

Feedback and Robust Local Readout of a Fixed Four-State Quantum Source

Jeremy Rodgers
Independent Researcher
www.everythingequation.com

27 October 2026

<https://doi.org/10.5281/zenodo.23721114>

Abstract

For a fixed hidden label among four equatorial qubit preparations, we study its common-sign target with independent conditional copy stock and a fixed decision deadline. Adaptive and randomized nonadaptive equatorial readers have equal minimax risk through three copies for every shared efficiency and visibility. At four copies, exact rational policies and continuous nonadaptive certificates establish strict feedback advantages in both ideal and specified noisy parents. The nominal one-percent copy minima lie in [13, 14] for both classes; under the original asymmetric fixed-gain/phase uncertainty each lies in [16, 17]. A robust repeated-angle minimax frontier extends to arbitrary common-positive frozen-angle lotteries, and convex-source theorems identify when a single reader suffices globally. Compact policy graphs, exact dual certificates and retained-record witnesses are supplied. Copy savings are distinct from total apparatus savings, and no new Born-selection law is inferred.

1 Task, contributions and operational contract

This paper develops local quantum-reader decision problems within Shadow Theory’s pre-existing source-readout programme [9]. The results below quantify target access under the declared preparation, measurement and retained-record conditions. Let $h \in \{X+, X-, Y+, Y-\}$ label $\rho_{X\pm} = (I \pm X)/2$ and $\rho_{Y\pm} = (I \pm Y)/2$. The target is the common sign. One label remains fixed for a complete experiment: the source provides $\rho_h^{\otimes B}$, rather than redrawing the axis on each specimen. The preparation controller is unavailable. A trial destructively measures one fresh copy and all commanded settings, outcomes, nulls and specified timings remain retained. Every fixed deadline requires a sign; no postselection or abstention is allowed. Randomized nonadaptive protocols may choose a joint distribution of entire schedules before outcomes and retain their seed. Adaptive protocols choose settings from past accessible records and ready coins. Neither class has coherent quantum memory returning between specimens. The effects specified below and independent ready stock define the statistical comparison; they do not identify unread detector baths. Equatorial phase control is realized by the declared two-mode phase rotation, mixing and detection interface, with phase, clock and reference costs separate from specimen count. The results give constructive source distinctions inaccessible to separate-mode intensities, together with restrictions under which feedback, randomization and calibration actually change attainable readout risk. Adaptive multiple-copy discrimination and minimax quantum decision theory have established treatments[5, 7]. The known-binary comparison in the former does not supply a bound for the present grouped target or for all arbitrary preselected measurement schedules. Here the hidden axis stays fixed across the complete conditional-copy stock. Standard quantum discrimination, minimax symmetrization and representation methods have antecedents cited throughout; the contributions are certified portions of the declared-task frontiers, rational constructions and independently checkable

continuous certificates. There is no worldwide novelty or empirical validation claim. Ordinary Born probabilities are assumptions.

1.1 Decision risks and the retained classical interface

Write $\mathcal{C}_{\text{ad}}$ and $\mathcal{C}_{\text{na}}$ for the two policy classes above. A policy Π has an independent, retained seed S , commands θ_t and a terminal sign D_B . In $\mathcal{C}_{\text{ad}}$, θ_t may depend on $(S, o_1, \dots, o_{t-1})$; in $\mathcal{C}_{\text{na}}$ its entire schedule depends on S alone. The terminal decoder may use all recorded outcomes in either class. Its label risk and minimax value are

$$r_h(\Pi; \eta, v) = \Pr_{h, \Pi; \eta, v} \{D_B \neq R(h)\}, \quad e_B^{\mathcal{C}}(\eta, v) = \inf_{\Pi \in \mathcal{C}} \max_h r_h(\Pi; \eta, v).$$

For a prior w , Bayes risk instead averages the label risks with weights w_h ; a Bayes-to-minimax equality is used only when a stated symmetry or least-favourable comparison proves it. In the robust task, the maximum additionally ranges over one fixed nuisance tuple for the complete experiment, and one policy is chosen before that tuple is revealed.

The calibrated shared-noise comparisons use the conditional trial kernels

$$K_{\theta, h}(\pm) = \frac{\eta}{2}(1 \pm v b_h \cdot n_{\theta}), \quad K_{\theta, h}(0) = 1 - \eta,$$

where b_h is the pure preparation's equatorial Bloch vector and $n_{\theta} = (\cos \theta, \sin \theta)$. Conditional on the fixed label and admitted nuisance, the fresh specimen and ready apparatus stock produce this kernel at each history. Equal one-copy marginals alone would not establish that sequential law.

The retained decision record contains commanded settings, ternary outcomes, ready seeds and declared timings. Timings included in this statistical interface are determined by those records or have a source-independent conditional law; additional source-dependent arrival-time information would require an enlarged kernel and a new comparison. The theorem does not identify unread baths, time-resolved absorbing instruments or residual quantum outputs merely from their terminal effects.

Organization and conventions. Sections 2 and 3 give the small-horizon continuous-menu optima and attainable four-copy policies. The nonadaptive certificates prove the cross-class separation. Sections 5 and 6 then treat the fixed asymmetric uncertainty; larger-horizon bounds and compact policies follow. The appendices explicitly restrict access to the X/Y menu. Angles are in radians, probabilities and Bloch gains are dimensionless, and physical pulse times and angular rates retain the units stated in their interaction models. All certificate paths below are relative to `supplement/readers/`; reproduction expands compressed inputs in a disposable tree.

2 Continuous equatorial access: exact small-horizon optima and an improved fixed reader

Permit every equatorial setting θ , with $\sigma_{\theta} = \cos \theta X + \sin \theta Y$, and keep the same shared calibrated efficiency and visibility:

$$E_{\pm|\theta} = \frac{\eta}{2}(I \pm v \sigma_{\theta}), \quad E_{0|\theta} = (1 - \eta)I, \quad 0 \leq \eta, v \leq 1.$$

The source label $h \in \{X+, X-, Y+, Y-\}$ is fixed, its controller is unread, and independently prepared copies are supplied conditional on h . A policy selects one setting per fresh copy and always decides the sign target at the fixed deadline. All actual settings, signs, nulls and timing records remain retained. The expanded class includes randomized nonadaptive schedules and their terminal decoders,

as well as nonanticipating adaptive policies. No collective POVM, readable source register or returning environment is added.

A phase rotation followed by a Hadamard beam splitter makes the bisector $\theta = \pi/4$ a physically specified extension of the reader menu. For example, apply $U_\theta = \exp(+i\theta Z/2)$ before the Hadamard and computational-path detection. Its sign effects are $U_\theta^\dagger(I \pm X)U_\theta/2 = (I \pm \sigma_\theta)/2$. A phase Hamiltonian $H_z = \hbar\Omega_z Z/2$ realizes this command with signed pulse area $\int \Omega_z(t) dt = -\theta$ and constant-drive duration $|\theta|/|\Omega_z|$. This charges phase-reference availability, preparation, calibration and switching time; the following optima count copies rather than total apparatus cost.

2.1 A posteriori formula and exact one/two-copy theorem

The full continuous menu is closed under X/Y exchange and global sign reversal. These operations permute source labels transitively, transform settings as $\theta \mapsto \pi/2 - \theta$ or reverse outcomes, and transform the final sign consistently. The shared η, v noise law is invariant. Consequently averaging a uniform-prior Bayes-optimal policy over this group makes all four source risks equal. Uniform-prior Bayes lower bounds and randomized minimax optima therefore agree in both the adaptive and the randomized-nonadaptive continuous classes. This argument does not transfer unchanged to unequal axis responses.

Lemma 2.1 (Exact one-step posterior risk). *For unnormalized nonnegative hypothesis weights $w = (w_{X+}, w_{X-}, w_{Y+}, w_{Y-})$, put*

$$W = \sum_h w_h, \quad d = w_{X+} - w_{X-} + w_{Y+} - w_{Y-}, \quad b = (w_{X+} + w_{X-}, w_{Y+} + w_{Y-}).$$

The optimal unnormalized Bayes loss with one remaining equatorial measurement is

$$L_1(w) = \frac{1}{2}\{W - (1 - \eta)|d| - \eta \max(|d|, v\|b\|_2)\}.$$

A reader parallel to b attains it; when $b = 0$ any reader suffices.

Proof. The plus-minus target mass difference on a sign outcome $r = \pm 1$ is $\eta[d + rvb \cdot n_\theta]/2$, while on a null it is $(1 - \eta)d$. For a two-target decision problem the Bayes error is half the total mass minus half the sum of absolute target-mass differences. The identity $|d + c| + |d - c| = 2 \max(|d|, |c|)$ reduces the optimization to maximizing $|b \cdot n_\theta|$, whose maximum is $\|b\|_2$. This is an exact optimization over the entire circle, not a grid scan. $\square$

Theorem 2.2 (Exact continuous optima at one and two copies). *For the shared-noise continuous equatorial class, the randomized adaptive and randomized nonadaptive minimax optima coincide at horizons one and two:*

$$e_1^{\text{ad}} = e_1^{\text{na}} = \frac{1}{2} - \frac{\eta v}{2\sqrt{2}}, \quad e_2^{\text{ad}} = e_2^{\text{na}} = \frac{1}{2} - \frac{\eta(2 - \eta)v}{2\sqrt{2}}.$$

Repeated use of the fixed bisector $\theta = \pi/4$, followed by majority of the nonnull signs and a fair tie decision, attains both optima.

Proof. At the uniform prior, $d = 0$ and $b = (1/2, 1/2)$, proving the one-copy statement through Lemma 2.1. For two copies, let the first angle be θ . After its sign r , the unnormalized posterior is $w_{a,s} = \eta(1 + rvs n_{\theta,a})/8$. Its total mass is $\eta/2$, its target difference is $\eta rv(\cos \theta + \sin \theta)/4$, and its vector $b = (\eta/4, \eta/4)$ is independent of that outcome and angle. Thus $|d| \leq v\|b\|_2$ and the bisector is optimal

on either sign branch. The null branch retains the uniform prior scaled by $1 - \eta$ and also admits an optimal bisector. Adding all three optimal branch losses gives

$$L_2(\theta) = \frac{1}{2} - \frac{\eta v \sqrt{2}}{4} - \frac{\eta(1-\eta)v}{4} |\cos \theta + \sin \theta|.$$

Its minimum uses $|\cos \theta + \sin \theta| = \sqrt{2}$, giving the formula. At $\eta = 1$ the first angle is immaterial to the uniform Bayes value; a repeated bisector still attains the minimax optimum. When $\eta v = 0$ all target information vanishes. Repeated bisector records have identical laws for $X+$ and $Y+$, and identical laws for $X-$ and $Y-$, so its four source risks already agree. The symmetry argument upgrades the Bayes lower bound to minimax equality. Since this attaining protocol is nonadaptive, the two admitted classes have equal optima. $\square$

The one-copy bisector is the familiar Breidbart direction for the unitarily equivalent BB84 subset/bit task[4]; it is not claimed as a new measurement discovery. The addition relative to the declared source-readout programme is the explicit noisy two-copy continuous-class theorem and its operational comparison. At $\eta = v = 1$, both errors are $(1 - 1/\sqrt{2})/2 \simeq 0.1464466$, below the earlier X/Y value $1/4$. At $\eta = 4/5, v = 9/10$, the exact errors are $1/2 - (9/50)\sqrt{2}$ and $1/2 - (27/125)\sqrt{2}$.

3 Continuous equatorial readers: exact equality through three copies

We establish equality of adaptive and randomized nonadaptive continuous-angle readers through three copies for every shared efficiency and visibility, then construct genuinely angle-adaptive four-copy readers. The latter are attainable upper bounds; universal nonadaptive comparison requires the separate lower certificates. No finite angular scan is used as a surrogate for the continuous class.

3.1 The complete local class and exact three-copy theorem

Use the four fixed source labels, their common sign target and the shared calibrated effects

$$E_{\pm|\theta} = \frac{\eta}{2}(I \pm v\sigma_\theta), \quad E_{0|\theta} = (1 - \eta)I, \quad 0 \leq \eta, v \leq 1.$$

All actual angles, signs and nulls remain retained. The label controller is inaccessible; copies are independent conditional on that one fixed label. Adaptive settings are nonanticipating functions of accessible records and private coins. Nonadaptive schedules may be jointly randomized in advance, with arbitrary terminal decoders and their seed retained. A fixed deadline requires a sign decision, with neither postselection nor abstention. The full equatorial menu and common detector law are closed under axis exchange and global source-sign reversal. Uniform-prior Bayes risk is a minimax lower bound; symmetrization promotes an attaining Bayes reader to minimax. The repeated bisector already has equal error on all four labels.

Let $q_n(\beta)$ be fair-tie majority error for n independent signs of positive bias β . In particular

$$q_0 = \frac{1}{2}, \quad q_1 = q_2 = \frac{1}{2} - \frac{\beta}{2}, \quad q_3 = \frac{1}{2} - \frac{3\beta}{4} + \frac{\beta^3}{4}.$$

Theorem 3.1 (Continuous equality through three copies). *For the complete local equatorial class, all $0 \leq \eta, v \leq 1$,*

$$e_3^{\text{ad}} = e_3^{\text{na}} = \sum_{n=0}^3 \binom{3}{n} \eta^n (1 - \eta)^{3-n} q_n(v/\sqrt{2}).$$

Repeated measurement at $\theta = \pi/4$, followed by majority of the nonnull signs and a fair tie, attains it. Equivalently,

$$e_3 = \frac{1}{2} - \frac{v(3\eta - 3\eta^2 + \frac{3}{2}\eta^3)}{2\sqrt{2}} + \frac{\eta^3 v^3}{8\sqrt{2}}.$$

At $\eta = 4/5, v = 9/10$ this is $1/2 - (4023/15625)\sqrt{2} \simeq 0.135879605669$. The theorem adds no assertion about four or more copies.

3.2 A continuous two-control inequality

The following lemma supplies the load-bearing continuous optimization. It covers negative as well as positive direction components.

Lemma 3.2. *Let (c, s) and (p, q) be unit vectors, $u \in [0, 1]$, and put*

$$A = c + s, \quad B = p + q, \quad C = cp + sq, \quad Q = c^2 p^2 + s^2 q^2.$$

Then

$$\sum_{r=\pm 1} \max\{|A + rB|, \sqrt{2 + 2ruC + u^2 Q}\} \leq \sqrt{2}(3 - u/2).$$

Equality is attained when both directions are the positive bisector.

Proof. There are four branches of the two maxima. If both choose absolute linear terms, their sum is $2 \max(|A|, |B|) \leq 2\sqrt{2}$. If both choose square roots, Cauchy–Schwarz and $Q \leq 1$ give $2\sqrt{2 + u^2}$. This is at most the claimed bound because $2(3 - u/2)^2 - 4(2 + u^2) = 10 - 6u - 7u^2/2 \geq 1/2$.

For a mixed branch, changing (p, q) to its negative if necessary reduces the required estimate to

$$M = |A + B| + \sqrt{2 - 2uC + u^2 Q} \leq \sqrt{2}(3 - u/2).$$

Write $(c, s) = (\cos \theta, \sin \theta)$, $(p, q) = (\cos \phi, \sin \phi)$ and set $\alpha = \theta - \pi/4$, $\beta = \phi - \pi/4$, $x = |\cos((\alpha + \beta)/2)|$, $y = |\cos((\alpha - \beta)/2)|$. Trigonometric identities give

$$\frac{M}{\sqrt{2}} = G(x, y) = 2xy + R, \quad R^2 = H^2 + u^2 x^2 (1 - x^2), \quad H = 1 + u/2 - uy^2.$$

Here $x, y \in [0, 1]$, $R \geq H \geq 1 - u/2 \geq 1/2$. If $x \geq y$, then $\partial_y G = 2x - 2uyH/R \geq 0$, since $xR \geq uyH$. If $y \geq x$, then $\partial_x G = 2y + u^2 x(1 - 2x^2)/R \geq 0$. For $x^2 \leq 1/2$ this is immediate. Otherwise $2yR \geq y \geq x \geq u^2 x(2x^2 - 1)$. Thus increasing the smaller coordinate to the larger cannot decrease G . On the diagonal, with $X = x^2 = y^2$,

$$G = 2X + \sqrt{(1 + u/2)^2 - 2uX}.$$

Its derivative is $2 - u/\sqrt{(1 + u/2)^2 - 2uX} \geq 0$ because the square root is at least $1 - u/2$ and $u \leq 1$. Hence its maximum is $G(1, 1) = 3 - u/2$. This proves both mixed branches and the lemma. $\square$

Proof of Theorem 3.1. First set $\eta = 1$. It suffices to lower-bound the uniform Bayes risk of every deterministic adaptive policy: independent random seeds only average such risks. Let the first direction be (c, s) . Its two equiprobable outcomes have posteriors related by global source-sign reversal, so their optimal continuation values agree. Let (p, q) be the second direction on the first-plus branch. After its outcome r , the unnormalized posterior weights have target mass difference $v(A + rB)/8$ and axis mass vector $((1 + rv^2 cp)/8, (1 + rv^2 sq)/8)$.

The preceding exact one-remaining-copy formula, obtained from $|d+z|+|d-z|=2\max(|d|,|z|)$, optimizes the third direction over the whole circle. Adding the two second outcomes and the equal first-sign branches gives three-copy Bayes loss

$$\frac{1}{2} - \frac{v}{8} \sum_{r=\pm 1} \max\{|A+rB|, \sqrt{2+2rv^2C+v^4Q}\}.$$

Lemma 3.2, with $u=v^2$, bounds this below $1/2 - 3v/(4\sqrt{2}) + v^3/(8\sqrt{2}) = q_3(v/\sqrt{2})$. The repeated bisector attains this value, including arbitrary terminal decision rules in the comparison.

For $0 \leq \eta \leq 1$, condition on the entire erasure pattern of the three trials. The pattern is independent of the source label and signs. A policy conditioned on a nominated pattern with n clicks induces an admitted n -copy adaptive ideal-efficiency reader; giving the controller that pattern in advance can only make the lower bound more permissive. The preceding $n=1, 2$ optima and the just-proved $n=3$ optimum therefore lower-bound each conditional risk by $q_n(v/\sqrt{2})$. Averaging over the binomial click count yields the claimed lower bound. The fixed bisector attains each conditional value simultaneously. Its four errors are equal, so the Bayes lower bound is an attained minimax value for both classes. $\square$

The conditional-erasure argument uses independent sign-blind erasures. It does not assert that a fresh reset removes arbitrary correlations, or that an unknown future pattern is physically accessible to the deployed reader. Advance pattern disclosure is used only as a lower-bound relaxation.

3.3 Constructive four-copy angle adaptation

At four copies we provide explicit frozen policy tables, not numerical claims of globally optimal adaptive controls. Both the ideal and nominal policies use forty declared measurement nodes and eighty-one terminal ternary records, including impossible ideal-null branches with harmless defaults. Each direction is an exactly unit rational pair

$$n(t) = ((1-t^2)/(1+t^2), 2t/(1+t^2)), \quad t \in \mathbb{Q}.$$

Numerical search selected candidate angles; rational conversion and exact full-word verification establish the delivered performance. Search failure or floating convergence is not an optimality certificate.

The base policy first uses a fixed angle θ , and on a first nonnull outcome uses ϕ . Its third setting depends on whether the second sign agrees with the first, disagrees, or is null. A first-null branch uses the bisector candidate. At the last node a rational approximation to the posterior axis-mass direction is used. Where the nominated posterior gives no possible Bayes-risk improvement from the last measurement, its setting is instead held at the previous setting. The published leaf decisions remain specified on every record, including zero-probability ideal-null defaults. This is a flat nominal-risk choice, not a claim that every leaf decision is identical under a changed noise model. The final trial still occurs and is fully archived. This saves potential actuator motion; it does not silently convert to an optional-stopping or reduced-copy experiment.

Approximate design angles, in radians, are:

Parent	θ	ϕ	same sign	opposite sign	null
Ideal	.17468079	1.39611555	.01091177	.78539702	.79955276
Nominal	.75070933	1.09092805	1.28465859	.49454527	.79086328

The JSON files, not these rounded displays, define the actual rational commands, all branches and terminal decisions. Root symmetrization uses two fair bits: exchange X and Y in every setting, and/or reverse every recorded sign and the final decision. This group acts transitively on the four labels and respects the complete admitted continuous class and common noise law. Averaging the base policy's

four exact errors therefore gives the same error on each label. The seed is available to the decoder and counted as controller stock; it does not reveal the hidden source label.

For the ideal parent $(\eta, v) = (1, 1)$ the resulting minimax attainable upper error is strictly less than 0.030605 (actual exact-rational value approximately 0.0306041631615). For the nominal parent $(4/5, 9/10)$ it is strictly less than 0.103471 (approximately 0.103470608280). These upper bounds are proved by summing each of the 3^4 retained-record probabilities for each source label. They do not by themselves exclude better nonadaptive readers or establish the adaptive optimum. The separate continuous nonadaptive certificate must be used for any cross-class separation.

3.4 Verification and limits

The independent standard-library checker verifies 630 rational instances of the mixed geometry inequality using exact signs in $\mathbb{Q}(\sqrt{2})$, then checks all 81 words for each policy, normalization on each source, all root symmetry permutations and the stated rational upper bounds. The analytical inequality proves the entire continuous three-copy theorem; finite tests alone do not do so. Frozen rational policy tables provide independently checkable attainability without rerunning the exploratory optimizer. Optional exploratory scripts use SciPy; core certification uses only the Python standard library.

The equality theorem and policy constructions extend the specified source-readout decision problem. The fixed inaccessible source label, conditional-copy independence, common calibrated detector law, controller coins, phase reference, clock and complete retained records remain premises. No complete residual-instrument equivalence, Born-selection law, consciousness criterion, trine-resource theorem or universal physical cost advantage is claimed by these decision-risk results.

4 Certified continuous nonadaptive obstructions

4.1 Four ideal copies: an all-angle lower certificate

Theorem 4.1 (Continuous nonadaptive four-copy lower bound). *In the ideal parent $\eta = v = 1$, every randomized nonadaptive four-copy continuous equatorial reader, with arbitrary terminal decision and retained settings, has worst-label sign error at least $1/32$.*

Proof. For a fixed schedule, an optimal uniform-prior decoder can be chosen odd on four sign bits. Its Walsh expansion has coefficients only on singletons and triples. There are $2^8 = 256$ such odd decoders; exact enumeration divides them into seven orbits under independent outcome reversals and permutations of the four measurements, of respective sizes 8, 64, 32, 48, 64, 32, 8. For coefficients α_i and β_I , its target correlation is

$$C = \frac{1}{2} \left\{ \sum_i \alpha_i (c_i + s_i) + \sum_{|I|=3} \beta_I \left(\prod_{i \in I} c_i + \prod_{i \in I} s_i \right) \right\}.$$

Eliminate the fourth measurement analytically. With

$$\begin{aligned} A &= \alpha_4 + \beta_{124} c_1 c_2 + \beta_{134} c_1 c_3 + \beta_{234} c_2 c_3, \\ B &= \alpha_4 + \beta_{124} s_1 s_2 + \beta_{134} s_1 s_3 + \beta_{234} s_2 s_3, \\ F_0 &= \sum_{i=1}^3 \alpha_i (c_i + s_i) + \beta_{123} (c_1 c_2 c_3 + s_1 s_2 s_3), \end{aligned}$$

its largest possible correlation is $(F_0 + \sqrt{A^2 + B^2})/2$. This maximization covers the full fourth circle, not a discretized setting.

For the first three directions use the rational circle

$$(c_i, s_i) = \epsilon_i \left(\frac{1 - t_i^2}{1 + t_i^2}, \frac{2t_i}{1 + t_i^2} \right), \quad -1 \leq t_i \leq 1, \quad \epsilon_i \in \{-1, 1\}.$$

The hemisphere signs are explicitly enumerated: suppressing them after choosing a decoder representative would not establish a full-circle bound. Together with seven decoder orbits this gives 56 cube calculations.

The delivered certificate partitions each cube into a finite prefix-free set of dyadic octree boxes. On a box, exact endpoint formulas enclose each c_i, s_i ; c_i additionally has maximum one when its interval contains zero. Rational interval products and sums give an upper bound on F_0 and intervals for A, B . An integer-square-root enclosure bounds $\sqrt{\max A^2 + \max B^2}$ upwards. Every resulting correlation bound is at most 15/16. The exact verifier checks all 143633 leaves, all orbit and sign variants, prefix-freeness and total cube volume one; the largest octree depth is eight. These finite exact inequalities prove the continuum bound. Floating calculations only constructed a candidate partition and are not accepted as proof of any leaf.

Thus uniform-prior error is at least $(1 - 15/16)/2 = 1/32$ for every schedule. A randomized schedule is an average of such experiments with its retained seed and cannot lower this average Bayes bound. Worst-label error is no smaller than uniform-prior error, proving the claim. $\square$

The delivered `research10/nonadaptive/four_certificate.json.gz` contains the finite proof object. The reproduction runner expands it without changing the underlying JSON content (the expanded hash is verified); `research10/nonadaptive/verify_four.py` recomputes every accepted inequality with standard-library rational arithmetic and integer square roots. The exploratory partition generator is retained in the historical evidence archive and is not required for verification. The bound 1/32 is deliberately sufficient for discrimination and is not claimed as the exact four-copy nonadaptive optimum. Floating exploratory optimization suggests a minimum near 0.03740607, but that number carries no global optimality claim. The certificate concerns the ideal parent only; it is not the principal nominal noisy four-copy optimum.

4.2 A noisy full-continuum separation at the principal benchmark

The ideal separation is not the only surviving regime. A second exact certificate treats the principal shared detector parent, retaining all nulls and all measurement settings.

Theorem 4.2 (Nominal continuous nonadaptive four-copy lower bound). *At $\eta = 4/5$, $v = 9/10$, every randomized nonadaptive four-copy continuous equatorial reader has worst-label sign error at least $13/125 = 0.104$.*

Proof. Condition on the entire source-independent erasure pattern and maximize the terminal decision separately on each pattern. Denote by C_k the optimal uniform-prior target correlation for the nominated k nonnull directions. For a single fixed direction i , put $a_i = c_i + s_i$; then

$$C_1(i) = \frac{v}{2}|a_i|, \quad C_2(i, j) = \frac{v}{2} \max(|a_i|, |a_j|).$$

For three specified directions,

$$C_3(1, 2, 3) = \frac{1}{8} \sum_{r_2, r_3 = \pm 1} \left| v \sum_{i=1}^3 r_i a_i + v^3 r_1 r_2 r_3 (c_1 c_2 c_3 + s_1 s_2 s_3) \right|, \quad r_1 = 1.$$

The preceding exact three-copy theorem gives global upper correlations $\overline{C}_1 = \overline{C}_2 = v/\sqrt{2}$ and $\overline{C}_3 = v(6 - v^2)/(4\sqrt{2})$.

Fix the first three directions, and grant separate fourth-direction optimization to the different erasure patterns. This relaxation need not be a realizable nonadaptive schedule; it only increases the upper correlation used for the lower error bound. With $z = 1 - \eta$, the resulting upper bound is

$$\begin{aligned} U = & \eta z^3 \{C_1(1) + C_1(2) + C_1(3) + \overline{C}_1\} \\ & + \eta^2 z^2 \{C_2(1, 2) + C_2(1, 3) + C_2(2, 3) + 3\overline{C}_1\} \\ & + \eta^3 z \{C_3(1, 2, 3) + 3\overline{C}_3\} + \eta^4 C_4^{\text{free}}(1, 2, 3). \end{aligned}$$

For C_4^{free} , enumerate the same odd four-bit decoders as in Theorem 4.1. Multiply every triple coefficient by v^2 in F_0, A, B , then its free-fourth correlation is $v(F_0 + \sqrt{A^2 + B^2})/2$. The first-three conditional term can also be capped by $\overline{C}_3$.

The original objective is invariant under simultaneous schedule/decoder outcome reversals and permutations. Apply these symmetries before choosing the relaxed fourth direction. The lower-click optimum terms are invariant under outcome reversals. Thus the seven decoder orbits and eight hemisphere sign choices again cover all original schedules. Relaxing first and then silently imposing a symmetry of the relaxed formula would not be justified.

The delivered `research10/nonadaptive/nominal_certificate.json.gz`, expanded by the reproduction runner, bounds $U \leq 99/125$ on all 56 cube variants. It contains 4230653 leaves, with maximum depth eleven. Every inequality is checked using outward dyadic intervals with scale 2^{60} . Interval sums, rational scalar products and general products use integer floor and ceiling. Integer square roots are rounded upwards. No floating number is used to accept a bound. Exact endpoint formulas for the rational circle give the starting intervals.

For clarity, at the stated η, v let c_1, c_2, c_3 denote upper interval numerators of the three braces above, and let f_0, A, B be upper/intermediate numerators for the free-fourth decoder expression. All are expressed at the same dyadic scale. The total upper numerator is rounded from

$$\frac{20c_1 + 80c_2 + 320c_3 + 576\{f_0 + \sqrt{\max A^2 + \max B^2}\}}{3125}.$$

The constants enclose $\overline{C}_1 = 9\sqrt{2}/20$ and $\overline{C}_3 = 4671\sqrt{2}/8000$ upwards. Prefix-freeness and exact cube-volume one establish full coverage, not merely a collection of point checks. The decoder orbit enumeration and all hemisphere signs are checked again. Hence every schedule has average error at least $(1 - 99/125)/2 = 13/125$; retaining its random seed and mixing schedules cannot improve that average lower bound. Worst-label risk is no smaller. $\square$

Combining this universal nonadaptive bound with the separately certified four-copy adaptive error < 0.103471 proves strict feedback benefit in the principal noisy parent over *all* continuous nonadaptive equatorial schedules. This is stronger than comparing selected angle allocations. The certificate is not an exact nonadaptive optimum; numerical exploration near 0.10476238 remains separately labelled.

For the selected tolerance $519/5000 = 0.1038$, the exact copy minima are four adaptive versus five randomized nonadaptive. The three-copy adaptive equality gives a lower risk $\simeq 0.13587961$, the four-copy adaptive construction passes, and the universal four-copy nonadaptive lower bound fails the target. Five repeated bisectors attain $1/2 - (11631897/39062500)\sqrt{2} \simeq 0.07888035$. This post hoc tolerance demonstrates an attained copy separation; it is not the original one-percent target.

The strict separation also persists on a specified open parameter box. For a frozen policy, or uniformly for every nonadaptive competitor, changing shared parameters by $(\delta\eta, \delta v)$ changes the one-trial record law by TV at most $\kappa = |\delta\eta| + |\delta v|/2$. Sequential coupling bounds any four-copy target risk change by 4κ . Thus the strict nominal gap of more than 0.000529 remains positive whenever $8\kappa < 0.000529$. In particular $|\delta\eta|, |\delta v| \leq 10^{-5}$ leaves a certified gap of more than 0.000409. The adaptive

policy stays frozen; inaccessible parameters do not select a different policy. This is a shared-noise robustness statement, not a solution of the separate asymmetric fixed-unknown-gain problem.

`verify_nominal_integer.py` carries the efficient exact replay; `verify_nominal.py` supplies a separately written Fraction interval reference. 961 selected leaf bounds are cross-checked with that reference. An independent audit checks every leaf with a root-free final squared allowance. The full certificate is large because it covers continuous settings rigorously, rather than claiming a finite grid solves the continuum. It is compressed in the supporting package.

4.3 The resulting ideal feedback separation and its limits

The separately constructed adaptive four-copy policy has a certified ideal error strictly below $1/32$. The preceding lower bound therefore covers *every* continuous nonadaptive competitor, rather than only the old X/Y menu or selected fixed-angle designs. This is a genuine continuous-access feedback separation under the stated preparation and record contract.

At the explicitly selected tolerance 0.031, the exhibited adaptive policy passes with four copies. Every nonadaptive reader with at most four copies fails by the $1/32$ bound and monotonicity. A five-copy repeated bisection has error $1/2 - (43/128)\sqrt{2} \simeq 0.02491263$, so five suffice nonadaptively. Thus the adaptive four-copy construction and exact nonadaptive five-copy minimum give a concrete task separation at that post hoc illustrative tolerance. The separate three-copy adaptive theorem gives the same ideal lower value $1/2 - (5/16)\sqrt{2} > 0.031$, so four is also the exact adaptive minimum. Thus the exact copy minima at this selected tolerance are four adaptive versus five randomized nonadaptive copies. The tolerance was chosen after the frontier calculation and is not the original one-percent target. Full large-horizon continuous optima, exact four-copy optimum values and general robust adaptation remain separate questions.

5 A robust frozen reader without X/Y symmetry

Admit unequal contrast transfer, a common unknown phase offset and a broader nominal equatorial direction. One frozen setting and one parameter-independent decoder suffice for a sharp restricted minimax theorem. The restriction is consequential: it optimizes the common-positive, repeated-single-angle class, not all equatorial adaptive policies or arbitrary collective measurements.

5.1 Physical parent and uncertainty contract

The source remains the four pure preparations $\rho_{X\pm}, \rho_{Y\pm}$ and the nominated binary sign target. Conditional on its fixed inaccessible label, fresh copies and apparatus stock are independent. Before the phase reader apply a fixed, unknown Pauli channel with Bloch multipliers $(g_X, g_Y, g_X g_Y)$, where $0 \leq g_X, g_Y \leq 1$. It has the explicit nonnegative Pauli probabilities

$$(p_I, p_X, p_Y, p_Z) = \frac{1}{4}((1 + g_X)(1 + g_Y), (1 + g_X)(1 - g_Y), (1 - g_X)(1 + g_Y), (1 - g_X)(1 - g_Y)).$$

Their sum is one, so the channel is completely positive and trace preserving for every admitted gain pair. A coherent two-mode reader then measures the actual equatorial angle $\theta + \delta$, followed by independent sign-blind erasure. Thus

$$E_{\pm|\theta} = \frac{\eta}{2} [I \pm (g_X \cos(\theta + \delta)X + g_Y \sin(\theta + \delta)Y)],$$

$$E_{0|\theta} = (1 - \eta)I.$$

The channel precedes the phase measurement; changing this order can change the model. All effects are positive because their Bloch-vector length is at most one. The phase rotation uses the specified

$H_z = \hbar\Omega_z Z/2$ coupling, calibrated beam splitter and terminal detector. Phase and offset are in radians, Ω_z is in radians per second and pulse durations are in seconds. The Pauli/environment ready stock and erasure port are physical resources, not free mathematical randomness.

The unknown parameters are fixed throughout a deployment block and obey

$$\eta \in [\eta_0, 1], \quad g_X \in [a, 1], \quad g_Y \in [b, 1], \quad |\delta| \leq \Delta, \quad 0 \leq \Delta < \pi/4, \quad 0 < \eta_0 \leq 1, \quad 0 < a, b \leq 1.$$

No prior distribution over these nuisance parameters is assumed. The controller fixes one nominal angle before learning their values. The admitted common-positive class is $\theta \in [\Delta, \pi/2 - \Delta]$, so both axes have nonnegative target-aligned signal at every admitted offset. Every trial at the fixed deadline is counted; the decoder returns the majority of the nonnull signs, with an independent fair guess on a tie, including the all-null record. Nulls and all settings remain archived. No hidden label or calibration purifier is read.

Under source label $X \pm$ the nonnull sign has mean $\pm\beta_X$, where $\beta_X = g_X \cos(\theta + \delta)$; under $Y \pm$ it has mean $\pm\beta_Y$, where $\beta_Y = g_Y \sin(\theta + \delta)$. The four labels need not have equal signal strengths. The proof below uses a hard binary pair for its lower bound rather than incorrectly reusing uniform-prior minimax equality from the symmetric experiment.

5.2 Exact minimax angle and full record risk

Define the majority error for n successful signs of bias β by

$$q_n(\beta) = \sum_{k < n/2} \binom{n}{k} \left(\frac{1+\beta}{2}\right)^k \left(\frac{1-\beta}{2}\right)^{n-k} + \frac{\mathbf{1}_{2|n}}{2} \binom{n}{n/2} \left(\frac{1-\beta^2}{4}\right)^{n/2},$$

and the deadline risk, including all erasures, by

$$Q_B(\eta, \beta) = \sum_{n=0}^B \binom{B}{n} \eta^n (1-\eta)^{B-n} q_n(\beta).$$

The empty-sign case is $q_0 = 1/2$. The integer k counts correct signs under a positive-target preparation.

Theorem 5.1 (Sharp robust repeated-angle frontier). *For a fixed nominal angle in the declared common-positive class, the minimum over every terminal decoder of the worst error over source labels and fixed nuisance parameters is exactly*

$$Q_B(\eta_0, m(\theta)), \quad m(\theta) = \min\{a \cos(\theta + \Delta), b \sin(\theta - \Delta)\}.$$

The same fair-tie majority decoder attains this optimum for all parameters. The robust maximum signal is attained at

$$\tan \theta_* = \frac{a \cos \Delta + b \sin \Delta}{b \cos \Delta + a \sin \Delta},$$

$$m_* = \frac{ab \cos(2\Delta)}{\sqrt{a^2 + b^2 + 2ab \sin(2\Delta)}}.$$

Consequently $Q_B(\eta_0, m_*)$ is the exact robust optimum over all fixed angles and terminal decoders in this class. The setting and decoder do not depend on an inaccessible realized nuisance parameter.

Proof. For the symmetric binary pair of source signs with known bias β , the likelihood ratio of a record is $[(1+\beta)/(1-\beta)]^{n_+ - n_-}$; null factors cancel. Thus fair-tie majority has equal error on the two signs and is minimax, with error $Q_B(\eta, \beta)$. The formula remains valid at the endpoints by continuity. Increasing

the bias cannot increase its optimal error: a larger-bias sign channel can simulate a smaller positive bias by an independent sign flip. Increasing the efficiency also cannot increase the error: success counts are coupled monotonically, and the optimal binary Bayes errors q_n are nonincreasing in n because an extra sign can be ignored. This is a direct experiment-degradation argument, not an assumed global four-label symmetry.

For every admitted parameter and source axis, the bias is at least $m(\theta)$, so the same majority policy has error no larger than the displayed bound. Conversely, if the minimum is supplied by the X term, choose $g_X = a$, $\delta = \Delta$ and $\eta = \eta_0$. The two admitted labels $X+$, $X-$ alone force every decoder's worst error to be at least $Q_B(\eta_0, m(\theta))$. If the minimum comes from Y, use $g_Y = b$, $\delta = -\Delta$ and the Y sign pair. The adversary need not make both axes worst at the same offset; one hard pair suffices.

On the admitted angle interval, the X floor is decreasing and the Y floor increasing. Their unique intersection maximizes their minimum. Solving their equality gives the stated tangent. Substitution gives m_* , and $\Delta < \pi/4$ puts the intersection inside the admitted interval. No assertion is made that an outside-quadrant, multilevel decoder or an angle-varying protocol belongs to this restricted theorem. $\square$

The same bound extends to deterministic unknown efficiencies $\eta_t \geq \eta_0$ when their erasure events are independent of one another, source label, sign and control, while the gains and phase offset remain fixed. Condition on the erasure schedule; the retained signs still have one constant bias and success-count coupling applies. This does not license arbitrary history-dependent losses or varying per-trial biases. In particular, adding a weak sign to a majority of strong signs need not improve its error; the fixed-bias premise is essential to the efficiency argument.

5.3 Exact algebraic certificates and quantitative gain

For rational $a, b, t = \tan \Delta$, the chosen nominal tangent is the rational $(a + bt)/(b + at)$, and

$$m_*^2 = \frac{a^2 b^2 (1 - t^2)^2}{(1 + t^2)[(a^2 + b^2)(1 + t^2) + 4abt]} \in \mathbb{Q}.$$

The nominal bisector instead has squared worst bias $\min(a, b)^2(1 - t)^2/[2(1 + t^2)]$. For rational efficiency the risk is computed exactly in $\mathbb{Q}(\sqrt{m_*^2})$: every value is represented as $A + C\sqrt{m_*^2}$ with rational coefficients. Comparisons with 1/100 use integer squaring after checking signs, not a floating-point threshold. Integer-square-root enclosures provide independently checkable 96-bit rational lower/upper bounds.

For $\eta_0 = 4/5$, $a = 9/10$, $b = 4/5$, and $\tan \Delta = 1/100$, the robust setting is $\tan \theta_* = 908/809$, with $m_*^2 = 129574081296/369773223625$. The exact comparisons give

Reader in the same uncertainty class	trials	worst error
Optimized frozen angle	16	0.0102368487113...
Optimized frozen angle	17	0.00844219845891...
Nominal bisector	18	0.0107238296666...
Nominal bisector	19	0.00905666129798...

Monotonicity in budget proves exact one-percent minimum budgets seventeen and nineteen within the respective repeated-angle classes. This is a two-trial improvement with one setting in each implementation, not a comparison against a restricted X/Y baseline. For the larger uncertainty example $\eta_0 = .9$, $a = .8$, $b = .6$, $\tan \Delta = .1$, the corresponding minimum budgets are thirty and thirty-nine. Complete rational/algebraic certificates, including adjacent failing budgets, are supplied.

The isotropic special case also clarifies the earlier access restriction. At $g_X = g_Y = .9$, $\Delta = 0$ and $\eta = .8$, a repeated bisector reaches one-percent error at fourteen copies. Thus the axis-restricted

nineteen-versus- twenty-one X/Y threshold is not an unrestricted equatorial optimum. The continuous-reader results supply separate adaptive comparisons; this section does not claim that a repeated bisector is globally optimal.

5.4 Calibration and physical/resource limits

The theorem's intervals are admitted premises until certified by a trusted calibration experiment or an independent device model. Four trusted source preparations at the actual frozen setting can directly bound the four sign/null kernels without inferring the entire channel or nuisance parameters uniquely. Such a calibration must pay for its probes, confidence coverage, source trust and stability into deployment. A control model claiming a lower gain or bounded phase over an interval of settings needs additional justified coverage; samples at one angle do not certify all angles for an arbitrary detector.

The frozen majority controller needs a running signed count and deadline counter, plus a fair terminal tie bit; it has no setting-feedback branch. The angle itself must be prepared and maintained, with its phase reference and accuracy charged. Each physical realization must state whether the phase is applied separately to every fresh copy or configured once in a held passive reader. Their timing costs differ. The source copies are provided independently, not cloned from an unknown specimen, and the ordered archive is retained even when the decision uses only a count.

Robust quantum discrimination is established literature [6]; its uncertain-state model is not the detector model optimized here. The ordinary degradation argument has classical antecedents[2], and its special case is proved above. Our additions are the explicit within-programme robust angle frontier, physical Pauli-channel realization and exact algebraic operating region. No worldwide priority, complete retained-output equivalence, empirical detector validation, Born-selection law or source ontology is asserted.

The executable tests compare seventy-two independent ordered ternary-path sums, 1,500 rational geometry inequalities and 180 algebraic comparison/ enclosure cases. These supplement, rather than replace, the analytical minimax proof. The parameter-independent decoder, phase/gain uncertainty, retained nulls and access restrictions are part of every numerical result.

6 Robust readers: a convex-hull theorem and a global copy lower bound

We remove the repeated-angle restriction at one and two copies. The larger-horizon robust optimum remains unresolved; no fixed-angle optimum is silently promoted to it.

6.1 A general antipodal composite-source theorem

Let S be a nonempty compact subset of the equatorial Bloch disk. A hidden $r \in S$ and sign $s \in \{+1, -1\}$ are fixed throughout an experiment, and the source supplies independent copies of $\rho_{s,r} = (I + sr \cdot \sigma)/2$. The reader cannot inspect the source label or nuisance register. On each copy it may choose any equatorial projective direction, followed by independent sign-blind erasure of fixed unknown efficiency $\eta \in [\eta_0, 1]$, $0 < \eta_0 \leq 1$. It must decide s at a fixed deadline, retaining all settings and nulls. Randomized policies and arbitrary terminal decision rules are admitted. The worst error takes a supremum over (s, r, η) , without imposing a nuisance prior.

Let z be the closest point to the origin in $\text{conv } S$ and $m = \|z\|$. Compactness ensures existence. The supporting-plane property is

$$r \cdot z \geq m^2 \quad (r \in S). \quad (1)$$

Theorem 6.1 (Global robust equality through two copies). *In this complete equatorial policy class, the adaptive and randomized nonadaptive minimax values at horizons one and two are exactly*

$$e_1 = \frac{1 - \eta_0 m}{2}, \quad e_2 = \frac{1 - \eta_0(2 - \eta_0)m}{2}.$$

For $m > 0$, repeating the single direction z/m and deciding by majority of nonnull signs with a fair tie attains both values. For $m = 0$, an uninformed fair decision attains $1/2$ and neither class improves it. The same statement for a general three-dimensional Bloch set requires full Bloch-sphere projective access, not merely the equatorial menu.

Proof. The minimum-distance projection inequality gives (1), so measuring z/m has target-aligned bias at least m on every admitted source. At one copy its worst error is at most $(1 - \eta_0 m)/2$. At two copies majority with a fair tie has sign bias m whenever at least one copy survives erasure, giving the stated upper bound. These bounds are valid for unknown fixed r and unknown fixed η , and use one controller policy for all their values.

For the lower bound, choose a finitely supported distribution μ on S with barycentre z ; planar Carathéodory gives at most three support points. This distribution is a mathematical least-favourable comparison, not a new prior assumed by the robust protocol. Choose a fair source sign and efficiency η_0 . Each sampled r remains fixed for all copies. Define the two conditional target mixtures

$$\Omega_s^{(n)} = \int_S \rho_{s,r}^{\otimes n} d\mu(r).$$

At one copy $\Omega_+^{(1)} - \Omega_-^{(1)} = z \cdot \sigma$, whose trace distance is m . At two copies, expanding each tensor product gives

$$\Omega_+^{(2)} - \Omega_-^{(2)} = \frac{1}{2} \{ (z \cdot \sigma) \otimes I + I \otimes (z \cdot \sigma) \}. \quad (2)$$

All quadratic terms cancel between opposite signs. Its eigenvalues are $m, 0, 0, -m$, so the two-copy trace distance is again m . This calculation uses the mixture of tensor powers with fixed r ; it never substitutes the tensor power of a one-copy mixture.

An orthogonal auxiliary null flag simulates the complete declared classical outcome kernel. For any fixed adaptive policy this simulation extends inductively over its histories and independent coins. We may therefore lower-bound the policy by allowing arbitrary joint measurement of the auxiliary states. At two copies the orthogonal zero-, one- and two-click blocks have total probability $(1 - \eta_0)^2$, $2\eta_0(1 - \eta_0)$ and η_0^2 . Their trace distances are $0, m, m$. Thus the total trace distance is $\eta_0(2 - \eta_0)m$; the one-copy value is $\eta_0 m$. Binary equal-prior minimum error $(1 - D)/2$ supplies the matching Bayes lower bounds. A robust worst error cannot be smaller than any valid comparison-prior average. This proves minimax equality. If $m = 0$, the same comparison makes both differences vanish. $\square$

This establishes a concrete interface where feedback and additional settings are unnecessary: through two copies, the closest point of the positive-target Bloch hull supplies a sufficient robust reader direction. It is a statistical decision result, not complete retained-instrument identification. Correlated accessible preparation records, a readable purifier, non-independent copies or a different erasure parent change the task. The trace-distance decision principle is established quantum theory; the composite-source cancellation and its application to the specified reader task are proved here.

6.2 Applying the theorem to asymmetric fixed calibration

The preceding detector kernels have the following equivalent statistical Bloch set for a positive sign:

$$S = \{g(\cos \delta, -\sin \delta) : g \in [a, 1], |\delta| \leq \Delta\} \cup \{g(\sin \delta, \cos \delta) : g \in [b, 1], |\delta| \leq \Delta\}.$$

Indeed the X source gives $g_X \cos(\theta + \delta)$ and the Y source gives $g_Y \sin(\theta + \delta)$ when paired with n_θ . Absorbing fixed nuisance into this set is a representation of the admitted classical kernels and decision risks. It neither identifies complete residual channels nor makes the nuisance accessible to the controller. Define

$$u = (a \cos \Delta, -a \sin \Delta), \quad v = (-b \sin \Delta, b \cos \Delta).$$

They are admitted positive-target vectors: choose gain floors, with $\delta = +\Delta$ on X and $\delta = -\Delta$ on Y. For the lower-bound prior this nuisance is correlated with the fixed hidden axis, and remains fixed throughout that experiment. This is allowed in the supremum over source label and nuisance; it is not a per-node adversary or a claim that a physical random calibration law follows that prior.

Put $c = u \cdot v = -ab \sin 2\Delta$. The closest point on their segment is

$$z = pu + (1 - p)v, \quad p = \frac{b^2 - c}{a^2 + b^2 - 2c}, \quad \|z\| = m_* = \frac{ab \cos 2\Delta}{\sqrt{a^2 + b^2 + 2ab \sin 2\Delta}}.$$

Here $0 < p < 1$, and $z \cdot u = z \cdot v = \|z\|^2$. The preceding optimized direction has signal at least m_* on the whole uncertainty set, so the supporting-plane inequality also holds for all its other admissible vectors. Consequently z is closest in the whole hull, not only this segment. The global robust optima are therefore

$$e_1 = (1 - \eta_0 m_*)/2, \quad e_2 = (1 - \eta_0(2 - \eta_0)m_*)/2,$$

over every continuous equatorial adaptive or randomized nonadaptive reader. No common-positive restriction on competitors remains at these horizons. The attaining reader retains the robust angle

$$\tan \theta_* = \frac{a \cos \Delta + b \sin \Delta}{b \cos \Delta + a \sin \Delta}.$$

For $\eta_0 = 4/5, a = 9/10, b = 4/5, \tan \Delta = 1/100$, one obtains $p = 654464/1478945$, $m_*^2 = 129574081296/369773223625$ and errors approximately 0.26321655725 and 0.21585986870. These are global small-horizon values, not a claim about the one-percent threshold or unrestricted collective readers in a different physical parent.

7 Local-effect certificates for larger copy budgets

Retaining local partial-transpose invariance gives exact all-angle comparison bounds without enumerating a continuum of policy trees. The relaxation is a lower-bound device, not a physically available collective reader.

7.1 Which effects every admitted reader produces

Rotate the original equatorial plane to the real X/Z Bloch plane. This is a mathematical common basis choice, not a new source preparation or accessible reference. A local outcome effect is then a real symmetric 2×2 matrix. For a fixed adaptive history, its joint probability on independent copies is computed by a tensor product of those effects. Although the direction on a later copy depends on earlier outcomes, each complete-history effect remains such a product. The effect E for terminal decision plus is the sum over its histories and ready-coin branches, so

$$0 \preceq E \preceq I, \quad E^{\Gamma_A} = E \quad \text{for every subset } A \text{ of copies.}$$

Here Γ_A transposes the local tensor factors in A . The assertion requires classical feedback and destructively measured separate copies, as in the admitted model. It does not cover a coherent

returning controller, retained quantum memory that jointly interacts with later specimens, or an arbitrary collective measurement.

Condition on an entire sign-blind erasure pattern. Its probability is independent of the fixed source label and measurement directions. Granting that pattern in advance is a permissible lower-bound relaxation. With n nonnull copies, all resulting terminal effects still have the displayed properties. Null histories may guide the controller but do not reveal the source. No click postselection is performed in deployment: final bounds average over all patterns, including the all-null record.

Let $\Omega_{\pm}$ be the two comparison target mixtures on these n copies, and put $\Delta_n = \Omega_+ - \Omega_-$. In both applications below they are real mixtures of fixed-source tensor powers and are permutation invariant. A formal permutation average of E preserves positivity, partial-transpose invariance and $\text{Tr}(\Delta_n E)$; no claim that every such averaged effect is locally attainable is needed. Global sign reversal is conjugation by $Y^{\otimes n}$, under which Δ_n changes sign. Replacing E by

$$\frac{1}{2}\{E + I - Y^{\otimes n} E Y^{\otimes n}\}$$

again preserves the relevant constraints and objective. Thus the actual reader's correlation is bounded by a relaxation with the following exact parameterization.

Lemma 7.1 (Odd symmetric real Pauli representation). *Let B_{ab} be the sum of all distinct tensor strings with a copies of X , b copies of Z and $n - a - b$ copies of I . Every permutation- and local-transpose-invariant terminal effect, after the sign symmetrization above, has form*

$$E = \frac{1}{2}I + \sum_{a+b \text{ odd}} y_{ab} B_{ab}, \quad |y_{ab}| \leq \frac{1}{2}.$$

Only pairs $a, b \geq 0$, $a + b \leq n$ occur. Its positivity is equivalent to positivity of its ordinary qubit Schur blocks and their complements.

Proof. A Pauli string containing Y changes sign under transpose on that factor. Invariance under every local transpose therefore eliminates all such coefficients. Permutation invariance makes the remaining coefficient depend only on the X and Z counts. Sign symmetrization removes every even-degree nonidentity string and fixes the identity coefficient to $1/2$. Orthogonality of individual Pauli strings gives $y_{ab} = 2^{-n} \text{Tr}(EP)$ for any one string P of those counts. Because P is traceless with equally many $+1$ and -1 eigenvalues and $0 \preceq E \preceq I$, its absolute trace pairing is at most 2^{n-1} . This proves the coefficient bound. The Schur transformation is unitary; positivity is therefore equivalent to positivity in all irreducible blocks. $\square$

Permutation and partial-transpose invariant discrimination relaxations are established methods[3]. The present contribution is their explicit grouped-target/fixed-nuisance application and the rational certificate interface below. The representation is a calculation of an operator relaxation, not a quantum memory resource construction.

7.2 Exact block generation without an exponential tensor matrix

Use the generating identity

$$(I + xX + zZ)^{\otimes n} = \sum_{a,b} x^a z^b B_{ab}.$$

In the Schur sector $K \in \{0, \dots, \lfloor n/2 \rfloor\}$, put $s = n - 2K$. The block is

$$(1 - x^2 - z^2)^K \text{Sym}^s \begin{pmatrix} 1+z & x \\ x & 1-z \end{pmatrix}.$$

This follows from the simultaneous determinant/symmetric-power qubit representation; the normalized Dicke basis is essential. In row k , column ℓ , the polynomial before multiplication by the determinant factor is

$$\sqrt{\frac{\binom{s}{k}}{\binom{s}{\ell}}} \sum_q \binom{k}{q} \binom{s-k}{\ell-q} x^{k+\ell-2q} (1+z)^{s-k-\ell+q} (1-z)^q,$$

where $\max(0, k + \ell - s) \leq q \leq \min(k, \ell)$. Consequently every block basis entry B_{ab}^K is an integer times the single rational square root shown. The standard-library implementation expands the polynomial exactly with integer coefficients and encloses its normalization with directed rational square-root bounds. Blocks have size at most $n + 1$; the number of coefficient variables grows quadratically.

For a given comparison ensemble, write $g_{ab} = \text{Tr}(\Delta_n B_{ab})$. The relaxed correlation optimum has primal form

$$\sup_y \sum_{a+b \text{ odd}} g_{ab} y_{ab}, \quad -\frac{1}{2}I \preceq \sum_{ab} y_{ab} B_{ab}^K \preceq \frac{1}{2}I \quad \text{for all } K.$$

The objective equals $\text{Tr}(\Delta_n E)$ because $\text{Tr} \Delta_n = 0$. Equal-prior average error is $(1 - \text{Tr}(\Delta_n E))/2$. The coefficient restriction from Lemma 7.1 is valid for all actual effects. It can safely be used to bound a certificate residual even when an SDP proposal omitted it.

In fact the lower-bound comparison can grant arbitrary separate-copy POVMs with classical feedback on these real auxiliary states. Replacing every local effect F by $(F + F^\top)/2$ preserves its conditional probability on every real source state, preserves positivity and preserves the POVM sum. The resulting real local effects again produce transpose-invariant terminal effects. Thus the relaxation is not exploiting the limitation to two projective outcomes. This is an auxiliary comparison only: it introduces no new physical detector law, and coherent memory passed between specimens remains excluded.

7.3 Residual-corrected dual certificate

Theorem 7.2 (An exactly checkable correlation upper bound). *For every Schur block choose arbitrary positive semidefinite matrices A_K, C_K . Put*

$$r_{ab} = g_{ab} - \sum_K \text{Tr}[(A_K - C_K) B_{ab}^K].$$

Every admitted local terminal effect satisfies

$$\text{Tr}(\Delta_n E) \leq U_n, \quad U_n = \frac{1}{2} \sum_K \text{Tr}(A_K + C_K) + \frac{1}{2} \sum_{ab} |r_{ab}|.$$

No exact dual feasibility, solver status or eigenvector accuracy is assumed.

Proof. Write $F_K = \sum y_{ab} B_{ab}^K$. Pairing $F_K \preceq I/2$ with $A_K \succeq 0$ and $-F_K \preceq I/2$ with $C_K \succeq 0$ gives $\sum_K \text{Tr}[(A_K - C_K) F_K] \leq \frac{1}{2} \sum_K \text{Tr}(A_K + C_K)$. The difference from the required objective is $\sum r_{ab} y_{ab}$. The coefficient bound $|y_{ab}| \leq 1/2$ proves the result. $\square$

The shipped certificate encodes each dual matrix as $LL^\top$ for an explicit rational factor L . Its positivity is therefore exact. The dual cost is rational. The coefficient residuals use outward rational intervals for every square root and, in the robust application, for the exact objective coefficients. An upper bound on each absolute residual completes the proof. Floating SDP solutions only propose factors, whose entries are rounded to rationals; no numerical feasibility claim survives without this verification. Even poor or corrupted factors remain harmless: their residual is charged and may simply make the bound too weak to exclude a horizon.

7.4 Nominal one-percent interval: [13,14]

For the shared nominal parent, rotate the source axes to X and Z and use the uniform four-label prior. For odd $a + b$, exact Pauli orthogonality gives

$$g_{ab} = \begin{cases} \binom{n}{a} v^a, & b = 0, \\ \binom{n}{b} v^b, & a = 0, \\ 0, & a, b > 0. \end{cases}$$

Define the auxiliary shortened states $\rho_{a\pm}^{(v)} = (I \pm v\sigma_a)/2$ for $a = X, Z$. Measuring them sharply reproduces the nominal visibility-dependent classical kernel; this is a lower-bound representation and does not identify the actual retained preparation or detector channel. The positive mixture is $\Omega_+^{(n)} = \frac{1}{2}(\rho_{X+}^{(v)\otimes n} + \rho_{Z+}^{(v)\otimes n})$, and the hidden axis stays fixed. It is not a power of its one-copy average.

For a horizon B , the all-null error is $1/2$. All other erasure patterns are bounded by their n -click certificates $e_n \geq \max(0, (1 - U_n)/2)$. Hence

$$e_B \geq \sum_{n=0}^B \binom{B}{n} \eta^n (1 - \eta)^{B-n} \max(0, (1 - U_n)/2).$$

The maximum with zero is harmless and permits arbitrary weak proposals. The delivered rational Gram certificates for $n = 1, \dots, 12$ at $\eta = 4/5, v = 9/10$ prove the exact lower bound reproduced by `research11/bounds/ppt_certificate.py` in the reader supplement and recorded in its generated `research11/bounds/results.json`; its decimal display is

$$e_{12} \geq L_{12} > \frac{603}{50000} = 0.01206 > 0.01, \quad L_{12} \simeq 0.012064933110195356.$$

Here L_{12} denotes the full exact rational value in that JSON; its decimal is only a display. Thus twelve or fewer copies fail for both local protocol classes. The specified fourteen-copy bisector attainment remains valid, giving [13, 14] for each class. Neither exact threshold nor a one-percent adaptive/nonadaptive separation follows yet. The exploratory SDP value at thirteen does not itself exclude thirteen and is not an attaining physical protocol.

7.5 Original asymmetric robust interval: [16,17]

Retain the original uncertainty set and its fixed nuisance throughout each experiment. Use the convex-hull valid comparison prior with

$$u = (a \cos \Delta, -a \sin \Delta), \quad w = (-b \sin \Delta, b \cos \Delta), \quad p = \frac{b^2 + ab \sin 2\Delta}{a^2 + b^2 + 2ab \sin 2\Delta}.$$

It selects positive vector u with probability p and w otherwise, then samples a fair source sign. The nuisance offset is correlated with the sampled axis in this mathematical prior but is not redrawn on later copies. Every prior point is admitted by the robust supremum. Therefore its average error lower-bounds that supremum for any one fixed policy.

After a common real-plane rotation, write these vectors as

$$r_u = (ac, -as), \quad r_w = (bc, bs), \quad c = \frac{1-t}{\sqrt{2(1+t^2)}}, \quad s = \frac{1+t}{\sqrt{2(1+t^2)}}, \quad t = \tan \Delta.$$

Their inner product remains $-ab \sin 2\Delta$. The exact objective is

$$g_{ij} = 2 \binom{n}{i} \binom{n-i}{j} \{p(r_{u,x})^i (r_{u,z})^j + (1-p)(r_{w,x})^i (r_{w,z})^j\}, \quad i+j \text{ odd}.$$

For rational a, b, t , these coefficients are rational multiples of $\sqrt{2(1+t^2)}$. The same exact dual verifier therefore applies with interval objective coefficients; no new physical noise law is introduced.

At $a = 9/10, b = 4/5, t = 1/100, \eta_0 = 4/5$, the fifteen-copy certificate proves

$$e_{15}^{\text{robust}} \geq L_{15}^{\text{robust}} > \frac{129}{12500} = 0.01032 > 0.01, \quad L_{15}^{\text{robust}} \simeq 0.010321278963318586.$$

Again the full exact rational is reproduced by `research11/bounds/robust_ppt_certificate.py` in its generated `research11/bounds/robust_results.json`; the decimal is only a display. The original seventeen-copy robust frozen reader supplies the upper endpoint. Thus both global robust thresholds lie in [16, 17]. This does not determine whether sixteen is attainable, whether feedback helps there, or whether a stronger nonadaptive schedule already suffices. All source/controller, preparation, phase, erasure and archive conditions of the original robust task remain in force.

7.6 Verification and practical limits

The two delivered `research11/bounds/` entry points verify the full nominal and robust certificates, respectively. An additional independent rational tensor-basis check verifies the normalized symmetric-power identity on 420 small-sector entries; its source and output accompany the supplement. These finite identity checks support implementation inspection, while the displayed representation and residual-corrected dual proofs supply the universal argument. Solver messages such as `optimal_inaccurate` are retained as proposal metadata; they are not proofs. The efficient exact core uses Python's standard library. Optional proposal generation requires NumPy, CVXPY and an SDP solver.

The computational improvement is material: a polynomial-size local-effect relaxation gives reproducible all-angle bounds without enumerating angular grid policies or exponential outcome trees. It does not establish that this relaxation is tight or that an SDP optimizer can be implemented as a reader. Its physical role is a comparison certificate. No Born-selection, consciousness, trine-resource or unique-source conclusion follows.

8 Frozen-angle randomization and preparation convexification

8.1 A common least-favourable comparison for every horizon

For general positive gain floors $a, b \leq 1$, efficiency floor $\eta_0 > 0$ and $0 \leq \Delta < \pi/4$, let

$$u = (a \cos \Delta, -a \sin \Delta), \quad v = (-b \sin \Delta, b \cos \Delta), \quad c = u \cdot v = -ab \sin 2\Delta.$$

The admitted hard comparison fixes X at its worst positive offset and Y at its worst negative offset, with gain floors and efficiency η_0 . A mathematical comparison prior may correlate its fixed nuisance with the fixed source axis; it is a lower-bound device for the supremum over source/nuisance pairs, not a new physical calibration law. Put

$$p = \frac{b^2 - c}{a^2 + b^2 - 2c}, \quad z = pu + (1 - p)v, \quad m = \|z\| = \frac{ab \cos 2\Delta}{\sqrt{a^2 + b^2 + 2ab \sin 2\Delta}}.$$

Then $z = mn_{\theta_*}$, where

$$\tan \theta_* = \frac{a \cos \Delta + b \sin \Delta}{b \cos \Delta + a \sin \Delta}.$$

The direction θ_* lies in the common-positive interval $I = [\Delta, \pi/2 - \Delta]$.

Let $q_n(\beta)$ be the error of majority with a fair tie after n independent signs of positive bias β , including $q_0 = 1/2$. Define the full deadline risk, with all erasures counted, by

$$Q_B(\eta, \beta) = \sum_{n=0}^B \binom{B}{n} \eta^n (1-\eta)^{B-n} q_n(\beta).$$

For $k \geq 0$,

$$q'_{2k+1}(\beta) = -\frac{(2k+1)\binom{2k}{k}}{2^{2k+1}}(1-\beta^2)^k, \quad q_{2k} = q_{2k-1} \quad (k \geq 1).$$

Consequently $Q_B(\eta, \beta)$ is decreasing and convex on $\beta \in [0, 1]$. Its second derivative is nonnegative term by term; no numerical curvature test is used to establish the statement.

Theorem 8.1 (All-horizon random-frozen-angle optimum and spread penalty). *Before a block, choose one random nominal angle $\Theta \in I$, independently of the hidden source and nuisance. Retain the seed, repeat that angle for all B trials, and allow any randomized terminal decoder using the full retained record. The exact robust minimax error in this class is $Q_B(\eta_0, m)$, attained by the deterministic angle θ_* and fair-tie majority. Every such randomized reader has robust worst error at least*

$$\begin{aligned} & Q_B(\eta_0, m) + m[-\partial_\beta Q_B(\eta_0, m)]\mathbb{E}[1 - \cos(\Theta - \theta_*)] \\ & \geq Q_B(\eta_0, m) + \frac{2m[-\partial_\beta Q_B(\eta_0, m)]}{\pi^2}\mathbb{E}[(\Theta - \theta_*)^2]. \end{aligned}$$

This is not an optimum over varying-angle schedules, adaptive control, angles outside I , or collective quantum access.

Proof. Choose the fair sign and axis probabilities $p, 1-p$ of the fixed hard comparison above. Conditional on a seed angle θ , its two positive biases are $\beta_X = a \cos(\theta + \Delta)$ and $\beta_Y = b \sin(\theta - \Delta)$, both nonnegative. On any retained word, each axis component's positive/negative sign likelihood difference has the sign of the nonnull signed count. Therefore their positive mixture has the same sign. Majority with a fair tie is Bayes-optimal even for the axis mixture and even if every ordered sign/null record is accessible. Its conditional Bayes risk is

$$f(\theta) = pQ_B(\eta_0, \beta_X) + (1-p)Q_B(\eta_0, \beta_Y).$$

Any other decoder has at least this comparison-prior average error. Convexity in the bias gives

$$f(\theta) \geq Q_B(\eta_0, p\beta_X + (1-p)\beta_Y) = Q_B(\eta_0, m \cos(\theta - \theta_*)).$$

Because Q_B is decreasing and convex, its tangent at m lower-bounds the last expression by $Q_B(\eta_0, m) + m[-Q'_B(\eta_0, m)](1 - \cos(\theta - \theta_*))$. Average over the retained seed. The robust supremum is at least this fixed comparison average. Finally $1 - \cos x \geq 2x^2/\pi^2$ for $|x| \leq \pi$, by concavity of sine on $[0, \pi/2]$. The interval I ensures this range. The preceding frozen reader attains $Q_B(\eta_0, m)$ uniformly over the original uncertainty set, proving the equality as well as the quantitative penalties. $\square$

For the original example,

$$p = \frac{654464}{1478945}, \quad m^2 = \frac{129574081296}{369773223625}.$$

Exact quadratic-field arithmetic gives

Trials	Robust random-frozen optimum	One-percent status
15	.0124288300109...	fails
16	.0102368487113...	fails
17	.00844219845891...	passes

Hence seventeen is the exact minimum in the enlarged random-frozen class. A phase lottery is not a route to improving that threshold. Any improvement must use another interface, such as varying settings within a block, feedback, or additional information. This does not prove that a varying nonadaptive schedule needs seventeen, nor that adaptation is advantageous at one percent.

8.2 When convexification makes the frozen reader globally optimal

Theorem 8.2 (An all-horizon convex-source benchmark). *Let S be a nonempty compact subset of the equatorial Bloch disk for the fixed antipodal source family $\rho_{s,r} = (I + sr \cdot \sigma)/2$. Suppose its closest convex-hull point z is itself physically admissible: $z \in S$. With independent copies conditional on fixed (s, r) , independent sign-blind erasures $\eta \geq \eta_0$ and a fixed deadline, both full continuous local adaptive and randomized nonadaptive minimax errors equal $Q_B(\eta_0, \|z\|)$ for every B . One frozen direction $z/\|z\|$ attains it when $z \neq 0$; the $z = 0$ value is $1/2$.*

Proof. The supporting-plane inequality $r \cdot z \geq \|z\|^2$ for every $r \in S$ makes the frozen direction's sign bias at least $\|z\|$. Majority risk is decreasing in positive bias and efficiency, giving the upper bound. For the lower bound, the actual pair $\rho_{+,z}, \rho_{-,z}$ is admissible. Those states commute and are diagonal in the proposed measurement basis. Their independent tensor powers, including orthogonal erasure flags at η_0 , are a classical ternary likelihood problem. The likelihood ratio is monotone in the signed count, so even an arbitrary joint auxiliary-state measurement cannot beat fair-tie majority. This gives the same lower bound for every local protocol and its randomization. The zero vector has no sign information and gives $1/2$. $\square$

The assumption $z \in S$ is essential to this proof. In the original asymmetric reader, the positive-target set consists of two separated near-axis regions and does not contain the diagonal point z . Replacing it by its physical convex hull would admit a new independent-copy mixed source. That is a different, explicitly enlarged preparation class, not a permitted replacement of the original fixed hidden axis. It provides an exact benchmark and identifies the preparation distinction that a more advanced reader would have to exploit.

8.3 A retained-record witness of the preparation distinction

For the original hard support u, v , let $\rho_{s,z} = p\rho_{s,u} + (1-p)\rho_{s,v}$. The actual fixed-axis mixture and independent copies of that averaged one-copy state differ at two copies:

$$\begin{aligned}\Omega_s^{(2)} &= p\rho_{s,u}^{\otimes 2} + (1-p)\rho_{s,v}^{\otimes 2}, \\ \Omega_s^{(2)} - \rho_{s,z}^{\otimes 2} &= \frac{p(1-p)}{4}[(u-v) \cdot \sigma] \otimes [(u-v) \cdot \sigma].\end{aligned}$$

Put $V = p(1-p)\|u-v\|^2$. The trace distance is exactly $V/2$. Since $z \cdot (u-v) = 0$, measuring each copy along $n_\perp = (u-v)/\|u-v\|$ gives sign-product expectation V in the fixed-axis mixture and zero in the independent averaged state. Coding a retained null as zero changes the difference to $\eta^2 V$ under the same independent sign-blind erasure law. No click postselection is used. For the original parameters,

$$V = \frac{134898283296}{369773223625} \simeq .364813552408, \quad \eta_0^2 V \simeq .233480673541.$$

This is a standard covariance identity given an explicit operational readout, not a new physical law or a solution of the sign target itself. It makes the distinction between the two preparation laws experimentally meaningful under trusted preparation and the specified access contract. The comparison-prior stock is not asserted to be the actual unknown source in every deployment.

9 Compact, exactly certified adaptive readers at thirteen and fourteen copies

The four-copy trees establish a continuous-angle feedback separation, but their exhaustive retained-history representation was not a scalable synthesis method for the one-percent task. We develop a finite-horizon alpha-vector construction and a separate exact policy verifier. The completed result is a compact feasible fourteen-copy reader and a checked thirteen-copy candidate above the target; it does not settle that endpoint.

9.1 A feasible policy-graph interface

Use the original four fixed labels $h \in \{X+, X-, Y+, Y-\}$, independent copies conditional on that one label, and the shared calibrated $\eta = 4/5, v = 9/10$ detector. An action a has an exact rational unit direction $n_a = (c_a, s_a)$. In outcome order $(+, -, 0)$ its conditional kernels are

$$K_{a,h} = \left(\frac{\eta}{2}(1 + vx_{a,h}), \frac{\eta}{2}(1 - vx_{a,h}), 1 - \eta \right), \quad x_{a,h} \in (c_a, -c_a, s_a, -s_a).$$

A level- t node stores one action and three child indices at level $t - 1$. Level zero stores a sign decision. Thus the controller follows a finite outcome-labelled directed acyclic graph; it needs no access to the source label, no continuously updated posterior and no inaccessible calibration parameter. Every null follows its own declared child and consumes a trial.

For node i define its four-component error vector $\alpha_{t,i}$. The exact recursion is

$$\begin{aligned} \alpha_{0,i}(h) &= \mathbf{1}\{d_i \neq R(h)\}, \\ \alpha_{t,i}(h) &= \sum_{o \in \{+, -, 0\}} K_{a_i,h}(o) \alpha_{t-1,j_{i,o}}(h). \end{aligned}$$

For a hidden-label prior b , its expected error is $b \cdot \alpha$. The prior is used offline for design, not assumed as a distribution in the final worst-label certificate.

Proposition 9.1 (A rational feasible-reader certificate). *Every finite graph of the above type defines an admitted adaptive reader. Bottom-up rational evaluation gives its exact four source-label errors. For the full continuous equatorial class, averaging the graph over axis exchange and global sign reversal gives equal error on every label, equal to $\frac{1}{4} \sum_h \alpha_{B,\text{root}}(h)$. Hence that mean is an attainable minimax upper bound. No belief-point coverage or policy-search optimality assumption is required.*

Proof. Condition on a fixed source label. Independence of each fresh specimen and ready detector stock makes the three terms in the recursion the exhaustive conditional cases for the next outcome. Induction gives every node's true error probability, including nulls and the forced terminal decision. Axis exchange acts by swapping (c_a, s_a) at every command. Global sign reversal acts by reversing virtual signs and the final decision. Both operations preserve the admitted common-noise law and continuous action class; their four compositions act transitively on the labels. Two root fair bits, retained for the controller/decoder, therefore equalize the four errors to their mean. The root bits are independent of the unavailable source label. No quantum residual-channel equivalence is claimed. $\square$

The base menu has thirteen rational directions; swapped commands are also admitted, whether or not they coincide with a base-menu entry. This is an attaining construction within continuous access, not an equality between a finite grid optimum and the continuous optimum.

9.2 Point-based synthesis and exact separation from its numerics

The offline construction specializes point-based POMDP value iteration [8] to a static four-label terminal-cost problem. For a collection of previous feasible policy vectors and a design belief b , choose separately for each action/outcome a child minimizing $b \cdot (K_{a,o} \odot \alpha)$, sum its four conditional errors, then select an action of least design cost. The resulting object is an actual policy graph, even if floating-point comparisons select a suboptimal child. Duplicate action/child triples are merged, and unreachable nodes are removed when freezing a nominated root.

Beliefs were sampled both from Dirichlet distributions and actual random measurement trajectories, with the four source symmetries included. The pilot used 21461 distinct design beliefs and thirteen rational actions. Their half-angle parameters were rounded offline to rational numbers with denominator 10000 before forming exact unit-circle coordinates. Numerical candidate generation never supplies the final risk certificate. Nor do we import the literature’s discounted infinite-horizon coverage bound into this finite-horizon reader. The guarantee here is feasibility and exact evaluation.

The independent standard-library verifier avoids relying on large exact fractions for its primary pass. At scale $S = 2^{80}$ it propagates integer lower and upper error intervals. For common denominator D and nonnegative kernel numerators N_o summing to D , it uses

$$L = \left\lfloor \frac{\sum_o N_o L_o}{D} \right\rfloor, \quad U = \left\lceil \frac{\sum_o N_o U_o}{D} \right\rceil.$$

Every operation rounds outward. Induction gives interval width at most $2t/S$ at level t ; terminal values are exact. The root’s group mean is then averaged as a rational number. An additional all-Fraction evaluator reproduces the label costs, and ordered-word enumeration checks small subtrees independently. Malformed directions and child indices are rejected.

9.3 Completed operating points and remaining one-percent gap

Copies	Internal graph nodes	Certified attainable error	One percent
13	600	0.01102649347305383...	fails
14	697	0.008749337644325401...	passes

The displayed digits summarize exact rational intervals reproduced by `research11/nominal/verify_nominal.py` from the supplied `research11/nominal/policy_14.json` and `policy_13.json`. Paths are relative to the reader supplement; no floating comparison decides pass/fail. The fourteen-copy protocol improves the specified fixed-bisector error 0.00886634707138012... at the same copy budget. Its modest error improvement does not change the attained one-percent budget. The main useful addition is a compact, independently evaluable reader and reusable primal certificate interface for larger-horizon searches.

An additional frozen thirteen-copy child-pool candidate in `research12/nominal/policy_13_pool.json` has certified symmetrized error 0.011009824045220384..., evaluated by the supplied outward-integer recursion. It also fails one percent. This supporting refinement does not change the fourteen-copy attainment or the unresolved threshold.

Neither the unsuccessful thirteen-copy candidate nor an optimizer’s failure proves a universal thirteen-copy obstruction. A separate lower-bound theorem is needed to constrain all continuous

readers. We make no claim of global optimality over either the thirteen-action construction class or the full continuous class, and no adaptive/nonadaptive separation at one percent.

A separate numerical search examined every two-setting count allocation at budgets twelve through fourteen, optimizing its two angles. The tested solutions collapsed near the fixed bisector and did not improve attainment. This is exploratory evidence only; it does not certify that every two-setting schedule is worse, let alone every continuous nonadaptive protocol.

9.4 Explicit controller and retained-record resources

The fourteen-copy graph has 697 internal nodes and two terminal decisions. A global node pointer therefore needs ten bits; two independent root symmetry bits are additional ready/controller stock. A direct ROM upper account uses four bits per base action and three ten-bit child pointers per internal node, plus two terminal decision bits: at most 23700 bits, excluding phase-command constants, decoder/control arithmetic, physical clock, references and actuator state. This is an attaining upper account, not a controller minimum. There are at most 26 physical phase commands after the root axis transform. A conservative archive uses five command-ID bits and two outcome bits per trial, plus root-seed and timing metadata. All histories remain retained even though feedback uses only the node pointer.

Exact mathematical rational commands do not assert infinite physical phase precision. A deployment must separately bound implementation error, preserve preparation/detector stability, charge phase controls and references, and validate the actual frozen reader if the nominal model is untrusted. No hardware, energy, clock-speed or thermodynamic advantage follows from the ROM count or the error improvement alone. The separate calibration route can consume this finite policy as a declared object without identifying its entire residual quantum instrument.

10 Physical resources and remaining boundaries

The one-percent minima remain [13, 14] nominally and [16, 17] under the original fixed-unknown asymmetric parent. Failed thirteen- and sixteen-copy candidate searches are not lower bounds. A separate finite-policy polynomial certificate and continuous lower dynamic programme are retained in the technical supplement; their conservative costs do not settle these endpoints. At the selected four-copy tolerances, the proved copy separation is exact; it must not be reported as a one-percent separation. For repeated per-copy phase pulses a timing model can be $T = B(\tau + |\theta|/\Omega_p)$; a separately assumed once-set held actuator instead gives $T = B\tau + |\theta|/\Omega_s$. The two physical parents cannot exchange their costs. Adaptive switching and feedback waits require their own path-dependent model. Append-only archives are distinct from causal controller state; source stock, ready randomness, trusted calibration, references, precision and actual latency remain charged. No energy, thermodynamic or hardware optimum follows from copy count. Calibration controls deployment only with a justified stability contract; terminal effects do not identify a complete residual quantum instrument. The source-readout distinction established here is operational and relative to admitted interventions. Unrestricted operational surrogates, readable preparation records, collective access or changed stock can alter the task; these results do not select a unique ontology or constitute a consciousness criterion.

A Adaptive phase readers: exact scope, optimum and genuine feedback gain

All risk values in this appendix are restricted to the X/Y measurement menu. The local e_B^{ad} and e_B^{na} notation therefore refers to that restricted class, rather than the continuous-menu values in the

main text. A difference between these values is a change of admitted access. The fixed-schedule analysis optimized one fixed X/Y allocation and its terminal decoder. We enlarge the admitted class to nonanticipating sequential choices between those same two readers, allowing classical randomization. This changes two resources: a schedule may already be randomized before any observation, and feedback may select later settings using observed records. We separate these effects. The ideal gain below comes entirely from schedule randomization; the specified-noise gain survives comparison with every randomized nonadaptive schedule and therefore is a genuine feedback advantage.

A.1 Declared operational experiment and symmetry

The source is the fixed-schedule preparation family $\rho_{X\pm} = (I \pm X)/2$, $\rho_{Y\pm} = (I \pm Y)/2$ with the nominated binary sign target. A label h is held fixed and an admitted independent stock $\rho_h^{\otimes B}$ supplies one copy per tick. The inaccessible preparation controller retains h ; it is not supplied to the decoder. Every trial chooses $a \in \{X, Y\}$ on a fresh copy, with the same terminal effects

$$E_{\pm|a} = \frac{\eta}{2}(I \pm v\sigma_a), \quad E_{0|a} = (1 - \eta)I.$$

The phase rotation, beam splitter, calibrated detector, timing and unread baths retain the fixed-schedule physical contract. The experiment always outputs a sign at the fixed deadline B , includes all nulls, and neither postselects clicks nor claims a full retained-quantum-channel equivalence. A policy may use earlier accessible records and independent ready coins. No collective POVM, another basis, returning environment or readable preparation record is admitted. Stopping early and retaining an irrevocable decision can be represented by ignoring subsequent records, but no saved-copy or random-stopping resource optimum is asserted.

For equal axis efficiencies and visibilities, the experiment is invariant under the four-element group generated by exchanging X/Y and reversing all signs. Sign reversal also reverses the target and final decision. This group acts transitively on the four source labels. Consequently the uniform source prior is least favourable for the complete randomized policy class: every minimax risk is at least its optimal uniform-prior Bayes risk, while averaging a Bayes optimal policy over the group makes each of its four errors equal to that Bayes risk. The group transformation is selected before the experiment and translated into actual settings, record signs and terminal decisions; it is not knowledge of the hidden source. Two independent fair root-control bits suffice for this group choice; a terminal posterior tie may additionally consume a fresh fair decision bit. This standard symmetrization argument applies both to adaptive policies and to the smaller randomized-nonadaptive class.

A.2 An all-horizon ideal theorem

Theorem A.1 (Ideal adaptive minimax value and its source of improvement). *With $\eta = v = 1$, the optimal uniform-prior Bayes risk and randomized adaptive minimax risk are equal, and*

$$e_0^{\text{ad}} = \frac{1}{2}, \quad e_1^{\text{ad}} = \frac{1}{4}, \quad e_B^{\text{ad}} = 2^{-B} \quad (B \geq 2).$$

For every $B \geq 2$ a randomized nonadaptive schedule attains the same value: choose equally between the allocations $(1, B-1)$ and $(B-1, 1)$ and use uniform-prior Bayes terminal decisions, with symmetric tie handling. Therefore ideal feedback supplies no further gain over randomized schedules.

Proof. Use unnormalized nonnegative hypothesis weights; terminal Bayes loss is $\min(w_{X+} + w_{Y+}, w_{X-} + w_{Y-})$. Every arbitrary adaptive tree is covered by choosing either X or Y at its root, summing the optimum losses of its two outcome branches, and recursing. Randomizing an action cannot improve a minimum of these two conditional losses.

The nontrivial reachable states, up to axis and sign symmetries, have three forms. Opposite-target two-state weights a, b on different axes have remaining- t loss $\min(a, b)/2^t$: the next two choices give $\min(a, b/2)$ or $\min(a/2, b)$ for one step, and induction gives the same halving relation for every t . A pair of equal weights b, b on opposite signs of one axis has loss b at $t = 0$ and zero for $t \geq 1$, by measuring that axis. Finally consider $A(a, b) = (a, 0, b, b)$ with $a \geq 2b$. Its loss is b at $t = 0, 1$, and $b/2^{t-1}$ for $t \geq 2$. Indeed, measuring X gives branches $A(a, b/2)$ and the equal Y pair $(b/2, b/2)$; measuring Y gives one single-target branch and the opposite-target pair $(a/2, b)$. The preceding formulas prove the claimed values by induction. This exhausts the successors of these types, so it is a lower bound for every decision tree as well as an attainable recursion.

Start with weights $(1, 1, 1, 1)$. Either first reader gives two symmetric $A(1, 1/2)$ branches. The formulas just proved, divided by total prior mass four, give $1/4$ for $B = 1, 2$ and 2^{-B} for $B \geq 3$. The symmetry argument then supplies the minimax equality.

For a fixed allocation with positive counts, the only opposite-target ambiguous corners have likelihoods $a = 2^{-n_Y}$ and $b = 2^{-n_X}$. Their uniform-prior Bayes contribution is $\min(a, b)/2$. At $(1, B - 1)$ this is 2^{-B} . Axis/sign symmetrization equalizes the four risks and retains a nonadaptive schedule, proving attainment in the smaller class. $\square$

Thus a one-percent target needs seven ideal fresh copies in the enlarged randomized class, rather than eight in the fixed-schedule single-allocation class (and twelve for balanced fixed schedules). At $B = 8$, the enlarged optimum is $1/256$, rather than $1/130$. This gain is explicitly attributed to admitted schedule randomization, not feedback. The target is still not perfectly identifiable from any finite record.

There is also a horizon-consistent feedback realization useful for erasures: choose an initial axis fairly, read it until its first opposing sign, then read the other axis on the next click. That other-axis result determines the sign. If all available clicks agree, return their common sign; if the first opposing sign is the last click, use a fair guess; with no clicks also guess fairly. Conditional on $N \geq 2$ clicks, an initially aligned reader makes no error. An initially cross-axis reader errs only on an all-equal sequence or a first flip at the last click, each of probability 2^{1-N} and error $1/2$. Averaging the initial axis gives 2^{-N} . For $N = 1$ it gives $1/4$ and for $N = 0$, $1/2$. The erasure extension is developed separately, preserving its additional independence and access premises.

A.3 Exact noisy Bellman certificates

For rational calibrated η, v , put every outcome likelihood over a common positive denominator D . Let $k_{a,o,h}$ be its nonnegative integer numerator. For unnormalized integer weights w , define

$$V_0(w) = \min(w_{X+} + w_{Y+}, w_{X-} + w_{Y-}), \quad V_t(w) = \min_{a=X,Y} \frac{1}{D} \sum_{o \in \{+, -, 0\}} V_{t-1}(w \odot k_{a,o}).$$

Homogeneity permits dividing w by its integer gcd and storing one rational value per resulting state. The minimum covers every adaptive action, not only the actions chosen by the exported policy. The uniform Bayes lower bound is $V_B(1, 1, 1, 1)/4$. A policy selecting a minimizing action and a terminal Bayes sign attains it; group averaging supplies the matching per-label minimax upper bound. This proves the recursion for the full declared policy class, not only a numerical ansatz or one nominated prior policy.

The exported JSON certificate stores all recursively required normalized states, both action values, the selected action and terminal losses. A checker verifies every Bellman equality bottom-up using exact fractions, then checks the root value. The producer separately evaluates the policy's four individual errors. There is no floating convergence criterion. Nominally impossible histories have a declared default

X action and fair terminal guess, so the policy remains defined before an actual-device perturbation is applied.

For $\eta = 4/5, v = 9/10$, the common denominator is $D = 25$ and the X-reader numerators for $+, -, 0$ are respectively $(19, 1, 10, 10)$, $(1, 19, 10, 10)$ and $(5, 5, 5, 5)$; exchange axes for Y. The following exact comparisons use identical copy budgets, reader effects, null retention and independent stock. The nonadaptive comparator includes all fixed counts and arbitrary root schedule randomization, not just the fixed-schedule single-schedule minimax design.

B	Adaptive minimax	Randomized nonadaptive minimax	Best counts
4	25709/156250	27167/156250	$(1, 3), (3, 1)$
8	4098018401/61035156250	38954017/488281250	$(3, 5), (5, 3)$

At eight copies these are approximately 0.06714193348 and 0.07977782682. Their strict difference is a genuine feedback gain. The fixed single-schedule value 0.08270910464 is a third, narrower comparison and is not used to exaggerate that gain. The specified-noise model also supplies a copy threshold. Its adaptive errors at budgets eighteen and nineteen are respectively

$$\frac{64382397399725181469211}{5820766091346740722656250} > \frac{1}{100}, \quad \frac{136397472235992257022563}{14551915228366851806640625} < \frac{1}{100}.$$

The best randomized-nonadaptive errors at budgets twenty and twenty-one are respectively

$$\frac{107414411331608425103}{9313225746154785156250} > \frac{1}{100}, \quad \frac{230516488566312459173}{23283064365386962890625} < \frac{1}{100}.$$

The latter minimizing counts are $(9, 11)$ at twenty and $(9, 12)$ at twenty-one, with their axis swaps. Extra observations can always be ignored, so risks are nonincreasing in budget. These adjacent certified comparisons prove that the one-percent threshold is exactly nineteen adaptive copies versus twenty-one randomized-nonadaptive copies in this calibrated noise model. They do not inherit the separate ideal-erasure robustness contract.

The source-copy savings are not automatically total apparatus savings: the adaptive reader consumes additional causal control capability.

To justify the nonadaptive optimum, group symmetry again reduces minimax to uniform-prior Bayes loss. A randomized mixture of fixed schedules has Bayes loss at least the smallest Bayes loss of one component. For each count pair, ordered-record exchangeability permits the exact histogram likelihood calculation. Minimizing its uniform-prior loss over all $n_X + n_Y = B$ counts, then symmetrizing its minimizing schedule, supplies matching lower and upper bounds for the whole randomized-nonadaptive class.

B One reader across unknown erasure and calibration parameters

For already calibrated readers, a generic perturbation allowance is possible. Here an explicit feedback policy yields an exact minimax law under unknown erasure efficiency, followed by a stronger model-specific calibration certificate when visibility is imperfect. The two results have different strength: erasure-only minimax optimality is proved; the visibility-box result certifies one frozen policy without claiming robust-minimax optimality. Neither depends on an ETH theorem.

B.1 A policy independent of the nuisance parameter

Choose an initial axis $A \in \{X, Y\}$ with a fair independent coin. Query that axis until its nonnull signs include both $+$ and $-$. Then query the other axis. The first nonnull result on that other axis commits its sign; further records are retained but ignored by the terminal classifier. If the deadline arrives

before a switch, return the common sign; if there was no click, or a switch without a subsequent click, return an independent fair guess. Thus a flip at the last successful trial is explicitly an inconclusive internal state, resolved by a declared decision coin. There is always a terminal decision; nulls are not postselected away.

Let e_n be the ideal adaptive minimax error for the X/Y-restricted four-label sign task established in the preceding appendix, not the continuous-menu optimum:

$$e_0 = \frac{1}{2}, \quad e_1 = \frac{1}{4}, \quad e_n = 2^{-n} \quad (n \geq 2).$$

The same value is the optimal uniform-prior Bayes error. The policy just defined attains it simultaneously at every click horizon. To see attainment, condition on the initial axis. If it is the true axis there is no error. If it is the other axis, the observed signs are fair. For $n \geq 2$, errors are possible only when all n signs agree, or the first opposing sign appears at position n . Each event has probability 2^{1-n} , with conditional error $1/2$. The fair initial axis gives error 2^{-n} for every source label. The $n = 0, 1$ cases give the displayed exceptional values.

Theorem B.1 (Exact unknown-erasure minimax frontier). *Fix a deadline of B trials. Visibility is one. Trial t is independently erased with probability $1 - \eta_t$, regardless of source label, axis and ideal result, where the deterministic efficiencies satisfy $0 \leq \eta_t \leq 1$. For known efficiencies, the minimum worst-case sign error among all adaptive X/Y policies and terminal decoders is*

$$E_B(\eta_1, \dots, \eta_B) = \mathbb{E}[e_N], \quad N = \sum_{t=1}^B Z_t, \quad Z_t \sim \text{Bernoulli}(\eta_t) \text{ independently.}$$

One policy, specified above without any knowledge of the efficiencies, attains this value. In particular, for a common efficiency and $B \geq 1$,

$$E_B(\eta) = (1 - \eta/2)^B - \frac{(1 - \eta)^B}{2} - \frac{B\eta(1 - \eta)^{B-1}}{4}. \quad (3)$$

For the robust class of all deterministic vectors $\eta_t \in [\eta_0, 1]$, unknown when the policy is fixed, the exact robust minimax error is $E_B(\eta_0)$. This includes a fixed unknown common efficiency.

Proof. For a lower bound use the uniform prior on the four labels and give the controller the entire independent erasure schedule in advance. Conditional on that schedule, only N informative trials are available; the ideal uniform-prior Bayes lower bound is e_N . Revealing independent schedule information cannot increase the optimum risk, so every original policy's worst-case risk is at least $\mathbb{E}[e_N]$.

The proposed policy depends on the successful signs and its initial coin, not on the erased-trial locations or the number of future successful trials. Conditional on the independent schedule it attains e_N on every label. This proves the matching upper bound. For common η , the binomial average of 2^{-N} is $(1 - \eta/2)^B$; correcting its $N = 0$ and $N = 1$ terms yields (3).

The sequence e_N is nonincreasing. Couple success indicators through independent uniforms as $Z_t = \mathbf{1}\{U_t \leq \eta_t\}$ and $Z_t^0 = \mathbf{1}\{U_t \leq \eta_0\}$. Then $N \geq N_0$ pointwise, so the same policy has risk at most $E_B(\eta_0)$. The admitted vector with every efficiency equal to η_0 supplies the matching lower bound. $\square$

The independence and sign-blindness conditions are consequential. A nuisance depending on hidden source data, a measured sign before erasure, or the past record is not covered by this theorem merely because each unconditional click rate exceeds η_0 . Correlated detector baths and preparation copies are also outside its parent. The theorem does permit unknown nonstationary deterministic efficiencies under independent per-trial erasure stock.

For seven trials and every $\eta_t \geq 49/50$,

$$E_7(49/50) = \frac{897410666811}{10^{14}} = 0.00897410666811 < 0.01.$$

This is a completed one-percent operating region requiring no estimate of the actual efficiency beyond the declared lower bound. Six ideal trials have optimal error $1/64 > 0.01$, so seven is also the exact robust minimum trial budget in this class. The comparison with the fixed-schedule analysis is explicit: its fixed deterministic schedule required eight ideal trials; randomization of the setting and the declared feedback controller expand that access class. The present result does not attribute the whole gain to feedback alone.

B.2 A fixed-policy certificate for imperfect visibility

Keep the same policy unchanged when the actual, fixed unknown pair is (η, v) in a rational rectangle. Each trial has the specified effects $E_{\pm|a} = \eta(I \pm v\sigma_a)/2$, $E_0 = (1 - \eta)I$ and fresh independent noise stock conditional on this pair. A history's probability is a product of terms $1 - \eta$, $\eta/2$, or $\eta(1 \pm v)/2$. Consequently every source-specific error is a polynomial in η, v of separate degree at most B . Axis randomization and sign covariance make all four risks equal to their mean polynomial $p_B(\eta, v)$. This avoids optimizing a different policy after learning an inaccessible nuisance parameter.

Proposition B.2 (Exact calibration-box certificate). *Affinely map a rational rectangle $[a, b] \times [c, d]$ to $[0, 1]^2$ and write the fixed policy's risk as $\sum_{ij} A_{ij} t^i s^j$. At degrees m, n no smaller than its respective degrees, set*

$$\beta_{kl} = \sum_{i \leq k, j \leq l} A_{ij} \frac{\binom{k}{i}}{\binom{m}{i}} \frac{\binom{l}{j}}{\binom{n}{j}}.$$

Then $\max_{kl} \beta_{kl}$ is a rigorous uniform error upper bound on that rectangle. Any evaluated point supplies a lower bound on the fixed policy's worst error. Apply the same construction on a finite subdivision and take the maximum of its upper bounds; if an upper bound equals an attained point value, the fixed-policy worst error is certified exactly. Every coefficient and comparison is rational.

Proof. The identities $t^i = \sum_{k=i}^m \binom{k}{i} \binom{m}{i}^{-1} \binom{m}{k} t^k (1 - t)^{m-k}$ and their s counterparts give the stated coefficients. The product Bernstein basis is nonnegative and sums to one on the unit square, so the polynomial is a convex combination of its coefficients. Subdivision covers the original rectangle. Attainment supplies the reverse inequality whenever the two bounds coincide. $\square$

This uses an established polynomial-enclosure technique, not a new theory of Bernstein polynomials[1]. Its addition here is a checked full-policy risk consumer, with the unknown parameters held fixed throughout the experiment. The implementation uses exact finite-state controller recurrences rather than accepting a grid scan as a proof. Its stored tuples may retain redundant history beyond a minimal logical-state quotient.

The delivered certificates have zero enclosure gaps on the following boxes:

B	admitted η	admitted v	worst fixed-policy error
7	[.98, 1]	{1}	0.00897410666811
7	[.99, 1]	[.999, 1]	0.009482646058652287...
8	[.98, 1]	[.999, 1]	0.005800538606142910...
8	[.79, .81]	[.89, .91]	0.11822395802359147...

All entries are exact rational calculations, with full fractions in `research8/robust/results.json`. The worst point in each displayed box is its lower-left corner, established by a matching Bernstein

upper bound rather than assumed from monotonicity. The seven-copy visibility box proves a new robust one-percent guarantee. It is not a claim that this particular policy is minimax for every visibility box: at lower visibility an adaptive optimizer can choose a substantially different policy within the same X/Y menu.

For the same seven-copy policy on $[\cdot99, 1] \times [\cdot999, 1]$, applying a generic conditional total-variation allowance to its ideal error gives approximately 0.0790373. The new model-specific upper bound is about 0.00948265. This improvement is in a certified error allowance under a stronger specified noise parent, not a measured device gain. Complete retained quantum outputs or arbitrary purifiers are not certified by these classical terminal-record polynomials.

C Reproducibility and assistance disclosure

The clean supplement supplies exact policies, continuous cube/Gram certificates and reviewed standard-library verifiers. The accompanying methods supplement also includes the finite-menu and physical-resource checks, with exact arithmetic distinguished from floating diagnostic comparisons. Exploratory generators, historical outputs and audit receipts remain in the separate research-preservation package; they are not required to verify the supplied mathematical certificates. Numerical solvers propose objects only; final interval/rational verification establishes the reported bounds. AI assistance was used for review and computational checking.

References

- [1] Mohamed Amin Ben Sassi and Sriram Sankaranarayanan. Bernstein inequalities for constrained polynomial optimization problems, 2015. URL <https://arxiv.org/abs/1509.01156>.
- [2] David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2): 265–272, 1953. doi: 10.1214/aoms/1177729032. URL <https://projecteuclid.org/journals/annals-of-mathematical-statistics/volume-24/issue-2/Equivalent-Comparisons-of-Experiments/10.1214/aoms/1177729032.full>.
- [3] J. Calsamiglia, J. I. de Vicente, R. Muñoz-Tapia, and E. Bagan. Local discrimination of mixed states. *Physical Review Letters*, 105:080504, 2010. doi: 10.1103/PhysRevLett.105.080504. URL <https://arxiv.org/html/1004.5522v1>.
- [4] Sarah Croke and Stephen M. Barnett. Difficulty of distinguishing product states locally. *Physical Review A*, 95:012337, 2017. doi: 10.1103/PhysRevA.95.012337. URL <https://journals.aps.org/prapdf/10.1103/PhysRevA.95.012337>.
- [5] Giacomo Mauro D’Ariano, Massimiliano Federico Sacchi, and Jonas Kahn. Minimax quantum state discrimination. *Physical Review A*, 72:032310, 2005. doi: 10.1103/PhysRevA.72.032310. URL <https://arxiv.org/html/quant-ph/0504048v5>.
- [6] Noam Elron and Yonina C. Eldar. Quantum detection with uncertain states. *Physical Review A*, 72:032338, 2005. doi: 10.1103/PhysRevA.72.032338. URL <https://journals.aps.org/prapdf/10.1103/PhysRevA.72.032338>.
- [7] B. L. Higgins, A. C. Doherty, S. D. Bartlett, G. J. Pryde, and H. M. Wiseman. Multiple-copy state discrimination: Thinking globally, acting locally. *Physical Review A*, 83:052314, 2011. doi: 10.1103/PhysRevA.83.052314. URL <https://arxiv.org/html/1012.3525v2>.

- [8] Joelle Pineau, Geoff Gordon, and Sebastian Thrun. Point-based value iteration: An anytime algorithm for pomdps. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 1025–1030, 2003. URL <https://robots.stanford.edu/papers/Pineau03a.pdf>.
- [9] Jeremy Rodgers. Source-readout non-equivalence: Descent and equivariant reconstruction obstructions, 2026. URL <https://doi.org/10.5281/zenodo.21369967>.