TIES WORTH KEEPING IN DUELING SUBMODULAR MINIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Many learning systems must choose a set but observe only which of two candidates is preferred. For submodular objectives, existing analyses combine the cost of reconstructing a Lovász subgradient with optimization. Under sigmoid feedback, we characterize the worst-case tradeoff between expected comparison count and mean squared error for reconstructing a canonical subgradient. This cost depends on which coordinates of the iterate are equal, with quadratic dimension dependence when all coordinates differ and linear dependence when all coincide. A Shapley identity yields an unbiased estimator. For regular, locally unbiased estimators, an electrical network argument gives a matching structural lower bound that covers arbitrary adaptive comparisons. Rounding the iterate creates ties at controlled bias, yielding a family of algorithms that improves the optimization bound over both the unquantized method and prior work at intermediate budgets. Experiments support the predicted dependence on tie structure and the optimization benefits of quantization. Across two objective families, one fitted constant describes the measured product of query cost and mean squared error within 7%; controlled optimization experiments reduce median error by up to a factor of 19 relative to the unquantized method.

1 INTRODUCTION

Machine learning systems increasingly decide on sets, and the signal that steers them is often a preference. A recommender picks a bundle for a user; a developer selects retrieved facts and tool calls for an answer. In either case feedback may be only a noisy verdict on which of two candidate sets is better, because a calibrated score for one set can be hard to obtain and trust while a direct comparison can be cheaper and more stable. When lower values are preferable and the set objective has diminishing returns, this is submodular minimization from noisy pairwise comparisons, the sigmoid dueling model of Sho & Ito (2026). Its two difficulties are that the feedback carries only relative information and the objective is discrete.

The standard way through the discreteness uses the Lovász extension to turn the set function into a convex one with the same minimum, then applies projected stochastic subgradient descent. Its descent direction consists of value gaps between nested sets. Yet no comparison delivers a gap; it reports only which of two sets is better, and reports even that unreliably. The only way across is repetition, and these repetitions are the sensing cost the learner pays for one direction. Prior work follows this route directly. It spends a fixed number of comparisons on each pair, converts win rates back into gaps, and hands the assembled direction to the optimizer. This approach yields the first rate for the problem, but that rate falls short of the square-root convergence familiar from stochastic convex optimization. The analysis does not explain why because it charges the sensing cost together with the optimization it feeds. A basic question has therefore never been asked. What does it cost, in its own right, to reconstruct a direction out of comparisons?

We answer this question by precisely characterizing the worst-case tradeoff between the expected number of comparisons and the accuracy of a canonical Lovász direction. Its dimension dependence is quadratic when every coordinate is distinct, linear when all are tied, and otherwise determined by the tie-block sizes (Figure 1). Equal coordinates leave their internal order free, allowing a Shapley wide-pair estimator to produce one unbiased random estimate of the block-averaged direction at constant expected duel cost. Averaging independent draws improves accuracy, with the tie partition

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.

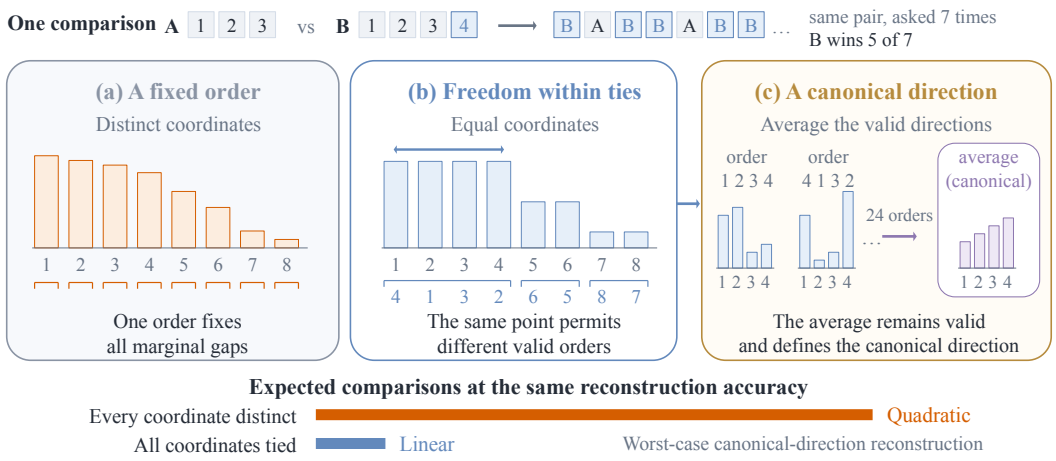

Figure 1: Why ties change the price of a direction. A comparison returns a noisy winner. (a) Distinct coordinates fix one sorted chain and its marginal gaps. (b) Equal coordinates permit multiple orders inside each block. (c) Two of the 24 orders are shown; averaging all 24 greedy directions gives the canonical subgradient. The four-coordinate example uses a cardinality function capped at two with unequal modular weights. The bars compare the quadratic and linear dimension dependence of the worst-case expected-query cost at fixed MSE for regular, locally unbiased canonical reconstruction. Intermediate costs depend on the tie-block sizes.

governing the resulting comparison cost. For locally unbiased reconstruction, a matching electrical-network lower bound establishes the same dependence on dimension and tie structure and covers arbitrary adaptive set comparisons. Rounding the iterate therefore creates ties at controlled bias and yields a one-parameter algorithm family that improves the tradeoff between dimension and budget, with square-root budget decay at its unquantized endpoint. The sensing law thus provides both a precise reconstruction limit and a design principle for optimization.

Contributions.

- **The hidden sensing bottleneck.** We separate preference-based submodular minimization into a sensing layer and an optimization layer, show that the bottleneck earlier analyses charged to the optimizer belongs to sensing, and identify the iterate's tie partition as what sets the sensing cost.
- **The price of a direction.** We give a tight expected-query–MSE law for canonical-direction reconstruction. A wide-pair estimator attains it, and a matching locally unbiased lower bound covers arbitrary adaptive choices of sets.
- **The gain from ties.** We turn the law into a design principle, rounding the iterate to manufacture the ties that make sensing cheap, and obtain a family of algorithms whose envelope improves the dimension–budget tradeoff of the best published rate.
- **Reconstruction limits and optimization directions.** Locally unbiased reconstruction of the exact greedy direction has quadratic sensing cost. We also identify the limit of a uniform-sensitivity Assouad certificate and construct a chord oracle that yields contraction when a suitable chord is available.

2 Background

Submodular minimization under comparison feedback. Submodular objectives cover graph cuts, coverage problems, and many structured losses. With value access, their minimization is settled, beginning with the ellipsoid method (Grötschel et al., 1981), followed by combinatorial algorithms that removed the reliance on numerical separation (Iwata et al., 2001; Schrijver, 2000), and later by cutting-plane methods that reduced the query count (Lee et al., 2015). When values cannot

be read, the Lovász extension still turns the set function into a convex one with the same minimum (Lovász, 1983; Bach, 2013), and stochastic subgradient descent provides the standard optimization route (Chakrabarty et al., 2017). In parallel, learning from pairwise comparisons developed from dueling bandits (Yue et al., 2012; Yue & Joachims, 2011) to tight regret bounds (Komiyama et al., 2015) and broad surveys (Bengs et al., 2021), then moved beyond finite arms to continuous domains with known comparison-complexity limits (Jamieson et al., 2012). These limits echo the information barriers for stochastic convex optimization with noisy first-order oracles (Nemirovsky & Yudin, 1983; Raginsky & Rakhlin, 2011; Agarwal et al., 2012). A comparison returns a random winner whose probability is determined by the value gap passed through a link function, classically the logistic link of Bradley & Terry (1952). When value access is replaced by pairwise comparisons, the Lovász convex geometry survives, but the numerical oracle that supplies its descent direction disappears.

From noisy winners to a Lovász direction. Dueling-bandit methods typically identify a preferred arm or control regret over a fixed family. In continuous domains, Saha et al. (2021) estimate normalized gradient directions from noisy comparisons for smooth convex optimization. Here the Lovász extension is generally nonsmooth, and the oracle compares sets. A greedy Lovász direction contains the marginal gaps along an entire sorted chain of sets, whereas a comparison carries only a winner compressed through the link. Finding a good set and reconstructing a complete direction are therefore different goals, and general comparison-complexity results do not price this conversion. Sho & Ito (2026) gave the first algorithm for it. Their method proceeds one coordinate at a time, spends a fixed number of comparisons on each adjacent pair in the sorted chain, converts the win rate back into a value gap through a bias-corrected logit (Firth, 1993), and hands the assembled direction to projected stochastic gradient descent.

The unpriced sensing layer. The resulting end-to-end rate does not determine how much of the reconstruction cost is unavoidable and how much can be reduced by changing the estimator. This distinction matters because ties leave the internal ordering free, allowing different greedy directions to be averaged into a valid subgradient. It remains to determine how this freedom affects the comparisons needed for accurate reconstruction. We characterize this dependence for the canonical direction obtained by uniform averaging, with matching bounds for locally unbiased reconstruction under arbitrary adaptive set comparisons. The resulting cost law motivates rounding query points to create ties, trading controlled approximation bias for lower sensing cost.

3 SETTING

We optimize a bounded submodular function from pairwise comparisons. Let $[n] = \{1, \dots, n\}$ and let $F : 2^{[n]} \to \mathbb{R}$ have range diameter at most $R > 0$,

$$\operatorname{osc}(F) := \max_{A \subseteq [n]} F(A) - \min_{A \subseteq [n]} F(A) \le R,$$

and satisfy

$$F(A) + F(B) \ge F(A \cup B) + F(A \cap B) \qquad \text{for all } A, B \subseteq [n].$$

The learner observes only duels. The diameter condition bounds every query gap and excess by R. Whenever a base polytope is used, we normalize $F(\emptyset) = 0$ by subtracting $F(\emptyset)$, preserving duels, minimizers and R. Write $g(S) = \sum_{i \in S} g_i$ and define $B(F) = \{g \in \mathbb{R}^n : g(S) \le F(S)\ \forall S \subseteq [n],\ g([n]) = F([n])\}$. Every $g \in B(F)$ satisfies $\|g\|_1 \le 2R$.

The Lovász extension defines a convex objective on $\mathbb{R}^n$, which we optimize over the hypercube. For $x \in \mathbb{R}^n$, let π order the coordinates as $x_{\pi_1} \ge \cdots \ge x_{\pi_n}$, and define prefixes $S_k = \{\pi_1, \dots, \pi_k\}$ with $S_0 = \emptyset$. The extension is

$$f(x) = F(\emptyset) + \sum_{k=1}^{n} x_{\pi_k} \big(F(S_k) - F(S_{k-1})\big).$$

A classical theorem states that f is convex exactly for submodular F (Lovász, 1983; Bach, 2013). Throughout, ∂f denotes the ambient subdifferential, while minima and minimizers of f are taken over $[0,1]^n$; projection handles this constraint. Moreover, $\min_{x \in [0,1]^n} f(x) = \min_{S \subseteq [n]} F(S)$, and threshold rounding of any $x \in [0,1]^n$ preserves the Lovász value in expectation.

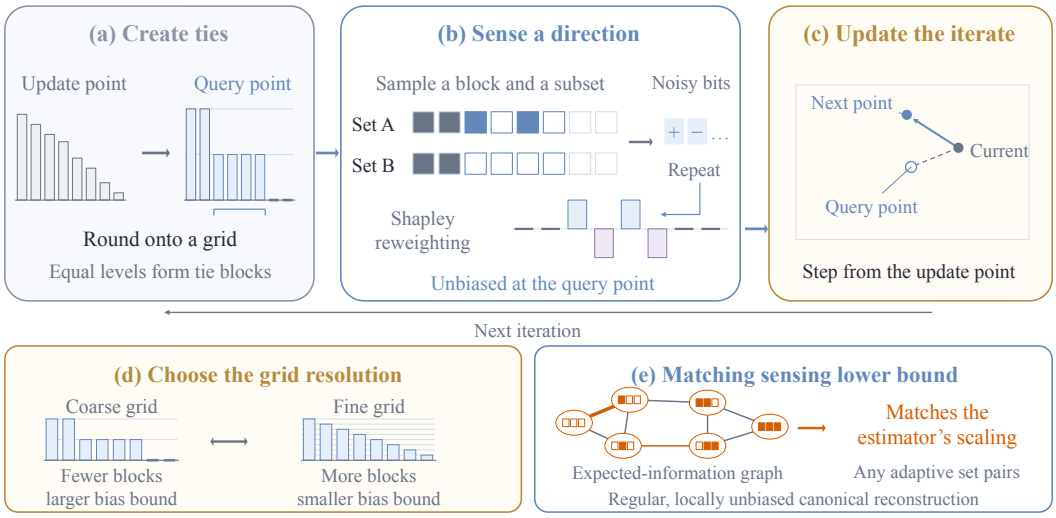

Figure 2: Overview of sensing and optimization. (a) Rounding creates tie blocks at a displaced query point. (b) A sampled wide pair, repeated comparisons, and Shapley reweighting yield an unbiased direction estimate there. (c) The estimate updates the original iterate. (d) Grid resolution balances the sensing term against quantization bias. (e) For regular, locally unbiased canonical reconstruction, the expected-information graph gives a lower bound with the same dimension and tie dependence as the estimator, up to constants depending on b, R. The loop illustrates an iteration of the method in Section 5.

The dueling oracle returns a noisy comparison through the sigmoid link. At round t, the learner may choose (A_t, B_t) using past queries and responses and its own randomization. Given the history $\mathcal{H}_{t-1}$ and the current query, the oracle returns $Y_t \in \{\pm 1\}$ with

$$\Pr(Y_t = 1 \mid \mathcal{H}_{t-1}, A_t, B_t) = \sigma\big(b(F(A_t) - F(B_t))\big), \qquad \sigma(u) = \frac{1}{1+e^{-u}},$$

where $b > 0$ is the transfer parameter. Each response uses fresh independent noise. Our guarantees assume known b and a valid known diameter bound R.

For an integer budget $T \geq 0$, the goal is to use at most T oracle queries and output a possibly randomized set $\widehat{S}_T \subseteq [n]$ with small expected excess value.

$$\mathbb{E}\big[F(\widehat{S}_T) - \min_{S \subseteq [n]} F(S)\big].$$

Expectations are over the oracle responses and the algorithmic randomization.

For the intermediate sensing problem, we allow an almost surely finite stopping time N and measure cost by $Q_\theta = \mathbb{E}_\theta N$, where θ is the set-value table. Local unbiasedness means that the estimator's mean and its derivative agree with the target at θ. We use regular estimators, with the stopped-likelihood differentiation condition given in Appendix A.5. The optimization algorithm enforces the hard budget T separately.

4 The Price of a Direction

4.1 Achievability with One Sampled Wide Pair

We fix x and study the comparisons needed to estimate its canonical direction at a given accuracy. Ties mean equal coordinates of x; the oracle still returns a binary winner.

Normalize $F(\varnothing) = 0$, so $f(x) = \max_{g \in B(F)} \langle g, x \rangle$ and $\langle q, x \rangle = f(x)$ for $q \in \partial f(x)$. Every $g \in B(F)$ has $\|g\|_1 \leq 2R$. At a point with levels $\lambda_1 > \cdots > \lambda_J$ and tie blocks $B_1, \ldots, B_J$, each ordering respecting the levels gives a greedy subgradient. Averaging uniformly over the orderings inside each block remains a subgradient at x and defines the canonical direction.

Proposition 1 (Canonical block subgradient). *Let $H_j = B_1 \cup \cdots \cup B_{j-1}$ and $h_j(U) = F(H_j \cup U) - F(H_j)$ for $U \subseteq B_j$. Let q_x be the concatenation over blocks of the Shapley vectors $\phi(h_j)$. Then $q_x \in \partial f(x)$.*

Sampling estimators for Shapley vectors (Shapley, 1953) are well studied (Castro et al., 2009; Lundberg & Lee, 2017). Here their exact representation in set values lets us sample one subset and reweight its value to estimate the whole block.

Lemma 2 (Whole-block estimation from a sampled wide pair). *Let $|B| = m$ and $h : 2^B \to \mathbb{R}$ with $h(\varnothing) = 0$. Then $\phi(h) = \sum_{U \subseteq B} h(U)\, c_m(U)$, where for $k = |U|$*

$$[c_m(U)]_i = \frac{1}{k\binom{m}{k}}\ (i \in U), \qquad [c_m(U)]_i = -\frac{1}{(m-k)\binom{m}{k}}\ (i \notin U),$$

and $\|c_m(U)\|_2^2 = a_{m,k}/\binom{m}{k}^2$ with $a_{m,k} = 1/k + 1/(m-k)$ for $1 \le k \le m-1$ and $1/m$ at $k \in \{0, m\}$. Draw K with probability $w_{m,k} \propto \sqrt{a_{m,k}}$, then U uniformly among K-subsets, so that $p_m(U) = w_{m,|U|}/\binom{m}{|U|}$. If Z_U is conditionally unbiased for $h(U)$ with $\mathbb{E}[Z_U^2 \mid U] \le V$, then $\widehat{\phi} = Z_U\, c_m(U)/p_m(U)$ satisfies

$$\mathbb{E}\widehat{\phi} = \phi(h), \qquad \mathbb{E}\|\widehat{\phi}\|_2^2 \le V\, C_m^2, \qquad C_m := \sum_{k=0}^{m} \sqrt{a_{m,k}},$$

and $2\sqrt{m} \le C_m \le 6\sqrt{m}$. This law minimizes the uniform second-moment bound $V \sum_k a_{m,k}/w_{m,k}$ among size distributions.

The value $h_j(U) = F(H_j \cup U) - F(H_j)$ is accessible through the *wide pair* $(H_j \cup U,\, H_j)$. An unbiased scalar estimate Z_U, multiplied by $c_m(U)/p_m(U)$, estimates the whole block. The weights are positive inside the sampled subset and negative outside it. Each draw therefore contributes to every coordinate of the block, and averaging over subsets recovers its individual Shapley marginals. Exact unbiasedness matters because the importance weights amplify residual bias. We randomize the truncation level following McLeish (2011), Rhee & Glynn (2015) and Lyne et al. (2015), with multilevel generalizations (Blanchet & Glynn, 2015; Vihola, 2018), and estimate each power by a product of independent draws (Hoeffding, 1948).

Lemma 14 in Appendix A gives such a Z for a duel of gap $|\Delta| \le R$, exactly unbiased with $\mathbb{E}Z = \Delta$ and $\mathbb{E}Z^2 \le V_{b,R} := 8s(1+s)/(b^2(1-s)^2)$ at $c_{b,R} = 4/(1-s)$ expected responses, where $s = \sigma(bR)$ is the saturation.

Theorem 3 (Sensing upper bound). *At a point with tie blocks of sizes $m_1, \ldots, m_J$, sampling block j with probability proportional to C_{m_j} and applying Lemma 2 yields an unbiased estimator of q_x with*

$$\mathbb{E}\|\widehat{q}_x\|_2^2 \;\le\; V_{b,R}\Big(\sum_{j=1}^{J} C_{m_j}\Big)^2 \;\le\; 36\, V_{b,R}\Big(\sum_{j=1}^{J} \sqrt{m_j}\Big)^2 \;\le\; 36\, V_{b,R}\, Jn,$$

at $O_{b,R}(1)$ expected duels.

Each draw has constant expected query cost; the block structure controls its second moment. Averaging r independent draws divides the MSE by r and multiplies expected cost by r. For a given accuracy, the resulting cost bound is quadratic in n at distinct coordinates and linear under a complete tie.

4.2 Impossibility from the graph-Laplacian structure of duel information

To test whether this dependence is unavoidable, we allow arbitrary adaptive set comparisons. A duel has natural parameter $b(\theta_A - \theta_B)$ in the table of 2^n set values and Fisher information along $e_A - e_B$. Expected information therefore forms a weighted graph Laplacian, with total edge weight at most $b^2 Q_\theta/4$. The following lemma turns this structure into a lower bound for estimating a linear target.

Lemma 4 (Graph-flow Cramér–Rao bound for stopped duels). *Let the target be $L\theta \in \mathbb{R}^d$ with columns ℓ_A satisfying $\sum_A \ell_A = 0$, which is forced because pairwise differences cannot identify a common shift. Consider a parameter-independent adaptive protocol on an open parameter domain,*

stopping almost surely with $0 < Q_\theta = \mathbb{E}_\theta N < \infty$. Any regular, locally unbiased estimator with finite second moment satisfies $\mathbb{E}_\theta\|\widehat{\tau} - L\theta\|_2^2 \geq \frac{1}{b^2 Q_\theta}(\sum_A \|\ell_A\|_2)^2$. Such an estimator exists only if the target is identifiable from the expected-information graph.

Idea. The matrix information inequality bounds the MSE from below by a minimum flow energy. The coefficient demands and total edge weight give the stated bound on this energy. Proof in Appendix A.

For the Shapley target, node $U \subseteq B$ has coefficient demand $c_m(U)$, so $\sum_U \|c_m(U)\|_2 = \sum_k \binom{m}{k} \cdot \sqrt{a_{m,k}}/\binom{m}{k} = C_m$. The coefficient sum governing the lower bound is the same quantity used in the sampling law.

Theorem 5 (Sensing lower bound)**.** *There is a strictly submodular F_0 and a full-dimensional open submodular neighbourhood of it on which every regular estimator locally unbiased for the canonical subgradient q_x, using a stopping time N of arbitrary adaptive duels with $0 < Q_\theta = \mathbb{E}_\theta N < \infty$, satisfies*

$$\mathbb{E}\|\widehat{q}_x - q_x\|_2^2 \;\geq\; \frac{1}{2b^2 Q_\theta}\Big(\sum_{j=1}^{J} C_{m_j}\Big)^2 \;\geq\; \frac{2Jn}{b^2 Q_\theta} \quad \textit{(equal-sized blocks).}$$

Idea. Take $F_0(S) = \frac{R}{2} + \rho|S|(n-|S|)$, whose elementary submodularity inequalities all have slack, so the local model is full-dimensional in the set values. Interior demands belong to one block and the shared prefix endpoints have disjoint supports, so the total demand is at least $\frac{1}{\sqrt{2}}\sum_j C_{m_j}$. Proof in Appendix A.

Combining the bounds with $2\sqrt{m} \leq C_m \leq 6\sqrt{m}$ gives the sensing law.

Corollary 6 (Sensing complexity)**.** *On the open submodular model above, the optimal expected-query–MSE product for regular locally unbiased estimation of the canonical Lovász subgradient at a point with tie blocks $m_1, \ldots, m_J$ is $\Theta_{b,R}\big((\sum_j \sqrt{m_j})^2\big)$, hence $\Theta_{b,R}(Jn)$ for balanced blocks.*

The bounds match in dimension and tie structure. Their displayed constants differ by the saturation factor $64s(1+s)/(1-s)^3$, comparing the upper bound on cost times second moment with the MSE lower bound.

Corollary 7 (Exact greedy vectors cost n^2)**.** *On the preceding full-dimensional open submodular model, take every block to be a singleton. Any regular locally unbiased estimator of the exact greedy vector g^π from arbitrary adaptive duels of expected cost Q_θ has $\mathbb{E}_\theta\|\widehat{g} - g^\pi\|_2^2 \geq 2n^2/(b^2 Q_\theta)$.*

These lower bounds concern regular, locally unbiased canonical reconstruction. Proposition 16 in Appendix A also shows that uniform mixing minimizes the coefficient sum among fixed, parameter-independent mixtures of orderings within a block. Arbitrary subgradient selections and biased estimators remain outside the lower bound's scope.

The cost at distinct coordinates motivates choosing query points with ties. Section 5 creates these points by rounding the iterate and controls the resulting approximation error.

5 One curve of algorithms

The sensing law suggests choosing query points with fewer tie blocks (Figure 2). We maintain an update point y_t and query $x_t = Q_L(y_t)$, rounding each coordinate to the nearest multiple of $1/L$. This creates $J \leq \min\{L+1, n\}$ blocks. The estimator remains unbiased for q_{x_t}. The displacement between the two points adds at most R/L to the descent bound; the bound $\|g\|_1 \leq 2R$ makes this error independent of n (Remark 9).

At x_t, sample a block with probability proportional to C_{m_j} and a subset U from Lemma 2. Estimate the gap of $(H_j \cup U, H_j)$ using Lemma 14, reweight to obtain $\widehat{q}_t$, and take a projected step on y_t. Average the quantized iterates and apply threshold rounding for the output. The grid constrains the query points, while the update iterates remain in the continuous cube. Successive query points can occupy different grid locations, so their average can have more distinct coordinate values than any individual query point. Algorithm 1 in Appendix D gives the fixed round count and hard cap.

Theorem 8 (The algorithm curve)**.** *For every integer $L \geq 1$ and $T \geq 2c_{b,R}$, Algorithm 1 with $K = \lfloor T/(2c_{b,R}) \rfloor$, $\eta = 1/(6\sqrt{V_{b,R} J_L K})$, and $J_L = \min\{L+1, n\}$ uses at most T duels on every path and satisfies*

$$\mathbb{E}\big[F(\widehat{S}) - \min_S F\big] \;\leq\; C_{b,R}\, n \sqrt{\frac{J_L}{T}} + \frac{R}{L} + R e^{-c'_{b,R} T}.$$

Here $c'_{b,R} = (1-s)/16$. For $K = 0$ the algorithm returns a fixed set with excess at most R. Define $Q_\infty(y) = y$, $J_\infty = n$, and $R/\infty = 0$. Optimizing over L, including $L = \infty$, gives

$$\mathbb{E}\big[F(\widehat{S}) - \min_S F\big] \;=\; O_{b,R}\Big(R \wedge \min\big\{n^{2/3} T^{-1/3},\; n^{3/2} T^{-1/2}\big\}\Big).$$

Remark 9 (Why the bias is cheap). Quantization perturbs the query point by at most $1/(2L)$ in ℓ_∞. Because every Lovász subgradient satisfies $\|g\|_1 \leq 2R$, the function is $2R$-Lipschitz for the ℓ_∞ norm, so the price is R/L with no factor of n.

The three error terms account for direction estimation, quantization, and budget overflow. Coarser grids reduce the sensing term through J_L, while increasing the bias R/L.

Two sources of improvement. At $L = \infty$, quantization bias vanishes and $J_L = n$ gives $C_{b,R} n^{3/2}/\sqrt{T}$, the guarantee of adjacent-prefix projected SGD. The exactly unbiased primitive permits $K = \Theta(T)$ steps, improving the budget exponent of the $O(n^{7/5} T^{-2/5})$ sigmoid bound of Sho & Ito (2026) from $2/5$ to $1/2$ for fixed b, R.

Finite quantization improves the dimension dependence by replacing n with J_L in the sensing term. The interior choice $L^\star \asymp (R\sqrt{T}/(C_{b,R} n))^{2/3}$ gives $n^{2/3} T^{-1/3}$. The two branches exchange at $T \asymp n^5$; in $n^2 \ll T \ll n^5$, finite quantization has smaller order than both the unquantized and published bounds. Section 6 compares this guarantee with minimax lower bounds to identify the remaining optimization gap.

6 The limit and its boundary

6.1 Where the minimax rate sits

Canonical reconstruction is one route to minimization. An optimizer may use other directions, so its overall query complexity remains a separate question. We compare the curve of Section 5 with the minimax excess $\mathfrak{R}_T(n, b, R)$.

Theorem 10 (Lower bounds)**.** *For every $n \geq 1$, $b, R > 0$, and integer $T \geq 1$,*

$$\mathfrak{R}_T \;\geq\; \frac{1}{8} \min\Big\{R,\; \frac{n}{b\sqrt{T}}\Big\}.$$

For $T \geq 4/(b^2 R^2)$, also $\mathfrak{R}_T \geq 1/(4b\sqrt{T})$ by a two-point reduction on $F_\pm(S) = R/2 \pm \varepsilon \mathbf{1}\{1 \in S\}$ with $\varepsilon = 1/(b\sqrt{T})$. The first bound is Assouad (Tsybakov, 2009; Yu, 1997) on the modular family $F_\mu(S) = R/2 + \varepsilon\langle \mu, \mathbf{1}_S \rangle$ with $\varepsilon = \min\{R/(2n),\, 1/(b\sqrt{2T})\}$.

For fixed b, R, Theorem 8 gives an upper bound of order $n^{3/2} T^{-1/2}$ for $T \gtrsim n^5$, against a lower bound of order $nT^{-1/2}$, a factor $\sqrt{n}$ up to constants. In $n^2 \lesssim T \lesssim n^5$, the ratio of the displayed upper and lower bounds is of order $n^{-1/3} T^{1/6}$, reaching constant order near n^2. Appendix C gives matching orders for the low-budget regime and the modular subclass at every budget.

6.2 A uniform-sensitivity Assouad certificate

One route to resolving the gap is to strengthen the lower bound. We examine how large a certificate can be when it combines a uniform bound on comparison sensitivity with the divergence decomposition and Pinsker's inequality.

Let $\{F_\mu\}_{\mu \in \{\pm 1\}^n}$ be any family of submodular $F_\mu : 2^{[n]} \to [0, R]$ with minimizers $S^\star(\mu) = \{i : \mu_i = -1\}$ satisfying, for $a, \varepsilon > 0$, the **Hamming excess** condition

$$F_\mu(S) - F_\mu(S^\star(\mu)) \;\geq\; a\varepsilon\, d_H\big(S, S^\star(\mu)\big) \qquad \text{for all } \mu, S,$$

This lower-bounds the loss by a coordinatewise Hamming loss. Write $\nu^{(i)}$ for ν with its ith sign flipped. For a pair (A, B) write $D_\nu(A, B) = F_\nu(A) - F_\nu(B)$ and define the *flip sensitivity*

$$\mathrm{Sens}_{A,B}(\nu) = \sum_{i=1}^{n} \big[D_\nu(A, B) - D_{\nu^{(i)}}(A, B)\big]^2, \qquad \Lambda = \sup_{A,B,\nu} \mathrm{Sens}_{A,B}(\nu).$$

Theorem 11 (Uniform-sensitivity certificate)**.** *For $T > 0$ and an adaptive algorithm with stopping time N and $\sup_\mu \mathbb{E}_\mu N \le T$, the divergence decomposition (Lattimore & Szepesvári, 2020) and Pinsker/Jensen give the certificate*

$$\sup_\mu \mathbb{E}_\mu \big[F_\mu(\widehat{S}) - F_\mu(S^\star(\mu))\big] \;\ge\; \mathcal{C}_T := \frac{a\varepsilon n}{2}\left[1 - \frac{b}{4}\sqrt{\frac{T\Lambda}{n}}\right]_+ .$$

Here $[u]_+ = \max\{u, 0\}$. Every such family has $\Lambda \ge 4a^2\varepsilon^2 n$, and consequently $\mathcal{C}_T \le n/(4b\sqrt{T})$.

Idea. Hamming separation forces a lower bound on Λ. Substituting it into this specified certificate bounds the value the certificate can deliver. Proof in Appendix A.

The modular family attains the sensitivity floor. It has $a = 1$ and $\mathrm{Sens}_{A,B}(\nu) = 4\varepsilon^2|A \triangle B|$, so $\Lambda = 4\varepsilon^2 n$ meets the floor $4a^2\varepsilon^2 n$ with equality. This family achieves the certificate's order in the non-saturated budget regime.

Remark 12 (Information retained by the certificate). The ceiling applies to $\mathcal{C}_T$, which uses a single worst-case sensitivity Λ. Exact divergences, trajectory-dependent sensitivities, and other reductions retain information outside this certificate and provide directions for sharper lower bounds.

6.3 Biased, minimizer-directed oracles

A complementary route is to improve the optimizer. Descent can follow from a direction separating the iterate from every minimizer, without reconstructing the canonical subgradient.

Write $m = \min f$, $\mathcal{X}^\star = \arg\min f$ and $p(y) = \Pi_{\mathcal{X}^\star}(y)$. The vector $G_y = \frac{f(y)-m}{\|y-p(y)\|_2^2}(y - p(y))$, set to zero on $\mathcal{X}^\star$, satisfies $f(y) - m \le \langle G_y, y - x^\star\rangle$ at every global minimizer $x^\star$, with $\|G_y\|_2 \le 2R$ (Proposition 15, Appendix A). Its norm is $O(R)$, small enough for the $\widetilde{O}(n)$ second-moment target. Its definition uses the unknown optimal face; the task is to obtain suitable directions from few comparisons.

A chord oracle with a halving certificate. For a candidate set T, the sensing law gives an estimator adapted to the direction from $\mathbf{1}_S$ to $\mathbf{1}_T$. Write $H_S(T) := F(S \cap T) + F(S \cup T) - 2F(S)$.

Theorem 13 (Chord identity and contraction)**.** *For any sets S, T, the one-sided derivative of f at $\mathbf{1}_S$ toward $\mathbf{1}_T$ equals $H_S(T)$, and submodularity gives $H_S(T) \le F(T) - F(S)$, so $H_S(S^\star) \le -[F(S) - m]$. Averaging greedy vectors uniformly within the blocks of the ordered partition $(S \cap T,\ S\setminus T,\ T\setminus S,\ [n]\setminus(S\cup T))$, omitting empty blocks, produces $q^{S\to T} \in \partial f(\mathbf{1}_S)$ with $\langle q^{S\to T}, \mathbf{1}_T - \mathbf{1}_S\rangle = H_S(T)$ and, by Theorem 3 with $J \le 4$, an unbiased wide-duel estimator of second moment $O_{b,R}(n)$. Finally, if $H_S(T) \le -[F(S)-m]+\delta$ then returning $S\cap T$ or $S\cup T$ with equal probability gives $\mathbb{E}[F(S^+) - m] \le \frac{1}{2}[F(S) - m] + \frac{\delta}{2}$.*

The construction estimates a subgradient adapted to the chosen chord using at most four blocks, with second moment $O_{b,R}(n)$. A sufficiently negative chord yields the stated contraction. Selecting such a chord at low query cost remains open. Appendix B gives the moment and descent conditions for narrowing the dimension gap and the barriers faced by natural constructions.

7 Experiments

We evaluate sensing and optimization on decomposable and graph-cut objectives and public image graph cuts under simulated sigmoid feedback. We measure sensing by the product of expected comparison count and mean squared reconstruction error, and optimization by the normalized objective gap to the minimum. We compare quantization levels, including the unquantized variant $L = \infty$, and use the sigmoid comparison method of Sho & Ito (2026) as a baseline on the public image tasks. Full experimental protocols appear in Appendix D.

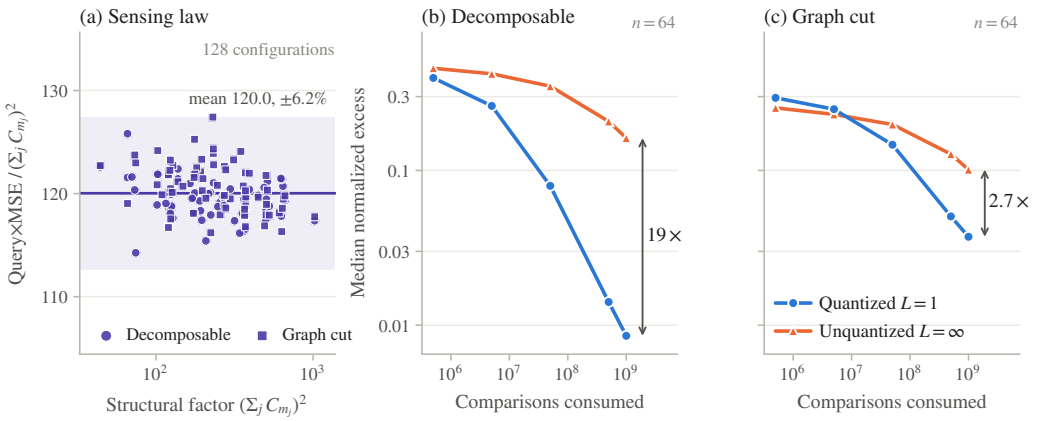

Figure 3: Sensing and quantization. (a) Query–MSE product divided by $(\sum_j C_{m_j})^2$ for 128 sensing configurations, with the sweep mean and its $\pm 6.2\%$ band. (b,c) Median conditional expected normalized excess at $n = 64$, using 32 matched replicas at the first three budgets and 128 at the final two, with 95% bootstrap intervals for the latter. Table 1 gives values.

The sensing law (Corollary 6). Across 128 sensing configurations spanning two objective families, we use exact targets to estimate query cost and MSE. The ratio of estimated $\mathbb{E}[N]\,\mathbb{E}\|\widehat{q}_x - q_x\|_2^2$ to $(\sum_j C_{m_j})^2$ stays within 6.2% of one fitted constant (Figure 3a). The partition factor thus captures most of the measured variation across the tested configurations, supporting the law across intermediate tie profiles.

Controlled end to end. At $n = 64$ and approximately 10^9 comparisons, $L = 1$ lowers median excess relative to $L = \infty$ by factors of 19 and 2.7 on decomposable and graph-cut objectives, respectively (Figure 3b,c). These gains are consistent with the tradeoff in Theorem 8, where ties reduce the sensing error term at the cost of approximation bias.

Application-derived graph cuts. Experiments on Middlebury image objectives obtained through SLIC contraction also show gains from quantization (Szeliski et al., 2008; Achanta et al., 2012). In Table 3, coarse quantization achieves lower error than the unquantized method, showing that finer query resolution need not improve optimization on these objectives. In the five selected image cases at a higher budget (Table 4), coarse quantization also achieves lower median error than the baseline.

At the nominal Flower configuration, finite levels form fewer tie blocks with comparable update counts (Appendix D.7). The trajectories therefore realize the tie structure that makes sensing cheaper. In the displayed Flower cases, coarse quantization achieves lower mean error than both the baseline and initialization, demonstrating optimization gains in these configurations.

8 Conclusion

We isolate the sensing problem in dueling submodular minimization by pricing the worst-case expected comparisons needed to reconstruct a canonical Lovász direction at a given accuracy. This cost is determined by the iterate's tie partition, growing quadratically with the dimension at distinct coordinates and linearly under a complete tie. A Shapley wide-pair estimator and a matching electrical-network lower bound establish this scaling for regular, locally unbiased estimation of the canonical direction under arbitrary adaptive set comparisons. Rounding the iterate manufactures ties at controlled bias and yields a one-parameter algorithm family that improves on the best published tradeoff between dimension and budget. One fitted constant describes measured query cost times MSE within 7% across 128 sensing configurations spanning two objective families; controlled experiments show up to $19\times$ lower median error than the unquantized method, alongside consistent quantization gains on the tested public image-labeling tasks.

AI USE STATEMENT

We used generative AI tools to assist with drafting, editing, and presentation during manuscript preparation. The authors reviewed and revised all AI-assisted material and take responsibility for the accuracy, claims, and final content of this work.

ETHICS STATEMENT

This work is theoretical, and its experiments use simulated and public data only. The public benchmarks comprise image-labeling energies and email and citation graph structures; we use only numeric node identifiers, graph structure, and benchmark-supplied labels. No human subjects were involved and no new data were collected.

REPRODUCIBILITY STATEMENT

Complete proofs and all assumptions appear in the paper and appendix. The experimental appendix describes the datasets, instance construction, implementation settings, randomization protocol, and evaluation procedure needed to reproduce the study. Reported statistics use the seed counts specified with each experiment.

REFERENCES

Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurélien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012. doi: 10.1109/TPAMI.2012.120.

Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of stochastic convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012. doi: 10.1109/TIT.2011.2182178.

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

Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. *Journal of Machine Learning Research*, 22(7): 1–108, 2021.

Jose H. Blanchet and Peter W. Glynn. Unbiased Monte Carlo for optimization and functions of expectations via multi-level randomization. In *Proceedings of the 2015 Winter Simulation Conference (WSC)*, pp. 3656–3667. IEEE Press, 2015.

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

Javier Castro, Daniel Gómez, and Juan Tejada. Polynomial calculation of the Shapley value based on sampling. *Computers & Operations Research*, 36(5):1726–1730, 2009.

Deeparnab Chakrabarty, Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. Subquadratic submodular function minimization. In *ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 1220–1231, 2017. doi: 10.1145/3055399.3055419.

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

Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981. doi: 10.1007/BF02579273.

Wassily Hoeffding. A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325, 1948. doi: 10.1214/aoms/1177730196.

Satoru Iwata, Lisa Fleischer, and Satoru Fujishige. A combinatorial strongly polynomial algorithm for minimizing submodular functions. *Journal of the ACM*, 48(4):761–777, 2001. doi: 10.1145/502090.502096.

Kevin G. Jamieson, Robert D. Nowak, and Benjamin Recht. Query complexity of derivative-free optimization. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 2672–2680, 2012.

Junpei Komiyama, Junya Honda, Hisashi Kashima, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in dueling bandit problem. In *Conference on Learning Theory (COLT)*, volume 40 of *Proceedings of Machine Learning Research*, pp. 1141–1154, 2015.

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

Yin Tat Lee, Aaron Sidford, and Sam Chiu-wai Wong. A faster cutting plane method and its implications for combinatorial and convex optimization. In *IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 1049–1065, 2015. doi: 10.1109/FOCS.2015.68.

László Lovász. Submodular functions and convexity. In Achim Bachem, Bernhard Korte, and Martin Grötschel (eds.), *Mathematical Programming The State of the Art*, pp. 235–257. Springer, 1983. doi: 10.1007/978-3-642-68874-4_10.

Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, 2017.

Anne-Marie Lyne, Mark Girolami, Yves Atchadé, Heiko Strathmann, and Daniel Simpson. On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statistical Science*, 30(4):443–467, 2015. doi: 10.1214/15-STS523.

Don McLeish. A general method for debiasing a Monte Carlo estimator. *Monte Carlo Methods and Applications*, 17(4):301–315, 2011. doi: 10.1515/mcma.2011.013.

Arkadi S. Nemirovsky and David B. Yudin. *Problem Complexity and Method Efficiency in Optimization*. John Wiley & Sons, New York, 1983.

Maxim Raginsky and Alexander Rakhlin. Information-based complexity, feedback and dynamics in convex programming. *IEEE Transactions on Information Theory*, 57(10):7036–7056, 2011. doi: 10.1109/TIT.2011.2154375.

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. PMLR, 2021. URL https://proceedings.mlr.press/v139/saha21b.html.

Alexander Schrijver. A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory, Series B*, 80(2):346–355, 2000.

Lloyd S. Shapley. A value for n-person games. In *Contributions to the Theory of Games II*, pp. 307–317. Princeton University Press, 1953.

Kaien Sho and Shinji Ito. Submodular function minimization with dueling oracle. In *International Conference on Learning Representations (ICLR)*, 2026. URL https://openreview.net/forum?id=BeMtzSH1d7.

Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall Tappen, and Carsten Rother. A comparative study of energy minimization methods for Markov random fields with smoothness-based priors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(6):1068–1080, 2008. doi: 10.1109/TPAMI.2007.70844. URL https://vision.middlebury.edu/MRF/.

Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York, 2009.

Matti Vihola. Unbiased estimators and multilevel Monte Carlo. *Operations Research*, 66(2):448–462, 2018. doi: 10.1287/opre.2017.1670.

Bin Yu. Assouad, Fano, and Le Cam. In David Pollard, Erik Torgersen, and Grace L. Yang (eds.), *Festschrift for Lucien Le Cam: Research Papers in Probability and Statistics*, pp. 423–435. Springer, New York, 1997.

Yisong Yue and Thorsten Joachims. Beat the mean bandit. In *International Conference on Machine Learning (ICML)*, pp. 241–248, 2011.

Yisong Yue, Josef Broder, Robert Kleinberg, and Thorsten Joachims. The K-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012. doi: 10.1016/j.jcss.2011.12.028.

A PROOFS

A.1 THE CANONICAL BLOCK SUBGRADIENT (PROPOSITION 1)

Every permutation listing B_1 first, then B_2, and so on is consistent with the nonincreasing order of x, so its greedy vector lies in $\partial f(x)$. Averaging over independent uniform permutations inside the blocks produces, in block B_j, exactly the Shapley vector of h_j; the concatenation is a convex combination of subgradients and hence a subgradient. □

A.2 WHOLE-BLOCK ESTIMATION (LEMMA 2)

The identity $\phi(h) = \sum_U h(U)c_m(U)$ follows by collecting, for each U, the Shapley weight with which $h(U)$ appears as a predecessor value and as a successor value. Unbiasedness of $\widehat{\phi}$ is immediate. Define $\psi_m(U) = c_m(U)/p_m(U)$. Direct computation gives $\|c_m(U)\|_2^2 = a_{m,k}/\binom{m}{k}^2$, so the design second moment is

$$\mathbb{E}\|\psi_m(U)\|_2^2 = \sum_{k=0}^{m} \binom{m}{k} \frac{a_{m,k}/\binom{m}{k}^2}{w_{m,k}/\binom{m}{k}} = \sum_{k=0}^{m} \frac{a_{m,k}}{w_{m,k}} = C_m^2,$$

and Cauchy–Schwarz shows $w_{m,k} \propto \sqrt{a_{m,k}}$ minimizes this among size distributions. For interior k, $a_{m,k} = m/(k(m-k)) \geq 4/m$, giving $C_m \geq 2(m-1)/\sqrt{m} + 2/\sqrt{m} = 2\sqrt{m}$. For the upper bound, set $M = \lfloor m/2 \rfloor$. Using symmetry, $a_{m,k} \leq 2/k$ for $1 \leq k \leq M$ and $\sum_{k=1}^{M} k^{-1/2} \leq 2\sqrt{M}$ gives $C_m \leq 6\sqrt{m}$. Finally $\sum_U c_m(U) = 0$, since for coordinate i the positive terms contribute $\sum_{k\geq 1} \binom{m-1}{k-1}/(k\binom{m}{k}) = 1$ and the negative terms contribute -1. □

A.3 THE SCALAR PRIMITIVE (LEMMA 14)

The scalar estimator uses bR to control random truncation and b to scale its output.

Lemma 14 (Exactly unbiased value differences at $O(1)$ expected cost)**.** *Let $Y \sim \mathrm{Bernoulli}(\sigma(b\Delta))$ with $|\Delta| \leq R$, put $s = \sigma(bR)$ and $\alpha = (1+s)/2$. Draw K_0, K_1 independently from $r_k = (1-\alpha)\alpha^{k-1}$, let I_0 indicate that K_0 fresh responses are all zero and I_1 that K_1 fresh responses are all one, and set $Z = \big(-I_0/(K_0 r_{K_0}) + I_1/(K_1 r_{K_1})\big)/b$. Then*

$$\mathbb{E}Z = \Delta, \qquad \mathbb{E}Z^2 \leq V_{b,R} := \frac{8s(1+s)}{b^2(1-s)^2}, \qquad \mathbb{E}[K_0 + K_1] = \frac{4}{1-s}.$$

With $p = \sigma(b\Delta)$, $\alpha = (1+s)/2 > s$ and $r_k = (1-\alpha)\alpha^{k-1}$, $\mathbb{E}[-I_0/(K_0 r_{K_0})] = -\sum_k (1-p)^k/k = \log p$ and symmetrically $\mathbb{E}[-I_1/(K_1 r_{K_1})] = \log(1-p)$, so $\mathbb{E}Z = b^{-1}\log\frac{p}{1-p} = \Delta$. With $\rho = s/\alpha$, each unsigned summand $L_p = I/(Kr_K)$ satisfies $\mathbb{E}L_p^2 \leq \sum_k s^k/(k^2 r_k) = \frac{\alpha}{1-\alpha}\sum_k \rho^k/k^2 \leq \frac{\alpha}{1-\alpha}\frac{\rho}{1-\rho} = \frac{2s(1+s)}{(1-s)^2}$, and $(a-b)^2 \leq 2a^2 + 2b^2$ gives the stated $V_{b,R}$. The expected draw count is $2/(1-\alpha) = 4/(1-s)$. □

A.4 SENSING UPPER BOUND (THEOREM 3)

Choosing block j with probability $p_j = C_{m_j}/\sum_\ell C_{m_\ell}$ and dividing the block estimator by p_j is unbiased for the concatenation, with second moment at most $V_{b,R}\sum_j C_{m_j}^2/p_j = V_{b,R}(\sum_j C_{m_j})^2$. Applying $C_m \le 6\sqrt{m}$ and $(\sum_j \sqrt{m_j})^2 \le J\sum_j m_j = Jn$ gives the displayed forms. The expected duel count is $c_{b,R} = 4/(1-s)$ by Lemma 14. Independent repetition and averaging divide the MSE by the number of repetitions and multiply expected cost by the same factor. □

A.5 GRAPH-FLOW CRAMÉR–RAO (LEMMA 4)

Stopped experiment and regularity. The original set-value table θ ranges over an open domain; the linear target $L\theta$ annihilates common shifts. Include the protocol's auxiliary randomization in the transcript W. Its query and stopping rules are fixed functions of the history, independent of the unknown parameter. Encode responses as $Y_r \in \{0,1\}$ and put $a_r = e_{A_r} - e_{B_r}$, $p_r = \sigma(ba_r^\top\theta)$. The event $\{N \ge r\}$ and the query a_r are determined before observing Y_r. Almost-sure termination is assumed throughout the domain. The stopped score is

$$S_\theta = b\sum_{r=1}^{N} a_r(Y_r - p_r).$$

The mean $m(\vartheta) = \mathbb{E}_\vartheta\widehat{\tau}$ is locally unbiased at θ when $m(\theta) = L\theta$ and $Dm(\theta) = L$. We call it regular if, for every parameter direction v, the stopped likelihood ratio $\mathcal{L}_h = dP_{\theta+hv}/dP_\theta$ satisfies

$$\lim_{h\to 0}\mathbb{E}_\theta\left[(1+\|\widehat{\tau}\|_2)\left|\frac{\mathcal{L}_h - 1}{h} - v^\top S_\theta\right|\right] = 0. \tag{1}$$

This condition justifies $Dm(\theta) = \mathbb{E}_\theta[\widehat{\tau}S_\theta^\top]$. A sufficient condition is $\mathbb{E}_\theta[(1+\|\widehat{\tau}\|_2)(1+N)e^{\epsilon N}] < \infty$ for some $\epsilon > 0$. Indeed, each log-response likelihood has directional derivative bounded by $2b\|v\|_\infty$, so $|\partial_h \log\mathcal{L}_h| \le 2b\|v\|_\infty N$ and $\mathcal{L}_h \le e^{2b|h|\|v\|_\infty N}$; dominated convergence proves equation 1.

Expected information and flow energy. The scores stopped at $N \wedge M$ form an L^2-Cauchy sequence because $\|a_r\|_2^2 \le 2$ and $Q_\theta < \infty$. Martingale orthogonality therefore yields $\mathbb{E}_\theta S_\theta = 0$ and

$$I_\theta = \mathbb{E}_\theta S_\theta S_\theta^\top = \mathbb{E}_\theta\sum_{r=1}^{N} b^2 p_r(1-p_r)a_r a_r^\top = \sum_e \omega_e a_e a_e^\top, \qquad \sum_e \omega_e \le b^2 Q_\theta/4.$$

The weights combine orientations and repeated visits to each unordered edge; self-comparisons have zero information. Regular local unbiasedness gives $\mathbb{E}_\theta[(\widehat{\tau} - L\theta)S_\theta^\top] = L$. Consequently $\ker I_\theta \subseteq \ker L$ is necessary for a finite-second-moment estimator. When this holds, expanding the covariance of $\widehat{\tau} - L\theta - LI_\theta^\dagger S_\theta$ gives $\mathrm{Cov}_\theta(\widehat{\tau}) \succeq LI_\theta^\dagger L^\top$. The trace is the minimum vector-flow energy $\min_{\mathrm{div}\,J=\ell}\sum_{e:\omega_e>0}\|J_e\|_2^2/\omega_e$. For every feasible flow, $\sum_A \|\ell_A\|_2 \le 2\sum_e \|J_e\|_2$. Weighted Cauchy–Schwarz gives

$$\sum_e \frac{\|J_e\|_2^2}{\omega_e} \ge \frac{(\sum_A \|\ell_A\|_2)^2}{4\sum_e \omega_e} \ge \frac{(\sum_A \|\ell_A\|_2)^2}{b^2 Q_\theta}.$$

Regularity of the construction. Keep the known tuning $s = \sigma(bR)$ fixed under parameter perturbations. Choose $\epsilon > 0$ with $\alpha e^\epsilon < 1$. For an all-one term,

$$\mathbb{E}\left[\frac{I_1}{K_1 r_{K_1}}(1+K_1)e^{\epsilon K_1}\right] = \sum_{k\ge 1}\frac{1+k}{k}(pe^\epsilon)^k < \infty.$$

The all-zero term replaces p by $1-p$; both are at most $s < \alpha$. The other independent geometric length has a finite exponential moment, establishing the sufficient regularity condition for Z. Finite block/subset importance weights preserve it for $\widehat{q}_x$. The estimator is unbiased throughout the parameter neighbourhood, and hence locally unbiased at every point there. Thus the upper and lower bounds use the same stopped statistical model. □

A.6 SENSING LOWER BOUND (THEOREM 5)

Take $F_0(S) = \frac{R}{2} + \rho|S|(n-|S|)$ with $\rho = R/(16n^2)$. Since $|S|(n-|S|)$ is a concave function of the cardinality it is submodular, and every elementary submodularity inequality has slack exactly 2ρ; hence all perturbations $F_\theta = F_0 + \theta$ with $\|\theta\|_\infty < \rho/4$ remain submodular and strictly inside $[0, R]$, so the model is full-dimensional in the set values. For the target q_x, the coefficient demand at node $H_j \cup U$ is the embedding of $c_{m_j}(U)$ into the coordinates of B_j, and $\sum_{U \subseteq B_j} \|c_{m_j}(U)\|_2 = C_{m_j}$ as computed above. Interior nodes are unique to a block; the only shared nodes are consecutive prefix endpoints, where the two demands have disjoint coordinate supports so that $\|u + v\|_2 = \sqrt{\|u\|_2^2 + \|v\|_2^2} \geq (\|u\|_2 + \|v\|_2)/\sqrt{2}$. The total demand is therefore at least $\frac{1}{\sqrt{2}} \sum_j C_{m_j}$, and Lemma 4 applies. Corollary 7 is the case $m_j \equiv 1$, where $C_1 = 2$ and the total demand is $2 + (n-1)\sqrt{2} \geq \sqrt{2}n$. □

A.7 THE ALGORITHM CURVE (THEOREM 8)

Since every $g \in B(F)$ has $\|g\|_1 \leq 2R$, f is $2R$-Lipschitz in ℓ_∞ and $\|Q_L(y) - y\|_\infty \leq 1/(2L)$, so for $q_t \in \partial f(x_t)$,

$$f(x_t) - f(x^\star) \leq \langle q_t, x_t - x^\star \rangle \leq \langle q_t, y_t - x^\star \rangle + \tfrac{R}{L}.$$

The deterministic projected-gradient inequality gives $\sum_t \langle \widehat{q}_t, y_t - x^\star \rangle \leq \frac{n}{2\eta} + \frac{\eta}{2} \sum_t \|\widehat{q}_t\|_2^2$; substituting Theorem 3 with $J \leq J_L$ and optimizing η yields $\frac{1}{K} \sum_t \mathbb{E}[f(x_t) - f(x^\star)] \leq 6n\sqrt{V_{b,R} J_L / K} + R/L$. Convexity gives $f(\bar{x}) \leq K^{-1} \sum_t f(x_t)$. Threshold rounding preserves the Lovász value in expectation, so the same bound holds for the returned set of this unrestricted K-round process.

For the hard budget, Algorithm 1 draws both geometric lengths before sending the round's queries and returns $\varnothing$ if their cost would exceed the remaining budget. Couple it to the unrestricted process with the same randomness, continuing the latter after any abort. Their outputs agree whenever the unrestricted total cost is at most T, and every excess is between zero and R. Thus the capped risk is at most the unrestricted risk plus R times the overflow probability; the unbiased-gradient analysis above is not conditioned on completion. With $d = 1 - \alpha = (1-s)/2$ and $c_{b,R} = 2/d$, this probability is

$$\begin{aligned} \Pr\left(\sum_{t=1}^{K} (K_{0,t} + K_{1,t}) > T \right) &= \Pr(\mathrm{Bin}(T, d) < 2K) \\ &\leq \Pr(\mathrm{Bin}(T, d) \leq Td/2) \leq e^{-Td/8} = e^{-(1-s)T/16}. \end{aligned}$$

For $K \geq 1$, $K \geq T/(4c_{b,R})$ gives the displayed leading constant with $C_{b,R} = 12\sqrt{c_{b,R} V_{b,R}}$. When $K = 0$, returning the fixed set incurs excess at most R. Minimizing $C_{b,R} n \sqrt{L/T} + R/L$ over real L gives $L^\star = (2R\sqrt{T}/(C_{b,R} n))^{2/3}$ and the rate $n^{2/3} T^{-1/3}$ whenever $1 \lesssim L^\star \lesssim n$; when $L^\star \gtrsim n$ the optimum is at $L = \infty$, where $J_L = n$, the bias vanishes and the bound is $C_{b,R} n^{3/2}/\sqrt{T}$. □

A.8 MINIMAX LOWER BOUNDS (THEOREM 10)

For the two-point bound, take $F_\pm(S) = R/2 \pm \varepsilon \mathbf{1}\{1 \in S\}$ with $\varepsilon = 1/(b\sqrt{T}) \leq R/2$. Every comparison's two natural parameters differ by at most $2b\varepsilon$. Since the logistic log-partition has second derivative at most $1/4$, transcript KL is at most $b^2 T \varepsilon^2 / 2 = 1/2$. Pinsker gives total variation at most $1/2$; the average probability of returning the wrong membership of element 1 is at least $1/4$, yielding excess at least $\varepsilon/4$.

For the modular family in the theorem, the excess is ε times Hamming error and $\Lambda = 4\varepsilon^2 n$. The certificate proved below gives risk at least $\varepsilon n(1 - b\varepsilon\sqrt{T}/2)/2 \geq \varepsilon n/4$ for $\varepsilon = \min\{R/(2n), 1/(b\sqrt{2T})\}$. This is at least $\frac{1}{8} \min\{R, n/(b\sqrt{T})\}$. Both families lie in $[0, R]$ and are submodular. □

A.9 UNIFORM-SENSITIVITY CERTIFICATE (THEOREM 11)

Step 1, sensitivity floor. Write $\delta = a\varepsilon$, fix μ, and put $S_0 = S^\star(\mu)$, $\phi(\nu) = D_\nu(\overline{S}_0, S_0)$. Hamming excess gives $\phi(\mu) \geq \delta n$ and $\phi(-\mu) \leq -\delta n$. For $p \in [0,1]^n$, let $\widetilde{\phi}(p) = \mathbb{E}_p \phi(X)$, where the signs

X_i are independent with $\Pr_p(X_i = 1) = p_i$. Then $\partial_i \widetilde{\phi}(p) = \mathbb{E}_p[X_i(\phi(X) - \phi(X^{(i)}))]$. Vector Jensen gives $\|\nabla \widetilde{\phi}(p)\|_2^2 \le \Lambda$. Integrating along the segment between the two opposite cube vertices corresponding to $\mu, -\mu$ yields $2\delta n \le \sqrt{n\Lambda}$, hence $\Lambda \ge 4\delta^2 n$. The range assumption also gives $\delta n \le R$.

Step 2, a uniform information bound. The sigmoid family has log-partition function $A(u) = \log(1 + e^u)$ with $A'' \le 1/4$, hence $\mathrm{kl}(\sigma(u)\|\sigma(u')) \le (u - u')^2/8$. Let P_μ be the stopped transcript law. The algorithm's query and stopping kernels are parameter-independent conditional on the history. The divergence decomposition, first stopped at $N \wedge M$ and then passed to the limit, gives

$$\mathrm{KL}(P_\mu \| P_{\mu^{(i)}}) \le \frac{b^2}{8} \mathbb{E}_\mu \sum_{r=1}^{N} [D_\mu(A_r, B_r) - D_{\mu^{(i)}}(A_r, B_r)]^2.$$

Summing over i and using $\mathbb{E}_\mu N \le T$ gives $\sum_i \mathrm{KL}(P_\mu \| P_{\mu^{(i)}}) \le b^2 T \Lambda / 8$.

Step 3, evaluate the certificate. Average over the uniform prior on μ. Coordinatewise testing and the Hamming excess condition give

$$\sup_\mu \mathbb{E}_\mu \mathrm{excess} \ge \frac{\delta}{2} \sum_{i=1}^{n} \left(1 - 2^{-n} \sum_\mu \mathrm{TV}(P_\mu, P_{\mu^{(i)}}) \right).$$

Pinsker and Jensen over both μ and i bound the average total variation by $(b/4)\sqrt{T\Lambda/n}$. Combining with the nonnegativity of excess yields $\mathcal{C}_T$. The certificate satisfies

$$\mathcal{C}_T \le \frac{\delta n}{2} [1 - b\delta\sqrt{T}/2]_+ \le \frac{n}{4b\sqrt{T}},$$

where the last inequality maximizes the quadratic over $\delta \ge 0$. For the modular family, $\Lambda = 4\delta^2 n$, so the sensitivity floor is attained. □

A.10 Geometric existence and the chord identity (Proposition 15 and Theorem 13)

Proposition 15 (Geometric existence). *Let $m = \min f$, $\mathcal{X}^\star = \arg\min f$, $p(y) = \Pi_{\mathcal{X}^\star}(y)$, $\delta(y) = f(y) - m$ and $r(y) = \|y - p(y)\|_2$. Set $G_y = 0$ when $r(y) = 0$ and $G_y = \frac{\delta(y)}{r(y)^2}(y - p(y))$ otherwise. Then G_y satisfies $\delta(y) \le \langle G_y, y - x^\star \rangle$ for every global minimizer $x^\star$, and $\|G_y\|_2 \le 2R$.*

For Proposition 15, the projection theorem gives $\langle y - p, x^\star - p \rangle \le 0$ and hence $\langle y - p, y - x^\star \rangle \ge \|y - p\|_2^2$; multiplying by δ/r^2 gives the separation inequality, and $\delta = f(y) - f(p) \le 2Rr$ gives $\|G_y\|_2 = \delta/r \le 2R$.

For Theorem 13, partition into $B_1 = S \cap T$, $B_2 = S \setminus T$, $B_3 = T \setminus S$, $B_4 = [n] \setminus (S \cup T)$, omitting empty blocks from the order and all sampling sums. The greedy order B_1, B_2, B_3, B_4 is consistent with $\mathbf{1}_S$ and sorts the directional coordinates decreasingly inside each tie block, so it attains the maximum of $\langle q, \mathbf{1}_T - \mathbf{1}_S \rangle$ over $\partial f(\mathbf{1}_S)$. The sums over B_1, B_2, B_3 are $F(S \cap T)$, $F(S) - F(S \cap T)$ and $F(S \cup T) - F(S)$, and since the direction is -1 on B_2 and $+1$ on B_3 the inner product telescopes to $F(S \cap T) + F(S \cup T) - 2F(S)$. Submodularity gives $H_S(T) \le F(T) - F(S)$. Averaging inside the $J \le 4$ nonempty blocks uses Proposition 1, and its estimator cost is Theorem 3 with $\sum_j C_{m_j} \le 6 \sum_j \sqrt{m_j} \le 12\sqrt{n}$. Finally $\frac{1}{2}[F(S \cap T) + F(S \cup T)] = F(S) + \frac{1}{2} H_S(T)$ gives the contraction. □

A.11 Other selections from the subdifferential

Theorem 5 constrains estimation of the *canonical* element of $\partial f(x)$, whereas an algorithm needs only some element of that face. The following comparison concerns a single block and fixed linear selections obtained by mixing its greedy orderings.

Proposition 16 (Minimum coefficient demand among fixed mixtures). *Let B have size $m \ge 1$, and let $v = (v(U))_{U \subseteq B}$ be an arbitrary, unnormalized set-value table. For a permutation π of B, put $S_k^\pi = \{\pi_1, \ldots, \pi_k\}$ and define the greedy vector by $[g^\pi(v)]_{\pi_k} = v(S_k^\pi) - v(S_{k-1}^\pi)$. Let M^π be its*

coefficient matrix, so that $g^\pi(v) = M^\pi v$. For any probability distribution λ on these permutations that is fixed independently of v, set $M^\lambda = \sum_\pi \lambda_\pi M^\pi$. Then

$$D(M^\lambda) := \sum_{U \subseteq B} \|M^\lambda_{\cdot,U}\|_2 \geq C_m.$$

Equality is attained by the uniform distribution, whose coefficient matrix has columns $c_m(U)$. The sum includes the empty-set column; no normalization $v(\varnothing) = 0$ is imposed.

Proof. For a permutation σ of the elements of B, define simultaneous row and subset relabeling by $[\sigma \cdot M]_{i,U} = M_{\sigma^{-1}(i),\sigma^{-1}(U)}$. The functional D is convex and invariant under this action, and $\sigma \cdot M^\pi = M^{\sigma\pi}$. As σ ranges over all $m!$ element permutations, $\sigma\pi$ ranges over all greedy orderings for each fixed π. Consequently

$$\overline{M} := \frac{1}{m!}\sum_\sigma \sigma \cdot M^\lambda = \frac{1}{m!}\sum_\pi M^\pi, \qquad D(\overline{M}) \leq \frac{1}{m!}\sum_\sigma D(\sigma \cdot M^\lambda) = D(M^\lambda).$$

Collecting coefficients in the uniform average gives $\overline{M}_{\cdot,U} = c_m(U)$, as in Lemma 2; in particular the empty-set column is $-\mathbf{1}_B/m$ and the full-set column is $\mathbf{1}_B/m$. Hence

$$D(\overline{M}) = \sum_{k=0}^{m} \binom{m}{k} \frac{\sqrt{a_{m,k}}}{\binom{m}{k}} = C_m,$$

Thus uniform mixing attains the minimum coefficient demand C_m. □

B BARRIERS TO ALTERNATIVE ESTIMATION STRATEGIES

This appendix states the oracle conditions for the further direction in Section 6 and identifies structural barriers faced by several natural approaches. The examples below characterize the source of each obstruction and the properties required of an alternative construction.

B.1 ORACLE REQUIREMENTS

Chord selection and accuracy. Finding a sufficiently negative chord is itself two smaller submodular minimizations, a restriction and a contraction problem. The contraction preserves the accuracy of the input minimizer. If T comes from a black-box minimizer of accuracy δ, the recursion $a \mapsto \frac{a}{2} + \frac{\delta}{2}$ has fixed point exactly δ, so it cannot bootstrap a weak minimizer into a strong one. It helps only if near-optimal chords come from a mechanism cheaper than the problem itself.

Minimizer-directed oracle conditions. The open question is whether one can construct a random G_y from $\widetilde{O}(1)$ expected duels, with $\mathbb{E}\|G_y\|_2^2 = \widetilde{O}_{b,R}(n)$ and $f(y) - f(x^\star) \leq \mathbb{E}\langle G_y, y - x^\star\rangle + \widetilde{O}_{b,R}(n/\sqrt{T})$ at every global minimizer $x^\star$ *only*. Demanding a zero-slack inequality against every point would make $\mathbb{E}G_y$ an exact subgradient. At a strictly ordered interior iterate this identifies the greedy target, whose regular locally unbiased estimation is covered by Corollary 7. An amortized form bounding the accumulated slack by $\widetilde{O}_{b,R}(n\sqrt{T})$ suffices equally. The following constructions illustrate the structural barriers to meeting these conditions.

B.2 THRESHOLD ROUNDING AND COARSENING

The correlation term is a symmetric Bregman energy. For any random X with $\mathbb{E}X = y$, $\mathbb{E}f(X) = f(y)$ and any selection $Q_X \in \partial f(X)$, the slack in the threshold approach is $\Gamma(y) = \mathbb{E}\langle Q_X, X - y\rangle = \frac{1}{2}\mathbb{E}\langle Q_X - Q_{X'}, X - X'\rangle \geq 0$ for an independent copy, and equivalently $\Gamma(y) = f(y) - \langle \mathbb{E}Q_X, y\rangle$, the support-function suboptimality of the accessible averaged field. For threshold vertices $X = \mathbf{1}_{\{i:y_i \geq \Theta\}}$, $\Theta \sim \mathrm{Unif}[0,1]$, and the canonical Shapley selection $Q_X = q_X$, nonnegativity also holds *coordinatewise*. For $0 < y_i < 1$, write $\mu_i^+ = \mathbb{E}[q_{X,i} \mid X_i = 1]$ and $\mu_i^- = \mathbb{E}[q_{X,i} \mid X_i = 0]$. Then $\gamma_i(y) = \mathbb{E}[q_{X,i}(X_i - y_i)] = y_i(1-y_i)(\mu_i^+ - \mu_i^-) \geq 0$; at $y_i \in \{0,1\}$ this contribution is zero. To see the sign, let P_i be the predecessors of i in a uniformly random

permutation. At a vertex S containing i, its canonical marginal uses predecessors $P_i \cap S$; at a vertex T excluding i, it uses $T \cup P_i$. The inclusion $P_i \cap S \subseteq T \cup P_i$ and diminishing returns give the same ordering of the expected marginals, hence $\mu_i^+ \geq \mu_i^-$. Thus the canonical threshold field has no coordinatewise cancellation between opposite signs.

Exact pointwise cancellation recovers the greedy vector. If $\Gamma(y) = 0$ at a strictly ordered y then $Q_X = g^\pi$ almost surely. Vanishing of the nonnegative integrand forces both subgradient slacks to vanish, so every selected vector is a common subgradient at every chain vertex, whence $q(P_k) = F(P_k)$ for all k and differencing gives g^π. Consequently, any correction that stays in $B(F)$ and cancels the slack exactly must rebuild the vector that Corollary 7 makes expensive.

The averaged field can vanish away from the optimum. On $\{1,2\}$ with $F(\varnothing) = F(\{1,2\}) = 0$, $F(\{1\}) = R/3$, $F(\{2\}) = R$ (submodular, range $[0,R]$), at $y = (\frac{3}{4}, \frac{1}{4})$ the averaged threshold field vanishes while $f(y) - \min f = \Gamma(y) = R/6$. The averaged field alone therefore does not distinguish this interior point from a minimizer.

Limits of unbiased coarsening. For an integer $L \geq 1$, draw $U \sim \mathrm{Unif}[0,1)$, shared by all coordinates, and set $X_L = Q_{L,U}(y)$ with $[Q_{L,U}(y)]_i = \frac{1}{L}\lfloor Ly_i + U \rfloor$. Write $\bar{q}_L(y) = \mathbb{E}[q_{X_L}]$ for the averaged canonical direction and $\Gamma_L(y) = \mathbb{E}\langle q_{X_L}, X_L - y\rangle$ for its residual. This quantization preserves both the iterate and the Lovász value in expectation and has residual $\Gamma_L(y) \leq 2R/L$. On the symmetric cut $F(\{1\}) = F(\{2\}) = R$, $F(\varnothing) = F(\{1,2\}) = 0$ with $y = (d,0)$ one gets $\Gamma_L(y) = Rd(1-Ld)_+$, whose supremum is exactly $R/(4L)$.

For a positive integer M, draw independent $\Theta_r \sim \mathrm{Unif}[0,1]$ and average their threshold vertices as $X^{(M)} = M^{-1}\sum_{r=1}^M \mathbf{1}_{\{i:y_i \geq \Theta_r\}}$. Using the canonical direction $q_{X^{(M)}}$ at this averaged point gives residual $\mathbb{E}\langle q_{X^{(M)}}, X^{(M)} - y\rangle = Rd(1-d)^M$ on the same symmetric cut. Its maximum over $d \in [0,1]$ is $\frac{R}{M+1}(\frac{M}{M+1})^M \geq \frac{R}{e(M+1)}$. For random M independent of the thresholds, Jensen gives the same lower bound with M replaced by $\mathbb{E}M$. These lower bounds quantify the residual left by unbiased coarsening.

Unsigned couplings preserve nonnegativity. For the canonical Shapley threshold field and any joint law of (Θ, Θ') with uniform marginals, $\mathbb{E}\langle Q(\Theta), X(\Theta) - X(\Theta')\rangle \geq 0$, by a bathtub rearrangement using the membership monotonicity above; and no nonnegative mixture of pairwise cross-threshold residuals can be negative. Identity coupling drives the residual to zero but pairs each subgradient with its own vertex, giving no update at the fractional point.

Two-scale affine extrapolation violates the optimality conditions. On $F = x_3 - x_1x_2x_3$ (submodular, with $\varnothing, \{1\}, \{2\}, \{1,2\}, \{1,2,3\}$ all minimizers) at $y = (\frac{4}{5}, \frac{1}{2}, \frac{1}{5})$, itself optimal, the dithered fields are $\bar{q}_1 = (-\frac{2}{15}, -\frac{17}{60}, \frac{5}{12})$ and $\bar{q}_2 = (0, -\frac{1}{5}, \frac{1}{5})$. For $G_\lambda = \bar{q}_2 + \lambda(\bar{q}_2 - \bar{q}_1)$ one has $\langle G_\lambda, y\rangle = -\frac{3}{50} + \frac{21}{200}\lambda$ and $\langle G_\lambda, y - e_1\rangle = -\frac{3}{50} - \frac{17}{600}\lambda$, which would require $\lambda \geq \frac{4}{7}$ and $\lambda \leq -\frac{36}{17}$ at once. For every λ one of the two is at most $-\frac{3}{50}$. The extrapolated vector leaves the normal cone of the minimizer *lattice*; controlling the residual in the single direction y therefore does not preserve optimality.

B.3 Sign updates and entropic geometry

Sign-based follow-the-leader cycles. Submodular minimization is exactly $\min_S F(S) = \frac{1}{2}(F([n]) - \min_{q \in B(F)} \|q\|_1)$, every sign linearization has a two-block best response, and external regret against $[-1,1]^n$ certifies the primal gap. Since the response is chosen after the action, exact implicit be-the-leader is equivalent to $\arg\min_{v \in B(F)} \|Q + v\|_1$, which at $Q = 0$ is the original problem; and pure follow-the-leader cycles. To specify the iteration, let Q_t be the cumulative response vector and choose $S_t = \{i : Q_{t,i} < 0\}$. After receiving q^{S_t}, update $Q_{t+1} = Q_t + q^{S_t}$. On $F = x_1 + x_2 + x_3 - x_1x_2 - 2x_1x_2x_3$ with responses $q^{\{1\}} = (1,-1,0)$, $q^{\{2\}} = (-1,1,0)$ and $Q_0 = (-\delta, 2\delta, 3\delta)$, $0 < \delta < \frac{1}{2}$, the iteration alternates $\{1\}, \{2\}, \ldots$ forever. The first update gives $Q_1 = (1-\delta, 2\delta - 1, 3\delta)$ and the second returns to Q_0. The ergodic average is $(\frac{1}{2}, \frac{1}{2}, 0)$, with constant excess $\frac{1}{2}$. This example shows that pure follow-the-leader can cycle away from ties.

Entropic geometry. Entropic mirror descent on the cube does not give a logarithmic diameter. The binary-entropy regularizer has range $n \log 2$ and is only $4/n$-strongly convex for ℓ_1, and in the general online linear-loss model, ℓ_1-bounded gradients on the cube force minimax regret $\Omega(\min\{K, \sqrt{nK}\})$ by Khintchine.

B.4 Other estimation constructions

A suboptimal local minimum of the multilinear extension. On $\{1, 2, 3\}$, take the K_3 cut with a positive unary term, $F(S) = \frac{R}{4}[|S|(3 - |S|) + |S|]$. This is submodular with range diameter R. Its multilinear extension is $M(x) = \frac{R}{4}(3\sum_i x_i - 2\sum_{i<j} x_i x_j)$. For $z \geq 0$ and $0 < t := \sum_i z_i < 1$,

$$M(\mathbf{1} - z) - M(\mathbf{1}) = \frac{R}{4}\Big(t - 2\sum_{i<j} z_i z_j\Big) \geq \frac{R}{4} t(1 - t) > 0.$$

Thus $\mathbf{1}$ is a strict local minimum relative to the cube, yet $M(\mathbf{1}) = 3R/4 > M(0) = 0$.

A sign-grouping rule targets Banzhaf marginals. Draw independent uniform signs $\varepsilon_i \in \{\pm 1\}$, set $A = \{i : \varepsilon_i = 1\}$, and use Lemma 14 to obtain Z conditionally unbiased for $F(A) - F([n] \setminus A)$. Conditioning on ε_i and pairing complementary sets gives

$$\mathbb{E}[\varepsilon_i Z] = 2^{-(n-1)} \sum_{U \subseteq [n] \setminus \{i\}} [F(U \cup \{i\}) - F(U)] =: \beta_i(F),$$

the Banzhaf marginal. For the submodular function $F(S) = -R\mathbf{1}\{S = \{1, 2, 3\}\}$ on three elements, each $\beta_i = -R/4$, so $\sum_i \beta_i = -3R/4 \neq F(\{1, 2, 3\}) = -R$. The mean of this particular rule therefore lies outside $B(F)$, the subdifferential at a complete tie.

A two-point smoothing benchmark. For $0 < h < 1/2$, let $D_h = [h, 1 - h]^n$ and $f_h(y) = \mathbb{E}_{v \sim \mathrm{Unif}(\mathbb{B}_2^n)} f(y + hv)$, where $\mathbb{B}_2^n$ is the unit ball. Draw u uniformly on its sphere and independently threshold-round $y \pm hu$ to $A_\pm$. Lemma 14 estimates their gap with $\mathbb{E}[Z \mid y, u] = f(y + hu) - f(y - hu)$, $\mathbb{E}[Z^2 \mid y, u] \leq V_{b,R}$ and expected cost $c_{b,R}$. Sphere symmetry and the divergence theorem give, for $G = nuZ/(2h)$,

$$\mathbb{E}[G \mid y] = \nabla f_h(y), \qquad \mathbb{E}[\|G\|_2^2 \mid y] \leq \frac{n^2 V_{b,R}}{4h^2}.$$

Run $K \geq 1$ projected SGD steps on D_h with $\eta = 2h/\sqrt{nV_{b,R}K}$ and threshold-round their average to $\widehat{S}$. Convexity gives $f \leq f_h$, while shrinking the cube and smoothing each cost at most $2Rh$ by the $2R$ Lipschitz bound. The projected-SGD inequality therefore gives

$$\mathbb{E}[F(\widehat{S}) - \min F] \leq 4Rh + \frac{n^{3/2}\sqrt{V_{b,R}}}{2h\sqrt{K}}.$$

With $h = K^{-1/4}/4$, this is $O_{n,b,R}(K^{-1/4})$. For a hard budget, fix $K = \lfloor T/(2c_{b,R}) \rfloor$ and use the pre-drawn-length cap and empty-set fallback from Theorem 8; the same coupling adds at most $Re^{-(1-s)T/16}$. When $K = 0$, return $\varnothing$. The resulting $O_{n,b,R}(T^{-1/4})$ is an upper bound for this smoothing-and-SGD construction.

Nonnegative gap estimates. No almost-surely finite procedure produces a nonnegative unbiased estimate of a value difference, since at $\Delta = 0$ nonnegativity forces the output to vanish on every finite transcript and the transcript laws are mutually absolutely continuous across Δ; so the mass-proportional recursions of exact-value gradient samplers cannot be imported directly.

C Two settled regimes

Two consequences of Theorem 10 admit matching upper and lower orders.

Corollary 17 (The low-budget regime is settled)**.** $\mathfrak{R}_T(n, b, R) = \Theta(R)$ *whenever* $T \lesssim n^2/(b^2R^2)$.

Algorithm 1 Theoretical fixed-K, hard-capped quantized block-Shapley SGD

Input: integer budget $T \geq 0$, b, R, level L, step size η when $K \geq 1$.
Set $s = \sigma(bR)$, $\alpha = (1+s)/2$, $c_{b,R} = 4/(1-s)$, $K = \lfloor T/(2c_{b,R}) \rfloor$.
if $K = 0$ **then**
  Output $\varnothing$ and stop.
end if
Initialize $y_1 = (\frac{1}{2}, \ldots, \frac{1}{2})$ and used budget $B = 0$.
for $t = 1, \ldots, K$ **do**
  Quantize $x_t = Q_L(y_t)$; use $x_t = y_t$ when $L = \infty$.
  Read off the tie blocks $B_1, \ldots, B_J$ of x_t, $J \leq \min\{L+1, n\}$.
  Sample block j with probability $\propto C_{m_j}$; sample $U \subseteq B_j$ from p_{m_j}.
  Draw independent $K_0, K_1 \sim \mathrm{Geom}(1-\alpha)$ on $\{1, 2, \ldots\}$.
  if $B + K_0 + K_1 > T$ **then**
    Output $\varnothing$ and stop.
  end if
  Use these lengths to duel $(H_j \cup U,\ H_j)$ (Lemma 14); set $B \leftarrow B + K_0 + K_1$ and $\widehat{q}_t = Z_U\, c_{m_j}(U) / \big(p_{m_j}(U)\ \Pr[j]\big)$ on B_j, zero elsewhere.
  $y_{t+1} = \Pi_{[0,1]^n}(y_t - \eta \widehat{q}_t)$.
end for
Let $\bar{x} = K^{-1} \sum_t x_t$; draw $\theta \sim \mathrm{Unif}[0,1]$; output $\{i : \bar{x}_i \geq \theta\}$.

Indeed the modular lower bound of Theorem 10 is then $\Theta(R)$, while returning any set achieves R. The lower and upper bounds thus match at order R below this saturation threshold, with the remaining $\sqrt{n}$ gap confined to larger budgets.

Corollary 18 (The modular subclass is settled at every budget)**.** *Restricted to modular F,* $\mathfrak{R}_T^{\mathrm{mod}} = \Theta_{b,R}\big(R \wedge n/(b\sqrt{T})\big)$.

For modular $F(S) = g(S)$, every canonical direction equals g. Average K independent estimators from Theorem 3 at a fully tied point. Their average $\bar{g}$ has $\mathbb{E}\|\bar{g} - g\|_2^2 \leq 36 V_{b,R} n/K$. The set $\widehat{S} = \{i : \bar{g}_i < 0\}$ satisfies $F(\widehat{S}) - \min F \leq \|\bar{g} - g\|_1$, so its expected excess is at most $6n\sqrt{V_{b,R}/K}$. The hard-cap coupling used for Theorem 8, with $K = \lfloor T/(2c_{b,R}) \rfloor$, gives an $O_{b,R}(n/(b\sqrt{T}))$ upper bound; use the fixed-set fallback for small budgets and cap the bound by R. The modular lower bound of Theorem 10 matches this dependence on n, T. Here the direction is independent of the query point, so all sensing can take place under a complete tie.

D Experimental protocol

Theoretical and experimental variants. Algorithm 1, with the step size in Theorem 8, has a fixed planned K and a pathwise hard cap, under which an overflowing round sends no queries and returns $\varnothing$. The experiments preceding Section D.7 use the fixed-round implementation of the proposed method, which runs its planned K without an overflow guard or fallback, with the stated theoretical or empirically calibrated step size. Its factor-two margin gives high-probability budget control. The theorem applies to the fixed-K, capped version with its prescribed step size; Section D.7 reports a separate budget-stopped diagnostic.

Risk and descriptive summaries. Optimization results report the conditional expectation $\mathbb{E}_\theta[F(\{i : \bar{x}_i \geq \theta\}) \mid \bar{x}] - F^\star$. We evaluate this expectation exactly by the Lovász level-set sum; medians and bootstrap intervals are over objective and oracle replicas, without an additional draw of the final threshold. The across-replica *mean* estimates the expected-risk quantity in the theorem; medians and median intervals describe typical performance across replicas. The fixed-round experiments retain these median summaries; the additional diagnostic reports means and confidence intervals as its prespecified primary summaries.

D.1 SYNTHETIC INSTANCE FAMILIES

Sensing instances. The sensing experiment fixes one objective for each pair (family, dimension) and reuses it over every ordered block-size profile. For the decomposable family, $F(S) = g(|S|) - \langle c, \mathbf{1}_S \rangle$. The n cardinality increments of g are independent $\mathrm{Unif}[0,1]$ draws sorted in decreasing order, and $c_i \overset{\text{iid}}{\sim} \mathrm{Unif}[0,1]$. Thus g is concave by construction and F is submodular. For the graph-cut family, a dense symmetric matrix is formed as $W = (U + U^\top)/2$ from an iid $\mathrm{Unif}[0,1]$ matrix U, its diagonal is set to zero, and $h_i \overset{\text{iid}}{\sim} \mathrm{Unif}[-1/2, 1/2]$ in $F(S) = \mathrm{cut}_W(S) + \langle h, \mathbf{1}_S \rangle$. Since $n \le 20$, all 2^n set values are enumerated. Each objective is scaled by its exact range diameter to $R = 1$ and then shifted so that $F(\varnothing) = 0$.

End-to-end decomposable instances. With $\alpha = 0.9$, $k^\star = \lfloor n/2 \rfloor$, and $A = 0.5/[\alpha(k^\star)^{\alpha-1}]$, it sets $g(k) = Ak^\alpha$ and draws a fresh $c_i \overset{\text{iid}}{\sim} \mathrm{Unif}[0,1]$ for every replica. At cardinality k, sorting c identifies both the minimizing and maximizing set, so scanning the descending and ascending prefix sums gives the exact minimum and range diameter. The objective is scaled to diameter $R = 1$ and shifted by $F(\varnothing)$.

End-to-end graph-cut instances. The variables form the closest-to-square four-neighbour grid. Edge weights are iid $\mathrm{Unif}[0.5, 1.5]$; the unary term is $-0.75 + \mathrm{Unif}[-1,1]$ on the left half and $+0.75 + \mathrm{Unif}[-1,1]$ on the right half. Max-flow capacities are multiplied by 10^5 and rounded to integers. Writing $\widetilde{m}$ for the resulting numerical minimum and $u = \sum_e w_e + \sum_i h_i^+$, it is scaled by $1/(u - \widetilde{m})$ and shifted by $F(\varnothing)$. This gives a numerical normalization at resolution 10^{-5}.

D.2 THE SENSING EXPERIMENT

For each ordered composition, coordinates are assigned consecutively to blocks of the specified sizes. One estimator draw samples block j with probability proportional to C_{m_j}, samples a size k with probability proportional to $\sqrt{a_{m_j,k}}$, samples U uniformly among the k-subsets of that block, duels $(H_j \cup U, H_j)$ through Lemma 14, and reweights by $1/(p_{m_j}(U)\Pr[j])$. Each configuration separately enumerates the canonical block-Shapley direction q_x and records $\|\widehat{q}_x - q_x\|_2^2$, the Bernoulli responses consumed, and their expected-query–MSE product. Each configuration uses 20 000 independent draws at $b = R = 1$.

The core protocol has $n \in \{8, 12, 16, 20\}$, $J \in \{2, 4, 8\}$, both families, and six ordered profiles (balanced, maximally imbalanced, and four seeded random positive compositions). Endpoint runs use $n \in \{8, 16\}$ and $J \in \{1, n\}$. This gives 160 configuration rows and 128 distinct family–dimension–ordered-profile combinations after duplicate endpoint profiles are removed. No optimization is run. Normalization and q_x are exact finite enumerations, while MSE and expected query cost are Monte Carlo estimates.

Across the 128 sensing configurations, the sweep gives $\mathbb{E}[N] \cdot \mathbb{E}\|\widehat{q}_x - q_x\|_2^2/(\sum_j C_{m_j})^2$ in $[114.259, 127.420]$, with mean 120.031 and maximum relative deviation 6.16% from that mean. At $b = R = 1$ the cost–moment upper bound of Theorem 3 carries the constant $c_{b,R}V_{b,R} = 2.08 \times 10^3$ and the lower bound of Theorem 5 the constant $1/2$, whose ratio is the factor $64s(1+s)/(1-s)^3 = 4.16 \times 10^3$ named in Section 4. The measured 120.031 sits a factor 17.3 below the upper constant and 240 above the lower one, so the sweep places the constant inside the proved window without narrowing it. The measured $J = n$-to-$J = 1$ product ratios are 5.634 and 9.606 for the decomposable family and 5.570 and 9.458 for graph cut at $n = 8$ and 16, respectively; the corresponding finite-n ratios of $(\sum_j C_{m_j})^2$ are 5.796 and 9.974.

D.3 THE END-TO-END EXPERIMENT

The fixed-round implementation is run at nominal budgets $T \in \{10^6, 10^7, 10^8, 10^9\}$; for $n = 64$ and $L \in \{1, 2, \infty\}$, an additional $T = 2 \times 10^9$ configuration extends the largest-budget axis. The empirically calibrated step size uses $G^2 = \bar{V}[J_L C_{\lceil n/J_L \rceil}]^2/b^2$, where $\bar{V} = 8.05$ and $J_L = \min\{L+1, n\}$; it uses an all-ceiling balanced-profile surrogate rather than the realized tie profile. The theoretical step size uses the worst-case theorem constant. The grid is $L \in \{2, 3, 4, 6, 8, 16, \infty\}$

Table 1: Controlled quantization ablation at n=64. Entries are median conditional expected normalized excess at the displayed duel targets, using 32 replicas in the first three budget columns and 128 in the final two. The comparison is between quantization levels of the same algorithm. Selected comparisons with the baseline appear in Section D.5.

	Decomposable					Graph cut				
duels	$5{\cdot}10^5$	$5{\cdot}10^6$	$5{\cdot}10^7$	$5{\cdot}10^8$	10^9	$5{\cdot}10^5$	$5{\cdot}10^6$	$5{\cdot}10^7$	$5{\cdot}10^8$	10^9
L=1	**0.396**	**0.261**	**0.080**	**0.014**	**0.009**	0.295	0.249	**0.147**	**0.050**	**0.037**
L=∞	0.457	0.421	0.350	0.208	0.162	**0.254**	**0.230**	0.198	0.128	0.101

Table 2: Median conditional expected normalized excess over 32 replicas at n=64 with the empirically calibrated step size, over the full quantization grid. Best level in each row in bold; blank entries were not run on the graph-cut family. Trivial baselines are 0.3725 and 0.0909.

T	L=1	L=2	L=3	L=4	L=6	L=8	L=16	L=∞
Decomposable								
10^6	**0.3961**	0.4018	0.4207	0.4189	0.4248	0.4274	0.4295	0.4575
10^7	**0.2613**	0.3270	0.3456	0.3491	0.3596	0.3628	0.3801	0.4212
10^8	**0.0795**	0.1567	0.1907	0.2024	0.2142	0.2332	0.2663	0.3504
10^9	**0.0137**	0.0344	0.0452	0.0558	0.0695	0.0797	0.1044	0.2053
Graph cut								
10^6	0.2946	0.2163		**0.2162**		0.2334		0.2537
10^7	0.2488	**0.1823**		0.1827		0.1951		0.2304
10^8	0.1466	**0.1269**		0.1302		0.1340		0.1981
10^9	**0.0518**	0.0591		0.0631		0.0666		0.1290

at $n \in \{16, 32, 64\}$ on the decomposable family and $L \in \{2, 4, 8, \infty\}$ at $n \in \{32, 64\}$ on graph cut, with the coarsest level L=1 added at n=64 on both. Ordinary configurations use 32 replicas. For $n = 64$, $L \in \{1, 2, \infty\}$, and nominal $T \in \{10^9, 2 \times 10^9\}$, four seeds of 32 replicas give 128 objective instances per cell. Medians are computed after pooling the individual replicas, rather than averaging batch medians. Figure 3 reports 95% percentile-bootstrap intervals for the medians at the two largest budgets, where individual replicas are retained. Confidence statements use 20 000 paired bootstrap resamples and only replica indices that share the same generated objective across levels. These replicas quantify variability under the fixed synthetic generators.

Fixed-round budget accounting. The experimental round count is $K = \lfloor T/(2c_{b,R}) \rfloor$, where the factor two is the tail margin that keeps the realized duel count below T with high probability in the absence of an overflow guard. Its expected count is $Kc_{b,R}$, approximately $T/2$ up to flooring. The reported consumption summaries over the grid lie between $0.49991\,T$ and $0.50035\,T$, so the displayed duel target is the realized consumption up to that tolerance.

Table 1 gives the level comparison. On decomposable objectives $L = 1$ has lower median excess than $L = \infty$ at all five budgets, and the ratio widens with the budget. On graph cut $L = \infty$ is lower at the two smallest budgets and $L = 1$ at the three largest.

Table 2 gives the full quantization grid behind the two levels reported in Section 7.

D.4 Public image graph cuts

Image objectives. We study quantization on six Middlebury image objectives. Table 3 compares all five levels at $Q \in \{10^5, 10^6, 10^7\}$, and Table 4 examines five image cases at $Q = 10^8$. Complete public results, construction details, and the definitions of the displayed cases accompany the data.

The baseline. We implement the sigmoid schedule in Algorithm 2 of Sho & Ito (2026). For our normalized objectives with $R = 1$, write

$$p_b = \sigma(b), \qquad C_b = \frac{|2p_b - 1|}{24b\,p_b^2(1 - p_b)^2}.$$

With displayed budget Q, the implementation uses

$$k = \max\left\{1, \left\lceil b^{2/5} C_b^{2/5} (Q/n)^{1/5} \right\rceil\right\}, \quad N = \left\lfloor \frac{Q}{nk} \right\rfloor, \quad \eta_0 = b^{6/5} C_b^{1/5} n^{2/5} Q^{-2/5}. \tag{2}$$

The original algorithm gives schedules before integer rounding; the ceiling and floor above make the repetitions integral and fit complete updates within Q. Every reported configuration has $N \geq 1$. Starting at $\frac{1}{2}\mathbf{1}$, each update stably sorts the iterate, queries each adjacent Lovász-prefix pair k times, and estimates its marginal by $b^{-1}\log[(s+1/2)/(k-s+1/2)]$, where s is the number of wins. Projected SGD uses η_0 and averages the N iterates preceding the updates. Uniform threshold rounding is evaluated by its exact conditional expectation. Here Q is the baseline's oracle cap, with nkN actual duels; the unused remainder is left unspent in Sections D.4 and D.5. Section D.7 applies the stated step multipliers and consumes the remainder while discarding the incomplete terminal update. The theoretical budget T and the fixed-round method's nominal input $2Q$ have the separate meanings specified below.

These experiments evaluate fixed public objectives under simulated logistic comparison feedback. Every response is generated independently from the Bernoulli oracle in Section 3; repeated calls may be sampled through their exact binomial sufficient statistics. Exact min-cut solutions and objective values are used only for offline evaluation and are not exposed to either optimizer.

Middlebury binary MRFs. We use the Flower, Person, and Sponge binary energies from the Middlebury MRF benchmark (Szeliski et al., 2008). For each image, deterministic SLIC superpixels (Achanta et al., 2012) with compactness 10 target 128 and 256 regions. Because SLIC's target is approximate, both target and realized dimensions are stored. The realized pairs are $(73, 176)$ for Flower, $(55, 124)$ for Person, and $(105, 209)$ for Sponge. Aggregating integer unary and boundary terms gives the exact restriction of the original pixel energy to labelings that are constant within each superpixel. The distributed label field is an all-one placeholder, so evaluation uses energy rather than semantic labels. Integer s–t min-cut gives the offline reference minimum. Dividing the raw objective by $R_{\mathrm{coef}} = \sum_i |h_i| + \sum_e w_e$ certifies $|F(S) - F(T)|/R_{\mathrm{coef}} \leq 1$. The two discretizations of each image are treated as separate workload conditions.

Public-workload protocol. The image grid uses $b = R = 1$, displayed duel targets $Q \in \{10^5, 10^6, 10^7\}$, $L \in \{1, 2, 4, 8, \infty\}$, the sensing-calibrated step size of Section D.5, and 32 stochastic replicas with seed 20260824. The fixed-round implementation receives nominal budget $2Q$ so that its expected consumption targets Q; Sho–Ito receives cap Q, and every realized count is retained. Here Q is a consumption target; realized counts can differ across methods, and the synthetic tolerance quoted above was not established for every public cell. The primary metric is conditional expected normalized excess per fixed image objective.

The initialization reference. The optimization tables use the conditional expected normalized excess of the untrained iterate $y_1 = (\frac{1}{2}, \ldots, \frac{1}{2})$, namely $\frac{1}{2}[F(\varnothing) + F([n])] - \min_S F$ under the same scaling. This reference requires no duels. Lower excess indicates progress beyond initialization.

Results. Table 3 compares levels under the specified sensing-calibrated schedule. On Middlebury, the preference for $L = 1$ across the tested tasks and budgets shows that finer query resolution is not necessary for lower optimization error. The initialization reference distinguishes an improvement over another level from progress on the objective itself. Table 4 compares image cases with the baseline at a higher budget.

D.5 Sensing-calibrated public-workload extension

The sensing-calibrated schedule used throughout the public suite sets $\bar{V} = 8.6$ in the empirical step of Section D. The maximum observed MSE coefficient after division by $(\sum_j C_{m_j})^2$ across the 160 sensing rows is approximately 8.568, which we round upward to 8.6. This empirical MSE calibration does not certify the uniform second-moment assumption needed for the theorem's step size. At $Q = 10^7$, a step-sensitivity grid multiplies this primary step by 0.5 and 2. The budget extension runs $L \in \{1, 2, \infty\}$ and Sho–Ito at $Q = 10^8$. All cells retain 32 replicas with seed 20260824, with one fixed level and step grid shared across tasks.

Table 3: Quantization on six Middlebury image objectives at the sensing-calibrated step. Median conditional expected normalized excess over 32 replicas (lower is better). Bold marks the minimum among the five levels within one fixed objective and budget. The final column is the initialization reference, which costs no duels.

fixed objective (n)	Q	L=1	L=2	L=4	L=8	L=∞	init.
Middlebury binary MRFs							
Flower 73	$1{\cdot}10^5$	**.4060**	.4348	.4566	.4610	.4697	.4845
	$1{\cdot}10^6$	**.2741**	.3579	.3814	.4086	.4420	.4845
	$1{\cdot}10^7$	**.1423**	.2237	.2577	.2804	.3909	.4845
Flower 176	$1{\cdot}10^5$	**.4597**	.4688	.4668	.4800	.4901	.4793
	$1{\cdot}10^6$	**.3870**	.4239	.4379	.4380	.4742	.4793
	$1{\cdot}10^7$	**.2457**	.3231	.3527	.3717	.4547	.4793
Person 55	$1{\cdot}10^5$	**.4319**	.4559	.4609	.4713	.4891	.4937
	$1{\cdot}10^6$	**.3183**	.3750	.4036	.4154	.4499	.4937
	$1{\cdot}10^7$	**.1183**	.2004	.2556	.2868	.3998	.4937
Person 124	$1{\cdot}10^5$	**.4684**	.4694	.4799	.4792	.4843	.4916
	$1{\cdot}10^6$	**.4220**	.4365	.4520	.4532	.4878	.4916
	$1{\cdot}10^7$	**.2675**	.3487	.3827	.3962	.4576	.4916
Sponge 105	$1{\cdot}10^5$	**.4791**	.4818	.4913	.4821	.4811	.4962
	$1{\cdot}10^6$	**.4240**	.4399	.4570	.4593	.4882	.4962
	$1{\cdot}10^7$	**.2745**	.3554	.3799	.4149	.4581	.4962
Sponge 209	$1{\cdot}10^5$	**.4778**	.4889	.4855	.4915	.4892	.4951
	$1{\cdot}10^6$	**.4494**	.4686	.4775	.4768	.4920	.4951
	$1{\cdot}10^7$	**.3598**	.4205	.4358	.4486	.4828	.4951

Table 4: Image case studies at $Q = 10^8$. Median conditional expected normalized excess over 32 replicas; lower is better. Initialization requires no duels. Bold marks the lowest method median. The last column gives the 95% bootstrap interval for the difference of medians, $L = 1$ minus the baseline.

fixed objective (n)	L=1	L=2	L=∞	Sho–Ito	init.	Difference interval
Flower 176	**.1286**	.1822	.3859	.1405	.4793	$[-0.0152, -0.0069]$
Person 55	**.0294**	.0625	.2552	.0326	.4937	$[-0.0064, -0.0015]$
Person 124	**.0946**	.1674	.3989	.1018	.4916	$[-0.0109, -0.0019]$
Sponge 105	**.0980**	.1732	.4099	.0994	.4962	$[-0.0058, +0.0037]$
Sponge 209	**.1830**	.2816	.4623	.2097	.4951	$[-0.0313, -0.0225]$

This sensitivity grid varies the proposed method's step size while the baseline uses equation 2. Section D.7 applies the same multipliers to each method's base step size and reports a Flower case study.

Uncertainty estimates. For each displayed image objective, we use 20 000 bootstrap resamples drawn independently within each method to estimate the difference of medians. The intervals describe oracle variability for that objective and use no adjustment for multiple comparisons.

In the image cases of Table 4, $L = 1$ has lower median error than both the baseline and initialization. These comparisons illustrate how coarse query points can support accurate solutions under a fixed comparison budget.

D.6 A MISSPECIFIED LINK

Both the analysis and the estimators assume the logistic link. We evaluate sensitivity to this assumption by replacing the oracle with a probit or linear link whose slope at zero matches the logistic one. This preserves the local comparison slope near a tie while changing the tails. The algorithm retains its logistic design, including geometric truncation, importance weights, step size, and the supplied values of b and R.

Table 5: Median conditional expected normalized excess at n=64 under a misspecified link, against the logistic link the method is designed for.

Family	Method	Duels	Logistic	Probit	Linear
Decomposable	Fixed-round, L=1	$5{\cdot}10^7$	0.0795	0.0788	0.0816
Decomposable	Fixed-round, L=1	$5{\cdot}10^8$	0.0137	0.0135	0.0137
Graph cut	Fixed-round, L=1	$5{\cdot}10^7$	0.1466	0.1501	0.1474
Graph cut	Fixed-round, L=1	$5{\cdot}10^8$	0.0518	0.0542	0.0503

Table 6: Independent validation of the calibrated second-moment prediction. Each row has 30 calibration and 60 validation targets. Deviation is $|\mathbb{E}\|\widehat{q}_x\|^2/\text{prediction} - 1|$, using the offline exact design moment.

bR	Calibrated a_b	Maximum deviation	Within 7%
0.25	4.66593	0.68%	60/60
1	8.24700	3.24%	60/60
3	51.50521	9.21%	54/60

Table 5 reports the outcome over 32 replicas per cell. Every entry stays within 5% of its logistic counterpart on both families and at both budgets. Matching the local slope makes these alternatives comparable near a tie. Performance is stable across the tested links and budgets.

D.7 Independent validation and optimization case studies

Budget and grid. This separate diagnostic completes updates until exactly Q simulated duels have been spent. If the final direction, or the baseline's final prefix gradient, cannot fit, only the remaining feedback is generated and that incomplete update is discarded. Only iterates from completed updates enter the average; if none complete, the initial quantized point is used. This protocol studies budget-stopped dynamics, whereas Theorem 8 analyzes fixed K with the hard-cap empty-set fallback and the theorem step size.

With normalized $R = 1$, we compare $L \in \{1, 2, \infty\}$ and Sho–Ito at $bR \in \{0.25, 1, 3\}$, with the same step multipliers $\{0.5, 1, 2\}$ for all four methods. Multiplier 1 is the nominal configuration. The optimization examples reported here use Flower at $n = 73$, 16 oracle seeds, and $Q \in \{10^5, 10^6, 10^7\}$. The complete grid, task construction details, and case definitions accompany the data.

Using constants a_b calibrated separately for each bR before optimization, the proposed levels set their base step to

$$G_L^2 = \frac{a_b}{b^2}\left[J_L C_{\lceil n/J_L \rceil}\right]^2, \qquad \eta_0 = \sqrt{\frac{n}{G_L^2 \bar{K}}}.$$

Here $J_L = \min\{L+1, n\}$, $J_\infty = n$, and $\bar{K} = \max\{1, \lfloor Q/c_{b,1} \rfloor\}$. The actual number of completed updates is random for these levels. The baseline uses equation 2. Each method multiplies its own η_0 by the grid value to compare sensitivity around the respective base schedules.

Independent sensing validation. Calibration uses 90 targets from the two sensing families at $n = 6$; validation uses 180 targets at $n \in \{8, 10\}$, with disjoint instance seeds and random streams. Each target has 1000 draws. The equal-target mean of the estimated $b^2\mathbb{E}\|\widehat{q}_x\|_2^2/(\sum_j C_{m_j})^2$ determines a_b separately within each calibration stratum. We then keep a_b fixed for validation. The prediction for a validation target is $a_b(\sum_j C_{m_j})^2/b^2$. These constants are empirical estimates, without a certified uniform second-moment bound.

Table 6 gives the independent check. The exact design second moments are within 7% of prediction on 174/180 targets and within 9.22% on all targets. This comparison concerns new targets and a second-moment prediction, separately from the fit of the query cost times MSE over 128 sensing configurations in Figure 3a. Finite-sample estimates are noisier. At $bR = 3$, observed second moments range from 0.567 to 1.520 times prediction. Bootstrap bands for the heavy-tailed squared

Table 7: Flower case study across feedback strengths, budgets, and step multipliers. Entries are means over 16 seeds; initialization excess is 0.48449. The last column is the 95% paired bootstrap interval for the mean excess of $L = 1$ minus the baseline.

bR	Q	Multiplier	$L = 1$	Sho–Ito	Paired interval
0.25	10^5	1	0.45439	0.46680	$[-0.04007, 0.01817]$
0.25	10^6	0.5	0.37843	0.38811	$[-0.03237, 0.01004]$
0.25	10^7	0.5	0.22992	0.25893	$[-0.04541, -0.01261]$
0.25	10^7	1	0.25282	0.27736	$[-0.03922, -0.00917]$
0.25	10^7	2	0.33622	0.34938	$[-0.02229, -0.00250]$
1	10^7	0.5	0.13354	0.13564	$[-0.01169, 0.00703]$

roulette samples are nominal empirical bands, without a coverage guarantee. Validation also records exact design MSE, actual query costs, and sample variances.

Optimization summaries. The primary summary is across-replica mean conditional expected excess; medians are descriptive. Reported 95% percentile bootstrap intervals use $10\,000$ resamples of oracle seeds within each fixed configuration. Paired differences resample matched runs with the same fixed instance, bR, seed, target budget, and realized Q, with no cost-selected subset. Intervals are marginal, without adjustment for multiple comparisons, and do not support inference over a population of tasks. The common reference, requiring no duels, is the conditional expected excess of $\frac{1}{2}\mathbf{1}$; the actual quantized initialization can differ and is recorded separately. Exact references and held-out labels remain evaluation-only.

Flower case study. Table 7 examines Flower across several feedback strengths, budgets, and step multipliers. In these configurations, $L = 1$ has lower mean error than both the baseline and initialization, illustrating that coarse comparisons can support progress on the objective.

Tie formation during optimization. The diagnostics record actual block counts and sizes, completed updates, quantized motion, stagnation, and constant outputs. At the nominal Flower configuration, the mean numbers of blocks are 2.00, 2.94, and 71.80 for $L = 1, 2, \infty$, respectively, with approximately $672\,450$ completed updates for each. The mean longest stagnant fraction of a run is 0.48% for $L = 1$ and 1.68% for $L = 2$. Thus the finite levels create ties during trajectories that continue to move, connecting the estimator's tie structure to actual optimization trajectories.

D.8 Cost

Each round of Algorithm 1 sorts the iterate in $O(n \log n)$ time to form its tie blocks and performs $O(n)$ vector operations. The cost of simulating a wide comparison depends on the representation of the objective. Separately, the scalar primitive consumes $4/(1 - s)$ Bernoulli draws in expectation by Lemma 14, independently of n and J for fixed b, R. The experiment records the realized query count alongside every MSE draw.