2) > \psi(s_1)$; thus ψ is strictly increasing on all of $\mathbb{R} \supseteq \mathrm{range}(f_1)$.

(i) $\nabla f_0(\bar{x}) = 0$. $h_0'(\tau) = 3\alpha(\varepsilon\ell)^3 \tanh^2(\tau/\ell)\,\mathrm{sech}^2(\tau/\ell)/\ell$, so $h_0'(0) = 0$: $\bar{x}$ is an exact critical point of f_0.

(ii) $\|\nabla f_1(\bar{x})\| = \varepsilon$, since $g'(0) = \varepsilon$.

$\bar{x}$ is not a local extremum of f_0. Along e_1, $s \mapsto \alpha(\varepsilon\ell)^3 \tanh^3(s/\ell)$ is strictly increasing and vanishes at $s = 0$, so $\mathrm{sign}\big(f_0(\bar{x} + s e_1) - f_0(\bar{x})\big) = \mathrm{sign}(s)$: both signs occur in every neighborhood. Moreover $\nabla^2 f_0(\bar{x}) = h_0''(0)\, e_1 e_1^\top = 0$ (as $h_0''(0) \propto \tanh(0) = 0$): $\bar{x}$ is a degenerate, monkey-saddle-type critical point.

Derivative bounds (upper bounds suffice). With $\mathrm{sech}^2, \tanh^2 \leq 1$ and the triangle inequality,

$$|(\tanh)'| \leq 1, \quad |(\tanh)''| = |2\tanh\mathrm{sech}^2| \leq 2, \quad |(\tanh)'''| = |-2\,\mathrm{sech}^4 + 4\tanh^2\mathrm{sech}^2| \leq 6;$$

$$|(\tanh^3)'| = |3\tanh^2\mathrm{sech}^2| \leq 3, \quad |(\tanh^3)''| = |6\tanh\mathrm{sech}^4 - 6\tanh^3\mathrm{sech}^2| \leq 12,$$

$$|(\tanh^3)'''| = |6\,\mathrm{sech}^6 - 42\tanh^2\mathrm{sech}^4 + 12\tanh^4\mathrm{sech}^2| \leq 60.$$

Only these upper bounds enter — membership needs one in-class instance — and the true maxima (e.g. 6 for $|(\tanh^3)'''|$) are smaller.

Membership. By Equation (G.2),

$$f_1: \quad \sup|g''| \leq \tfrac{2\varepsilon}{\ell}, \qquad \sup|g'''| \leq \tfrac{6\varepsilon}{\ell^2}, \qquad \text{gap} = \varepsilon\ell;$$

$$f_0: \quad \sup|h_0''| \leq 12\alpha\varepsilon^3\ell, \qquad \sup|h_0'''| \leq 60\alpha\varepsilon^3, \qquad \text{gap} = \alpha(\varepsilon\ell)^3.$$

With $\ell = \Delta/\varepsilon$: for f_1, the gradient-Lipschitz constant is $\leq 2\varepsilon^2/\Delta \leq L_1$ iff $2\varepsilon^2 \leq L_1\Delta$, the Hessian-Lipschitz constant is $\leq 6\varepsilon^3/\Delta^2 \leq L_2$ iff $6\varepsilon^3 \leq L_2\Delta^2$, and the gap is Δ; for f_0, the two

Lipschitz constants are $\le 12\alpha\varepsilon^2\Delta \le L_1$ and $\le 60\alpha\varepsilon^3 \le L_2$ and the gap is $\alpha\Delta^3 \le \Delta$, all three by the choice of α (so f_0 constrains neither ℓ nor the feasibility). The two conditions on f_1 read $\varepsilon \le \frac{1}{\sqrt{2}}\sqrt{L_1\Delta}$ and $\varepsilon \le 6^{-1/3}(L_2\Delta^2)^{1/3}$; since $6^{-1/3} \approx 0.550 \le 1/\sqrt{2}$, the single hypothesis of the theorem,

$$\varepsilon \le 6^{-1/3}\min\big(\sqrt{L_1\Delta},\,(L_2\Delta^2)^{1/3}\big),$$

suffices, and then $f_0, f_1 \in \mathcal{F}(L_1, L_2, \Delta)$.

Indistinguishability. ψ is strictly increasing on $\mathbb{R} \supseteq \mathrm{range}(f_1)$ and $f_0 = \psi \circ f_1$; by Lemma G.2, f_0 and f_1 induce identical transcript and output distributions for every randomized adaptive comparison algorithm under every blind tie rule.

Impossibility. Let V be complete for exact critical points, so $\mathbb{P}[V \text{ accepts } f_0] \ge 2/3$ because $\nabla f_0(\bar{x}) = 0$. Were V sound against ε-gradient points it would satisfy $\mathbb{P}[V \text{ accepts } f_1] \le 1/3$ because $\|\nabla f_1(\bar{x})\| = \varepsilon$. But Lemma G.2 forces the two probabilities equal, so $2/3 \le p \le 1/3$ — a contradiction. Hence the completeness class of every sound verifier must exclude the critical point $\bar{x}$ of f_0.

No clash with the positive results of this paper. The excluded $\bar{x}$ is a degenerate, non-extremal critical point (f_0 increases strictly through $\bar{x}$ along e_1, and $\nabla^2 f_0(\bar{x}) = 0$), hence not a strict local minimum: it carries no value margin at any scale, so it lies in the accept class $M^+(m_0)$ of Theorem 6.1 for no $m_0 > 0$, and its exclusion is consistent with value-margin verification (whose $\beta = \infty$ limit is noiseless). The verifiable core is order structure — value margins — not criticality. □

The one-dimensional cartoon is $t \mapsto t^3$ against $t \mapsto t$: both increase through 0, so every comparison agrees, yet one is critical there and the other is not.

G.4 Selection is impossible

The same reparametrization device defeats a task far weaker than testing: merely selecting a stationary candidate from a list.

Theorem G.6 (Selection is impossible, noiseless oracle)**.** *Fix $n \ge 1$, $L_1, L_2, \Delta > 0$, $\varepsilon > 0$, and $c' > 1$ satisfying*

$$c'^2\varepsilon^2 \le \frac{M_1^2}{16M_2}\,\Delta L_1 \qquad \text{and} \qquad c'^{3/2}\varepsilon^{3/2} \le \frac{M_1^{3/2}}{8\sqrt{M_3}}\,\Delta\sqrt{L_2},$$

where M_1, M_2, M_3 are the (universal) derivative sup-norms of a fixed smooth step profile; this window is nonempty for every $c' > 1$. There exist $f_0, f_1 \in \mathcal{F}(L_1, L_2, \Delta)$ with identical comparison behavior ($f_1 = \phi \circ f_0$, ϕ strictly increasing) and points x_1, x_2 with

$$\begin{aligned} &\text{under } f_0: && \|\nabla f_0(x_1)\| = \varepsilon, && \|\nabla f_0(x_2)\| = 2c'\varepsilon; \\ &\text{under } f_1: && \|\nabla f_1(x_1)\| \ge 2c'(1-\vartheta)\varepsilon > c'\varepsilon, && \|\nabla f_1(x_2)\| \le \tfrac{1+\vartheta}{2}\,\varepsilon \le \varepsilon, \end{aligned}$$

where $\vartheta = \max\big((8c')^{-3}, (64c')^{-1}\big) \le 1/8$. Hence for the selection *problem — given $J \ge 2$ candidate points with the promise that at least one is ε-stationary, output one that is $c'\varepsilon$-stationary — every randomized adaptive comparison algorithm, at any budget, fails with probability $\ge 1/2$ on one of f_0, f_1: the promise holds for both (via x_1 for f_0, x_2 for f_1), but the correct output sets $\{x_1\}$ and $\{x_2\}$ are disjoint. The extension to any $J \ge 2$ is by padding with dummies that are $c'\varepsilon$-non-stationary under both functions.*

Proof. Let $\sigma(z) = 1/(1+e^{-z})$ be the logistic function (the link of Definition 2.1), with $\|\sigma'\|_\infty = 1/4$ and $\|\sigma''\|_\infty = 1/(6\sqrt{3})$. Fix a C^∞ nondecreasing smoothstep $R\colon \mathbb{R} \to [0,1]$ with $R \equiv 0$ on $(-\infty, 0]$ and $R \equiv 1$ on $[1, \infty)$, and set $M_k := \sup_{\mathbb{R}}|R^{(k)}|$. A single ramp $\tau \mapsto hR((\tau - p)/W)$ has k-th derivative sup hM_k/W^k and support $[p, p+W]$; ramps on disjoint windows have identically zero cross-terms.

Step 0 (reduction to one dimension). Write $f_0(x) = g(t(x))$ with $t(x) = \langle e_1, x - \bar{x}\rangle$ and $f_1 = \phi \circ f_0$. By Equation (G.2), every derivative norm of f_0, f_1 equals that of the 1-d profiles g and $h := \phi \circ g$, and $\mathrm{range}(f_0) = \mathrm{range}(g)$.

Step 1 (the profile: three ramps on disjoint windows $p_1 < p_c < p_2$).

$$g = h_1 R\Big(\tfrac{\tau - p_1}{W}\Big) + R_c R\Big(\tfrac{\tau - p_c}{W_c}\Big) + h_2 R\Big(\tfrac{\tau - p_2}{W}\Big), \qquad h_1 = \frac{\varepsilon W}{M_1}, \quad h_2 = \frac{2c'\varepsilon W}{M_1}.$$

The peak slopes are ε (ramp 1), $s_{\text{small}} \coloneqq R_c M_1 / W_c$ (the connector), and $2c'\varepsilon$ (ramp 2). In value space, from the bottom: ramp 1 sweeps the band $I_1 = [v_{\text{mid}} - R_c/2 - h_1,\ v_{\text{mid}} - R_c/2]$, the connector sweeps a gap of height R_c centered at v_{mid}, and ramp 2 sweeps $I_2 = [v_{\text{mid}} + R_c/2,\ v_{\text{mid}} + R_c/2 + h_2]$. Let x_1 be the steepest point of ramp 1 (at level $a_1 \in I_1$) and x_2 that of ramp 2 (level $a_2 \in I_2$); then $a_1 < a_2$ and every value in $I_1 \cup I_2$ is at distance $\geq R_c/2$ from v_{mid}. Off the three windows g is constant.

Step 2 (reparametrization and plateau lemma). Set

$$\phi'(v) = \frac{1}{4c'} + \Big(2c' - \frac{1}{4c'}\Big)\sigma\Big(\frac{v_{\text{mid}} - v}{w_0}\Big), \qquad w_0 \coloneqq \frac{R_c}{6\ln(8c')}, \qquad \phi(v) = a_1 + \int_{a_1}^{v} \phi'.$$

Since $2c' - 1/(4c') > 0$ and $\sigma > 0$, $\phi'(v) \geq 1/(4c') > 0$ everywhere, so ϕ is strictly increasing on $\mathbb{R} \supseteq \text{range}(f_0)$. The *plateau lemma* states: on I_1, $\phi' \in [2c'(1-\vartheta_1),\ 2c']$ with $\vartheta_1 = (8c')^{-3}$; on I_2, $\phi' \in [1/(4c'),\ (1/(4c'))(1+\vartheta_2)]$ with $\vartheta_2 = 1/(64c')$. Indeed, on I_1 the argument $u = (v_{\text{mid}} - v)/w_0 \geq (R_c/2)/w_0 = 3\ln(8c')$, so $\sigma(u) \geq 1 - e^{-u} \geq 1 - (8c')^{-3}$ and hence $\phi' \geq 2c'(1-(8c')^{-3})$; on I_2, $u \leq -3\ln(8c')$, so $\sigma(u) \leq (8c')^{-3}$ and $\phi' \leq 1/(4c') + 2c'(8c')^{-3} = (1/(4c'))(1 + 1/(64c'))$, using $2c'(8c')^{-3} = 1/(256c'^2)$. Write $\vartheta \coloneqq \max(\vartheta_1, \vartheta_2) \leq 1/8$. The logistic's dynamic range $8c'^2$ (from $1/(4c')$ up to $2c'$) forces the tail to beat $\vartheta/(8c'^2)$; the width $w_0 = R_c/(6\ln 8c')$ delivers a tail $(8c')^{-3}$, which does, whereas a c'-independent width $R_c/6$ would leave a tail $e^{-3} \approx 0.05$ that fails for large c'.

Step 3 (Lemma G.2 applies). ϕ is strictly increasing on $\text{range}(f_0)$, so by Lemma G.2 the pair f_0 and $f_1 = \phi \circ f_0$ is comparison-indistinguishable: identical output distributions for every randomized adaptive comparison algorithm under the blind tie-rule coupling.

Step 4 (class membership, region by region). By the chain rule,

$$h'' = \phi''(g)\, g'^2 + \phi'(g)\, g'', \qquad h''' = \phi'''(g)\, g'^3 + 3\phi''(g)\, g' g'' + \phi'(g)\, g'''.$$

Everywhere $\|\phi'\|_\infty \leq 2c'$; off the connector $|\phi''|, |\phi'''|$ are at most $(8c')^{-3}$ times their connector peaks — the *leak*. We bound $|h''|$ and $|h'''|$ on each region.

- *Ramp* 1 ($\phi' \leq 2c'$, the binding amplification): by Equation (G.2), $|h''| \leq 2c'\varepsilon M_2/(M_1 W) + \text{leak} \leq L_1/2 + \text{leak}$ for $W \geq 4M_2 c'\varepsilon/(M_1 L_1)$, and $|h'''| \leq 2c'\varepsilon M_3/(M_1 W^2) + \text{leak} \leq L_2/2 + \text{leak}$ for $W \geq 2\sqrt{M_3 c'\varepsilon/(M_1 L_2)}$.
- *Ramp* 2 ($\phi' \leq (1/(4c'))(1+\vartheta_2)$ attenuates): the binding term is g itself, $|g''| = 2c'\varepsilon M_2/(M_1 W) \leq L_1$ and $|g'''| = 2c'\varepsilon M_3/(M_1 W^2) \leq L_2$ — weaker than ramp 1's constraints.
- *Connector* (ϕ'', ϕ''' peak here, but $|g'| \leq s_{\text{small}}$): the peaks are $|\phi''| \leq c'/(2w_0)$ and $|\phi'''| \leq c'/(3\sqrt{3}\, w_0^2)$. Choosing W_c so that $s_{\text{small}} \leq \frac{1}{2}\min\{\sqrt{L_1 w_0/c'},\ (L_2 w_0^2/c')^{1/3}\}$ gives $(c'/(2w_0))\, s_{\text{small}}^2 \leq L_1/8$ and $(c'/(3\sqrt{3}\, w_0^2))\, s_{\text{small}}^3 \leq L_2/24$; the mixed term $3\phi'' g' g''$ carries an extra factor $w_0/R_c = 1/(6\ln 8c') < 1$. In total $|h''| \leq L_1$ and $|h'''| \leq L_2$ on the connector — the crux, discharged because ϕ'' is supported exactly where $|g'|$ is tiny.
- *Far field*: g is constant and all derivatives vanish.

Take $W \coloneqq \max\{4M_2 c'\varepsilon/(M_1 L_1),\ 2\sqrt{M_3 c'\varepsilon/(M_1 L_2)}\}$.

Step 5 (range and feasibility). Set $R_c \coloneqq h_1 = \varepsilon W/M_1$, so every leak term is $\leq (8c')^{-3}$ times its peak and is absorbed in the slack above. Then

$$\text{range}(f_0) = h_1 + R_c + h_2 = (2 + 2c')\frac{\varepsilon W}{M_1} \leq \frac{4c'\varepsilon W}{M_1},$$

$$\text{range}(f_1) = 2c'\, h_1 + \Big(c' R_c + \frac{R_c}{8c'}\Big) + \frac{h_2}{4c'} = \Big(3c' + \frac{1}{8c'} + \frac{1}{2}\Big)\frac{\varepsilon W}{M_1} \leq \frac{4c'\varepsilon W}{M_1},$$

the connector contributing $\int \phi'$ over the value-gap, namely $R_c/(4c') + (2c' - 1/(4c'))R_c/2$ (using $\int_{-a}^{a} \sigma = a$). Both ranges are $\le \Delta$ once $4c'\varepsilon W/M_1 \le \Delta$; substituting the two branches of W gives exactly the feasibility window of the statement,

$$c'^2\varepsilon^2 \le \frac{M_1^2}{16M_2}\,\Delta L_1 \quad \text{and} \quad c'^{3/2}\varepsilon^{3/2} \le \frac{M_1^{3/2}}{8\sqrt{M_3}}\,\Delta\sqrt{L_2},$$

which is nonempty for every $c' > 1$ (both left sides $\to 0$ as $\varepsilon \to 0$). Hence $f_0, f_1 \in \mathcal{F}(L_1, L_2, \Delta)$.

Step 6 (gradient facts; promise; disjointness). Disjoint supports give the exact peak slopes $\|\nabla f_0(x_1)\| = \varepsilon$ and $\|\nabla f_0(x_2)\| = 2c'\varepsilon$. By the plateau lemma ($\vartheta \le 1/8$),

$$\|\nabla f_1(x_1)\| = \phi'(a_1)\,\varepsilon \ge 2c'(1-\vartheta)\varepsilon > c'\varepsilon \quad (\text{since } 2(1-\vartheta) > 1),$$
$$\|\nabla f_1(x_2)\| = \phi'(a_2)\,2c'\varepsilon \le \frac{1+\vartheta}{2}\,\varepsilon \le \varepsilon.$$

The promise — at least one candidate is ε-stationary — holds for both instances (x_1 for f_0, x_2 for f_1), while the correct output sets — candidates that are $c'\varepsilon$-stationary — are $\{x_1\}$ for f_0 and $\{x_2\}$ for f_1, and these are disjoint.

Step 7 (impossibility, $J = 2$). By Step 3 the selector's output distribution p over the two candidates is identical on f_0 and f_1. The unique correct candidate is x_1 on f_0 and x_2 on f_1; since $p(x_1)+p(x_2) \le 1$, one of $p(x_1), p(x_2)$ is $\le 1/2$, so on the corresponding instance the selector returns the correct candidate with probability $\le 1/2$ and fails with probability $\ge 1/2$ — at any query budget.

Step 8 (padding to $J \ge 2$). For $J > 2$ candidates, add one further ramp of slope $2c'\varepsilon$ at a low level (inside the I_1 region, where $\phi' \approx 2c'$) and place all $J - 2$ dummy candidates in its core. Then $\|\nabla f_0(d_j)\| = 2c'\varepsilon > c'\varepsilon$ and $\|\nabla f_1(d_j)\| \approx 4c'^2\varepsilon > c'\varepsilon$: every dummy is $c'\varepsilon$-non-stationary under both functions, and the correct sets are unchanged. The extra ramp tightens the feasibility window to $c'^4\varepsilon^2 \lesssim \Delta L_1$ and $c'^3\varepsilon^{3/2} \lesssim \Delta\sqrt{L_2}$ — still nonempty — and Step 7 applies verbatim. □

The selector faces two worlds that produce identical transcripts but whose unique correct answers are different points; no rule can output the right one in both. In particular the natural repair — return the candidate of least value — fails too:

Corollary G.7 (The min-value heuristic fails)**.** *In the construction of Theorem G.6, ϕ preserves the order of candidate values and places x_1 at the lower value under both f_0 and f_1. Hence returning the least-value candidate outputs x_1 in both worlds — correct for f_0, but wrong for f_1, where $\|\nabla f_1(x_1)\| > c'\varepsilon$. The natural "return the minimum-value iterate" repair is provably unsound.*

Proof. Use the construction in the proof of Theorem G.6. Since $a_1 < a_2$ and the reparametrization ϕ is increasing, $f_0(x_1) < f_0(x_2)$ and $f_1(x_1) < f_1(x_2)$; with the $J - 2$ dummies placed at levels above a_1, the candidate of least value is x_1 under both f_0 and f_1. Returning it is correct on f_0 (where $\|\nabla f_0(x_1)\| = \varepsilon \le c'\varepsilon$) but wrong on f_1 (where $\|\nabla f_1(x_1)\| \ge 2c'(1-\vartheta)\varepsilon > c'\varepsilon$). Hence no minimum-value rule solves the selection problem. □

H Further discussion and open problems

H.1 The dimension cost of adaptive verification

At constant confidence, the verification bounds have dimension factors $[1, n^2]$ for adaptive methods and $[n, n^2]$ for nonadaptive methods in the regime $\beta\varepsilon^2/L_1 \le 1/64$. The first question is whether every adaptive verifier also requires $\Omega(nL_1^2/(\beta^2\varepsilon^4))$ comparisons. A complementary question is whether the quadratic upper bound can be improved.

The proof of Theorem 4.2 identifies where the present method stops. It averages pairwise likelihood overlaps over two independent directions in the hard ensemble. For fixed queries, spherical symmetry controls the off-diagonal terms. With adaptive queries, the observed history also changes the query direction, and the likelihood tilt must remain inside the averaging argument. Taking a worst-case query bound first loses this cancellation and recovers only the dimension-free two-point estimate. A

useful extension would control information gained per query under the evolving mixture rather than replace that mixture by its worst pair. Remark E.4 describes the obstruction in the proof's notation.

The confidence and dimension contributions also need to be distinguished. Corollary E.5 gives a maximum of a logarithmic confidence term and a linear dimension term. It does not give their product. The crossover occurs near $n = 1 + 849\ln(1/\delta)$; for $\delta = 1/16$ this is about 2356.

H.2 Search: the gap budget and higher smoothness

The search construction has its own budget: $\Delta_{\mathrm{inst}} = 10\varepsilon^2/L_1$ for $\mathcal{F}(L_1)$ and $\Delta_{\mathrm{inst}} = (32/3)\varepsilon^{3/2}/\sqrt{L_2}$ for $\mathcal{F}(L_1, L_2)$. A larger allowed budget enlarges the class and preserves the lower bound, but this monotonicity does not establish a rate growing with Δ. Gap-dependent lower bounds in gradient and stochastic-gradient models (Carmon et al., 2020; 2019; Arjevani et al., 2023) motivate asking for an analogous dependence under comparison access.

A direct-sum construction would need a new information argument. Simply chaining $\Delta/\Delta_{\mathrm{inst}}$ wells and charging the present single-well cost at each one would formally suggest $\Omega(n\Delta L_1^3/(\beta^2\varepsilon^6))$ in the gradient-Lipschitz class. That accounting is not justified: an unrestricted comparison can span different parts of a construction, so independent per-well charges need not add. Nor can the confidence-score upper bound of El Bakkali et al. (2026a) settle this question for raw BT bits, because its observation assumptions differ. A valid result needs to track the information of the full comparison transcript.

Higher smoothness raises a separate question. The ε^{-4} and ε^{-3} exponents in Theorem 3.2 arise from different admissible cap radii. They do not yet separate the complexities of $\mathcal{F}(L_1)$ and $\mathcal{F}(L_1, L_2)$. An upper bound of order $\mathrm{poly}(n)/\varepsilon^3$ over $\mathcal{F}(L_1, L_2)$ would establish such a separation in accuracy; a stronger lower bound over $\mathcal{F}(L_1, L_2)$ could show that the current cubic-cap construction is not sharp. Either direction requires more than a comparison of the present lower bounds.

H.3 Approaching the calibration boundary

For $\rho < c'$, Theorem 5.2 settles how the unknown-scale cost grows near the boundary in the regime $\beta_{\min} \le \beta^\dagger$: it is of order $L_1^2/(\beta_{\min}^2\varepsilon^4(c'-\rho)^4)$, up to a factor polynomial in n. The lower bound comes from two instances whose logit potentials differ by a single step of height proportional to $(c'-\rho)^2$, not from instances with proportional logits; proportional logits give no uniform bound, because their difference is unbounded over the query space. The saturated regime remains open. There the tester of Proposition F.3 costs order $n^2\rho^2 L_1\Delta\log(2n/\delta)/(\varepsilon^2(c'-\rho)^2)$ for all large $\beta_{\min}$, while the lower bound of Theorem 5.2 decays like $\beta_{\min}^{-2}$. A lower bound that grows as $\rho \to c'$ uniformly in $\beta_{\min}$ would need to exploit the gap budget, since without it the amplitude grows linearly in $\beta_{\min}$ (Remark C.3).

Remark H.1 (Why an exact budget changes identifiability). The unknown-scale impossibility uses the class $f(\bar{x}) - \inf f \le \Delta$. If the budget is fixed exactly at $\Delta > 0$, the observable potential $\beta(f - f(\bar{x}))$ has infimum $-\beta\Delta$. Full knowledge of that potential therefore identifies β. This is an identification statement at unbounded information, not a uniform finite-query guarantee. The two scaled witnesses of Theorem 5.1 have different budgets and do not establish impossibility for the exactly calibrated class.

Remark H.2 (Finite-scale verification and the noiseless oracle). The known-scale tester changes its probe radius as β grows, so its finite-scale guarantee cannot be transferred by simply taking a pointwise limit at fixed pairs. For the noiseless sign oracle, monotone transformations preserve all answers (Lemma G.2), and gradient-norm verification is impossible (Theorem G.3). These are distinct observation models; no convergence of the unknown-scale experiment to the noiseless experiment is claimed by letting $\rho \to \infty$.

H.4 Value margins and gradient guarantees

Theorem 6.1 treats m_0 as an independent promise parameter. Its sample count therefore carries no dimension or smoothness factor. The specialization in Corollary F.7 is different: choosing $r = \varepsilon/(2\sqrt{n}L_1)$ supplies a constant mass of descent directions when $\|\nabla f(\bar{x})\| \ge \varepsilon$, but the compatible margin is at most $\varepsilon^2/(8nL_1)$. Inserting that margin into the ordinal count recovers a

factor of order $n^2 L_1^2/(\beta^2\varepsilon^4)$ in the noise-limited regime. The ordinal result consequently does not bypass the dimension cost of gradient verification, and its completeness promise remains the value-margin class.

Remark H.3 (The smoothness-imposed margin ceiling). For $f \in \mathcal{F}(L_1)$, sphere averaging gives $\mathbb{E}_\mu[\mathrm{gap}] \le (L_1/2)r^2 =: m_{\mathrm{s}}$: the linear Taylor term averages to zero and the remainder is at most $(L_1/2)r^2$. Thus $\mathcal{M}^+(m_0) \cap \mathcal{F}(L_1)$ is nonempty only if $0 < m_0 \le m_{\mathrm{s}}$. At the radius in Corollary F.7, this is $m_{\mathrm{s}} = \varepsilon^2/(8nL_1)$. No such ceiling is imposed in the unrestricted measurable class of Theorem 6.1.

H.5 Other observation models

The packing and two-point arguments use the BT link through its global quadratic KL bound. They extend to another link if an appropriate global bound is established. A local Fisher-information calculation alone is insufficient to guarantee that bound. The nonadaptive mixture proof uses the logistic score more specifically, and its extension would require a separate analysis. Likewise, the upper bounds use inverse-link estimation and thus require knowledge of the response law.

Bounded adversarial comparison noise (Lobanov et al., 2024; Smirnov et al., 2024), uniform-margin feedback, and multiway preferences (Saha et al., 2022) change the information supplied by a query. The transcript chain rule still provides a starting point, but the relevant per-query divergence, invariances, and calibration assumptions must be derived for each model. The present rates and exact threshold should not be transferred solely by matching a noise parameter.