When Do Nonlinear Preferences Slow Submodular Minimization?

Anonymous authors
Paper under double-blind review

Abstract

We study minimization of $[0,1]$-valued submodular functions on n elements from noisy comparisons between sets that differ by at most one element, which we call 1-local queries. For every fixed sigmoid strength $b > 0$, we establish the minimax rate $\Theta_b(\min\{1, n^{3/2}/\sqrt{T}\})$. Our algorithm uses sequential unbiased log-odds estimates and at most T comparisons; the previous upper bound was $O_b(n^{7/5}T^{-2/5})$. More generally, let ψ be the inverse of the odd part of a known transfer function ρ. Under regularity and shape conditions on ψ, we prove the minimax rate $\Theta_\rho(\min\{1, n\psi_{\mathrm{cl}}(\sqrt{n/T})\})$ jointly in n and T, where ψ_{cl} extends ψ by the constant one above its domain. A curvature bound for antithetic multilevel estimation gives the upper bound for power, logarithmic, and infinite-order flatness. Within this class, the $T^{-1/2}$ rate is equivalent to the existence of unbiased marginal estimators with uniformly bounded expected cost and second moment. Beyond this class, an analytic transfer function with positive derivative at zero can have minimax rate $\Theta(T^{-1/6})$ on two elements because 1-local queries encounter flatness away from ties. For arbitrary transfer functions and fixed $n \geq 2$, the inverse's modulus of continuity characterizes the rate.

1 Introduction

Submodular function minimization with a dueling oracle uses noisy pairwise comparisons in place of function values. A transfer function ρ maps the difference between two function values to the mean comparison response. Under 1-local queries, Sho & Ito (2026) give the optimal $O(n^{3/2}/\sqrt{T})$ rate for a linear transfer function and an $O(n^{7/5}T^{-2/5})$ upper bound for a sigmoid. Here n is the ground-set size, T is the number of comparisons, and a 1-local query compares sets that differ by at most one element. Does the sigmoid model require a slower rate, or can a different algorithm close the gap?

Optimal rates for sigmoid preferences. Using only 1-local queries, we match the sigmoid lower bound of Sho & Ito (2026), which applies to all algorithms with $|A \triangle B| = O(1)$ (their Restriction 1). This confirms their conjectured $T^{-1/2}$ dependence. For every fixed strength $b > 0$, the minimax rate is $\Theta_b(\min\{1, n^{3/2}/\sqrt{T}\})$. We estimate each marginal from a random number of comparisons, obtaining unbiased estimates with uniformly bounded expected cost and second moment. A cap on the complete optimization run enforces the budget of T comparisons. The same rate holds under the inverse regularity conditions in Theorem 2, which also cover probit preferences.

A joint minimax rate. We then determine how the transfer function changes the optimal dependence on n and T. Let $\Gamma(t) = [\rho(t) - \rho(-t)]/2$ be its odd part, and let $\psi = \Gamma^{-1}$. Let ψ_{cl} agree with ψ on $[0, \Gamma(1)]$ and equal one for larger arguments. Write $R^{(1)}_{n,T}(\rho)$ for the minimax expected suboptimality under 1-local queries. Under the inverse regularity and shape conditions in Assumption 3,

$$R^{(1)}_{n,T}(\rho) = \Theta_\rho\left(\min\left\{1, n\psi_{\mathrm{cl}}\left(\sqrt{n/T}\right)\right\}\right). \tag{1}$$

The constants depend on the fixed transfer function, but not on n or T.

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.

Odd part Γ	Previous upper bound	Optimal rate
Linear	$O(n^{3/2}T^{-1/2})$	$\Theta(n^{3/2}T^{-1/2})$
Sigmoid	$O_b(n^{7/5}T^{-2/5})$	$\Theta_b(n^{3/2}T^{-1/2})$
$\lambda \operatorname{sgn}(t)\lvert t\rvert^q$, $q \geq 1$	—	$\Theta_q(n^{1+1/(2q)}(\lambda^2 T)^{-1/(2q)})$
$\Gamma(t) = ct/\log(e/t)^a$	—	$\Theta_\rho(n^{3/2}T^{-1/2}\log(T/n)^a)$
$\operatorname{sgn}(t)e^{-\lvert t\rvert^{-a}}$	—	$\Theta_a(n/\log(e+T/n)^{1/a})$

Table 1: Expected suboptimality under 1-local queries, capped at one. The logarithmic expressions describe the small-error regime. Here $0 < \lambda < 1$ and $a > 0$. The fourth row specifies $\Gamma(t)$ for small positive t, with a continuation satisfying Assumption 3. The previous upper bounds and the linear and sigmoid lower bounds are from Sho & Ito (2026).

The lower bound uses modular (additive) functions with n unknown marginal signs, each comparison giving information about at most one sign. In the small-error regime, these hard instances have comparison signals of order $\sqrt{n/T}$. Applying ψ gives the corresponding marginal size, and errors across coordinates accumulate. The upper bound attains this scale under Assumption 3.

Unbiased estimation and the $T^{-1/2}$ rate. Within Assumption 3, a scalar marginal admits an unbiased estimator with uniformly bounded expected comparison cost and second moment if and only if $\psi(u) = O(u)$ near zero. This is also exactly the condition for the $T^{-1/2}$ rate (Theorem 5). When the inverse grows faster, we allow bias and choose the estimator according to the budget. Balancing its bias, second moment, and cost attains Equation (1).

Flatness away from ties. For general transfer functions, positive slope at zero does not ensure the $T^{-1/2}$ rate. We construct two submodular functions whose minimizers are the empty and full sets, respectively. Distinguishing them using 1-local queries requires querying singleton sets, placing every informative value gap near a nonzero flat point of the transfer function (Figure 1). An analytic transfer function with $\rho'(0) > 0$ then has minimax rate $\Theta(T^{-1/6})$ on two elements. The modulus of continuity of the inverse characterizes the budget dependence for every fixed $n \geq 2$ (Theorem 6). This result explains the role of the assumptions away from zero in the joint rate theorem.

Proof strategy. We estimate marginal differences along a chain of sets to form a subgradient of the Lovász extension. For the general rate, an antithetic multilevel estimator compares the inverse applied to a pooled sample mean with the average of its values on two half-sample means. A curvature term bounds this difference, and its normalized values telescope across levels. The resulting bounds on bias, second moment, and expected cost yield the joint rate. Independent runs, each with a comparison cap, enforce the total budget. Adaptive testing arguments give the matching lower bound.

Our experiments illustrate how biased marginal estimates can change the learned minimizer and how changing nonoptimal set values can reduce the information available from 1-local queries.

1.1 Related Work

Submodular minimization. The Lovász extension reduces submodular minimization to convex optimization (Bach, 2013). Ito (2019) studies noisy function evaluations and analyzes stochastic subgradient methods, including the effect of observing several values from the same noisy function. The dueling SFM model of Sho & Ito (2026) replaces value observations with pairwise comparisons. Their Table 1 also lists an $O(n^{4/3}T^{-1/3})$ upper bound for general transfer functions. For the regular inverses in Theorem 2, we obtain the optimal $O(n^{3/2}T^{-1/2})$ rate. For power transfer functions with $q > 3/2$, Theorem 4 rules out a $T^{-1/3}$ rate under 1-local queries. Zhang et al. (2021) obtain high-probability guarantees for stochastic discrete convex minimization with noisy value observations.

Optimization from comparisons. Jamieson et al. (2012) relate comparison reliability to optimization rates for smooth, strongly convex functions. Saha et al. (2021; 2025) study dueling convex optimization, with the latter giving rates governed by the transfer function near zero. Their continuous-domain algorithms can make value differences small by querying nearby points in a

Lipschitz objective. In SFM, changing one element can produce a large value difference, as our example in Section 5 demonstrates. Our algorithms use the known transfer function, whereas those of Saha et al. (2025) do not require it as input.

Unbiased estimation. Sequential unbiased estimation under binomial sampling has a classical literature (Girshick et al., 1946; DeGroot, 1959). Mendo (2025) gives sequential estimators of odds and log odds. Randomized multilevel estimators provide unbiased estimates of nonlinear functions of expectations (Rhee & Glynn, 2015; Blanchet & Glynn, 2015; Blanchet et al., 2019; Wang & Wang, 2023), including under positive Hölder regularity of the derivative (Blanchet et al., 2019; Wang & Wang, 2023). El Bakkali et al. (2026) use randomized expansions of analytic inverses to estimate value differences from comparisons in nonsmooth continuous optimization. Our analysis uses antithetic corrections under the stated Dini or shape conditions. The curvature estimate also gives sharp bounds for finitely many levels when uniform unbiased estimation is impossible. Bernoulli-factory results connect regularity to simulation cost (Nacu & Peres, 2005); our impossibility argument permits real-valued estimates with a bounded second moment and relates the estimation condition to the minimax optimization rate.

2 Model and Reduction

Let $[n] = \{1, \dots, n\}$. All logarithms are natural unless a base is specified. A function $f : 2^{[n]} \to [0, 1]$ is submodular if $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for all $A, B \subseteq [n]$. Write $\mathcal{F}_n$ for this class. A query (A, B) returns $O \in \{-1, 1\}$ with

$$\mathbb{P}_f(O = 1 \mid A, B) = \frac{1 + \rho(f(A) - f(B))}{2}. \tag{2}$$

The known transfer function $\rho : [-1, 1] \to (-1, 1)$ is continuous and strictly increasing, with $\rho(0) = 0$. Write $\theta = f(A) - f(B)$ for the value gap; a tie has $\theta = 0$. Each response uses fresh randomness conditional on the chosen pair. A query (A, B) is 1-local if $|A \triangle B| \leq 1$. We restrict algorithms to these queries. Algorithms may be adaptive and randomized, but their total number N of comparisons must satisfy $N \leq T$ almost surely. If $\widehat{S}_{\mathcal{A}}$ denotes the output of algorithm $\mathcal{A}$, the minimax expected suboptimality over these algorithms is

$$R^{(1)}_{n,T}(\rho) = \inf_{\mathcal{A}:\, N \leq T \text{ a.s.}} \sup_{f \in \mathcal{F}_n} \mathbb{E}_f[f(\widehat{S}_{\mathcal{A}}) - f^\star], \qquad f^\star = \min_S f(S). \tag{3}$$

The superscript (1) denotes the restriction to 1-local queries. Only comparisons are charged. Arithmetic and preprocessing of the known transfer function are outside the query count. An implementation needs an evaluable transfer function and the regularity and second-moment bounds specified in the constructions below. Constants indexed by ρ may depend on this function and its chosen extensions, but not on n or T.

Symmetrization. Let $\Gamma(t) = [\rho(t) - \rho(-t)]/2$, $a_\rho = \Gamma(1)$, and $\psi = \Gamma^{-1}$ on $[-a_\rho, a_\rho]$. One comparison can produce a sign with mean $\Gamma(\theta)$: choose a uniform sign R, reverse the query if $R = -1$, and return $Y = RO$. Thus the upper bounds can use the odd transfer function Γ. The lower bounds allow adaptive choices of the original query orientation (Lemma 7). For $u \geq 0$, set $\psi_{\mathrm{cl}}(u) = \psi(u)$ when $u \leq a_\rho$, and $\psi_{\mathrm{cl}}(u) = 1$ otherwise.

Lovász extension. For $w \in [0, 1]^n$, let π sort its coordinates in decreasing order, using a fixed tie rule. Put $B_0 = \varnothing$ and $B_j = \{\pi(1), \dots, \pi(j)\}$. The Lovász extension $\widehat{f}$ is convex and has a subgradient with coordinates

$$g_{\pi(j)}(w) = f(B_j) - f(B_{j-1}), \qquad \widehat{f}(w) = f(\varnothing) + \sum_{j=1}^{n} w_{\pi(j)} g_{\pi(j)}(w). \tag{4}$$

Each marginal is accessible through a 1-local query. We estimate these marginals, take projected stochastic subgradient steps, and average the pre-update iterates to obtain $\bar{w}$. The minimum of $\widehat{f}$ is $f^\star$. For an independent U uniform on $[0, 1]$, threshold rounding gives $\mathbb{E}_U f(\{i : w_i \geq U\}) = \widehat{f}(w)$. A bound for the averaged iterate therefore gives the same expected suboptimality after rounding.

From marginal estimation to optimization. Suppose each fresh marginal estimate has absolute bias at most β, second moment at most M, and expected comparison cost at most ν, uniformly over the marginal value, with $M, \nu > 0$. For an expected budget $u \geq 2n\nu$, take $K = \lfloor u/(n\nu) \rfloor$ iterations, starting at $w_1 = \mathbf{1}/2$ with step size $\eta = (2\sqrt{MK})^{-1}$. Each iteration estimates the n chain marginals and projects $w_t - \eta \widehat{g}_t$ onto $[0,1]^n$. The expected cost is at most u, and Lemma 8 gives

$$\mathbb{E}[\widehat{f}(\bar{w}) - f^\star] \leq n\beta + n\sqrt{\frac{n\nu M}{2u}}. \tag{5}$$

Unbiased estimates with bounded νM give the $T^{-1/2}$ rate after enforcing the comparison cap. For singular inverses, the estimators below trade bias against this product.

3 Optimal Rates for Regular Transfer Functions

Bounded sample size. Applying a nonlinear inverse to a sample mean generally introduces bias. The following standard polynomial argument for Bernoulli sampling explains the obstruction to unbiased estimation with bounded sample size; see also Nacu & Peres (2005, Section 2).

Proposition 1 (Bounded stopping times)**.** *If an integrable statistic Z uses iid $\mathrm{Ber}(p)$ observations and a stopping time $N \leq k$ almost surely, then $p \mapsto \mathbb{E}_p Z$ is a polynomial of degree at most k. It cannot equal a nonpolynomial inverse on an interval.*

The proof sums over the finitely many possible observation histories. Thus exact inverse estimation may require an unbounded stopping time, even when its expected cost is finite. Optimization can also use biased estimates, as in Section 4.

An explicit sigmoid estimator. For $\rho_b(\theta) = \tanh(b\theta/2)$, a response is positive with probability $p = \sigma(b\theta)$, where $\sigma(x) = (1+e^{-x})^{-1}$. Thus $\theta = [\log p - \log(1-p)]/b$. Using independent response sequences, count comparisons through the first positive response and through the first negative response. These counts satisfy $G_+ \sim \mathrm{Geom}(p)$ and $G_- \sim \mathrm{Geom}(1-p)$. With $H_m = \sum_{j=1}^{m} j^{-1}$ and $H_0 = 0$, the inverse-binomial identity $\mathbb{E}H_{G_+ - 1} = \sum_{j \geq 1} (1-p)^j/j = -\log p$ (Mendo, 2025) yields the estimator

$$Z_b = \frac{H_{G_- - 1} - H_{G_+ - 1}}{b}. \tag{6}$$

This estimator satisfies

$$\mathbb{E}Z_b = \theta, \quad \mathbb{E}Z_b^2 \leq 1 + \frac{\pi^2}{6b^2}, \quad \mathbb{E}(G_+ + G_-) \leq 4\cosh^2(b/2). \tag{7}$$

These bounds hold uniformly over $\theta \in [-1, 1]$. The waiting times also have uniformly exponential tails for fixed b. The harmonic identity uses the same log-odds series as the Russian-roulette construction of El Bakkali et al. (2026); Appendix A.4 gives the moment calculations.

A single run for sigmoid preferences. Submodularity and the range $f(S) \in [0,1]$ imply $\|g(w)\|_1 \leq 2$ for every greedy subgradient. Thus the cost of estimating an entire chain can be bounded more sharply than by summing n worst-case scalar costs. For $n \geq 2$, let C_t count the comparisons used to form the vector $\widehat{g}_t$ of independent harmonic estimates, and let $\mathcal{F}_{t-1}$ contain the history before iteration t. Then

$$\mathbb{E}[C_t \mid \mathcal{F}_{t-1}] \leq \nu_{n,b} := 4n + 8\sinh^2(b/2), \qquad \mathbb{E}[\|\widehat{g}_t\|_2^2 \mid \mathcal{F}_{t-1}] \leq V_{n,b} := 2 + \frac{n\pi^2}{6b^2}. \tag{8}$$

Take $K = \lfloor T/(2\nu_{n,b}) \rfloor$ iterations from $w_1 = \mathbf{1}/2$, with step size $\eta = \sqrt{n}/(2\sqrt{V_{n,b}K})$ when $K \geq 1$. At each iteration, sort the coordinates, estimate the chain marginals using Equation (6), and take a projected subgradient step. Cap the run at T comparisons. Return the pre-update average if it completes and the initialization otherwise, also using the initialization when $K = 0$. A conditional exponential-moment bound makes the interruption probability exponentially small in T for fixed b. Together with Lemma 8, this gives the optimal $O_b(\min\{1, n^{3/2}/\sqrt{T}\})$ rate. Appendix A.5 proves the bound and gives the larger iteration counts used in the experiments.

Regular inverses. The same upper bound holds whenever each marginal admits an unbiased estimator with uniformly bounded expected cost and second moment. The following regularity condition on the inverse ensures such an estimator through multilevel sampling. For a continuously differentiable function h, write $\omega_{h'}(t) = \sup_{|x-y|\le t} |h'(x) - h'(y)|$, with x, y in its domain. Its derivative is Dini-continuous if $\int_0^1 \omega_{h'}(t)\, dt/t < \infty$.

Theorem 2 (Regular transfer functions)**.** *Suppose ψ has an odd C^1 extension to $[-1, 1]$ with a Dini-continuous derivative, and $\psi'(0) > 0$. Then, for all $n, T \ge 1$,*

$$R^{(1)}_{n,T}(\rho) = \Theta_\rho\Big(\min\{1, n^{3/2}/\sqrt{T}\}\Big). \tag{9}$$

In particular this holds for every fixed sigmoid strength $b > 0$ and for the probit transfer function $\rho(\theta) = 2\Phi_{\mathrm{N}}(b\theta) - 1$, where Φ_{N} is the standard normal distribution function.

A derivative with any positive Hölder exponent satisfies the Dini condition. In this general case, the multilevel construction in Section 4 gives uniform finite bounds on expected cost and second moment. Applying Equation (5) and the budget conversion in Proposition 11 yields the upper bound. The modular lower-bound argument of Sho & Ito (2026) for linear and sigmoid preferences extends to these transfer functions (Appendix E.2). In particular, the sigmoid comparison complexity for expected suboptimality ϵ is $\Theta_b(n^3/\epsilon^2)$ in the nontrivial regime.

4 A Joint Minimax Rate

When the transfer function is flat at zero, small differences in function values produce even smaller differences in comparison probabilities. Its inverse can then have an unbounded slope. We impose regularity away from zero to control estimation at other gaps, and a shape condition near zero to bound the multilevel corrections.

Assumption 3 (Conditions on the inverse)**.** The inverse ψ satisfies either the hypothesis of Theorem 2, or both of the following: on some $(0, u_0]$, ψ' is nonnegative, nonincreasing, and convex; on every compact interval avoiding zero, ψ is C^1 with a Dini-continuous derivative. At the attainable endpoints the latter condition includes one-sided extendability.

These conditions permit $\psi'(0+) = \infty$. Where three derivatives exist, the inequalities near zero are $\psi' \ge 0$, $\psi'' \le 0$, and $\psi''' \ge 0$. Under these conditions, the modulus of continuity of ψ defined in Section 5 is comparable to the inverse at zero.

Theorem 4 (Joint minimax rate)**.** *Under Assumption 3, for every integer $n, T \ge 1$,*

$$R^{(1)}_{n,T}(\rho) = \Theta_\rho\Big(\min\{1, n\psi_{\mathrm{cl}}(\sqrt{n/T})\}\Big). \tag{10}$$

The upper bound has a fixed comparison budget and applies to every $f \in \mathcal{F}_n$. The lower bound applies to all adaptive randomized algorithms using 1-local queries and holds already on modular functions.

For $\rho(\theta) = \lambda \operatorname{sgn}(\theta)|\theta|^q$, Equation (10) gives the exponent $1/(2q)$ in Table 1. The same theorem retains logarithmic factors and covers infinite-order flatness.

Multilevel estimation. Sample means can lie outside the domain $[-a_\rho, a_\rho]$ of ψ, so we use an extension $\Psi : [-1, 1] \to \mathbb{R}$ from Lemma 13. Write $\overline{Y}_m = m^{-1}\sum_{i=1}^m Y_i$ for the mean of m iid symmetrized signs. At level $\ell \ge 1$, form two independent means A_ℓ, B_ℓ of $2^{\ell-1}$ signs and define

$$\Delta_0 = \Psi(Y_1), \qquad \Delta_\ell = \Psi((A_\ell + B_\ell)/2) - \tfrac{1}{2}[\Psi(A_\ell) + \Psi(B_\ell)]. \tag{11}$$

Fix a maximal level L. Draw $J \in \{0, \dots, L\}$ independently with probabilities q_ℓ, and return $Z_L = \Delta_J/q_J$. The expectations telescope, giving $\mathbb{E}Z_L = \mathbb{E}\Psi(\overline{Y}_{2^L})$. This approximates $\psi(\mathbb{E}Y_1)$ with a uniform bias bound that tends to zero as L increases.

The curvature estimate. The antithetic correction vanishes for a linear inverse. To retain this cancellation, we write $\Psi = \chi + \Phi$, where χ has a Dini-continuous derivative and Φ is the odd

extension of a concave function ϕ whose derivative is nonnegative, nonincreasing, and convex. The quantity $D(t) = 2\phi(t) - \phi(2t)$ measures its departure from linearity. Writing $h(t) = \phi(t)/t$, we have

$$\frac{D(t)}{t} = 2[h(t) - h(2t)].$$

At the scales $s_\ell = 2^{-\ell/2}$, $2s_\ell = s_{\ell-2}$. Summing separately over even and odd levels therefore telescopes. Together with the pointwise bound in Lemma 14, this gives

$$\left|\Phi((x+y)/2) - \tfrac{1}{2}[\Phi(x) + \Phi(y)]\right| \le 2D(|x-y|/2), \qquad \sum_{\ell=1}^{L} \frac{D(s_\ell)}{s_\ell} \le 4\frac{\phi(s_L)}{s_L}.$$

The first inequality controls each correction. The second lets us sum their contributions without an extra factor depending on the number of levels.

Bias, cost, and second moment. Let v_ℓ bound the second moment of Δ_ℓ uniformly over the true marginal. For these bounds, the choice $q_\ell \propto \sqrt{v_\ell/2^\ell}$ minimizes the product of expected cost $\nu_L = \sum_\ell q_\ell 2^\ell$ and second-moment bound $M_L = \sum_\ell v_\ell/q_\ell$. By Lemma 15, we can choose

$$\sqrt{v_\ell} = 4D(s_\ell) + 4s_\ell\omega_{\chi'}(s_\ell) \qquad (\ell \ge 1).$$

Multiplying by the square root of the level cost, $2^{\ell/2} = 1/s_\ell$, gives $4D(s_\ell)/s_\ell + 4\omega_{\chi'}(s_\ell)$. The curvature terms telescope, and the Dini condition makes the sum of the regular terms finite. Using $s + \phi(s) \asymp_\rho \psi(s)$ from Lemma 13, we obtain

$$\nu_L M_L = \left(\sum_{\ell=0}^{L} \sqrt{2^\ell v_\ell}\right)^2 \lesssim_\rho \left(\frac{\psi(s_L)}{s_L}\right)^2 \qquad (s_L \le a_\rho), \tag{12}$$

and the bias is $O_\rho(\psi(s_L))$.

From the estimator to the minimax rate. Apply Equation (5) with bias $O_\rho(\psi(s_L))$ and the product $\nu_L M_L$ from Equation (12). Choosing $2^L \asymp u/n$ makes both terms of order $n\psi(\sqrt{n/u})$ in the nontrivial budget regime. To enforce the fixed budget, reserve blocks of sizes $1, 2, 4, \ldots$ with total at most T. In each block of size B, start an independent run with expected cost at most $B/4$ and interrupt it at B comparisons. Return the output of the last completed run, or the initialization if none completes. Markov's inequality bounds each interruption probability by $1/4$. Independence from previous outputs yields an error recursion that preserves the rate (Proposition 11). Appendices C to E give the estimator analysis and matching adaptive lower bound.

Theorem 5 (Unbiased estimation and the $T^{-1/2}$ rate)**.** *Within Assumption 3, the following are equivalent:*

- *(i) A scalar marginal admits an unbiased estimator using comparisons of that fixed pair, with uniformly bounded expected cost and second moment, even allowing adaptive query orientation.*
- *(ii)* $\psi(u) = O(u)$ *as* $u \downarrow 0$.
- *(iii)* $\Gamma(t) \asymp t$ *as* $t \downarrow 0$.
- *(iv) The minimax risk* $R^{(1)}_{n,T}(\rho)$ *satisfies Equation (9).*

When the inverse has finite slope at zero, the multilevel estimator can use infinitely many possible levels while retaining finite expected cost and second moment. Conversely, an information bound for sequential estimation shows that any unbiased estimator with uniform bounds ν, M must satisfy $\nu M \ge c_\rho[\psi(s)/s]^2$. An unbounded inverse slope therefore rules out these simultaneous uniform bounds. The estimator with finitely many levels still attains the slower optimal rate in Theorem 4. Appendix F proves the equivalence.

5 Flatness Away from Ties

The assumptions away from zero matter because two sets that differ by one element need not have close values. Consider the odd analytic transfer function

$$\rho_*(t) = \frac{t^5}{5} - \frac{t^3}{6} + \frac{t}{16}, \qquad \rho_*'(t) = (t^2 - 1/4)^2. \tag{13}$$

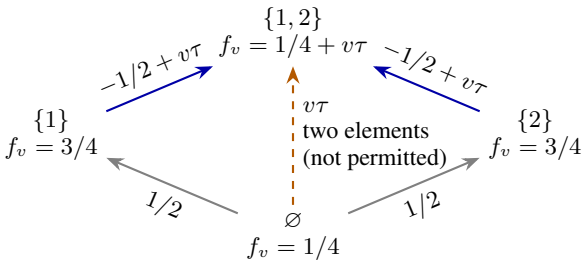

Figure 1: 1-local queries in Equation (15). Arrows query the destination against the source and are labeled by the value gap. The gray edges reveal no information about v. The blue edges have gaps near the flat point $-1/2$. Comparing the two candidate minimizers directly would give gap $v\tau$, but changes two elements and is not permitted.

It is strictly increasing on $[-1, 1]$, takes values in $(-1, 1)$, and satisfies $\rho'_*(0) = 1/16$. Nevertheless,

$$R^{(1)}_{2,T}(\rho_*) = \Theta(T^{-1/6}). \tag{14}$$

To see the obstruction, take $v \in \{-1, 1\}$ and $0 < \tau \le 1/8$ and set

$$f_v(\varnothing) = \tfrac{1}{4}, \quad f_v(\{1\}) = f_v(\{2\}) = \tfrac{3}{4}, \quad f_v(\{1, 2\}) = \tfrac{1}{4} + v\tau. \tag{15}$$

Both functions are submodular. Their minimizers are opposite endpoints, but every informative 1-local query has a gap near $1/2$ or $-1/2$. The change in the comparison signal is cubic:

$$\rho_*(-1/2 + \tau) - \rho_*(-1/2 - \tau) = \tfrac{2}{3}\tau^3 + \tfrac{2}{5}\tau^5. \tag{16}$$

Thus each observation has Kullback–Leibler (KL) divergence $O(\tau^6)$ between the two instances. Taking $\tau \asymp T^{-1/6}$ leaves a constant probability of choosing the wrong endpoint and yields the lower bound in Equation (14). A direct comparison between the endpoints would observe their small gap near zero, but changes two elements and is not permitted.

The modulus of continuity of ψ on $[0, a_\rho]$ is

$$\Omega_\rho(s) = \max\{b - a : 0 \le a \le b \le 1,\ \Gamma(b) - \Gamma(a) \le s\}. \tag{17}$$

This is the largest difference in positive value gaps compatible with a signal difference of at most s.

Theorem 6 (Arbitrary transfer functions at fixed dimension)**.** *Under the model of Section 2, for every fixed $n \ge 2$,*

$$R^{(1)}_{n,T}(\rho) = \Theta_{n,\rho}(\Omega_\rho(T^{-1/2})). \tag{18}$$

Consequently a uniform $O_{n,\rho}(T^{-1/2})$ risk is possible if and only if $\Gamma(b) - \Gamma(a) \ge c(b - a)$ for all $0 \le a \le b \le 1$ and some $c > 0$.

For the lower bound, fix $0 \le a < b \le 1$, put $d = b - a$, and define

	$\varnothing$	$\{1\}, \{2\}$	$\{1, 2\}$
f_0	0	b	d
f_1	d	b	0

Both functions are submodular, with opposite endpoint minimizers. Each hypercube edge exchanges gaps a and b, up to sign. If $\Gamma(b) - \Gamma(a) \le s$, each comparison has KL divergence $O_\rho(s^2)$, while a wrong endpoint decision costs at least d. Taking $s \asymp_\rho T^{-1/2}$ and maximizing over intervals gives the modulus in Equation (17).

For the upper bound, estimate the marginal on each edge of a spanning tree of the subset hypercube, reconstruct relative function values, and minimize them. This uses $2^n - 1$ edges, so the constants can grow exponentially with dimension. For Equation (13), $\Omega_{\rho_*}(s) \asymp s^{1/3}$ near zero, which completes Equation (14). Under Assumption 3, this modulus is comparable to ψ_{cl}. Appendix G proves these claims.

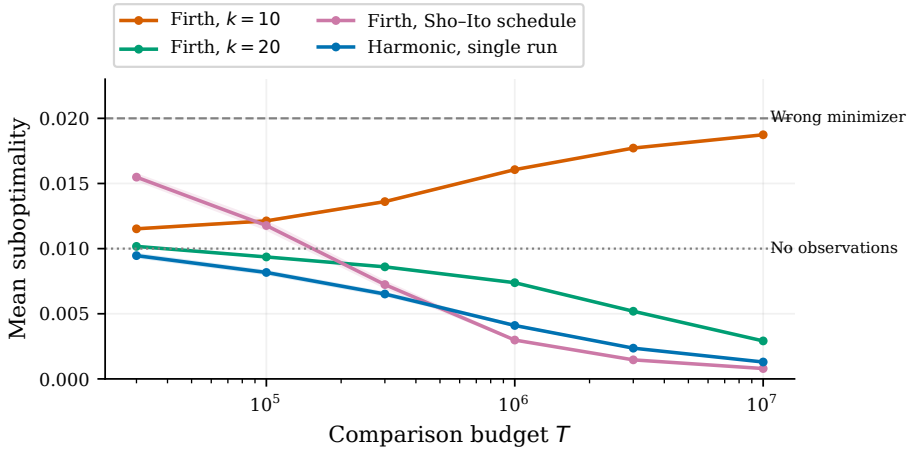

Figure 2: Marginal estimation on a fixed objective f and $f^{\mathrm{rev}}(S) = f([3] \setminus S)$. Curves average 300 runs per budget, balanced between the two functions; bands are 95% bootstrap intervals. The best rule without observations has suboptimality 0.01. Fixed-k Firth with $k = 10$ approaches the wrong minimizer; the larger batch, Sho–Ito schedule, and harmonic estimator recover the correct minimizer.

6 Synthetic Experiments

We use two synthetic experiments to illustrate how estimation bias can change the minimizer and how 1-local queries can limit information. All optimization runs obey a fixed comparison budget. We report $\widehat{f}(\bar{w}) - f^\star$, the suboptimality of the Lovász extension, which equals the expected suboptimality after threshold rounding conditional on $\bar{w}$. Appendix J gives the protocols, rounded-set results, and additional controls.

Estimation bias can change the minimizer. Consider the three-element cardinality function f with values $0, 1, 1-c, 1-2c$, where $c = 0.49$. Its unique minimizer is empty, and its greedy marginals sum to 0.02. Under sigmoid preferences with $b = 2$, we apply Firth's bias-reduced log-odds estimator (Firth, 1993; Sho & Ito, 2026) with fixed batch size $k = 10$. Its expected marginal estimates sum to -0.0254. They are subgradients of a different submodular function whose unique minimizer is the full set. Proposition 20 shows that the threshold-rounded averaged iterate converges in probability to this wrong minimizer.

We balance f and $f^{\mathrm{rev}}(S) = f([3] \setminus S)$, exchanging the two endpoints. In Figure 2, fixed-k Firth approaches suboptimality 0.02 at the wrong minimizer as the budget increases. At $T = 10^7$, its mean suboptimality is 0.0187, and threshold rounding selects the wrong endpoint with probability 0.937. The harmonic estimator improves with budget. Increasing the fixed batch to 20 removes the reversal, and the growing-batch schedule of Sho & Ito (2026) also recovers. Their $T^{-2/5}$ guarantee balances estimation bias against the number of optimization steps available under a fixed comparison budget.

Changing only nonoptimal values. For $v \in \{-1, 1\}$, set $f_v(\varnothing) = 1/4$, $f_v(\{1\}) = f_v(\{2\}) = 1/4 + a$, and $f_v(\{1,2\}) = 1/4 + v\tau$. Fix $\tau = 0.05$ and vary a. The optimal endpoint and the loss from choosing the other endpoint stay fixed. We use $8\rho_*$ from Equation (13), whose derivative at zero matches the linear control $\rho(t) = t/2$.

Every informative 1-local query observes the same pair of Bernoulli distributions, up to orientation. We can therefore compute the optimal equal-prior testing error for this family by a binomial likelihood-ratio test. In Figure 3, raising a from 0.10 to 0.50 increases the comparisons needed for 90% accuracy under equal priors from 3,126 to approximately 14.5 million.

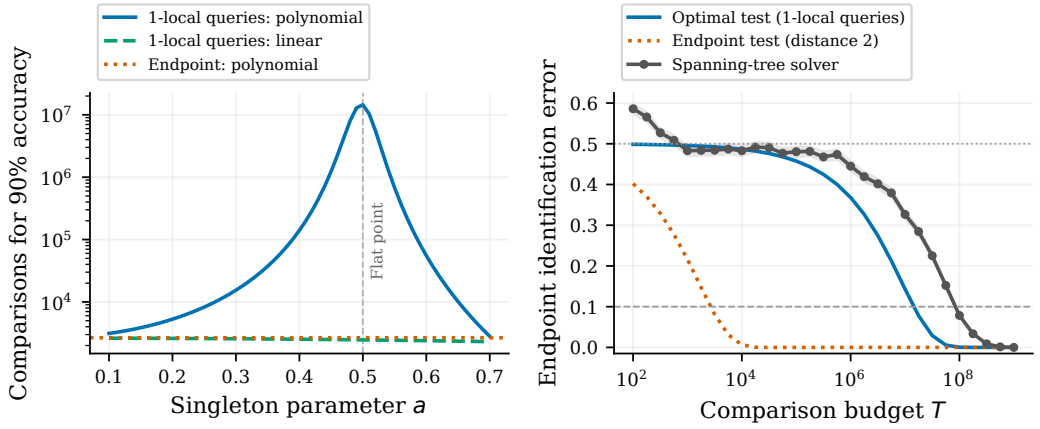

Figure 3: Information in the permitted comparisons. Left: binomial testing budgets as only the singleton values change. The polynomial and linear transfer functions have the same slope at zero. Right: testing error at $a = 0.50$, compared with 4,000 runs per budget of the spanning-tree solver, which is given only the transfer function and dimension. Bands are 95% bootstrap intervals. The dashed horizontal line marks 10% error (90% accuracy). The endpoint query has $|A \triangle B| = 2$ and is outside the allowed class; the two testing curves use knowledge of the family.

Moving past the flat point reduces the cost again. A direct endpoint comparison needs 2,663 comparisons throughout; its query has $|A \triangle B| = 2$ and is outside the allowed class. The generic spanning-tree solver also learns the endpoint, using more observations than the test with known family parameters. Appendix J.3 shows how the effect persists under perturbations with strictly positive derivative everywhere.

7 Discussion

Estimation bias and flatness of the transfer function affect optimization in different ways. Fixed-batch inversion can preserve each marginal's sign while changing their relative magnitudes enough to change the minimizer. For sigmoid preferences, sequential unbiased estimation removes this distortion and attains the optimal rate under 1-local queries. Our lower bounds use instances whose permitted comparisons carry little information. A one-element change can have a large value gap, so this obstruction can arise away from ties. The inverse near zero determines the joint rate under Assumption 3; its modulus of continuity determines the rate at fixed dimension for arbitrary transfer functions.

A sharp joint dependence on dimension and budget for arbitrary transfer functions remains open. Other directions include unrestricted queries, unknown transfer functions, and robustness to misspecification.

Reproducibility Statement

The appendices give the proofs, estimator constructions, and budget arguments. Appendix J specifies the synthetic instances, comparison models, algorithm schedules, random seeds, and statistical summaries. The accompanying package contains executable code, per-run numerical arrays, figure scripts, and finite consistency checks.

AI Use Statement

The authors formulated the research problem, directed the investigation, and selected among candidate approaches. OpenAI ChatGPT (GPT-5.6 Sol) was used under the authors' guidance to assist in developing the estimators, formulating mathematical claims, developing and drafting proofs, and

deriving mathematical consequences. OpenAI ChatGPT and Anthropic Claude Fable 5.1 were also used in later stages for literature discovery and review, proof auditing, manuscript drafting and editing, figure design and editing, experimental design and analysis, LaTeX and code assistance, and terminology and consistency checks. References and claims identified through AI-assisted literature search were checked by the authors against the original sources, and issues identified during proof audits were reviewed and corrected. The authors reviewed all AI-assisted content, including the experimental code and results, and take full responsibility for the final manuscript and its correctness.

REFERENCES

Francis Bach. Learning with submodular functions: A convex optimization perspective. *Foundations and Trends in Machine Learning*, 6(2–3):145–373, 2013. https://arxiv.org/abs/1111.6453.

Jose H. Blanchet and Peter W. Glynn. Unbiased Monte Carlo for optimization and functions of expectations via multilevel randomization. In *Winter Simulation Conference*, pp. 3656–3667, 2015. https://web.stanford.edu/~glynn/papers/2015/BlanchetG15.html.

Jose H. Blanchet, Peter W. Glynn, and Yanan Pei. Unbiased multilevel Monte Carlo: Stochastic optimization, steady-state simulation, quantiles, and other applications. *arXiv preprint arXiv:1904.09929*, 2019. https://arxiv.org/abs/1904.09929.

Morris H. DeGroot. Unbiased sequential estimation for binomial populations. *Annals of Mathematical Statistics*, 30(1):80–101, 1959. doi:10.1214/aoms/1177706361.

Taha El Bakkali, El Mahdi Chayti, and Omar Saadi. Nonsmooth optimization with zeroth order comparison feedback. *arXiv preprint arXiv:2602.05622*, 2026. https://arxiv.org/abs/2602.05622v2.

David Firth. Bias reduction of maximum likelihood estimates. *Biometrika*, 80(1):27–38, 1993. doi:10.1093/biomet/80.1.27.

M. A. Girshick, Frederick Mosteller, and L. J. Savage. Unbiased estimates for certain binomial sampling problems with applications. *Annals of Mathematical Statistics*, 17(1):13–23, 1946. doi:10.1214/aoms/1177731018.

Shinji Ito. Submodular function minimization with noisy evaluation oracle. In *Advances in Neural Information Processing Systems*, volume 32, 2019. https://proceedings.neurips.cc/paper/2019/hash/58d2d622ed4026cae2e56dffc5818a11-Abstract.html.

Kevin G. Jamieson, Robert D. Nowak, and Benjamin Recht. Query complexity of derivative-free optimization. In *Advances in Neural Information Processing Systems*, volume 25, 2012. https://papers.nips.cc/paper_files/paper/2012/hash/e6d8545daa42d5ced125a4bf747b3688-Abstract.html.

Luis Mendo. Estimating odds and log odds with guaranteed accuracy. *Statistical Papers*, 66, article 26, 2025. doi:10.1007/s00362-024-01639-w. https://arxiv.org/abs/2404.17705.

Serban Nacu and Yuval Peres. Fast simulation of new coins from old. *Annals of Applied Probability*, 15(1A):93–115, 2005. doi:10.1214/105051604000000549.

Chang-Han Rhee and Peter W. Glynn. Unbiased estimation with square root convergence for SDE models. *Operations Research*, 63(5):1026–1043, 2015. doi:10.1287/opre.2015.1404.

Aadirupa Saha, Tomer Koren, and Yishay Mansour. Dueling convex optimization. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9245–9254, 2021. https://proceedings.mlr.press/v139/saha21b.html.

Aadirupa Saha, Tomer Koren, and Yishay Mansour. Dueling convex optimization with general preferences. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 52552–52564, 2025. https://proceedings.mlr.press/v267/saha25a.html.

Kaien Sho and Shinji Ito. Submodular function minimization with dueling oracle. In *International Conference on Learning Representations*, 2026. https://proceedings.iclr.cc/paper_files/paper/2026/file/45d0b8e7ec7b52d960a1c7b43fdccac9-Paper-Conference.pdf.

Tianze Wang and Guanyang Wang. Unbiased multilevel Monte Carlo methods for intractable distributions: MLMC meets MCMC. *Journal of Machine Learning Research*, 24(249):1–40, 2023. https://www.jmlr.org/papers/v24/22-1468.html.

Haixiang Zhang, Zeyu Zheng, and Javad Lavaei. Stochastic $L^\natural$-convex function minimization. In *Advances in Neural Information Processing Systems*, volume 34, 2021. https://proceedings.neurips.cc/paper/2021/hash/6c81c83c4bd0b58850495f603ab45a93-Abstract.html.

A Auxiliary Facts and Unbiased Estimation

A.1 Symmetrization and comparison of information

Lemma 7. *One response to a 1-local query simulates a sign of mean $\Gamma(\theta)$. Let P_x be the original response law for a value gap x. There is $K_\rho < \infty$ such that*

$$\mathrm{KL}(P_x \| P_y) \le K_\rho |\Gamma(x) - \Gamma(y)|^2, \qquad x, y \in [-1, 1]. \tag{19}$$

The same bound holds after reversing the query orientation.

Proof. Let R be an independent uniform sign. Query (A, B) for $R = 1$ and (B, A) for $R = -1$, and return $Y = RO$. If $\theta = f(A) - f(B)$, then

$$\mathbb{E}Y = \tfrac{1}{2}\rho(\theta) - \tfrac{1}{2}\rho(-\theta) = \Gamma(\theta).$$

The compared sets and the raw comparison count are unchanged. Write

$$\eta_\rho = \tfrac{1}{2}\min\{1 + \rho(-1), 1 - \rho(1)\} > 0.$$

Every original response probability belongs to $[\eta_\rho, 1 - \eta_\rho]$. If $x \ge y$, monotonicity gives

$$0 \le \rho(x) - \rho(y) \le \rho(x) - \rho(y) + \rho(-y) - \rho(-x) = 2[\Gamma(x) - \Gamma(y)].$$

Thus the corresponding success probabilities differ by at most $|\Gamma(x) - \Gamma(y)|$. For Bernoulli laws, $\mathrm{kl}(p, q) \le (p - q)^2/[q(1 - q)]$, which proves the bound with $K_\rho = [\eta_\rho(1 - \eta_\rho)]^{-1}$. Oddness of Γ gives the same result for $-x, -y$. ■

A.2 Lovász reduction and biased projected descent

We recall the reduction to fix the rounding and normalization conventions. Put $f_0(S) = f(S) - f(\varnothing)$ and, for a vector g, write $g(S) = \sum_{i \in S} g_i$. The base polytope of f_0 consists of vectors satisfying $g([n]) = f_0([n])$ and $g(S) \le f_0(S)$ for every S. The greedy vector in Equation (4) satisfies these conditions. Indeed, when the elements of S appear along the greedy chain, diminishing returns bounds each chain marginal by the corresponding marginal along the chain containing only elements of S; summing proves the inequality.

For decreasingly sorted w, summation by parts gives

$$\langle g, w \rangle = \sum_{j=1}^{n-1} (w_{\pi(j)} - w_{\pi(j+1)}) g(B_j) + w_{\pi(n)} g([n]).$$

For any base-polytope vector, each term is bounded by replacing $g(B_j)$ with $f_0(B_j)$, and the greedy vector attains equality. Thus $\widehat{f}(w)$ is a maximum of linear functions plus $f(\varnothing)$ and is convex; its greedy maximizer is a subgradient. The threshold set $\{i : w_i \ge U\}$ equals a chain set almost surely, with probabilities $1 - w_{\pi(1)}$, $w_{\pi(j)} - w_{\pi(j+1)}$, and $w_{\pi(n)}$. Therefore its expected value equals $\widehat{f}(w)$. Every such convex combination is at least $f^\star$, and the indicator vector of a minimizing set attains $f^\star$. These arguments also cover coordinate ties, using any fixed ordering of tied elements.

Lemma 8 (Biased stochastic subgradient bound)**.** *At iteration t, suppose a vector $\widehat{g}_t$ satisfies, conditionally on the past,*

$$\mathbb{E}\|\widehat{g}_t\|_2^2 \le nM, \qquad \|\mathbb{E}\widehat{g}_t - g(w_t)\|_\infty \le \beta.$$

Initialize $w_1 = \frac{1}{2}\mathbf{1}$, use $w_{t+1} = \mathrm{Proj}_{[0,1]^n}(w_t - \eta\widehat{g}_t)$, and average $\bar{w} = K^{-1}\sum_{t=1}^K w_t$. With $\eta = (2\sqrt{MK})^{-1}$,

$$\mathbb{E}[\widehat{f}(\bar{w}) - f^\star] \le \frac{n\sqrt{M}}{2\sqrt{K}} + n\beta. \tag{20}$$

Proof. Fix a minimizing indicator vector $w^\star$. Nonexpansiveness of projection yields

$$\|w_{t+1} - w^\star\|_2^2 \le \|w_t - w^\star\|_2^2 - 2\eta\langle \widehat{g}_t, w_t - w^\star \rangle + \eta^2 \|\widehat{g}_t\|_2^2.$$

Conditional expectation and convexity give

$$\mathbb{E}\langle \widehat{g}_t, w_t - w^\star \rangle \ge \mathbb{E}[\widehat{f}(w_t) - f^\star] - n\beta,$$

since $\|w_t - w^\star\|_1 \le n$. Sum over t, use $\|w_1 - w^\star\|_2^2 = n/4$, and apply convexity at the average. This gives $n/(8\eta K) + \eta nM/2 + n\beta$, whose first two terms sum to the displayed bound at the stated step size. A fresh uniform threshold then preserves the expected suboptimality. ■

A.3 Bounded stopping times

Proof of Proposition 1. Pad the observations to length k, without changing the output, and include an independent internal random seed. For every binary sequence $x \in \{0,1\}^k$, average the output over the seed conditional on that sequence and call the result z_x. Integrability at any one interior parameter $p_0 \in (0,1)$ implies $|z_x| < \infty$ for every x, since every sequence has positive probability at p_0. Hence

$$\mathbb{E}_p Z = \sum_{x \in \{0,1\}^k} z_x p^{\sum_i x_i} (1-p)^{k - \sum_i x_i},$$

a polynomial of degree at most k. An equivalent argument sums over the finite stopped histories. The seed and stopping rule do not depend on the unknown p. ■

A.4 Harmonic identities and unbiased sigmoid estimation

Write $\mathrm{Li}_2(z) = \sum_{j \geq 1} z^j / j^2$ for the dilogarithm.

Lemma 9. *For $G \sim \mathrm{Geom}(q)$ on $\{1, 2, \ldots\}$,*

$$\mathbb{E} H_{G-1} = -\log q, \qquad \mathrm{Var}(H_{G-1}) = \mathrm{Li}_2(1-q).$$

Proof. Write $x = 1 - q$. Since $H_{G-1} = \sum_{j \geq 1} \mathbf{1}\{G > j\}/j$, Tonelli's theorem gives $\mathbb{E} H_{G-1} = \sum_{j \geq 1} x^j / j = -\log(1-x)$. To calculate the second moment, note that

$$H_m^2 - H_{m-1}^2 = \frac{2H_{m-1}}{m} + \frac{1}{m^2}.$$

The generating function $\sum_{r \geq 0} H_r x^r = -\log(1-x)/(1-x)$ follows by interchanging two nonnegative sums. Integrating it termwise gives

$$\sum_{m \geq 1} \frac{H_{m-1}}{m} x^m = \frac{1}{2} \log^2(1-x).$$

Therefore $(1-x) \sum_{m \geq 0} H_m^2 x^m = \log^2(1-x) + \mathrm{Li}_2(x)$, and multiplication by the geometric mass $q = 1 - x$ proves the claim after subtracting the squared mean. ■

Apply Lemma 9 to the two independent streams in Equation (6). Their mean difference is $\log p - \log(1-p) = b\theta$. Their variances add:

$$\mathrm{Var}(Z_b) = \frac{\mathrm{Li}_2(p) + \mathrm{Li}_2(1-p)}{b^2}.$$

The identity

$$\mathrm{Li}_2(p) + \mathrm{Li}_2(1-p) = \frac{\pi^2}{6} - \log p \log(1-p) \leq \frac{\pi^2}{6}$$

follows by differentiating both sides and evaluating the limit at $p \downarrow 0$. Adding the squared mean $\theta^2 \leq 1$ gives the second-moment bound in Equation (7). The mean cost is $p^{-1} + (1-p)^{-1} = 4\cosh^2(b\theta/2)$. For $|\theta| \leq 1$, both p and $1-p$ are at least $\sigma(-b) > 0$, so a union bound on the two geometric tails gives a uniform exponential tail. These calculations establish all scalar claims in Section 3.

A.5 Sigmoid optimization with one capped run

We first bound cost and second moment over the entire greedy chain. For a greedy subgradient g, let $P = \{i : g_i > 0\}$. The base-polytope inequalities in Appendix A.2 give

$$g(P) \leq f(P) - f(\varnothing) \leq 1 - f(\varnothing), \qquad g([n]) = f([n]) - f(\varnothing).$$

Consequently,

$$\|g\|_1 = 2g(P) - g([n]) \leq 2 - f(\varnothing) - f([n]) \leq 2, \qquad \|g\|_2^2 \leq \|g\|_1 \leq 2.$$

Here $\|g\|_2^2 \le \|g\|_1$ uses $|g_i| \le 1$. Convexity and the value at zero imply $\sinh^2(bt/2) \le t\sinh^2(b/2)$ for $0 \le t \le 1$. Summing the scalar cost identity along the chain gives

$$\sum_{i=1}^{n} 4\cosh^2(bg_i/2) = 4n + 4\sum_{i=1}^{n}\sinh^2(bg_i/2) \le 4n + 8\sinh^2(b/2) = \nu_{n,b}.$$

Likewise, summing the scalar second moments gives $\mathbb{E}[\|\widehat{g}_t\|_2^2 \mid \mathcal{F}_{t-1}] \le \|g(w_t)\|_2^2 + n\pi^2/(6b^2) \le V_{n,b}$. This proves Equation (8).

The cost also admits an exponential-moment bound that holds conditionally at every iteration. For $0 < z < \log(1+e^{-b})$ and $|t| \le 1$, define

$$\ell_z(t) = -\log\!\left(1 - 4(e^z-1)e^{-2z}\cosh^2(bt/2)\right), \qquad \mathcal{L}_{n,b}(z) = (n-2)\ell_z(0) + 2\ell_z(1). \tag{21}$$

Proposition 10. *Let $n \ge 2$, $b > 0$, and $K \ge 1$ be a fixed integer. Starting at $w_1 = \mathbf{1}/2$, run projected subgradient descent with independent harmonic estimates and step size $\eta = \sqrt{n}/(2\sqrt{V_{n,b}K})$. Cap the run at T raw comparisons. If all K iterations complete, return their pre-update average; otherwise return w_1. For the returned vector $\widehat{w}$ and every $0 < z < \log(1+e^{-b})$,*

$$\mathbb{E}[\widehat{f}(\widehat{w}) - f^\star] \le \frac{\sqrt{nV_{n,b}}}{2\sqrt{K}} + e^{-zT + K\mathcal{L}_{n,b}(z)}. \tag{22}$$

Threshold rounding gives the same expected suboptimality bound for the returned set.

Proof. For a fixed marginal t, put $p = \sigma(bt)$. Independence of the two geometric waiting times gives

$$\mathbb{E}e^{z(G_+ + G_-)} = \frac{p(1-p)e^{2z}}{(1-(1-p)e^z)(1-pe^z)} = e^{\ell_z(t)}.$$

The stated range of z lies in the domain of finiteness uniformly over $|t| \le 1$. As a function of $t \in [0,1]$, $\ell_z(t)$ is increasing and convex: $\cosh^2(bt/2)$ is convex, and its composition with $x \mapsto -\log(1-4(e^z-1)e^{-2z}x)$ preserves convexity and monotonicity on this domain. The chord between zero and one therefore gives

$$\ell_z(g_i) = \ell_z(|g_i|) \le \ell_z(0) + |g_i|[\ell_z(1) - \ell_z(0)].$$

The coordinate estimates use independent observations conditional on the current chain. Hence $\|g\|_1 \le 2$ implies

$$\log\mathbb{E}[e^{zC_t} \mid \mathcal{F}_{t-1}] = \sum_{i=1}^{n}\ell_z(g_i(w_t)) \le \mathcal{L}_{n,b}(z).$$

Consider the ideal run of K complete iterations, and let $S_K = \sum_{t=1}^{K} C_t$. Iterated conditioning and Markov's inequality yield $\mathbb{P}(S_K > T) \le e^{-zT + K\mathcal{L}_{n,b}(z)}$. Applying Lemma 8 to the ideal run with $M = V_{n,b}/n$ and $\beta = 0$ bounds its expected suboptimality by the first term in Equation (22). Couple the capped and ideal runs using the same observations. They agree on $\{S_K \le T\}$; on interruption the returned vector has suboptimality at most one. Thus the suboptimality of the capped run is at most that of the ideal run plus $\mathbf{1}\{S_K > T\}$. Taking expectations proves the claim. ■

Choosing the iteration count. For a chosen interruption probability $\delta \in (0,1)$, any z in the above interval gives the valid count

$$K = \left\lfloor \frac{zT - \log(1/\delta)}{\mathcal{L}_{n,b}(z)} \right\rfloor$$

when this is at least one. The bound in Equation (22) is then $\sqrt{nV_{n,b}}/(2\sqrt{K}) + \delta$. The choice of z can be optimized before observing any responses. Taking $\delta = 1/(4\sqrt{T})$ makes the interruption term decay at the $T^{-1/2}$ rate. For a simpler rate guarantee, take $K = \lfloor T/(2\nu_{n,b})\rfloor$. Since $\ell_z(t)/z \to 4\cosh^2(bt/2)$ as $z \downarrow 0$, there is $z_b > 0$, depending only on b, for which $\mathcal{L}_{n,b}(z_b) \le (3/2)z_b\nu_{n,b}$ for all $n \ge 2$. The interruption probability is then at most $e^{-z_bT/4}$. Because $\nu_{n,b} = O_b(n)$ and $V_{n,b} = O_b(n)$, this yields $O_b(\min\{1, n^{3/2}/\sqrt{T}\})$, using the initialization when $K = 0$. For $n = 1$, replace $\mathcal{L}_{n,b}(z)$ by $\ell_z(1)$; the same proof applies.

B From Expected Cost to a Fixed Budget

Proposition 11 (From expected cost to a fixed budget)**.** *Suppose algorithms have, uniformly over instances, expected cost at most u and expected loss at most $A\min\{1, r(u)\}$, where losses lie in $[0,1]$. If r is nonincreasing, $r(u/2) \le c_0 r(u)$ with $c_0 < 4$, and $r(u) \ge c\min\{1, u^{-s}\}$ for some $s < 2$, then an algorithm using at most T comparisons attains expected loss $O(\min\{1, r(T)\})$.*

Proof. Replace r by $\min\{1, r\}$; the scaling condition persists, increasing c_0 to at least one if necessary. Let $k = \lfloor \log_2(T+1) \rfloor$ and $B_j = 2^j$ for $0 \le j < k$. Their sum is $2^k - 1 \le T$. Initialize an arbitrary incumbent. In block j, start an independent ideal algorithm of expected cost at most $B_j/4$, with its fixed initialization and fresh oracle randomness. Stop before any call that would exceed B_j calls in that block. Replace the incumbent only on completion.

Fix an instance. Let F_j denote interruption in block j, ℓ_j the ideal run's loss, and ℓ_{old} the previous incumbent's loss. Markov's inequality gives $\mathbb{P}(F_j) \le 1/4$, and fresh randomness makes F_j independent of ℓ_{old}. If r_j is expected incumbent loss after the block, then

$$r_j = \mathbb{E}[\ell_j \mathbf{1}\{F_j^c\}] + \mathbb{E}[\ell_{\mathrm{old}} \mathbf{1}\{F_j\}] \le A r(B_j/4) + \tfrac{1}{4} r_{j-1}.$$

Nonnegativity bounds the first term by the ideal run's expected loss, and independence bounds the second by $\mathbb{P}(F_j) r_{j-1}$. Iteration gives

$$r_{k-1} \le A \sum_{i=0}^{k-1} 4^{-i} r(B_{k-1}/(4 \cdot 2^i)) + 4^{-k}.$$

For $T \ge 1$, $B_{k-1} > T/4$, so monotonicity and four applications of the dyadic scaling bound yield

$$r(B_{k-1}/(4 \cdot 2^i)) \le c_0^{i+4} r(T).$$

The geometric series converges since $c_0 < 4$. The remainder is $O(T^{-2})$, absorbed by the assumed lower bound with $s < 2$. Finally, loss is at most one, and the construction preserves every pathwise query restriction. ∎

Rates covered by the proposition. If F is nonnegative, increasing, concave, $F(0) = 0$, and $F(t) \ge t$, then

$$r(u) = \min\{1, nF(\sqrt{n/u})\}$$

satisfies $r(u/2) \le \sqrt{2} r(u)$ and $r(u) \ge \min\{1, u^{-1/2}\}$ for $n \ge 1$. Thus the proposition applies with absolute scaling constants.

Reduction using unbiased estimates. Suppose each scalar marginal θ has an estimator satisfying $\mathbb{E}Z = \theta$, $\mathbb{E}Z^2 \le M$, and expected comparison cost at most ν, uniformly over θ, with $M, \nu < \infty$. Estimate all n chain marginals independently at each of K iterations. Conditional expectation bounds the ideal total cost by $nK\nu$ even when individual costs depend on current marginals. Lemma 8 gives expected suboptimality at most $n\sqrt{M}/(2\sqrt{K})$. Choosing $K = \lfloor u/(n\nu) \rfloor$ whenever $u \ge 2n\nu$, and a fixed output otherwise, gives expected suboptimality at most $C\sqrt{M\nu}\, n^{3/2}/\sqrt{u}$. The restart procedure yields

$$R^{(1)}_{n,T}(\rho) \le C \min\{1, \sqrt{M\nu}\, n^{3/2}/\sqrt{T}\}. \tag{23}$$

Here $M\nu \ge c_\rho > 0$ can always be ensured by increasing the bounds. For sigmoid, use $M = 1 + \pi^2/(6b^2)$ and $\nu = 4\cosh^2(b/2)$.

C A Decomposition of the Inverse Transfer Function

We prove the estimator bounds for a class of transfer functions that contains those satisfying Assumption 3.

Definition 12. A transfer function belongs to $\mathcal{C}$ if ψ has an extension $\Psi : [-1,1] \to \mathbb{R}$ of the form

$$\Psi(u) = \chi(u) + \Phi(u), \qquad \Phi(u) = \mathrm{sgn}(u)\phi(|u|), \tag{24}$$

where χ is an odd C^1 function with a Dini-continuous derivative, $\phi : [0,\infty) \to [0,\infty)$ is continuous with $\phi(0) = 0$, and ϕ' exists on $(0,\infty)$ and is nonnegative, nonincreasing, and convex. The choice $\phi = 0$ is permitted. We also require $\Gamma(t) \le L_\Gamma t$ for sufficiently small $t \ge 0$ and some finite L_Γ.

Lemma 13. *Every transfer function in Assumption 3 belongs to $\mathcal{C}$. For a transfer function in $\mathcal{C}$, put $F(s) = s + \phi(s)$. Then*

$$F(s) \asymp_\rho \psi(s) \quad (0 \le s \le a_\rho), \qquad \min\{1, nF(s)\} \asymp_\rho \min\{1, n\psi_{\mathrm{cl}}(s)\} \quad (n \ge 1, s \ge 0). \tag{25}$$

Proof. In the regular case set $\chi = \Psi$ and $\phi = 0$. Since $\psi'(0) > 0$, $\psi(s) \ge cs$ near zero, equivalently $\Gamma(t) \le t/c$ near zero.

In the singular case, choose $0 < u_1 < u_0$ with $\psi'(u_1) > 0$. Such a point exists because ψ is strictly increasing and its nonnegative derivative is nonincreasing near zero. Set $\phi = \psi$ on $[0, u_1]$ and extend it linearly beyond u_1 with slope $\psi'(u_1)$. Its derivative is nonnegative, nonincreasing, and convex: at the joining point its slope can only jump upwards to zero. Thus ϕ is continuously differentiable away from zero. Set $\chi = \psi - \Phi$ on the attainable signed interval. It vanishes near zero. On a compact interval separated from zero, ψ' is Dini-continuous by assumption; the extended derivative of ϕ is locally Lipschitz, since a finite convex function is locally Lipschitz in the interior of its domain. Their difference is Dini-continuous. A finite covering, together with subadditivity of moduli, gives the condition across the junctions. Extend χ linearly beyond the attainable endpoints and use odd reflection. The one-sided endpoint condition ensures that this extension is C^1 with a Dini-continuous derivative on $[-1, 1]$. Concavity gives $\psi(s) \ge [\psi(u_1)/u_1]s$ for $s \le u_1$, which implies the required upper bound on Γ near zero.

For an arbitrary member of $\mathcal{C}$, the bound on Γ near zero gives $\psi(s) \ge c_\rho s$ there. Strict positivity and compactness extend this inequality to $[0, a_\rho]$. Let $K_\chi = \|\chi'\|_\infty$. Then $|\chi(s)| \le K_\chi s$, and

$$\psi(s) \le K_\chi s + \phi(s), \qquad \phi(s) \le \psi(s) + K_\chi s.$$

Combining these inequalities proves $F \asymp_\rho \psi$ on the attainable positive interval. For $s > a_\rho$, monotonicity gives $F(s) \ge F(a_\rho) > 0$, while $\psi_{\mathrm{cl}}(s) = 1$. After capping at one, the remaining comparison follows uniformly over $n \ge 1$. ■

D Bounds for Antithetic Multilevel Estimation

D.1 A pointwise curvature bound

Lemma 14. *For Φ as in Definition 12, put $D(t) = 2\phi(t) - \phi(2t)$. Then $D \ge 0$ is nondecreasing, $D(ct) \le c^2 D(t)$ for $c \ge 1$, and, with $d = |x - y|/2$,*

$$\left|\Phi((x+y)/2) - \tfrac{1}{2}[\Phi(x) + \Phi(y)]\right| \le 2D(d) \qquad (x, y \in \mathbb{R}). \tag{26}$$

For $s_\ell = 2^{-\ell/2}$ and $L \ge 1$,

$$\sum_{\ell=1}^{L} \frac{D(s_\ell)}{s_\ell} \le 4\frac{\phi(s_L)}{s_L}. \tag{27}$$

Proof. We provide the representation argument, including nonsmooth derivatives. The convex, nonnegative, nonincreasing function ϕ' has a nondecreasing right derivative $(\phi')'_+ \le 0$. Let $k(t) = -(\phi')'_+(t)$. It is nonnegative and nonincreasing, and $k(t) \to 0$ as $t \to \infty$; otherwise ϕ' would eventually become negative. Let μ_ϕ be the nonnegative measure with tail $\mu_\phi((t, \infty)) = k(t)$, and set $\ell_0 = \lim_{t\to\infty} \phi'(t) \ge 0$. Writing $(x)_+ = \max\{x, 0\}$, local absolute continuity of convex functions and Tonelli's theorem give

$$\phi'(t) = \ell_0 + \int_t^\infty k(v)\, dv = \ell_0 + \int_{(0,\infty)} (a - t)_+ \, \mu_\phi(da).$$

This remains valid for an infinite total measure near zero. Integrating from zero and using $\phi(0) = 0$ gives

$$\Phi(x) = \ell_0 x + \int_{(0,\infty)} Q_a(x)\, \mu_\phi(da), \qquad Q_a(x) = \mathrm{sgn}(x) \begin{cases} a|x| - |x|^2/2, & |x| \le a, \\ a^2/2, & |x| > a. \end{cases} \tag{28}$$

The integral is finite at every x because $\phi(|x|)$ is finite. The integration from zero is justified by first integrating from $\epsilon > 0$ and then using monotone convergence and continuity of ϕ at zero.

For $t \geq 0$, define $D_a(t) = 2Q_a(t) - Q_a(2t)$. Direct calculation gives

$$D_a(t) = \begin{cases} t^2, & 0 \leq t \leq a/2, \\ 2at - t^2 - a^2/2, & a/2 \leq t \leq a, \\ a^2/2, & t \geq a. \end{cases}$$

Thus $D_a(t) \geq \frac{1}{2}\min\{t^2, a^2\}$, D_a is nondecreasing, and $D_a(t)/t^2$ is nonincreasing. Since $Q_a'(x) = (a - |x|)_+$ is 1-Lipschitz, cancellation of its linear Taylor terms bounds the absolute antithetic correction by $d^2/2$. The bound $|Q_a| \leq a^2/2$ also bounds it by a^2. Hence it is at most $\min\{d^2/2, a^2\} \leq 2D_a(d)$. Integrate against μ_ϕ; the linear part cancels. This proves Equation (26) and both properties of D.

For the sum, put $h(t) = \phi(t)/t$. Concavity makes h nonincreasing. Since $2s_\ell = s_{\ell-2}$,

$$\frac{D(s_\ell)}{s_\ell} = 2[h(s_\ell) - h(s_{\ell-2})].$$

Telescoping yields

$$\sum_{\ell=1}^{L} \frac{D(s_\ell)}{s_\ell} = 2[h(s_L) + h(s_{L-1}) - h(1) - h(\sqrt{2})] \leq 4h(s_L).$$

The formula also holds for $L = 1$, with the repeated $h(1)$ terms cancelling. This proves Equation (27). ■

D.2 Bias, moments, and level allocation

Lemma 15. *Let Y_i be iid signs with mean $\mu \in [-a_\rho, a_\rho]$ and define Equation (11). Set $s_\ell = 2^{-\ell/2}$ and $K_\chi = \|\chi'\|_\infty$. For $L \geq 0$,*

$$\left|\mathbb{E}_\mu \Psi(\overline{Y}_{2^L}) - \psi(\mu)\right| \leq K_\chi s_L + 2\phi(s_L). \tag{29}$$

For $\ell \geq 1$,

$$\sup_\mu \|\Delta_\ell\|_{L^2(P_\mu)} \leq 4D(s_\ell) + 4s_\ell \omega_{\chi'}(s_\ell). \tag{30}$$

Proof. A nonnegative concave function vanishing at zero is subadditive. Its odd extension satisfies $|\Phi(x) - \Phi(y)| \leq 2\phi(|x - y|)$: on one side of zero use subadditivity, and across zero bound each of $\phi(|x|), \phi(|y|)$ by $\phi(|x - y|)$. Since $\mathbb{E}|\overline{Y}_m - \mu| \leq m^{-1/2}$, Jensen's inequality for ϕ and the Lipschitz bound for χ give Equation (29).

For a level $\ell \geq 1$, let $m = 2^{\ell-1}$ and $d = (A_\ell - B_\ell)/2 = m^{-1}\sum_{i=1}^m X_i$, where $X_i = (Y_i - Y_i')/2$. The X_i are independent and centered, with $\mathbb{E}X_i^2 = \mathbb{E}X_i^4 = (1-\mu^2)/2 \leq 1/2$. Expanding the fourth moment gives

$$\mathbb{E}d^4 = m^{-4}\left[m\mathbb{E}X_1^4 + 3m(m-1)(\mathbb{E}X_1^2)^2\right] \leq \frac{3}{4m^2} = 3s_\ell^4.$$

By Lemma 14, the absolute antithetic correction for Φ is at most $2D(|d|)$. Monotonicity and quadratic scaling of D imply

$$\mathbb{E}|\Delta_\ell^\Phi|^2 \leq 4D(s_\ell)^2 \mathbb{E}\max\{1, (|d|/s_\ell)^4\} \leq 16D(s_\ell)^2.$$

For the regular component, cancellation of the first-order Taylor terms gives $|\Delta_\ell^\chi| \leq |d|\omega_{\chi'}(|d|)$. A modulus of continuity on an interval is subadditive, hence

$$|d|\omega_{\chi'}(|d|) \leq 2s_\ell \omega_{\chi'}(s_\ell)\max\{1, (|d|/s_\ell)^2\}.$$

Its second moment is at most $16s_\ell^2\omega_{\chi'}(s_\ell)^2$. Minkowski's inequality proves Equation (30). If either bound vanishes, the corresponding increment vanishes almost surely. ■

Set

$$a_0 = \max\{1, |\Psi(-1)|, |\Psi(1)|\}, \qquad a_\ell = 4D(s_\ell) + 4s_\ell\omega_{\chi'}(s_\ell) \quad (\ell \geq 1), \tag{31}$$

$$v_\ell = a_\ell^2, \qquad c_\ell = 2^\ell, \tag{32}$$

$$Q_L = \sum_{\ell=0}^{L} \sqrt{v_\ell / c_\ell}, \qquad W_L = \sum_{\ell=0}^{L} \sqrt{v_\ell c_\ell}, \qquad q_\ell = \frac{\sqrt{v_\ell / c_\ell}}{Q_L}. \tag{33}$$

Omit zero-weight levels. Sample J with these probabilities independently of the observations and return $Z_L = \Delta_J / q_J$, using exactly c_J raw comparisons after symmetrization.

Lemma 16. *The estimator above satisfies*

$$\mathbb{E}_\mu Z_L = \mathbb{E}_\mu \Psi(\overline{Y}_{2^L}), \tag{34}$$

$$\nu_L := \mathbb{E} c_J = W_L / Q_L, \qquad \sup_\mu \mathbb{E}_\mu Z_L^2 \le M_L := Q_L W_L, \tag{35}$$

$$\nu_L M_L = W_L^2 \le C_\rho [F(s_L)/s_L]^2, \qquad \sup_\mu |\mathbb{E}_\mu Z_L - \psi(\mu)| \le C_\rho F(s_L). \tag{36}$$

Moreover $\nu_L \le 2^L$.

Proof. Both half means at level ℓ have the distribution of $\overline{Y}_{2^{\ell-1}}$, so $\mathbb{E}\Delta_\ell = \mathbb{E}\Psi(\overline{Y}_{2^\ell}) - \mathbb{E}\Psi(\overline{Y}_{2^{\ell-1}})$. Summing proves Equation (34). The moment and cost identities follow directly:

$$\mathbb{E} c_J = \sum q_\ell c_\ell = W_L / Q_L, \qquad \mathbb{E}_\mu Z_L^2 = \sum \frac{\mathbb{E}_\mu \Delta_\ell^2}{q_\ell} \le Q_L W_L.$$

For any other positive allocation p_ℓ, Cauchy–Schwarz gives

$$\left(\sum p_\ell c_\ell\right)\left(\sum v_\ell / p_\ell\right) \ge \left(\sum \sqrt{v_\ell c_\ell}\right)^2;$$

Thus Equation (33) minimizes the product of expected cost and the second-moment bound for these choices of v_ℓ.

The Dini condition implies $S_\chi := \sum_{\ell \ge 1} \omega_{\chi'}(s_\ell) < \infty$ by integrating on the geometric intervals $[s_\ell, s_{\ell-1}]$. Lemma 14 now gives

$$W_L \le a_0 + 16 \frac{\phi(s_L)}{s_L} + 4 S_\chi \le C_\rho \frac{F(s_L)}{s_L}.$$

The same bound for $L = 0$ follows by increasing C_ρ. The bias is Equation (29). Finally, every realized cost is at most 2^L, which bounds its expectation. ■

E Proof of the Joint Minimax Rate

E.1 The upper bound, including all budget choices

Fix a transfer function in $\mathcal{C}$. For an ideal expected budget $u \ge 16n$, choose a power $m = 2^L$ satisfying

$$\frac{u}{16n} < m \le \frac{u}{8n};$$

changing the strict lower inequality to a weak one at endpoints has no effect. Estimate each chain marginal independently using Z_L and use $K = \lfloor u/(n\nu_L) \rfloor$ steps. Since $\nu_L \le m \le u/(8n)$, $K \ge 8$ and $K \ge u/(2n\nu_L)$. Conditional expectation bounds the total comparison cost by u. Lemmas 8 and 16 imply

$$\begin{aligned} \mathbb{E}[\widehat{f}(\bar{w}) - f^\star] &\le C\, n \sqrt{n \nu_L M_L / u} + C_\rho n F(m^{-1/2}) \\ &\le C_\rho n F(\sqrt{n/u}). \end{aligned} \tag{37}$$

Indeed $m \asymp u/n$, and concavity gives $F(cs) \le cF(s)$ for $c \ge 1$. The same conclusion after capping at one holds for $u < 16n$ by using a fixed set: $nF(\sqrt{n/u}) \ge n\sqrt{n/u} > 1/4$.

Choose between the fixed output and the ideal run using their known upper bounds on expected suboptimality. The resulting risk is at most $C_\rho \min\{1, nF(\sqrt{n/u})\}$. Threshold rounding uses no comparisons and preserves this expectation. This rate satisfies the hypotheses of Proposition 11, as shown in Appendix B; applying it yields the upper bound under a fixed budget. Lemma 13 replaces F with ψ_{cl}.

Enforcing the comparison budget. For each block $B = 1, 2, 4, \ldots$ fitting the geometric allocation, set $u = B/4$ and compute the preceding L, bounds, K, and step size; use a fixed set when the small-budget branch applies. Initialize each run at $w = \mathbf{1}/2$. At each iteration sort w, estimate all chain marginals by Equation (33), and then update. Accumulate only pre-update iterates of fully completed steps. If the run reaches its prescribed K steps, replace the incumbent vector by their average. Apply the restart procedure to the bounded Lovász loss and round the final incumbent once; this gives the same expected loss as retaining a rounded output from each independent completed run. If it exceeds the block allowance, retain the earlier completed output. A finite level's full batch may be rejected before drawing its observations when it does not fit; for an outcome-dependent scalar stopping rule, interrupt at the next raw comparison. Both conventions realize an interruption whenever the corresponding ideal run would exceed the block.

E.2 An adaptive modular lower bound

We use the modular family of Ito (2019), also used for the dueling lower bound of Sho & Ito (2026), and choose its marginal gap through the inverse transfer function.

Lemma 17 (Lower bound from modular functions)**.** *For every transfer function in the model, there are $c > 0$ and $c_\rho > 0$ such that*

$$R^{(1)}_{n,T}(\rho) \geq c \min\{1, n\psi_{\mathrm{cl}}(c_\rho\sqrt{n/T})\}. \tag{38}$$

Without a scaling property of the inverse, the constant inside ψ_{cl} must be retained.

Proof. For $v \in \{-1, 1\}^n$ and $0 < \tau \leq 1/(2n)$, define

$$f_v(S) = \tfrac{1}{2} + \tau \sum_{i \in S} v_i. \tag{39}$$

These functions are modular and take values in $[0, 1]$. Their minimizer includes exactly the negative coordinates. From an output S define $\widehat{v}_i = -1$ for $i \in S$ and $\widehat{v}_i = 1$ otherwise. Then, with d_H denoting Hamming distance,

$$f_v(S) - \min_U f_v(U) = \tau d_H(v, \widehat{v}).$$

Let $v^{(i)}$ flip bit i. A 1-local query involving i has gap $\pm\tau v_i$; flipping the bit changes its success probability by exactly $\Gamma(\tau)$. All other queries have the same law. If N_i counts queries involving i, the adaptive KL chain rule, with the internal seed included, gives

$$\mathrm{KL}(P_v \| P_{v^{(i)}}) \leq K_\rho \Gamma(\tau)^2 \mathbb{E}_v N_i.$$

An algorithm that stops early can be padded by uninformative self-comparisons, so the transcript has a common horizon. Pathwise $\sum_i N_i \leq T$. Pairing hypotheses that differ in one bit, the sum of their bit-error probabilities is at least one minus their total variation distance $\mathrm{TV}(P_v, P_{v^{(i)}})$. Pinsker's inequality and Jensen's inequality therefore give

$$\begin{aligned} 2^{-n} \sum_v \mathbb{E}_v d_H(v, \widehat{v}) &\geq \frac{n}{2}\left(1 - \frac{1}{n2^n} \sum_{v,i} \mathrm{TV}(P_v, P_{v^{(i)}})\right) \\ &\geq \frac{n}{2}\left(1 - \sqrt{\frac{K_\rho T \Gamma(\tau)^2}{2n}}\right). \end{aligned}$$

Choose

$$\tau = \min\left\{\frac{1}{2n}, \psi_{\mathrm{cl}}(c_\rho\sqrt{n/T})\right\}.$$

In each branch $\Gamma(\tau) \leq c_\rho\sqrt{n/T}$, including when the inverse is clipped. Choose c_ρ so the square root is at most $1/2$. The average Hamming error is then at least $n/4$. Multiplication by τ and $n\tau \geq \frac{1}{2}\min\{1, n\psi_{\mathrm{cl}}(c_\rho\sqrt{n/T})\}$ prove the claim. The uniform-prior Bayes risk is a lower bound on worst-case risk. ■

For a transfer function in $\mathcal{C}$, Equation (25) and concavity imply

$$\min\{1, n\psi_{\mathrm{cl}}(c_\rho s)\} \gtrsim_\rho \min\{1, nF(c_\rho s)\} \gtrsim_\rho \min\{1, nF(s)\} \gtrsim_\rho \min\{1, n\psi_{\mathrm{cl}}(s)\},$$

reducing c_ρ to at most one if needed. Combine this with Equation (38) and the upper bound to prove Theorem 4. In the regular case $F(s) = s$, proving Theorem 2. This also shows precisely where a constant inside the inverse is removed and which hypothesis permits it.

E.3 Power transfer functions with constants uniform in their amplitude

For $q > 1$, set $\alpha = 1/q$ and $\Psi(u) = \lambda^{-\alpha}\,\mathrm{sgn}(u)|u|^\alpha$ on $[-1, 1]$. It is a valid extension and has Hölder constant at most $2\lambda^{-\alpha}$. In Equation (11),

$$\sup_\mu \mathbb{E}\Delta_\ell^2 \le 4\lambda^{-2\alpha}2^{-\alpha\ell} \quad (\ell \ge 1), \qquad \mathbb{E}\Delta_0^2 = \lambda^{-2\alpha},$$

by the Hölder bound and the second moment of $A_\ell - B_\ell$. The bias is at most $2\lambda^{-\alpha}2^{-\alpha L/2}$. With these bounds, Equation (33) gives $\nu_L M_L \le C_q\lambda^{-2\alpha}(2^L)^{1-\alpha}$. The upper-bound proof in Appendix E.1 therefore gives

$$R^{(1)}_{n,T}(\rho) \le C_q \min\{1, \lambda^{-1/q}n^{1+1/(2q)}T^{-1/(2q)}\}.$$

For $q = 1$, $Z = Y/\lambda$ has unit cost and second moment λ^{-2}, so Equation (23) proves the same endpoint.

For the matching bound, use Equation (39). If $z = \lambda\tau^q \le 1/2$, the response KL after a sign flip is

$$z \log \frac{1+z}{1-z} \le 4z^2.$$

The same adaptive argument with $\tau = c_q \min\{n^{-1}, (n/(\lambda^2 T))^{1/(2q)}\}$ gives risk at least $c_q \min\{1, n^{1+1/(2q)}(\lambda^2 T)^{-1/(2q)}\}$. Thus the constants can be chosen independent of $0 < \lambda < 1$, as stated in Table 1. The uniformity assertion is specific to this direct power calculation; the general theorem uses constants that depend on ρ.

F Unbiased Estimation and an Information Lower Bound

Lemma 18 (Information lower bound for sequential estimation)**.** *Consider a sequential scalar estimator using one fixed set pair, with possibly adaptive orientation, and suppose* $\sup_\theta \mathbb{E}_\theta N \le \nu$ *and* $\sup_\theta \mathbb{E}_\theta Z^2 \le M$*. If its bias is uniformly at most* β*, then for* $0 < s \le a_\rho$*,*

$$\nu M \ge c_\rho \frac{(\psi(s) - 2\beta)_+^2}{s^2}. \tag{40}$$

In particular, such an unbiased estimator requires $\psi(s) = O(s)$ *near zero.*

Proof. Compare the parameters $\theta = 0$ and $\theta = \psi(s)$. For every adaptive orientation, Lemma 7 bounds the KL divergence for one comparison by $K_\rho s^2$. For the transcript stopped at $N \wedge m$, the KL chain rule gives at most $K_\rho s^2 \mathbb{E}_{\psi(s)}(N \wedge m) \le K_\rho \nu s^2$. The stopping decision and the independent seed add no conditional KL. The full stopped transcript is the limit of these finite prefixes. The variational characterization of relative entropy, or monotone convergence of relative entropy on increasing transcript sigma-fields, therefore gives

$$\mathrm{KL}(P_{\psi(s)} \| P_0) \le K_\rho \nu s^2.$$

For any two laws P, Q and output whose second moments are at most M, Cauchy–Schwarz on square-root densities yields

$$\begin{aligned} |\mathbb{E}_P Z - \mathbb{E}_Q Z| &\le \left(\int Z^2(\sqrt{dP} + \sqrt{dQ})^2\right)^{1/2} \left(\int (\sqrt{dP} - \sqrt{dQ})^2\right)^{1/2} \\ &\le 2\sqrt{M\,\mathrm{KL}(P\|Q)}. \end{aligned}$$

The last step uses $\int(\sqrt{dP} - \sqrt{dQ})^2 \le \mathrm{KL}(P\|Q)$; it follows, for example, from Jensen's inequality applied to the Hellinger affinity. Uniform bias gives $|\mathbb{E}_{\psi(s)}Z - \mathbb{E}_0 Z| \ge (\psi(s) - 2\beta)_+$. Combining the bounds proves Equation (40), with $c_\rho = (4K_\rho)^{-1}$. ∎

Proof of Theorem 5. We work in $\mathcal{C}$, which contains every transfer function satisfying Assumption 3. Since ϕ is concave and vanishes at zero,

$$\lim_{s\downarrow 0} \phi(s)/s = \phi'(0+) \in [0,\infty].$$

The comparisons in Lemma 13 show that $\psi(s) = O(s)$ is equivalent to $\phi'(0+) < \infty$. Combined with the lower linear bound $\psi(s) \geq c_\rho s$, it is also equivalent to $\Gamma(t) \asymp t$ near zero.

If $\phi'(0+) < \infty$, then Lemma 14 bounds $\sum_{\ell\geq 1} D(s_\ell)/s_\ell$ by $4\phi'(0+)$. The Dini sum is finite, so $W_\infty < \infty$ in Equation (33); also $0 < Q_\infty \leq W_\infty < \infty$. Use the infinite allocation $q_\ell = \sqrt{v_\ell/c_\ell}/Q_\infty$. It has finite expected cost W_∞/Q_∞ and uniform second moment at most $Q_\infty W_\infty$. Moreover,

$$\sum_{\ell\geq 0} \mathbb{E}|\Delta_\ell| \leq \sum_{\ell\geq 0} \sqrt{v_\ell} \leq W_\infty < \infty.$$

Thus its expectation equals the telescoping sum. The law of large numbers and bounded continuity of Ψ give $\mathbb{E}\Psi(\overline{Y}_{2^L}) \to \Psi(\mu) = \psi(\mu)$, proving exactness. Lemma 18 with $\beta = 0$ proves the converse, even for adaptive orientations of the original queries.

Finally, Theorem 4 and $\psi(s) \asymp s$ give Equation (9). Conversely, Equation (9) in dimension one, together with the lower bound in Theorem 4, implies $\psi(T^{-1/2}) \leq C_\rho T^{-1/2}$ for all sufficiently large integer T. Monotonicity and the bounded ratios between adjacent values $T^{-1/2}$ extend this inequality to every sufficiently small positive s. This proves all four equivalences. ■

Dependence on the estimation error. For the estimator with finitely many levels, the bias bound is a constant times $\psi(s_L)$ and the product of expected cost and second-moment bound is at most a constant times $[\psi(s_L)/s_L]^2$. Equation (40) supplies a lower bound at any separation s for which $\psi(s) > 2\beta$. To express this as a function of a prescribed error ϵ, one must keep constant multiples of ϵ inside Γ. For power transfer functions this gives order ϵ^{2-2q}, with the amplitude factor. For infinitely flat transfer functions, these constants can change an exponential, so they cannot generally be removed from the argument of Γ.

G The Modulus of Continuity of the Inverse

G.1 Elementary properties

Lemma 19. *The function $\Omega = \Omega_\rho$ is nondecreasing and subadditive, with $\Omega(0) = 0$, $\Omega(s) = 1$ for $s \geq a_\rho$, and*

$$\Omega(cs) \leq \lceil c\rceil \Omega(s), \qquad \Omega(s/c) \geq \Omega(s)/\lceil c\rceil \quad (c \geq 1). \tag{41}$$

The modulus of continuity of ψ on $[-a_\rho, a_\rho]$ is at most $2\Omega(s)$. For $0 < s \leq 1$, $\Omega(s) \geq s/2$.

Proof. The maximization in Equation (17) is over a compact nonempty set, so its maximum exists. Strict increase gives $\Omega(0) = 0$; the other initial claims are immediate. Given $[a,b]$ with $\Gamma(b) - \Gamma(a) \leq s + t$, split at c with $\Gamma(c) = \Gamma(a) + s$ if this value is internal. The two subintervals have respective widths at most $\Omega(s)$ and $\Omega(t)$. If no split is needed, the full interval already has width at most $\Omega(s)$. This proves subadditivity, hence Equation (41).

Oddness reduces an interval not crossing zero to the positive case. If $x \leq 0 \leq y$ and $\Gamma(y) - \Gamma(x) \leq s$, then both y and $-x$ are at most $\Omega(s)$, proving the factor two. Finally, partition $[0, a_\rho]$ into at most $\lceil a_\rho/s\rceil$ intervals of length at most s. The lengths of their images under ψ sum to one, so

$$\Omega(s) \geq 1/\lceil a_\rho/s\rceil \geq s/(a_\rho + s) \geq s/2.$$

If $s \geq a_\rho$, the conclusion follows from $\Omega(s) = 1$. ■

G.2 A lower bound for an arbitrary interval

Fix $s > 0$ and choose $0 \leq a \leq b \leq 1$ attaining $d = b - a = \Omega(s)$. Define functions on two elements by

	$\varnothing$	$\{1\},\{2\}$	$\{1,2\}$
f_0	0	b	d
f_1	d	b	0

Both are $[0,1]$-valued and submodular, since $2b \geq d$. Under the two hypotheses, every hypercube edge changes its gap between a and b, between their negatives, or with the direction reversed. Lemma 7 bounds the KL divergence per comparison by $K_\rho s^2$, so any adaptive transcript of at most T comparisons has KL at most $K_\rho T s^2$.

Use the output to decide hypothesis one only when it is $\{1,2\}$. A wrong decision incurs suboptimality at least d: under f_0 the chosen endpoint has suboptimality d, and under f_1 every other set has value at least d. Taking $s = c_\rho T^{-1/2}$ makes total variation at most $1/2$, hence the average error probability is at least $1/4$. Thus

$$R^{(1)}_{2,T}(\rho) \geq \tfrac{1}{4}\Omega(c_\rho T^{-1/2}) \gtrsim_\rho \Omega(T^{-1/2}),$$

where Equation (41) removes the constant without any shape assumption. Adding dummy elements embeds the same construction in every $n \geq 2$. Queries on dummy coordinates give no additional information, and an output maps to its intersection with the original two elements.

G.3 The spanning-tree upper bound

Estimate a marginal θ by m symmetrized comparisons, with mean $\overline{Y}_m$, and return

$$\widehat{\theta} = \psi(\mathrm{Proj}_{[-a_\rho, a_\rho]}\, \overline{Y}_m).$$

Projection does not increase distance from the true mean. By Lemma 19,

$$\begin{aligned}\mathbb{E}|\widehat{\theta} - \theta| &\leq 2\mathbb{E}\Omega(|\overline{Y}_m - \Gamma(\theta)|) \\ &\leq 2\Omega(m^{-1/2})\mathbb{E}[1 + \sqrt{m}\,|\overline{Y}_m - \Gamma(\theta)|] \leq 4\Omega(m^{-1/2}).\end{aligned} \tag{42}$$

The middle inequality follows from subadditivity, using $\Omega(x) \leq (1 + x/s)\Omega(s)$.

Take any rooted spanning tree of the hypercube of subsets, with $N_{\mathrm{tree}} = 2^n - 1$ edges. For $T \geq 2N_{\mathrm{tree}}$, estimate each oriented edge with $m = \lfloor T/N_{\mathrm{tree}} \rfloor$ comparisons, so the raw total is at most T. Reconstruct relative values along root-to-vertex paths and return a minimizer of the reconstructed values. The error in each reconstructed value is bounded by the sum of the N_{tree} absolute edge errors. If $\widetilde{f}$ denotes these relative values and $S^\star$ is optimal, then

$$f(\widehat{S}) - f(S^\star) \leq |f(\widehat{S}) - f(\varnothing) - \widetilde{f}(\widehat{S})| + |\widetilde{f}(S^\star) - f(S^\star) + f(\varnothing)|.$$

Using Equation (42), $m \geq T/(2N_{\mathrm{tree}})$, and Equation (41), the risk is at most

$$8N_{\mathrm{tree}}\Omega(\sqrt{2N_{\mathrm{tree}}/T}) \leq 8N_{\mathrm{tree}}\lceil\sqrt{2N_{\mathrm{tree}}}\rceil\Omega(T^{-1/2}). \tag{43}$$

For $T < 2N_{\mathrm{tree}}$, a fixed output has risk at most one, absorbed by a dimension-dependent multiple of $\Omega(T^{-1/2}) \geq T^{-1/2}/2$. This proves Equation (18), with the lower bound above.

The inequality in Theorem 6 is equivalent to Lipschitz continuity of ψ, or $\Omega(s) = O(s)$. One direction follows directly from Equation (17). For the other, apply $b - a \leq \Omega(\Gamma(b) - \Gamma(a))$. A $T^{-1/2}$ rate and the lower bound first yield $\Omega(T^{-1/2}) = O(T^{-1/2})$ on integer horizons; monotonicity fills the intervals between adjacent values. Compactness extends a small-s estimate to larger s. This proves the stated necessary-and-sufficient criterion.

G.4 The analytic example and the modulus of its inverse

The derivative in Equation (13) is nonnegative and vanishes only at $\pm 1/2$, so ρ_* is strictly increasing. Also $\rho_*(1) = 23/240 < 1$, and oddness puts its entire range inside $(-1,1)$. For $0 \leq x < y \leq 1$, with $d = y - x$,

$$\rho_*(y) - \rho_*(x) = \int_x^y (t + 1/2)^2(t - 1/2)^2\, dt \geq \frac{1}{4}\int_x^y (t - 1/2)^2\, dt \geq \frac{d^3}{48}.$$

The last integral is minimized by centering an interval of width d at $1/2$, where it equals $d^3/12$. Thus $\Omega_{\rho_*}(s) \leq (48s)^{1/3}$, capped at one. Conversely, direct polynomial expansion gives

$$\rho_*(1/2 + \tau) - \rho_*(1/2 - \tau) = \tfrac{2}{3}\tau^3 + \tfrac{2}{5}\tau^5.$$

For small enough s, choose $\tau = cs^{1/3} \le 1/8$ to make the signal at most s; this gives $\Omega_{\rho_*}(s) \ge 2\tau \gtrsim s^{1/3}$. For the remaining compact range of s, adjust the constants. Hence $\Omega_{\rho_*}(s) \asymp \min\{1, s^{1/3}\}$, and Theorem 6 proves Equation (14).

For the concrete family Equation (15), submodularity is $3/2 \ge 1/2 + v\tau$, which holds for the stated range. The value range is contained in $[0, 1]$. The changing edges have gaps $-1/2 + v\tau$ and their negatives. Equation (16) and probabilities bounded away from zero and one give a KL bound for one comparison $C\tau^6$. The KL of any adaptive transcript is at most $CT\tau^6$. Under $v = 1$, outputting $\{1, 2\}$ costs τ; under $v = -1$, every other output costs at least τ. Choosing a sufficiently small constant in $\tau = cT^{-1/6}$ gives a total variation bound $1/2$ and average suboptimality at least $\tau/4$. This direct argument explains the bottleneck used in the experiments.

G.5 Comparison with the inverse at zero

For $\mu, \nu \in [0, a_\rho]$, Equation (24) and the modulus bound in Lemma 15 imply

$$|\psi(\nu) - \psi(\mu)| \le K_\chi |\nu - \mu| + 2\phi(|\nu - \mu|) \le C_\rho F(|\nu - \mu|).$$

The first comparison in Equation (25) thus gives $\Omega(s) \lesssim_\rho \psi_{\mathrm{cl}}(s)$ after handling $s > a_\rho$ by clipping. Intervals starting at zero give $\Omega(s) \ge \psi_{\mathrm{cl}}(s)$. This proves the interpretation following Theorem 6.

H Examples and Boundary Conditions

Regular sigmoid and probit. For sigmoid, $\psi(u) = 2\,\mathrm{arctanh}(u)/b$ on a compact interval strictly inside $(-1, 1)$. For probit, $\psi(u) = \Phi_{\mathrm{N}}^{-1}((1+u)/2)/b$ on its compact attainable interval. Each has a smooth odd extension to $[-1, 1]$ and a positive derivative at zero, satisfying the regular case of Assumption 3. A linear extension at the attained endpoints is already C^1 with a Lipschitz derivative, though a smooth continuation is also possible.

Power transfer functions with logarithmic corrections. Suppose near zero $\Gamma(t) = ct^q \log(e/t)^a$, with $q > 1$ and fixed $a \in \mathbb{R}$, and suppose the inverse is C^1 with a Dini-continuous derivative away from zero. Its inverse is asymptotic to a positive constant times $u^{1/q} \log(1/u)^{-a/q}$. This follows by substitution: for $t = u^{1/q} \log(1/u)^{-a/q}$, $\log(e/t) \sim q^{-1} \log(1/u)$, and the remaining constant is absorbed in the leading multiplier. Direct differentiation and

$$\psi' = \frac{1}{\Gamma'}, \qquad \psi'' = -\frac{\Gamma''}{(\Gamma')^3}, \qquad \psi''' = \frac{3(\Gamma'')^2 - \Gamma'\Gamma'''}{(\Gamma')^5}$$

show $\psi' > 0$, $\psi'' < 0$, and $\psi''' > 0$ sufficiently near zero. Indeed the leading coefficients are those of t^q, for which $\Gamma'' > 0$ and $3(\Gamma'')^2 - \Gamma'\Gamma'''$ has positive coefficient $c^2q^2(q-1)(2q-1)$; all log derivatives are lower by relative factors $O(1/\log(1/t))$. Thus Assumption 3 applies and gives the corresponding exact logarithmic factor in Equation (10).

Logarithmic correction to a linear transfer function. For $\Gamma(t) = ct/\log(e/t)^a$, $a > 0$, write $z = \log(e/t)$. Then $\Gamma' = c(z^{-a} + az^{-a-1})$ and $\Gamma'' = ca\,t^{-1}[z^{-a-1} + (a+1)z^{-a-2}] > 0$. Differentiating again gives a negative leading term $-ca\,t^{-2}z^{-a-1}$, so $\Gamma''' < 0$ for sufficiently small t. The inverse identities above give the required signs. Substitution yields $\psi(u) \asymp u \log(1/u)^a$, proving the fourth row of Table 1 when the inverse is regular away from zero.

Infinite-order flatness. For $\rho(t) = \mathrm{sgn}(t)\exp(-|t|^{-a})$, $a > 0$, the positive inverse is $\psi(u) = \log(1/u)^{-\alpha}$ with $\alpha = 1/a$. Put $z = \log(1/u)$. Then

$$\psi' = \alpha u^{-1} z^{-\alpha-1}, \quad \psi'' = \alpha u^{-2} z^{-\alpha-2}(\alpha + 1 - z),$$

and

$$\psi''' = \alpha u^{-3} z^{-\alpha-3}\big[2z^2 - 3(\alpha+1)z + (\alpha+1)(\alpha+2)\big].$$

For sufficiently large z, the signs are positive, negative, and positive, respectively. Away from zero the inverse is smooth, so Assumption 3 holds. The clipped inverse expression is comparable to $\min\{1, n/\log(e + T/n)^{1/a}\}$, as in Table 1. When deriving comparison complexity from this rate, constants inside Γ must be retained because they change the exponential dependence on the error.

Unbiased estimation without a Dini-continuous derivative. For small $u > 0$, set $\phi(u) = \int_0^u [1 - 1/\log(e/v)]\, dv$. On a sufficiently small interval,

$$\phi'(0+) = 1, \qquad \phi''(u) = -\frac{1}{u \log(e/u)^2} < 0, \qquad \phi'''(u) = \frac{\log(e/u) - 2}{u^2 \log(e/u)^3} > 0.$$

Extend ϕ linearly after an interior point. For its odd extension, the modulus of continuity of the derivative is at least $1/\log(e/u)$, whose Dini integral diverges. Scaling this extension so it equals one at an attainable endpoint $a_\rho < 1$ gives a valid inverse transfer function. Nevertheless $\phi'(0+) < \infty$, so Theorem 5 gives an unbiased estimator with uniformly bounded expected cost and second moment. The curvature condition therefore covers examples beyond inverses with a Dini-continuous derivative.

I Bias in Marginal Estimates Can Change the Minimizer

For $k \in \mathbb{N}$ and $b > 0$, define

$$\mu_{k,b}(\theta) = \sum_{j=0}^{k} \binom{k}{j} \sigma(b\theta)^j (1 - \sigma(b\theta))^{k-j} \frac{1}{b} \log \frac{j + 1/2}{k - j + 1/2}.$$

The summand's transformed value is strictly increasing in j, so binomial stochastic monotonicity makes $\mu_{k,b}$ strictly increasing. Binomial symmetry makes it odd. Here $\mu_{k,b}$ is the expected marginal estimate produced by the Firth estimator used in Sho & Ito (2026), building on Firth (1993).

Proposition 20 (Convergence to the wrong minimizer)**.** *Let $c \in (0, 1/2)$, $a = \mu_{k,b}(1)$, and $d = \mu_{k,b}(c)$, and suppose $\gamma = 2d - a > 0$. Define a three-element cardinality function f_c by values $0, 1, 1 - c, 1 - 2c$. Run projected stochastic subgradient descent on $[0, 1]^3$ from $\mathbf{1}/2$, using coordinatewise fixed-k Firth estimates and step size $h/\sqrt{K}$ for a fixed $h > 0$. Let $\widehat{S}_K$ be a threshold rounding of the average of the K pre-update iterates. There is a constant C, depending only on k, b, h, such that*

$$\mathbb{P}(\widehat{S}_K \neq [3]) \le \frac{C}{\gamma\sqrt{K}}, \qquad \left|\mathbb{E} f_c(\widehat{S}_K) - (1 - 2c)\right| \le \frac{C}{\gamma\sqrt{K}}.$$

Thus the expected suboptimality converges to $1 - 2c > 0$, although $\varnothing$ is the unique minimizer.

Proof. The increments $(1, -c, -c)$ are nonincreasing, so f_c is submodular. Its nonempty values are positive. Define a surrogate cardinality function $\widetilde{f}$ with values $0, a, a - d, a - 2d$. Strict monotonicity and oddness of $\mu_{k,b}$ give $a > d > 0$. The surrogate is therefore submodular, and the conditional mean Firth vector is a greedy subgradient of its Lovász extension, for the same chain and tie rule. Its unique minimizer is $[3]$, with value $-\gamma$, and every other set has surrogate suboptimality at least γ.

Each scalar Firth estimate has absolute value at most $B_{k,b} = \log(2k+1)/b$. The projected subgradient calculation in Lemma 8, applied to the surrogate, gives expected suboptimality at most $C/\sqrt{K}$, where $C = 3/(8h) + 3hB_{k,b}^2/2$. Threshold rounding preserves this expectation. The surrogate gap then bounds $\mathbb{P}(\widehat{S}_K \neq [3])$. Since f_c takes values in $[0, 1]$ and $f_c([3]) = 1 - 2c$, the second inequality follows. ∎

For $c = 0.49$, $b = 2$, and $k = 10$, evaluation of the finite binomial sum gives

$$\mu_{10,2}(1) = 0.9559738403\ldots, \qquad \mu_{10,2}(c) = 0.4906887303\ldots, \qquad \gamma = 0.0254036203\ldots.$$

The true suboptimality at the wrong endpoint is 0.02. At $k = 20$, the surrogate's total increment is instead $\mu_{20,2}(1) - 2\mu_{20,2}(c) = 0.0152923384\ldots > 0$, and its unique minimizer is empty. These signs are also verified by interval evaluation of the finite sums. Complementation, $f_c^{\mathrm{rev}}(S) = f_c([3] \setminus S)$, preserves submodularity and exchanges the true and surrogate endpoints. On the balanced pair, the minimum expected suboptimality without observations is 0.01.

J Experimental Protocol and Additional Results

J.1 Implementation and statistical summaries

All comparisons are simulated. The optimizer uses only comparison responses and the known transfer function; function values enter the oracle simulator and evaluation. A binomial draw simulates a batch of independent comparisons, and a geometric draw simulates a waiting time. Each draw is charged for every underlying comparison. When a harmonic run reaches its cap, its partial output is discarded and the initialization is returned.

The subgradient methods start at $\mathbf{1}/2$ and return an average of pre-update iterates on completion. We store $\widehat{f}(\bar{w}) - f^\star$ and the suboptimality of a separately threshold-rounded set. Their expectations agree, but the former integrates out rounding randomness. We also store comparisons, completed iterations, iterations contributing to the returned average, and interruptions.

The release contains `run.py`, `engine.py`, `report.py`, and `validate.py`, with compressed per-run arrays and a file-hash manifest. Both suites use seed 20260926 and NumPy `SeedSequence` with distinct suite, configuration, and run indices. Exact package versions, column definitions, and commands are included in the release. Computation uses binary64 arithmetic.

Intervals are percentile 95% bootstrap intervals from 2,000 resamples, stratified by the two alternatives in each balanced experiment. They describe oracle and algorithm randomness on the specified family and are pointwise intervals without a multiple-comparison adjustment.

J.2 A fixed objective with a wrong surrogate minimizer

We use $c = 0.49$, $b = 2$, and the function in Proposition 20. Each arm and budget has 300 runs, half on f_c and half on f_c^{rev}. The six budgets are

$$30{,}000, \quad 100{,}000, \quad 300{,}000, \quad 1{,}000{,}000, \quad 3{,}000{,}000, \quad 10{,}000{,}000.$$

The objective is fixed throughout. For fixed $k \in \{10, 20\}$, Firth uses $K = \lfloor T/(3k) \rfloor$ and $\eta = (2\sqrt{M_b K})^{-1}$, where $M_b = 1 + \pi^2/(6b^2)$. The main harmonic arm uses the optimized count in Appendix A.5, with $\delta = 1/(4\sqrt{T})$ and the step size based on $V_{n,b}$. An additional harmonic arm matches the $k = 10$ arm's iteration count and step size. Every arm uses the same hard cap.

Published schedule. For $p_b = \sigma(b)$, put $C(b) = |(2p_b - 1)/(24bp_b^2(1-p_b)^2)|$. The implementation of Sho & Ito (2026, Algorithm 2) uses

$$k(T) = \max\{1, \lceil b^{2/5} C(b)^{2/5} (T/n)^{1/5} \rceil\}, \quad K = \lfloor T/(nk(T)) \rfloor, \quad \eta = b^{6/5} C(b)^{1/5} (n/T)^{2/5}.$$

At $T = 30{,}000$, 300,000, and 3,000,000, it selects $k = 10, 16, 25$ on the three-element instance. The corresponding mean marginal sums are approximately -0.0254, 0.00746, and 0.01920. Thus the growing batch eventually removes the reversal on this fixed family.

For an averaged vector w, the probability of selecting the full set is $\min_i w_i$, and that of selecting the empty set is $1 - \max_i w_i$. These give the wrong-endpoint diagnostic without additional rounding noise. At the largest budget, the probabilities are approximately 0.937 for fixed $k = 10$, 0.146 for fixed $k = 20$, 0.039 for the published schedule, and 0.065 for the single-run harmonic method. The published schedule has mean suboptimality 0.00080, compared with 0.00130 for the single-run harmonic method. The matched-iteration harmonic arm has suboptimality 0.00166 and wrong-endpoint probability 0.083. All 9,000 runs respect the cap; none of the harmonic runs interrupts.

J.3 Fixed-gap information benchmarks

Use

$$f_v(\varnothing) = 1/4, \qquad f_v(\{1\}) = f_v(\{2\}) = 1/4 + a, \qquad f_v(\{1,2\}) = 1/4 + v\tau.$$

We take $\tau = 0.05$ and $a \in [0.10, 0.70]$. The range lies in $[0, 1]$, and submodularity follows from $2a \geq v\tau$. The optimum is empty for $v = 1$ and full for $v = -1$; the other endpoint has suboptimality τ. Define $\bar{\rho} = 8\rho_*$, so $\bar{\rho}'(0) = 1/2$ and $|\bar{\rho}(1)| < 1$.

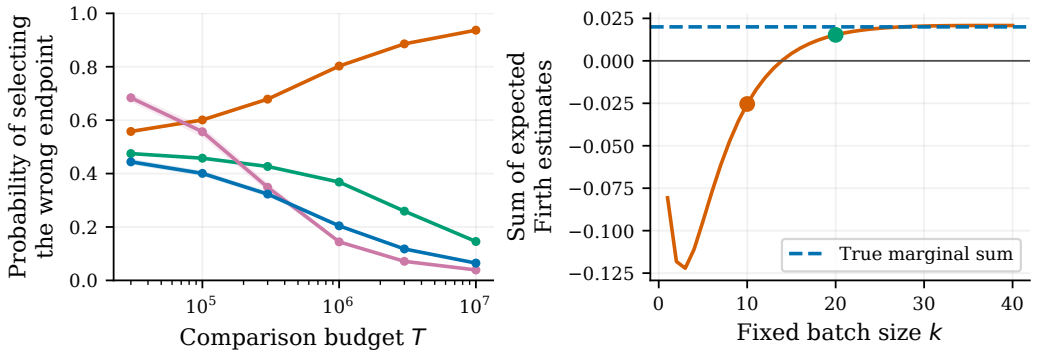

Figure 4: Controls for the experiment in Figure 2. Left: the conditional probability that threshold rounding returns the wrong endpoint, averaged over runs. Colors match Figure 2. Right: the sum of expected Firth marginal estimates, computed from the binomial probabilities at each batch size. The true sum is 0.02; the negative sign at $k = 10$ becomes positive at $k = 20$.

Optimal testing benchmark. For known (a, τ), comparisons between empty and singleton sets have a law independent of v. Each singleton-to-full comparison, after a known orientation change, has success probability in the pair

$$p_- = \frac{1 + \bar{\rho}(a - \tau)}{2}, \qquad p_+ = \frac{1 + \bar{\rho}(a + \tau)}{2}.$$

The two informative edges yield equivalent observations. Any adaptive procedure using 1-local queries can therefore be simulated from T draws of this coin and independent randomness. Querying one informative edge T times attains the resulting bound. For equal priors, its optimal error is

$$e_T(p_-, p_+) = \frac{1}{2} \sum_{j=0}^{T} \min\left\{ \binom{T}{j} p_-^j (1 - p_-)^{T-j}, \binom{T}{j} p_+^j (1 - p_+)^{T-j} \right\}.$$

We evaluate this expression through binomial tails at the likelihood-ratio threshold. Integer bisection finds the smallest T with $e_T \le 0.1$. No curve scale is fitted.

The left panel of Figure 3 uses 61 equally spaced values of a. At $a = 0.10, 0.25, 0.40, 0.50, 0.60, 0.70$, the budgets for the polynomial transfer function are respectively 3,126, 8,263, 140,194, 14,475,033, 55,359, and 2,804. The linear control $\rho(t) = t/2$ has the same derivative at zero and requires approximately 2,300–2,600 comparisons across this range. The direct endpoint test uses probabilities $(1 \pm \bar{\rho}(\tau))/2$ and requires 2,663 comparisons, independently of a. Its query changes two elements.

Spanning-tree solver. The spanning-tree solver is given the transfer function and $n = 2$. It estimates the three edges $(\varnothing, \{1\})$, $(\varnothing, \{2\})$, and $(\{1\}, \{1, 2\})$ with $\lfloor T/3 \rfloor$ comparisons each, inverts clipped empirical means, reconstructs relative values, and returns a minimizing set. It uses neither a nor τ. We run 4,000 balanced trials at each of the six singleton parameters above and 29 logarithmically spaced budgets from 100 to 10^9. Figure 3 shows $a = 0.50$. Its mean endpoint error is 0.152 at approximately 5.62×10^7 comparisons and 0.0785 at 10^8. All six parameters and the returned-set suboptimalities are retained in the release.

Strictly positive derivatives. To examine nearby regular transfer functions, set

$$\rho_\kappa(t) = 8\big(\rho_*(t) + \kappa t^3\big), \qquad 0 \le \kappa \le 0.01.$$

Each has slope $1/2$ at zero and range inside $(-1, 1)$. For $\kappa > 0$, $\rho_\kappa'(t) = 8[(t^2 - 1/4)^2 + 3\kappa t^2] > 0$ everywhere. At $a = 1/2$,

$$\rho_\kappa(1/2 + \tau) - \rho_\kappa(1/2 - \tau) = 8\left[\frac{3\kappa}{2}\tau + \left(\frac{2}{3} + 2\kappa\right)\tau^3 + \frac{2}{5}\tau^5\right].$$

The fixed-accuracy comparison scale is consequently $(\kappa\tau + \tau^3)^{-2}$. For fixed $\kappa > 0$, sufficiently small gaps give the usual quadratic dependence on τ^{-1}, while the cubic term governs larger gaps. Figure 5 evaluates the same binomial benchmark for $\kappa \in \{0, 10^{-4}, 10^{-3}, 10^{-2}\}$ and 21 gaps in $[0.01, 0.1]$.

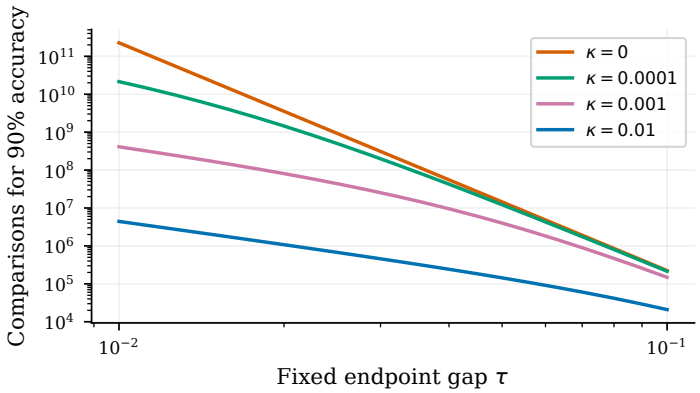

Figure 5: Testing with 1-local queries near the flat point and under regular perturbations. All transfer functions have derivative $1/2$ at zero; every $\kappa > 0$ makes the derivative positive everywhere. The comparison cost remains large near the former flat point, with a transition as τ^2 crosses the scale of κ.

K Numerical Consistency Checks

The released `validate.py` verifies array shapes and finiteness, stored comparison counts, completed versus retained iterations, and the submodularity and minimizers of the experimental families. It also checks the conditional exponential-moment schedules, the binomial likelihood-ratio calculation against direct finite sums, the strict Firth signs by interval arithmetic, and the telescoping identity using finite binomial sums. The largest telescoping residual is below 10^{-15}. The release includes the executed checks, their results, and hashes of the code and numerical arrays.