REFINED LABEL COMPARISONS FOR CONSTRAINED NON-MONOTONE k-SUBMODULAR MAXIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Non-monotone k-submodular maximization under support constraints requires controlling both feasibility conflicts and label-dependent losses. We develop Random-root, a minimax randomized greedy method using support budgets, and a continuous algorithm that refines these comparisons through label-specific replacement losses. Under a matroid, the continuous algorithm improves the approximation coefficient from $\sqrt{2}-1$ to $\alpha_2 \approx 0.4218$, up to any accuracy loss $\tau > 0$. For every $k \geq 2$ and fixed accuracy, the baseline uses $O(nk\rho)$ value queries for n elements and rank ρ. Refinements attain $0.4366 - \tau$ for three to five labels and a uniform $0.4259 - \tau$ guarantee for all $k \geq 3$. More generally, for orthant-submodular objectives with r-wise monotonicity, Random-root yields a closed-form approximation guarantee on p-systems, while continuous optimization attains $\alpha_r - \tau$ under matroid constraints, with $\alpha_r > 1/(1+\sqrt{r})$ for $r \geq 2$. For $r \geq 3$, the extra-label refinement saturates at $k = r+1$: all $k > r$ share the same exact replacement-loss envelope.

1 INTRODUCTION

Offline expert-data allocation trades feature coverage against separation of incompatible examples. Submodular objectives filter image–caption datasets (Das et al., 2026), while expert routing accounts for gradient conflicts within experts (Yang et al., 2025). Our model (Section 2) captures this trade-off through diminishing coverage returns and potentially negative label marginals.

Multi-label optimization also models sensor placement and topic-dependent influence allocation (Ohsaka and Yoshida, 2015; Zhou et al., 2025). The k-submodular framework allows diminishing returns and negative label marginals (Huber and Kolmogorov, 2012; Ward and Živný, 2016). Selection can block optimal elements; labeling changes their comparison value. We study how these losses determine local randomization under full-domain orthant submodularity and r-wise monotonicity, including k-submodularity at $r = 2$.

For non-monotone k-submodular matroid maximization, the continuous algorithm improves $\sqrt{2}-1$ (Chen et al., 2026b) to $\alpha_2 - \tau$, where $\alpha_2 \approx 0.4218$ and $\tau > 0$ is arbitrarily small. For all $k \geq 2$, the baseline matches Chen et al. (2026b)'s $O(nk\rho)$ value-query bound at fixed accuracy (n elements, rank ρ). The refined coefficient is $\alpha_2^+ > 0.4366$ for three to five labels; a competing capacity constraint lowers it from six labels onward.

Random-root's adaptive uniform prefix solves the local minimax problem, yielding the p-system factor $R(p, r)$ in (2). For matroids and $r \geq 3$, one extra label saturates the replacement envelope: all $k > r$ share a stronger coefficient (Theorem 2.2).

Related work. Support charging builds on Fisher et al. (1978); Mestre (2006); Feldman et al. (2017). For k-submodular maximization, Sakaue (2017) studies monotone objectives under matroids; Ye et al. (2025) give p-system factors $1/(p+1)$ for monotone and $1/(p+2)$ for normalized real-valued objectives. Huang et al. (2026) obtain $1/(r+1)$ for nonnegative OSM, r-wise matroid maximization. Root improves the latter two bounds at $r = 2$ and $r > 1$, respectively, and extends Chen et al. (2026b)'s $r = 2$ matroid coefficient to p-systems and larger r (Table 1).

Continuous greedy attains the optimal $1 - 1/e$ for monotone submodular matroid maximization (Călinescu et al., 2011); measured continuous greedy treats non-monotone objectives (Feldman

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.

Table 1: Selected guarantees on the common class of normalized nonnegative objectives. Continuous guarantees allow an additive loss τ in the approximation coefficient.

Method	Objective / constraint	Guarantee
Ye et al. (2025)	k-submodular / p-system	$1/(p+2)$
Huang et al. (2026)	OSM, r-wise / matroid	$1/(r+1)$
Zhou et al. (2025)	k-submodular / matroid	$1/3-\tau$
Chen et al. (2026b)	k-submodular / matroid	$\sqrt{2}-1$
Random-root	OSM, r-wise / p-system	$R(p,r)$
Continuous baseline, $k \geq r$	OSM, r-wise / matroid	$\alpha_r-\tau$
Continuous, $r \geq 3,\ k > r$	OSM, r-wise / matroid	$\alpha_r^+-\tau$
Continuous, $r=2,\ k \geq 3$	k-submodular / matroid	$\beta_{2,k}-\tau$

et al., 2011). We use Zhou et al. (2025)'s categorical extension and modified comparison trajectory for k-submodularity. Iwata et al. (2016); Oshima (2021) develop label-specific comparisons for unconstrained k-submodular maximization. Chen et al. (2026b) use proportional top-two randomization and local replacement and addition bounds. We optimize weight-dependent replacement comparisons, classify their label-count envelopes, and couple them to a decreasing-weight potential.

Sequential updates reduce continuous greedy's discretization and sampling costs (Badanidiyuru and Vondrák, 2014). Stochastic continuous submodular optimization uses variance reduction (Hassani et al., 2020) and Lyapunov analysis of bias and noise (Chen et al., 2026a). Our query bound uses two-label protection and slack-weighted selection errors.

Technical overview. Support budgets determine Root's optimal prefix. Retaining inactive elements' optimal labels yields losses reduced by asymmetric distributions. Matroid exchange aggregates the rowwise bounds: local slack pays for weight decay, and sequential updates preserve the potential. Two approximately largest labels protect all unchosen labels; the other $r-2$ are arbitrary.

2 Problem formulation and main results

Let V contain n elements and let $[k]=\{1,\dots,k\}$. A labeling $s \in \Omega = \{0,\dots,k\}^V$ has support $\mathrm{supp}(s)=\{e : s_e \neq 0\}$. Write $s \preceq t$ if t preserves every positive label of s, and $s+e_j$ for assigning label j to an inactive element e. For disjoint supports, $s+t$ combines the two labelings. For $f : \Omega \to \mathbb{R}$, put $\Delta_{e,j}f(s) = f(s+e_j)-f(s)$. Orthant submodularity (OSM) requires $\Delta_{e,j}f(s) \geq \Delta_{e,j}f(t)$ when $s \preceq t$ and $e \notin \mathrm{supp}(t)$. For integer $1 \leq r \leq k$, r-wise monotonicity means

$$\sum_{j\in J} \Delta_{e,j} f(s) \geq 0 \qquad (J \subseteq [k],\ |J|=r). \tag{1}$$

Both conditions hold throughout Ω; $r=1$ gives monotonicity, and OSM with $r=2$ characterizes k-submodularity (Ward and Živný, 2016). Larger supplied orders r weaken the restriction on negative marginals.

Feasible supports form a nonempty downward-closed family $\mathcal{I}$, a p-system if any two maximal independent subsets of each $W \subseteq V$ differ in size by at most a factor $p \geq 1$. Matroids have $p=1$. Set $\rho = \max_{S\in\mathcal{I}} |S|$ and $\mathrm{OPT} = \max_{\mathrm{supp}(s)\in\mathcal{I}} f(s)$. For $F = f - f(\mathbf{0})$ and $M = \mathrm{OPT} - f(\mathbf{0}) \geq 0$, a coefficient a means $\mathbb{E}F(Y) \geq aM$, giving the usual multiplicative guarantee when $f(\mathbf{0}) \geq 0$. Value queries are exact on all labelings, including infeasible supports; independence queries test membership in $\mathcal{I}$. Their counts are Q_f, Q_I. Outputs are feasible and empty at rank zero.

Theorem 2.1 (Root guarantee)**.** *For a full-domain real-valued OSM, r-wise monotone objective and a p-system, Algorithm 1 returns a feasible $\widehat{s}$ with*

$$\mathbb{E}F(\widehat{s}) \geq R(p,r)M, \qquad R(p,r) = \frac{2}{p+1+\sqrt{(p-1)^2+4pr}}. \tag{2}$$

For $\rho \geq 1$, it uses $O(nk\rho)$ value queries and arithmetic operations, $O(n\rho)$ independence queries, and $O(n+k)$ numerical records.

The continuous coefficients use integral (15) and the envelopes in Sections 4–5.

Theorem 2.2 (Continuous guarantees)**.** *Under the same objective assumptions and a matroid constraint, for every $k \geq r \geq 2$ and $0 < \tau < a_{r,k}$ there is a finite randomized algorithm returning feasible Y with*

$$\mathbb{E}F(Y) \geq (a_{r,k}-\tau)M, \qquad a_{r,k} = \begin{cases} \alpha_r, & k = r, \\ \alpha_r^+, & r \geq 3,\ k > r, \\ \beta_{2,k}, & r = 2,\ k \geq 3. \end{cases}$$

Every extra-label coefficient exceeds $\alpha_r > R(1,r)$, the baseline attainable up to any prescribed loss for all $k \geq r$. For $r = 2$, $\beta_{2,k} = \alpha_2^+ > 0.4366$ when $3 \leq k \leq 5$, while $0.4259 < \beta_{2,k} < \alpha_2^+$ when $k \geq 6$. The baseline satisfies $\alpha_2 > 0.421795$.

Sections 3–6 prove these guarantees and the baseline query bound (Theorem 6.2).

An offline data-allocation example. For $k \geq 2$, $S_j = \{e : s_e = j\}$ contains the examples assigned to expert j. Features $z_{e\ell}$, expert weights $a_{j\ell}$ and caps $b_{j\ell}$, all nonnegative, define the normalized, nonnegative, monotone submodular coverage $C_j(A) = \sum_{\ell=1}^{m} a_{j\ell} \min\{b_{j\ell}, \sum_{e\in A} z_{e\ell}\}$. With an undirected conflict graph $w_{uv} = w_{vu} \geq 0$, $w_{uu} = 0$, and $\eta \geq 0$, define

$$f(s) = \sum_{j=1}^{k} C_j(S_j) + \frac{\eta}{2}\sum_{j=1}^{k} \mathrm{cut}_w(S_j), \qquad \mathrm{cut}_w(A) = \sum_{u\in A,\, v\notin A} w_{uv}. \tag{3}$$

Writing $d_e = \sum_{v\in V} w_{ev}$, the graph term equals $\eta\sum_{e\in\mathrm{supp}(s)} d_e/2 - \eta\sum_j \sum_{\{u,v\}\subseteq S_j} w_{uv}$: a degree-weighted selection reward minus within-expert conflict penalties. A cardinality budget or disjoint-source upper quotas constrain the support by a matroid.

The surrogate is normalized and nonnegative throughout Ω. Put $w(e,A) = \sum_{v\in A} w_{ev}$. An inactive example has marginal

$$\Delta_{e,j} f(s) = C_j(S_j \cup \{e\}) - C_j(S_j) + \eta\big(d_e/2 - w(e,S_j)\big).$$

The terms are nonincreasing in S_j, proving OSM. Monotone coverage and disjointness give, for $i \neq j$,

$$\Delta_{e,i} f(s) + \Delta_{e,j} f(s) \geq \eta\big(d_e - w(e, S_i \cup S_j)\big) \geq 0.$$

Thus the guarantees apply with $r = 2$, despite possible negative marginals. For two examples u, v, set $w_{uv} = \eta = 1$ and $C_j(A) = \min\{1, |A|\}$. After assigning u to expert 1, assigning v to experts 1 or 2 has marginal $-1/2$ or $3/2$, respectively: its label determines whether the objective falls or rises.

Comparison rows. For $d \geq r$, define $\mathcal{C}_r^d = \{g \in \mathbb{R}^d : \sum_{j\in J} g_j \geq 0 \text{ for all } J \subseteq [d],\ |J| = r\}$. Every marginal row lies in $\mathcal{C}_r = \mathcal{C}_r^k$. For $g \in \mathcal{C}_r^d$, write $m(g) = \max_j g_j$ and $\mathcal{A}(g) = \{a \in \mathcal{C}_r^d : a \leq g\}$. For a full row g, a represents marginals at a compatible comparison state: OSM gives $a \leq g$ and r-wise monotonicity gives $a \in \mathcal{C}_r$. For $\Delta_d = \{q \in \mathbb{R}_{\geq 0}^d : \sum_j q_j = 1\}$, a distribution $q \in \Delta_k$ has gain $h(q,g) = q \cdot g$. Both algorithms use $\mathcal{A}(g)$. Inequalities are coordinatewise; $x_+ = \max\{x, 0\}$, $\mathbf{e}_j$, $\mathbf{1}$ and $\mathbf{1}[\cdot]$ denote positive parts, basis vectors, the all-ones vector and indicators.

3 Random-root on p-systems

Random-root selects a feasible element of largest label marginal and randomizes its label. For an observed row g and label distribution q, the gain and support-budgeted comparison loss are

$$h(q,g) = q \cdot g, \qquad D_{\mathrm{R}}^p(g,q) = \max_{a\in\mathcal{A}(g)} \{p\, m(g) - q \cdot a\}. \tag{4}$$

The budget $p\,m(g)$ pays for displaced support; $q \cdot a$ is the expected comparison gain.

Index labels by decreasing marginal, breaking ties by original index. Write $G_t = \sum_{j=1}^{t} g_j$ and let $L_s(g)$ sum the s smallest entries, with $L_0 = 0$. Uniform sampling from the top $t \le r$ labels gives $h = G_t/t$ and $q \cdot a \ge -L_{r-t}/t$. Choose

$$t^*(g; p, r) = \min \arg \min_{1 \le t \le r,\ G_t > 0} \frac{tp\,m(g) + L_{r-t}(g)}{G_t}. \tag{5}$$

Rank selection computes this positive-gain optimum of D_{R}^p/h in $O(k)$ time (Appendix A.3). Write $\mathrm{Top}_t(g)$ for the first t labels and $E(s) = \{e \notin \mathrm{supp}(s) : \mathrm{supp}(s) + e \in \mathcal{I}\}$.

Algorithm 1 Random-root greedy

Require: Value oracle for f, independence oracle for $\mathcal{I}$, $p \ge 1$, $1 \le r \le k$

$s \leftarrow \mathbf{0}$
while $E(s) \neq \varnothing$ **do**
  $e \in \arg\max_{u \in E(s)} \max_{j \in [k]} \Delta_{u,j} f(s)$ ▷ fixed tie order
  $g \leftarrow (\Delta_{e,j} f(s))_{j \in [k]}$; **if** $m(g) = 0$ **then return** s
  $J \sim \mathrm{Unif}(\mathrm{Top}_{t^*(g;p,r)}(g))$
  $s \leftarrow s + e_J$
end while
return s

Theorem 3.1 (Local minimax)**.** *For every nonzero $g \in \mathcal{C}_r$, the uniform prefix in* (5) *minimizes $\Lambda_p(g, q) := D_{\mathrm{R}}^p(g, q)/h(q, g)$ over distributions with $h(q, g) > 0$, with minimum equal to the selected ratio. Moreover,*

$$\sup_{g \in \mathcal{C}_r \setminus \{0\}} \inf_{q : h(q,g) > 0} \Lambda_p(g, q) = C_{p,r} := \frac{p - 1 + \sqrt{(p-1)^2 + 4pr}}{2}. \tag{6}$$

For the selected ratio λ and $m = m(g)$, the hidden row $b_j = \min\{g_j, pm - \lambda g_j\}$ lies in $\mathcal{A}(g)$ (Appendix A.1) and gives $pm - q \cdot b \ge \lambda q \cdot g$ for every distribution q.

Worst-row value in Theorem 3.1. Put $C = C_{p,r}$ and $\theta = C + 1 - p > 0$, so $C\theta = pr$. The top-one and top-r prefixes give complementary bounds. If $G_r \le \theta m$, the first has ratio at most $p + L_{r-1}/m \le p - 1 + G_r/m \le C$. If $G_r > \theta m$, the second has ratio $prm/G_r < C$. Balancing them gives $C(C + 1 - p) = pr$, hence (6). For equality, take $g^* = (1, c, \ldots, c)$ with $c = p/(1 + C)$. Every prefix ratio equals C since $tp + (r - t)c = C(1 + (t - 1)c)$ for $1 \le t \le r$. □

Proof of Theorem 2.1. Fix an optimal labeling o with support O. Let s^t, S_t denote the algorithmic labeling and support after t of its $T \le \rho$ insertions. At step t, the selected element, row and action are e_t, g_t, q_t; write $M_t = m(g_t)$ and $h_t = q_t \cdot g_t > 0$.

Give $u \in O$ deadline t when selected or first blocked by S_t. Let C_t contain those with deadline t, $D_t = \bigcup_{i \le t} C_i$, $b_t = |C_t|$ and $B_t = |D_t|$. The deadline lemma (Lemma A.1) gives $B_t \le pt$ and nonincreasing M_t. With $M_{T+1} = 0$, summation by parts gives pathwise

$$\sum_{t=1}^{T} b_t M_t = \sum_{t=1}^{T} B_t (M_t - M_{t+1}) \le p \sum_{t=1}^{T} t (M_t - M_{t+1}) = p \sum_{t=1}^{T} M_t. \tag{7}$$

Let o^t carry the algorithmic labels on S_t and original labels on $O \setminus D_t$. Deleting C_t from o^{t-1} gives $\bar{o}^t$. Each deleted element was addable before step t, so OSM bounds the deletion loss by $b_t M_t$. The insertion row $a_t = (\Delta_{e_t,j} F(\bar{o}^t))_j$ lies in $\mathcal{A}(g_t)$. Inserting the sampled label J_t gives $F(o^{t-1}) - F(o^t) \le b_t M_t - a_{t,J_t}$. Conditional expectations before each label draw, telescoping, (7) and Theorem 3.1 yield

$$\begin{aligned} M - \mathbb{E} F(o^T) &\le \mathbb{E} \sum_t (p M_t - q_t \cdot a_t) \\ &\le C_{p,r} \mathbb{E} \sum_t h_t = C_{p,r} \mathbb{E} F(s^T). \end{aligned}$$

Each element of $O \setminus D_T$ remains addable with zero maximum marginal at termination. OSM therefore gives $F(o^T) \le F(s^T)$. Hence $(1 + C_{p,r})\mathbb{E}F(s^T) \ge M$, proving (2). See Appendix A.2 for the stopping-time formalization and resource bounds. □

At $r = 1$, a best label gives $1/(p + 1)$; at $p = 1, r = 2$, the guarantee is $\sqrt{2} - 1$.

4 Replacement comparisons and continuous optimization

For a matroid and $k \ge r \ge 2$, the comparison trajectory retains inactive base elements' optimal labels. Replacing a label gives a finer loss than the support budget:

$$D_{\mathrm{E}}(g, q) = \max_{a \in \mathcal{A}(g),\, j \in [k]} (a_j - q \cdot a) \le D^1_{\mathrm{R}}(g, q). \tag{8}$$

The inequality follows from $a_j \le m(g)$ and (4). We optimize the slack $h - \lambda H_{\mathrm{E}}$, where $H_{\mathrm{E}} = h + D_{\mathrm{E}}$ and λ is the comparison weight. Matroid exchange aggregates these bounds in the potential.

4.1 The local saving

For $k = r$, sort $g_1 = m \ge g_2 = v \ge \cdots$, and put $s = \sum_i g_i$, $Q = \max_{i \ge 2} q_i$ and $H^1_{\mathrm{R}} = h + D^1_{\mathrm{R}}$. Solving the hidden-row problem gives the exact saving.

Proposition 4.1 (Exact diagonal charge difference)**.** *For every $q \in \Delta_r$,*

$$H^1_{\mathrm{R}}(g, q) = m + s\|q\|_\infty, \tag{9}$$

$$H_{\mathrm{E}}(g, q) = \max\{v + sq_1, m + sQ\}, \tag{10}$$

$$H^1_{\mathrm{R}}(g, q) - H_{\mathrm{E}}(g, q) = \min\{m - v, s(q_1 - Q)_+\}. \tag{11}$$

Proof. Writing $a = g - d$ identifies $\mathcal{A}(g)$ with $d \ge 0$, $\sum_i d_i \le s$. Consequently $\min_a q \cdot a = h - s\|q\|_\infty$. For fixed j, $\max_a(a_j - q \cdot a) = g_j - h + s \max_{i \ne j} q_i$, since d's largest coefficient is $\max_{i \ne j} q_i \ge 0$. Maximizing over j gives (10). If $q_1 \le Q$, both charges are $m + sQ$. Otherwise their difference is the smaller of $m - v$ and $s(q_1 - Q)$. The calculation includes $s = 0$. □

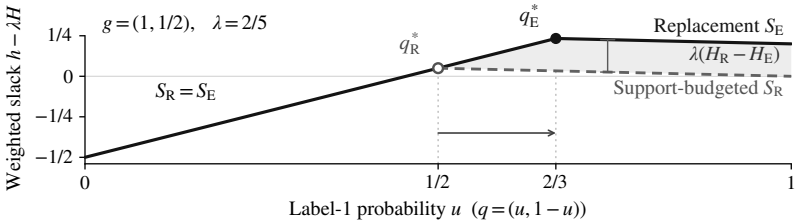

Figure 1: Label-specific replacement changes the optimal action. For $g = (1, 1/2)$, $p = 1$, $k = r = 2$ and $\lambda = 2/5$, the curves show $h - \lambda H$ over $q = (u, 1 - u)$. The maximizer shifts from Root's ratio optimizer $(1/2, 1/2)$ to $(2/3, 1/3)$.

For a nonnegative row with $s > 0$, let $u^{(t)}$ be uniform on the first t labels. The shifted prefix $q^{(t)} = \delta \mathbf{e}_1 + (1 - \delta)u^{(t)}$, with $\delta = (m - v)/s$, balances the charges for $t \ge 2$, attaining the full reduction $m - v$. A uniform prefix with $t \ge 2$ has $q_1 = Q$ and zero reduction. For $g = (1, t)$, $0 < t \le 1$, the best-label and proportional actions have slack difference $t[(1 - t)/(1 + t) - \lambda]$. Hence an optimal action is

$$q^*(\lambda) = \begin{cases} (1, 0), & 0 \le \lambda \le (1 - t)/(1 + t), \\ (1, t)/(1 + t), & \lambda > (1 - t)/(1 + t). \end{cases} \tag{12}$$

Decreasing weight can switch the optimal rule from proportional randomization to the best label.

For $k \ge r$, the sparse baseline uses r labels including the largest entries, with caps $b_i = [g_i]_+$ and $S = \sum_i b_i$. The sufficient charge $H_+(g, q) = \max_{i \ne j}(b_j + Sq_i)$ bounds every comparison label,

including those outside the action's support. Minimizing the shifted-prefix maximum of $h - \lambda H_+$ over normalized nonnegative r-rows defines $\mathcal{K}_r(\lambda)$ (Lemmas B.2 and B.3). The exact full-row envelope is

$$K_{r,k}(\lambda) = \inf_{\substack{g \in \mathcal{C}_r \\ m(g)=1}} \max_{q \in \Delta_k} \{h(q,g) - \lambda H_{\mathrm{E}}(g,q)\}. \tag{13}$$

Thus $K_{r,r} = \mathcal{K}_r$; for $k > r$, the full model refines the r-label relaxation.

4.2 Continuous lifting

Let $x_{ej} \geq 0$, $\sum_j x_{ej} \leq 1$ be label probabilities. Draw rows independently, leaving e inactive with the remaining probability. The categorical extension is $\mathcal{F}(x) = \mathbb{E}F(X_x)$. Complete an optimum to a labeled base o using nonnegative maximum marginals. Its comparison extension is $\mathcal{Z}(x) = \mathcal{F}(z(x))$, where

$$z_{ej}(x) = x_{ej} + \mathbf{1}[e \in O](1 - \textstyle\sum_l x_{el})\mathbf{1}[j = o_e], \qquad O = \mathrm{supp}(o).$$

Both extensions are row-affine; inactive comparison elements retain their optimal labels.

For $g_e = \nabla_e \mathcal{F}(x)$ and $a_e = \nabla_e \mathcal{F}(z(x))$, full-domain OSM and the averaged cone inequalities give $a_e \in \mathcal{A}(g_e)$. The charge bounds label replacement and insertion outside O, the latter by $\max_j a_j \geq 0$.

Proposition 4.2 (Envelope transfer)**.** *Suppose a charge H_e and an action q_e satisfy, with $h_e = q_e \cdot g_e$,*

$$a_j - q_e \cdot a \leq H_e - h_e, \qquad -q_e \cdot a \leq H_e - h_e \quad (a \in \mathcal{A}(g_e),\ j \in [k]).$$

Along $x'_e = q_e$ on a base B and zero off B, the potential $\Phi = (1-\lambda)\mathcal{F} + \lambda\mathcal{Z}$ satisfies

$$\Phi' \geq \sum_{e \in B} (h_e - \lambda H_e) + \lambda'(\mathcal{Z} - \mathcal{F}), \qquad \lambda \geq 0. \tag{14}$$

If each local score is at least $K(\lambda) \max_j g_{ej}$, $K(\lambda) \geq 0$, a maximum-weight base for these maxima and $\lambda' = -K(\lambda)$ give $\Phi' \geq 0$. For positive continuous K, $x(0) = 0$ and $\lambda(1) = 0$ imply

$$\int_0^{\lambda(0)} \frac{du}{K(u)} = 1, \qquad \mathcal{F}(x(1)) \geq \lambda(0) M. \tag{15}$$

The drift bound holds almost everywhere on absolutely continuous trajectories with these directions.

For an exact sequential block, hold λ_i fixed, greedily construct a base by current marginal maxima, and update each selected row once by δq_e, keeping all states legal. If $\lambda_i \geq \lambda_{i+1} \geq 0$ and $\lambda_i - \lambda_{i+1} \leq \delta K(\lambda_i)$, the block and endpoint weight change do not decrease Φ.

Proof. The comparison direction is $z'_e = q_e - \mathbf{e}_{o_e}$ for $e \in B \cap O$ and q_e for $e \in B \setminus O$. Thus $\mathcal{F}' = \sum_B h_e$ and $\mathcal{Z}' \geq \sum_B (h_e - H_e)$, proving (14). Appending missing comparison labels and using OSM gives

$$\mathcal{Z} - \mathcal{F} \leq \sum_{e \in O} (1 - \textstyle\sum_j x_{ej}) g_{e,o_e} \leq \sum_{e \in O} \max_j g_{ej}.$$

Nonnegative maxima give a maximum-weight base with sum at least the gap; the clock makes the drift nonnegative for either gap sign. Integrating from $\Phi(0) = \lambda(0)M$ to $\lambda(1) = 0$ proves the assertion.

For a finite block, adapt the sequential-update analysis of Badanidiyuru and Vondrák (2014) to the row-affine potential. Maintain a base of the direction prefix and unused original comparison elements. A selected element in this base matches itself; otherwise it exchanges with an unused original element. The partner was addable, so greediness bounds its current maximum by the selected maximum; OSM gives the same bound at block end. Exact row-affine increments give a fixed-weight increase of at least $\delta K(\lambda_i) \sum_{u \in O} \max_j g_{uj}(x^{\mathrm{end}})$. The endpoint change subtracts $(\lambda_i - \lambda_{i+1})(\mathcal{Z} - \mathcal{F})$. The clock inequality pays this cost if the gap is positive; otherwise the change helps. □

For $K = \mathcal{K}_r$, (15) defines α_r. At $r = 2$, Lemma B.3 gives

$$\mathcal{K}_2(\lambda) = \begin{cases} \dfrac{(1-\lambda)^2}{1+\lambda}, & 0 \le \lambda \le \sqrt{5}-2, \\ 2\sqrt{2(1-\lambda)} - 2, & \sqrt{5}-2 \le \lambda \le \frac{1}{2}. \end{cases} \tag{16}$$

Substitution in (15) gives $\alpha_2 \approx 0.4218$. In the same potential, the fixed proportional rule $g_+/\sum_i [g_i]_+$ on two-label rows yields $a_{\mathrm{prop}} \approx 0.4185$. Optimizing the action at each weight raises the envelope on $[0, \sqrt{5}-2)$ and the unit-time coefficient to α_2 (Appendix B.2).

5 Exact envelopes across label counts

Write $\mathcal{K}_r^+ = K_{r,r+1}$ and $\gamma_r = 4/(3+\sqrt{8r+1})$. For $r = 2$, let M_k denote the branch attained by one-heavy rows $(1, t, \dots, t)$, defined below. The exact off-diagonal classification is (Theorem B.10)

$$K_{r,k}(\lambda) = \begin{cases} \mathcal{K}_r^+(\lambda), & k > r \ge 3, \quad 0 \le \lambda \le \gamma_r, \\ \min\{\mathcal{K}_2^+(\lambda), M_k(\lambda)\}, & r = 2,\ k \ge 4, \quad 0 \le \lambda \le \frac{1}{2}. \end{cases} \tag{17}$$

Also $K_{2,3} = \mathcal{K}_2^+$ on $[0, \gamma_2]$. The unit-time coefficient $\beta_{2,k}$ of $K_{2,k}$ equals α_2^+ for $3 \le k \le 5$ and decreases strictly from $k = 6$ onward (Corollary B.15).

At $k = r+1$, the comparison cone becomes $\{a : \sum_i a_i \ge \max_i a_i\}$. For a nonnegative sorted observation $g_1 = 1 \ge g_2 = t \ge \cdots$ and $S = \sum_i g_i$, the exact charge is

$$H_{\mathrm{E}} = \max\{1 + (S-1)\max_{i \ge 2} q_i, \quad t + (1-t)q_1 + (S-1)\max_i q_i\}. \tag{18}$$

Ordered capacity filling and minimization over the row give (Theorem B.7)

$$\mathcal{K}_r^+(\lambda) = \begin{cases} \dfrac{(1-\lambda)^2}{1+(r-1)\lambda}, & 0 \le \lambda \le \eta_r, \\ 1 - \lambda - \dfrac{r}{\left(\sqrt{2-\lambda} + \sqrt{r(1-\lambda)}\right)^2}, & \eta_r \le \lambda \le \gamma_r, \end{cases} \tag{19}$$

where $\eta_r \in (0, \gamma_r)$ is the unique solution of $(1-\eta_r)^3 = r\eta_r^2(2-\eta_r)$. Its unit-time integral in (15) defines α_r^+.

For $k > r$, retain the top $r+1$ positive-part caps b on labels T, with smallest cap w, and choose $\max_{i \in T} q_i \ge 1/r$. In this local space, a minimizer of $q \cdot u$ can have $\max_i u_i \ge w$, giving $w - \min_{u \in \mathcal{A}(b)} q \cdot u \le D_{\mathrm{E}}(b, q)$ (Lemma B.11). Every full comparison row restricts to $\mathcal{A}(b)$ and has exterior targets at most w, so this charge protects all labels. For $r \ge 3$, the remaining capped boundary $(1, 1, c, \dots, c)$ has nonnegative slack slope $[r-2-(r-1)\lambda]/r$ in c on the clock range. Its minimum occurs where capped and unrestricted actions coincide, so the cap preserves the envelope (Theorem B.12). A matching upper bound at $g = (1, t, c, \dots, c)$ mixes comparison rows $(1, -(r-1)c, c, \dots, c)$ and $(-(r-1)c, t, c, \dots, c)$, targeting labels 1 and 2. The mixture equalizes all label payoffs (Theorem B.9).

For $r = 2$, two-target replacement risk competes with probability capacity: the lower-bound proof completes coordinatewise probability lower bounds to a distribution. Put $a = \lambda/(1-\lambda)$ and $N = k-1$. For a mixed comparison with target mass u on the best label, concavity makes equal remaining masses the tightest constraint:

$$M_k(\lambda) = (1-\lambda) \min_{0 \le u \le 1} \left\{ 1 - \frac{Na(1-u)(1+au)}{N + a(1-u)} \right\}. \tag{20}$$

The minimizer gives the one-heavy family (Appendix B.5.2). This branch stays above $\mathcal{K}_2^+$ for $k = 4, 5$, becomes active on part of the integration interval from $k = 6$, and decreases with k, producing Figure 2(b)'s plateau and decline.

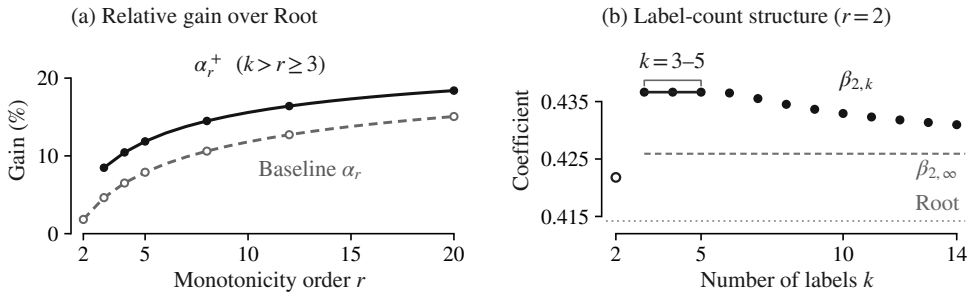

Figure 2: Approximation coefficients before accuracy loss τ. (a) Relative gain over $R(1, r)$: baseline α_r and extra-label α_r^+, shared by all $k > r$ for $r \geq 3$. Baseline gain increases with r (Theorem E.1). (b) For $r = 2$, α_2 applies at $k = 2$, α_2^+ at $k = 3$–5; $\beta_{2,k}$ decreases for $k \geq 6$ toward $\beta_{2,\infty} > 0.4259$. Points evaluate exact scalar definitions; strict bounds are certified in Corollary E.2.

6 Finite implementation and oracle complexity

Algorithm 2 uses whole-row updates and a decreasing rational clock (λ_i, δ_i). Weighting selection errors by residual caps and the envelope permits coarser estimates. Write $F_A(s) = F(A+s) - F(A)$ and $\mathcal{I}_A = \mathcal{I}/\operatorname{supp}(A)$. At $V \geq M$ and initial weight $0 \leq a \leq \gamma_r$, preprocessing returns prefix A and fixed caps $m_e = \max_j \Delta_{e,j} F(A)$. Its comparison base O_A contains A and satisfies

$$\mathbb{E}[(1-a)F(A) + aF(O_A)] \geq aM, \qquad U(A) := \max_{B \text{ base of } \mathcal{I}_A} \sum_{e \in B} m_e \leq 3V.$$

Feasible insertions reach $U(A) = 0$ at full rank, giving at most ρ iterations (Lemma C.1).

Put $E_A(B) = \{e \notin B : B + e \in \mathcal{I}_A\}$. SampledDirection estimates an element and label action. Collapse samples one label per row, preserving its mass; Compress applies it in sparse runs and is the identity otherwise. Round labels the base sampled by $\mathcal{R}$, with randomness hidden from optimization (Appendices C.3, C.4 and D).

Let $D = F(\textsc{Greedy}(F, \mathcal{I}))$, so $M/(r+1) \leq D \leq M$. The baseline uses $\mathcal{V}(D) = \{2^j D : 0 \leq j \leq \lceil \log_2(r+1) \rceil\}$ and Lemma D.5's clock, with unused duration at the largest weight. Full-row versions use $\mathcal{V}(D) = \{(r+1)D\}$ and Lemma C.2's clock with Appendix C.4's parameters.

Algorithm 2 Continuous optimization with replacement comparisons

Require: $F, \mathcal{I}, k, r, \tau$; envelope K and its finite implementation
$s_g \leftarrow \textsc{Greedy}(F, \mathcal{I})$; $D \leftarrow F(s_g)$; $Y \leftarrow s_g$
if $D = 0$ **then return 0**
$\mathcal{C} \leftarrow \textsc{Clock}(K, \tau)$
for $V \in \mathcal{V}(D)$ **do**
  $(A, m) \leftarrow \textsc{Preprocess}(F, \mathcal{I}, V)$; $x \leftarrow 0$; $\mathcal{R} \leftarrow \varnothing$
  for $(\lambda_i, \delta_i) \in \mathcal{C}$ **do**
    $B \leftarrow \varnothing$
    while $|B| < \operatorname{rank}(\mathcal{I}_A)$ **do**
      $(e, q) \leftarrow \textsc{SampledDirection}(F_A, x, E_A(B), \lambda_i; m)$
      $x_e \leftarrow x_e + \delta_i q$; $B \leftarrow B + e$
    end while
    $\mathcal{R} \leftarrow \textsc{Accumulate}(\mathcal{R}, B, \delta_i)$; $x \leftarrow \textsc{Compress}(x)$
  end for
  $x \leftarrow \textsc{Collapse}(x)$; $Z \leftarrow \textsc{Round}(A, x, \mathcal{R})$
  Choose $Y \in \arg\max_{S \in \{Y, A, Z\}} F(S)$
end for
return Y

The baseline uses a sparse action requiring only two distinguished labels.

Lemma 6.1 (Two-label protection)**.** *Let $g \in \mathcal{C}_r^k$, and let T contain r labels, two of which each dominate every coordinate outside T. Put $b = (g_T)_+$, $S = \sum_T b_i$, and let q be a distribution supported on positive coordinates of b. For $H = \max_{i \neq j \in T}(b_j + Sq_i)$ and $h = q \cdot b$, every $a \in \mathcal{C}_r^k$ with $a \leq g$ satisfies $a_z - q \cdot a \leq H - h$ for $z \in [k]$, and $-q \cdot a \leq H - h$.*

Proof. The deficits $d_i = b_i - a_i \geq 0$ on T sum to at most S by the r-coordinate inequality. For $z \in T$, $a_z - q \cdot a \leq b_z - h + S \max_{i \neq z} q_i \leq H - h$. Also $-q \cdot a \leq S \max_i q_i - h \leq H - h$, since $b \geq 0$. For $z \notin T$, choose i maximizing q_i and one of the two dominating coordinates $j \neq i$. Then $a_z \leq g_z \leq g_j \leq b_j$, so $a_z - q \cdot a \leq b_j - h + Sq_i \leq H - h$. □

Only two labels need approximate maxima; the other $r - 2$ are arbitrary. Inflating sampled caps on T restores domination, with slack loss bounded by the estimation errors.

Theorem 6.2. *Under the matroid assumptions of Section 2, for $k \geq r \geq 2$ and $0 < \tau < \alpha_r$, the baseline implementation of Algorithm 2 returns a feasible Y satisfying $\mathbb{E}F(Y) \geq (\alpha_r - \tau)M$. For $\rho \geq 1$, every execution uses*

$$Q_f = O\left(\frac{k\rho(n+r)}{\tau^3\sqrt{r}} \log(r+1) \log \frac{2r}{\tau}\right) \tag{21}$$

value queries, $Q_I = O(n\rho \log(r+1)/(\tau\sqrt{r}))$ independence queries, and $O(nr + k + 1/(\tau\sqrt{r}))$ numerical records with block-boundary compression.

Query complexity analysis. Let μ_e be the largest current mean marginal and m_e its residual cap. A tournament returns E with $\mathbb{E}[\mu_E + \epsilon_s m_E] \geq \max_e \mu_e$, charging $\mathcal{K}_r(\lambda)\epsilon_s m_E$ to the potential. Since $M/(r+1) \leq D \leq M$, the scales $V = 2^j D$ include a fixed branch with $M \leq V < 2M$ and cap sum at most $3V < 6M$. Maximizing over feasible branch outputs retains its guarantee. The clock bounds the total expected selection loss by

$$W = \sum_j \delta_j \mathcal{K}_r(\lambda_j) \leq 4\alpha_r, \qquad \mathcal{E}_{\text{sel}} \leq 6\epsilon_s MW \leq 24\epsilon_s \alpha_r M.$$

Thus $\epsilon_s = \Theta(\tau/\alpha_r)$ suffices. The clipped-mean estimator (Lemma D.1) and sampled action yield (21) over $O(1/(\tau\sqrt{r}))$ blocks and $O(\log(r+1))$ scales. Appendix D gives the error, rounding and finite-randomness bounds and counts every scale's queries.

The stronger extra-label coefficients use full-row estimates with $0 < \tau < a_{r,k}$. For $r \geq 3$ their value-query bound is $O(nk\rho r^{7/2}\tau^{-3} \log(2nkr\rho/\tau))$; for $r = 2, k \geq 4$, the fixed-label bound is $O(nk\rho\tau^{-3} \log(2nk\rho/\tau))$, with a rational local LP. Appendix C proves these finite transfers.

Proof of Theorem 2.2. For the chosen envelope K with unit-time coefficient a and $0 < \tau < a$, the clock starts at $\lambda_0 \geq a - \tau/2$. Preprocessing on the fixed valid branch gives $\mathbb{E}\Phi(0) \geq \lambda_0 M$ (Lemma C.1 and Appendix D.3). Conditional sampling and quantization incur total expected loss at most $\tau M/4$ (Appendices C.4 and D). Row collapse preserves both potential components in conditional expectation; at each pre-update history, Proposition 4.2's exact row increments pay the clock decrease. Telescoping yields $\mathbb{E}\mathcal{F}_A(x_{\text{end}}) \geq (\lambda_0 - \tau/4)M$, where $\mathcal{F}_A(x) = \mathbb{E}F(A + X_x)$ includes the prefix. After terminal collapse, hidden support merges follow the swap-rounding law conditional on the optimization history and terminal labels, without decreasing expectation (Appendix C.3). Since the retained output lies in $[0, M]$, bounded-bit coupling costs at most $\tau M/4$. Thus $\mathbb{E}F(Y) \geq (a - \tau/2 - \tau/4 - \tau/4)M = (a - \tau)M$. □

7 Conclusion

We developed Random-root and a continuous refinement for constrained non-monotone k-submodular maximization and its r-wise orthant-submodular generalization. Support budgets yield minimax randomization on p-systems, while label-specific replacement envelopes determine label distributions and comparison-weight decay under matroid constraints. The exact envelope classification reveals saturation at $k = r+1$ for $r \geq 3$, and a three-to-five-label coefficient plateau followed by a capacity-driven decline for $r = 2$. Two-label protection and envelope-weighted estimation errors yield a finite implementation of the k-submodular baseline, attaining $\alpha_2 - \tau$ for every $k \geq 2$ with $O(nk\rho)$ value queries at fixed accuracy.

REPRODUCIBILITY STATEMENT

The main text and Appendices A–E provide complete proofs, finite algorithms and coefficient certificates. The supplementary code verifies the numerical bounds in Appendix E, checks the exact comparison in Figure 1, and reproduces both figures.

AI USE STATEMENT

Generative AI tools assisted with literature discovery, research discussions, proof development and checking, notation consistency, drafting and language editing, and code for computational checks and figure preparation. The authors reviewed the AI-assisted material and take responsibility for the final content.

REFERENCES

Ashwinkumar Badanidiyuru and Jan Vondrák. Fast algorithms for maximizing submodular functions. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1497–1514. SIAM, 2014. doi: 10.1137/1.9781611973402.110.

Manuel Blum, Robert W. Floyd, Vaughan Pratt, Ronald L. Rivest, and Robert E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7(4):448–461, 1973. doi: 10.1016/S0022-0000(73)80033-9. URL https://people.csail.mit.edu/rivest/pubs/BFPRT73.pdf.

Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011. doi: 10.1137/080733991.

Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science*, pages 575–584. IEEE Computer Society, 2010. doi: 10.1109/FOCS.2010.60.

Shengminjie Chen, Xiaoming Sun, Wenguo Yang, Jialin Zhang, and Zihan Zhao. DR-submodular maximization with stochastic biased gradients: Classical and quantum gradient algorithms. In *International Conference on Learning Representations*, 2026a. URL https://openreview.net/pdf?id=0CZAimzcVr.

Siyuan Chen, Shengminjie Chen, Suixiang Gao, Zheyu Jiang, Chenhao Wang, and Wenguo Yang. Breaking the $1/3$ barrier for k-submodular maximization under matroid and knapsack constraints: A proportional top-2 randomized framework. arXiv:2609.15677v1, 2026b. URL https://arxiv.org/abs/2609.15677.

Arnav M. Das, Gantavya Bhatt, Sahil Verma, Yiping Wang, Viswa Virinchi Muppirala, and Jeff Bilmes. Matched data, better models: Target aligned data filtering with sparse autoencoders. In *International Conference on Learning Representations*, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/hash/8291acb610ddc260d554da5a3d273ae9-Abstract-Conference.html.

Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006. URL https://www.jmlr.org/papers/v7/evendar06a.html.

Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 570–579. IEEE Computer Society, 2011. doi: 10.1109/FOCS.2011.46.

Moran Feldman, Christopher Harshaw, and Amin Karbasi. Greed is good: Near-optimal submodular maximization via greedy optimization. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 758–784. PMLR, 2017. URL https://proceedings.mlr.press/v65/feldman17b.html.

Marshall L. Fisher, George L. Nemhauser, and Laurence A. Wolsey. An analysis of approximations for maximizing submodular set functions—II. In *Polyhedral Combinatorics*, volume 8 of *Mathematical Programming Studies*, pages 73–87. Springer, 1978. doi: 10.1007/BFb0121195.

Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Zebang Shen. Stochastic conditional gradient++: (Non)Convex minimization and continuous submodular maximization. *SIAM Journal on Optimization*, 30(4):3315–3344, 2020. doi: 10.1137/19M1304271.

Haifeng Huang, Qian Liu, Yang Zhou, and Min Li. The approximation algorithm and fast algorithm for constrained or-submodular maximization problem. *Journal of Combinatorial Optimization*, 51(5):46, 2026. doi: 10.1007/s10878-026-01421-8.

Anna Huber and Vladimir Kolmogorov. Towards minimizing k-submodular functions. In *Combinatorial Optimization*, volume 7422 of *Lecture Notes in Computer Science*, pages 451–462. Springer, 2012. doi: 10.1007/978-3-642-32147-4_40. URL https://arxiv.org/abs/1309.5469v1. Revised author version: arXiv:1309.5469v1 (2013).

Satoru Iwata, Shin-ichi Tanigawa, and Yuichi Yoshida. Improved approximation algorithms for k-submodular function maximization. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 404–413. SIAM, 2016. doi: 10.1137/1.9781611974331.ch30.

Ariel Kulik, Thiago Oliveira, Roy Schwartz, and Mohit Singh. If it is good then drop it—a spiteful Poisson process for submodular maximization, 2026. URL https://arxiv.org/abs/2608.05062v1. arXiv:2608.05062v1.

Julián Mestre. Greedy in approximation algorithms. In *Algorithms—ESA 2006*, volume 4168 of *Lecture Notes in Computer Science*, pages 528–539. Springer, 2006. doi: 10.1007/11841036_48.

Naoto Ohsaka and Yuichi Yoshida. Monotone k-submodular function maximization with size constraints. In *Advances in Neural Information Processing Systems 28*, pages 694–702, 2015.

Hiroki Oshima. Improved randomized algorithm for k-submodular function maximization. *SIAM Journal on Discrete Mathematics*, 35(1):1–22, 2021. doi: 10.1137/19M1277692.

Shinsaku Sakaue. On maximizing a monotone k-submodular function subject to a matroid constraint. *Discrete Optimization*, 23:105–113, 2017. doi: 10.1016/j.disopt.2017.01.003.

Justin Ward and Stanislav Živný. Maximizing k-submodular functions and beyond. *ACM Transactions on Algorithms*, 12(4):47:1–47:26, 2016. doi: 10.1145/2850419.

Longrong Yang, Dong Shen, Chaoxiang Cai, Fan Yang, Tingting Gao, Di Zhang, and Xi Li. Solving token gradient conflict in mixture-of-experts for large vision-language model. In *International Conference on Learning Representations*, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/hash/532ce4fcf853023c4cf2ac38cbc5d002-Abstract-Conference.html.

Hanlu Ye, Min Li, Yang Zhou, and Qian Liu. Maximizing a k-submodular function with p-system constraints. *Mathematical Foundations of Computing*, 8(3):448–459, 2025. doi: 10.3934/mfc.2024018.

Huanjian Zhou, Lingxiao Huang, and Baoxiang Wang. Improved approximation algorithms for k-submodular maximization via multilinear extension. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=EPHsIa0Ytg.

A Random-root proofs and implementation

A.1 Local minimax

For $g \in \mathcal{C}_r$, the maximum $m(g)$ is nonnegative. If $m(g) = 0$, all entries are nonpositive, and the exact-r inequalities force every entry to vanish. Thus a nonzero row has $m(g) > 0$. Its comparison set $\mathcal{A}(g) = \{a \in \mathcal{C}_r : a \le g\}$ contains g and is compact. Indeed, an r-subset J containing coordinate i gives $a_i \ge -\sum_{j \in J \setminus \{i\}} g_j$, while $a_i \le g_i$. For every distribution q, $q \cdot a \le q \cdot g \le m(g)$, so $D^p_{\mathrm{R}}(g, q) \ge 0$.

Prefix optimality in Theorem 3.1. Write $m = m(g) > 0$ and sort g as in (5). If q is uniform on its top $t \le r$ coordinates, adding the bottom $r - t$ coordinates to the same exact-r inequality gives

$$\sum_{j=1}^{t} a_j \ge -L_{r-t}(g), \qquad q \cdot g = G_t/t.$$

The two groups are disjoint under the fixed label order, including ties. For $G_t > 0$, the ratio in (5) is therefore an upper bound on $\Lambda_p(g, q)$.

Let λ be the minimum of these ratios. The $t = 1$ candidate is available, and every candidate is nonnegative because $G_t + L_{r-t} \ge 0$ and $G_t \le tpm$. Define the comparison row

$$b_j = \min\{g_j, pm - \lambda g_j\}. \tag{22}$$

To verify $b \in \mathcal{C}_r$, fix an r-set H and let $I \subseteq H$ be the coordinates using the second term in this minimum. Put $t = |I|$. For $t = 0$, use $g \in \mathcal{C}_r$. Otherwise $\lambda \ge 0$ gives

$$\sum_{j \in H} b_j = tpm - \lambda \sum_{j \in I} g_j + \sum_{j \in H \setminus I} g_j \ge tpm - \lambda G_t + L_{r-t}.$$

If $G_t > 0$, the right side is nonnegative by definition of λ. If $G_t \le 0$, it is at least $tpm - (\lambda+1)G_t \ge 0$ because $L_{r-t} \ge -G_t$. Hence $b \in \mathcal{A}(g)$, and every distribution satisfies $pm - q \cdot b \ge \lambda q \cdot g$. This proves the exact minimum and the optimality of the prescribed uniform prefix. The main text proves the worst-row value (6) from this characterization. □

For $r = 2$, the two candidate ratios reduce to

$$\min\left\{p + \frac{g_k}{g_1},\ \frac{2pg_1}{g_1 + g_2}\right\}, \tag{23}$$

with the second omitted when $g_1 + g_2 = 0$. For $r = 1$, only the best label is selected. A deterministic label rule has worst local coefficient $p + r - 1$: at $g = \mathbf{1}$, place $-(r-1)$ on its selected coordinate and 1 elsewhere in the hidden row. This row lies in $\mathcal{A}(g)$. The best-label rule attains the coefficient by the $t = 1$ bound. For $r > 1$, $C_{p,r} < p + r - 1$, since substituting $p + r - 1$ into $x(x + 1 - p) - pr$ gives $r(r-1) > 0$.

A.2 Deadline charging and resource bounds

Use the deadline notation in the main-text proof of Theorem 2.1, so $B_t = |D_t| = \sum_{i \le t} b_i$. The selected maximum $M_t = m(g_t)$ is the largest marginal over all elements addable to S_{t-1} and all labels.

Lemma A.1. *On every trajectory, $B_t \le pt$, and the positive maxima M_t are nonincreasing.*

Proof. The set D_t is independent, being a subset of O. In the restriction to $Y_t = S_t \cup D_t$, the set S_t is maximal independent: every element of $D_t \setminus S_t$ was blocked by an earlier support and remains blocked by downward closure. Extend D_t to a maximal independent subset R_t of Y_t. The p-system property gives $B_t \le |R_t| \le p|S_t| = pt$. The candidate family shrinks and OSM decreases each remaining label marginal along compatible extensions, so M_t is nonincreasing. The summation-by-parts argument in the main text then gives (7), including the empty trajectory through empty sums. □

At each executed step, the pre-draw history $\mathcal{H}_t$ fixes e_t, g_t, q_t, b_t and a_t. Thus

$$\mathbb{E}[F(o^{t-1}) - F(o^t) \mid \mathcal{H}_t] \le b_t M_t - q_t \cdot a_t.$$

Set the loss and gain increments to zero after stopping and sum to the fixed horizon ρ. Conditional expectation then justifies the main-text telescoping bound and $\mathbb{E}\sum_t h_t = \mathbb{E}F(s^T)$.

Resource bounds in Theorem 2.1. For $\rho \ge 1$, at most ρ rounds evaluate marginal rows: each inserts an element or terminates before full rank. Each round checks at most n candidates and uses $O(k)$ value queries per marginal row, giving $O(nk\rho)$ value queries. Feasibility checks use $O(n\rho)$ independence queries, including the final termination check. Finding row maxima takes $O(nk\rho)$ arithmetic operations, and the prefix choices take $O(k\rho)$ by Proposition A.3. The labeling and $O(k)$ marginal data require $O(n+k)$ numerical records. □

Matroid query comparison. The matroid algorithm of Chen et al. (2026b, Algorithm 1 and Theorem 4.1) has ρ iterations, each querying k label marginals for the remaining elements. This gives $O(\sum_{t=1}^{\rho} k(n-t+1)) = O(nk\rho)$ value calls. Greedy construction of a maximum-score contraction base requires $O(n)$ independence tests per iteration, hence $O(n\rho)$ in total. For $\rho \ge 1$, these are also the query orders in Theorem 2.1.

A.3 Exact prefix selection in linear time

Use decreasing marginal and increasing original label index as a total order. Its reversal lists $\ell_1 \le \cdots \le \ell_k$. Set $m = m(g) > 0$, $A_t = tpm + L_{r-t}(g)$, $\lambda_t = A_t/G_t$ when $G_t > 0$, and $d = \min\{r, |\{j : g_j > 0\}|\}$.

Lemma A.2 (Single switch)**.** *The earliest minimizer in* (5) *is in* $\{1, \dots, d\}$*. For* $1 \le t < d$*, the predicate*

$$A_t g_{t+1} \le G_t(pm - \ell_{r-t}) \tag{24}$$

changes at most once from false to true. Its first true index is the earliest minimizer; if it is always false, that minimizer is d*.*

Proof. The exact-r inequality gives $G_t + L_{r-t} \ge 0$, so $A_t \ge 0$. Also $A_{t+1} - A_t = pm - \ell_{r-t} \ge 0$. If $g_t \le 0$ and $G_t > 0$, then $G_{t-1} \ge G_t > 0$ and $A_{t-1} \le A_t$, so $\lambda_{t-1} \le \lambda_t$. Such a prefix cannot be the earliest minimizer.

For $t < d$, put $D_{t+1} = (pm - \ell_{r-t})/g_{t+1}$. These values are nondecreasing: their nonnegative numerators increase and their positive denominators decrease. The identity

$$\lambda_{t+1} = \frac{G_t\lambda_t + g_{t+1}D_{t+1}}{G_t + g_{t+1}}$$

shows that (24) is equivalent to $\lambda_t \le D_{t+1}$ and to $\lambda_t \le \lambda_{t+1}$. Once it holds and $t+1 < d$, $\lambda_{t+1} \le D_{t+1} \le D_{t+2}$, so the next predicate holds as well. Thus the ratios strictly decrease before the first true predicate and are nondecreasing afterwards. □

Proposition A.3 (Linear prefix implementation)**.** *The prefix in* (5)*, its gain and insertion-loss bound can be computed using* $O(k)$ *exact operations and* $O(k)$ *scratch records. Given a uniform rank in that prefix, its canonical label can be recovered within the same bounds. Under common uniform-rank draws, the selected label equals that of the sorted implementation.*

Proof. Scan for m and d. If $d = 1$, select the best label and use linear selection to obtain L_{r-1}. Otherwise, locate the first true predicate in (24), taking d if none is true.

For an answer interval $[a, b] \subseteq [1, d]$, only descending ranks $a+1, \dots, b$ and ascending ranks $r-b+1, \dots, r-a$ can be queried. Retain these two unsorted windows together with their preceding sums G_a and L_{r-b}. At $[a, b] = [1, d]$, deterministic selection (Blum et al., 1973) isolates both windows and their preceding sums m, L_{r-d} in $O(k)$ operations.

At the midpoint $t = \lfloor (a+b)/2 \rfloor$, linear selection within the windows gives g_{t+1} and ℓ_{r-t}. Add the preceding sums to the window sums before g_{t+1} and through ℓ_{r-t}, respectively, to recover G_t and L_{r-t}. A true predicate replaces b by t; a false one replaces a by $t+1$. The new windows are

subsets of the old ones, and adding discarded preceding entries to the sums preserves the invariant. Lemma A.2 therefore identifies the earliest minimizing index.

Each window has $b - a$ entries, and this width at least halves per step, so selection and summation cost $O(d)$ in total. A final partition of the original row recovers the top t labels, G_t and L_{r-t}, giving gain G_t/t and insertion-loss bound L_{r-t}/t. Selecting a supplied uniform rank in this prefix, under the fixed total order, returns the same label as the sorted algorithm. Initialization and final selection cost $O(k)$, and the windows use $O(k)$ records. □

For rational $p = P/D$, prefix comparisons use cross multiplication of $[tPm + DL_{r-t}]/[DG_t]$, with $G_t > 0$. Thus the algorithm uses exact arithmetic and comparisons; the square root defining $C_{p,r}$ enters only its guarantee.

A.4 Bounded random bits

For any $0 < \epsilon < 1$, a dyadic approximation to each uniform prefix draw gives coefficient $R(p, r) - \epsilon$ with the same oracle bounds. The empty ground set returns $\mathbf{0}$; for $r = 1$ every prefix is a singleton and needs no random bits. For $n \geq 1$ and $r \geq 2$, apply the dyadic sampling argument in Appendix C.3 with $Q = n$ and $s = r$. By Lemma A.2, the chosen prefix contains only positive marginals. The dyadic draws stay within that prefix, so both procedures satisfy $0 \leq F(Y) \leq M$ on every trajectory. The expected loss is therefore at most ϵM, using at most

$$n \left\lceil \log_2 \frac{n(r-1)}{\epsilon} \right\rceil$$

fair bits; singleton prefixes need none.

B Replacement comparisons and exact label envelopes

The common comparison domain and positive-support reduction lead to the diagonal baseline and the exact envelope with one additional label. Matching upper bounds and exterior-label protection then establish saturation and the pairwise fixed-label classification. All minimizations below permit signed observations and comparison rows.

Use the model notation of Section 2 and the losses in (4) and (8). Fix a nonzero row g, put $m = \max g$ and $h = q \cdot g$, and write $H_\chi = h + D_\chi$ for $\chi \in \{\mathrm{R}, \mathrm{E}\}$, with p fixed in the root case. Taking the comparison row $a = g$ gives $H_\mathrm{E} \geq m > 0$ and $H^p_\mathrm{R} \geq pm$. The comparison set is compact by Appendix A.1, so these charges are continuous in q. Consequently

$$\theta_\chi := \max_{q \in \Delta_k} \frac{h}{H_\chi} = \frac{1}{1 + \gamma_\chi}, \qquad \gamma_\chi := \min_{q : h > 0} \frac{D_\chi}{h}.$$

A best label gives positive gain, so the maximum is positive and attained at $h > 0$, which also proves attainment of the minimum. Factoring $h - \lambda H_\chi = H_\chi(h/H_\chi - \lambda)$ yields

$$\operatorname{sign} \max_q \{h - \lambda H_\chi\} = \operatorname{sign}(\theta_\chi - \lambda).$$

Since $D^p_\mathrm{R} \geq D^1_\mathrm{R}$, (8) makes the optimal replacement slack at least the optimal Root slack for $p \geq 1$ and $\lambda \geq 0$. Every comparison row has a nonnegative maximum, so $D_\mathrm{E}(g, q) \geq -q \cdot a$ for all $a \in \mathcal{A}(g)$. The zero row contributes zero gain and loss.

B.1 Local values and normalized envelopes

Proposition 4.1 gives the exact diagonal charges and proves their difference in the main text. For general $p \geq 1$, $H^p_\mathrm{R} - H_\mathrm{E} = (H^1_\mathrm{R} - H_\mathrm{E}) + (p-1)m$. We now optimize the replacement charge and determine its worst observed row.

Write the local value as

$$V_{r,k}(g, \lambda) = \max_{q \in \Delta_k} \{q \cdot g - \lambda H_\mathrm{E}(g, q)\}. \tag{25}$$

Its normalized worst-row value is $K_{r,k}$ from (13).

Lemma B.1 (Positive support and nonnegative observations)**.** *For nonzero $g \in \mathcal{C}_r^k$ and $0 \le \lambda < 1$, every optimal action in* (25) *is supported on the positive coordinates of g, and*

$$V_{r,k}(g, \lambda) \ge V_{r,k}(g_+, \lambda).$$

The infima over normalized signed and nonnegative observations coincide. These conclusions also hold with $q_1 \ge 1/r$ at a best coordinate.

Proof. Let $I = \{i : g_i > 0\}$. Summing all r-subset inequalities gives $\sum_i a_i \ge 0$ for every $a \in \mathcal{A}(g)$. Since $a_j \le g_j \le 0$ outside I, also $\sum_I a_i \ge 0$. Transfer the mass $\epsilon = \sum_{j \notin I} q_j$ uniformly to I, obtaining q'. Then

$$(q' - q) \cdot a = \frac{\epsilon}{|I|} \sum_{i \in I} a_i - \sum_{j \notin I} q_j a_j \ge 0.$$

If $\epsilon > 0$, observed gain strictly increases while replacement loss does not increase. Thus $(1-\lambda)q \cdot g - \lambda D_{\mathrm{E}}(g, q)$ strictly improves. For $b = g_+$, take a positive-support optimal action. Its gains on b and g agree, and $\mathcal{A}(g) \subseteq \mathcal{A}(b)$, which proves the value comparison. Clipping preserves the maximum coordinate, and nonnegative observations are admissible signed ones. The transfer only increases the probability of a best coordinate. □

Sequential filling at fixed charge. For observations $b_1 \ge \cdots \ge b_d$ and probability bounds $0 \le q_i \le c_i(H)$, $\sum_i q_i = 1$, assume the capacities are nonnegative and sum to at least one. Moving mass from a later coordinate to an earlier unsaturated one cannot decrease $q \cdot b$. Hence a maximizer fills the capacities in order, with at most one partly filled coordinate. If the capacities are affine in H, the resulting value $q \cdot b - \lambda H$ is affine between the breakpoints $\sum_{i=1}^{t} c_i(H) = 1$; its maximum on a compact feasible interval is attained at a breakpoint or an endpoint. The filled vector has $q_1 = \min\{1, c_1(H)\}$, so it also meets a lower bound $q_1 \ge u$, $0 \le u \le 1$, exactly when $c_1(H) \ge u$.

B.2 The baseline envelope and the diagonal case

For $r \ge 2$, define

$$\beta_r = \frac{1}{1 + \sqrt{r/2}}, \qquad \nu_r = \frac{1}{2 + \sqrt{2r+1}}, \qquad D_r(x) = 2 + (r-2)x.$$

The baseline envelope is

$$\mathcal{K}_r(x) = \begin{cases} (1-x)^2/[1 + (r-1)x], & 0 \le x \le \nu_r, \\ \dfrac{2}{r}\left[\sqrt{2(1-x)(r + D_r(x))} - D_r(x)\right], & \nu_r \le x \le \beta_r. \end{cases} \tag{26}$$

It gives a sufficient local bound for every $k \ge r$ and the exact replacement envelope when $k = r$. Its unit-time coefficient α_r is defined by (15).

Fix a full row $g \in \mathcal{C}_r^k$ and let T be its current r largest coordinates, with deterministic ties. Write $b_i = [g_i]_+$ for $i \in T$, $S = \sum_T b_i$ and $h = q \cdot g = q \cdot b$ for a distribution supported on positive coordinates in T. Define

$$H(q, b) = \max_{i,j \in T,\ i \ne j} \{b_j + S q_i\}.$$

Lemma B.2. *For every full comparison row $a \le g$, $a \in \mathcal{C}_r^k$, and every comparison label $j \in [k]$, including $j \notin T$,*

$$a_j - q \cdot a \le H(q, b) - h, \qquad -q \cdot a \le H(q, b) - h.$$

Proof. Apply Lemma 6.1 to the top-r set T: every selected coordinate dominates every unselected one, so the required two coordinates exist. □

The top-r set is chosen separately for each element and state.

The shifted-prefix oracle. For $0 \le \lambda < 1$, the zero row has value zero. Otherwise sort the positive entries of b as $m = b_1 \ge b_2 = v \ge \cdots \ge b_d > 0$; let $v = 0$ if $d = 1$ and $W = S - m + v$. For each $1 \le t \le d$ set

$$q_1^{(t)} = \frac{m - v}{S} + \frac{W}{tS}, \qquad q_i^{(t)} = \frac{W}{tS}\ (2 \le i \le t), \qquad H_t = m + \frac{W}{t}.$$

All remaining probabilities are zero. Choose t maximizing $q^{(t)} \cdot b - \lambda H_t$. Writing $\delta = (m - v)/S$, these candidates are $q^{(t)} = \delta \mathbf{e}_1 + (1 - \delta)u^{(t)}$, where $u^{(t)}$ is uniform on the first t labels. For $t \ge 2$, the action equalizes the two terms in (10) and obtains the full charge reduction $m - v$. At $r = 2$ the candidates are the best label and the proportional top-two action. Sorting and prefix sums cost $O(k \log k)$.

On this r-coordinate row, (10) identifies H with the replacement charge, so Lemma B.1 restricts the optimum to positive coordinates. At fixed charge bound H, their capacities are $(H - v)/S$ for the first and $(H - m)/S$ for the others. The search interval is $H_d \le H \le H_1$: its lower endpoint is where all d capacities sum to one, and the upper endpoint admits the best label. Sequential filling gives precisely the breakpoints H_t and the displayed actions. This also covers $d = 1$, when $W = 0$ and both endpoints equal m.

Worst-row calculation.

Lemma B.3. *For every nonnegative $b \in \mathbb{R}^r$ and $0 \le \lambda \le \beta_r$, the prefix oracle satisfies $h - \lambda H \ge \mathcal{K}_r(\lambda) \max_i b_i$. This coefficient is exact for the sufficient charge H.*

Write $V_r(b, \lambda) = \max_{q \in \Delta_r} \{q \cdot b - \lambda H(q, b)\}$ for a nonnegative normalized row b. Continuity and compactness ensure that its minimum exists. In this proof set $D = 2 + (r-2)\lambda$ and $L = 1 + (r-1)\lambda$. The first-branch equality row in the envelope below is $(1, c, \ldots, c)$ with $c = (1 - \lambda)/L$ for $0 \le \lambda \le \nu_r$. On the second branch $\nu_r \le \lambda \le \beta_r$, it is $(1, t, c, \ldots, c)$ with

$$1 + t = D\sqrt{\frac{2}{(1 - \lambda)(r + D)}}, \qquad c = \frac{(1 - \lambda)(1 + t)}{D}.$$

When $r = 2$, the second row has only its first two coordinates.

Proof. We first consider $r \ge 3$. Sort b as $1 = b_1 \ge b_2 = t \ge b_3 \ge \cdots \ge b_r \ge 0$, and write $P_\ell = \sum_{i=1}^{\ell} b_i$. The prefix oracle above gives

$$V_r(b, \lambda) = \max_{1 \le \ell \le r} W_\ell, \qquad W_\ell = \frac{1 - t}{S} - \lambda + \frac{S - 1 + t}{\ell S}(P_\ell - \lambda S).$$

Zero coordinates may be included in this list: they add feasible prefix distributions, and the preceding oracle argument shows that they cannot improve on the optimum supported on positive coordinates. The singleton row is also covered, since $S = 1, t = 0$ makes the last term vanish for every ℓ.

Fix S and t. Replacing $b_3, \ldots, b_r$ by their common average $c = \frac{S - 1 - t}{r - 2}$ keeps $0 \le c \le t$. Since the average of any initial segment of a sorted list is at least the average of the entire list, this replacement weakly decreases every P_ℓ for $\ell \ge 3$. The coefficient of P_ℓ is nonnegative. Thus it simultaneously weakly decreases all W_ℓ and does not change W_1, W_2. It is enough to minimize over rows $(1, t, c, \ldots, c)$ with $0 \le c \le t \le 1$.

Put $d = r - 2$, $S = 1 + t + dc$, and $\Delta = 1 + t - 2c - \lambda S$. For $\ell \ge 2$ one has $P_\ell = \ell c + 1 + t - 2c$, so

$$W_\ell = \frac{1 - t}{S} - \lambda + \frac{S - 1 + t}{S} c + \frac{S - 1 + t}{\ell S}\Delta.$$

Consequently only W_1, W_2, W_r need be retained. Explicitly,

$$W_1 = 1 - \lambda(S + t),$$
$$W_2 = \frac{1 + t}{2} + \frac{(1 - t)^2}{2S} - \frac{\lambda}{2}(S + 1 + t),$$
$$W_r = (1 - \lambda)\frac{S - 1 + t}{r} + \frac{1 - t}{S} - \lambda.$$

We next show that a minimizing row can always be chosen with $\Delta = 0$.

If $\Delta \geq 0$, then $W_2 \geq W_r$, and both W_1, W_2 are nonincreasing as c increases with t fixed. For W_2 this follows from $\partial_S W_2 = -(1-t)^2/(2S^2) - \lambda/2 \leq 0$. Increase c until $\Delta = 0$ or $c = t$. If the latter occurs with $\Delta > 0$, the row is one-heavy. On that boundary,

$$\Delta = 1 - t - \lambda[1 + (r-1)t], \qquad W_1 - W_2 = \frac{S - 1 + t}{2S}\Delta \geq 0.$$

Its value is $W_1 = 1 - \lambda(1 + rt)$, which is nonincreasing in t. Increasing $t = c$ to $(1-\lambda)/L \leq 1$ reaches $\Delta = 0$ without increasing the value. These weak inequalities also cover $\lambda = 0$.

If $\Delta \leq 0$, then $W_r \geq W_2$. Moreover $\lambda S \geq 1 + t - 2c \geq 1 - t$, so $W_1 - W_2 = \frac{S-1+t}{2S}(1 - t - \lambda S) \leq 0$. The value is therefore W_r. At fixed S, it is nonincreasing in t, because $\partial_t W_r = (1-\lambda)/r - 1/S \leq 0$. In terms of S, t, the condition $\Delta \leq 0$ is $t \leq DS/r - 1$. Increase t to $\min\{1, DS/r - 1\}$, changing $c = (S - 1 - t)/d$ accordingly. Feasibility is preserved: increasing t preserves $c \leq t$, and $D \leq r$ implies $DS/r - 1 \leq S - 1$, so c stays nonnegative. If the first endpoint reached is $DS/r - 1$, then $\Delta = 0$. Otherwise $t = 1$ and $S \geq 2r/D$. Along $t = 1$ the value is $(1-\lambda)S/r - \lambda$, so decreasing S to $2r/D$ does not increase it. This endpoint lies in $[2, r]$, since $2 \leq D \leq r$, and it has $\Delta = 0$. Thus all cases reduce to that surface.

On $\Delta = 0$, put $z = 1 + t$. Then

$$c = \frac{(1-\lambda)z}{D}, \qquad S = \frac{rz}{D}, \qquad \frac{D}{L} \leq z \leq 2.$$

The lower bound is exactly the constraint $c \leq t$. Here $W_2 = W_r \geq W_1$, and the value becomes

$$Az + \frac{B}{z} - B, \qquad A = \frac{(1-\lambda)(r+D)}{rD} > 0, \quad B = \frac{2D}{r} > 0.$$

Its unconstrained minimum is at $z_0 = \sqrt{B/A} = D\sqrt{2/[(1-\lambda)(r+D)]}$. This is increasing in λ, and $z_0(\beta_r) = 2$, so the upper endpoint does not bind before β_r. The lower endpoint D/L is decreasing. Their equality is equivalent to

$$2L^2 = (1-\lambda)(r+D), \quad \text{or} \quad (2r-3)\lambda^2 + 4\lambda - 1 = 0,$$

whose positive root is ν_r. The minimum is consequently at $z = D/L$ in the first branch and at $z = z_0$ in the second. Substitution gives the two stated formulas and attaining rows, which satisfy all ordering constraints.

When $r = 2$, write the row as $(1, t)$, $0 \leq t \leq 1$. The prefix oracle has only two candidates, with values

$$W_1 = 1 - \lambda(1 + 2t), \qquad W_2 = \frac{1 + t^2}{1 + t} - \lambda(1 + t).$$

Their difference has the sign of $1 - t - \lambda(1 + t)$ (and vanishes also at $t = 0$). Consequently the first candidate dominates up to $t_c = (1-\lambda)/(1+\lambda)$, where its value is nonincreasing in t. It suffices to minimize the second candidate for $t \geq t_c$. With $z = 1 + t$, this is

$$\min_{2/(1+\lambda) \leq z \leq 2} \left[(1-\lambda)z + \frac{2}{z} - 2 \right].$$

This is the preceding scalar minimization with $D = 2$, $L = 1 + \lambda$, $A = 1 - \lambda$ and $B = 2$. Its endpoint–stationary-point comparison therefore gives the same two branches, switching at $\nu_2 = \sqrt{5} - 2$ and ending at $\beta_2 = 1/2$, with the stated attaining two-coordinate rows. □

Fixed proportional actions. For $k = r = 2$, fix the action to $q^{\mathrm{prop}} = g_+ / \sum_i [g_i]_+$ on every nonzero row. On a normalized nonnegative row $g = (1, t)$, $0 \leq t \leq 1$, its exact charge is $H_{\mathrm{E}} = 1 + t$ and its slack is W_2 above. Minimizing over all $t \in [0, 1]$ gives

$$K_{\mathrm{prop}}(\lambda) = 2\sqrt{2(1-\lambda)} - 2, \qquad 0 \leq \lambda \leq \tfrac{1}{2}, \tag{27}$$

attained at $t = \sqrt{2/(1-\lambda)} - 1 \in [\sqrt{2} - 1, 1]$. For a signed normalized row $(1, t)$ with $-1 \leq t < 0$, clipping selects $(1, 0)$ and (10) gives $H_{\mathrm{E}} = 1$. Its slack $1 - \lambda$ cannot lower this envelope; the zero row

has zero slack. Thus (27) also covers all signed rows. Using the same potential in Proposition 4.2, the unit-time coefficient a_{prop} is the unique solution of $T_{\mathrm{prop}}(a_{\mathrm{prop}}) = 1$, where, for $0 \le x < 1/2$,

$$T_{\mathrm{prop}}(x) = \int_0^x \frac{du}{K_{\mathrm{prop}}(u)} = \frac{1}{2}\left[\sqrt{2} - w + \log\frac{\sqrt{2}-1}{w-1}\right], \qquad w = \sqrt{2(1-x)}.$$

This clock increases from zero to infinity, with $0.41849 < a_{\mathrm{prop}} < 0.41850$. The optimized envelope $\mathcal{K}_2$ is strictly larger than K_{prop} for $0 \le \lambda < \nu_2$ and agrees with it thereafter. Consequently $a_{\mathrm{prop}} < \alpha_2 \approx 0.4218$.

Corollary B.4 (Diagonal exactness)**.** *For $0 \le \lambda \le \beta_r$, $K_{r,r}(\lambda) = \mathcal{K}_r(\lambda)$.*

Proof. Lemma B.1 reduces the lower bound to nonnegative observations. On the diagonal, (10) identifies H with the exact replacement charge, so Lemma B.3 gives the bound and its equality rows attain it. □

B.3 One additional label

When $k = r + 1$, the cone is described by its total sum and largest coordinate. This gives an exact signed charge and a different prefix rule.

Theorem B.5 (Exact charge)**.** *Let $k = r + 1$ and sort a nonzero normalized observation as $1 = g_1 \ge g_2 = t \ge \cdots \ge g_{r+1}$. Put $S = \sum_i g_i$, $Q = \max_i q_i$, and $Q_- = \max_{i \ge 2} q_i$. Then $S \ge 1$, $t \ge 0$, and*

$$H_{\mathrm{E}}(g, q) = \max\{1 + (S-1)Q_-,\ t + (1-t)q_1 + (S-1)Q\}. \tag{28}$$

Both terms are attained by feasible comparison rows.

Proof. For $k = r + 1$, the cone condition is equivalent to

$$\sum_i a_i \ge \max_i a_i. \tag{29}$$

Applying it to g gives $S \ge 1$; its r coordinates outside the first have nonnegative sum, so $t \ge 0$.

Suppose a comparison maximum exceeds t. Only its first coordinate can attain that maximum. Raising this coordinate to 1 preserves (29) and increases the target loss at rate $1 - q_1 \ge 0$. The other coordinates have nonnegative sum and caps g_i. Their total deficit from these caps is at most $S - 1$. Assigning this entire deficit to an outside coordinate with probability Q_- gives total sum and maximum both equal to 1. Its charge is $1 + (S-1)Q_-$, which also bounds every comparison of this type. This construction remains feasible when $t = 1$, although the case with maximum greater than t is then empty.

Now consider a comparison maximum $v \le t$. Necessarily $v \ge 0$. The coordinate caps become $\min(g_i, v)$, and at least two caps equal v. Write $D(v) = \sum_i \min(g_i, v) - v$. Feasibility of this maximum requires $D(v) \ge 0$. Conversely, subtracting $D(v)$ from one largest-probability coordinate leaves another coordinate at v, and gives a feasible comparison row with sum and maximum v. The maximum charge at this height is therefore $q \cdot g + v - q \cdot \min(g, v) + Q\, D(v)$. On an interval with h coordinates strictly above v, its derivative is $1 - \sum_{g_i > v} q_i + (h-1)Q \ge 0$. Consequently its largest value is at $v = t$, which is feasible because $D(t) = S - 1 \ge 0$. Substitution gives the second term of (28). The two cases exhaust every comparison maximum. □

Proposition B.6 (Exact signed prefix oracle)**.** *Under the hypotheses of Theorem B.5, for $0 \le \lambda < 1$ an optimum in* (25) *uses positive observed coordinates only and can be chosen with $q_1 \ge q_2 \ge \cdots \ge q_{r+1}$. If $S = 1$, the best-label action is optimal and $V_{r,r+1}(g, \lambda) = 1 - \lambda$. If $S > 1$, let d_+ denote the number of positive coordinates. An optimum is among the following d_+ candidates, indexed by $1 \le \ell \le d_+$:*

$$D_\ell = (S-1) + (\ell - 1)(S - t), \tag{30}$$

$$q_1^{(\ell)} = \frac{S - 1 + (\ell - 1)(1 - t)}{D_\ell}, \qquad q_i^{(\ell)} = \frac{S-1}{D_\ell} \quad (2 \le i \le \ell), \qquad q_i^{(\ell)} = 0 \quad (i > \ell), \tag{31}$$

$$H_\ell = 1 + \frac{(S-1)^2}{D_\ell}. \tag{32}$$

Choose the candidate maximizing $q^{(\ell)} \cdot g - \lambda H_\ell$. Sorting followed by one prefix-sum scan takes $O(k \log k)$ arithmetic operations. The sum S in these expressions is the actual signed sum.

Proof. Lemma B.1 restricts optimal actions to positive coordinates. If $q_1 < Q_-$, the first term in (28) dominates the second: their difference is $(1-t)(1-q_1) \geq 0$. Transfer probability from outside maximizers to coordinate 1 until $q_1 \geq Q_-$. The gain cannot decrease and the charge cannot increase. Sorting the remaining probabilities in the observation order then increases gain without changing either charge term. Thus an optimum can be chosen sorted, in which case

$$H_{\mathrm{E}}(g, q) = \max\{1 + (S-1)q_2,\ t + (S-t)q_1\}. \tag{33}$$

When $S = 1$ this charge equals 1 for every probability vector, proving the asserted best-label rule. This includes $t = 0$, because $S \geq 1$ and $g_2 \leq 0$ then force every coordinate outside the first to be zero.

Assume $S > 1$. All charges in (33) are at most S. For a prescribed bound $1 \leq H \leq S$, the positive-coordinate capacities are

$$q_1 \leq \frac{H-t}{S-t}, \qquad q_i \leq \frac{H-1}{S-1} \quad (i \geq 2). \tag{34}$$

The first capacity is at least every other capacity, since their difference is $(1-t)(S-H)/[(S-t)(S-1)]$. Both denominators are positive. Sequential filling therefore yields a sorted optimizer on the feasible interval $H_{d_+} \leq H \leq S$. Solving the breakpoint equations gives exactly (30)–(32); they include both endpoints, with $H_1 = S$. □

The worst observed row. Let $\eta_r \in (0,1)$ solve

$$(1-\eta_r)^3 = r\eta_r^2(2-\eta_r), \tag{35}$$

and put $\gamma_r = 4/(3+\sqrt{8r+1})$.

Theorem B.7 (Exact envelope for $k = r+1$)**.** *For every integer $r \geq 2$ and $0 \leq \lambda \leq \gamma_r$,*

$$\min_{\substack{g \in \mathcal{C}_r^{r+1} \\ \max_i g_i = 1}} V_{r,r+1}(g,\lambda) = \mathcal{K}_r^+(\lambda) = \begin{cases} \dfrac{(1-\lambda)^2}{1+(r-1)\lambda}, & 0 \leq \lambda \leq \eta_r, \\ 1-\lambda-\dfrac{r}{(\sqrt{2-\lambda}+\sqrt{r(1-\lambda)})^2}, & \eta_r \leq \lambda \leq \gamma_r. \end{cases} \tag{36}$$

The switch satisfies $0 < \eta_r < \gamma_r < 1$, and the two expressions agree there. Every minimum is attained by a nonnegative row $(1, t, c, \ldots, c)$ with $r-1$ copies of c.

Proof. Lemma B.1 reduces the minimization to nonnegative rows. The singleton row has value $1-\lambda$ and will also be covered by the boundary argument below. For any other row, hold S and t fixed and replace coordinates $3, \ldots, r+1$ by their common average c. This preserves each candidate charge and probability in (31)–(32), while weakly decreasing every prefix sum of observed values. It therefore weakly decreases the optimal value. We may minimize over $(1, t, c, \ldots, c)$ with $0 \leq c \leq t \leq 1$.

Put $s = t + (r-1)c$, so $S = 1+s$ and $s/r \leq t \leq \min\{s, 1\}$, with $0 \leq s \leq r$. For prefix lengths at least 2, the candidate value is a ratio of affine functions of the length, with a positive denominator. Its maximum is at an endpoint. Thus only lengths $1, 2, r+1$ are needed, with values

$$\begin{aligned} A &= 1-\lambda(1+s), \\ B &= \frac{1+s-t+st-\lambda(1+2s-t+s^2)}{1+2s-t}, \\ C &= \frac{s(1+s)+r(1-t)-\lambda[s+r(1+s-t)+s^2]}{s+r(1+s-t)}. \end{aligned} \tag{37}$$

These expressions also equal $1-\lambda$ at $s = t = 0$. For $s > 0$, differentiation at fixed s gives

$$\partial_t B = \frac{s^2(2-\lambda)}{(1+2s-t)^2} > 0, \qquad \partial_t C = -\frac{rs[r-(1-\lambda)s]}{[s+r(1+s-t)]^2} \leq 0. \tag{38}$$

The case $\lambda = 0$ has value 1 for every normalized observation; henceforth suppose $0 < \lambda \le \gamma_r$.

Define $c_0 = (1-\lambda)/[1+(r-1)\lambda]$, $s_0 = rc_0$ and $s_1 = 2r/[2+(r-1)\lambda]$. If $s \le s_0$, the candidate A is at least $(1-\lambda)^2/[1+(r-1)\lambda]$. For $s \ge s_0$, direct subtraction gives

$$B - A = \frac{s[\lambda s + \lambda + (1-\lambda)t - 1]}{1 + 2s - t} \ge 0.$$

At $t = s/r$, the sign of $B - C$ is the opposite of $[1+(r-1)\lambda]s - r(1-\lambda)$. At the upper endpoint it is positive when $s < 1$, whereas at $t = 1$, $B - C = \{2r - [2+(r-1)\lambda]s\}/[2(r+1)]$. It follows from (38) that for $s_0 \le s \le s_1$ the minimum over t occurs at $B = C$. For $s > s_1$, it occurs at $t = 1$; there $C(s,1) = [(1-\lambda)s + 1 - (r+1)\lambda]/(r+1)$ is increasing in s. This last region cannot improve on $s = s_1$.

For the crossing calculation, set $a = 2-\lambda$, $b = (r-1)\lambda - (r-2)$ and $c_1 = (2-\lambda)/[2+(r-1)\lambda]$. Solving $B = C$ with $s = t + (r-1)c$ gives

$$t(c) = \frac{c\{(r-1)[(r-1)\lambda + 2]c + b\}}{a - bc}, \qquad c_0 \le c \le c_1. \tag{39}$$

This is an exact parameterization of the feasible crossing. To check its domain explicitly, $a - bc \ge \min\{a, a-b\} > 0$ for $0 \le c \le 1$, because $a - b = r(1-\lambda) > 0$, and

$$\begin{aligned} t - c &= \frac{rc\{[1+(r-1)\lambda]c - (1-\lambda)\}}{a - bc}, \\ t - 1 &= \frac{[2+(r-1)\lambda][(r-1)c + 1](c - c_1)}{a - bc}. \end{aligned} \tag{40}$$

Thus $c \le t \le 1$ is equivalent to $c_0 \le c \le c_1$ on the positive crossing. Conversely, these inequalities imply $s \ge rc_0 = s_0$ and $s \le 1 + (r-1)c_1 = s_1$. The endpoints are feasible, and $c_1 - c_0 = r\lambda/([2+(r-1)\lambda][1+(r-1)\lambda]) > 0$.

Since $a - b = r(1-\lambda) > 0$, the common value on this curve is

$$F_r(c,\lambda) = 1 - \lambda - \frac{r}{a/c + (a-b)/(1-c)}. \tag{41}$$

The denominator is strictly convex and, by Cauchy–Schwarz, is at least $(\sqrt{a} + \sqrt{a-b})^2$, with equality uniquely at

$$c_* = \frac{\sqrt{a}}{\sqrt{a} + \sqrt{a-b}}. \tag{42}$$

The condition $c_* \le c_1$ is equivalent to $r\lambda^2 \le (2-\lambda)(1-\lambda)$, hence to $\lambda \le \gamma_r$; at γ_r, $c_* = c_1 > c_0$. Also c_* strictly increases with λ, while c_0 strictly decreases from 1. They therefore cross once in $(0, \gamma_r)$, at the root η_r of $P_r(\lambda) := (1-\lambda)^3 - r\lambda^2(2-\lambda)$. The constrained minimum is at c_0 before this crossing and at c_* afterward.

Substitution at c_0 gives the first branch in (36), also attaining the lower bound for $s \le s_0$. The equality case of Cauchy–Schwarz gives the second branch. The same minimizing parameter occurs at the switch, proving continuity and agreement there. The constructed nonnegative rows attain the minima, and Lemma B.1 extends the lower bound to every signed observation. □

The integrated coefficient.

Corollary B.8. *The function $\mathcal{K}_r^+$ is continuous, strictly decreasing from 1 to 0 on $[0, \gamma_r]$, and satisfies*

$$-(r+1)(y-x) \le \mathcal{K}_r^+(y) - \mathcal{K}_r^+(x) \le -(y-x) \qquad (0 \le x \le y \le \gamma_r). \tag{43}$$

There is a unique $\alpha_r^+ \in (0, \gamma_r)$ such that

$$\int_0^{\alpha_r^+} \frac{d\lambda}{\mathcal{K}_r^+(\lambda)} = 1, \tag{44}$$

and $\alpha_r^+ > \alpha_r$ for every integer $r \ge 2$.

Proof. For normalized nonnegative observations, $1 \le H_{\mathrm{E}}(g,q) \le S \le r+1$. Every fixed-action payoff therefore satisfies the two increment inequalities in (43); taking the maximum over actions and then the minimum over observations preserves them. Formula (36) gives $\mathcal{K}_r^+(0) = 1$. On its second branch, the equation $\mathcal{K}_r^+(\lambda) = 0$ is equivalent to $(r-1)\lambda^2 + 3\lambda - 2 = 0$ in $(0,1)$, giving the endpoint γ_r. The Lipschitz bound implies $\mathcal{K}_r^+(\lambda) \le (r+1)(\gamma_r - \lambda)$, so its reciprocal integral diverges at that endpoint. Strict monotonicity of the integral proves existence and uniqueness of α_r^+.

The baseline positive interval lies inside the adjacent-label interval: substituting $(r-2)\beta_r^2 + 4\beta_r - 2 = 0$ gives $(r-1)\beta_r^2 + 3\beta_r - 2 = \beta_r(\beta_r - 1) < 0$, so $\beta_r < \gamma_r$. Lemma B.2 and the baseline prefix envelope give $\mathcal{K}_r^+ \ge \mathcal{K}_r$ on $[0, \beta_r]$. The baseline switch satisfies $P_r(\nu_r) = \nu_r(\nu_r - 3)(\nu_r - 1)/2 > 0$, where $(2r-3)\nu_r^2 + 4\nu_r - 1 = 0$ was used to eliminate r. Thus $\eta_r > \nu_r$. In the baseline envelope proof, for $\nu_r < \lambda < \beta_r$ the strictly convex expression $Az + B/z - B$ is minimized strictly inside its feasible interval, so its value is strictly below the one-heavy boundary value $(1-\lambda)^2/[1+(r-1)\lambda]$. Consequently $\mathcal{K}_r^+ > \mathcal{K}_r$ on $(\nu_r, \min\{\eta_r, \beta_r\})$. Also $\alpha_r > \nu_r$ by (74) and $\int_0^{\nu_r} 1/\mathcal{K}_r < 1/2$. Integrating the reciprocal envelopes up to α_r gives a strict inequality below 1, and hence $\alpha_r^+ > \alpha_r$. □

B.4 An upper bound independent of the additional labels

Two hidden rows give a matching upper bound for the adjacent-label envelope even when the action can use every additional label.

Theorem B.9 (Two-row upper bound)**.** *For every integer $k > r \ge 2$ and $0 \le \lambda \le \gamma_r$,*

$$K_{r,k}(\lambda) \le \mathcal{K}_r^+(\lambda). \tag{45}$$

Proof. At $\lambda = 0$ every normalized row has value 1. For $0 < \lambda \le \gamma_r$, choose the minimizing parameters from Theorem B.7: use $c = c_0$ on the first branch and $c = c_*$ on the second, and let $t = t(c)$ as in (39). Put $d = r - 1$, $s = t + dc$ and $Z = 1 + 2s - t$. Extend the observed row to any number of labels by setting $g = (1, t, c, \dots, c) \in \mathbb{R}^k$. Consider two comparison rows and their designated targets,

$$a^{(1)} = (1, -dc, c, \dots, c), \quad j_1 = 1; \qquad a^{(2)} = (-dc, t, c, \dots, c), \quad j_2 = 2.$$

Both rows lie below g. Each has one negative coordinate $-dc$; all its other coordinates are at least c, so every r-coordinate sum is nonnegative. They therefore belong to $\mathcal{A}(g)$ for every $k > r$.

Mix the first comparison with weight

$$\pi = \frac{\lambda s - (1-\lambda)(1-t)}{\lambda Z} \tag{46}$$

and the second with weight $1 - \pi$. This is a probability distribution. Indeed $c, t \ge c_0$ and $\lambda r c_0 = (1-\lambda)(1-c_0)$ imply that its numerator is nonnegative. Also $\pi \le s/Z < 1$, since $Z - s = 1 + dc > 0$. For a fixed comparison (a, j), the payoff of selecting label i is $(1-\lambda)g_i + \lambda a_i - \lambda a_j$. After mixing the two comparisons, these per-label payoffs are

$$\begin{aligned} P_1 &= 1 - \lambda - \lambda s(1-\pi), \\ P_2 &= (1-\lambda)t - \pi\lambda(1 + s - t), \\ P_i &= c - \lambda t - \pi\lambda(1-t) \qquad (i \ge 3). \end{aligned}$$

Substitution of (46) gives

$$P_1 = P_2 = B = \frac{1 + s - t + st - \lambda(1 + 2s - t + s^2)}{Z}.$$

To verify the remaining labels without any dependence on k, set $a_0 = 2 - \lambda$ and $b = d\lambda - (d-1)$. Direct subtraction gives

$$Z(P_i - B) = -\big[t(a_0 - bc) - c\{d(d\lambda + 2)c + b\}\big] = 0 \qquad (i \ge 3),$$

where the final equality is precisely (39). Thus every label has the same mixed payoff, $B = \mathcal{K}_r^+(\lambda)$. For every probability vector q, its minimum payoff over all admissible comparisons is at most this particular mixture's payoff. Maximizing over q proves (45). On the first branch $t = c = c_0$ and $\pi = 0$, so the construction reduces to the single second comparison and remains valid at the switch.

□

At $r = 2$ and $\lambda = \gamma_2 = (\sqrt{17} - 3)/2$, the two rows specialize to Oshima's comparison witnesses for unconstrained three-label maximization (Oshima, 2021, Theorem 3.5 and Lemma 3.6).

B.5 Saturation and the pairwise fixed-label envelope

Theorem B.10 (Complete off-diagonal envelopes)**.** *For every $k > r \geq 3$ and $0 \leq \lambda \leq \gamma_r$, $K_{r,k}(\lambda) = \mathcal{K}_r^+(\lambda)$. For $r = 2$, $k \geq 4$ and $0 \leq \lambda \leq 1/2$, put $a = \lambda/(1-\lambda)$ and $N = k - 1$. Define*

$$B_N(a) = \begin{cases} \dfrac{Na}{N+a}, & a^2 \leq N(1-a), \\ N\big(\sqrt{N+1+a} - \sqrt{N}\big)^2, & a^2 \geq N(1-a). \end{cases} \tag{47}$$

Then the exact pairwise envelope is

$$K_{2,k}(\lambda) = \min\{\mathcal{K}_2^+(\lambda), M_k(\lambda)\}, \qquad M_k(\lambda) = \frac{1 - B_N(a)}{1+a}. \tag{48}$$

The two expressions for B_N agree at the switch. The infima are attained by nonnegative observations, although the optimization permits signed observations and signed comparison rows.

The matching upper bounds use $(1, t, c, \ldots, c)$ from Theorem B.9 and, in the pairwise case, the one-heavy rows $(1, t, \ldots, t)$ identified below. It remains to prove the universal lower bound. For $r \geq 3$ it has a sparse, explicit action. For fixed k in the pairwise case, it uses a finite linear program.

B.5.1 A capped prefix for exterior labels

Two reductions will be useful. An optimal action can be sorted in the observation order. Indeed, swapping $q_i < q_j$ when $g_i \geq g_j$ increases observed gain. For a comparison with $a_i \geq a_j$ it increases $q \cdot a$; otherwise swap a_i, a_j as well. This preserves the cone, the maximum and the upper bounds, since $a_i < a_j \leq g_j \leq g_i$. Taking the worst comparison proves that the loss cannot increase.

Lemma B.1 reduces the lower bound to nonnegative observations, also when a best coordinate must satisfy $q_1 \geq 1/r$.

Lemma B.11 (Exterior safety)**.** *Let $b \geq 0$ have $r+1$ coordinates, with smallest cap w. If $\max_i q_i \geq 1/r$, its exact $(r+1)$-coordinate replacement loss bounds every additional comparison label whose cap is at most w.*

Proof. Put $Q = \max_i q_i$. A minimizer of $q \cdot a$ over $\mathcal{A}(b)$ can be chosen with $\max a \geq w$. Otherwise, writing $v = \max a < w$, the cone gives $\sum_i (v - a_i) \leq rv$, hence $q \cdot a \geq (1 - rQ)v \geq (1 - rQ)w$. The row with $-(r-1)w$ at a largest-probability coordinate and w elsewhere is admissible and attains the latter value. Consequently $w - \min_{a \in \mathcal{A}(b)} q \cdot a \leq D_{\mathrm{E}}(b, q)$. Restrict a full comparison row to the chosen $r+1$ coordinates. An internal target is covered by $D_{\mathrm{E}}(b, q)$; an exterior target is at most w and is covered by the preceding inequality. This includes $w = 0$. □

Normalize $b_1 = 1 \geq b_2 = t \geq \cdots \geq b_{r+1} \geq 0$ and write $s = \sum_i b_i - 1$. For sorted q, Theorem B.5 gives $H_{\mathrm{E}}(b, q) = \max\{1 + sq_2, t + (1 + s - t)q_1\}$. Let $V_{\mathrm{cap}}(b, \lambda)$ maximize $q \cdot b - \lambda H_{\mathrm{E}}(b, q)$ subject to $q_1 \geq 1/r$.

Theorem B.12 (Capped-prefix envelope)**.** *For $r \geq 3$ and $0 \leq \lambda \leq \gamma_r$, $\inf_b V_{\mathrm{cap}}(b, \lambda) = \mathcal{K}_r^+(\lambda)$. For $r = 2$ and $0 \leq \lambda \leq 1/2$, the infimum is*

$$Q_2(\lambda) = \min\{\mathcal{K}_2^+(\lambda), 1 - 2\lambda\} = \inf_{k \geq 4} K_{2,k}(\lambda). \tag{49}$$

Selecting the top $r+1$ positive-part caps and applying this rule therefore gives the stated lower bounds for every signed full row.

Proof. At fixed s, t, averaging the last $r-1$ caps decreases every ordered action's gain and preserves its charge and the constraint $q_1 \geq 1/r$. It suffices to consider $(1, t, c, \ldots, c)$, with $s = t + (r-1)c$. The unrestricted optimum is $\max\{A, B, C\}$ from (37). The actions for A, B have first probability at least $1/2$. The only unresolved case is $C \geq A, B$ and its action having $q_1 < 1/r$, equivalently $s > r(r-1)(1-t)$. Set $u = 1/r$, $v = 1/r - (r-1)(1-t)/(rs)$, and distribute the remaining mass

equally among the last $r-1$ labels. The inequality just given ensures that the resulting probabilities are nonnegative and ordered. This action lies on the segment between the two-prefix and last-prefix actions, on which both charges agree. Its value is

$$W(s,t) = \frac{2-\lambda}{r} + \frac{r-2-(r-1)\lambda}{r(r-1)}s - \frac{r-2+(r-1)^2\lambda}{r(r-1)}t + \frac{t^2-t}{s}. \tag{50}$$

Along that segment the objective's slope in its first probability is $Z = 1-c+(t-c)(1+s-t)/s - \lambda(1+s-t)$. Thus $C \geq B$ is equivalent to $Z \leq 0$; it also implies $A \leq B$, since $\lambda(1+s-t) \geq 1-c \geq 1-t$. For fixed $z = t/s \in [1/r, 1]$, $W(s, zs)$ is affine in s, and

$$(r-1)Z = r-2+rz-(r-1)\lambda + s(1-z)[rz-2-(r-1)\lambda].$$

Its constant term is positive for $\lambda < 1$. Hence the difficult region at fixed z is an interval: its lower endpoints come from $Z = 0$ and $s = r(r-1)/[1+r(r-1)z]$, and its upper endpoint is $s = 1/z$. At the first two endpoints W equals an unrestricted optimum, at least $\mathcal{K}_r^+$. At the last endpoint, $t = 1$ and $Z \leq 0$ requires $c \geq c_1 = (2-\lambda)/[2+(r-1)\lambda]$. There

$$W(1+(r-1)c, 1) = \{2-(r+1)\lambda + [r-2-(r-1)\lambda]c\}/r.$$

For $r \geq 3$, $\lambda \leq \gamma_r \leq 1/2 \leq (r-2)/(r-1)$, so this is minimized at c_1, where $B = C$ again. This proves the lower bound. The adjacent-label equality rows attain it with a best-label or two-prefix action and hence meet the cap. The singleton case and $\lambda = 0$ follow directly.

For $r = 2$ the last displayed value decreases in c, giving the lower bound $1-2\lambda$ at $c = 1$. The row $(1,1,1)$ attains that value under $q_1 \geq 1/2$; the adjacent-label equality rows attain the other branch. Exterior safety and positive clipping give the full-row lower bound. Conversely the canonical rows attain $\mathcal{K}_2^+$ for every $k \geq 3$, and the constant k-coordinate row has exact value $1-2\lambda+\lambda/k$: uniform q has loss $1-1/k$, and a unit comparison on a least-probability target gives the reverse inequality. Letting $k \to \infty$ proves (49). For $r \geq 3$, the upper bound in Theorem B.9 now proves the first assertion of Theorem B.10 on its entire interval, for every $r \geq 3$. □

The action is explicitly computable. For $s = 0$, select the first label. For $s > 0$, the capacities in (34) have first entry $c_1(H) = (H-t)/(1+s-t)$. Sequential filling on its feasible interval has $q_1 = c_1(H)$, so the cap restricts that interval to $H \geq H_{\mathrm{cap}} := t + (1+s-t)/r$. Retain the candidates in (31) satisfying $q_1 \geq 1/r$. All prefixes of length at most r satisfy this condition. If the final prefix violates it, necessarily $d_+ = r+1$ and $H_{r+1} < H_{\mathrm{cap}} \leq H_r$. Add the filled vector at this new endpoint: $q_1 = 1/r$, $q_i = v$ for $2 \leq i \leq r$, and $q_{r+1} = 1-1/r-(r-1)v$, with v as above. It is nonnegative and sorted because it lies on the final feasible capacity segment. These are all breakpoints and endpoints of the restricted interval. Sorting and prefix sums give $O(k \log k + r)$ exact arithmetic operations and $O(k)$ records.

Let $\lambda_\dagger \in (1/3, 1/2)$ be the unique root of $\lambda^4+12\lambda^2-16\lambda+4 = 0$. Then $Q_2 = \mathcal{K}_2^+$ through $\lambda_\dagger$ and $Q_2 = 1-2\lambda$ afterward. On the first adjacent-label branch, $(1-2\lambda)(1+\lambda)-(1-\lambda)^2 = \lambda(1-3\lambda)$. On the second, isolating its positive radical and squaring gives the quartic. Its derivative is negative on $[0, 1/2]$ and its values at $1/3, 1/2$ have opposite signs; also $\eta_2 < 1/3$. Thus Theorem B.10 already holds in the initial pairwise interval. We prove the remaining interval next.

B.5.2 The pairwise lower bound

Throughout this subsection $k \geq 4$ and $\lambda_\dagger \leq \lambda \leq 1/2$. Write $K_+ = \mathcal{K}_2^+(\lambda)$, $K = \min(K_+, M_k)$ and use scalar variables $a = \lambda/(1-\lambda)$, $\kappa = K/(1-\lambda)$ and $b = 1-\kappa$. Here a denotes a scalar, not a comparison row. Define $\vartheta(p) = ap/(1+ap)$, $f_0(p) = p\vartheta(p)$ and $B_+(a) = (1+a)^2/(1+\sqrt{1+a/2})^2$. On this interval $b = \max\{B_+(a), B_N(a)\}$: the first term controls two-target replacement risk, and the second is a probability-capacity requirement. Thus $K = [1-\max\{B_+(a), B_N(a)\}]/(1+a)$.

For sorted nonnegative observations $g_1 = 1 \geq \cdots \geq g_k = c$, the full comparison game and its restriction to defect depth at most c both admit an ordered optimizer $q_1 \geq \cdots \geq q_k$: the coordinate swaps above preserve the cone and the lower bound $-c$ on every comparison coordinate. For these actions, truncation preserves the worst comparison loss, which is attained among the target–row pairs

$$g_j\mathbf{e}_j, \qquad c\mathbf{1} - 2c\mathbf{e}_i + (g_j - c)\mathbf{e}_j \quad (i \neq j), \qquad \text{target } j. \tag{51}$$

A pairwise comparison has at most one negative coordinate. For defect depth $d \in [0, c]$, set all other non-target coordinates to d and the target to g_j; its loss is affine in d, so an endpoint suffices. For a deeper defect, its coordinate must be the unique minimum cap k. Its depth coefficient is $2q_k + q_j - 1 \le 0$, because another coordinate has probability at least q_k. This proves the truncation claim and hence equality of the two games' optimal row values.

Fix any mixture of the target–row pairs in (51), and let p_j be its probability of target j. Define lower probabilities $\ell_1 = (ap_1 + \kappa)/(1 + ap_1)$ and $\ell_i = \vartheta(p_i)$ for $i \neq 1$. Concavity of ϑ shows that $\sum_i \ell_i \le 1$ for every p exactly when

$$\kappa \le \min_{0 \le u \le 1} \phi_N(u), \qquad \phi_N(u) = 1 - \frac{Na(1-u)(1+au)}{N + a(1-u)}. \tag{52}$$

Indeed, fix $p_1 = u$ and average the other N probabilities. Equivalently, $B_N(a) = \max_{0 \le u \le 1}[1 - \phi_N(u)]$. For $a > 0$, set $z = a(1-u)$, $w = N + z$ and $L = 1 + a$. Then

$$1 - \phi_N(u) = \frac{Nz(L-z)}{N+z} = N\left(2N + L - w - \frac{N(N+L)}{w}\right), \qquad N \le w \le N + a.$$

The expression $w + N(N+L)/w$ is minimized at $\widehat{w} = \min\{N + a, \sqrt{N(N+L)}\}$. The interior choice applies exactly when $a^2 \ge N(1-a)$. Substitution gives both branches of (47), including their equality at the switch. The one-heavy parameter is $t = (1+au)/[1 + a(1-u)/N] = N(N+L)/w - N$, and

$$\mu\phi_N(u) = \frac{N\mu(1+t^2) - (N - 1 + \lambda)t}{N + t}. \tag{53}$$

where $\mu = 1 - \lambda$. Thus the minimizing parameter is $t = \max(t_M, t_*)$, with $t_M = N/(N+a)$ and $t_* = \sqrt{N^2 + N(1+a)} - N$. The calculation holds for every $0 < a \le 1$; at $a = 0$ the capacity is zero. In particular, $b(N+a) - Na \ge 0$, $a - b \le a/k$, and

$$a\sum_i p_i \ell_i - a + b = 1 - \sum_i \ell_i \ge 0. \tag{54}$$

It suffices to complete every ℓ to a distribution $q \ge \ell$ with

$$R_p(q) := \sum_j p_j \max\{1 - q_j, 2\max_{i \neq j} q_i\} \le b/a. \tag{55}$$

Choose

$$q_i = \max\{\ell_i, \theta\}, \qquad \sum_i q_i = 1, \tag{56}$$

which is possible by (52). The analysis below verifies (55) for this distribution. To see why (55) yields the lower bound, the payoff of (56) against this mixture is at least

$$\sum_i [(\mu + \lambda p_i)q_i - \lambda p_i]g_i - \lambda c \sum_j p_j (2\max_{i \neq j} q_i + q_j - 1)_+.$$

All coefficients outside coordinate 1 are nonnegative and its first coefficient is at least K. Using $g_1 = 1$ and $g_i \ge c$, the expression is at least $(1-c)K + c[\mu - \lambda R_p(q)] \ge K$. Finite-dimensional minimax therefore gives the universal lower bound.

We use the following scalar inequality to analyze the completed distribution; its proof follows the construction analysis.

Lemma B.13 (A scalar completion inequality)**.** *Let $1/2 < a \le 1$, $m \ge 2$, $B_+(a) \le b \le a$ and $b \ge (m+1)a/(m+1+a)$. For $u, v \ge 0$, $u + v \le 1$, set $w = (1-u-v)/m$, $x = 1 - b/(1+au)$ and $y = \vartheta(v)$. Then $aT_A - a + b \ge 0$ and $aT_B - a + b \ge 0$, where*

$$T_A = ux + (1-2x)v + mf_0(w), \qquad T_B = u(1-2y) + (1-2x)v + mf_0(w).$$

The choice $b = \max(B_+, B_N)$ meets these assumptions with $m = k - 2$. Indeed $K_+ \ge 1 - 2\lambda$ on the present interval, while $\phi_N(u) > (1-a)$ follows from

$$\phi_N(u) - (1-a) = \frac{a\{Nu(1-a+au) + a(1-u)\}}{N + a(1-u)} > 0.$$

Thus $B_N < a$, $B_+ \le b \le a$, and capacity gives the endpoint bound.

Completing the target probabilities. For the distribution (56), if $Q = \max_i \ell_i < 1/3$, then $\max_i q_i \le 1/3$ and $1 - R_p(q) = p \cdot q \ge p \cdot \ell$; (54) proves (55). If $Q \ge 1/3$, filling leaves the largest bound unchanged because $\theta \le 1/k < 1/3$. If the second largest bound is raised, then $\theta \le (1-Q)/(k-1) \le (1-Q)/2$, so the largest coordinate's reward remains Q; all other rewards can only increase. It is enough to prove

$$G_0(p) := \sum_j p_j \min(\ell_j, 1 - 2\max_{i \ne j} \ell_i) \ge (a-b)/a. \tag{57}$$

Suppose first that the anchor has a largest lower bound $x = \ell_1 \ge 1/3$, and write $u = p_1$, $d = 1-2x$. If $d \ge 0$ and every exterior lower bound is at most d, then its reverse penalty is absent as well, since $d \le (1-x)/2$; use (54). Otherwise merge all exterior target masses satisfying $\vartheta(p_i) > d$ into one mass v. Their total reward is dv; the largest exterior lower bound increases, which only decreases the anchor reward. Average the other $m = k-2$ masses, including the zeros created by merging. Their average remains below the truncation threshold, and convexity of f_0 decreases their reward. Thus $G_0(p) \ge u\min(x, 1-2\vartheta(v)) + dv + mf_0(w) = \min(T_A, T_B)$. If $d < 0$, take $v = 1-u, w = 0$. The merged bound may exceed the anchor; the scalar lemma needs no ordering assumption. It proves (57) in either case.

Now let a non-anchor bound be largest, with target mass v and $y = \vartheta(v) \ge 1/3$, and let $u = p_1$, $x = \ell_1 < y$. Then $av \ge 1/2$ and $y \le 1/2$. Every other non-anchor bound is at most $(1-y)/2$: its mass is at most $1-v$, and $a(1-v)(1+2av) \le 1$, since this expression decreases for $v \ge 1/(2a)$ and its endpoint value is $2a - 1 \le 1$. For the remaining cases we use

$$\kappa = 1 - b < 1 - 4a/5. \tag{58}$$

To verify it, put $\xi = \sqrt{1 + a/2}$. The derivative of $B_+(a)/a$, expressed as a function of ξ, is $(2\xi^2-1)(2\xi^2-2\xi-1)/[(\xi-1)^2(\xi+1)^4] < 0$ on $\sqrt{5}/2 \le \xi \le \sqrt{6}/2$. Hence $B_+(a)/a \ge B_+(1) = 40 - 16\sqrt{6} > 4/5$, the last sign being $2401 > 2400$.

If $2x + y > 1$, no other non-anchor bound exceeds $1-2y$. Otherwise the total target mass would force $\kappa > \kappa_c = (1-y)/2 - (1+y)[a - y/(1-y) - (1-2y)/(2y)]/2$, whereas

$$\kappa_c - (1 - 4a/5) = \frac{3-5y}{10}\left(a - \frac{y}{1-y}\right) + \frac{20y^3 + 11y^2 - 20y + 5}{20y(1-y)} > 0.$$

Here the cubic, with $z = y - 2/5$, equals $31(z - 4/155)^2 + 3/155 + (4 + 20z)z^2 > 0$ on $z \in [-1/15, 1/10]$. This contradicts (58). Thus the anchor and largest non-anchor rewards are $u(1-2y)$ and $v(1-2x)$, and all other rewards are untruncated. Averaging the remaining masses reduces the bound to T_B.

Finally, if $2x + y \le 1$, all rewards are nonnegative and the largest non-anchor contributes vy. Put $z = av \ge 1/2$. For $k \ge 6$, $aG_0 \ge z^2/(1+z) \ge 1/6 \ge a/k \ge a - b$. For $k = 4, 5$, it suffices to prove $G_0 \ge 1/5$ by (58). Let j be the largest non-anchor and put $X_i = ap_i$ for $i \ne j$. Since $X_i \le a - z \le 1 - z \le z$ and $d = 1 - 2y = (1-z)/(1+z)$, each remaining reward is at least $X_i/(1+z)$, including the anchor because $\ell_1 \ge \vartheta(p_1)$. Cauchy–Schwarz gives

$$aG_0 \ge \frac{z^2 + \sum_{i \ne j} X_i^2}{1+z} \ge \frac{z^2 + (a-z)^2/4}{1+z} > \frac{a}{5}.$$

The strict inequality follows from

$$5a^2 - (14z+4)a + 25z^2 = (1-a)(14z - 1 - 5a) + (z - \tfrac{1}{2})(25z - \tfrac{3}{2}) + \tfrac{1}{4} > 0,$$

using $a \le 1$ and $z \ge 1/2$. This proves (55) and the universal pairwise lower bound.

Auxiliary inequalities.

Lemma B.14 (Three-mass domination)**.** *For $1/2 \le a \le 1$, any real $m \ge 2$, and $X, Y, Z \ge 0$ with $X + Y + Z = a$, put $D_m = (1+X)[2Y + mZ/(m+Z)]/(1+2Y)$. Then*

$$D_m \le \max\left\{B_+(a), \frac{(m+1)a}{m+1+a}\right\}. \tag{59}$$

Proof. Set $\ell(a) = (8a+3)/14 + 1/800$. We prove the stronger bound $D_m \le \max\{\ell(a), (m+1)a/(m+1+a)\}$ and first verify $\ell(a) < B_+(a)$. For $q = \sqrt{1+a/2} \ge 1$ and $s = q - 9/8$, direct expansion gives

$$5600(q+1)^2[B_+(a) - \ell(a)] = (16000q^2 + 23200q + 8343)s^2 - \frac{819}{4}s + \frac{227}{64}.$$

The quadratic coefficient is at least 47543, so this expression is at least $227/64 - (819/4)^2/(4 \cdot 47543) = 2530375/760688 > 0$.

Let B be the maximum of D_m on the compact triangle. It is at most a: replacing $mZ/(m+Z)$ by Z leaves, at fixed X, a linear-fractional function of Y, whose endpoints are bounded by $(1+X)(a-X) \le a$ and $2Y(1+a-Y)/(1+2Y) \le a$; the latter inequality is $2(Y-1/2)^2 + a - 1/2 \ge 0$. Suppose $B > \max\{\ell(a), (m+1)a/(m+1+a)\}$, and set $E = B(1+2Y)/(1+X) - 2Y - mZ/(m+Z)$ with $Z = a - X - Y$. At a maximizing point $E = 0$, and all derivatives tangent to its face vanish. We exclude the three edges and the interior.

On $Y = 0$, put $n = m+1$. Clearing positive denominators in $na/(n+a) - D_m$ gives

$$a^2 + [m(n+a)(1-a) - na]X + m(n+a)X^2 \ge \tfrac{1}{2}(nX - a)^2 + \tfrac{1}{2}a^2 > 0,$$

since $a \le 1$ and $m(n+a) \ge n^2/2$. On $Z = 0$, direct maximization gives $D_m \le 2 + a - \sqrt{3+2a} \le (8a+3)/14$; the final inequality follows from $196(3+2a) - (6a+25)^2 = (2a-1)(37-18a) \ge 0$. These bounds include all vertices.

At an interior maximum let $t = m/(m+Z) \in (0,1)$, $d = 1+X$, and $h = m(1-t)$. The equations $E_Y = E_X = 0$ give $B/d = 1 - t^2/2$ and $1 + 2Y = 2dt^2/(2-t^2)$. Together with $E = 0$, they imply $d = (1-h)(2-t^2)/t^4$ and $h = [4 - (2a+3)t^4]/[2(2-t^3)]$. All denominators are positive, hence

$$B = \frac{(2-t^2)^2[(2a+3)t - 2]}{4t(2-t^3)} \le \frac{8a+3}{14}.$$

For completeness, $28t(2-t^3)$ times the claimed nonnegative difference is $(2-2a)P_0(t) + (2a-1)P_1(t)$, where $P_0 = 28(1-t)^3(t+1)(t+2)$ and

$$P_1 = (t - 5/6)^2 Q(t) + 425/3888 + (751/1296)(1-t),$$
$$Q(t) = 3463/108 + (571/12)(1-t) + (199/3)t(1-t) + 35t(1-t^2).$$

Every displayed summand is nonnegative on $[0,1]$, and $P_1 > 0$.

On the remaining edge $X = 0$, take $Y, Z > 0$, $z = Z/(m+Z)$ and $h = mz$. Stationarity and $E = 0$ give

$$B = \frac{1 + 2z - z^2}{2}, \qquad a = \frac{B - hz}{(1-z)^2}, \qquad 2z \le h \le B. \tag{60}$$

Thus $a \ge B/(1-z)$, and $a \le 1$ implies $0 < z \le 2 - \sqrt{3} < 1/3$. The assumed violation of the explicit capacity bound implies

$$E_z(h) := 4h^2 + 4h(z^2 - z - 1) - z^4 + 4z^3 - 4z^2 + 1 \ge 0,$$

since $E_z = 4(1-z)^2[B(m+1+a) - (m+1)a]$. For $1/7 \le z \le 1/3$, the cubic $z^3 - 13z^2 + 9z - 1$ is increasing and is $8/343 > 0$ at $1/7$. Both endpoint values $E_z(2z) = -(z+1)(z^3 - 13z^2 + 9z - 1)$ and $E_z(B) = -2z(1-z)(1 + 2z - z^2)$ would therefore be negative. Convexity on $2z \le h \le B$ gives a contradiction, so $z < 1/7$. Because $E_z(B) < 0$, its smaller root bounds h from above:

$$h \le \frac{1 + z - z^2 - \sqrt{z(2 + 3z - 6z^2 + 2z^3)}}{2} \le \frac{1 + z - z^2 - (7/5)\sqrt{z}}{2}.$$

Here $2 + 3z - 6z^2 + 2z^3 \ge 2 > 49/25$. Substituting this bound into (60) gives a lower bound $a_0 \le a$. For $v = \sqrt{z}$, expansion yields

$$\begin{aligned} 5600(1-z)^2[\ell(a_0) - B] &= 2800v^8 - 9600v^6 + 9207v^4 + 2240v^3 - 814v^2 + 7 \\ &\ge 7800v^4 + 2240v^3 - 814v^2 + 7 > 0. \end{aligned}$$

The first inequality uses $v^2 < 1/7$ and $9207 - 9600/7 > 7800$. For the last, with $s = v - 1/7$, the quartic equals

$$\left(7800v^2 + \frac{31280}{7}v + \frac{14874}{49}\right)s^2 - \frac{1532}{343}s + \frac{401}{2401} \geq 300s^2 - 5|s| + \frac{1}{6} \geq \frac{7}{48} > 0.$$

Consequently $B < \ell(a_0) \leq \ell(a)$, the final contradiction. □

Proof of Lemma B.13. For T_A, put $X = au$, $Y = av$ and $Z = amw$. The exact identity

$$aT_A - a + b = b\frac{1+2Y}{1+X} - 2Y - \frac{mZ}{m+Z}$$

and Lemma B.14 prove the claim. For T_B, set $A = a(u+v)$ and $s = \sqrt{b}$. At fixed A, minimize $1/(A+2-D)+b/D$ over $1 \leq D \leq 1+A$. The unconstrained minimizer is $D = (A+2)s/(1+s)$; its upper endpoint never binds. For $A \leq A_0 = 1/s - 1$ the lower endpoint gives $u = 0$, where $T_B = T_A$. Otherwise the minimized value is

$$F(A) = b + 4s - m - 2A - \frac{2(s+1)^2}{A+2} + \frac{m^2}{m+a-A}.$$

Its third derivative is positive and $F'(a) = -1 + 2(s+1)^2/(a+2)^2 < 0$, since $b \leq a$ and $(a+2)^2 - 2(1+\sqrt{a})^2 = a^2 + 2(\sqrt{a}-1)^2 > 0$. Convexity of F' places the minimum of F on $[A_0, a]$ at an endpoint. The first is already covered; the second equals $[ab+4(a+1)\sqrt{b}-2(a+1)^2]/(a+2) \geq 0$, exactly the condition $b \geq B_+(a)$. If $A_0 > a$, only the lower-endpoint case occurs. □

Attainment. The two-special-coordinate upper bound is already proved in Theorem B.9; its best-label or two-prefix action has first probability at least $1/2$, so exterior safety gives equality on its extended row. For M_k, set $t_M = N/(N+a)$, $t_* = \sqrt{N^2 + N(1+a)} - N$, $t = \max(t_M, t_*)$, and use $g = (1, t, \ldots, t)$ and $q = (u, (1-u)/N, \ldots, (1-u)/N)$, where $u = (N-1)/(2N-1)$ if $t_* \leq t_M$, and $u = [N - (N-1)t]/(N+t)$ otherwise. The exact loss is $\max\{1-u, t(N-1+u)/N, 2tu\}$, obtained by separating nonnegative comparisons from a single negative coordinate. Since $t_M \geq N/[2(N-1)]$, substitution gives (53). For a matching upper bound, mix the nonnegative comparisons $g_j \mathbf{e}_j$ with first probability $p_1 = [(N+a)t - N]/[a(N+t)] \in [0, 1/k]$ and other probabilities $(1-p_1)/N$. The identity $(\mu + \lambda p_1) = (\mu + \lambda p_i)t$ makes every pure action's payoff equal to M_k. Choosing the family with the smaller value completes Theorem B.10.

B.5.3 THE THREE-TO-FIVE-LABEL PLATEAU

Set $\beta_{2,3} = \alpha_2^+$. For $k \geq 4$, define $\beta_{2,k}$ and $\beta_{2,\infty}$ by

$$\int_0^{\beta_{2,k}} \frac{d\lambda}{K_{2,k}(\lambda)} = 1, \qquad \int_0^{\beta_{2,\infty}} \frac{d\lambda}{Q_2(\lambda)} = 1. \tag{61}$$

The envelopes are positive below $1/2$, nonincreasing, and 3-Lipschitz: every normalized full charge lies in $[1, 3]$, and the capped envelope has the same property. Their reciprocal integrals reach one before $1/2$, since $K_{2,k} \leq \mathcal{K}_2^+ \leq (1-\lambda)^2/(1+\lambda)$ and $\int_0^{1/2}(1+\lambda)/(1-\lambda)^2\, d\lambda = 2 - \log 2 > 1$.

Corollary B.15 (Three to five labels)**.** *For $k = 4, 5$, $K_{2,k} = \mathcal{K}_2^+$ throughout $[0, 1/2]$. Consequently $\beta_{2,4} = \beta_{2,5} = \alpha_2^+$, also the three-label coefficient. For every $k \geq 6$, $\beta_{2,\infty} < \beta_{2,k} < \alpha_2^+$.*

Proof. The first adjacent-label branch is at most $1 - 2\lambda < M_k$ for $\lambda > 0$. On the second branch put $z = \sqrt{1 + a/2}$, $b = b_+ = (2z^2-1)^2/(z+1)^2$, $E = b(N+a) - Na$ and $C = N(1-a) - b$. For $N = 4$,

$$4NE - C^2 = \frac{(3-2z^2)(2z^2-1)^2(2z^2+16z+13)}{(z+1)^4} \geq 0.$$

Thus $Na^2u^2 + aCu + E \geq 0$ for every real u, which is exactly the capacity condition for K_+; hence $M_5 \geq K_+$. Furthermore $\partial_N \phi_N = -a^2(1-u)^2(1+au)/[N+a(1-u)]^2 \leq 0$, so $M_4 \geq M_5$. At $\lambda = 0$ all envelopes equal one. For $k = 6$, direct substitution at $\lambda = 3/7$ gives

$$M_6(3/7) = \frac{32}{161} < \frac{4}{7} - \frac{14}{19 + 4\sqrt{22}} = \mathcal{K}_2^+(3/7).$$

The strict inequality is equivalent to $120\sqrt{22} > 557$, since $120^2 \cdot 22 - 557^2 = 6551 > 0$. Corollary E.2 gives $\alpha_2^+ > 0.4366 > 3/7$, so continuity and $M_k \le M_6$ for $k \ge 6$ make the integral inequality strict. Finally $M_k > 1 - 2\lambda$ for finite k, while $\mathcal{K}_2^+ > 1 - 2\lambda$ above $\lambda_\dagger$. The bound on $\beta_{2,\infty}$ in Corollary E.2 exceeds $\lambda_\dagger$, so strictness on a positive-length interval proves $\beta_{2,k} > \beta_{2,\infty}$.

For $0 < \lambda < 1/2$, the displayed derivative is strictly negative at the minimizing $u < 1$, so $M_{k+1} < M_k$. Since $\beta_{2,k} < \alpha_2^+$ for $k \ge 6$, the branch M_k is active on a positive-length interval below $\beta_{2,k}$; thus $\beta_{2,k+1} < \beta_{2,k}$. Finally, $K_{2,k} \to Q_2 > 0$ on $[0, \alpha_2^+]$, so dominated convergence in (61) gives $\beta_{2,k} \downarrow \beta_{2,\infty}$. □

C FINITE MATROID TRANSFER

Work in the matroid setting of Section 2, with $k \ge r \ge 2$, and remove loops. Using the comparison state of Proposition 4.2, we complete Algorithm 2's preprocessing, rational clock, rounding, and full-row local implementations. Appendix D gives the two-label implementation.

C.1 RESIDUAL CAPS AND SIGNED PREPROCESSING

GREEDY repeatedly inserts a maximum-feasible-marginal label until its support is a base. Its value D satisfies $M/(r+1) \le D \le M$. Indeed, exchange deletes a comparison element of marginal at most the chosen maximum m. The r-sum condition bounds the inserted label's comparison marginal below by $-(r-1)m$, so the comparison loses at most rm when greedy gains m. Telescoping gives $M - D \le rD$. If $D = 0$, the empty labeling is optimal.

For a prefix A, let $m_e(A) = \max_i \Delta_{e,i} F(A)$ and let $U(A)$ be the maximum sum of these caps over a base of the contraction by A. Zero rows need no sampling. For a nonzero row $g \in \mathcal{C}_r$, put $m = m(g)$, let T index its r largest entries, and set $b_i = (g_i)_+$ for $i \in T$. Write $S = \sum_T b_i$, $Q = \sum_T b_i^2$, and sample with $q_i = b_i/S$ on T and $q_i = 0$ elsewhere. For $a \in \mathcal{A}(g)$, with $h = Q/S$ and $\kappa_r = (1 + \sqrt{r})/2$,

$$m \le \kappa_r h, \qquad q \cdot a \ge h - m \ge -(\kappa_r - 1)h.$$

The first inequality follows from $Q - 2mS/(1+\sqrt{r}) \ge (S - \sqrt{r}\, m)^2/(r-1) \ge 0$. For the second, $\sum_T a_i \ge 0$ and $a_i \le g_i$ give $Sq \cdot a = m \sum_T a_i - \sum_T (m - b_i) a_i \ge Q - m \sum_T g_i \ge Q - mS$.

Residual-marginal preprocessing for ordinary nonnegative set functions appears in Kulik et al. (2026); the labelled drift below supplies the signed transfer used here.

Lemma C.1 (Residual preprocessing)**.** *Fix $V \ge M$ and $0 \le a \le \gamma_r$. While $U(A) > 3V$, choose a maximum-h contraction base, a uniform element of that base, and a proportional top-r label. In at most ρ insertions this gives a prefix and comparison base satisfying*

$$\mathbb{E}[(1-a)F(A) + aF(O_A)] \ge aM, \qquad U(A) \le 3V. \tag{62}$$

Every realization of the procedure, including one run at $V < M$, uses $O(nk\rho)$ value and $O(n\rho)$ independence calls.

Proof. Maintain a feasible comparison base containing A, starting from an optimal labeled base. At residual rank d, a basis-exchange bijection matches a uniform element of the chosen base to a uniform comparison element, matching common elements to themselves. On the fixed-label submodular restriction, backward marginal summation gives

$$\sum_{u \in O_A \setminus A} [F(O_A) - F(O_A \setminus \{u\})] \le F(O_A) - F(A) \le M.$$

Here $F(A) \ge 0$, since only nonnegative-gain labels are inserted, while $F(O_A) \le M$ follows from feasibility. Thus mean deletion loss is at most M/d, without any lower bound on the signed comparison value. The selected mean gain satisfies $\bar{h} \ge U(A)/(\kappa_r d)$. The potential drift is at least $(1 - a\kappa_r)\bar{h} - aM/d$, and while the loop continues it exceeds $[1 - (4/3)a\kappa_r]\bar{h}$. This is positive throughout $a \le \gamma_r$: the polynomial $P(x) = (r-1)x^2 + 3x - 2$ has positive root γ_r, and its monotonicity for $x \ge 0$ and $P(3/(4\kappa_r)) = 1/4 > 0$ give $\gamma_r \kappa_r < 3/4$. Summing the conditional

drifts proves the potential inequality. At termination $U(A) \le 3V$; a full-rank prefix has $U(A) = 0$. Each insertion uses exact marginal rows and two weighted matroid-greedy scans, proving the bounds even at an invalid scale. □

Taking $V = (r+1)D$ gives $U(A) \le 3(r+1)M$. Condition on the resulting prefix and write $m_e = m_e(A)$. Every compatible marginal observation lies in $[-(r-1)m_e, m_e]$, and the sum of these fixed caps over any residual base is at most $U(A)$. These pathwise bounds will charge sampling errors.

Condition on the preprocessing history, including A and its comparison base O_A. In the contracted matroid use $\mathcal{F}_A(x) = \mathbb{E}F(A + X_x)$ and $\mathcal{Z}_A(x) = \mathcal{F}_A(z(x))$, using the comparison-state construction of Section 4 with the residual comparison labels. Write $\Phi_\lambda(x) = (1-\lambda)\mathcal{F}_A(x) + \lambda\mathcal{Z}_A(x)$; then $\Phi_a(0) = (1-a)F(A) + aF(O_A)$. The row-affine and exchange proof of Proposition 4.2 applies to this fixed residual comparison base: it uses its feasibility, not its optimality. Its terminal value is the value of the full labeling including A. Averaging over preprocessing therefore preserves (62)'s lower bound aM.

C.2 A CLOCK WHOSE LOSS SCALES WITH ITS COEFFICIENT

The gain can pay for finite implementation error, but a separate clock argument also reduces the number of sampled directions needed at a fixed tolerance. It uses equal changes in comparison weight and unequal block durations. The construction needs only monotonicity of the envelope.

Lemma C.2 (Relative-error clock)**.** *Let K be continuous and nonincreasing on $[0, b]$, with $K(0) = 1$ and $K > 0$ on $[0, b)$. Suppose $\alpha = u(1) < b \le 1$, where $u' = K(u)$ and $u(0) = 0$. For integers $J \ge 2$, $G \ge 1$ and $H \ge J^3$, there is a rational clock with levels $\ell_j = ja/J$, J positive durations totaling one, and initial coefficient*

$$\alpha(1 - 1/J) - 1/G \le a \le \alpha. \tag{63}$$

All durations are multiples of $1/H$. Unused time may be placed at any chosen level ℓ_s, with

$$\sum_{j=1}^{J} \delta_j K(\ell_j) \le a + J/H + K(\ell_s). \tag{64}$$

For $K = \mathcal{K}_r$, the clock can be constructed by exact rational comparisons with no objective or independence queries.

Proof. For a candidate $a \in G^{-1}\mathbb{Z} \cap [0, 1]$, write $h = a/J$ and $\ell_j = jh$. If $a > 0$, choose the smallest positive integer $t_j \le H$ satisfying $(t_j/H)K(\ell_j) \ge h$. Reject a candidate if a level is outside $[0, b)$, a required integer does not exist, or $\sum_j t_j > H$. Accept zero with $t_j = 1$. Feasibility decreases with a: both the numerator h and the arguments of the nonincreasing K increase. Select the largest accepted grid candidate. Set $\delta_j = t_j/H$ and add all remaining duration to δ_s. Thus every duration is positive,

$$\sum_{j=1}^{J} \delta_j = 1, \qquad \ell_j - \ell_{j-1} \le \delta_j K(\ell_j). \tag{65}$$

The algorithm executes these pairs in order $j = J, J-1, \dots, 1$.

For $a > 0$ with $(J+1)h < b$, monotonicity of $1/K$ gives

$$S(a) := \sum_{j=1}^{J} \frac{h}{K(jh)} \le \int_h^{(J+1)h} \frac{dx}{K(x)} \le T((J+1)h), \qquad T(x) = \int_0^x \frac{dv}{K(v)}.$$

Put $\sigma = J/H \le J^{-2}$ and $a_0 = J\,u(1-\sigma)/(J+1)$. Then $S(a_0) \le 1 - \sigma$. Ceiling the durations adds less than $J/H = \sigma$, so every grid candidate at most a_0 is accepted. The trajectory u is concave, because its nonnegative velocity $K(u)$ decreases. Hence

$$a \ge a_0 - 1/G \ge \frac{J}{J+1}(1 - J^{-2})\alpha - 1/G = (1 - 1/J)\alpha - 1/G.$$

Conversely, every accepted positive candidate satisfies

$$T(a) = \sum_{j=1}^{J} \int_{(j-1)h}^{jh} \frac{dx}{K(x)} \le \sum_{j=1}^{J} \frac{h}{K(jh)} \le 1,$$

and therefore $a \le \alpha$. This is also immediate at $a = 0$. Before padding, minimality of the integers gives $\sum_j (t_j/H)K(\ell_j) \le a + J/H$, also at $a = 0$. The extra duration is at most one, proving (64). □

The full-row implementation below takes $H = J^3$ and $s = 1$. The envelope predicates below implement the test $K(x) \ge y$. In the duration search take $y = Hh/t_j$. Nested integer bisections require $O(J \log H \log(G+1))$ scalar predicate calls. Levels have denominator dividing JG and durations have denominator H; integer lengths are polynomial in $\log r, \log J, \log G, \log H$.

Exact scalar predicates. For $\mathcal{K}_r(x) \ge y$, require $(r-2)x^2+4x-2 \le 0$. At equality only $y = 0$ is accepted. The switch is determined by $(2r-3)x^2+4x-1$; before it, test $(1-x)^2 \ge y[1+(r-1)x]$. After it, with $D_r(x) = 2 + (r-2)x$, test $(ry + 2D_r(x))^2 \le 8(1-x)(r + D_r(x))$. All searches use nonnegative x, y.

For $\mathcal{K}_r^+(x) \ge y$, require $(r-1)x^2+3x-2 \le 0$ and $x < 1$. The sign of $(1-x)^3 - rx^2(2-x)$ selects the first branch, with the same rational test as above. On the second branch, put $f = 1-x-y$, $a_0 = 2-x$, $R = r(1-x)$ and $z = r - f(a_0 + R)$. The test is $f > 0$ and either $z \le 0$ or $4f^2 a_0 R \ge z^2$. Isolating the positive square root before squaring gives this sign condition. Thus the relative clock also gives

$$\alpha_r^+(1 - 1/J) - 1/G \le a \le \alpha_r^+. \tag{66}$$

For the pairwise envelope $K_{2,k} = \min\{\mathcal{K}_2^+, M_k\}$, use the preceding predicate and, with $N = k-1$, $a_0 = x/(1-x)$ and $b_0 = 1 - y/(1-x)$, test

$$Na_0^2u^2 + a_0[N(1-a_0) - b_0]u + b_0(N + a_0) - Na_0 \ge 0 \quad (0 \le u \le 1).$$

Endpoints and an interior convex vertex suffice. For the uniform sparse pairwise envelope Q_2, combine the $\mathcal{K}_2^+$ predicate with $1 - 2x \ge y$. These predicates evaluate the local formulas using rational arithmetic.

C.3 Compression, support rounding, and finite random bits

COLLAPSE replaces each nonzero row independently by $z_e \mathbf{e}_{J_e}$, where $z_e = \sum_j x_{ej}$ and $\Pr(J_e = j \mid x) = x_{ej}/z_e$; zero rows are unchanged. Whole-row affinity applied successively gives

$$\mathbb{E}[\mathcal{F}(\overline{x}) \mid x] = \mathcal{F}(x), \qquad \mathbb{E}[\mathcal{Z}(\overline{x}) \mid x] = \mathcal{Z}(x).$$

COMPRESS applies this operation after a completed block and its comparison-weight decrease; it is the identity for the dense fixed-k pairwise LP. The next block may adapt to the collapsed labels because its comparison is conditional on the current state. Terminal collapse is also valid if all rows were retained.

Each direction base B_j contributes $\delta_j \mathbf{1}_{B_j}$ to fractional support. The durations sum to one, so the final support is in the residual base polytope; intermediate row sums are legal because partial durations do not exceed one. Use symmetric-exchange swap rounding (Chekuri et al., 2010). A two-base merge takes at most ρ exchanges, each found with at most 2ρ independence tests. ACCUMULATE merges each new weighted base into a reservoir using an independent random stream hidden from the optimizer. Conditional on the full optimization and compression history, the bases, weights, and terminal labels are fixed, and this reservoir has the usual weighted swap-rounding law. For fixed labels ℓ, $H_\ell(S) = F(A + \ell(S))$ is a possibly signed submodular set function. Its multilinear extension is convex along each exchange direction because its mixed partials are nonpositive. Every mean-preserving swap therefore preserves or increases expected value. After terminal collapse, ROUND returns $A + \ell(R)$, where R is the reservoir base and ℓ_e is row e's surviving label; at residual rank zero it returns A. Only the reservoir and the next direction base need be stored.

Retain the best of the rounded output, empty labeling, greedy solution, and preprocessing prefix. Each retained improvement lies in $[0, M]$.

Dyadic sampling and adaptive coupling. Consider an adaptive procedure with at most $Q \geq 1$ categorical draws, each supported on at most $s \geq 2$ outcomes, where Q, s are integers. For $0 < \varepsilon < 1$, round all but the last supported outcome's probability down to multiples of 2^{-b} and assign the remainder to the last, with $b = \lceil \log_2(Q(s-1)/\varepsilon) \rceil$. Every dyadic outcome remains in the original support. Inverse cumulative sampling uses b fair bits per nontrivial draw. Each conditional law changes in total variation by at most $(s-1)2^{-b}$. Couple the exact and dyadic procedures while their histories agree. The probability of a first disagreement is at most $Q(s-1)2^{-b} \leq \varepsilon$, so a final reward in $[0, M]$ on every supported trajectory loses at most εM in expectation. Deterministic draws need no bits; a procedure with no draws is unchanged. The probabilities can be computed by bounded exact-real comparisons with dyadic rationals, without further oracle queries. Here Q counts all categorical draws used in sampled contexts, compression, preprocessing, and support swaps.

For rational inputs, arithmetic and bit costs include the lengths of oracle answers and conditional probabilities. Rational upper bounds on sample-budget logarithms and a rational tolerance in $(\tau/2, \tau)$ give finite valid budgets with the same asymptotic bounds.

C.4 Finite full-row implementations

Theorem C.3. *For the assumptions of this appendix and $0 < \tau < a_{r,k}$, a finite algorithm returns feasible Y with $\mathbb{E}F(Y) \geq (a_{r,k} - \tau)M$. For $r \geq 3$, and also for $(r,k) = (2,2), (2,3)$, it uses*

$$Q_f = O\left(nk\rho r^{7/2}\tau^{-3}\log\frac{2nkr\rho}{\tau}\right), \quad Q_I = O\left(\frac{n\rho}{\tau\sqrt{r}}\right),$$

and $O(nr+k+1/(\tau\sqrt{r}))$ numerical records. For $0 < \tau < \alpha_r$, the baseline $\alpha_r - \tau$ has these bounds for every $k \geq r$. For $r = 2, k \geq 4$, attaining $\beta_{2,k} - \tau$ instead uses $Q_f = O(nk\rho\tau^{-3}\log(2nk\rho/\tau))$, $Q_I = O(n\rho/\tau)$, and $O(nk + k^3 + \tau^{-1})$ numerical records. For every $k \geq 3$ and $0 < \tau < \beta_{2,\infty}$, the sparse pairwise rule attains $\beta_{2,\infty} - \tau$ with the same query orders and $O(n + k + \tau^{-1})$ records. LP arithmetic, context construction, and oracle-answer lengths are charged separately.

Proof. Use Lemma C.1 with $V = (r+1)D$, so $U \leq 3(r+1)M$. Its valid weight range includes all coefficients in the statement: they do not exceed γ_r, and the pairwise coefficients are below $1/2 < \gamma_2$.

Condition on the history before each selection. For every candidate, draw fresh categorical contexts shared across its complete marginal row. The empirical row lies in $\mathcal{C}_r^k$ on every outcome. Raise it by $\eta m_e \mathbf{1}$. On the simultaneous success event for one selection, $g_e \leq \bar{g}_e \leq g_e + 2\eta m_e \mathbf{1}$ for every candidate. Zero-cap candidates require no samples. The local construction determines a gain-error factor χ and an action-support bound s.

For the baseline and adjacent envelopes, and the capped-prefix rule at $r \geq 3$, apply the local optimizer to the raised row. Its positive-support action has at most $s = r + 1$ entries and true-gain discrepancy at most $\chi\eta m_e$ with $\chi = 2$.

For the fixed-k and uniform pairwise refinements, positively clip the raised row and round its positive entries upward to multiples of ηm_e, obtaining b_e. Its positive support has true-gain discrepancy at most $3\eta m_e$; take $\chi = 3$. The dense fixed-k implementation has $s = k$ and solves the rational LP $\max_q\{(1-\lambda)q \cdot b_e - \lambda D_{\mathrm{E}}(b_e, q)\}$. Its complete comparisons are $b_j\mathbf{e}_j$ and $d_i\mathbf{1} - 2d_i\mathbf{e}_i + (b_j - d_i)\mathbf{e}_j$ for $i \neq j$, where $d_i = \min_{h \neq i} b_h$, as proved in the pairwise local analysis. There are $k+1$ variables and $O(k^2)$ inequalities. A complete rational LP solver supplies the optimizer; after scaling by m_e, coefficient lengths are polynomial in the clock and accuracy parameter lengths. An optimal action avoids zero caps when $b_e \neq 0$; use any label when $b_e = 0$, whose true row is zero on success. The uniform rule instead applies the capped-prefix optimizer to the top three coordinates of b_e, with $s = 3$. The dense implementation retains fractional rows until terminal collapse; the uniform rule uses block-boundary compression.

Both constructed rows dominate g_e, so every true comparison remains admissible. Select the addable element of greatest constructed cap, which dominates the true maximum of its exchange-matched partner. Round all but a largest-cap supported action coordinate down to a dyadic denominator D_q, assigning it the remainder. This moves at most $(s-1)/D_q$ mass, cannot reduce surrogate gain, and

adds at most $r(s-1)m_e/D_q$ to any true comparison loss. Set

$$\eta = \frac{\tau}{48\chi(r+1)}, \qquad d_{\text{fail}} = \frac{\tau}{48(r+1)^2\rho}, \qquad D_q \geq \frac{48r(s-1)(r+1)}{\tau}.$$

Hoeffding's inequality and a union bound over nk coordinates require

$$O\left(r^2\eta^{-2}\log\frac{2nk}{d_{\text{fail}}}\right) = O\left(r^4\tau^{-2}\log\frac{2nkr\rho}{\tau}\right)$$

contexts per candidate. Proposition 4.2's exchange matching and whole-row affinity bound the conditional expected error in a block of duration δ by

$$\delta U\{\chi\eta + r(s-1)/D_q + (r+1)\rho d_{\text{fail}}\}.$$

A failed selection has deficit at most $(r+1)U$ by the structural interval bounds. Multiplying by its conditional failure probability and summing over at most ρ selections gives the last term. Each of the three terms contributes at most $\delta\tau M/16$, so summing durations bounds the local loss by $3\tau M/16$.

Put $\bar{a} = \bar{\gamma}_r := 4/(3+\lfloor\sqrt{8r+1}\rfloor)$ for $\chi = 2$, and $\bar{a} = 1/2$ for the pairwise refinements with $\chi = 3$. Choose $J = \max\{2, \lceil 4\bar{a}/\tau\rceil\}$ and $G \geq 4/\tau$. Apply Lemma C.2 with the corresponding envelope and its exact predicates above. Each coefficient is at most $\bar{a}$, so the clock loss is at most $\tau/2$ and $J = O(1/(\tau\sqrt{r}))$. Execute blocks in decreasing weight order, compressing explicit sparse actions only at completed block boundaries. Terminal collapse and independent support swaps preserve signed expected value by Section C.3.

For finite-bit sampling, a compressed row has at most $s+1$ active labels from its old label and new action; a dense row has at most k. Including inactivity, preprocessing and support swaps, categorical support is bounded by $s_0 = \max\{\rho, s+2, 2\}$ for compressed states and $s_0 = \max\{\rho, k+1, 2\}$ for dense states. The dyadic coupling with $\varepsilon = \tau/8$ therefore adds at most $\tau M/8$. Together with the clock and local errors, the total loss is less than τM.

There are at most $J\rho$ selections, each scanning at most n rows of k marginals. Multiplying by the context count proves the value-query bounds; for $r = 2$ the per-candidate count is $O(\tau^{-2}\log(2nk\rho/\tau))$. Independence scans and support merges cost $O(Jn\rho)$, including preprocessing within these bounds. With streamed sampled rows and support merges, compressed states use $O(nr + k + J)$ numerical records. Dense states use $O(nk)$ records, $O(k^3)$ for the explicit local constraint matrix, and $O(J)$ for the clock. At $r = 2$, the sparse rule uses $O(n + k + \tau^{-1})$ records. □

D Query bounds for the baseline continuous coefficient

This appendix proves Theorem 6.2. Fresh samples at the current state permit block-boundary label compression. The two-label interface attains $\alpha_r - \tau$ for every $k \geq r \geq 2$; Theorem C.3 supplies the separate full-row budgets for the stronger extra-label coefficients. Use the notation of Appendix C. At a fixed prefix A, the known cap $m_e = \max_i \Delta_{e,i}F(A)$ bounds every compatible observation by $-(r-1)m_e \leq X_{ei} \leq m_e$. Its mean row lies in $\mathcal{C}_r^k$, so $\mu_e = \max_i \mathbb{E}X_{ei} \in [0, m_e]$. Zero caps need no samples. Every observation below uses fresh randomness conditional on the current history, and all budgets are deterministic.

D.1 Truncated means and selection charged to the returned cap

Lemma D.1. *Let $X \in [-(r-1)m, m]$, $m > 0$, and $g = \mathbb{E}X$. For $0 < \epsilon \leq 1/4$ and $0 < p < 1$, put $\eta_0 = \epsilon/(6+\epsilon)$ and take $N \geq \lceil r\eta_0^{-2}\log(2/p)\rceil$ independent observations. Writing $\widehat{d}$ for their mean after transformation $Z = 1 - X/m$, the estimate*

$$U = m\max\left\{-1,\ 1 - \max\left(0, \frac{\widehat{d}-\eta_0}{1+\eta_0}\right)\right\} \tag{67}$$

satisfies $\max(g, -m) \leq U \leq \max(g, -m) + \epsilon m$ with probability at least $1-p$.

Proof. Normalize $m = 1$ and write $d = \mathbb{E}Z$. Since $0 \le Z \le r$, $\mathrm{Var}(Z) \le rd$. Bernstein's inequality gives

$$\Pr\{|\widehat{d} - d| > \eta_0(d+1)\} \le 2\exp\left[-\frac{N\eta_0^2(d+1)^2}{2rd + (2/3)r\eta_0(d+1)}\right] \le 2e^{-N\eta_0^2/r}.$$

The last inequality uses $2d/(d+1)^2 \le 1/2$ and $2\eta_0/[3(d+1)] \le 1/2$. On the complementary event, the expression before truncation at -1 dominates $g = 1 - d$, so $U \ge \max(g, -1)$. If $U = -1$, equality holds. Otherwise $\widehat{d} < 2 + 3\eta_0$ and $d < (2 + 4\eta_0)/(1 - \eta_0)$. For $\widehat{d} \ge \eta_0$ this gives $U - g \le 2\eta_0(d+1)/(1+\eta_0) \le 6\eta_0/(1-\eta_0) = \epsilon$. For $\widehat{d} < \eta_0$, instead $U - g = d < 2\eta_0/(1-\eta_0) \le \epsilon$. Replacing g by $\max(g, -1)$ only reduces this excess. □

Median elimination applies to the clipped means $v_i = \max(g_i, -m)$, using Lemma D.1 rather than an unbiased estimate of v_i. To select within am of their maximum, round h uses accuracy $b_h m/2$ and failure $d_h/4$ per surviving label, where $b_h = (a/4)(3/4)^h$ and $d_h = p/2^{h+1}$. The best survivor is not underestimated except with probability $d_h/4$. If every $b_h m$-optimal label is discarded, at least half the labels were overestimated by more than $b_h m/2$. Markov's inequality bounds this event by $d_h/2$. Summing the losses and failures proves the selection guarantee, following the median-elimination argument of Even-Dar et al. (2006). The observation count is $O(kra^{-2}\sum_{h\ge 0}(8/9)^h[\log(2/p) + h + 1]) = O(kra^{-2}\log(2/p))$. Odd survivor counts change only an absolute constant. Selecting to half the desired accuracy and independently estimating the selected clipped mean estimates the maximum with the same order of samples. The true row maximum is nonnegative, so truncation does not change it; the final estimate may be clipped to $[0, m]$.

The next comparison charges an error to the cap of the returned element, which may be much smaller than the largest available cap.

Lemma D.2 (Asymmetric comparison)**.** *Suppose $0 \le \mu_H \le H$, $0 \le \mu_L \le L \le H$, $0 \le c < 1$, and $0 < e < 1$. Put $z = (c+e)/(1+c)$ and $p = ez/(1+c)$. If $L \le zH$, return the high-cap candidate. Otherwise estimate both means to error $eH/2$ with joint failure at most p, clip to their cap intervals, and return the low-cap candidate exactly when $\widehat{\mu}_L - \widehat{\mu}_H > c(H - L) + eH$. The returned candidate J satisfies*

$$\mathbb{E}[\mu_J + (c+3e)m_J] \ge \max\{\mu_H + cH, \mu_L + cL\}.$$

Proof. Let the maximum on the right be V. On success, a low-cap decision gives $\mu_L - \mu_H > c(H - L)$; a high-cap decision gives $\mu_L - \mu_H \le c(H - L) + 2eH$. Thus $\mu_J + (c+2e)m_J \ge V$ in either case. The observation-free case satisfies the same inequality because $(1+c)L \le (c+e)H$. On every outcome the returned cap is at least zH, while $V \le (1+c)H$ and the returned augmented mean is nonnegative. Its expected deficit is at most $p(1+c)H$. The added em_J pays this deficit since $ezH = p(1+c)H$. When $H = 0$, both means vanish and either candidate suffices. □

Lemma D.3 (Returned-cap selection)**.** *For N_c candidates and $0 < \epsilon_s < 1$, a selector returns E with*

$$\mathbb{E}[\mu_E + \epsilon_s m_E] \ge \max_e \mu_e, \qquad Q_{\mathrm{select}} = O\left(N_c kr\epsilon_s^{-2}\log\frac{2}{\epsilon_s}\right). \tag{68}$$

It uses $O(N_c + k)$ numerical records. Its guarantee is conditional on the prior history, and the budget holds on every outcome.

Proof. Use a fixed balanced bracket. At level h, let $e_h = (\epsilon_s/12)(3/4)^h$, $c_0 = 0$, and $c_{h+1} = c_h + 3e_h$. Apply Lemma D.2 to adjacent representatives, advancing an unpaired one without observations. Conditional on the representatives, the lemma bounds the expected augmented winner by the greater augmented input. The expectation of a maximum dominates the maximum of expectations; an unpaired representative gains a nonnegative cap coefficient. Induction therefore bounds each representative's expected augmented value below by the largest original mean in its subtree. Since $c_h < \epsilon_s$, the final representative satisfies the assertion.

There are $O(N_c/2^{h+1})$ matches at level h, including odd brackets: $N_{h+1} = \lceil N_h/2 \rceil$ and every nonempty level has $2^h < N_c$. The maximum estimator above uses normalized accuracy at least $e_h/2$

because each cap is at most H. Its failure logarithm is $O(\log(2/\epsilon_s) + h + 1)$, so total observations are at most a constant times

$$N_c k r \epsilon_s^{-2} \sum_{h \geq 0} (8/9)^h [\log(2/\epsilon_s) + h + 1].$$

Requested normalized accuracies above $1/4$ may use the $1/4$ routine; an absolute tolerance at least half the cap needs only the cap midpoint. The bracket stores $O(N_c)$ representatives and one maximum-mean call stores $O(k)$ entries; observations are accumulated and discarded. □

D.2 Two approximately maximal labels suffice

Lemma 6.1 supplies the two-label protection inequality. One distinguished label would not suffice: take $r = 3$, $k = 4$, $g = (1, 1, 0, 0)$, $T = \{1, 3, 4\}$ and $q = \mathbf{e}_1$. The claimed charge would be zero, but $v = (0, 1, 0, 0)$ with target two has loss one.

At a selected element with cap $m > 0$, find a label i within ηm of the largest clipped mean and then a distinct label j within ηm of the largest remaining clipped mean, each with failure probability $d/4$. Add any $r - 2$ unused labels. Independently estimate each of these r clipped means by Lemma D.1, to accuracy ηm and failure $d/(2r)$, obtaining U_l. Set $b_l = (U_l + 3\eta m)_+$ and use the shifted-prefix oracle on b. Round all but one supported probability down to denominator D_q and give the remainder to that supported coordinate. A zero surrogate row uses a fixed legal label.

Lemma D.4 (Sampled action)**.** *Let $V_e = \nabla_e \Phi_\lambda(x) \cdot q_e$ be the fixed-weight potential derivative along the sampled action at element e. For $0 < \eta \leq 1/4$, the preceding rule satisfies*

$$\mathbb{E}[V_e \mid E = e, \mathcal{H}] \geq \mathcal{K}_r(\lambda)\mu_e - \left(4\eta + \frac{r^2}{D_q} + (r+1)d\right) m_e, \tag{69}$$

using $O(kr\eta^{-2}\log(2r/d))$ scalar observations. Here $0 \leq \lambda \leq \beta_r$ and $\mathcal{H}$ fixes the current state.

Proof. The simultaneous success event has probability at least $1 - d$. For $r = 2$, no mean lies below $-m$. For $r \geq 3$, the largest mean outside any chosen label i is at least $-g_i/(r-1) \geq -m/(r-1) \geq -m/2$. With $\eta \leq 1/4$, neither approximately maximal returned clipped mean equals its truncation floor $-m$. Thus both searches approximate true means. After raising, both distinguished coordinates dominate every unselected true mean. Also $b_l \geq g_l$ on T. If $g_l \leq -m$, success gives $U_l \leq -m + \eta m$, hence $b_l = 0$. Every positive action coordinate therefore satisfies $0 \leq b_l - g_l \leq 4\eta m$.

Raise only the coordinates in T in a proof-only full row, leaving the others unchanged. This row remains in the cone, dominates the actual row and satisfies Lemma 6.1; its maximum on T is at least μ_e. The sufficient envelope of Lemma B.3 consequently gives $V_e \geq \mathcal{K}_r(\lambda)\mu_e - 4\eta m_e$ before quantization. At most r/D_q mass moves under quantization. The coefficients of the action in the potential derivative lie in an interval of length rm_e, so its additional loss is at most $r^2 m_e/D_q$. On failure every legal action still has $V_e \geq -rm_e$, while $\mathcal{K}_r(\lambda)\mu_e \leq m_e$. This proves (69). Two searches cost $O(kr\eta^{-2}\log(2/d))$ observations; the r final estimates cost $O(r^2\eta^{-2}\log(2r/d))$, absorbed since $r \leq k$. □

Combining Lemmas D.3 and D.4 charges a selected cap with

$$e(\lambda) = \mathcal{K}_r(\lambda)\epsilon_s + 4\eta + r^2/D_q + (r+1)d. \tag{70}$$

In particular, the coefficient on selection error remains $\mathcal{K}_r(\lambda)$, rather than being replaced by one. The exchange partner may depend on the returned element and later updates: its block-end maximum is pathwise bounded by the largest addable mean at the pre-selection history. Multiplying (68) by the fixed nonnegative $\mathcal{K}_r(\lambda)$ therefore proves the conditional comparison needed for sequential basis exchange.

D.3 Deterministic scale guesses and signed preprocessing

Let D be the greedy value from Section C.1, so $M/(r+1) \leq D \leq M$. If $D = 0$, return $\mathbf{0}$. Otherwise run a separate continuation for each $V_j = 2^j D$, $0 \leq j \leq \lceil \log_2(r+1) \rceil$, and retain the best feasible output. The first $V_j \geq M$ is a fixed index j_* for the input and satisfies $M \leq V_{j_*} < 2M$.

For each guess V, run the preprocessor of Lemma C.1 with fixed comparison weight $0 \le a \le \beta_r < \gamma_r$. At j_* it gives $\mathbb{E}\Phi_a(0) \ge aM$ and $U(A) \le 3V_{j_*} < 6M$. Every branch's preprocessor uses at most ρ insertions, $O(nk\rho)$ value calls and $O(n\rho)$ independence calls. All branch outputs are feasible, so their maximum dominates the fixed valid branch.

D.4 A CLOCK WITH BOUNDED CUMULATIVE SELECTION WEIGHT

Let $\bar\beta_r = 2/(2 + \lfloor\sqrt{2r}\rfloor) \ge \beta_r$, and put $c_- = 1/[2(1 + \lceil\sqrt{r}\rceil)]$. Choose integers $J \ge 4\bar\beta_r/\tau$, $G \ge 4/\tau$ and $H \ge \max\{J^3, J/c_-\}$. Apply Lemma C.2 with $K = \mathcal{K}_r$ and place unused time at the largest weight, $\ell_J = a$.

Lemma D.5. *The resulting positive rational durations sum to one and satisfy*

$$\ell_j - \ell_{j-1} \le \delta_j \mathcal{K}_r(\ell_j), \qquad \alpha_r - \tau/2 \le a \le \alpha_r, \qquad W := \sum_j \delta_j \mathcal{K}_r(\ell_j) \le 4\alpha_r. \tag{71}$$

Taking the least admissible integers gives $J = O(1/(\tau\sqrt{r}))$. Clock construction uses no objective queries.

Proof. The common clock gives $a \ge \alpha_r - \alpha_r/J - 1/G \ge \alpha_r - \tau/2 > \alpha_r/2$ and the first inequality in (71). For $c_r = 1/(1+\sqrt{r})$, the first-branch equality-row candidate in Lemma B.3 gives, throughout the positive interval, $\mathcal{K}_r(x) \le (1-x)^2/[1 + (r-1)x] \le 1/[1 + (r-1)x]$. Since $\alpha_r \ge c_r$ and $c_r[1 + (r-1)c_r/2] = 1/2$, we have $\mathcal{K}_r(a) \le 2c_r \le 2\alpha_r$. As $J/H \le c_- \le \alpha_r/2$, (64) yields $W \le a + J/H + \mathcal{K}_r(a) \le 7\alpha_r/2 < 4\alpha_r$. The rational envelope predicates in Section C.2 implement all integer comparisons. Finally $\bar\beta_r = \Theta(r^{-1/2})$ and $\tau < \alpha_r < \beta_r$, giving the stated block count. □

D.5 GLOBAL ERROR, FINITE RANDOMNESS AND QUERY ACCOUNTING

Proof of Theorem 6.2. Use the same clock and parameters in every scale continuation:

$$\epsilon_s = \frac{\tau}{384\bar\beta_r}, \qquad \eta = \frac{\tau}{384}, \qquad d = \frac{\tau}{96(r+1)}, \qquad D_q \ge \frac{96r^2}{\tau}, \tag{72}$$

where D_q is dyadic. They lie in the estimator domains because $\tau < \alpha_r < \beta_r \le \bar\beta_r$. On the fixed valid branch from Section D.3, preprocessing gives $\mathbb{E}\Phi_a(0) \ge aM$ and the pathwise cap-sum bound $U \le 6M$. Construct each direction base sequentially with Lemmas D.3 and D.4. Exchange matches every selected element to a distinct original comparison element that was addable at selection. Whole-row affinity and (70) bound a block's expected deficit by $\delta_j U\, e(\ell_j)$ after the clock pays its endpoint weight drop. The selection bound is conditional before selecting the element; the label bound is conditional after selection. This order also covers adaptive comparison partners.

Summing durations and retaining their envelope weights gives

$$\sum_j \delta_j\, e(\ell_j) = \epsilon_s W + 4\eta + r^2/D_q + (r+1)d \le \tau/96 + \tau/96 + \tau/96 + \tau/96 = \tau/24.$$

The total expected local loss is at most $\tau M/4$. Compress labels independently after complete blocks if desired; row affinity preserves both potential components in conditional expectation. After each compression, subsequent estimates use fresh samples conditional on the updated state. The fractional support is the weighted sum of its direction bases, with positive weights δ_j. Independent mean-preserving base merges and terminal label choices preserve the expected signed submodular value by Section C.3. Retain the empty, greedy and prefix labelings, with their evaluated values.

Fix a deterministic upper bound on all categorical draws in one continuation, including preprocessing, sampled product contexts, compression and support swaps. The bounded dyadic sampler of Section C.3 couples them to ideal draws with total discrepancy at most $\tau/4$. Both retained outputs lie in $[0, M]$, so this costs at most $\tau M/4$. It follows that the finite valid branch has expected gain at least $(\alpha_r - \tau/2 - \tau/4 - \tau/4)M$. The final feasible maximum dominates the fixed valid branch j_* on every outcome, so $\mathbb{E}F(Y) \ge \mathbb{E}F(Y_{j_*}) \ge (\alpha_r - \tau)M$.

Since $r\bar\beta_r^2 = O(1)$, the per-insertion selection cost is $O(nk\tau^{-2}\log(2\bar\beta_r/\tau))$. The label cost is $O(kr\tau^{-2}\log(2r/\tau))$, so their sum is $O(k(n+r)\tau^{-2}\log(2r/\tau))$. Every scalar observation costs

two original value queries in a fresh product context. There are at most $J\rho$ insertions per branch and $O(\log(r+1))$ branches. The initial reference, all preprocessors and final value comparisons cost $O(nk\rho\log(r+1))$, absorbed in (21). Candidate scans, preprocessing and support merges use $O(Jn\rho\log(r+1))$ independence calls. These counts hold on every statistical outcome and prove the theorem. □

Streaming the branches and support merges uses $O(nr+k+J)$ numerical records with block-boundary label compression. At most $r+1$ positive labels occur in a row during a block, so a product-context draw has at most $r+2$ outcomes including inactivity. State coordinates have denominator dividing HD_q, since durations are multiples of $1/H$ and compression preserves row sums. The common sampling and arithmetic conventions are in Section C.3. Constructing an explicit product context costs $O(n)$ arithmetic per observation beyond the value-query count.

E Coefficient gains and finite certificates

Let $\mathcal{K}_r, \nu_r, \beta_r$ be the envelope, switch and endpoint in Appendix B.2, and let α_r be its unit-time coefficient. Define

$$c_r = \frac{1}{1+\sqrt{r}}, \quad A_r = (1+\sqrt{r})\alpha_r, \quad \Gamma_r = A_r - 1, \quad T_r(x) = \int_0^x \frac{dt}{\mathcal{K}_r(t)}.$$

These scalar formulas extend to real $r \geq 2$.

E.1 Monotonicity and limit of the relative gain

Integrating the reciprocal first branch gives

$$\theta_r := T_r(\nu_r) = \frac{r\nu_r}{1-\nu_r} + (r-1)\log(1-\nu_r). \tag{73}$$

The switch identity $(2r-3)\nu_r^2 + 4\nu_r - 1 = 0$ implies

$$1 + (r-1)\nu_r = \frac{(1-\nu_r)^2}{2\nu_r}, \qquad \mathcal{K}_r(\nu_r) = 2\nu_r. \tag{74}$$

Since this branch is strictly decreasing and positive, $0 < \theta_r < \nu_r/\mathcal{K}_r(\nu_r) = 1/2$.

Theorem E.1. *The relative coefficient $A_r = (1+\sqrt{r})\alpha_r$ is strictly increasing on $[2,\infty)$, and*

$$\lim_{r\to\infty} A_r = C := \sqrt{2} - \frac{e^{-3/2}}{\sqrt{2}} \approx 1.2564.$$

In particular, $\alpha_r - c_r \sim (C-1)/\sqrt{r}$.

Proof. Put $s = \sqrt{r} \geq \sqrt{2}$ and normalize the envelope by the Random Root coefficient:

$$E_s(u) = (1+s)\mathcal{K}_{s^2}\left(\frac{u}{1+s}\right). \tag{75}$$

Its switch is $a(s) = (1+s)/(2+\sqrt{2s^2+1})$, and its endpoint is $b(s) = (1+s)/(1+s/\sqrt{2})$. Both are strictly increasing, $a(s) < 1$, and $b(s) < \sqrt{2}$. For example,

$$a'(s) = \frac{2\sqrt{2s^2+1}+1-2s}{\sqrt{2s^2+1}(2+\sqrt{2s^2+1})^2} > 0.$$

We show that, at every fixed $0 \leq u < b(s)$, the active formula for $E_s(u)$ strictly increases with s.

On the first branch, $E_s(u) = (s+1-u)^2/[(s+1)(1+(s-1)u)]$, and

$$\partial_s E_s(u) = \frac{(s+1-u)\{s(2u^2-2u+1)+1-u\}}{(s+1)^2(1+(s-1)u)^2} > 0. \tag{76}$$

Indeed, the active branch has $u \leq a(s) < 1$, and $2u^2 - 2u + 1 = 2(u-1/2)^2 + 1/2 > 0$.

On the second branch, write $y = E_s(u)$ and $w = y + 2u$. Squaring its explicit positive-root formula yields

$$Q_s(u, w) := w^2 + \left(\frac{8}{s} + \frac{8(1-u)}{s^2}\right) w - 8 - \frac{16}{s} + \frac{8(u^2-1)}{s^2} = 0. \tag{77}$$

The selected w is the positive root. Since $0 < u < b(s) < \sqrt{2}$, $Q_s(u, 2) = -4 + 8(u-1)^2/s^2 < 0$, so $w > 2$. Moreover, $\partial_w Q_s = 2w + 8(s+1-u)/s^2 > 0$, and the derivative at fixed u, w factors as

$$\partial_s Q_s = -\frac{8}{s^3}\left[(s+2-2u)(w-2) + 2(u-1)^2\right] < 0. \tag{78}$$

The sign follows from $s + 2 - 2u > 2 - \sqrt{2} > 0$ and $w > 2$. Implicit differentiation therefore gives $\partial_s E_s(u) = \partial_s w = -(\partial_s Q_s)/(\partial_w Q_s) > 0$. The two active formulas agree at the moving switch. Since $a(s)$ is strictly increasing, a fixed u crosses at most one switch as s grows; continuity and the positive derivatives on both sides prove the pointwise strict comparison throughout the common positive domain.

Now the normalized clock is $\widetilde{T}_s(u) = \int_0^u dv/E_s(v)$, with $\widetilde{T}_s(A_{s^2}) = 1$. For $s_2 > s_1$ and $0 < u < b(s_1)$, the preceding pointwise comparison yields $\widetilde{T}_{s_2}(u) < \widetilde{T}_{s_1}(u)$. Apply this at $u = A_{s_1^2} < b(s_1)$. The new clock is still below one, so its unique unit-time point satisfies $A_{s_2^2} > A_{s_1^2}$.

For the limit, the distance from α_r to the envelope zero β_r is controlled by a chord and a tangent. Put $d_r = \beta_r - \nu_r$ and $\sigma_r = 2\nu_r/d_r$. On the second branch, let $P_r(x) = 2(1-x)(r + D_r(x))$. Differentiation gives

$$\mathcal{K}_r''(x) = -\frac{8r}{P_r(x)^{3/2}} < 0, \qquad \mathcal{K}_r(\beta_r) = 0, \qquad \mathcal{K}_r'(\beta_r) = -2.$$

Strict concavity therefore yields, for $\nu_r < x < \beta_r$,

$$\sigma_r(\beta_r - x) < \mathcal{K}_r(x) < 2(\beta_r - x). \tag{79}$$

In particular, the envelope is positive on this interval and $T_r(x) \to \infty$ at its right endpoint. Together with $\theta_r < 1/2$, this proves $\nu_r < \alpha_r < \beta_r$. Integrating (79) up to α_r gives

$$\sigma_r(1-\theta_r) < \log\frac{d_r}{\beta_r - \alpha_r} < 2(1-\theta_r). \tag{80}$$

Put $h = r^{-1/2}$. The explicit switch and endpoint formulas give $\nu_r/h \to 1/\sqrt{2}$, $\beta_r/h \to \sqrt{2}$, $d_r/h \to 1/\sqrt{2}$, and $\sigma_r \to 2$. Moreover,

$$\theta_r = h^{-2}\left[\frac{\nu_r}{1-\nu_r} + \log(1-\nu_r)\right] - \log(1-\nu_r) \longrightarrow \frac{1}{4},$$

since the bracket is $\nu_r^2/2 + O(\nu_r^3)$. Thus (80) implies $(\beta_r - \alpha_r)/d_r \to e^{-3/2}$ and hence $\alpha_r/h \to C$. Since $c_r/h \to 1$, both asserted limits follow. □

Consequently,

$$A_2 \le A_r < C, \qquad A_2 = (1+\sqrt{2})\alpha_2, \tag{81}$$

with equality on the left exactly at $r = 2$. Thus A_2 is the largest uniform multiplier in $\alpha_r \ge A_2/(1+\sqrt{r})$.

For a finite algorithm with coefficient $\alpha_r - \tau$, the relative gain above c_r is

$$\Gamma_r - (1+\sqrt{r})\tau. \tag{82}$$

Thus setting $\tau = (1-\xi)(\alpha_r - c_r)$, for $0 < \xi < 1$, retains the fraction ξ of the baseline improvement. Combining $\alpha_r - c_r = \Theta(r^{-1/2})$ with Theorem 6.2 gives $\widetilde{O}(k\rho(n+r)r/(1-\xi)^3)$ value queries. Here the tilde suppresses logarithmic factors in r and $(1-\xi)^{-1}$.

E.2 NUMERICAL COEFFICIENT BOUNDS

Corollary E.2. *The baseline and extra-label coefficients satisfy*

$$0.421795 < \alpha_2 < 0.421796, \qquad 1.018304 < A_2 < 1.018305,$$
$$\alpha_2^+ > 0.4366, \qquad \beta_{2,\infty} > 0.4259.$$

The minimum relative gain is therefore approximately 1.83%.

Proof. For the two-label endpoint, put $\nu = \sqrt{5} - 2$ and $\theta = (\sqrt{5} - 1)/2 + \log(3 - \sqrt{5})$. On the second branch, $\mathcal{K}_2(x) = 2\sqrt{2(1-x)} - 2$. With $w = \sqrt{2(1-x)}$, elementary integration gives

$$T_2(x) - \theta = \frac{1}{2}\left[\sqrt{5} - 1 - w + \log\frac{\sqrt{5} - 2}{w - 1}\right].$$

Rational enclosures of the square roots and logarithms give

$$T_2(0.421795) < 1 < T_2(0.421796), \qquad T_2\left(\frac{1.018304}{1+\sqrt{2}}\right) < 1 < T_2\left(\frac{1.018305}{1+\sqrt{2}}\right).$$

Strict monotonicity of T_2 proves both intervals. The accompanying certificate verifies these inequalities directly from the displayed formula.

For either $K = \mathcal{K}_2^+$ or $K = Q_2$, start at $\ell_0 = 0$ and take the largest multiple of $1/G$ in the positive domain of K satisfying

$$\ell_{j+1} - \ell_j \le K(\ell_{j+1})/J, \qquad j = 0, \dots, J-1.$$

The previous level is feasible, so the levels are nondecreasing. Monotonicity of K gives

$$\int_0^{\ell_J} \frac{dx}{K(x)} = \sum_{j=0}^{J-1} \int_{\ell_j}^{\ell_{j+1}} \frac{dx}{K(x)} \le \sum_{j=0}^{J-1} \frac{\ell_{j+1} - \ell_j}{K(\ell_{j+1})} \le 1.$$

Thus ℓ_J is a lower bound on the unit-time coefficient.

The exact adjacent-label predicate in Appendix C gives the first endpoint below. For Q_2, additionally impose $1 - 2\ell_{j+1} \ge J(\ell_{j+1} - \ell_j)$:

$$K = \mathcal{K}_2^+, \quad J = 8192, \quad G = 2^{32}: \qquad \ell_J = \frac{1875261745}{4294967296} > 0.4366,$$
$$K = Q_2, \quad J = 16000, \quad G = 2^{40}: \qquad \ell_J = \frac{234144369483}{549755813888} > 0.4259.$$

Both computations use rational comparisons and integer bisection; the accompanying verifiers recompute the stated numerators. □