

Technical Companion: Approximation and Accounting for Quantum Sources

Jeremy Rodgers
Independent Researcher
www.everythingequation.com

27 October 2026

Abstract

This companion preserves alternate exact encodings, finite-record approximation, output-specific bath approximation and causal resource accounting. Exact minima must be used only with their declared tasks. Historical feasible constructions remain valid but are superseded as best known memory intervals or channel costs where explicitly noted.

This companion develops approximation and resource-accounting results within Shadow Theory’s pre-existing source-readout programme [10]. The stationary classical completion used for comparison is established in Rodgers [9].

1 A noisy delayed trine process: exact qutrit realizations

The formerly unresolved positive-noise case admits an exact three-dimensional quantum source in a nonempty, explicitly certified regime. This improves the two-qubit upper bound. The result preserves every finite output word, and its cone constraints are certified on an entire parameter interval by exact rational arithmetic. A separate analytic construction gives a narrower interval without computer-assisted positivity checks. Both constructions use ordinary quantum instruments; neither derives the Born rule or an actual-event selection law.

1.1 Process, source class and baseline bounds

For $0 \leq \lambda \leq 1$, write $r = 1 - \lambda$ and

$$T_{aa}^\lambda = \frac{1-r}{3}, \quad T_{ab}^\lambda = \frac{1+r/2}{3} \quad (a \neq b). \quad (1)$$

All parameters and probabilities are dimensionless. Interleave two independent stationary chains with this symmetric doubly stochastic transition matrix. The stationary word law is

$$P_\lambda(y_1 \cdots y_n) = \frac{1}{9} \prod_{t=1}^{n-2} T_{y_t, y_{t+2}}^\lambda, \quad n \geq 2, \quad P_\lambda(y_1) = 1/3. \quad (2)$$

The source class is a time-homogeneous quantum instrument on $\mathbb{C}^d$: CP maps Φ_j with $\sum_j \Phi_j$ trace preserving, a stationary density matrix ω , and word probabilities $\text{Tr}[\Phi_{y_n} \cdots \Phi_{y_1}(\omega)]$. Outputs are emitted into a retained classical archive. The current symbol need not occupy an orthogonal sector of the internal memory. Fresh Markov environments are admitted explicitly; the cost d counts internal memory, not archives, purifying environments, or a hidden history-dependent controller.

The eigenvalues of T^λ are $1, -r/2, -r/2$. For $\lambda < 1$ the past-two/future-two cut matrix

$$H_{ab,cd} = \frac{1}{9} T_{ac}^\lambda T_{bd}^\lambda \quad (3)$$

has ordinary rank nine. Classical finite HMMs factor this matrix through their hidden states, so require at least nine. Explicitly, the pair source has states (a, b) , transition $(a, b) \mapsto (b, j)$ with probability T_{aj}^λ , uniform stationary weights $1/9$, and emitted symbol j at its update; it attains the full law with nine states. A d -dimensional instrument gives factors $A_{ab} = \Phi_b \Phi_a(\omega)$ and $B_{cd} = \Phi_c^* \Phi_d^*(I)$, with $H_{ab,cd} = \text{Tr}(A_{ab} B_{cd})$, in the d^2 -dimensional real space of Hermitian matrices, so $d \geq 3$.

A four-dimensional construction exists for every λ . With equatorial trine Bloch vectors n_a , let $\rho_a = (I + n_a \cdot \sigma)/2$ and $E_j = (I - r n_j \cdot \sigma)/3$. Then

$$\Phi_j^{(4)}(X) = \text{Tr}_1[(E_j \otimes I)X] \otimes \rho_j \quad (4)$$

is CP, has effects $E_j \otimes I$ summing to identity, and fixes $I_4/4$ after summing j . It sends $\rho_a \otimes \rho_b$ to $T_{aj}^\lambda \rho_b \otimes \rho_j$, establishing every word of (2). Thus

$$3 \leq d_{\min}(\lambda) \leq 4 \quad (0 \leq \lambda < 1), \quad (5)$$

whereas at $\lambda = 1$ the law is iid uniform and $d_{\min} = 1$. The rank argument and quantum-process framework have established antecedents [5, 8]; no new rank method is claimed.

1.2 A general exact encoded-pair construction

Let $x_0, x_1, x_2 \in \mathbb{R}^2$ have zero sum, and let a diagonal metric $g = \text{diag}(g_1, g_2)$ obey

$$\sum_{\mu=1}^2 g_\mu x_{a\mu} x_{j\mu} = \begin{cases} 1 & a = j, \\ -1/2 & a \neq j. \end{cases}$$

Let $A_1, A_2, B_1, B_2, C_{11}, C_{12}, C_{21}, C_{22}$ be any basis of traceless Hermitian 3×3 matrices. Use the contractions

$$A_a = \sum_\mu x_{a\mu} A_\mu, \quad B_b = \sum_\nu x_{b\nu} B_\nu, \quad C_{ab} = \sum_{\mu,\nu} x_{a\mu} x_{b\nu} C_{\mu\nu}, \quad C_{\nu,j} = \sum_\eta x_{j\eta} C_{\nu\eta}.$$

For nonzero positive coefficients u, v, w , define candidate pair states

$$S_{ab} = I_3/3 + u A_a + v B_b + w C_{ab}, \quad R_j = I_3/3 + v B_j. \quad (6)$$

Specify a Hermiticity-preserving complex-linear map by its values on the Hermitian basis:

$$\Phi_j(I_3/3) = R_j/3, \quad (7)$$

$$\Phi_j(A_\mu) = -\frac{r}{3u} g_\mu x_{j\mu} R_j, \quad (8)$$

$$\Phi_j(B_\nu) = \frac{u}{3v} A_\nu + \frac{w}{3v} C_{\nu,j}, \quad (9)$$

$$\Phi_j(C_{\mu\nu}) = -\frac{ru}{3w} g_\mu x_{j\mu} A_\nu - \frac{r}{3} g_\mu x_{j\mu} C_{\nu,j}. \quad (10)$$

Lemma 1.1 (Exact update and full-law realization). *If all S_{ab} are positive and all maps in (10) are CP, then they form a quantum instrument with stationary state $I_3/3$ and generate (2) exactly. In particular,*

$$\Phi_j(S_{ab}) = T_{aj}^\lambda S_{bj}. \quad (11)$$

For $r > 0$ this attains the minimum dimension three.

Proof. All basis matrices except I are traceless, so the candidates have trace one. Expand (10) on (6), writing $q_{aj} = \sum_{\mu} g_{\mu} x_{a\mu} x_{j\mu}$. The constant and a terms give $(1 - r q_{aj}) R_j / 3$; the b and ab terms give $(1 - r q_{aj})(u A_b + w C_{bj}) / 3$. This is $T_{aj}^{\lambda} S_{bj}$, proving the update.

Summing the trace of the maps on the basis yields trace preservation: the $I/3$ traces sum to one, the A traces vanish because $\sum_j x_j = 0$, and the B, C images are traceless. Summing $\Phi_j(I/3) = R_j/3$ gives $I/3$, proving stationarity. The first output j has probability $1/3$ and leaves state R_j . A direct application gives $\Phi_k(R_j) = S_{jk}/3$, so the second output is independent uniform and leaves S_{jk} . Every later update is (11); induction proves all finite words. This does not infer temporal independence from a reset: it uses the explicitly specified CP instrument at every step. Finally (3) gives the matching dimension-three lower bound. $\square$

For positive u, v, w , the nine candidate matrices span the full Hermitian qutrit space. The functions $1, x_{a1}, x_{a2}$ form a basis of functions on three labels; their tensor products provide the nine independent coordinate functions used above. Thus the construction is not silently collapsing an unobserved classical coordinate or omitting an internal controller.

1.3 An analytic high-noise interval

Choose the ordinary unit trine coordinates $x_a = (\cos(2\pi a/3), \sin(2\pi a/3))$, $g_1 = g_2 = 1$, and an orthonormal traceless Hermitian basis with $\text{Tr}(F_k F_l) = \delta_{kl}$, split as the A, B, C basis above. Put

$$r = \epsilon^3, \quad u = \epsilon^2, \quad v = \epsilon, \quad w = \epsilon^3, \quad 0 < \epsilon \leq 1/7. \quad (12)$$

The state perturbation has exact squared Frobenius norm $\epsilon^2 + \epsilon^4 + \epsilon^6 < 1/9$, so every S_{ab} is positive. For the unnormalized Choi matrix J_j of (10), the depolarizing branch $X \mapsto \text{Tr}(X)I/9$ has Choi matrix $I_9/9$. Orthonormality of the input basis makes squared Frobenius norms add, giving

$$\begin{aligned} \|J_j - I_9/9\|_F^2 &= \frac{\epsilon^2}{3} + \frac{\epsilon^2}{9}(1/3 + \epsilon^2) + \frac{2(\epsilon^2 + \epsilon^4)}{9} + \frac{2(\epsilon^4 + \epsilon^6)}{9} \\ &= \frac{16\epsilon^2}{27} + \frac{5\epsilon^4}{9} + \frac{2\epsilon^6}{9}. \end{aligned} \quad (13)$$

At $\epsilon = 1/7$, eighty-one times this expression is $117471/117649 < 1$; the expression increases with $\epsilon \geq 0$. The operator norm is at most the Frobenius norm, so J_j is strictly positive and the map is CP by the Choi criterion [1]. Lemma 1.1 therefore proves

$$\boxed{d_{\min}(\lambda) = 3 \quad \text{for } 342/343 \leq \lambda < 1.} \quad (14)$$

This is a complete analytic interval proof, independent of the sharper finite certificate below.

1.4 A larger interval certified entirely by rational arithmetic

Take rational trine coordinates

$$x_0 = (1, 0), \quad x_1 = (-1/2, 1/2), \quad x_2 = (-1/2, -1/2), \quad (g_1, g_2) = (1, 3).$$

Let E_{pq} denote matrix units and set

$$\begin{aligned} A_1 &= E_{01} + E_{10}, & A_2 &= -iE_{01} + iE_{10}, \\ B_1 &= E_{02} + E_{20}, & B_2 &= -iE_{02} + iE_{20}, \\ C_{11} &= E_{12} + E_{21}, & C_{12} &= -iE_{12} + iE_{21}, \\ C_{21} &= \text{diag}(1, -1, 0), & C_{22} &= \text{diag}(1, 1, -2). \end{aligned}$$

These are orthogonal, with squared Frobenius norms two except $\|C_{22}\|_F^2 = 6$. This is a different, explicitly specified frame from the preceding analytic construction. Put

$$r = \epsilon^3, \quad u = \alpha\epsilon^2, \quad v = \beta\epsilon, \quad w = \gamma\epsilon^3, \quad (\alpha, \beta, \gamma) = (3/8, 3/4, 3/4), \quad 0 < \epsilon \leq h := 7/20. \quad (15)$$

The maps and states are now degree-three matrix polynomials in ϵ , with Gaussian rational coefficients.

Here are explicit coefficients, making the positivity certificate reconstructible without fitting numerical samples. Superscript $\vee$ denotes the Hilbert–Schmidt dual basis ($F^\vee = F/2$, except $C_{22}^\vee = C_{22}/6$). Define

$$A_j^{g,\vee} = \sum_{\mu} g_{\mu} x_{j\mu} A_{\mu}^{\vee}, \quad C_{j,\nu}^{g,\vee} = \sum_{\mu} g_{\mu} x_{j\mu} C_{\mu\nu}^{\vee}.$$

In the convention $J_j = \sum_{pq} E_{pq} \otimes \Phi_j(E_{pq})$, write $J_j(\epsilon) = \sum_{l=0}^3 \epsilon^l J_{jl}$. Then

$$J_{j0} = I_9/9, \quad (16)$$

$$J_{j1} = \frac{\beta}{3} I \otimes B_j - \frac{1}{9\alpha} (A_j^{g,\vee})^T \otimes I + \frac{\alpha}{3\beta} \sum_{\nu} (B_{\nu}^{\vee})^T \otimes A_{\nu}, \quad (17)$$

$$J_{j2} = -\frac{\beta}{3\alpha} (A_j^{g,\vee})^T \otimes B_j + \frac{\gamma}{3\beta} \sum_{\nu} (B_{\nu}^{\vee})^T \otimes C_{\nu,j} - \frac{\alpha}{3\gamma} \sum_{\nu} (C_{j,\nu}^{g,\vee})^T \otimes A_{\nu}, \quad (18)$$

$$J_{j3} = -\frac{1}{3} \sum_{\nu} (C_{j,\nu}^{g,\vee})^T \otimes C_{\nu,j}. \quad (19)$$

For any degree-three matrix polynomial $P(\epsilon) = \sum_l \epsilon^l P_l$, its Bernstein representation on $[0, h]$ is

$$P(ht) = \sum_{k=0}^3 \binom{3}{k} t^k (1-t)^{3-k} M_k, \quad M_k = \sum_{l=0}^k \frac{\binom{k}{l}}{\binom{3}{l}} h^l P_l. \quad (20)$$

The coefficients outside M_k are nonnegative and sum to one. Thus positive control matrices certify positivity on the whole interval, not at sampled values.

Theorem 1.2 (Exact high-noise qutrit interval). *For the noisy delayed trine law,*

$$\boxed{d_{\min}(\lambda) = 3 \quad 7657/8000 \leq \lambda < 1.} \quad (21)$$

The concrete states and maps are (6), (10), and (15).

Proof. For the twelve Choi Bernstein control matrices from (19)–(20), and the thirty-six state control matrices from (6), exact rational LDL^T certificates establish strict positivity. The certificate uses the realification

$$\mathcal{R}(Z) = \begin{pmatrix} \operatorname{Re} Z & -\operatorname{Im} Z \\ \operatorname{Im} Z & \operatorname{Re} Z \end{pmatrix};$$

a Hermitian matrix is positive definite precisely when its realification is. In the fixed computational-basis ordering, the certified bounds on every LDL pivot are:

Bernstein index k	0	1	2	3
Choi control matrices (all j)	1/9	> 9/100	> 3/50	> 1/100
State control matrices (all a, b)	> 1/10	> 1/10	> 1/10	> 1/10

These are pivot bounds, not eigenvalue bounds. They are verified with exact `Fraction` arithmetic in `supplement/memory/three/memory/exact_qutrit_certificate.py`; the explicit rational control matrices and all pivots are recorded in `supplement/memory/expected_outputs/three/memory/qutrit_bernstein_matrices.json` and `supplement/memory/expected_outputs/three/memory/qutrit_certificate.json`. For clarity the finite algebraic verification is the standard recurrence

$$D_k = M_{kk} - \sum_{l < k} L_{kl}^2 D_l, \quad L_{ik} = \frac{M_{ik} - \sum_{l < k} L_{il} L_{kl} D_l}{D_k}, \quad L_{kk} = 1.$$

It checks $D_k > 0$ and reconstructs every entry of LDL^T exactly. This computer-assisted part involves only 48 explicitly defined rational matrices; no numerical eigensolver or tolerance certifies their signs. Bernstein convexity proves positivity throughout $0 \leq \epsilon \leq 7/20$.

The same code independently verifies 197 exact polynomial identities by applying the Choi polynomial directly: all 27 reset identities through degree six, and trace-preservation and stationarity through degree three. The universal symbolic proof of those identities is already Lemma 1.1. It follows that this is a stationary CP instrument reproducing every word. The ordinary-rank lower bound gives minimality. Since $r = \epsilon^3$, the interval is $\lambda = 1 - r \geq 1 - (7/20)^3 = 7657/8000$, excluding the iid endpoint from the dimension-three assertion. $\square$

1.5 Finite observations versus exact memory requirements

The exact dimension remains at least three for every $\lambda < 1$, even arbitrarily close to the iid endpoint. This does not create a finite-data discontinuity. Let U_n be the uniform iid law on n trits and $n \geq 2$. Define

$$d(r) = \frac{1-r}{3} \log(1-r) + \frac{2(1+r/2)}{3} \log(1+r/2).$$

With $0 \log 0 = 0$, direct conditional-likelihood calculations give

$$D(P_{\lambda,n} \| U_n) = (n-2)d(r), \quad \chi^2(P_{\lambda,n} \| U_n) = (1+r^2/2)^{n-2} - 1. \quad (22)$$

Indeed the likelihood ratio is $\prod_{t=1}^{n-2} 3T_{y_t, y_{t+2}}^\lambda$. The KL chain rule gives one equal row contribution $d(r)$ per link. Under U_n , summing the newest symbol in the squared likelihood contributes $3 \sum_j (T_{aj}^\lambda)^2 = 1+r^2/2$, independently of a ; recursion proves the chi-square identity. Therefore

$$\text{TV}(P_{\lambda,n}, U_n) \leq \min \left\{ 1, \sqrt{(n-2)d(r)/2}, \frac{1}{2} \sqrt{(1+r^2/2)^{n-2} - 1} \right\}. \quad (23)$$

The second bound is Pinsker and the third is Cauchy–Schwarz. Since $d(r) = r^2/4 + O(r^3)$, fixed finite records approach the iid law continuously. These are established divergence inequalities evaluated exactly for this process, not new information measures. The comparison concerns the complete declared symbol sequence, including its ordering; extra retained apparatus channels would require their joint law as well.

1.6 What the construction establishes

The original trine phase geometry supplied a four-dimensional realization; its zero-noise optimality proof did not settle noisy qutrit models. The present construction supplies such models explicitly and proves exact minimality on a full interval. It introduces a hierarchical encoding of the two individual labels and their joint Fourier components, with amplitudes that shrink at different rates. This permits the update maps to remain inside the CP cone while retaining the complete order-two law. The sharper

interval uses a rational qutrit frame and a convex polynomial certificate; it does not claim the endpoint is optimal.

The finite circular phase representation draws on the geometric phase and angular-register descriptions in [6, abstract and Sections 2–3] and [7, pp. 64–68]. The noisy qutrit encoding, update maps and interval certificates are the constructions established here.

The CP-realization problem, nonorthogonal quantum process memories, quantum dimension advantages and lower bounds are established research areas [3–5, 8, 11]. The contribution is the explicit full-law construction, interval certificate and matching dimension bound for this declared delayed process. Targeted primary-literature checks do not establish publication priority. No claim is made about reduced complete environment memory, a Hamiltonian mechanism, CSL dynamics, Born-rule selection, actual-event affinity, or conscious perspective.

2 Exact finite-record accuracy and a memoryless comparator

An exact resource minimum does not alone establish that the resource is needed at a finite experimental accuracy. This section evaluates that distinction for the same full trine law, without changing its parameters or introducing an alternative Born rule. It is an analytical addition to the completion theorem’s rank-based certificates. It is not a claimed optimum over all approximate hidden-source realizations.

Let $P_{\lambda,N}$ be the N -symbol stationary word law in the main resource manuscript, and let $U_N(y_1, \dots, y_N) = 3^{-N}$. The latter is generated by the one-dimensional quantum instrument $\Phi_j(x) = x/3$, or a one-state classical source with fresh uniform trit emissions. Both have the same specified output alphabet and the same record length. Fresh randomness or its dilation still costs resources and is not contained in the one-dimensional memory count.

Theorem 2.1 (Exact binomial reduction of the finite-record discrepancy). *For $N \geq 2$ set $k = N - 2$ and $R = \sum_{t=1}^k \mathbf{1}\{Y_t = Y_{t+2}\}$. Then*

$$\text{TV}(P_{\lambda,N}, U_N) = \text{TV}(\text{Bin}(k, \lambda/3), \text{Bin}(k, 1/3)).$$

The count R is a sufficient statistic for this binary experiment. Thus even retaining every symbol gives no stronger test for these two models than retaining that count. This is an assertion about the declared classical word record, not about their microscopic output environments.

Proof. For each word with r repeats,

$$\frac{P_{\lambda,N}(y)}{U_N(y)} = \lambda^r \left(\frac{3 - \lambda}{2} \right)^{k-r}.$$

Conditional on every positive-probability preceding history, the next repeat indicator has probability $\lambda/3$. Iterating conditional probabilities proves that these k indicators are iid Bernoulli, despite the overlapping symbol comparisons. Under U_N their parameter is $1/3$. The likelihood ratio is constant on each fibre of R ; summing absolute signed masses fibre by fibre therefore gives the equality in TV, with no data-processing slack. $\square$

For $0 < \lambda < 1$ put $b = (3 - \lambda)/2$ and $r_* = \lfloor k \log b / (\log b - \log \lambda) \rfloor$. Since the likelihood ratio decreases with r , the exact error is

$$F_{k,\lambda/3}(r_*) - F_{k,1/3}(r_*),$$

where F is the binomial CDF. Equality at an integer crossing contributes zero signed mass, so either threshold convention is valid. At $\lambda = 0$ the result is $1 - (2/3)^k$, and at $\lambda = 1$ it is zero. For $N \leq 2$ both word laws agree.

Two closed-form bounds are also available:

$$\text{TV}(P_{\lambda,N}, U_N) \leq \min \left\{ 1, \frac{1}{2} \sqrt{\left(1 + \frac{(1-\lambda)^2}{2}\right)^k - 1} \right\},$$

$$\text{TV}(P_{\lambda,N}, U_N) \leq \sqrt{1 - \left[\frac{\sqrt{\lambda} + 2\sqrt{(3-\lambda)/2}}{3} \right]^{2k}}.$$

The first follows from $\text{TV} \leq \frac{1}{2} \sqrt{\chi^2}$ and the binomial likelihood second moment; the second follows by Cauchy–Schwarz from the product Bhattacharyya affinity. In particular, the one-state comparator meets TV tolerance ϵ whenever

$$k \leq \frac{\log(1 + 4\epsilon^2)}{\log(1 + (1-\lambda)^2/2)}.$$

This is a sufficient finite-horizon resource statement, not an exact minimum-memory frontier. The exact CDF expression gives the sharper decision for any specified (N, λ, ϵ) .

What this changes. The completion theorem established that nine classical hidden states remain necessary for every exact non-iid trine law. The present result shows quantitatively when a one-state source already approximates its complete finite passive record. Thus an exact quantum/classical dimension separation may coexist with a memoryless approximation at the nominated tolerance. Resolving the exact qutrit question does not remove this operational qualification.

A common fixed downstream reader, including a reader with an explicitly retained erasure symbol, contracts this discrepancy. For a source-independent erasure probability $1 - \eta$ on each emitted label, retaining the original un-erased word as well makes the joint discrepancy exactly the original TV: the channel $y \mapsto (y, \text{erased record})$ preserves it. Keeping only the erased report can reduce it. These statements compare the same admitted classical records on both sides; neither grants access to a microscopic bath or to a quantum source’s hidden memory.

Verification. The compact release supplies `supplement/memory/checks/verify_consolidated_constants.py`, which independently compares ordered-word enumeration and the binomial TV formula in twenty exact rational cases. The analytical proof above supplies the all-horizon statement. The historical diagnostic `check_finite.py`, preserved in the separate evidence archive, enumerated all words of lengths two through eight at six rational noise parameters (42 cases), and evaluated additional binomial CDF and retained-erasure examples. Those archived diagnostics are not advertised as extra checks run by the compact release’s reproduction command. These are mathematical calculations, not empirical data.

3 Finite observable accuracy and an explicit rank test

The exact source separations have an accessible finite-accuracy consequence against qubit memories and classical sources with at most eight causal states. This is a known spectral-rank technique newly specialized to the delayed law; it does not improve the difficult three-versus-four quantum boundary. It supplies a controlled, finite-data interface for the established quantum and classical memory claims.

Proposition 3.1 (Observable approximation floor). *Let H be a normalized joint distribution of a past and a future block, viewed as a matrix. Let $\gamma_c(H) = (\sum_{i>c} \sigma_i(H)^2)^{1/2}$. For any complete competing law P' factoring through a causal cut with matrix rank at most c ,*

$$\text{TV}(P', P) \geq \gamma_c(H)/\sqrt{2}.$$

A classical c -state source has this rank bound; a quantum memory of Hilbert dimension d has rank at most d^2 . Every future-accessed, history-correlated system belongs to that cut. The statement allows nonstationary competitors and extra null symbols.

Proof. The classical joint table is a sum of at most c past/future product matrices. Quantum subnormalized conditional states and future effects pair in the d^2 -dimensional real Hermitian space. Pad the target with zeros for extra outcomes, preserving its nonzero singular values. The matrix difference has zero total sum; its positive and negative entries each sum to $\delta = \text{TV}(P', P)$. Consequently its squared Frobenius norm is at most $2\delta^2$. The least Frobenius error of any rank- c approximant is $\gamma_c(H)$, proving the assertion. The latter identity follows by truncating the singular-value decomposition for attainment; for the lower bound, project onto the orthogonal complement of the approximant's column space and minimize the discarded quadratic mass. $\square$

For the four-record trine cut, put $r = 1 - \lambda$. The singular values of $H = T \otimes T/9$ are

$$1/9, \quad r/18 \text{ (four times)}, \quad r^2/36 \text{ (four times)}.$$

Every qubit-cut competitor therefore obeys

$$\text{TV}(P', P_{\lambda,4}) \geq \frac{r\sqrt{1+r^2}}{18\sqrt{2}},$$

while any at-most-eight-state classical competitor obeys

$$\text{TV}(P', P_{\lambda,4}) \geq \frac{r^2}{36\sqrt{2}}.$$

These are outer bounds, not attained optimal approximation frontiers. They vanish at the iid endpoint, as required. They do not exclude a qutrit, because its ordinary rank ceiling is already nine.

3.1 Independent four-record trials and finite confidence

Assume N independently prepared trials from one fixed complete four-record law, with every trial and null retained. The preparation/reset mechanism, clock and independence are physical conditions, not inferred from a reset flag. Pool all nontarget past pairs into one extra row and all nontarget future pairs into one extra column; this coarse-graining preserves the rank bound and gives a finite complete table. Let $\hat{H}$ be its empirical frequency matrix and

$$e_N(\alpha) = \frac{1 + \sqrt{\log(1/\alpha)}}{\sqrt{N}}.$$

Then

$$\Pr\{\|\hat{H} - H\|_F > e_N(\alpha)\} \leq \alpha.$$

Indeed $\mathbb{E}\|\hat{H} - H\|_F^2 = (1 - \|H\|_F^2)/N \leq 1/N$; replacing one trial changes the norm by at most $\sqrt{2}/N$. Bounded differences gives $\Pr\{f - \mathbb{E}f > u\} \leq e^{-Nu^2}$, and Jensen gives the mean bound. The concentration step can be derived by exposing one independent trial at a time and bounding each conditional centered moment-generating function by the exponential of one eighth its squared range; optimizing its exponential Markov bound yields the displayed inequality.

Distance to the rank- c matrix set is one-Lipschitz, so the test

$$\gamma_c(\hat{H}) > e_N(\alpha)$$

has false-rejection probability at most α for every admitted rank- c source. Against the exact target its rejection probability is at least $1 - \beta$ whenever $\gamma_c(H) > e_N(\alpha) + e_N(\beta)$. This supplies a stated sufficient trial budget. It is conservative and is not claimed sample-optimal. It certifies an observable rank obstruction, not a particular microscopic channel or the full infinite word law.

Spectral identification, singular-value-tail thresholding and finite confidence are established techniques. Du Roy de Chaumaray et al., *JMLR* 25(415), 1–37 (2024), Eqs. (5)–(6), Theorem 4, uses the same tail statistic for nonparametric hidden Markov order estimation under a different, mixing single-sequence sampling model. The present contribution is the exact trine specialization, joint classical/quantum cut interpretation and complete-null four-record interface. Original least-rank approximation is Eckart–Young (1936). Hsu–Kakade–Zhang (2012), Lemma 8, already gives the same empirical probability-vector/Frobenius concentration method for independently sampled observation triples. No first spectral memory test is claimed. Neither this test nor its target constitutes an experimental distinction between Shadow Theory and ordinary quantum theory.

At the new noise endpoint, the calculated qubit-cut TV floor is approximately 0.00961570 and the eight-state classical floor is 0.00111370. For one-percent false-rejection risk and at least 99-percent power against the exact target, the sufficient independent four-record trial budgets are 214,080 and 15,958,893, respectively. Each trial costs four emitted records plus its independently prepared stock. These conservative sufficiency budgets are not practical hardware estimates or sample-optimal claims. The supplied synthetic million-trial table rejects the qubit rank cap but does not reject the eight-state cap at the stated threshold, as expected.

4 A robust dimension witness for arbitrary four-record approximations

The low-noise support argument excludes qutrit memories on a small exact-noise interval. That does not by itself quantify approximation by an arbitrary qutrit process whose perturbed law need not remain in that one-parameter family. The same support argument gives an observable matrix witness and a nonzero finite-record approximation floor. The constants are deliberately conservative; they are not proposed experimental accuracy targets.

Matrix witness. Let $C \geq 0$ be a 9-by-9 column-stochastic matrix, indexed by pairs $(a, b), (c, d) \in \{0, 1, 2\}^2$. Define

$$\mu = \min_{a,b} \sum_{c,d} C_{ab,cd} > 0, \quad L = \max_{a,b} \left[\sum_d C_{ab,ad} + \sum_c C_{ab,cb} \right],$$

and let $\sigma_2, \sigma_5 > 0$ be its second and fifth singular values. If C has a complex Hermitian PSD factorization of size at most three, then

$$\boxed{\frac{27L}{\mu} \geq \frac{\sigma_2^2 \sigma_5^2}{36(\sigma_2 + 18\sqrt{6})^2}}. \quad (24)$$

The indices, leakage and singular values are ordinary classical matrix data. No particular quantum update architecture or reference frame is assumed in this lower bound.

Proof. Since $\sigma_5 > 0$, a factorization cannot be supported on dimension two. Congruence normalization of its row factors gives a POVM $\sum_{a,b} A_{ab} = I_3$ and density matrices b_{cd} with $C_{ab,cd} = \text{Tr}(A_{ab} b_{cd})$. The row sums imply $\text{Tr} A_{ab} \geq \mu/9$, hence a largest eigenvalue at least $\mu/27$. Put $R_a = \sum_d b_{ad}$ and $Q_b = \sum_c b_{cb}$. Both have trace three, and the leakage implies

$$\lambda_{\min}(R_a + Q_b) \leq e := 27L/\mu.$$

We prove the required approximate plane/line dichotomy directly. Fix any $s > e$; all distances below are on the same normalized factorization.

Case 1: some R_a or Q_b has second eigenvalue at least s . By interchanging the two indices if necessary, take such an R_a . Let P project onto its two largest eigenvectors and let z be a unit vector in the complementary line. For each b , choose a unit n_b with $\langle n_b, (R_a + Q_b)n_b \rangle \leq e$. Positivity and $R_a \succeq sP$ imply

$$\|Pn_b\|^2 \leq e/s, \quad \langle n_b, Q_b n_b \rangle \leq e.$$

Choose the phase so $\langle z, n_b \rangle \geq 0$. Since $e/s < 1$,

$$\|n_b - z\|^2 = 2(1 - \sqrt{1 - \|Pn_b\|^2}) \leq 2e/s.$$

The bound $\|Q_b\| \leq \text{Tr } Q_b = 3$ and the triangle inequality after applying $\sqrt{Q_b}$ give

$$\langle z, Q_b z \rangle \leq \left(\sqrt{e} + \sqrt{3}\sqrt{2e/s} \right)^2 =: \ell = e(1 + \sqrt{6/s})^2.$$

Every preparation obeys $b_{cd} \preceq Q_d$, hence $\text{Tr}[(I - P)b_{cd}] \leq \ell$ for this common plane. The projection bound $\|\rho - P\rho P\|_1 \leq 2\sqrt{\text{Tr}[(I - P)\rho]}$ therefore bounds every preparation-compression error by $2\sqrt{\ell}$. For completeness, purify ρ , apply $P \otimes I$, and compute the trace norm of the difference between the original rank-one projector and the projected, unnormalized rank-one operator: for discarded mass x , its norm is $\sqrt{4x - 3x^2} \leq 2\sqrt{x}$. Partial trace contracts this Hermitian trace norm, giving the stated bound also for mixed ρ .

Replace each b_{cd} by $Pb_{cd}P$ and apply the same POVM. The resulting real matrix $C^{(4)}$ has rank at most four, since Hermitian operators on a two-dimensional plane have real dimension four. A POVM contracts the trace norm of a Hermitian difference: if $D = D_+ - D_-$ is its positive/negative decomposition, then $\sum_{ab} |\text{Tr}(A_{ab}D)| \leq \text{Tr } D_+ + \text{Tr } D_- = \|D\|_1$. Consequently each of the nine column errors has Euclidean norm at most $2\sqrt{\ell}$. Thus

$$\sigma_5 \leq \|C - C^{(4)}\|_{\text{op}} \leq \|C - C^{(4)}\|_{\text{F}} \leq 6\sqrt{e}(1 + \sqrt{6/s}).$$

The first inequality follows from the singular-value lower bound for rank-four approximation, which applies to every real approximant.

Case 2: all R_a, Q_b have second eigenvalue below s . Let P_a and Q'_b be their top-eigenvector projectors. The mass outside each such line is below $2s$. Since $b_{ab} \preceq R_a, Q_b$,

$$D(b_{ab}, P_a) \leq \sqrt{2s}, \quad D(b_{ab}, Q'_b) \leq \sqrt{2s},$$

where D denotes half trace distance. Here $D(\rho, |v\rangle\langle v|) \leq \sqrt{1 - \langle v, \rho v \rangle}$: it is equality for pure states and follows for mixed states by convexity and Jensen's inequality. Fix P_0 . Comparing through b_{0b} gives $D(P_0, Q'_b) \leq 2\sqrt{2s}$, whence $D(b_{ab}, P_0) \leq 3\sqrt{2s}$ for all nine preparations. Replace every preparation by this same pure state. The resulting matrix $C^{(1)}$ has identical normalized columns and rank at most one. Each column difference is zero-sum and has classical total variation at most $3\sqrt{2s}$ by POVM contraction. For a zero-sum real vector v , its positive and negative parts each have mass $\|v\|_1/2$, so $\|v\|_2^2 \leq 2(\|v\|_1/2)^2$. Each column thus has Euclidean norm at most $6\sqrt{s}$, and

$$\sigma_2 \leq \|C - C^{(1)}\|_{\text{op}} \leq \|C - C^{(1)}\|_{\text{F}} \leq 18\sqrt{s}.$$

These two cases exhaust the possible normalized factors.

If the right side of (24) exceeded e , choose $s < \sigma_2^2/324$ sufficiently close to that upper endpoint. Then both displayed inequalities would fail, since the limiting first bound is $6\sqrt{e}(1 + 18\sqrt{6}/\sigma_2) < \sigma_5$. This choice also satisfies $s > e$: column stochasticity gives $\sigma_2, \sigma_5 \leq \|C\|_{\text{F}} \leq 3$, and the assumed bound on e is less than $\sigma_2^2/(24 \cdot 324)$. This contradiction proves the necessary inequality. $\square$

4.1 A floor for every qutrit approximation of the trine record law

Let $P_{\lambda,4}$ be the specified four-symbol delayed trine law, with

$$H_{ab,cd} = P_{\lambda,4}(a, b, c, d) = \frac{T_{ac}^\lambda T_{bd}^\lambda}{9}, \quad C_\lambda = 9H.$$

Suppose another complete four-record law P' is generated through a quantum causal cut of dimension at most three. All history-correlated systems used again in the future belong to that cut; independent fresh ports and an untouched retained archive obey the same convention as the accounting chapter. The competing law may have a different update rule, lack stationarity, or lie outside the trine family. Extra null/loss outcomes can remain in its complete alphabet, with zero target probability.

Put $\delta = \text{TV}(P', P_{\lambda,4}) < 1/9$. Restrict its joint table H' to the target trit pairs and let $h'_{cd} = \sum_{ab} H'_{ab,cd} \geq 1/9 - \delta$. The column-normalized matrix $C'_{ab,cd} = H'_{ab,cd}/h'_{cd}$ still has PSD rank at most three. Define

$$d = \frac{36\delta}{1 - 9\delta}, \quad a = \frac{1 - \lambda}{2} - d, \quad e_* = \frac{27}{1 - d} \frac{2\lambda/3 + 18\delta}{1 - 9\delta}.$$

Whenever $d < 1$, $a > 0$, the following necessary condition holds:

$$e_* \geq \frac{a^4}{36(a + 18\sqrt{6})^2}. \quad (25)$$

Thus failure of this scalar inequality excludes *every* qutrit-cut approximation at that tolerance, rather than merely a fitted trine member.

Derivation. Column normalization and total variation give

$$\|C' - C_\lambda\|_{\text{op}} \leq \sum_{ab,cd} |C'_{ab,cd} - (C_\lambda)_{ab,cd}| \leq d.$$

Indeed the two contributions are at most the joint-table ℓ^1 error 2δ and the marginal ℓ^1 error 2δ , divided by $1/9 - \delta$. Thus row sums are at least $1 - d$, and both relevant singular values are at least a . The unnormalized leakage functional takes values in $[0, 2]$ and has target expectation $2\lambda/27$. Its change is at most 2δ ; dividing by the smallest column mass bounds $L(C')$ by $(2\lambda/3 + 18\delta)/(1 - 9\delta)$. Apply (24); its right side increases with each positive singular value. This gives (25).

For explicit simple consequences, exact rational constants in the proof with $s = 1/2000$ give

$$\text{TV}(P', P_{0,4}) > 10^{-9}, \quad (26)$$

and, uniformly for $0 \leq \lambda \leq 10^{-8}$,

$$\text{TV}(P', P_{\lambda,4}) > 5 \times 10^{-10}. \quad (27)$$

At the first boundary the squared rank-four approximation error is below 0.216, and at the second below 0.188; the rank-one squared error is 0.162, while the required singular value exceeds 0.49 in both cases. These loose bounds are verified with fractions in the accompanying code.

What changes. The earlier exact dimension statement was robust only along a specified noise curve. The new witness supplies a finite-output accuracy obstruction against all ordinary qutrit-cut realizations, including time-dependent alternatives and complete null/loss records. It is a quantitative outer bound, not a sharp approximate-memory frontier, a large practical separation, or a resource theorem for nonlinear ontic assignments. Larger horizons inherit the same floor by retaining the four-record marginal. It compares the declared classical record law; it identifies no microscopic source or conscious perspective.

4.2 An explicit robust version of the support witness

The noiseless growing lower bound is sensitive to small forbidden-event probabilities. A quantitative replacement is possible, with a poor large-grid margin that should be stated explicitly. Suppose a row-stochastic matrix $\tilde{T}$ factors through r classical states as AB , with both A and B row stochastic. Such a normalization is available for every nonnegative factorization of a stochastic matrix after zero factors are removed. Set $e = \max_i \tilde{T}_{ii}$ and $\mu = \min_{i \neq j} \tilde{T}_{ij}$. For $\tau > 0$ put $S_i = \{\ell : A_{i\ell} \geq \tau\}$. If $S_i \subseteq S_j$, then

$$\tilde{T}_{ij} \leq e/\tau + r\tau. \quad (28)$$

Indeed, on S_i the sum of $B_{\ell j}$ is at most e/τ because $A_{j\ell} \geq \tau$, and outside S_i the contribution is at most $r\tau$. Choosing $\tau = \sqrt{e/r}$ shows that $\mu > 2\sqrt{re}$ forces an antichain; the case $e = 0$ follows by a limit or the exact support argument.

For an entrywise perturbation of the phase-circle transition by at most δ , write

$$\mu_m = \frac{1 - \cos(2\pi/m)}{m} = \frac{2 \sin^2(\pi/m)}{m}.$$

The sufficient antichain condition is $\mu_m - \delta > 2\sqrt{r\delta}$, or equivalently

$$\delta < \delta_{m,r}^* := \frac{\mu_m^2}{(\sqrt{r} + \mu_m + \sqrt{r})^2}. \quad (29)$$

If a normalized observed triple distribution p' lies at TV distance η from the exact law, with $\eta < m^{-2}$, its conditional transition in any middle-symbol fibre obeys

$$\max_{i,j} \left| \frac{p'(i, y, j)}{\sum_k p'(i, y, k)} - T_{ij} \right| \leq \frac{2\eta}{m^{-2} - \eta}.$$

Both the single-entry and row-mass probability errors are at most η ; the inequality follows by adding numerator and denominator errors. Therefore, choosing $r = r_m - 1$, the original classical sector lower bound mr_m survives whenever

$$\eta < \frac{\delta_{m,r}^*}{m^2(2 + \delta_{m,r}^*)}. \quad (30)$$

This is a rigorous finite-data statement, but its radius scales as roughly m^{-8}/r and is extremely small at $m = 180, 360, 720$. It is not proposed as a practical large-grid experiment. The trine rank margin is substantially more usable.

5 Sharp bath approximation costs with and without retained purification

Exact Choi rank counts even arbitrarily weak Kraus directions. It does not by itself say what apparatus savings become possible at a prescribed error. For the three-charge qutrit instruments there is a constructive, sharp answer at three useful budgets. The answer differs depending on whether the purifying bath is unread or is part of the compared retained output. This is an application of established spectral approximation and Stinespring methods, not a new general theory of approximate quantum channels. The standard maximum-trace eigenvalue principle and minimum-pure-Stinespring dimension are antecedents [2, 12]; the proofs needed here are supplied below. The added result is exact saturation for the declared charge-orbit channels, with separate retained and unread-bath tasks.

5.1 The typed channel and its three trace-scaled components

Let $S = \mathbb{C}^3$ be the reusable memory, $A = \mathbb{C}^3$ the registered pointer, and B an initially pure, fresh purifying bath. The reduced output channel is

$$\mathcal{A}(X) = \sum_{j,k=0}^2 |j\rangle\langle j|_A \otimes K_{jk} X K_{jk}^\dagger.$$

Only A and the final memory belong to this reduced output; an arbitrary inert reference is retained throughout. In the retained-purification task we instead compare the full isometries into $A \otimes S \otimes B$. The two metrics must not be substituted for one another.

Write $\omega = e^{2\pi i/3}$, $Z = \text{diag}(1, \omega, \omega^2)$ and $X|p\rangle = |p+1\rangle$. The rank-one charge blocks give Kraus matrices of the form

$$K_{jk} = X^j K_{0k} Z^{-j}, \quad |K_{0k}\rangle\rangle = \alpha_k |f_k\rangle|1\rangle + \beta_k |f_{k+1}\rangle|2\rangle, \quad (31)$$

where $|f_k\rangle = 3^{-1/2} \sum_p \omega^{pk} |p\rangle$ and the vectorization orders input before output. Put $\nu_k = |\alpha_k|^2 + |\beta_k|^2 > 0$, and assume $\sum_k \nu_k = 1$. The three vectors at fixed j are orthogonal. Moreover

$$\sum_j K_{jk}^\dagger K_{jk} = \nu_k I_3. \quad (32)$$

Indeed $K_{0k}^\dagger K_{0k}$ is a weighted sum of two conjugate Fourier projectors, so each computational diagonal entry is $\nu_k/3$. Twirling it by Z kills the off-diagonal entries and gives (32). Consequently

$$\mathcal{Q}_k(X) = \nu_k^{-1} \sum_j |j\rangle\langle j| \otimes K_{jk} X K_{jk}^\dagger$$

is CPTP with Choi rank three, and $\mathcal{A} = \sum_k \nu_k \mathcal{Q}_k$. The three component Choi supports are mutually orthogonal, including the pointer factor. The normalized Choi state $\rho_{\mathcal{A}} = J(\mathcal{A})/3$ therefore has eigenvalues $\nu_k/3$, each with multiplicity three.

5.2 Exact approximation theorem

Relabel the weights $\nu_{(1)} \geq \nu_{(2)} \geq \nu_{(3)}$ and put $p_m = \sum_{i=1}^m \nu_{(i)}$, $m = 1, 2, 3$. A pure bath of dimension b means the minimum pure Stinespring environment of the specified reduced channel; all fresh apparatus factors purifying that channel are included.

Theorem. Among all CPTP channels with the same input and reduced output spaces and pure bath dimension at most $3m$,

$$\inf \frac{1}{2} \|\tilde{\mathcal{A}} - \mathcal{A}\|_\diamond = 1 - p_m. \quad (33)$$

The infimum is attained by the actual channel

$$\mathcal{A}_m = p_m^{-1} \sum_{i=1}^m \nu_{(i)} \mathcal{Q}_{(i)}. \quad (34)$$

It retains the same pointer alphabet and needs no added null branch. The normalization defines a new fixed trace-preserving instrument; it is not postselection on original trials.

For the full retained pure Stinespring output, allow any bath-only isometric identification of the competitor's at-most- $3m$ dimensional bath with the target bath. Then the corresponding optimum is

$$\inf \frac{1}{2} \|\mathcal{W} - \mathcal{V}\|_\diamond = \sqrt{1 - p_m}, \quad (35)$$

where $\mathcal{V}(\rho) = V\rho V^\dagger$ and similarly for W . The dimension counts are intrinsic bath dimensions; the isometric embedding only defines a common comparison space.

Proof of the reduced statement. For a channel of bath dimension at most b , its Choi state has rank at most b . If P projects onto its support, the variational characterization of trace distance gives

$$\frac{1}{2}\|\rho_{\mathcal{A}} - \rho_{\tilde{\mathcal{A}}}\|_1 \geq \text{Tr}[(I - P)\rho_{\mathcal{A}}] \geq 1 - \sum_{i=1}^b \lambda_i(\rho_{\mathcal{A}}).$$

The last step follows by diagonalizing $\rho_{\mathcal{A}}$: for any rank- b projector its diagonal weights lie in $[0, 1]$ and sum to b , so the maximum weighted eigenvalue sum selects the largest b eigenvalues. Evaluating the diamond norm on a maximally entangled input gives the same lower bound for half diamond distance. For $b = 3m$ this is $1 - p_m$.

For $m = 3$ attainment is exact and immediate. Otherwise both $\mathcal{A}_m$ and the normalized omitted component $\mathcal{R}_m$ are CPTP, and $\mathcal{A} = p_m \mathcal{A}_m + (1 - p_m) \mathcal{R}_m$. The half diamond distance of any two channels is at most one, hence $\frac{1}{2}\|\mathcal{A} - \mathcal{A}_m\|_\diamond \leq 1 - p_m$. The preceding lower bound is equal, proving (33). Orthogonal Choi supports explain why no slack remains.

Proof of the retained statement. Use a bath basis $|j, k\rangle$ and define

$$V = \sum_{j,k} |j\rangle_A \otimes K_{jk} \otimes |j, k\rangle_B, \quad W_m = p_m^{-1/2} \sum_{j,k \text{ retained}} |j\rangle_A \otimes K_{jk} \otimes |j, k\rangle_B.$$

Both are isometries by (32), and

$$V^\dagger W_m = \sqrt{p_m} I_3. \quad (36)$$

Every pure input with any inert reference therefore gives pure outputs of overlap $\sqrt{p_m}$, whose half trace distance is $\sqrt{1 - p_m}$. Convexity handles mixed inputs. This proves the upper bound uniformly over references.

For any competing bath of dimension at most b , apply its isometry and the target isometry to a maximally entangled input. The competitor output has Schmidt rank at most b across bath versus pointer-memory-reference. After any bath-only identification it lies in some bath subspace of dimension at most b . Its squared overlap with the target is at most the target mass in that subspace, at most the largest- b eigenvalue sum of its bath marginal. The target pure state's nonzero bath eigenvalues are those of $\rho_{\mathcal{A}}$, so this sum is p_m for $b = 3m$. Pure-state trace distance is thus at least $\sqrt{1 - p_m}$. This proves (35). $\square$

5.3 An exact, independently checked qutrit example

At the fixed-seed construction's critical parameter $t = \sqrt{14}/21$, $r = 4\sqrt{14}/63$, the three positive charge blocks have traces

$$(\nu_0, \nu_1, \nu_2) = (1/9, 1/3, 5/9).$$

For example their first diagonal entries are $(2/27, 2/27, 14/27)$ and off-diagonal entries are $c_0 = t(\omega - 2)/9$, $c_1 = c_2 = -t(8 + 5\omega)/9$. Their determinants vanish; each is nonzero and has rank one. Thus the following are sharp resource points for this specified channel:

pure bath budget	optimal reduced half diamond error	optimal retained half diamond error
3	4/9	2/3
6	1/9	1/3
9	0	0

No exact answer is claimed at the intervening budgets four, five, seven or eight. The same theorem applies to any other certified member with the charge form (31), using its actual block weights; it does not require the particular numerical example.

Repeated use. For the displayed streaming implementations, fresh independent pure bath ports and the same internal dynamics give $V_n^\dagger W_{m,n} = p_m^{n/2} I$: contract the last pair of isometries using (36) and induct. Their full retained distance is therefore exactly $\sqrt{1 - p_m^n}$ for every input and reference. This is the error of these explicitly declared implementations, not the optimum over globally compressed horizon baths or a grant of fresh preparation independence in an unrelated physical source.

Practical meaning and limits. The earlier count nine was an exact one-step apparatus cost. It now has a precision-sensitive interpretation: for this channel a bath of size six cannot achieve reduced error below $1/9$, and cannot achieve retained-pure error below $1/3$; both boundaries are constructively attained. The comparison keeps the same source space, original preparation, pointer, clock tick and arbitrary inert reference. It compares alternative admitted interactions. Access to a returning bath or controller changes the reusable-memory ledger. An initially mixed apparatus requires its purification resource to be declared; the pure-bath budget here cannot silently be replaced by its unpurified marginal dimension.

This does not optimize over all qutrit sources with the same passive word law, compare a three-dimensional source with a four-dimensional one under arbitrary interventions, or certify a laboratory realization. It does not turn an unread-bath bound into equality of complete retained instruments. The two separate, explicit statements above provide the appropriate interfaces. The construction supplies no Born selection, mass-noise admission, actual-event mechanism or conscious perspective.

`supplement/memory/five/audit/check_bath_approximation.py` independently reconstructs all nine Kraus matrices, proves their orbit trace identities and Choi orthogonality with exact symbolic arithmetic, and checks the scalar isometry overlaps. Its arbitrary-reference samples are numerical diagnostics, not the universal proof. All parameters and outputs are in `supplement/memory/expected_outputs/five/audit/bath_approximation_results.json`.

6 Finite-horizon resource accounting: what survives beyond the memory box

An exact streaming advantage survives the expanded accounting, but it is an advantage in *reusable causal memory*, not a claim that an entire closed apparatus has four dimensions. For the noiseless trine process and any declared horizon $n \geq 4$, four quantum dimensions and nine classical states are the exact causal-memory minima when every history-correlated controller, reference and environment that can influence the future is included at the causal cut. In contrast, the retained archive and an isometric implementation of its classicalization require resources growing with the horizon. These resources can be calculated exactly for the supplied quantum construction. This corrects a possible overreading of the earlier four-versus-nine result.

The calculation concerns ordinary finite-dimensional instruments and exact channels. It assumes Born probabilities, does not select actual events, and does not assign hardware size, energy, temperature or a physical reset cost to a Hilbert dimension. The preparation, interface and accuracy must be the same for every comparison; the finite-accuracy comparator elsewhere in this companion changes the appropriate conclusion when an exact law is unnecessary.

6.1 Declared task and phase encoding

For the emitted-symbol construction, use phases $\theta_j = 2\pi j/m$, $m \geq 3$,

$$|\psi_j\rangle = \frac{|0\rangle + e^{i\theta_j}|1\rangle}{\sqrt{2}}, \quad |\psi_j^\perp\rangle = \frac{|0\rangle - e^{i\theta_j}|1\rangle}{\sqrt{2}}, \quad E_j = \frac{2}{m} |\psi_j^\perp\rangle \langle \psi_j^\perp|,$$

with $\rho_j = |\psi_j\rangle\langle\psi_j|$ and

$$T_{ij} = \text{Tr}(E_j \rho_i) = \frac{1 - \cos(\theta_i - \theta_j)}{m}.$$

Starting from $\tau = I_4/4$, the iterated instrument

$$\Psi_j(X) = \text{Tr}_1[(E_j \otimes I_2)X] \otimes \rho_j \quad (37)$$

has the full stationary word law

$$p(w_1 \cdots w_n) = \frac{1}{m^2} \prod_{t=1}^{n-2} T_{w_t, w_{t+2}}, \quad n \geq 2. \quad (38)$$

The circular phase representation is informed by [6, abstract and Sections 2–3] and [7, pp. 64–68]; the stochastic process and instrument are defined above, and their resource costs are proved below.

The task here retains all n outputs in their original order, their settings and timing metadata, and every declared null/loss outcome. The ideal alphabet contains no loss; a failed physical implementation must add it and changes the target law. The archive is append-only: it is not consulted to decide future source transitions. Future fresh stock is initially uncorrelated with the history and memory. Any violation must be explicitly added to the causal memory or to the task’s intervention interface.

There are two distinct output tasks. A *coherent archive* retains the amplitudes between different output words. A *classical archive* is diagonal in the word basis, including when the input memory is entangled with an arbitrary retained reference. The passive word probabilities agree, but the channels and their resource requirements differ.

6.2 An exact single-step dilation

Let S swap the two memory qubits and set

$$L_j = \sqrt{2/m} |\psi_j\rangle\langle\psi_j^\perp|, \quad K_j = S(L_j \otimes I_2).$$

Then $\Psi_j(X) = K_j X K_j^\dagger$ and $\sum_j K_j^\dagger K_j = I_4$. Thus

$$V = \sum_{j=0}^{m-1} |j\rangle_A \otimes K_j : \mathbb{C}^4 \longrightarrow \mathbb{C}_A^m \otimes \mathbb{C}_M^4 \quad (39)$$

is an isometry. A fresh blank m -level output register and a fixed unitary extension implement the coherent-output step with no additional discarded system. This does *not* say that the output register is free.

For comparison the dephased complete-record channel is

$$\mathcal{C}_1(X) = \sum_j |j\rangle\langle j|_A \otimes K_j X K_j^\dagger.$$

Its Choi rank is exactly m : each nonzero output block has rank one. Consequently a globally isometric implementation with a pure initially independent environment requires an environment of dimension at least m , and dimension m suffices by copying the label,

$$F : |j\rangle_A \mapsto |j\rangle_A |j\rangle_E.$$

The environment must be retained elsewhere or traced; tracing implements classicalization. Conditioning on one actual outcome is a different operation from replacing this complete channel by a pure branch.

The memory channel with its record discarded, $\Lambda(X) = \sum_j K_j X K_j^\dagger$, has Choi rank exactly three, independent of m . Indeed the three functions $1, e^{i\theta_j}, e^{-i\theta_j}$ are linearly independent for $m \geq 3$, and

$$L_j = \frac{1}{\sqrt{2m}} \begin{pmatrix} 1 & -e^{-i\theta_j} \\ e^{i\theta_j} & -1 \end{pmatrix}$$

span a three-dimensional operator space. Multiplication by S and tensoring with I_2 preserve this span dimension. Dropping the record can therefore reduce a dilation cost from m to three, but solves a weaker task.

These minimal-environment statements are standard Choi–Kraus–Stinespring theory, applied here explicitly. If a pure-environment dilation has dimension d_E , its environment-basis Kraus representation has at most d_E terms, so $\text{rank } J(\mathcal{C}) \leq d_E$; a spectral decomposition of the positive Choi matrix attains equality. A primary textbook statement is Watrous, *The Theory of Quantum Information*, Corollary 2.27, printed pp.89–90, author’s chapter PDF pp.16–17: <https://cs.uwaterloo.ca/~watrous/TQI/TQI.double.2.pdf>.

6.3 Exact horizon cost and the obstruction to indefinite reuse

For a word w , let $K_w = K_{w_n} \cdots K_{w_1}$. Define the complete classical-record channel

$$\mathcal{C}_n(X) = \sum_w |w\rangle\langle w|_A \otimes K_w X K_w^\dagger. \quad (40)$$

This is a channel on arbitrary input states and hence includes all initially retained input references. It is stronger than requesting a distribution only for $I_4/4$.

Horizon dilation theorem. For the noiseless phase-circle instrument, $n \geq 2$, put

$$N_n = m^2(m-1)^{n-2}.$$

Then $\mathcal{C}_n$ has Choi rank exactly N_n . Its minimum pure initial environment dimension is N_n . For the fixed stationary input $I_4/4$, the final memory–classical-archive state also has rank N_n . Conversely the coherent n -output isometry requires no additional bath beyond its output registers, while producing a different retained-output channel.

Proof. After the first two steps,

$$K_{(a,b)} = \frac{2}{m} |\psi_a \psi_b\rangle \langle \psi_a^\perp \psi_b^\perp|.$$

Every further step shifts the two states and contributes a scalar whose squared modulus is $T_{w_t, w_{t+2}}$. Therefore K_w has operator rank one and is nonzero precisely when $w_{t+2} \neq w_t$ at every later step. There are m^2 first pairs and $m-1$ later choices, giving N_n positive words. Each word block of the Choi matrix is the nonzero rank-one matrix $|K_w\rangle\langle K_w|$; distinct archive blocks are orthogonal. Their ranks add, proving the Choi claim and the pure-environment minimum. On $I_4/4$ every word block is a positive scalar times $\rho_{w_{n-1}} \otimes \rho_{w_n}$, so the fixed-output state has the same rank. Finally the ordered iteration of V is an isometry before dephasing and introduces no bath. $\square$

The lower bound is sharp for the declared *final complete channel*:

$$X \mapsto \sum_w |w\rangle_A |w\rangle_E \otimes K_w X$$

is its Stinespring isometry. It is not a claim that N_n is the minimum physical bath for every source having only the same passive output law, or that independently optimizing each intermediate cut gives

a minimum protocol. It is a concrete finite-horizon implementation and lower bound for (40). The distinction between a channel's dilation rank and a whole strategy's memory is essential; see Bisio et al., *Phys. Rev. A* 85, 032333 (2012), DOI [10.1103/PhysRevA.85.032333](https://doi.org/10.1103/PhysRevA.85.032333), Definitions 3–4 and Theorem 3, <https://arxiv.org/abs/1112.3853>.

A fixed reusable environment of dimension D , initially pure and with no reset, export or new environment, cannot realize these dephased channels at arbitrarily large n : it must satisfy $D \geq N_n$. For $m > 2$,

$$n \leq 2 + \left\lfloor \frac{\log(D/m^2)}{\log(m-1)} \right\rfloor$$

whenever a horizon $n \geq 2$ is feasible. This is a finite-dimensional isometric obstruction, not a claim about the speed or thermodynamics of erasure. If an initially mixed environment has rank r and output dimension D , a Kraus expansion gives only $Dr \geq N_n$; its purifying stock cannot be silently declared free. If the apparatus exports the purifier or outcome copies, those registers belong in the resource ledger. If a fundamental nonunitary law is introduced, this dilation accounting no longer describes that law and must not be used to rule it out.

The raw n -symbol archive has dimension m^n for either realization. Its positive-word subspace has dimension N_n , allowing an exactly reversible lossless word code of that capacity. A controller implementing such global compression is additional circuitry; it is not produced by the rank bound. With stationary input, the required archive entropy is

$$H(W_1 \cdots W_n) = 2 \log_2 m + (n-2)h_m, \quad h_m = - \sum_j T_{0j} \log_2 T_{0j}.$$

This follows by the chain rule from the order-two law and its uniform pair distribution. For the noiseless trine, $h_3 = 1$ and all $N_n = 9 \cdot 2^{n-2}$ positive words are equiprobable. The purifying environment for the stationary output can have exactly this entropy; entropy in bits is not an energy estimate.

6.4 An explicitly budgeted classical comparator

For $m = 3$, the nine-state orthogonal-pair source achieves the exact classical minimum for this law. A declared quantum-channel embedding of that classical update on $\mathbb{C}^m \otimes \mathbb{C}^m$ is

$$D_j = \text{diag}(T_{0j}, \dots, T_{m-1,j}), \quad \mathcal{Q}_j(X) = \text{Tr}_1[(D_j \otimes I_m)X] \otimes |j\rangle\langle j|.$$

It preserves the same passive word law from I_{m^2}/m^2 . It retains second-register coherence in this extension rather than gratuitously dephasing the entire memory. Its independent Kraus operators are

$$F_{j,a} = \sqrt{T_{aj}} S_m(|j\rangle\langle a| \otimes I_m), \quad a \neq j.$$

Thus its one-step complete-record Choi rank is $m(m-1)$, and for $n \geq 2$ its horizon Choi rank is $(m-1)^2 N_n$. To see the latter, a word map has pure final memory $|w_{n-1}, w_n\rangle$ and effect proportional to $D_{w_1} \otimes D_{w_2}$, of rank $(m-1)^2$; archive ranks again add. These environment counts are exact for these specified CP updates, not a proved optimum over every nine-state classical HMM and every physical extension of its diagonal-state behavior.

For the trine, the following ledger separates resources rather than hiding them in a single memory number. Its dilation rows concern the specified final channel; they must not be multiplied with a per-step memory optimum and called a simultaneously optimized closed protocol. The simple streaming implementation uses independent fresh per-step ports; global compression needs its own controller.

Resource, exact horizon $n \geq 4$	Quantum shift	Canonical classical pair
Reusable causal memory	4 dimensions	9 states
Raw ordered symbol archive	3^n dimensions	3^n dimensions
Positive-word archive capacity	$N_n = 9 \cdot 2^{n-2}$	same N_n
Pure bath for the specified final complete channel	N_n	$4N_n$
Pure bath for one specified recorded step	3	6
Purifier of stationary mixed memory	4	9
Clock, switch routing, calibration	separately supplied	separately supplied

The purifier row refers to each stated stationary mixed-memory preparation; it is not evidence that the two process models have the same response to arbitrary interventions on those different purifiers. Bath ranks compare specific channels, not an optimization of all closed realizations.

For example $n = 16$ needs $N_n = 147456$ positive-word labels (18 binary storage positions suffice for their lossless fixed-length code), despite a reusable quantum memory of dimension four. Repeated independent per-step baths of dimension three give a simple, larger upper bound 3^n ; finite-horizon Kraus compression reduces the final-channel bath requirement to N_n . Neither count is the same as a device's number of elementary components.

6.5 A causal-cut separation that includes correlated controllers

There is a precise sense in which the memory advantage is not an artifact of drawing a small box around the source. At the cut between outputs two and three, define M_{live} to include *all* jointly history-correlated systems that can influence future source operations: memory, feedback controller, reused environment, and any history-bearing phase or clock register. A history-bearing part of a future tape belongs here, even if physically outside the nominal source. Independent fresh stock is a declared preparation resource, while an append-only archive and environments never accessed again are retained-history resources, not reusable causal memory.

Expanded causal-memory proposition. For the exact noiseless trine four-output law, a quantum M_{live} needs dimension at least four, and a classical M_{live} needs at least nine states. Both bounds are attainable while using the explicitly budgeted fresh resources and archives above. They therefore hold for every exact horizon $n \geq 4$ under this interface.

Proof. Future operations, including independent fresh ancillas, define positive effects B_{cd} on M_{live} . The unnormalized state after outputs a, b is $A_{ab} \geq 0$. Hence

$$H_{ab,cd} = p(a, b, c, d) = \text{Tr}(A_{ab}B_{cd}) = \frac{T_{ac}T_{bd}}{9}.$$

For trine T , $\text{rank } T = 3$, so $\text{rank } H = 9$. A classical D -state cut factors H through D nonnegative coordinates, giving $D \geq 9$. A quantum cut has $d^2 \geq 9$, so $d \geq 3$. To exclude $d = 3$, use the zeros whenever $a = c$ or $b = d$. Positive matrices with zero trace pairing have orthogonal ranges. Put $R_a = \text{span}_d \text{ran } B_{ad}$ and $C_b = \text{span}_c \text{ran } B_{cb}$. The nonzero range of A_{ab} is orthogonal to $R_a + C_b$, hence $\dim(R_a + C_b) \leq 2$. If one R_a has dimension two, all C_b lie in it, so every effect is supported on a common plane; compressing all factors would give rank at most four, contradicting nine. The same holds for a two-dimensional C_b . Otherwise all these nonzero spaces are lines. Each nonzero B_{ab} lies in both R_a and C_b , forcing all lines equal and rank at most one, again a contradiction. Thus $d \geq 4$. The two-qubit shift and nine-state pair source attain the respective bounds. $\square$

Certified noisy-interval corollary. The same expanded interface gives exact causal-cut minima

noise interval	quantum dimensions	classical states
$0 \leq \lambda \leq 3 \times 10^{-8}$	4	9
$7657/8000 \leq \lambda < 1$	3	9.

For the first row use the quantitative PSD-rank exclusion proved in the main resource manuscript and the four-dimensional shift. This companion retains the narrower rational-encoding interval in the second row so its stated construction remains self-contained; the main manuscript extends qutrit attainment to the larger interval beginning at $\lambda_{\dagger}$. For the second use the exact qutrit instrument in Section 1; ordinary rank nine of the same four-time cut excludes a quantum dimension below three and a classical size below nine. Fresh independent ancillas and fixed clocks still define effects on M_{live} , so including correlated controllers does not change these factorization arguments. These are exact-law statements for $n \geq 4$; the middle interval and minimum approximate-memory frontier are not asserted.

This is the earlier proved rank/support argument with the causal interface made explicit, not a new lower-bound method. It closes an accounting loophole: moving history into a controller or returning bath does not evade the bound when that system is counted in M_{live} . Allowing a pre-correlated future tape and then excluding it from M_{live} would invalidate the factorization, not disprove the proposition. Allowing coherent archive feedback changes the interface and must include the interacting archive in the cut.

This restricted four-versus-nine separation remains useful even when common archive storage is large. In logarithmic memory capacity it is $\log_2(9/4) \simeq 1.170$ bits. Adding an identical separately budgeted storage term does not remove that difference, but the reusable-memory fraction of the total storage can become small as n grows. It is not an optimization of total closed hardware, and no energy or speed ordering follows.

6.6 Phase, calibration and clock accounting

The instrument is fixed at every tick. It does not call a hidden classical controller that remembers the previous output and prepares a time-dependent phase state; the preparation and update are already included in K_j . Nevertheless the fixed m -outcome coupler, its calibration and output ports are resources. Increasing m does not make those resources disappear merely because the reusable memory remains four-dimensional.

A modest symmetry result limits what phase reference is actually needed in this abstract construction. Let $U = \text{diag}(1, e^{2\pi i/m})$, $R = U \otimes U$, and let $P|j\rangle = |j + 1 \bmod m\rangle$. The chosen phases of L_j give $RK_jR^\dagger = K_{j+1}$, hence

$$VR = (P \otimes R)V. \quad (41)$$

The stationary state is invariant. The isometry therefore respects simultaneous discrete rotation of the memory and relabelling of the ports. In particular, the common azimuth is not a separate time-varying source variable.

There is a symmetry-preserving unitary extension with a fresh output blank $|s\rangle = m^{-1/2} \sum_j |j\rangle$, invariant under P . To prove existence, the invariant subspace $|s\rangle \otimes \mathbb{C}^4$ and the image of V carry equivalent representations of the finite group C_m by (41). Their orthogonal complements have the same representation multiplicities, so choose a unitary intertwiner on each complement and combine it with V on the input subspace. This is a finite-group constructive existence argument, not an assertion of reference-free implementation under continuous $U(1)$ or of a particular energy-conserving Hamiltonian. Physical phase stability and gate precision still require calibration at the declared accuracy.

A clock supplies ticks separated by a chosen Δt in seconds and routes the next fresh output register. Its precision, reset and control law are independent implementation inputs. A finite routing

register with $n + 1$ orthogonal positions is one possible explicitly budgeted schedule; no minimum clock dimension or thermodynamic cost is proved. If that clock acquires history correlations used to modify the source, it joins M_{live} . An unused archival timestamp does not become hidden controller memory.

6.7 Noise, accuracy and what remains established

The exact zero pattern is load bearing. For the earlier noisy construction $E_j^\lambda = [I - (1 - \lambda)\mathbf{n}_j \cdot \boldsymbol{\sigma}]/m$ with $0 < \lambda \leq 1$, each effect has rank two. The *same four-dimensional shift channel*, which need not be the minimum realization of the noisy law, then has one-step recorded Choi rank $2m$ and unrecorded rank four. For $n \geq 2$, every word is positive and the word map has pure final memory but an input effect proportional to $E_{w_1}^\lambda \otimes E_{w_2}^\lambda$, of rank four. Consequently its complete-channel Choi rank is $4m^n$, although its stationary fixed-output state without an initial retained reference has rank only m^n . This is a concrete reason not to replace a retained-reference channel task by one fixed-input density matrix when counting environment.

These ranks are exact channel identities, not robust physical costs at finite tolerance. At $\lambda = 1$ the output law is IID uniform and a one-dimensional source suffices; the four-dimensional shift's dilation ranks are plainly not process-wide minima. The exact and finite-accuracy realization results elsewhere in this companion must be used for the corresponding task and tolerance. No discontinuous exact rank jump is offered as an experimentally discontinuous hardware demand.

The rational encoded-pair qutrit has a larger specified bath. The exact rational certificate in Section 1 establishes that all three 9-by-9 Choi matrices of its qutrit branch maps are *strictly* positive throughout $7657/8000 \leq \lambda < 1$; the Bernstein control matrices are positive definite. This independently determines that construction's dilation ranks, rather than borrowing the four-dimensional shift's formula.

If $J(\Phi) > 0$ and $J(\Theta) > 0$, choose $c > 0$ with $J(\Phi) \geq cI \otimes I$. Complete positivity and strict positivity of $\Theta(I)$ give

$$J(\Theta \circ \Phi) = (\text{id} \otimes \Theta)(J(\Phi)) \geq cI \otimes \Theta(I) > 0.$$

Here $\Theta(I) > 0$ follows, for example, by partially tracing a strictly positive Choi matrix. Induction proves that every word branch of the specified qutrit source has Choi rank nine. There are 3^n positive words; orthogonal classical archive blocks therefore give exact recorded-channel Choi rank $9 \cdot 3^n$. On stationary input $I_3/3$, every conditional word state is strictly positive, giving fixed-output rank $3 \cdot 3^n$ instead.

For these specified implementations on the rational encoded-pair interval $7657/8000 \leq \lambda < 1$, the comparable ledger is

Same passive trine word law	Qutrit	Shift	Classical pair
Reusable memory dimension	3	4	9
One-step complete-record Choi rank	27	6	9
Horizon complete-record Choi rank, $n \geq 2$	$9 \cdot 3^n$	$4 \cdot 3^n$	$9 \cdot 3^n$
Stationary fixed-output rank, reference omitted	$3 \cdot 3^n$	3^n	3^n
Raw classical symbol archive	3^n	3^n	3^n

The classical column uses the explicit orthogonal-pair CP extension above; for positive noise its effects D_j have rank three and its word effects $D_{w_1} \otimes D_{w_2}$ have rank nine. It is again a specified implementation, not a universal optimum over classical physical extensions. The source with one-step rank nine and the mixed-support history source in the main manuscript improve on this rational construction. This table preserves the distinct, exactly specified channels above; it is not the current process-wide frontier.

The lower memory dimension and higher minimum pure bath for these particular channels form a real resource tradeoff. The two complete memory/bath channels are different; only the declared passive symbol task is the same. This is not a lower bound on the bath of every possible qutrit realization, nor a proof that the four-dimensional source minimizes total hardware. It is precisely why one should report a resource vector before choosing a cost function. An implementation allowed only finite output accuracy can use a different comparator, as quantified in Section 2.

The completed benefit is an explicit audit of where the earlier compression does and does not apply: a sharp exact reusable-causal-memory separation, exact finite-horizon bath requirements for its complete channel, a proof that finite unreset environments cannot produce indefinitely many exact classical records in the isometric model, and a symmetry-respecting fixed implementation. No statement here closes the measurement programme's outstanding physical obligations or identifies a conscious perspective.

Verification. The historical diagnostic `research3/accounting/verify_accounting.py`, preserved in the separate evidence archive, checks Kraus completeness, Choi ranks, word probabilities, conditional states and discrete covariance. It tests noiseless grids $m = 3, 4, 5$ through horizons 7,5,5 respectively, and the positive trine shifts at $\lambda = 0.2, 1$ through horizon four. The canonical classical channels are tested separately. Numerical compositions of the qutrit certificate at $\epsilon = 0, 0.2, 0.35$ through horizon four independently check strict Choi ranks and the new tradeoff. All 7,198 checks pass; the maximum completeness or state/probability error is 1.26×10^{-14} . Exact proofs provide the universal claims; floating-point tests diagnose implementation mistakes. These archived diagnostic counts are not additional checks executed by the compact release's nine-job memory reproduction command.

6.8 A conditional normal form for a mixed-support equality profile

The following reduces a candidate equality branch to a finite spectral form. It is not an existence proof for the mixed support profiles.

Lemma. Let $\Phi : M_2 \rightarrow M_3$ be CP, trace preserving, of linear rank three, with faithful $\Phi(I_2)$. If two distinct pure input rays have singular nonzero outputs, then after an invertible output congruence its image is the full diagonal algebra. There are a density matrix $\rho \succ 0$ and three scalars λ_i , at least two distinct and of unit modulus, with $|\lambda_i| \leq 1$, such that in a suitable input basis

$$\Phi \begin{pmatrix} x & y \\ z & w \end{pmatrix} = \sqrt{\rho} U \text{diag}(x + w + \lambda_i y + \bar{\lambda}_i z)_{i=1}^3 U^\dagger \sqrt{\rho}.$$

Its Choi rank is three or four, according as the remaining scalar has unit modulus or strict modulus below one.

Proof. The Hermitian kernel is one-dimensional and contains a nonzero traceless matrix because the map is trace preserving. Diagonalize it. The two diagonal input projectors then have the same output ρ , which is faithful. Writing the off-diagonal response as C gives $\Phi(X) = (x + w)\rho + yC + zC^\dagger$. Choi positivity is equivalent to

$$\begin{pmatrix} I & K \\ K^\dagger & I \end{pmatrix} \succeq 0, \quad K = \rho^{-1/2} C \rho^{-1/2}, \quad \|K\| \leq 1.$$

For a normalized pure input, its output is congruent to $I + \sin \theta \text{Re}(e^{i\phi} K)$. A zero eigenvalue requires $\sin \theta = 1$ and a unit vector v with $\text{Re}\langle v, e^{i\phi} K v \rangle = -1$. Equality in the contraction and Cauchy–Schwarz inequalities gives $Kv = -e^{-i\phi} v$. A unit-circle eigenvector of a contraction is reducing: $K^\dagger v = -e^{i\phi} v$, since its saturated norm implies $(I - K^\dagger K)v = 0$. Distinct pure rays give distinct phases and therefore

orthogonal reducing eigenvectors. The remaining orthogonal space is one-dimensional, so K is normal. Diagonalizing it gives the displayed form. Linear rank three makes the three diagonal response functionals independent. The Choi matrix then splits into three positive 2-by-2 blocks $\begin{pmatrix} 1 & \lambda_i \\ \lambda_i & 1 \end{pmatrix}$ under congruence. Two have rank one and the third has rank one or two, proving the last statement. $\square$

For a non-trace-preserving branch with faithful input effect on its two-dimensional input support, an invertible input whitening makes it trace preserving without changing ranks or distinct pure rays. In a putative equality profile $(a, b) = (1, 2)$, two other output planes project to distinct rays in this branch's input plane, and equality requires their posterior outputs to have rank two. Thus the lemma applies to that branch. It is only a necessary normal form: its output-congruenced diagonal structure is not literal commutativity of the physical output states, and no global full-law realization follows from it alone.

7 Retained reference access separates two optimal passive generators

The source-readout programme already distinguishes a passive output law from an intervention or an accessible retained reference. Here this distinction has a sharp, quantitative example even after jointly minimizing memory and one-step pure-environment dimension. It is an explicit application of established channel discrimination, not a new general identifiability ceiling.

Use the minimum-Kraus family, writing its parameter as $h \in (0, 2)$:

$$r(h)^2 = \frac{h^3(2-h)}{2(h+1)(3h+2)}, \quad a(h) = \frac{1}{2} \left(1 + \sqrt{\frac{2-h}{3h+2}} \right), \quad t(h)^2 = \frac{h(2-h)}{2(h+1)(3h+2)}.$$

Its nine pair states and CP reset maps, defined in the main resource manuscript, generate the entire stationary delayed-trine law with $r = 1 - \lambda$. Each realization has reusable dimension three and recorded one-step Choi rank nine, both minimal at the indicated parameter. Its stationary memory is $I_3/3$. The maximum of $r(h)$ occurs at $h_* = (1 + \sqrt{17})/4$; every interior $0 < r < r_{\max}$ has two distinct parameters on opposite sides of that maximum.

Theorem 7.1 (An exact access-dependent operational separation). *Let $h_1 \neq h_2$ have $r(h_1) = r(h_2) > 0$. The two admitted source instruments generate identical passive classical word laws, attain the same one-step Pareto point $(d, k) = (3, 9)$, and have full-history cost $6 \cdot 3^n$ for $n \geq 2$ (the one-step rank is nine). Nevertheless their single-use pointer channels on the common qutrit input obey*

$$\frac{1}{2} \|\mathcal{M}_{h_1} - \mathcal{M}_{h_2}\|_{\diamond} = \frac{|h_1 - h_2|}{3}.$$

The distance is attained either by one fixed known Fourier-basis input, or by retaining and reading a common maximally entangled purifier of the stationary local input $I_3/3$. It is invisible to the passive pointer archive alone at every finite horizon. No claim equates the two full retained quantum instruments.

Proof. Let P_j be the three fixed orthogonal Fourier projectors, indexed once consistently with the source covariance. The effects computed from the Choi blocks are

$$E_j(h) = \frac{2-h}{6} I + \frac{h}{2} P_j.$$

This follows from the fixed phase $\zeta = -t\omega^2$, eigenvalues $(2-h)/6, (2-h)/6, (1+h)/3$ and covariance under Z^j . Thus

$$\mathcal{M}_h = (h/2)\mathcal{M}_P + (1-h/2)\mathcal{D},$$

where $\mathcal{M}_P$ is the projective Fourier measurement with classical pointer and $\mathcal{D}(\rho) = \text{Tr}(\rho)I_A/3$. For any positive joint input ρ_{SR} of trace one put $\sigma_j = \text{Tr}_S[(P_j \otimes I)\rho_{SR}]$; $\sigma_j \geq 0$ and $\sum_j \sigma_j = \rho_R$. Then

$$\sum_j \|\sigma_j - \rho_R/3\|_1 \leq \sum_j \left[\frac{2}{3} \text{Tr} \sigma_j + \frac{1}{3} \sum_{k \neq j} \text{Tr} \sigma_k \right] = \frac{4}{3}.$$

The standard stabilization of channel trace-distance discrimination permits restriction to positive joint inputs. Hence half the diamond distance between $\mathcal{M}_P$ and $\mathcal{D}$ is at most $2/3$. A Fourier eigenstate gives deterministic versus uniform labels and attains this value. Scaling by $|h_1 - h_2|/2$ proves the theorem's distance.

For the common purification $|\Omega\rangle = 3^{-1/2} \sum_i |i\rangle_S |i\rangle_R$, the output pointer/reference block is $E_j(h)^T/3$. Its difference is $(h_1 - h_2)(P_j^T - I/3)/6$. Since $\|P_j - I/3\|_1 = 4/3$, the sum of the three block norms gives the same half trace distance $|h_1 - h_2|/3$. One can attain it by measuring the reference in the conjugate Fourier basis and comparing that outcome with the pointer. All preparations, reference basis and decoding are fixed independently of which candidate source is used. The equal passive laws and resource counts follow from the original family theorem. $\square$

For equal prior probabilities of the two candidate generators, one use therefore achieves optimal success probability

$$p_{\text{succ}} = \frac{1}{2} + |h_1 - h_2|/6.$$

This is discrimination between two conventional quantum implementations, not a test of Shadow Theory against standard quantum mechanics. The purifier must be physically available and measurable; retaining an inaccessible reference does not make this decision possible. Supplying a joint purifier is an original preparation contract shared by both models; its marginal agrees with the stationary state but its operational access differs from a pointer-only experiment.

7.1 An algebraic pair and a reproducible certificate

Take $h_1 = 1$, $r^2 = 1/20$, so $\lambda = 1 - 1/\sqrt{20}$. The unique root $h_2 \in (3/2, 8/5)$ of

$$10h_2^3 - 10h_2^2 - 7h_2 - 2 = 0$$

is the second parameter. Indeed

$$r(h)^2 - \frac{1}{20} = -\frac{(h-1)(10h^3 - 10h^2 - 7h - 2)}{20(h+1)(3h+2)}.$$

The cubic changes sign across that interval and its derivative is positive there, proving the root's existence and uniqueness. Rational bisection certifies it without a floating-point feasibility claim. The code also rebuilds the effects and checks the passive full-word update from the state-basis inverse, with explicit entangled-reference and ordinary-input distance calculations. The universal conclusion follows from the proof.

This example strengthens the concrete operational interface: equality of all passive histories and joint dimension optimality still do not identify the generator's reference correlations. It establishes an access-dependent distinction between the specified generators, not an additional ontological source.

8 Current-status overlay

The analytic and rational encoded-pair constructions remain valid sufficient examples, not the best current qutrit feasibility endpoint. The exact dimension-three interval now starts at $\lambda_{\dagger} = 1 -$

$\sqrt{71 - 17\sqrt{17}}/4$. For complete channels retaining final memory, the current universal bounds are $48, 145, 437 \cdot 3^{n-4}$; the explicit mixed source at .99 attains $50, 152, 456 \cdot 3^{n-4}$. Earlier specified-channel costs remain valid for those channels, but are not universal process minima. The current one-step qutrit optimum is nine on the certified high-noise interval. These overlays identify the current process-wide status; historical tables continue to describe their explicitly specified channels rather than different updated channels.

References

- [1] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975. doi: 10.1016/0024-3795(75)90075-0.
- [2] Ky Fan. On a theorem of weyl concerning eigenvalues of linear transformations I. *Proceedings of the National Academy of Sciences*, 35(11):652–655, 1949. doi: 10.1073/pnas.35.11.652.
- [3] Farzad Ghafari, Nora Tischler, Jayne Thompson, Mile Gu, Lynden K. Shalm, Varun B. Verma, Sae Woo Nam, Raj B. Patel, Howard M. Wiseman, and Geoff J. Pryde. Dimensional quantum memory advantage in the simulation of stochastic processes. *Physical Review X*, 9:041013, 2019. doi: 10.1103/PhysRevX.9.041013. URL <https://journals.aps.org/prx/abstract/10.1103/PhysRevX.9.041013>.
- [4] Mile Gu, Karoline Wiesner, Elisabeth Rieper, and Vlatko Vedral. Quantum mechanics can reduce the complexity of classical models. *Nature Communications*, 3:762, 2012. doi: 10.1038/ncomms1761. URL <https://arxiv.org/abs/1102.1994>.
- [5] Alex Monràs and Andreas Winter. Quantum learning of classical stochastic processes: The completely-positive realization problem. *Journal of Mathematical Physics*, 57:015219, 2016. doi: 10.1063/1.4936935. URL <https://arxiv.org/abs/1412.3634>.
- [6] Jason D. Padgett. Holographic mathematics: A dual representation of algebra, trigonometry, probability, and calculus in quantum information holography, . Undated supplied PDF manuscript; abstract and Sections 2–3.
- [7] Jason D. Padgett. Quantum information holography, . Undated supplied manuscript; pages 64–68.
- [8] Paul M. Riechers and Thomas J. Elliott. Identifiability and minimality bounds of quantum and post-quantum models of classical stochastic processes, 2025. URL <https://arxiv.org/html/2509.03004v1>. Version 1.
- [9] Jeremy Rodgers. Exact markov extensions of three-time readouts: Hidden state counts, reversibility obstructions, and nonidentifiability, 2026. Version 1.0, 6 September 2026; supplied manuscript.
- [10] Jeremy Rodgers. Source–readout non-equivalence: Descent and equivariant reconstruction obstructions, 2026. URL <https://doi.org/10.5281/zenodo.21369967>.
- [11] Rishi Sundar and Thomas J. Elliott. Quantum dimension reduction of hidden markov models. *Physical Review A*, 113:062461, 2026. doi: 10.1103/plt9-cm1x. URL <https://arxiv.org/html/2601.16126v2>. Version 2, 1 July 2026.
- [12] John Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018. doi: 10.1017/9781316848142. URL <https://cs.uwaterloo.ca/~watrous/TQI/TQI.pdf>.