RISK IDENTIFICATION AND EVIDENCE COMPLEXITY FOR VERIFIER-GUIDED SELECTION

Anonymous authors
Paper under double-blind review

ABSTRACT

A fixed verifier-guided selector can have different risks under routing laws with identical low-order statistics. We study the cost of estimating its worst compatible risk when each paid query restores a retained state and returns a fresh bounded loss. The feasible domain contains all nonnegative routing laws with the specified moments. For top-b selection and moments through order r, the worst-case record count scales as k^b while the unseen mass is bounded by $O(k^{-d})$, where $d = \lfloor r/2 \rfloor$. We construct one positive law with prescribed heavy records at every geometric scale, uniformly in the number of scales and bounded product biases. The resulting acquisition rates match away from $b = 2d$ and at critical breadth when $\log(1/\delta)$ dominates $\Lambda = 1 + \log(1/\epsilon)$. At fixed confidence, the general critical bounds are $\Omega(\epsilon^{-2}\Lambda^2)$ and $O(\epsilon^{-2}\Lambda^3)$; the balanced top-two lower bound improves by a factor $\log \Lambda$. The lower bound concerns scalar output. The upper algorithm learns a near-optimal law. A specified mean family admits cheaper scalar estimation at fixed confidence. We also characterize the full space of exactly identified losses. A supplied distance bound to this space reduces the effective query depth even when the approximating loss is unknown, under the same paid observation interface.

1 INTRODUCTION

Verifier-guided search scores candidate continuations and retains only a few of them. When two verifiers are available, a routing mechanism chooses which one scores each candidate family. The downstream risk depends on how those choices vary together. Individual routing frequencies, and even all pairwise frequencies, need not determine it.

We study an auditor that knows both scoring alternatives and the released routing statistics. It can restore any legal retained state and purchase a fresh continuation loss. The target is the worst risk over all routing laws consistent with the release. This separates two sources of uncertainty. Missing routing dependence leaves several compatible risks. Finite continuation data make their endpoints costly to estimate. Repeated observations reduce the second uncertainty while preserving the first.

Our central question is how this evidence cost grows as the desired error decreases and the number of candidates increases. Deep records have small feasible mass but many possible identities. At a critical balance, many depth scales contribute comparable difficulty. We construct one nonnegative routing law that places prescribed heavy records at every scale while preserving all released moments. This simultaneous construction yields scalar lower bounds on the complete feasible domain. Together with an adaptive upper algorithm, it gives matching rates away from the critical regime and gives the critical bounds at fixed confidence. A comparison within the same domain shows why estimating one risk value can be cheaper than learning a near-optimal routing law. A final result isolates the part of an unknown continuation that creates the costly tail.

1.1 RELATED WORK

Moment-based partial identification studies the range of target values compatible with observed restrictions (Manski, 2003). Finite-moment optimization expresses such extrema as linear programs over nonnegative measures (Prékopa, 1990). Here the target is a downstream loss after a fixed selector. The moment law is first projected through the retained record. This projection determines

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.

both which risks are compatible and where paid observations can be obtained. The risk interval uses the classical moment framework; its selector-specific geometry is the input to our acquisition analysis.

The distinction between global consistency and local restrictions also appears in graphical-model marginal polytopes (Wainwright & Jordan, 2008) and disjunctive formulations (Balas, 1985). We adapt these representations by completing each retained-record fiber with a nonnegative conditional law. The exact affine basis describes which continuation losses are identified by the released moments. Positivity imposes further constraints, as the finite separating certificate demonstrates. These roles separate the transferred optimization machinery from the particular selector domain.

Limited independence bounds the mass of rare Boolean events (Schmidt et al., 1995; Benjamini et al., 2012). Sharp all-ones probabilities are studied by Peled et al. (2011); Berend et al. (2026). Decision-tree arguments show why low-order information determines shallow leaves (Ben-Eliezer et al., 2009; Hatami & Hoza, 2024). The probability construction here requires prescribed records at many depths to coexist under one law. Its constants are uniform in the number of scales and in bounded product biases. This simultaneous requirement links rare-event constructions to the critical evidence lower bound.

Approximate independence can be repaired to exact independence (Alon et al., 2003; O'Donnell & Zhao, 2018). Our construction combines positive moment repair with a cancellation of the leading contributions from all target cylinders. The remaining mixed-degree error is summable uniformly over scales. Gaussian quadrature (Golub & Welsch, 1969) supplies the classical moment identities; the joint cancellation and scale-uniform control make them suitable for the prescribed positive selector laws.

For an unconstrained distribution on finitely many records, the upper risk is the largest record mean. Best-mean estimation is therefore a natural statistical comparison. Osogami et al. (2025) obtain confidence-asymptotic optimality for Gaussian arms at fixed arm count and accuracy. Our rates vary accuracy and candidate count at fixed confidence. Decoupled exploration (Rouyer & Seldin, 2020), adaptive change of measure (Kaufmann et al., 2016), and mixture information (Lee & Valiant, 2021) supply sampling and testing tools. We adapt them to the actual retained-record response and combine them with the simultaneous positive witnesses. The distinction between a scalar estimate and a returned feasible law is then quantified on one specified mean family inside the complete domain.

2 Problem Background: From Selection to a Risk Audit

Consider three candidate families competing for two continuation slots. Each family has two known scoring alternatives. Bit $X_i = 0$ chooses its lower score l_i; $X_i = 1$ chooses its higher score u_i. The selector keeps the two greatest chosen scores. Its retained record $T(X)$ stores their identities and continuation states. Figure 1 shows one such selector. If the first two bits are high, the third bit cannot change the retained record; both assignments give the leaf $11*$.

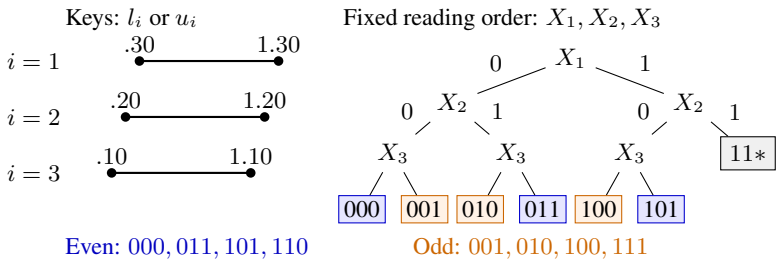

Figure 1: A balanced assignment audit with three families and two survivors. Each segment joins two possible scores. The tree groups assignments by their retained record; $11*$ allows either value of the third bit.

A routing law P assigns probabilities to the eight bit vectors. For each retained record t, let $\theta_t \in [0,1]$ be the mean downstream loss after continuing from that state. Grouping assignments that leave the same record gives

$$\begin{aligned} q_t &= \Pr_P\{T(X) = t\} = \sum_{x:T(x)=t} P(x), \\ R(P,\theta) &= \sum_x P(x)\theta_{T(x)} = \sum_t q_t\theta_t. \end{aligned} \tag{1}$$

Thus routing determines the weights q_t; continuations determine the losses θ_t.

Suppose the release says that every bit is fair and every pair is independent. These are the moment constraints

$$\mathbb{E}_P X_i = \tfrac{1}{2}, \qquad \mathbb{E}_P X_i X_j = \tfrac{1}{4} \quad (i \neq j).$$

They permit uniform laws on assignments with an even or odd number of high bits, denoted by P_+ and P_-. If the continuation fails only at record 000, then

$$R(P_+,\theta) = P_+(000) = \tfrac{1}{4}, \qquad R(P_-,\theta) = P_-(000) = 0.$$

In fact, every compatible law has the form $P_a(x) = 2^{-3}[1 + a(-1)^{x_1+x_2+x_3}]$, $-1 \le a \le 1$. Its risk is $(1+a)/8$, so the full compatible interval is $[0, 1/4]$. This is uncertainty about dependence, even when all losses are known.

In a live audit those losses are unknown. A paid query restores one requested record and observes a fresh bounded loss. Repetition estimates θ_t; it supplies no missing routing moments. Conversely, releasing the complete routing law fixes q but still leaves the losses to learn. Finite-order independent hash families provide one way to implement a moment-constrained router (Thorup & Zhang, 2012). The audit concerns the information released about that router. We now specify its compatible laws and the observations used to estimate their risk endpoints.

For the general audit, let K be the number of candidate families, $[K] = \{1, \dots, K\}$, and $X \in \{0,1\}^K$ their routing bits. Family i has known keys $l_i \le u_i$. The selector keeps the b greatest chosen keys. A key consists of its score and a fixed identity priority, ordered lexicographically to resolve ties. For each family and bit, the continuation state is fixed before routing. The record $T(X) \in \mathcal{T}$ stores the retained identities, bits and states in order, with their common context. A fixed kernel K generates the downstream outcome Y: $\mathcal{L}(Y \mid X = x) = \mathsf{K}(\cdot \mid T(x))$. Hence $\theta_t = \mathbb{E}_{\mathsf{K}(\cdot|t)}[\mathrm{loss}(Y)]$ depends only on the record, with $0 \le \mathrm{loss}(Y) \le 1$.

Write ψ for the released statistics, including normalization, and m for their feasible expectations. With $\Delta(A)$ the probability simplex on A, define

$$\mathcal{P}(m) = \{P \in \Delta(\{0,1\}^K) : \mathbb{E}_P \psi(X) = m\}, \qquad \mathcal{Q}(m) = \{T_\# P : P \in \mathcal{P}(m)\}.$$

Here $T_\# P$ is the record distribution q in (1). All its coordinates must come from one nonnegative assignment law. Optimizing the risk over this complete domain gives the sharp identified interval

$$I(m,\theta) = [R_-(m,\theta), R_+(m,\theta)], \qquad R_- = \min_{q \in \mathcal{Q}(m)} q^\top\theta, \quad R_+ = \max_{q \in \mathcal{Q}(m)} q^\top\theta. \tag{2}$$

Compactness gives attaining laws, and their mixtures attain every intermediate risk. The balanced order-r class $\mathcal{U}_r^K$ specifies $\mathbb{E}_P \prod_{i\in A} X_i = 2^{-|A|}$ for all $|A| \le r$. Product moments instead specify $\prod_{i\in A} p_i$ for known $p_i \in [p_0, 1-p_0]$, where $0 < p_0 \le 1/2$. Exact randomization or a finite assignment release can supply these moments. Estimated population moments require their own confidence set and error allowance.

The auditor receives the keys, moments and state-restoration map. Given its full history $\mathcal{H}_{j-1}$, it requests any legal record t_j and buys one response $Z_j \in [0,1]$ with $\mathbb{E}[Z_j \mid \mathcal{H}_{j-1}, t_j] = \theta_{t_j}$. A rare record can be requested directly. A reused saved terminal label is not a fresh response. The unrestricted class permits every fixed $\theta \in [0,1]^{\mathcal{T}}$; it imposes no prior on these means.

An (ϵ,δ)-faithful audit returns $[\widehat{a}, \widehat{z}]$ with

$$\Pr\{\widehat{a} \le R_- \le R_+ \le \widehat{z}, \quad \max(R_- - \widehat{a}, \widehat{z} - R_+) \le \epsilon\} \ge 1 - \delta. \tag{3}$$

Thus ϵ controls estimation error at the endpoints. It does not require the compatible risks themselves to be close. A scalar audit estimates only R_+ to absolute error ϵ. We count paid responses: lower bounds allow adaptive queries and almost-sure stopping with expected cost, while upper bounds give deterministic budgets. The next step is to determine which records can carry enough feasible mass to affect these endpoints.

3 RECORD GEOMETRY AND THE ACQUISITION THRESHOLD

The audit can spend less on records whose compatible mass is small. We therefore need both a record count and a uniform mass bound. Order the families by decreasing upper key. For a record t, a discarded family whose upper key could beat its last survivor must have been low. Its fiber therefore fixes a prefix: $T^{-1}(t) = \{x : x_{[d_t]} = z_t\}$, where d_t is its depth and z_t its fixed pattern. Let $C_b(x)$ return the first b high positions, or all high positions if fewer than b occur, and write $\mathcal{S}_{K,b} = \{S \subseteq [K] : |S| \le b\}$.

Proposition (selector geometry). *For every legal geometry T, let $g_T(S) = T(\mathbf{1}_S)$ and $G_T(t, S) = \mathbf{1}\{g_T(S) = t\}$. Then*

$$T = g_T \circ C_b, \qquad \mathcal{Q}_T = G_T \mathcal{Q}^{\mathrm{can}}_{K,b,r}, \qquad \mathcal{Q}^{\mathrm{can}}_{K,b,r} = (C_b)_{\#}\mathcal{U}^K_r. \tag{4}$$

A query at $g_T(S)$ implements a canonical query at S at the same cost. The canonical map is itself a legal selector. Hence its audit cost equals the worst cost over known legal geometries. For balanced moments, uniform assignment law U, and the bth greatest lower key $l_{[b]}$,

$$D_T := \max_t d_t = \#\{i : u_i \ge_{\mathrm{lex}} l_{[b]}\}, \qquad \sup_{\theta \in [0,1]^{\mathcal{T}}} (R_+ - R_-) = P_U(d_{T(X)} > r). \tag{5}$$

If $c_{\mathcal{Q}}(A) = \sup_{q \in \mathcal{Q}} q(A)$ denotes the capacity of a record set A, then, for $k \ge 4b$,

$$\#\{t : d_t \le k\} \le \sum_{j=0}^{b} \binom{k}{j} = O_b(k^b), \qquad c_{\mathcal{Q}}\{t : d_t > k\} \le C_{b,r} k^{-d}, \quad d = \lfloor r/2 \rfloor. \tag{6}$$

The tail is zero at $k = K$. For prescribed bounded product biases, the same count exponent and tail exponent hold, with constants also depending on the common marginal bound p_0.

After the bth high, every unread upper key is below a retained high. Changing an unread bit leaves the record fixed. Each shallow fiber has its prescribed mass 2^{-d_t}. For $r < K$, the two laws $2^{-K}[1 \pm (-1)^{X_1 + \cdots + X_{r+1}}]$ agree on these shallow fibers and occupy disjoint deep fibers, attaining the ambiguity formula. The count bound enumerates at most b high positions. The tail bound uses the even-moment inequality for limited independence (Schmidt et al., 1995; Benjamini et al., 2012); the appendix gives the selector-specific reductions and exact attaining laws.

A shell ending near depth k_h has $M_h = O(k_h^b)$ possible records and mass bound $w_h = O(k_h^{-d})$. Their squared evidence scale is $M_h w_h^2 = O(k_h^{b-2d})$. It decreases for $b < 2d$, stays constant for $b = 2d$, and increases for $b > 2d$. A lower bound needs actual compatible laws attaining the requisite masses. The next theorem constructs the simultaneous laws at the critical threshold.

4 POSITIVE MULTISCALE LAWS AND ACQUISITION COMPLEXITY

At critical breadth, a lower bound must accumulate evidence from many scales. Separate large-record capacities do not supply one statistical experiment: the selected records must coexist under one feasible law.

Theorem 1 (Simultaneous cylinders under product moments)**.** *Fix $r \ge 2$, $d = \lfloor r/2 \rfloor$ and $0 < p_0 \le 1/2$. There are $B_0, \gamma > 0$, depending only on r, p_0, such that for every integer $B \ge B_0$ and $L \ge 1$ the following holds. A prelude of B bits is followed by blocks of sizes $N_h = B2^h$, $0 \le h < L$. Choose means $p_i \in [p_0, 1 - p_0]$ and any $2d$ specified high positions A_h in each block. Let E_h fix all earlier bits to zero, the current block to exactly those highs, and leave future bits free. One nonnegative probability law P satisfies*

$$\mathbb{E}_P \prod_{i \in A} X_i = \prod_{i \in A} p_i \quad (|A| \le r), \qquad P(E_h) \ge \gamma N_h^{-d} \quad (0 \le h < L). \tag{7}$$

The constants are uniform in L, the selected positions, and the prescribed mean vector. Independent tail extension preserves the result.

For balanced $r = 2, 3$ and dyadic $B \ge 4$, put $\gamma^*_{r,L} = \max_{P \in \mathcal{U}^{B2^L}_r} \min_{h<L} N_h P(E_h)$. Then

$$\gamma^*_{2,L} \downarrow \tfrac{1}{4}, \qquad \gamma^*_{3,L} \downarrow \tfrac{1}{8}, \qquad 0 \le \gamma^*_{r,L} - 2^{-r} \le 2^{1 - \lfloor L/2 \rfloor}. \tag{8}$$

An explicit infinite binary-linear construction has $P(E_h) = 2^{-r}N_h^{-1}$ for every h. Its finite prefixes use at most $2^{r-1}B2^L$ equally weighted seed outcomes. For balanced $r = 4, 5$ and even B, valid explicit parameters are $B_0 = 2^{192}$ and $\gamma = 2^{-24}, 2^{-25}$, respectively.

Each cylinder is contained in a genuine top-$2d$ retained record. The theorem preserves its single-scale polynomial mass at every scale. A mixture of L unrelated constructions would divide each guaranteed mass by L. At orders two and three, nested binary spans give the construction; block-wise permutations realize every prescribed tuple. The optimal simultaneous constants are half the corresponding single-cylinder limits (Berend et al., 2026).

At general order, Gaussian quadrature creates a central moment deficit that cancels the leading moments of all rare cylinders. A decaying frame and mixed-degree summation keep the remaining error uniform in L. A fixed positive count reference repairs this residual away from the target cylinders. Appendix B explains why the leading cancellation is needed and distinguishes this repair from the classical positive-mixture method of O'Donnell & Zhao (2018). Appendix C.4 gives the complete numerical fourth-order construction; its conservative cutoff is an existence bound, separate from the finite-K calculation below.

Write $u = \log(8/\delta)$, $\log_+ x = \max(0, \log x)$, and

$$k_\epsilon = \min\{K, \epsilon^{-1/d}\}, \qquad \Lambda_K = 1 + \log k_\epsilon.$$

A scalar audit estimates R_+ within ϵ; a faithful audit returns the interval in (3). Their minimax expected query costs are denoted by N^{sc} and N^{int}.

Theorem 2 (Finite-size acquisition under a fixed moment specification)**.** *Fix $b \geq 1$, $r \geq 2$, $d = \lfloor r/2 \rfloor$, and $0 < p_0 \leq 1/2$. For every known vector $p \in [p_0, 1-p_0]^K$, prescribe its product moments through order r. For $K \geq \max(b, r)$, sufficiently small $\epsilon > 0$, and $0 < \delta \leq 1/16$, minimax costs over known legal geometries and all continuation means satisfy the following bounds, up to constants depending only on b, r, p_0:*

	lower bound on N^{sc}	*upper bound on N^{int}*
$b < 2d$	$\epsilon^{-2}u$	$\epsilon^{-2}u$
$b = 2d$	$\epsilon^{-2}\Lambda_K^2 u$	$\epsilon^{-2}(\Lambda_K^3 + \Lambda_K^2 u)$
$b > 2d$	$\epsilon^{-2}k_\epsilon^{b-2d}u$	$\epsilon^{-2}k_\epsilon^{b-2d}u$.

(9)

Upper budgets are deterministic and use the original bounded-response conditional-mean interface. Scalar lower bounds hold on a canonical common-overlap geometry with fresh Bernoulli responses and adaptive stopping. For balanced $b = 2$, $r \in \{2, 3\}$, the lower bound strengthens to

$$N^{\mathrm{sc}} \geq c\epsilon^{-2}\Lambda_K^2 u[1 + \log_+(\Lambda_K/u)]. \tag{10}$$

The two columns are comparable: a faithful upper endpoint satisfies $0 \leq \widehat{z} - R_+ \leq \epsilon$ on its success event and hence is a valid scalar estimate. At fixed K, the precision power is always two. When K reaches order $\epsilon^{-1/d}$, the supercritical expression becomes $\epsilon^{-b/d}u$. For pairwise moments this gives precision powers two for top-one, two with critical logarithms for top-two, and three for top-three. The asymptotic powers do not specify a numerical query count at a fixed candidate budget.

The common calculation partitions queried records into blocks of sizes M_h and total-mass bounds w_h. Set

$$Z = \sum_h (M_h w_h^2)^{1/3}, \qquad \mathsf{H} = \sum_h \sqrt{M_h} w_h.$$

Constrained sampling followed by fresh evaluation costs $O(\epsilon^{-2}(Z^3 + u\mathsf{H}^2) + \epsilon^{-1}\mathsf{H}\sqrt{M_{\max}}u)$. At critical breadth, the simultaneous theorem supplies masses $p_h \asymp N_h^{-d}$ among $M_h \asymp N_h^{2d}$ candidate records. Raise one unknown mean per block from a common flat baseline by $a_h = 4\epsilon/(Lp_h)$. Averaging the stopped-information bound over the M_h possible changed records gives

$$\begin{aligned} cu &\leq C \sum_h \frac{a_h^2}{M_h} \mathbb{E}_0 N(\mathcal{C}_h) \\ &= \frac{16C\epsilon^2}{L^2} \sum_h \frac{\mathbb{E}_0 N(\mathcal{C}_h)}{M_h p_h^2} \leq \frac{C'\epsilon^2}{L^2} \mathbb{E}_0 N. \end{aligned}$$

Single-record witnesses suffice away from critical breadth. The extra balanced low-order logarithm instead uses the ratio-four construction and stopped-mixture argument in Appendix D.10. Thus the statistical theorem uses the probability construction together with separate sampling and testing arguments.

For $u \geq \Lambda_K$, the critical bounds match. At fixed confidence and large enough K, put $\Lambda = 1 + \log(1/\epsilon)$. The general critical lower and upper bounds have orders $\epsilon^{-2}\Lambda^2$ and $\epsilon^{-2}\Lambda^3$. The balanced top-two lower bound additionally contains $\log \Lambda$:

$$\Omega(\epsilon^{-2}\Lambda^2 \log \Lambda) \leq N_{2,2}^{\mathrm{can}} \leq O(\epsilon^{-2}\Lambda^3). \tag{11}$$

Here $N_{2,2}^{\mathrm{can}}$ is the scalar cost on the complete canonical pairwise domain and unrestricted mean cube. The upper algorithm first learns a near-optimal feasible law and then evaluates it. At critical breadth, $M_h w_h^2 \asymp 1$ makes its Z^3 term cubic in the number of scales. A scalar estimate need not identify such a law; the following comparison makes this distinction quantitative within the original positive domain.

On $\mathcal{Q} = \mathcal{Q}_{K,2,2}^{\mathrm{can}}$, fix an integer $s_0 \geq 8$, $L = \lfloor \log_2(1/\epsilon)/8 \rfloor$, $s_j = 4^j s_0$, and $K = 2s_L - 1$. Let A_j be the record-measurable event that the first $2s_j - 1$ bits are zero. Block h consists of the next $6s_{h-1}$ bits; let $\mathcal{C}_h$ contain the records with both highs in its relative positions $2, \ldots, 2s_{h-1} + 2$. For these disjoint sets the restricted family is

$$\theta_t = v_t^0 + \sum_{h=1}^{L} \lambda_h \vartheta_t \mathbf{1}_{\{t \in \mathcal{C}_h\}}, \qquad 0 \leq \vartheta_t \leq 1,$$

$$v_t^0 = \tfrac{1}{4} + \frac{\sum_{j=0}^{L} \mathbf{1}_{A_j}(t)}{4(L+1)}, \qquad \lambda_h = \frac{1}{8(L+1)s_{h-1}}.$$

The background v^0 and scales λ_h are known. A quadratic majorant bounds the objective on every assignment; one compatible law simultaneously attains each block's maximizing record. These two certificates give, over the complete domain,

$$R_+(\theta) = R_+(v^0) + \sum_{h=1}^{L} p_h \lambda_h \max_{t \in \mathcal{C}_h} \vartheta_t, \qquad p_j = (2s_j)^{-1}, \quad |\mathcal{C}_h| p_h^2 \asymp 1.$$

Appendix D.12 uses actual record responses to estimate this sum in $O(\epsilon^{-2}L^2 u[1 + \log_+(L/u)]^2)$ queries, whereas returning a near-optimal feasible law costs $\Theta(\epsilon^{-2}(L^3 + L^2 u))$. Let $F = \arg\max_{q \in \mathcal{Q}} q^\top v^0$. The specified λ_h ensure $R_+(\theta) = \max_{q \in F} q^\top \theta$ throughout this family. This equality fails for unrestricted means: the independent fair record law q^U has $q^U(A_j) = 2^{-(2s_j - 1)} < p_j$, so

$$R_+(1 - v^0) \geq 1 - (q^U)^\top v^0 > 1 - R_+(v^0) = \max_{q \in F} q^\top (1 - v^0).$$

This strict inequality makes the role of the mean class explicit: the block-maximum identity holds on the specified family, whereas arbitrary record means require optimization over the full positive domain.

5 Structured continuations without changing the query interface

The preceding bounds price arbitrary record means. To identify which part of a continuation creates the costly tail, we now separate losses whose expectations are fixed by the moments from an unknown residual. All observations remain paid responses at original records.

Fix $b \geq 1$, $r \geq 2$, $K \geq \max(2b, r)$ and $d = \lfloor r/2 \rfloor$. The domain is $\mathcal{Q} = \mathcal{Q}_{K,b,r}^{\mathrm{can}}$. Let $\mathcal{V}$ be the space of record functions g for which $g(C_b(x))$ is a multilinear polynomial of degree at most r. A multilinear polynomial uses each bit with exponent at most one. This is the full identified class:

$$\begin{aligned} R_+(\theta) = R_-(\theta) &\iff (q - q')^\top \theta = 0 \quad (\forall q, q' \in \mathcal{Q}) \\ &\iff \theta \in \operatorname{span}(\mathcal{Q} - \mathcal{Q})^\perp = \mathcal{V}. \end{aligned}$$

Lemma 7 proves the last space identity. The residual model thus measures departure from the full identified class. Write $\|v\|_\infty = \max_t |v_t|$. For a specified number $0 \le \eta \le 1/8$, define

$$\begin{aligned} \Theta_\eta &= \{\theta \in [0,1]^{\mathcal{S}_{K,b}} : \inf_{g \in \mathcal{V}} \|\theta - g\|_\infty \le \eta\}, \\ k_\eta &= \min\{K, \max(1, (\eta/\epsilon)^{1/d})\}, \qquad \Lambda_\eta = 1 + \log k_\eta. \end{aligned} \tag{12}$$

The parameter η specifies a model class; neither an approximating function g nor its coefficients are supplied to the auditor. Thus the guarantees below require this distance bound as input. Let $\mathcal{N}_\eta^{\mathrm{sc}}$ and $\mathcal{N}_\eta^{\mathrm{int}}$ denote the minimax expected costs of an ϵ-accurate upper support value and an ϵ-faithful interval, respectively, uniformly on this class.

Theorem 3 (Residual-dependent acquisition with unknown visible component)**.** *For fixed b, r, sufficiently small $\epsilon > 0$, and $0 < \delta \le 1/16$, put $u = \log(8/\delta)$. There are positive constants depending only on b, r such that, uniformly in K and η,*

$$\begin{aligned} b \ne 2d: &\quad \mathcal{N}_\eta^{\mathrm{sc}} \asymp_{b,r} \mathcal{N}_\eta^{\mathrm{int}} \asymp_{b,r} \epsilon^{-2} u \max\{1, k_\eta^{b-2d}\}, \\ b = 2d: &\quad c_{b,r}\epsilon^{-2}\Lambda_\eta^2 u \le \mathcal{N}_\eta^{\mathrm{sc}} \le \mathcal{N}_\eta^{\mathrm{int}} \le C_{b,r}\epsilon^{-2}(\Lambda_\eta^3 + \Lambda_\eta^2 u). \end{aligned} \tag{13}$$

For $b = 2$ and $r \in \{2, 3\}$, the lower bound is additionally

$$\mathcal{N}_\eta^{\mathrm{sc}} \ge c\epsilon^{-2}\Lambda_\eta^2 u[1 + \log_+(\Lambda_\eta/u)]. \tag{14}$$

The upper bounds have deterministic query budgets and apply to bounded responses with the specified conditional means. The lower bounds use fresh Bernoulli responses and allow adaptive stopping. Rational inputs admit an implementation with an assignment domain of size $2^{\bar{k}}$, where $\bar{k} \le C_{b,r} k_\eta$. In particular, for $\eta \le c\epsilon$ with fixed c, both the query cost and the prefix optimization dimension are independent of K up to constants depending on b, r, c.

The structural component is unknown but has fixed-dimensional observable coordinates. Set $z_i = 1 - X_i$. Its complete basis is

$$\{1\} \cup \{z_A = \prod_{i \in A} z_i : 1 \le |A| \le r,\ \max A \le b + |A| - 1\}.$$

Every basis expectation is fixed by the released moments, and every basis function can be recovered from the retained record. Boolean interpolation on at most 2^{b+r-1} actual anchor records learns an approximating component $\widehat{g}$ with a stability constant depending only on b, r. Its expectation $\mu(\widehat{g})$ is then known exactly. On the paid pilot's confidence event, $\|\theta - \widehat{g}\|_\infty \le C_{b,r}(\eta + \epsilon)$. Because every feasible law has the same expectation of $\widehat{g}$, the omitted tail costs only $C_{b,r}(\eta + \epsilon)k^{-d}$. The proof samples the residual by a bounded affine transformation of one actual response, without assuming that a small residual is noiseless. The matching lower bounds use genuine alternatives inside Θ_η.

For $\eta \le \epsilon/4$, an explicit implementation avoids optimization altogether. Sample $n = \lceil 8\epsilon^{-2}\log(2/\delta)\rceil$ fresh losses under independent fair assignments and average them to $\widehat{v}$. The interval $[\widehat{v} - \epsilon/4 - 2\eta, \widehat{v} + \epsilon/4 + 2\eta] \cap [0,1]$ is faithful for every K. Its proof uses that the fair-law risk lies between the extrema and within 2η of each. At $\eta = 0$, both extrema coincide although their common value remains unknown. A different structured class has a parameter rate even when its number of unknown coefficients grows with K. In the canonical common-overlap geometry, order the lower keys by index and let $\mathcal{L}(S)$ be the first $b - |S|$ indices in $[b] \setminus S$. Suppose a continuation draws one retained family/state pair uniformly and returns a fresh loss, without revealing the chosen pair. With unknown $v_{i,x} \in [0,1]$, its means are

$$\theta_S = \frac{1}{b}\left(\sum_{i \in S} v_{i,1} + \sum_{i \in \mathcal{L}(S)} v_{i,0}\right), \qquad \mathcal{N}_{\mathrm{add}} \asymp_{b,r} \epsilon^{-2} u. \tag{15}$$

Here $\mathcal{N}_{\mathrm{add}}$ is the faithful-interval minimax cost on this class. Record swaps identify high-state mean differences and a fixed number of low-state contrasts. Graded sampling then uses the same tail bound as the general theorem. Appendix E.1 gives all charged queries and a uniform risk-functional error bound; no component identity or direct component query is supplied. The admissible coefficient ranges do not shrink with K.

6 COMPUTING AND CERTIFYING A FINITE AUDIT

The statistical results specify how much continuation evidence is needed. A usable audit must also optimize over compatible routing laws and verify its numerical answers. This is a separate cost: a low observation budget does not by itself make the feasible domain easy to represent.

For the product-moment classes above, restrict to a prefix of length k. Let M_k evaluate its released monomials and C_k map assignments to their stopped records. A dummy coordinate collects the unobserved tail and receives a known value during each optimization call. The exact prefix domain is

$$\mathcal{P}_k = \{\pi \geq 0 : M_k \pi = m_k\}, \qquad \mathcal{D}_k = C_k \mathcal{P}_k, \qquad \pi \in \mathbb{R}^{2^k}.$$

Restriction and independent tail extension give both inclusions in this description. A linear call with record objective c has checkable lower and upper certificates. A feasible assignment law π and a dual vector y with $M_k^\top y \geq C_k^\top c$ give

$$c^\top C_k \pi \leq \max_{\widetilde{\pi} \geq 0 : M_k \widetilde{\pi} = m_k} c^\top C_k \widetilde{\pi} \leq \max_{\widetilde{\pi} \geq 0 : M_k \widetilde{\pi} = m_k} y^\top M_k \widetilde{\pi} = y^\top m_k.$$

Their difference certifies that call's optimization error. In the audit, these calls support constrained sampling; fresh observations then evaluate the selected laws. Numerical optimization, sampling error and omitted tail mass each receive a separate tolerance. The final error budget combines all three terms.

Rational inputs give arithmetic cost $2^{O_{b,r,p_0}(k)} \operatorname{poly}(K, B_{\mathrm{in}}, B_{\mathrm{resp}}, \epsilon^{-1}, u)$, where B_{in} and B_{resp} bound input and response bit lengths. For unrestricted means, $k = \min(K, C_{b,r,p_0} \epsilon^{-1/d})$. The implementation is exponential in this effective prefix length. Theorem 3 instead gives $k \leq C_{b,r} k_\eta$; if $\eta = O(\epsilon)$, that dimension remains bounded. For $\eta \leq \epsilon/4$, the explicit average in the previous section avoids optimization.

Preserving positivity is essential when seeking a smaller representation. At $K = 4, b = r = 2$, the complete affine equations, every individual capacity, and the exact depth profile still admit a false record vector. Its distance from the true domain is $1/24$ in total variation. The sixteen-assignment certificate in Appendix A.8 exhibits both a separating quadratic and an attaining positive law. Thus these local constraints can enlarge the set over which risk is maximized. A tractable exact formulation needs additional control of common positivity. The finite-bit algorithm and composed error bounds are in Appendix F. The experiments next use exact finite optimization to show what different information releases identify.

7 FINITE AUDITS: WHAT THE ROUTING INFORMATION IDENTIFIES

We test the identification mechanism on frozen verifier-guided search trees. For each question, all eight assignments of three routing bits have saved losses. Exact optimization then changes the compatible routing law while holding those losses fixed. This asks how much risk ambiguity remains after releasing routing statistics.

The main panel has 512 questions, balanced between MATH and GSM8K (Hendrycks et al., 2021; Cobbe et al., 2021). Qwen2.5-Math-1.5B-Instruct generates the trees (Qwen Team, 2024). Qwen2.5-Math-PRM-7B and RLHFlow Llama3.1-8B PRM score the same original prefixes (Qwen Team, 2025; Xiong et al., 2024). Trees, path seeds and both complete scoring panels are validated before main labels are opened. We report the earlier 1,024-question development panel and the external 64-question panel separately. The latter changes both questions and generator.

For question i, let $f_i(x)$ be its saved loss under assignment x. With $\mathcal{P}_r$ the fair order-r class on three bits, its identified width is

$$W_i(r) = \max_{P \in \mathcal{P}_r} \mathbb{E}_P f_i - \min_{P \in \mathcal{P}_r} \mathbb{E}_P f_i.$$

The first four columns of Table 1 average these widths over questions. Order zero supplies normalization alone; order one fixes the marginals, order two also fixes the pairs, and order three fixes the uniform joint law. The nested feasible classes explain why the columns decrease. On the main panel, the mean width falls from 11.523 to 5.273 to 1.465 pp. Complete-joint width zero means that a fixed loss table has one expected risk; that risk need not be zero.

Table 1: Risk ambiguity on fixed loss tables, in percentage points. The first four columns average per-question widths under nested moment constraints. The fifth uses one common pairwise-balanced law. The final column counts positive, negative and zero even–odd contrasts. All three panels use their original labels and equal-task weights.

	Question-dependent laws				Common law	Contrast signs
Panel	None	Singleton	Pairwise	Joint	Pairwise	$+ / - / 0$
Main 512	11.523	5.273	1.465	0.000	0.098	14/16/482
Earlier 1,024	30.371	11.328	2.832	0.000	0.049	57/57/910
External 64	9.375	2.344	0.391	0.000	0.391	0/1/63

The common-law column answers a different question. Suppose one unknown law is used independently of question identity, with the same PRM meaning for each routing bit. The auditor must then optimize the panel average under that one law. Let π_i be the normalized question weights and $\Delta_i = \mathbb{E}_{P_+} f_i - \mathbb{E}_{P_-} f_i$ the even–odd risk contrast. Since these parity laws are the endpoints of $\mathcal{P}_2$,

$$W_{\text{common}} = \left| \sum_i \pi_i \Delta_i \right| \le \sum_i \pi_i |\Delta_i| = W_{\text{question}}.$$

A question-dependent law can choose its worse parity separately for every question. A common law must make the same choice throughout the panel. Table 1 records 14 positive, 16 negative and 482 zero contrasts in the main panel. The positive and negative contributions largely cancel, leaving common-law width 0.098 pp despite mean per-question width 1.465 pp. The external panel has only one nonzero contrast, so its two widths coincide.

Sampling uncertainty across questions is a further issue. The main and external signed-risk confidence intervals both contain zero. The main absolute-contrast materiality decision also changes under the post-outcome format diagnostic: its one-sided lower bound falls from 1.039 to 0.997 pp, across the registered one-pp threshold. The original labels and both analyses remain in Appendix G.2. The identification calculations condition on these labels and preserve both evaluations.

The same mechanism has exact finite checks beyond three bits. A 32-family, width-four shared-leaf construction has fourth-order risk interval $[0, 31/17290]$, of width 0.1793 pp. Seed enumeration at $K = 16, 64, 128, 256$ also verifies the first two moments and simultaneous cylinder masses $1/(4N_h)$ and $1/(8N_h)$. These checks connect the identification example to the positive construction used in the rate proof. Acquisition-rate evaluation varies accuracy and candidate count under fresh responses. The same-cache timings measure replay computation. The complete protocol, rational endpoint checks and timing comparator are in Appendix G.

8 Conclusion and outlook

The risk of a fixed selector depends on its routing law and continuation losses. Partial moments restrict the former; paid responses estimate the latter. Their interaction creates the acquisition threshold: the number of possible deep records competes with their feasible probability mass. Simultaneous positive constructions make this comparison valid across scales. They yield matching noncritical rates and the fixed-confidence critical bounds on the original domain. The finite audits show why risk identification and fresh-response estimation must remain distinct.

The scalar–law separation motivates a sharper fixed-confidence analysis of the critical regime. A natural question is whether joint estimation across depths attains the scalar lower bound on the full positive domain, or whether stronger testing constructions raise that bound. Every record response must remain charged. A second direction is to find compact positive-domain representations for certified risk optimization. Structured continuations offer one route when their residual bound can be justified. Finally, experiments with increasing candidate counts and fresh continuations can test the predicted budget regimes. These studies should report both frozen-oracle and label-sensitivity analyses, with generation, scoring and auditing costs recorded separately.

AI USE STATEMENT

Generative AI assisted manuscript writing and revision, literature search, and the development of experimental code. It also assisted with theoretical models and conjectures, construction and proof development, experiment design and orchestration, and interpretation of results. Verification included primary-source reference checks, exact-arithmetic replay, code and build checks, and independent AI-assisted proof reviews. The authors are responsible for the final text, mathematical claims, experimental results and accompanying artifacts.

REPRODUCIBILITY STATEMENT

Section 2 specifies the moment constraints and paid observation interface. The mathematical assumptions and complete derivations appear in Appendices A–F. Appendix G specifies the frozen trees, model revisions, scoring and evaluation procedures, cohort definitions and finite identification calculations. Recorded outcomes and rational certificates support independent replay. Figure scripts, source manifests and build instructions identify the data and source files used in the manuscript.

ETHICS STATEMENT

The experiments use existing mathematics benchmarks and pretrained generators and verifiers. Their losses are defined by the stated evaluation procedures. A risk audit consequently depends on the validity of its loss and observation model. The original evaluation and the documented format-sensitivity analysis are both retained. Applications to consequential decisions require validation of the verifiers, labels and routing assumptions for the intended setting.

REFERENCES

Srinivas M. Aji and Robert J. McEliece. The generalized distributive law. *IEEE Transactions on Information Theory*, 46(2):325–343, 2000. doi: 10.1109/18.825794.

Noga Alon, Oded Goldreich, and Yishay Mansour. Almost k-wise independence versus k-wise independence. *Information Processing Letters*, 88(3):107–110, 2003. doi: 10.1016/S0020-0190(03)00359-4.

Egon Balas. Disjunctive programming and a hierarchy of relaxations for discrete optimization problems. *SIAM Journal on Algebraic and Discrete Methods*, 6(3):466–486, 1985. doi: 10.1137/0606047.

Ido Ben-Eliezer, Shachar Lovett, and Ariel Yadin. Polynomial Threshold Functions: Structure, Approximation and Pseudorandomness. arXiv:0911.3473, 2009. URL https://arxiv.org/abs/0911.3473. Section 2.2 and Claim 12.

Itai Benjamini, Ori Gurel-Gurevich, and Ron Peled. On K-wise independent distributions and boolean functions, 2012. URL https://arxiv.org/abs/1201.3261.

Daniel Berend, Philip A. Ernst, Aryeh Kontorovich, and Rishi Kumar. Exact expressions for the maximal probability that all k-wise independent bits are 1. *Journal of Theoretical Probability*, 39 (3), 2026. doi: 10.1007/s10959-026-01487-4. URL https://link.springer.com/article/10.1007/s10959-026-01487-4. Article 58.

Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, Massachusetts, 1997. ISBN 9781886529199. URL https://www.dynamic-ideas.com/books/q3ggobj67w504tsa95ae456yvndd2d.

Andrej Bogdanov, Krishnamoorthy Dinesh, Yuval Filmus, Yuval Ishai, Avi Kaplan, and Akshayaram Srinivasan. Bounded indistinguishability for simple sources. In *Innovations in Theoretical Computer Science*, volume 215, pp. 26:1–26:18, 2022. doi: 10.4230/LIPIcs.ITCS.2022.26.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Rina Dechter. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence*, 113 (1–2):41–85, 1999. doi: 10.1016/S0004-3702(99)00059-4.

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(39):1079–1105, 2006. URL https://www.jmlr.org/papers/v7/evendar06a.html.

David A. Freedman. On tail probabilities for martingales. *The Annals of Probability*, 3(1):100–118, 1975. doi: 10.1214/aop/1176996452. URL https://doi.org/10.1214/aop/1176996452.

Gene H. Golub and John H. Welsch. Calculation of gauss quadrature rules. *Mathematics of Computation*, 23(106):221–230, 1969. doi: 10.1090/S0025-5718-69-99647-1. URL https://www.ams.org/journals/mcom/1969-23-106/S0025-5718-69-99647-1/S0025-5718-69-99647-1.pdf.

Pooya Hatami and William Hoza. Paradigms for Unconditional Pseudorandom Generators. *Foundations and Trends in Theoretical Computer Science*, 16(1–2):1–210, 2024. doi: 10.1561/0400000109. URL https://williamhoza.com/research/papers/HH/fnt-tcs.pdf. Section 2.1.3, Proposition 2.1.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*,

volume 1, 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/be83ab3ecd0db773eb2dc1b0a17836a1-Abstract-round2.html.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. doi: 10.1080/01621459.1963.10500830. URL https://www.cs.rpi.edu/academics/courses/spring06/random/hoefding.pdf.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016. URL https://jmlr.org/papers/v17/kaufman16a.html.

Leonid G. Khachiyan. A polynomial algorithm in linear programming. *Soviet Mathematics Doklady*, 20:191–194, 1979.

Jasper C. H. Lee and Paul Valiant. Uncertainty about uncertainty: Optimal adaptive algorithms for estimating mixtures of unknown coins. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms*, pp. 414–433. SIAM, 2021. URL https://www.cs.purdue.edu/homes/pvaliant/cleaning.pdf.

Charles F. Manski. *Partial Identification of Probability Distributions*. Springer Series in Statistics. Springer, New York, 2003. ISBN 978-0-387-00454-9. doi: 10.1007/b97478. URL https://link.springer.com/book/10.1007/b97478.

Ryan O'Donnell and Yu Zhao. On closeness to k-wise uniformity. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018)*, volume 116 of *Leibniz International Proceedings in Informatics*, pp. 54:1–54:19, 2018. doi: 10.4230/LIPIcs.APPROX-RANDOM.2018.54.

Takayuki Osogami, Junya Honda, and Junpei Komiyama. Optimal estimation of the best mean in multi-armed bandits. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/hash/04bb165d5fb90aef582f39a58f4bf072-Abstract-Conference.html.

Ron Peled, Ariel Yadin, and Amir Yehudayoff. The maximal probability that k-wise independent bits are all 1. *Random Structures & Algorithms*, 38(4):502–525, 2011. doi: 10.1002/rsa.20329. URL https://arxiv.org/abs/0801.0059v3.

András Prékopa. Sharp bounds on probabilities using linear programming. *Operations Research*, 38(2):227–239, 1990. doi: 10.1287/opre.38.2.227.

Qwen Team. Qwen2.5-Math: The world's leading open-sourced mathematical LLMs, 2024. URL https://qwenlm.github.io/blog/qwen2.5-math/.

Qwen Team. Qwen2.5-Math-PRM-7B model card, 2025. URL https://huggingface.co/Qwen/Qwen2.5-Math-PRM-7B.

Arjun Ramachandra and Karthik Natarajan. Tight probability bounds with pairwise independence, 2022. URL https://arxiv.org/abs/2006.00516v6. Version 6, November 2022.

Chloé Rouyer and Yevgeny Seldin. Tsallis-INF for decoupled exploration and exploitation in multi-armed bandits. In *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pp. 3227–3249, 2020. URL https://proceedings.mlr.press/v125/rouyer20a.html.

Sivan Sabato, Anand D. Sarwate, and Nathan Srebro. Auditing: Active learning with outcome-dependent query costs. In *Advances in Neural Information Processing Systems*, volume 26, 2013. URL https://papers.nips.cc/paper_files/paper/2013/hash/40008b9a5380fcacce3976bf7c08af5b-Abstract.html.

Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff–hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995. doi: 10.1137/S089548019223872X.

Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley and Sons, 1986. URL https://ir.cwi.nl/pub/10206.

Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems*, volume 27, pp. 828–836, 2014. URL https://proceedings.neurips.cc/paper_files/paper/2014/hash/f8d84caae48546d0934d637ab54f7086-Abstract.html.

Mikkel Thorup and Yin Zhang. Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation. *SIAM Journal on Computing*, 41(2):293–331, 2012. doi: 10.1137/100800774.

Joel A. Tropp. Freedman's inequality for matrix martingales. *Electronic Communications in Probability*, 16:262–270, 2011. doi: 10.1214/ECP.v16-1624. URL https://tropp.caltech.edu/papers/Tro11-Freedmans-Inequality.pdf.

Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1–2):1–305, 2008. doi: 10.1561/2200000001.

Wei Xiong, Hanning Zhang, Nan Jiang, and Tong Zhang. An implementation of generative PRM, 2024. URL https://huggingface.co/RLHFlow/Llama3.1-8B-PRM-Deepseek-Data.

APPENDIX

Appendices A–C establish the selector domain and the simultaneous positive laws. Appendix D derives the paid-observation bounds. Appendix E treats structured means under that same experiment. Appendices F–G give finite computation and experimental details.

CONTENTS AND CONTRIBUTION TYPES

$\star$ denotes a result or construction developed for the selector problem. $\triangle$ denotes a refinement or transfer of existing theory. $\circ$ denotes a finite certificate, protocol or empirical calculation. The classifications concern each section's substantive step; the sources of its ingredients are specified below and at their use.

$\triangle$ A Selector fibers and positive moment completion . . . 15
$\triangle$ A.1 The selector information profile . . . 15
$\triangle$ A.2 Known losses and prescribed moments . . . 16
$\triangle$ A.3 Proof of the information profile . . . 17
$\triangle$ A.4 Proof of positive fiber completion . . . 19
$\triangle$ A.5 The signed-state audit . . . 20
$\triangle$ A.6 A common record statistic . . . 20
$\triangle$ A.7 The full identified-loss space . . . 22
$\circ$ A.8 A finite positivity certificate . . . 23
$\star$ B Multiscale cancellation and positive repair . . . 24
$\star$ C Simultaneous positive moment constructions . . . 25
$\triangle$ C.1 Count identities and a positive reference . . . 25
$\star$ C.2 Uniform mixed-degree error bounds . . . 27
$\star$ C.3 Simultaneous cylinders with prescribed biases . . . 30
$\star$ C.4 Explicit fourth- and fifth-order constants . . . 32
$\star$ C.5 Sharp simultaneous constants at orders two and three . . . 37
$\star$ D Paid observation bounds and acquisition rates . . . 38
$\triangle$ D.1 Endpoint padding and exact queries . . . 38
$\triangle$ D.2 Stopped change of measure . . . 39
$\triangle$ D.3 Finite-profile constrained sampling . . . 41
$\triangle$ D.4 Evidence profiles and lower certificates . . . 43
$\triangle$ D.5 Stopped-mixture information . . . 45
$\triangle$ D.6 Weighted block-maximum estimation . . . 48
$\star$ D.7 Exact record-value acquisition . . . 50
$\circ$ D.8 Finite-size parameters . . . 52
$\star$ D.9 General record acquisition rates . . . 52
$\star$ D.10 Product maxima on the full positive domain . . . 53
$\star$ D.11 The stronger pairwise critical lower bound . . . 56
$\star$ D.12 Scalar and feasible-law output costs . . . 59
$\star$ E Structured continuation means . . . 62
$\triangle$ E.1 Selected-family interactions and additive reduction . . . 66
$\triangle$ F Finite-bit implementation and certificates . . . 70
$\triangle$ F.1 Combined moment, representation and sampling errors . . . 73
$\circ$ G Experimental protocol and finite identification . . . 75
$\circ$ G.1 Trees, scoring and observations . . . 75
$\circ$ G.2 Risk contrasts and oracle sensitivity . . . 75
$\circ$ G.3 Sufficient-state auditing and its comparator . . . 76
$\circ$ G.4 Finite certificates and identification calculations . . . 77
$\circ$ G.5 Common-law intervals for finite panels . . . 78

Contribution provenance

The simultaneous-cylinder theorem and its scale-uniform error control are the central probability construction. The acquisition theorems use these common positive witnesses to obtain scalar information bounds. The structured theorem transfers the same observation model to an unknown component in the full identified-loss space. The scalar–law comparison constructs a positive domain and a specified mean family on which the two output requirements have different costs.

Several ingredients have separate origins. Shallow-leaf preservation under limited independence is standard (Ben-Eliezer et al., 2009; Hatami & Hoza, 2024); Appendix A derives the selector-specific fibers, attaining laws and affine basis. Fiber completion adapts finite-moment linear programming and disjunctive convexification (Prékopa, 1990; Balas, 1985; Wainwright & Jordan, 2008). The count construction uses classical Gaussian quadrature (Golub & Welsch, 1969) and positive moment-repair ideas (Alon et al., 2003; O'Donnell & Zhao, 2018). Its additional step is cancellation and joint error control for all prescribed cylinders.

The statistical upper proof adapts decoupled exploration (Rouyer & Seldin, 2020). Stopped change of measure (Kaufmann et al., 2016), adaptive mixture information (Lee & Valiant, 2021), and action elimination (Even-Dar et al., 2006) provide the testing and selection tools. Their use here preserves the charged record response, conditioning history and common positive domain. The finite-bit implementation uses classical rational linear programming (Khachiyan, 1979; Schrijver, 1986). Finite examples and experimental calculations illustrate these identification and computation statements.

Notation

The record map is T, and record t has depth d_t. We use N for paid observations and τ_{stop} for an adaptive stopping time. There are L geometric scales. The effective-depth and confidence logarithms are $\Lambda_K = 1 + \log\min(K, \epsilon^{-1/d})$ and $u = \log(8/\delta)$. The profile quantity is $\mathsf{H} = \sum_h \sqrt{M_h} w_h$. Hermite polynomials are He_j; the Hellinger distance is Hel. The standardized count vector is V. Constants are uniform in the parameters quantified in their statements.

For the finite parity audit, write

$$a(y) = P_+(Y = y) - P_-(Y = y), \qquad C_{\mathrm{out}} = \tfrac{1}{2}\sum_y |a(y)|, \qquad \Delta = \sum_y a(y)\,\mathrm{loss}(y), \tag{16}$$

where P_+, P_- are the even/odd laws. The first two quantities describe output sensitivity; Δ is the signed contrast for the specified loss.

A Selector fibers and positive moment completion

A.1 The selector information profile

For each feasible retained record t, the cutoff fixes a prefix F_t in decreasing upper-key order and leaves the other coordinates free. Write $d_t = |F_t|$, and let m_r specify every fair marginal of order at most r. The following result determines the information ambiguity of the selector; a particular fixed loss can have smaller ambiguity.

Proposition 4 (Sharp selector information)**.** *For $1 \le b \le K$ and $0 \le r < K$, let U be the uniform assignment law, v the $(r+1)$st largest upper key, and*

$$n_{\mathrm{fixed}}(v) = |\{i : l_i > v\}|, \qquad n_{\mathrm{switch}}(v) = |\{i : l_i \le v < u_i\}|.$$

Then

$$\begin{aligned} \mathcal{A}_r(T) &:= \sup_{P,Q \in \mathcal{P}(m_r)} \mathrm{TV}(T_{\#}P, T_{\#}Q) = P_U(d_T > r) \\ &= \Pr\{\mathrm{Binomial}(n_{\mathrm{switch}}(v), 1/2) \le b - 1 - n_{\mathrm{fixed}}(v)\}, \end{aligned} \tag{17}$$

$$\sup_T \mathcal{A}_r(T) = 2^{-r} \sum_{j=0}^{\min(b-1,r)} \binom{r}{j}. \tag{18}$$

One pair of globally compatible laws attains the fixed-geometry bound simultaneously on every contributing fiber. A fixed binary retained-record loss attains its TV difference. Any common retained-information suffix contracts this difference. For $r \geq K$, the ambiguity is zero.

Remark (decision-tree interpretation and attribution). Read bits in decreasing upper-key order and stop as soon as the retained record is determined. This is an oblivious decision tree: its variable order is fixed, and its leaves are exactly the prefix subcubes in (21). Every leaf of depth $j \leq r$ has probability 2^{-j} under every fair r-wise law, as under U. The resulting common-leaf-mass upper bound (22) is the classical finite-independence decision-tree argument (Ben-Eliezer et al., 2009, Sec. 2.2, Claim 12); an explicit general statement appears in Hatami & Hoza (2024, Proposition 2.1). The selector-specific content is the exact heterogeneous-key formula, its simultaneous attainment on all deep fibers by one prefix-parity pair, and the ensuing record-count and capacity inputs to acquisition. Parity provides a standard bounded-indistinguishability example (Bogdanov et al., 2022, Section 1). TV contraction is applied directly in the proof below. The proof is in Section A.3. Taking $r = K - 1$ recovers the even/odd intervention, including $K = b = 1$. With $a = |\{i : l_i \leq \min_j u_j\}|$,

$$c_{\mathrm{root}} = 2^{1-a} \sum_{j=K-b}^{a-1} \binom{a-1}{j} \leq \Gamma_{K,b} := 2^{1-K} \sum_{j=0}^{b-1} \binom{K-1}{j}. \tag{19}$$

Thus $|\Delta| \leq C_{\mathrm{out}} \leq c_{\mathrm{root}}$ for the fixed loss and suffix in the model. The formula needs $O(K)$ key comparisons after the $2K$ root scores are supplied; exact integer arithmetic has its own bit cost. General suffix-kernel attainment is distinct from attainment by a particular cumulative-score beam.

A.2 Known losses and prescribed moments

For arbitrary feasible unary and pair moments m, retain $\psi(x) = (1, (x_i)_i, (x_i x_j)_{ij \in E})$. Let $\mathcal{C}(G) = \mathrm{conv}\{(x_i, x_i x_j)_{i \in V(G), ij \in E(G)} : x \in \{0,1\}^{V(G)}\}$ be the correlation polytope. For a fiber with fixed assignment z_t on F_t and free set U_t, let A_t be the linear moment-extension map obtained by substituting $x_{F_t} = z_t$ into ψ.

Proposition 5 (Cutoff-conditioned completion)**.** *A record law q belongs to $\mathcal{Q}(m)$ if and only if there are free-coordinate moment vectors (y^t, z^t) such that*

$$(q_t, y^t, z^t) \in \mathrm{cone}(\{1\} \times \mathcal{C}(E[U_t])), \qquad \sum_t A_t(q_t, y^t, z^t) = m. \tag{20}$$

For known θ, minimizing or maximizing $\sum_t q_t \theta_t$ gives $R_{\pm}(m, \theta)$ with one common attaining assignment law; an optimum needs at most $\mathrm{rank}(\psi)$ atoms. Dual separation on a fiber is quadratic Boolean optimization on U_t. The free sets are nested; fixing a fiber changes unary fields but not the remaining pair coefficients.

This specializes finite-moment LP duality and disjunctive convexification (Prékopa, 1990; Balas, 1985; Wainwright & Jordan, 2008); it does not assume conditional independence. Section A.4 gives the aggregation equations and primal/dual certificate. There are at most $2^b \binom{K}{b}$ candidate records. With a supplied elimination order of width w, each pricing problem takes $\mathrm{poly}(K + |E|) 2^{w+1}$ arithmetic operations (Dechter, 1999). Dense cases and the number of generated assignment columns can remain exponential.

Audit consequence. For a fixed signed dependence contrast, equal sufficient states share the same suffix kernel; their signed masses can be summed, and zero sums need no further measurement. This aggregation uses the distributive law underlying standard message passing (Aji & McEliece, 2000, Section II). For parity, only singleton root fibers survive. They can be found among $2 \sum_{j=0}^{b-1} \binom{K-1}{j} = O_b(K^{b-1})$ assignments, giving $O_b(K^b)$ key comparisons without a full-cube replay. Labels are needed only on the nonzero signed output support, when unresolved labels may vary independently. The comparator must use the same model-input cache. This measures a fixed contrast; it does not supply unknown suffix values for free.

A.3 Proof of Proposition 4

All comparisons use the complete identity-resolved keys from the model. A retained state includes its family bit even when its two score endpoints coincide. At order $r = 0$, m_0 prescribes only normalization; the same proof applies with an empty leading-r prefix.

Fix a proposed record t and let c_t be its lowest retained key. Every retained coordinate is fixed to its recorded bit. For a discarded coordinate i, $l_i \geq c_t$ makes the record infeasible; $l_i < c_t \leq u_i$ forces $X_i = 0$; and $u_i < c_t$ leaves X_i free. Thus every nonempty fiber is exactly

$$T^{-1}(t) = \{x : x_{F_t} = z_t\}, \qquad U_t = \{i : u_i < c_t\}, \qquad F_t = [K] \setminus U_t. \tag{21}$$

In decreasing upper-key order, F_t is the first d_t coordinates. It contains at most b ones, because a fixed high state must be retained. The fibers are disjoint and cover the cube. Enumerating a retained identity set and its bits gives at most $2^b\binom{K}{b}$ candidates; the cutoff conditions decide their feasibility.

If $d_t \leq r$, every law in $\mathcal{P}(m_r)$ assigns the fiber mass 2^{-d_t}. Write $p_t = P(T = t)$ and $q_t = Q(T = t)$. Then

$$\begin{aligned} \mathrm{TV}(T_\# P, T_\# Q) &= \tfrac{1}{2}\sum_t (p_t + q_t - 2\min\{p_t, q_t\}) \\ &= 1 - \sum_t \min\{p_t, q_t\} \\ &\leq 1 - \sum_{t: d_t \leq r} 2^{-d_t} = P_U(d_T > r). \end{aligned} \tag{22}$$

Let B be the leading $r + 1$ upper-priority coordinates and set

$$P_\pm(x) = 2^{-K}\left(1 \pm \prod_{i \in B}(-1)^{x_i}\right). \tag{23}$$

For $A \subseteq [K]$, $|A| \leq r$, choose $j \in B \setminus A$. Writing $\chi_B(x) = \prod_{i \in B}(-1)^{x_i}$ gives

$$P_\pm(x) \geq 0, \qquad \sum_{x: x_A = a} \chi_B(x) = \sum_{x_{[K]\setminus(A\cup\{j\})}} \prod_{i \in B\setminus\{j\}} (-1)^{x_i} \underbrace{\sum_{x_j=0}^{1} (-1)^{x_j}}_{0} = 0,$$

$$P_\pm(X_A = a) = 2^{-K}2^{K-|A|} = 2^{-|A|}.$$

The case $A = \varnothing$ also verifies normalization. Put $\chi_B(z_t) = \prod_{i \in B}(-1)^{z_t(i)}$ when $B \subseteq F_t$. Summing the free bits in (23) gives

$$P_\pm(T = t) = \begin{cases} 2^{-d_t}, & d_t \leq r, \\ 2^{-d_t}(1 \pm \chi_B(z_t)), & d_t > r. \end{cases} \tag{24}$$

Consequently, for $\theta_t = \mathbf{1}\{d_t > r, \chi_B(z_t) = 1\}$,

$$\begin{aligned} \mathrm{TV}(T_\# P_+, T_\# P_-) &= \tfrac{1}{2}\sum_{t: d_t > r} 2^{1-d_t}|\chi_B(z_t)| = P_U(d_T > r), \\ \mathbb{E}_{P_+}\theta_T - \mathbb{E}_{P_-}\theta_T &= \sum_{t: P_+(T=t) > P_-(T=t)} [P_+(T = t) - P_-(T = t)] \\ &= \mathrm{TV}(T_\# P_+, T_\# P_-). \end{aligned}$$

Let $v = u_{(r+1)}$ and let N_v count realized keys strictly above v. Under the independent fair law U,

$$\begin{aligned} N_v &= \sum_{i: l_i > v} 1 + \sum_{i: l_i \leq v < u_i} X_i \overset{d}{=} n_{\mathrm{fixed}}(v) + \mathrm{Binomial}(n_{\mathrm{switch}}(v), 1/2), \\ \{d_T > r\} &= \{c_T \leq v\} = \{N_v \leq b - 1\}, \\ P_U(d_T > r) &= 2^{-n_{\mathrm{switch}}(v)} \sum_{j=0}^{n_{\mathrm{switch}}(v)} \binom{n_{\mathrm{switch}}(v)}{j} \mathbf{1}\{j \leq b - 1 - n_{\mathrm{fixed}}(v)\}. \end{aligned}$$

This form includes thresholds outside the binomial support.

For the geometry-uniform bound, consider the leading r coordinates. Whenever at least b are high, those bits determine the record. Let $t(a)$ denote it for a prefix $a \in \{0,1\}^r$ with $|a| \geq b$. The common submeasure and its mass are

$$\mu_t = 2^{-r} \sum_{a:|a|\geq b} \mathbf{1}\{t(a) = t\}, \qquad \sum_t \mu_t = 2^{-r} \sum_{j=b}^{r} \binom{r}{j}.$$

Thus $P(T = t), Q(T = t) \geq \mu_t$ and

$$\mathrm{TV}(T_\# P, T_\# Q) \leq 1 - \sum_t \mu_t = 2^{-r} \sum_{j=0}^{\min(b-1,r)} \binom{r}{j}.$$

For equality choose strict common-overlap keys, with all highs above all lows. The first b highs determine the record. Until they appear, changing the next unread bit can change the record. Thus, for $r < K$,

$$\{d_T > r\} = \left\{ \sum_{i=1}^{r} X_i < b \right\},$$

$$\mathrm{TV}(T_\# P_+, T_\# P_-) = P_U(d_T > r) = 2^{-r} \sum_{j=0}^{\min(b-1,r)} \binom{r}{j}.$$

This equals one when $b > r$. For $r \geq K$, every atom probability is prescribed, so $P = Q$ and the ambiguity is zero.

A common retained-information suffix is a Markov kernel K. For $D_t = P(T = t) - Q(T = t)$, $\sum_t D_t = 0$. Therefore

$$|P(Y \in A) - Q(Y \in A)| = \left| \sum_t D_t \mathsf{K}(A \mid t) \right| \leq \sum_{t:D_t>0} D_t = \tfrac{1}{2} \sum_t |D_t|,$$

$$|\mathbb{E}_P \operatorname{loss}(Y) - \mathbb{E}_Q \operatorname{loss}(Y)| = \left| \sum_t D_t \theta_t \right| \leq \tfrac{1}{2} \sum_t |D_t|, \qquad 0 \leq \theta_t \leq 1.$$

Taking the supremum over A proves output TV contraction.

At $r = K-1$, (23) is the full even/odd parity pair. Its deep fibers have $d_t = K$, hence are singletons. Let $u_* = \min_i u_i$ and $a = |\{i : l_i \leq u_*\}|$. Then $n_{\mathrm{fixed}}(u_*) = K - a$ and $n_{\mathrm{switch}}(u_*) = a - 1$. For $Z \sim \mathrm{Binomial}(a-1, 1/2)$, $a - 1 - Z$ has the same law. Hence

$$\begin{aligned} P_U(d_T = K) &= \Pr\{K - a + Z \leq b - 1\} \\ &= \Pr\{a - 1 - Z \geq K - b\} = 2^{1-a} \sum_{j=K-b}^{a-1} \binom{a-1}{j}. \end{aligned}$$

proving (19). It is zero exactly when at least b low endpoints exceed the smallest upper endpoint. The geometry-uniform bound is $\Gamma_{K,b}$ by (18).

If arbitrary laws P, Q agree on every proper coordinate marginal, their difference satisfies $D(x) + D(x \oplus e_i) = 0$ for every i. Successively flipping the coordinates with $x_i = 1$ gives

$$\begin{aligned} D(x) &= (-1)^{|x|} D(0), \qquad \mathrm{TV}(P, Q) = 2^{K-1} |D(0)|, \\ \sum_{x:T(x)=t} D(x) &= \begin{cases} 0, & d_t < K, \\ D(x_t), & T^{-1}(t) = \{x_t\}, \end{cases} \\ \mathrm{TV}(T_\# P, T_\# Q) &= \tfrac{1}{2} |D(0)| \, |\{t : d_t = K\}| \\ &= P_U(d_T = K)\, 2^{K-1} |D(0)| = c_{\mathrm{root}} \, \mathrm{TV}(P, Q). \end{aligned}$$

This is contraction along the specified indistinguishable direction, not a statement that adding root families improves accuracy.

A.4 Proof of Proposition 5

The fiber-specific construction below adapts disjunctive convexification (Balas, 1985) to the marginal representation of each conditional law (Wainwright & Jordan, 2008). For each fiber, its unnormalized free moments are $y_i^t = q_t \mathbb{E}[X_i \mid T = t]$ and $z_{ij}^t = q_t \mathbb{E}[X_i X_j \mid T = t]$. The correlation cone makes these the moments of one law on the entire free cube, and forces all free moments to zero if $q_t = 0$. The extension map has constant coordinate q_t and remaining coordinates

$$M_i^t = \begin{cases} z_t(i) q_t, & i \in F_t, \\ y_i^t, & i \in U_t, \end{cases}$$

$$M_{ij}^t = \begin{cases} z_t(i) z_t(j) q_t, & i, j \in F_t, \\ z_t(i) y_j^t, & i \in F_t,\ j \in U_t, \\ z_t(j) y_i^t, & j \in F_t,\ i \in U_t, \\ z_{ij}^t, & i, j \in U_t. \end{cases} \tag{25}$$

Thus $\sum_t A_t(q_t, y^t, z^t) = m$ imposes $\sum_t q_t = 1$, $\sum_t M_i^t = m_i$, and $\sum_t M_{ij}^t = m_{ij}$ for every supplied pair.

For a feasible P, put $q_t = P(T = t)$ and disintegrate on every $q_t > 0$:

$$P(x) = \sum_t q_t \mathbf{1}\{x_{F_t} = z_t\} \nu_t(x_{U_t}), \qquad \nu_t = P(X_{U_t} \in \cdot \mid T = t).$$

Its free moment vector lies in $q_t \mathcal{C}(E[U_t])$. Conversely, membership in this cone supplies coefficients $\nu_t(u) \geq 0$, $\sum_u \nu_t(u) = 1$, with

$$(q_t, y^t, z^t) = q_t \sum_{u \in \{0,1\}^{U_t}} \nu_t(u) (1, (u_i)_i, (u_i u_j)_{ij \in E[U_t]}).$$

For $q_t = 0$, all coordinates are zero and ν_t is arbitrary. Define the global P by the same displayed mixture. Its components have disjoint record fibers, so

$$P(x) \geq 0, \quad \sum_x P(x) = \sum_t q_t = 1, \quad P(T = t) = q_t,$$

$$\mathbb{E}_P \psi(X) = \sum_t A_t(q_t, y^t, z^t) = m,$$

$$\mathbb{E}_P \theta_{T(X)} = \sum_t q_t \theta_t.$$

Write $A_{\cdot x} = \psi(x)$ and $c_x = \theta_{T(x)}$. The primal set $\{p \geq 0 : Ap = m\}$ is nonempty and contained in the probability simplex. Finite LP duality (Bertsimas & Tsitsiklis, 1997, Chapter 4) gives

$$\min_{p \geq 0 : Ap = m} c^\top p = \max_{\lambda : A^\top \lambda \leq c} m^\top \lambda.$$

Both optima are attained. A basic optimum uses linearly independent columns of A, hence at most $\operatorname{rank}(A)$ atoms. The lower-endpoint dual is

$$\max_\lambda \lambda^\top m \quad \text{subject to} \quad h_\lambda(x) := \lambda_0 + \sum_i \lambda_i x_i + \sum_{ij \in E} \lambda_{ij} x_i x_j \leq \theta_{T(x)} \quad \text{for all } x. \tag{26}$$

The upper endpoint uses the corresponding majorant. On a fiber,

$$h_\lambda(x) = a_t + \sum_{i \in U_t} \left(\lambda_i + \sum_{j \in F_t : ij \in E} \lambda_{ij} z_t(j) \right) x_i + \sum_{ij \in E[U_t]} \lambda_{ij} x_i x_j, \tag{27}$$

where a_t collects terms using only fixed coordinates. Conditioning shifts unary fields and takes a principal subgraph; it does not change surviving pair coefficients. The free sets are nested by (21).

For completeness, eliminate a free bit x_i by replacing all current terms containing it with the table

$$g_i(x_{B_i \setminus \{i\}}) = \min_{x_i \in \{0,1\}} \sum_{\alpha : i \in \alpha} g_\alpha(x_\alpha), \qquad B_i = \bigcup_{\alpha : i \in \alpha} \alpha.$$

The distributive identity $\min_{x_i,x_J}(g(x_i,x_J)+h(x_J)) = \min_{x_J}(\min_{x_i} g(x_i,x_J)+h(x_J))$ preserves the minimum at each step. A supplied width-w order has $|B_i| \le w+1$, so each table uses at most 2^{w+1} assignments (Dechter, 1999). This optimizes the dual polynomial; it imposes no factorization on the unknown P. Maximization uses $\max h_\lambda = -\min(-h_\lambda)$. For an upper endpoint, a feasible rational pair satisfies

$$p \ge 0, \quad Ap = m, \quad A^\top \lambda \ge c, \qquad c^\top p \le R_+ \le m^\top \lambda.$$

Equality of the two displayed objectives certifies the endpoint. The lower certificate reverses the dual inequality. An elimination trace certifies the all-assignment dual constraint. Its length may be exponential. Likewise, the 2^K possible columns guarantee finite exact column generation, without a polynomial iteration bound.

A.5 The signed-state audit consequence

Let $s \sim s'$ mean that two sufficient states have the same downstream kernel H. At any cut, grouping each equivalence class C gives

$$\sum_s w(s) H(B \mid s) = \sum_C \underbrace{\left(\sum_{s \in C} w(s)\right)}_{w(C)} H(B \mid C).$$

Terms with $w(C) = 0$ vanish. Applying this identity at every cut preserves the signed output measure. For otherwise unconstrained binary labels,

$$\Delta(\ell) = \sum_y a(y)\ell(y), \qquad \Delta(\ell + e_y) - \Delta(\ell) = a(y) \quad \text{if } \ell(y) = 0.$$

Thus every label with $a(y) \ne 0$ is necessary for exact recovery; the labels on this support are also sufficient.

For parity, any free coordinate cancels within its fiber. In a singleton fiber every high is retained, so there are at most b highs. If there are exactly b, the family with smallest upper endpoint must be among them; otherwise that discarded coordinate would be free. Enumerate high subsets of size below b, and size-b subsets containing this family. Their number is

$$\begin{aligned}\sum_{j=0}^{b-1} \binom{K}{j} + \binom{K-1}{b-1} &= \sum_{j=0}^{b-1} \left[\binom{K-1}{j} + \binom{K-1}{j-1}\right] + \binom{K-1}{b-1} \\ &= 2\sum_{j=0}^{b-1} \binom{K-1}{j}.\end{aligned}$$

Apply the selector and retain a candidate exactly when every discarded upper endpoint exceeds its cutoff. Its discarded bits are then uniquely low. These and only these candidates are singleton fibers. Fixed-width top-b selection takes $O_b(K)$ comparisons per candidate, proving the initialization cost.

State sufficiency is essential. With three families, width two, and endpoints $(10,11),(8,9),(1,2)$, the third family is never retained and the retained parity laws coincide. A suffix that remembers its discarded score can nevertheless recover full parity and produce risk difference one. Thus text-only merging is valid only when it preserves the declared common suffix kernel.

A.6 A universal record statistic for all score geometries

Let $[K]$ be ordered by strictly decreasing upper keys, including the fixed tie rule. Define

$$\mathcal{S}_{K,b} = \{S \subseteq [K] : |S| \le b\}, \qquad C_b(x) = \begin{cases} \{\text{indices of the first } b \text{ ones in } x\}, & |x| \ge b, \\ \{i : x_i = 1\}, & |x| < b. \end{cases}$$

The statistic is equivalent to the retained record of a common-overlap selector. Write $\mathcal{Q}^{\mathrm{can}}_{K,b,r} = (C_b)_{\#}\mathcal{U}^K_r$.

Lemma 6 (Canonical factorization and exact oracle transfer)**.** *For every legal top-b score geometry T on the ordered families,*

$$T = g_T \circ C_b, \qquad g_T(S) = T(\mathbf{1}_S), \qquad G_T(t, S) = \mathbf{1}\{g_T(S) = t\}. \tag{28}$$

Consequently

$$\mathcal{Q}_T = G_T \mathcal{Q}^{\mathrm{can}}_{K,b,r}, \qquad R_\pm(\mathcal{Q}_T, \theta) = R_\pm(\mathcal{Q}^{\mathrm{can}}_{K,b,r}, G_T^\top \theta). \tag{29}$$

For exact queries, fresh Bernoulli queries, or the class of all history-adapted $[0,1]$ responses with the specified conditional mean, and for scalar-value, faithful-interval, or feasible-record-law output,

$$\sup_{T \text{ legal}} \mathcal{N}_T(\epsilon, \delta) = \mathcal{N}^{\mathrm{can}}_{K,b,r}(\epsilon, \delta). \tag{30}$$

A polynomial-size linear formulation for the canonical family yields one for every heterogeneous record projection by adjoining $q = G_T z$. Conversely, a uniform formulation for all legal geometries includes the canonical family. The same equivalence holds for exact linear optimization and for additive-error value oracles with the same error tolerance.

Proof. Fix $S = C_b(x)$. If $|S| < b$, then $x = \mathbf{1}_S$. Otherwise let $j = \max S$. The total-key order gives

$$x_{[j]} = \mathbf{1}_S|_{[j]}, \qquad \#\{i \le j : X_i = 1,\ u_i \ge u_j\} = b, \qquad \max_{i>j} u_i < u_j.$$

Each unread family's key is at most its upper key, which is below the b observed upper keys. Replacing $x_{j+1:K}$ by zeros therefore preserves the retained identities and states: $T(x) = T(\mathbf{1}_S) = g_T(S)$. For every law P and every record t,

$$(T_\# P)_t = \sum_{S : g_T(S) = t} P\{C_b = S\} = [G_T (C_b)_\# P]_t.$$

Consequently

$$\begin{aligned} \{T_\# P : P \in \mathcal{U}_r^K\} &= G_T \mathcal{Q}^{\mathrm{can}}_{K,b,r}, \\ R_\pm(\mathcal{Q}_T, \theta) &= \mathrm{ext}_{z \in \mathcal{Q}^{\mathrm{can}}_{K,b,r}} (G_T z)^\top \theta = R_\pm(\mathcal{Q}^{\mathrm{can}}_{K,b,r}, G_T^\top \theta). \end{aligned}$$

To run a canonical algorithm, replace its request for S by one request for $g_T(S)$. Since $(G_T^\top \theta)_S = \theta_{g_T(S)}$, exact responses and fresh Bernoulli kernels agree with the simulated experiment. For bounded responses, condition on the entire simulated history:

$$\mathbb{E}[Z \mid S, \mathcal{F}_{n-1}] = (G_T^\top \theta)_S, \qquad N_T = N_{\text{simulated}} \quad \text{pathwise.}$$

A returned law $\widehat{z}$ maps to the feasible law $G_T \widehat{z}$ with identical value. Hence $\mathcal{N}_T \le \mathcal{N}^{\mathrm{can}}_{K,b,r}$ for all T. A common-overlap T is a legal geometry and has the canonical record statistic, proving the reverse inequality. For an extended formulation $\{(z, w) : H(z, w) \le a\}$ of the canonical domain, adjoining $q = G_T z$ projects exactly onto $\mathcal{Q}_T$. An optimization call uses the objective $G_T^\top \theta$; an additive error e remains e because the two objective values are equal. The reverse reductions use the legal common-overlap geometry. The matrix G_T has $\sum_{j=0}^{b} \binom{K}{j}$ columns. □

Truncated projections. For $b \le k \le K$, let $T_{\le k}$ return a determined record or $\perp$ when $d_T > k$. On the first k bits, retain the canonical partial states S with $|S| < b$ separately. When $|S| = b$, the preceding proof already determines T. When $|S| < b$, $x_{[k]} = \mathbf{1}_S$ is fully known, so comparison of the bth prefix key with u_{k+1} decides whether T is fixed; if not, return $\perp$. At $k = K$ there is no unread key. This defines an explicit map $G_{T,k}$ and

$$\{(T_{\le k})_\# P : P \in \mathcal{U}_r^K\} = G_{T,k} \mathcal{Q}^{\mathrm{can}}_{k,b,\min(r,k)}. \tag{31}$$

For one inclusion, restrict any $P \in \mathcal{U}_r^K$ to $[k]$. For the other, let $P_k \in \mathcal{U}^k_{\min(r,k)}$ and append independent balanced tail bits. For $A \subseteq [K]$ with $|A| \le r$,

$$\mathbb{E}_{P_k \otimes U_{K-k}} \prod_{i \in A} X_i = \mathbb{E}_{P_k} \prod_{i \in A \cap [k]} X_i \, \mathbb{E}_{U_{K-k}} \prod_{i \in A \setminus [k]} X_i = 2^{-|A \cap [k]|} 2^{-|A \setminus [k]|} = 2^{-|A|}.$$

The extension is feasible and has the required prefix law, proving (31). Partial states must remain separate until $G_{T,k}$ is applied. Both finite optimization and record acquisition therefore use this same universal positive polytope.

A.7 THE EXACT AFFINE CONSTRAINTS FOR EVERY BREADTH AND MOMENT ORDER

Fix $1 \le b \le K$ and $0 \le r \le K$. Let $z_i = 1 - X_i$ and $z_A = \prod_{i \in A} z_i$. Define

$$\mathcal{B}_{K,b,r} = \{\varnothing\} \cup \{A \subseteq [K] : 1 \le |A| \le r,\ \max A \le b + |A| - 1\}, \quad M = \sum_{j=0}^{b} \binom{K}{j}.$$

For a function on the Boolean cube, measurability below means constancy on every fiber of the canonical statistic C_b.

Lemma 7 (Complete affine hull)**.** *The space of record-measurable multilinear polynomials of degree at most r is exactly*

$$\mathcal{V}_{K,b,r} = \operatorname{span}\{z_A : A \in \mathcal{B}_{K,b,r}\}, \qquad d_{K,b,r} := \dim \mathcal{V}_{K,b,r} = 1 + \sum_{j=1}^{r} \binom{\min(K, b+j-1)}{j}. \tag{32}$$

Write $f_A(C_b(x)) = z_A(x)$ for the induced record functions. Then

$$\begin{aligned} \operatorname{aff}(\mathcal{Q}^{\mathrm{can}}_{K,b,r}) &= \left\{ q \in \mathbb{R}^M : \sum_S q_S f_A(S) = 2^{-|A|} \quad (A \in \mathcal{B}_{K,b,r}) \right\}, \\ \dim \mathcal{Q}^{\mathrm{can}}_{K,b,r} &= M - d_{K,b,r}. \end{aligned} \tag{33}$$

If $K \ge b + r - 1$, then $d_{K,b,r} = \binom{b+r}{r}$.

Proof. For nonempty $A \in \mathcal{B}$, set $j = \max A$. If the bth high occurs before j, then

$$|[j-1] \setminus A| = j - |A| \le b - 1 \quad \Longrightarrow \quad C_b(x) \cap A \ne \varnothing \quad \Longrightarrow \quad z_A(x) = 0.$$

Otherwise the record determines every bit in A. In both cases $z_A(x) = \mathbf{1}\{A \cap C_b(x) = \varnothing\}$. The empty monomial is constant, so every displayed monomial is record-measurable.

Conversely, expand a measurable polynomial as $f(z) = \sum_{|A| \le r} c_A z_A$. Take a set A with $j = \max A > b + |A| - 1$ and put $A' = A \setminus \{j\}$. Because the expansion is multilinear, define

$$h(z_1, \dots, z_{j-1}) = (\partial_{z_j} f)(z_1, \dots, z_{j-1}, 0, \dots, 0).$$

For every $S \subseteq [j-1]$ with $|S| \le j - b - 1$, there are at least b highs among the first $j-1$ coordinates at $z = \mathbf{1}_S$. The record is unchanged when z_j changes, so $h(\mathbf{1}_S) = 0$. Since $|A'| \le j - b - 1$, Boolean coefficient inversion gives

$$c_A = [z_{A'}]h = \sum_{S \subseteq A'} (-1)^{|A'| - |S|} h(\mathbf{1}_S) = 0. \tag{34}$$

Only the sets in $\mathcal{B}$ remain. Squarefree monomials are linearly independent on the Boolean cube. For each size j, the allowed sets are precisely the j-subsets of $[\min(K, b + j - 1)]$, proving (32). If $K \ge b + r - 1$, Pascal's identity gives

$$1 + \sum_{j=1}^{r} \binom{b+j-1}{j} = 1 + \sum_{j=1}^{r} \left\{ \binom{b+j}{j} - \binom{b+j-1}{j-1} \right\} = \binom{b+r}{r}.$$

Let A_r be the assignment-to-moment matrix including normalization, and B the assignment-to-record incidence matrix. The uniform assignment law $p_0 = 2^{-K}\mathbf{1}$ is strictly positive. For every nonzero $v \in \ker A_r$,

$$0 < t < 2^{-K}/\|v\|_\infty \quad \Longrightarrow \quad p_0 \pm tv \ge 0, \qquad A_r(p_0 \pm tv) = A_r p_0.$$

For feasible p, p', $A_r(p - p') = 0$, so $\operatorname{span}(\mathcal{Q} - \mathcal{Q}) \subseteq B \ker A_r$. Conversely, $Bv = [B(p_0 + tv) - B(p_0 - tv)]/(2t)$ for the displayed choice of t. Thus

$$\begin{aligned} \operatorname{span}(\mathcal{Q} - \mathcal{Q}) &= B \ker A_r, \\ u \in (B \ker A_r)^\perp &\iff \langle B^\top u, v \rangle = 0 \quad (v \in \ker A_r) \\ &\iff B^\top u \in \operatorname{row}(A_r). \end{aligned}$$

Every record has a nonempty fiber, so $B^\top$ is injective. Its vectors in $\mathrm{row}(A_r)$ are exactly the measurable polynomials characterized above. Therefore, writing $q_0 = Bp_0$,

$$q \in \mathrm{aff}(\mathcal{Q}) \iff \langle q - q_0, f_A \rangle = 0 \quad (A \in \mathcal{B}),$$
$$\langle q_0, f_A \rangle = \mathbb{E}_{p_0} \prod_{i \in A} (1 - X_i) = 2^{-|A|},$$
$$\dim \mathrm{aff}(\mathcal{Q}) = M - |\mathcal{B}|.$$

These are (33). □

At $b = r = 2$, $K \geq 3$, the dimension is six, recovering the earlier affine constraints. They do not characterize positive completion: the exact $K = 4$ certificate in Appendix A.8 satisfies this whole affine hull and every individual capacity but has distance $1/24$ from $\mathcal{Q}$. The lemma removes all possible missing affine equalities from that question; the remaining issue is the projected positivity domain.

A.8 A FINITE CERTIFICATE FOR THE MISSING POSITIVITY CONSTRAINTS

The full affine hull does not determine the positive record domain. This already occurs for $K = 4$ and $b = r = 2$. Let $\mathcal{Q} = \mathcal{Q}^{\mathrm{can}}_{4,2,2}$ and define the probability vector

$$\bar{q}_{12} = \bar{q}_{13} = \bar{q}_{23} = \tfrac{1}{4}, \qquad \bar{q}_{\varnothing} = \tfrac{1}{6}, \qquad \bar{q}_{\{4\}} = \tfrac{1}{12}, \qquad \bar{q}_S = 0 \text{ otherwise.}$$

It satisfies every affine equation of $\mathcal{Q}$ and every individual record-capacity constraint. Its total mass at each record depth is also realized by a feasible law. Nevertheless,

$$\inf_{q \in \mathcal{Q}} \mathrm{TV}(\bar{q}, q) = \tfrac{1}{24}. \tag{35}$$

Here subscripts ij abbreviate the record $\{i, j\}$.

To verify the affine claim, list assignments in lexicographic order $0000, 0001, \ldots, 1111$ and use the signed vector

$$\bar{P} = \tfrac{1}{12}(2, 1, 0, 0, 0, 0, 1, 2, 0, 0, 1, 2, 1, 2, 1, -1).$$

Its moment sums and record projection are

$$\sum_x \bar{P}(x) = 1, \qquad \left(\sum_x x_i \bar{P}(x) \right)_{i=1}^{4} = \tfrac{1}{2}(1, 1, 1, 1),$$

$$\left(\sum_x x_i x_j \bar{P}(x) \right)_{i<j} = \tfrac{1}{4}(1, 1, 1, 1, 1, 1), \qquad (C_2)_{\#} \bar{P} = \tfrac{1}{12}(3e_{12} + 3e_{13} + 3e_{23} + 2e_{\varnothing} + e_{\{4\}}).$$

Here e_S is the unit vector at record S. The strictly positive uniform assignment law makes the affine image of these moment equations equal to $\mathrm{aff}(\mathcal{Q})$, as proved in Lemma 7. For the depth-two and depth-three nonzero records,

$$q_{12} = \tfrac{1}{4}, \qquad q_{13} \leq P(X_1 = X_3 = 1) = \tfrac{1}{4}, \quad q_{23} \leq P(X_2 = X_3 = 1) = \tfrac{1}{4}.$$

The even-parity law on the first three bits attains all three bounds. Every depth-four record is a single assignment and has capacity $1/6$: if D counts disagreements with that assignment, then

$$\mathbf{1}\{D = 0\} \leq (D - 2)(D - 3)/6 \qquad (D \in \{0, 1, 2, 3, 4\}),$$
$$\mathbb{E}_P D = 4/2 = 2, \qquad \mathbb{E}_P D(D - 1) = 4 \cdot 3/4 = 3,$$
$$\mathbb{E}_P[(D - 2)(D - 3)/6] = \{\mathbb{E}_P D(D - 1) - 4\mathbb{E}_P D + 6\}/6$$
$$= (3 - 8 + 6)/6 = 1/6.$$

Equality is attained by the uniformly lifted count law $P(D = 0) = 1/6$, $P(D = 2) = 1/2$, $P(D = 3) = 1/3$. Its first two factorial moments are 2 and 3, respectively. Thus every nonzero coordinate of $\bar{q}$ respects its exact capacity.

Set $f(x) = \mathbf{1}\{C_2(x) = \{1,3\}\} + \mathbf{1}\{x = 0000\}$. The two events are disjoint. The quadratic polynomial

$$g(x) = 1 - (x_1 + x_3)/2 - x_4 - x_1x_2/2 + x_1x_3 + x_1x_4/2 \\ - x_2x_3/2 + x_2x_4/2 + x_3x_4/2$$

majorizes f on every assignment. Its sixteen slacks, in the same lexicographic order, are

$$(g - f)(x) = \tfrac{1}{2}(0,0,1,0,2,1,0,0,1,0,0,0,0,0,0,1) \geq 0.$$

The linear coefficients of g sum to -2 and its pair coefficients sum to $3/2$. Hence every pairwise-balanced law satisfies

$$\mathbb{E}_P g = 1 - 2 \cdot \tfrac{1}{2} + \tfrac{3}{2} \cdot \tfrac{1}{4} = \tfrac{3}{8}, \qquad \bar{q}_{13} + \bar{q}_{\varnothing} = \tfrac{1}{4} + \tfrac{1}{6} = \tfrac{5}{12}.$$

Since f is an indicator,

$$\mathrm{TV}(\bar{q}, q) \geq \bar{q}(f) - q(f) \geq 5/12 - 3/8 = 1/24 \qquad (q \in \mathcal{Q}).$$

For attainment, take a uniform even-parity law on the first three bits and an independent fair fourth bit. Its record vector q^* has $q^*_{12} = q^*_{13} = q^*_{23} = 1/4$ and $q^*_{\varnothing} = q^*_{\{4\}} = 1/8$. Thus

$$\mathrm{TV}(\bar{q}, q^*) = \tfrac{1}{2}\left(\left|\tfrac{1}{6} - \tfrac{1}{8}\right| + \left|\tfrac{1}{12} - \tfrac{1}{8}\right|\right) = \tfrac{1}{24}.$$

Both vectors have depth masses $(1/4, 1/2, 1/4)$ at depths two, three, and four. This supplies a positive-distance counterexample to the affine-plus-capacity representation. It does not rule out a different polynomial-size extended formulation of the complete domain.

B Multiscale cancellation and positive repair

A single rare record can be made heavy under a finite-order moment constraint. The statistical lower bound needs a different object: every prescribed tuple of records at different depths must be heavy under one common nonnegative law. The difficulty is visible already for fair bits. Let U be the uniform assignment law, let σ_h fix the target cylinder in block h and use fair independent future bits, and set

$$\mu = (1 - \sum_h v_h)U + \sum_h v_h \sigma_h, \qquad v_h = \gamma N_h^{-d}, \quad N_h = B2^h.$$

Choose $\gamma B^{-d}/(1 - 2^{-d}) \leq 1$, so that $\sum_h v_h \leq 1$. For a set A of $2d$ coordinates in the zero remainder of block h, earlier cylinder laws average its Walsh product to zero and later laws fix it to one. Hence

$$\widehat{\mu}(A) := \mathbb{E}_\mu \prod_{i \in A} (-1)^{X_i} = \sum_{j \geq h} v_j, \qquad \sum_{|A| = 2d} \widehat{\mu}(A)^2 \geq \sum_h \binom{N_h - 2d}{2d} v_h^2 \geq c_d \gamma^2 L \tag{36}$$

for $B \geq 8d$, since

$$\binom{N_h - 2d}{2d} v_h^2 = \frac{\gamma^2}{(2d)!} \prod_{j=0}^{2d-1} \frac{N_h - 2d - j}{N_h} \geq \frac{\gamma^2}{2^{2d}(2d)!}.$$

Increasing B therefore does not make this correction input small uniformly in the number of scales.

The proof cancels these degree-$2d$ terms before repairing residual moments. Count lifting first replaces the individual bit equations by factorial moments of group counts. Gaussian quadrature then supplies a bounded central law with a prescribed deficit:

$$\tau(f) = \Gamma(f) - d! a^{2d} \sum_h \hom_{2d}(\bar{f})(w_h).$$

Here Γ uses Gaussian large-group coordinates, $\bar{f}$ averages the selected-group coordinates, and w_h points toward the hth rare cylinder. The decaying frame distributes this deficit among arbitrarily many directions while keeping each coordinate bounded. The cylinder measures restore this leading tensor. At odd order the reflected law has weight $\rho^{-(2d+1)}$ and reflection ratio $-\rho$. Thus both

the degree-$2d+1$ term and the degree-$2d$ term times a selected linear factor have multiplier $1-\rho^{-(2d+1)}\rho^{2d+1}=0$; Lemma 11 bounds the terms that remain.

Counts must be integers. Mean-preserving adjacent-grid rounding changes higher moments. Lemma 10 sums this error over all mixed indices rather than bounding each index and multiplying by the number of groups. Its layer-cake estimate gives a total $O(B^{-1}(1+\log B)^{r-2})$ error; the remaining binomial and cylinder errors are $O(B^{-1/2})$.

The last repair uses a fixed positive count reference R. For its orthonormal products ϕ_α, put $e_\alpha=\mu(\phi_\alpha)$. If $\mu\ge\lambda\mathsf{R}$, then

$$\mu'=\mu-\mathsf{R}\sum_{1\le|\alpha|\le r}e_\alpha\phi_\alpha,\qquad \mu'\ge\left(\lambda-\sup_\alpha\|\phi_\alpha\|_\infty\sum_\alpha|e_\alpha|\right)\mathsf{R}.$$

This is why the proof controls an ℓ_1 sum: it bounds the negative part of this particular correction pointwise. The reference has positive prelude count, so its support misses every prescribed cylinder. The correction fixes the moments without changing those cylinder masses. Finally, independent thinning produces the specified heterogeneous biases.

This repair is distinct from the positive-mixture construction of O'Donnell & Zhao (2018, Theorem 1). Their intrinsic Walsh-energy bound is dimension-free and their mixture preserves all events up to one factor. The ℓ_1 condition above is sufficient for our fixed-reference repair, not necessary for every possible repair. The common obstacle exhibited by (36) is the uncancelled leading energy. Appendix C.4 supplies every numerical constant for a complete fourth-order specialization and its fifth-order extension.

C Simultaneous positive moment constructions

C.1 Count identities and a positive reference

For $n\in\mathbb{N}$ write $(z)_0=1$ and $(z)_j=z(z-1)\cdots(z-j+1)$. Define

$$H_{n,j,h}(z)=\frac{(z)_h(n-z)_{j-h}}{(n)_j},\qquad H_{n,j,h}(z)=H_{n,j+1,h}(z)+H_{n,j+1,h+1}(z)\quad(0\le h\le j<n).\tag{37}$$

For integer $z\in[0,n]$, this is the probability of a specified j-bit pattern with h ones under uniform sampling without replacement. For $j<n$, the recurrence follows from

$$\begin{aligned}H_{n,j+1,h}(z)+H_{n,j+1,h+1}(z)&=\frac{(z)_h(n-z)_{j-h}}{(n)_{j+1}}\big[(n-z-j+h)+(z-h)\big]\\&=\frac{(z)_h(n-z)_{j-h}(n-j)}{(n)_j(n-j)}=H_{n,j,h}(z).\end{aligned}$$

Lemma 8 (Count lifting and positive moment repair)**.** *Let a probability law ρ specify group counts $Z_l\in\{0,\dots,n_l\}$. Uniform independent lifting within each group gives*

$$\mathbb{E}\prod_{i\in A}X_i=\mathbb{E}_\rho\prod_l\frac{(Z_l)_{|A\cap B_l|}}{(n_l)_{|A\cap B_l|}}.\tag{38}$$

Let ν be a finite-support probability measure, μ a probability measure, and $\lambda\ge0$. Let $f_\alpha\in L^1(\mu)$ be a finite orthonormal family in $L^2(\nu)$ with $\nu f_\alpha=0$. If $\mu\ge\lambda\nu$, put $e_\alpha=\mu f_\alpha$. Then

$$\mu'=\mu-\nu\sum_\alpha e_\alpha f_\alpha,\qquad\sum_\alpha|e_\alpha|\|f_\alpha\|_{\infty,\nu}\le\lambda\quad\Longrightarrow\quad\mu'\ge0,\ \mu'1=1,\ \mu' f_\alpha=0.\tag{39}$$

Here $\|f\|_{\infty,\nu}=\max_{\operatorname{supp}\nu}|f|$. The correction vanishes outside $\operatorname{supp}\nu$.

Proof. Put $a_l=|A\cap B_l|$. Conditional on the counts, the groups are independent and

$$P(X_i=1\text{ for every }i\in A\mid(Z_l)_l)=\prod_l\frac{\binom{n_l-a_l}{Z_l-a_l}}{\binom{n_l}{Z_l}}=\prod_l\frac{(Z_l)_{a_l}}{(n_l)_{a_l}},$$

$$\mathbb{E}\prod_{i\in A}X_i=\mathbb{E}_\rho\prod_l\frac{(Z_l)_{a_l}}{(n_l)_{a_l}}.$$

The ratio is zero when $Z_l < a_l$ and one when $a_l = 0$. For the second statement,

$$\mu' 1 = \mu 1 - \sum_\alpha e_\alpha \nu f_\alpha = 1,$$

$$\mu' f_\beta = e_\beta - \sum_\alpha e_\alpha \nu (f_\alpha f_\beta) = 0,$$

$$\mu' \geq \nu \left(\lambda - \sum_\alpha |e_\alpha| \|f_\alpha\|_{\infty,\nu} \right) \geq 0.$$

□

Let $G \sim N(0,1)$ and let Z_d have the positive d-node Gaussian quadrature law. With monic probabilists' Hermite polynomial He_d,

$$\mathbb{E}g(Z_d) - \mathbb{E}g(G) = -d![z^{2d}]g \qquad (\deg g \leq 2d+1). \tag{40}$$

The quadrature rule is classical (Golub & Welsch, 1969). Its nodes are the zeros of He_d, its law is symmetric, and it matches Gaussian polynomials of degree at most $2d-1$. Write

$$g(z) = az^{2d+1} + b\mathrm{He}_d(z)^2 + q(z), \qquad \deg q \leq 2d-1, \quad b = [z^{2d}]g.$$

Symmetry and exactness give

$$\begin{aligned} \mathbb{E}g(Z_d) - \mathbb{E}g(G) &= a \underbrace{\left(\mathbb{E}Z_d^{2d+1} - \mathbb{E}G^{2d+1}\right)}_{0} + b\big(0 - \mathbb{E}\mathrm{He}_d(G)^2\big) + \underbrace{\big(\mathbb{E}q(Z_d) - \mathbb{E}q(G)\big)}_{0} \\ &= -b\, d!. \end{aligned}$$

Lemma 9 (A uniformly positive finite count reference)**.** *Fix $r \geq 2$ and $p_0 \in (0, 1/2]$. For $p \in [p_0, 1-p_0]$ and $n \geq n_0(r, p_0)$, there is a law $\nu_{n,p}$ on $r+1$ grid points $z = (c - np)/\sqrt{np(1-p)}$, $c \in \{1, \ldots, n-1\}$, whose integer-count pushforward matches $\mathrm{Bin}(n,p)$ through order r. Its standardized nodes are bounded and separated, its weights are at least $a_0 > 0$, and its orthonormal polynomials through degree r have uniformly bounded coefficients. In particular $\phi_0 = 1$, $\phi_1 = z$, and $\|\phi_k\|_{\infty,\nu} \leq a_0^{-1/2}$.*

Proof. Let $(\xi_l, \pi_l)_{l=0}^r$ be the positive Gaussian quadrature rule. With $q = 1-p$, choose

$$c_{l,n} = \mathrm{round}(np + \sqrt{npq}\,\xi_l), \quad z_{l,n} = (c_{l,n} - np)/\sqrt{npq}, \quad V_n = (z_{l,n}^j)_{0 \leq j,l \leq r}.$$

The rounding error satisfies $|z_{l,n} - \xi_l| \leq (2\sqrt{npq})^{-1}$. Since $pq \geq p_0(1-p_0)$, increasing n_0 ensures $1 \leq c_{l,n} \leq n-1$ and $|z_{l,n} - z_{k,n}| \geq |\xi_l - \xi_k|/2$ for $l \neq k$. If $Z_n = (\mathrm{Bin}(n,p) - np)/\sqrt{npq}$ and $\mu_n = (\mathbb{E}Z_n^j)_{j=0}^r$, the partition expansion gives, uniformly in p,

$$\mathbb{E}Z_n^j = \sum_{\pi \in \mathfrak{P}_j^{\geq 2}} \frac{(n)_{|\pi|}}{n^{j/2}} \prod_{B \in \pi} \mathbb{E}\left[\left(\frac{X-p}{\sqrt{pq}} \right)^{|B|} \right] = \mathbb{E}G^j + O_{r,p_0}(n^{-1/2}).$$

Here $X \sim \mathrm{Bernoulli}(p)$ and the sum is over set partitions without singleton blocks; at $j = 0$ use the empty partition. Each pair block contributes $\mathbb{E}(X-p)^2/(pq) = 1$. For even j, the pair partitions contribute $(j-1)!!(n)_{j/2} n^{-j/2} = (j-1)!! + O_r(n^{-1})$. Every other partition has $|\pi| \leq (j-1)/2$, so its contribution is $O_{r,p_0}(n^{-1/2})$. The number of partitions depends only on r. With $V_\infty = (\xi_l^j)_{j,l}$, these bounds yield

$$\begin{aligned} V_n &= V_\infty + O(n^{-1/2}), \qquad \det V_\infty \neq 0, \\ w_n &= V_n^{-1} \mu_n = \pi + O(n^{-1/2}), \qquad \min_l w_{l,n} \geq \tfrac{1}{2} \min_l \pi_l =: a_0. \end{aligned}$$

The law $\nu_{n,p}(z_{l,n}) = w_{l,n}$ has the required standardized moments; the inverse affine map recovers the integer-count law. The polynomial Gram matrices satisfy

$$M_n = V_n \,\mathrm{diag}(w_n) V_n^\top \longrightarrow V_\infty \,\mathrm{diag}(\pi) V_\infty^\top \succ 0.$$

Their smallest eigenvalues are bounded away from zero. Gram–Schmidt in this fixed dimension gives uniformly bounded coefficients. The first two matched moments give $\phi_1 = z$, and

$$1 = \sum_l w_l \phi_k(z_l)^2 \geq a_0 \max_l \phi_k(z_l)^2.$$

□

C.2 MIXED-DEGREE SUMMATION AND THE ODD-ORDER CYLINDER CORRECTION

The next two lemmas isolate the estimates whose constants must be uniform in the number of geometric blocks. Throughout this subsection, He_j denotes the probabilists' Hermite polynomial, normalized by $\exp(tz-t^2/2)=\sum_{j\ge 0}\mathrm{He}_j(z)t^j/j!$. A bounded noise family means $\sup_\nu \|\xi_\nu\|_\infty \le C$, with one fixed C independent of the numbers of groups and noise coordinates.

Lemma 10 (Dimension-uniform mixed interpolation error)**.** *Fix $r \ge 2$, $\kappa > 1$ and positive constants c, C. Let $I \subset \mathbb{N}_0$ be finite, with group sizes $cB2^i \le n_i \le CB2^i$, $i \in I$, where $B \ge 2$. Suppose $V = \mathcal{K}\xi$, where the coordinates ξ_ν are independent and bounded, $\mathcal{K}\mathcal{K}^\top = I$, and*

$$|\mathcal{K}_{i\nu}| \le C\kappa^{-|i-\ell(\nu)|}, \qquad \sup_\nu \sum_i |\mathcal{K}_{i\nu}| \le C, \qquad \#\{\nu : \ell(\nu) = j\} \le C.$$

Assume $|V_i| \le C$ and $\mathbb{E}\,\mathrm{He}_j(\xi_\nu) = 0$ for $1 \le j \le r-2$. For $0 \le k \le r$, let $\phi_{i,k}$ be polynomials of degree at most k with uniformly bounded coefficients, $\phi_{i,0} = 1$, $\phi_{i,1}(z) = z$, and $|\mathbb{E}\phi_{i,k}(G)| \le Cn_i^{-1/2}$ for $k \ge 1$, where G is standard normal. Let conditionally independent coordinate randomizations satisfy

$$\mathbb{E}[\phi_{i,k}(V_i^{\mathrm{lat}}) \mid V] = \phi_{i,k}(V_i) + r_{i,k}(V_i), \qquad r_{i,0} = r_{i,1} = 0,$$

with $\|r_{i,k}\|_\infty \le Cn_i^{-1}$ and $\mathrm{Lip}(r_{i,k}) \le Cn_i^{-1/2}$ on the relevant bounded interval. Put $\phi_\alpha = \prod_i \phi_{i,\alpha_i}$. Then

$$\sum_{1\le|\alpha|\le r} |\mathbb{E}\phi_\alpha(V^{\mathrm{lat}}) - \mathbb{E}\phi_\alpha(V)| \le C_r B^{-1}(1+\log B)^{r-2}. \tag{41}$$

The constant depends on the displayed fixed constants and r, but not on $|I|$ or the number of noise coordinates.

Proof. For a nonempty residual set $A \subseteq I$ and degrees $k_i \ge 2$, put

$$f_A = \prod_{i\in A} r_{i,k_i}(V_i), \quad P_A = \prod_{i\in A} n_i^{-1}, \quad L_A = 1 + \sum_{i\in A}\log n_i, \quad \eta_\nu = \min\{1, \sum_{i\in A}\sqrt{n_i}\kappa^{-|i-\ell(\nu)|}\}.$$

If V' replaces only ξ_ν by an independent copy, then

$$\begin{aligned} |V_i - V_i'| &\le C|\mathcal{K}_{i\nu}|, \\ |f_A(V) - f_A(V')| &\le \sum_{i\in A} |r_{i,k_i}(V_i) - r_{i,k_i}(V_i')| \prod_{j\in A\setminus\{i\}} \|r_{j,k_j}\|_\infty \\ &\le C_r P_A \sum_{i\in A} \sqrt{n_i}\kappa^{-|i-\ell(\nu)|}. \end{aligned}$$

The separate bound $2\|f_A\|_\infty \le C_r P_A$ gives $\mathrm{osc}_\nu(f_A) \le C_r P_A \eta_\nu$. For $0 < t \le 1$, the implication

$$\eta_\nu \ge t \implies \exists i \in A : \quad |i - \ell(\nu)| \le \frac{\log(|A|\sqrt{n_i}/t)}{\log\kappa}$$

and the bounded multiplicity of each noise position give

$$\begin{aligned} &\#\{\nu : \eta_\nu \ge t\} \le C_r[L_A + \log(1/t)], \\ &\sum_{\nu_1,\dots,\nu_v} \min_{j\le v} \eta_{\nu_j} = \int_0^1 \#\{\nu : \eta_\nu \ge t\}^v \, dt \\ &\qquad \le C_r \int_0^\infty (L_A + x)^v e^{-x}\, dx = C_r \sum_{j=0}^v \binom{v}{j} L_A^{v-j} j! \le C_r' L_A^v, \quad 1 \le v \le r-2. \end{aligned} \tag{42}$$

Let $\mathcal{R}_\nu$ substitute an independent copy of ξ_ν in f_A and write $\Delta_\nu = I - \mathcal{R}_\nu$. For a set U of distinct noise indices and positive integers m_ν with $\sum_{\nu\in U} m_\nu \le r-2$, independence and the zero Hermite

means imply

$$\mathbb{E}\left[f_A \prod_{\nu\in U} \mathrm{He}_{m_\nu}(\xi_\nu)\right] = \mathbb{E}\left[(\prod_{\nu\in U} \Delta_\nu f_A) \prod_{\nu\in U} \mathrm{He}_{m_\nu}(\xi_\nu)\right],$$
$$\|\prod_{\nu\in U} \Delta_\nu f_A\|_\infty \le 2^{|U|-1} \min_{\nu\in U} \mathrm{osc}_\nu(f_A) \le C_r P_A \min_{\nu\in U} \eta_\nu. \tag{43}$$

For example, every term containing $\mathcal{R}_\nu$ has expectation zero after integrating the original ξ_ν. Expanding the product of Δ_ν therefore leaves only the original term in expectation. The norm bound follows by keeping any one difference and bounding each of the remaining difference operators by two.

The algebraic identity $\mathcal{K}\mathcal{K}^\top = I$ gives the generating function identity

$$\begin{aligned}\sum_\beta \mathrm{He}_\beta(V)\frac{t^\beta}{\beta!} &= \exp(t^\top \mathcal{K}\xi - \|t\|_2^2/2)\\ &= \prod_\nu \exp\big((\mathcal{K}_{\cdot\nu}^\top t)\xi_\nu - (\mathcal{K}_{\cdot\nu}^\top t)^2/2\big)\\ &= \prod_\nu \sum_{j\ge 0} \mathrm{He}_j(\xi_\nu)\frac{(\mathcal{K}_{\cdot\nu}^\top t)^j}{j!}.\end{aligned}$$

At total degree v, the column-sum bound, the bounded noise supports and (43) give

$$\begin{aligned}\sum_{|\beta|=v} \frac{|\mathbb{E}[f_A \mathrm{He}_\beta(V)]|}{\beta!} &\le C_r P_A \sum_{|m|=v} \frac{\min_{\nu:m_\nu>0}\eta_\nu}{\prod_\nu m_\nu!} \prod_\nu \|\mathcal{K}_{\cdot\nu}\|_1^{m_\nu}\\ &\le \frac{C_r' P_A}{v!} \sum_{\nu_1,\dots,\nu_v} \min_j \eta_{\nu_j} \le C_r'' P_A L_A^v.\end{aligned} \tag{44}$$

The coefficient counting uses that an ordered v-tuple with multiplicities m has $v!/\prod m_\nu!$ orderings. For $v=0$, use $|\mathbb{E}f_A| \le C_r P_A$.

We now count all original mixed polynomial indices, including those whose Hermite expansion has constant terms. Conditional independence gives

$$\mathbb{E}\phi_\alpha(V^{\mathrm{lat}}) - \mathbb{E}\phi_\alpha(V) = \sum_{\substack{\varnothing\ne A\subseteq \mathrm{supp}\,\alpha\\ \alpha_i\ge 2\ (i\in A)}} \mathbb{E}\left[\prod_{i\in A} r_{i,\alpha_i}(V_i) \prod_{j\notin A} \phi_{j,\alpha_j}(V_j)\right]. \tag{45}$$

If $|A| = q$, then

$$\sum_{i\in A} \alpha_i \ge 2q, \qquad \sum_{j\notin A} \alpha_j \le r-2q, \qquad q \le \lfloor r/2\rfloor. \tag{46}$$

Expand $\phi_{j,k} = \sum_{v=0}^k b_{j,k,v}\,\mathrm{He}_v$. The nonconstant coefficients are bounded by a fixed constant, while $|b_{j,k,0}| \le C n_j^{-1/2}$. For any specified nonzero Hermite multi-index β, there are at most $(r+1)^r$ choices of positive degrees on its support. Summing over additional coordinates assigned Hermite degree zero costs at most

$$\prod_{j\notin A}\left(1 + \sum_{k=1}^r |b_{j,k,0}|\right) \le \exp(Cr\sum_j n_j^{-1/2}) \le C_r. \tag{47}$$

There are also at most r^r choices of residual degrees on a fixed A. Since $\beta! \le r!$, equations (44)–(47) give the following bound on the contribution of a fixed A:

$$C_r \sum_{v=0}^{r-2q} \sum_{|\beta|=v} |\mathbb{E}[f_A \mathrm{He}_\beta(V)]| \le C_r' P_A \sum_{v=0}^{r-2q} L_A^v \le C_r'' P_A L_A^{r-2q}.$$

Finally, for $1 \le q \le \lfloor r/2 \rfloor$,

$$\sum_{|A|=q} P_A L_A^{r-2q} \le C_r B^{-q}(1+\log B)^{r-2q} \sum_{i_1<\cdots<i_q} 2^{-\sum_j i_j}(1+\sum_j i_j)^{r-2q}$$
$$\le C_r' B^{-q}(1+\log B)^{r-2q}.$$

For $v = r-2q$, the series is bounded by $[\sum_{i\ge 0} 2^{-i}(1+i)^v]^q < \infty$, since $1+\sum_j i_j \le \prod_j (1+i_j)$. Hence the total is at most

$$C_r \sum_{q=1}^{\lfloor r/2 \rfloor} B^{-q}(1+\log B)^{r-2q} \le C_r' B^{-1}(1+\log B)^{r-2},$$

as asserted in (41). □

Lemma 11 (Mixed remainders of the biased cylinder pair)**.** *Fix $r \in \{2d, 2d+1\}$, $b = 2d$, $p \in [1/2, 1-p_0]$, $q = 1-p$ and $\rho = q/p$. Use a prelude and geometric blocks, each separated into a selected group of b positions and a large remainder. Standardize a group of size n by $(C - np)/\sqrt{npq}$. On every large group use the polynomials of Lemma 9; on every selected group use the standardized* $\mathrm{Bin}(b,p)$ *orthonormal polynomials of degrees at most b. All multi-indices below have total degree at most r and selected-group degrees at most b. Let w_h be the large-group zero-count vector through block h, divided by $\sqrt{N_h}$, with zero future coordinates.*

Define σ_h^- by fixing all earlier groups and the current large remainder to count zero, the current selected group to count b, and using independent reference laws on all future groups. At odd order define σ_h^+ by using full counts in all earlier groups and the current large remainder, count $\mathrm{Bin}(b, 1-\rho)$ *in the current selected group, and the same future reference laws. Put*

$$(D_h, c_e, s_e) = \begin{cases} (\sigma_h^-, 1, 1), & r = 2d, \\ (\sigma_h^- + \rho^{-(2d+1)}\sigma_h^+, 1+\rho^{-1}, 1+\rho^{-(2d+1)}), & r = 2d+1. \end{cases}$$

For $\phi_\alpha = \prod_i \phi_{i,\alpha_i}$, average all selected groups under their reference laws to obtain $\overline{\phi_\alpha}$, and let $H_{2d,\overline{\phi_\alpha}}$ be its homogeneous part of degree $2d$ in the large variables. Then

$$D_h(\phi_\alpha) = c_e N_h^d H_{2d,\overline{\phi_\alpha}}(w_h) + R_{\alpha h}, \qquad \sum_{1\le|\alpha|\le r} |R_{\alpha h}| \le C_{r,p_0}(h+1)^r N_h^{d-1/2}. \tag{48}$$

Proof. All standardized selected counts are bounded uniformly in h. The two large vectors are $\sqrt{N_h}w_h$ and $-\rho\sqrt{N_h}w_h$. Earlier selected standardized counts have the same reflection ratio. For the current selected group, with $t_* = 1-\rho$, its exact mean obeys

$$\begin{aligned} t_* - p &= 1 - q/p - p = -q^2/p = -\rho(1-p), \\ \mathbb{E}_{\sigma_h^+} z_{A_h} &= -\rho \mathbb{E}_{\sigma_h^-} z_{A_h}, \qquad 1 - \rho^{-(2d+1)}\rho^{2d+1} = 0. \end{aligned} \tag{49}$$

The polynomial of selected degree one is precisely z, with no constant term. Polynomials of selected degree greater than b are not included; their count support has only $b+1$ points.

A positive index on a future group has zero expectation under its independent reference, so $D_h(\phi_\alpha) = 0$. The leading homogeneous term also vanishes: a future linear factor vanishes at w_h, while taking a constant term of a future polynomial of degree at least two leaves degree at most $r-2 < 2d$. A future selected factor is removed by its zero reference mean. It remains to consider at most $2h+3$ nonfuture groups.

Let s be the total selected polynomial degree and expand the remaining polynomial into homogeneous large-variable terms. When $s = 0$, the even term of degree $2d$ has multiplier $1 + \rho^{-(2d+1)}\rho^{2d} = c_e$ at odd order and multiplier one at even order. A large term of degree $2d+1$ has multiplier zero by (49). When $s = 1$, the only possible large term of degree $2d$ occurs at odd order, and its additional selected linear factor gives the same zero multiplier. If $s \ge 2$, the large degree is at most $r - s \le 2d - 1$. Thus, after extracting the displayed degree-$2d$ term, the degree bound is

selected degree s	0	1	≥ 2
remaining large degree	$\le 2d-1$	$\le 2d-1$	$\le r-s \le 2d-1$.

All coefficients have fixed bounds, and every nonfuture large standardized count has magnitude at most $C_{p_0}\sqrt{N_h}$. Consequently,

$$|R_{\alpha h}| \le C_{r,p_0} \sum_{v=0}^{2d-1} N_h^{v/2} \le C'_{r,p_0} N_h^{d-1/2}, \qquad \sum_{1\le|\alpha|\le r} |R_{\alpha h}| \le \binom{2h+3+r}{r} C'_{r,p_0} N_h^{d-1/2},$$

which gives (48). □

C.3 Simultaneous rare records with nonuniform product moments

The construction uses Gaussian quadrature (Golub & Welsch, 1969) within the finite count-moment framework (Prékopa, 1990). The proof below derives the groupwise count lifting and combines it with the simultaneous cylinder correction of Lemma 11.

We prove Theorem 1. Set $b = 2d$.

Proof. We first construct common bias $p = 1 - p_0 \ge 1/2$. Set $q = 1 - p$ and $\rho = q/p \in (0, 1]$. All constants below depend only on r, p_0. Split each block into its selected group of size b and remainder of size $n_h = N_h - b$. Together with the prelude there are $2L + 1$ groups. The prelude and remainders are called large groups. For $B \ge 2b$ and each fixed $v > 0$,

$$\sum_{i \text{ large}} n_i^{-v} \le B^{-v} + 2^v B^{-v} \sum_{h\ge 0} 2^{-vh} = \left(1 + \frac{2^v}{1 - 2^{-v}}\right) B^{-v}. \tag{50}$$

Standardize a group count as $z_i = (c_i - n_i p)/\sqrt{n_i pq}$.

Use $\nu_{n_i,p}$ from Lemma 9 on every large group, with uniform constants a_0, a_1, C independent of the number of groups. Use the standardized pushforward of the exact $\mathrm{Bin}(b, p)$ law on selected groups and reduce a_0 to cover their positive masses. Write R for the product reference, with orthonormal polynomials $\phi_{i,k}$ and $\phi_{i,0} = 1$. Use $k \le r$ on large groups and $k \le b$ on selected groups. Uniform support separation and positivity bound all polynomial coefficients, and $|\phi_{i,k}| \le a_0^{-1/2}$ on reference support. In particular, $\phi_{i,1} = z_i$.

Retain every product $\phi_\alpha = \prod_i \phi_{i,\alpha_i}$ with $1 \le |\alpha| \le r$ and the stated coordinate-degree restrictions. If a probability measure μ contains $\lambda\mathsf{R}$, put $e_\alpha = \mu(\phi_\alpha)$. Then

$$\mu' = \mu - \mathsf{R} \sum_{1\le|\alpha|\le r} e_\alpha \phi_\alpha, \qquad \sum_\alpha |e_\alpha| \le \tfrac{\lambda}{2} a_0^{r/2} \implies \mu' \ge \tfrac{\lambda}{2}\mathsf{R}. \tag{51}$$

Equation (51) is Lemma 8 with $\|\phi_\alpha\|_{\infty,\mathsf{R}} \le a_0^{-r/2}$. A selected group has only $b + 1$ count values: polynomials of larger degree reduce to degree at most b on this support, without increasing total degree. No nonexistent selected-group polynomial of degree $b + 1$ is used at odd order.

Let w_h be the large-group standardized zero-count vector of the hth cylinder, divided by $\sqrt{N_h}$. Its prelude coordinate is $-\sqrt{p/q}\sqrt{B/N_h}$, its remainder-j coordinate for $j \le h$ is $-\sqrt{p/q}\sqrt{n_j/N_h}$, and future coordinates are zero. Put $W = (w_h)_h$. Give the prelude position zero, remainder h position $h + 1$, and column h position $h + 1$. With weights $\kappa^{|i-j|}$, $\kappa = 2^{1/4}$, geometric sums bound both absolute row and column norms of W by C_{p_0}. Choose $a > 0$ sufficiently small that $a^2\|WW^\top\|_{\infty,\kappa} < 1/4$, uniformly in L, and define

$$\mathsf{C} = I - a^2 WW^\top, \qquad V = aWZ + \mathsf{C}^{1/2}\Xi = \mathcal{K}\xi.$$

Here V is the standardized count vector; T denotes the selector. The coordinates Z_h are independent d-node Gaussian quadrature variables and Ξ_i are independent $(r + 1)$-node Gaussian quadrature variables, independent of Z. The convergent square-root power series and submultiplicativity of the weighted norms give

$$\mathcal{K}\mathcal{K}^\top = I, \quad |\mathcal{K}_{i\nu}| \le C\kappa^{-|i-\ell(\nu)|}, \quad \sup_\nu \sum_i |\mathcal{K}_{i\nu}| \le C, \quad |V_i| \le C. \tag{52}$$

Here each integer position has at most two noise coordinates, and $\ell(\nu)$ is the position of noise ξ_ν. The matrix identity is algebraic; it does not assert that the non-Gaussian V has covariance I when $d = 1$.

Let Γ use independent Gaussian large groups and independent selected reference groups, and let τ replace the large vector by V. For a polynomial f of total degree at most r, first average its selected variables under their reference laws, obtaining $\bar{f}$. Let $H_{2d,\bar{f}}$ be its homogeneous large-variable part of degree $2d$. For each Z_h, (40) gives $\mathbb{E}Z_h^{2d} - \mathbb{E}G^{2d} = -d!$, with all lower moments equal. In degree at most $r \le 2d+1$, at most one noise coordinate can contribute such a deficit. Its coefficient at degree $2d$ is $a^{2d}H_{2d,\bar{f}}(w_h)$. Therefore

$$\tau(f) = \Gamma(f) - d!a^{2d}\sum_h H_{2d,\bar{f}}(w_h). \tag{53}$$

At degree $2d+1$, a single odd noise power vanishes by symmetry, and a $2d$th-power deficit multiplied by another first-degree noise vanishes by independence and centering. This verifies the identity at odd order as well, including $d = 1$.

Conditionally on V, independently round each large coordinate to the adjacent standardized count-grid points, retaining the independent selected references. For sufficiently large B all these counts are legal. Write τ_{lat} for this distribution. On a grid cell of width $s_i = (n_i pq)^{-1/2}$, the linear-interpolation remainder satisfies

$$|r_{i,k}(z)| \le s_i^2 \sup|\phi''_{i,k}|/8, \qquad |r'_{i,k}(z)| \le s_i \sup|\phi''_{i,k}|, \qquad r_{i,0} = r_{i,1} = 0.$$

The interpolation is continuous at the cell boundaries. Thus the supremum and global Lipschitz assumptions of Lemma 10 hold.

Write $\phi_{i,k}(z) = \sum_{j=0}^k c_{i,k,j}z^j$ and let Z_i^{bin} be a standardized $\mathrm{Bin}(n_i, p)$ variable. The binomial expansion in Lemma 9 gives $|\mathbb{E}G^j - \mathbb{E}(Z_i^{\mathrm{bin}})^j| \le C_{r,p_0}n_i^{-1/2}$ for $j \le r$. Since the reference matches these moments and the coefficients are bounded,

$$|\mathbb{E}\phi_{i,k}(G)| = \left|\sum_{j=0}^k c_{i,k,j}\{\mathbb{E}G^j - \mathbb{E}(Z_i^{\mathrm{bin}})^j\}\right| \le C_{r,p_0}n_i^{-1/2}, \qquad 1 \le k \le r.$$

Index the large prelude by $i = 0$ and the large remainder of block h by $i = h+1$. For $B \ge 2b$, their sizes satisfy

$$n_0 = B, \quad n_{h+1} = B2^h - b, \qquad (B/4)2^i \le n_i \le B2^i.$$

This is the geometric indexing required by the summation lemma, with constants independent of the number of blocks. The noise variables match Gaussian moments through at least $r-2$. Equation (52) checks the remaining frame hypotheses. Lemma 10 therefore gives

$$\sum_\alpha |\tau_{\mathrm{lat}}(\phi_\alpha) - \tau(\phi_\alpha)| \le C_{r,p_0}B^{-1}(1+\log B)^{r-2}.$$

Nonzero selected-reference indices have expectation zero under both laws and contribute nothing. The same zero-Hermite coefficient estimate and independence give

$$\sum_\alpha |\Gamma(\phi_\alpha) - \mathsf{R}(\phi_\alpha)| \le \prod_{i \text{ large}}(1 + Crn_i^{-1/2}) - 1 \le \exp\left(Cr\sum_{i \text{ large}} n_i^{-1/2}\right) - 1 \le C_{r,p_0}B^{-1/2}. \tag{54}$$

Use D_h, c_e, s_e from Lemma 11. At odd order its current selected count is

$$\mathrm{Bin}(b, t_*), \qquad t_* = (2p-1)/p = 1-\rho, \qquad \beta_* = \rho^{-(2d+1)}. \tag{55}$$

The selected mean obeys $\mathbb{E}_{\sigma_h^+}z_{A_h} = -\rho\mathbb{E}_{\sigma_h^-}z_{A_h}$, as required by (49).

Take $\omega = 1/2$, $\gamma_0 = (1-\omega)d!a^{2d}/c_e$, $v_h = \gamma_0 N_h^{-d}$, $P_0 = s_e\sum_h v_h$, and

$$\mu = (\omega - P_0)\mathsf{R} + (1-\omega)\tau_{\mathrm{lat}} + \sum_h v_h D_h.$$

Its mass and reference coefficient satisfy

$$\mu(1) = (\omega - P_0) + (1-\omega) + s_e\sum_h v_h = 1, \qquad \lambda = \omega - P_0 \ge \frac{1}{2} - \frac{s_e\gamma_0 B^{-d}}{1-2^{-d}} \ge \frac{1}{4}$$

for sufficiently large B. Since $-(1-\omega)d!a^{2d}+\gamma_0 c_e=0$, the two leading tensors cancel. For every retained α, $\mathsf{R}(\phi_\alpha)=0$ and

$$e_\alpha=(1-\omega)(\tau_{\mathrm{lat}}-\tau)(\phi_\alpha)+(1-\omega)(\Gamma-\mathsf{R})(\phi_\alpha)+\sum_h v_h R_{\alpha h},$$

$$\sum_\alpha |e_\alpha| \le C\left\{\frac{(1+\log B)^{r-2}}{B}+B^{-1/2}+\sum_{h\ge 0}(h+1)^r N_h^{-1/2}\right\}$$
$$\le C'\{B^{-1/2}+(1+\log B)^{r-2}/B\}.$$

The last step uses $\sum_{h\ge0}(h+1)^r N_h^{-1/2}=B^{-1/2}\sum_{h\ge0}(h+1)^r 2^{-h/2}<\infty$. Choose B_0 depending only on r,p_0 so that all earlier reference conditions hold and this bound is at most $a_0^{r/2}/8$. Then $\sum_\alpha |e_\alpha|\le a_0^{r/2}/8\le(\lambda/2)a_0^{r/2}$, so (51) gives a nonnegative probability with every required moment. Its correction has positive prelude count, disjoint from each E_h, so $P(E_h)\ge v_h$ remains.

Uniform count lifting (38) gives, for $|S|\le r$,

$$\mathbb{E}\prod_{i\in S}Y_i=\mathbb{E}\prod_j\frac{(C_j)_{|S\cap B_j|}}{(n_j)_{|S\cap B_j|}}=\prod_j p^{|S\cap B_j|}=p^{|S|}.$$

For the desired heterogeneous means draw independent $G_i\sim\mathrm{Bernoulli}(p_i/p)$ and set $X_i=Y_iG_i$. Independence gives

$$\mathbb{E}\prod_{i\in S}X_i=p^{|S|}\prod_{i\in S}\frac{p_i}{p}=\prod_{i\in S}p_i,\qquad P_X(E_h)\ge P_Y(E_h)\prod_{i\in A_h}\frac{p_i}{p}\ge\gamma_0 N_h^{-d}(p_0/p)^b.$$

Thus $\gamma=\gamma_0(p_0/p)^b$ has the required value. Appending independent bits with their prescribed means preserves moments and the target high-record cylinders. □

For $d=2$, the common identity (53) reads

$$\mathbb{E}V_iV_jV_kV_l=\delta_{ij}\delta_{kl}+\delta_{ik}\delta_{jl}+\delta_{il}\delta_{jk}-2a^4\sum_h w_{h,i}w_{h,j}w_{h,k}w_{h,l}.$$

At $r=5$, $1-\beta_*\rho^5=0$ and $c_e=1+\rho^{-1}$, so $v_h=(1-\omega)2a^4N_h^{-2}/c_e$ cancels that same fourth-order tensor.

C.4 An explicit fourth-order construction at every scale

The following specialization supplies numerical parameters for the entire multiscale construction. Its cutoff is deliberately conservative; it is an existence threshold, not an estimate of the smallest feasible prefix.

Proposition 12 (Explicit balanced fourth-order parameters)**.** *Set $B_0=2^{192}$ and $\gamma=2^{-24}$. For every even integer $B\ge B_0$, every $L\ge1$, and every choice of four specified high positions in each block of size $N_h=B2^h$, $0\le h<L$, there is a four-wise independent fair law on the prelude and all blocks such that*

$$P(E_h)\ge 2^{-24}N_h^{-2}\qquad(0\le h<L).$$

Here E_h fixes the prelude and previous blocks to zero, fixes the current block to the specified four highs, and leaves future bits free. The construction and every constant below are uniform in L. For fair fifth-order moments the same B_0 is valid with $\gamma=2^{-25}$.

Proof. We retain the count construction but give its reference, frame and residual bounds explicitly. A large group has an even size $n\ge256$. Define

$$m_n=\lceil\sqrt{n}/2\rceil,\quad t_n=4m_n^2/n,\quad \beta_n=3-2/n,$$
$$1\le t_n\le(1+2/\sqrt{n})^2\le 81/64,\quad 383/128\le\beta_n<3.$$

Put reference mass C_n at count $n/2$, mass $A_n/2$ at each of $n/2 \pm m_n$, and mass $D_n/2$ at each of $n/2 \pm 2m_n$, where

$$A_n = \frac{4t_n - \beta_n}{3t_n^2}, \qquad D_n = \frac{\beta_n - t_n}{12t_n^2}, \qquad C_n = 1 - A_n - D_n. \tag{56}$$

Since $2m_n \le \sqrt{n} + 2 < n/2$, all five counts lie strictly between zero and n. For the standardized count $z = 2(c - n/2)/\sqrt{n}$, the five nodes are $0, \pm\sqrt{t_n}, \pm 2\sqrt{t_n}$. Their separation is at least one, and their absolute values are at most $9/4$. Substituting (56) gives

$$\begin{aligned}
\mathbb{E}z = \mathbb{E}z^3 = 0, \qquad \mathbb{E}z^2 &= \frac{4t_n - \beta_n + \beta_n - t_n}{3t_n} = 1, \\
\mathbb{E}z^4 &= \frac{4t_n - \beta_n + 4(\beta_n - t_n)}{3} = \beta_n, \\
D_n/2 &\ge \frac{383/128 - 81/64}{24(81/64)^2} = \frac{884}{19683} > \frac{1}{32}, \\
A_n/2 &\ge \frac{1}{6(81/64)^2} = \frac{2048}{19683} > \frac{1}{32}, \\
C_n &= \frac{4t_n^2 - 5t_n + \beta_n}{4t_n^2} \ge \frac{383/128 - 1}{4(81/64)^2} = \frac{2040}{6561} > \frac{1}{32}.
\end{aligned} \tag{57}$$

These are precisely the first four standardized binomial moments. Use this reference on each large group and the exact $\mathrm{Bin}(4, 1/2)$ law on each selected group. Write R for their product and $\phi_{i,k}$ for their orthonormal polynomials, with $\phi_{i,0} = 1$. On a large group,

$$\phi_{i,1}(z) = z, \qquad \phi_{i,2}(z) = \frac{z^2 - 1}{\sqrt{2 - 2/n_i}}, \qquad \max_{z \in \operatorname{supp} \mathsf{R}_i} |\phi_{i,k}(z)| < 6.$$

Indeed, $1 \ge \mathsf{R}_i\{z\}\phi_{i,k}(z)^2 > |\phi_{i,k}(z)|^2/32$. For nodes $z_0, \dots, z_4$, Lagrange interpolation gives

$$\|\operatorname{coeff} \phi_{i,k}\|_1 \le \sum_{j=0}^{4} |\phi_{i,k}(z_j)| \frac{\prod_{l \ne j}(1 + |z_l|)}{\prod_{l \ne j} |z_j - z_l|} < 5 \cdot 6\,(13/4)^4/4 < 2^{12}.$$

The denominator four is the smallest product of the four node distances. The same bound holds for the selected-group polynomials, whose nodes are $-2, -1, 0, 1, 2$ and minimum mass is $1/16$. Symmetry makes every odd polynomial odd. Consequently, for a standard normal G,

$$\mathbb{E}\phi_{i,k}(G) = 0 \ (1 \le k \le 3), \qquad |\mathbb{E}\phi_{i,4}(G)| \le 2^{13}/n_i. \tag{58}$$

Indeed the only differing moment through degree four is $\mathbb{E}G^4 - \beta_{n_i} = 2/n_i$.

Index the prelude by $i = 0$ and the large remainder of block h by $i = h + 1$. Thus $n_0 = B$ and $n_{h+1} = B2^h - 4$. Define

$$W_{ih} = -\sqrt{n_i/N_h}\, \mathbf{1}\{i \le h + 1\}, \qquad a = 1/64, \qquad \mathsf{C} = I - a^2 W W^\top.$$

The position of column h is $h + 1$. For $\kappa = 2^{1/4}$, use the maximum of the weighted absolute row and column sums as a matrix norm; weights are $\kappa^{|i-j|}$. For $u = 2^{-1/4} < 7/8$, the weighted row and column sums are at most $\sqrt{2}/(1 - u)$ and $\sqrt{2} + 1/(1 - u)$, respectively. Hence

$$\begin{aligned}
\|W\|_\kappa \le 16, \qquad a^2\|WW^\top\|_\kappa &\le (16/64)^2 = 1/16, \\
\|\mathsf{C}^{1/2}\|_\kappa \le 1 + \sum_{j \ge 1} \left| \binom{1/2}{j} \right| 16^{-j} &= 2 - \sqrt{15/16} < 2.
\end{aligned}$$

The absolutely convergent square-root series therefore defines $\mathsf{C}^{1/2}$. Let the independent variables Z_h be fair signs, and let independent Ξ_i have masses $1/6, 2/3, 1/6$ at $-\sqrt{3}, 0, \sqrt{3}$. Put

$$V = aWZ + \mathsf{C}^{1/2}\Xi = \mathcal{K}\xi. \quad \mathcal{K}\mathcal{K}^\top = I, \quad \|\mathcal{K}\|_\kappa \le 16/64 + 2 < 3, \quad |V_i| \le 2\|\mathcal{K}\|_\kappa \le 6. \tag{59}$$

There are at most two noise variables at each integer position, and $|\mathcal{K}_{i\nu}| \le 3\kappa^{-|i-\ell(\nu)|}$. Let τ be the law using V on the large groups and independent selected references. Let G instead use

independent standard normal large groups. For a polynomial f of total degree at most four, average its selected coordinates to get $\bar{f}$, and write $\mathrm{hom}_4(\bar{f})$ for its homogeneous part of degree four. Since $\mathbb{E}Z_h^4 - 3 = -2$ and every Ξ_i matches the first four Gaussian moments, expansion of the independent noises gives

$$\tau(f) = \mathsf{G}(f) - 2a^4 \sum_h \mathrm{hom}_4(\bar{f})(W_{\cdot h}). \tag{60}$$

All products of degree at most three agree; at degree four the displayed term is the sum of the individual fourth cumulants.

Round each large count conditionally and independently to its two adjacent integer counts, preserving its conditional mean. Write τ_{lat} for this law and $r_{i,k}(v) = \mathbb{E}[\phi_{i,k}(V_i^{\mathrm{lat}}) \mid V_i = v] - \phi_{i,k}(v)$. The grid spacing is $2/\sqrt{n_i}$, and all relevant nodes lie in $[-7, 7]$. The coefficient bound above implies $\sup_{[-7,7]} |\phi''_{i,k}| \le 2^{12} \cdot 12 \cdot 7^2 < 2^{25}$. With grid width $s_i = 2/\sqrt{n_i}$, linear interpolation gives $\|r_{i,k}\|_\infty \le s_i^2 2^{25}/8 = 2^{24}/n_i$ and $\mathrm{Lip}(r_{i,k}) \le s_i 2^{25} = 2^{26}/\sqrt{n_i}$. In particular,

$$r_{i,0} = r_{i,1} = 0, \qquad \|r_{i,k}\|_\infty \le C_R/n_i, \qquad \mathrm{Lip}(r_{i,k}) \le C_R/\sqrt{n_i}, \quad C_R = 2^{27}. \tag{61}$$

The following calculation makes the number-of-scales dependence explicit. Fix one residual $R_i = r_{i,k}(V_i)$ and set

$$A_i = 1 + \log n_i, \qquad \eta_{i\nu} = \min\{1, 6\sqrt{n_i}\kappa^{-|i-\ell(\nu)|}\}.$$

Replacing ξ_ν changes V_i by at most $12\kappa^{-|i-\ell(\nu)|}$, so the change in R_i is at most

$$\min\left\{\frac{2C_R}{n_i}, \frac{12C_R}{\sqrt{n_i}}\kappa^{-|i-\ell(\nu)|}\right\} = \frac{2C_R\eta_{i\nu}}{n_i}.$$

For $0 < t \le 1$, counting integer positions and using $\log\kappa > 1/8$ yields

$$\begin{aligned} \#\{\nu : \eta_{i\nu} \ge t\} &\le 128[A_i + \log(1/t)], \\ \sum_\nu \eta_{i\nu} \le 2^8 A_i, \qquad \sum_{\nu,\mu} \min(\eta_{i\nu}, \eta_{i\mu}) &\le 2^{17} A_i^2. \end{aligned} \tag{62}$$

For example, the last sum is the integral of the squared first count: $128^2 \int_0^1 [A_i + \log(1/t)]^2 dt = 128^2(A_i^2 + 2A_i + 2) \le 2^{17} A_i^2$. Centering of ξ_ν and of $\xi_\nu^2 - 1$, followed by one or two independent replacements, gives

$$\begin{aligned} |\mathbb{E}R_i\xi_\nu| &\le 4C_R\eta_{i\nu}/n_i, \\ |\mathbb{E}R_i(\xi_\nu^2 - 1)| &\le 6C_R\eta_{i\nu}/n_i, \\ |\mathbb{E}R_i\xi_\nu\xi_\mu| &\le 16C_R \min(\eta_{i\nu}, \eta_{i\mu})/n_i \quad (\nu \ne \mu). \end{aligned}$$

The column sums of $\mathcal{K}$ are at most three. The total absolute coefficient of $\xi_\nu^2 - 1$ in the quadratic sums below is at most $\|\mathcal{K}_{\cdot\nu}\|_1^2 \le 9$; that of $\xi_\nu\xi_\mu$ is at most $2\|\mathcal{K}_{\cdot\nu}\|_1\|\mathcal{K}_{\cdot\mu}\|_1 \le 18$. Expanding V_j, $V_j^2 - 1$, and V_jV_l in these centered noise polynomials therefore bounds their summed correlations by

$$\begin{aligned} \sum_j |\mathbb{E}R_iV_j| &\le 2^{12} C_R A_i/n_i, \\ \sum_j |\mathbb{E}R_i(V_j^2 - 1)| + \sum_{j<l} |\mathbb{E}R_iV_jV_l| &\le \frac{C_R}{n_i}\left(54\sum_\nu \eta_{i\nu} + 288\sum_{\nu<\mu} \min(\eta_{i\nu}, \eta_{i\mu})\right) \\ &\le (54 \cdot 2^8 + 144 \cdot 2^{17}) C_R A_i^2/n_i < 2^{25} C_R A_i^2/n_i. \end{aligned}$$

At total degree four, a single residual has only zero, one, or two remaining polynomial degrees. Their polynomials are $1, V_j, (V_j^2 - 1)/\sqrt{2 - 2/n_j}$, or V_jV_l; the denominator is at least one. Thus the single-residual terms for fixed i sum to at most $[3 + 2 \cdot 2^{12} A_i + 2^{25} A_i^2] C_R/n_i \le 2^{27} C_R A_i^2/n_i$. The only two-residual term is $r_{i,2}r_{j,2}$, whose sum is at most $C_R^2(\sum_i n_i^{-1})^2/2$. Write $A = 1 + \log B$.

Since $n_{h+1} \ge B2^h/2$ and $A_{h+1} \le A + h\log 2$,

$$\sum_i n_i^{-1} \le \frac{1}{B} + \frac{2}{B}\sum_{h\ge 0} 2^{-h} = \frac{5}{B},$$

$$\sum_i \frac{A_i^2}{n_i} \le \frac{A^2}{B} + \frac{2}{B}\sum_{h\ge 0} 2^{-h}(A + h\log 2)^2$$

$$= \frac{5A^2 + 8A\log 2 + 12(\log 2)^2}{B} \le \frac{25A^2}{B} \le \frac{32A^2}{B}.$$

Consequently, for products ϕ_α of total degree at most four,

$$\sum_{1\le|\alpha|\le 4} |\tau_{\mathrm{lat}}(\phi_\alpha) - \tau(\phi_\alpha)| \le 2^{59}(1+\log B)^2/B + 25\cdot 2^{53}/B^2 \le 2^{64}(1+\log B)^2/B. \quad (63)$$

Positive selected indices contribute zero under both laws. Equation (58) also gives $\sum_\alpha |\mathsf{G}(\phi_\alpha) - \mathsf{R}(\phi_\alpha)| \le 2^{13}\sum_i n_i^{-1} \le 5\cdot 2^{13}/B < 2^{16}/B$: at this total degree, only a single fourth-degree large index can contribute.

Let σ_h fix all preceding counts and the current large count to zero, fix the selected count in block h to four, and use independent references on future groups. A positive future polynomial index has zero expectation. Among the other indices, remove the homogeneous large-variable term of degree four. Each remaining monomial has large variable degree at most three, including indices with positive selected degree. Every large zero count has standardized magnitude at most $\sqrt{N_h}$; each selected standardized count has magnitude at most two. The product coefficient norm is at most $(2^{12})^4$, so each remainder has magnitude at most $2^{48}\cdot 2^4 N_h^{3/2} \le 2^{64}N_h^{3/2}$. There are at most $\binom{2h+7}{4} \le 2^{12}(h+1)^4$ nonfuture indices. Hence

$$\sigma_h(\phi_\alpha) = N_h^2 \hom_4(\overline{\phi_\alpha})(W_{\cdot h}) + \mathcal{E}_{\alpha h},$$
$$|\mathcal{E}_{\alpha h}| \le 2^{64}N_h^{3/2}, \qquad \sum_\alpha |\mathcal{E}_{\alpha h}| \le 2^{76}(h+1)^4 N_h^{3/2}. \quad (64)$$

The needed sum follows from

$$\sum_{h\ge 0}(h+1)^4 2^{-h/2} \le \left.\frac{1+11x+11x^2+x^3}{(1-x)^5}\right|_{x=3/4} = 16240 < 2^{15}.$$

Set $v_h = a^4 N_h^{-2}$, $v = \sum_h v_h$ and $\lambda = 1/2 - v$. Since $v \le (4/3)a^4B^{-2}$, one has $\lambda \ge 1/4$. The probability measure

$$\mu = \lambda\mathsf{R} + \tfrac{1}{2}\tau_{\mathrm{lat}} + \sum_h v_h\sigma_h$$

has mass $\lambda + 1/2 + v = 1$. The degree-four coefficient is $-(1/2)2a^4 + v_h N_h^2 = 0$. For $e_\alpha = \mu(\phi_\alpha)$,

$$e_\alpha = \tfrac{1}{2}(\tau_{\mathrm{lat}} - \tau)(\phi_\alpha) + \tfrac{1}{2}(\mathsf{G} - \mathsf{R})(\phi_\alpha) + \sum_h v_h\mathcal{E}_{\alpha h},$$

$$\sum_\alpha |e_\alpha| \le 2^{63}\frac{(1+\log B)^2}{B} + \frac{2^{15}}{B} + 2^{-24}2^{76}B^{-1/2}\sum_{h\ge 0}(h+1)^4 2^{-h/2}$$

$$\le 2^{63}\frac{(1+\log B)^2}{B} + \frac{2^{15}}{B} + 2^{67}B^{-1/2}.$$

In particular,

$$\sum_{1\le|\alpha|\le 4} |e_\alpha| \le 2^{70}\left[B^{-1/2} + \frac{(1+\log B)^2}{B}\right] < 2^{-25} < 2^{-13} \qquad (B \ge 2^{192}). \quad (65)$$

For the numerical inequality it suffices to use $\log 2 < 1$ at $B = 2^{192}$; both functions in the brackets decrease thereafter. In particular, $2^{70}(2^{-96} + 193^2 2^{-192}) < 2^{-25}$. On the reference support, $|\phi_\alpha| \le (\sqrt{32})^4 = 2^{10}$ for $|\alpha| \le 4$. Orthonormality now proves

$$\mu' = \mu - \mathsf{R}\sum_{1\le|\alpha|\le 4} e_\alpha\phi_\alpha, \qquad \mu' \ge (\lambda - 2^{10}\sum_\alpha |e_\alpha|)\mathsf{R} \ge (\lambda/2)\mathsf{R}, \quad \mu'(\phi_\alpha) = 0. \quad (66)$$

The mass remains one. Reference prelude counts are strictly positive, so this correction is disjoint from every E_h and $\mu'(E_h) \geq v_h$. Finally choose a uniform subset of the stipulated count in each group. For any set S of at most four original bit positions, count lifting gives

$$\mathbb{E}\prod_{j\in S} X_j = \mathbb{E}_{\mu'}\prod_i \frac{(C_i)_{|S\cap\mathcal{B}_i|}}{(n_i)_{|S\cap\mathcal{B}_i|}} = 2^{-|S|}.$$

Here $\mathcal{B}_i$ ranges over both the large and the selected groups, and its size is n_i (four for a selected group). This is the required balanced law with the asserted unchanged cylinder masses. For order five, let Y have this fourth-order law, draw an independent fair bit J, and set $X_i = Y_i \oplus J$. For every nonempty set S of at most five indices,

$$\mathbb{E}\prod_{i\in S}(-1)^{X_i} = \mathbb{E}\prod_{i\in S}(-1)^{Y_i}\,\mathbb{E}(-1)^{|S|J} = 0 :$$

the last factor is zero at odd cardinality, and at even cardinality $|S| \leq 4$ the first factor is zero. The branch $J = 0$ retains every cylinder, so $P_X(E_h) \geq P_Y(E_h)/2 \geq 2^{-25}N_h^{-2}$. □

For $b = 4$, $r \in \{4, 5\}$, $\epsilon \leq 2^{-1024}$, $\delta \leq 1/16$, and $K \geq \lceil \epsilon^{-1/2} \rceil$, the construction also makes the basic critical scalar lower bound quantitative:

$$\mathcal{N}^{\mathrm{sc}}_{4,r}(\epsilon, \delta; K) \geq 2^{-80}\epsilon^{-2}[1 + \log(1/\epsilon)]^2 \log(8/\delta). \tag{67}$$

Here $\mathcal{N}^{\mathrm{sc}}_{4,r}$ is the worst-case expected number of fresh Bernoulli record observations for estimating the upper support value within ϵ with probability $1 - \delta$ on the balanced canonical domain. To check the constants, set $B = 2^{192}$, $t = \log_2(1/\epsilon) \geq 1024$, $L = \lfloor t/4 \rfloor$ and $p_h = \gamma_r N_h^{-2}$, where $\gamma_4 = 2^{-24}$ and $\gamma_5 = 2^{-25}$. Set $M_h = \binom{N_h}{4}$. Then

$$B2^L \leq 2^{192+t/4} \leq 2^{t/2} = \epsilon^{-1/2},$$

$$M_h p_h^2 = \binom{N_h}{4}\gamma_r^2 N_h^{-4} \geq \gamma_r^2/384 \geq 2^{-59},$$

$$a_h = \frac{4\epsilon}{Lp_h} \leq \frac{4\epsilon\, 2^{25} B^2 2^{2(L-1)}}{L} \leq \frac{2^{409}\sqrt{\epsilon}}{L} \leq \frac{2^{-103}}{L} < \frac{1}{4}.$$

For every tuple of one candidate per block, raise only those means from $1/2$ by a_h. One law from the proposition raises the support value by at least $\sum_h p_h a_h = 4\epsilon$. All other means stay $1/2$. Let N_h^{obs} count queries in block h under the common flat experiment, denoted by expectation $\mathbb{E}_0$. Threshold the estimated value at $1/2 + 2\epsilon$. The two successful-output sets are disjoint. For a fixed candidate tuple $I = (I_h)_h$, each query to I_h contributes

$$\mathrm{kl}(1/2, 1/2 + a_h) = -\tfrac{1}{2}\log(1 - 4a_h^2) \leq \frac{2a_h^2}{1 - 4a_h^2} \leq \tfrac{8}{3}a_h^2 \leq 3a_h^2.$$

All other queries have zero KL. Apply Lemma 15 with these response laws and average over uniform independent I_h. The flat experiment is independent of I, so $\mathbb{E}_I\mathbb{E}_0 N^{\mathrm{obs}}_{I_h} = \mathbb{E}_0 N_h^{\mathrm{obs}}/M_h$. With $u = \log(8/\delta)$, this gives

$$\begin{aligned}\frac{u}{4} &\leq (1 - \delta)\log(1/\delta) - \log 2 \leq \mathrm{kl}(1 - \delta, \delta)\\ &\leq 3\sum_h \frac{a_h^2}{M_h}\mathbb{E}_0 N_h^{\mathrm{obs}} \leq \frac{64\epsilon^2}{L^2}\sum_h \frac{\mathbb{E}_0 N_h^{\mathrm{obs}}}{M_h p_h^2} \leq 2^{65}\epsilon^2 \mathbb{E}_0 N/L^2.\end{aligned}$$

The first inequality uses $\delta \leq 1/16$. Since $L \geq t/4 - 1 \geq [1 + \log(1/\epsilon)]/16$,

$$\mathbb{E}_0 N \geq 2^{-67}\epsilon^{-2}L^2 u \geq 2^{-75}\epsilon^{-2}[1 + \log(1/\epsilon)]^2 u,$$

which implies (67). The finite cutoff is explicit but conservative; it is not substituted for the finite-K bounds at ordinary candidate counts.

C.5 AN EXACT SIMULTANEOUS-CYLINDER CONSTANT AT ORDERS TWO AND THREE

The following result concerns joint laws themselves. The events are mutually exclusive cylinders, and the same law realizes their masses.

We prove the sharp part of Theorem 1. Throughout, $r = 2, 3$, $B = 2^a \geq 4$, and the selected pairs are fixed consistently across finite prefixes.

Proof. Let $V_h = \operatorname{span}_{\mathbb{F}_2}\{e_1, \ldots, e_{a+h+1}\}$. Choose B distinct nonzero prelude vectors spanning V_0, including $e_1, e_2, e_1 + e_2$. Such a set exists: start with a basis and $e_1 + e_2$, then add unused vectors, since $a + 2 \leq 2^a$. Inductively, the vectors assigned before block h are N_h distinct nonzero elements spanning V_h. Exactly $N_h - 1$ nonzero vectors of V_h remain unused. Assign any $N_h - 2$ of them to the unselected positions of block h. Assign to its two selected positions

$$e_{a+h+2}, \qquad e_{a+h+2} + e_1.$$

These are distinct and outside V_h. After the block, $2N_h = N_{h+1}$ assigned vectors span V_{h+1}. Thus all assigned vectors v_i are distinct and nonzero in the countable binary direct sum.

Let $U_1, U_2, \ldots$ be independent fair bits and first put $X_i = \langle v_i, U\rangle$ over $\mathbb{F}_2$. For distinct i, j, the map $U \mapsto (\langle v_i, U\rangle, \langle v_j, U\rangle)$ has rank two. Its four outputs have equal numbers of preimages, so $P(X_i = u, X_j = v) = 1/4$ for $u, v \in \{0, 1\}$. All bits before block h vanish exactly when $U|_{V_h} = 0$. On that event all unselected positions also vanish, and both selected positions equal U_{a+h+2}. Therefore

$$E_h = \{U|_{V_h} = 0,\ U_{a+h+2} = 1\}, \qquad P(E_h) = 2^{-(a+h+2)} = \frac{1}{4N_h}.$$

This is one infinite product-space probability construction, not a collection of incompatible finite maximizers.

For order three draw one more independent fair bit G and set $X_i = G + \langle v_i, U\rangle$. Every coefficient vector is $(1, v_i)$. A sum of one or three such vectors has first coordinate one. A sum of two has nonzero last coordinates because $v_i \neq v_j$. Thus every set of at most three columns is linearly independent, and each associated output pattern has probability $2^{-|S|}$. Since the prelude contains $e_1, e_2, e_1 + e_2$, its three associated bits sum to G. The event that the prelude vanishes forces $G = 0$. Consequently every E_h now has exactly half its previous mass. At L blocks the ordinary columns span V_L, of dimension $a + L + 1$. The three special augmented prelude columns sum to $(1, 0)$, so the augmented rank is $a + L + 2$. Thus the seed-to-output maps are injective, proving the exact support sizes and seed counts.

For $0 \leq h < L$, the events E_j, $h \leq j < L$, are disjoint and their union F_h has at most two ones among the first $2N_h$ bits. Thus

$$P(F_h) = \sum_{j=h}^{L-1} P(E_j) \geq \frac{2\gamma^*_{r,L}}{N_h}[1 - 2^{-(L-h)}] \tag{68}$$

for a maximizing finite law, which exists by compactness. At order two, S counting those $2N_h$ bits satisfies

$$\mathbb{E}S = \sum_{i=1}^{2N_h} \tfrac{1}{2} = N_h, \qquad \operatorname{Var} S = \sum_i \tfrac{1}{4} + 2\sum_{i<j}(\tfrac{1}{4} - \tfrac{1}{4}) = N_h/2.$$

Since $\{S \leq 2\} \subseteq \{|S - N_h| \geq N_h - 2\}$,

$$P(F_h) \leq P(S \leq 2) \leq \frac{\mathbb{E}(S - N_h)^2}{(N_h - 2)^2} = \frac{N_h}{2(N_h - 2)^2}, \qquad \gamma^*_{2,L} \leq \frac{N_h^2}{4(N_h - 2)^2[1 - 2^{-(L-h)}]}.$$

At order three, F_h additionally forces the first prelude bit to zero. Conditional on this bit being zero, the remaining bits are pairwise balanced: each pattern on it and two other bits has mass $1/8$, and the conditioning event has mass $1/2$. Their sum S' has conditional mean $(2N_h - 1)/2$ and variance $(2N_h - 1)/4$. Thus

$$\begin{aligned} P(F_h) &= \tfrac{1}{2} P(F_h \mid X_{\mathrm{pre}} = 0) \\ &\leq \frac{1}{2}\,\frac{(2N_h - 1)/4}{((2N_h - 1)/2 - 2)^2} = \frac{2N_h - 1}{2(2N_h - 5)^2}. \end{aligned}$$

Substitution in (68) gives

$$\gamma^*_{3,L} \le \frac{N_h(2N_h-1)}{4(2N_h-5)^2[1-2^{-(L-h)}]}.$$

All denominators are positive because $N_h \ge B \ge 4$. For $N \ge 4$ and $0 \le y \le 1/2$,

$$\left(1+\frac{12}{N}\right)(N-2)^2 - N^2 = \frac{8(N-4)^2+20(N-4)}{N} \ge 0,$$

$$\left(1+\frac{21}{N}\right)(2N-5)^2 - 2N(2N-1) = \frac{66(N-4)^2+133(N-4)+1}{N} \ge 0,$$

$$(1-y)^{-1} = 1 + \frac{y}{1-y} \le 1+2y.$$

Choose $h = \lfloor L/2 \rfloor$ and $y = 2^{-(L-h)}$. Then

$$\gamma^*_{2,L} - \tfrac{1}{4} \le \tfrac{1}{4}\{(1+12/N_h)(1+2y)-1\} \le 6/N_h + y/2 \le 2^{1-h},$$
$$\gamma^*_{3,L} - \tfrac{1}{8} \le \tfrac{1}{8}\{(1+21/N_h)(1+2y)-1\} \le 21/(4N_h) + y/4 \le 2^{1-h}.$$

The last inequalities use $N_h \ge 4 \cdot 2^h$ and $y \le 2^{-h}$. Restricting any feasible $(L+1)$-block law to its first L blocks shows $\gamma^*_{r,L+1} \le \gamma^*_{r,L}$; the explicit law supplies the matching lower constant for every L. □

Bit flips reduce the marginal single-cylinder problem to an all-zero atom. For $2N$ pairwise fair bits with count C, the integer-valued majorant and its expectation are

$$\mathbf{1}\{C=0\} \le \frac{(C-N)(C-N-1)}{N(N+1)}, \qquad P(C=0) \le \frac{N/2}{N(N+1)} = \frac{1}{2(N+1)}.$$

Uniform count lifting attains equality with counts $0, N, N+1$ of masses $1/[2(N+1)], 1/2, N/[2(N+1)]$. For $2N-1$ bits, the majorant $(C-N)^2/N^2$ gives $P(C=0) \le 1/(2N)$, attained with counts $0, N$ of masses $1/(2N), 1-1/(2N)$. Conditioning a three-wise law on its first bit being zero therefore gives the upper bound $1/(4N)$ at length $2N$. It is attained by $(J, Y_1 \oplus J, \ldots, Y_{2N-1} \oplus J)$, where Y has the preceding pairwise law and J is independent fair; the parity calculation used above verifies three-wise independence. Thus the normalized marginal optima are $N/[2(N+1)]$ and $1/4$. Their large-scale limits are twice those in (8); minimizing over scales at fixed B instead gives ratios $2B/(B+1)$ and two. These atom constants concern cross-scale compatibility of the laws, not the scalar sampling rate.

D PAID OBSERVATION BOUNDS AND ACQUISITION RATES

Throughout this section, $\mathcal{D} \subseteq \Delta_M$ is nonempty and compact, $r_\pm(c) = \operatorname{ext}_{q\in\mathcal{D}} q^\top c$, and $0 \le c \le 1$. The filtration $\mathcal{F}_n$ contains the complete observable history and used private randomness. A fresh query I_n returns $Y_n \in [0,1]$ with

$$\mathbb{E}[Y_n \mid \mathcal{F}_{n-1}, I_n = i] = c_i. \tag{69}$$

All queries below are charged; a known deterministic coordinate is the only exception. The likelihood lemmas use the specified full response kernels, not only the conditional means.

D.1 ENDPOINT PADDING AND EXACT OBSERVATIONS

Lemma 13 (Common endpoint accounting)**.** *Let $0 \le R_- \le R_+ \le 1$, $\alpha, \tau_-, \tau_+ \ge 0$, $\beta > 0$, and $0 < \delta < 1$. Put $v_\pm = \operatorname{ext}_{\mathcal{D}} q^\top c$ and suppose $v_\pm - \tau_- \le R_\pm \le v_\pm + \tau_+$. If feasible q^-, q^+ and estimates $\widehat{m}_-, \widehat{m}_+$ satisfy*

$$\begin{gathered} v_- \le c^\top q^- \le v_- + \alpha, \qquad v_+ - \alpha \le c^\top q^+ \le v_+, \\ |\widehat{m}_\pm - c^\top q^\pm| \le \beta, \end{gathered} \tag{70}$$

then the clipped interval

$$[L,U] = [\widehat{m}_- - \alpha - \beta - \tau_-,\ \widehat{m}_+ + \alpha + \beta + \tau_+] \cap [0,1] \tag{71}$$

contains $[R_-, R_+]$ and has both endpoint excesses at most $\alpha + 2\beta + \tau_- + \tau_+$. For each evaluation, freeze $q^\pm$ before sampling and draw $\Pr(I_n = i \mid \mathcal{F}_{n-1}) = q_i^\pm$ afresh at every round. The mean of $m = \lceil \log(8/\delta)/(2\beta^2) \rceil$ responses has conditional failure at most $\delta/4$.

Proof. Set $\mu_\pm = c^\top q^\pm$. The frozen-law evaluation has $\mathbb{E}[Y_n \mid \mathcal{F}_{n-1}] = \sum_i q_i^\pm c_i = \mu_\pm$. Conditional Hoeffding and iterated expectation give

$$\mathbb{E}[e^{s(Y-\mu_\pm)} \mid \mathcal{F}] \le e^{s^2/8}, \qquad \Pr(|\widehat{m}_\pm - \mu_\pm| > \beta \mid \mathcal{F}_{\text{train}}) \le 2e^{-2m\beta^2} \le \delta/4.$$

Before clipping, the premises give

$$\begin{aligned} L \le \mu_- - \alpha - \tau_- \le v_- - \tau_- \le R_-, \\ U \ge \mu_+ + \alpha + \tau_+ \ge v_+ + \tau_+ \ge R_+, \\ R_- - L \le (v_- + \tau_+) - (\mu_- - \beta - \alpha - \beta - \tau_-) \le \alpha + 2\beta + \tau_- + \tau_+, \\ U - R_+ \le (\mu_+ + \beta + \alpha + \beta + \tau_+) - (v_+ - \tau_-) \le \alpha + 2\beta + \tau_- + \tau_+. \end{aligned}$$

Clipping improves both inequalities. The fallback $[0, 1]$ is required only when the *padded* endpoints are reversed. No independence between the training and evaluation success events is assumed. □

Lemma 14 (Exact capacity and common-baseline hitting)**.** *For $l \le c \le u$,*

$$r_\pm(l) \le r_\pm(c) \le r_\pm(u), \qquad r_\pm(u) - r_\pm(l) \le \max_{q \in \mathcal{D}} q^\top (u - l). \tag{72}$$

For an initially unrestricted cube $c \in [0, 1]^M$, each coordinate is available at fixed first-query price $a_i \ge 0$ and returns only its exact value; subsequent repeats may be cached. There is no initial cache or extra prerequisite. Among deterministic policies terminating and faithful for every c, the zero-failure minimax price is

$$\min_{S:\, \sup_{q \in \mathcal{D}} q(S^c) \le \epsilon} \sum_{i \in S} a_i. \tag{73}$$

For one randomized policy with the same initial information, valid at $c = 0$ and every relevant $c = e_i$, with failure at most $0 \le \delta < 1/2$, put $w_i = \max_{q \in \mathcal{D}} q_i$. Then

$$w_i > \epsilon \implies \Pr_0(i \textit{ queried}) \ge 1 - 2\delta, \qquad \mathbb{E}_0 \operatorname{cost} \ge (1 - 2\delta) \sum_{i : w_i > \epsilon} a_i. \tag{74}$$

Proof. Let $C = \max_q q^\top (u - l)$. For every q, $q^\top l \le q^\top c \le q^\top u \le q^\top l + C$. Taking minima or maxima gives (72). For a fixed queried set S, fill unqueried coordinates by zero and one, obtaining $l \le c \le u$ and $u - l = \mathbf{1}_{S^c}$. Then $[r_-(l), r_+(u)]$ covers the true interval with each excess at most $C = \sup_q q(S^c)$, proving sufficiency in (73). On a zero-response path querying S, the alternative $c' = \mathbf{1}_{S^c}$ has the same transcript; hence

$$\sup_q q(S^c) = r_+(c') \le U \le \epsilon.$$

For (74), couple the algorithm's private seed in worlds $0, e_i$. On $\{i \text{ not queried}\}$ the outputs coincide, but

$$U \le \epsilon \quad (c = 0), \qquad U \ge w_i > \epsilon \quad (c = e_i).$$

Thus $\Pr_0(i \text{ not queried}) \le \Pr_0(\text{fail}) + \Pr_{e_i}(\text{fail}) \le 2\delta$. Sum the paid indicators. This adapts the all-negative and single-positive construction of Sabato et al. (2013, Theorem 4.2) to the common-baseline expected query bound above. □

D.2 Stopped information

Lemma 15 (Full-response and censored change of measure)**.** *Fix $0 < \delta < 1/2$. Two worlds $0, u$ have the same joint law of initial information and private seed. The common history-adapted policy and output rule stop at $\tau_{\text{stop}} < \infty$ almost surely in both worlds. At round n, write A_n for the action and $\mathfrak{h}_{n-1}$ for the realized observable history. The transcript includes all actions, complete*

responses, stopping and output, with an absorbing state after stopping. Let $d_a(h) = \mathrm{KL}(K_a^0(\cdot \mid h) \| K_a^u(\cdot \mid h))$. *If exact responses coincide, then*

$$\mathrm{KL}(P_0^{\mathrm{tr}} \| P_u^{\mathrm{tr}}) = \mathbb{E}_0 \sum_{n \le \tau_{\mathrm{stop}}} d_{A_n}(\mathfrak{h}_{n-1}). \tag{75}$$

For disjoint successful-output sets, each reached with probability at least $1 - \delta$ *in its own world,*

$$\mathrm{kl}(1-\delta, \delta) \le \mathbb{E}_0 \sum_{n \le \tau_{\mathrm{stop}}} d_{A_n}(\mathfrak{h}_{n-1}). \tag{76}$$

Let the finite or countable set D_u *contain all distinguishing exact atoms; all other exact responses coincide. In the request–response filtration, censor before the response to the first requested atom in* D_u*, retaining that request and a censor marker. Write* σ *for that request time (*∞ *if none) and* $p = \Pr_0(\sigma \le \tau_{\mathrm{stop}})$. *Then, for the fresh actions,*

$$\mathbb{E}_0 \sum_{n \le \tau_{\mathrm{stop}},\, n < \sigma} d_{A_n}(\mathfrak{h}_{n-1}) \ge g_\delta(p), \quad g_\delta(p) = \begin{cases} \mathrm{kl}(1-\delta-p, \delta), & p < 1-2\delta, \\ 0, & p \ge 1-2\delta. \end{cases} \tag{77}$$

Proof. Append a common absorbing symbol after stopping. Given the initial information and private seed, the conditional action, stopping and output kernels are identical in the two worlds. Their conditional KL contributions are zero. For a deterministic cap m, the chain rule gives

$$\begin{aligned} \mathrm{KL}(P_0^{\tau_{\mathrm{stop}} \wedge m} \| P_u^{\tau_{\mathrm{stop}} \wedge m}) &= \sum_{n=1}^{m} \mathbb{E}_0 \big[\mathbf{1}_{\{n \le \tau_{\mathrm{stop}}\}} \, \mathrm{KL}\big(K_{A_n}^0(\cdot \mid \mathfrak{h}_{n-1}) \,\|\, K_{A_n}^u(\cdot \mid \mathfrak{h}_{n-1})\big) \big] \\ &= \mathbb{E}_0 \sum_{n \le \tau_{\mathrm{stop}} \wedge m} d_{A_n}(\mathfrak{h}_{n-1}). \end{aligned}$$

This also holds with infinite terms. The increasing capped histories generate the complete stopped transcript, including its output. Relative entropy on these increasing sigma-fields and monotone convergence of the nonnegative cost sum give

$$\mathrm{KL}(P_0^{\tau_{\mathrm{stop}} \wedge m} \| P_u^{\tau_{\mathrm{stop}} \wedge m}) \uparrow \mathrm{KL}(P_0^{\mathrm{tr}} \| P_u^{\mathrm{tr}}) = \mathbb{E}_0 \sum_{n \le \tau_{\mathrm{stop}}} d_{A_n}(\mathfrak{h}_{n-1}).$$

For $E = \{\widehat{O} \in \mathsf{G}_0\}$, $P_0 E \ge 1 - \delta$ and $P_u E \le \delta$. Binary KL increases in its first argument and decreases in its second when the first exceeds the second. Data processing therefore gives $\mathrm{KL}(P_0^{\mathrm{tr}} \| P_u^{\mathrm{tr}}) \ge \mathrm{kl}(P_0 E, P_u E) \ge \mathrm{kl}(1-\delta, \delta)$. In the censored experiment, the request at σ is observed but its response is replaced by a common marker. Its divergence is zero. The same capped chain rule therefore sums only over $n < \sigma$. The event $E_c = \{\tau_{\mathrm{stop}} < \sigma, \widehat{O} \in \mathsf{G}_0\}$ is observable, and

$$P_0 E_c \ge 1 - \delta - p, \qquad P_u E_c \le \delta, \qquad \mathrm{KL}(P_0^{\mathrm{cens}} \| P_u^{\mathrm{cens}}) = \mathbb{E}_0 \sum_{n \le \tau_{\mathrm{stop}}, n < \sigma} d_{A_n}(\mathfrak{h}_{n-1}).$$

If $p < 1-2\delta$, then $1-\delta-p > \delta$, so binary KL monotonicity gives $\mathrm{kl}(P_0 E_c, P_u E_c) \ge \mathrm{kl}(1-\delta-p, \delta)$. Otherwise use nonnegativity. This proves (77), the censored form of the stopped likelihood argument in Kaufmann et al. (2016, Lemma 1; Appendix A, Lemma 19). □

In the following two consequences all exact responses coincide. Every fresh full-response conditional KL is bounded by the displayed Bernoulli KL (in particular the full response may be that Bernoulli variable). For $0 \le a_i \le 1/4$,

$$\mathrm{kl}(1/2, 1/2 + a_i) = -\tfrac{1}{2}\log(1 - 4a_i^2) \le 3a_i^2, \qquad \tfrac{1}{8}u \le 3 \sum_i a_i^2 \mathbb{E}_0 N_i, \quad u = \log(8/\delta). \tag{78}$$

Indeed, $-\frac{1}{2}\log(1-4a_i^2) \le 2a_i^2/(1-4a_i^2) \le 8a_i^2/3$. Under the disjoint-output premise and $\delta \le 1/4$,

$$\mathrm{kl}(1-\delta, \delta) = (1-2\delta)\log\frac{1-\delta}{\delta} \ge \tfrac{1}{2}\log\frac{3}{4\delta} \ge \tfrac{1}{8}\log\frac{8}{\delta}.$$

Let the blocks B_h be disjoint, and sample uniform locations $I_h \in B_h$, $|B_h| = M_h$, once before the experiment. The entire null law is independent of these locations. Assume the full conditional divergence at action i is bounded by $3\sum_h a_h^2 \mathbf{1}\{i = I_h\}$ on every history. Then

$$\mathrm{KL}\left(P_0 \| \mathbb{E}_I P_I\right) \le \mathbb{E}_I \,\mathrm{KL}(P_0 \| P_I) \le 3 \sum_h \frac{a_h^2}{M_h} \mathbb{E}_0 N(B_h). \tag{79}$$

The first inequality is convexity of KL in its second argument. For the second, the complete null transcript is independent of I, so

$$\mathbb{E}_I \mathbb{E}_0 \sum_n 3 \sum_h a_h^2 \mathbf{1}\{A_n = I_h\} = 3 \sum_h a_h^2 \sum_{i \in B_h} \frac{\mathbb{E}_0 N_i}{M_h}.$$

Thus all counts in (79) belong to the same null world.

D.3 Finite-profile constrained sampling

Lemma 16 (Finite-profile constrained sampling)**.** *Let $\mathcal{D}$ be a nonempty compact convex subset of a finite probability simplex. Partition its coordinates into nonempty blocks B_h, with $|B_h| = M_h$ and $q(B_h) \le w_h \le 1$ for every $q \in \mathcal{D}$. Remove identically zero coordinates and empty blocks, so that each remaining $w_h > 0$. For positive a_h, define*

$$\mathcal{R}(q) = -\sum_h a_h \sum_{i \in B_h} \sqrt{q_i}, \qquad S = \sum_h a_h \sqrt{M_h w_h},$$
$$V = \left(\sum_h a_h^{-1/2} M_h^{1/4} w_h^{3/4} \right)^2, \quad \mathsf{H} = \sum_h \sqrt{M_h} w_h.$$

Write $M_{\max} = \max_h M_h$ and $B_\star = 1 + 2\mathsf{H}\sqrt{M_{\max}}$. Since $\sum_h w_h \ge 1$, we have $\mathsf{H} \ge 1$ and $B_\star \le 3\mathsf{H}\sqrt{M_{\max}}$. The fixed distribution $p_i^0 = w_h/(\sqrt{M_h}\mathsf{H})$ for $i \in B_h$ sums to one and uses only known coordinates and capacities.

We adapt decoupled Tsallis sampling (Rouyer & Seldin, 2020, Lemmas 1–2). For a deterministic horizon N, step size $\eta > 0$, and $G_0 = 0$, choose before each query

$$\begin{aligned}
q_n &= \arg\min_{q \in \mathcal{D}} \{G_{n-1} \cdot q + \mathcal{R}(q)/\eta\}, \\
D_n &= \sum_i q_{n,i}^{3/4} / \sqrt{a_{h(i)}}, & p_{n,i}^A &= \frac{q_{n,i}^{3/4}}{\sqrt{a_{h(i)}} D_n}, \\
p_{n,i} &= \tfrac{1}{2}(p_{n,i}^A + p_i^0), & \widehat{g}_n &= Y_n e_{I_n} / p_{n,I_n}, \qquad I_n \sim p_n, \\
G_n &= G_{n-1} + \widehat{g}_n.
\end{aligned} \tag{80}$$

Here $Y_n \in [0,1]$ has conditional mean c_{I_n} given the entire past and the chosen record, for fixed $c \in [0,1]^M$. Each round makes at most one real observation; a known dummy value is returned arithmetically. The exploration p_n need not be feasible, while q_n and their average are feasible.

For every $\alpha > 0$ and $0 < \delta < 1$, the choices derived below give an α-optimal feasible average with failure at most δ at cost (84)*. Two such runs and Lemma 13 give both risk endpoints.*

Proof. Every retained coordinate has a feasible positive witness. Their finite average gives $q^\circ \in \mathcal{D}$ with all coordinates positive. If an exact minimizer q_n had a zero coordinate, then for $q(t) = (1-t)q_n + tq^\circ$ its objective increment would satisfy

$$F_{n-1}(q(t)) - F_{n-1}(q_n) = -\frac{\sqrt{t}}{\eta} \sum_{i : q_{n,i} = 0} a_{h(i)} \sqrt{q_i^\circ} + O(t) < 0$$

for small $t > 0$, a contradiction. Here $F_n(q) = G_n \cdot q + \mathcal{R}(q)/\eta$.

We prove the bound also for predictable positive feasible q_n with certified first-order residual $\xi_n \ge 0$:

$$\max_{z \in \mathcal{D}} \langle \nabla F_{n-1}(q_n), q_n - z \rangle \le \xi_n.$$

Exact minimizers have $\xi_n = 0$. Put $A_n = \min_{\mathcal{D}} F_n$ and choose $z_n \in \arg\min F_n$. The identity

$$F_{n-1}(z) - F_{n-1}(q_n) = \langle \nabla F_{n-1}(q_n), z - q_n \rangle + D_{\mathcal{R}}(z, q_n)/\eta \geq D_{\mathcal{R}}(z, q_n)/\eta - \xi_n$$

uses no gradient at z, which may have zero coordinates. Since $A_{n-1} \leq F_{n-1}(q_n)$,

$$A_n - A_{n-1} \geq F_n(z_n) - F_{n-1}(q_n) \geq \widehat{g}_n \cdot z_n + D_{\mathcal{R}}(z_n, q_n)/\eta - \xi_n,$$

$$\sum_{n=1}^{N} \widehat{g}_n \cdot (q_n - q) \leq A_N - A_0 - G_N \cdot q + \sum_{n=1}^{N} \{\widehat{g}_n \cdot (q_n - z_n) - D_{\mathcal{R}}(z_n, q_n)/\eta + \xi_n\}$$

$$\leq \frac{S}{\eta} + \sum_{n=1}^{N} \{\widehat{g}_n \cdot (q_n - z_n) - D_{\mathcal{R}}(z_n, q_n)/\eta + \xi_n\}.$$

The last line uses $A_N \leq F_N(q)$ and $-S \leq \mathcal{R}(q) \leq 0$, which follow from $\sum_{i \in B_h} \sqrt{q_i} \leq \sqrt{M_h w_h}$. For the sole nonzero coordinate $i \in B_h$ of $\widehat{g}_n \geq 0$, drop the other nonnegative scalar divergences. A coordinate $z_i > q_{n,i}$ contributes nonpositively; for $0 \leq z_i \leq q_{n,i}$, $\mathcal{R}_i''(z_i) \geq a_h/(4q_{n,i}^{3/2})$. With $x = q_{n,i} - z_i \geq 0$, integration of this curvature gives

$$\widehat{g}_{n,i} x - \frac{D_{\mathcal{R}_i}(z_i, q_{n,i})}{\eta} \leq \widehat{g}_{n,i} x - \frac{a_h x^2}{8\eta q_{n,i}^{3/2}} \leq \sup_{x \geq 0} \left\{ \widehat{g}_{n,i} x - \frac{a_h x^2}{8\eta q_{n,i}^{3/2}} \right\} = \frac{2\eta \widehat{g}_{n,i}^2 q_{n,i}^{3/2}}{a_h}.$$

The mixed exploration bounds this term on every path:

$$\frac{Y_n^2 q_{n,I_n}^{3/2}}{a_{h(I_n)} p_{n,I_n}^2} \leq 4 Y_n^2 D_n^2 \leq 4V, \qquad \sum_n \widehat{g}_n \cdot (q_n - q) \leq S/\eta + 8\eta V N + \sum_n \xi_n. \tag{81}$$

Here $p_{n,i} \geq q_{n,i}^{3/4}/(2\sqrt{a_h} D_n)$ and $D_n \leq \sum_h a_h^{-1/2} M_h^{1/4} w_h^{3/4} = \sqrt{V}$.

Fix a true minimizer $q^\star$ and put $d_n = q_n - q^\star$ and $X_n = (c - \widehat{g}_n) \cdot d_n$. Since $\mathbb{E}[\widehat{g}_{n,i} \mid \mathcal{F}_{n-1}] = p_{n,i} c_i / p_{n,i} = c_i$, we have $\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] = 0$. The comparator is fixed by c, so no union over feasible laws is needed. Since both q_n and $q^\star$ are nonnegative with block masses at most w_h,

$$\sum_{i \in B_h} d_{n,i}^2 \leq \sum_{i \in B_h} (q_{n,i}^2 + (q_i^\star)^2) \leq 2 w_h^2.$$

Using $p_n \geq p^0/2$ first bounds the uncentered second moment:

$$\begin{aligned} \mathbb{E}[X_n^2 \mid \mathcal{F}_{n-1}] &= \mathbb{E}[(\widehat{g}_n \cdot d_n)^2 \mid \mathcal{F}_{n-1}] - (c \cdot d_n)^2 \\ &\leq \sum_i \frac{d_{n,i}^2}{p_{n,i}} \leq 2\mathsf{H} \sum_h \frac{\sqrt{M_h}}{w_h} \sum_{i \in B_h} d_{n,i}^2 \leq 4\mathsf{H}^2, \\ |X_n| &\leq 1 + \max_i |d_{n,i}|/p_{n,i} \leq 1 + 2\mathsf{H}\sqrt{M_{\max}} = B_\star. \end{aligned} \tag{82}$$

This includes the sampling variability of the estimated gradient, rather than only its conditional response noise.

Scalar Freedman concentration (Tropp, 2011, Theorem 1.1, p. 263), originating in Freedman (1975, Theorem (1.6)), states that if $X_n \leq B$ and $\sum_n \mathbb{E}[X_n^2 \mid \mathcal{F}_{n-1}] \leq v$ almost surely, then

$$\Pr\left\{ \sum_n X_n \geq t \right\} \leq \exp\left\{ -\frac{t^2}{2(v + Bt/3)} \right\}.$$

Use $v = 4\mathsf{H}^2 N$, $B = B_\star$ and $u = \log(8/\delta)$. For $t = \sqrt{2vu} + 2Bu/3$, $t^2 - 2u(v + Bt/3) = (2Bu/3)\sqrt{2vu} \geq 0$. Thus the failure is at most $e^{-u} \leq \delta$. Finally, $Nc \cdot (\bar{q} - q^\star) = \sum_n \widehat{g}_n \cdot d_n + \sum_n X_n$. Substitution in (81) gives

$$c \cdot \bar{q} - \min_{\mathcal{D}} c \cdot q \leq \frac{S}{\eta N} + 8\eta V + \bar{\xi} + \sqrt{\frac{8\mathsf{H}^2 u}{N}} + \frac{2B_\star u}{3N}, \qquad \bar{q} = N^{-1} \sum_n q_n. \tag{83}$$

Here $\bar{\xi} = N^{-1}\sum_n \xi_n$ is controlled on every path. For a finite rational moment polytope, positive rational mixtures of vertices and a linear optimization oracle, with certified gradient approximations, verify any positive residual tolerance in finite computation. No polynomial arithmetic bound is asserted.

Take $\eta = \sqrt{S/(8VN)}$. For objective tolerance α, choose $\bar{\xi} \le \alpha/4$ and

$$N = \left\lceil \max\left\{ \frac{512SV}{\alpha^2},\ \frac{128\mathsf{H}^2 u}{\alpha^2},\ \frac{8B_\star u}{3\alpha} \right\} \right\rceil.$$

The four contributions in (83) satisfy

$$\frac{S}{\eta N} + 8\eta V = 2\sqrt{\frac{8SV}{N}} \le \alpha/4, \quad \bar{\xi} \le \alpha/4, \quad \sqrt{\frac{8\mathsf{H}^2 u}{N}} \le \alpha/4, \quad \frac{2B_\star u}{3N} \le \alpha/4.$$

For $a_h = (w_h/M_h)^{1/6}$, $a_h\sqrt{M_h w_h} = a_h^{-1/2} M_h^{1/4} w_h^{3/4} = (M_h w_h^2)^{1/3}$. Hence $S = \sqrt{V} = Z := \sum_h (M_h w_h^2)^{1/3}$. Consequently the deterministic training budget is

$$O\Big(\alpha^{-2}[Z^3 + \mathsf{H}^2 u] + \alpha^{-1}\mathsf{H}\sqrt{M_{\max}}u + 1\Big). \tag{84}$$

There is no restriction on how small δ may be.

For arbitrary finite profiles, the earlier generic bound also remains available. Taking expectations before applying Freedman gives expected gap at most $S/(\eta N) + 8\eta V + \mathbb{E}\bar{\xi}$. For scalar-value output, constant-success optimization, fresh evaluation, and median repetition give $O(\alpha^{-2} Z^3 \log(2/\delta))$ observations. This fallback can be preferable when the range term in (84) is large. Its evaluation cost is absorbed by $Z^3 \ge 1$.

□

D.4 A COMMON EVIDENCE PROFILE AND CHECKABLE LOWER WITNESSES

Let $\mathcal{D} = A^\top \mathcal{Q} \subseteq \Delta(J)$ be a nonempty complete physical domain; $A = I$ gives the record experiment. Each original coordinate has an unrestricted unknown Bernoulli mean v_j. The lower bounds below use these actual responses and the single baseline $v^0 = \mathbf{1}/2$. Artificial known dummies used for upper truncation are excluded from them.

Partition queried coordinates into blocks of counts M_h and valid mass bounds $w_h \in (0, 1]$. Remove identically zero coordinates. Write

$$Z = \sum_h (M_h w_h^2)^{1/3}, \qquad \mathsf{H} = \sum_h \sqrt{M_h} w_h, \qquad u = \log(8/\delta).$$

A lower-bound witness must belong to the same complete moment domain. For disjoint unknown-coordinate sets J_h of sizes m_h, a *balanced positive witness table* consists of nonnegative weights ω_s, tuples $(j_{s,h})_h$, and genuine laws $P^s \in \mathcal{P}(m)$ satisfying

$$\sum_s \omega_s = 1, \qquad \sum_{s: j_{s,h} = j} \omega_s = m_h^{-1}, \quad (A^\top T_\# P^s)_{j_{s,h}} \ge p_h > 0. \tag{85}$$

This condition is checked on original assignment probabilities and moments; it is stronger than separately bounding record capacities.

Lemma 17 (Given-domain evidence certificates)**.** *For $0 < \epsilon \le 1/16$, $0 < \delta \le 1/16$, the feasible-law sampler with certified optimization residuals returns a faithful interval in*

$$O\Big(\epsilon^{-2}(Z^3 + \mathsf{H}^2 u) + \epsilon^{-1}\mathsf{H}\sqrt{M_{\max}}u + 1\Big) \tag{86}$$

queries. The blocks may describe the full domain or its exact projection with a known-zero tail dummy of capacity at most $\epsilon/4$. Let $c_j = \max_{z \in \mathcal{D}} z_j$ in the original domain. Every uniformly valid scalar estimator of R_+, including adaptive stopping, obeys

$$\mathbb{E}_0 N \ge c\epsilon^{-2} u \max\left\{ 1, \sum_{j: c_j \ge 16\epsilon} c_j^2 \right\}. \tag{87}$$

If a table (85) *has* $\mathsf{H}_p = \sum_h \sqrt{m_h} p_h \geq 16\epsilon\sqrt{\max_h m_h}$*, then*

$$\mathbb{E}_0 N \geq c\epsilon^{-2} u \mathsf{H}_p^2. \tag{88}$$

All lower comparisons use the same flat-mean baseline and actual unknown coordinates, excluding artificial dummies. The upper is deterministic in observation count; its finite numerical implementation is proved in Appendix F.

Proof. Run Lemma 16 with costs v and $1-v$, tolerance $\alpha = \epsilon/4$, and failure allowance $\delta/4$ per run. Freeze the two feasible averages and evaluate each with fresh responses, using $\beta = \epsilon/4$ and failure allowance $\delta/4$. For a known-zero tail dummy, $\tau_- = 0$ and $\tau_+ \leq \epsilon/4$. The explicit endpoint bounds in Lemma 13 give

$$\Pr(\text{any training or evaluation failure}) \leq 4(\delta/4) = \delta, \qquad R_- - L,\ U - R_+ \leq \alpha + 2\beta + \tau_+ \leq \epsilon.$$

Each evaluation uses $\lceil 8\epsilon^{-2}\log(32/\delta)\rceil$ observations. Since $\mathsf{H} \geq 1$ and $\log(32/\delta) \leq 2u$, this cost is absorbed in (86). All samples use original coordinates, including when $A \neq I$.

For every nonnegative $a \in [0,1/4]^J$ with $\sup_{z\in\mathcal{D}} z^\top a \geq 4\epsilon$, the alternative $v^0 + a$ has its upper endpoint at least 4ϵ above the baseline. The successful scalar-output intervals are therefore disjoint. Write $\kappa_\delta = \mathrm{kl}(1-\delta,\delta) \geq u/8$. Here $a_j \leq 1/4$ and $\delta \leq 1/16$, as required by the common Bernoulli bound (78). Its stopped-test application is

$$\kappa_\delta \overset{(76)}{\leq} \sum_j \mathrm{kl}(1/2, 1/2 + a_j)\mathbb{E}_0 N_j \overset{(78)}{\leq} 3\sum_j a_j^2 \mathbb{E}_0 N_j. \tag{89}$$

For the constant bump $a = 4\epsilon\mathbf{1}$,

$$\sup_{z\in\mathcal{D}} z^\top a = 4\epsilon, \qquad \kappa_\delta \leq 48\epsilon^2 \mathbb{E}_0 N.$$

For $c_j \geq 16\epsilon$, use $a = (4\epsilon/c_j)e_j$ and a maximizing law in the compact domain:

$$\sup_{z\in\mathcal{D}} z^\top a = 4\epsilon, \qquad \kappa_\delta \leq \frac{48\epsilon^2}{c_j^2}\mathbb{E}_0 N_j \quad\Longrightarrow\quad \mathbb{E}_0 N_j \geq \frac{\kappa_\delta c_j^2}{48\epsilon^2}.$$

Summation is valid because every comparison uses the same baseline.

For table row s, raise $j_{s,h}$ by a_h, where $0 \leq a_h \leq 1/4$ and $\sum_h p_h a_h \geq 4\epsilon$. Its actual compatible law supplies the required endpoint increase. Average (89) with weights ω_s:

$$\begin{aligned}
\kappa_\delta &\leq 3\sum_s \omega_s \sum_h a_h^2 \mathbb{E}_0 N_{j_{s,h}} \\
&= 3\sum_h a_h^2 \sum_{j\in J_h} \mathbb{E}_0 N_j \underbrace{\sum_{s: j_{s,h}=j} \omega_s}_{1/m_h} \\
&= 3\sum_h \frac{a_h^2}{m_h}\mathbb{E}_0 N(J_h).
\end{aligned} \tag{90}$$

Choose $a_h = 4\epsilon\sqrt{m_h}/\mathsf{H}_p$. The hypothesis gives

$$a_h \leq \tfrac{1}{4}, \qquad \sum_h p_h a_h = 4\epsilon, \qquad \kappa_\delta \leq \frac{48\epsilon^2}{\mathsf{H}_p^2}\mathbb{E}_0 N \quad\Longrightarrow\quad \mathbb{E}_0 N \geq \frac{u\mathsf{H}_p^2}{384\epsilon^2}.$$

Lemma 15 includes the finite-prefix limit for almost-sure stopping. Infinite expected cost already satisfies the bounds. □

The stronger every-tuple property used by the multiscale construction implies a small witness table. Order each J_h and put $j_h(t) = 1 + \lfloor m_h t \rfloor$ for $t \in [0,1)$. Partition at every j/m_h; on each interval I_s the tuple is constant. Set $\omega_s = |I_s|$. Then

$$\sum_{s: j_{s,h}=j} \omega_s = |[(j-1)/m_h, j/m_h)| = m_h^{-1}, \qquad \#\{s\} \leq 1 + \sum_h (m_h - 1).$$

Choose a positive law for each tuple. For a finite rational moment model, a feasible row has a basic witness with at most the number of independent moment equations (including normalization) plus the number of selected-mass constraints many atoms. Such a supplied witness is checked by direct moment and routing evaluation; finding it remains a separate problem.

For bank depths $\nu(j)$, suppose

$$\#\{j : \nu(j) \le k\} \le C_0 k^a, \qquad \sup_{z \in \mathcal{D}} z\{\nu > k\} \le C_0 k^{-d} \quad (k < K),$$

with zero tail at K. Put $k_* = \min(K, \epsilon^{-1/d})$ and $\Lambda_* = 1 + \log k_*$. Use a cutoff $\bar{k} = \min\{K, \lceil (4C_0/\epsilon)^{1/d} \rceil\}$. For geometric shells ending at $k_h \le \bar{k}$, the count and tail bounds give

$$M_h \le C k_h^a, \quad w_h \le C k_h^{-d}, \quad M_{\max} \le C \max\{1, k_*^a\},$$
$$Z \le C \sum_h k_h^{(a-2d)/3}, \qquad \mathsf{H} \le C \sum_h k_h^{(a-2d)/2}.$$

Each geometric sum is bounded, proportional to Λ_*, or dominated by its last term according to the sign of its exponent. Thus

$$Z^3 + u\mathsf{H}^2 \le C \begin{cases} u, & a < 2d, \\ \Lambda_*^3 + u\Lambda_*^2, & a = 2d, \\ u k_*^{a-2d}, & a > 2d. \end{cases}$$

Since $\epsilon k_*^d \le 1$, the remaining range term satisfies

$$\epsilon^{-1}\mathsf{H}\sqrt{M_{\max}}u \le C\epsilon^{-2}u \begin{cases} \epsilon \max\{1, k_*^{a/2}\} \le 1, & a < 2d, \\ \Lambda_* \epsilon k_*^d \le \Lambda_*^2, & a = 2d, \\ k_*^{a-2d} \epsilon k_*^d \le k_*^{a-2d}, & a > 2d. \end{cases}$$

Substituting these three bounds into (86) yields

$$\mathcal{N}_{\mathcal{D}} \le C \begin{cases} \epsilon^{-2}u, & a < 2d, \\ \epsilon^{-2}(\Lambda_*^3 + \Lambda_*^2 u), & a = 2d, \\ \epsilon^{-2} k_*^{a-2d} u, & a > 2d. \end{cases} \tag{91}$$

The flat lower bound matches the first case on every such domain. In the last case, matching holds if, at dyadic scales up to a constant multiple of k_*, at least $c_1 k^a$ actual coordinates have certified $c_j \ge c_2 k^{-d}$. Choose a scale comparable to $\min(K, (c_2/(16\epsilon))^{1/d})$ and apply (87). At $a = 2d$, require a balanced positive table on consecutive dyadic shells from a fixed initial scale through $c_3 k_*$, with $m_h \asymp k_h^{2d}$ and $p_h \gtrsim k_h^{-d}$. Choose fixed $c_3 > 0$ small enough that $16\epsilon\sqrt{m_{\max}} \le \mathsf{H}_p$. These shells retain $\Theta(\Lambda_*)$ scales whenever that quantity is unbounded; bounded cases use the flat lower bound. Equation (88) then gives $c\epsilon^{-2}\Lambda_*^2 u$. Thus it matches for $u \gtrsim \Lambda_*$; the finer low-confidence scalar rate needs additional structure, as in the original stopped-mixture proof.

D.5 Stopped-mixture information

Use the common stopped experiment of Lemma 15, with all initial information and the private seed independent of the latent parameters and included in the filtration, an absorbing state after stopping, and $\mathrm{Hel}^2(P, Q) = 1 - \int \sqrt{dP\, dQ}$. The following cutoff calculation is the adaptation of Lee & Valiant (2021, Sec. 5.2, Proposition 5.1) used below. The conditional information decomposition also follows their Section 5.1, Lemma 5.2, with the charged response kernel specified here.

Lemma 18 (Stopped mixture, local coordinates, and signed outputs)**.** *Let $J \ge 1$, $0 < \rho \le 1/2$, $0 < c \le 1/2$, $0 < a \le 1/3$, $\tau_j = 2^{-j}$, $\tau_* = 2^{-J}$, and $\eta_j = c\tau_*/\tau_j$. Both priors give the latent σ a null mass $1 - \rho$, independent uniform nonzero signs, and absolute-amplitude probabilities*

	$\lvert\sigma\rvert = \tau_j$	$\lvert\sigma\rvert = 3\tau_j/2$
ν_+	$\rho(1-\eta_j)/(2J)$	$\rho(1+\eta_j)/(2J)$
ν_-	$\rho(1+\eta_j)/(2J)$	$\rho(1-\eta_j)/(2J).$

(92)

Queries return fresh $\mathrm{Bernoulli}(1/2+a\sigma)$ *observations. Every common rule stopping almost surely at each supported parameter and at zero satisfies*

$$\mathrm{Hel}^2(P_+,P_-)\le C\frac{\rho a^2\tau_*^2}{J}\mathbb{E}_0\tau_{\mathrm{stop}}\le 2C\frac{\rho a^2\tau_*^2}{J}\mathbb{E}_{\bar\nu}\tau_{\mathrm{stop}}, \tag{93}$$

where $\bar\nu$ *is either prior or their midpoint.*

For finitely many independent latent coordinates with local response kernels, let $P_{[i]}$ *switch only coordinate* i *from its minus to plus prior. If each prior ratio lies in* $[1/3,3]$*, then*

$$\mathrm{KL}(P_-\|P_+)=\sum_i \mathrm{KL}(P_-\|P_{[i]})\le 16\sum_i \mathrm{Hel}^2(P_-,P_{[i]}). \tag{94}$$

Here a response kernel may depend on the history and the requested coordinate's latent, but on no other latent. Actions remain arbitrary history-adapted choices. For a real square-integrable output Z *and* $z_0\in\mathbb{R}$*,*

$$|\mathbb{E}_P Z-\mathbb{E}_Q Z|^2\le 4\,\mathrm{Hel}^2(P,Q)\{\mathbb{E}_P(Z-z_0)^2+\mathbb{E}_Q(Z-z_0)^2\}. \tag{95}$$

If $|\mathbb{E}_\theta Z-f(\theta)|\le\zeta$ *and* $\mathrm{Var}_\theta Z\le v$*, then*

$$\mathbb{E}_\nu(Z-z_0)^2\le v+2\mathbb{E}_\nu(f(\theta)-z_0)^2+2\zeta^2. \tag{96}$$

All stopping costs are interpreted in the extended sense.

Proof. Let P_d mix $1/2\pm d$ coins with equal weights, $0<d\le 1/4$. Under the null let $X_s=2Y_s-1$ and $S_t=\sum_{s\le t}X_s$. For the latent sign $z\in\{-1,1\}$, the likelihood and its equal mixture are

$$\begin{aligned}
L_t^z&=\prod_{s\le t}(1+2zdX_s)=(1-4d^2)^{t/2}e^{z\,\mathrm{arctanh}(2d)S_t},\\
L_t&=(L_t^+ + L_t^-)/2=(1-4d^2)^{t/2}\cosh(\mathrm{arctanh}(2d)S_t),\\
b_t&=d\,\mathbb{E}[z\mid X_{1:t}]=d\frac{L_t^+-L_t^-}{L_t^++L_t^-}=d\tanh(\mathrm{arctanh}(2d)S_t).
\end{aligned}$$

Since $\mathrm{kl}(1/2,1/2+b)\le 8b^2/3$ for $|b|\le 1/4$ and $\mathrm{arctanh}(2d)\le 8d/3$, every bounded stopping time U obeys

$$\mathrm{KL}(P_0^U\|P_d^U)=\mathbb{E}_0\sum_{t<U}\mathrm{kl}(1/2,1/2+b_t)\le\frac{8}{3}d^2\left(\frac{8d}{3}\right)^2\mathbb{E}_0\sum_{t<U}S_t^2.$$

The null increment satisfies $\mathbb{E}_0[S_{t+1}^4-S_t^4\mid\mathcal{F}_t]=6S_t^2+1$. Put $A=\mathbb{E}_0S_U^4$ and $B=\mathbb{E}_0U^2$. Summing this identity with the predictable indicators $\mathbf{1}\{t<U\}$, then applying Cauchy–Schwarz and Doob to the stopped martingale, gives

$$\begin{aligned}
A&=6\mathbb{E}_0\sum_{t<U}S_t^2+\mathbb{E}_0U\\
&\le 6\sqrt{B}\big(\mathbb{E}_0\max_{t\le U}|S_t|^4\big)^{1/2}+B\le\frac{32}{3}\sqrt{AB}+B\le\frac{1}{2}A+\frac{521}{9}B.
\end{aligned}$$

Here $\mathbb{E}_0U\le B$ because U is integer-valued, and Doob gives $\mathbb{E}_0\max_{t\le U}|S_t|^4\le(4/3)^4A$. The last step is $(32/3)\sqrt{AB}\le A/2+(512/9)B$. For $B=0$, $U=A=0$; otherwise $A\le 1042B/9<128B$. Since $\mathrm{Hel}^2\le\mathrm{KL}/2$ and $\mathbb{E}_0\sum_{t<U}S_t^2=(A-\mathbb{E}_0U)/6\le A/6$, we obtain

$$\mathrm{Hel}^2(P_0^U,P_d^U)\le 480d^4\mathbb{E}_0U^2. \tag{97}$$

Append the latent scale only for the comparison experiment. Its masses are identical under both priors. At scale j, put $d_j=a\tau_j$, $m_j=\lceil d_j^{-2}\rceil$, $U_j=\tau_{\mathrm{stop}}\wedge m_j$. Retain the prefix through U_j and reveal the true magnitude and sign only on $\{\tau_{\mathrm{stop}}\ge m_j\}$. The original transcript is a Markov postprocessing: after advice, independent responses at the revealed parameter simulate its

remainder. Let A_j, B_j be the advice laws at magnitudes $d_j, 3d_j/2$, and $\widetilde{A}_j, \widetilde{B}_j$ their prefix laws. The removed affinity on the advice event is bounded by its mean probability, whence

$$\mathrm{Hel}^2(A_j, B_j) \le \mathrm{Hel}^2(\widetilde{A}_j, \widetilde{B}_j) + \tfrac{1}{2}\{P_{d_j}(\tau_{\mathrm{stop}} \ge m_j) + P_{3d_j/2}(\tau_{\mathrm{stop}} \ge m_j)\}.$$

For a prefix-measurable event E, the L_2 triangle inequality yields

$$\sqrt{P_d(E)} \le \sqrt{P_0(E)} + \sqrt{2\,\mathrm{Hel}^2(P_d^{U_j}, P_0^{U_j})}, \quad P_d(E) \le 2P_0(E) + 4\,\mathrm{Hel}^2(P_d^{U_j}, P_0^{U_j}).$$

Apply the event bound with $E = \{\tau_{\mathrm{stop}} \ge m_j\}$. The triangle inequality gives $\mathrm{Hel}^2(\widetilde{A}_j, \widetilde{B}_j) \le 2\,\mathrm{Hel}^2(\widetilde{A}_j, P_0^{U_j}) + 2\,\mathrm{Hel}^2(P_0^{U_j}, \widetilde{B}_j)$. Equation (97) applies because $3d_j/2 \le 3a/4 \le 1/4$, and yields

$$\mathrm{Hel}^2(A_j, B_j) \le C\{d_j^4 \mathbb{E}_0(\tau_{\mathrm{stop}} \wedge m_j)^2 + P_0(\tau_{\mathrm{stop}} \ge m_j)\}.$$

For $M_j = (A_j + B_j)/2$, write $dA_j/dM_j = 1 + z$, $dB_j/dM_j = 1 - z$. Opposite tilts of size η_j have densities $1 \pm \eta_j z$, so

$$\begin{aligned}\mathrm{Hel}^2((1+\eta_j z)M_j, (1-\eta_j z)M_j) &= \mathbb{E}_{M_j}[1 - \sqrt{1-\eta_j^2 z^2}] \\ &\le \eta_j^2 \mathbb{E}_{M_j} z^2 \le 2\eta_j^2 \mathbb{E}_{M_j}[1 - \sqrt{1-z^2}] \\ &= 2\eta_j^2\,\mathrm{Hel}^2(A_j, B_j).\end{aligned}$$

Both inequalities use $x/2 \le 1 - \sqrt{1-x} \le x$ for $0 \le x \le 1$. Data processing gives the original conditional transcript bound

$$\mathrm{Hel}^2(P_{j,+}, P_{j,-}) \le C\eta_j^2\{d_j^4 \mathbb{E}_0(\tau_{\mathrm{stop}} \wedge m_j)^2 + P_0(\tau_{\mathrm{stop}} \ge m_j)\}. \tag{98}$$

Write $x_j = d_j^{-2}$; then $x_{j+1} = 4x_j$, $x_j \ge 16$, $x_j \le m_j < x_j + 1$, $m_{j+1} \ge 4m_j - 3 \ge 2m_j$, and $d_j^2 m_j \le 17/16$. For integer $t \ge 2$, with $j_0 = \min\{j : m_j > t\}$ when this set is nonempty,

$$\sum_{m_j \le t} m_j \le 2t, \quad \sum_{m_j \le t} d_j^2 m_j^2 \le 17t/8, \quad t^2 \sum_{j \ge j_0} d_j^2 \le \frac{4t^2}{3(t-1)} \le 8t/3.$$

Empty sums vanish; $t = 0, 1$ follow directly. Thus

$$\sum_j d_j^2 \min(t, m_j)^2 \le 5t, \qquad \sum_j m_j \mathbf{1}\{t \ge m_j\} \le 2t. \tag{99}$$

Since $\eta_j^2 d_j^4 = c^2 a^2 \tau_*^2 d_j^2$ and $\eta_j^2 = c^2 a^2 \tau_*^2 x_j \le c^2 a^2 \tau_*^2 m_j$,

$$\begin{aligned}\mathrm{Hel}^2(P_+, P_-) &\le \frac{\rho}{J} \sum_j \mathrm{Hel}^2(P_{j,+}, P_{j,-}) \\ &\le \frac{C\rho c^2 a^2 \tau_*^2}{J} \mathbb{E}_0 \left[\sum_j d_j^2 \min(\tau_{\mathrm{stop}}, m_j)^2 + \sum_j m_j \mathbf{1}\{\tau_{\mathrm{stop}} \ge m_j\}\right] \\ &\le C' \frac{\rho a^2 \tau_*^2}{J} \mathbb{E}_0 \tau_{\mathrm{stop}}.\end{aligned}$$

The first inequality reveals a common scale with masses ρ/J and a common null branch of mass $1 - \rho$. If the expected stopping time is infinite, the bound holds in the extended sense. Otherwise apply the calculation first to $\tau_{\mathrm{stop}} \wedge m$ and let $m \to \infty$. Almost-sure termination gives convergence of the stopped transcript laws; data processing on their increasing prefix sigma-fields gives convergence of Hellinger distance, and the cost sum converges monotonically. The common null mass gives $\mathbb{E}_{\bar{\nu}} \tau_{\mathrm{stop}} \ge \mathbb{E}_0 \tau_{\mathrm{stop}}/2$.

At a finite transcript leaf z, the action sequence is fixed and independent priors with local kernels give

$$L_{\nu_i,i}(z) = \int \prod_{n: A_n = i} K_i(Y_n \mid \mathfrak{h}_{n-1}, \sigma_i)\, \nu_i(d\sigma_i), \qquad p_\nu(z) = k(z) \prod_i L_{\nu_i,i}(z).$$

The factor $k(z)$ contains the common initial, action and stopping kernels. Integration factors because the latents are independent; the realized histories inside each kernel are held fixed. Changing only prior i therefore gives

$$\frac{p_{[i]}(z)}{p_-(z)} = \frac{L_{\nu_i^+,i}(z)}{L_{\nu_i^-,i}(z)}, \qquad \frac{p_+(z)}{p_-(z)} = \prod_i \frac{p_{[i]}(z)}{p_-(z)}.$$

Taking $-\log$ and expectation under the same P_- proves the KL sum. The prior-ratio bound persists after integration, so $r = p_{[i]}/p_- \in [1/3, 3]$ and $\mathbb{E}_- r = 1$. Hence

$$r - 1 - \log r \le \frac{(r-1)^2}{2\min(1,r)} \le 4(\sqrt{r}-1)^2, \qquad \mathrm{KL}(P_- \| P_{[i]}) \le 8\,\mathrm{Hel}^2(P_-, P_{[i]}).$$

Here $\mathrm{KL}(P_- \| P_{[i]}) = \mathbb{E}_-[r - 1 - \log r]$ and $\mathbb{E}_-(\sqrt{r}-1)^2 = 2\,\mathrm{Hel}^2(P_-, P_{[i]})$. The weaker constant 16 suffices. At cap m, the identity holds for every coordinate. Increasing-prefix entropy convergence and the finite number of coordinates give $\lim_m \sum_i \mathrm{KL}(P_-^m \| P_{[i]}^m) = \sum_i \mathrm{KL}(P_- \| P_{[i]})$. The cost limit is the monotone limit from Lemma 15.

For densities p, q,

$$\mathbb{E}_P Z - \mathbb{E}_Q Z = \int (Z - z_0)(\sqrt{p} - \sqrt{q})(\sqrt{p} + \sqrt{q}).$$

Cauchy–Schwarz gives

$$\begin{aligned} |\mathbb{E}_P Z - \mathbb{E}_Q Z|^2 &\le \underbrace{\int (\sqrt{p} - \sqrt{q})^2}_{2\,\mathrm{Hel}^2(P,Q)} \int (Z - z_0)^2 (\sqrt{p} + \sqrt{q})^2 \\ &\le 4\,\mathrm{Hel}^2(P,Q)\{\mathbb{E}_P(Z - z_0)^2 + \mathbb{E}_Q(Z - z_0)^2\}. \end{aligned}$$

Finally, the law of total variance and the uniform bias bound give

$$\mathbb{E}_\nu(Z - z_0)^2 = \mathbb{E}_\nu \operatorname{Var}_\theta Z + \mathbb{E}_\nu(\mathbb{E}_\theta Z - z_0)^2 \le v + 2\mathbb{E}_\nu(f - z_0)^2 + 2\zeta^2.$$

□

D.6 Weighted block-maximum estimation

A bounded response satisfying (69) can be replaced by $B = \mathbf{1}\{U \le Y\}$ using a fresh private uniform variable:

$$P(B = 1 \mid \mathcal{F}, I = i) = \mathbb{E}[P(U \le Y \mid \mathcal{F}, I = i, Y) \mid \mathcal{F}, I = i] = \mathbb{E}[Y \mid \mathcal{F}, I = i] = c_i.$$

This is a Bernoulli response at the same cost. The subroutines below use disjoint fresh observations and randomness.

Lemma 19 (Weighted maxima: mean-square and exponential-moment versions)**.** *For $M \ge 1$ records of means $c_i \in [0,1]$, weight $0 < p \le 1$, target $f = p \max_i c_i$, and $0 < \alpha \le p$, there are estimators $X_2, X_{\exp}$ with deterministic budgets satisfying*

$$f - \alpha \le \mathbb{E}X_2 \le f, \quad \mathbb{E}(X_2 - f)^2 \le 5\alpha^2/4, \quad \mathrm{cost}(X_2) \le 294 M p^2 \alpha^{-2}, \tag{100}$$

and, for absolute $C, C_0 > 0$,

$$\begin{aligned} &f - \alpha \le \mathbb{E}X_{\exp} \le f, \qquad \mathbb{E}(X_{\exp} - f)^2 \le C\alpha^2, \qquad \mathrm{cost}(X_{\exp}) \le C M p^2 \alpha^{-2}, \\ &\mathbb{E}e^{t(X_{\exp} - \mathbb{E}X_{\exp})} \le e^{C_0 t^2 \alpha^2} \quad (t \in \mathbb{R}). \end{aligned} \tag{101}$$

Known deterministic arms are allowed and require no original observation.

Proof. Use a knockout bracket with $2^{\lceil \log_2 M \rceil} < 2M$ slots for $M > 1$; empty slots are byes and cost nothing. Put $w_k = (3/4)^{k-1}/4$. At round k sample each paired survivor n_k times; let D_k be the drop in the largest surviving true mean. Given the past, its best survivor has a fixed opponent or

a bye. In a played match write these arms as i, j and put $\Delta = c_i - c_j \geq 0$. If i survives, $D_k = 0$; otherwise $D_k \leq \Delta$ and $\widehat{c}_i \leq \widehat{c}_j$. Hence

$$D_k \leq |\widehat{c}_i - c_i| + |\widehat{c}_j - c_j|, \qquad \|D_k\|_2 \leq \frac{1}{2\sqrt{n_k}} + \frac{1}{2\sqrt{n_k}},$$
$$\Pr(D_k > t \mid \mathcal{F}_{k-1}) \leq \mathbf{1}\{\Delta > t\} \Pr(\widehat{c}_i - \widehat{c}_j - \Delta \leq -\Delta \mid \mathcal{F}_{k-1}) \leq e^{-n_k t^2}.$$

The last step is Hoeffding's inequality for $2n_k$ independent bounded summands with signed weights $1/n_k$. For $M = 1$, omit selection. In all cases evaluate the selected record I with $n_e = \lceil \eta^{-2} \rceil$ independent samples and return $X = p\overline{Y}_I$.

For X_2, take $\eta = \alpha/p$ and $n_k = \lceil (\eta w_k)^{-2} \rceil$. If $R = \max_i c_i - c_I$, then

$$\|R\|_2 \leq \sum_k \|D_k\|_2 \leq \eta \sum_k w_k \leq \eta, \qquad \mathbb{E}X = f - p\mathbb{E}R,$$
$$\mathbb{E}[\overline{Y}_I - c_I \mid \mathcal{F}_{\text{sel}}] = 0, \qquad \mathbb{E}[(\overline{Y}_I - c_I)^2 \mid \mathcal{F}_{\text{sel}}] \leq (4n_e)^{-1},$$
$$\mathbb{E}(X - f)^2 = p^2\mathbb{E}R^2 + p^2\mathbb{E}(\overline{Y}_I - c_I)^2 \leq \alpha^2 + \alpha^2/4.$$

There are at most $2M/2^{k-1}$ paired survivors. Hence

$$\begin{aligned}\text{cost}_{\text{sel}} &\leq \sum_{k \geq 1} \frac{2M}{2^{k-1}} [1 + 16\eta^{-2}(4/3)^{2(k-1)}] \\ &= 4M + 32M\eta^{-2} \sum_{k \geq 0} (8/9)^k \leq 292M\eta^{-2}, \qquad n_e \leq 2\eta^{-2}.\end{aligned}$$

This proves (100); it is the paid mean-square adaptation of PAC elimination (Even-Dar et al., 2006, Section 3.2).

For $X_{\exp}$, take $\eta = \alpha/(3p)$ and $n_k = \lceil (1+k)(\eta w_k)^{-2} \rceil$. For $x \geq 1$,

$$\begin{aligned}\Pr(R \geq \eta x) &\leq \Pr\{\exists k : D_k > \eta w_k x\} \leq \sum_{k \geq 1} e^{-n_k \eta^2 w_k^2 x^2} \\ &\leq \sum_{k \geq 1} e^{-(1+k)x^2} = \frac{e^{-2x^2}}{1 - e^{-x^2}} \leq 2e^{-2x^2}.\end{aligned} \tag{102}$$

The finite sum of executed w_k is less than one, so the event inclusion also holds at equality. The cost calculation now uses $\sum_{k \geq 1}(1+k)(8/9)^{k-1} = 90$, so cost $\leq CM\eta^{-2}$. Tail integration and conditional evaluation give

$$\mathbb{E}R = \int_0^\infty \Pr(R > t)\, dt \leq \eta \left[1 + \int_1^\infty 2e^{-2x^2}\, dx\right] \leq 2\eta,$$
$$\begin{aligned}\Pr(|X - f| > \alpha x) &\leq \Pr(R > \eta x) + \Pr(|\overline{Y}_I - c_I| > 2\eta x) \\ &\leq 2e^{-2x^2} + 2e^{-8n_e\eta^2 x^2} \leq 4e^{-2x^2} \quad (x \geq 1).\end{aligned}$$

For all $x \geq 0$ the last probability is at most $e^2 e^{-2x^2}$. In particular, $f - \alpha \leq \mathbb{E}X \leq f$ and $\mathbb{E}(X - f)^2 \leq C\alpha^2$. For an independent X', put $D = X - X'$. This variable is symmetric and

$$\Pr(|D| > \alpha x) \leq 2e^2 e^{-x^2/2},$$
$$\mathbb{E}D^{2m} = \int_0^\infty 2mt^{2m-1} \Pr(|D| > t)\, dt \leq 2e^2 m!(2\alpha^2)^m,$$
$$\mathbb{E}e^{tD} = 1 + \sum_{m \geq 1} \frac{t^{2m}\mathbb{E}D^{2m}}{(2m)!} \leq 1 + 2e^2 \sum_{m \geq 1} \frac{(2t^2\alpha^2)^m}{m!} \leq e^{4e^2 t^2 \alpha^2},$$
$$\mathbb{E}e^{t(X - \mathbb{E}X)} \leq \mathbb{E}e^{t(X - X')} \qquad \text{(conditional Jensen)}.$$

Here $(2m)! \geq (m!)^2$ and $1 + A(e^z - 1) \leq e^{Az}$ for $A \geq 1, z \geq 0$. This proves (101). No signed statistic is used as a Bernoulli input: each selected arm is a real or known deterministic arm. □

D.7 EXACT RECORD-VALUE ACQUISITION

Proposition 20 (Exact record-value acquisition)**.** *Fix $b \geq 1$, $r \geq 2$, $0 \leq \delta < 1/2$, and $d = \lfloor r/2 \rfloor$. For $K \geq \max(r, 2b)$ and sufficiently small ϵ, the minimax expected exact-record query count for an (ϵ, δ)-faithful interval, over legal root geometries and unrestricted $\theta \in [0,1]^{\mathcal{T}}$, is*

$$N^*_{b,r}(K, \epsilon, \delta) = \Theta_{b,r,\delta}\big(\min\{K, \epsilon^{-1/d}\}^{\,b}\big). \tag{103}$$

The upper bound is deterministic and allows heterogeneous keys; the lower bound allows adaptive randomized queries. With only fair singleton information and $\epsilon < 1/2$, the scale is $\Theta_{b,\delta}(K^b)$.

Keep the prescribed moment vector fixed and write $\mathcal{Q} = \mathcal{Q}(m_r)$ and $R_\pm(\theta) = R_\pm(m_r, \theta)$. Unknown record values may vary freely in $[0,1]$; structural restrictions are treated in Theorem 3.

Read bits in decreasing upper-key order. Once b highs have appeared, their keys dominate every unseen upper endpoint, so the record is known. Stop at this event or after k bits. The stopping tree has

$$\sum_{j=b}^{k} \binom{j-1}{b-1} + \sum_{h=0}^{b-1} \binom{k}{h} = \sum_{h=0}^{b} \binom{k}{h} \tag{104}$$

leaves, with out-of-range binomial coefficients zero. Every record with $d_t \leq k$ is fixed on at least one leaf, by (21); distinct records require distinct leaves. Equation (104) therefore bounds their number. These records can be enumerated from the known keys and prefix patterns, without knowing their suffix values. Heterogeneous low keys may reveal a record earlier, which only improves the bound.

Let H_k count highs in the leading k bits. The same argument gives $\{d_T > k\} \subseteq \{H_k < b\}$. Put $d = \lfloor r/2 \rfloor$. Set $W_i = 1 - 2X_i$ and $S_k = \sum_{i \leq k} W_i = k - 2H_k$. The even-moment argument for limited independence (Schmidt et al., 1995; Benjamini et al., 2012) takes the following form:

$$\begin{aligned}
\mathbb{E}_P S_k^{2d} &= \sum_{i_1,\ldots,i_{2d} \leq k} \mathbb{E}_P \prod_{a=1}^{2d} W_{i_a} \\
&= \#\{(i_1, \ldots, i_{2d}) : \text{each index occurs an even number of times}\} \\
&\leq \underbrace{(2d-1)!!}_{\text{pairings of } [2d]}\; k^d.
\end{aligned}$$

The middle equality uses $2d$-wise fairness. Each surviving tuple admits a pairing of equal indices; counting all pairings can only overcount. For $k \geq 4b$, Markov's inequality now gives

$$\begin{aligned}
\sup_{q \in \mathcal{Q}} q\{d_T > k\} &\leq \sup_P P(H_k \leq b-1) \\
&= \sup_P P(S_k \geq k - 2b + 2) \\
&\leq \frac{\sup_P \mathbb{E}_P S_k^{2d}}{(k-2b+2)^{2d}} \\
&\leq \frac{(2d-1)!!\, k^d}{(k-2b+2)^{2d}} \leq 2^{2d}(2d-1)!!\, k^{-d}.
\end{aligned} \tag{105}$$

where $C_r = 2^{2d}(2d-1)!!$ suffices. At $k = K$ the record tail is exactly zero.

Choose $k = \min\{K, \max(4b, \lceil (C_r/\epsilon)^{1/d} \rceil)\}$ and query all records with $d_t \leq k$. If $k < K$, their missing-evidence capacity is at most ϵ; if $k = K$, all values are known. Writing $A_k = \{t : d_t \leq k\}$, the endpoint-width identity (72) reads

$$R_+(\theta_{A_k}, \mathbf{1}_{A_k^c}) - R_+(\theta_{A_k}, \mathbf{0}_{A_k^c}) \leq c_{\mathcal{Q}}(A_k^c) \leq \epsilon, \qquad |A_k| \leq \sum_{h=0}^{b} \binom{k}{h} = O_{b,r}(k^b).$$

The analogous identity for R_- gives the faithful interval and the claimed $O_{b,r}(\min\{K, \epsilon^{-1/d}\}^b)$ cost.

The BCH bound of Benjamini et al. (2012, Theorem 34, Eq. (38)) gives, for $n = 2^m - 1$ and odd order $2d + 1 \le n$, a fair $(2d + 1)$-wise law with an all-one atom of mass at least $\frac{1}{2}(n + 1)^{-d}$. For any $j \ge r$, set $m = \lceil \log_2(j + 1) \rceil$ and restrict this law to j coordinates. If $r = 2d$, the chosen odd length is at least $2d + 1$; if $r = 2d + 1$, it is at least r. Thus the imported construction applies in both cases, and restriction preserves every required moment. Let P_n be the imported law and P_j its restriction. Then

$$P_j(X_{1:j} = \mathbf{1}) \ge P_n(X_{1:n} = \mathbf{1}) \ge \tfrac{1}{2}(2^m)^{-d} \ge 2^{-d-1}(j + 1)^{-d}.$$

Coordinate flips preserve fair moments and send $\mathbf{1}$ to any specified z. For $j < r$, use $P_j = U_j$. The uniform positive constant

$$\gamma_r = \min\left\{2^{-d-1}, \min_{0 \le j < r} 2^{-j}(j + 1)^d\right\}$$

therefore works for all j. Extend with U_{K-j}, independently:

$$\mathbb{E}_{P_j \otimes U_{K-j}} \prod_{i \in A} X_i = \mathbb{E}_{P_j} \prod_{i \in A \cap [j]} X_i \, 2^{-|A \setminus [j]|} = 2^{-|A|} \qquad (|A| \le r).$$

The resulting full-law cylinder bound is

$$\max_{P \in \mathcal{P}(m_r)} P(X_{1:j} = z) \ge \gamma_r (j + 1)^{-d} \quad \text{for every pattern } z \text{ and } j \le K. \tag{106}$$

These classical coding-based independent spaces supply admissible alternatives within the stated domain of unknown laws.

In strict common-overlap geometry, a record whose bth high is at upper-priority position j fixes exactly the first j bits. Its earlier $b - 1$ high positions may be chosen arbitrarily. Thus

$$|\mathcal{C}_j| = \binom{j - 1}{b - 1}, \qquad w_t := \max_{q \in \mathcal{Q}} q_t \ge \gamma_r (j + 1)^{-d} \quad (t \in \mathcal{C}_j). \tag{107}$$

Every witness is a globally compatible assignment law. Different target records may use different extremizers: the prescribed moment class stays fixed, and the algorithm must handle every unknown value vector.

Lemma 14 gives $\Pr_0(t \text{ queried}) \ge 1 - 2\delta$ whenever $w_t > \epsilon$. Apply this single common-baseline bound to the record capacities in (107). Take $c_r = (\gamma_r / 2^{d+1})^{1/d}$ and $J = \min\{K, \lfloor c_r \epsilon^{-1/d} \rfloor\}$. For sufficiently small ϵ, $J \ge 2b$. For every $j \le J$,

$$w_t \ge \gamma_r (j + 1)^{-d} \ge \gamma_r (2J)^{-d} \ge 2\epsilon > \epsilon.$$

Let I_t indicate that record t is queried on the zero baseline. The hitting bound and $N \ge \sum_t I_t$ give

$$\begin{aligned} \mathbb{E}_0 N &\ge \sum_{j=b}^{J} \sum_{t \in \mathcal{C}_j} \Pr_0(I_t = 1) \\ &\ge (1 - 2\delta) \sum_{j=b}^{J} \binom{j - 1}{b - 1} = (1 - 2\delta) \binom{J}{b} \\ &\ge \frac{1 - 2\delta}{b!} (J/2)^b = \Omega_{b,r,\delta}(\min\{K, \epsilon^{-1/d}\}^b). \end{aligned} \tag{108}$$

All expected counts are in the same zero-answer baseline. Together with the upper bound this proves (103). In particular, accuracy below a constant times K^{-d} makes the cost saturate at order K^b.

With only fair singleton information, put probability $1/2$ at a point in a nonempty fiber and $1/2$ at its bitwise complement. Every record therefore has capacity at least $1/2$. For $\epsilon < 1/2$,

$$(1 - 2\delta)|\mathcal{T}| \le N^*_{b,1}(K, \epsilon, \delta) \le N_{\text{all}} = |\mathcal{T}|, \qquad |\mathcal{T}| = \sum_{h=0}^{b} \binom{K}{h} = \Theta_b(K^b)$$

in common-overlap geometry. The upper bound is achieved by querying every record.

For $r = 2, 3$, the theorem specializes to $\Theta_{b,\delta}(\min\{K, \epsilon^{-1}\}^b)$; this is not a separate acquisition theorem. Cantelli's inequality improves the tail constant to $k/[k + (k - 2b + 2)^2]$, and the classical fair-pair atom capacity is $[2\lceil (j+1)/2 \rceil]^{-1}$ (Ramachandra & Natarajan, 2022; Berend et al., 2026). These refinements change constants, not the query exponent or proof structure.

Each alternative changes a suffix value while preserving the keys and every supplied moment. It is realizable by a common retained-information suffix assigning an outcome to each record. A particular beam with shared leaves, or a richer cache oracle, may disallow these independent alternatives; the lower bound does not automatically transfer to that stronger information resource. The deterministic upper bound only needs restoration of the complete state and remains valid for smaller value classes. Root measurement, generation, scoring, judging and arithmetic costs must be added according to their actual interfaces.

D.8 Finite-size parameters for the critical lower bound

For $b = 2$ and $r \in \{2, 3\}$, the stronger lower bound in Theorem 2 has the following precise finite-size version. For any fixed integer $s_0 \geq 2$, replace Λ_K in Eq. (10) by

$$L = \min\left\{ \left\lfloor \log_4 \frac{K - \mathbf{1}\{r = 3\} + 1}{2s_0} \right\rfloor, \ \left\lfloor \tfrac{1}{8} \log_2(1/\epsilon) \right\rfloor \right\} \geq 1. \tag{109}$$

The bound holds on a legal geometry with exactly K families, and the small-ϵ threshold is uniform in K, δ. Section D.9 verifies all amplitudes and padding under this single threshold.

D.9 Record rates from the common lemmas

We prove Theorem 2. Fix the known product vector $p \in [p_0, 1 - p_0]^K$. All constants in this first part depend only on b, r, p_0. The balanced low-order improvement is proved afterward.

Put $d = \lfloor r/2 \rfloor$, $u = \log(8/\delta)$, $k_\epsilon = \min(K, \epsilon^{-1/d})$, and $\Lambda_K = 1 + \log k_\epsilon$. For $S_k = \sum_{i \leq k} X_i$ and $\mu_k = \sum_{i \leq k} p_i$, the degree-$2d$ centered moment agrees with the independent Bernoulli moment. In its expansion, a coordinate appearing once has zero expectation; every surviving term uses at most d coordinates and has absolute expectation at most one. Thus, for $k \geq 2(b-1)/p_0$,

$$\mathbb{E}(S_k - \mu_k)^{2d} \leq C_r \sum_{j=1}^{d} \binom{k}{j} \leq C_r' k^d, \qquad P(S_k < b) \leq \frac{\mathbb{E}(S_k - \mu_k)^{2d}}{(\mu_k - b + 1)^{2d}} \leq \frac{C_r' k^d}{(p_0 k/2)^{2d}}.$$

At $k = K$ the unobserved record tail is zero. The prefix projection (31) extends to the prescribed biases: if P_k has the required prefix moments, set $P = P_k \otimes \bigotimes_{i>k} \mathrm{Bernoulli}(p_i)$. Then

$$\mathbb{E}_P \prod_{i \in A} X_i = \mathbb{E}_{P_k} \prod_{i \in A \cap [k]} X_i \prod_{i \in A \setminus [k]} p_i = \prod_{i \in A} p_i \qquad (|A| \leq r).$$

Restriction gives the reverse inclusion. Choose

$$k = \min(K, \lceil C_{b,r,p_0} \epsilon^{-1/d} \rceil), \qquad \tau_- = 0, \quad \tau_+ = \epsilon/4, \quad \alpha = \beta = \epsilon/4.$$

Lemma 13 gives endpoint excess $\alpha + 2\beta + \tau_+ = \epsilon$ with evaluation charged. For geometric cutoffs $k_h \asymp 2^h$, the actual priority blocks have $M_h = O_b(k_h^b)$ and $w_h = O_{b,r,p_0}(k_h^{-d})$. With $H = \lceil \log_2 k \rceil$,

$$Z \leq C \sum_{h=0}^{H} 2^{(b-2d)h/3}, \qquad \mathsf{H} \leq C \sum_{h=0}^{H} 2^{(b-2d)h/2}.$$

The sum $\sum_{h=0}^{H} x^h$ equals $(x^{H+1} - 1)/(x - 1)$ for $x \neq 1$ and $H + 1$ for $x = 1$, giving

	$b < 2d$	$b = 2d$	$b > 2d$
Z^3	$O(1)$	$O(\Lambda_K^3)$	$O(k_\epsilon^{b-2d})$
H^2	$O(1)$	$O(\Lambda_K^2)$	$O(k_\epsilon^{b-2d})$.

Since $\epsilon k_\epsilon^d \leq 1$, the range term in (84) is absorbed: respectively,

$$\epsilon^{-1} k_\epsilon^{b/2} \leq \epsilon^{-2}, \quad \epsilon^{-1} \Lambda_K k_\epsilon^d \leq \epsilon^{-2} \Lambda_K, \quad \epsilon^{-1} k_\epsilon^{b-d} \leq \epsilon^{-2} k_\epsilon^{b-2d}.$$

This proves all upper bounds for the specified fixed vector p.

For individual lower bounds, first construct a heavy zero assignment with any prescribed n marginal probabilities. Apply Theorem 1 with one scale, use those n bits as the prelude, and choose arbitrary admissible probabilities in an auxiliary block of size n. Marginalizing that block gives

$$P(X_1 = \cdots = X_n = 0) \geq \gamma n^{-d}$$

and every prescribed product moment through order r. For a depth-j record with b specified highs, apply this construction to its $j - b$ zero positions and tensor independent Bernoulli bits on the high positions and the tail. After restoring their original order, this is a feasible law for the given p. Its record mass is at least $\gamma(j-b)^{-d}\prod_{i \text{ high}} p_i \geq \gamma p_0^b j^{-d}$ for sufficiently large j. There are $\binom{j-1}{b-1} \geq c_b j^{b-1}$ such depth-j records. The common-null individual bound (87), with mean shifts $4\epsilon/(c_{b,r,p_0} j^{-d}) \leq 1/4$, yields

$$\mathbb{E}_0 N \geq c\epsilon^{-2} u \sum_{j_0 \leq j \leq \min(K, c'\epsilon^{-1/d})} j^{b-1-2d}.$$

Writing $J_* = \lfloor \min(K, c'\epsilon^{-1/d}) \rfloor$, if $b > 2d$ and $J_* \geq 2j_0$, then $\sum_{j=j_0}^{J_*} j^{b-1-2d} \geq \sum_{j=\lceil J_*/2 \rceil}^{J_*} (J_*/2)^{b-1-2d} \geq c_{b,r} J_*^{b-2d}$. At bounded cutoffs, or $b < 2d$, the subfamily $\theta_t = \mu$ for every record supplies $c\epsilon^{-2}u$ and absorbs the bounded geometric factor into the constant.

At $b = 2d$, let $k_c = \min(K, c\epsilon^{-1/d})$ and $L = \lfloor \log_2(k_c/B_0) \rfloor$, with a fixed admissible integer B_0. If $L \geq 1$, set $N_h = B_0 2^h$, $M_h = \binom{N_h}{2d}$ and $p_h = \gamma N_h^{-d}$. The simultaneous theorem applies to the given product vector and gives $M_h p_h^2 \geq c_{d,p_0} > 0$. Choose the constant c so that $a_h = 4\epsilon/(Lp_h) \leq 1/4$. Raising one uniformly chosen record per block from the common flat mean $1/2$ increases the support value by at least $\sum_h p_h a_h = 4\epsilon$. Averaging the stopped comparison over these independent record choices gives

$$cu \leq C \sum_h \frac{a_h^2}{M_h} \sum_{t \in \mathcal{C}_h} \mathbb{E}_0 N_t = \frac{16C\epsilon^2}{L^2} \sum_h \frac{\mathbb{E}_0 N(\mathcal{C}_h)}{M_h p_h^2} \leq C'\epsilon^2 L^{-2} \mathbb{E}_0 N.$$

The averages are under the same flat experiment, which is independent of the hidden tuple, and the candidate classes are disjoint. For growing k_ϵ, $L \geq c'\Lambda_K$ after a fixed threshold. Below it, the flat-mean lower bound absorbs the bounded Λ_K^2. This proves the basic critical bound for every fixed prescribed vector. At balanced orders two and three, the nested linear construction also provides these witnesses; independent blockwise permutations implement any prescribed high-position tuple and preserve all preceding zeros.

D.10 Product maxima on the full positive domain

The following selector-specific construction is used both in the stronger critical lower bound and in the scalar–policy separation. Fix integers $s_0 \geq 2$, $L \geq 1$ and set

$$s_j = 4^j s_0, \quad N_j = 2s_j - 1, \quad p_j = (2s_j)^{-1}, \qquad M_h = s_{h-1}(2s_{h-1} + 1). \tag{110}$$

Use N_L common-overlap families, with smaller indices preferred among both high and low keys, and width two. Let S_j count highs in the first N_j families and $A_j = \{S_j = 0\}$. This event is measurable from the retained record. Block h adds $6s_{h-1}$ families; call its first bit G_h. Let $\mathcal{C}_h$ be the high–high records whose two relative positions belong to $\{2, \ldots, 2s_{h-1} + 2\}$.

Lemma 21 (Full-domain product maximum)**.** *There is a nonempty face F of the complete fair-pairwise domain whose projection onto these candidates is $\prod_{h=1}^{L} p_h \Delta_{M_h}$. Take*

$$\begin{aligned} a_0 &= \frac{1}{4}, \qquad \beta_j = \frac{1}{4(L+1)}, \qquad \lambda_h = \frac{1}{8(L+1)s_{h-1}}, \\ \ell_\vartheta(t) &= a_0 + \sum_{j=0}^{L} \beta_j \mathbf{1}_{A_j}(t) + \sum_{h=1}^{L} \lambda_h \vartheta_t \mathbf{1}_{\{t \in \mathcal{C}_h\}}. \end{aligned} \tag{111}$$

where every $\vartheta_t \in [0,1]$. All means lie in $[1/4, 17/32]$, and over the entire moment class,

$$R_+(\ell_\vartheta) = a_0 + \sum_{j=0}^{L} \beta_j p_j + \sum_{h=1}^{L} p_h \lambda_h \max_{t \in \mathcal{C}_h} \vartheta_t. \tag{112}$$

Moreover $q(\mathcal{C}_h) \le 4p_h$ for every feasible q, and $M_h p_h^2 = (2s_{h-1}+1)/(64 s_{h-1}) \ge 1/32$.

Proof. Pairwise fairness gives

$$\mathbb{E}S_j = (2s_j - 1)/2, \quad \operatorname{Var} S_j = (2s_j-1)/4, \quad \mathbb{E}(S_j - s_j)^2 = (2s_j-1)/4 + 1/4 = s_j/2.$$

Since $\mathbf{1}_{A_j} \le (S_j - s_j)^2/s_j^2$, $P(A_j) \le p_j$. Equality holds exactly when $S_j \in \{0, s_j\}$ almost surely. Let F impose equality for every j. The initial prefix law putting mass $1/(2s_0)$ at zero and the remaining mass uniformly on count s_0 has

$$\mathbb{E}S_0 = (1 - 1/(2s_0))s_0 = N_0/2, \qquad \mathbb{E}(S_0)_2 = (1-1/(2s_0))s_0(s_0-1) = (N_0)_2/4.$$

Suppose an old prefix has $2s-1$ bits, count in $\{0, s\}$, and zero probability $1/(2s)$. To prescribe any new candidate record, choose a new-block assignment x with $4s$ highs, gate zero, and those first two highs. At most $2s$ zeros have been prescribed, so this is possible. Let $H = \{i : x_i = 1\}$ and $Z = \{i : x_i = 0\}$, of sizes $4s$ and $2s$. For a set A, let $\mathsf{U}_{A,k}$ be the uniform law on bit vectors with exactly k highs in A. With S_{old} the old high count, define the conditional new-block law explicitly by

$$\mathsf{K}_s(\cdot \mid x_{\mathrm{old}}) = \begin{cases} \frac{1}{4}\delta_0 + \frac{1}{4}\delta_x + \frac{1}{2}(\mathsf{U}_{H,2s} \otimes \delta_{\mathbf{1}_Z}), & S_{\mathrm{old}} = 0, \\ \mathsf{U}_{H,2s} \otimes \mathsf{U}_{Z,s}, & S_{\mathrm{old}} = s. \end{cases}$$

All coefficients are nonnegative and sum to one. For a new bit i,

$$\mathbb{E}[X_i \mid S_{\mathrm{old}} = 0] = \begin{cases} 1/4 + (1/2)(2s/(4s)), & i \in H, \\ 1/2, & i \in Z, \end{cases} \qquad \mathbb{E}[X_i \mid S_{\mathrm{old}} = s] = 1/2.$$

Consequently, for every old bit j and new bit i, $\mathbb{E}[(2X_j - 1)(2X_i - 1)] = \mathbb{E}[(2X_j - 1)\mathbb{E}(2X_i - 1 \mid X_{\mathrm{old}})] = 0$. For a uniform fixed-count subset of size k in n coordinates, $\mathbb{E}X_iX_j = k(k-1)/[n(n-1)]$ for distinct indices. At $k = n/2$, $4\mathbb{E}X_iX_j - 2\mathbb{E}X_i - 2\mathbb{E}X_j + 1 = -1/(n-1)$. The zero-history mixture therefore has ZZ correlation one and HH correlation $1/4 + 1/4 - 1/[2(4s-1)] = (2s-1)/(4s-1)$; its HZ correlation is $1/4 - 1/4 = 0$. The full table is

	ZZ	HH	HZ
old zero	1	$(2s-1)/(4s-1)$	0
old nonzero	$-1/(2s-1)$	$-1/(4s-1)$	$0.$

The ZZ and HH pairs satisfy, respectively,

$$\frac{1}{2s} - \frac{2s-1}{2s}\frac{1}{2s-1} = 0, \qquad \frac{2s-1}{2s(4s-1)} - \frac{2s-1}{2s}\frac{1}{4s-1} = 0.$$

The kernel preserves the entire old marginal. Its total count is zero on the first zero-history branch and $4s$ on every other branch. Both the all-zero event and the prescribed-record branch therefore have mass $(1/(2s))(1/4) = 1/(8s)$. The other active zero-history branch has its gate high; nonzero old histories already have two higher-priority highs. For prescribed records $t_1, \dots, t_L$, form the joint law

$$P(x) = P^{(0)}(x_{1:N_0}) \prod_{h=1}^{L} \mathsf{K}_{s_{h-1}}^{t_h}(x_{N_{h-1}+1:N_h} \mid x_{1:N_{h-1}}).$$

Induction using the displayed conditional moments gives $P \in \mathcal{U}_2^{N_L}$, $P(A_h) = p_h$, and $P(T = t_h) = p_h$ on each candidate block. Hence this one law belongs to F and realizes the entire prescribed tuple.

On F,

$$\{T \in \mathcal{C}_h\} = A_{h-1} \cap \{G_h = 0\} \cap A_h^c.$$

Indeed a nonzero new block contains $4s_{h-1}$ highs among $6s_{h-1}$ bits, so with a zero gate its second high occurs by position $2s_{h-1}+2$. Since $\mathbf{1}_{A_{h-1}} = 1 - S_{h-1}/s_{h-1}$ on F, pairwise fairness with G_h gives, with $s = s_{h-1}$,

$$P(A_{h-1}, G_h = 0) = \mathbb{E}[(1 - S_{h-1}/s)(1 - G_h)] = \frac{1}{2} - \frac{2s-1}{4s} = \frac{1}{4s} = \frac{p_{h-1}}{2}.$$

Subtract $P(A_h) = p_h = p_{h-1}/4$ to obtain $q(\mathcal{C}_h) = p_h$. The constructed laws attain every vertex of the product of simplices; convexity proves the exact projection. Without imposing F, $\{T \in \mathcal{C}_h\} \subseteq A_{h-1}$ gives $q(\mathcal{C}_h) \le p_{h-1} = 4p_h$.

Set $m_h = \max_{t \in \mathcal{C}_h} \vartheta_t$ and

$$P_j(x) = \frac{(S_j - s_j)^2}{s_j^2}, \qquad D_h(x) = \frac{(S_{h-1} - s_{h-1})^2 - (S_{h-1} - s_{h-1})(1 - G_h)}{s_{h-1}}.$$

The integer inequality $u^2 - u(1 - G_h) \ge 0$ yields, pointwise,

$$D_h \ge \mathbf{1}_{A_{h-1}, G_h = 0} + s_{h-1}\mathbf{1}_{A_{h-1}}, \quad \mathbf{1}_{\{T \in \mathcal{C}_h\}} \le \mathbf{1}_{A_{h-1}, G_h = 0} - \mathbf{1}_{A_h}, \quad \mathbf{1}_{A_j} \le P_j.$$

Define $c_0 = \beta_0 - s_0\lambda_1 m_1$, $c_j = \beta_j - \lambda_j m_j - s_j\lambda_{j+1}m_{j+1}$ for $1 \le j < L$, and $c_L = \beta_L - \lambda_L m_L$. For $0 < j < L$, $c_j \ge [1/4 - 1/(8s_{j-1}) - 1/8]/(L+1) \ge 0$; the two endpoint coefficients are nonnegative by the same bounds. Consequently the displayed pointwise inequalities give

$$\ell_\vartheta \le a_0 + \sum_h \lambda_h m_h D_h + \sum_j c_j \mathbf{1}_{A_j} \le a_0 + \sum_h \lambda_h m_h D_h + \sum_j c_j P_j.$$

For $s = s_{h-1}$, fairness yields

$$\mathbb{E}[(S_{h-1} - s)(1 - G_h)] = (2s-1)/4 - s/2 = -1/4, \quad \mathbb{E}D_h = \frac{s/2 + 1/4}{s} = \frac{1}{2} + \frac{p_{h-1}}{2}.$$

Since $\mathbb{E}P_j = p_j$, the majorant expectation is

$$a_0 + \sum_j \beta_j p_j + \sum_h \lambda_h m_h \left(\frac{1}{2} + \frac{p_{h-1}}{2} - s_{h-1}p_{h-1} - p_h\right) = a_0 + \sum_j \beta_j p_j + \sum_h \lambda_h m_h p_h,$$

where $s_{h-1}p_{h-1} = 1/2$ and $p_{h-1} = 4p_h$. A law in F prescribing each maximizing record attains that value. The coefficient bound $a_0 + \sum_j \beta_j + \max_h \lambda_h \le 17/32$ proves legality. □

For a top-two geometry with fair moments through order $2d$, add a gate whose high key is above every old key and whose low key is below every old key. Give gate-high records known loss zero and all other records their old losses. The complete fair-$(2d+1)$ endpoints satisfy

$$R_{\pm}^{\mathrm{gate}}(\ell) = \tfrac{1}{2} R_{\pm}(\ell). \tag{113}$$

For every feasible joint law, conditioning on gate bit $J = 0$ gives a fair-$2d$ old law, because $\mathbb{E}[\chi_S(X) \mid J = 0] = \mathbb{E}[\chi_S(X)(1 + \chi_{\{J\}})] = 0$ for $1 \le |S| \le 2d$. Conversely, for any fair-$2d$ law Y, draw an independent fair J and set $X_i = Y_i \oplus J$. For $e \in \{0,1\}$ and $1 \le |S| + e \le 2d+1$,

$$\mathbb{E}[\chi_S(X)(-1)^{eJ}] = \mathbb{E}\chi_S(Y)\,\mathbb{E}(-1)^{(|S|+e)J} = 0.$$

If $|S| + e$ is odd, the second factor is zero. Otherwise $1 \le |S| \le 2d$ and the first factor is zero. This attains every conditional old law. The gate is absent from top two on $J = 0$ and recognizable on $J = 1$; each unknown gated record is simulated by exactly one old-record query.

Take a prescribed pattern on 32 bits and let D be its Hamming distance. Choose $D \in \{0, 13, 14, 19, 20\}$ with respective masses $(31/17290, 2/13, 3/7, 6/19, 1/10)$ and then choose the flipped coordinates uniformly. Its factorial moments are

$$\big(\mathbb{E}(D)_j\big)_{j=0}^4 = (1, 16, 248, 3720, 53940) = \big((32)_j 2^{-j}\big)_{j=0}^4.$$

Count lifting therefore gives four-wise fairness after the prescribed bit flips. Consecutive integer roots make each of $(D-13)(D-14)$ and $(D-19)(D-20)$ nonnegative, giving the pointwise bound

$$\mathbf{1}\{D = 0\} \le \frac{(D-13)(D-14)(D-19)(D-20)}{69160} \quad (D = 0, \dots, 32)$$

Its fair-four-wise expectation is

$$\frac{\mathbb{E}[(D)_4 - 60(D)_3 + 1424(D)_2 - 15808D + 69160]}{69160}$$
$$= \frac{53940 - 60 \cdot 3720 + 1424 \cdot 248 - 15808 \cdot 16 + 69160}{69160} = \frac{31}{17290}.$$

The shell law attains this bound because the four nonzero shells are roots. For an ordinary cumulative-score realization, take $K = 32, b = 4$, with both upper and lower keys decreasing in family index and every upper key above every lower key. Families 1–28 have correct terminal leaves; families 29–32 have incorrect leaves. These labels are fixed and shared across all records. If $D_{\mathrm{key}} = \max_i u_i - \min_i l_i$, assign continuation scores -1 and $-(D_{\mathrm{key}} + 2)$ to correct and incorrect leaves, respectively. Any retained correct family then wins the final cumulative-score comparison. Consequently failure is exactly $E_0 = \{X_{1:28} = 0,\ X_{29:32} = 1\}$. The shell law attains its maximal probability. A uniform odd-parity law on the first 28 bits, independently extended to the remaining four, is fair-four-wise and makes failure impossible. Therefore

$$I(m_4, \theta) = [0, 31/17290], \qquad |I(m_4, \theta)| = 310/1729\,\mathrm{pp} \approx 0.1793\,\mathrm{pp}. \tag{114}$$

This controlled tree uses a fixed continuation and shared leaf labels; its width does not require independently varying record values. The shell law establishes a finite shared-leaf identification width; it is separate from the uniform-in-L higher-order construction.

D.11 The stronger pairwise critical lower bound

For each block in Lemma 21, reserve a baseline record and partition $2m_h = 2\lfloor (M_h - 1)/2 \rfloor$ others into pairs. Put their means at $b_h \pm a_h\sigma_{hi}$, where $b_h = a_0 + \sum_{j<h} \beta_j + \lambda_h/2$ and $0 < a_h$, $(3/4)a_h \le \lambda_h/2$, and $2a_h \le 1/3$. The first amplitude condition covers the entire prior support $|\sigma_{hi}| \le 3/4$; the second is the paid strong-coin simulation condition. All other candidate means equal b_h. This is a legal instance of (111). If $p_h a_h = A$, its endpoint, minus a known constant, is

$$F = A \sum_{h=1}^{L} W_h, \qquad W_h = \max_{i \le m_h} |\sigma_{hi}|.$$

Several pairs may be active simultaneously; the full-domain identity remains exact for each realization. Use independent priors (92) with $\rho_h = 1/(2m_h)$ and $c = 1/4$. For one pair let $G_\pm$ be its absolute-amplitude CDF. On $[\tau_j, 3\tau_j/2)$, $G_- - G_+ = \rho_h\eta_j/J$, and elsewhere they agree. The mean-value theorem and Bernoulli's inequality give

$$\frac{m_h\rho_h\eta_j}{2J} \le m_h(1-\rho_h)^{m_h-1}\frac{\rho_h\eta_j}{J} \le G_-(\tau_j)^{m_h} - G_+(\tau_j)^{m_h} \le \frac{m_h\rho_h\eta_j}{J}.$$

Here $(1-\rho_h)^{m_h-1} \ge 1-(m_h-1)\rho_h \ge 1/2$. Integrating the tail difference and using $m_h\rho_h = 1/2$, $\tau_j\eta_j = \tau_*/4$, yields

$$\Delta_h := \mathbb{E}_+ W_h - \mathbb{E}_- W_h = \sum_j \frac{\tau_j}{2}\{G_-(\tau_j)^{m_h} - G_+(\tau_j)^{m_h}\}, \qquad \frac{\tau_*}{32} \le \Delta_h \le \frac{\tau_*}{16}. \tag{115}$$

Moreover $W_h^2 \le \sum_i |\sigma_{hi}|^2$ and $\sum_{j\ge1}\tau_j^2 = 1/3$ yield

$$\mathrm{Var}_\pm W_h \le \mathbb{E}_\pm W_h^2 \le 31/(96J). \tag{116}$$

Indeed, $\eta_j \le 1/2$ gives, for either prior,

$$\mathbb{E}\sigma_{hi}^2 = \frac{\rho_h}{2J}\sum_j \tau_j^2\{1 \mp \eta_j + \tfrac{9}{4}(1 \pm \eta_j)\} \le \frac{31\rho_h}{16J}\sum_j \tau_j^2 \le \frac{31\rho_h}{48J}.$$

Multiplication by m_h gives (116).

Lemma 22 (A common-baseline information budget)**.** *Use the product-maximum geometry of Lemma 21, with the coefficient choice* (111)*, $L \ge 1$ blocks, and*

$$s_h = 4^h s_0, \quad M_h = s_{h-1}(2s_{h-1}+1), \quad p_h = \frac{1}{8s_{h-1}}, \quad \lambda_h = \frac{1}{8(L+1)s_{h-1}}, \quad s_0 \in \mathbb{N},\ s_0 \ge 2.$$

Reserve one baseline record per block and form $m_h = \lfloor (M_h - 1)/2 \rfloor$ pairs of the other records. Set $b_h = a_0 + \sum_{j<h} \beta_j + \lambda_h/2$. Pair i in block h has full response kernels $Y \mid (\mathcal{F}, \sigma_i, I = \pm) \sim \mathrm{Bernoulli}(b_h \pm a_h \sigma_i)$, where $0 < a_h$, $(3/4)a_h \le \lambda_h/2$, $2a_h \le 1/3$, and $p_h a_h = A$. Initial information and private randomness are independent of every pair latent. Draw the pair latents independently from the two priors in (92)*, with $\rho_h = 1/(2m_h)$, $c = 1/4$, $J \ge 1$, and $\tau_* = 2^{-J}$. Let P_- and P_+ be the complete transcript laws when all pairs use the respective priors. For any common almost surely terminating adaptive algorithm,*

$$\mathrm{KL}(P_- \| P_+) \le C \frac{A^2 \tau_*^2}{J} \mathbb{E}_- N, \tag{117}$$

where N counts actual record observations. The constant is absolute. Transcripts retain all actions, responses, stopping decisions, and the independent private seed. If $\mathbb{E}_- N = \infty$, the bound is understood in the extended sense.

Proof. Treat each complete pair as one latent coordinate. The two arms in a pair share its latent and are not independent prior coordinates. Distinct pairs have independent latents, and each observation depends only on its requested pair's latent. Thus the independent-coordinate part of Lemma 18 applies.

For pair i in block h, simulate all other pairs using their prior-minus latents and independent response streams. These are independent external randomness for the single-pair experiment. Its real-coordinate prior-minus transcript, after marginalizing this simulation randomness, is exactly the common full baseline P_-.

A request with mean $b_h \pm a_h \sigma_i$ is simulated from one stronger observation $X \sim \mathrm{Bernoulli}(1/2 + 2a_h\sigma_i)$. With probability $1/2$, return X for the first arm and $1 - X$ for the second; otherwise return an independent $\mathrm{Bernoulli}(2b_h - 1/2)$ observation. Indeed the resulting means are

$$\frac{1}{2}(1/2 \pm 2a_h\sigma_i) + \frac{1}{2}(2b_h - 1/2) = b_h \pm a_h \sigma_i.$$

The known-noise parameter belongs to $[0, 1]$ because $b_h \in [1/4, 17/32]$, and $2a_h \le 1/3$. Purchase the stronger observation on every pair request, including when it is discarded. Hence there is exactly one stronger-coin observation per original request to this pair.

Let N_i count that pair's original observations, and let $\mathbb{E}_{0,i}$ refer to the simulation with only its latent fixed at zero. Data processing and (93) give

$$\begin{aligned} \mathrm{Hel}^2(P_-, P_{[i]}) &\le C \frac{\rho_h a_h^2 \tau_*^2}{J} \mathbb{E}_{0,i} N_i \\ &\le 2C \frac{\rho_h a_h^2 \tau_*^2}{J} \mathbb{E}_- N_i. \end{aligned}$$

The second inequality uses $\mathbb{E}_- N_i \ge (1 - \rho_h)\mathbb{E}_{0,i} N_i \ge \mathbb{E}_{0,i} N_i/2$; all other pairs remain under the same minus prior.

Each single-pair prior ratio lies in $[1/3, 3]$. Applying (94) and summing pair counts within each block yields

$$\mathrm{KL}(P_- \| P_+) \le C \sum_h \frac{\rho_h a_h^2 \tau_*^2}{J} \mathbb{E}_- N_h.$$

Here every count is under the same product-minus experiment. Since

$$m_h \ge M_h/4, \qquad \rho_h = \frac{1}{2m_h} \le \frac{2}{M_h}, \qquad a_h = \frac{A}{p_h}, \qquad M_h p_h^2 \ge \frac{1}{32},$$

we have $\rho_h a_h^2 \le 2A^2/(M_h p_h^2) \le 64A^2$. Queries to known records contribute cost and no parameter information, so $\sum_h N_h \le N$. This proves (117). Almost-sure stopping and infinite expected costs are covered by the common stopped-experiment convention and the kernel lemma; no expectation-dependent pathwise budget is introduced. □

Proof of Theorem 2: the stronger critical lower bound. We first work at order two on the core geometry with $K_{\mathrm{core}} = 2s_0 4^L - 1$. We will verify the uniform amplitude conditions and the prescribed

number of roots below. Let $\mathcal{C} = \sup_\ell \mathbb{E}_\ell N$. For $\mathcal{C} = \infty$, $\mathcal{C} \ge c\epsilon^{-2}L^2u[1 + \log_+(L/u)]$ holds in the extended real order. In the rest of the proof assume $\mathcal{C} < \infty$.

Take every candidate parameter to be $\vartheta_t = 1/2$. For every tuple containing one unknown record i_h in each block, increase exactly that record's actual mean by

$$\zeta_h = \frac{4\epsilon}{Lp_h}.$$

These are legal instances when $\zeta_h \le \lambda_h/2$. The full-class identity (112) gives an endpoint increase of exactly 4ϵ. Thresholding an ϵ-accurate scalar estimate midway between these two endpoints therefore distinguishes the flat baseline from every tuple alternative with both errors at most δ. All queried means are uniformly separated from zero and one. The stopped Bernoulli change-of-measure bound gives

$$c_1 u \le C \sum_h \zeta_h^2 \mathbb{E}_0 N_{i_h}.$$

This is Lemma 15 for the displayed threshold test, with $\mathrm{kl}(1-\delta, \delta) \ge c_1 u$ for $\delta \le 1/16$. Average over all tuples, uniformly and independently within blocks. The flat transcript is independent of the tuple. Hence

$$\mathbb{E}_I \sum_h \zeta_h^2 \mathbb{E}_0 N_{I_h} = \frac{16\epsilon^2}{L^2} \sum_h \frac{\mathbb{E}_0 N(\mathcal{C}_h)}{M_h p_h^2} \le \frac{512\epsilon^2}{L^2} \mathbb{E}_0 N,$$

which yields

$$\mathcal{C} \ge \mathbb{E}_0 N \ge c_2 \epsilon^{-2} L^2 u. \tag{118}$$

Fix an integer $J \ge 1$ for the moment and use the independent rare-pair priors in Lemma 22. After subtracting a known offset, the true full-class endpoint is exactly

$$F = A \sum_{h=1}^{L} W_h, \qquad W_h = \max_{i \le m_h} |\sigma_{hi}|.$$

Equations (115)–(116), with $c = 1/4$, give independent block variables in $[0, 3/4]$ satisfying

$$\frac{\tau_*}{32} \le \Delta_h := \mathbb{E}_+ W_h - \mathbb{E}_- W_h \le \frac{\tau_*}{16}, \qquad \mathrm{Var}_\pm(W_h) \le \frac{31}{96J}. \tag{119}$$

Write $D = \sum_h \Delta_h$ and choose $A = 4\epsilon/D$. The two prior means of the true endpoint are then exactly 4ϵ apart. Put $t = \epsilon/(2A) = D/8$, so $L\tau_*/256 \le t \le L\tau_*/128$. Bernstein's inequality gives

$$\begin{aligned} P_\pm\{|F - \mathbb{E}_\pm F| > \epsilon/2\} &\le 2\exp\left\{-\frac{t^2}{31L/(48J) + t/2}\right\} \\ &\le 2\exp\{-LJ\tau_*^2/65536\}. \end{aligned} \tag{120}$$

Indeed $J2^{-J} \le 1/2$ and $t \le L\tau_*/128$ give

$$\frac{31L}{48J} + \frac{t}{2} \le \frac{L}{J}\left(\frac{31}{48} + \frac{J\tau_*}{256}\right) \le \frac{2L}{3J}, \qquad \frac{t^2}{31L/(48J) + t/2} \ge \frac{3LJ\tau_*^2}{131072} \ge \frac{LJ\tau_*^2}{65536}.$$

If the final bound in (120) is at most δ, thresholding the estimator at the midpoint of the two prior means has each error at most 2δ. A testing inequality then yields

$$\mathrm{KL}(P_- \| P_+) \ge \log \frac{1}{8\delta} \ge \frac{u}{7}. \tag{121}$$

For a test event E, put $\alpha = P_-(E)$ and $\beta = P_+(E^c)$. Jensen and Cauchy–Schwarz give

$$e^{-\mathrm{KL}(P_- \| P_+)/2} \le \int \sqrt{dP_- dP_+} \le \sqrt{\alpha(1-\beta)} + \sqrt{(1-\alpha)\beta} \le \sqrt{2(\alpha+\beta)} \le \sqrt{8\delta}.$$

Also $\log(1/(8\delta)) \ge \log(8/\delta)/7$ for $\delta \le 1/16$. Combining Lemma 22, $A \asymp \epsilon/(L\tau_*)$, and $\mathbb{E}_- N \le \mathcal{C}$ gives

$$\mathcal{C} \ge c_3 \epsilon^{-2} L^2 J u. \tag{122}$$

Let $R = L/u$. If $R \geq 2^{32}$, take $J = \lfloor \frac{1}{4} \log_2 R \rfloor \geq 8$. Then $\tau_*^2 \geq R^{-1/2}$ and

$$LJ\tau_*^2 \geq uJ\sqrt{R} \geq 65536u.$$

The concentration error in (120) is at most $2e^{-u} = \delta/4$, so (122) applies. Moreover J is at least an absolute positive constant times $\log R$. If $R < 2^{32}$, the factor $1 + \log_+ R$ is bounded by $1 + 32 \log 2$, and (118) applies directly. These two cases prove (10) on the core. This argument does not yield an additional $\log L$ factor uniformly over arbitrarily small δ.

In the multiscale case, $J \leq \lfloor \frac{1}{4} \log_2 L \rfloor$, so $\tau_*^{-1} \leq L^{1/4}$. The bound $D \geq L\tau_*/32$ gives $A \leq 128\epsilon/(L\tau_*)$. Since $p_h \lambda_h/2 = 1/[128(L+1)s_{h-1}^2]$, the condition

$$65536\epsilon\tau_*^{-1}s_{L-1}^2 \leq 1 \tag{123}$$

is sufficient because, even after doubling A for the odd-order reduction,

$$\frac{2A}{p_h\lambda_h/2} \leq 32768\frac{L+1}{L}\epsilon\tau_*^{-1}s_{h-1}^2 \leq 65536\epsilon\tau_*^{-1}s_{L-1}^2 \leq 1.$$

It also suffices for the flat bumps, taking $\tau_* = 1$ for that construction. By (109),

$$s_{L-1}^2 \leq s_0^2\epsilon^{-1/2}, \qquad L \leq \tfrac{1}{8}\log_2(1/\epsilon).$$

Thus its left side is at most $65536s_0^2\epsilon^{1/2}[\log_2(1/\epsilon)/8]^{1/4} \to 0$. Choose ϵ_0 once, depending only on s_0. Every prior sample and every flat alternative then belongs to the legal $\vartheta \in [0,1]$ family, with actual Bernoulli means in the fixed bounded range of (111). This threshold does not deteriorate as δ decreases or K increases.

Put $g = \mathbf{1}\{r = 3\}$. For $r = 3$, keep exactly the same L, J and apply the gate reduction (113) after doubling the pairwise amplitudes and the flat bumps. One added gate halves the full-class endpoint and preserves its unknown-record query interface. The resulting gaps are again 4ϵ; information and legality change only by absolute constants. The number of active families is $K_{\mathrm{core}} + g \leq K$.

For either order put $n = K_{\mathrm{core}} + g$ and use a common-overlap canonical selector on K families, with the active coordinates first. Set $\rho(S) = S \cap [n]$ and $\pi = g_{T_{\mathrm{active}}} \circ \rho$. Lemma 6 gives pointwise

$$\pi(C_2^K(x)) = T_{\mathrm{active}}(x_{1:n}), \qquad \widetilde{\ell}(S) = \ell(\pi(S)).$$

Restriction of a balanced order-r law preserves its active moments; independent balanced extension attains every active law. Hence $R_\pm(\mathcal{Q}_{K,2,r}^{\mathrm{can}}, \widetilde{\ell}) = R_\pm(\mathcal{Q}_{\mathrm{active}}, \ell)$. Every requested full record S is simulated by one query at $\pi(S)$, with its actual conditional Bernoulli law unchanged. Thus the same lower bound holds on this common-overlap geometry; no fully inactive keys or unchanged full-record identities are required. The minimum of the available-depth and accuracy-depth cutoffs in (109) is comparable to $1 + \log \min(K, \epsilon^{-1})$ after a fixed threshold. Below that threshold this factor is bounded, and the flat-mean test supplies the same bound after decreasing its constant. Thus the argument proves (10) with Λ_K for every stated K. □

The direct scalar comparison is proved next in Appendix D.12, using this same product-maximum geometry and the common paid block primitive.

D.12 Direct scalar estimation on the ratio-four family

This comparison uses the full positive domain and the original record queries. The family of continuation means, rather than the feasible routing laws, is restricted. The proof uses the block primitive already established in Lemma 19.

Proposition 23 (Scalar output and feasible-law output)**.** *Fix $s_0 \geq 8$ and use the ratio-four geometry and entire mean family ℓ_ϑ in* (111)*, with*

$$L = \lfloor \log_2(1/\epsilon)/8 \rfloor, \quad K = 2s_0 4^L - 1, \quad u = \log(8/\delta).$$

For sufficiently small ϵ and $0 < \delta \leq 1/16$, estimating the upper support value to error ϵ on this family has a deterministic query upper bound

$$N_{\mathrm{scalar}} \leq C\epsilon^{-2}L^2u[1 + \log_+(L/u)]^2. \tag{124}$$

Returning a feasible $\widehat{q} \in \mathcal{Q}^{\mathrm{can}}_{K,2,2}$ with $R_+(\ell) - \widehat{q}^\top \ell \le \epsilon$ and probability $1-\delta$ has minimax expected cost

$$N_{\mathrm{law}} = \Theta(\epsilon^{-2}(L^3 + L^2 u)). \tag{125}$$

Consequently $N_{\mathrm{scalar}}/N_{\mathrm{law}} \to 0$ if $u = o(L)$. If $u \ge L$, both contracts have cost $\Theta(\epsilon^{-2} L^2 u)$.

Proof. On block $\mathcal{C}_h$ the actual mean is $c_{hi} = b_h^0 + \lambda_h \vartheta_i$, where $b_h^0 = a_0 + \sum_{j<h} \beta_j$ is known. The full-domain certificate (112) gives

$$R_+(\ell) = v_0 + \sum_h p_h \max_i c_{hi}, \qquad v_0 = a_0 + \sum_{j=0}^{L} \beta_j p_j - \sum_h p_h b_h^0, \quad C_h := M_h p_h^2 \asymp 1. \tag{126}$$

Every c_{hi} is observed by one original record query. The Bernoulli conversion in Lemma 19 allows the proof to use independent restarts under the original conditional-mean contract.

We first price a general weighted sum $F = v_0 + \sum_h f_h$, where $f_h = p_h \max_i c_{hi}$ and $C_h = M_h p_h^2$. For every $\alpha > 0$, that lemma supplies $X_h(\alpha)$ with bias in $[-\alpha, 0]$, mean-square error $O(\alpha^2)$, centered exponential moment at most $\exp(Ct^2\alpha^2)$, and cost at most $C C_h \alpha^{-2}$. When $\alpha \ge p_h$, use zero with no query. Set

$$S = \sum_h \sqrt{C_h}, \quad \gamma = \frac{\epsilon}{A\sqrt{uS}}, \quad \alpha_{hj} = 2^{-j}\gamma C_h^{1/4}, \quad J = \left\lceil \log_2 \max(1, \sqrt{L/u}) \right\rceil,$$

with a sufficiently large absolute A. Let $m_{hj} = \mathbb{E} X_h(\alpha_{hj})$. One independent coarse invocation per block gives $X_{\mathrm{coarse}} = \sum_h X_h(\alpha_{h0})$. Since $\sum_h \alpha_{h0}^2 = \gamma^2 S$ and $\sum_h C_h \alpha_{h0}^{-2} = S/\gamma^2$, the product of the centered exponential-moment bounds gives

$$\mathbb{E} e^{t(X_{\mathrm{coarse}} - \mathbb{E} X_{\mathrm{coarse}})} \le e^{Ct^2\gamma^2 S}, \qquad \Pr(|X_{\mathrm{coarse}} - \mathbb{E} X_{\mathrm{coarse}}| > \epsilon/8) \le \delta/4,$$
$$N_{\mathrm{coarse}} \le C\epsilon^{-2} u S^2.$$

If $J \ge 1$, use independent differences $\Delta_{hj} = X_h(\alpha_{hj}) - X'_h(\alpha_{h,j-1})$. Draw (h, j) with probability $\pi_{hj} = 4^{-(j-1)}/(LQ)$, where $Q = \sum_{j=1}^{J} 4^{-(j-1)}$, and return $Z = \Delta_{hj}/\pi_{hj}$. Then

$$\mathbb{E} Z = d_* := \sum_h (m_{hJ} - m_{h0}),$$
$$\mathbb{E} Z^2 \le V := C\gamma^2 LQJS, \qquad \mathbb{E}\,\mathrm{cost}(Z) \le D_* := \frac{CJS}{LQ\gamma^2}, \qquad V D_* \le C J^2 S^2.$$

These estimates use the complete second moment:

$$\mathbb{E}\Delta_{hj}^2 \le 2\mathbb{E}(X_h(\alpha_{hj}) - f_h)^2 + 2\mathbb{E}(X'_h(\alpha_{h,j-1}) - f_h)^2 \le C 4^{-(j-1)} \gamma^2 \sqrt{C_h},$$
$$\mathbb{E} Z^2 = \sum_{h,j} \frac{\mathbb{E}\Delta_{hj}^2}{\pi_{hj}} \le C\gamma^2 LQ \sum_{h,j} \sqrt{C_h} = C\gamma^2 LQJS,$$
$$\mathbb{E}\,\mathrm{cost}(Z) \le \sum_{h,j} \frac{4^{-(j-1)}}{LQ} \frac{C C_h 4^j}{\gamma^2 \sqrt{C_h}} \le \frac{CJS}{LQ\gamma^2}.$$

Let $N_\lambda = \lfloor \lambda \rfloor + \mathrm{Bernoulli}(\lambda - \lfloor \lambda \rfloor)$ be independent and set $W = \lambda^{-1} \sum_{j=1}^{N_\lambda} Z_j$. Since $\mathbb{E} N_\lambda = \lambda$ and $\mathrm{Var}\, N_\lambda \le \lambda$,

$$\mathbb{E} W = d_*, \qquad \mathrm{Var}\, W = \mathrm{Var}\, Z/\lambda + (\mathrm{Var}\, N_\lambda) d_*^2/\lambda^2 \le V/\lambda.$$

Choose $\lambda = 1024V/\epsilon^2$. Then $\Pr(|W - d_*| > \epsilon/8) \le 64V/(\lambda\epsilon^2) \le 1/16$. Interrupt each replicate before its next actual query would exceed $\lceil C'\lambda D_* \rceil$; return a fixed fallback on interruption. The uncapped cost has expectation at most λD_*, so the interruption probability is at most $1/C'$. Take $C' \ge 16$. Each capped replicate fails with probability at most $1/8$; for m independent replicates, Hoeffding bounds median failure by $\exp(-2m(1/2 - 1/8)^2) = \exp(-9m/32)$. Thus $m = O(u)$ gives failure at most $\delta/4$ and deterministic budget

$$m\lceil C'\lambda D_* \rceil \le Cu(\epsilon^{-2} V D_* + 1) \le C' u(\epsilon^{-2} J^2 S^2 + 1).$$

For $J = 0$, take the correction to be zero. In both cases,

$$0 \le \sum_h (f_h - m_{hJ}) \le 2^{-J}\gamma \sum_h C_h^{1/4} \le 2^{-J}\gamma\sqrt{LS} \le \epsilon/A.$$

On the two success events, $|v_0 + X_{\mathrm{coarse}} + \widehat{d}_* - F| \le \epsilon/8 + \epsilon/8 + \epsilon/A \le \epsilon$. For (126), $S \asymp L$, proving (124). The random truncation changes the output on a controlled failure event; unbiasedness of the final capped estimate is not required.

For the law lower bound, let $\theta^0 = \ell_0$ and $v_{\mathrm{base}} = R_+(\theta^0)$. Draw independent $\Theta_i \sim \mathrm{Bernoulli}(1/M_h)$ in each block and set

$$a_h = 16\epsilon/(Lp_h), \qquad \ell_i = \theta_i^0 + a_h\Theta_i.$$

These means belong to the stated family because

$$\frac{a_h}{\lambda_h} = 1024\epsilon s_{h-1}^2 \frac{L+1}{L} \le 2048 s_0^2\sqrt{\epsilon} \longrightarrow 0,$$

where $s_{h-1}^2 \le s_0^2 4^{2L} \le s_0^2\epsilon^{-1/2}$. Put

$$A_h = \{\sum_{i\in\mathcal{C}_h} \Theta_i > 0\}, \quad l = \sum_h a_h p_h \mathbf{1}_{A_h}, \quad R = \sum_h a_h \sum_{i\in\mathcal{C}_h} \Theta_i \widehat{q}_i.$$

Since $P(A_h) = 1 - (1 - 1/M_h)^{M_h} \ge 1 - e^{-1} \ge 1/2$, the exact value and feasibility give

$$R_+(\ell) = v_{\mathrm{base}} + l, \quad 0 \le l \le 16\epsilon, \quad \mathbb{E}l \ge 8\epsilon,$$
$$\{\text{success}\} \Longrightarrow R \ge l - \epsilon, \qquad \mathbb{E}R \ge \mathbb{E}l - \epsilon - 16\epsilon\delta \ge 6\epsilon \ge 3\epsilon.$$

Let $P_{i,0}, P_{i,1}$ fix $\Theta_i = 0, 1$ and retain the same prior on all other latents. Each distinguishing response has KL at most $a_h^2/[c_i^1(1-c_i^1)] \le 16a_h^2$, because both means stay in $[1/4, 17/32]$. Lemma 15, with the other latents included as common independent randomness and then marginalized, gives $\mathrm{KL}(P_{i,0}^{\mathrm{tr}}, P_{i,1}^{\mathrm{tr}}) \le 16a_h^2\mathbb{E}_{i,0}N_i$. Since $0 \le \widehat{q}_i \le w_h := 4p_h$, Pinsker's inequality gives

$$\mathbb{E}_{i,1}\widehat{q}_i \le \mathbb{E}_{i,0}\widehat{q}_i + 2\sqrt{2}\, w_h a_h \sqrt{\mathbb{E}_{i,0}N_i}.$$

For $\pi_h = 1/M_h$, the prior mixture identities are

$$\mathbb{E}(\Theta_i\widehat{q}_i) = \pi_h\mathbb{E}_{i,1}\widehat{q}_i, \qquad \mathbb{E}_{i,0}\widehat{q}_i \le \frac{\mathbb{E}\widehat{q}_i}{1-\pi_h}, \qquad \mathbb{E}_{i,0}N_i \le \frac{\mathbb{E}N_i}{1-\pi_h}.$$

Summing, using $\sum_{i\in\mathcal{C}_h}\widehat{q}_i \le w_h$, and applying Cauchy–Schwarz gives, with $n_h = \mathbb{E}N(\mathcal{C}_h)$,

$$\mathbb{E}\sum_{i\in\mathcal{C}_h}\Theta_i\widehat{q}_i \le \frac{w_h}{M_h - 1} + 2\sqrt{2}\, w_h a_h\sqrt{n_h/(M_h-1)},$$
$$\mathbb{E}R \le \frac{64\epsilon}{L}\sum_h \frac{1}{M_h - 1} + \frac{C\epsilon^2}{L^2}\sum_h \frac{\sqrt{n_h}}{p_h\sqrt{M_h - 1}}$$
$$\le \epsilon + C'\epsilon^2 L^{-3/2}\sqrt{\mathbb{E}N}.$$

Here $M_h \ge 136$ gives $64L^{-1}\sum_h (M_h - 1)^{-1} \le 64/135 < 1$; also $p_h\sqrt{M_h - 1} \ge 1/8$ and $\sum_h \sqrt{n_h} \le \sqrt{L\mathbb{E}N}$. Comparison with $\mathbb{E}R \ge 3\epsilon$ proves $\mathbb{E}N \ge c\epsilon^{-2}L^3$.

For the confidence term, instead put exactly one needle at an independent uniform $I_h \in \mathcal{C}_h$ in every block. Under the baseline θ^0, append these same proof-only independent indices. For the event

$$E = \{(\theta^0)^\top\widehat{q} + \sum_h a_h\widehat{q}_{I_h} \ge v_{\mathrm{base}} + 15\epsilon\},$$

under the augmented null, I_h is independent of the algorithm and $\mathbb{E}_0\widehat{q}_{I_h} = M_h^{-1}\mathbb{E}_0\widehat{q}(\mathcal{C}_h)$. Consequently,

$$\overline{P}_0(E) \le \frac{\sum_h a_h w_h/M_h}{15\epsilon} = \frac{64}{15L}\sum_h \frac{1}{M_h} \le \frac{8}{255}, \qquad \overline{P}_1(E) \ge 1 - \delta.$$

The second inequality is the feasible-law success guarantee, since the alternative value is $v_{\mathrm{base}}+16\epsilon$. Binary data processing and the stopped chain rule give

$$cu \le \mathrm{kl}(\overline{P}_0(E), \overline{P}_1(E)) \le \mathrm{KL}(\overline{P}_0, \overline{P}_1) \le 16 \sum_h \frac{a_h^2}{M_h} \mathbb{E}_0 N(\mathcal{C}_h) \le C\epsilon^2 L^{-2} \mathbb{E}_0 N.$$

For scalar output, the null and alternative values are v_{base} and $v_{\mathrm{base}}+16\epsilon$. Testing the scalar estimate at their midpoint gives the same $\epsilon^{-2}L^2u$ lower bound. Finally, the common finite-profile sampler applied to $1-\ell$ returns a feasible near-optimal law with

$$Z^3 = O(L^3), \quad \mathsf{H} = O(L), \quad \sqrt{M_{\max}} = O(K) = O(\epsilon^{-1/4}),$$
$$N \le C\{\epsilon^{-2}(L^3 + L^2u) + \epsilon^{-1}LKu\} \le C'\epsilon^{-2}(L^3 + L^2u).$$

The last inequality uses $\epsilon K \le C\epsilon^{3/4} \le CL$. This proves (125). If $u = o(L)$, put $x = u/L$; the ratio is bounded by $Cx[1+\log(1/x)]^2 \to 0$. If $u \ge L$, then $J = 0$, the scalar upper bound is $O(\epsilon^{-2}L^2u)$, and the shared confidence lower bound matches it. □

E Structured continuation means under the original record experiment

The unrestricted mean cube permits independent variation of its coordinates; it does not place a probability distribution on the unknown means. This section restricts the mean functions while retaining the complete allocation domain and exactly the same paid record observations.

We prove Theorem 3 with the notation of Section 5.

Proof. Write $m = \min(K, b+r-1)$ and use the basis index set $\mathcal{B} = \mathcal{B}_{K,b,r}$ from Lemma 7. For every $A \in \mathcal{B}$, its record function is $f_A(S) = \mathbf{1}\{A \cap S = \varnothing\}$. Every $g = \sum_{A\in\mathcal{B}} c_A f_A$ has the same expectation under every compatible law:

$$q^\top g = \mu(g) := \sum_{A\in\mathcal{B}} 2^{-|A|} c_A \quad (q \in \mathcal{Q}), \qquad R_+(\theta) - R_-(\theta) \le 2\eta \quad (\theta \in \Theta_\eta). \tag{127}$$

A minimizing g exists because $\mathcal{V}$ is finite-dimensional. Writing $\theta = g + e$ with $\|e\|_\infty \le \eta$ gives

$$R_\pm(\theta) = \mu(g) + \mathrm{ext}_{q\in\mathcal{Q}}\, q^\top e, \qquad |q^\top e| \le \sum_t q_t |e_t| \le \eta.$$

Subtracting the extrema proves the second inequality in (127).

We construct a stable interpolant using finitely many actual records. For $B \subseteq [m]$, let $x_i^B = \mathbf{1}\{i \in [m] \setminus B\}$ and $t_B = C_b(x^B)$. Only sets $B \subseteq A$ for some $A \in \mathcal{B}$ are used. Merge coincident records into an anchor set $\mathcal{J}$; its cardinality is at most $2^m \le 2^{b+r-1}$. For an arbitrary table y on the anchors, set

$$(\mathcal{L}y)(S) = \sum_{A\in\mathcal{B}} f_A(S) \sum_{B\subseteq A} (-1)^{|A|-|B|} y_{t_B}, \qquad \kappa_{b,r} = 1 + \sum_{a=1}^{r} 2^a \binom{b+a-1}{a}. \tag{128}$$

Indeed, for $g \in \mathcal{V}$,

$$g(t_B) = \sum_{D\in\mathcal{B}: D\subseteq B} c_D,$$
$$\sum_{B\subseteq A} (-1)^{|A|-|B|} g(t_B) = \sum_{D\in\mathcal{B}: D\subseteq A} c_D \sum_{D\subseteq B\subseteq A} (-1)^{|A|-|B|} = c_A. \tag{129}$$

The inner sum in (129) is $(1-1)^{|A\setminus D|}$, hence $\mathcal{L}(g|_{\mathcal{J}}) = g$. For an arbitrary anchor table,

$$|(\mathcal{L}y)(S)| \le \sum_{A\in\mathcal{B}} \sum_{B\subseteq A} |y_{t_B}| \le \sum_{A\in\mathcal{B}} 2^{|A|} \|y\|_\infty \le \kappa_{b,r} \|y\|_\infty.$$

At every anchor purchase

$$n_0 = \left\lceil \frac{(64\kappa_{b,r})^2}{2\epsilon^2} \log \frac{8|\mathcal{J}|}{\delta} \right\rceil$$

responses and let $\widehat{\theta}_j$ be their average. Each response has its prescribed mean conditional on the full past. The conditional Hoeffding bound and a union bound give

$$\Pr\left\{ \max_{j \in \mathcal{J}} |\widehat{\theta}_j - \theta_j| > \frac{\epsilon}{64\kappa_{b,r}} \right\} \le 2|\mathcal{J}| \exp\left\{ -2n_0 \left(\frac{\epsilon}{64\kappa_{b,r}} \right)^2 \right\} \le \delta/4.$$

Write the complementary event as

$$\mathcal{E}_0 = \{ \max_{j \in \mathcal{J}} |\widehat{\theta}_j - \theta_j| \le \epsilon/(64\kappa_{b,r}) \}, \qquad \Pr(\mathcal{E}_0) \ge 1 - \delta/4.$$

Let $\widehat{g} = \mathcal{L}\widehat{\theta}$ and $e = \theta - \widehat{g}$. On $\mathcal{E}_0$, for any admissible approximant g,

$$\begin{aligned} \|e\|_\infty &\le \|\theta - g\|_\infty + \|\mathcal{L}((g-\theta)|_{\mathcal{J}})\|_\infty + \|\mathcal{L}((\theta - \widehat{\theta})|_{\mathcal{J}})\|_\infty \\ &\le (1 + \kappa_{b,r})\eta + \epsilon/64 =: a_\eta. \end{aligned} \tag{130}$$

The pilot costs $O_{b,r}(\epsilon^{-2}u)$. For every pilot outcome, not just on $\mathcal{E}_0$, $\|\widehat{g}\|_\infty \le \kappa_{b,r}$ because every anchor average belongs to $[0,1]$.

Choose an integer cutoff $\bar{k} \ge \min(K, \max(m, 4b))$ such that

$$a_\eta \sup_{q \in \mathcal{Q}} q\{d_t > \bar{k}\} \le \epsilon/32, \qquad \bar{k} \le C_{b,r} k_\eta. \tag{131}$$

The capacity bound supplies the choice

$$\bar{k} = \min\left\{ K, \max\left(m, 4b, \left\lceil (32 C_{b,r} a_\eta/\epsilon)^{1/d} \right\rceil \right) \right\}.$$

If $\bar{k} < K$, then $a_\eta C_{b,r} \bar{k}^{-d} \le \epsilon/32$. If $\bar{k} = K$, the tail is empty. In both cases $a_\eta \le C_{b,r}(\eta + \epsilon)$ gives the second bound in (131). Let $\Pi_{\bar{k}}$ retain every record of depth at most $\bar{k}$ and merge the other records into one dummy. Its exact domain is $\mathcal{D}_{\bar{k}} = \Pi_{\bar{k}} \mathcal{Q}$. Assign the dummy residual mean zero, and denote this projected vector by e'. On $\mathcal{E}_0$, for every $q \in \mathcal{Q}$,

$$\left| q^\top \theta - \mu(\widehat{g}) - (\Pi_{\bar{k}} q)^\top e' \right| = \left| \sum_{t : d_t > \bar{k}} q_t e_t \right| \le a_\eta q\{d_t > \bar{k}\}.$$

Taking minima or maxima over the same laws, with $\widetilde{R}_\pm = \operatorname{ext}_{z \in \mathcal{D}_{\bar{k}}} z^\top e'$, gives

$$|R_\pm(\theta) - \mu(\widehat{g}) - \widetilde{R}_\pm| \le a_\eta \sup_{q \in \mathcal{Q}} q\{d_t > \bar{k}\} \le \epsilon/32. \tag{132}$$

The identity $q^\top \widehat{g} = \mu(\widehat{g})$ holds on the full original domain, including outside any attaining face.

Let $\mathcal{H}_0$ be the complete pilot history and put $D = 1 + 2\kappa_{b,r}$. For a requested retained record, buy one fresh response Y and return to the sampler

$$W = \frac{Y - \widehat{g}_t + \kappa_{b,r}}{D} \in [0,1], \qquad \mathbb{E}[W \mid \mathcal{H}_0, \mathcal{F}_{n-1}, I_n = t] = \frac{e_t + \kappa_{b,r}}{D}.$$

The dummy returns $\kappa_{b,r}/D$ without a query. Thus $c = (e' + \kappa_{b,r}\mathbf{1})/D \in [0,1]$ and $\widetilde{R}_\pm = DR_\pm(\mathcal{D}_{\bar{k}}, c) - \kappa_{b,r}$. Apply Lemma 16 to this compact convex probability domain for c and $1 - c$. Use training and fresh-evaluation tolerances $\alpha' = \beta' = \epsilon/(8D)$. After inverting the affine map, the feasible values $m_\pm$ and their estimates satisfy, with $\alpha = \beta = \epsilon/8$,

$$\widetilde{R}_- \le m_- \le \widetilde{R}_- + \alpha, \quad \widetilde{R}_+ - \alpha \le m_+ \le \widetilde{R}_+, \quad |\widehat{m}_\pm - m_\pm| \le \beta.$$

Allocate failure at most $\delta/8$ to each of the two training and two evaluation calls, conditional on every $\mathcal{H}_0$. For $t = \epsilon/32$, return

$$[\mu(\widehat{g}) + \widehat{m}_- - \alpha - \beta - t,\ \mu(\widehat{g}) + \widehat{m}_+ + \alpha + \beta + t] \cap [0,1].$$

Lemma 13, with two-sided tail allowance t, gives outer coverage and endpoint excess $\alpha+2\beta+2t = 7\epsilon/16 \le \epsilon$. The total failure is at most $\Pr(\mathcal{E}_0^c) + 4\delta/8 \le 3\delta/4$. Every sampler call costs at most one fresh original response.

For dyadic prefix shells, $M_h \le C_{b,r} 2^{bh}$ and $w_h \le C_{b,r} 2^{-dh}$. Write $Z = \sum_h (M_h w_h^2)^{1/3}$ and $\mathsf{H} = \sum_h \sqrt{M_h} w_h$. With $H = \lceil \log_2 \bar{k} \rceil$, the known dummy adds only a constant, and

$$Z \le C_{b,r} \sum_{h=0}^{H} 2^{(b-2d)h/3}, \qquad \mathsf{H} \le C_{b,r} \sum_{h=0}^{H} 2^{(b-2d)h/2}.$$

For a geometric ratio $\rho \ne 1$, the sum is $(\rho^{H+1}-1)/(\rho-1)$; at $\rho = 1$ it is $H+1$. Since $\bar{k} \le C_{b,r} k_\eta$, these give

	$b < 2d$	$b = 2d$	$b > 2d$
Z^3	$O_{b,r}(1)$	$O_{b,r}(\Lambda_\eta^3)$	$O_{b,r}(k_\eta^{b-2d})$
H^2	$O_{b,r}(1)$	$O_{b,r}(\Lambda_\eta^2)$	$O_{b,r}(k_\eta^{b-2d})$.

The remaining range term is $C_{b,r}\epsilon^{-1}\mathsf{H}\bar{k}^{b/2}u$. The cutoff also gives $\epsilon\bar{k}^d \le C_{b,r}(\eta+\epsilon) \le C_{b,r}$, so

$$\epsilon \mathsf{H} \bar{k}^{b/2} \le C_{b,r} \mathsf{H} \bar{k}^{b/2-d} \le C'_{b,r} \begin{cases} 1, & b < 2d, \\ \Lambda_\eta^2, & b = 2d, \\ k_\eta^{b-2d}, & b > 2d. \end{cases}$$

Adding the pilot proves the upper bounds. The exact prefix domain uses at most $2^{\bar{k}}$ assignment atoms by restriction and independent extension of balanced laws; the finite-bit construction applies to the transformed responses as well.

We next prove the lower bounds in the same observation experiment. All constant mean vectors belong to Θ_0, so the common constant shift test gives $c\epsilon^{-2}u$ for every η. For $\eta > 0$, the class also contains the entire cube $\{\frac{1}{2}\mathbf{1} + v : \|v\|_\infty \le \eta\}$. If a record has capacity $c_t \ge 4\epsilon/\eta$, compare the common baseline $\theta^0 = \mathbf{1}/2$ with $\theta^t = \theta^0 + (4\epsilon/c_t)e_t$. Both lie in Θ_η, and their upper support values are separated by 4ϵ. Midpoint thresholding has both testing errors at most δ. The worlds differ only at t, with $a_t = 4\epsilon/c_t \le 1/8$:

$$\begin{aligned} \mathrm{kl}(1-\delta, \delta) &\le \mathrm{kl}(1/2, 1/2 + a_t)\mathbb{E}_0 N_t \\ &= -\tfrac{1}{2}\log(1-4a_t^2)\mathbb{E}_0 N_t \le 3a_t^2 \mathbb{E}_0 N_t, \\ \mathbb{E}_0 N_t &\ge c\epsilon^{-2} u c_t^2. \end{aligned} \tag{133}$$

The heavy-atom construction and record multiplicities in Appendix D.7 give $c_t \ge c_{b,r} j^{-d}$ for at least $c_{b,r} j^{b-1}$ records of depth j beyond a fixed initial depth. Sum (133) up to a constant multiple of $\min(K, (\eta/\epsilon)^{1/d})$. For $b > 2d$,

$$\mathbb{E}_0 N \ge c_{b,r}\epsilon^{-2} u \sum_{j \le c_{b,r} k_\eta} j^{b-1-2d} \ge c'_{b,r}\epsilon^{-2} u k_\eta^{b-2d}.$$

When the summation range is bounded, the constant-shift test supplies the same bound after adjustment of the fixed constant.

For $b = 2d$, the simultaneous-cylinder construction gives disjoint candidate classes of sizes $M_h \asymp N_h^{2d}$ and masses $p_h \ge c_{b,r} N_h^{-d}$ under a common positive witness for every tuple. Choose L scales with

$$B 2^L \le K, \qquad 2^{dL} \le (\eta/\epsilon)^{1/2}, \qquad a_h = \frac{4\epsilon}{L p_h}.$$

If η/ϵ is sufficiently large, L can be chosen with $1 + L \asymp_{b,r} \Lambda_\eta$ and $\max_h a_h \le C_{b,r}\sqrt{\epsilon\eta}/L \le \eta$. Raise one uniformly selected coordinate by a_h in every block. All these mean vectors belong to Θ_η and a genuine witness raises the endpoint by at least 4ϵ. Averaging over uniformly chosen tuples gives

$$cu \le C \sum_h \frac{a_h^2}{M_h} \sum_{t \in \mathcal{C}_h} \mathbb{E}_0 N_t = \frac{16 C \epsilon^2}{L^2} \sum_h \frac{\mathbb{E}_0 N(\mathcal{C}_h)}{M_h p_h^2} \le \frac{C_{b,r}\epsilon^2}{L^2} \mathbb{E}_0 N.$$

Here $M_h p_h^2 \geq c_{b,r} > 0$, the classes are disjoint, and every expectation is under the same flat baseline. For $d = 1$ the balanced nested linear construction supplies these witnesses; for $d \geq 2$ use Theorem 1. Bounded η/ϵ or bounded L is again covered by the constant shift. This proves the middle lower entry of (13).

For the strengthened order-two and order-three lower bound, retain the ratio-four family ℓ and its independent pair priors from Section D.10, but use actual record means

$$\theta = \tfrac{1}{2}\mathbf{1} + \eta(\ell - \tfrac{1}{2}\mathbf{1}). \tag{134}$$

Every realization belongs to Θ_η and every actual mean lies in $[7/16, 9/16]$. The optional odd-order gate is transformed by the same formula. Choose the number of geometric scales using relative precision ϵ/η:

$$L = \min\left\{ \left\lfloor \log_4 \frac{K - \mathbf{1}\{r=3\} + 1}{2s_0} \right\rfloor, \left\lfloor \tfrac{1}{8}\log_2(\eta/\epsilon) \right\rfloor \right\}, \qquad s_0 \geq 2 \text{ fixed.}$$

For L and η/ϵ beyond fixed thresholds, $L \asymp \Lambda_\eta$. In the transformed pair means, let the actual amplitude be $\widetilde{a}_h$ and choose $p_h \widetilde{a}_h = A$ with $A \asymp \epsilon/(L\tau_*)$ as in the proof of the stronger critical bound. Write $\widetilde{b}_h = 1/2 + \eta(b_h - 1/2)$. The mixture simulation of Lemma 22 has the required actual mean:

$$\tfrac{1}{2}(1/2 \pm 2\widetilde{a}_h\sigma) + \tfrac{1}{2}(2\widetilde{b}_h - 1/2) = \widetilde{b}_h \pm \widetilde{a}_h\sigma.$$

Since $\widetilde{b}_h \in [7/16, 9/16]$, its nuisance coin has parameter $2\widetilde{b}_h - 1/2 \in [3/8, 5/8]$. One stronger coin is charged for each simulated pair request. Using $\rho_h \leq 2/M_h$ and $p_h\widetilde{a}_h = A$ in that lemma gives

$$\begin{aligned} \mathrm{KL}(P_-, P_+) &\leq C \sum_h \frac{\rho_h \widetilde{a}_h^2 \tau_*^2}{J} \mathbb{E}_- N_h \\ &\leq \frac{2CA^2\tau_*^2}{J} \sum_h \frac{\mathbb{E}_- N_h}{M_h p_h^2} \leq C' \frac{\epsilon^2}{L^2 J} \mathbb{E}_- N. \end{aligned} \tag{135}$$

The last line uses $M_h p_h^2 \geq 1/32$ and $A \asymp \epsilon/(L\tau_*)$. The amplitude condition is now $\widetilde{a}_h \leq \eta\lambda_h/2$. It follows from

$$65536(\epsilon/\eta)\tau_*^{-1} s_{L-1}^2 \leq 1, \qquad s_{L-1}^2 \leq s_0^2(\eta/\epsilon)^{1/2}, \quad \tau_*^{-1} \leq L^{1/4},$$

for sufficiently small ϵ/η, including the fixed extra factor for the odd-order gate. Put $D_* = \sum_h (\mathbb{E}_+ W_h - \mathbb{E}_- W_h)$ and $A = 4\epsilon/(g_r D_*)$, where $g_2 = 1$ and $g_3 = 1/2$ is the odd gate mass. The prior means are 4ϵ apart, and $L\tau_*/32 \leq D_* \leq L\tau_*/16$. The same independent W_h satisfy $\mathrm{Var}_\pm W_h \leq 31/(96J)$, so (120) gives

$$P_\pm\left\{ \left| g_r A \sum_h (W_h - \mathbb{E}_\pm W_h) \right| > \epsilon/2 \right\} \leq 2e^{-LJ\tau_*^2/65536}.$$

For $L/u \geq 2^{32}$, take $J = \lfloor \frac{1}{4}\log_2(L/u) \rfloor$. Then $LJ\tau_*^2 \geq 65536u$. Midpoint testing and (121) give

$$u/7 \leq \mathrm{KL}(P_-, P_+) \leq C\frac{\epsilon^2}{L^2 J}\mathbb{E}_- N, \qquad \mathbb{E}_- N \geq c\epsilon^{-2} L^2 J u.$$

When its thresholds fail, Λ_η is bounded or the extra logarithm is bounded, and the preceding lower bounds suffice. This proves (14) using the transformed record responses. □

There is an optimization-free implementation whenever $\eta \leq \epsilon/4$. Let $q^0 = (C_b)_\# U$ for independent fair assignments and put $v_0 = (q^0)^\top \theta$. Since $q^0 \in \mathcal{Q}$ and $\theta = g + e$ with $g \in \mathcal{V}$, $\|e\|_\infty \leq \eta$,

$$R_- \leq v_0 \leq R_+, \qquad 0 \leq v_0 - R_- \leq 2\eta, \qquad 0 \leq R_+ - v_0 \leq 2\eta.$$

Draw fresh records from q^0, purchase one actual response for each, and average $n = \lceil 8\epsilon^{-2}\log(2/\delta) \rceil$ responses to obtain $\widehat{v}_0$. With probability at least $1 - \delta$, $|\widehat{v}_0 - v_0| \leq \epsilon/4$. Return

$$[\widehat{v}_0 - \epsilon/4 - 2\eta, \widehat{v}_0 + \epsilon/4 + 2\eta] \cap [0, 1]. \tag{136}$$

The two inequalities relative to v_0 give outer coverage. Moreover,

$$R_- - (\widehat{v}_0 - \epsilon/4 - 2\eta) \leq \epsilon/2 + 2\eta \leq \epsilon, \quad (\widehat{v}_0 + \epsilon/4 + 2\eta) - R_+ \leq \epsilon/2 + 2\eta \leq \epsilon.$$

Thus the displayed interval is faithful for every K, without an anchor pilot or an optimization call. At $\eta = 0$ both target endpoints coincide, but their common value still has to be estimated. For larger specified η, the pilot-and-truncation construction supplies the bound in Theorem 3. A reference implementation of this optimization-free estimator is `scripts/structured_mean_audit.py`; it charges every callback and checks the specified residual and query-budget conditions before sampling.

More general known feature models also connect to the same experiment. Suppose $\theta_t = \phi_t^\top \beta$, the feature span has dimension p, and p legal anchor records give an invertible row matrix F. Write $L_t = \phi_t^\top F^{-1}$ and $\kappa_\Phi = \max_t \|L_t\|_1$. Estimate each anchor mean to precision $\epsilon/(4\kappa_\Phi)$ with failure at most δ/p. Then

$$\widehat{\theta}_t = L_t \widehat{\theta}_{\mathrm{anchor}}, \qquad \|\widehat{\theta} - \theta\|_\infty \le \kappa_\Phi \|\widehat{\theta}_{\mathrm{anchor}} - \theta_{\mathrm{anchor}}\|_\infty \le \epsilon/4.$$

Because every q is a probability vector,

$$|R_\pm(\widehat{\theta}) - R_\pm(\theta)| \le \sup_{q \in \mathcal{Q}} |q^\top(\widehat{\theta} - \theta)| \le \|\widehat{\theta} - \theta\|_\infty \le \epsilon/4.$$

This gives an observation bound $O(p\kappa_\Phi^2 \epsilon^{-2} \log(2p/\delta))$, with optimization charged separately. It is the usual finite linear-model estimation argument (Soare et al., 2014, Proposition 1) applied to record features. Family-additive means use features $\phi_{t,(i,x)} = b^{-1}\mathbf{1}\{(i,x) \in t\}$ and rank at most $2K$. Their anchor condition number can depend on K. The original-record family/state estimator in Appendix E.1 avoids this generic condition-number dependence by graded record contrasts. The residual theorem above obtains its stronger cutoff using the particular moment-visible feature space, whose anchor dimension and stability bound are uniform in K.

E.1 Selected-family interactions and a paid additive reduction

We use the canonical record set $\mathcal{S}_{K,b} = \{S \subseteq [K] : |S| \le b\}$ and the complete domain $\mathcal{Q}^{\mathrm{can}}_{K,b,r}$. A query remains a fresh bounded response at a chosen original record S. A coefficient observation below is implemented from such a response rather than supplied by a new channel.

For $1 \le s \le b$, define the selected-family interaction class

$$\mathcal{M}_s = \left\{ \theta \in [0,1]^{\mathcal{S}_{K,b}} : \theta_S = \sum_{J \subseteq S,\ |J| \le s} \beta_J \text{ for some real } (\beta_J)_{|J| \le s} \right\}. \tag{137}$$

The unique coefficients are

$$\beta_J = \sum_{U \subseteq J} (-1)^{|J|-|U|} \theta_U. \tag{138}$$

On records $|J| \le s$, inversion follows from

$$\sum_{U \subseteq J} (-1)^{|J|-|U|} \sum_{D \subseteq U} \beta_D = \sum_{D \subseteq J} \beta_D (1-1)^{|J \setminus D|} = \beta_J.$$

At $s = b$ the class is the entire mean cube. At $s = 1$ it is $\theta_S = \theta_\varnothing + \sum_{i \in S}(\theta_{\{i\}} - \theta_\varnothing)$. For example, $\theta_\varnothing = 1/2$ and independently chosen $\beta_{\{i\}} \in [-1/(4b), 1/(4b)]$ give a K-dimensional box in $\mathcal{M}_1$ with every record mean in $[1/4, 3/4]$. The parameter range has no K-dependent shrinking radius. Degree in selected-family indicators differs from degree in the original assignment bits: $\mathbf{1}\{i \in C_b(X)\}$ can depend on a long prefix. Thus $\mathcal{M}_s$ is not the space of losses whose expectation is fixed by the released order-r moments.

Corollary 24 (Interaction order under unchanged record access)**.** *Fix $b \ge 1$, $r \ge 2$ and $1 \le s \le b$, and set $d = \lfloor r/2 \rfloor$. For $0 < \epsilon \le 1/16$, $0 < \delta \le 1/16$ and $K \ge \max(b, r)$, put $k_* = \min\{K, \lceil C_{b,r,s} \epsilon^{-1/d} \rceil\}$, $\Lambda_* = 1 + \log k_*$ and $u = \log(8/\delta)$. On $\mathcal{M}_s$, a faithful interval can be obtained with deterministic original-record query budget bounded by a constant depending only on b, r, s times*

$$\begin{cases} \epsilon^{-2} u, & s < 2d, \\ \epsilon^{-2}(\Lambda_*^3 + \Lambda_*^2 u), & s = 2d, \\ \epsilon^{-2} k_*^{s-2d} u, & s > 2d. \end{cases} \tag{139}$$

When $s < 2d$, the minimax rate is $\Theta_{b,r,s}(\epsilon^{-2}u)$, uniformly in K. The lower bound already holds for scalar upper-endpoint estimation with Bernoulli responses and arbitrary adaptive stopping.

Proof. Put $C_s = \sum_{j=0}^{s} \binom{b}{j}$, $a_s = 2^s$, and $\mathcal{J}_s = \{J \subseteq [K] : |J| \le s\}$. For $q \in \mathcal{Q}^{\mathrm{can}}_{K,b,r}$ write

$$z_J(q) = \sum_{S \supseteq J} q_S, \qquad D_J(q) = z_J(q)/C_s, \qquad D_\dagger(q) = 1 - C_s^{-1} \sum_{J \in \mathcal{J}_s} z_J(q).$$

Here $\dagger$ is a known dummy. Since $\sum_J z_J(q) = \sum_S q_S \sum_{j=0}^{s} \binom{|S|}{j} \le C_s$, the coordinates are nonnegative and sum to one. The feature domain is the exact linear image $\mathcal{D}_s = \{D(q) : q \in \mathcal{Q}^{\mathrm{can}}_{K,b,r}\}$. Set $v_J = (\beta_J + a_s)/(2a_s)$ and $v_\dagger = 1/2$. Equation (138) gives $|\beta_J| \le 2^{|J|} \le a_s$, and

$$\begin{aligned} D(q)^\top v &= \frac{1}{2a_s C_s} \sum_J z_J(q)\beta_J + \frac{1}{2C_s} \sum_J z_J(q) + \frac{1}{2} D_\dagger(q) \\ &= \frac{q^\top \theta}{2a_s C_s} + \frac{1}{2}. \end{aligned} \tag{140}$$

Both extrema are therefore affine rescalings of extrema over this positive image; the law q is common to all feature coordinates.

To observe v_J, draw U uniformly from the subsets of J, purchase one original response Z_U, and return

$$W_J = 2^{|J|}(-1)^{|J|-|U|} Z_U, \qquad Y_J = (W_J + a_s)/(2a_s). \tag{141}$$

Conditionally on the complete past $\mathcal{F}$,

$$\mathbb{E}[W_J \mid \mathcal{F}, J] = 2^{-|J|} \sum_{U \subseteq J} 2^{|J|}(-1)^{|J|-|U|}\theta_U = \beta_J, \qquad \mathbb{E}[Y_J \mid \mathcal{F}, J] = v_J.$$

The bound $|W_J| \le 2^{|J|} \le a_s$ gives $Y_J \in [0,1]$. Each observation costs one fresh record response, including the query to $\varnothing$ when it is drawn. The dummy costs no response because its value is known. Different features may query the same physical record; the finite-profile upper bound requires the displayed conditional means, not independent unknown parameters.

Give nonempty J depth $\nu(J) = \max J$, and give $\varnothing, \dagger$ depth zero. There are at most $2 + \sum_{j=1}^{s} \binom{k}{j} = O_s(k^s)$ features up to depth k. If $\nu(J) > k$ and $J \subseteq S$, the canonical record cannot be determined by the first k assignment bits. Hence, writing d_T for record depth,

$$\begin{aligned} \sum_{J:\nu(J)>k} D_J(q) &= C_s^{-1} \sum_S q_S \#\{J \subseteq S : |J| \le s,\ \nu(J) > k\} \\ &\le q\{d_T > k\} \le C_{b,r} k^{-d}. \end{aligned}$$

The tail vanishes at $k = K$. At dyadic scales the feature domain has $M_h \le C_s 2^{sh}$ and $w_h \le C_{b,r} 2^{-dh}$. Equation (91) applies to these counts, capacities and bounded conditional responses at tolerance $\epsilon' = \epsilon/(2a_s C_s)$. The affine identity (140) multiplies each endpoint excess by $2a_s C_s$. Since a_s, C_s are fixed, this proves (139). For the lower bound, the subfamily $\theta_S = \mu$ belongs to every $\mathcal{M}_s$. Both endpoints equal μ and every physical response can be Bernoulli(μ). Compare $\mu_0 = 1/2$ and $\mu_1 = 1/2 + 4\epsilon \le 3/4$. Midpoint testing and Lemma 15 give

$$\mathrm{kl}(1-\delta, \delta) \le \mathrm{kl}(\mu_0, \mu_1)\mathbb{E}_0 N \le 48\epsilon^2 \mathbb{E}_0 N.$$

Thus $\mathbb{E}_0 N \ge c\epsilon^{-2}\log(1/\delta)$ for every query policy. □

The additive case admits uniform acquisition of the risk of every feasible law. For an arbitrary vector in the original mean cube define

$$(\mathcal{A}\theta)_S = \theta_\varnothing + \sum_{i \in S} (\theta_{\{i\}} - \theta_\varnothing).$$

This vector can leave $[0,1]^{\mathcal{S}_{K,b}}$, but each nonconstant coefficient is a difference of original means and has absolute value at most one.

Lemma 25 (Paid uniform acquisition of the additive component)**.** *Fix $b \ge 1$, $r \ge 2$ and $K \ge \max(b, r)$. For every $0 < e \le 1/16$ and $0 < \delta \le 1/16$, a deterministic-budget procedure querying only $\varnothing$ and singleton records returns a known vector $\widehat{f}$ such that, uniformly on the full mean cube,*

$$\Pr_\theta \left\{ \sup_{q \in \mathcal{Q}^{\mathrm{can}}_{K,b,r}} |q^\top(\widehat{f} - \mathcal{A}\theta)| \le e \right\} \ge 1 - \delta, \qquad N \le C_{b,r} e^{-2} \log(8/\delta). \tag{142}$$

Proof. Enlarge the tail constant so that it applies to every positive integer cutoff, and put $k = \min\{K, \lceil (8bC_{b,r}/e)^{1/d} \rceil\}$. Partition $[k]$ into $I_0 = \{1\}$ and $I_h = (2^{h-1}, 2^h] \cap [k]$ for $h \geq 1$, omitting empty blocks. Let $M_h = |I_h|$ and set

$$w_0 = b, \qquad w_h = b \min\{1, C_{b,r} 2^{-d(h-1)}\} \quad (h \geq 1).$$

For $\mu_i(q) = q\{S : i \in S\}$, the depth implication above and $|S| \leq b$ give

$$\sum_{i \in I_h} \mu_i(q) \leq w_h, \qquad \sum_{i > k} \mu_i(q) \leq e/8. \tag{143}$$

The latter sum vanishes if $k = K$.

Allocate failure probability $\delta/2$ to the empty record and $\delta_h = \delta 2^{-h-2}$ to block h. Define

$$\ell_h = \log(2M_h/\delta_h), \quad D = \sum_h (M_h \ell_h w_h^2)^{1/3}, \quad e_h = \frac{e}{4D} \left(\frac{M_h \ell_h}{w_h} \right)^{1/3}.$$

For $e_h < 1$, take $n_h = \lceil \ell_h/(2e_h^2) \rceil$ observations at each singleton in I_h, and let $\widehat{\theta}_i$ be their average. For $e_h \geq 1$, use $\widehat{\theta}_i = 1/2$ without any observation. For $e_h < 1$, conditional Hoeffding (Hoeffding, 1963) gives

$$\Pr\{\exists i \in I_h : |\widehat{\theta}_i - \theta_i| > e_h\} \leq 2M_h e^{-2n_h e_h^2} \leq 2M_h e^{-\ell_h} = \delta_h.$$

For $e_h \geq 1$, the error is at most $1/2 \leq e_h$ without sampling. Estimate $\theta_\varnothing$ to accuracy $e_{\mathrm{base}} = e/[8(b+1)]$ with failure at most $\delta/2$ using $\lceil \log(4/\delta)/(2e_{\mathrm{base}}^2) \rceil$ additional responses. Write this estimate as $\widehat{\theta}_\varnothing$ and put

$$\widehat{f}_S = \widehat{\theta}_\varnothing + \sum_{i \in S \cap [k]} (\widehat{\theta}_i - \widehat{\theta}_\varnothing).$$

Let $\Delta_i = \widehat{\theta}_i - \theta_i$ and $\Delta_0 = \widehat{\theta}_\varnothing - \theta_\varnothing$. The exact error and its bound on the common event are

$$\begin{aligned} q^\top(\widehat{f} - \mathcal{A}\theta) &= \left(1 - \sum_{i \leq k} \mu_i(q)\right) \Delta_0 + \sum_{i \leq k} \mu_i(q) \Delta_i - \sum_{i > k} \mu_i(q)(\theta_i - \theta_\varnothing), \\ |q^\top(\widehat{f} - \mathcal{A}\theta)| &\leq (b+1)e_{\mathrm{base}} + \sum_h e_h \sum_{i \in I_h} \mu_i(q) + \sum_{i > k} \mu_i(q) |\theta_i - \theta_\varnothing| \\ &\leq e/8 + \sum_h w_h e_h + e/8 = e/2 \leq e. \end{aligned}$$

Its probability is at least $1 - \delta$ because $\sum_{h \geq 0} \delta_h = \delta/2$.

The number of singleton responses satisfies

$$\sum_h M_h n_h \mathbf{1}_{\{e_h < 1\}} \leq \sum_h M_h \ell_h/(2e_h^2) + k = 8D^3/e^2 + k. \tag{144}$$

For $u = \log(8/\delta)$, the definitions give $M_h \leq 2^h$, $\ell_h \leq u + C(h+1)$ and $w_h \leq C_{b,r} 2^{-dh}$ after a fixed enlargement. Therefore

$$\begin{aligned} D &\leq C_{b,r} \sum_{h \geq 0} 2^{-(2d-1)h/3} [u + h + 1]^{1/3} \\ &\leq C'_{b,r} u^{1/3} \sum_{h \geq 0} 2^{-(2d-1)h/3} (h+2)^{1/3} \leq C''_{b,r} u^{1/3}. \end{aligned}$$

Since $k \leq C_{b,r} e^{-1/d} \leq C_{b,r} e^{-2}$ and the baseline cost is $O_b(e^{-2}u)$, (144) proves the budget. □

For $b = r = 2$, write $\mathcal{Q}_K = \mathcal{Q}_{K,2,2}^{\mathrm{can}}$ and $h_{\mathcal{Q}_K}(v) = \max_{q \in \mathcal{Q}_K} q^\top v$. Let

$$\Gamma_{ij} = \theta_{\{i,j\}} - \theta_{\{i\}} - \theta_{\{j\}} + \theta_\varnothing \quad (i < j), \qquad \Gamma_S = 0 \quad (|S| < 2). \tag{145}$$

Then $\theta = \mathcal{A}\theta + \Gamma$ exactly. The pilot event gives

$$\left| h_{\mathcal{Q}_K}(\theta) - h_{\mathcal{Q}_K}(\widehat{f} + \Gamma) \right| \le \sup_{q \in \mathcal{Q}_K} |q^\top (\mathcal{A}\theta - \widehat{f})| \le e \tag{146}$$

on its common event. The complete domain is unchanged, and $\widehat{f}$ is an acquired background. A fresh query to a uniformly chosen $U \subseteq \{i, j\}$, transformed to $4(-1)^{2-|U|} Z_U$, has conditional mean Γ_{ij} and lies in $[-4, 4]$. It costs one original observation. The remaining critical problem is joint estimation on this same positive domain with its background retained.

A supplied structural certificate $\max_{i<j} |\Gamma_{ij}| \le \eta$ implies $|q^\top \Gamma| \le \eta \sum_{i<j} q_{ij} \le \eta$. The pilot followed by optimization of $\widehat{f}$ estimates each endpoint within $e + \eta$. Padding by $e + \eta$ gives an outer interval with excess at most $2(e + \eta)$. This bound charges acquisition of the unknown additive component. Certifying η from unknown pair records requires additional observations; here it is an input restriction rather than a free empirical test.

Return to fixed $b \ge 1$, $r \ge 2$ and $\mathcal{Q} = \mathcal{Q}^{\mathrm{can}}_{K,b,r}$. The same parametric conclusion also holds for a concrete family/state continuation, rather than only for additive selected-high indicators. Consider the canonical common-overlap geometry with lower keys in index order. Let $v_{i,x} \in [0, 1]$ be unknown means, and suppose the continuation chooses one of the b retained family/state pairs uniformly and returns its fresh loss, without revealing which pair was chosen. For a high record S, let $\mathcal{L}(S)$ be the first $b - |S|$ indices in $[b] \setminus S$. There are enough such indices because $|[b] \setminus S| \ge b - |S|$. Its original response mean is

$$\theta_S = \frac{1}{b} \left(\sum_{i \in S} v_{i,1} + \sum_{i \in \mathcal{L}(S)} v_{i,0} \right). \tag{147}$$

When $K = b$ there are only 2^b records, so direct estimation costs $O_b(\epsilon^{-2} u)$. Suppose $K > b$ and put $S^\circ = [b]$, $c = b + 1$, $m = \theta_{S^\circ}$, and $S^{(i)} = (S^\circ \setminus \{b\}) \cup \{i\}$ for $i > b$. The following contrasts use only original record means:

$$\begin{aligned} &d_b = 0, \qquad d_i = \theta_{S^{(i)}} - m \quad (i > b), \\ &d_i = \theta_{S^{(c)}} - \theta_{(S^\circ \setminus \{i\}) \cup \{c\}} \quad (i < b), \qquad e_i = \theta_{S^\circ \setminus \{i\}} - m \quad (i \le b). \end{aligned} \tag{148}$$

In each swap the unchanged high terms cancel, while $S^\circ \setminus \{i\}$ has the sole low index i. Hence $b d_i = v_{i,1} - v_{b,1}$ and $b e_i = v_{i,0} - v_{i,1}$. Using $|S| + |\mathcal{L}(S)| = b$ and $m - \sum_{j \le b} d_j = v_{b,1}$ in (147) gives

$$\theta_S = m - \sum_{j \le b} d_j + \sum_{i \in S} d_i + \sum_{i \in \mathcal{L}(S)} (d_i + e_i). \tag{149}$$

The terms m, d_i, e_i with $i \le b$ use at most $2b + 1$ anchor records. Estimate each anchor mean within $\zeta = e/[8(1+9b)]$, with total failure at most $\delta/2$. Then $|\widehat{m} - m| \le \zeta$, $|\widehat{d}_i - d_i| \le 2\zeta$ and $|\widehat{e}_i - e_i| \le 2\zeta$ for $i \le b$. For $b < i \le k$, apply the schedule (143)–(144) to the original means $\theta_{S^{(i)}}$ instead of singleton means, assigning failure probability $\delta/2$ to these estimates and setting $\widehat{d}_i = \widehat{\theta}_{S^{(i)}} - \widehat{m}$. The block bounds remain valid after removing $[b]$ from each block. For $i > k$ set $\widehat{d}_i = 0$; choose $k \ge \min(K, b)$ by a fixed increase in the cutoff constant. If $\widehat{\theta}$ is reconstructed from (149), then on the common event, for every $q \in \mathcal{Q}$,

$$\begin{aligned} |q^\top (\widehat{\theta} - \theta)| &\le \zeta + 2b\zeta + 2b\zeta + b\zeta + 4b\zeta + \sum_h w_h e_h + \sum_{i > k} \mu_i(q) |d_i| \\ &\le (1 + 9b)\zeta + e/4 + e/8 \le e/2. \end{aligned}$$

The tail bound uses (143) and $|d_i| \le 1$. The first five terms respectively account for m, the initial sum over $[b]$, the selected high anchors, the shared offset of nonanchor high means, and the low terms. The finite anchor budget is $O_b(e^{-2} u)$, and the other query budget is (144). Hence the uniform functional is acquired in $O_{b,r}(e^{-2} u)$ original responses. Setting all $v_{i,x} = \mu$ gives the matching lower bound. No bank identity or direct bank query is used; the improvement follows from the shared additive mean structure.

Algorithm 1 Finite-precision audit with charged optimization error

Require: Rational keys/moments, fresh-record backend, ϵ, δ.

1: Build the exact prefix atom LP, certify omitted capacity $t_k \le \epsilon/8$, and form its trimmed domain with $\zeta = \epsilon/8$.
2: Obtain block counts/caps and rational factor-two approximations to $a_h = (w_h/M_h)^{1/6}$; use the horizon in Appendix F.
3: Run the sampler for θ' with response Y, and for $1-\theta'$ with response $1-Y$. Compute a feasible rational q_n with certified first-order residual $\xi \le \epsilon/8$, draw the actual rational exploration law p_n, and update $G_n = G_{n-1} + Y_n e_{I_n}/p_{n,I_n}$. The dummy is 0 for θ' and 1 for its complement, and costs no query.
4: Freeze both averages and evaluate the original loss, with dummy 0, on fresh samples to error $\beta = \epsilon/4$; training error is $\alpha = \epsilon/4$.
5: Return $[\widehat{m}_- - \alpha - \zeta - \beta, \widehat{m}_+ + \alpha + \zeta + \beta + t_k] \cap [0,1]$. Use $[0,1]$ only on a charged failure branch or reversed endpoints.

F FINITE-BIT IMPLEMENTATION AND AUDIT CERTIFICATES

The statistical upper bound optimizes over a known domain. For balanced or bounded product moments, rational linear programming implements these calls and supplies checkable primal and dual certificates (Khachiyan, 1979; Schrijver, 1986).

Theorem 26 (Finite-precision faithful audit)**.** *Fix $1 \le b \le K$, $r \ge 2$, $d = \lfloor r/2 \rfloor$, rational $0 < \epsilon < 1$ and $0 < \delta \le 1/16$, encoded in B_{in}. For rational keys, known balanced moments, and fresh Bernoulli or bounded rational record responses, set $k = \min(K, \lceil C_{b,r}\epsilon^{-1/d} \rceil)$. Algorithm 1 has the observation order in* (86) *and arithmetic/random-bit cost*

$$2^{O_{b,r}(k)}\,\mathrm{poly}(K, B_{\mathrm{in}}, B_{\mathrm{resp}}, \epsilon^{-1}, \log(1/\delta)), \tag{150}$$

where B_{resp} bounds response bit length. It returns a faithful interval and rational certificates for all numerical optimization steps. For known rational product biases in $[p_*, 1-p_*]$ with fixed $p_* > 0$, replace $C_{b,r}$ and the time-bound constants by C_{b,r,p_*}. Internal backend generation costs are represented by the separately charged fresh queries.

Proof. Let $n = 2^k$, $B_{S,x} = \prod_{i \in S} x_i$ for $|S| \le \min(r,k)$, and $m_{k,S} = 2^{-|S|}$, including normalization. The one-hot map C assigns each prefix its stopped record or a zero dummy. Restriction and independent extension give exactly

$$\mathcal{P}_k = \{\pi \ge 0 : B\pi = m_k\}, \quad \mathcal{D}_k = C\mathcal{P}_k, \quad \pi \otimes U_{K-k} \in \mathcal{U}_r^K \quad (\pi \in \mathcal{P}_k).$$

Every mixed moment of the extension factors as $\mathbb{E}_\pi X_{S \cap [k]} 2^{-|S \setminus [k]|} = 2^{-|S|}$. In the given upper-key order, b observed highs outrank every unread upper key, so $\{T_k = \perp\} \subseteq \{H_k < b\}$. The dummy is a probability coordinate, not a zero column of C. Its tail can be certified by the degree-$2d$ majorant

$$\mathbf{1}\{H_k < b\} \le \frac{(H_k - k/2)^{2d}}{(k/2 - b + 1)^{2d}}, \qquad \mathbb{E}(H_k - k/2)^{2d} \le 2^{-2d}(2d-1)!!\,k^d.$$

For $k \ge 4(b-1)$, the denominator is at least $(k/4)^{2d}$, so

$$\sup_{\pi \in \mathcal{P}_k} \Pr_\pi(T_k = \perp) \le 2^{2d}(2d-1)!!\,k^{-d}.$$

At $k = K$ there is no dummy. Choose an integer $C_{b,r} \ge \max(1, 4(b-1))$ with $C_{b,r}^d \ge 8\,4^d(2d-1)!!$; the displayed bound then supplies $t_k \le \epsilon/8$ whenever $k < K$. Empty record labels are removed. The uniform prefix center $q^0 = C(\mathbf{1}/n)$ satisfies $q_i^0 \ge 1/n$. Choose capacities at most one (normalization certifies this bound). Rational capacities have certificates

$$B^\top y_h \ge C^\top \mathbf{1}_{B_h}, \quad m_k^\top y_h \le w_h, \qquad w_h \ge M_h/n. \tag{151}$$

For every $\pi \in \mathcal{P}_k$, the certificate gives $\mathbf{1}_{B_h}^\top C\pi \le y_h^\top B\pi = m_k^\top y_h \le w_h$. The lower bound follows by evaluating at the uniform prefix law. If not supplied, obtain these certificates by finite LPs.

For rational $\zeta = \epsilon/8$, use $\mathcal{D}_{k,\zeta} = (1-\zeta)\mathcal{D}_k + \zeta q^0$. Then $q_i \ge \mu := \zeta/n$ and $\mathcal{D}_{k,\zeta} \subseteq \mathcal{D}_k$. For $c \in [0,1]^M$ and $q \in \mathcal{D}_k$, $|c^\top[(1-\zeta)q + \zeta q^0 - q]| = \zeta|c^\top(q^0 - q)| \le \zeta$. Taking extrema bounds the trimming error by ζ. Choose rational a_h between $(w_h/M_h)^{1/6}$ and twice that quantity. For $F(q) = G^\top q + \mathcal{R}(q)/\eta$, let

$$\nabla^2 F(q) = \operatorname{diag}\left(\frac{a_{h(i)}}{4\eta q_i^{3/2}}\right), \quad L_F = \frac{a_{\max}}{4\eta\mu^2}, \qquad \frac{a_{\min}}{4\eta} I \preceq \nabla^2 F(q) \preceq L_F I.$$

These are explicit finite bounds on the entire trimmed domain.

For tolerance $\xi > 0$, set $\Delta = \min(L_F, \xi^2/(16L_F))$. Start at q^0. At iteration j, approximate $\nabla F(q^{(j)})$ by a rational g_j to ℓ_∞ error $\Delta/8$, minimize $g_j^\top z$ over $\mathcal{D}_{k,\zeta}$ by LP, and update

$$q^{(j+1)} = (1-\gamma_j)q^{(j)} + \gamma_j z_j, \qquad \gamma_j = 2/(j+2).$$

Let q^* minimize F and $h_j = F(q^{(j)}) - F(q^*)$. The LP choice, gradient error and convexity give

$$\begin{aligned}\langle \nabla F(q^{(j)}), z_j - q^{(j)}\rangle &\le \langle \nabla F(q^{(j)}), q^* - q^{(j)}\rangle + \|\nabla F(q^{(j)}) - g_j\|_\infty \|z_j - q^*\|_1 \\ &\le -h_j + \Delta/4, \\ h_{j+1} &\le h_j + \gamma_j(-h_j + \Delta/4) + \tfrac{1}{2}L_F\gamma_j^2\|z_j - q^{(j)}\|_2^2 \\ &\le (1-\gamma_j)h_j + \gamma_j\Delta/4 + L_F\gamma_j^2.\end{aligned}$$

Here probability vectors have ℓ_1 distance at most two and squared ℓ_2 distance at most two. Since $\gamma_0 = 1$, $h_1 \le \Delta/4 + L_F$. Induction uses $(j+1)(j+3) \le (j+2)^2$ to obtain $h_j \le \Delta/4 + 4L_F/(j+2)$ for $j \ge 1$. Thus $J = \max(1, \lceil 8L_F/\Delta\rceil)$ iterations give $h_J \le \Delta$. Let $W(q) = \max_z \langle \nabla F(q), q - z\rangle$ and let z attain it. For $q = q^{(J)}$ and $0 < s \le 1$, smoothness gives

$$F(q^*) \le F(q + s(z-q)) \le F(q) - sW(q) + L_F s^2 \quad \Longrightarrow \quad W(q) \le h_J/s + L_F s.$$

With $s = \sqrt{\Delta/L_F}$ this is $W(q^{(J)}) \le 2\sqrt{\Delta L_F} \le \xi/2$.

For a checkable certificate recompute g with error at most $\xi/8$ and an optimal rational dual y satisfying $B^\top y \le C^\top g$. Then

$$W(q) \le g^\top q - \zeta g^\top q^0 - (1-\zeta)m_k^\top y + \xi/4 \le W(q) + \xi/2 \le \xi. \tag{152}$$

Every LP returns a basic feasible prefix law π and dual y; verify $B\pi = m_k$, $\pi \ge 0$, $B^\top y \le C^\top g$, and $g^\top C\pi = m_k^\top y$. These rational checks prove LP optimality. Gradient enclosures use squared rational comparisons for $\sqrt{q_i}$; its positive lower bound μ ensures finite precision. The retained joint witness is the same rational convex combination of LP atoms, mixed with $\zeta\mathbf{1}/n$; its feasibility is checked exactly.

For $f_i = q_i^{3/4}/\sqrt{a_{h(i)}}$ and $s_i = w_{h(i)}/\sqrt{M_{h(i)}}$, choose positive dyadic $\widetilde{f}_i \in [f_i, 2f_i]$ and $\widetilde{s}_i \in [s_i, 2s_i]$. Fourth and second powers certify these intervals. Normalize them and average to obtain the actual rational sampling law p. The upper bounds on the denominators give

$$p_i = \frac{1}{2}\left(\frac{\widetilde{f}_i}{\sum_j \widetilde{f}_j} + \frac{\widetilde{s}_i}{\sum_j \widetilde{s}_j}\right) \ge \frac{1}{4}(p_i^A + p_i^0).$$

Use this actual p in $\widehat{g} = Ye_I/p_I$: $\mathbb{E}[\widehat{g}_i \mid \mathcal{F}] = p_i c_i/p_i = c_i$. For $d_n = q_n - q^*$, the stability and martingale bounds from Lemma 16 become

$$\frac{Y^2 q_I^{3/2}}{a_{h(I)} p_I^2} \le 16V, \qquad \sum_i \frac{d_{n,i}^2}{p_i} \le 4\mathsf{H}\sum_h \frac{\sqrt{M_h}}{w_h}\, 2w_h^2 = 8\mathsf{H}^2.$$

The corresponding range is $B_* = 1 + 4\mathsf{H}\sqrt{M_{\max}}$. The first-order residual adds $N\xi$ pathwise to the regret bound $S/\eta + 32\eta VN$. Freedman therefore gives

$$c^\top \bar{q} - \min_{\mathcal{D}_{k,\zeta}} c^\top q \le S/(\eta N) + 32\eta V + \xi + 4\mathsf{H}\sqrt{u/N} + 2B_* u/(3N). \tag{153}$$

Take rational upper enclosures $S \le \bar{S} \le 2S$, $V \le \bar{V} \le 2V$, $\mathsf{H} \le \overline{\mathsf{H}} \le 2\mathsf{H}$, $B_* \le \bar{B} \le 2B_*$. A rational $\eta \in [\sqrt{\bar{S}/(32\bar{V}N)}, 2\sqrt{\bar{S}/(32\bar{V}N)}]$ suffices. For $t = \epsilon/8$, take

$$N = \left\lceil \max\left\{1800\bar{S}\bar{V}/t^2,\ 144\overline{\mathsf{H}}^2 u/t^2,\ 2\bar{B}u/t\right\}\right\rceil, \qquad \xi = \epsilon/8.$$

Writing $\eta = x\sqrt{\bar{S}/(32\bar{V}N)}$, $1 \le x \le 2$, gives

$$\frac{S}{\eta N} + 32\eta V \le \sqrt{\frac{32\bar{S}\bar{V}}{N}}(x^{-1} + x) \le \frac{5}{2}\sqrt{\frac{32\bar{S}\bar{V}}{N}} \le t/3.$$

The other two nonresidual terms are at most $t/3$ each. Thus $\alpha = t + \xi = \epsilon/4$. For the rational weights, $S \le 2Z$ and $\sqrt{V} \le Z$, hence $SV \le 2Z^3$. Use costs $(\theta_{\mathrm{real}}, 0)$ and $(1 - \theta_{\mathrm{real}}, 1)$, including the known dummy complement. Freeze their averages and purchase $n_e = \lceil u/(2\beta^2)\rceil$ evaluations each in the original loss, where the dummy is zero, $\beta = \epsilon/4$. Return

$$[\widehat{m}_- - \alpha - \zeta - \beta, \widehat{m}_+ + \alpha + \zeta + \beta + t_k] \cap [0, 1]. \tag{154}$$

Trimming enlarges the feasible optimization gap from α to $\alpha + \zeta$. Truncation gives $v_\pm \le R_\pm \le v_\pm + t_k$, because the original loss is nonnegative. Lemma 13 therefore bounds both excesses by $\alpha + \zeta + 2\beta + t_k \le \epsilon$. Reversed padded endpoints can occur only outside this success event; on that branch return $[0, 1]$. Training and evaluation count every repeated real query; the dummy is free.

The bounds $M \le n$, $w_h \ge M_h/n$, $a_{\min} \ge 1/n$ and $a_{\max} \le 2$ give $S \ge 1/n$ and $V \le n(\sum_h M_h^{1/4})^2 \le n^3$. Hence

$$\eta^{-1} \le \sqrt{32\bar{V}N/\bar{S}} \le 8n^2\sqrt{N}, \qquad L_F = \frac{a_{\max}n^2}{4\eta\zeta^2} \le 4n^4\sqrt{N}\zeta^{-2}.$$

Thus J is polynomial in n, N, ϵ^{-1}. Each LP uses the fixed $0/1$ prefix moment matrix; its basic primal weights have $O(r + D_r(k)\log D_r(k))$ bits by determinant bounds. Convex combinations over J iterations preserve polynomial bit length. Use the fixed mesh $\rho = 2^{-\lceil \log_2(4/\mu)\rceil}$ and round f_i, s_i upward by integer squared/fourth-power comparisons. If their integer weights are u_i, v_i, the actual probability is $(u_i \sum_j v_j + v_i \sum_j u_j)/(2\sum_j u_j \sum_j v_j)$. Its bit length does not inherit previous mixture denominators. The FW average has exact weights $2(j+1)/(J(J+1))$ on its J basic vertices. Their denominator lengths add within one run. Each new run uses the same fixed moment matrix, so previous mixture denominators do not enter its basic vertices. Since $p_i \ge 1/(4n^2)$, $\|G_n\|_\infty \le 4n^2N$, and accumulation of bounded-bit responses also has polynomial bit length. Each gradient enclosure and LP dual certificate consequently has polynomial bit cost in the same parameters.

A rational draw uses a common denominator D, fair random bits to sample an integer in $[0, 2^{\lceil \log_2 D\rceil})$, and rejection above D. Each attempt succeeds with probability at least $1/2$. Over at most $R = 2N + 2n_e$ draws, cap each at $\lceil \log_2(2R/\delta)\rceil$ attempts; on exhaustion return $[0, 1]$. Coupling with unlimited exact draws gives $\Pr(\text{some exhaustion}) \le R\,2^{-\lceil \log_2(2R/\delta)\rceil} \le \delta/2$. Choose the integer $u_{\mathrm{num}} = \lceil \log_2(16/\delta)\rceil$ in place of u in the algorithm; the four concentration failures total at most $6e^{-u_{\mathrm{num}}} < \delta/2$. This proves both deterministic termination and the faithful guarantee. As $n = 2^k$ and all other quantities are polynomial in $n, K, B_{\mathrm{in}}, B_{\mathrm{resp}}, \epsilon^{-1}, u$, the claimed bit-work bound follows.

For rational product biases $p_i \in [p_*, 1 - p_*]$, set $\lambda_k = \prod_{i \le k} \min(p_i, 1 - p_i) \ge p_*^k$ and use the product prefix center. Replace the balanced numerical floors throughout by $q_i^0 \ge \lambda_k$, $w_h \ge M_h\lambda_k$, $\mu = \zeta\lambda_k$, $p_{n,i} \ge \lambda_k/(4n)$ and $\|G_N\|_\infty \le 4nN/\lambda_k$. The basic-weight bit bound additionally includes the denominators of all supplied prefix biases. For $\nu_k = \sum_{i \le k} p_i$, use the majorant $(H_k - \nu_k)^{2d}/(\nu_k - b + 1)^{2d}$. In the centered product expansion, a term with an index appearing once has expectation zero. Every surviving term uses at most d distinct indices and has absolute expectation at most one. Thus

$$\mathbb{E}(H_k - \nu_k)^{2d} \le C_d \sum_{j=1}^{d} \binom{k}{j} \le C_d' k^d.$$

Enlarge the cutoff constant so that $k < K$ implies $k \ge 2(b-1)/p_*$. Then $\nu_k - b + 1 \ge p_* k/2$, giving the $O_{b,r,p_*}(k^{-d})$ tail. Mixed prefix–tail moments extend independently. All inverse floors remain $2^{O_{p_*}(k)}$, proving the stated extension. If each backend response has an additional worst-case generation time T_{resp}, add $(2N + 2n_e)T_{\mathrm{resp}}$ to the auditor's bit-work bound. □

For a fully known bounded rational loss, two exact prefix LPs and the same tail certificate already compute endpoints to tolerance ϵ in the time (150), without statistical queries. For a general correlated moment release, prefix extension must be certified separately; the next composition result explicitly charges any outer-domain error. A known physical-routing depth profile likewise does not by itself supply a small prefix representation.

F.1 Combined moment, representation and observation errors

Let $\varphi : \Omega \to \mathbb{R}^s$ be the finite supplied feature map, $\mathcal{C} = \operatorname{conv} \varphi(\Omega)$, and $\mathcal{P}(m) = \{P \in \Delta(\Omega) : \mathbb{E}_P \varphi = m\}$. A known stochastic map B from assignments to observed banks gives $\mathcal{D}(m) = B^\top \mathcal{P}(m)$. In the selector model, $B_{xj} = A_{T(x),j}$; direct record acquisition uses $A = I$. For compact sets in the same probability simplex define

$$h_{\to}(F, G) = \sup_{z \in F} \inf_{z' \in G} \mathrm{TV}(z, z').$$

Proposition 27 (Conditional composition of a faithful audit)**.** *A calibration pilot returns a convex compact set $\mathcal{M}$ with $\Pr(m \in \mathcal{M}) \geq 1 - \delta_m$ and $\mathcal{M} \cap \mathcal{C} \neq \varnothing$. Let*

$$\mathcal{D}_{\mathcal{M}} = \bigcup_{m' \in \mathcal{M} \cap \mathcal{C}} \mathcal{D}(m'), \qquad \overline{\mathcal{D}} \supseteq \mathcal{D}_{\mathcal{M}},$$

where the outer domain is compact and convex. Suppose certificates give

$$h_{\to}(\overline{\mathcal{D}}, \mathcal{D}_{\mathcal{M}}) \leq \gamma_{\mathrm{rep}}, \qquad \sup_{m' \in \mathcal{M} \cap \mathcal{C}} h_{\to}(\mathcal{D}_{\mathcal{M}}, \mathcal{D}(m')) \leq \gamma_m.$$

Conditionally on every pilot history, use fresh continuation observations to obtain an (ϵ_v, δ_v)-faithful interval for $\overline{\mathcal{D}}$. The same output is

$$(\epsilon_v + \gamma_{\mathrm{rep}} + \gamma_m,\ \delta_v + \delta_m)\text{-faithful for } I(m, v). \tag{155}$$

In particular, the finite-bit procedure may allocate $\epsilon_v = \alpha + \zeta + 2\beta + t_k$ as in (154)*; its numerical residual is included in α. Its prefix representation and omitted-capacity certificate must be valid for the entire working domain used after the pilot.*

Proof. The union is the image of the compact convex set $\{P \in \Delta(\Omega) : \mathbb{E}_P \varphi \in \mathcal{M}\}$. On $m \in \mathcal{M}$, every $z \in \overline{\mathcal{D}}$ admits $z' \in \mathcal{D}_{\mathcal{M}}$ and $z'' \in \mathcal{D}(m)$ with

$$\mathrm{TV}(z, z'') \leq \mathrm{TV}(z, z') + \mathrm{TV}(z', z'') \leq \gamma_{\mathrm{rep}} + \gamma_m =: \gamma.$$

For $0 \leq v \leq 1$, $|v^\top (z - z'')| \leq \mathrm{TV}(z, z'')$. Taking extrema and using $\mathcal{D}(m) \subseteq \mathcal{D}_{\mathcal{M}} \subseteq \overline{\mathcal{D}}$ gives

$$0 \leq R_-(m, v) - R_-(\overline{\mathcal{D}}, v) \leq \gamma_{\mathrm{rep}} + \gamma_m,$$
$$0 \leq R_+(\overline{\mathcal{D}}, v) - R_+(m, v) \leq \gamma_{\mathrm{rep}} + \gamma_m.$$

For a successful continuation interval $[L, U]$, these inequalities give

$$L \leq R_-(\overline{\mathcal{D}}, v) \leq R_-(m, v), \qquad U \geq R_+(\overline{\mathcal{D}}, v) \geq R_+(m, v),$$
$$R_-(m, v) - L \leq \epsilon_v + \gamma, \qquad U - R_+(m, v) \leq \epsilon_v + \gamma.$$

If $\mathcal{H}_0$ is the pilot history, the total failure probability is at most

$$\Pr(m \notin \mathcal{M}) + \mathbb{E}\big[\mathbf{1}_{\{m \in \mathcal{M}\}} \Pr(\text{continuation failure} \mid \mathcal{H}_0)\big] \leq \delta_m + \delta_v.$$

The fresh conditional mean contract includes the entire pilot history. If error certificates depend on that history, a prescribed tolerance requires their sum to be within its budget on the claimed success event. No sampling law is conditioned on coverage instead of the actual history. □

The directed defect is the exact worst upper-endpoint enlargement over bounded losses, rather than only a norm sufficient condition:

$$h_{\to}(F, \mathcal{D}(m)) = \sup_{0 \leq v \leq 1} \{\sup_{z \in F} v^\top z - R_+(m, v)\}. \tag{156}$$

For fixed z, both optimization sets below are compact and convex, and the objective is bilinear. Finite convex minimax gives

$$\min_{z' \in \mathcal{D}(m)} \mathrm{TV}(z, z') = \min_{z' \in \mathcal{D}(m)} \max_{v \in [0,1]^J} v^\top (z - z') \\ = \max_{v \in [0,1]^J} \min_{z' \in \mathcal{D}(m)} v^\top (z - z') = \max_{v \in [0,1]^J} \{v^\top z - R_+(m, v)\}.$$

Taking $\sup_{z \in F}$ and interchanging the two suprema proves (156). Replacing v by $1 - v$ controls the lower endpoint.

Set $E = \mathrm{span}(\mathcal{C} - \mathcal{C})$. Suppose a feasible center $\widehat{m}$ has the certified relative moment ball

$$\widehat{m} + \{u \in E : \|u\|_\infty \le r_0\} \subseteq \mathcal{C},$$

and $\mathcal{M} \cap \mathcal{C}$ lies within distance $\eta < r_0$ of $\widehat{m}$. Then

$$\gamma_m \le \frac{2\eta}{r_0 + \eta}. \tag{157}$$

For $m \in \mathcal{M} \cap \mathcal{C}$ and $u \in E$ with $\|u\|_\infty \le r_0 - \eta$, $\|m + u - \widehat{m}\|_\infty \le \eta + (r_0 - \eta) = r_0$. Thus m has a feasible relative ball of radius $r_0 - \eta$. For a law $P' \in \mathcal{P}(m')$, put $s = \|m' - m\|_\infty \le 2\eta$ and choose P^c at the feasible moment $m - (r_0 - \eta)(m' - m)/s$. If $s > 0$,

$$P = \frac{r_0 - \eta}{r_0 - \eta + s} P' + \frac{s}{r_0 - \eta + s} P^c \in \mathcal{P}(m), \qquad \mathrm{TV}(P', P) \le \frac{s}{r_0 - \eta + s} \le \frac{2\eta}{r_0 + \eta}.$$

Indeed the numerator of its moment is $(r_0 - \eta)m' + sm - (r_0 - \eta)(m' - m) = (r_0 - \eta + s)m$. Stochastic pushforward contracts TV, giving (157). For $s = 0$, use $P = P'$.

A supplied full-support law P^0 at $\widehat{m}$ gives a finite radius certificate. Find a rational linear map $G : E \to \mathbb{R}^\Omega$ with $\mathbf{1}^\top G u = 0$ and $\Phi G u = u$ on E, where $\Phi_{i,x} = \varphi_i(x)$. Then

$$r_0^{-1} = \max_x \max_{u \in E, \|u\|_\infty \le 1} |(Gu)_x| / P^0_x$$

is checked by two rational LPs per assignment. For $\|u\|_\infty \le r_0$,

$$|(Gu)_x| \le P^0_x, \qquad \mathbf{1}^\top (P^0 + Gu) = 1, \qquad \Phi(P^0 + Gu) = \widehat{m} + u.$$

Hence $P^0 + Gu$ certifies the whole relative ball. For rational features and P^0, this uses polynomial work in the explicit assignment representation, which can be exponential in K. When $E = \{0\}$, all feasible moments coincide and $\gamma_m = 0$.

For a rational moment box centered at $\widehat{m}$ with radius η, optimization over all compatible laws has the finite upper dual

$$\min_{a,y} \{a + y^\top \widehat{m} + \eta \|y\|_1 : a + y^\top \varphi(x) \ge (Bv)_x \quad \forall x \in \Omega\}. \tag{158}$$

To see the objective, attach nonnegative multipliers y^+, y^- to the upper and lower moment inequalities. Their contribution is

$$(y^+ - y^-)^\top \widehat{m} + \eta \mathbf{1}^\top (y^+ + y^-).$$

For fixed $y = y^+ - y^-$, minimizing over this decomposition gives $y^\top \widehat{m} + \eta \|y\|_1$. The normalization multiplier is a; finite LP duality gives (158). For a known rational objective, including the actual gradient approximation used by the numerical algorithm, feasible joint atoms and rational pointwise checks certify primal values, the majorant and their gap. Unknown true loss values are controlled by conditional observations, not treated as known rational LP inputs. For a supplied polyhedral outer domain, a finite relaxation-defect certificate lists, for each of its vertices, a compatible law within the claimed TV tolerance; convexity extends the bound to the domain. Neither certificate's discovery nor its length is assumed polynomial without a specified representation. Theorem 26 provides that representation for known product moments.

For known m, take $\mathcal{M} = \{m\}$ and omit calibration; for known v, omit continuation acquisition. When both quantities are unknown, their observation fees add and their errors compose by (155). The shared-count necessity example is kept with its full model and proof in the technical notes.

Table 2: Risk contrasts in percentage points. The format diagnostic is post-outcome for the main cohort. External intervals are descriptive; the external cohort changes both generator and questions.

Condition	n	Mean $\lvert\Delta\rvert$	Two-sided 95% interval
Main, frozen oracle	512	1.465	[0.957, 1.972]
Main, format diagnostic	512	1.416	[0.916, 1.916]
External, 7B generator	64	0.391	[0.000, 1.187]

G Experimental protocol and finite identification

This appendix specifies the experiment supporting Section 7. Mathematical certificate checks, retrospective calculations, and fresh-to-study question cohorts are reported separately.

G.1 Fixed trees, scoring, and observations

The main generator is Qwen2.5-Math-1.5B-Instruct, revision `aafeb0fc6f22cbf0eaeed126eff8be45b0360a35`. Sampling uses physical batch one, temperature 0.8, top-$p = 1$, top-$k = 50$, and the official chain-of-thought instruction. The first three segments have 256-token caps and stop at the first double newline; the final segment has a 1,024-token cap. The potential branching is $(3, 3, 3, 1)$ and beam widths are $(2, 2, 2, 1)$. Generation seeds are deterministic hashes of question identity and node path. EOS is absorbing and preserves path identity and score; there are no invented descendants or repeated score increments. Registered failures and capped generations remain in the evaluation.

Both scoring panels use the frozen FP32/eager/no-TF32 function. Qwen reads the final process marker's positive class; Llama uses its alternating user/assistant-plus renderer and FP32 two-tag readout. Probabilities are clipped to $[10^{-6}, 1 - 10^{-6}]$ before taking logarithms and summed along nonempty steps. Neutral EOS-only transitions incur no new PRM evaluation. A same-valued parameter-storage adapter matched the canonical scorer on 32 registered prefixes per model before scoring resumed. This changed storage, not the scoring function.

Question lists and 32 balanced batches are fixed without labels. Generation resumes according to resources and file integrity, not correctness. Main correctness is read after all 512 trees, both full panels, and the original independent validator have completed. Independent final reconstruction checks 8,192 beam/one-shot assignment outcomes and their terminal-label bindings. The one-shot comparator scores every potential terminal in the same fixed pool; it generally costs more than beam search and is not an equal-compute accuracy comparator.

The external condition fixes 32 MATH and 32 GSM8K questions excluded from identifiable earlier study and reservation lists. It uses Qwen2.5-Math-7B-Instruct revision `ef9926d75ab1d54532f6a30dd5e760355eb9aa4d`, retaining the selector and PRMs. A four-question label-blind check is reused byte-for-byte and counted once. The external labels are opened only after complete validation. The exclusion is defined by the study's question lists; model pretraining and verifier-training corpora have separate provenance.

G.2 Risk contrasts, intervals, and oracle sensitivity

Table 2 reports equal-task means and approximate stratified Welch intervals. The registered materiality criterion is a one-sided 95% lower bound above one percentage point (pp) for $\mathbb{E}|\Delta|$. The frozen-oracle lower bound is 1.039 pp, but a post-outcome format-equivalence diagnostic lowers it to 0.997 pp. Thus the narrow materiality decision is not robust to that diagnostic. The frozen signed interval is $[-0.624, 0.429]$ pp and includes zero. Neither the original nor diagnostic analysis establishes a better assignment law. This sensitivity illustrates the distinction between sampling uncertainty and truth identification in Section 6.

Structural and risk endpoints differ: 64/512 main questions have $C_{\text{out}} > 1/4$, 81/512 have zero root capacity, and 335/512 have changed text distributions with zero error contrast. The independent external condition uses 32 new questions per task and a fixed 7B generator. Its signed contrast is

Table 3: The separate registered structural family in the main cohort. B1 and B2 use adjusted one-sided lower bounds; B3 uses an adjusted two-sided interval.

Endpoint	Estimate	Registered bound/interval
B1: $C_{\mathrm{out}} > 1/4$ (%)	12.500	lower 8.057
B2: mean $(C_{\mathrm{out}} - 1/4)_+$ (pp)	3.418	lower 2.532
B3: $c_{\mathrm{root}} = 0$ (%)	15.820	[10.337, 22.677]

-0.391 pp with interval $[-1.187, 0.406]$ pp. It is not pooled with the main cohort or interpreted as a causal model-size comparison.

For task $h \in \{1, 2\}$ with n_h questions, sample mean $\bar{x}_h$ and sample variance s_h^2, the equal-task estimate and variance are

$$\widehat{\mu} = \frac{\bar{x}_1 + \bar{x}_2}{2}, \qquad v = \sum_{h=1}^{2} \frac{s_h^2}{4n_h}, \qquad \nu = \frac{v^2}{\sum_{h=1}^{2} (s_h^2/(4n_h))^2/(n_h - 1)}.$$

Registered intervals use the corresponding Student quantile and are clipped to the mathematical support. These Welch intervals are approximate. At zero empirical variance, a bounded-variable Hoeffding fallback avoids a zero-width population interval. A 20,000-replicate question-level stratified bootstrap, seed 593091, is a sensitivity analysis; its main absolute-contrast interval is $[0.977, 2.002]$ pp.

The risk materiality test and the three-endpoint structural family are separate. The latter uses $\alpha = 0.05/3$ and within-task Clopper–Pearson intervals for binary endpoints, with the error budget divided between tasks. There is no joint either-one-passes claim. The main positive/negative/zero contrast counts are $14/16/482$; the external counts are $0/1/63$.

The format diagnostic was prompted by an answer receipt: `\text{C, F, H}` was parsed as a symbol, while `C, F, H` was parsed as a set and judged wrong. The question asks for comma-separated letters, and exact rational slopes independently confirm C, F, H. The diagnostic removes a whole-list font wrapper and whitespace from comma-separated single uppercase letters, preserving order. Applied to all main labels, it changes 14 terminal labels on one question and makes that question's contrast zero. The original oracle and registered analysis are preserved. The rule was fixed before external execution and changes no external question. It is neither a replacement confirmatory test nor a complete human audit of oracle accuracy.

G.3 Sufficient-state auditing and its comparator

A merge key contains active paths, required family states, cumulative scores, and generator identities. Equal visible text is insufficient because different path seeds can have different continuations. PRM input caching instead uses model identity and exact input tokens. The full-replay comparator takes the union of queries from all eight assignments with this same input cache; it does not sum uncached runs.

Every live block starts with an empty cache and uses the same resident models and trees. In AB order, replay and audit take 112.100366 and 75.765306 seconds; in BA order, 112.883933 and 75.828746 seconds. The saved unrounded audit times sum to 151.59405142394826 seconds, which rounds to the 151.59 seconds in Table 4; summing the previously rounded block times would instead give 151.60. Common loading and hashing take 55.11 seconds and warmup 0.42 seconds. The benchmark makes 2,260 neural forwards including warmup and no generation calls. All actual probabilities equal the frozen panel, and all contrast coefficients equal independent replay. The table reports the two observed paired timings on shared hardware. A fixed four-question implementation check precedes the 16-question benchmark.

Full-tree labeling covers 13,310 terminal texts in the main cohort. Selected-text deduplication requires 1,812 labels and nonzero contrast support requires 852 on this already exposed table. These are exact requirements for the saved finite loss table. Their units are label counts; the runtime measurements above concern the separate replay benchmark.

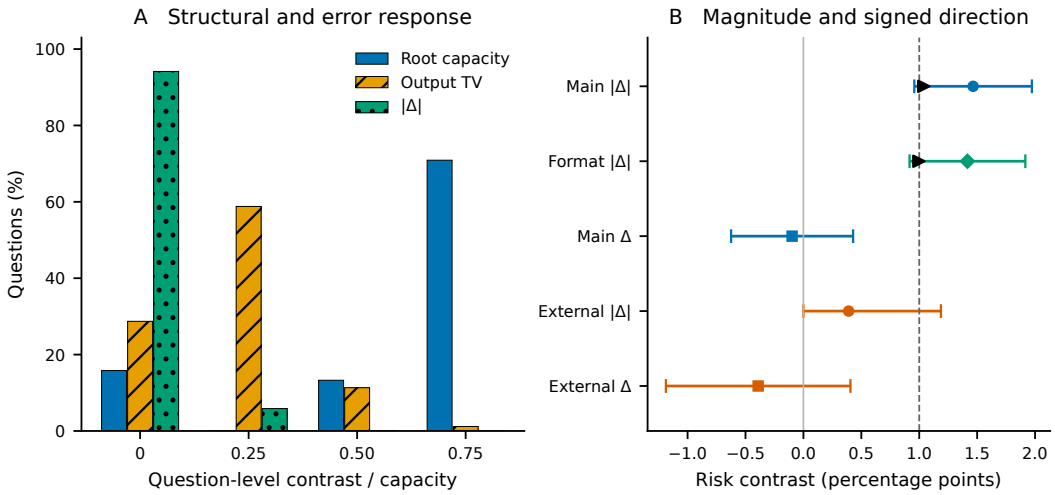

Figure 2: Frozen main-cohort distributions and interval sensitivity. Left: empirical distributions across 512 questions; bars are not population-error intervals. Right: contrasts with two-sided Welch intervals; triangles mark the original and post-outcome diagnostic one-sided lower bounds, and the dashed line is one pp.

Table 4: Same-cache audit on sixteen fixed questions. Counts are per block; times combine the two execution orders. These are fixed-protocol measurements on shared hardware.

Quantity	Full replay	Signed audit	Reduction
Neural forward calls	662	467	29.46%
PRM input tokens	199,328	128,858	35.35%
Total time (seconds)	224.98	151.59	32.62%

G.4 Finite certificates and identification calculations

For prescribed pairwise completion, 48 controlled instances with $3 \leq K \leq 10$, widths two to four, nonuniform correlated moments, heterogeneous keys, and endpoint ties produce 96 rational primal/dual endpoint certificates. A second implementation validates them without importing the solver or calling an LP optimizer; four corrupted controls are rejected. A complementary implementation adds 56 endpoints on 28 instances through $K = 12$. The fixed-geometry information profile is checked on 105 root tables and 504 profile cells, separately from 138 attaining settings. These are finite checks of implementation and constructions, not new empirical samples or substitutes for the proofs.

Retrospective calculations. The earlier 1,024-question BF16-conditioned development panel is distinct from the main and external cohorts. In that panel, 417 questions have changed text distributions and zero risk contrast; nonzero contrast support requires 1,203 labels versus 3,126 distinct selected texts. With parsing-aware answer equivalence, the stronger comparison is 758 versus 2,134 answer comparisons, with parsing cost still present. Development repair attempts did not establish reliable held-out accuracy improvement.

Optimizing the saved risk tables under progressively richer fair assignment moments gives Table 1. The complete-joint column fixes the uniform cube law $P(x) = 1/8$; it does not fix either of the two parity interventions. For three bits, fair pairwise laws form the parity line. Each question's width therefore equals its previously measured absolute parity contrast. This is a recalculation of the same risk effect. The main minimum loss-identification orders $0, 1, 2, 3$ occur on 453, 7, 22, and 30 questions, respectively. In the earlier panel, independently optimized stagewise terms are strictly looser than a common-world bound on 227 singleton and 121 pairwise instances, with mean excess

Table 5: Examples of the cost of replacing one common law by local conditional consistency. All methods in a row use identical fair pairwise moments, scores, and fixed losses. Exact endpoints have global probability and all-world dual certificates.

Instance	Pair-local upper	Pair and all triples	Exact interval
$K = 5, b = 2$	$9/10$	$25/28$	$[1/2, 25/28]$
$K = 6, b = 2$	$31/48$	$19/30$	$[1/4, 5/8]$

widths 13.01 and 4.13 pp. These calculations motivate maintaining one compatible law; they do not measure generalization or an additional natural-data effect.

Finite candidate counts and truncation. For $K = 32$, the canonical top-one, top-two, and top-three selectors have $\sum_{j=0}^{b} \binom{32}{j} = 33, 529$, and $5{,}489$ records, respectively. Under balanced pairwise moments, the prefix count $S_k = \sum_{i \le k} X_i$ has mean $k/2$ and variance $k/4$. For $k > 2b-2$, Cantelli's inequality therefore gives

$$P(S_k < b) = P(S_k \le b - 1) \le \frac{k/4}{k/4 + (k/2 - b + 1)^2} = \frac{k}{k + (k - 2b + 2)^2}.$$

For top-two at $k = 8, 16, 32$, these bounds are $2/11$, $4/53$, and $8/233$. At $k = K$ every assignment determines a full retained record, so the unobserved record tail is exactly zero even if $P(S_K < b) > 0$. This distinguishes the count event from the truncation error in the finite-K acquisition theorem. The counts and tail bounds are algebraic certificates, not measured fresh-response query budgets.

G.5 Common-law intervals for finite panels

This subsection derives the common-law intervals in Table 1, together with their nonfair and finite-batch extensions. It uses the parity representation of the compatible laws and linear optimization of the fixed loss tables.

Fix question weights $\pi_i \ge 0$, $\sum_i \pi_i = 1$, and saved losses $f_i(x) = \theta_{i,T_i(x)} \in [0, 1]$. The assignment bits retain the same verifier meanings across questions. Write

$$\bar{f}(x) = \sum_i \pi_i f_i(x), \qquad V(P) = \sum_x P(x) \bar{f}(x).$$

One law P is used independently of the question. For a sealed batch with empirical law P_N, this is the standardized risk obtained by independently pairing that batch with the fixed panel. Historical question-dependent routing has a different estimand.

For $K = r+1$, let P_0 be any distribution with the supplied proper marginals, and set $\chi(x) = (-1)^{|x|}$. For $d = P - P_0$, equality of the proper marginals implies

$$\begin{aligned} 0 &= \sum_{x_j \in \{0,1\}} d(x) = d(x) + d(x \oplus e_j), \\ d(x) &= (-1)^{|x|} d(0) = d(0)\chi(x), \\ \sum_{x_j \in \{0,1\}} \chi(x) &= 0. \end{aligned}$$

The second line follows by flipping the nonzero coordinates of x. The third verifies the converse. Nonnegativity now gives

$$\begin{aligned} P_\tau = P_0 + \tau\chi \ge 0 \quad &\iff \quad \tau \ge -P_0(x)\ (\chi(x) = 1), \quad \tau \le P_0(x)\ (\chi(x) = -1) \\ &\iff \quad -a \le \tau \le b_0, \\ a = \min_{\chi=1} P_0(x), \quad b_0 &= \min_{\chi=-1} P_0(x). \end{aligned}$$

This argument includes arbitrary compatible nonfair moments. With $\mu_i = P_0 \cdot f_i$, $D_i = \chi \cdot f_i$, $\bar{\mu} = \sum_i \pi_i \mu_i$, and $\bar{D} = \sum_i \pi_i D_i$,

$$V(P_\tau) = \sum_i \pi_i \sum_x (P_0(x) + \tau\chi(x)) f_i(x) = \bar{\mu} + \tau\bar{D}.$$

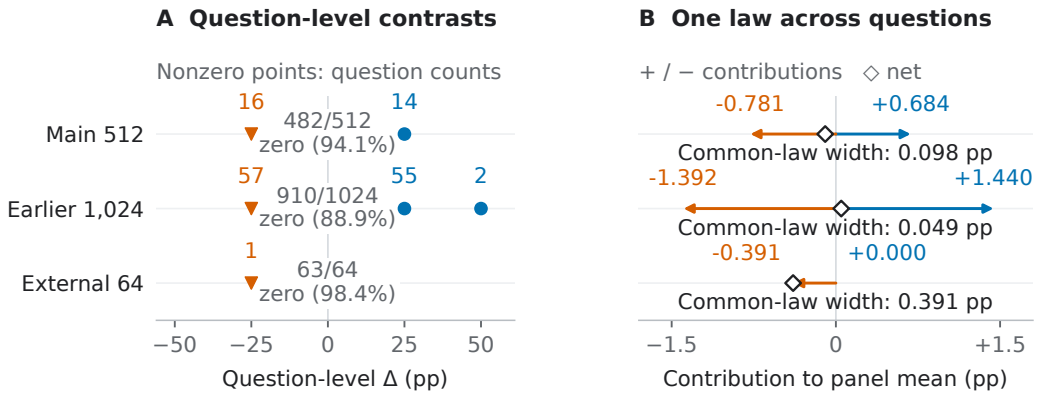

Figure 3: How per-question ambiguity aggregates across a fixed panel. (A) Nonzero even–odd contrasts, with question counts above the markers; gray labels give the zero-contrast counts and percentages. (B) Positive and negative contributions to the panel mean. The diamond marks their signed sum; its absolute value is the common-law width. The sum of their magnitudes is the mean question-dependent width. These finite-table quantities keep cohorts separate and use no fresh samples.

Minimizing and maximizing this affine function on $[-a, b_0]$ yields

$$I_{\text{common}} = \bar{\mu} + [\min\{-a\bar{D}, b_0\bar{D}\}, \max\{-a\bar{D}, b_0\bar{D}\}],$$
$$W_{\text{common}} = (a+b_0)|\bar{D}|, \qquad W_{\text{group}} = (a+b_0)\sum_g \left|\sum_{i\in g}\pi_i D_i\right|, \tag{159}$$

For the displayed group formula, each group is separately constrained by the same supplied proper marginals; an overall aggregate report alone does not imply this groupwise information. Taking singleton groups under that contract gives the original mean width. Writing $D_g = \sum_{i\in g}\pi_i D_i$, each group chooses $\tau_g \in \{-a, b_0\}$ according to the sign of D_g. For any partition refinement $g = \bigcup_{h\subseteq g} h$,

$$W_{\text{common}} = (a+b_0)\left|\sum_g D_g\right| \le (a+b_0)\sum_g |D_g| \le (a+b_0)\sum_g\sum_{h\subseteq g}|D_h|.$$

The choice between common and group-specific laws is part of the release contract. Observed cancellation does not determine that choice.

For exact counts, start from an integer feasible table n_0. The same calculation gives

$$n = n_0 + t\chi, \qquad -\min_{\chi=1} n_0 \le t \le \min_{\chi=-1} n_0, \qquad n \in \mathbb{Z}^{2^K} \iff t \in \mathbb{Z}.$$

Both extreme values of t are integers. They therefore attain the two linear-risk extrema with actual finite batches. Intermediate risks can be discrete. For general $K > r+1$, integer attainability requires a separate check.

In the fair three-bit case, $\Delta_i = \mathbb{E}_{\text{even}} f_i - \mathbb{E}_{\text{odd}} f_i = D_i/4$, so (159) gives the widths in Section 7. The common-law intervals on the main, earlier, and external panels are $[313, 315]/2048$, $[1424, 1425]/2048$, and $[29, 30]/256$. Their original question-dependent widths remain $15/1024$, $29/1024$, and $1/256$. On the earlier panel, allowing one law for each of the two prespecified tasks gives the intermediate width $1/256$. The existing main format diagnostic changes the common-law interval to $[311, 314]/2048$ and the question-dependent width to $29/2048$; the external diagnostic leaves both unchanged. All values are rational calculations on the saved outcomes. Figure 3 displays the corresponding sign distribution and its contribution to the common-law width. Population signed intervals and the diagnostic's post-outcome status are reported above.