CERTIFYING SHARED GUMBEL COUPLINGS: GEOMETRY, QUERY COMPLEXITY, AND MARGINAL ERROR

Anonymous authors
Paper under double-blind review

ABSTRACT

Correct categorical marginals do not determine how a sampler couples outputs across inputs. We study certification of the shared Gumbel coupling through categorical queries to reusable hidden seeds. Certificates must be valid for every sampler satisfying a stated marginal promise and succeed efficiently under the Gumbel reference. We characterize exactly when a fixed finite input family permits inverse-accuracy certification: for full-support targets p, q on $n \geq 3$ labels, the ordered ratios $r_i = q_i/p_i$ must satisfy $r_{(2)} = r_{(n-1)}$. At most three inputs then suffice. Outside this class, every fixed finite family has inverse-square complexity, but accuracy-dependent queries attain a faster rate. For a fixed alphabet of at least four labels and distinct targets bounded away from zero, we prove the uniform optimal order $\log(1/\delta)[\rho^{-1} + D\rho^{-4/3}]$, where $D = \log(r_{(n-1)}/r_{(2)})$. The order holds simultaneously for reference expected seeds, reference expected calls, and a deterministic call budget valid under every sampler; lower bounds allow adaptive queries and stopping. We also determine the tolerance transition for fixed nonexceptional targets. For a structured critical-tolerance family, weighted orbit arguments yield stronger certificate bounds and an $\epsilon^{-11/7}$ score-energy lower bound, while leaving the sharp critical law open.

1 INTRODUCTION

Two implementations can sample the same categorical distribution at every input and still disagree about what happens when the input changes while the randomness is held fixed. This distinction matters for shared-randomness sampling and Gumbel-max counterfactual models: the object of interest is a joint law across inputs, rather than a collection of separate marginals (Moulton & Jiang, 2018; Oberst & Sontag, 2019; Lorberbom et al., 2021; Daliri et al., 2025). It also creates a verification problem. A sampler can pass every ordinary single-input distribution check and implement the wrong coupling.

We ask how many categorical queries are needed to certify the shared Gumbel joint law at a target pair. The verifier receives handles to independent hidden seeds and may reuse each handle at several inputs. It observes categorical outputs, not the latent exponentials or numerical seed. The reported total-variation bound must cover the true error for every sampler in the promised class; efficient certification is required at the Gumbel reference.

Directly estimating the target joint table costs inverse-square accuracy for a fixed alphabet. Auxiliary inputs can be more informative because their known marginals impose additional moment constraints on the shared response map. We identify precisely when these constraints determine every target event on the reference support. For categorical Gumbel races, this happens exactly when all but possibly the smallest and largest changes in log weights coincide.

The main results connect this exact classification to a quantitative complexity law. Write ρ for the desired TV accuracy, δ for the failure probability, and D for the span of the interior log-ratios.

- **Finite-query classification.** A fixed finite family permits $O(\rho^{-1}\log(1/\delta))$ certification exactly when $D = 0$. The constructive direction uses at most the two targets and one

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.

maximum anchor. If $D > 0$, every fixed finite family requires inverse-square accuracy, even though its size and locations may be chosen for the targets.

- **Accuracy-dependent query complexity.** When the input family may depend on accuracy, the optimal cost is $\Theta(\log(1/\delta)[\rho^{-1} + D\rho^{-4/3}])$, uniformly on each fixed compact target family with distinct targets. A reference predictor with fresh audits attains this order in actual sampler calls. A common-seed change of measure gives the matching lower bound under adaptation, interleaving, and stopping.
- **Marginal tolerance.** We determine the full tolerance transition for each fixed pair with $D > 0$. For a structured $D = 0$ family at tolerance $\rho/3$, weighted paths give finite-sample certificate upper bounds and a universal score-energy lower bound.

The improvement to inverse-$4/3$ complexity requires accuracy-dependent query locations. The finite-family classification and the optimal-rate theorem thus address different query regimes. Our bounds count sampler calls; preprocessing, numerical precision, and wall-clock time are separate computational costs.

Relation to existing work. Shared Gumbel races and maximum anchors provide the sampling geometry (Moulton & Jiang, 2018). Their restrictions on target transitions are central to Gumbel-max structural causal models (Oberst & Sontag, 2019). Since different causal mechanisms can share the same marginals (Lorberbom et al., 2021), our model takes the Gumbel mechanism as its reference rather than inferring it from observational data.

Sampler testing also uses access models richer than independent samples (Chakraborty & Meel, 2019; Bhattacharyya et al., 2024). The huge-object model separates sampling an object from querying its coordinates (Goldreich & Ron, 2022); here the coordinates are categorical inputs with known marginal expectations. Finite-moment identification and convex testing supply general tools (McFadden, 2005; Polyanskiy & Wu, 2026; Goldenshluger et al., 2015). We use prediction with unbiased residual correction (Angelopoulos et al., 2023) to evaluate affine statistics from categorical queries. The new results are the exact Gumbel input classification and the matching geometric seed-and-call rates, including the cost of evaluating the residual.

2 Model and the certification objective

Let $[n] = \{1, \dots, n\}$ and let Δ_{n-1} be the probability simplex. A shared-seed sampler consists of categorical random variables $X_w = F(w, U)$ on one probability space, indexed by $w \in \Delta_{n-1}$. Repeated normalized inputs return the same variable on the same seed. The exact-marginal class satisfies $\Pr_F(X_w = i) = w_i$ for every input. In the tolerance class, a known $\tau \geq 0$ instead guarantees

$$\sup_{w \in \Delta_{n-1}} \mathrm{TV}(\mathcal{L}_F(X_w), w) \leq \tau. \tag{1}$$

Exact target marginals are imposed separately when needed. The marginal conditions are assumptions supplied to the verifier. Alternative response maps may be discontinuous in the input. All queried interactions must be measurable; our lower-bound samplers have jointly measurable response maps.

Under the reference, $E_1, \dots, E_n$ are independent unit exponentials and

$$G_w = \arg\min_{i: w_i > 0} E_i / w_i. \tag{2}$$

The same E is used for every input on one seed, and independent seeds use independent copies. Thus $G_w \sim w$. For fixed targets p, q, put $g = \mathcal{L}(G_p, G_q)$ and $d(F) = \mathrm{TV}(\mathcal{L}_F(X_p, X_q), g)$.

A certificate returns $\widehat{B} \in [0, 1]$. It is *honest* and *reference-complete at accuracy* ρ if, respectively,

$$\Pr_F\{d(F) \leq \widehat{B}\} \geq 1 - \delta \quad \text{for every allowed } F, \qquad \Pr_G\{\widehat{B} \leq \rho\} \geq 1 - \delta. \tag{3}$$

Reference completeness does not promise equally fast success for every sampler with $d(F) = 0$ but different auxiliary couplings.

The verifier can open independent seeds, revisit old seeds, interleave their queries, choose inputs from previous observations and private randomness, and stop sequentially. Opening a seed means

making its first actual query; unused handles have no cost. We study three optimal resources, each minimized over honest, reference-complete procedures: expected opened seeds under G, expected actual calls under G, and a deterministic upper bound on actual calls valid under every sampler and every private-coin outcome. The lower bounds permit all the freedoms above. The constructions use finitely many full-support inputs chosen before sampling for the requested accuracy, with adaptive execution inside that potential pool.

Throughout, n is fixed, $0 < \delta \le 1/4$, and TV is half the ℓ^1 norm. Constants with a subscript indicate their allowed dependence. We use $\mathcal{C}$ when the same order holds for all three resources.

3 An exact classification of finite fast certificates

3.1 The finite-moment criterion

For a fixed finite input family $S \supseteq \{p, q\}$, let $g_S = \mathcal{L}((G_w)_{w\in S})$, and define the space

$$\mathcal{H}_S = \operatorname{span}\{1, \mathbf{1}\{G_w = i\} : w \in S,\ i \in [n]\} \subseteq L^2(g_S). \tag{4}$$

An equality in this space is required only on positive-probability reference rows. The certificate itself must also handle rows that are impossible under the reference.

Theorem 3.1 (Affine observability)**.** *For a fixed finite S, the following are equivalent: every target-atom indicator belongs to $\mathcal{H}_S$; the prescribed probe marginals uniquely determine the target law among laws supported on* $\operatorname{supp} g_S$*; and an honest, reference-complete certificate exists with $O_S(\rho^{-1}\log(1/\delta))$ independent rows. If these conditions fail, every such certificate on S needs $\Omega_S(\rho^{-2}\log(1/\delta))$ rows for sufficiently small ρ. The same lower bound holds for reference expected opened seeds under adaptive querying and stopping when every possible query lies in S.*

The positive direction is constructive. Let Y be the actual target-pair unit vector, and let V be an affine reconstruction, evaluated on the actual probe outputs. Exact marginals give $\mathbb{E}_F V = g$ even on alternatives. Consequently

$$d(F) \le \mathbb{E}_F R, \qquad R = \tfrac{1}{2}\|Y - V\|_1, \qquad R = 0 \quad G\text{-a.s.} \tag{5}$$

The finite-row residual is bounded, so all-zero residual observations give an inverse-accuracy certificate. When reconstruction fails, an orthogonal bounded tilt preserves every probe marginal while changing a target atom. Its KL divergence is quadratic in the target change, yielding the inverse-square lower bound. Appendix A proves both directions and extends the finite-row alternatives to full-input samplers.

3.2 The Gumbel geometry

For full-support targets, set $r_i = q_i/p_i$ and sort the ratios with multiplicity. Define

$$D = \log \frac{r_{(n-1)}}{r_{(2)}}. \tag{6}$$

Theorem 3.2 (Full-support pair classification)**.** *For $n \ge 3$, a fixed finite family with an inverse-accuracy certificate exists if and only if $D = 0$. In that case, with $c = r_{(2)}$, at most the three inputs*

$$p, \quad q, \quad h_i = \frac{\max(p_i, q_i/c)}{\sum_j \max(p_j, q_j/c)} \tag{7}$$

suffice. When $D > 0$, no fixed finite input family has that rate.

The criterion has a direct logit interpretation. If $p = \operatorname{softmax}(u)$ and $q = \operatorname{softmax}(v)$, all but at most one smallest and one largest coordinate of $v - u$ must agree. A common shift does not matter. Three labels always satisfy the criterion; four labels are the first nontrivial case.

To see the construction, suppose there is a low extreme a, a high extreme b, and middle labels I with ratio c. A transition $G_p = i, G_q = j$, $i \ne j$, is possible exactly when $r_i < r_j$. Moreover, for each $j \in I$,

$$\mathbf{1}\{G_p = G_q = j\} = \mathbf{1}\{G_h = j\}. \tag{8}$$

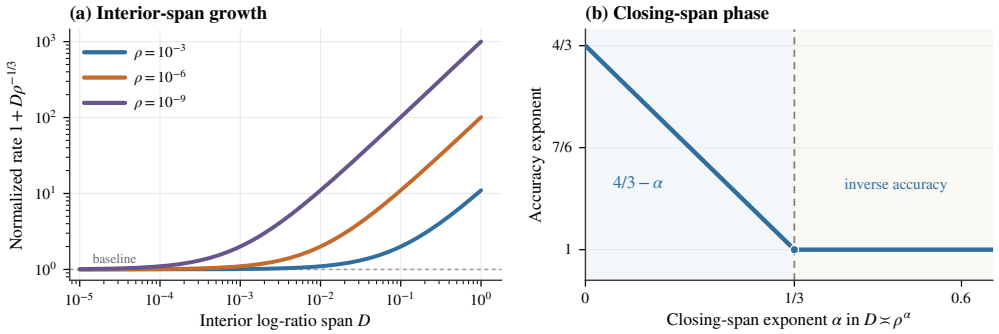


Figure 1: Analytic rate in Theorem 4.1, with fixed alphabet and confidence. Left: the normalized rate $\rho\mathcal{C}/\log(1/\delta)$ has order $1+D\rho^{-1/3}$. Right: along $D \asymp \rho^{\alpha}$, the accuracy exponent is $\max\{1, 4/3-\alpha\}$. The curves evaluate the rate expression with unit constants, rather than measured sample counts or runtimes.

This maximum-anchor identity is established sampling geometry (Moulton & Jiang, 2018). Subtracting it from target row and column indicators recovers transitions (a, j) and (j, b); subtracting these from row a recovers (a, b). The resulting residual in (5) satisfies $0 \le R \le 3$, independently of n. Therefore N all-zero rows certify $d(F) \le 3\log(1/\delta)/N$ using at most $3N$ calls.

For example, take $p = (1, 1, 1, 1)/4$. The target $q = (.1, .25, .25, .4)$ has $D = 0$ and anchor $h = (5, 5, 5, 8)/23$. Changing the middle target coordinates to $.2, .3$ gives $D = \log(3/2) > 0$, ruling out inverse-accuracy certification from any fixed finite pool. For necessity, we use two disjoint competing label pairs. Away from the finitely many query boundaries, a perturbation correlates their winner bits while preserving each probe marginal. It is therefore invisible to the probes' marginals but changes the target joint law.

4 Optimal Seed and Actual-Call Complexity

Theorem 4.1 (Uniform geometric rate). *Fix $n \ge 4$ and $0 < \eta < 1/n$. There is $\rho_0 > 0$, depending only on n, η, such that for distinct targets with $p_i, q_i \ge \eta$ and $0 < \rho \le \rho_0$, exact-marginal certification has*

$$\mathcal{C} = \Theta_{n,\eta}\Big(\log(1/\delta)\,[\rho^{-1} + D\rho^{-4/3}]\Big). \tag{9}$$

The lower bound permits arbitrary adaptive hidden-seed policies and stopping. The upper bound uses full-support queries and has a deterministic call cap of this order under every sampler. If $p = q$, exact marginals and repeated-input consistency determine the target law with zero queries.

For a fixed pair with $D > 0$, the optimal rate is inverse-$4/3$. The crossover is $D \asymp \rho^{1/3}$: below this scale, the inverse-accuracy term dominates. A potential input pool held fixed as $\rho \downarrow 0$ still has the inverse-square lower bound of Theorem 3.1.

4.1 The Origin of the $4/3$ Exponent

Appendix B gives the finite constructions and uniform estimates. The proof starts with a maximum-anchor path that splits target events along consecutive interior ratio gaps. Rectangular integrals of single-input winner indicators approximate the resulting table coordinates by affine statistics. With smoothing width h, these statistics have centered reference residuals with

$$\mathrm{Var}_G(Z) \le C_{n,\eta}(h^3 + Dh^2), \qquad |Z| \le C_n(1 + D/h). \tag{10}$$

Finite quadrature preserves these orders. The nominal expectation of each affine term is known from the marginal promise, so the centered residual estimates the actual target discrepancy under every alternative. Variance-adaptive bounded-mean concentration (Maurer & Pontil, 2009) gives a sufficient row count of order

$$\log(1/\delta)\left[\frac{h^3 + Dh^2}{\rho^2} + \frac{1 + D/h}{\rho}\right]. \tag{11}$$

Setting $h \asymp \rho^{1/3}$ yields (9).

The matching lower bound uses a signed change in the latent race law. Two nearby comparison coordinates carry a checkerboard perturbation; the remaining coordinates retain enough width to make its constants uniform as $D \downarrow 0$. A single response-changing branch repairs marginal defects simultaneously at every input. The resulting alternative has target TV greater than ρ and one-seed reverse KL at most $C_{n,\eta}\rho^{4/3}/D$ in the regime where this term is relevant. A separate exact-marginal mixture supplies the inverse-accuracy lower bound. Both comparisons use a common response mapping under two augmented seed laws, so they apply to arbitrary auxiliary queries.

If the one-seed reverse KL is d, revealing the full latent seed at its first query gives augmented-history KL $d\mathbb{E}_G M_T$ at a finite call horizon. Revisits contribute no additional latent information. Binary data processing for the event of reporting at most ρ, followed by a limit in the horizon, yields

$$\mathbb{E}_G M \geq \frac{\mathrm{kl}(1-\delta, \delta)}{d}. \tag{12}$$

This proves the seed lower bound under adaptation and stopping. Each opened seed requires a call, so it also lower-bounds the other two resources.

4.2 Query-efficient evaluation of affine statistics

A row contains a target observation Y and an affine term $V = c_0 + \sum_j a_j I_j$ over potentially many auxiliary inputs. A categorical pilot predicts the indicators $\widehat{I}_j$ from the same hidden seed. After completing that prediction, draw a fresh private index J with $\pi_j = |a_j| / \sum_k |a_k|$ and query its actual indicator. Then

$$\widetilde{V} = c_0 + \sum_j a_j \widehat{I}_j + \frac{a_J}{\pi_J}(I_J - \widehat{I}_J) \tag{13}$$

is conditionally unbiased for V given the complete response map and pilot. This holds for every sampler. Prediction error affects the reference variance; calibration follows from the fixed response map and independent audit draw.

A short gate checks whether a target winner survives a prescribed weight decrease. On the stable branch, aligned quadrature groups cancel exactly under Gumbel, permitting a constant-cost group audit. Only with probability $O(h)$ is a precision pilot with $O_n(\log(1/h))$ calls needed. The prediction error is chosen small enough to preserve (10). The expected pilot overhead per row is bounded. For a deterministic all-sampler call budget, cap the number of expensive branches and report 1 if an additional branch would exceed the cap. A reference tail bound limits this event to half the failure probability. The resulting procedure controls both reference expected seeds and calls under every sampler.

5 Marginal uncertainty and a critical extension

5.1 A sharp tolerance transition for fixed nonexceptional pairs

Theorem 5.1 (Known marginal tolerance)**.** *Fix full-support p, q with $D > 0$. For sufficiently small ρ and a known allowance $\tau \in [0, 1]$ in (1), all three resources have order*

$$\mathcal{C} = \Theta_{p,q,n}\Big(\log(1/\delta)\min\{\rho^{-2}, \max\{\rho^{-4/3}, \tau^2\rho^{-4}\}\}\Big). \tag{14}$$

The lower bound remains valid when the alternative's two target marginals are exact. The upper bound allows their marginal errors as well.

Small errors retain the inverse-$4/3$ rate: the transition begins at $\tau \asymp \rho^{4/3}$. Between that scale and ρ, the cost is $\tau^2\rho^{-4}\log(1/\delta)$; at tolerance of order ρ or larger, direct target estimation attains the optimal inverse-square order. The affine coefficient mass makes the effect transparent. Marginal error adds bias $O(\tau/h)$, so a useful bandwidth must satisfy $h \gtrsim \tau/\rho$. Optimizing the free-bandwidth certificate and combining it with direct estimation proves the upper bound. The corresponding partially repaired common-seed alternatives prove the lower bound; see Appendix B.

The constants in Theorem 5.1 depend on the fixed pair with $D > 0$. We next analyze a structured $D = 0$ family at the critical allowance.

5.2 WEIGHTED PATHS AT THE CRITICAL ALLOWANCE

Let $n \geq 4$, $p_i \geq \eta$, and $q_i \propto p_i \exp(-a\epsilon \mathbf{1}\{i = 1\} + b\epsilon \mathbf{1}\{i = n\})$, where $a, b > 0$ are fixed and $b/(a + b)$ is irrational with bounded continued-fraction partial quotients. These targets have $D = 0$. Impose exact target marginals and the all-input allowance $\tau = \rho/3$.

For a mean-zero seed score H, write $\mathcal{B}_H(w)_i = \mathbb{E}_G[H\mathbf{1}\{G_w = i\}]$. For a nonconstant middle-label detector $c \in \{-2, 1\}^{n-2}$, let $C_c = \sum_{j=2}^{n-1} c_j \mathbf{1}\{G_p = G_q = j\}$. Define $K_*(\epsilon)$ as the infimum of $\mathbb{E}_G H^2$ over these detectors and scores satisfying target nullity, $\mathbb{E}_G[HC_c] = 3$, and $\sup_w \|\mathcal{B}_H(w)\|_{\mathrm{TV}} \leq 1$.

Theorem 5.2 (Critical certificate and energy bounds)**.** *Uniformly over the fixed compact data above, for sufficiently small ϵ and $0 < \rho \leq c_0\epsilon$, there is an honest, reference-complete certificate with each of the three resource bounds at most*

$$C \log(1/\delta) \max\{\epsilon^{1/4}\rho^{-5/4},\ \epsilon^{1/2}\rho^{-7/5},\ \epsilon^{11/7}\rho^{-2}\}. \tag{15}$$

The strict-cap score energy satisfies

$$c\epsilon^{-11/7} \leq K_*(\epsilon) \leq C\epsilon^{-2}. \tag{16}$$

The orbit estimate weights each step by its actual height. For orbit length N and J occupied strips, most steps remain in a fixed strip, and only $O(j + 1)$ steps reach strip j. The accumulated thin-event error is $C\epsilon\sqrt{K}(N + J^{3/2})$. Predetermined finite orbit pools and convex separation give bounded affine residuals with variance

$$V_{N,J} \leq C\left[\frac{\epsilon}{JN^4} + \epsilon^2 + \frac{\epsilon^2 J^3}{N^2}\right], \qquad 3 - \kappa_N \asymp N^{-1}. \tag{17}$$

The margin $\rho - \kappa_N\rho/3$ is positive. Exact local prediction, fresh audits, and a gate-count cap turn this residual into a finite-sample certificate. Optimizing N, J gives (15); the strict-cap weighted inequality gives (16). Appendix C proves the estimates and the full call accounting.

The three upper-bound terms dominate, respectively, above $\epsilon^{5/3}$, between $\epsilon^{25/14}$ and $\epsilon^{5/3}$, and below $\epsilon^{25/14}$. Matching statistical lower bounds for the middle and deep regimes remain open. An $\epsilon^{-11/7}$ energy construction is also open; the energy lower bound excludes $K_* = O(\epsilon^{-3/2})$ without fixing the optimal exponent. The certificate proof controls the joint small-parameter regime at finite sample sizes.

6 DISCUSSION AND COMPUTATIONAL EVIDENCE

Auxiliary inputs turn marginal information into constraints on the target joint law. Their value is determined by how accurately affine statistics of the auxiliary outputs reconstruct target events. The finite-pool classification identifies exact reconstruction; the uniform rate quantifies the cost of approximation. The upper and lower bounds use the same full-input sampler class.

The implementation supports exact finite-law checks, finite affine plans, and compressed execution for supplied critical input pools. At $N = 10^6$, the three-label reference experiments give tighter equal-call bounds for affine residual certificates than for direct estimation. The four-label counted-query experiments favor direct estimation. The critical weighted-orbit pool generator and global optimizer remain unimplemented. These experiments assess finite plans, rather than the asymptotic critical rates; Appendix D gives the implementation details and results.

The uniform bounds assume a fixed alphabet and coordinates bounded away from zero. Growing alphabets and vanishing coordinates require separate analysis. Certifying the all-input marginal promise from finite observations requires additional input regularity. The certificates concern target joint laws, not literal seed alignment or identification of a causal mechanism.

Reproducibility statement. The appendices contain the proofs, parameter choices, and stopped-query comparisons. Appendix D lists executable interfaces, reproduction commands, and saved results. Figure 1 plots the analytic rate expression.

AI USE STATEMENT

Large language models were used extensively for mathematical exploration, proof drafting and checking, literature search, code development, and manuscript editing. Model outputs were revised during the research process. The authors are responsible for the claims, proofs, references, and code.

REFERENCES

Anastasios N. Angelopoulos, Stephen Bates, Clara Fannjiang, Michael I. Jordan, and Tijana Zrnić. Prediction-powered inference. *Science*, 382(6671):669–674, 2023. doi: 10.1126/science.adi6000. URL https://doi.org/10.1126/science.adi6000.

Rishiraj Bhattacharyya, Sourav Chakraborty, Yash Pote, Uddalok Sarkar, and Sayantan Sen. Testing self-reducible samplers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 7952–7960, 2024. doi: 10.1609/aaai.v38i8.28632. URL https://doi.org/10.1609/aaai.v38i8.28632.

Sourav Chakraborty and Kuldeep S. Meel. On testing of uniform samplers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 7777–7784, 2019. doi: 10.1609/aaai.v33i01.33017777. URL https://doi.org/10.1609/aaai.v33i01.33017777.

Majid Daliri, Christopher Musco, and Ananda Theertha Suresh. Coupling without communication and drafter-invariant speculative decoding. In *2025 IEEE International Symposium on Information Theory (ISIT)*, pp. 1–6, 2025. doi: 10.1109/ISIT63088.2025.11195663. URL https://doi.org/10.1109/ISIT63088.2025.11195663.

Alexander Goldenshluger, Anatoli Juditsky, and Arkadi Nemirovski. Hypothesis testing by convex optimization. *Electronic Journal of Statistics*, 9(2):1645–1712, 2015. doi: 10.1214/15-EJS1054. URL https://doi.org/10.1214/15-EJS1054.

Oded Goldreich and Dana Ron. Testing distributions of huge objects. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics*, pp. 78:1–78:19, 2022. doi: 10.4230/LIPIcs.ITCS.2022.78. URL https://doi.org/10.4230/LIPIcs.ITCS.2022.78.

Guy Lorberbom, Daniel D. Johnson, Chris J. Maddison, Daniel Tarlow, and Tamir Hazan. Learning generalized Gumbel-max causal mechanisms. In *Advances in Neural Information Processing Systems*, volume 34, pp. 26792–26803, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/e143c01e314f7b950daca31188cb5d0f-Abstract.html.

Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Conference on Learning Theory (COLT 2009)*, 2009. URL https://www.cs.mcgill.ca/~colt2009/papers/012.pdf.

Daniel L. McFadden. Revealed stochastic preference: A synthesis. *Economic Theory*, 26(2): 245–264, 2005. doi: 10.1007/s00199-004-0495-3. URL https://doi.org/10.1007/s00199-004-0495-3.

Ryan Moulton and Yunjiang Jiang. Maximally consistent sampling and the Jaccard index of probability distributions. In *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*, pp. 1094–1103, 2018. doi: 10.1109/ICDMW.2018.00159. URL https://doi.org/10.1109/ICDMW.2018.00159.

Michael Oberst and David Sontag. Counterfactual off-policy evaluation with Gumbel-Max structural causal models. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4881–4890. PMLR, 2019. URL https://proceedings.mlr.press/v97/oberst19a.html.

Yury Polyanskiy and Yihong Wu. Dualizing Le Cam's method for functional estimation I: General theory. *The Annals of Statistics*, 54(1):1–24, 2026. doi: 10.1214/25-AOS2498. URL https://doi.org/10.1214/25-AOS2498.

Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society: Series B*, 86(1):1–27, 2024. doi: 10.1093/jrsssb/qkad009. URL https://doi.org/10.1093/jrsssb/qkad009.

A FINITE-POOL OBSERVABILITY AND THE PAIR CLASSIFICATION

This appendix proves the finite-dimensional criterion and the local obstruction. All statements assume exact input marginals. A universal sampler is a family $X_s = F(s, U)$, indexed by probability vectors on $[n]$, with $\Pr(X_s = i) = s_i$. Equal inputs denote the same random variable; measurability in U is required at every fixed input. The reference family is

$$G_s = \arg\min_{i:s_i>0} E_i/s_i, \qquad E_1, \dots, E_n \text{ independent } \mathrm{Exp}(1).$$

Ties have probability zero. Fix a finite set S containing the target inputs $P = (p^1, \dots, p^m)$. Each of N independent runs reveals the entire vector $(X_s : s \in S)$. Write g_S for its reference law, $\mathcal{C}_S = \mathrm{supp}(g_S)$, and g_P for the reference target law. A report $\widehat{B} \in [0, 1]$ is honest and reference-complete at (ρ, δ) when, respectively,

$$\Pr_F\{\mathrm{TV}(\mu_F^P, g_P) \le \widehat{B}\} \ge 1 - \delta \qquad \text{for every universal exact-marginal } F, \tag{18}$$

$$\Pr_G\{\widehat{B} \le \rho\} \ge 1 - \delta. \tag{19}$$

Independent randomness in the report is allowed. Throughout, $0 < \delta < 1/2$ and $\kappa_\delta = (1 - 2\delta)\log((1-\delta)/\delta)$.

Every law on the finite transcript space with marginals s, $s \in S$, extends to a universal exact-marginal sampler. Indeed, generate its vector with one seed component and return the designated coordinate at each input in S. At every other input use inverse-CDF sampling with one independent uniform component. This construction respects repeated inputs. No collision constraint is imposed outside S.

A.1 THE AFFINE CRITERION

Define the subspace

$$\mathcal{H}_S = \mathrm{span}\{1, \mathbf{1}\{G_s = i\} : s \in S,\ i \in [n]\} \subseteq L^2(g_S).$$

Equality in this space means equality on every positive-probability reference transcript.

Theorem A.1 (Fixed-pool dichotomy)**.** *For fixed S and P, the following are equivalent:*

1. *Every target-atom indicator $\mathbf{1}\{G_P = z\}$ belongs to $\mathcal{H}_S$.*
2. *Every law supported on $\mathcal{C}_S$ with the prescribed individual input marginals has target law g_P.*
3. *Honest, reference-complete certificates exist with $N = O_S(\rho^{-1}\log(1/\delta))$ for all sufficiently small $\rho > 0$.*

If these conditions fail, there are $v_S, \rho_S > 0$ such that every such certificate satisfies

$$N \ge \frac{v_S \kappa_\delta}{4\rho^2}, \qquad 0 < \rho \le \rho_S. \tag{20}$$

The pool is fixed independently of accuracy and confidence.

Proof. Suppose the first condition holds. For each reference target atom choose coefficients with $\mathbf{1}\{G_P = z\} = b_z + \sum_{s\in S} A_{z,s,G_s}$. Evaluate the same affine expression V_z on arbitrary transcripts. Combine target outcomes outside $\mathrm{supp}(g_P)$ into a category $\perp$ and set $V_\perp = (g_P)_\perp = 0$. Let Y be the one-hot vector of this coarsened target outcome. This coarsening preserves TV against g_P. Under every exact-marginal law,

$$\mathbb{E}V = g_P, \qquad \mathrm{TV}(\mu_F^P, g_P) \le \mathbb{E}R, \qquad R = \tfrac{1}{2}\|Y - V\|_1. \tag{21}$$

The first identity uses only individual input marginals; the inequality is Jensen's inequality. Since the transcript space is finite, $0 \le R \le L$ for some finite L, and $R = 0$ under the reference. No positivity or mass-conservation property of V on nonreference transcripts is needed.

If $L = 0$, report zero. Otherwise, report one after any positive observed residual, and report $\min\{1, L(1-\delta^{1/N})\}$ if all residuals vanish. When the true TV exceeds the latter number below one, (21) implies $\Pr_F(R > 0) > 1 - \delta^{1/N}$; the probability of all zero residuals is therefore less than δ. This proves honesty. Under the reference, the report is surely at most $L\log(1/\delta)/N$. Taking expectations of the affine identities under a law supported on $\mathcal{C}_S$ also proves the second condition.

Conversely, let $f = \mathbf{1}\{G_P = z\} \notin \mathcal{H}_S$ and set $h = f - \Pi_{\mathcal{H}_S} f$. Orthogonality gives

$$\mathbb{E}_g h = 0, \quad \mathbb{E}_g[h\mathbf{1}\{G_s = i\}] = 0, \quad \mathbb{E}_g(hf) = \mathbb{E}_g h^2 =: v_S > 0.$$

The law $g_t = g_S(1 + th)$ is positive for $0 < t \le (2\|h\|_\infty)^{-1}$. It preserves every input marginal and changes atom z by tv_S, disproving the second condition. Furthermore,

$$\mathrm{TV}((g_t)_P, g_P) \ge tv_S, \qquad D_{\mathrm{KL}}(g_t\|g_S) \le \chi^2(g_t\|g_S) = t^2 v_S.$$

Extend g_t to a universal sampler as above and take $t = 2\rho/v_S$, which is valid when $\rho \le \rho_S := v_S/(4\|h\|_\infty)$. For $A = \{\widehat{B} \le \rho\}$, honesty and completeness give $\Pr_t(A) \le \delta$ and $\Pr_g(A) \ge 1 - \delta$. KL additivity, data processing, and binary relative entropy yield

$$\frac{4N\rho^2}{v_S} \ge D_{\mathrm{KL}}(g_t^{\otimes N}\|g_S^{\otimes N}) \ge \mathrm{kl}(\Pr_t(A), \Pr_g(A)) \ge \kappa_\delta.$$

The same argument includes independent reporting randomness, whose law is unchanged. It proves (20) and rules out the third condition whenever the first fails. □

A.2 Exactly which full-support pairs are observable?

Theorem A.2 (Full-support pair classification). *Let $n \ge 3$ and let p, q have full support. Put $r_i = q_i/p_i$ and order these ratios with multiplicity. A fixed finite pool containing p, q affinely reconstructs every reference pair atom, or equivalently permits the inverse-accuracy rate in Theorem A.1, if and only if*

$$r_{(2)} = r_{(n-1)}. \tag{22}$$

When this holds, at most three distinct inputs suffice: p, q, and

$$h_i = \frac{\max(p_i, q_i/c)}{\sum_k \max(p_k, q_k/c)}, \qquad c = r_{(2)}. \tag{23}$$

If $r_{(2)} < r_{(n-1)}$, then for every fixed finite $S \supseteq \{p, q\}$ there is $k_S > 0$ such that even honesty restricted to exact-marginal alternatives preserving every probability $\Pr(X_s = X_t = i)$, $s, t \in S$, requires

$$N \ge k_S\kappa_\delta/\rho^2, \qquad 0 < \rho \le k_S/2. \tag{24}$$

The negative statement allows auxiliary inputs on the simplex boundary.

Proof of sufficiency. Write $T_i = E_i/p_i$. Distinct winners $G_p = i, G_q = j$ require $T_i < T_j$ and $T_j/r_j < T_i/r_i$, hence $r_i < r_j$. Conversely, if $r_i < r_j$, choosing $T_i = 1$, $1 < T_j < r_j/r_i$, and all other coordinates sufficiently large gives these winners on an open set of positive probability. Every diagonal is possible as well.

Under (22), put $J = \{j : r_j = c\}$. There is at most one low label a with $r_a < c$ and at most one high label b with $r_b > c$. Write $I_i^s = \mathbf{1}\{G_s = i\}$ and $D_i = \mathbf{1}\{G_p = G_q = i\}$. Extremal nesting gives $D_a = I_a^q$ and $D_b = I_b^p$. For $j \in J$,

$$\frac{h_i}{h_j} = \max\left\{\frac{p_i}{p_j}, \frac{q_i}{q_j}\right\}, \qquad D_j = I_j^h.$$

The second identity follows directly from the simultaneous race inequalities. This is the maximum-anchor construction of Moulton & Jiang (2018), shared here by all middle labels. The remaining possible atom indicators are

$$\begin{aligned}
&\mathbf{1}\{G_p = a, G_q = j\} = I_j^q - I_j^h, && j \in J,\\
&\mathbf{1}\{G_p = j, G_q = b\} = I_j^p - I_j^h, && j \in J,\\
&\mathbf{1}\{G_p = a, G_q = b\} = I_a^p - I_a^q - \sum_{j\in J}(I_j^q - I_j^h).
\end{aligned} \tag{25}$$

These follow by subtracting diagonals from columns or rows and then subtracting the other cells from row a. Omit identities mentioning an absent extreme. Every other atom is zero on the reference support. If only the low extreme exists, $h = p$; if only the high extreme exists, $h = q$. With neither extreme, normalization gives $p = q = h$. Thus duplicate-input consistency covers all degenerate cases, and Theorem A.1 applies. □

Proof of necessity and the lower bound. A strict middle gap provides disjoint pairs $A = \{a, a'\}$ and $B = \{b, b'\}$ with $\max_A r_i < \min_B r_i$. Choose λ strictly between these numbers and positive latent centers

$$e^*_a = p_a\lambda, \quad e^*_{a'} = p_{a'}\lambda, \quad e^*_b = q_b, \quad e^*_{b'} = q_{b'}, \quad e^*_k = c_k > \max(\lambda p_k, q_k) \quad (k \notin A \cup B).$$

At this center p has precisely the tied minimizing pair A and q has precisely the tied minimizing pair B. For each probe $s \in S$, exclude equality of two supported scores at the center unless the compared pair is A or B. These exclusions are finitely many proper affine hyperplanes in the open parameter domain $(\lambda, (c_k)_{k\notin A\cup B})$: an A–B comparison has a nonzero λ coefficient, and a comparison involving an extra label has a nonzero coefficient in a free c_k. Choose parameters outside their union. Each probe then has a singleton minimizer, or exactly the minimizing pair A, or exactly the minimizing pair B.

Use independent scaled coordinates $Z_a = E_a/p_a$, $Z_{a'} = E_{a'}/p_{a'}$, $Z_b = E_b/q_b$, $Z_{b'} = E_{b'}/q_{b'}$, and $Z_k = E_k$ otherwise. By the strict score gaps, a sufficiently small product box

$$\mathcal{B} = I \times I \times J' \times J' \times \prod_{k\notin A\cup B} K_k$$

around the center has the following property: every probe is constant on the box or depends only on one of $U = \mathbf{1}\{Z_a < Z_{a'}\}$ and $V = \mathbf{1}\{Z_b < Z_{b'}\}$. In fact a tied A comparison at the center has equal coefficients of $Z_a, Z_{a'}$, and likewise for B. The targets read U and V, respectively. Unsupported labels cannot win.

Conditional on $\mathcal{B}$, independence of all scaled coordinates persists. Thus U, V are independent with means $\alpha, \zeta \in (0, 1)$. Equal intervals and positive coordinate densities ensure both orderings have positive probability; no exchangeability of scaled exponentials is needed. Set

$$H = \mathbf{1}_{\mathcal{B}}(U - \alpha)(V - \zeta), \qquad k_S = \Pr_G(\mathcal{B})\alpha(1 - \alpha)\zeta(1 - \zeta) > 0.$$

Then $|H| \le 1$, $\mathbb{E}_G H = 0$, and $\mathbb{E}_G H^2 = k_S$. Single-probe indicators depend on at most one bit on the box. The same is true of $\mathbf{1}\{G_s = G_t = i\}$: two probes depending on different bits take values in the disjoint label sets A, B, so their diagonal indicator vanishes. Centering the unused bit therefore gives

$$\mathbb{E}_G[H\mathbf{1}\{G_s = i\}] = \mathbb{E}_G[H\mathbf{1}\{G_s = G_t = i\}] = 0.$$

Tilt the latent law by $1+\theta H$, $|\theta| \le 1/2$, retain its transcript on S, and apply the finite-family extension. All required marginals and diagonals are preserved. The four target cells $(a, b), (a, b'), (a', b), (a', b')$ change by $\theta k_S(1, -1, -1, 1)$, and all others are unchanged. Consequently,

$$\mathrm{TV}(\mu^{(p,q)}_{F_\theta}, g_{(p,q)}) = 2|\theta|k_S, \qquad D_{\mathrm{KL}}(\mathbb{P}_\theta \| \mathbb{P}_G) \le \theta^2 k_S.$$

The KL inequality is on the latent space and hence also bounds transcript KL. Choosing $\theta = \rho/k_S$ and applying the binary testing argument from Theorem A.1 proves (24). The tilted transcript law is supported on $\mathcal{C}_S$ but changes the target law, so the same theorem excludes affine reconstruction. □

Proposition A.3 (Adaptive stopping within a fixed pool)**.** *Fix the finite allowed input set S. A measurable procedure may choose queries in S adaptively, use independent private randomness, interleave and revisit seeds, and stop sequentially. A new seed is an independent draw and is counted as opened on its first sampler call. Assume the procedure terminates almost surely under the reference, is honest, and is reference-complete in the sense of (18)–(19); assign report one to nontermination under an alternative. Let K be the number of opened seeds. If affine observability fails, the constants in Theorem A.1 satisfy*

$$\mathbb{E}_G K \ge \frac{v_S \kappa_\delta}{8\rho^2}, \qquad 0 < \rho \le \rho_S.$$

For the obstructed pairs in Theorem A.2, its construction also gives

$$\mathbb{E}_G K \geq \frac{k_S \kappa_\delta}{2\rho^2}, \qquad 0 < \rho \leq k_S/2,$$

even with honesty restricted to its diagonal-preserving alternatives.

Proof. For a bounded, mean-zero tilt $d\mathbb{P}_t/d\mathbb{P}_0 = 1 + tH$ with $|tH| \leq 1/2$, the inequality $-\log(1+x) \leq -x + 2x^2$ gives

$$D_{\mathrm{KL}}(\mathbb{P}_0 \| \mathbb{P}_t) \leq 2t^2 \mathbb{E}_0 H^2. \tag{26}$$

This bounds the reverse KL $d = D_{\mathrm{KL}}(g_S \| (g_t)_S)$ of the full S-response vector, by data processing when the tilt is defined on the latent race space.

For the KL calculation, reveal a seed's entire S-response vector when that seed is first opened. The procedure still chooses queries from its actual observations. Conditional on the augmented past, a newly opened vector has law g_S under the reference and $(g_t)_S$ under the alternative. Each opening contributes d to reverse conditional KL; revisits contribute zero because their answers are already determined. Private random choices have the same conditional kernel in both experiments and contribute zero as well. Thus, for the augmented history through at most m calls, including a stopping decision and report when these occur, the chain rule gives reverse KL at most $d\mathbb{E}_G K_m$, where K_m counts openings by that horizon. The case $\mathbb{E}_G K = \infty$ is immediate, so assume it is finite.

Let A_m be the event that the procedure has terminated with report at most ρ by this horizon. Binary data processing gives

$$d\mathbb{E}_G K_m \geq \mathrm{kl}(\Pr_G(A_m), \Pr_t(A_m)).$$

The events increase to $A = \{\widehat{B} \leq \rho\}$; nontermination cannot belong to A, since the stated accuracy ranges have $\rho < 1$. Almost-sure reference termination, completeness, and honesty at an alternative with target error 2ρ yield $\Pr_G(A) \geq 1 - \delta$ and $\Pr_t(A) \leq \delta$. Taking the limit, using monotone convergence and lower semicontinuity of binary KL, proves $d\mathbb{E}_G K \geq \kappa_\delta$.

For the general affine obstruction use $H = h$, $t = 2\rho/v_S$, and $\mathbb{E}_g h^2 = v_S$ from Theorem A.1; then $d \leq 8\rho^2/v_S$. For the pair construction use its latent H, $t = \rho/k_S$, and $\mathbb{E}_G H^2 = k_S$, giving $d \leq 2\rho^2/k_S$. Both alternatives have target error at least 2ρ and preserve the promises already established above. Only their common response-vector experiment on S is used: the universal extensions outside S need not share a reference response map. □

Corollary A.4 (An alphabet-independent certificate)**.** *Under (22), the three-input affine construction has a residual $0 \leq R \leq 3$ on every transcript and $R = 0$ under the reference. Thus $N \geq \lceil 3\log(1/\delta)/\rho \rceil$ rows suffice, with at most $3N$ sampler calls.*

Proof. Suppose both extremes exist and evaluate the diagonal identities and (25) on arbitrary outputs, with $P_i = \mathbf{1}\{X_p = i\}$, $Q_i = \mathbf{1}\{X_q = i\}$, and $H_i = \mathbf{1}\{X_h = i\}$. Their affine table V has total mass one on every assignment. Its only potentially negative entries are $Q_j - H_j$, $P_j - H_j$, and

$$V_{ab} = P_a + Q_b - 1 + \sum_{j \in J} H_j.$$

If the anchor output is middle, the last entry is nonnegative and the negative mass is at most two. If the anchor output is extreme, only V_{ab} can be negative, with magnitude at most one. Hence $\|V\|_1 \leq 5$ and $\frac{1}{2}\|Y - V\|_1 \leq 3$ for the actual one-hot target vector Y (including the collapsed impossible-outcome category). With only a low extreme the nonzero entries are $V_{aa} = Q_a$, $V_{jj} = P_j$, and $V_{aj} = Q_j - P_j$; their total mass is one and negative mass at most one, giving $R \leq 2$. The high-only case follows by exchanging the targets. With no extremes, $p = q$ and $R = 0$. The zero-residual certificate in Theorem A.1 now proves the assertion. □

A.3 Optimality of the inverse-accuracy rate

Proposition A.5 (An inverse-accuracy lower bound)**.** *Fix distinct full-support p, q on $n \geq 2$ labels. There is $d_{p,q} > 0$ such that every honest, reference-complete certificate using N iid rows on any finite*

pool containing the targets satisfies

$$N \geq \frac{\log((1-\delta)/\delta)}{-\log(1-2\rho/d_{p,q})} \geq \frac{d_{p,q}}{4\rho}\log\frac{1-\delta}{\delta}, \qquad 0 < \rho \leq d_{p,q}/4.$$

The pool may depend on ρ, δ; the constant does not depend on it.

Proof. Let Q be a universal exact-marginal sampler whose outputs at the two distinct targets are independent, using the finite-family extension. Set $d_{p,q} = \mathrm{TV}(p \otimes q, g_{(p,q)})$. This number is positive: distinct normalized targets have two unequal ratios, and an off-diagonal transition from a larger ratio to a smaller one has reference probability zero and positive product probability. On each seed select the entire sampler G with probability $1-\lambda$ and Q with probability λ, sharing this choice across every input. The resulting exact-marginal sampler has target error $\lambda d_{p,q}$. Set $\lambda = 2\rho/d_{p,q} \leq 1/2$. With probability $(1-\lambda)^N$ every row uses G; conditional on this event all observations, including any reporting randomness, have their reference law. Hence an incorrect report at most ρ has probability at least $(1-\lambda)^N(1-\delta)$. Honesty bounds this by δ. Taking logarithms and using $-\log(1-\lambda) \leq 2\lambda$ proves the claim. □

For distinct pairs satisfying (22), the three-input construction and Proposition A.5 give the optimal order $\Theta_{p,q}(\rho^{-1}\log(1/\delta))$ for $\delta \leq 1/4$. At $p = q$, repeated-input consistency and the exact marginal promise determine the target law without queries. For pairs failing the criterion, the inverse-square obstruction is *pool-specific*: its constant and accuracy threshold depend on the inputs in S. Accuracy-dependent pools require the uniform lower bounds and approximation constructions developed in Appendix B, which yield the $\rho^{-4/3}$ rate.

B Uniform minimax rates with exact input marginals

This appendix proves uniform seed and call bounds as the interior ratio span tends to zero. At each accuracy, the construction uses a finite set of full-support inputs. Calibration identities hold for arbitrary sampler responses; the low-variance estimates use the Gumbel reference law. Alternative response maps may be discontinuous in the input.

B.1 Model, statement, and proof organization

Let Δ_{n-1} be the probability simplex. A universal sampler is a jointly measurable map $F(s,U) \in [n]$, with one hidden seed U shared across inputs. Repeated queries at the same input and seed return the same label. Different opened seeds are independent. In the exact class,

$$\Pr\{F(s,U) = i\} = s_i \qquad (s \in \Delta_{n-1},\ i \in [n]). \tag{27}$$

The reference is the exponential-race representation of shared Gumbel noise: $G_s = \arg\min_{i:s_i>0} E_i/s_i$, where $E_1, \ldots, E_n$ are independent unit exponentials. Fix measurable tie breaking throughout. For two targets, write $\mu_F = \mathcal{L}(F(p,U), F(q,U))$ and $g = \mu_G$.

A certificate reports $\widehat{B} \in [0,1]$. Honesty and reference completeness mean, respectively,

$$\Pr_F\{\mathrm{TV}(\mu_F, g) \leq \widehat{B}\} \geq 1-\delta \qquad \text{for every sampler in the stated class,} \tag{28}$$

$$\Pr_G\{\widehat{B} \leq \rho\} \geq 1-\delta. \tag{29}$$

We allow measurable adaptive queries, independent private randomness, interleaving and revisiting old seeds, and sequential stopping. A seed is counted when its first actual sampler call occurs. Reference termination is almost sure; nontermination under an alternative is assigned report 1. The three resources are the expected reference number of opened seeds, the expected reference number of actual calls, and a deterministic cap on actual calls valid for all sampler responses and private coin outcomes. Optimization of these three resources is understood separately.

Theorem B.1 (Uniform exact-marginal resource law)**.** *Fix $n \geq 4$ and $0 < \eta_0 < 1/n$. There are constants $c, C, \rho_* > 0$, depending only on n, η_0, with the following property. Suppose $p \neq q$ and $p_i, q_i \geq \eta_0$ for every i. Order $r_i = q_i/p_i$ increasingly, breaking ties by any fixed rule, and set*

$$D = \log(r_{n-1}/r_2), \qquad \Psi(\rho,\delta,D) = \log(1/\delta)\big(\rho^{-1} + D\rho^{-4/3}\big).$$

For $0 < \rho \le \rho_*$ and $0 < \delta \le 1/4$, every honest, reference-complete certificate for the class (27) has expected reference seed count at least $c\Psi$. There is such a certificate using only full-support inputs and having an all-sampler deterministic call cap at most $C\Psi$. Consequently all three optimal resources have order Ψ, uniformly in these targets, accuracies, and confidence levels.

At $p = q$, repeated-input equality and exact marginals determine the target table without observations. The theorem's uniformity is restricted to coordinates bounded away from zero. The resource bounds count oracle calls; preprocessing, arithmetic, and bit complexity are separate costs.

We first prove the adaptive lower bound and construct the required full-input alternatives. We then build the finite affine statistic and its randomized evaluator, followed by the confidence and truncation arguments.

B.2 A STOPPED HIDDEN-SEED COMPARISON

Lemma B.2 (Latent entropy under arbitrary stopping)**.** *Suppose seed laws P, Q use the same measurable response map, and $d_* = D_{\mathrm{KL}}(P\|Q) < \infty$. For a policy in the preceding model, terminating almost surely under P, let T count opened seeds. Then*

$$D_{\mathrm{KL}}\big(\mathcal{L}_P(\widehat{B})\|\mathcal{L}_Q(\widehat{B})\big) \le d_*\mathbb{E}_P T. \tag{30}$$

If the reference target law has distance greater than ρ from the alternative target law, honesty and completeness imply $d_*\mathbb{E}_P T \ge \kappa_\delta$, where $\kappa_\delta = \mathrm{kl}(1-\delta, \delta)$.

Proof. Include the complete private random tape W, whose law is identical under both hypotheses, and reveal for this proof the whole latent seed at first use. This enlarges the experiment without changing the procedure. Given $W, Z_1, \dots, Z_{j-1}$, whether seed j is ever opened is determined before its draw. This remains true if the intervening execution revisits old seeds arbitrarily or never terminates. Record an absorbing empty symbol for every unopened slot. The entropy chain rule for the first m slots gives

$$D_{\mathrm{KL}}(P|_{\mathcal{F}_m}\|Q|_{\mathcal{F}_m}) = d_* \sum_{j=1}^{m} \Pr_P(T \ge j) = d_*\mathbb{E}_P(T \wedge m).$$

Empty slots contribute zero. Relative entropy converges along increasing sigma-fields; monotone convergence on the right and data processing to the report prove (30), including infinite expected cost. The argument uses relative-entropy convergence rather than optional stopping of the likelihood ratio. The event $\{\widehat{B} \le \rho\}$ has probabilities at least $1-\delta$ and at most δ under the reference and alternative. Binary data processing proves the last assertion. For $\delta \le 1/4$, $\kappa_\delta \asymp \log(1/\delta)$ with absolute constants. □

B.3 UNIFORM LOWER ALTERNATIVES AND SIMULTANEOUS MARGINAL REPAIR

Put $U_0 = 1 - (n-1)\eta_0$, $\ell_0 = \log(U_0/\eta_0)$, and $D_0 = 2\ell_0$. Thus $|\log r_i| \le \ell_0$ and $D \le D_0$. Constants in the exact-rate proof depend only on n, η_0 unless specified otherwise.

The inverse-accuracy term. Because $p \ne q$, $r_n > r_1$. The reference cell $(n, 1)$ is impossible: the two winning inequalities would require both $E_n/E_1 < p_n/p_1$ and $E_n/E_1 > q_n/q_1$. An independent target coupling assigns this cell probability $p_n q_1 \ge \eta_0^2$. It has a measurable universal exact-marginal extension: use independent uniforms for the two distinct targets and a third common inverse-CDF uniform at all other inputs. Let its target distance from g be $d \ge \eta_0^2$. Mixing it with Gumbel with probability $\lambda = 2\rho/d$ gives distance exactly 2ρ. Use an augmented seed containing the branch indicator and both response mechanisms. Under the reference the branch indicator is zero, and under the alternative it is Bernoulli(λ); all other components have the same law. For $\rho \le \eta_0^2/4$ the reverse seed entropy is at most $-\log(1-\lambda) \le 4\rho/\eta_0^2$. Lemma B.2 proves the uniform lower bound $c\kappa_\delta/\rho$.

A quotient box with only one shrinking side. Assume $D > 0$. Use the low pair $\{1, 2\}$ and high pair $\{n-1, n\}$. Write the exponential vector as $e^u x$, where

$$x = \big(p_1 e^{d_1}, p_2 e^{-d_1}, (e^{z_l})_{l \notin \{1,2,n-1,n\}}, e^t q_{n-1} e^{d_2}, e^t q_n e^{-d_2}\big). \tag{31}$$

The ordering of unused coordinates is fixed. Set

$$\eta = \min\{1, D/2\}, \qquad t_0 = -\tfrac{1}{2}(\log r_2 + \log r_{n-1}), \qquad t_\pm = t_0 \pm \eta/2,$$
$$A_l = \log\max\{p_l, e^{t_+} q_l\} + 1, \quad Q = [t_-, t_+] \times \prod_{l \notin \{1,2,n-1,n\}} [A_l, A_l + 1]. \tag{32}$$

An empty product has volume one. In particular $\mathrm{vol}(Q) = \eta$, independently of n. At $d_1 = d_2 = 0$, the p winner is in the low pair and the q winner is in the high pair. The log gap to the other pair is at least $\eta/2$ and the gap to unused labels is at least one. This persists for $|d_1|, |d_2| \le \eta/16$.

The Jacobian of the logarithmic coordinates (u, d_1, d_2, t, z) has absolute determinant 4. Integrating the common scale u therefore gives quotient density

$$f(d_1, d_2, t, z) = \frac{4(n-1)! \prod_l x_l}{(\sum_l x_l)^n}. \tag{33}$$

Indeed the scale integral is $\int_{\mathbb{R}} e^{nu} \exp(-e^u \sum_l x_l)\, du = (n-1)!/(\sum_l x_l)^n$. On the indicated closed box,

$$x_- := \eta_0 e^{-\ell_0 - 1} \le x_l \le U_0 e^{\ell_0 + 3} =: x_+.$$

Consequently

$$\min_l \frac{x_l}{\sum_j x_j} \ge a_* := \frac{x_-}{n x_+}, \qquad f \ge b_* := \frac{4(n-1)! x_-^n}{(n x_+)^n} > 0. \tag{34}$$

These constants stay positive as D tends to zero.

The checkerboard tilt. For $0 < w \le \eta/16$, define

$$H_w = \frac{b_*}{f} \mathbf{1}\{|d_1|, |d_2| < w,\ (t, z) \in Q\} \operatorname{sign}(d_1) \operatorname{sign}(d_2), \qquad \alpha = b_* \eta. \tag{35}$$

Density cancellation yields

$$|H_w| \le 1, \qquad \mathbb{E}_G H_w = 0, \qquad \mathbb{E}_G H_w^2 \le 4\alpha w^2. \tag{36}$$

For $0 < \theta \le 1/2$, tilt the quotient density by $1 + \theta H_w$. The signs of d_1, d_2 determine the winners within the two target pairs. Each of their four cells changes by $\theta\alpha w^2(+1, -1, -1, +1)$, with a suitable common ordering of the four cells. Thus both target marginals stay exact and the target distance is $2\theta\alpha w^2$.

Lemma B.3 (Uniform probe control)**.** *Let m_s be the marginal at an arbitrary input s under this tilt, and put $J = \binom{n-2}{2}$. Then*

$$\sup_s \mathrm{TV}(m_s, s) \le \frac{8\alpha J}{\eta} \theta w^3. \tag{37}$$

Moreover, $m_{s,l} \ne s_l$ implies $s_l \ge a_*$.

Proof. Regard the low pair, high pair, and each unused label as $n - 2$ blocks. At $d_1 = d_2 = 0$, each supported block's best log score is affine in (t, z): the low score is constant, the high score has coefficient one in t, and each singleton score has coefficient one in its own z_l. A difference between distinct block scores has a coordinate coefficient 1 or -1. Its slab of absolute width $2w$ occupies at most the fraction $4w/\eta$ of Q, since every side of Q has length at least η. There are at most J block comparisons. Outside their slabs, the winning block stays fixed for all $(d_1, d_2) \in (-w, w)^2$: each block score moves by at most w. Its output then depends on at most one of d_1, d_2, so integration against the sign of the other coordinate vanishes.

For any label subset $A \subseteq [n]$, subtract $1/2$ from $\mathbf{1}\{G_s \in A\}$ on the remaining integral. The constant integrates to zero, and the centered magnitude is at most $1/2$. The absolute marginal change for this subset is therefore at most

$$\theta b_* \frac{1}{2} (2w)^2 \frac{4Jw}{\eta} \mathrm{vol}(Q) = \frac{8\alpha J}{\eta} \theta w^3.$$

Taking the supremum over subsets proves (37), including inputs on the simplex boundary. Finally, if label l wins anywhere on the perturbation support, the inequalities $s_j \le s_l E_j / E_l$ imply $s_l \ge E_l / \sum_j E_j \ge a_*$. A smaller nominal coordinate can never be affected by the perturbation. □

One branch repairs every input. Write $K_m = 8\alpha J/\eta$ and take

$$\lambda = K_m \theta w^3/a_*, \qquad r^{\mathrm{rep}}_{s,l} = \frac{s_l - (1-\lambda)m_{s,l}}{\lambda}. \tag{38}$$

Assume $\lambda \le 1/2$. These coordinates sum to one. If $m_{s,l} - s_l \le 0$, their nonnegativity is immediate; otherwise $s_l \ge a_*$ and the numerator is at least $\lambda a_* - (m_{s,l} - s_l) \ge 0$. At p, q, the repair probabilities are exactly p, q.

On the repair branch draw a pair (T_p, T_q) with the tilted target joint law, and an independent uniform U. At the two targets return this pair; at every other input use inverse-CDF sampling with probabilities r^{rep}_s and the same U. Mix this response map with the tilted race using branch probabilities $\lambda, 1-\lambda$, shared across inputs. Every marginal is now s, while both branches have the same tilted target table. Repeated-input equality is respected. Joint measurability follows from measurability of the integrated race indicators $s \mapsto m_s$, inverse-CDF sampling, and the two singleton target cases. The construction is independent of the probe family and query policy.

For the entropy comparison, use augmented seeds (B, E, T_p, T_q, U) with the same response map under both hypotheses. Under the reference $B = 0$ surely and the race vector has its reference law; under the alternative B is Bernoulli(λ) and the vector has the tilted law. The remaining variables have the same independent repair law under both hypotheses. Since $-\log(1+x) \le -x + x^2$ for $|x| \le 1/2$, their reverse entropy satisfies

$$D_{\mathrm{KL}}(P_G \| P_F) \le 2\lambda + \theta^2 \mathbb{E}_G H_w^2. \tag{39}$$

Choose $\theta = \rho/(b_* \eta w^2)$, so the target distance is 2ρ. The required feasibility conditions are

$$w \le \eta/16, \qquad \rho \le b_* \eta w^2/2, \qquad \rho w \le a_* \eta/(16J). \tag{40}$$

Under these conditions,

$$D_{\mathrm{KL}}(P_G \| P_F) \le \frac{16 J \rho w}{a_* \eta} + \frac{4\rho^2}{b_* \eta w^2}. \tag{41}$$

Let $A_0 = \max\{2, D_0\}$ and $c_0 = \min\{16^{-3}, (b_*/2)^3, (a_*/(16J))^{3/4}\}$. Because $\eta \ge D/A_0$, the choice $w = \rho^{1/3}$ is feasible whenever $\rho \le c_0 D^3/A_0^3$. Equation (41) is then at most $C\rho^{4/3}/D$. Lemma B.2 gives the lower term $c\kappa_\delta D \rho^{-4/3}$. Outside this range, $D\rho^{-4/3} \le A_0 c_0^{-1/3} \rho^{-1}$, so the mixture lower bound already controls both terms. Combining the two lower bounds proves the lower half of Theorem B.1, also at $D = 0$.

B.4 Four-cycle reduction and the single-switch anchor path

For a seed, write $P_i = \mathbf{1}\{X_p = i\}$, $Q_i = \mathbf{1}\{X_q = i\}$, and $Y_{ij} = P_i Q_j$. For $2 \le i \le n-1$, query the normalized maximum-anchor weights $\max\{p_l/p_i, q_l/q_i\}$ and let D_i be its label-i indicator. Also put $D_1 = Q_1$ and $D_n = P_n$. Under the reference,

$$D_i = Y_{ii} \quad \text{for all } i. \tag{42}$$

For an interior label, winning its maximum anchor is equivalent to winning both targets, by intersecting all the race inequalities. At an extreme ratio one target's inequalities contain the other's, proving the endpoint identities. Every anchor has positive weights.

Put $U_i = P_i - D_i$, $V_j = Q_j - D_j$, and define the affine table

$$\begin{aligned} W^0_{ii} &= D_i, & W^0_{1j} &= V_j \quad (2 \le j < n), \\ W^0_{in} &= U_i \quad (2 \le i < n), & W^0_{1n} &= U_1 - \sum_{j=2}^{n-1} V_j, \end{aligned} \tag{43}$$

with all other entries zero. It has total sum one on every categorical assignment. Define

$$\mathcal{I} = \{(i,j) : 2 \le i < j \le n-1,\ r_i < r_j\}, \qquad C_{ij} = e_{ij} - e_{1j} - e_{in} + e_{1n}. \tag{44}$$

The reference support rule says that an off-diagonal cell (i,j) is possible only when $r_i < r_j$. This follows by dividing its two winning inequalities. Subtract the internal cells from their rows and columns, and use (42); the remaining entries are exactly those of (43). Hence

$$Y - W^0 = \sum_{(i,j) \in \mathcal{I}} Y_{ij} C_{ij} \qquad (G\text{-a.s.}). \tag{45}$$

This also covers ratio ties. The cycles are independent because each has a unique internal entry. The identity holds on the reference support.

Let $s(c)$ normalize $\max(p, q/c)$, and put $T_* = \min_l E_l/p_l$, $S_* = \min_l E_l/q_l$. Then

$$\frac{E_l}{\max(p_l, q_l/c)} = \min\{E_l/p_l, cE_l/q_l\}, \qquad \min_l \frac{E_l}{\max(p_l, q_l/c)} = \min\{T_*, cS_*\}. \tag{46}$$

Thus this path selects only the two target winners. If they differ, it switches exactly once at $c_* = T_*/S_* \in (r_{G_p}, r_{G_q})$, from the q winner below the switch to the p winner above it. With $\delta_k = \log(r_{k+1}/r_k)$, define

$$\mathcal{T} = \{(i,j,k) : (i,j) \in \mathcal{I},\ i \le k < j,\ \delta_k > 0\}, \qquad R_{ijk} = \mathbf{1}\{G_{s(r_{k+1})} = i, G_{s(r_k)} = j\}.$$

The single-switch property proves the disjoint decomposition

$$Y_{ij} = \sum_{k:(i,j,k)\in\mathcal{T}} R_{ijk}, \qquad |\mathcal{T}| \le S_n := \binom{n-1}{3}, \qquad \sum_{\mathcal{T}} \delta_k \le B_n D, \tag{47}$$

where $B_n = \lfloor (n-2)^2/4 \rfloor$. Indeed gap k occurs $(k-1)(n-k-1)$ times, and only $2 \le k \le n-2$ occur. The outer gaps therefore require no non-affine coefficients.

B.5 A TWO-BLOCK INTEGRAL AND FINITE AFFINE QUADRATURE

Fix a triple $(i,j,k) \in \mathcal{T}$, set $\Delta = \delta_k$, and split the labels into $L_k = [k]$ and $H_k = \{k+1, \dots, n\}$. Define

$$a = \min_{l \in L_k \setminus \{i\}} \log \frac{E_l/p_l}{E_i/p_i}, \qquad b = \min_{l \in H_k \setminus \{j\}} \log \frac{E_l/q_l}{E_j/q_j},$$
$$t = \log \frac{r_{k+1} E_j/q_j}{E_i/p_i}. \tag{48}$$

Both minima are nonempty. For parameters v, w, β, use normalized weights

$$\omega_l(v,w,\beta) = \begin{cases} p_i, & l = i, \\ p_l e^v, & l \in L_k \setminus \{i\}, \\ (q_j/r_{k+1}) e^\beta, & l = j, \\ (q_l/r_{k+1}) e^{\beta + w}, & l \in H_k \setminus \{j\}. \end{cases} \tag{49}$$

Let $J(v,w,\beta)$ and $K(v,w,\beta)$ denote its label-j and label-i indicators, respectively. These are defined for any response map. Write $\mathsf{H}(x) = \mathbf{1}\{x > 0\}$. Direct race comparisons give, under G,

$$J(v,w,\beta) = \mathsf{H}(b-w)\mathsf{H}(\beta - t)\mathsf{H}(a - v + \beta - t),$$
$$K(v,w,\beta) = \mathsf{H}(a-v)\mathsf{H}(t-\beta)\mathsf{H}(t + b - \beta - w),$$
$$R_{ijk} = \mathsf{H}(a)\mathsf{H}(b)\mathbf{1}\{0 < t < \Delta\}. \tag{50}$$

The inputs $(0,0,0)$ and $(0,0,\Delta)$ are precisely the two slice endpoints in (47).

For $0 < h \le 1$, let $S_h(x) = h^{-1} \int_0^h \mathsf{H}(x - v)\, dv$ and define

$$F_a(w) = \frac{1}{h}\int_0^h [J(v,w,\Delta) - J(v,w,0)]\, dv + \frac{1}{h}\int_0^\Delta [J(h,w,\beta) - J(0,w,\beta)]\, d\beta,$$
$$F_b(v) = \frac{1}{h}\int_0^h [K(v,w,0) - K(v,w,\Delta)]\, dw + \frac{1}{h}\int_0^\Delta [K(v,h,\beta) - K(v,0,\beta)]\, d\beta,$$
$$F_{ab} = \frac{1}{h}\int_0^h F_a(w)\, dw, \qquad \mathcal{V}_h = F_a(0) + F_b(0) - F_{ab}. \tag{51}$$

We use these integrals to derive the reference approximation. Averaging J over v gives $\mathsf{H}(b - w)\mathsf{H}(\beta - t)S_h(a + \beta - t)$. Its continuous derivative in β is canceled by the second integral in F_a; its jump at $\beta = t$ remains. Averaging K over w and using the reversed β direction gives the corresponding identity for F_b. Explicitly, writing $C_\Delta(t) = \mathbf{1}\{0 < t < \Delta\}$,

$$F_a(w) = S_h(a)\mathsf{H}(b-w)C_\Delta(t), \qquad F_b(v) = \mathsf{H}(a-v)S_h(b)C_\Delta(t),$$
$$R_{ijk} - \mathcal{V}_h = C_\Delta(t)[\mathsf{H}(a) - S_h(a)][\mathsf{H}(b) - S_h(b)]. \tag{52}$$

The last expression is in $[0,1]$ and supported on $0 < a, b < h$, $0 < t < \Delta$.

Lemma B.4 (Local affine approximation)**.** *The coefficient mass in* (51) *is at most* $6(1+\Delta/h)$*, and*

$$\mathbb{E}_G(R_{ijk}-\mathcal{V}_h)^2 \le (k-1)(n-k-1)e^{-3}\Delta h^2. \tag{53}$$

If $\Delta = J_*d$ *and* $h = Md$ *for positive integers* J_*, M*, replacing each integral by its length-weighted midpoint sum, with the same inner sum at every outer midpoint, gives a finite affine statistic* V^d *with the same coefficient bound and*

$$|V^d - \mathcal{V}_h| \le 13d/h \qquad (G\text{-a.s.}). \tag{54}$$

Proof. Each of $F_a(0), F_b(0), F_{ab}$ has coefficient mass at most $2(1+\Delta/h)$. For the moment bound, if $a, b \in (0,h)$, some low rival α and high rival β attain gaps in this interval. Conditional on $\log E_i$, the independent variables $\log E_\alpha, \log E_j, \log E_\beta$ have densities bounded by e^{-1}. Their affine transformation to the two individual gaps and t has determinant of absolute value one. Each choice of rivals therefore contributes at most $e^{-3}\Delta h^2$. Sum over the $(k-1)(n-k-1)$ choices and use (52).

For quadrature, every scalar integrand in a single integral is a monotone Boolean function of its integration variable. A midpoint sum of mesh d has error at most d for such a function. The four scalar terms in F_a or F_b give error at most $4d/h$. In F_{ab}, these inner errors average to at most $4d/h$; its exact outer integrand is a constant of magnitude at most one times $\mathsf{H}(b-w)$, giving a further d/h. Summing the three errors proves (54). □

For the exact-rate construction put $H = \rho^{1/3}$. Choose ρ_* so that $H \le 1/4$ and $H\sqrt{D_0} \le 13$. For each positive internal gap use

$$J_k = \left\lceil 13\sqrt{\delta_k}/H^2 \right\rceil, \quad d_k = \delta_k/J_k, \quad M_k = \lceil H/d_k \rceil, \quad h_k = M_k d_k. \tag{55}$$

Then $H \le h_k \le 2H$ and $d_k \le h_k^2\sqrt{\delta_k}/13$. Lemma B.4 gives, for each slice,

$$L_{ijk} := 6(1+\delta_k/h_k), \qquad \mathrm{Var}_G(R_{ijk} - V^d_{ijk}) \le 2(1+B_n e^{-3})\delta_k h_k^2. \tag{56}$$

For every exact-marginal sampler, the mean of this finite statistic equals its known coefficient-weighted sum of input probabilities. Calibration is exact because it uses the finite sum; quadrature error enters only the reference approximation.

B.6 Exact finite groups on stable branches

We now evaluate each V^d with bounded expected reference calls, despite its potentially large coefficient list. In this subsection suppress the slice index and use the grid $\Delta = J_*d$, $h = Md$, with $v_l = w_l = (l+1/2)d$ and $\beta_j = (j+1/2)d$.

Write $k(w,\beta) = K(0,w,\beta)$, $f(x) = k(0,x)$, and $C^d = F^d_a(0) - F^d_{ab}$. Expanding the midpoint sums expresses C^d as the following groups, each with coefficient $1/M^2$:

$$\begin{aligned} A_{lm} &= J(v_l,0,\Delta) - J(v_l,0,0) - J(v_l,w_m,\Delta) + J(v_l,w_m,0),\\ B_{jm} &= J(h,0,\beta_j) - J(0,0,\beta_j) - J(h,w_m,\beta_j) + J(0,w_m,\beta_j). \end{aligned} \tag{57}$$

The indices are $l, m < M$ and $j < J_*$. Each group has absolute value at most two and uses at most four queries. Under Gumbel every group vanishes on both $b \le 0$ and $b \ge h$: in the first case all its J terms vanish; in the second they do not depend on w.

The remaining term $F^d_b(0)$ has two exact representations:

$$\begin{aligned} F^d_b(0) = f(0) - f(\Delta) &+ \frac{1}{M}\sum_{l<M}\big([k(w_l,0)-f(0)] - [k(w_l,\Delta)-f(\Delta)]\big)\\ &+ \frac{1}{M}\sum_{j<J_*}[k(h,\beta_j)-f(\beta_j)], \end{aligned} \tag{58}$$

$$\begin{aligned} F^d_b(0) = \frac{1}{M}&\sum_{l<M}\big([k(w_l,0)-f(w_l)] - [k(w_l,\Delta)-f(\Delta+w_l)]\big)\\ &+ \frac{1}{M}\sum_{j<J_*}[k(h,\beta_j)-f(\beta_j+h)]. \end{aligned} \tag{59}$$

The first is the midpoint definition with added and subtracted endpoint terms. To verify the second for arbitrary response maps, put $f_a = f((a+1/2)d)$ and use the exact cancellation

$$\sum_{l=0}^{M-1} (f_l - f_{J_*+l}) + \sum_{j=0}^{J_*-1} (f_{M+j} - f_j) = 0. \tag{60}$$

Write both differences as differences of partial sums of the same sequence. The cancellation follows from repeated-input consistency: identical indices represent identical normalized inputs and outputs.

Under G, if $b \geq h$ then $k(w, \beta) = f(\beta) = \mathsf{H}(a)\mathsf{H}(t - \beta)$ on this grid. Every pair in (58) vanishes and its baseline $f(0) - f(\Delta)$ equals R_{ijk}. If $b \leq 0$, then $k(w, \beta) = f(\beta + w)$; every pair in (59) vanishes, and $R_{ijk} = 0$. Together with (57), either representation is

$$V^d = B_{\text{base}} + \sum_g \gamma_g B_g, \qquad \sum_g |\gamma_g| b_g = 4 + 3\Delta/h \leq L, \tag{61}$$

where $|B_g| \leq b_g$, $b_g = 2$ for a four-term group and $b_g = 1$ for a pair, and $L = 6(1 + \Delta/h)$. Both baseline indicators are queried under every sampler.

Given either representation, sample a fresh private group index with $\pi_g = |\gamma_g| b_g / M_{\text{grp}}$, where M_{grp} is the mass in (61), and return

$$\widetilde{V} = B_{\text{base}} + \gamma_G B_G / \pi_G. \tag{62}$$

Conditional on all potential input outputs, its coin expectation is exactly V^d for every sampler. On a correct stable reference branch every sampled group is zero, so this audit has no conditional variance and returns R_{ijk} exactly.

B.7 A FULL-SUPPORT PRECISION PILOT AND UNBIASED CORRECTIONS

Lemma B.5 (Finite categorical pilot)**.** *For $0 < \epsilon \leq 1$, there is a finite adaptive pilot using only positive inputs, of deterministic depth $O_n(1 + \log(1/\epsilon) + \log(1 + \log(1/\epsilon)))$, that produces Boolean predictions $\widehat{I}_i(s)$ at every specified finite set of inputs. For each fixed input and label separately,*

$$\Pr_G\{\widehat{I}_i(s) \neq \mathbf{1}\{G_s = i\}\} \leq C_n \epsilon. \tag{63}$$

The predictions are defined and Boolean under arbitrary sampler responses.

Proof. First describe a binary-support pilot. Put $T_\epsilon = \log(2(n-1)/\epsilon)$ and $k_\epsilon = \lceil \log_2(2T_\epsilon/\epsilon) \rceil$. For each $i < n$, a query on $\{i, n\}$ with weights $(e^t, 1)$ compares $D_i = \log(E_i/E_n)$ with t. Query both endpoints $\pm T_\epsilon$. When they bracket D_i, make k_ϵ bisections and return the midpoint $\widehat{D}_i$. Otherwise use a fixed bounded default. An output outside the binary pair also invokes that default. Put $\widehat{D}_n = 0$ and predict a winner by minimizing $\widehat{D}_i - \log s_i$ over its support, with fixed tie breaking.

Every D_i is logistic, so the probability that any coordinate lies outside the initial interval is at most ϵ. Otherwise $\max_i |\widehat{D}_i - D_i| \leq \epsilon$. A wrong prediction at a fixed s requires an unordered pair with

$$|\log(E_i/E_j) - \log(s_i/s_j)| \leq 2\epsilon.$$

The logistic density is at most $1/4$. A union over pairs bounds the error by $(1 + \binom{n}{2})\epsilon$ for each fixed input separately.

To make every actual query positive, let $B_\epsilon = (n-1)(2 + 2^{k_\epsilon} - 1)$ bound the number of all potential binary nodes. Replace a node b by $b^\zeta = (1-\zeta)b + \zeta u_n$, where u_n is uniform and $\zeta = \epsilon/(2B_\epsilon)$. Run the same decision tree, retaining the bounded default for outputs outside the original pair.

For two normalized race inputs x, y at TV distance d, their common-race disagreement probability is at most $2d/(1+d)$. Indeed, put $m_i = \min(x_i, y_i)$ and $w_i = \max(x_i, y_i)$. The event $E_i/m_i < \min_{j \neq i} E_j/w_j$ forces agreement at i and has probability at least $m_i / \sum_j w_j$. These events are disjoint over i, giving agreement probability at least $(1-d)/(1+d)$. Thus the chance that any original and softened node disagree is at most $2B_\epsilon\zeta = \epsilon$. Outside this event, induction down the tree gives identical predictions. This adds only ϵ to the error bound and does not change its depth or finiteness. □

For a finite affine expression $V = c_0 + \sum_r a_r I_r$, let $L = \sum_r |a_r|$ after discarding zero coefficients. After the pilot, sample a fresh independent index Q with probabilities $|a_r|/L$ and return

$$\widetilde{V} = c_0 + \sum_r a_r \widehat{I}_r + L \operatorname{sign}(a_Q)(I_Q - \widehat{I}_Q). \tag{64}$$

If $L = 0$, return c_0. There is at most one actual correction query. Conditional on the outputs at the finite union of all possible inputs, the pilot is fixed and the correction index is independent. Consequently

$$\mathbb{E}_Q(\widetilde{V} \mid \mathcal{X}) = V \tag{65}$$

under every sampler. Also $|\widetilde{V} - c_0| \le 2L$, and under the reference

$$\mathbb{E}_G \operatorname{Var}_Q(\widetilde{V} \mid E) \le L^2 \sum_r \frac{|a_r|}{L} \Pr_G(I_r \ne \widehat{I}_r) \le C_n L^2 \epsilon. \tag{66}$$

The variance uses a coefficient-weighted average of the per-input error bounds, avoiding a union bound over quadrature terms.

B.8 Full-support gates and the uniform row bounds

For each slice (i, j, k) use its bandwidth h_k from (55). In a first gate input, give every low label weight H^4, label j weight one, and each high rival l weight q_l/q_j. A high-rival output selects the lower representation (59). If the first output is j, query again after multiplying all high-rival weights by e^{h_k}; another j selects the upper representation (58). All other outcomes mark the slice for a pilot. Both inputs have full support.

Under Gumbel the two stable decisions imply, respectively, $b \le 0$ and $b \ge h_k$. Increasing high rivals after a first winner j cannot create a low winner. The marked event $\mathcal{R}_{ijk}$ is therefore contained in $\{0 < b < h_k\}$ or the event that the first winner is low. A union over individual shifted logistic gaps and the exact first-gate marginal give

$$\Pr_G(\mathcal{R}_{ijk}) \le (n - k - 1)h_k/4 + kH^4 \le 2nH. \tag{67}$$

The small positive low weights can cause a pilot, but never an incorrect stable reference decision.

Run all gates on a row. On stable slices use fresh group audits (62). If any slice is marked, run one shared pilot from Lemma B.5 with precision $\epsilon = H^4$ and use (64) on each marked slice. Condition on all outputs in the finite union of the gates, pilot tree, and possible audit inputs. Each branch has the same conditional coin expectation, V^d_{ijk}. Branch selection therefore preserves the known mean of this fixed affine statistic under every sampler.

Define the same pilot counterfactually on every reference row. Multiplying its error indicator by a marked-event indicator can only decrease its expectation. Stable branches have zero conditional variance. Total variance, (56), and (66) give

$$\begin{aligned} \operatorname{Var}_G(R_{ijk} - \widetilde{V}_{ijk}) &\le C_n \delta_k h_k^2 + C_n L_{ijk}^2 H^4 \\ &\le C_{n,\eta_0}(H^3 + \delta_k H^2). \end{aligned} \tag{68}$$

Here $H \le 1$, $h_k \in [H, 2H]$, and $\delta_k^2 \le D_0 \delta_k$. The pilot error remains unconditional in this calculation.

Define the audited affine table and residual table by

$$\widetilde{W} = W^0 + \sum_{(i,j,k)\in\mathcal{T}} \widetilde{V}_{ijk} C_{ij}, \qquad Z = Y - \widetilde{W}. \tag{69}$$

Every entry of $\widetilde{W}$ has an exactly known all-sampler mean m_z, obtained from its fixed underlying affine expansion. Its entries sum to one on every response and coin outcome. Under Gumbel, $Z = \sum_{\mathcal{T}}(R_{ijk} - \widetilde{V}_{ijk})C_{ij}$. Using $\operatorname{Var}(\sum_{l=1}^m X_l) \le m \sum_l \operatorname{Var} X_l$ handles all correlations from the shared pilot. We obtain

$$\operatorname{Var}_G Z_z \le \sigma_H^2 := C_{n,\eta_0}(H^3 + DH^2), \qquad \operatorname{osc}(Z_z) \le R_H := C_n(1 + D/H). \tag{70}$$

The oscillation bound is valid on every sampler response and private coin outcome: use $|\widetilde{V}_{ijk}| \le 2L_{ijk}$, (47), and the bounded entries of W^0. In particular, all target cells, including reference-impossible cells, remain in the reported bound.

B.9 HONEST CERTIFICATION AND A DETERMINISTIC CALL CAP

Honesty follows from bounded-variable concentration; reference variance is used separately to prove completeness. For $N \geq 2$ iid rows and $K = n^2$ target cells, write

$$\widehat{d}_z = \overline{Z}_z - (g_z - m_z), \qquad S_z^2 = \frac{1}{N(N-1)} \sum_{r<s} (Z_{r,z} - Z_{s,z})^2,$$
$$q_z = \sqrt{2S_z^2 u/N} + \frac{7R_H u}{3(N-1)}, \qquad \widehat{B} = \min\left\{1, \frac{1}{2}\sum_z (|\widehat{d}_z| + q_z)\right\}. \tag{71}$$

Here $u = \log(16K/\delta)$. We use the standard empirical Bernstein inequality for iid variables in an interval of length R: with the unbiased sample variance S^2, the two-sided mean error is at most $\sqrt{2S^2u/N} + 7Ru/[3(N-1)]$ except with probability $4e^{-u}$. We also use its sample-standard-deviation bound $S \leq \sigma + R\sqrt{2u/(N-1)}$, except with probability e^{-u}. These are the bounded-variable empirical Bernstein inequalities of Maurer & Pontil (2009, Theorems 4 and 10), after rescaling the interval.

Conditional audit calibration and exact marginals give $\mathbb{E}_F \widehat{d}_z = \mu_F(z) - g_z$. A union over the K cells proves honesty of (71) except with probability $4Ke^{-u}$. Under G, the standard-deviation events and (70) give

$$q_z \leq \sigma_H \sqrt{2u/N} + \frac{13R_H u}{3(N-1)}. \tag{72}$$

Choose

$$N = 2 + \left\lceil u \left[\frac{8K^2\sigma_H^2}{\rho^2} + \frac{26KR_H}{3\rho} \right] \right\rceil. \tag{73}$$

Each term in (72) is at most $\rho/(2K)$. On the simultaneous event, $|\widehat{d}_z| \leq q_z$ under the reference, so $\widehat{B} \leq \sum_z q_z \leq \rho$. Its failure probability is at most $5Ke^{-u} < \delta/2$. Substituting $H = \rho^{1/3}$ yields

$$N = O_{n,\eta_0}\left(\log(1/\delta)[\rho^{-1} + D\rho^{-4/3}]\right). \tag{74}$$

Each row has at most $C_0 = n + 8S_n$ non-pilot calls: the targets and diagonal anchors, and a bounded number of gate, baseline, and group or scalar-audit calls per slice. The union of marked events has probability at most $C_g H$, where $C_g = 2nS_n$. The shared pilot has deterministic depth $P_H = O_n(1 + \log(1/H) + \log(1 + \log(1/H)))$. Thus the untruncated expected reference call count is at most $N(C_0 + C_g H P_H) = O_n(N)$.

For a deterministic all-sampler cap, stop and report 1 before running more than

$$B_{\text{pilot}} = \lceil (e-1)C_g NH + \log(2/\delta) \rceil \tag{75}$$

pilots. In the untruncated reference experiment the marked-row indicators are iid. If their sum is M, then $\mathbb{E}_G e^M \leq \exp((e-1)C_g NH)$. Markov's inequality bounds the probability of truncation by $\delta/2$. Reporting 1 preserves honesty under every sampler: couple truncated and untruncated executions until truncation, and note that a truncation report is always valid. Combining the failure budgets proves completeness at level $1 - \delta$.

The deterministic call cap is

$$C_0 N + B_{\text{pilot}} P_H = O_{n,\eta_0}(N). \tag{76}$$

Indeed HP_H and ρP_H are uniformly bounded, while $N \geq c\log(1/\delta)/\rho$. At $D = 0$ the slice set is empty and $Y = W^0$ under the reference. Use the same report with reference variance zero and a constant raw range. It uses $O_n(\log(1/\delta)/\rho)$ rows and a fixed bounded number of calls per row, with no pilot.

Every executed input has full support: this is immediate for targets, anchors, local groups, and gates, and follows from the softening construction for pilots. All potential inputs form a finite set selected before any data are collected. Finally, each opened seed requires one call, and a deterministic call cap bounds the reference expectation. The lower bound already proved therefore applies to all three resources and completes Theorem B.1.

B.10 KNOWN MARGINAL TOLERANCE FOR A FIXED OBSTRUCTED PAIR

For a fixed pair with $D > 0$, the same construction yields a sharp marginal-tolerance transition. Constants in this section may depend on the targets. Define the class

$$\mathfrak{F}_\tau = \left\{ F : \sup_{s \in \Delta_{n-1}} \mathrm{TV}(\mathcal{L}(F(s,U)), s) \le \tau \right\}.$$

The bound τ is supplied to the certificate in advance and applies to the targets as well as all auxiliary inputs. The same result holds for the subclass additionally required to have exact target marginals.

Theorem B.6 (Fixed-pair tolerance law)**.** *Fix distinct positive targets with $r_2 < r_{n-1}$ and $n \ge 4$. There is $\rho_* > 0$, depending on the pair, such that for $0 < \rho \le \rho_*$, $0 < \delta \le 1/4$, and every known $0 \le \tau \le 1$, all three optimal resources over $\mathfrak{F}_\tau$ have order*

$$\log(1/\delta)\, \Phi(\rho, \tau), \qquad \Phi(\rho,\tau) = \min\{\rho^{-2}, \max\{\rho^{-4/3}, \tau^2 \rho^{-4}\}\}. \tag{77}$$

The lower bound permits the full adaptive stopping model, and its witnesses have exact target marginals. The upper bound has a deterministic call cap and uses only full-support inputs.

Proof. Lower bound. Fix the box and constants of Section B.3. Use the tilt (35) without repair and choose $\theta = \rho/(\alpha w^2)$. Whenever $w \le w_0 := \eta/16$ and $\rho \le \alpha w^2/2$, the target distance is 2ρ, the target marginals are exact, and

$$\sup_s \mathrm{TV}(m_s, s) \le A\rho w, \qquad D_{\mathrm{KL}}(P_G \| P_\theta) \le 4\rho^2/(\alpha w^2), \quad A = 8J/\eta. \tag{78}$$

For sufficiently small fixed ρ_*, if $\tau \ge A\rho^{4/3}$, choose $w = \min\{w_0, \tau/(A\rho)\} \ge \rho^{1/3}$. Feasibility follows after requiring $\rho_* \le \min\{w_0^3, (\alpha/2)^3\}$. This gives a member of $\mathfrak{F}_\tau$. Lemma B.2 yields

$$\mathbb{E}_G T \ge \frac{\alpha \kappa_\delta}{4} \min\{w_0^2 \rho^{-2}, \tau^2/(A^2\rho^4)\}.$$

For every τ, the exact-marginal subclass also gives $c\kappa_\delta \rho^{-4/3}$. Writing $x = \rho^{-4/3}$, $y = \tau^2\rho^{-4}$, and $v = \rho^{-2}$, with $x \le v$, combine these by $\max\{x, \min\{v, y\}\} = \min\{v, \max\{x, y\}\}$. When $\tau < A\rho^{4/3}$ the exact-class bound alone suffices. This proves the required lower bound uniformly in τ.

Free-bandwidth upper bound. Repeat the finite slice construction with an arbitrary small bandwidth $H = h$, rather than fixing it to $\rho^{1/3}$. Choose $h_0 > 0$ so that the grid inequalities hold for $0 < h \le h_0$. Since the pair is fixed, the resulting audited residual coordinates have reference variance at most Ch^2, all-response oscillation at most C/h, and underlying affine coefficient masses $L_z \le \Lambda/h$. The gate probability is at most $C_g h$, and the shared pilot depth is P_h as before. These bounds follow from (70) with D fixed and hold for any bandwidth in the stated range, independently of the desired statistical accuracy.

Let m_z be the nominal finite affine mean and write $e_z = \mathbb{E}_F \widetilde{W}_z - m_z$. Conditional audit calibration still holds under every alternative. Its underlying fixed affine expansion therefore implies

$$|e_z| \le \tau L_z \le \Lambda\tau/h. \tag{79}$$

This bound follows from the marginal guarantee at each fixed input in the affine expansion. Modify (71) by adding τL_z inside each cell's summand. Since $\mathbb{E}_F \widehat{d}_z = \mu_F(z) - g_z - e_z$, the empirical Bernstein event and (79) prove honesty. For any stochastic accuracy $a > 0$, taking

$$N = 2 + \left\lceil C \log(16K/\delta) \left(h^2/a^2 + 1/(ha) \right) \right\rceil \tag{80}$$

makes the reference report at most $a + K\Lambda\tau/(2h)$ except with probability at most $\delta/2$, by the same calculation as (72)–(73).

Set $A_1 = \max\{K\Lambda, 1\}$ and $h = \max\{\rho^{1/3}, A_1\tau/\rho\}$. If $h \le h_0$, use $a = \rho/2$. The bias is at most $\rho/2$ and

$$N = O\left(\log(1/\delta)[h^2/\rho^2 + 1/(h\rho)]\right) = O\left(\log(1/\delta) \max\{\rho^{-4/3}, \tau^2\rho^{-4}\}\right).$$

Reduce h_0 to at most one. This branch has $\tau \le \rho$, so its last maximum equals Φ. Cap the number of pilots at $\lceil (e-1)C_g N h + \log(2/\delta) \rceil$ and report 1 on truncation. The preceding exponential-moment argument gives truncation probability at most $\delta/2$. It also gives an $O(N)$ deterministic call cap: hP_h is bounded, and $h \ge \rho^{1/3}$ implies $N \ge c \log(1/\delta) h^{-4}$, so the extra $\log(2/\delta)P_h$ is absorbed by N.

If the chosen bandwidth exceeds h_0, use direct target estimation. For N iid target pairs put $e_N = \sqrt{\log(2K/\delta)/(2N)}$ and report $\min\{1, \mathrm{TV}(\widehat{\mu}, g) + K e_N/2\}$. A coordinate Hoeffding union bound proves honesty, and under G the report is at most $K e_N$. Thus $N = \lceil K^2 \log(2K/\delta)/(2\rho^2) \rceil$ suffices and costs $2N$ calls. Require $\rho_* \le h_0^3$. The direct branch then implies $\tau > h_0 \rho / A_1$, and consequently $\Phi \ge \min\{1, (h_0/A_1)^2\} \rho^{-2}$. Its cost has the required order as well. The upper construction is honest over the larger tolerance class, while the lower witnesses have exact target marginals, proving both versions stated in the theorem. □

For a fixed obstructed pair, tolerances of order at most $\rho^{4/3}$ preserve the exact-marginal rate, whereas any fixed positive tolerance eventually yields the direct-estimation rate.

C Certification at the critical marginal allowance

This appendix studies targets whose two outer logit gaps vanish, with marginal tolerance equal to one third of the target accuracy. We derive score-energy bounds and a finite-sample certificate. Constants may depend on the fixed number of labels, the compactness bound, and the two arithmetic parameters; they are uniform in $\epsilon, \rho, \delta, N, J$.

C.1 Target family and score energy

Fix $n \ge 4$ and an anchor $h \in \Delta_{n-1}$ with $h_i \ge \eta > 0$. This is equivalent, up to a fixed change of the compactness constant, to the main text's parametrization by p: its maximum anchor satisfies $h_i \ge p_i e^{-b\epsilon} \ge \eta e^{-b}$ for $\epsilon \le 1$; conversely $p_i \ge h_i e^{-b\epsilon}$. Let T_i be independent exponential random variables of rates h_i, and write $Y_i = \log T_i$. For raw logits $\ell \in \mathbb{R}^n$, set

$$w_i(\ell) = \frac{h_i e^{\ell_i}}{\sum_k h_k e^{\ell_k}}, \qquad G_\ell = \arg\min_i (Y_i - \ell_i).$$

Adding a constant to all logits changes neither the input nor the output. The same race is used at every input. Zero-weight inputs mean that the corresponding competitors are absent. Define

$$p = w(-b\epsilon e_n), \qquad q = w(-a\epsilon e_1), \qquad a, b > 0. \tag{81}$$

Assume that $\theta = b/(a+b)$ is irrational of bounded continued-fraction type. Equivalently, there is $c_\theta > 0$ such that $\|k\theta\|_{\mathbb{R}/\mathbb{Z}} \ge c_\theta / k$ for every positive integer k. The equivalent bounded-type condition for b/a will be used below. We restrict ϵ to a sufficiently small fixed interval.

Let $\mathcal{J} = \{2, \dots, n-1\}$ and let $c \in \{-2, 1\}^{\mathcal{J}}$ be nonconstant. On an actual target row define $C_c = \sum_{j \in \mathcal{J}} c_j 1_{\{X_p = X_q = j\}}$. For a square-integrable score H on the reference seed space, possibly with additional seed randomness, write

$$B_H(w) = \mathbb{E}_G[H e_{G_w}], \qquad K(H) = \mathbb{E}_G H^2, \qquad \|v\|_{\mathrm{TV}} = \tfrac{1}{2} \sum_i |v_i|.$$

The score is target-null if $\mathbb{E}_G H = 0$ and $B_H(p) = B_H(q) = 0$. Its all-input cap is $S_H = \sup_{w \in \Delta_{n-1}} \|B_H(w)\|_{\mathrm{TV}}$. Define

$$\begin{aligned} K_{\mathrm{all,c}}(\epsilon) &= \inf\{K(H) : H \text{ target-null}, \quad \mathbb{E}_G H C_c = 3,\ S_H \le 1\}, \\ K_*(\epsilon) &= \min_{c \text{ nonconstant}} K_{\mathrm{all,c}}(\epsilon). \end{aligned} \tag{82}$$

The gain is exactly three, rather than a free normalization parameter. The cap includes every input, including boundary inputs.

Theorem C.1 (Strict-cap energy interval)**.** *Uniformly on the target family above,*

$$c_1 \epsilon^{-11/7} \le K_*(\epsilon) \le C_1 \epsilon^{-2}. \tag{83}$$

The same bounds hold for each nonconstant detector separately. In particular, an energy law of order $\epsilon^{-3/2}$ is impossible under the strict all-input cap. The two exponents in (83) are not asserted to match.

The infimum in (82) is attained. Indeed, the constraints are closed affine hyperplanes or closed convex sets in $L^2(G)$: evaluation against each bounded categorical feature is a continuous linear functional. Their intersection is weakly closed, and a bounded minimizing sequence has a weakly convergent subsequence. Nonemptiness follows from the construction below. Strict convexity gives uniqueness for a fixed detector, up to equality almost surely.

C.2 A FULL CHECKERBOARD GIVES THE UPPER BOUND

For four labels, select middle labels j_-, j_+ with detector coefficients $-2, 1$. Remove the common log-time u and introduce quotient coordinates:

$$Y_1 = u + x, \quad Y_n = u + y, \quad Y_{j_+} = u + (-z)_+, \quad Y_{j_-} = u + z_+.$$

The reference law in these coordinates has a positive continuous density $r(x, y, z)$. On each fixed compact box this density is bounded above and below uniformly in h. This follows either by integration over u, or from the log-ratio density displayed in (86) below. Put $\alpha = a\epsilon/4$, $\beta = b\epsilon/4$, and fix $L = 1/4$. Let f_x^- and f_x^+ be uniform probability densities on $[-\alpha, 0]$ and $[0, L]$, respectively; define $f_y^\pm$ with β in place of α. Set

$$k(x, y, z) = (f_x^+ - f_x^-)(x)(f_y^+ - f_y^-)(y)(1_{[1,2]}(z) - 1_{[-2,-1]}(z)), \qquad H = k/r. \tag{84}$$

The signed density k has total mass zero. On the positive z slice the active middle label is j_+, and on the negative slice it is j_-. For either slice, the four (x, y) rectangles induce target cells $(j, j), (1, j), (j, n), (1, n)$ with signs $+, -, -, +$. The two target decreases exceed the respective negative widths, so these cell identifications hold on the entire rectangles. Thus the signed target table is

$$C^{\mathrm{cyc}}_{j_+} - C^{\mathrm{cyc}}_{j_-}, \qquad C^{\mathrm{cyc}}_j = e_{jj} - e_{1j} - e_{jn} + e_{1n}.$$

It has zero target marginals, and the detector gain is $1 - (-2) = 3$. Moreover, r is bounded below on the support and

$$\mathbb{E}_G H^2 = \int k^2/r \le C(\alpha^{-1} + L^{-1})(\beta^{-1} + L^{-1}) \le C\epsilon^{-2}.$$

It remains to prove the all-input cap one. A normalized signed product $(f_x^+ - f_x^-)(f_y^+ - f_y^-)$ is the expectation of the random rectangle $(\delta_{x_+} - \delta_{x_-}) \otimes (\delta_{y_+} - \delta_{y_-})$, with $x_- < x_+$ and $y_- < y_+$. At a fixed input the extreme race-time intervals are $[A_-, A_+]$ and $[D_-, D_+]$, in their natural increasing order. If they do not overlap, the four-corner output is zero. If they overlap, let E be the extreme label with larger lower endpoint, F the one with smaller upper endpoint, and set $l = \max(A_-, D_-)$, $v = \min(A_+, D_+)$. A middle competitor with time m and label j gives the four-corner output

$$0 \quad (m < l), \qquad e_j - e_E \quad (l < m < v), \qquad e_F - e_E \quad (m > v).$$

All these vectors, over all middle labels and times, belong to $\{e_i - e_E : i \in [n]\}$, a set of TV diameter one. Pairing z with $-z$ therefore gives TV at most one, even when the winning middle label and its time change between the slices. Integrating the rectangle and slice parameters preserves the bound. If an extreme is absent, its signed factor integrates to zero; if all middle labels are absent the two slices cancel. This proves the bound on the closed simplex.

For $n > 4$, keep the selected four labels and append quotient coordinates $Y_i - u \in [1, 2]$ for every remaining label, with uniform density on each such interval. These labels are dominated at the anchor and both targets. At an arbitrary input they can instead be included among the middle competitors in the diameter argument. The quotient density is again uniformly positive on the fixed box. This proves the upper bound for every nonconstant detector without requiring the unused labels to have zero reference rates.

C.3 BINARY PROXIES AND THIN-EVENT COMPARISON

We now prove the lower bound. Four-clock statements below use labels $1, 2, 3, 4$, with middle detector $(-2, 1)$; exchanging the outer labels and the target names lets us assume $a > b$, and we put $d = a - b > 0$. For a logit query u, the binary proxy G_u^{bin} chooses between the two smallest anchor log-times, using their shifted log-times at u. It is exact at the anchor and at every single-coordinate decrease $-te_i$, $t \ge 0$: either the original winner remains the winner, or its only possible replacement is the original runner-up.

Lemma C.2 (Thin-event increment). *Let $K = \mathbb{E}_G H^2$. If $0 < s \le r \le r_0$, $\|u\|_\infty, \|v\|_\infty \le r$, and $\|u - v\|_\infty \le s$, then, for $\mathcal{R}_H(u) = \mathbb{E}_G H(e_{G_u^{\mathrm{bin}}} - e_{G_u})$,*

$$\|\mathcal{R}_H(u) - \mathcal{R}_H(v)\|_{\mathrm{TV}} \le C\sqrt{rsK}. \tag{85}$$

Also $\Pr(Y_{(3)} - Y_{(1)} \le 2r) \le Cr^2$.

Proof. Actual and proxy outputs agree outside $T_r = \{Y_{(3)} - Y_{(1)} \le 2r\}$. If a winner changes between u and v, a pair difference is in an interval of length $2s$. On T_r, two additional competitors are within $O(r)$ of the anchor minimum. If the changing pair belongs to the anchor top three, one pair difference and a third-label difference occupy area $O(sr)$. Otherwise a three-dimensional set of volume $O(sr^2)$ suffices: the fourth label must be within $O(r)$ too, by its winning inequality. The same argument for a proxy change uses only the anchor top two. The requisite densities are uniformly bounded. More explicitly, the joint log-ratio density of $k+1$ selected clocks, relative to base b, is

$$k! \frac{h_b \prod_{i \ne b} h_i e^{z_i}}{(h_b + \sum_{i \ne b} h_i e^{z_i})^{k+1}} = k! \frac{\prod_{i \ne b} x_i}{(1 + \sum_{i \ne b} x_i)^{k+1}}, \quad x_i = (h_i/h_b)e^{z_i}. \tag{86}$$

It is bounded for $k = 1, 2, 3$. A finite union bound gives probability Crs for the support of the error increment; its pointwise TV is at most two. Cauchy–Schwarz proves (85). The same density calculation on the top-three event gives Cr^2. □

Symmetrize the signed measure by first conditioning H on the quotient $z = (Y_1 - Y_4, Y_2 - Y_4, Y_3 - Y_4)$. Write q_h for its density and let σ exchange coordinates $2, 3$. Replace H by

$$H^\sharp(z) = \frac{H(z)q_h(z) - H(\sigma z)q_h(\sigma z)}{2q_h(z)}. \tag{87}$$

Formula (86) gives $q_h(\sigma z)/q_h(z) = [Q_h(z)/Q_h(\sigma z)]^4$, where $Q_h(z) = h_4 + \sum_{i<4} h_i e^{z_i}$. Compactness of the rates bounds this ratio and its reciprocal. Consequently the new energy is at most CK, and $B^\sharp(\ell) = [B(\ell) - \sigma B(\sigma \ell)]/2$. Target fields stay zero; field caps stay valid on any set closed under σ. For every σ-invariant quotient event T,

$$|\mu^\sharp|(T) \le \mathbb{E}_G[|H| 1_T], \qquad \mu^\sharp = H^\sharp G. \tag{88}$$

This follows from conditional Jensen and $|\mu - \sigma_* \mu| \le |\mu| + |\sigma_* \mu|$.

The path argument uses the following proxy identities. Let E_{ij} be the event that the unordered anchor top two are i, j, and put $P_{ij}(t) = \mu^\sharp(E_{ij} \cap \{Y_i - Y_j \le \epsilon t\})$. On E_{ij}, label i wins precisely when its pair difference is below the corresponding logit difference. This gives an edge representation of the proxy field. Antisymmetry pairs edges $12, 13$ with opposite signs, pairs edges $24, 34$ with opposite signs, kills edge 14, and makes the middle-edge profile even. Denote the first profile by $f = P_{12}$ and the second, after absorbing the middle-edge constant at zero, by V. If the antisymmetric anchor is $(0, -s, s, 0)$, zero fields at the two targets give

$$f(0) = V(b), \quad f(-a) = V(0), \quad f(0) - f(-a) = s, \qquad \int_{\mathbb{R}} |f'(t)|^2 dt \le C\epsilon K. \tag{89}$$

To bound the derivative, disintegrate the signed measure over $Y_1 - Y_2$. If its marginal density is p, fiberwise Cauchy–Schwarz bounds the square of the signed density by p times the fiber energy. The density p is bounded by (86); changing variables to ϵt gives the factor ϵ.

Direct summation of the edges gives the following external-cut ($O = \{1, 4\}$) values at unit logits:

$$\begin{array}{c|c} A_t = (t + b, 0, b, 0) & f(t + b) - f(t) + s \\ C_t = (0, 0, -t, 0) & s - f(t) + V(-t) \\ D_t = (b - a, b, b - a - t, 0) & -s - f(t) + V(b - a - t). \end{array} \tag{90}$$

One can also verify the table from the identity that the proxy cut at $\epsilon(x + z, z, w, 0)$, for $x + z = y + w$, is $f(x) - f(y) - V(z) + V(w)$. The comparisons $A_t^0 = (t + b, 0, 0, 0)$ and $D_t^0 = (0, 0, b - a - t, 0)$ are at distance $O(\epsilon)$ from the respective physical queries. At A_t^0 both actual and proxy external cuts

vanish by middle antisymmetry. At C_t, D_t^0 the proxy is exact when $t \geq 0$. If the actual field cap is S, Lemma C.2, applied at the radius of each individual comparison, therefore gives

$$f(t+b) \leq f(t) + (S-s) + C\epsilon\sqrt{K}\sqrt{1+t_+}, \qquad t \geq -a, \tag{91}$$

$$f(x-d) \leq f(x) + 2(S-s) + C\epsilon\sqrt{K}\sqrt{1+x}, \qquad x \geq d. \tag{92}$$

For the second line combine the upper cut at C_x with the lower cut at D_{x-d}. The permitted arguments satisfy $C\epsilon(1+t_+) \leq r_0$. If the start of a step is in $[-a, a]$, all queries in its derivation have norm at most $R\epsilon$, for $R = 10(a+b+1)$. Its error can alternatively be bounded by Cm_T, where

$$T_\epsilon = \{Y_{(3)} - Y_{(1)} \leq 2R\epsilon\}, \quad m_T = \mathbb{E}_G[|H|1_{T_\epsilon}], \quad K_T = \mathbb{E}_G[H^2 1_{T_\epsilon}]. \tag{93}$$

Indeed, the direct proxy error is supported on this event and has pointwise TV at most one; use (88).

C.4 Weighted balanced paths

Lemma C.3 (Residues and paths). *There are $\gamma > 0, N_0$ such that the following holds for integers $N \geq N_0$, $1 \leq J \leq \gamma N$. Write $kb = q_k a + r_k$, $0 < r_k < a$, for $N \leq k < 2N$. For every $0 \leq j < J$, there are paths with steps $+b, -d$ from $-a$ to $ja + r_k$ and from $(j+1)a - r_k$ to 0, each of length k. Every negative step has a nonnegative endpoint. All but $C(j+1)$ step starts lie in $[-a, a]$, and every start has height at most $C(j+1)$. In each strip $[ja, (j+1)a]$, at least cN disjoint intervals of length at most C/N contain both a low endpoint $ja + r_k$ and a high endpoint $(j+1)a - r_{k'}$.*

Proof. The negative-step budgets are $q_k - j - 1$ and $q_k + j + 1$, respectively; the positive-step budgets are their complements in k. All four are at least cN when γ is small. In the low path first take the fixed number of positive steps that brings $-a$ into $[d, a)$. While both budgets remain, take a positive step below d and a negative step at or above d. The alternating part stays in $[0, a]$. For the high path first descend while above a, then apply the same rule. When one budget is exhausted, use the remaining steps. The specified endpoints and positive step lengths show that the initial descent and final monotone part have at most $C(j+1)$ steps. A terminal negative part cannot cross zero because its prescribed endpoint is nonnegative. The height bound follows as well.

For the net assertion, let q be the largest continued-fraction denominator of b/a not exceeding N. Bounded partial quotients imply $q \geq cN$. The first q residues are within $1/q$ of a translated permutation of the uniform grid of mesh $1/q$ (use the convergent p/q and $q|b/a - p/q| < 1/q$). Thus every block of N consecutive residues has covering radius C/N on the circle. Reflection preserves this property. Partition a strip into intervals of a sufficiently large fixed multiple of $1/N$ and discard the boundary intervals. Each remaining interval contains both residue types, and there are at least cN of them. Irrationality precludes coincidence between oppositely prescribed residues, since coincidence would imply $(k + k')b/a \in \mathbb{Z}$. □

Theorem C.4 (Weighted and localized path inequalities). *Suppose a four-clock score is target-null, has detector gain three, energy K, and all-input cap at most $S = 1 + \Delta$, $\Delta \geq 0$. For N, J as in Lemma C.3 with $C\epsilon(J+1) \leq r_0$,*

$$1 \leq C\left[N\Delta + \sqrt{K}\left\{\frac{\sqrt{\epsilon}}{N\sqrt{J}} + \epsilon N + \epsilon J^{3/2}\right\}\right], \tag{94}$$

$$1 \leq C\left[N(\Delta + m_T) + \frac{\sqrt{\epsilon K}}{N\sqrt{J}} + \epsilon\sqrt{K}J^{3/2}\right]. \tag{95}$$

For general n, these statements hold after grouping the middle labels according to the two detector values. The event in (93) then uses the resulting four group clocks, while m_T, K_T, K refer to the original score.

Proof. If $N\Delta \geq 1$, both inequalities are immediate after enlarging C. It therefore suffices to treat $N\Delta < 1$. Write $x = B_2(0), y = B_3(0)$. Since $-2x + y = 3$ and $|x|, |y| \leq S$, the nonnegative quantities $S + x, S - y$ satisfy $2(S+x) + (S-y) = 3\Delta$. Antisymmetrization has anchor magnitude $s = (y - x)/2$, whence

$$0 \leq S - s \leq \tfrac{3}{2}\Delta, \qquad s \geq 1 - \tfrac{1}{2}\Delta.$$

Sum (91) and (92) along the paths in Lemma C.3. The sum of square-root height weights is at most $C[N + (j+1)^{3/2}]$. Thus every low endpoint has value at most $f(-a) + E$ and every high endpoint has value at least $f(0) - E$, where either of the following choices is valid:

$$E = C\{N\Delta + \epsilon\sqrt{K}(N + J^{3/2})\}, \qquad E = C\{N(\Delta + m_T) + \epsilon\sqrt{K}J^{3/2}\}.$$

In a cell containing both endpoint types, absolute continuity and Cauchy–Schwarz give derivative energy at least $cN(s - 2E)_+^2$. There are cN disjoint cells in each of J strips. Equation (89) consequently gives

$$(s - 2E)_+^2 JN^2 \le C\epsilon K, \qquad s \le 2E + \frac{C\sqrt{\epsilon K}}{N\sqrt{J}}.$$

This proves both inequalities: there are $\Theta(JN)$ cells, and each contributes the additional factor N through its derivative energy.

For the grouping assertion, retain the outer clocks and take minima over the two nonempty middle groups. These are independent exponential clocks whose rates stay in a fixed compact interval. Condition the original score on them. Inputs multiplying all weights within a group by the same factor realize the four-clock fields. Grouping decreases TV; conditioning decreases energy and the localized absolute mass. It preserves zero target fields and the detector gain, proving the claim without a restriction on the original score representation. □

To finish the lower bound in Theorem C.1, set $\Delta = 0$ and take

$$N \asymp \epsilon^{-3/14}, \qquad J \asymp \epsilon^{-1/7}.$$

All three coefficients multiplying $\sqrt{K}$ in (94) are $O(\epsilon^{11/14})$. The integer, ratio, and radius requirements hold for sufficiently small ϵ. Thus $K \ge c\epsilon^{-11/7}$.

Corollary C.5 (Conditional localization)**.** *For each fixed D, a strict-cap score with $K \le D\epsilon^{-11/7}$ necessarily satisfies*

$$m_T \asymp_D \epsilon^{3/14}, \qquad K_T \asymp_D \epsilon^{-11/7}.$$

There are $u_D, v_D > 0$ such that

$$\mathbb{E}_G[H^2 1_{T_\epsilon \cap \{|H| \ge u_D \epsilon^{-25/14}\}}] \ge v_D \epsilon^{-11/7}.$$

This is a necessary structure conditional on the indicated energy budget; it is not an existence theorem at that budget.

Proof. Choose $J = \lfloor \alpha\epsilon^{-1/7} \rfloor$ with α small and $N = \lceil \beta\epsilon^{-3/14} \rceil$ with β large, depending on D. The last two terms on the right of (95), including its outside constant, are at most $1/2$. Thus $m_T \ge c_D/N$. Lemma C.2 and Cauchy–Schwarz give $m_T \le C\epsilon\sqrt{K_T} \le C_D\epsilon^{3/14}$, and also the asserted lower bound for K_T. The absolute mass below amplitude $u_D\epsilon^{-25/14}$ is at most $Cu_D\epsilon^{3/14}$. Choose u_D so this is less than half the required mass. Cauchy–Schwarz on the remaining part of T_ϵ, whose probability is at most $C\epsilon^2$, proves the tail claim. □

C.5 A FINITE HOMOGENEOUS CONSEQUENCE

The weighted proof uses finitely many actual queries. Retain the inputs in (90) and their comparisons for every step of the paths, close the list under middle exchange, and include p, q, h. In particular, retain both comparison families A_t^0, D_t^0 used in the thin-increment argument. Do this for every two-group partition of the middle labels and both orientations of the outer labels. All step heights belong to $\{ub + va : |u|, |v| \le CN,\ -a \le ub + va \le C(J+1)\}$. There are at most CNJ such heights: for each of $O(N)$ possible u, there are $O(J)$ possible integers v. Hence the resulting scoring pool $\mathcal{S}_{N,J}$ has size $M \le C_n NJ$ and raw-logit radius $C\epsilon(J+1)$. This describes a deterministic pool, fixed before the sampler is queried.

For a target-null score let D_H be the TV of its signed target table, and let $S_H = \max_{w \in \mathcal{S}_{N,J}} \|B_H(w)\|_{\mathrm{TV}}$. Under the reference law the only target cells are the diagonal, first row, and last column. The zero-marginal condition therefore implies that the signed target table is $\sum_{j \in \mathcal{J}} \beta_j C_j^{\mathrm{cyc}}$. Put $u = \sum_{\beta_j > 0} \beta_j$, $v = -\sum_{\beta_j < 0} \beta_j$. Direct summation of the entries gives

$$D_H = 2\max(u, v) + \min(u, v), \qquad \max(u, v) \le S_H. \tag{96}$$

Here $\beta_j = B_H(h)_j$, by the reference anchor identity $1_{\{G_p = G_q = j\}} = 1_{\{G_h = j\}}$ for middle j.

Proposition C.6 (Finite weighted remainder). *For admissible N, J, there are fixed $c, C > 0$ such that*

$$D_H \le \kappa_N S_H + a_{N,J}\sqrt{K}, \qquad \kappa_N = 3 - c/N, \qquad a_{N,J} = C\left\{\frac{\sqrt{\epsilon}}{N^2\sqrt{J}} + \epsilon + \frac{\epsilon J^{3/2}}{N}\right\}. \quad (97)$$

Choose c small enough that $2 < \kappa_N < 3$ for every admissible N. For every target event $D \subseteq [n]^2$ there are single-input payoffs f_w, u, v and a constant c_D such that the actual-row statistic

$$T_D = 1_D(X_p, X_q) - c_D - u(X_p) - v(X_q) - \sum_{w \in \mathcal{S}_{N,J}} f_w(X_w) \quad (98)$$

is target-null under G and satisfies

$$\sum_w \operatorname{osc}(f_w) \le \kappa_N, \qquad \mathbb{E}_G T_D^2 \le V_{N,J} := C\left\{\frac{\epsilon}{JN^4} + \epsilon^2 + \frac{\epsilon^2 J^3}{N^2}\right\}, \qquad |T_D| \le B_0. \quad (99)$$

The last bound holds on every categorical row, including reference-null rows, with B_0 independent of N, J, ϵ.

Proof. If $D_H \le 2S_H$, the first assertion is immediate after decreasing c. Otherwise both sign groups in (96) are nonempty. Group and orient them so the middle anchor entries are $-m, l$, where $m = \max(u, v)$ and $l = \min(u, v)$. Antisymmetrization yields $s = (m + l)/2 \le S_H$ and $D_H = m + 2s \le S_H + 2s$. The path proof, without normalizing the gain, gives

$$s \le C_2 N(S_H - s) + C\epsilon\sqrt{K}(N + J^{3/2}) + \frac{C\sqrt{\epsilon K}}{N\sqrt{J}}.$$

Solve for s and use $D_H \le S_H + 2s$. The cap coefficient is $3 - 2/(C_2 N + 1)$ and the remainder is the one in (97). This derivation explains why the constant ϵ term survives in the finite remainder.

Work now in the finite-dimensional reference row space. Let $\mathcal{V}$ be the target-null subspace, and P its orthogonal projection. In the gauge $f_w(n) = 0$, the oscillation ball $\sum_w \operatorname{osc}(f_w) \le \kappa_N$ is compact and convex. Its image under $f \mapsto P\sum_w f_w(X_w)$ has support function $\kappa_N S_H$ at $H \in \mathcal{V}$: this is exactly the duality between TV on zero-sum vectors and oscillation. Since $|\langle H, 1_D\rangle| \le D_H$, finite-dimensional separation and (97) imply that $P1_D$ is within L^2 distance $a_{N,J}$ of this image. The orthogonal complement of $\mathcal{V}$ consists of a constant and target payoffs. This proves (98) and the variance bound.

To check the bound on all rows, shift each f_w to have minimum zero, absorbing the change into c_D. Then their sum lies in $[0, 3]$ on every row. The reference target cells satisfy $g_{ii} \ge c_0$ and $g_{1j} \ge c_0\epsilon$ for every $j \ge 2$, including $j = n$. To verify them directly, take T_i in a fixed small interval and every other clock in a fixed larger interval for the diagonal bound. For middle j, take $T_j = t \in [1, 2]$, $T_1 \in [e^{-a\epsilon}t, t]$, and every other clock above 3. This realizes $(G_p, G_q) = (1, j)$ with probability at least $c\epsilon$. For $(1, n)$, take $T_1 = t \in [1, 2]$, $T_n \in [t, e^{a\epsilon/2}t]$, and all middle clocks above 3. Then $T_1 < e^{b\epsilon}T_n$ and $T_n < e^{a\epsilon}T_1$, as required. The compact rate bounds make all constants uniform. Our admissibility conditions imply $V_{N,J} \le C\epsilon$: $J \le \gamma N$ gives $J^3/N^2 \le CJ$, and $\epsilon J \le C$. Consequently $|\mathbb{E}_G[T_D \mid X_p = i, X_q = j]| \le \sqrt{V_{N,J}/g_{ij}}$ is uniformly bounded on the first-row and diagonal cells. Absorb c_D into v and impose $u(1) = 0$. The first-row cells, together with $(1, 1)$, bound every $v(j)$; the diagonal cells then bound every $u(i)$. The displayed expression for T_D is therefore uniformly bounded on arbitrary categorical rows. This includes rows outside the reference support. □

C.6 Executing the finite residual with a bounded number of calls

The pilot below uses only full-support inputs. Both the finite scoring pool and every possible pilot query are fixed before observations.

Lemma C.7 (Exact reference pilot and one fresh audit). *Suppose a finite scoring pool has M inputs with raw logits of norm at most $R \le R_0$. There is a deterministic pilot which predicts all pool responses exactly under G, uses two calls on its good branch, and at most $2 + P$ calls otherwise, where $P \le C_n \log(M + 1)$. The reference probability of the latter branch is at most $C_n R$. Every*

potential query has full support and raw-logit range at most $4R$; the entire potential input family has cardinality $O_n(M)$. For arbitrary sampler responses the pilot always returns a prescribed prediction. A fresh randomized audit of (98) then gives a statistic $\widetilde{T}_D$ whose conditional mean is the true residual, whose reference variance is unchanged, and whose absolute value is bounded by a fixed constant.

Proof. Query the anchor and call its reported label i. Query the anchor with coordinate i decreased by $2R$ in logit. If the result is still i, predict i throughout the pool. Under G this is exact, since $Y_j - Y_i \geq 2R$ for all $j \neq i$. The bad branch has probability at most $C_n R$, by the bounded density of each pair log-ratio and a union bound.

On the bad branch put $d = 2R$ and query, for each $j \neq i$, the input that boosts only coordinate j of h by e^d. Declare j active exactly when this query reports j, and always include i. Under G, activation is equivalent to $Y_j - Y_i < d$. An inactive label cannot beat i anywhere in the scoring pool. For each active pair j, k and each scored threshold $t = \ell_j^w - \ell_k^w \in [-d, d]$, use the full-support input with raw logits

$$u_j = d + \max(t, 0), \qquad u_k = d + \max(-t, 0), \qquad u_l = 0 \quad (l \notin \{j, k\}).$$

Both boosted active labels have adjusted log-time below Y_i, while every other label has log-time at least Y_i. Thus this query returns j precisely when $Y_j - Y_k < t$, and otherwise returns k. Binary search locates the pair difference among its distinct sorted thresholds in at most $\lceil \log_2(M+1) \rceil$ calls. The resulting pair comparisons identify every pool winner by a tournament, without further calls. The $n - 1$ activation calls are absorbed into P. Under an arbitrary response map the comparisons need not be consistent; interpret an unexpected output by a fixed search branch and use a fixed tournament even for inconsistent comparisons. The procedure still terminates and returns categorical predictions. All raw coordinates of a pair query lie in $[0, 2d]$, so their range is at most $4R$. Gate and activation queries have range $2R$. Fix the binary-search trees in advance. There are at most M nodes per pair, and the pair query depends only on its pair and threshold, not on the active set or the anchor winner. Together with the at most $2n$ gates and activation queries this gives $O_n(M)$ potential inputs.

Put $\omega_w = \operatorname{osc}(f_w)$ and $\Omega = \sum_w \omega_w \leq \kappa_N$. Independently of the seed and pilot, draw one W with probability $\pi_w = \omega_w / \Omega$, omitting zero oscillations, and query its actual output. If $\widehat{X}_w$ are the pilot predictions, use

$$\widetilde{T}_D = 1_D(X_p, X_q) - c_D - u(X_p) - v(X_q) - \sum_w f_w(\widehat{X}_w) - \frac{f_W(X_W) - f_W(\widehat{X}_W)}{\pi_W}.$$

When $\Omega = 0$ no audit is needed. Conditional on the seed, the pilot, and all target responses, the last term averages to the difference between actual and predicted payoffs. This proves conditional unbiasedness for every sampler. Under G it is identically zero, and the pilot predictions are exact, so $\widetilde{T}_D = T_D$. The fresh correction is bounded by Ω; all other terms are bounded by the proof of Proposition C.6. This also proves a uniform bound on every response path. □

The statistical class has exact marginals at p, q and all-input marginal TV at most $\tau = \rho/3$. A bad sampler means one with target TV greater than ρ. For each target event D, exact target marginals and the oscillation bound give

$$\mathbb{E}_F \widetilde{T}_D \geq \Pr_F((X_p, X_q) \in D) - g(D) - \kappa_N \tau.$$

At least one of the finitely many events D of a bad sampler therefore has mean at least $s = \rho - \kappa_N \tau = c\rho/(3N) > 0$. Keeping all 2^{n^2} events is sufficient here since n is fixed. They can share the pilot; each needs at most one fresh audit per row.

We use a finite-horizon test based on bounded-mean betting factors (Waudby-Smith & Ramdas, 2024). Suppose $|Z| \leq H_0$, $\mathbb{E}_G Z = 0$, and $\mathbb{E}_G Z^2 \leq V$. Set

$$\lambda = \frac{1}{32} \min\left\{ \frac{s}{V + s^2}, \frac{1}{H_0 + s} \right\}, \qquad e = 1 - \lambda(Z - s) > 0.$$

If $\mathbb{E}_F Z \geq s$, then $\mathbb{E}_F e \leq 1$. Also, for $|x| \leq 1/32$, Taylor's theorem gives $(1+x)^{-1/2} \leq 1 - x/2 + x^2$. Using $x = \lambda(s - Z)$ shows

$$\mathbb{E}_G e^{-1/2} \leq 1 - \lambda s/2 + \lambda^2(V + s^2) \leq 1 - \lambda s/4.$$

Let $F_0 = 2^{n^2}$ and $L = 2F_0/\delta$. On $Q = \lceil 6 \log L/(\lambda s) \rceil$ independent rows, report the target bound ρ only if every event's product of factors is at least L. Otherwise report one. Under a bad sampler Markov's inequality for its offending event bounds the probability of reporting ρ by $1/L \le \delta$. Under G, Markov applied to the inverse square root product gives failure probability at most $L^{1/2} e^{-Q\lambda s/4} \le 1/L$ per event. A union bound gives reference failure at most $\delta/2$. Correlations between event statistics within a row do not affect either argument. The number of rows obeys

$$Q \le C \log(1/\delta) \left[\frac{\epsilon}{JN^2\rho^2} + \frac{\epsilon^2 N^2}{\rho^2} + \frac{\epsilon^2 J^3}{\rho^2} + \frac{N}{\rho} \right]. \tag{100}$$

To obtain a deterministic total-call cap, stop and report one if more than $C(RQ + \log(1/\delta))$ rows enter the bad pilot branch. Under G these branch indicators are independent Bernoulli variables of mean at most CR; the elementary exponential-moment bound for their sum makes the additional reference failure probability at most $\delta/2$. Aborting only removes certifications, so honesty is preserved. The cap and the expected reference call count are bounded, respectively, by

$$C\{Q + PRQ + P \log(1/\delta)\}, \qquad CQ(1 + PR). \tag{101}$$

Opened seeds are at most Q; (101) includes both pilot and audit calls.

C.7 Optimization and the remaining gap

Theorem C.8 (Uniform critical certificate upper bound)**.** *For sufficiently small ϵ, $0 < \rho \le c_0\epsilon$, and $0 < \delta \le 1/4$, under exact target marginals and the all-input allowance $\tau = \rho/3$, each of the following resources admits the upper bound*

$$C \log(1/\delta) \max \left\{ \epsilon^{1/4}\rho^{-5/4}, \quad \epsilon^{1/2}\rho^{-7/5}, \quad \epsilon^{11/7}\rho^{-2} \right\} : \tag{102}$$

expected reference opened seeds, expected reference actual calls, and a deterministic all-sampler total-call cap. The certificate has reference success probability at least $1 - \delta$ and is honest over the entire promised class.

Proof. Use (100) with the following parameter scales:

range of ρ	N	J
$\epsilon^{5/3} \le \rho \le c_0\epsilon$	$(\epsilon/\rho)^{1/4}$	$(\epsilon/\rho)^{1/4}$
$\epsilon^{25/14} \le \rho \le \epsilon^{5/3}$	$\epsilon^{1/2}\rho^{-2/5}$	$\epsilon^{-1/2}\rho^{1/5}$
$0 < \rho \le \epsilon^{25/14}$	$\epsilon^{-3/14}$	$\epsilon^{-1/7}$.

(103)

Take a sufficiently large fixed multiplier for N, a sufficiently small one for J, and round to integers with $J \ge 1$. Choose c_0 small enough that the top endpoint also permits $N \ge N_0$. Then $J \le \gamma N$ and all path and radius conditions hold. Fixed multipliers and integer rounding change only constants.

In the first range the first and fourth terms of (100) have size $\epsilon^{1/4}\rho^{-5/4}$, and the other terms are no larger. For the second range the first, third, and fourth terms have size $\epsilon^{1/2}\rho^{-7/5}$. The second is no larger precisely when $\rho \ge \epsilon^{25/14}$. In the third range the first three terms are $O(\epsilon^{11/7}\rho^{-2})$, and the fourth is no larger. The displayed maximum has these three respective dominant ranges: its adjacent crossover points are $\epsilon^{5/3}$ and $\epsilon^{25/14}$.

Finally, these choices have $N \le C\epsilon^{-1/4}$, $J \le C\epsilon^{-1/6}$ and $M \le C\epsilon^{-5/14}$. Thus $R \le C\epsilon^{5/6}$ and $PR \le C\epsilon^{5/6}\log(1/\epsilon)$ is bounded. The additive $P\log(1/\delta)$ in (101) is absorbed by Q, since $Q \ge c\rho^{-1}\log(1/\delta)$ whereas $\rho \le c_0\epsilon$. The pilot and bounded test therefore convert the row bound into all three asserted resource bounds. □

The certificate follows from finite homogeneous separation and the query executor. The known energy upper bound is $K_* \le C\epsilon^{-2}$, from the checkerboard construction. Whether the weighted lower bound can be attained by a score of energy $O(\epsilon^{-11/7})$ remains open. Matching statistical lower bounds for the middle and deep regimes of (102) also remain open. They require hard samplers with both the necessary amplitude control and the all-input marginal guarantee.

D COMPUTATION AND REPRODUCIBILITY

This appendix documents exact finite-plan checks, controlled finite-law comparisons, and counted-query tests. The experiments assess finite certificates and query costs; the asymptotic rates are established by the proofs of Theorems 4.1 and 5.2.

D.1 EXECUTABLE COMPONENTS

The accompanying `finite_lp.py` contains exact rational categorical law calculations, affine residual plans, confidence bounds, and query executors. `verification_experiment.py` generates the saved controlled comparisons. Relevant public interfaces are:

Interface	Implemented scope
`gumbel_joint_three`	Exact finite shared-race laws on three labels.
`three_label_residual_plan`, `three_label_residual_certificate`	Finite affine reconstruction and a bounded nonnegative residual certificate.
`three_label_target_certificate`, `three_label_hybrid_certificate`	Direct target-law estimation and a confidence-split hybrid.
`shifted_group_pair_run`, `local_block_pair_run`	Shared-seed pilots, fresh audits, gate truncation, and counted calls for supplied finite plans.
`critical_affine_pair_plan`, `critical_affine_pair_run`	Compilation and checked compressed execution for a supplied critical input pool.

The minimal critical pool $\{p, q, h\}$ uses an analytic coefficient solution. For larger supplied pools, candidate coefficients from SLSQP undergo exact rational checks of their linear constraints, target orthogonality, oscillation budget, and reference variance. A repaired candidate is retained only when its variance is no larger than the analytic baseline's. These checks certify the returned finite plan, not the optimizer's optimality or attainment of the theoretical variance rate. Generation of the weighted orbit pools in Appendix C and the full reciprocal conic optimizer remain unimplemented.

Before sampling, the checked executor verifies binary-rational comparison thresholds of the actual floating query weights and computes the resulting gate and radius bounds. It uses fresh integer-valued private audits and conservative product accumulation. It assumes exact target marginals, the stated auxiliary allowance, a fixed response map per seed, repeated-input consistency, and independent private coins. These are properties required of the interface, rather than guarantees inferred from the queried labels.

D.2 CONTROLLED FINITE-LAW COMPARISONS

Appendix A's affine principle applies to any fixed finite target tuple. The three-label experiments use the unnormalized target rows

$$P_{\text{edge}} = ((1,1,0),(1,0,1),(0,1,2)), \qquad P_{\text{cyclic}} = ((1,2,3),(3,1,2),(2,3,1)),$$

normalized individually. Their six-input affine plans are fixed in advance. Reference finite laws are computed exactly. For a mixture alternative, each seed uses the shared Gumbel race with probability $1-\lambda$ and a common-uniform inverse CDF with probability λ. Both branches have exact marginals at every input. Additional cases isolate a correct target pair with an independent anchor, and a tilted race with unknown marginal errors.

Multinomial counts from the exact laws simulate the rows. Residual, direct, and hybrid methods share validation rows where their observation models coincide. An additional direct baseline receives independent target observations to match the total output-call budget. Marginal information is matched to each method's stated contract. Sample sizes are $10, 10^2, \ldots, 10^6$, confidence is $\delta = .05$, and the saved run uses $2{,}000$ replicates with random seed 20260907. Coverage is evaluated separately at each fixed horizon.

The independent-anchor case has the correct target law but population residual $3/4$: reference completeness is specific to the full shared mechanism, including auxiliary outputs. The hybrid

Table 1: Median upper confidence bounds at $N = 10^6$ and $\delta = .05$, from `verification_experiment.json`. Residual bounds are smaller at the reference; direct bounds are smaller for the displayed alternatives. The last row includes a marginal-error correction.

Setting	True TV	Residual	Direct, equal calls	Hybrid
Edge, reference	0	.00002696	.00208882	.00003320
Cyclic, reference	0	.00002696	.00239116	.00003320
Edge, $\lambda = .01$	.00516667	.00763013	.00660140	.00732465
Cyclic, $\lambda = .01$	.00385965	.00673539	.00570686	.00656744
Correct pair, independent anchor	0	.75384056	.00229676	.00291960
Tilted race, unknown marginals	.08333333	.09150419	.08517842	.08567477

combines residual and direct bounds with a confidence split, removing this population floor. The split can increase the finite-sample bound relative to either component. The saved finite-law run had no observed undercoverage; uniform honesty comes from the analytic guarantees under the sampler assumptions.

D.3 Counted-query implementation checks

Shifted-group and local-block tests query hidden-seed callbacks and count target, pilot, and audit calls. Each setting uses three replicates and floating-point statistical arithmetic. These runs check execution and call accounting, rather than estimate asymptotic exponents or coverage frequencies.

The shifted-group targets are $p = (.4, .4, .1, .1)$ and $q = (.1, .1, .4, .4)$, with widths $.5, .25$. At $N = 16384$, saved median reference bounds are $.232346, .259210$, compared with equal-call direct bounds $.007232, .009442$. Median calls per started seed are $26.385, 25.442$. The direct baseline is tighter in every tested setting. Rank-permuted alternatives reach the call cap and receive the conservative report 1.

The local confirmation uses $p = (1/4, 1/4, 1/4, 1/4)$, $q = (1/8, .245, .255, 3/8)$, width .12, and random seed 20260910. At $N = 16384$, local-cycle and equal-call direct bounds are $.014869, .006348$; at $N = 65536$, they are $.003993, .003577$. Calls per seed are approximately 10.8. All 18 confirmation runs finish without overflow. The direct baseline is tighter in every saved run.

D.4 Reproduction

Python with NumPy, SciPy, and Matplotlib reproduces the comparisons. Exact checks can be run independently of the full simulations:

```
python -B finite_lp.py --self-test --example triangle --exact
python -B verification_experiment.py
python geometry_phase.py
```

Additional checks run in Python as follows:

```
import finite_lp as f
f.audit_self_test()
f.shifted_group_self_test()
f.local_block_self_test()
f.critical_affine_pair_self_test()
f.critical_affine_pair_execution_self_test()
```

The group/local modes use `-mode groups` or `-mode local`, respectively, with `-replicates 3 -group-confidence betting` added to the experiment command. The saved longer-horizon confirmation also sets the sample sizes, width, and random seed specified above; exact commands are provided in the accompanying README. `geometry_phase.py` evaluates the analytic rate expression without fitted parameters or experimental data. The JSON files contain the settings and saved results. Mathematical queries use exact real weights; floating-weight checks apply to the finite plans accepted by the implementation. The theory counts sampler calls separately from preprocessing, bit complexity, and end-to-end runtime.