ONE REFERENCE BIT AND THE STATISTICAL COST OF PAIRED-ONLY AUDITS

Anonymous authors
Paper under double-blind review

ABSTRACT

Paired audits may retain only whether two responses differ, discarding the responses themselves. We quantify the statistical cost of this record constraint in a two-hypothesis binary response model. Each acquired response costs one query and is independently flipped with probability $\eta \in (0, 1/2)$; comparisons are exact XORs of already acquired responses, not additional measurements. Under a common raw-response budget, we optimize the sampling allocation, pairing graph, and causal adaptive policy. For fixed designs, we derive the exact finite-budget minimax risk and show that the optimal error exponent is half the raw-record exponent. For causal policies whose actions depend only on released differences, we prove that an anytime, noise-rate-free rule—sample the less certain latent bit—is minimax optimal at every horizon. Its exponent is strictly between the fixed-design and raw-record exponents. Retaining one already acquired noisy reference response instead orients any connected difference record, reconstructs all raw responses, and attains the raw-record exponent without another query. Thus adaptivity recovers part, but not all, of the information lost by paired-only storage. The retained bit is an observed noisy response, not a true label or an oracle.

1 INTRODUCTION

A controlled evaluation compares responses before and after an intervention. In vulnerability evaluation, one may change a root-cause condition while holding incidental cues fixed, or change the cues while preserving the root cause. LLM4Vuln separates knowledge, context, and prompting factors (Sun et al., 2024); semantics-oriented transformations provide another way to stress vulnerable–patched contrasts (Huang et al., 2026). These experiments motivate a finite statistical question: can the observations distinguish a designated response rule from a specified alternative? The answer depends not only on which inputs are queried but also on which responses survive in the audit record.

Consider an informative response $\theta \oplus E$ and a control response U, where $\theta \in \{0, 1\}$ is the hypothesis and E, U are independent Bernoulli noise with rate $\eta < 1/2$. A stored difference is $Z = \theta \oplus E \oplus U$. Repeatedly querying informative inputs against this same hidden control cannot average away U: the error approaches η. Retaining the actual control response permits $Z \oplus U = \theta \oplus E$. This example is elementary, but it leaves the essential design question unanswered. Could a better paired-only audit refresh controls, choose a different graph, or allocate later queries adaptively and remove the disadvantage?

We optimize those alternatives under a common budget of fresh raw responses. Edges in a pairing graph reuse the acquired bits; a cycle does not supply another independent measurement. The graph is a representation of the record, not a source of observations. Its connected record determines the raw vector only up to a global complement. One retained raw reference fixes that ambiguity, whereas a paired-only procedure must learn the informative direction and the hidden reference direction together.

The binary problem admits an exact answer. Among all fixed designs, the best allocation splits evidence into nearly balanced odd-sized groups, with error exponent $\beta_\eta/2$, where $\beta_\eta = -\log(2\sqrt{\eta(1-\eta)})$. For causal adaptive designs, a rate-free policy samples whichever latent direction is less certain. We prove its optimality at every finite horizon and determine its exponent. This exponent is strictly between $\beta_\eta/2$ and β_η, the rate obtained after retaining one existing reference

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

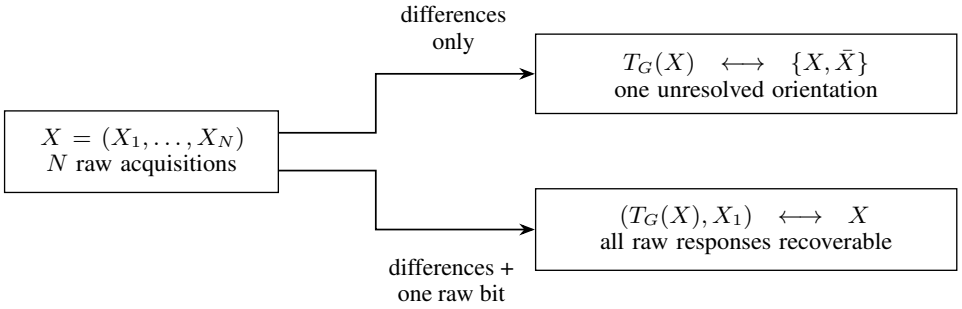

Figure 1: Two records of the same N acquired responses. For a connected pairing graph, differences identify only the unordered pair $\{X, \bar{X}\}$. Keeping one existing noisy response X_1 selects the orientation and reconstructs X. The arrows are storage maps, not extra acquisitions.

response. The principal result is this optimized comparison of record constraints, not the observation that information loss can hurt.

This problem connects noisy computation to controlled sensing and feedback coding. Its adaptive exponent is a known expression (Nitinawarat et al., 2013; Burnashev, 2022); the contribution is the finite-budget characterization and exact comparison under the paired-record constraint. Section 7 distinguishes the respective decision problems.

2 A FIXED ACQUISITION BUDGET AND THREE RECORDS

Let $h^\star$ be a known designated binary rule and h^1 a declared alternative on a finite query pool $\mathcal{Q}$. Normalize by the designated label and write $s(q) = h^1(q) \oplus h^\star(q)$. Assume that a *control* q_0 with $s(q_0) = 0$ and an *informative query* q_1 with $s(q_1) = 1$ are available. Query types with the same discrepancy are statistically equivalent in the model. Their realization as valid interventions is a separate premise; a finite static program example is given in Appendix D.

With $N \geq 2$ acquisitions, the normalized raw responses are

$$X_v = \theta s(q_v) \oplus E_v, \qquad E_1, \ldots, E_N \overset{\text{iid}}{\sim} \text{Bern}(\eta), \qquad 0 < \eta < \tfrac{1}{2}. \tag{1}$$

Every new response costs one, including repeated controls. The rule $h^\star$ is $\theta = 0$ and the alternative is $\theta = 1$. At least one control is required. A graph G on the acquired vertices records $T_G(X) = (X_u \oplus X_v)_{\{u,v\} \in E(G)}$. Query types and edges are public; the raw responses are not. Edge values have no additional noise.

Fixed, adaptive, and retained-reference experiments. A fixed design selects every query type and the graph before seeing responses. A causal adaptive design may select later queries and edges using only previously released differences, public actions, and independent randomness. The number of raw acquisitions is still exactly N; stopping times are not optimized. A retained-reference experiment supplies a connected difference record and one already acquired X_v. It reveals neither the noiseless label nor the noise realization separately.

For a test $\phi \in \{0, 1\}$, risk is $\max\{\mathbb{P}_0(\phi = 1), \mathbb{P}_1(\phi = 0)\}$. Tests may randomize. Let $R^\star_{\text{fix}}(N)$, $R^\star_{\text{ad}}(N)$, and $R^\star_{\text{keep}}(N)$ be the corresponding optimal risks. Fixed designs are deterministic; randomizing among them cannot improve the binary equalizer optimum below. Retained-reference designs may also adapt, but controls cannot reveal θ once raw responses are available.

Lemma 1 (What the graph forgets)**.** *A graph record determines the raw vector up to an independent complement on each connected component. A spanning forest is equivalent to the full graph. A connected star dominates every graph on the same acquired responses by deterministic post-processing, and one raw reference per component reconstructs the raw vector.*

Proof. $T_G(x) = T_G(x')$ holds exactly when $x \oplus x'$ is constant on every edge, hence on every component. Pathwise xor recovers every edge from a spanning forest. A star supplies $x_1 \oplus x_v$ for every v; xoring two such values recovers any requested pair. Adding x_1 recovers every x_v. □

Thus a lower bound for the best connected record applies to every fixed graph, including disconnected and densely connected designs. The lemma is an elementary incidence-map identity and an instance of comparison by post-processing (Blackwell, 1953); the statistical cost of its missing orientation is what remains to be determined.

3 THE EXACT OPTIMUM WITHOUT ADAPTATION

Write $e_n = e_n(\eta)$ for the minimax error of n independent binary symmetric observations with a fair tie, and set $e_0 = 1/2$:

$$e_n = \sum_{j>n/2} \binom{n}{j} \eta^j (1-\eta)^{n-j} + \frac{\mathbf{1}_{\{n \text{ even}\}}}{2} \binom{n}{n/2} [\eta(1-\eta)]^{n/2}. \tag{2}$$

An allocation consists of I informative responses and K controls, with $I + K = N$ and $K \geq 1$.

Theorem 2 (Fixed-design risk and allocation)**.** *A connected record has exact binary minimax risk*

$$r(I, K) = e_I + e_K - 2e_I e_K = \frac{1 - (1-2e_I)(1-2e_K)}{2}. \tag{3}$$

Let $M = 2\lfloor N/2 \rfloor$. The minimum over all fixed graphs and allocations is $r(i_M, k_M)$, where

$$(i_M, k_M) = \begin{cases} (2q+1, 2q+1), & M = 4q+2, \\ (2q-1, 2q+1), & M = 4q, \end{cases} \tag{4}$$

with $q \geq 0$ in the first case and $q \geq 1$ in the second. At odd N, add the remaining response to either group. Moreover, $R^\star_{\text{keep}}(N) = e_{N-1}$.

Why the two errors combine by parity. Reorder the raw vector as (z, w), separating informative and control coordinates, and let p_n be the all-zero binary product law. An orbit contains (z, w) and $(\bar{z}, \bar{w})$. Its likelihood difference factors as

$$P_0 - P_1 = [p_I(z) - p_I(\bar{z})]\,[p_K(w) - p_K(\bar{w})]. \tag{5}$$

The likelihood-ratio decision is therefore the xor of the two group-majority decisions. Their independent errors have probabilities e_I and e_K, giving (3). Complementing the informative group swaps the hypotheses and the decision, so the two conditional errors are equal. The Bayes-optimal test is an equalizer and hence minimax.

Why exactly balanced counts can be suboptimal. Fair ties imply $e_{2j} = e_{2j-1}$: the extra response at an even count has zero marginal decision value. For odd counts define $b_{2j+1} = 1 - 2e_{2j+1}$. A central-binomial identity gives

$$b_{2j+1} = (1-2\eta) \sum_{i=0}^{j} \binom{2i}{i} [\eta(1-\eta)]^i.$$

Write $a_i = \binom{2i}{i}[\eta(1-\eta)]^i$ and $S_j = \sum_{i=0}^{j} a_i$. Then

$$\frac{a_{i+1}}{a_i} = \frac{2(2i+1)}{i+1} \eta(1-\eta) < 4\eta(1-\eta) < 1.$$

Since a_{j+1} decreases while S_j increases, the ratio $S_{j+1}/S_j = 1 + a_{j+1}/S_j$ decreases. Thus the odd-count advantages are log-concave in their index. If two odd groups differ by more than two, transferring two responses from the larger to the smaller cannot decrease their product in (3). For an even total, an even–even allocation can first be reduced to two odd groups without changing risk, and its two unused responses assigned without increasing risk. For an odd total, reducing the even group by one leaves two odd groups. Optimizing these groups proves (4); adding back the unused response changes no error. For example, at $N = 4$ and $\eta = 0.1$, the $1 + 3$ split has error 0.1224, whereas the seemingly balanced $2 + 2$ split has error 0.18.

A retained reference reconstructs the raw vector. Controls then have a hypothesis-independent law, and $N - 1$ informative responses attain e_{N-1}. Even an adaptive raw-record experiment with a

compulsory control can be simulated from $N-1$ informative responses and independent control noise, so it cannot do better.

For clarity, the binomial exponent (Chernoff, 1952) needs only elementary bounds here. The square-root likelihood bound gives $e_n \le \frac{1}{2}e^{-n\beta_\eta}$. One central term supplies

$$e_{2r} \ge \frac{e^{-2r\beta_\eta}}{2(2r+1)}, \qquad e_{2r+1} \ge \frac{\sqrt{\eta/(1-\eta)}}{2r+2}e^{-(2r+1)\beta_\eta}.$$

Thus $-\log e_n = n\beta_\eta + O(\log(n+1))$ at fixed η. Since $e_{\min(I,K)} \le r(I,K) \le 2e_{\min(I,K)}$, the nearly balanced odd groups yield exponent $\beta_\eta/2$, whereas the $N-1$ reconstructed informative responses yield β_η. The loss survives optimization over graphs and control refreshes; it is not an artifact of reusing one uncalibrated anchor forever.

4 An optimal adaptive policy at every horizon

Adaptation need not commit half the budget to each direction. It can use the released differences to decide which direction still needs evidence. The useful state is especially small, but it is not a binary memoryless channel.

Two independent latent bits, one parity decision. Start with a hidden control response $U \sim \mathrm{Bern}(\eta)$ and introduce the equal prior on θ. Then $V = \theta \oplus U$ is uniform and independent of U, while $\theta = U \oplus V$. Every subsequent control difference observes U through fresh noise, and every informative difference observes V. The four states are

(U,V)	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
θ	0	1	1	0
prior mass	$(1-\eta)/2$	$(1-\eta)/2$	$\eta/2$	$\eta/2$.

The decision groups these states by parity; it does not identify all four. Under the two values of θ, marginal observations share a persistent latent bit. A no-adaptation-gain theorem for two conditionally memoryless hypotheses therefore does not apply (Nitinawarat et al., 2013, Proposition 1).

Let A and B denote signed posterior log odds in units $\log((1-\eta)/\eta)$, initially $(A,B) = (1,0)$. Observing difference y adds $1-2y$ to the queried coordinate. Causal actions introduce no additional likelihood factor, so the posterior of (U,V) remains a product. With $a = |A|$, $b = |B|$, $\rho = (1-\eta)/\eta$, define

$$q_a = (1+\rho^a)^{-1}, \qquad p_a = \frac{(1-\eta)\rho^a + \eta}{1+\rho^a}.$$

The terminal parity error is $V_0(a,b) = q_a + q_b - 2q_aq_b$. For t remaining acquisitions, Bellman's equation is

$$\begin{aligned} V_t(a,b) = \min\{&p_aV_{t-1}(a+1,b) + (1-p_a)V_{t-1}(|a-1|,b),\\ &p_bV_{t-1}(a,b+1) + (1-p_b)V_{t-1}(a,|b-1|)\}. \end{aligned} \tag{6}$$

The optimum is $V_{N-1}(1,0)$. This recursion evaluates risk; the optimal acquisition rule is simpler.

Theorem 3 (Less-certain-side sampling)**.** *For every $N \ge 2$ and $0 < \eta < 1/2$, sample the coordinate with smaller absolute signed count, breaking action ties by a fixed rule. At the end, output zero if $AB > 0$, one if $AB < 0$, and use a fair decision tie if $AB = 0$. This policy attains $R^\star_{\mathrm{ad}}(N)$ over all causal paired-only graph policies. Its actions and terminal decision need neither η nor the horizon.*

Executing the policy without a statistical solver. The policy can be implemented using two signed integer counters. The acquisition interface holds the reference while forming differences; the auditor receives only the released record.

The two signed counters use $O(\log N)$ bits. Each acquisition needs one magnitude comparison and one integer update. The terminal tie uses independent randomness. The same policy works for every permitted noise rate and budget.

Table 1: Rate-free paired-only protocol. Every response in the first two rows counts toward N; a comparison does not acquire another response.

Stage	Operation
Initialize	Acquire a control X_1 without revealing it. Set $(A, B) = (1, 0)$ and the acquisition count to one.
Acquire	If $\|A\| \le \|B\|$, acquire a control; otherwise acquire an informative response X_v. Release $y = X_v \oplus X_1$.
Update	Add $1-2y$ to A after a control or to B after an informative query. Increment the acquisition count.
Decide	After the prescribed N acquisitions, return $\mathbf{1}_{\{AB<0\}}$, using a fair coin when $AB = 0$.

Why the greedy action is globally optimal. Put $c = \frac{1}{2}\log((1-\eta)/\eta)$ and $\beta = \log\cosh c = \beta_\eta$. Under an auxiliary measure Q, make each newly released star difference an independent fair coin. For any fixed policy π, let R^{B}_π be its Bayes risk under the equal hypothesis prior, with an optimal terminal decision. Summing the smaller mass of the two parity classes gives

$$R^{\mathrm{B}}_\pi(N) = \tfrac{1}{2}e^{-\beta N}\mathbb{E}_Q \cosh\big(c(|A_{N-1}| - |B_{N-1}|)\big). \tag{7}$$

To derive the identity, put $p = \eta(1-\eta)$. The joint mass of a transcript y and a latent state is

$$P(U = u, V = v, y) = \tfrac{1}{2}p^{N/2}e^{c[(1-2u)A+(1-2v)B]}.$$

Summing within the two parity classes gives $p^{N/2}\cosh(c(A+B))$ and $p^{N/2}\cosh(c(A-B))$. Their smaller value is $p^{N/2}\cosh(c(|A|-|B|))$. Summing over the 2^{N-1} equally weighted transcripts under Q gives (7). The equalizer argument below makes this Bayes optimum minimax.

Let L_1, L_2 average neighboring signed states along the first or second coordinate:

$$L_1 f(a,b) = \tfrac{1}{2}[f(a-1,b) + f(a+1,b)], \quad L_2 f(a,b) = \tfrac{1}{2}[f(a,b-1) + f(a,b+1)].$$

Set $f_0(a,b) = \cosh(c(|a|-|b|))$ and $f_{t+1} = \min\{L_1 f_t, L_2 f_t\}$. The claim is $\Delta_t = (L_1 - L_2)f_t \le 0$ on $0 \le a \le b$. At $t = 0$, the hyperbolic-cosine addition identity gives zero for $1 \le a \le b$, zero at $(0,0)$, and $-\sinh c \sinh(cb) < 0$ at $a = 0 < b$.

Inductively, $f_t = L_1 f_{t-1}$ in the ordered quadrant. Commutativity of the coordinate averages, evenness, and exchange symmetry give all boundary cases:

$$\Delta_t(a,b) = \begin{cases} 0, & a = b, \\ \frac{1}{2}[\Delta_{t-1}(a-1,b) + \Delta_{t-1}(a+1,b)], & 1 \le a \le b-2, \\ \frac{1}{2}\Delta_{t-1}(a-1,a+1), & 1 \le a = b-1, \\ \Delta_{t-1}(1,b), & a = 0,\ b \ge 2, \\ 0, & (a,b) = (0,1). \end{cases} \tag{8}$$

Every right-hand side is nonpositive. This proves optimality for every horizon, including reflecting and diagonal states; it is not a myopic information-gain argument. The action is the same at every backward step and for every $c > 0$.

Finally, complement every informative difference and leave controls unchanged. Under the two hypotheses this transformation reflects B while preserving A, both magnitudes and the chosen actions. It reverses the terminal parity decision, with fair randomization on ties. The conditional errors are therefore equal. Bayes optimality under the equal prior provides a minimax lower bound, and this equalizer attains it.

Why an adaptive graph cannot improve on the star. Condition on the policy's independent random seed. Give a virtual simulator every difference $Z_v = X_v \oplus X_1$, with $Z_1 = 0$. At any prefix, an edge requested by the original policy is exactly $Z_u \oplus Z_v$. Feeding this value back reproduces that policy's observable history. Induction on the acquisitions therefore reproduces its next query, every later edge request, and its final decision pathwise. The virtual star can reveal more differences than

the original graph; its simulator simply ignores the extras. This proves that its optimum is a lower bound, including for a graph that later joins previously disconnected components.

The first query is chosen before any observation. If it is a control, its raw response is U; if informative, it is $\theta \oplus U$, where $U \sim \mathrm{Bern}(\eta)$. With $V = \theta \oplus U$, the fair-prior law is the same independent pair (U, V) in either case. A later query of the same type as the first observes U through fresh noise; the opposite type observes V. Hence choosing an informative first query only exchanges the action names. The smaller-magnitude rule first observes the opposite type, because its initial magnitudes are $(1, 0)$, so it includes both a control and an informative response when $N \geq 2$. Starting with a control is therefore admissible and attains the lower bound.

This reduction would fail if actions could inspect an unreleased raw bit. A device could reveal X_1 by choosing a tree when $X_1 = 0$ and that tree plus an extra edge when $X_1 = 1$. The public topology would then be an additional observation channel. Requiring causal actions to use only released differences prevents precisely this leakage; it is part of the experiment, not an implementation convention.

A five-acquisition example makes the gain concrete. At $N = 5$, the adaptive risk is $8\eta^2(1-\eta)^2$, below the fixed optimum by $\eta(1-\eta)(1-2\eta)^2$. At $\eta = 0.1$, the errors are 0.0648 and 0.1224. Appendix B gives the explicit branches and their probabilities.

5 The exact value of adaptation

An optimal finite-horizon policy does not by itself determine the limiting error rate. The remaining issue is the exponential cost of rare confidence imbalances under the fair measure in (7).

Theorem 4 (Strict separation of optimized experiments)**.** *For $c = \frac{1}{2}\log((1-\eta)/\eta) > 0$ and $\beta = \log\cosh c$, the three limits exist and satisfy*

$$\begin{aligned} \lim_{N\to\infty} -\tfrac{1}{N}\log R^\star_{\mathrm{fix}}(N) &= \tfrac{1}{2}\beta, \\ \lim_{N\to\infty} -\tfrac{1}{N}\log R^\star_{\mathrm{ad}}(N) &= \beta - \log\cosh(c/3), \\ \lim_{N\to\infty} -\tfrac{1}{N}\log R^\star_{\mathrm{keep}}(N) &= \beta. \end{aligned} \tag{9}$$

The adaptive exponent lies strictly between the other two for every $0 < \eta < 1/2$.

Here is the main argument for the nontrivial middle limit. Sort the absolute state as (a_t, b_t) with $0 \leq a_t \leq b_t$, initially $(0, 1)$. Under Q, the greedy policy moves the smaller coordinate as a fair walk reflected at zero. The maximum b_t never decreases. On entering level $b_t = B$, the smaller coordinate is $B - 1$; the process remains at that level until its reflected walk steps from B to $B + 1$.

Write $g_t = \mathbb{E}_Q \cosh(c(b_t - a_t))$. For a discount $z \in (0, 1)$, let $\kappa = \operatorname{arcosh}(1/z)$ and τ_B be the first entrance to level B. Solving the one-dimensional walk equations gives

$$\mathbb{E}_Q z^{\tau_B} = \frac{\cosh\kappa}{\cosh(\kappa(B-1))\cosh(\kappa B)}, \quad G_B(B-1, 0; z) = \frac{2\cosh^2\kappa}{\cosh(\kappa(B+1))}, \tag{10}$$

where G_B is the discounted number of visits to zero before leaving level B. To see the entrance formula, the discounted hitting function satisfies $h(j) = z[h(j-1) + h(j+1)]/2$, $h(0) = zh(1)$ and $h(B+1) = 1$. Its solution is $h(j) = \cosh(\kappa j)/\cosh(\kappa(B+1))$. Starting at $B - 1$ and multiplying these transforms for levels 1 through $B - 1$ gives the telescoping entrance expression in (10); the first entrance time is $\tau_1 = 0$.

For the occupation formula, let $L = B + 1$. The discounted probability of hitting zero before L, starting from j, is $\sinh(\kappa(L-j))/\sinh(\kappa L)$. At zero, reflection first moves to one, so the renewal equation yields

$$G_B(0, 0; z) = \left(1 - z\frac{\sinh(\kappa(L-1))}{\sinh(\kappa L)}\right)^{-1} = \frac{\sinh(\kappa L)}{z\sinh\kappa\cosh(\kappa L)}.$$

Multiply this by the hitting transform at $j = L - 2$ and use $z = 1/\cosh\kappa$ to obtain the second formula in (10). It includes $B = 1$, whose entry state is already zero. These two kernels separate the cost of reaching a high-confidence level from the cost of returning the other coordinate to uncertainty.

The entrance cost is at most $C_z e^{-2\kappa B}$. From the entry point $B-1$, the discounted first-hit probability of $j \le B-1$ is at most $e^{-\kappa(B-1-j)}$, obtained by removing the upper killing boundary. Subsequent occupation is bounded by $(1-z)^{-1}$. For $j = B$ this occupation bound alone suffices. Setting $d = B - j$ yields

$$\sum_{t\ge 0} g_t z^t \le C'_z \sum_{B\ge 1} e^{-2\kappa B} \sum_{d=0}^{B} e^{(c-\kappa)d} < \infty \quad \text{if } 3\kappa > c. \tag{11}$$

If $c \le \kappa$, the inner sum grows at most linearly; otherwise its growth is $e^{(c-\kappa)B}$. Nonnegative coefficients, followed by $z \uparrow z_c = \operatorname{sech}(c/3)$, imply $\limsup_t g_t^{1/t} \le z_c^{-1}$.

A converse must hold at every sufficiently large horizon, not only along a subsequence. Count just the visits to $a_t = 0$. Multiplying the two kernels in (10) gives their weighted generating function

$$F_0(z) = \sum_{B\ge 1} \frac{2\cosh^3\kappa\,\cosh(cB)}{\cosh(\kappa(B-1))\cosh(\kappa B)\cosh(\kappa(B+1))}. \tag{12}$$

Set $r(z) = (1-\sqrt{1-z^2})/z$, with $r(0) = 0$ and the square-root branch equal to one at zero. A level term is

$$\frac{8z^{-3}[(e^c r^3)^B + (e^{-c} r^3)^B]}{(1+r^{2(B-1)})(1+r^{2B})(1+r^{2(B+1)})}.$$

For $B \ge 2$, subtract $8z^{-3}(e^c r^3)^B$. The remainder is bounded on compact disks by a constant times $|z|^{-3}(e^{cB}|r|^{5B-2} + e^{-cB}|r|^{3B})$. Because $r(z_c) = e^{-c/3}$, one can choose $R > z_c$ below one with $e^c r(R)^5 < 1$. The remainder series then converges normally on that disk, including its removable singularity at zero. The finite $B = 1$ adjustment is analytic too. Thus

$$F_0(z) = 8z^{-3}\frac{e^c r(z)^3}{1-e^c r(z)^3} + J(z),$$

where J is analytic beyond z_c. The leading term has a simple pole at z_c. The other possible roots $r = e^{-c/3}e^{\pm 2\pi i/3}$ map to $z = 2r/(1+r^2)$ of strictly larger modulus. Shrinking R excludes them. Since $r'/r = (z\sqrt{1-z^2})^{-1}$, the principal part is

$$F_0(z) = \frac{K_c}{1-z/z_c} + J_1(z), \qquad K_c = \frac{8\sqrt{1-z_c^2}}{3z_c^3} > 0.$$

Cauchy's bound on a circle of radius $R_1 > z_c$ in the analytic remainder gives $[z^t]F_0 = K_c z_c^{-t} + O(R_1^{-t})$. These visits are included in g_t, proving the matching lower bound at every sufficiently large horizon. The level-kernel derivations and analytic boundary details are given in Appendix C.

Together with (7), this proves the middle limit. Strictness follows from $\cosh c > \cosh^2(c/3) > 1$. In particular, keeping the reference and adapting the queries are not interchangeable ways of spending the same acquisition budget. The former removes the latent orientation; the latter allocates noisy evidence toward learning it.

Corollary 5 (Acquisitions needed for small error)**.** *Let* $N_j(\varepsilon) = \min\{N \ge 2 : R_j^\star(N) \le \varepsilon\}$ *for* $j \in \{\mathrm{fix}, \mathrm{ad}, \mathrm{keep}\}$. *At each fixed* $0 < \eta < 1/2$*,*

$$\lim_{\varepsilon\downarrow 0} \frac{N_{\mathrm{fix}}(\varepsilon)}{N_{\mathrm{keep}}(\varepsilon)} = 2, \qquad \lim_{\varepsilon\downarrow 0} \frac{N_{\mathrm{ad}}(\varepsilon)}{N_{\mathrm{keep}}(\varepsilon)} = \frac{\beta}{\beta - \log\cosh(c/3)}.$$

The latter ratio tends to $3/2$ *as* $\eta \downarrow 0$*, and to* $9/8$ *as* $\eta \uparrow 1/2$*, after the small-error limit.*

Proof. Any extra acquisition may be ignored, so each optimized risk is nonincreasing. If its positive exponent is E_j, then, for every $\delta \in (0, E_j)$ and all sufficiently large N, $e^{-(E_j+\delta)N} \le R_j^\star(N) \le e^{-(E_j-\delta)N}$. Inverting and letting $\delta \downarrow 0$ gives $N_j(\varepsilon) \sim E_j^{-1}\log(1/\varepsilon)$. The noise limits follow from $\log\cosh c \sim c$ as $c \to \infty$ and $\log\cosh c \sim c^2/2$ as $c \to 0$. □

This is a sample-cost interpretation of the exponents, not a finite-budget ratio or a statement uniform in the noise. Even the optimal adaptive policy retains an acquisition penalty relative to reference retention. Its size depends on the noise regime; the strict separation should not be read as a universal factor of two for adaptation.

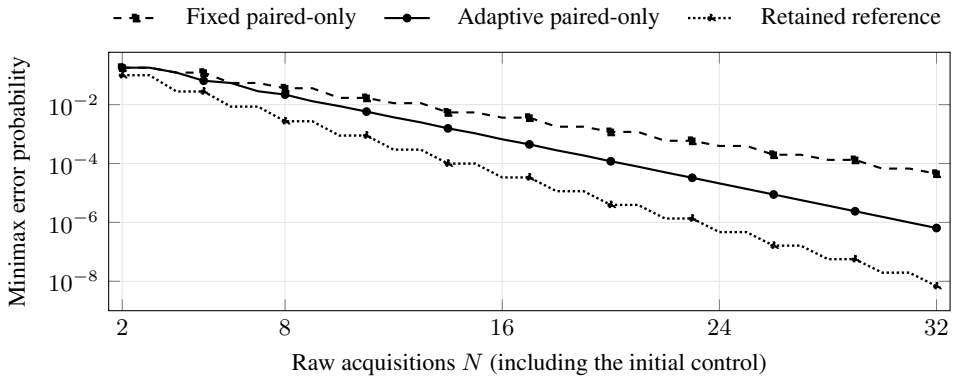

Figure 2: Exact binary risk at $\eta = 0.1$, counting the initial control in N. The curves use the optimal fixed allocation, the optimal adaptive policy, and retention of one existing reference. Discrete plateaus reflect finite-sample ties. Lines connect exact values; no fitted slope is used to establish Theorem 4.

Table 2: Error exponents per raw acquisition, evaluated from Theorem 4.

Noise η	Fixed paired-only	Adaptive paired-only	Retained reference
0.1	0.255413	0.445220	0.510826
0.2	0.111572	0.196686	0.223144
0.3	0.043588	0.077239	0.087177

6 Exact finite verification

The experiments test the model, finite formulas, policy and proof identities rather than estimate asymptotic slopes. Every equality of risks is checked with rational arithmetic. Figure 2 compares the three optimized experiments at a common raw-response budget; Table 2 evaluates the proved exponents.

For the fixed experiment, direct projection of all raw strings agrees with a separately implemented component-orbit calculation on every graph with two through five vertices, every nonconstant binary discrepancy, and three noise rates: 95,004 settings. All allocations through $N = 64$ corroborate the odd-group optimum at 189 budget–rate conditions. These checks include disconnected graphs and designs with indistinguishable orbit laws, rather than retaining only separating examples.

For the adaptive experiment, exact Bellman evaluation covers 93 root conditions through $N = 32$ and 10,245 states. An independent four-state decision tree, which does not use the absolute-log-odds reduction, agrees through $N = 8$. Enumerating ordered raw-noise strings separately checks both conditional errors of the rate-free online policy. Exact signed-lattice tests exercise every case of the greedy induction, including the diagonal, the reflecting boundary and the adjacent diagonal.

The exponent proof is corroborated by 36 discounted-walk systems solved as rational linear equations, independently of the producer's closed forms. Three causal graph rules reproduce all 1,524 declared histories under the virtual-star simulation. Deliberate corruptions of risks, policy actions and walk kernels are rejected. The public verifiers use explicit failure checks, including when Python assertions are disabled. These finite checks corroborate the written proofs; they do not replace them.

7 Relation to the closest work

Controlled sensing and feedback coding. The adaptive exponent also has the form

$$\beta - \log\cosh(c/3) = -\log\Big(\eta^{1/3}(1-\eta)^{2/3} + \eta^{2/3}(1-\eta)^{1/3}\Big). \tag{13}$$

Nitinawarat et al. (2013, Section V, Eq. (26)) obtain this achievable rate for identifying one of three locations with three sensing actions. Their fixed-design exponent is $2\beta/3$, rather than our $\beta/2$.

Burnashev (2022, Theorems 1–2) establishes the same exponent for three-message binary symmetric feedback coding and an optimal most-probable-message partition. Thus neither the expression nor finite-horizon optimality of feedback strategies in general is new.

Our experiment instead decides the parity of four latent states with a prescribed unequal prior and two coordinate queries. Its terminal loss adds the posterior masses within each parity class, whereas three-message decoding selects the largest individual posterior. Those Bellman problems are different. We supply a direct all-horizon proof for the parity loss and compare its optimum with retaining an already acquired noisy reference. Appendix E gives the action and loss comparison. Conditional memorylessness holds given (U, V) but not given θ alone, so the simple-binary no-adaptation-gain result of Nitinawarat et al. (2013) does not apply.

Noisy computation and sequential testing. Noisy parity and other Boolean functions are classical decision-tree problems (Feige et al., 1994). Later work sharpens noisy-query bounds and their dependence on error and noise (Zhu et al., 2023; Gu et al., 2025; 2026). Our two-latent-bit reduction lies within this noisy-computation perspective; we make no claim of a new general query model. Sequential experimental design (Chernoff, 1959) and active sequential hypothesis testing (Naghshvar & Javidi, 2013) optimize a different resource, allowing observation-dependent stopping. Here N is fixed, and the retention restriction applies throughout acquisition.

Evaluation and comparison of experiments. Behavioral tests (Ribeiro et al., 2020) and shortcut diagnostics (Geirhos et al., 2020) motivate declared alternative response patterns. They do not establish that current benchmarks discard raw responses. Comparison of statistical experiments explains the direction of the retention inequality (Torgersen, 1991); it does not alone supply the exact finite-budget optimum or its exponent. Appendix E locates the wider evaluation, identification and testing literature without treating application labels or classical design primitives as additional contributions.

8 Interpretation and limits

The distinction is between preserving a contrast and preserving enough information to orient it. A paired-only procedure can consistently distinguish the two declared rules, but must devote acquisitions to both latent directions. Adaptation reallocates that effort when their realized uncertainty differs. One already acquired reference response removes the orientation ambiguity instead. Its value is informational, not the quality of an extra label.

The binary factor-of-two comparison is not universal across alternative classes. For a fixed finite class, Appendix D gives matching exponents governed by agreement and disagreement margins, with relative efficiency ranging sharply from one half to one. An observation may orient one alternative while discriminating another. Nor does the equal-rate result apply unchanged to more reliable controls: a certified heterogeneous-noise example exceeds the homogeneous half-exponent. These extensions delimit the result rather than enlarge the main policy theorem.

The claims require independent common-rate flips, available controls, exact differences of acquired responses, and causal actions that use no hidden raw bit. Choosing the graph according to such a bit would itself leak the reference and change the experiment. Arbitrary compression and variable stopping are also outside scope. The finite program example supplies declared response signatures, not validation of real source transformations, learned vulnerability detectors, or internal semantic reasoning. The prevalence of paired-only logging is not measured here. The exact comparison concerns acquisition efficiency within this response model.

Design implication. Record design is part of the statistical experiment. With one retained acquired response, a connected paired record is equivalent to the raw record; without it, query allocation cannot remove the quotient by global complement. Audit reports should therefore state both the raw acquisition budget and the released record, rather than count pairwise comparisons as independent observations. This recommendation concerns the information retained by the declared experiment and does not validate the interventions that induce its response signatures. The statistical comparison begins only after those signatures have been justified.

AI USE STATEMENT

Generative AI tools were used for language polishing and grammar checking. The authors reviewed the resulting text for clarity and grammatical correctness.

REPRODUCIBILITY STATEMENT

All assumptions and complete proofs supporting the main results appear in the main text and appendices of this PDF. The separate supplementary document contains auxiliary structural results and additional exact finite checks. No external dataset or trained model is used.

REFERENCES

Md Basim Uddin Ahmed, Nima Shiri Harzevili, Jiho Shin, Hung Viet Pham, and Song Wang. SecVulEval: Benchmarking LLMs for real-world C/C++ vulnerability detection. *arXiv preprint arXiv:2505.19828*, 2025. URL https://arxiv.org/abs/2505.19828.

Dana Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988. doi: 10.1023/A:1022821128753.

Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. URL https://arxiv.org/abs/1907.02893.

Anthony C. Atkinson and Valerii V. Fedorov. The design of experiments for discriminating between two rival models. *Biometrika*, 62(1):57–70, 1975. doi: 10.1093/biomet/62.1.57.

Ruqi Bai, Yao Ji, Zeyu Zhou, and David I. Inouye. From invariant representations to invariant data: Provable robustness to spurious correlations via noisy counterfactual matching. *arXiv preprint arXiv:2505.24843*, 2025. URL https://arxiv.org/abs/2505.24843.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24 (2):265–272, 1953. doi: 10.1214/aoms/1177729032.

Marat V. Burnashev. On the reliability function for a BSC with noiseless feedback at zero rate. *Problems of Information Transmission*, 58(3):203–216, 2022. doi: 10.1134/S0032946022030012. URL https://doi.org/10.1134/S0032946022030012.

Saikat Chakraborty, Rahul Krishna, Yangruibo Ding, and Baishakhi Ray. Deep learning based vulnerability detection: Are we there yet? *IEEE Transactions on Software Engineering*, 48(9): 3280–3296, 2022. doi: 10.1109/TSE.2021.3087402.

Yizheng Chen, Zhoujie Ding, Lamya Alowain, Xinyun Chen, and David Wagner. DiverseVul: A new vulnerable source code dataset for deep learning based vulnerability detection. In *Proceedings of the 26th International Symposium on Research in Attacks, Intrusions and Defenses*, pp. 654–668. ACM, 2023. doi: 10.1145/3607199.3607242.

Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952. doi: 10.1214/aoms/1177729330.

Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3): 755–770, 1959. doi: 10.1214/aoms/1177706205.

D. R. Cox. Tests of separate families of hypotheses. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pp. 105–123. University of California Press, 1961. URL https://projecteuclid.org/euclid.bsmsp/1200512162.

Xueying Du, Geng Zheng, Kaixin Wang, Jiayi Feng, Wentai Deng, Mingwei Liu, Xin Peng, Tao Ma, and Yiling Lou. Vul-RAG: Enhancing LLM-based vulnerability detection via knowledge-level RAG. *arXiv preprint arXiv:2406.11147v1*, 2024. URL https://arxiv.org/abs/2406.11147v1.

Jiahao Fan, Yi Li, Shaohua Wang, and Tien N. Nguyen. A C/C++ code vulnerability dataset with code changes and CVE summaries. In *Proceedings of the 17th International Conference on Mining Software Repositories*, pp. 508–512, 2020. doi: 10.1145/3379597.3387501.

Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994. doi: 10.1137/S0097539791195877.

Michael Fu and Chakkrit Tantithamthavorn. LineVul: A transformer-based line-level vulnerability prediction. In *Proceedings of the 19th International Conference on Mining Software Repositories*, pp. 608–620, 2022. doi: 10.1145/3524842.3528452.

Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. Evaluating models' local decision boundaries via contrast sets. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 1307–1323. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.findings-emnlp.117.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020. doi: 10.1038/s42256-020-00257-z.

Sally A. Goldman and Michael J. Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995. doi: 10.1006/jcss.1995.1003.

Yuzhou Gu, Xin Li, and Yinzhan Xu. Tight bounds for noisy computation of high-influence functions, connectivity, and threshold. In *Proceedings of the 38th Conference on Learning Theory*, volume 291 of *Proceedings of Machine Learning Research*, pp. 2540–2591, 2025. URL https://proceedings.mlr.press/v291/gu25a.html.

Yuzhou Gu, Xin Li, and Yinzhan Xu. A unified lower bound on the noisy query complexity of boolean functions. In *Proceedings of the Thirty-Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pp. 2940–2962. PMLR, 2026. URL https://proceedings.mlr.press/v336/gu26a.html.

Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2007.01434.

Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R. Bowman, and Noah A. Smith. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 107–112. Association for Computational Linguistics, 2018. doi: 10.18653/v1/N18-2017.

Te Sun Han and Ryo Nomura. First- and second-order hypothesis testing for mixed memoryless sources. *Entropy*, 20(3):174, 2018. doi: 10.3390/e20030174.

Feiyang Huang, Yuqiang Sun, Fan Zhang, Ziqi Yang, Han Liu, and Yang Liu. Do fine-tuned LLMs understand vulnerabilities? an investigation into the semantic trap. *arXiv preprint arXiv:2601.22655*, 2026. URL https://arxiv.org/abs/2601.22655.

Pritish Kamath, Akilesh Tangella, Danica J. Sutherland, and Nathan Srebro. Does invariant risk minimization capture invariance? In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 4069–4077. PMLR, 2021. URL https://proceedings.mlr.press/v130/kamath21a.html.

Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger (eds.), *Complexity of Computer Computations*, pp. 85–103. Plenum Press, 1972. doi: 10.1007/978-1-4684-2001-2_9.

Divyansh Kaushik, Eduard Hovy, and Zachary C. Lipton. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=Sklgs0NFvr.

Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, Tony Lee, Etienne David, Ian Stavness, Wei Guo, Berton Earnshaw, Imran Haque, Sara M. Beery, Jure Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. WILDS: A benchmark of in-the-wild distribution shifts. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5637–5664. PMLR, 2021. URL https://proceedings.mlr.press/v139/koh21a.html.

David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (REx). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 5815–5826. PMLR, 2021. URL https://arxiv.org/abs/2003.00688.

Erich L. Lehmann and Joseph P. Romano. *Testing Statistical Hypotheses*. Springer, third edition, 2005. doi: 10.1007/0-387-27605-X.

Yansong Li, Paula Branco, Alexander M. Hoole, Manish Marwah, Hari Manassery Koduvely, Guy-Vincent Jourdan, and Stephan Jou. SV-TrustEval-C: Evaluating structure and semantic reasoning in large language models for source code vulnerability analysis. In *2025 IEEE Symposium on Security and Privacy (SP)*, pp. 3014–3032. IEEE, 2025. doi: 10.1109/SP61157.2025.00191.

Zhen Li, Deqing Zou, Shouhuai Xu, Xinyu Ou, Hai Jin, Sujuan Wang, Zhijun Deng, and Yuyi Zhong. VulDeePecker: A deep learning-based system for vulnerability detection. In *Network and Distributed System Security Symposium*, 2018. doi: 10.14722/ndss.2018.23158.

Zhen Li, Deqing Zou, Shouhuai Xu, Hai Jin, Yawei Zhu, and Zhaoxuan Chen. SySeVR: A framework for using deep learning to detect software vulnerabilities. *IEEE Transactions on Dependable and Secure Computing*, 19(4):2244–2258, 2022. doi: 10.1109/TDSC.2021.3051525.

László Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4): 383–390, 1975. doi: 10.1016/0012-365X(75)90058-8.

R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3428–3448. Association for Computational Linguistics, 2019. doi: 10.18653/v1/P19-1334.

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, pp. 220–229, 2019. doi: 10.1145/3287560.3287596.

Bernard M. E. Moret and Henry D. Shapiro. On minimizing a set of tests. *SIAM Journal on Scientific and Statistical Computing*, 6(4):983–1003, 1985. doi: 10.1137/0906067.

Mohammad Naghshvar and Tara Javidi. Active sequential hypothesis testing. *The Annals of Statistics*, 41(6):2703–2738, 2013. doi: 10.1214/13-AOS1144.

Sirin Nitinawarat, George K. Atia, and Venugopal V. Veeravalli. Controlled sensing for multi-hypothesis testing. *IEEE Transactions on Automatic Control*, 58(10):2451–2464, 2013. doi: 10.1109/TAC.2013.2261188.

Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri. Asleep at the keyboard? assessing the security of GitHub Copilot's code contributions. In *IEEE Symposium on Security and Privacy*, pp. 754–768, 2022. doi: 10.1109/SP46214.2022.9833571.

Kexin Pei, Yinzhi Cao, Junfeng Yang, and Suman Jana. DeepXplore: Automated whitebox testing of deep learning systems. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pp. 1–18, 2017. doi: 10.1145/3132747.3132785.

Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: Identification and confidence intervals. *Journal of the Royal Statistical Society: Series B*, 78(5):947–1012, 2016. doi: 10.1111/rssb.12167.

Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. Hypothesis only baselines in natural language inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pp. 180–191. Association for Computational Linguistics, 2018. doi: 10.18653/v1/S18-2023.

Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. Beyond accuracy: Behavioral testing of NLP models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4902–4912. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.acl-main.442.

Elan Rosenfeld, Pradeep Ravikumar, and Andrej Risteski. The risks of invariant risk minimization. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2010.05761.

Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *International Conference on Learning Representations*, 2020. URL https://arxiv.org/abs/1911.08731.

Sergio Segura, Gordon Fraser, Ana B. Sánchez, and Antonio Ruiz-Cortés. A survey on metamorphic testing. *IEEE Transactions on Software Engineering*, 42(9):805–824, 2016. doi: 10.1109/TSE.2016.2532875.

Ayumi Shinohara and Satoru Miyano. Teachability in computational learning. *New Generation Computing*, 8(4):337–347, 1991. doi: 10.1007/BF03037091.

Yuqiang Sun, Daoyuan Wu, Yue Xue, Han Liu, Wei Ma, Lyuye Zhang, Yang Liu, and Yingjiu Li. LLM4Vuln: A unified evaluation framework for decoupling and enhancing LLMs' vulnerability reasoning. *arXiv preprint arXiv:2401.16185*, 2024. URL https://arxiv.org/abs/2401.16185.

Yuchi Tian, Kexin Pei, Suman Jana, and Baishakhi Ray. DeepTest: Automated testing of deep-neural-network-driven autonomous cars. In *Proceedings of the 40th International Conference on Software Engineering*, pp. 303–314, 2018. doi: 10.1145/3180155.3180220.

Erik Torgersen. *Comparison of Statistical Experiments*. Cambridge University Press, 1991. doi: 10.1017/CBO9780511666353.

Antonio Torralba and Alexei A. Efros. Unbiased look at dataset bias. In *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1521–1528, 2011. doi: 10.1109/CVPR.2011.5995347.

Victor Veitch, Alexander D'Amour, Steve Yadlowsky, and Jacob Eisenstein. Counterfactual invariance to spurious correlations: Why and how to pass stress tests. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://arxiv.org/abs/2106.00545.

Yaqin Zhou, Shangqing Liu, Jingkai Siow, Xiaoning Du, and Yang Liu. Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://arxiv.org/abs/1909.03496.

Banghua Zhu, Ziao Wang, Nadim Ghaddar, Jiantao Jiao, and Lele Wang. On the optimal bounds for noisy computing. In *2023 IEEE International Symposium on Information Theory (ISIT)*, pp. 1788–1793, 2023. doi: 10.1109/ISIT54713.2023.10206881. URL https://doi.org/10.1109/ISIT54713.2023.10206881.

Sandra Zilles, Steffen Lange, Robert Holte, and Martin Zinkevich. Models of cooperative teaching and learning. *Journal of Machine Learning Research*, 12:349–384, 2011. URL https://jmlr.org/papers/v12/zilles11a.html.

A FIXED-DESIGN AND RETAINED-REFERENCE PROOFS

A.1 THE GRAPH QUOTIENT AND BINARY MINIMAX IDENTITY

Let B_G be the edge–vertex incidence matrix over $\mathbb{F}_2$. An edge has ones in its two endpoint columns, so $T_G(x) = B_G x$. Its kernel consists exactly of vectors constant on each connected component. If G has $c(G)$ components, every fiber has $2^{c(G)}$ raw vectors. A spanning forest has the same kernel and determines every edge by path summation. For a connected graph, the star record $S(x) = (x_1 \oplus x_v)_{v=2}^N$ and $T_G(x)$ determine each other, so they are equivalent experiments, not merely equal-rank linear maps. For an arbitrary graph, $T_G(x)$ is obtained from $S(x)$ by pairwise xor. This deterministic post-processing proves dominance simultaneously for all hypotheses and all response laws. Adding x_1 makes the connected map injective and gives $x_v = x_1 \oplus S_v(x)$.

For the binary specialization, reorder the coordinates so the discrepancy vector is $(\mathbf{1}_I, \mathbf{0}_K)$. Under hypothesis zero, the probability of (z, w) is $p_I(z)p_K(w)$. Under hypothesis one it is $p_I(\bar{z})p_K(w)$. A connected record has likelihoods

$$P_0(\{(z,w),(\bar{z},\bar{w})\}) = p_I(z)p_K(w) + p_I(\bar{z})p_K(\bar{w}),$$
$$P_1(\{(z,w),(\bar{z},\bar{w})\}) = p_I(\bar{z})p_K(w) + p_I(z)p_K(\bar{w}).$$

Their difference is the factored likelihood difference. Since $\eta < 1/2$, $p_n(u) > p_n(\bar{u})$ exactly when u has fewer than $n/2$ ones. Take group-majority decisions, breaking each tied majority by an independent fair coin, and output their xor. This is invariant under complementing the representative, and is a likelihood-ratio test with fair randomization whenever the likelihoods tie. Under hypothesis zero the two majority errors are independent, with probabilities e_I and e_K. Their xor has error $e_I(1 - e_K) + (1 - e_I)e_K$.

Under hypothesis one, complementing the informative group transports the law under zero and reverses the test decision; the two error probabilities are equal. For any test, maximum error is at least equal-prior average error. The likelihood-ratio test minimizes that average and has equal errors, so it attains the minimax risk. If $I = 0$ or $K = 0$, the two orbit laws coincide and risk is $1/2$, as the formula with $e_0 = 1/2$ also gives. This includes the nonseparating cases needed in the finite-alternative statement.

A.2 THE FINITE-BUDGET ALLOCATION OPTIMUM

First prove the tie plateau. Couple $2r - 1$ and $2r$ observations by adding one independent bit. The decisions can differ only when the first $2r - 1$ bits have $r - 1$ or r errors. Their probabilities are in ratio $(1 - \eta)/\eta$. The loss incurred by a final error in the first case is exactly offset by the gain from a final correct bit in the second case, after fair ties. Hence $e_{2r} = e_{2r-1}$.

Adding two observations to $2r - 1$ can change the strict-majority result only at those same boundary counts. Writing $p = \eta(1 - \eta)$ and collecting the two boundary contributions yields

$$e_{2r-1} - e_{2r+1} = (1 - 2\eta)\binom{2r-1}{r-1}p^r > 0 \qquad (r \geq 1).$$

Let $a_r = 1 - 2e_{2r+1}$ for $r \geq 0$. It is positive and increasing, and its increments $\Delta_r = a_r - a_{r-1}$ satisfy

$$\frac{\Delta_{r+1}}{\Delta_r} = \frac{2(2r+1)}{r+1}p < 4p < 1.$$

Thus $a_{r+1}/a_r = 1 + \Delta_{r+1}/a_r$ is strictly smaller than $a_r/a_{r-1} = 1 + \Delta_r/a_{r-1}$. This proves strict log-concavity. Whenever $s \geq r + 2$, it implies $a_{r+1}a_{s-1} > a_r a_s$. Under a fixed sum of odd group sizes, repeatedly moving two acquisitions from the larger group to the smaller therefore maximizes the product at equal or closest sizes.

It remains to show that an optimum uses these odd group sizes. At even N, an allocation with both groups even, $(2u, 2v)$, is strictly improved by $(2u - 1, 2v + 1)$: the first reliability is unchanged by the plateau and the second increases. Therefore an optimum has two odd groups. At $N = 4q + 2$ the equal odd split is available; at $N = 4q$ the closest odd sizes are $2q - 1$ and $2q + 1$. At odd N exactly one group is even. Removing one acquisition from that group leaves its risk unchanged, reducing to

budget $N-1$. Conversely, adding one acquisition to either odd group of an optimal $N-1$ design produces the same risk. These arguments establish every case of the odd-group split above, including $N = 2, 3$.

In the retained-reference experiment all X_v are reconstructed. Controls are independent of the binary hypothesis and convey no information about it, so the exact risk with I informative responses is e_I. With at least one common control, $I \le N-1$. The majority risk is nonincreasing, and $I = N-1$ is feasible. This proves $R^\star_{\mathrm{keep}}(N) = e_{N-1}$; it does not require that every optimum use that allocation, because ties may create others. The same lower bound holds for adaptive raw-record designs: prepare $N-1$ independent informative responses and an independent stream of control noise. On each informative action reveal the next informative response; on each control action reveal fresh control noise. Because at least one control is compulsory, this simulates the entire adaptive raw transcript using at most $N-1$ informative responses. Control noise and the policy random seed are hypothesis independent. Post-processing cannot improve the binary testing risk of those $N-1$ responses. Thus adaptive allocation does not improve the retained-reference optimum either.

A.3 Uniform binomial-tail bounds and exponents

The likelihood-ratio Chernoff bound gives $e_n \le \frac{1}{2}(2\sqrt{\eta(1-\eta)})^n$ by summing the square root of the two product-law likelihoods (Chernoff, 1952). A matching bound up to a polynomial follows from one central term. For $n = 2r$, the tied term gives

$$e_{2r} \ge \frac{1}{2}\binom{2r}{r}[\eta(1-\eta)]^r \ge \frac{1}{2(2r+1)}e^{-2r\beta_\eta}.$$

For $n = 2r+1$, the first strict-majority term and $\binom{2r+1}{r+1} \ge 2^{2r+1}/(2r+2)$ give

$$e_{2r+1} \ge \frac{\sqrt{\eta/(1-\eta)}}{2r+2}e^{-(2r+1)\beta_\eta}.$$

At fixed η, these bounds establish $-\log e_n = \beta_\eta n + O(\log(n+1))$ uniformly in $n \ge 1$. For positive I, K and $m = \min\{I, K\}$, monotonicity and $0 \le e_n \le 1/2$ imply

$$e_m \le e_I + e_K - 2e_I e_K \le 2e_m.$$

The same inequalities hold for a zero-sized group using $e_0 = 1/2$. The optimal odd groups have sizes $N/2 + O(1)$, while the retained-reference design has $N-1$ informative acquisitions. The binary exponent conclusion follows.

B The optimal causal policy

B.1 Exact finite-horizon adaptive optimization

Under the equal prior on θ, let U be the initially acquired hidden control response and $V = \theta \oplus U$. The joint prior is $P(U=u, V=v) = P(U=u)/2$, so U, V are independent. An informative difference is $V \oplus E$ and a control difference is $U \oplus F$. For any realized adaptive transcript, its likelihood conditional on (U, V) is a product of factors involving only U or only V: the chosen actions are functions of previous observations and introduce no hypothesis-dependent likelihood factor. The posterior therefore factorizes. Its signed log odds are integers times $\log\rho$, beginning at $(1, 0)$.

If the signed odds for one coordinate equal $a \ge 0$, its probability of being zero is $\rho^a/(1+\rho^a)$. The next observation is zero with probability p_a and updates the odds to $a+1$; observation one updates to $a-1$. Reflecting a negative coordinate interchanges the labels of its two possible observations, and its absolute coordinate follows the same transition rule. At coordinate zero the two observations both lead to absolute value one. Thus absolute odds suffice for future Bayes risk and action choice. The two coordinates are interchangeable, allowing a sorted state representation without altering the recurrence.

At a terminal state, posterior bit-error probabilities relative to the two most likely signs are q_a and q_b. Their parity is wrong exactly when one bit is wrong, so terminal risk is $q_a + q_b - 2q_aq_b$. Conditional

expectation of the optimal remaining risk gives the Bellman recurrence above. Backward induction proves optimality over all randomized and deterministic endpoint-selection rules in this declared adaptive star model: at each state a random mixture of the two actions cannot beat their minimum.

For minimax rather than Bayes optimality, select a deterministic minimizing action using a fixed tie rule on the absolute state. Complement every informative difference and leave control differences unchanged. This transformation maps the conditional transcript law for $\theta = 0$ to that for $\theta = 1$: it reflects the signed V odds, leaves the U odds unchanged, and preserves the absolute state at every step. The chosen actions are therefore identical. At the end, the posterior parity decision reverses, with fair ties preserving the complementary decision distribution. Both conditional errors equal the Bayes error, hence the Bayes-optimal policy is an equalizer and realizes the minimax value.

After the initial anchor and $N - 1 - t$ endpoints, $a + b \leq N - t$. All states therefore lie in $\{(t, a, b) \in \mathbb{Z}^3_{\geq 0} : t + a + b \leq N\}$, with additional parity restrictions only reducing the count. There are $O(N^3)$ states and two actions per state. This counts arithmetic states, not constant-time bit complexity. Rational noise yields rational transition probabilities and values, retained exactly in the certificate. The recurrence evaluates risk at a specified rate. The optimal policy proved next does not need that rate.

For an independent small oracle, retain four unnormalized posterior masses indexed by $(U, V) \in \{0, 1\}^2$, initially proportional to $(1 - \eta, 1 - \eta, \eta, \eta)$. At each decision node, evaluate both query types and both observation branches directly by multiplying these masses by their binary-channel probabilities. At depth zero, choose the smaller sum of masses over the two parity classes. Summing branch risks and minimizing actions gives the exact decision-tree Bayes risk without an absolute-log-odds reduction. The exact finite calculation agrees with the reduced recursion through $N = 8$ at all three rates. A separate replay of all ordered raw-noise strings verifies both conditional errors for its selected policy. The next argument establishes finite-horizon greedy optimality and the limiting exponent analytically; the finite agreement alone would not do so.

B.2 Causal graph policies reduce to one virtual reference

A causal paired-only policy chooses each next query type and every requested pair from the previously released pair values, previously chosen actions, and randomness independent of the hypothesis and response noise. It may request edges between any already acquired vertices, including previously disconnected components; requests introduce no new response noise or acquisition cost. No action may depend on an unreleased raw response. All actions and requested pair values form the transcript. The raw-response budget is fixed at N.

Condition on the policy's random seed. After acquiring its first response X_1, give a virtual simulator the differences $Z_v = X_v \oplus X_1$ for every subsequent acquisition, with $Z_1 = 0$. At every prefix, each pair requested by the original policy is exactly $Z_u \oplus Z_v$. Feeding this value to the policy reproduces its entire history, hence its next query, its graph choices and its final decision, pathwise. Extra virtual differences can be ignored. Thus the optimal virtual-star risk is a lower bound for every causal graph policy, even one that adaptively joins or refreshes reference groups.

The first query is chosen before any observation. If informative, write its response as $\theta \oplus U$; if a control, as U, where $U \sim \mathrm{Bern}(\eta)$. With $V = \theta \oplus U$, the latent prior under a fair hypothesis is always $U \sim \mathrm{Bern}(\eta)$ independent of uniform V. Later queries of the same type as the first observe U through fresh noise; queries of the opposite type observe V. Hence the two choices of first type yield identical optimal risk after relabeling actions. The greedy policy proved below queries the opposite type at its first endpoint, so it includes a control and an informative acquisition whenever $N \geq 2$. Starting with a control makes the achieving star admissible even when only control-centered pairings are available. Consequently the lower bound is attained and the adaptive optimum is valid over all causal graph policies in this class, not just precommitted stars.

The measurability condition is essential. For example, a device could choose between a tree and that tree plus one extra edge according to the unreleased first raw bit. The known graph topology would then itself disclose that bit. Such a raw-dependent action is a different retention channel, not a counterexample to the causal paired-only theorem.

B.3 A RATE-INDEPENDENT OPTIMAL ADAPTIVE POLICY

Write $c = \frac{1}{2}\log((1-\eta)/\eta) > 0$, $\beta = \log\cosh c$, and $p = \eta(1-\eta)$. Let $A = 1 + \#\{\text{control zeros}\} - \#\{\text{control ones}\}$ and $B = \#\{\text{informative zeros}\} - \#\{\text{informative ones}\}$. These are signed posterior log odds in units $2c$. Their initial values are $(1, 0)$, reflecting the prior of the hidden anchor and the uniform prior of its parity with the hypothesis.

A fair-measure identity. Fix any deterministic endpoint-selection policy and a length-$t = N - 1$ transcript y. Conditional actions add no likelihood factors. Including the prior on (U, V), its unnormalized four-state mass is

$$P(U = u, V = v, y) = \tfrac{1}{2}p^{N/2}\exp\{c[(1-2u)A + (1-2v)B]\}.$$

Summing over the two parity classes gives $p^{N/2}\cosh(c(A+B))$ and $p^{N/2}\cosh(c(A-B))$. Their minimum is $p^{N/2}\cosh(c(|A| - |B|))$. Let Q instead generate every next observed difference by a fair independent coin, regardless of the action. All 2^t transcripts then have probability 2^{-t}, and the Bayes error of this policy with its optimal terminal decision is exactly

$$R^{\mathrm{B}}_{\pi}(N) = \tfrac{1}{2}e^{-\beta N}\mathbb{E}_Q\cosh(c(|A_t| - |B_t|)).$$

The fair measure is a change of measure for analysis, not the physical response distribution. Randomizing policies cannot lower the minimum of these expectations.

Greedy optimality for every horizon. On the signed lattice $\mathbb{Z}^2$, define the coordinate averaging operators

$$(L_1 f)(a,b) = \tfrac{1}{2}[f(a-1,b) + f(a+1,b)], \qquad (L_2 f)(a,b) = \tfrac{1}{2}[f(a,b-1) + f(a,b+1)].$$

Put $f_0(a,b) = \cosh(c(|a| - |b|))$ and $f_{t+1} = \min\{L_1 f_t, L_2 f_t\}$. All f_t are even in each coordinate and symmetric under exchange. We show

$$\Delta_t(a,b) := (L_1 - L_2)f_t(a,b) \le 0 \qquad (0 \le a \le b).$$

For $t = 0$, $\Delta_0 = 0$ when $1 \le a \le b$, by the hyperbolic-cosine addition identity. For $a = 0 < b$, it is $-\sinh c \sinh(cb) < 0$; at $(0,0)$ it is zero. Suppose the inequality holds at $t-1$. Then $f_t = L_1 f_{t-1}$ throughout the ordered quadrant. Commutativity of the averaging operators, together with evenness and exchange symmetry at its boundaries, gives

$$\Delta_t(a,b) = \begin{cases} 0, & a = b, \\ \frac{1}{2}[\Delta_{t-1}(a-1,b) + \Delta_{t-1}(a+1,b)], & 1 \le a \le b-2, \\ \frac{1}{2}\Delta_{t-1}(a-1,a+1), & 1 \le a = b-1, \\ \Delta_{t-1}(1,b), & a = 0,\ b \ge 2, \\ 0, & (a,b) = (0,1). \end{cases}$$

Every entry on the right is nonpositive. This proves the induction, including the reflecting and diagonal cases. Therefore it is optimal at every horizon to query a coordinate of smaller absolute odds; at equality either coordinate is optimal. A fixed tie rule produces a single policy valid for every $c > 0$ and every horizon. Its actions depend only on the two integer confidence magnitudes, not on η or the remaining budget. The terminal parity decision is also rate independent. The preceding sign-reflection equalizer argument converts Bayes optimality to minimax optimality.

B.4 A COMPLETE FIVE-ACQUISITION ADAPTIVE CALCULATION

Acquire one hidden control and two informative differences Z_1, Z_2. If they disagree, acquire two more informative endpoints and output their majority, with a fair tie. If they agree, acquire a control C_1; when $C_1 = 0$, output Z_1 and acquire one unused final endpoint. When $C_1 = 1$, acquire a second control C_2 and output $Z_1 \oplus C_2$. Every group in the calculation below refers to fresh independent raw noise, even when compared to the same anchor. It suffices to condition on $\theta = 0$: complementing all informative differences transports the law for $\theta = 0$ to that for $\theta = 1$, preserves all policy branches, and reverses the final decision. Thus the policy is an equalizer.

Set $r = 2\eta(1-\eta)$, $s = 1 - r = \eta^2 + (1-\eta)^2$, and $q = \eta^2/s$. The first two informative differences disagree exactly when their endpoint noises differ, an event of probability r independent of the anchor.

On this branch, the final two informative differences have fair-majority error η relative to $\theta \oplus U$ and are independent of U. Their error relative to θ is therefore $\eta + (1-2\eta)\eta = r$. This branch contributes r^2.

The first two informative differences agree with probability s. Conditional on agreement, their common endpoint-noise bit is one with probability q, independently of the anchor. On this branch the first control is $C_1 = U \oplus F_1$. It is zero with probability s, in which case $P(U = 1 \mid C_1 = 0) = q$. Outputting Z_1 then errs with probability $2q(1-q)$. When $C_1 = 1$, an event of probability r, the policy outputs $Z_1 \oplus C_2$; the anchor cancels and the error is the xor of the common endpoint-noise bit and the fresh second control noise, with probability $q + \eta - 2q\eta$. The agreement branch contributes

$$s\big[s\,2q(1-q) + r(q+\eta-2q\eta)\big] = r^2.$$

Adding the two branches gives $2r^2 = 8\eta^2(1-\eta)^2$. The exact fixed-design optimum at $N = 5$ equals that at $N = 4$, namely $\eta + (1-2\eta)e_3$, with $e_3 = 3\eta^2 - 2\eta^3$. Subtraction gives $\eta(1-\eta)(1-2\eta)^2 > 0$. The optional final unused endpoint in the $C_1 = 0$ branch ensures the fixed total of five acquisitions and cannot alter the decision.

C The adaptive error exponent

C.1 Reflected-walk and coefficient arguments

Under the fair measure and greedy policy, write the ordered absolute state as (a, b) with $0 \le a \le b$. It starts at $(0, 1)$. At $a = 0 < b$ it moves to $(1, b)$; at $0 < a < b$ it moves to $(a-1, b)$ or $(a+1, b)$ with equal probabilities. At $a = b > 0$, the next state is $(a-1, b)$ or $(a, b+1)$ with equal probabilities. Consequently the maximum b never decreases. On first entering level $b = B$, the smaller coordinate is $B - 1$. Until entering level $B + 1$, that coordinate is a simple random walk reflected at zero and killed when stepping from B to $B + 1$.

Let $g_t = \mathbb{E}_Q \cosh(c(b_t - a_t))$. We prove

$$\lim_{t\to\infty} \frac{1}{t} \log g_t = \log\cosh(c/3).$$

Together with the fair-measure identity, this gives $E_{\mathrm{ad}} = \beta - \log\cosh(c/3)$.

Discounted level kernels. Fix $0 < z < 1$ and put $\kappa = \operatorname{arcosh}(1/z) > 0$. If τ_B is the first entry time to level B, then $\tau_1 = 0$. The reflected-walk equations $h(j) = z[h(j-1) + h(j+1)]/2$, $h(0) = zh(1)$ and $h(B+1) = 1$ give

$$\mathbb{E}[z^{\tau_{B+1} - \tau_B}] = \frac{\cosh(\kappa(B-1))}{\cosh(\kappa(B+1))}, \qquad H_B(z) := \mathbb{E}[z^{\tau_B}] = \frac{\cosh\kappa}{\cosh(\kappa(B-1))\cosh(\kappa B)}.$$

The second identity is the telescoping product of the first, using the strong Markov property. The discounted expected number of visits to state zero within level B, starting at $B - 1$, is

$$G_B(B-1, 0; z) = \frac{2\cosh^2\kappa}{\cosh(\kappa(B+1))}.$$

For completeness, put $L = B + 1$. The discounted probability of hitting zero before L from a is $\sinh(\kappa(L-a))/\sinh(\kappa L)$. Starting at zero, the renewal equation gives

$$G_B(0, 0; z) = \left(1 - z\frac{\sinh(\kappa(L-1))}{\sinh(\kappa L)}\right)^{-1} = \frac{\sinh(\kappa L)}{z\sinh\kappa\cosh(\kappa L)}.$$

Multiply these expressions at $a = L - 2$, and use $z = 1/\cosh\kappa$, to obtain the zero-state Green formula. This also holds at $B = 1$, where the starting point is zero.

Upper exponential growth bound. Let $F(z) = \sum_{t\ge 0} g_t z^t$. Decompose discounted occupation by levels. From the entry point $B - 1$, the discounted first-hit probability of a state $j \le B - 1$ before being killed is at most $e^{-\kappa(B-1-j)}$: removing the upper killing boundary yields this usual one-sided walk solution. Subsequent discounted occupation is at most $(1-z)^{-1}$. For $j = B$ the latter bound suffices. Since $H_B(z) \le C_z e^{-2\kappa B}$, it follows that

$$F(z) \le C'_z \sum_{B\ge 1} e^{-2\kappa B} \sum_{d=0}^{B} e^{(c-\kappa)d} < \infty \quad \text{when } \kappa > c/3.$$

If $c \le \kappa$ the inner sum grows at most linearly; otherwise its geometric growth is $e^{(c-\kappa)B}$. Both cases justify the stated threshold. With $z_c = \operatorname{sech}(c/3)$, nonnegative coefficients imply $\limsup g_t^{1/t} \le z_c^{-1}$ by taking $z \uparrow z_c$.

A matching lower bound at every sufficiently large horizon. Count only visits with $a = 0$ and define $g_t^0 = \mathbb{E}_Q[\mathbf{1}_{\{a_t=0\}} \cosh(cb_t)] \le g_t$. The two exact kernels give its generating function

$$F_0(z) = \sum_{t\ge 0} g_t^0 z^t = \sum_{B\ge 1} \frac{2\cosh^3\kappa \cosh(cB)}{\cosh(\kappa(B-1))\cosh(\kappa B)\cosh(\kappa(B+1))}.$$

It remains to control coefficients, since a radius-of-convergence argument alone would yield only a subsequence lower bound. Use the analytic function

$$r(z) = \frac{1-\sqrt{1-z^2}}{z}, \qquad r(0) = 0,$$

with the square-root branch equal to one at zero. For real $z \in (0,1)$, $r = e^{-\kappa}$, and its Taylor coefficients are nonnegative. Rewriting a level term in the zero-occupation generating function gives

$$\frac{8z^{-3}[(e^c r^3)^B + (e^{-c} r^3)^B]}{(1+r^{2(B-1)})(1+r^{2B})(1+r^{2(B+1)})}.$$

For $B \ge 2$, subtracting $8z^{-3}(e^c r^3)^B$ leaves terms bounded on compact disks by a constant times $|z|^{-3}(e^{cB}|r|^{5B-2} + e^{-cB}|r|^{3B})$. All apparent singularities at zero are removable. Since $r(z_c) = e^{-c/3} < 1$, choose a real $R > z_c$, still below one, with $e^c r(R)^5 < 1$. The subtracted remainder series converges normally on $|z| < R$; the finite $B = 1$ adjustment is analytic there after decreasing R if needed. Therefore

$$F_0(z) = 8z^{-3}\frac{e^c r(z)^3}{1-e^c r(z)^3} + J(z),$$

where J is analytic on a disk of radius larger than z_c.

The leading term has a simple pole at z_c. There is no other pole of this modulus. Indeed $e^c r^3 = 1$ requires $r = r_0 e^{2\pi i \ell/3}$ with $r_0 = e^{-c/3}$. The two non-real choices, if on the chosen branch, give $z = 2r/(1+r^2)$ with modulus $2r_0/\sqrt{1+r_0^4-r_0^2} > 2r_0/(1+r_0^2) = z_c$. Taking R smaller excludes these poles. Since $r'/r = 1/(z\sqrt{1-z^2})$, subtracting the principal part of the pole decomposition gives

$$F_0(z) = \frac{K}{1-z/z_c} + J_1(z), \qquad K = \frac{8\sqrt{1-z_c^2}}{3z_c^3} > 0,$$

where J_1 is analytic on a disk of radius greater than z_c. On a fixed circle of radius $R_1 > z_c$ inside that disk, Cauchy's coefficient bound gives $g_t^0 = K z_c^{-t} + O(R_1^{-t})$. Thus $\liminf g_t^{1/t} \ge z_c^{-1}$. Combined with the upper bound, this proves the claimed growth limit and the exact risk exponent for every fixed $0 < \eta < 1/2$.

Strict separation of the three experiments. Clearly $0 < \log\cosh(c/3) < \log\cosh c$. More sharply, for $x = \cosh(c/3) > 1$, $\cosh c - \cosh^2(c/3) = x(x-1)(4x+3) > 0$. Hence $\log\cosh(c/3) < \beta/2$, proving $\beta/2 < E_{\mathrm{ad}} < \beta$. The lower comparison concerns all fixed pairing graphs and allocations; adaptive optimality covers the causal paired-only graph policies defined above by the virtual-reference reduction. Variable stopping, raw-dependent actions and arbitrary record-compression mechanisms are not covered.

D Extensions and counterexamples

D.1 A fixed finite alternative class

Put $D_h = \{q : s_h(q) = 1\}$, $d_h(n) = \sum_{q\in D_h} n_q$, and

$$\Gamma_N^{\pm} = \max_{n:\sum n_q = N,\ n_0 \ge 1} \min_h \min\{d_h(n), N - d_h(n)\}.$$

Then

$$e_{\Gamma_N^{\pm}} \le R^{\star}_{\mathrm{pair}}(N) \le \min\{1/2, 2He_{\Gamma_N^{\pm}}\}.$$

With probability allocations on the query pool,

$$\gamma_{\pm} = \max_p \min_h \min\{p(D_h), 1 - p(D_h)\}, \qquad \gamma_0 = \max_p \min_h p(D_h),$$

the exponents are $\beta_\eta\gamma_{\pm}$ and $\beta_\eta\gamma_0$, and $\gamma_0/2 \le \gamma_{\pm} \le \gamma_0$. Both endpoints occur. This is not an exact finite composite-minimax formula. The max-min linear program is a classical design primitive.

The four certified classes, omitting the common control, are:

Disagreement sets	γ_0	$\gamma_{\pm}$
$\{1\}$	1	1/2
$\{1\}, \{2\}$	1/2	1/2
$\{1,2\}, \{1,3\}, \{2,3\}$	2/3	1/2
$\{1,2\}, \{1,3\}, \{1,4\}, \{2,3,4\}$	3/5	1/2

Rate-free detector. For an orbit representative x, let $t_h(x) = \min\{\mathrm{wt}(x \oplus \delta_h), N - \mathrm{wt}(x \oplus \delta_h)\}$, where $\delta_h = (s_h(q_v))_v$. Reject if an alternative is closer than the designated rule. Otherwise, when b alternatives tie it, accept with probability 2^{-b}. This is exactly a union of pairwise likelihood-ratio tests with independent fair ties, using $O(HN)$ bit comparisons and no knowledge of η. The binary optimum and finite-class brackets also hold uniformly over $0 < \eta \le \bar{\eta} < 1/2$, at worst-case rate $\bar{\eta}$.

Fix an integer allocation n. Define its balanced margin $g(n) = \min_h \min\{d_h(n), N - d_h(n)\}$. For any graph and any composite test, restricting the risk to $h^\star$ and an alternative attaining $g(n)$ gives a lower bound of $e_{g(n)}$. Since $g(n) \le \Gamma_N^{\pm}$ and e_j is nonincreasing, every design has risk at least $e_{\Gamma_N^{\pm}}$.

For achievability choose an allocation with $g(n) = \Gamma_N^{\pm}$ and a connected graph. Each alternative h has a pairwise likelihood-ratio test with both errors equal to $r(d_h, N - d_h) \le 2e_{\Gamma_N^{\pm}}$. Use independent fair randomization for each pairwise tie, and reject $h^\star$ if any of these tests does. A false rejection is a union of H events; under alternative h, acceptance of the composite test implies acceptance of its own pairwise test. Thus maximum error is at most $2He_{\Gamma_N^{\pm}}$. Random guessing proves the additional $1/2$ cap. The union test need not be exactly minimax for $H > 1$; the theorem claims a finite bound and a matching exponent, not a finite composite closed form.

Let $L = m + 1$ be the fixed number of query types. The objective in the continuous margins is continuous on the compact simplex, so an optimum exists. Given an optimal p, reserve one control response and allocate the remaining $N - 1$ responses by flooring $(N-1)p_q$ and distributing the residual arbitrarily. For each set D_h, its count differs from $(N-1)p(D_h)$ by at most L. Thus $\Gamma_N^{\pm} \ge (N-1)\gamma_{\pm} - L - 1$. Conversely, n/N is a probability allocation, so $\Gamma_N^{\pm} \le N\gamma_{\pm}$. Therefore $\Gamma_N^{\pm}/N \to \gamma_{\pm}$. The uniform tail bounds prove the exponent, including the case $\gamma_{\pm} = 0$: then the relevant risk has a nonvanishing lower bound or, directly, zero exponent. The raw-record proof replaces $\min\{d_h, N - d_h\}$ by d_h throughout. Finite H is essential for the union factor to remain subexponential.

For completeness the four values in the four-class table above have the following exact certificates. In each allocation the first coordinate is the common control.

- For the binary family, raw allocation $(0, 1)$ attains $\gamma_0 = 1$ and paired allocation $(1/2, 1/2)$ attains $\gamma_{\pm} = 1/2$. Probability mass and the two-sided margin give the matching upper bounds.
- For the disjoint family, $(0, 1/2, 1/2)$ attains both values $1/2$. The two disjoint disagreement masses have sum at most one, proving the raw upper bound; the paired upper bound is always $1/2$.

- For the triangle family, $(0, 1/3, 1/3, 1/3)$ gives raw margin $2/3$. The average of the three disagreement masses is at most $2/3$, proving optimality. The paired allocation $(1/4, 1/4, 1/4, 1/4)$ makes each disagreement mass exactly $1/2$.
- For the four-probe family, raw allocation $(0, 2/5, 1/5, 1/5, 1/5)$ attains $3/5$. If every disagreement mass were at least t, the sum of the first three inequalities would give $3p_1 + p_2 + p_3 + p_4 \geq 3t$, and the fourth gives $p_2 + p_3 + p_4 \geq t$. Since $p_1 + p_2 + p_3 + p_4 \leq 1$, the former implies $3 - 2(p_2 + p_3 + p_4) \geq 3t$; combining yields $t \leq 3/5$. Paired allocation $(1/6, 1/3, 1/6, 1/6, 1/6)$ makes all four masses $1/2$.

These witnesses also explain the range of exponent ratios without numerical linear programming. The executable finite-class tests enumerate all integer allocations at the declared small budgets and replay an explicit composite test; they do not claim to compute exact composite minimax risk.

D.2 A FINITE STATIC PROGRAM INSTANCE

The syntax has array length 2, input indices $\{0, 1, 2, 3\}$, guard $i < 2$ or $i \leq 2$, and four unused annotation bits z. The designated root predicate is existential enabled out-of-bounds access. The four shortcuts read the annotation bits. No memory access is executed. These five queries realize the four-probe class:

Query	Root	Annotation bits	Discrepancy bits
0	0	0000	0000
1	0	1110	1110
2	1	0110	1001
3	0	0101	0101
4	1	1100	0011

The guarded-access family has index domain $\{0, 1, 2, 3\}$ and array bounds $0 \leq i < 2$. The strict guard enables exactly indices $0, 1$; the non-strict guard enables $0, 1, 2$. The existential unsafe predicate is therefore false for the former and true for the latter. Annotation bits occur in neither guard nor bounds, so all cue-only edits preserve that predicate. Switching guards reverses it. These exhaust the declared syntax; no claim is made for a richer source language. The cues in the finite-syntax table above give discrepancy columns with query supports $\{1, 2\}$, $\{1, 3\}$, $\{1, 4\}$, and $\{2, 3, 4\}$, respectively. This is exactly the four-probe family already certified above.

For rate-free likelihood ordering, write $f_j = \eta^j(1-\eta)^{N-j} + \eta^{N-j}(1-\eta)^j$. For $0 \leq j < \lfloor N/2 \rfloor$,

$$f_j - f_{j+1} = (1 - 2\eta)\left[\eta^j(1-\eta)^{N-j-1} - \eta^{N-j-1}(1-\eta)^j\right] > 0.$$

The inequality follows from $N - j - 1 > j$ and $(1-\eta)/\eta > 1$. Orbit likelihoods are exactly $f_{t_h(x)}$, so comparing them is equivalent to comparing distances in the reverse order. If any alternative has smaller distance, its pairwise test rejects with probability one. Otherwise every farther alternative accepts with probability one and every tied alternative accepts with probability $1/2$. Independent tie coins give composite acceptance probability 2^{-b}, matching the union test exactly.

To justify the uniform noise corollary, couple each Bernoulli error as $\mathbf{1}\{V_i \leq \eta\}$ for independent uniform V_i. Fair-majority error is nondecreasing under this coordinatewise coupling. The expression $r(I, K) = e_I + e_K - 2e_I e_K$ is nondecreasing in each error because its partial derivatives are $1 - 2e_K$ and $1 - 2e_I$, both nonnegative. Thus every pairwise error is bounded by its value at $\bar{\eta}$. The same union argument yields the finite-class upper bound, while testing the fixed rate $\bar{\eta}$ supplies the lower bound even for a test designed for unknown noise. The allocation and all likelihood-order decisions are unchanged as the rate varies. Uniformity fails without a fixed gap below $1/2$, because rates approaching $1/2$ make the two binary laws arbitrarily close.

D.3 HETEROGENEOUS CONTROLS AND ACQUISITION COSTS

Anchors have noise $a \in (0, 1/2)$ and cost $c \geq 1$, endpoints noise η and cost one. With n anchors, I informative endpoints and C controls, cost is $B = cn + I + C$. If $\beta_a \leq c\beta_\eta$, the optimal exponent

over fixed star schedules is $\beta_\eta/2$; one retained anchor permits β_η. The condition is sufficient, not necessary. Outside it, $(a, \eta, c) = (.01, .2, 1)$ with one fresh anchor and four informative endpoints per block exceeds the homogeneous half-exponent. Its per-ordered-string mass at weight j is

$$q_j = (1-a)\eta^j(1-\eta)^{4-j} + a\eta^{4-j}(1-\eta)^j,$$

and the block coefficient is $Z = \sum_{j=0}^{4} \binom{4}{j} \sqrt{q_j q_{4-j}}$. The strict certified inequality is $Z^2 < [2\sqrt{\eta(1-\eta)}]^5$.

Write the star observations as

$$Z_{gj} = \theta \oplus U_g \oplus E_{gj}, \qquad W_{g\ell} = U_g \oplus F_{g\ell},$$

where $U_g \sim \mathrm{Bern}(a)$ and all $E, F \sim \mathrm{Bern}(\eta)$ independently. Revealing $(U_g)_g$ lets the auditor recover independent informative observations $\theta \oplus E_{gj}$. The U and F variables have laws independent of θ, so the enhanced experiment has exact risk $e_I(\eta)$. Consequently the original record has risk at least this value.

For the second enhanced experiment reveal $V_g = \theta \oplus U_g$. Then $Z_{gj} \oplus V_g = E_{gj}$ is pure noise independent of all remaining observations and θ, whereas $W_{g\ell} \oplus V_g = \theta \oplus F_{g\ell}$. The enhanced experiment is therefore n independent binary symmetric observations with error a and C with error η, plus irrelevant independent noise. Its Bayes/minimax risk is another lower bound for the original experiment. In particular, revealing V is a legitimate hypothesis-dependent random variable in an enhanced statistical experiment: conditional on each hypothesis it is generated jointly with the original raw noises, and discarding it recovers exactly the original record.

To control the exponent uniformly over changing schedules, consider arbitrary independent binary symmetric observations with rates in the fixed set $\{a, \eta\}$. Let their joint laws be P_0, P_1 and write

$$J = -\log \sum_x \sqrt{P_0(x)P_1(x)}, \qquad L(x) = \log \frac{P_0(x)}{P_1(x)}.$$

The normalized square-root law $Q(x) = e^J \sqrt{P_0(x)P_1(x)}$ makes every observed bit independently fair. Thus L has mean zero and variance at most $M\ell^2$, where M is the number of observations and $\ell = \max\{\log((1-a)/a), \log((1-\eta)/\eta)\}$. Chebyshev's inequality yields $Q\{|L| \le 2\ell\sqrt{M}\} \ge 3/4$. Since

$$\frac{1}{2}\sum_x \min\{P_0(x), P_1(x)\} = \frac{1}{2} e^{-J} \mathbb{E}_Q e^{-|L|/2},$$

this error lies between $(3/8)e^{-J-\ell\sqrt{M}}$ and $(1/2)e^{-J}$ for $M \ge 1$. For $M = 0$ it is $1/2$. This gives a uniform sublinear remainder, not just a separate limit for each fixed schedule.

For the two enhanced experiments, $J_1 = \beta_\eta I$ and $J_2 = \beta_a n + \beta_\eta C$. With $c \ge 1$, all observation counts are at most B. Hence every fixed schedule satisfies

$$-\log R \le \min\{J_1, J_2\} + O(\sqrt{B}),$$

where the constant depends only on a, η. If $\beta_a \le c\beta_\eta$, then

$$\min\{J_1, J_2\} \le \tfrac{1}{2}(J_1 + J_2) \le \tfrac{1}{2}\beta_\eta(I + C + cn) = \tfrac{1}{2}\beta_\eta B.$$

This proves the converse uniformly for arbitrary schedules and allocations in the declared star model.

For achievability, one anchor together with I informative and C control endpoints is a single connected quotient. Its uncompressed control block has one coordinate with error a and C with error η. The likelihood difference factors as in the factored likelihood difference, with that heterogeneous control likelihood in place of p_K. If $e_{a,C}$ is its binary testing risk, the paired-only risk is $e_I + e_{a,C} - 2e_I e_{a,C}$. The uniform product bound above gives $-\log e_{a,C} = \beta_\eta C + O(\sqrt{C+1})$ because the one anchor contributes the constant β_a. Choosing $I, C = B/2 + O(1)$ gives exponent $\beta_\eta/2$. Retaining the actual anchor reconstructs the endpoints; taking $I = B - c + O(1)$ and no extra controls gives exponent β_η.

Certified failure outside the sufficient regime. At $(a, \eta, c) = (1/100, 1/5, 1)$, each refreshed block contains one anchor and four informative endpoints. Under $\theta = 0$, every ordered endpoint-difference

string of weight j has mass q_j in the heterogeneous-star statement; under $\theta = 1$ its mass is q_{4-j}. The two laws are interchanged by bitwise complementation. Their Chernoff coefficient is minimized at parameter $1/2$: the coefficient is symmetric about $1/2$ and log-convex. Independent block repetition therefore has exponent $-\log Z$ per block, or $-\log Z/5$ per acquisition. A finite-support Chernoff theorem supplies this exponent (Chernoff, 1952); alternatively, applying the square-root change of measure to independent block log-likelihood ratios gives the same subexponential lower bound as above.

The certificate stores rational lower and upper bounds for every $\sqrt{q_j q_{4-j}}$. Each bound is independently validated by squaring, avoiding floating-point comparisons. Summing with exact multiplicities $\binom{4}{j}$ produces $Z_{\mathrm{lo}} \le Z \le Z_{\mathrm{hi}}$. It also gives $b_{\mathrm{lo}} \le 2\sqrt{\eta(1-\eta)}$, and verifies $Z_{\mathrm{hi}}^2 < b_{\mathrm{lo}}^5$. Consequently $-\log Z/5 > \beta_\eta/2$ strictly. This disproves an unconditional homogeneous-rate extrapolation; it does not solve the heterogeneous schedule optimization problem.

E Context and distinctions

This section relates the response model to its motivating evaluations and statistical precedents.

E.1 What vulnerability evaluations observe

VulDeePecker uses code gadgets (Li et al., 2018), SySeVR constructs syntax- and semantics-based candidates (Li et al., 2022), Devign represents programs with graphs (Zhou et al., 2019), and LineVul studies transformer-based line-level prediction (Fu & Tantithamthavorn, 2022). These studies learn representations; our rules specify response functions.

Big-Vul links C/C++ changes to vulnerability descriptions (Fan et al., 2020); DiverseVul expands source-code vulnerability data (Chen et al., 2023); SecVulEval emphasizes context and statement-level localization (Ahmed et al., 2025). Data and evaluation choices affect the interpretation of detection results (Chakraborty et al., 2022). Our static instance in Appendix D is deliberately much smaller.

Security evaluation also concerns weaknesses in generated code (Pearce et al., 2022), retrieval-augmented detection (Du et al., 2024), and structural and semantic vulnerability reasoning (Li et al., 2025). LLM4Vuln separates knowledge, context and prompting (Sun et al., 2024); semantics-oriented transformations provide further controlled contrasts (Huang et al., 2026). We study the subsequent recording decision under declared response patterns.

E.2 Alternative rules, invariance and behavioral tests

Shortcut learning distinguishes intended from predictive features (Geirhos et al., 2020); cross-dataset evaluation exposes dataset bias (Torralba & Efros, 2011). Annotation artifacts (Gururangan et al., 2018), hypothesis-only baselines (Poliak et al., 2018), and controlled syntactic examples (McCoy et al., 2019) motivate explicit alternative rules, not claims about a predictor's internal reasoning.

Counterfactual augmentation challenges spurious associations (Kaushik et al., 2020); contrast sets probe local decision boundaries (Gardner et al., 2020); CheckList organizes behavioral tests (Ribeiro et al., 2020). Noisy counterfactual matching concerns imperfect invariant pairs (Bai et al., 2025). Their example-level uncertainty differs from exact xors of reused noisy responses.

Invariant causal prediction uses cross-environment stability for identification (Peters et al., 2016), while invariant risk minimization seeks invariant predictors (Arjovsky et al., 2019). Its recovery limitations depend on the formulation (Rosenfeld et al., 2021; Kamath et al., 2021). Counterfactual invariance formalizes robustness to spurious correlations (Veitch et al., 2021). Our discrepancies are specified rather than inferred.

Group distributionally robust optimization targets worst-group performance (Sagawa et al., 2020); risk extrapolation exploits variation in environment risks (Krueger et al., 2021). DomainBed studies model selection in domain generalization (Gulrajani & Lopez-Paz, 2021); WILDS standardizes natural distribution shifts (Koh et al., 2021). Our minimax comparison instead fixes one noise law and a declared hypothesis class.

Metamorphic testing relates executions without a direct oracle (Segura et al., 2016). DeepXplore combines differential testing and coverage-guided generation (Pei et al., 2017); DeepTest transforms inputs to test driving networks (Tian et al., 2018). Model cards emphasize documenting intended use and evaluation limits (Mitchell et al., 2019).

E.3 Identification versus a noisy retained record

Query learning supplies membership and equivalence information (Angluin, 1988). Teachability concerns finite informative presentations (Shinohara & Miyano, 1991); teaching dimension minimizes a target-identifying sample (Goldman & Kearns, 1995); cooperative teaching specifies shared teacher–learner assumptions (Zilles et al., 2011). Minimum test collections distinguish objects by noiseless outcomes (Moret & Shapiro, 1985). Our difficulty remains after a distinguishing query exists: its response is noisy and its orientation may be discarded.

Finite-class disagreement allocation uses classical Set Cover (Karp, 1972) and integral/fractional-cover distinctions (Lovász, 1975). Appendix D proves explicit allocation certificates; the optimization primitive is not new. The structural supplement's coding reductions are separate from the adaptive sampling theorem.

Noisy computation already treats unreliable queries and parity (Feige et al., 1994). Later bounds depend on the function, noise and error tolerance (Gu et al., 2025; 2026); OR, MAX, SEARCH and SORT illustrate different access and budget conventions (Zhu et al., 2023). Our comparison therefore keeps the asymmetric initial prior, compulsory first response, fixed horizon and retained reference explicit.

E.4 Decision objects behind the shared exponent

In the three-message problem of Burnashev (2022), the message is uniform, a binary channel input specifies a partition of three messages, and terminal Bayes loss is $1 - \max_i \pi_i$. Theorem 2 chooses the most probable message against the remaining pair. Its zero-rate exponent is the expression in Equation (13).

In our experiment, the posterior on (U, V) factorizes, the two available actions query one coordinate, and the decision concerns $U \oplus V$. If $u = \Pr(U = 1 \mid \mathcal{T})$ and $v = \Pr(V = 1 \mid \mathcal{T})$, the terminal loss is

$$\min\{u(1-v) + (1-u)v,\ uv + (1-u)(1-v)\}.$$

This is not the four-state identification loss $1 - \max\{uv, u(1-v), (1-u)v, (1-u)(1-v)\}$. For instance, at $u = v = 1/2$, these are $1/2$ and $3/4$. The prescribed initial prior is $u = \eta$, $v = 1/2$ and includes the paid initial control. Consequently, replacing the grouped loss by a maximum-posterior decoder would change the finite-budget optimization. The comparison here holds the acquisition budget fixed while varying which record is retained.

E.5 Statistical experiments and acquisition control

Experiment comparison orders the raw and difference records (Blackwell, 1953; Torgersen, 1991). Equalizer arguments connect their Bayes and minimax risks (Lehmann & Romano, 2005).

Chernoff's asymptotic efficiency supplies the error-exponent framework (Chernoff, 1952). Sequential design (Chernoff, 1959) and active sequential testing (Naghshvar & Javidi, 2013) also optimize stopping. Separate-family testing (Cox, 1961) and regression-model discrimination (Atkinson & Fedorov, 1975) are established design problems. Our deterministic budget is not an expected stopping cost.

Controlled sensing distinguishes fixed and causal actions (Nitinawarat et al., 2013). Our hidden reference makes the binary marginal experiment composite: expanding to (U, V) restores conditional memorylessness, but leaves a parity decision. Mixed-source testing likewise distinguishes persistent mixtures from resampling (Han & Nomura, 2018).