

Technical Supplement: Continuous Reader Optimization and Deployment Interfaces

Jeremy Rodgers
Independent Researcher
www.everythingequation.com

27 October 2026

Abstract

We give supporting methods and restricted benchmarks for common-sign readout of a fixed four-state qubit source: deterministic measurement allocation, an expanded rational menu, continuous-risk lower backups, fixed-nuisance polynomial certificates and physical timing/calibration interfaces. Every result retains its own measurement and stability contract. These methods support the associated reader frontiers without settling the remaining adjacent copy thresholds or a universal total-resource optimum.

This supplement develops optimization and deployment interfaces within Shadow Theory’s pre-existing source-readout programme [5]. It supplies independently checkable methods for the fixed-label four-state source. Copies are independent conditional on one label throughout a block; the preparation controller is inaccessible. The shared detector law is $E_{\pm|\theta} = \eta(I \pm v[\cos \theta X + \sin \theta Y])/2$, $E_{0|\theta} = (1 - \eta)I$. The asymmetric parent separately applies one fixed Pauli channel with multipliers $(g_X, g_Y, g_X g_Y)$ before a reader with fixed phase offset δ and independent sign-blind erasure. These response laws are premises, not inferred from equal one-copy marginals. The current one-percent copy intervals are [13, 14] nominally and [16, 17] for the specified asymmetric robust task, separately for both local policy classes. Supporting lower comparisons and unsuccessful candidates do not alter those status statements.

Contents

1	A physically admissible weighted completion-design frontier	2
1.1	Source, allowed interventions and retained outputs	2
1.2	An exact ideal fixed-schedule optimum	3
1.3	Exact rational noisy design with matching certificates	4
1.4	Controlled deviations and fair physical accounting	5
2	A larger rational measurement menu: the copy gap closes but feedback still helps	6
2.1	Declared settings, noise and symmetry	6
2.2	Exact adaptive and randomized-nonadaptive certificates	7
2.3	Implementation and limits	8
3	Nominal threshold: a checked obstruction and a continuous lower-backup method	9
3.1	The thirteen-copy operator comparison	9
3.2	A rational lower dynamic programme for the full circle	9
3.3	Exact demonstrations and scalability ceiling	11

4	Original robust task: fixed-nuisance certification and an unresolved sixteen-copy endpoint	11
4.1	A three-variable exact interface for a frozen reader	11
4.2	Exact Bernstein cover certificates	12
4.3	Why gain-endpoint checks alone are invalid	13
4.4	Searches and exact updated status	13
5	Uniform angular approximation	14
5.1	A certified angular-net transfer, not a grid optimum	14
5.2	Verification and exact status	14
6	From the continuous feedback gap to a physical resource certificate	15
6.1	A nonzero noise and control-error operating region	15
6.2	Actual phase travel, rather than a generic switch charge	16
6.3	A full-menu copy advantage with a conditional time advantage	18
6.4	Finite calibration covering all continuous angles	18
6.5	Reproducible checks and remaining limits	20

1 A physically admissible weighted completion-design frontier

The deterministic collision test identifies whether a target is exactly fixed by a record. A physical quantum reader also needs a quantitative error analysis when different targets can produce the same finite record. Here rational record likelihoods and exact optimization close that gap for one finite, explicitly specified phase-source experiment. This is a new programme-specific construction, independent of the ETH theorem; it demonstrates a substantive consumer of the weighted-completion interface. It is not a faster implementation of the ETH algorithm.

1.1 Source, allowed interventions and retained outputs

Use a two-mode one-excitation register with computational basis $|0\rangle, |1\rangle$ and equal populations. The finite preparation promise is

$$\rho_{X\pm} = (I \pm X)/2, \quad \rho_{Y\pm} = (I \pm Y)/2.$$

Equivalently their phases are $0, \pi, \pi/2, 3\pi/2$ in $(|0\rangle + e^{i\phi}|1\rangle)/\sqrt{2}$. All four preparations have the same separate-mode intensity readout $(1/2, 1/2)$. The target is only the sign class: $R(X+) = R(Y+) = +$ and $R(X-) = R(Y-) = -$. This target is an externally nominated function on the finite preparation set, not the eigenvalue of a quantum observable assigned to all states.

A source label h is chosen once and held fixed for an experiment. Conditional on h , the source supplies independent copies $\rho_h^{\otimes B}$, independent of the reader ready stock. This is a preparation promise, not a conclusion from matching one-copy marginals or from resetting a device. The source controller retaining h is declared preparation stock and is unavailable to the decoder; no generator-memory separation is inferred from this experiment.

The permitted experiment menu consists of a fixed, nonadaptive schedule: n_X copies measured with the X reader and n_Y with the Y reader. The two incompatible observables act on different copies. Neither is silently assumed jointly measurable on the same source. The Z intensity reader alone yields exactly uniform outcome sequences for all h and cannot improve worst-case sign error below $1/2$.

A phase-controlled beam splitter implements the ideal reader: rotate by $U_a = \exp(-i\theta_a Z/2)$ with $\theta_X = 0, \theta_Y = -\pi/2$ under a fixed phase convention, then apply a Hadamard beam splitter and detect the two terminal paths. The convention can equivalently be fixed by the effects below. A

controlled Hamiltonian $H_z = \hbar\Omega_z Z/2$ realizes the phase rotation in duration $|\theta_a|/|\Omega_z|$, with the sign of the drive declared. The beam splitter and detector have separately allotted durations; there is no instantaneous-pulse or speed optimum claim.

The calibrated noisy reader has terminal effects

$$E_{\pm}^{(a)} = \frac{\eta}{2}(I \pm v\sigma_a), \quad E_0^{(a)} = (1 - \eta)I, \quad 0 \leq \eta, v \leq 1, \quad (1)$$

where $\sigma_X = X$, $\sigma_Y = Y$. A realization measures the ideal Pauli observable, flips its reported sign with independent probability $(1 - v)/2$, and independently replaces it with a null with probability $1 - \eta$. Fresh independent coin/environment stock per copy is admitted and counted as apparatus resources. The ordered record retains every setting, terminal sign, null and time. Environments may remain retained but unread; no assertion equates their complete quantum channels.

This specifies effects and a conventional realization rather than importing a complete apparatus-response certificate from terminal counts. It is an ideal/specified-noise reader construction, not a solution of arbitrary detector mismatch. A separate phase-to-record model can supply perturbation budgets only when its preparation, access and output assumptions match. No delayed-trine resource bound is inferred from this copy-allocation task.

1.2 An exact ideal fixed-schedule optimum

Let $e^*(n_X, n_Y)$ be the minimum over all randomized terminal decoders of the maximum sign error over the four possible source labels. It includes the entire original population, with no click postselection. In this subsection $\eta = v = 1$.

Theorem 1.1 (Target-relative phase-reader allocation). *For $n_X, n_Y \geq 1$,*

$$e^*(n_X, n_Y) = \frac{1}{2^{n_X} + 2^{n_Y}}. \quad (2)$$

If either count is zero, $e^* = 1/2$. At total copy budget $B \geq 2$, the optimal permitted allocation is $(1, B - 1)$ or its swap, giving $e_B^* = 1/(2 + 2^{B-1})$. The minimum worst-case error for identifying the full four-valued label under the same fixed schedule is $2/(2^{n_X} + 2^{n_Y})$ when both counts are positive.

Proof. An $X+$ source gives only plus signs at the X reader and fair independent signs at the Y reader; exchange signs or axes for the other sources. Apart from the two opposing corners

$$A = (X \text{ all } +, Y \text{ all } -), \quad B = (X \text{ all } -, Y \text{ all } +),$$

the record determines the target sign. Corner A has probability $a = 2^{-n_Y}$ under $X+$ and $b = 2^{-n_X}$ under $Y-$. If a decoder chooses $+$ there with probability q , its worst error is at least $\max\{a(1 - q), bq\}$, whose minimum is $ab/(a + b)$. Choosing $q = a/(a + b)$ at A , $1 - q$ at B , and the uniquely determined sign elsewhere attains this bound on all four sources. If one axis is absent, the two opposite signs of that axis induce identical records, which forces worst error at least $1/2$; a fair guess attains it.

For fixed B and positive counts, maximize $2^{n_X} + 2^{B-n_X}$. This convex function has its maximum on the integer interval $1 \leq n_X \leq B - 1$ at an endpoint. For full-label discrimination, all four constant-sign corners have two possible labels, one from each axis. Assign the X label with probability $a/(a + b)$ at each corner. Each source then incurs error $2ab/(a + b)$. For a matching lower bound, use prior weights $\alpha/2$ on each X label and $(1 - \alpha)/2$ on each Y label with $\alpha = b/(a + b)$. At each corner the two weighted likelihoods are equal, contributing $ab/[2(a + b)]$ to every decoder's Bayes error. Summing four corners gives the claimed lower bound. $\square$

For eight source copies, the target-optimal schedule (1, 7) has error $1/130 \simeq 0.0076923$, while the balanced (4, 4) schedule has error $1/32 = 0.03125$. The comparison charges the same eight copies, one readout per copy, both calibrated axes and the same retained archive. It is not comparison against an optimal arbitrary POVM or adaptive protocol. For a one-percent sign-error target the optimal menu schedule needs eight copies, while balanced or nearly balanced schedules need twelve. Full-label identification optimized over the same menu needs nine. Thus target choice and experiment allocation both change the accessible-resource requirement. The eight-copy one-percent guarantee does not require an exact biased decision coin: with (1, 7), return the common Y sign if all seven Y results agree, and otherwise return the single X result. Its worst error is $1/128 < 1/100$. This deterministic decoder gives a nearby constructive repair when ideal rational randomization is unavailable; the sharper $1/130$ remains the randomized-class optimum. Extend the rule to every record by returning $+$ whenever a null occurs. Those records have zero nominal probability but must have a declared decision before applying an actual-device perturbation bound.

Every finite positive-count experiment still has a record compatible with opposite target values. On the admissible relation of source labels and possible outcome records, an exact support-only completion test correctly rejects zero-error sufficiency at every finite budget. A small weighted error is a different target and is not called an exact repair of that collision. This is why the likelihood/weight extension adds something the Boolean collision oracle alone cannot supply.

1.3 Exact rational noisy design with matching certificates

From (1), let

$$p_h^{(a)}(\pm) = \eta(1 \pm vr_a(h))/2, \quad p_h^{(a)}(0) = 1 - \eta,$$

where $r_X(X\pm) = \pm 1$, $r_Y(Y\pm) = \pm 1$ and the cross-axis values are zero. For rational η, v these are rational probabilities. A setting's histogram $k = (k_+, k_-, k_0)$ has probability

$$\frac{n!}{k_+!k_-!k_0!} p(+)^{k_+} p(-)^{k_-} p(0)^{k_0}.$$

The product of the two setting histograms is sufficient for this likelihood task: under the declared conditional product law, all permutations with the same histogram have identical likelihood ratios. The complete ordered physical archive is retained; the decoder's sufficient statistic does not erase or return it to the source. For fixed setting counts, there are only $\binom{n_X+2}{2} \binom{n_Y+2}{2}$ histograms. Their rational probabilities can already be computed in polynomial time here; the ETH algorithm is not needed or claimed to speed this simple special case.

Write the four likelihoods at record o as $P_{X+}(o), P_{X-}(o), P_{Y+}(o), P_{Y-}(o)$. For any $\alpha \in [0, 1]$, every decoder has worst-case error at least

$$L(\alpha) = \frac{1}{2} \sum_o \min\{\alpha P_{X+}(o) + (1 - \alpha) P_{Y+}(o), \alpha P_{X-}(o) + (1 - \alpha) P_{Y-}(o)\}. \quad (3)$$

This is the Bayes risk for an explicit prior, not an assumption that the source actually uses that prior. The function is concave piecewise linear. Its maximum occurs at an endpoint or at a rational crossing of its two affine record weights.

Exact construction and certificate. Enumerate those breakpoints and select a maximizing α_* . On non-tied records use the weighted likelihood-ratio decision. On tied records construct two policies minimizing and maximizing the difference between the X -axis and Y -axis average errors. Symmetrize under global sign reversal so the errors of the two signs within each axis agree. Every such policy attains $L(\alpha_*)$ at its nominated prior. If $0 < \alpha_* < 1$, optimality of the concave maximum makes

the two attainable error-difference endpoints bracket zero; mix them to equalize the axis errors. At $\alpha_* = 0$, choose the endpoint whose X error does not exceed its Y error; at $\alpha_* = 1$ choose the converse. One-sided maximality supplies the respective endpoint condition. Thus the constructed policy’s worst error equals the lower bound $L(\alpha_*)$. This proves exact minimax optimality for each schedule.

The implementation retains a simpler independently checkable certificate: all rational likelihoods, the prior α_* , and a rational decision probability $q(o)$ for each histogram. Recomputing each of the four errors gives an upper bound. Evaluating (3) gives a lower bound. Equality is an exact zero primal/dual gap; it requires no floating optimizer or assumption of numerical convergence. For any histogram absent from the union of nominal supports, the exported decoder returns +; this completes it on actual perturbed records.

With calibrated $\eta = 4/5$ and $v = 9/10$, enumerate every fixed schedule of at most B copies at unit cost:

Budget B	Certified optimal (n_X, n_Y)	Exact minimax error
2	(1, 1)	71/250
4	(2, 2)	229/1250
6	(3, 3)	94003/781250
8	(4, 4)	8077061/97656250
10	(4, 6) or (6, 4)	3252843423709/56955566406250

This table minimizes over a single deterministic X/Y count allocation and its minimax decoder. It does not optimize mixtures over different allocations; the randomized-schedule comparator is a broader class. In particular, at budget eight this fixed-allocation optimum error is 0.08270910464, while the formerly ideal-optimal (1, 7) schedule has error 726821937001/4803417968750 $\simeq$ 0.1513134900. Noise changes the optimal design. Balanced allocation is not asserted optimal at all budgets: the exact ten-copy result is asymmetric. The data are mathematical model outputs, not experimental measurements.

1.4 Controlled deviations and fair physical accounting

Suppose deterministic bounds $\delta_t \in [0, 1]$ control the total variation between actual and nominated conditional outcome laws at every matched past history. Sequential maximal coupling keeps the records matched at step t , conditional on matching so far, with probability at least $1 - \delta_t$. Consequently the complete classical-record distance is at most

$$\kappa_B = 1 - \prod_{t=1}^B (1 - \delta_t) \leq \sum_t \delta_t.$$

Independence is not needed under this stronger uniform conditional contract; one-copy marginal bounds do not imply it. The product bound is sharp as a record-distance statement: compare a nominal all-zero word with independent per-position flips of probabilities δ_t . Every fixed classifier’s error increases by at most κ_B . This is not a claim of robust minimax optimality and covers classical outputs, not residual quantum baths.

For example, the eight-copy ideal (1, 7) decoder retains a one-percent worst-case error guarantee whenever the actual conditional laws obey $\delta_t \leq 1/4000$ at every history: the exact rational bound $1/130 + 1 - (3999/4000)^8$ is strictly below $1/100$. This makes the finite design’s calibration allowance quantitative without inferring future stability from past calibration alone. The deterministic decoder also passes: replace $1/130$ by $1/128$ in that bound, obtaining approximately $0.0098107509 < 0.01$ without a randomized-decision port.

For the specific reader family, changing (η_0, v_0) to (η, v) changes each one-copy outcome law by at most

$$|\eta - \eta_0| + \eta_0|v - v_0|/2.$$

Trace-distance preparation error δ_S adds at most δ_S . These are explicit consumers for a certified calibration/drift budget. An unconditional one-copy marginal estimate alone does not justify the sequential conditional promise or independent fresh apparatus stock.

The resource optimized above is only source-copy/readout count within the declared menu. Each trial consumes an independent preparation and reader port, a calibrated phase setting and clock interval. The archived setting and three-outcome record remain append-only and unavailable to future source operations. Histogram counters and the final decoder are additional classical resources. Randomized decisions also consume independent ready stock. An exact rational coin can be admitted as a separately calibrated qubit prepared by a finite rotation: $H_c = \hbar\Omega_c Y/2$ for duration $2\arcsin\sqrt{q}/\Omega_c$, followed by a terminal computational measurement. Its preparation and parameter accuracy are not free. With only fair bits, exact rational rejection sampling has finite expected cost but unbounded stopping time; a finite deadline requires an explicit residual error margin. No total energy, hardware or thermodynamic optimum follows.

For example, a continuous equatorial bisector-basis measurement is outside this menu and can outperform some small-budget axis schedules. There is also a useful unrestricted lower comparison for every always-deciding measurement on B copies, collective or otherwise:

$$e_B \geq \frac{1 - \sqrt{1 - 2^{-B}}}{2} > 0.$$

Indeed restrict the four-label minimax problem to the opposite-target pair $X+, Y-$, whose B -copy squared overlap is 2^{-B} . Their trace distance is $\sqrt{1 - 2^{-B}}$, obtained in their two-dimensional span. Maximizing $\text{Tr } E(\rho - \sigma)$ over $0 \leq E \leq I$ gives this trace distance, so the optimal pair-average error is the displayed value; four-label worst error cannot be smaller. This is the established two-state trace-distance bound, not a new global optimum. An unambiguous task with abstentions has a different objective. Collective measurements and adaptive basis choices are also excluded from the present optimum. These limitations are essential: arbitrary quantum-state discrimination is established theory, and the new result is a sharply typed target-relative design example with exact noise certificates.

General discrimination and multiple-copy local/adaptive comparisons have established treatments[1, 3]. The result here is a target-relative weighted access-design example with exact noise certificates under the declared preparation and measurement restrictions. No worldwide novelty claim is made for BB84-family discrimination or the general idea of allocating measurements.

2 A larger rational measurement menu: the copy gap closes but feedback still helps

The axis-restricted result proved a nineteen-versus-twenty-one copy threshold inside its X/Y menu. That statement does not imply a separation after additional phase settings become available. We make a strict, exactly computable expansion of the same physical reader and compare both protocol classes in that same expanded menu. The result distinguishes a surviving error advantage from a copy threshold that no longer separates the classes.

2.1 Declared settings, noise and symmetry

Admit exactly four equatorial axes, listed by their unit direction vectors:

$$\mathcal{M} = \{(1, 0), (0, 1), (3/5, 4/5), (4/5, 3/5)\}.$$

These are the original X/Y settings and two additional phase rotations $\theta = \arctan(4/3)$ and $\arctan(3/4)$, implemented by the specified controlled two-mode phase rotation, beam splitter and terminal detector.

The effects are $E_{\pm|a} = \eta(I \pm v(c_a X + s_a Y))/2$ and $E_{0|a} = (1 - \eta)I$. We use the same known, fixed common parameters $\eta = 4/5$, $v = 9/10$ for every setting. No setting-dependent calibration or switching cost is silently inferred from that assumption. All four source preparations, unavailable source label, independent fresh copies, fixed deadline and retained signs/nulls/settings keep the specified contract. Different observables act on different copies. This is a finite physical menu, not a discretization claimed to equal the continuous equatorial class.

The menu is closed under exchanging X and Y. Global source-sign reversal is implemented in a protocol by reversing its recorded signs and final decision. These two symmetries act transitively on the source labels. Consequently the uniform-prior Bayes optimum is a lower bound for minimax risk and averaging an optimal policy over the four symmetries attains it with equal error on each label. Two fair root bits suffice; our implementation resolves terminal posterior ties deterministically as plus before applying that root symmetry. This avoids charging an unneeded additional decision coin. The root seed must remain available to the decoder and is a counted controller resource.

2.2 Exact adaptive and randomized-nonadaptive certificates

For this menu all likelihoods have denominator $D = 125$. For the X reader the plus numerators on $(X+, X-, Y+, Y-)$ are $(95, 5, 50, 50)$, and for the two new directions they are $(77, 23, 86, 14)$ and $(86, 14, 77, 23)$. Minus swaps each sign pair and every null numerator is 25. Let $k_{a,o}$ denote these four-vectors. The exact recurrence on integer prior weights is

$$V_0(w) = \min(w_1 + w_3, w_2 + w_4), \quad V_t(w) = \min_{a \in \mathcal{M}} \frac{1}{125} \sum_{o \in \{+, -, 0\}} V_{t-1}(w \odot k_{a,o}).$$

Homogeneity permits division by the integer gcd of a weight vector, giving an exact rational memoized computation. Every action is included at each node. Selecting a minimizing action and then a terminal Bayes decision attains the lower bound. The preceding symmetry promotes the root value $V_B(1, 1, 1, 1)/4$ to minimax equality over all randomized, nonanticipating protocols using this menu.

For nonadaptive protocols, conditional product likelihoods imply that an action sequence's Bayes risk depends only on its four action counts. Enumerate all $\binom{B+3}{3}$ count allocations and all ternary histograms within each setting. At a histogram the multiplicity-weighted contribution is the minimum of its two target likelihood sums, divided by four. This is an integer sum over 125^B , with exact multinomial multiplicities. A mixture of schedules has average Bayes risk at least the smallest component risk. Root symmetrization of a minimizing schedule gives the matching minimax upper bound, so this comparator includes every randomized nonadaptive schedule and terminal decision rule in the same menu.

Theorem 2.1 (Finite expanded-menu frontier). *For the specified four-setting menu and detector law, the exact minimax errors are:*

B	<i>Adaptive</i>	<i>Randomized nonadaptive</i>
4	2062489/19531250	414539/3906250
8	177227695861393/4768371582031250	35860640527277/953674316406250

Both inequalities are strict, so feedback retains a genuine advantage within this larger menu. Nevertheless both classes require exactly fourteen copies for one-percent worst-case sign error.

Proof. The recurrence and full count-allocation enumeration above give exact lower and attaining upper certificates. The saved four/eight-copy Bellman DAGs contain both terminal losses and every action value, and the checker verifies all equalities with rational arithmetic. Direct evaluation of the selected policy's four label errors supplies an additional upper check before group averaging. Nonadaptive histogram sums are independently compared with ordered records on small horizons.

For the copy threshold, at $B = 13$ the errors are

$$\frac{20475493238962971614182459}{1818989403545856475830078125}, \quad \frac{20862994861778108626565401}{1818989403545856475830078125},$$

respectively, both greater than $1/100$. At $B = 14$ they are

$$\frac{162424078504800687452994223}{18189894035458564758300781250}, \quad \frac{165670382020832988542373811}{18189894035458564758300781250},$$

both less than $1/100$. The exact computation enumerates all 560 and 680 nonadaptive count allocations at those horizons and covers all adaptive actions by the same Bellman recursion. Its code, inputs and exact values are delivered; large threshold DAGs can be regenerated instead of being duplicated as a large static archive. More copies cannot worsen either optimum because an extra record can be ignored. The adjacent exact comparisons therefore establish the minimum budgets. $\square$

At eight copies the errors are approximately 0.03716734 and 0.03760261, both substantially lower than the old X/Y values 0.06714193 and 0.07977783. At fourteen they are 0.00892936 and 0.00910783. The best nonadaptive fourteen-copy allocation uses seven copies at each new oblique setting and none at X or Y. Thus a feasible broader nonadaptive reader already beats the nineteen-copy restricted adaptive budget. The original theorem remains correct in its original class, but its copy saving is not a measurement-menu-independent advantage.

There is a constructive repair at a stricter declared target. The fourteen-copy gap gives an entire interval of tolerances for which the adaptive reader passes and every nonadaptive reader fails. For example, at error tolerance $9/1000$ the exact minimum budgets are fourteen adaptive versus fifteen randomized nonadaptive copies. The thirteen-copy adaptive lower bound already exceeds this tolerance, and the fourteen-copy nonadaptive lower bound does too. For attainment at fifteen use the explicit allocation $(0, 0, 7, 8)$ and its root-symmetrized counterpart; its exact error is

$$\frac{206141219809325120449786527454}{28421709430404007434844970703125} \simeq 0.00725294938 < 9/1000.$$

Thus an expanded-menu copy advantage survives on a specified accuracy frontier, although the original one-percent separation closes. The 0.9% illustration was selected within the proved risk-gap interval after calculation; it is not presented as a pre-registered experimental target or a universal advantage.

2.3 Implementation and limits

`research9/menu/expanded_reader.py` supplies the exact kernels, generic finite-menu Bellman calculation, direct policy risks, schedule optimization and certificate checker. Its public minimax interface requires exact axis-swap closure; it rejects an asymmetric menu rather than incorrectly reusing uniform-prior minimax equality. Nominally unreachable branches have declared defaults. `research9/menu/verify_menu.py` checks 48 normalized-posterior reference problems, 16 ordered-record/histogram comparisons, malformed menus, a forged certificate, the four/eight-copy certificates and the thirteen/fourteen-copy thresholds. The replay runtime is hardware dependent; no asymptotic speed claim follows from a particular execution time.

The added settings, feedback controller, root coins, phase accuracy and switching times are additional apparatus resources. The current theorem optimizes copies and error within the declared menu, not total physical cost. Unequal setting efficiencies, common phase drift and uncertain controls are separate physical parents requiring new bounds. These exact rational results supply a useful restriction and benchmark for the continuous-menu investigation, not a global continuous optimum. The methods extend established Bayesian/minimax decision theory and the previous programme's source-readout interface; no worldwide priority or new Born, trine or consciousness result is asserted.

3 Nominal threshold: a checked obstruction and a continuous lower-backup method

The nominal one-percent copy threshold remains [13, 14] for each admitted reader class. This route does not establish thirteen-copy feasibility or impossibility. It supplies an exact thirteen-copy risk lower certificate and a convergent continuous-action lower-backup construction, while identifying why the tested cheaper relaxations do not settle the endpoint.

3.1 The thirteen-copy operator comparison

The local-effect grouped-target block representation and residual-corrected Gram dual bound also admit a thirteen-click certificate. Rational Gram factorization and exact residual charging give, at thirteen trials,

$$e_{13} \geq L_{13} = 0.009440886794141703\dots$$

The exact rational in `research12/threshold/verification_results.json` is authoritative. It applies to both declared local classes and is below 0.01: it cannot exclude thirteen copies. The supplied baseline thirteen-copy policy has error 0.011026493473...; the additional child-pool policy `research12/nominal/policy_13_pool.json` improves its certified error to 0.011009824045220384..., still above 0.01. The outward rational bounds are reproduced by `research12/nominal/verify_nominal12.py`. Fourteen copies attain 0.008749337644.... These are not matching lower and upper certificates.

We also tested a stricter first-step relaxation: the first command is local and chosen before erasures, and each posterior continuation is relaxed to a general permutation/local-transpose-invariant effect. The posterior need not be sign symmetric, so all real Pauli coefficient degrees, including the identity, enter the continuation SDP. On nine sampled directions from zero to $\pi/2$, the smallest proposed bound was about 0.00949372508, at the bisector; the axis endpoints gave about 0.01001652347. These are floating explorations, not continuous certificates. Their failure to close the gap is recorded without treating a sample minimum or solver status as proof.

3.2 A rational lower dynamic programme for the full circle

A finite angular menu normally restricts the reader and supplies upper, not lower, bounds on the continuous optimum. The following outer-action construction reverses that inclusion in a controlled way. Let $w \in \mathbb{R}_+^4$ be an unnormalized label weight vector. Write $V_t(w)$ for minimum expected terminal-error mass with t trials remaining. Thus $V_0(w) = \min(w_{X+} + w_{Y+}, w_{X-} + w_{Y-})$. Every fixed continuation has a four-component loss vector in $[0, 1]^4$; V_t is their infimum. It is positively homogeneous, concave, and on the unit simplex is 1/2-Lipschitz in ℓ_1 distance. For normalized w , the last statement follows because any loss vector has range at most one and $\sum_h(w_h - w'_h) = 0$.

For a real pair $a = (c, s)$ define the nominal kernels

$$K_{\pm,h}(a) = \frac{\eta}{2}(1 \pm va_h), \quad K_{0,h}(a) = 1 - \eta, \quad a_h = (c, -c, s, -s)_h.$$

These are physical for $c^2 + s^2 = 1$. They are still nonnegative normalized classical label kernels throughout $[-1, 1]^2$, although points outside the unit disk are explicitly fictitious relaxations, not allowed phase commands. Define

$$Q_t(w, a) = \sum_o V_{t-1}(w \odot K_o(a)).$$

It is concave in a wherever these kernels remain nonnegative.

Take rational unit supporting normals

$$(c_j, s_j) = \left(\frac{1 - (j/N)^2}{1 + (j/N)^2}, \frac{2j/N}{1 + (j/N)^2} \right), \quad 0 \leq j < N,$$

with their three quarter-turn rotations. Intersect successive supporting lines $c_j c + s_j s = 1$ to obtain the $4N$ vertices of a circumscribed polygon P_N . Cardinal normals ensure $P_N \subset [-1, 1]^2$; all fictitious kernels are therefore valid classical kernels. Let G_M be the simplex grid whose coordinates are integer multiples of $1/M$.

Theorem 3.1 (Certified outer-action/grid lower backups). *Suppose $L_{t-1}(q) \leq V_{t-1}(q)$ at every $q \in G_M$. For each grid point w and polygon vertex a , decompose each outcome weight as*

$$z_o = w \odot K_o(a) = \sum_i \gamma_{o,i} q_{o,i}, \quad \gamma_{o,i} \geq 0, \quad q_{o,i} \in G_M, \quad \sum_i \gamma_{o,i} = \|z_o\|_1.$$

Then

$$L_t(w) := \min_{a \in \text{vert } P_N} \sum_{o,i} \gamma_{o,i} L_{t-1}(q_{o,i}) \leq V_t(w).$$

Starting from exact $L_0 = V_0$ on the grid gives valid full-circle adaptive Bayes-risk lower bounds at every horizon. They also lower-bound minimax risk and every randomized nonadaptive reader. They do not establish local attainment of any polygon vertex.

Proof. Homogeneity and concavity of the true V_{t-1} imply $\sum_i \gamma_{o,i} L_{t-1}(q_{o,i}) \leq V_{t-1}(z_o)$. Every physical circle point lies in the convex hull of the polygon vertices. Concavity of the true Q_t gives $\min_{a \in \text{vert } P_N} Q_t(w, a) \leq \min_{a \in S^1} Q_t(w, a) = V_t(w)$. Combining the two inequalities proves the backup. In particular, no concavity of a chosen piecewise-linear interpolant is asserted or needed. The argument first compares the true Q_t on the outer polygon, then bounds its vertex values individually. Random schedules and feedback coins are covered because a Bayes minimum over randomized policies cannot be smaller than the minimum over its deterministic components. $\square$

Proposition 3.2 (Explicit convergence allowance). *Use the cumulative-coordinate simplex triangulation implemented in `convex_grid_bound.py`. At normalized grid beliefs,*

$$0 \leq V_t(w) - L_t(w) \leq t \left(\frac{4}{M} + \frac{\eta v}{4N^2} \right).$$

Proof. In cumulative coordinates, each interpolating simplex vertex differs from the normalized child by less than $1/M$ in each cumulative coordinate. The four original probability coordinates consequently differ by at most $2/M$ each. The $\ell_1/2$ Lipschitz bound charges at most $4/M$ for the interpolation; child masses sum to one. For consecutive normals the half-gap has tangent at most $1/N$, so any outer vertex has radius at most $\sqrt{1 + N^{-2}} \leq 1 + 1/(2N^2)$. Project it radially to the circle. For every label the outcome kernels change in total variation by at most $\eta v/(4N^2)$. The same bound applies to any continuation's terminal $[0, 1]$ loss and hence to its optimal value. Combining the one-step bounds with the preceding grid-value error and inducting proves the claim. This is an explicit conservative allowance, not a claim that it is sharp. $\square$

This is a finite-label specialization of established belief-grid bounding methods, with a rational outer-action interface for this continuous reader. Lovejoy's grid-based upper/lower POMDP approximation[4] is a close antecedent (Operations Research 39(1), 1991, pp. 162–175, DOI [10.1287/opre.39.1.162](https://doi.org/10.1287/opre.39.1.162)); the theorem above is proved independently. No worldwide novelty claim is made.

3.3 Exact demonstrations and scalability ceiling

The standard-library implementation constructs rational intersections, checks every support inequality and kernel normalization, and verifies every child’s nonnegative exact barycentric reconstruction. At $(M, 4N) = (8, 16)$, the exact lower values at the uniform prior are

$$L_1 = 6/25, \quad L_2 = 199/1250, \quad L_3 = 18278/171875.$$

They are deliberately coarse and do not improve the existing small-horizon closed forms. Replays at $(4, 8)$ and $(8, 8)$ verify the same algorithm under coarser interpolation; exact results are supplied. There are $\binom{M+3}{3}$ belief points, $4N$ actions and three outcomes; a straightforward backup has $O(tNM^3)$ rational operations and $O(NM^3)$ cached interpolation storage. Each child needs at most four vertices. Bit complexity is additional and depends on the exact rational directions and horizon.

To force the conservative allowance below 10^{-4} at thirteen trials, one sufficient choice is $M > 1,040,000$ together with a correspondingly large angular resolution. A full uniform grid is therefore impractical as an immediate endpoint solver. Targeted refinements, stronger continuation bounds or a different structure could improve it, but were not established here. We record this limit rather than propose an enormous unperformed run. The branch’s current threshold statement remains [13, 14] for each class.

4 Original robust task: fixed-nuisance certification and an unresolved sixteen-copy endpoint

Throughout the generic domain statements, $0 < \eta_0 \leq 1$; the benchmark uses $\eta_0 = 4/5$. The gain floors and phase bound retain the stated physical-parent constraints.

The original robust one-percent intervals remain [16, 17], separately for adaptive and randomized nonadaptive continuous local readers. A larger endpoint-policy search did not attain sixteen copies. We do not count its very small endpoint-risk improvement as a substantive threshold advance. The useful supporting result is an exact low-dimensional representation of any frozen finite policy’s entire nuisance-dependent risk, together with a sound, strict-feasibility-complete certificate procedure. It does not prove that sixteen copies are impossible.

4.1 A three-variable exact interface for a frozen reader

A finite policy is a directed acyclic graph, with one rational unit direction (c, s) and three outcome children at each internal node, a fixed deadline B , and sign decisions at leaves. It may be compiled from a larger history tree. A fixed rational root mixture is allowed by taking the corresponding weighted sum of the deterministic-root polynomials below. Classical root sign symmetrization is treated in the same way. No axis-exchange symmetry is assumed for the asymmetric parent.

For a fixed source label, let g be its relevant gain (g_X or g_Y), and introduce

$$e = \eta, \quad u = \eta g \cos \delta, \quad w = \eta g \sin \delta.$$

These parameters remain fixed throughout the complete graph. They are a coordinate change for the specified physical parent, not a new dynamics or an adversary choosing fresh values at each node. For a positive X source, the three outcome probabilities at direction (c, s) are

$$(e + cu - sw)/2, \quad (e - cu + sw)/2, \quad 1 - e.$$

For a positive Y source replace $cu - sw$ by $su + cw$; the negative label reverses the signed terms.

Theorem 4.1 (Exact frozen-policy nuisance polynomial). *For every such B -trial reader, each label risk is a rational polynomial $q_h(e, u, w)$ of total degree at most B . Its coefficients are computed bottom-up by the same three-child recurrence as the risk itself, with polynomial rather than scalar arithmetic. The number of possible monomials per label and node is at most $\binom{B+3}{3}$. For gain floor $a_h > 0$, the exact physical parameter domain is*

$$K_h = \{(e, u, w) : e \in [\eta_0, 1], u \geq 0, |w| \leq tu, a_h^2 e^2 \leq u^2 + w^2 \leq e^2\}, \quad t = \tan \Delta < 1.$$

Consequently uniform certification of this frozen reader is equivalent to $q_h \leq \epsilon$ on each K_h . The unused other-axis gain need not be introduced into a label's risk calculation.

Proof. At a leaf the error polynomial is zero or one. Each kernel is affine in (e, u, w) . Multiplying the three child risks by their kernels and adding raises the degree by at most one, proving the assertion by induction. For the domain, a specified physical point plainly satisfies the stated inequalities. Conversely the positive lower radial bound and $|w| \leq tu$ force $u > 0$. Define $g = \sqrt{u^2 + w^2}/e$ and $\delta = \arctan(w/u)$. Then $g \in [a_h, 1]$ and $|\delta| \leq \Delta$, and the original coordinates are recovered. Every point of K_h is therefore admitted. Maximizing separately over labels is valid because the original robust objective is a supremum over label and one nuisance choice fixed for that experiment. This does not authorize changing nuisance values within a single trajectory. $\square$

This reduces certification from a superficial five-parameter search to four three-variable polynomial problems. It does not reduce the unknown policy's optimization dimension, assume that the gain floors alone are worst cases, or certify arbitrary command-dependent errors outside the specified parent. Coefficient growth, graph compilation, precision and the cost of the resulting polynomial optimization still have to be charged. The bound is a monomial count, not a practical polynomial-time robust-design claim.

4.2 Exact Bernstein cover certificates

For a rational box $Q = \prod_i [l_i, r_i]$, map its coordinates affinely to the unit cube. If $\tilde{q}(t) = \sum_\beta a_\beta t^\beta$ has coordinate degrees d_i , its tensor Bernstein coefficients are

$$b_k = \sum_{\beta \leq k} a_\beta \prod_i \frac{\binom{k_i}{\beta_i}}{\binom{d_i}{\beta_i}}.$$

The Bernstein basis functions are nonnegative and sum to one, hence $\min_k b_k \leq q \leq \max_k b_k$ on Q . All entries are rational for the compiled policy and rational box. This standard polynomial-enclosure method is used here as an exact frozen-reader certificate, not claimed as a new approximation technique.

Start with $Q_0 = [\eta_0, 1] \times [0, 1] \times [-t, t]$. Discard boxes whose interval radial or angular inequalities exclude K_h ; otherwise accept a box only when every Bernstein coefficient is strictly below ϵ . Bisect unresolved boxes, cycling through their nontrivial coordinates. A finite accepted cover proves the uniform risk guarantee; a cap or interrupted search reports *inconclusive*.

Proposition 4.2 (Soundness and strict-feasibility completeness). *Every finite cover accepted by this procedure is a valid original-parent risk certificate. For a fixed polynomial with $\max_{K_h} q_h < \epsilon$, the uncapped subdivision procedure terminates. No termination is asserted when the maximum equals the target or exceeds it.*

Proof. Pruning only removes boxes disjoint from the physical set by necessary inequalities. Accepted boxes have a polynomial upper bound below the target, so their union proves the claim. For completeness, the compact set K_h has a positive strict risk margin. Continuity supplies a neighborhood on which the

margin persists. Bernstein coefficient ranges converge to the point value as a box shrinks. Every point outside K_h violates at least one defining closed inequality strictly; the elementary interval exclusions eventually recognize a sufficiently small box around it. Compactness of the initial box then yields a finite subdivision depth at which every surviving box is either excluded or risk-certified. Cycling nontrivial coordinates ensures shrinking box diameters. The argument is a completeness statement for strict certification, not a complexity bound or unrestricted quantum-reader theorem. $\square$

`polynomial_certificate.py` implements exact polynomial compilation, Bernstein bounds, physical-domain exclusions and a capped cover search. Independent ordered-word enumeration verifies 108 polynomial identities on a three-trial genuinely adaptive graph; 500 exact grid evaluations check the Bernstein enclosures. The proof, rather than those finite tests, covers arbitrary degrees. Constant safe and unsafe examples exercise the stopping and inconclusive statuses. This package does not claim that the cover code certified a new sixteen-copy operating region.

4.3 Why gain-endpoint checks alone are invalid

Even a legal fixed reader can have an interior maximum in one source label’s gain. Set $\eta = 1$, $\delta = 0$, use a direction with $c = 4/5$, $s = 3/5$ for sixteen trials, and decide minus exactly when fourteen plus outcomes occur. For X+ the error is

$$q(g) = \binom{16}{14} p(g)^{14} [1 - p(g)]^2, \quad p(g) = (1 + 4g/5)/2.$$

On $g \in [9/10, 1]$ it is maximized at $g = 15/16$, where $q = .28914388658\dots$, larger than the endpoint values $.28471846580\dots$ and $.27452150946\dots$. The maximizer follows from $14/p - 2/(1 - p) = 0$, and all three comparisons are exact rational inequalities in the verifier. This is a counterexample to unjustified per-label gain-endpoint certification, not a low-error reader or a counterexample to the correctly scoped frozen-majority theorem. Other labels of this deliberately poor reader can have still larger errors.

4.4 Searches and exact updated status

The specified transpose-invariant operator relaxation, evaluated at sixteen copies under the existing fixed hard prior, gives only approximately $.00838194185$. That value is exploratory and too weak to exclude sixteen; we did not repeat the unchanged solve or call its optimizer a physical reader.

A new point-based policy proposal used 24 rational directions, including rational approximations to the optimized frozen direction and nearby commands, with 150 sampled hard-endpoint trajectories and a design-belief bank. It retained approximately 12,214 candidate nodes at level sixteen. A separate linear programme mixed two roots and used a fresh sign-symmetry coin. The resulting pruned graph has 1,467 nodes and floating endpoint errors approximately $.01023427595453$ on each hard axis. This is still above one percent. Its tiny change from $.0102368487113$ does not settle an adjacent threshold and is not advertised as a robust advance. The full original nuisance continuum was not certified for this unsuccessful candidate.

Four separately optimized two-setting nonadaptive schedules used splits $(1, 15)$, $(2, 14)$, $(4, 12)$, $(8, 8)$, the exact endpoint offset $\arctan(.01)$, and an optimized terminal Bayes-decoder prior. Their best observed endpoint worst errors remained near the frozen benchmark. Such searches cannot exclude all schedules, all feedback policies, or the original continuous uncertainty set. They remain reproducible exploratory evidence.

Both full continuous robust thresholds therefore remain in $[16, 17]$. Seventeen is still the exact minimum only in the established random-frozen common-positive-angle class. Further repetition of these same policy families is not supported by the current search outcomes; a new structural lower bound or substantially different control construction would be needed.

5 Uniform angular approximation

5.1 A certified angular-net transfer, not a grid optimum

Theorem 5.1 (Uniform finite-horizon setting-rounding bound). *Let a rational menu contain the circle points*

$$n_k = \left(\frac{M^2 - k^2}{M^2 + k^2}, \frac{2Mk}{M^2 + k^2} \right), \quad -M \leq k \leq M,$$

their negatives and their axis swaps, where $M \geq 1$. For the fixed shared η, v parent, write e_B^{cont} and e_B^{grid} for the minimax error in either the adaptive class or the randomized nonadaptive class, consistently on both sides. Then

$$e_B^{\text{cont}} \leq e_B^{\text{grid}} \leq e_B^{\text{cont}} + \epsilon_B, \quad \epsilon_B = 1 - \left(1 - \frac{\eta v}{2M} \right)^B \leq \frac{B\eta v}{2M}.$$

Thus an independently certified grid interval $[L, U]$ gives the continuous interval $[\max(0, L - \epsilon_B), U]$ for the same policy class.

Proof. Up to sign, every direction is $n(t) = ((1 - t^2)/(1 + t^2), 2t/(1 + t^2))$ for $t \in [-1, 1]$. Round t to k/M , within $1/(2M)$. The chord identity

$$\|n(t) - n(t')\|_2 = \frac{2|t - t'|}{\sqrt{(1 + t^2)(1 + t'^2)}} \leq \frac{1}{M}$$

implies source-specific one-trial outcome-law TV at most $\eta v \|n - n'\|_2 / 2 \leq \eta v / (2M)$. For an adaptive policy the rounded controller maintains its virtual nominal history and applies the rounded setting at each branch. Coupling the outcome records sequentially gives probability of any divergence at most ϵ_B . The terminal target error changes by at most that amount on every source label. This rounds every continuous policy to a legal grid policy, without converting feedback to a nonadaptive schedule. Root mixtures are rounded componentwise. The converse inclusion of classes gives the other inequality. $\square$

The coupling compares outcome histories and their logical policy commands. It does not claim TV closeness of literal physical setting archives: a different real angle can make those raw setting records disjoint. All actual settings remain retained, and this theorem certifies target-decision risk rather than full retained quantum-channel equivalence. It assumes the same calibrated noise model at the omitted and retained angles; unrelated setting-dependent losses do not satisfy it automatically.

The code constructs the rational menu and an interval-transfer adapter, but it does not solve a large grid or advertise its illustrative interval as a measured continuous optimum. Required grid size, policy search cost and coefficient precision must all be charged by any consumer. This keeps the continuous-angle guarantee distinct from an uncertified finite scan.

5.2 Verification and exact status

The all-posterior formula is also exposed as a reusable terminal-control certificate: rational weights specify a direction numerator, squared norm and an exact rational-plus-square-root loss. It can eliminate the last continuous setting search in a policy optimizer; it does not by itself solve the preceding controls or establish a new minimax prior.

The standard-library implementation provides exact arithmetic in $\mathbb{Q}(\sqrt{2})$, the fixed-reader risks, small-horizon theorem values, the rational circle menu and the transfer allowance. Independent formulas check 85 click-mixture sums, 125 two-copy ordered-record calculations and 720 exact rational chord-rounding cases, plus 114 arbitrary-posterior terminal checks with rational Pythagorean directions.

Analytical proofs, rather than those finite tests, cover every angle and the all-horizon transfer. The angular-net interface alone establishes no large-horizon continuous optimum, collective optimum or total energy/apparatus advantage. Exact continuous equalities through three copies and four-copy comparison certificates are supplied separately in the reader paper. The physical interaction remains the declared controlled reader under ordinary Born probabilities. No trine, Born-selection or consciousness theorem is added.

6 From the continuous feedback gap to a physical resource certificate

The continuous four-copy construction and universal nonadaptive lower certificate support a resource statement: a feedback advantage in the *same complete equatorial menu*, including a nonzero operating region of asymmetric gain, erasure and phase errors. We also replace a crude setting-switch count by an exact actuator-path account and give a finite calibration interface covering the whole circle. All these are restricted operational claims, not a universal apparatus, energy or thermodynamic optimum.

6.1 A nonzero noise and control-error operating region

Let e_A denote the delivered exact rational ideal risk of the frozen four-copy adaptive tree, including its two root symmetry bits. Its value is $e_A = 0.0306041631615\dots$. All forty action nodes and eighty-one ternary leaves remain specified, including nominally impossible ideal null histories. Keep the four fixed-label preparations and inaccessible preparation controller. At a fixed unknown detector tuple use the asymmetric parent

$$\eta \geq \eta_0, \quad g_X, g_Y \geq g_0, \quad \eta, g_X, g_Y \leq 1,$$

with the fixed Pauli channel of Bloch multipliers $(g_X, g_Y, g_X g_Y)$ before the phase reader. Every commanded angle has total phase error of absolute value at most $d < 1$ radian, including the common fixed offset and declared control inaccuracy. The phase-error bound must apply to every command used by the policy; it is not inferred from earlier calibration alone. Fresh apparatus stock and copies are independent conditional on the fixed hidden source and nuisance parameters. For the adaptive/nonadaptive comparison, η, g_X, g_Y and the common offset are fixed throughout the block. Any residual optical/control error or quantization is a response to the command and trial index fixed in advance, never a function of deployment outcomes or an uncounted feedback controller. Independent stateless fresh noise may be included if declared separately. The uniform TV upper bound tolerates broader conditional errors, but that fact does not license outcome-dependent hidden feedback in the nonadaptive comparator.

Theorem 6.1 (Frozen-tree operating certificate). *The same implementable four-copy policy, fixed before the nuisance parameters are revealed, has worst sign error at most*

$$e_A + 1 - (1 - \kappa)^4, \quad \kappa = \frac{(1 - \eta_0) + (1 - \eta_0 g_0) + \eta_0 g_0 d}{2}.$$

The coupling bound controls the complete retained classical deployment record, including commanded/logical setting identifiers, nulls and controller coins, for the same frozen policy. Unknown actual analog phases are part of the response kernel, not additional accessible phase logs. It is not a norm bound on the residual quantum instrument or its bath.

Proof. For a source sign along either axis the ideal signed mean is one nominated component n_h of the unit direction. In the perturbed apparatus it is $\mu_h = \eta g_h n'_h$. The Euclidean chord is at most the phase error, so

$$|n_h - \mu_h| \leq 1 - \eta g_h + \eta g_h d \leq 1 - \eta_0 g_0 + \eta_0 g_0 d.$$

For ternary rows parametrized by efficiency and signed mean,

$$p = (\tfrac{1}{2}(\eta + \mu), \tfrac{1}{2}(\eta - \mu), 1 - \eta),$$

the exact row distance is

$$\text{TV}(p, p') = \tfrac{1}{2}\{|\eta - \eta'| + \max(|\eta - \eta'|, |\mu - \mu'|)\}.$$

The displayed mean allowance dominates $1 - \eta_0$, giving row distance at most κ . Couple identical root coins and, until the first record mismatch, use the same history-dependent command in both experiments. The probability of no mismatch is at least $(1 - \kappa)^4$ by the conditional row bound. Every terminal event, including a wrong decision, therefore differs in probability by at most $1 - (1 - \kappa)^4$. The bound is uniform in the four source labels and preserves one fixed nuisance tuple throughout the tree. The nominal final minimax risk is e_A , which proves the assertion. $\square$

For the explicit operating box

$$\eta_0 = g_0 = 99999/100000, \quad d = 1/10000,$$

we obtain the exact rational $\kappa = 12999790001/200000000000000$ and

$$e_A + 1 - (1 - \kappa)^4 < 0.030865 < 31/1000 < 1/32.$$

The total phase allowance is about 0.00573 degrees. This is a narrow, precisely declared operating region near the ideal apparatus, *not* an operating certificate for the principal noisy benchmark $\eta = .8, v = .9$ (whose separate finite-horizon comparison is handled in the continuous-reader paper), nor for the specified larger asymmetric box $a = .9, b = .8, \eta_0 = .8$.

The uncertainty set includes the ideal apparatus. Hence every robust randomized-nonadaptive protocol still has worst error at least the universal ideal lower bound $1/32$ at four copies. A stronger pointwise statement is also available: for any fixed nuisance tuple, each nonnull unsharp equatorial measurement is an ideal equatorial projector at its effective Bloch direction followed by an independent sign flip, together with the sign-blind erasure. A nonadaptive schedule remains nonadaptive under this simulation. Thus adding these detector imperfections cannot improve the full-menu nonadaptive optimum. This is a statistical simulation, not an identification of inaccessible environments.

An explicit five-copy nonadaptive upper construction repeats the bisector and takes fair-tie majority. Its ideal error is $1/2 - 43\sqrt{2}/128 < 2492/100000$. Applying the same row comparison at five copies gives

$$e_{N,5} \leq 2492/100000 + 1 - (1 - \kappa)^5 < 0.025245 < 31/1000.$$

Together with the independent ideal three-copy adaptive lower bound in the continuous-reader paper, these results give exact minimum budgets of *four adaptive versus five nonadaptive* at the selected target $31/1000$ throughout this operating box. The target was selected transparently inside a certified gap, not presented as a preregistered experimental tolerance. The fixed unknown detector tuple is never learned by an inaccessible oracle. The error guarantee holds for the same published four-copy policy over the entire box.

6.2 Actual phase travel, rather than a generic switch charge

There are two distinct physical timing parents. In the serial pulse parent, each fresh copy receives $H_z = \hbar\Omega(t)Z/2$ with $|\Omega(t)| \leq \Omega_p$. A fresh-copy phase command costs its required rotation angle divided by Ω_p . In the held-actuator parent, a calibrated phase element obeys $\dot{q} = u, |u| \leq \Omega_s$, and its current setting is retained between trials. Its time is the total actual phase travel divided by Ω_s . This

is an additional apparatus assumption, not a consequence of a single qubit's Hamiltonian pulse. Both start at phase zero. No terminal reset is included; any required reset is a separately charged operation.

Equatorial *axes* have period π if the controller keeps track of an outcome-polarity relabelling. This does not permit an instantaneous physical π jump. A continuous shortest axis path is lifted to the actual phase coordinate, and its known polarity bit recodes the recorded signs. Hardware without that permission instead pays for its actual angle range. Rates are in radians per second; all timing expressions below are in seconds.

Theorem 6.2 (Nonadaptive schedule ordering). *For independent copies and time-homogeneous fixed nuisance parameters, a nonadaptive schedule's decision risk is unchanged by permuting its settings and correspondingly permuting the labelled outcome data supplied to its decoder. Let S be the set of required projective-axis settings and include the actuator's initial point q_0 . On the angle circle of circumference π , delete each consecutive gap in turn. If the remaining arc has length L and q_0 has distances u and $L - u$ from its endpoints, then the exact minimum open actuator travel over schedule orders is*

$$\mathcal{L}(S; q_0) = \min_{\text{gaps}} \{L + \min(u, L - u)\}.$$

Repeated settings require no additional travel when measured consecutively. In the common-positive quadrant with $q_0 = 0$, increasing angle order attains $\mathcal{L} = \max S$.

Proof. The conditional likelihood is a product of setting-labelled factors, so a permutation with the inverse data relabelling leaves every source-label error unchanged. The physical timestamps change and are retained; this is not equality of the original timed apparatus channels.

An optimal open actuator route can be shorter than a full circle, since starting at q_0 and visiting all settings in one circular direction already has length strictly below π for a nontrivial finite set. Its travelled support therefore lies in a proper closed arc containing all required points and q_0 . Shortening either unused endpoint to a required point shows that the omitted gap may be chosen consecutive. On an interval of length L , a route starting at coordinate u must visit both endpoints. It pays at least $L + \min(u, L - u)$, attained by visiting the nearer endpoint first and then traversing to the farther endpoint. Minimize over omitted gaps. Conversely every such candidate route is implementable by the bounded-speed actuator. In the common-positive quadrant, direct increasing order from zero attains the farthest point's unavoidable distance. $\square$

This theorem gives the nonadaptive comparator its best setting order; it cannot be charged an arbitrary back-and-forth order to manufacture a feedback advantage. It does not apply to arbitrary drift or correlated future stock. An adaptive schedule generally cannot be reordered without changing its observation-dependent actions.

For the actual forty-node adaptive tree, evaluate every ternary history and both root axis-swap controls. The sign-reversal root operation is implemented by logical sign and decision recoding and adds no phase motion. Rational outward arctangent intervals give worst total phase costs

$$\mathcal{L}_A < 139/25 = 5.56 \quad (\text{held actuator}), \quad \mathcal{P}_A < 47/10 = 4.7 \quad (\text{per-copy pulse}).$$

The enclosed values are respectively 5.5517277189120... and 4.6905651101989... radians. All 162 ternary-history/root-swap cases are covered, so rare perturbed nulls have not been discarded.

A finite command interface can use fifteen-bit indices for 32768 equally spaced levels on $[0, \pi/2]$. Nearest rounding has error at most $\pi/[4(32767)] < 1/40000$ radian, using $\pi < 22/7$. Allocate at most $1/20000$ to the fixed offset and at most $1/40000$ to additional optical phase error; the total remains inside $d = 1/10000$. The archived command identifiers and finite packets are identical in the reference and perturbed record comparison; quantization enters their response kernel. No TV comparison of disjoint exact-analog-angle logs is asserted. Quantization adds at most seven rounding

errors to a four-trial held path and four to the pulse sum. Thus $7/40000$ and $4/40000$ respectively still fit inside the conservative 5.56 and 4.7 radian caps. The assumed slew model concerns the commanded actuator trajectory; bounded optical phase error alone would not bound arbitrary unmodelled actuator wandering or settling time.

The policy already uses a useful physical tie-break: on twenty-one final nodes no nonnull sign can change the nominated decision, and it keeps the previous setting. Twenty of those nodes have the same decision on all three outcomes. One has a different default decision on the ideal-impossible null, but its null probability is angle-independent in this parent. The retained terminal trial remains present. The repair preserves ideal risk and avoids unnecessary held-actuator movement; the operating certificate directly covers the resulting frozen tree rather than assuming all residual records are unchanged.

6.3 A full-menu copy advantage with a conditional time advantage

Let $\tau > 0$ be the serial per-copy duration for preparation, propagation, detection and matched terminal handling, excluding the separately modelled phase operation. Let $\ell \geq 0$ upper-bound each of the three noninitial causal lookup/feedback delays. Sufficient adaptive time bounds are

$$T_{A,s} \leq 4\tau + 3\ell + \frac{139}{25\Omega_s}, \quad T_{A,p} \leq 4\tau + 3\ell + \frac{47}{10\Omega_p}.$$

Every nonadaptive protocol meeting the same target within the same continuous menu and operating box needs at least five fresh copies, so its time is at least 5τ , even before its optimized phase travel is charged. Consequently either

$$3\ell + \frac{139}{25\Omega_s} < \tau, \quad \text{or} \quad 3\ell + \frac{47}{10\Omega_p} < \tau$$

suffices for a strict clock advantage in its corresponding serial parent. These are conservative sufficient conditions, not exact hardware optima. The two rates refer to different controls and cannot be substituted for one another by reusing a single apparatus certificate.

For $\Omega_s\tau = 20$ and $\ell/\tau = 1/20$, the held implementation saves at least $143\tau/250 = 0.572\tau$ per experiment relative to the universal nonadaptive lower time. In the separately specified pulse parent with $\Omega_p\tau = 20$ the lower saving is $123\tau/200 = 0.615\tau$. These are illustrative declared rates, not measured device data. If an actual extra net setup/calibration time D is charged to the adaptive apparatus, R repeated uses certify a benefit only when R times that lower saving exceeds D . For the illustrative $D = 10\tau$, the sufficient counts are 18 and 17 uses, respectively. Reuse requires a calibration/reference hold contract over the whole block. The nonadaptive reader is not forced to calibrate settings it never uses. Failed sufficient inequalities do not prove that feedback is slower.

The constructive adaptive controller may store a seven-bit pointer over the 121 internal/terminal prefixes; the remaining deadline is encoded in the pointer's depth. Two root group bits remain separate. A direct finite-command ROM account uses at most $80 \cdot 15 = 1200$ angle bits for both root axis controls, plus 81 terminal decision bits; indexed tree arithmetic avoids a separate child-pointer table. Arithmetic workspace, clock precision, the held actuator's physical state, phase tracking and the reference remain additional resources. A fully explicit ternary-prefix archive can be retained independently. These are upper resource accounts, not minimum-controller theorems. They cannot be turned into a thermodynamic work, energy or apparatus-volume conclusion.

6.4 Finite calibration covering all continuous angles

A finite calibration route is available under an explicit structural promise: shared sign-blind efficiency, unbiased outcomes, and a time-homogeneous sinusoidal effect law. For each trusted positive source

$h = X+, Y+$ write

$$\mu_h(\theta) = A_h \cos \theta + B_h \sin \theta,$$

with the negative source having mean $-\mu_h(\theta)$. The asymmetric Pauli/phase parent with fixed gains and a common fixed offset has exactly this form about the known realized command angle. This calibration theorem is a separate structural promise: arbitrary command-dependent optical errors or nonlinear quantization residuals need not preserve the sinusoid. They require a separately validated uniform row-error allowance, additional command calibration, or a justified response model; four probe batches do not certify them automatically. Use M independent trusted probes for each of the four pairs

$$(X+, 0), (X+, \pi/2), (Y+, 0), (Y+, \pi/2).$$

The recorded variable $Z \in \{-1, 0, 1\}$ estimates the four coefficients; the pooled click fraction over $4M$ probes estimates η . Hoeffding and a union bound give the simultaneous predeployment event

$$\Pr \left(\max_h \{ |\hat{A}_h - A_h|, |\hat{B}_h - B_h| \} > u \text{ or } |\hat{\eta} - \eta| > w \right) \leq 8e^{-Mu^2/2} + 2e^{-8Mw^2}.$$

For completeness, the bounded-variable inequality follows by writing $K(t) = \log \mathbb{E} e^{t(Z - \mathbb{E}Z)}$. Under exponential tilting, $K''(t)$ is a variance and is at most one quarter of the squared support width; $K(0) = K'(0) = 0$ then bounds the moment generating function by $e^{t^2(b-a)^2/8}$. Chernoff optimization for an independent sample gives the two-sided bound $2e^{-2Mu^2/(b-a)^2}$. Apply widths two and one to the signed means and pooled clicks respectively, then union the five events. For $\alpha = 2^{-k}$ the exact conservative choice

$$M = \left\lceil \max \left\{ \frac{2(k+4)}{u^2}, \frac{k+2}{8w^2} \right\} \right\rceil$$

makes the failure probability at most α ; total consumed stock is $4M$. This bound covers *every* angle, because the coefficient error has absolute value at most $u(|\cos \theta| + |\sin \theta|) \leq \sqrt{2}u$. It cannot validate an arbitrary detector between unmeasured settings; without the sinusoidal/common-efficiency promise additional coverage is required. Trusted source preparation and independence are paid premises.

For any nominated valid reference kernel whose four coefficients and efficiency lie inside those same calibration intervals, the actual and reference kernels differ by at most

$$\kappa_{\text{cal}} = \min\{1, w + \max(w, \sqrt{2}u)\}$$

in every source/angle row. Thus the same policy at horizon B has retained classical commanded-setting/outcome record distance at most $1 - (1 - \kappa_{\text{cal}})^B$. It excludes an extra unread analog-phase log. This includes policies and angles selected from the calibration data: the event covers the whole continuous family before deployment. It does not identify the full residual quantum instrument or unread environment.

If calibration acceptance is A , a selected reference policy has verified reference error e and transfer allowance d_B , then

$$\Pr(A \cap \text{error}) \leq (e + d_B) \Pr(A) + \alpha.$$

An accepted-run conditional error bound must retain the corresponding $\alpha/\Pr(A)$ term, unless a separate acceptance-probability guarantee justifies a smaller expression. Bad calibration is charged once for a stable shared block; past samples do not prove future stability or a reference-reuse lifetime.

This is a completed finite interface, not a practical calibration-cost claim. For example, $k = 10, u = w = .01$ consumes 1,120,000 probes but does not by itself certify any advertised reader accuracy. A much tighter illustration, $k = 20, u = 1/22000, w = 1/88000$, consumes 92,928,000,000 probes. Its conservative coupling allowance certifies $\Pr(A \cap \text{error}) < .01$ for the nominal fourteen-copy bisector under the stated reference-acceptance contract; this is not an unqualified conditional-on-acceptance bound.

The stock cost plainly overwhelms a few saved copies. Better device-specific likelihood calibration can be pursued; it is not hidden inside the present resource claim.

Phase-reference preparation, conditional update, reset, holding and consumption remain separately declared resources. Established reference- degradation analyses may bound an average even when records are kept; that does not automatically supply a uniform conditional guarantee or the adaptive interaction model used here[2]. No reference lifetime is transferred without matching those contracts.

6.5 Reproducible checks and remaining limits

The delivered `research10/resources/verify_resources.py` and `research10/resources/policy_resources.py`, with the independent `research10/audit/check_resources.py`, check 84 exhaustive schedule-tour instances against independent permutation enumeration, 873 individual quadrant orders, 81 exact ternary row distances, the two exact calibration allocations, and the rational operating and clock inequalities. The policy cost evaluator uses the delivered tree and outward arctangent intervals rather than a floating-point maximum. The independent checker can verify these certificates without trusting the exploratory angle search. The nominal .8/.9 continuous one-percent threshold, the larger asymmetric robust frontier and universal total-resource optimization remain unresolved. The source-readout contribution is a typed full-menu reader, calibration and cost interface, not a Born-selector, source-ontology or consciousness mechanism.

References

- [1] Stephen M. Barnett and Sarah Croke. Quantum state discrimination. *Advances in Optics and Photonics*, 1(2):238–278, 2009. doi: 10.1364/AOP.1.000238. URL <https://opg.optica.org/aop/abstract.cfm?URI=aop-1-2-238>.
- [2] Stephen D. Bartlett, Terry Rudolph, Robert W. Spekkens, and Peter S. Turner. Degradation of a quantum reference frame. *New Journal of Physics*, 8:58, 2006. doi: 10.1088/1367-2630/8/4/058. URL <https://arxiv.org/html/quant-ph/0602069v2>.
- [3] 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>.
- [4] William S. Lovejoy. Computationally feasible bounds for partially observed markov decision processes. *Operations Research*, 39(1):162–175, 1991. doi: 10.1287/opre.39.1.162.
- [5] Jeremy Rodgers. Source-readout non-equivalence: Descent and equivariant reconstruction obstructions, 2026. URL <https://doi.org/10.5281/zenodo.21369967>.