ONE QUBIT, TWO CLASSICAL STATES: A CERTIFIED MEMORY GAP

Anonymous authors
Paper under double-blind review

ABSTRACT

A quantum model can beat every classical model we train without establishing a memory advantage: a better classical model may remain undiscovered. We develop certificates that rule out this possibility for finite binary event distributions. They prove when a single qubit achieves lower total-variation error than every homogeneous two-state edge-emitting hidden Markov model. The key is shared dynamics: one physically feasible recurrence must explain every temporal context. Covering its two-dimensional coefficient region yields classical error lower bounds checked with exact arithmetic. These constraints detect structure missed by rank tests, although satisfying them need not yield a valid classical model. Across six designed quantum constructions, the method certifies candidates left unresolved by the specified clock, rank and complete-model search policies on every construction. On five, it brackets the best classical error within an interval of width 0.025. The comparison concerns persistent-state memory; the qubit remains inexpensive to simulate classically.

1 INTRODUCTION

Given a target sequence distribution and a quantum generator, can any classical model with a specified state budget match its accuracy? Fitted baselines cannot settle this question: an untested model may fit better. A certificate must bound the error of the entire classical class, separating a memory limitation from an optimization failure.

We study finite binary event laws and homogeneous two-state edge-emitting hidden Markov models (HMMs). Our acoustic example fits a repeated qubit instrument to crossing patterns extracted by a classical speech filter. The instrument generates event sequences autonomously; its dynamics do not encode a time-varying waveform. Both generators are compared on the same preprocessed event law.

Every two-state generator's no-event matrix obeys a quadratic characteristic equation. Its trace and determinant form one physically feasible coefficient pair, which must explain every past and future context. We cover this two-dimensional region, relax each cell into linear probability constraints, and verify a lower error bound in every cell. The weakest cell supplies the bound for the whole classical class (Figure 1).

Shared feasible recurrences can exclude finite laws that satisfy ordinary rank-two constraints. Yet satisfying all such recurrences need not produce a nonnegative stochastic realization. We characterize the set reached by coefficient refinement and give quantitative examples separating these notions. These distinctions explain what the certificate gains over rank and why refinement can still leave a gap to the best classical model.

On a frozen law from 3,999 Speech Commands recordings, the noisy qubit's error is below the certified two-state floor. An earlier clock witness already settles that ordering. We therefore test additional decision value on six prospective quantum constructions. On every construction the complete pipeline validates candidates left unresolved by the specified clock, rank and complete-model certificates; five also resolve physical classical-error intervals within width 0.025.

Quantum instruments and finite-memory advantages are established (Monras et al., 2011; Budroni et al., 2021; Fanizza et al., 2024). Our contribution is the certificate, its refinement and realization

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

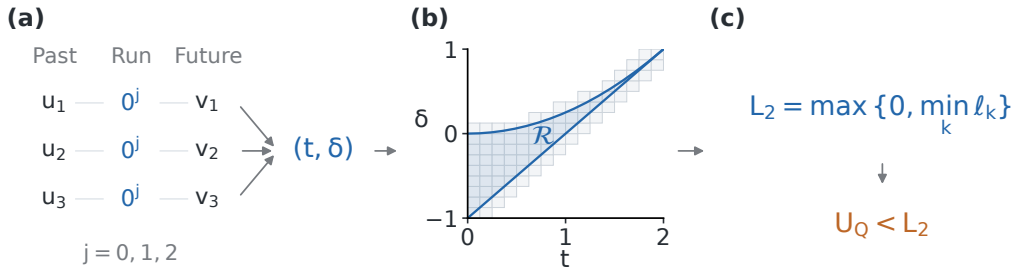


Figure 1: (a) Every context shares one feasible coefficient pair. (b) The retained 80 cells cover its attainable region. (c) The minimum checked cell bound gives a whole-class floor; candidate error below it excludes every two-state HMM. Contexts and workflow are schematic; the feasible region and cover are exact.

analysis, and its tested additional decisions. The comparison concerns persistent states: the qubit is cheaply simulated classically, four-state HMMs fit the audio better, and transfer and a separate hardware pilot fail.

2 Finite event laws and the memory comparison

Let p be a distribution on binary words of length n, with total-variation distance $\mathrm{TV}(p,q) = \frac{1}{2}\sum_w |p(w) - q(w)|$. For shorter words, $P_p(w)$ denotes the prefix probability obtained by summing trailing symbols; $P_p(\emptyset) = 1$. A finite target need not have an extension generated by any fixed small model.

The class $\mathcal{H}_2^{(n)}$ consists of homogeneous two-state edge-emitting HMM laws

$$q(w) = \pi B_{w_1} \cdots B_{w_n} \mathbf{1}, \qquad B_0, B_1 \geq 0, \quad (B_0 + B_1)\mathbf{1} = \mathbf{1}. \tag{1}$$

The matrices are 2×2, and π is an arbitrary fixed initial probability row. A step jointly emits a bit and chooses the next state. Reducible models, zero probabilities and nonstationary initial distributions are allowed. The same matrices act at every position; an additional persistent clock, mixture selector or history controller adds memory. Fresh randomness does not.

For a candidate law q_Q, let U_Q be a verified upper bound on $\mathrm{TV}(p, q_Q)$. We seek a floor $L_2 \leq \inf_{q\in\mathcal{H}_2^{(n)}} \mathrm{TV}(p,q)$. The decision is

$$U_Q < L_2. \tag{2}$$

A fitted HMM instead supplies a *physical upper bound* U_2 on the infimum. Thus $[L_2, U_2]$ brackets the best classical error, whereas $L_2 - U_Q$ validates the particular candidate. An upper witness in a recurrence relaxation need not be a physical HMM; we keep these quantities separate throughout.

2.1 The acoustic candidate

At $n = 8$, each bit records a positive crossing within a four-sample bin at 16 kHz. A causal fourth-order 750–1250 Hz Butterworth filter supplies the band-limited signal; eight bins span 2 ms. Crossing anchors overlap. We average their exact word frequencies first within each recording, then across recordings of each speaker, then across speakers. The resulting finite-corpus target is rational. Digit labels select recordings; the model does not predict them. Appendix C specifies filtering, boundaries and weights.

The candidate is a restricted hidden quantum Markov model (Monras et al., 2011; Srinivasan et al., 2018). A real orthogonal matrix acts on a memory qubit and an ancilla initialized to zero. Measuring the ancilla defines 2×2 Kraus operators K_0, K_1, with $\sum_y K_y^{\mathsf{T}} K_y = I$. At noise rate ϵ, the reported-symbol map is

$$\begin{aligned} \mathcal{M}_y^\epsilon(\rho) &= (1-\epsilon)\mathcal{D}_\epsilon(K_y \rho K_y^{\mathsf{T}}) + \epsilon \mathcal{D}_\epsilon(K_{1-y}\rho K_{1-y}^{\mathsf{T}}), \\ q_\epsilon(w) &= \mathrm{Tr}[\mathcal{M}_{w_n}^\epsilon \cdots \mathcal{M}_{w_1}^\epsilon(\rho_0)]. \end{aligned} \tag{3}$$

Here $\mathcal{D}_\epsilon$ preserves diagonals and multiplies off-diagonals by $1-\epsilon$; report flips are independent. The ancilla resets after each step, the memory restarts at the same learned ρ_0 for every block, and outputs do not control later operations. Both generators use the classical frontend. Persistent-state accounting excludes frontend and output storage; ancilla, control, reset and measurement remain implementation costs. Equation (3) is also a cheap classical simulation recipe.

3 A CERTIFICATE FROM SHARED FEASIBLE COEFFICIENTS

3.1 ONE PAIR MUST EXPLAIN EVERY CONTEXT

Proposition 1 (feasible coefficients and shared recurrence). The trace t and determinant δ of a nonnegative row-substochastic 2×2 matrix range over exactly

$$\mathcal{R} = \{(t,\delta) : 0 \le t \le 2, \quad t-1 \le \delta \le t^2/4\}. \tag{4}$$

Every law in $\mathcal{H}_2^{(n)}$ satisfies, with one pair in $\mathcal{R}$,

$$P(u00v) - tP(u0v) + \delta P(uv) = 0 \quad \text{for all } |u|+|v| \le n-2. \tag{5}$$

Proof sketch. Writing $B_0 = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, row sums give $ad-bc \ge ad-(1-a)(1-d) = t-1$, and $ad-bc \le ad \le t^2/4$. Conversely, diagonal entries $t/2$ and off-diagonal entries $\sqrt{t^2/4-\delta}$ realize every pair in (4); a nonnegative B_1 fills the remaining row mass. Multiplying $B_0^2 - tB_0 + \delta I = 0$ by πB_u and $B_v\mathbf{1}$ proves (5). Appendix A includes boundary cases and the complete certificate proof.

For example, $u=v=1$ gives $P(1001) - tP(101) + \delta P(11) = 0$. The pair explaining that silence run must also explain every other context. At eight symbols there are 769 such identities. Fitting unrelated coefficients to each one discards the shared-dynamics restriction.

3.2 COVERING THE UNKNOWN PAIR

For a rectangle $C = [t_-, t_+] \times [\delta_-, \delta_+]$, nonnegative prefix probabilities relax (5) to

$$\begin{aligned} P(u00v) - t_+ P(u0v) + \delta_- P(uv) &\le 0, \\ -P(u00v) + t_- P(u0v) - \delta_+ P(uv) &\le 0. \end{aligned} \tag{6}$$

These inequalities define a polytope $\mathcal{Q}_C$ in the word-probability simplex. If a finite family $\mathcal{F}$ covers $\mathcal{R}$, then every HMM2 law lies in at least one $\mathcal{Q}_C$. Hence

$$\inf_{q \in \mathcal{H}_2^{(n)}} \mathrm{TV}(p,q) \;\ge\; \min_{C\in\mathcal{F}} \mathrm{dist}_{\mathrm{TV}}(p, \mathcal{Q}_C). \tag{7}$$

The minimum is essential: the unknown classical model may lie in the weakest cell. A high bound from a subset of cells cannot certify the class.

Our acoustic confirmation uses a fixed 16×16 grid over $[0,2]\times[-1,1]$. Exact strict tests discard only cells disjoint from $\mathcal{R}$, retaining all boundaries and 80 cells. The prospective pipeline starts from the full region and adaptively bisects the cell with the weakest inherited bound. It retains a complete frontier at every stopping time; an unresolved cell keeps its valid earlier bound. Search order affects power, not validity.

3.3 CHECKING THE FLOOR WITHOUT TRUSTING OPTIMIZATION

Write a cell's rational inequalities as $A_C q \le 0$. A numerical LP supplies proposed dual multipliers, which we repair on denominator $D = 2^{40}$ to nonnegative integer vectors a, b, λ satisfying $a_j + b_j \le D/2$. With

$$\begin{aligned} \nu &= \min_j (a - b + A_C^{\mathsf{T}}\lambda)_j, \quad \ell_C(p) = \frac{(b-a)^{\mathsf{T}}p + \nu}{D}, \\ L_2 &= \max\{0, \min_{C\in\mathcal{F}} \ell_C(p)\}. \end{aligned} \tag{8}$$

Weak duality proves $\ell_C(p) \le \mathrm{dist}_{\mathrm{TV}}(p, \mathcal{Q}_C)$. The fixed grid scales A_C to integers; adaptive cells use exact rational arithmetic. Solver optimality is unnecessary. A separate reader reconstructs the

target, context rows, cover and dual inequalities. Acoustic candidate probabilities use outward-rounded interval arithmetic on the stored angles and declared channel. The proof establishes validity; finite implementation controls check its realization in code.

The pipeline fits a physical HMM upper, attempts an exact affine correction at one coefficient pair, and refines both symbol orientations. It takes the larger checked floor because complementation preserves HMM2. This floor and the physical upper bracket the HMM optimum; comparison with the candidate error gives the memory decision. A successful correction supplies a recurrence upper that prunes that orientation's search without improving the physical upper. The uncorrected orientation uses only the physical upper. Corrected laws need not be HMMs. Appendices A and D give the correction, coverage and policy.

4 What the certificate retains and misses

4.1 Refinement has a precise target

Let $E_0^{(n)}$ be the valid length-n laws admitting one pair in $\mathcal{R}$ that satisfies every zero-symbol recurrence (5). Let $E^{(m)}$ be the union of cell polytopes from nested rectangular partitions, retaining exactly the cells that intersect $\mathcal{R}$, with largest side length h_m. Write $d_m(p) = \mathrm{dist}_{\mathrm{TV}}(p, E^{(m)})$.

Proposition 2 (mesh limit and residual certificate). For fixed $n \geq 4$, if $h_m \to 0$, then

$$\mathcal{H}_2^{(n)} \subseteq E_0^{(n)} = \bigcap_m E^{(m)}, \qquad d_m(p) \uparrow \mathrm{dist}_{\mathrm{TV}}(p, E_0^{(n)}) \tag{9}$$

uniformly over valid p. Moreover, $d_m(p) \geq (r_\star(p) - 2h_m)_+/4$, where

$$r_\star(p) = \min_{(t,\delta)\in\mathcal{R}} \max_{|u|+|v|\leq n-2} |P_p(u00v) - tP_p(u0v) + \delta P_p(uv)|.$$

The proof uses nested compact relaxations and the fact that each cell's recurrence residual vanishes with its width. The residual is 4-Lipschitz in TV, giving the displayed finite-mesh bound. This is convergence of optimal cell relaxations to $E_0^{(n)}$. It supplies neither a general convergence rate nor convergence of arbitrary repaired witnesses to the physical HMM distance. Appendix B proves these statements and the following two separations.

4.2 Sharing and feasibility each add information

A compatible prefix-Hankel table has entries $H[u, v] = P_p(uv)$, with prefix and suffix sets chosen so that every concatenation fits within the horizon.

Proposition 3 (information beyond ordinary rank). Fix $n \geq 3$, $0 < b \leq 1/(4n)$, and $z = 1/2 + ib$. For a word of length n containing k zeros, define $p_{n,b}(w) = \mathrm{Re}[z^k \bar{z}^{n-k}]$. This is a strictly positive law. Every compatible prefix-Hankel table has rank at most two, each context separately admits feasible coefficients, and all contexts admit shared unrestricted coefficients. Nevertheless,

$$\frac{b^3}{4} \leq \mathrm{dist}_{\mathrm{TV}}(p_{n,b}, E_0^{(n)}) \leq 4n(n-1)^2 b^3. \tag{10}$$

The shared pair is $(1, 1/4 + b^2)$, outside $\mathcal{R}$. Separate context fits can instead choose feasible pairs. Thus neither sharing alone nor feasibility alone enforces their conjunction. A real two-dimensional factorization proves the rank bound; a discriminant polynomial of three survival probabilities proves the TV exclusion. A positive derivative-law construction supplies the upper bound. That upper concerns the recurrence set, not the physical HMM class. The constants depend on n.

4.3 Recurrence consistency still permits nonphysical laws

Let $E_1^{(n)}$ impose the corresponding shared feasible recurrence for symbol one.

Proposition 4 (positive-realization gap). At every fixed $n \geq 3$, there is a strictly positive law in $E_0^{(n)} \cap E_1^{(n)}$ whose compatible prefix-Hankel tables have rank at most two, yet which lies outside

$\mathcal{H}_2^{(n)}$. At $n = 3$, an explicit example and its closed TV neighborhood of radius $1/2000$ exclude the entire two-state class.

The construction is a signed difference of two Bernoulli product laws that remains positive through the chosen horizon. Its first four no-event probabilities identify a unique quadratic continuation. Any matching HMM must follow that continuation, which later becomes negative. This contradicts a stochastic realization without requiring the hypothetical HMM itself to be exchangeable.

Together, these propositions give the certificate an interpretable boundary. It preserves a shared physical restriction that rank can miss, but can still lose joint stochastic realizability. They do not establish general nesting between rank and recurrence relaxations. The finite constructions are mechanism examples, not quantum advantage examples; Proposition 3's families have no homogeneous qubit realization for $n \geq 5$. Physical quantum laws are tested separately below.

5 Evidence: an acoustic application and additional decisions

5.1 A frozen acoustic comparison

We fit eight model recipes on 3,000 Speech Commands v0.02 recordings using eight initializations per recipe (Warden, 2018). These include noisy/clean qubit objectives, an oscillator and HMMs with two, three and four states. Selection uses the declared training objective alone. All 64 fits, including capped runs, are retained. Development first failed its prospective test, then informed certificate design without changing the learned model. Parameters, frontend, noise and certificate were frozen before decoding the official confirmation set. Appendix C records the complete protocol.

Confirmation contains 3,999 recordings, 241 speakers and 3,046,061 overlapping anchors. At 1% partial dephasing and report flips, outward interval endpoints establish

$$\mathrm{TV}(p_{\mathrm{SC}}, q_{0.01}) \leq 0.187878493 < 0.204525522 \leq \inf_{q \in \mathcal{H}_2^{(8)}} \mathrm{TV}(p_{\mathrm{SC}}, q). \tag{11}$$

The finite-corpus ordering does not depend on resampling. As a stability diagnostic, 2,000 whole-speaker bootstrap replicates give a conditional 95% gap interval $[0.012555, 0.019540]$. They preserve the target's recording weights and use the frozen witnesses without refitting or reoptimization; this is not a population guarantee.

Table 1: The acoustic comparison and its scope. Quantum error and classical floors use conservative displayed endpoints; fitted HMM errors are rounded evaluations. The clock floor imports the published constants used by the saved witness (Budroni et al., 2021). HMM4 outperforms the qubit despite reaching its optimization cap. Appendix C gives all selected recipes.

Quantity on confirmation	TV error or floor	Interpretation
Noisy qubit	≤ 0.187878493	Verified achievable error
Earlier clock certificate	≥ 0.196644706	Already excludes all HMM2 candidates at the qubit's error
Shared-recurrence certificate	≥ 0.204525522	Strengthens the all-HMM2 floor
Fitted HMM2	0.365027	Achieved error; cannot establish a class-wide floor
Fitted HMM4	0.101113	Better achieved audio fit with a larger state budget

Equation (11) establishes a memory-budget comparison, but does not by itself establish the new certificate's decision value: the clock witness also succeeds. That question motivates the next study.

5.2 Does the complete method decide additional cases?

We fixed six new rational qubit constructions, two HMM2 controls and a 50-rate degradation grid before generating the target laws. All targets have horizon eight. A target is the construction's

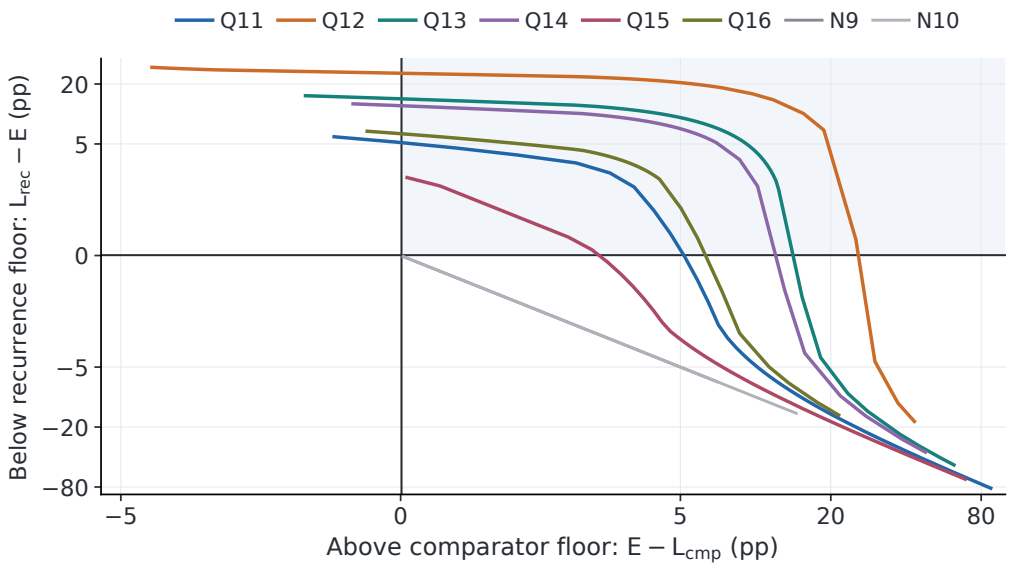


Figure 2: Decision margins for all 288 eligible quantum candidates and 96 HMM2 controls. Solid lines connect tested rates in increasing noise order; control curves overlap. The shaded region requires both margins to be at least 0.01 percentage points: 162/288 quantum candidates and 0/96 controls satisfy this criterion. Axes are symmetric-log, linear within ± 2 percentage points. L_{cmp} is the largest comparator floor. Both margins share E; their sum is fixed per construction. Exact fractions determine decisions. Grid points are dependent; connecting lines imply no between-rate guarantee.

base law; candidates vary dephasing and report flips together. Their exact rational errors $E_i(\eta) = \mathrm{TV}(p_i, q_{i,\eta})$ require no fitting. The 48 rates strictly between zero and one half are eligible; identity and uniform endpoints are retained as diagnostics.

Four fixed policies receive the target law: the complete recurrence pipeline, complete-model interval search, clock witnesses and rank witnesses. The complete-model comparator receives the same physical-fit initializer allowance as recurrence. Each producer has a nominal 120-second ceiling. A 600-second logical ceiling also charges target construction, target audit, proof readback and the full shared final check. The clock/rank/complete-model envelope sums separately budgeted arms; it is not one equally budgeted solver. Appendix D defines every construction, policy and charge.

For method floors $L_{i,a}$, an additional decision requires

$$L_{i,\mathrm{rec}} - E_i(\eta) \geq 10^{-4}, \qquad E_i(\eta) - \max\{L_{i,\mathrm{clock}}, L_{i,\mathrm{rank}}, L_{i,\mathrm{complete}}\} \geq 10^{-4}. \tag{12}$$

The first inequality validates the quantum candidate against every HMM2. The second places its error above all three comparator floors, showing why these certificates are inconclusive. A construction counts once if any eligible rate passes both margins. The registered endpoint requires at least three of six constructions, valid proofs, zero control floors, cost compliance and no exact-law or symbol-complement collision with the earlier panel.

All six quantum constructions pass. Both HMM2 controls have zero floors for all four methods. All target reconstructions and certificate checks succeed, and every method stays within its declared budget. Figure 2 shows all candidate decision margins; Figure 3 summarizes decisions, physical widths and costs by construction. Recurrence supplies 162 additional decisions on the 288 eligible quantum grid points. These are dependent design points on a nonuniform grid, not independent replications or a continuous-noise guarantee. No tested candidate has the reverse strict ordering.

Recurrence meets the physical-width target on five constructions; complete-model search meets it on none. Q16's width is 0.036503. Fully charged costs are 51.964–121.107 seconds for recurrence, 399.896–428.410 for complete-model search and about five each for clock/rank. All 32 arms finish

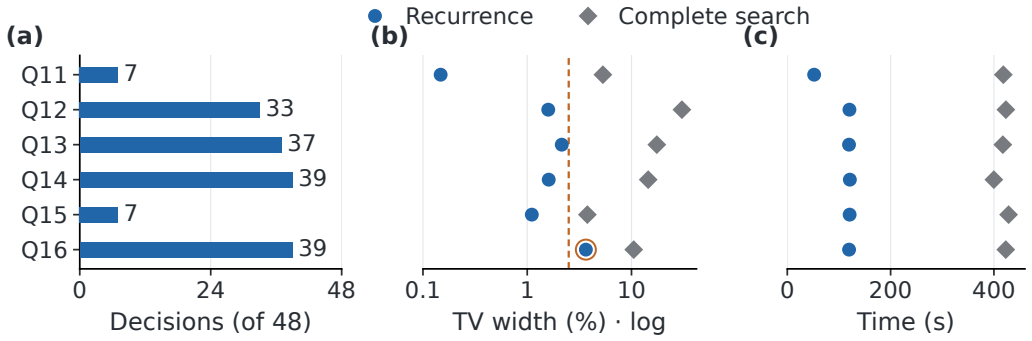


Figure 3: (a) Additional decisions among 48 eligible rates per construction. (b) Physical TV widths subtract the verified recurrence or complete-model floor from a feasible HMM2 upper, shown as percentages on a logarithmic axis; recurrence-only correction uppers are excluded. Recurrence meets the dashed 2.5% target on five of six constructions and complete-model search on none; the ring marks Q16's miss. (c) Total classical wall time in seconds includes construction and checking under the fixed policies.

within 600 seconds. Complete-model proof checking accounts for 284–313 seconds of its charged time.

This fixed-policy comparison neither isolates each component nor optimizes each algorithm's allocation under total cost. A general solver ranking does not follow. The constructions are designed examples, not a representative sample; exact noncollision establishes distinct inputs, not statistical independence.

Together, the six uncorrected orientations already supply 139 of the 162 added decisions. Pairing adds 23 decisions on Q14 and Q16 and the fifth physical-width success, on Q14. Four winning floors come from the uncorrected orientation. Removing correction would change the search budget, so this saved-output decomposition is not a correction-disabled experiment. Appendix D gives both floors and stage costs.

5.3 Boundaries that survive the positive results

The frozen qubit fails AudioMNIST transfer on 29,997 recordings from 60 speakers: error 0.138026 exceeds floor 0.121712, with conditional gap interval $[-0.025371, -0.011119]$. The fixed conversion and model are not retuned (Becker et al., 2023). A separate synthetic hardware pilot also fails: reset-order accuracy is 88.1592% against 90%, and memory error is 0.202496 in **L2** against 0.02. It is a four-bit toy, not the learned acoustic instrument.

A shared-engine factorial finds a positive scheduling effect and a negative clipping effect at 15 seconds; cheap upper witnesses leave all 192 physical intervals wider than 0.025. These results motivated the final pipeline without isolating correction's contribution. Appendix E summarizes supporting and adverse studies; the archive retains every outcome.

6 Relation to prior work

Repeated quantum instruments, their learning and memory/accuracy tradeoffs are established (Monras et al., 2011; Srinivasan et al., 2018; Adhikary et al., 2020; Banchi, 2024). The autonomous finite-block comparator follows existing memory-witness work (Vieira & Budroni, 2022). Budroni et al. (2021) certify sequence probabilities using parameter grids, and Vieira et al. (2024) give a convergent dimension hierarchy. Cayley–Hamilton witnesses also appear for normalization-preserving operations (Białecki et al., 2024). Our certificate uses a substochastic symbol map to build a shared-context TV union over two feasible coefficients. Neither characteristic polynomials nor global covers are new.

HMM algebra and positive-realization theory distinguish linear representations from stochastic or completely positive ones (Critch, 2012; Monras & Winter, 2014; Fanizza et al., 2024). Classical partial realization already recovers finite-state processes from finite windows under its stated assumptions (Huang et al., 2015). In contrast, our target can be an arbitrary finite law; the certificate bounds its distance from the entire homogeneous edge-emitting HMM2 class, including boundaries and nonstationary starts. Finite-window input alone is not the novelty.

Signed finite mixtures and root-collision geometry are also established (Janson et al., 2016; Batenkov & Yomdin, 2014). Propositions 2–4 characterize the particular recurrence relaxation: its exact limit, quantitative information beyond ordinary rank and persistent realization gap. State-emitting algebra, other-norm Hankel bounds and stationary unifilar KL bounds impose different conditions (Critch, 2012; Kotsalis & Shamma, 2015; Yang et al., 2023). For speech-derived processes, Sundar & Elliott (2026) compress a fitted HMM against greedy merging; our acoustic application compares an empirical law with the whole two-state class.

7 Discussion

The certificate distinguishes imperfect fitting from an error floor for all two-state HMMs. The acoustic law and six designed constructions demonstrate its use under the tested policies; usefulness on new targets remains unmeasured.

A continuous classical density-matrix representation cheaply simulates the qubit but falls outside the finite-state comparator. No speech-recognition, runtime or energy advantage follows. Restarting each block from a learned state does not establish an indefinitely stationary acoustic generator. Human-identity overlap across corpora and demographic generalization are unverified.

The supplement supplies the prospective study's mathematical proofs, exact candidate reconstruction and a generator for new rational targets (Appendix F). Historical acoustic and timing reproduction remain separate. The open problem is useful quantitative certification at larger classical memory budgets, where exact separation can have negligible approximation margins.

AI use statement

Generative AI tools assisted with code implementation, data processing, and manuscript preparation, including figures and references. They also supported literature review, hypothesis and model development, mathematical claims and proofs, experimental design, construction of synthetic processes, and interpretation of results. Synthetic laws were generated by the explicit mathematical constructions in Appendix D. Numerical findings come from recorded computations; checks include exact-arithmetic verification, source and input bindings, retained failure records, and saved-proof replay. These checks and AI-assisted internal critiques do not constitute independent human peer review.

Ethics statement

The acoustic experiments use existing Speech Commands and AudioMNIST recordings. The model generates binary acoustic-event patterns rather than recognizable speech or speaker identities. Human-identity overlap across the corpora and demographic generalization have not been assessed. Raw recordings are excluded from the proof packages; recording-specific redistribution rights for AudioMNIST and unified project-code licensing remain unresolved.

Reproducibility statement

Appendices A and B give the proofs; C and D specify protocols, constructions, comparisons and costs. Appendix F maps sources and replay commands. The supplement reconstructs all 400 candidate errors, checks every prospective-study method bound, and provides a minimal new-target generator. The full historical workflow retains an unbundled runtime. Portable checks establish saved mathematical results; they do not rerun optimization, audio or hardware, or authenticate historical times.

REFERENCES

Sandesh Adhikary, Siddarth Srinivasan, Geoff Gordon, and Byron Boots. Expressiveness and learning of hidden quantum markov models. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 4151–4161, 2020. URL https://proceedings.mlr.press/v108/adhikary20a.html.

Leonardo Banchi. Accuracy vs memory advantage in the quantum simulation of stochastic processes. *Machine Learning: Science and Technology*, 5:025036, 2024. doi: 10.1088/2632-2153/ad444a. URL https://arxiv.org/abs/2312.13473.

Dmitry Batenkov and Yosef Yomdin. Geometry and singularities of the prony mapping. *Journal of Singularities*, 10:1–25, 2014. doi: 10.5427/jsing.2014.10a. URL https://arxiv.org/abs/1301.1336.

Sören Becker, Johanna Vielhaben, Marcel Ackermann, Klaus-Robert Müller, Sebastian Lapuschkin, and Wojciech Samek. AudioMNIST: Exploring explainable artificial intelligence for audio analysis on a simple benchmark. *Journal of the Franklin Institute*, 2023. doi: 10.1016/j.jfranklin.2023.11.038. URL https://www.sciencedirect.com/science/article/pii/S0016003223007536.

Tomasz Białecki, Tomasz Rybotycki, Josep Batle, and Adam Bednorz. Quantum dimension witness with a single repeated operation. *Physical Review A*, 109:L020201, 2024. doi: 10.1103/PhysRevA.109.L020201. URL https://journals.aps.org/pra/abstract/10.1103/PhysRevA.109.L020201.

Costantino Budroni, Giuseppe Vitagliano, and Mischa P. Woods. Ticking-clock performance enhanced by nonclassical temporal correlations. *Physical Review Research*, 3:033051, 2021. doi: 10.1103/PhysRevResearch.3.033051. URL https://arxiv.org/abs/2005.04241.

Andrew Critch. Binary hidden markov models and varieties, 2012. URL https://arxiv.org/abs/1206.0500.

Marco Fanizza, Josep Lumbreras, and Andreas Winter. Quantum theory in finite dimension cannot explain every general process with finite memory. *Communications in Mathematical Physics*, 405:50, 2024. doi: 10.1007/s00220-023-04913-4. URL https://link.springer.com/article/10.1007/s00220-023-04913-4.

Qingqing Huang, Rong Ge, Sham Kakade, and Munther Dahleh. Minimal realization problems for hidden markov models, 2015. URL https://arxiv.org/abs/1411.3698v2. Version 2.

Svante Janson, Takis Konstantopoulos, and Linglong Yuan. On a representation theorem for finitely exchangeable random vectors, 2016. URL https://arxiv.org/html/1410.1777v3. Version 3.

Georgios Kotsalis and Jeff S. Shamma. Limits of performance for the model reduction problem of hidden markov models. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 4674–4679, 2015. doi: 10.1109/CDC.2015.7402948. URL https://doi.org/10.1109/CDC.2015.7402948.

Lvzhou Li and Daowen Qiu. A polynomial-time algorithm for the equivalence between quantum sequential machines. *arXiv preprint quant-ph/0604085*, 2006. URL https://arxiv.org/abs/quant-ph/0604085.

Alex Monras and Andreas Winter. Quantum learning of classical stochastic processes: The completely-positive realization problem. In *9th Conference on the Theory of Quantum Computation, Communication and Cryptography*, volume 27 of *Leibniz International Proceedings in Informatics*, pp. 99–109, 2014. doi: 10.4230/LIPIcs.TQC.2014.99. URL https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.TQC.2014.99.

Alex Monras, Almut Beige, and Karoline Wiesner. Hidden quantum markov models and non-adaptive read-out of many-body states. *Applied Mathematical and Computational Sciences*, 3: 93, 2011. URL https://arxiv.org/abs/1002.2337.

Siddarth Srinivasan, Geoff Gordon, and Byron Boots. Learning hidden quantum markov models. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 1979–1987, 2018. URL https://proceedings.mlr.press/v84/srinivasan18a.html.

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.

Lucas B. Vieira and Costantino Budroni. Temporal correlations in the simplest measurement sequences. *Quantum*, 6:623, 2022. doi: 10.22331/q-2022-01-18-623. URL https://arxiv.org/html/2104.02467v2.

Lucas B. Vieira, Simon Milz, Giuseppe Vitagliano, and Costantino Budroni. Witnessing environment dimension through temporal correlations. *Quantum*, 8:1224, 2024. doi: 10.22331/q-2024-01-10-1224. URL https://arxiv.org/html/2305.19175v2.

Pete Warden. Speech commands: A dataset for limited-vocabulary speech recognition, 2018. URL https://arxiv.org/abs/1804.03209.

Chengran Yang, Andrew J. P. Garner, Feiyang Liu, Nora Tischler, Jayne Thompson, Man-Hong Yung, Mile Gu, and Oscar Dahlsten. Provably superior accuracy in quantum stochastic modeling. *Physical Review A*, 108(2):022411, 2023. doi: 10.1103/PhysRevA.108.022411. URL https://journals.aps.org/pra/abstract/10.1103/PhysRevA.108.022411.

A CERTIFICATE CORRECTNESS AND VERIFIED COMPUTATION

This appendix proves the whole-class certificate and the operations used to refine it. Fix a horizon $n \geq 2$, and let Δ_n be the simplex of distributions on $\{0,1\}^n$. For $p \in \Delta_n$, write $P_p(w)$ for the probability of the prefix w, obtained by summing its length-n completions, and set $P_p(\emptyset) = 1$. All distances below use $\mathrm{TV}(p,q) = \frac{1}{2}\|p-q\|_1$.

A.1 FEASIBLE COEFFICIENTS AND CONTEXT IDENTITIES

The class $\mathcal{H}_2^{(n)}$ consists of laws generated by fixed nonnegative 2×2 matrices B_0, B_1, with $(B_0 + B_1)\mathbf{1} = \mathbf{1}$, and an arbitrary initial probability row π. For a word w, let B_w be its ordered matrix product, with $B_\emptyset = I$. Summing trailing outputs gives $P_q(w) = \pi B_w \mathbf{1}$ whenever $|w| \leq n$. The same matrices act at every position; the initial distribution need not be stationary.

Proof of the coefficient and recurrence statements in Proposition 1. Write $B_0 = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$, $t = a + d$, and $\delta = ad - bc$. Its row sums give $b \leq 1 - a$, $c \leq 1 - d$, and $a, d \in [0,1]$. Consequently,

$$0 \leq t \leq 2, \qquad \delta \geq ad - (1-a)(1-d) = t - 1, \qquad \delta \leq ad \leq t^2/4. \tag{A.1}$$

Conversely, for any pair in

$$\mathcal{R} = \{(t,\delta) : 0 \leq t \leq 2,\ t - 1 \leq \delta \leq t^2/4\}, \qquad D_0 = [0,2] \times [-1,1], \tag{A.2}$$

take $a = d = t/2$ and $b = c = \sqrt{t^2/4 - \delta}$. The lower determinant bound implies $b^2 \leq (1 - t/2)^2$, so both row sums are at most one. A nonnegative B_1 can place each remaining row mass in its first column. Thus every pair in $\mathcal{R}$ occurs in an admissible two-state model. Cayley–Hamilton now gives

$$\begin{aligned} 0 &= \pi B_u (B_0^2 - tB_0 + \delta I) B_v \mathbf{1} \\ &= P_q(u00v) - tP_q(u0v) + \delta P_q(uv) \end{aligned} \tag{A.3}$$

for every context with $|u| + |v| \leq n - 2$. This identity includes defective matrices and boundary parameters. □

Let $\mathcal{C}_n = \{(u,v) : |u| + |v| \leq n - 2\}$. For $|u| + |v| = k$, there are $k + 1$ splits and 2^k binary assignments, giving $\sum_{k=0}^{n-2}(k+1)2^k$ contexts. Define binary prefix-selection matrices V_j, $j = 0, 1, 2$, by $(V_j q)_{uv} = P_q(u0^j v)$, with length-$n$ words in lexicographic order. At $n = 8$, each matrix has 769 rows and 256 columns. Empty prefixes select every coordinate.

A.2 EXACT GEOMETRY AND CELL INCLUSION

For a closed rectangle $C = [t_l, t_u] \times [\delta_l, \delta_u] \subseteq D_0$, define

$$G_C = \begin{bmatrix} V_2 - t_u V_1 + \delta_l V_0 \\ -V_2 + t_l V_1 - \delta_u V_0 \end{bmatrix}, \qquad \mathcal{Q}_C = \{q \in \Delta_n : G_C q \leq 0\}. \tag{A.4}$$

If q satisfies every recurrence at $(t,\delta) \in C$, the first block is $(t - t_u)V_1 q + (\delta_l - \delta)V_0 q \leq 0$, and the second is $(t_l - t)V_1 q + (\delta - \delta_u)V_0 q \leq 0$. Thus the exact laws at those coefficients belong to $\mathcal{Q}_C$. Moreover, $C' \subseteq C$ implies $\mathcal{Q}_{C'} \subseteq \mathcal{Q}_C$: tightening endpoints adds nonnegative multiples of nonnegative prefix probabilities to the left sides of the inequalities.

A rectangle intersects the feasible region exactly when

$$C \cap \mathcal{R} \neq \emptyset \quad \Longleftrightarrow \quad \delta_u \geq t_l - 1 \quad \text{and} \quad \delta_l \leq t_u^2/4. \tag{A.5}$$

Necessity follows because both boundary functions increase on $[0,2]$. For sufficiency, set $t_* = \min(t_u, \delta_u + 1)$. The first inequality gives $t_* \geq t_l$. Both arguments of the minimum are nonnegative, and

$$t_u^2/4 \geq \delta_l, \qquad (\delta_u + 1)^2/4 \geq \delta_u \geq \delta_l.$$

Hence $t_*^2/4 \geq \delta_l$. Taking $\delta_* = \max(\delta_l, t_* - 1)$ gives $\delta_* \leq \delta_u$ and $\delta_* \leq t_*^2/4$, so $(t_*, \delta_*) \in C \cap \mathcal{R}$. Equality boundaries must therefore be retained.

For the original 16-bin cover, $i, j \in \{0, \dots, 15\}$ specify $t \in [i/8, (i+1)/8]$ and $\delta \in [(j-8)/8, (j-7)/8]$. Equation (A.5) discards a cell precisely when $j + 1 < i$ or $32(j-8) > (i+1)^2$, leaving 80 cells. Multiplying (A.4) by eight gives the integer matrix

$$A_{ij} = \begin{bmatrix} 8V_2 - (i+1)V_1 + (j-8)V_0 \\ -8V_2 + iV_1 - (j-7)V_0 \end{bmatrix}. \tag{A.6}$$

Since every model has a pair in $\mathcal{R}$, any complete retained cover $\mathcal{F}$ yields

$$\mathrm{dist}_{\mathrm{TV}}(p, \mathcal{H}_2^{(n)}) \geq \min_{C \in \mathcal{F}} \mathrm{dist}_{\mathrm{TV}}(p, \mathcal{Q}_C). \tag{A.7}$$

This proves the whole-class cover bound in Equation (7). No realization of a relaxed optimizer is needed.

Equation (A.8) records the optional coefficient clipping used in the earlier adaptive development study. The final prospective candidate pipeline leaves boxes unclipped; the complete-cover bounds (A.7) and (A.14) hold with or without clipping. The optional procedure intersects a box with D_0, rejects inverted intervals or either strict exclusion in (A.5), and performs four sweeps of

$$\begin{aligned} t_u &\leftarrow \min(t_u, \delta_u + 1, 2), \\ \delta_l &\leftarrow \max(\delta_l, t_l - 1, -1), \\ \delta_u &\leftarrow \min(\delta_u, t_u^2/4, 1), \\ t_l &\leftarrow \max\{t_l, 0, \lfloor \sqrt{4\max(\delta_l, 0)} \rfloor_{2^{-40}}\}. \end{aligned} \tag{A.8}$$

Each update uses the latest endpoints and is followed by the same exact emptiness checks. Every feasible point obeys $t \leq \delta + 1$, $\delta \geq t - 1$, $\delta \leq t^2/4$, and, when $\delta_l > 0$, $t \geq \sqrt{4\delta_l}$. Each update therefore preserves $C \cap \mathcal{R}$. For $\max(\delta_l, 0) = u/v$, the last bound is computed as $2^{-40}\,\mathrm{isqrt}(\lfloor 4u2^{80}/v \rfloor)$. It is at most the true root and less than 2^{-40} below it. Rounding upward could remove a feasible boundary point. Four sweeps need not produce the smallest enclosing rectangle; validity requires preservation of the feasible intersection.

A.3 Dual witnesses with exact repair

For any cell matrix G, the TV relaxation is the linear program

$$\begin{aligned} \text{minimize} \quad & \tfrac{1}{2}\mathbf{1}^\mathsf{T} s \\ \text{subject to} \quad & q - s \leq p, \quad -q - s \leq -p, \quad Gq \leq 0, \\ & \mathbf{1}^\mathsf{T} q = 1, \quad q, s \geq 0. \end{aligned} \tag{A.9}$$

Nonnegative multipliers a, b, y, with a free normalization multiplier ν, give the Lagrangian

$$\nu + (b-a)^\mathsf{T} p + (a - b + G^\mathsf{T} y - \nu\mathbf{1})^\mathsf{T} q + (\tfrac{1}{2}\mathbf{1} - a - b)^\mathsf{T} s. \tag{A.10}$$

If $a + b \leq \mathbf{1}/2$ and $a - b + G^\mathsf{T} y \geq \nu\mathbf{1}$, its infimum over nonnegative q, s is $\nu + (b-a)^\mathsf{T} p$. Weak duality proves this lower bound even when the solver has stopped early or its proposed primal is infeasible.

The fixed-grid implementation proposes multipliers with HiGHS, then repairs them at denominator $D = 2^{40}$. It floors nonnegative proposed magnitudes times D, clips the first numerator a_j to $D/2$, clips the second b_j to $D/2 - a_j$, and retains nonnegative integer context numerators y. For an integer matrix A, it sets

$$\nu_{\mathrm{num}} = \min_j (a - b + A^\mathsf{T} y)_j, \qquad \ell_C = \max\left\{0, \frac{\nu_{\mathrm{num}} + (b-a)^\mathsf{T} p}{D}\right\}. \tag{A.11}$$

Here a, b, y denote integer numerators. Dividing by D satisfies both dual constraints exactly. Arbitrary-precision integer arithmetic checks every coordinate; rounding can weaken the result but cannot invalidate it.

Adaptive boxes have rational endpoints. When the proposed multipliers apply to the unscaled matrix G_C, let Q be a common denominator and $M = QG_C$ its integer matrix. The corresponding repair is

$$\nu_{\mathrm{num}} = \min_j \left(a_j - b_j + \frac{(M^\mathsf{T} y)_j}{Q}\right), \qquad \ell_C = \max\left\{0, \frac{\nu_{\mathrm{num}} + (b-a)^\mathsf{T} p}{D}\right\}. \tag{A.12}$$

The normalization numerator may now be rational. The division by Q remains essential because the multipliers were defined for G_C. A separate reader reconstructs the matrix and verifies all multiplier signs, $a_j + b_j \le D/2$, normalization inequalities, and target dot products. Solver optimality is not a premise of either formula.

A.4 Complete adaptive coverage

Define the exact recurrence fiber and its union by

$$A_\theta = V_2 - tV_1 + \delta V_0, \qquad K_\theta = \{x \in \Delta_n : A_\theta x = 0\},$$
$$E_0^{(n)} = \bigcup_{\theta \in \mathcal{R}} K_\theta. \tag{A.13}$$

Every fiber is nonempty: the coefficient construction above supplies an HMM law. Equations (A.3)–(A.4) give $\mathcal{H}_2^{(n)} \subseteq E_0^{(n)}$ and $K_\theta \subseteq \mathcal{Q}_C$ for $\theta \in C \cap \mathcal{R}$.

Start from a finite root cover of $\mathcal{R}$, with valid nonnegative lower bounds. A solved node receives the larger of its inherited bound and its new bound ℓ_C. A split replaces a nonpoint box, optionally clipped by (A.8), by both closed children, split at the exact midpoint of one side. Each child inherits the parent's bound. In the implemented policy, the longer side is split, with trace chosen on a tie; pending nodes are selected by inherited bound and node identifier. These scheduling choices affect progress, but not validity.

Let $\mathcal{F}$ be the current frontier: every nonsplit leaf except one proved geometry-empty. It includes pending leaves, leaves pruned because their lower bound exceeds an incumbent, and leaves stopped by a point, depth, budget, target-width or solver condition. An unsuccessful or unstarted solve keeps its inherited bound. Then

$$L = \min_{C \in \mathcal{F}} b_C \le \mathrm{dist}_{\mathrm{TV}}(p, E_0^{(n)}) \le \mathrm{dist}_{\mathrm{TV}}(p, \mathcal{H}_2^{(n)}). \tag{A.14}$$

To prove (A.14), any optional clipping preserves the feasible intersection, both closed children cover the resulting parent box, and only empty intersections are deleted. Induction therefore shows that the frontier covers $\mathcal{R}$. For any $x \in E_0^{(n)}$, choose an admitting θ and a frontier cell containing it. The cell inclusion gives $x \in \mathcal{Q}_C$, and its inherited and direct bounds give $\mathrm{TV}(p, x) \ge b_C \ge L$. Taking the infimum proves the claim. A split must record both children before removing its parent. Child inheritance makes the global lower bound nondecreasing, without guaranteeing that a prescribed width will be reached in finite computation.

A.5 Exact correction and the two kinds of upper bound

A probability vector satisfying (A.13) supplies an upper bound on distance to $E_0^{(n)}$. The following construction converts an approximate fixed-coefficient solution into such a vector without assuming that its rounded coordinates satisfy the cell inequalities.

Fix rational $\theta \in \mathcal{R}$ and a strictly positive rational anchor $r \in K_\theta$. Choose independent rows A_I spanning A_θ, put $B = [\mathbf{1}^\mathsf{T}; A_I]$, and choose columns J for which $M = B[:, J]$ is square and invertible. The rows of B are independent: evaluating a row dependence on r removes its normalization coefficient, and independence of A_I removes the rest. Such columns therefore exist. For an exactly normalized rational law q, solve

$$Mh = \begin{bmatrix} 0 \\ -A_I q \end{bmatrix}, \quad d = E_J h, \quad z = q + d, \quad \tau = \max\left\{0, \max_j \frac{-z_j}{r_j}\right\}, \quad x = \frac{z + \tau r}{1 + \tau}, \tag{A.15}$$

where E_J inserts the coordinates into positions J. The solve gives $\mathbf{1}^\mathsf{T} d = 0$ and $A_I d = -A_I q$. Row spanning then gives $A_\theta z = 0$. Each coordinate of $z + \tau r$ is nonnegative, its mass is $1 + \tau$, and its recurrence residual is zero. Hence $x \in K_\theta$. The chosen τ is the smallest nonnegative mixing parameter with this property.

The displacement satisfies

$$\mathrm{TV}(q, x) \le \frac{\|d\|_1/2 + \tau\,\mathrm{TV}(q, r)}{1 + \tau} \le \frac{\|d\|_1/2 + \tau}{1 + \tau}, \qquad U_{E_0} = \tfrac{1}{2} \sum_j |p_j - x_j|. \tag{A.16}$$

The first inequality follows by subtracting q from (A.15); the second uses $\mathrm{TV}(q, r) \le 1$. The affine vector z can be signed, so $\|d\|_1/2$ need not be a TV distance between probability laws. Coordinate clipping of z would generally destroy the recurrence. If z is already nonnegative, no anchor mixing is needed.

An interior root $\lambda \in (0, 1)$ of $\lambda^2 - t\lambda + \delta = 0$ supplies an iid anchor $r(w) = \lambda^{\#0(w)}(1-\lambda)^{\#1(w)}$. Alternatively, a strictly positive rational B_0 with row sums less than one supplies a rational anchor: use $\pi = (1/2, 1/2)$ and put half of each row's remaining mass in each entry of B_1. All word probabilities are positive, normalization follows from row-stochasticity, and (A.3) gives the required recurrence at its exact trace and determinant. Boundary fibers need not admit a strictly positive anchor.

The implementation rounds a saved floating primal at denominator 2^{48}, using nearest-even rounding of its exact stored binary values, clips negative numerators to zero, and normalizes by their positive sum. This supplies the input law q in (A.15). The completed-law reader checks normalization, nonnegativity and every original recurrence directly, so the verified membership of the final x does not depend on numerical rank metadata. Exact row spanning and nonsingularity justify the construction itself.

Combining (A.14) with (A.16) gives $[L, U_{E_0}]$ for the recurrence distance. A physical upper U_2 instead requires an explicitly verified HMM law q_2, with $U_2 \ge \mathrm{TV}(p, q_2)$, and gives $[L, U_2]$ for the HMM distance. Even when the anchor is an HMM law, the corrected vector need not be one. For the saved development target the two upper endpoints are, respectively, 0.372757400752161 and 0.378052374754519; their common lower endpoint is 0.350828193264852. Thus only the recurrence interval has width at most 0.025. These are outward displays of separately checked exact quantities.

A.6 Target arithmetic and probability enclosures

For an empirical target, integer word counts, eligible-anchor totals, speaker counts and recording counts define the speaker-balanced distribution as exact rational sums. A Fraction-based reader reconstructs those sums and checks normalization. The certificate reader reconstructs every context, root-cover decision, child partition and inherited bound in addition to the dual checks above. A partial cover cannot certify the global minimum.

Quantum error bounds use 40-digit outward-rounded interval arithmetic applied to the stored angles and declared channel probabilities. The evaluator reconstructs the initial state, six rotations, both symbol maps and all branch probabilities. Interval subtraction and absolute value enclose every contribution to TV; conversion to binary floating-point endpoints is rounded outward. The resulting upper bounds concern the instrument defined by those stored values. They involve no sampled estimate of its output probabilities. The finite-corpus lower and upper bounds are deterministic; bootstrap diagnostics address a different question.

B Information retained and realization limits

The coefficient cover tests whether a finite law admits a shared feasible recurrence. This appendix identifies its refinement limit, the information that sharing and feasibility add to ordinary rank constraints, and the positive-realization conditions that remain outside the test. Every statement fixes a finite horizon. The constructed laws are mathematical controls, not quantum or acoustic constructions.

B.1 The limit of coefficient refinement

Fix $n \ge 4$. Use Δ_n, $P_p(w)$, $\mathcal{C}_n$, $\mathcal{R}$ and $\mathcal{Q}_C$ from Appendix A, and define

$$\begin{aligned} r_{uv}(p; t, \delta) &= P_p(u00v) - tP_p(u0v) + \delta P_p(uv), \\ E_0 &= \{p \in \Delta_n : \exists \theta \in \mathcal{R},\ r_{uv}(p; \theta) = 0\ \forall (u, v) \in \mathcal{C}_n\}. \end{aligned} \tag{B.1}$$

At each level m, partition $D_0 = [0, 2] \times [-1, 1]$ into finitely many closed rectangles with shared boundaries, and retain exactly those intersecting $\mathcal{R}$. Assume each partition refines the preceding

one. Write $\mathcal{F}_m$ for its retained cells, h_m for their largest side width, and

$$E^{(m)} = \bigcup_{C \in \mathcal{F}_m} \mathcal{Q}_C, \qquad d_m(p) = \mathrm{dist}_{\mathrm{TV}}(p, E^{(m)}), \qquad d_\infty(p) = \mathrm{dist}_{\mathrm{TV}}(p, E_0). \tag{B.2}$$

The exact intersection test is (A.5). Retaining only intersecting cells is a substantive condition: covering $\mathcal{R}$ with shrinking cells would not suffice if unrelated cells elsewhere in D_0 were also retained.

Proof of Proposition 2. Proposition 1 gives $\mathcal{H}_2^{(n)} \subseteq E_0$. A law in E_0 belongs to every cell relaxation containing its coefficient pair. A child cell strengthens its parent's inequalities, so

$$E_0 \subseteq E^{(m+1)} \subseteq E^{(m)}. \tag{B.3}$$

For $q \in \mathcal{Q}_C$, the two endpoint inequalities imply, at every common pair $(t, \delta) \in C$,

$$|r_{uv}(q; t, \delta)| \le (t_u - t_l) P_q(u0v) + (\delta_u - \delta_l) P_q(uv) \le 2h_m. \tag{B.4}$$

If $q \in \bigcap_m E^{(m)}$, choose at each level a containing cell and a pair $\theta_m \in C \cap \mathcal{R}$. Compactness of $\mathcal{R}$ supplies a convergent subsequence. Since $h_m \to 0$, continuity of the finitely many residuals makes all of them zero at the limit. Thus $\bigcap_m E^{(m)} = E_0$.

Each $E^{(m)}$ is compact, as a finite union of closed subsets of the simplex. The set E_0 is compact as the projection of a closed subset of $\Delta_n \times \mathcal{R}$. They are nonempty, since they contain HMM laws. Put $\eta_m = \max_{q \in E^{(m)}} \mathrm{dist}_{\mathrm{TV}}(q, E_0)$. If η_m did not tend to zero, a subsequence of maximizing laws would converge to a point a positive distance from E_0. Nesting and closedness put that limit in every $E^{(j)}$, contradicting their intersection. Nearest points and the triangle inequality therefore give

$$0 \le d_\infty(p) - d_m(p) \le \eta_m \longrightarrow 0 \quad \text{uniformly over } p \in \Delta_n. \tag{B.5}$$

Nesting makes $d_m(p)$ nondecreasing.

For the quantitative part, set

$$r_\star(p) = \min_{(t,\delta) \in \mathcal{R}} \max_{(u,v) \in \mathcal{C}_n} |r_{uv}(p; t, \delta)|. \tag{B.6}$$

Every prefix probability is an event probability. For $(t, \delta) \in D_0$,

$$|r_{uv}(p; t, \delta) - r_{uv}(q; t, \delta)| \le (1 + t + |\delta|)\, \mathrm{TV}(p, q) \le 4\, \mathrm{TV}(p, q). \tag{B.7}$$

For $q \in \mathcal{Q}_C$, select one pair in $C \cap \mathcal{R}$ and use it in every context. Equations (B.4) and (B.7) give $r_\star(p) \le 2h_m + 4\, \mathrm{TV}(p, q)$. Taking the infimum proves

$$d_m(p) \ge \frac{(r_\star(p) - 2h_m)_+}{4}. \tag{B.8}$$

This completes the proof. □

The limit is distance to E_0. The proof supplies no rate for (B.5), and it does not assert convergence of numerical dual witnesses: a valid repaired witness can have additional optimization slack. Equation (B.8) is a sufficient condition for a positive exact relaxation distance, not an error bound relative to the physical model class. The factor $2h_m$ uses a feasible pair; a cell's lower corner need not lie in $\mathcal{R}$.

B.2 Rank and a shared feasible recurrence

For $n \ge 3$ and $s \in \{0, 1\}$, let $E_s^{(n)}$ contain the laws for which one pair $(t_s, \delta_s) \in \mathcal{R}$ satisfies

$$P_p(ussv) - t_s P_p(usv) + \delta_s P_p(uv) = 0 \quad \text{whenever } |u| + |v| \le n - 2. \tag{B.9}$$

The pair is shared across contexts, but the two symbols may have different pairs. A compatible prefix-Hankel table is $[P_p(uv)]_{u \in U, v \in V}$ with $|uv| \le n$ for every entry. Every two-state HMM factors such a table through two dimensions. The next construction satisfies all these rank conditions while failing (B.9).

Proposition 3 (shared feasible coefficients retain information beyond rank). Fix $n \ge 3$, $0 < b \le 1/(4n)$, and $z = 1/2 + ib$. For a word w of length $\ell \le n$ containing k zeros, put

$$c(w) = z^k \bar{z}^{\ell - k}, \qquad P(w) = \mathrm{Re}\, c(w), \qquad p_{n,b}(w) = P(w) \quad (|w| = n). \tag{B.10}$$

This defines a strictly positive law. Every compatible prefix-Hankel table has rank at most two, and one has rank exactly two. Each context separately admits a feasible pair, and every context shares one unrestricted pair. Nevertheless, an explicit positive $q_{n,b} \in E_0^{(n)} \cap E_1^{(n)}$ satisfies

$$\frac{b^3}{4} \le \mathrm{dist}_{\mathrm{TV}}(p_{n,b}, E_0^{(n)}) \le \mathrm{TV}(p_{n,b}, q_{n,b}) \le 4n(n-1)^2 b^3. \tag{B.11}$$

The same lower bound holds for distance to $\mathcal{H}_2^{(n)}$. For fixed n, distance to either recurrence set $E_0^{(n)}$ or $E_0^{(n)} \cap E_1^{(n)}$ is therefore $\Theta(b^3)$; the upper bound in (B.11) is not an HMM upper bound.

Proof. Write $\theta = \arg z = \arctan(2b)$. For $\ell \le n$, $|\arg c(w)| \le \ell\theta \le 2nb \le 1/2$, so $P(w) \ge (7/8)|z|^\ell > 0$. Since $z + \bar{z} = 1$, the identity $P(w0) + P(w1) = P(w)$, starting at $P(\emptyset) = 1$, proves normalization and identifies these quantities as the prefix probabilities of $p_{n,b}$.

For every compatible pair of words,

$$P(uv) = \mathrm{Re}\, c(u)\, \mathrm{Re}\, c(v) - \mathrm{Im}\, c(u)\, \mathrm{Im}\, c(v). \tag{B.12}$$

This simultaneous factorization proves the rank upper bound. The minor indexed by $U = V = \{\emptyset, 0\}$ has determinant $P(00) - P(0)^2 = -b^2$, proving exact rank two. Both z and $\bar{z}$ solve $x^2 - x + (1/4 + b^2) = 0$, so both symbols share the unrestricted pair $(1, 1/4 + b^2)$, which violates $\delta \le t^2/4$. On the other hand, each individual symbol-s context admits

$$(t_{u,v,s}, \delta_{u,v,s}) = \left(\frac{P(ussv)}{P(usv)}, 0 \right) \in \mathcal{R}. \tag{B.13}$$

Indeed, the law depends only on symbol counts, so $P(ussv) = P((usv)s)$; its denominator is the sum of two positive extensions. The ratio is strictly between zero and one. The construction is therefore admitted if either sharing or feasibility is dropped.

To prove the lower bound, for an arbitrary candidate law put $m_j = P(0^j)$, $m_0 = 1$, and define

$$D = m_2 - m_1^2, \quad T = m_3 - m_1 m_2, \quad N = m_1 m_3 - m_2^2, \quad F = T^2 - 4DN. \tag{B.14}$$

The first two survival equations of a law in $E_0^{(n)}$ give $m_2 = tm_1 - \delta$, $m_3 = tm_2 - \delta m_1$. Substitution yields $T = tD$, $N = \delta D$, and hence

$$F = (t^2 - 4\delta)D^2 \ge 0. \tag{B.15}$$

No division by D is used, so singular cases are included. At $p_{n,b}$, $(m_1, m_2, m_3) = (1/2, 1/4 - b^2, 1/8 - 3b^2/2)$ and $F = -4b^6$.

Introduce centered coordinates $y_1 = m_1 - 1/2$, $y_2 = m_2 - m_1 + 1/4$, and $y_3 = m_3 - 3m_2/2 + 3m_1/4 - 1/8$. Expanding (B.14) gives

$$F = y_3^2 - 6y_1y_2y_3 + 4y_1^3y_3 + 4y_2^3 - 3y_1^2y_2^2. \tag{B.16}$$

The target coordinates are $(0, -b^2, 0)$. In a closed TV ball of radius $b^3/4$, each prefix mass changes by at most that radius, so $|y_1| \le b^3/4$, $|y_2 + b^2| \le b^3/2$, and $|y_3| \le 13b^3/16$. Using $b \le 1/4$, we have $y_2 \le -7b^2/8$, $|y_2| \le 9b^2/8$. Bounding the first three terms absolutely and dropping the last, nonpositive term gives

$$\begin{aligned} F &\le \frac{169}{256}b^6 + \frac{351}{256}b^8 + \frac{13}{256}b^{12} - \frac{343}{128}b^6 \\ &\le -\frac{2027763}{1048576}b^6 < 0. \end{aligned} \tag{B.17}$$

This contradicts (B.15) throughout the ball. Perturbed laws need not remain exchangeable or rank two. The inclusion $\mathcal{H}_2^{(n)} \subseteq E_0^{(n)}$ gives the HMM lower bound as well.

For the upper witness, let $r = 1/2 + b$, $f_w(r) = r^k(1-r)^{\ell-k}$, and define

$$Q(w) = f_w(r) - bf_w'(r), \qquad q_{n,b}(w) = Q(w) \quad (|w| = n). \tag{B.18}$$

The derivative acts on f_w, with b fixed. Differentiating the Bernoulli normalization and prefix identities proves those identities for Q. Also,

$$\frac{Q(w)}{f_w(r)} = 1 - b\left(\frac{k}{r} - \frac{\ell - k}{1 - r} \right) \ge 1 - \frac{bn}{r} \ge \frac{1}{2} > 0.$$

For a context put $A(r) = f_u(r)f_v(r)$. Then $Q(u0^j v) = (A - bA')r^j - bAjr^{j-1}$, which obeys the repeated-root recurrence with coefficients $(2r, r^2)$. Inserted ones give $(2(1-r), (1-r)^2)$. Both pairs lie in $\mathcal{R}$, since $0 < r < 1$ and $r^2 - (2r-1) = (1-r)^2 \geq 0$. Thus $q_{n,b} \in E_0^{(n)} \cap E_1^{(n)}$.

Write $f_w(1/2 + x) = \sum_{j=0}^{n} A_j(w)x^j$ for terminal words. The coefficients through degree two cancel between (B.10) and (B.18), leaving

$$p_{n,b}(w) - q_{n,b}(w) = \sum_{j=3}^{n} c_j A_j(w) b^j, \qquad c_j = j - 1 + \mathrm{Re}(i^j), \quad 0 \leq c_j \leq j.$$

Expanding the n factors $(1/2 \pm x)$ gives $\sum_{|w|=n} |A_j(w)| \leq \binom{n}{j} 2^j$. Consequently,

$$\begin{aligned} \mathrm{TV}(p_{n,b}, q_{n,b}) &\leq \tfrac{1}{2} \sum_{j=3}^{n} j \binom{n}{j} (2b)^j \\ &= nb\big[(1+2b)^{n-1} - 1 - 2(n-1)b\big] \\ &\leq 2n(n-1)^2 b^3 e^{2(n-1)b} < 4n(n-1)^2 b^3. \end{aligned} \tag{B.19}$$

The penultimate bound uses $\binom{m}{k} \leq m^k/k!$ and $e^x - 1 - x \leq x^2 e^x/2$; the last uses $2(n-1)b < 1/2$. Equations (B.17)–(B.19) prove the proposition. □

B.3 What the rank comparison does not imply

Rank and the original-symbol recurrence constraints have no general nesting. To see the other direction explicitly, choose $0 < a, b < 1$, $a \neq b$, put $c = 1 - a$, $d = 1 - b$, and consider

$$B_0 = \mathrm{diag}(a, a, b), \qquad B_1 = \begin{pmatrix} 0 & c & 0 \\ 0 & 0 & c \\ d & 0 & 0 \end{pmatrix}, \qquad \pi = (1, 0, 0). \tag{B.20}$$

Its length-n law belongs to $E_0^{(n)}$ because $B_0^2 - (a+b)B_0 + abI = 0$, and $(a+b, ab) \in \mathcal{R}$. These are annihilator coefficients, not the trace and determinant of the three-state matrix. For $n \geq 4$, prefixes $(\emptyset, 1, 11)$ and suffixes $(\emptyset, 0, 10)$ give

$$H = \begin{pmatrix} 1 & a & ca \\ c & ca & c^2 b \\ c^2 & c^2 b & c^2 da \end{pmatrix}, \qquad \det H = -c^4(b-a)^2 \neq 0. \tag{B.21}$$

Subtracting c times the first row from the second and c^2 times the first from the third verifies the determinant. Every two-state model factors this table through two dimensions, so the law is outside $\mathcal{H}_2^{(n)}$ despite its exact no-event recurrence. Together with Proposition 3, this establishes both directions for ordinary rank constraints and the original-symbol recurrence set. It does not compare the recurrence test with every positivity or moment criterion.

The upper witness in Proposition 3 is also outside $\mathcal{H}_2^{(n)}$. Its moments have $m_1 = 1/2$, $D = -b^2$, $T = 2rD$, and $N = r^2 D$. Since $D \neq 0$, the first two recurrence equations force any matching two-state B_0 to have trace $2r$ and determinant r^2. For a nonnegative $B_0 = \left(\begin{smallmatrix} a & c \\ d & e \end{smallmatrix}\right)$,

$$0 = (\mathrm{tr}\, B_0)^2 - 4 \det B_0 = (a - e)^2 + 4cd \tag{B.22}$$

forces $a = e = r$ and $cd = 0$. Thus $B_0 = rI + N_0$ with $N_0 \geq 0$, giving $P(0) = \pi B_0 \mathbf{1} \geq r > 1/2$, a contradiction. The cubic recurrence upper therefore supplies no matching physical upper.

Neither family in Proposition 3 has a homogeneous qubit realization for $n \geq 5$. We give the finite-equivalence argument to make this restriction explicit (Li & Qiu, 2006). A scalar series $\alpha M_w \beta$ of dimension D that vanishes on all words of length at most $D - 1$ vanishes identically. Indeed, let V_k span $M_w \beta$ for $|w| \leq k$. If $V_{k+1} = V_k$, that space is invariant under both symbol matrices and is the full reachable span. Otherwise its dimension increases. Starting from a nonzero β, stabilization occurs by $D - 1$, and α annihilates the resulting span. If $\beta = 0$, the assertion is immediate.

The formal continuation of (B.10) has real linear dimension two, using real and imaginary coordinates. A fixed instrument on Hilbert dimension d has real linear dimension at most d^2, from its

action on Hermitian matrices (Monras & Winter, 2014). Their difference has dimension at most $d^2 + 2$. Agreement on a length-n law implies agreement on every prefix, so $n \geq d^2 + 1$ would force agreement at all lengths. Yet with $\theta = \arg z > 0$ and $k = \lfloor \pi/(2\theta) \rfloor + 1$, the continuation has $P(0^k) = |z|^k \cos(k\theta) < 0$. A physical instrument cannot produce this value. The continuation of (B.18) also has dimension at most two, by propagating a product and its derivative, and $Q(0^k) = r^{k-1}(r - bk)$ is eventually negative. Thus either finite family requires $d^2 \geq n$ in any homogeneous quantum realization; the analogous necessary condition for an h-state classical HMM is $h \geq n$. These conditions assert no existence at the indicated dimensions. Time-dependent control changes the comparator.

B.4 A POSITIVE-REALIZATION OBSTRUCTION AFTER BOTH SYMBOL TESTS

Checking both symbols still leaves a gap between linear representations and stochastic realizations. The next family passes both shared feasible recurrences and every compatible rank-two Hankel test. Its obstruction comes from a continuation forced by its first four survival probabilities.

Proposition 4 (finite-horizon realization gap). For every $n \geq 3$, there is a strictly positive law in $E_0^{(n)} \cap E_1^{(n)}$ outside $\mathcal{H}_2^{(n)}$. Every compatible prefix-Hankel table has rank at most two, with one of rank exactly two. At $n = 3$, such a law has an excluded closed TV ball of radius $1/2000$.

Proof of the construction and nonrealizability. Choose $0 < a < b < 1$ and $(a/b)^{n+1} < r < (a/b)^n$. For a word of length $k \leq n$ containing z zeros, put

$$F(w) = \frac{a^z(1-a)^{k-z} - rb^z(1-b)^{k-z}}{1-r}, \qquad p_n(w) = F(w) \quad (|w| = n). \tag{B.23}$$

The ratio of the first product to the second is at least $(a/b)^k \geq (a/b)^n > r$, so all these values are positive. The binomial theorem gives terminal mass one, and $F(w0) + F(w1) = F(w)$ identifies the prefixes. Define

$$A_0 = \mathrm{diag}(a, b), \quad A_1 = \mathrm{diag}(1-a, 1-b), \quad \alpha = (1, -r)/(1-r), \quad \beta = \mathbf{1}. \tag{B.24}$$

Then $F(w) = \alpha A_w \beta$, a signed linear representation. Both coefficient pairs $(a + b, ab)$ and $(2 - a - b, (1-a)(1-b))$ lie in $\mathcal{R}$: for $x, y \in [0, 1]$, $xy - (x + y - 1) = (1-x)(1-y) \geq 0$ and $(x+y)^2/4 - xy = (x-y)^2/4 \geq 0$. Multiplying each symbol's quadratic identity by αA_u and $A_v \beta$ proves membership in both recurrence sets. The factorization $F(uv) = (\alpha A_u)(A_v \beta)$ proves every compatible rank upper bound.

Let $m_k = F(0^k) = (a^k - rb^k)/(1-r)$. The minor $m_0 m_2 - m_1^2 = -r(a-b)^2/(1-r)^2 < 0$ proves exact rank two and makes the system

$$\begin{pmatrix} m_1 & -m_0 \\ m_2 & -m_1 \end{pmatrix} \begin{pmatrix} t \\ \delta \end{pmatrix} = \begin{pmatrix} m_2 \\ m_3 \end{pmatrix} \tag{B.25}$$

invertible. Every matching two-state HMM must satisfy this system, since $n \geq 3$, so its unique coefficients are $(t, \delta) = (a + b, ab)$. Cayley–Hamilton then forces every later survival probability from m_0, m_1, including

$$\pi B_0^{n+1} \mathbf{1} = \frac{a^{n+1} - rb^{n+1}}{1-r} < 0. \tag{B.26}$$

This contradicts nonnegativity of any HMM, without imposing exchangeability on that hypothetical model. The finite law depends on n; its signed continuation is not a valid infinite process. □

The HMM parameter space is compact, and the map to a finite word law is continuous. The constructed law therefore has positive distance from $\mathcal{H}_2^{(n)}$. An explicit margin follows from a polynomial condition that avoids solving (B.25). For any two-state HMM, let $s_j = \pi B_0^j \mathbf{1}$, set $k = n - 3$, $m_i = s_{k+i}$, and define

$$D = m_0 m_2 - m_1^2, \qquad V = m_0 m_3^2 - 2 m_1 m_2 m_3 + m_2^3. \tag{B.27}$$

The matrix $[m_{i+j}]_{i,j=0}^2$ factors through two dimensions, hence its determinant is $Dm_4 - V = 0$. Nonnegativity of the unobserved continuation $m_4 = s_{n+1}$ gives

$$DV = D^2 m_4 \geq 0. \tag{B.28}$$

This includes $D = 0$, when the determinant identity also forces $V = 0$. For the signed construction, $D = -r(ab)^k(a-b)^2/(1-r)^2 < 0$ and $V = DF(0^{n+1}) > 0$.

If a candidate law is within TV distance ε of the target, set $l_i = \max(0, m_i - \varepsilon)$, $u_i = \min(1, m_i + \varepsilon)$, with $l_0 = u_0 = 1$ when $k = 0$. Each observed prefix mass lies in its interval. Nonnegativity then gives

$$D \le D_+ = u_0u_2 - l_1^2, \qquad V \ge V_- = l_0l_3^2 - 2u_1u_2u_3 + l_2^3. \tag{B.29}$$

Thus $D_+ < 0 < V_-$ excludes the entire closed TV ball by (B.28). For $n = 3$, $a = 1/4$, $b = 3/4$, $r = 2/81$, the target masses are $(1, 75/316, 63/1264, 27/5056)$. At $\varepsilon = 1/2000$, exact substitution gives

$$D_+ = -\frac{143599241}{24964000000} < 0, \qquad V_- = \frac{7300628629}{1972156000000000} > 0. \tag{B.30}$$

This proves the quantitative assertion of Proposition 4. No rank or recurrence condition is imposed on perturbed laws in that ball.

For the same signed representation, every product A_h is diagonal with entries in $[0, 1]$; its quadratic annihilator also has coefficients in $\mathcal{R}$. Thus adding all available contextual recurrences for repeated words does not remove this obstruction. Signed finite mixtures and the distinction between linear and positive realizations are established (Janson et al., 2016; Monras & Winter, 2014). Here the forced negative continuation excludes arbitrary two-state edge-emitting HMMs at a specified finite horizon. The construction supplies neither a quantum implementation nor a tolerance uniform in the horizon. Refining a cover of necessary recurrence conditions cannot by itself close this realization gap.

C Acoustic protocol and complete outcomes

C.1 Frontend, weights and corpus roles

Speech Commands v0.02 digit recordings are decoded as mono PCM16 at 16 kHz, with signed integers divided by 32,768 (Warden, 2018). Lengths and amplitudes are unchanged: there is no padding, cropping, clipping, energy gate or normalization. A causal Butterworth bandpass from 750 to 1250 Hz uses prototype order two, producing a fourth-order filter in two second-order sections, initialized to zero for each recording. A positive crossing between adjacent filtered samples $y_k \le 0 < y_{k+1}$ occurs at $k - y_k/(y_{k+1} - y_k)$. Each crossing anchors eight open-left, closed-right bins of four samples. The anchor itself is excluded, and any positive number of subsequent crossings in a bin produces one event. Eligibility requires the entire 32-sample block after trimming 1,600 samples at each recording edge. Right-sided searches determine boundary membership without a numerical tolerance adjustment.

For eligible recording r of speaker s, let $c_{sr}(w)$ count anchors producing word w, and let N_{sr} be the number of anchors. With S eligible speakers and R_s eligible recordings for speaker s, the exact target is

$$p(w) = \frac{1}{S}\sum_{s=1}^{S}\frac{1}{R_s}\sum_{r=1}^{R_s}\frac{c_{sr}(w)}{N_{sr}}. \tag{C.1}$$

Integer counts make this a rational distribution. Pooling anchors would change the estimand by giving more weight to recordings with more crossings. Overlapping anchors are dependent; digit labels select the recordings and are not prediction targets. A recording with no eligible anchors, and any resulting empty speaker, would be excluded and reported. Every selected recording is eligible.

Table C.1: Distinct data roles. Confirmation contains 3,046,061 overlapping anchors; transfer contains 12,352,293. Training and development belong to the official training partition. The three Speech Commands cohorts are disjoint by path, speaker identifier and waveform hash.

Cohort	Recordings	Speakers	Role
Speech Commands training	3,000	871	Fitting and selection
Speech Commands development	808	197	First failed test; subsequently exposed
Speech Commands confirmation	3,999	241	Frozen official-test evaluation
AudioMNIST transfer	29,997	60	Frozen external evaluation

The first prospective analysis fixed all fits before decoding development. Its older 642-context, 64-cell certificate gave floor 0.1866060913 against primary-qubit TV 0.1981519336: the gap was -0.0115458423, with conditional interval $[-0.0222756595, -0.0052683718]$. It failed. Applying the new recurrence certificate to this exposed cohort raised the floor to 0.2181172199 without changing the model. This is a post hoc development result.

For confirmation, the parameters, frontend, noise, 769 contexts, 16-grid cover, selection rules and scientific source were frozen before official-test decoding. A metadata-only proposal using unused official-validation recordings failed its minimum of 51 speakers: only 134 deduplicated recordings from 30 speakers remained. The subsequent official-test protocol was fixed before decoding. Its 4,107 digit candidates passed prior-exposure filters; lexical-first deduplication within same-label hash groups removed 108 copies, leaving 3,999 recordings without label conflicts or speaker/hash overlap with prior data. The prepared manifest records zero decoded test waveforms and zero fits. Compatibility keys containing `validation_` or `posthoc` in the inherited engine do not determine the cohort's role; the outer protocol and explicit official-test field do.

C.2 Models, optimization and selection

The qubit uses the real instrument in the main text. With memory as the inner index, six Givens rotations left-multiply I_4 in planes $(0,1),(0,2),(0,3),(1,2),(1,3),(2,3)$. Each block has cosine on its diagonal, minus sine above it, and sine below it. Rows $2x, 2x+1$, columns 0,1 define K_x. Six rotation angles and the initial-state angle give seven stored parameters. The two-parameter oscillator starts in $|0\rangle$ and uses

$$K_0 = \mathrm{diag}(1, \sqrt{1-\gamma}) R_y(\theta), \qquad K_1 = \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix} R_y(\theta). \tag{C.2}$$

Its parameters are θ and a logit for γ. Both quantum families apply partial dephasing and independent report flips at the declared rate; 1% dephasing multiplies off-diagonal entries by 0.99. A d-state HMM uses row softmaxes over $2d$ joint symbol/next-state outcomes and a separate initial softmax, storing $2d^2 + d$ logits: 10, 21 and 36 for $d = 2, 3, 4$. Removing softmax shifts leaves $2d^2 - 1$ stochastic coordinates. The certified HMM class also includes zero-probability boundary models, which finite-logit fits cannot attain.

The eight recipes each use seeds 66301–66308. Qubit angles and HMM logits start from independent normal draws with standard deviation 0.8; the four qubit recipes share starts seed by seed. The oscillator draws two values at this scale, then its angle uniformly on $[0.2, \pi - 0.2]$. There are no held-out warm starts. Cross-entropy minimizes $-\sum_w p_{\mathrm{tr}}(w) \log \max(q(w), 10^{-300})$; TV training minimizes

$$\frac{1}{2} \sum_w \left[\sqrt{(q(w) - p_{\mathrm{tr}}(w))^2 + 10^{-12}} - 10^{-6} \right]. \tag{C.3}$$

L-BFGS-B uses Torch float64 gradients, `maxiter=1000`, `ftol=1e-12`, `gtol=1e-8` and `maxls=50`. Selection minimizes final cross-entropy for CE recipes and final **unsmoothed** TV

for TV recipes, at the recipe's training noise. All 64 fits remain, including 14 capped fits. Convergence flags can reflect relative objective reduction and do not establish a stationary point or global optimum.

Table C.2: Training-only selections. Training TV is evaluated at training noise, so clean/noisy rows are not a common noisy evaluation. Both selected HMM3 and HMM4 reach the 1,000-iteration cap. The primary recipe has six solutions near TV 0.193575 and two poor solutions, 0.585807 and 0.561840, despite eight convergence flags. Every start, parameter vector, objective trace and termination message is in `results/memory-fidelity-v1/fit/`; immutable revision16 Tables B.1–B.3 give the complete training summaries.

Recipe	Selected seed	Training TV	Converged / capped
Qubit, CE / clean	66307	0.196696	8 / 0
Qubit, CE / 1%	66308	0.206623	8 / 0
Qubit, TV / clean	66308	0.179739	8 / 0
Qubit, TV / 1%	66301	0.193575	8 / 0
Oscillator, TV / 1%	66303	0.343166	8 / 0
HMM2, TV	66303	0.375317	8 / 0
HMM3, TV	66301	0.209202	2 / 6
HMM4, TV	66305	0.087656	0 / 8

C.3 Confirmation, secondary models and noise

The primary rule requires the noisy-TV qubit at 1% noise to have both its certified gap and bootstrap 2.5% gap quantile above 10^{-6}, with at least 90% eligible recordings. It passes on confirmation:

$$\mathrm{TV}(p_{\mathrm{SC}}, q_{0.01}) \le 0.187878493 < 0.204525522 \le \inf_{h \in \mathcal{H}_2} \mathrm{TV}(p_{\mathrm{SC}}, h). \tag{C.4}$$

These endpoints are rounded outward. All other model/noise comparisons are secondary.

Table C.3: Every selected model. Quantum/oscillator evaluations use 1% partial dephasing and report flips; HMMs have no added noise. These rounded achieved errors are distinct from certified classical floors. HMM4 fits Speech Commands better than the qubit despite its cap; the certificate is a comparison against HMM2.

Selected model	Development TV	Confirmation TV	AudioMNIST TV
Qubit, CE / clean	0.223500	0.208658	0.153740355
Qubit, CE / 1%	0.211603	0.198644	0.146730214
Qubit, TV / clean	0.202675	0.190984	0.139396329
Primary qubit, TV / 1%	0.198152	0.187878	0.138026097
Oscillator, TV / 1%	0.347425	0.333816	0.244598944
HMM2, TV	0.378279	0.365027	0.265773171
HMM3, TV	0.214038	0.208990	0.184163334
HMM4, TV	0.093146	0.101113	0.168626896

Each of 2,000 whole-speaker bootstrap replicates resamples equally weighted speakers, preserving their recording-normalized histograms. The reader evaluates all 80 frozen dual witnesses, takes their minimum clipped at zero, and subtracts the fixed candidate's TV on the same resampled law. It performs no fitting or LP reoptimization. Seeds are 72999 for confirmation, 71999 for the post hoc development certificate, 70999 for the failed original test, and 78999 for transfer. Percentile intervals describe conditional stability of the frozen analysis; they are not finite-sample guarantees for an unknown speaker population.

Table C.4: All three declared noise settings for the unchanged primary model. A negative lower-bound gap is an unresolved comparison, not proof of classical superiority. Changing dephasing and flips together does not isolate coherence loss.

Corpus	Noise	Certified gap, approx.	Conditional 95% interval
Confirmation	0%	0.022355931	[0.014921738, 0.027410104]
Confirmation	1%	0.016647030	[0.012555471, 0.019540439]
Confirmation	5%	−0.083501223	[−0.087175300, −0.081422408]
AudioMNIST	0%	−0.031352183	Secondary failure
AudioMNIST	1%	−0.016314455	[−0.025370796, −0.011119053]
AudioMNIST	5%	−0.091476489	Secondary failure

The older padded-clock certificate already separates confirmation: its floor is approximately 0.196644706. Recurrence's higher floor therefore strengthens this acoustic bound but is unnecessary for the validation decision. The fixed cross-condition comparison and its adverse cases are retained in Appendix E.

C.4 External transfer and hardware failures

AudioMNIST is pinned at repository revision `630d7dab4c040882834de6fa21baf9a60372accd` (Becker et al., 2023). Lexical-first Git-blob deduplication removes three same-label copies from 30,000 files. All remaining mono PCM16 files are decoded at 48 kHz and converted to float64/32,768, then resampled with SciPy's `resample_poly` with up/down factors 1/3, Kaiser parameter 5.0 and constant padding. The centered FIR uses future samples, so the full transfer frontend is noncausal. There is no subsequent quantization, clipping or normalization; the downstream filter and event extraction are unchanged. The selected inventory totals 5.35434 hours. No model or frontend is retuned. The two-state floor is approximately 0.121711641, below primary-qubit error 0.138026097; both primary gap requirements fail despite 100% eligibility. Speaker identifiers do not rule out human-identity overlap across corpora, and demographic performance is unassessed.

The separate Tuna-17 hardware pilot executes a four-bit synthetic oscillator with $\theta = \pi/2$, $\gamma = 3/4$, initially $|0\rangle$. It does not execute the learned acoustic instrument. Each histogram has 4,096 shots.

Table C.5: All pilot cases. Persistent-memory raw TV is 0.2841796875; **0.202496 is an L_2 error**, not TV. The streaming histogram is within L_2 0.017510 of the ideal cleared-memory law and 0.024974 of its measured histogram. These observations are consistent with memory degradation without isolating its mechanism.

Pilot case	Observation	Frozen criterion
Reset ordering	3,611/4,096 correct; 88.1592%	Accuracy >90%: fails
Persistent memory	Target L_2 error 0.202496	$L_2 < 0.02$: fails
Clear both qubits	Own-target L_2 error 0.015356	Diagnostic control

Raw-shot export was disabled. Counts are reconstructed uniquely from displayed rounding intervals and sum to 4,096; register order is reversed once into chronological order and checked by the deterministic control. An empty-source submission failed before execution and remains recorded. Even an ideal histogram at this shot count would fail the toy study's conservative 99% certificate: concentration radius 0.049156 exceeds its rank-tail floor 0.030035. More shots alone would not remove the observed bias.

C.5 VERIFICATION AND RETAINED RECORDS

The acoustic records are `results/memory-fidelity-v1/`, `results/event-context-memory-v1/`, `results/memory-confirmation-v1/` and `results/audiomnist-memory-transfer-v1/`; hardware sources, reconstructed counts and failed gates are in `results/memory-hardware-v1/`. The confirmation's separate rational reader rebuilds the target, 769 contexts, all 256 grid decisions and 80 integer witnesses. It checks supplied interval inclusions; the transcendental interval engine is not independently reimplemented. Noisy four-symbol Qiskit checks agree with direct recurrence to 3.89×10^{-16}. Two preliminary finite-difference gradient checks failed at 2×10^{-7}; a fourth-order check passes the same threshold with error 2.88×10^{-9}. Their failures and unchanged model/objective are preserved.

A separate same-host environment reruns extraction, all 64 fits and both Speech Commands evaluations from 7,807 recordings: 362 arrays are bitwise equal and 720 proof-array members exactly equal. The historical cohort selection is supplied rather than regenerated. This is retrospective reproduction, not a second confirmation or independent frontend implementation. Immutable revision16 Appendices B, C, F and I retain full acquisition, fit and execution records; the current reproduction guide distinguishes saved-proof replay from raw-audio reruns. No public release, independent external replication or completed recording-rights audit is implied.

D PROSPECTIVE CANDIDATE VALIDATION

This study tests whether the complete recurrence pipeline changes a candidate-validation decision under fixed computational policies. Six new rational qubit constructions and two known HMM2 nulls were fixed before evaluation. For each target, four methods certify lower bounds on its distance from the entire homogeneous two-state edge-emitting HMM class, and a fixed family supplies physical candidates with exact errors. The primary endpoint succeeds on all six quantum constructions, against a registered threshold of three. The separate physical-width target succeeds on five: Q16 remains a failure for that endpoint.

D.1 EXACT CONSTRUCTIONS AND CANDIDATE FAMILY

All laws have horizon eight and 256 probabilities in binary lexicographic order, with the earliest reported bit most significant. The target order is N9, Q11, Q12, Q13, N10, Q14, Q15, Q16. Write

$$c(t) = \frac{1 - t^2}{1 + t^2}, \qquad s(t) = \frac{2t}{1 + t^2}.$$

Starting from I_4, successively left-multiply Givens rotations with blocks

$$\begin{pmatrix} c(t) & -s(t) \\ s(t) & c(t) \end{pmatrix}$$

in the ordered planes $(0,1)$, $(0,2)$, $(0,3)$, $(1,2)$, $(1,3)$, $(2,3)$. For the resulting U, let $(K_y)_{ij} = U_{2y+i,j}$, $i, j \in \{0,1\}$, and $\rho_0 = \psi\psi^T$, where $\psi = (c(z), s(z))^T$. Orthogonality gives $\sum_y K_y^T K_y = I_2$. Table D.1 fixes every parameter.

Table D.1: Exact rational definitions of the six designed quantum processes. They are not fitted acoustic laws.

Target	Six ordered Givens parameters	z	Base rate e
Q11	$(2/7, -1/4, 3/5, 1/3, -3/7, 2/5)$	$1/5$	0
Q12	$(-3/7, 2/5, 1/6, -2/3, 3/8, 1/4)$	$-2/7$	0
Q13	$(4/9, 1/3, -2/5, 3/7, -1/4, 2/3)$	$2/9$	$1/100$
Q14	$(-2/9, 3/7, 2/5, -1/3, 4/7, 1/5)$	$-3/10$	$1/100$
Q15	$(3/8, -2/7, 1/3, 4/5, -1/6, 3/7)$	$3/11$	$1/25$
Q16	$(-4/7, 1/5, 3/8, -2/5, 1/3, 2/9)$	$-2/5$	$1/25$

Let D_γ preserve density-matrix diagonals and multiply off-diagonals by $1-\gamma$. The candidate instrument is

$$\begin{aligned} \gamma_\eta &= e + (1-e)\eta, & \phi_\eta &= e + (1-2e)\eta, \\ \mathcal{I}_{y,\eta}(\rho) &= (1-\phi_\eta) D_{\gamma_\eta}(K_y \rho K_y^T) + \phi_\eta D_{\gamma_\eta}(K_{1-y} \rho K_{1-y}^T). \end{aligned} \tag{D.1}$$

Thus $q_{i,\eta}(w) = \mathrm{tr}[\mathcal{I}_{w_8,\eta} \circ \cdots \circ \mathcal{I}_{w_1,\eta}(\rho_0)]$, and the target is $p_i = q_{i,0}$. Each candidate uses one persistent qubit. Equation (D.1) varies dephasing and report flips together; it is one modeled degradation family, not an isolated test of coherence.

For N9, take $\pi = (p, 1-p)$ and

$$B_0 = \begin{pmatrix} e_0 t_{00} & e_0(1-t_{00}) \\ e_1 t_{10} & e_1(1-t_{10}) \end{pmatrix}, \quad B_1 = \begin{pmatrix} (1-e_0)t_{01} & (1-e_0)(1-t_{01}) \\ (1-e_1)t_{11} & (1-e_1)(1-t_{11}) \end{pmatrix},$$

with $(p, e_0, e_1, t_{00}, t_{01}, t_{10}, t_{11}) = (4/11, 3/10, 5/7, 7/12, 3/8, 5/9, 4/7)$. For N10,

$$\pi = (2/9, 7/9), \quad B_0 = \begin{pmatrix} a & b \\ 0 & a \end{pmatrix}, \quad B_1 = \begin{pmatrix} c & 1-a-b-c \\ d & 1-a-d \end{pmatrix},$$
$$(a, b, c, d) = (4/11, 2/11, 1/11, 3/11).$$

Both are exact HMM2 laws. Their candidates use $B_y(\eta) = (1-\eta)B_y + \eta B_{1-y}$ and the same π.

The common fixed grid is

$$G = \{j/1000 : 0 \le j \le 40\} \cup \{1/20, 3/50, 2/25, 1/10, 3/20, 1/5, 3/10, 2/5, 1/2\}. \tag{D.2}$$

All laws and errors are rational and evaluated without candidate fitting. The endpoints $\eta = 0$ and $1/2$ respectively give the target and uniform law; they are diagnostic only. The 48 interior rates are eligible. All 400 rows, including both endpoints for every construction, remain in the evidence.

D.2 Registered decision and computational policies

For each target and arm $a \in \{P, D, C, K\}$, the reader verifies

$$L_{i,a} \le \inf_{h \in \mathcal{H}_2} \mathrm{TV}(p_i, h), \qquad E_i(\eta) = \tfrac{1}{2} \sum_w |p_i(w) - q_{i,\eta}(w)|.$$

An ordinary decision requires $E_i(\eta) < L_{i,a}$. With $M_i = \max(L_{i,C}, L_{i,K}, L_{i,D})$, an incremental decision requires

$$L_{i,P} - E_i(\eta) \ge 10^{-4}, \qquad E_i(\eta) - M_i \ge 10^{-4}. \tag{D.3}$$

The primary endpoint requires at least three of six Q constructions to satisfy (D.3) at an eligible rate, together with every gate in Table D.5. Each construction counts once. Margin equality is included; ties, failures and unresolved cases remain. Grid rows are dependent design points, with no statistical independence, monotone interpolation or continuous-noise guarantee. The secondary P-versus-D decision replaces M_i by $L_{i,D}$.

The numerical input payload contains only the horizon and target probabilities; policies are fixed. It supplies no generator parameters, candidates, null witnesses or other arms' outputs. Every producer has a 120-second process-group cap including startup, imports and serialization, with ten seconds

reserved for serialization. The arm order rotates P,D,C,K left by target index; the orientation order is original/complement at even indices and reversed at odd indices.

P: complete recurrence pipeline. A cold physical-HMM2 upper search has a nominal allowance of at most 15 seconds, including at most ten seconds for fitting, and zero cover splits. Its exact witness supplies U_H. One correction is attempted in the first scheduled orientation, with at most 60 seconds for the child and 30 seconds for its single LP. The anchor clamps emissions to $[2^{-12}, 1 - 2^{-12}]$, retains the oriented B_0 conditional transitions, uses initial distribution $(1/2, 1/2)$, and sets $B_{1,ij} = (1 - \sum_k B_{0,ik})/2$. This anchor's initial distribution and uniform remainder are separate from the unchanged physical witness. Correction uses the frozen SymPy DomainMatrix backend and denominator 2^{48}.

A corrected law supplies only its orientation's recurrence upper $U_{R,s}$. It need not be physically realizable and never replaces U_H. The other orientation retains the physical upper as a valid recurrence upper. The remaining working allowance is divided equally between orientations without lending; the second is clipped to the remaining overall working deadline. Each cover starts from $[0, 2] \times [-1, 1]$ with zero floor, uses weakest-bound scheduling without endpoint clipping, rejects exactly empty feasible intersections, and permits depth 32, 5,000 LP calls, two seconds per LP, dual denominator 2^{40}, 40-bit square-root bounds and width target zero. The physical floor is $L_P = \max_s L_s$. Pending or stopped leaves retain valid inherited bounds.

D: fitted complete-HMM2 search. The complete weighted physical cover receives the remaining working allowance, the same at-most-ten-second fitting initializer and a 200,000-split cap. D therefore includes a fitted feasible upper witness. P and D share producer and total caps but allocate work differently. The complete cover uses the seven coordinates in the N9 construction, including the closed cube boundary. Its lower bound on each cell is the maximum of zero, forward interval-TV deficits, backward normalized-row interval-TV deficits, and center TV minus the coupling radius. The latter is $r_p + 8 \max_i [r_{e_i} + c_{e_i} r_{t_{i0}} + (1 - c_{e_i}) r_{t_{i1}}]$, where c, r are cell centers and half-widths. The lowest-bound leaf is split along the largest weighted width, with weights $(1, 8, 8, 8, 8, 8, 8)$ and index ties. Closed children, retained leaves and upper-pruned cells cover the full cube; their weakest bound gives the global floor. The supplement's `study/comparator-bounds.md` derives these bounds and their implementation; revision16 Appendix O retains the original account.

The shared initializer evaluates a 193-point bank: $p \in \{0, 1/2, 1\}$, all other coordinates in $\{0, 1\}$, plus the all-half point. Eight Powell starts use $(1/2, 1/4, 3/4, 3/4, 1/4, 1/4, 3/4)$ and seven uniform points in [0.05, 0.95] from NumPy `default_rng(112001)`, with per-start caps of 4,000 evaluations and 200 iterations, `xtol=1e-6` and `ftol=1e-8`. The best visited point is retained regardless of termination status. Proposals are quantized to 2^{-24} with ties to even, then checked and evaluated exactly. Nonfinite values or excursions beyond the cube by more than 10^{-12} reject; smaller roundoff is clipped and recorded before exact checking. The time allowance covers the bank, fits and exact validation.

C and K: fixed certificates. C applies the padded-clock relaxation with exact dual certification and a nominal 15-second allowance. Its 240 context rows use both symbol labels and conditional first-tick bounds $P(u0^{\ell-1}1) \le c_\ell P(u)$. The padded constants for $\ell = 3, \ldots, 8$ are exactly 0.29642, 0.2502, 0.1484, 0.1484, 0.1057 and 0.1057 (Budroni et al., 2021). Its reader checks their application, without rederiving the published grid proof; LP dual repair uses denominator 2^{40}. K uses the same nominal allowance and visits prefix lengths 2–8 and their splits in fixed order, retaining the maximum certified nuclear- or Frobenius-norm floor for classical rank two. Its table and singular-threshold denominators are 2^{40} and 2^{30}, with at most 256 downward threshold steps. For each fixed-split probability table, the certified singular-value tails give the larger of $\frac{1}{2} \sum_{j>2} \underline{\sigma}_j$ and $\sqrt{\frac{1}{2} \sum_{j>2} \underline{\sigma}_j^2}$, after subtracting its rational quantization-error bound; the square root is rounded down on denominator 2^{80}. There are 28 tables. Unused cheap-arm time is not reassigned. These are specified clock/rank constructions and one complete-search implementation, not exhaustive optimizations over every possible certificate or solver.

D.3 COMPLETE NUMERICAL OUTCOMES

All six Q constructions satisfy (D.3). Of their 288 eligible rows, P certifies 185 and the C/K/D union certifies 23; 162 meet both incremental margins. No reverse-portfolio or D-versus-P margin decision occurs. Neither null certifies an eligible candidate.

Table D.2: Decisions among 48 eligible rates per target. P/D/C/K and their union use strict $E < L$; the last two columns use both 10^{-4} margins. Counts are descriptive, not independent trials.

Target	P	D	C	K	C/K/D union	Incremental	P-versus-D
N9	0	0	0	0	0	0	0
Q11	9	2	0	0	2	7	7
Q12	45	2	12	10	12	33	43
Q13	40	0	0	3	3	37	40
N10	0	0	0	0	0	0	0
Q14	42	0	0	3	3	39	42
Q15	7	0	0	0	0	7	7
Q16	42	0	0	3	3	39	42

For an exact compact description, define $J_{a:b} = \{j/1000 : a \le j \le b,\ j \in \mathbb{Z}\}$.

Table D.3: Exhaustive successful rate sets on the fixed grid. All other tested rates, including unsuccessful and diagnostic rows, remain in the complete CSV. Set notation supplies no claim between tested points.

Target	All rates satisfying (D.3)
Q11	$J_{3:9}$
Q12	$J_{13:40} \cup \{1/20, 3/50, 2/25, 1/10, 3/20\}$
Q13	$J_{4:40}$
Q14	$J_{4:40} \cup \{1/20, 3/50\}$
Q15	$J_{1:7}$
Q16	$J_{4:40} \cup \{1/20, 3/50\}$

Table D.4a: First 16 arms, retaining all four methods for N9 and Q11–Q13. Bounds are rounded displays; exact saved fractions decide comparisons. Widths are physical TV intervals and seconds include full charges. The continuation uses identical columns and units.

Target	Arm	Certified floor	Physical width	Charged seconds
N9	P	0.000000000	0.000000014	94.635
N9	D	0.000000000	0.000000014	401.498
N9	C	0.000000000	n/a	3.368
N9	K	0.000000000	n/a	3.260
Q11	P	0.065878232	0.001477491	51.964
Q11	D	0.014197989	0.053157734	418.229
Q11	C	0.000000000	n/a	4.810
Q11	K	0.006350043	n/a	4.727
Q12	P	0.297483032	0.015904159	120.302
Q12	D	0.008923773	0.304463419	423.156
Q12	C	0.040972750	n/a	4.944
Q12	K	0.034001851	n/a	4.998
Q13	P	0.156397880	0.021436856	119.483
Q13	D	0.002905392	0.174929344	417.280
Q13	C	0.000000000	n/a	4.892
Q13	K	0.014385412	n/a	4.940

Table D.4b: Remaining 16 arms; Tables D.4a–b retain all 32. Decimal bounds/widths are rounded displays, not outward proof endpoints; decisions use the saved exact fractions. Physical widths are $U_H - L_P$ for P and $U_D - L_D$ for D. C/K do not construct physical uppers. Costs include the full common and shared-final charges in (D.4), not producer time alone.

Target	Arm	Certified floor	Physical width	Charged seconds
N10	P	0.000000000	0.000032703	101.259
N10	D	0.000000000	0.000032703	380.306
N10	C	0.000000000	n/a	3.334
N10	K	0.000000000	n/a	3.267
Q14	P	0.128586053	0.016068314	121.107
Q14	D	0.000000000	0.144593093	399.896
Q14	C	0.000000000	n/a	5.059
Q14	K	0.007489311	n/a	5.094
Q15	P	0.026668818	0.011064529	120.671
Q15	D	0.000000000	0.037733347	428.410
Q15	C	0.000000000	n/a	5.210
Q15	K	0.002957670	n/a	5.198
Q16	P	0.068530893	0.036502701	119.589
Q16	D	0.000000000	0.105033594	422.923
Q16	C	0.000000000	n/a	5.449
Q16	K	0.005127916	n/a	5.286

P meets the preregistered physical-width threshold $1/40$ on Q11–Q15; D meets it on none of the six Q cases. **Q16 fails**, with physical width approximately 0.036502701. Its complemented recurrence gap, approximately 0.014143252, is a different interval and cannot change that outcome. All eight null floors are exactly zero; their small nonzero fitted physical uppers reflect withholding the known null witnesses. All eight corrections have independently verified results, with no correction timeout or fallback. Correction can improve only recurrence uppers, and this study does not isolate its causal contribution. P does not uniformly improve the fitted physical upper: Q14's D upper is slightly smaller.

D.4 What the saved orientations contribute

The following decomposition uses the frozen proof outputs without rerunning any policy. "Uncorrected" retains only the orientation that received the unchanged physical HMM upper; "paired" takes the maximum of both saved floors. Both use the same candidate errors, comparator envelope and margins in (12).

Table D.4c: Saved orientation outcomes. Floor percentages and seconds are rounded for display; all decisions use exact fractions. Correction was attempted only on the named orientation. This table is not a correction-disabled experiment.

Target	Corrected orientation	Original floor (%)	Complement floor (%)	Decisions: uncorrected / paired	Correction (s)
Q11	complement	6.59	0.02	7 / 7	40.47
Q12	original	9.24	29.75	33 / 33	42.10
Q13	complement	15.64	5.47	37 / 37	41.64
Q14	complement	4.18	12.86	17 / 39	37.11
Q15	original	1.14	2.67	7 / 7	43.95
Q16	complement	5.71	6.85	38 / 39	45.07

The uncorrected orientations already give at least one additional decision on all six constructions, totaling 139 of the paired method's 162 decisions. Pairing contributes 22 more on Q14 and one on Q16. The uncorrected floors resolve four physical widths; pairing adds Q14 as the fifth. On Q11,

Q12, Q13 and Q15, the winning floor comes from the uncorrected orientation. Correction consumes 37.11–45.07 seconds per quantum target and changes only the recurrence upper used for pruning, not the physical HMM upper. It may affect how quickly that orientation's cover is explored. Removing it would return time to search, so these saved outcomes cannot predict a correction-disabled pipeline's result.

Stage accounting also matters for the runtime comparison. D's recorded producer stage takes about 111 seconds, while its proof audit takes 284.41–312.64 seconds on the six quantum targets. The latter dominates D's charged total. The reported advantage is therefore in obtaining a checked result under the specified implementations and allocations. It is not an isolated search-speed comparison. The exact stage charges remain in the submitted arm table.

D.5 Timing, independent checks and retained failures

Target generation and independent candidate checking have 120- and 180-second watchdogs. A failed common stage prevents every method launch, not just that target's subset. Producers run serially; mathematical method audits follow the completed producer phase. Every logical arm has a 600-second total budget:

$$T_{i,a} = T_i^{\text{generation}} + T_i^{\text{target audit}} + T_{i,a}^{\text{producer}} + T_{i,a}^{\text{proof audit}} + T^{\text{final}}. \tag{D.4}$$

These measured stage charges include cold startup, imports, fitting, search, serialization and observed overhead. Common work and the entire shared final validation are charged to every arm even though physically performed once. Final validation has a 30-second reserve inside 600 seconds; its 2.166015-second observed cost includes evidence parsing, hashing and decisions. Subsequent timing-ledger serialization is excluded. A C/K/D portfolio costs the sum of three complete standalone arms, including three common/final charges. Its Q-case totals range from 410.049 to 438.818 seconds.

All 32 producers and independent method audits complete on their first attempts. The longest external producer execution is 111.432357 seconds, below its 120-second cap; charged producer-stage time additionally includes surrounding overhead. All logical totals are below 600 seconds, with maximum 428.409601 seconds. On Q cases P costs 51.964–121.107 seconds and D 399.896–428.410; C and K cost roughly five seconds. These are first-run observations of classical certificate computation, not repeated-run speed estimates or quantum execution times.

Table D.5: All 14 final gates. Exact gate identifiers, source bindings, clocks and failure predicates remain in the final ledger.

Acceptance gate	Outcome
Prepared sources, configuration and authorization bound	Pass
All eight target/candidate audits	Pass
All 400 candidate rows available	Pass
Freshness against historical and new laws	Pass
Both nulls: every verified floor zero	Pass
All 32 mathematical proof audits	Pass
All target stages within their allowances	Pass
All 32 producers within their caps	Pass
All 32 method audits timely	Pass
Recorded execution order serial	Pass
Every pre-final logical charge within 600 seconds	Pass
Final reader within its external cap	Pass
Shared final validation within 30 seconds	Pass
Every complete logical charge within 600 seconds	Pass

The target reader independently reconstructs physical channels, every candidate law, all prefixes and exact TV errors. It checks normalization, Kraus completeness, density validity and endpoint identities. The mathematical readers verify upper membership, orientation-specific correction premises, exact duals and complete covers. Freshness checks hash all 80 historical input files, compare the

eight new laws with 32 historical horizon-eight laws and with one another, and test symbol complements: 284 pairs and 568 predicates yield no collision. Collisions would receive no replacement.

The original final readback recomputes decisions and binds the clean recorded proof-reader executions to their sources, inputs and output hashes; it does not rerun all 32 proofs. Missing proofs remain unresolved. Only an external correction timeout or explicit solver time limit permits the registered timeout fallback to U_H; no-budget stages retain zero/fallback records. Non-timeout solver, arithmetic, source or input failures cannot be silently replaced. No external process timeout, orphan failure, method retry or correction fallback occurs in the accepted run. Stopped search leaves remain in the covers. The Q16 width failure and every unsuccessful candidate comparison remain part of the result.

D.6 Evidence files and replay

The immutable study is `results/candidate-validation-v1/`. Its `prepared/catalog.json`, `prepared/config.json` and `prepared/method-policy.json` fix constructions, grid, endpoints and policies. The prepared manifest binds the complete source closure. Within `prepared/source/`, the `iteration146-targets.py` and `iteration146-target-reader.py` pair constructs and independently checks the laws; `iteration146-methods.py`, `iteration146-upper.py` and `iteration146-method-reader.py` produce and check the certificates. `iteration146-study.py`, `iteration146-final-reader.py` and `iteration146-freshness.py` define supervision, costs, decisions and collisions.

The authoritative outcomes are `run/decision/readback.json` and `run/decision/final.json`. The unchanged tables `notes/iteration148-summary/arms.csv` and `notes/iteration148-summary/candidates.csv` retain all 32 arms with exact bounds/component costs and all 400 candidate rows with exact errors, margins and decisions. Correction statuses labeled `produced_pending_independent_audit` describe producer output; every accepted correction subsequently passes its independent reader.

The historical supplement's `audit_supplement.py` reruns all eight target readers and 32 method-proof readers, recomputes 400 decisions and the original cost ledger, and checks the separate learned-model proof. It performs no fitting, producer search, audio decoding or QPU execution. Its first full replay succeeds from a renamed read-only extraction while operating-system controls deny original-workspace reads, network access and extraction writes. The observed 2,783.543016-second replay duration is additional verification work and never replaces the original experimental times. The unbundled runtime pins CPython 3.12.13 including its binary identity, NumPy 2.5.3, SciPy 1.18.1, SymPy 1.14.0 and mpmath 1.3.0, requires absent flint, and rejects optimized Python. This is same-host relocation, not cross-platform validation or independent external replication. Appendix F distinguishes this historical archive from the submission ZIP's portable mathematical replay and new-target generator.

E Supporting evidence and limits

E.1 Evidence behind the final policy

The final pipeline uses weakest-bound scheduling without endpoint clipping, and constructs a fitted physical upper witness before attempting recurrence correction. Table E.1 gives the evidence relevant to those choices. Earlier comparisons use their own targets, budgets and upper policies; their outcomes are not pooled with the prospective candidate test.

Table E.1: Consolidated supporting studies. All entries refer to fixed policies and finite designed or exposed laws. "Stronger floor" means a better certified lower bound, not a better achieved fit. The archive retains each arm and failure; these studies do not supply a universal ranking of methods.

Study	Supported observation and limiting result	Immutable revision16 record
Clock comparison on 20 saved conditions	No recurrence-only decisions; nine clock-only, three both, eight neither. After grouping 14 distinct exact laws: zero, five, three and six. Clock already validates acoustic confirmation.	Appendix D, Table D.3
Four-event complete search	Continued recurrence and complete search each give the stronger floor on two of five targets; one zero tie. Partial grids contribute no floor.	Appendix O, Tables O.1–O.2
Fitted complete-search ablation	On two exposed eight-event laws, fitting improves physical uppers but alone slightly lowers median floors under 30-second budgets. Weighted branching helps; recurrence stays stronger in all four repeats.	Appendix Q, Table Q.1
Initialization/ branching factorial	Matched fixed-target comparisons separate feasible initialization from search order. They are timing repeats on reused laws, not independent target samples.	Appendix S, all 50 calls
Independent flow relaxation	The root floor is always zero. Two 30-second calls give development/cyclic floors 0.068162/0, so the two-target positivity endpoint fails.	Appendix U, root and pilot
48-law prospective panel	Under nominal 15-second method ceilings, recurrence exceeds complete-search floors by 0.058914 on average over 36 non-null laws. This older complete-search arm has **no fitted initializer**.	Appendix V, all 192 certified arms
Fresh 20-law adaptive panel	Adaptive minus uniform-recurrence mean floor is 0.015341 over 16 non-null laws. Four known nonmembers remain zero ties. Adaptive physical widths meet 0.025 on 0/20; complete search does so on 7/20 with better uppers.	Appendix Y, Tables Y.1–Y.3
Shared-engine clipping/scheduling	At 15 seconds, mean scheduling effect is +0.027865887, clipping effect −0.004672728, interaction +0.009241971. Cheap uppers leave **0/192** physical widths at most 0.025.	Appendix Z, Tables Z.2–Z.9
Three-state acoustic certificate	Development point gap is at least 0.0000030109148167756224, but conditional interval [−0.0145300252, 0.0004188704] fails stability. No refinement follows that bootstrap.	Appendix H, Equation H.7
Mixtures and nonspeech controls	The original recurrence certificate excludes two of four mixtures; fixed rank bounds detect all four. No mixture certifies noisy-qubit advantage. Only one of 17 nonspeech conditions passes recurrence comparison.	Appendices E and J

The clipping/scheduling factorial is the direct attribution experiment. One engine crosses breadth-first versus weakest-bound scheduling with endpoint clipping off versus on, using four fresh qubit constructions and two exact HMM2 nulls at horizons six and eight. Two allocations (5 and 15 seconds) and two timing repeats give 192 cold processes. Empty-intersection rejection, complete-cover arithmetic and all other controls are shared. At the primary 15-second allocation, the effects equally weight 16 qubit execution cases: four constructions, two dependent horizons and two repeats. Clipping changes the breadth-first mean by −0.009293713 and the weakest-bound mean by −0.000051742. Thus the positive interaction reduces a clipping penalty; it is not evidence of a pos-

itive clipping main effect. All four construction means for the bundled change are positive. At five seconds the averaged scheduling/clipping effects are +0.033770111/−0.003873039.

Each factorial arm receives the same five iid upper candidates, with zero rates $0, 1/4, 1/2, 3/4, 1$; no fitted upper or correction is supplied. Every null floor is zero. All 384 orientation proofs verify, including 220,987 repaired duals and 192 retained solver timeouts among 221,179 LP calls. The complete covers retain unfinished leaves. These results motivate the final search policy and richer upper stage, but do not identify exact correction as the cause of the final pipeline's gain. Appendix D evaluates that complete pipeline as a bundle.

E.2 What remains unestablished

The nonspeech panel fixes 592 one-second PCM16 waveforms in 17 conditions, with eight phases and paired noise seeds frozen before generation. Each condition weights waveforms equally and then eligible anchors equally within each waveform. Pure tones use periods 14, 16 and 20 samples, each exactly and at ś0.05 samples, plus 1100 Hz; noisy 1000/1100 Hz tones use input SNRs 20, 0 and −10 dB, alongside Gaussian-only noise. Noise precedes the common filter. The sole recurrence success, 1000 Hz at −10 dB, has qubit TV at most 0.521847268 against floor at least 0.526139474; HMM4 is better at 0.373813572. The other 16 fail this comparison, and clock bounds validate nine additional conditions. Boundary-sensitive exact tones retain the unchanged right-closed rule. There is no inference over sampled speakers or frontend-parameter intervals.

A separate retrospective acoustic calculation proves a continuous sufficient noise range using a channel coupling bound: the saved 32-grid development floor 0.250601570 validates [0, 2.65%] joint dephasing/report-flip noise, compared with [0, 1.25%] for the saved clock floor. Recurrence alone validates [1.65%, 2.65%] against those saved clock/rank scores. These modeled-channel intervals differ from Appendix D's finite grid and from measured device tolerances. All 51 evaluations and the continuity proof remain in immutable revision16 Appendix R and `results/noise-tolerance-v1/`.

The acoustic confirmation does not establish general transfer: AudioMNIST fails at all three declared noise settings, HMM4 fits Speech Commands better, and the synthetic recurrent hardware pilot fails both primary checks (Appendix C). Q16 passes incremental candidate validation but misses the physical-width target (Appendix D). The learned qubit's larger-state separation concerns its own exact law; its microscopic constructive approximation bound gives no useful device or sampling tolerance. All qubit laws here admit inexpensive density-matrix simulation on a classical computer. Persistent finite-state memory, certificate computation cost, and physical implementation cost are separate resources.

E.3 Archival access and compute scope

The complete historical report is `memory-manuscript-v16/manuscript.md` and its frozen PDF. The study directories, relative to `results/`, are `clock-transfer-comparison-v1`, `complete-h2-tv-benchmark-v1`, `stronger-complete-comparison-v1`, `factorial-complete-comparison-v1`, `flow-h2-pilot-v1`, `prospective-law-benchmark-v1`, `adaptive-breadth-v1`, `clipping-scheduling-v1`, `three-state-memory-v5`, `synthetic-memory-controls-v1` and `mixture-memory-v1`. The source map accompanying this revision identifies the exact retained claim dependencies; the archive holds full numerical tables, construction constants, execution failures and manifests. No old result is replaced by a later success.

The local acoustic/certificate environment uses an Apple M5 Max, 18 CPU cores, 128 GiB unified memory, macOS 26.4.1 and CPython 3.12.13, with one requested numerical CPU thread and no GPU or QPU for acoustic evaluation. Original recorded stages take 8.56 seconds for 64 fits, 8.05 for development evaluation/certification, and 5.49 for confirmation decoding/evaluation/certification; their scopes are separate from the fully charged Appendix D comparison. Exploratory work, failed runs and document generation are not covered by a complete project meter. Peak memory, energy and total project compute are not inferred. Speech Commands attribution and CC BY 4.0 notices, and AudioMNIST's repository MIT notice and citation, are retained; recording-specific rights, demographic fairness and public release licensing remain outside the completed verification.

F REPRODUCTION AND THE COMPLETE EVIDENCE RECORD

F.1 ANONYMOUS SUPPLEMENTARY MATERIAL

The submission ZIP contains standalone LaTeX source, all 400 candidate rows and 32 method summaries, all prospective-study method proofs, their readers, and a minimal certificate generator. It contains no raw recordings. The quick check is:

```
python3 -I -B study/verify_study.py --recurrence-only
```

This standard-library command independently reconstructs all eight targets and all 400 candidate errors from the rational definitions in Appendix D. It checks the 16 recurrence covers, repaired duals, physical HMM upper witnesses and exact membership of the corrected recurrence laws. Every proof target and pruning upper is bound to these reconstructed objects. Comparative decisions in this quick mode remain conditional on the supplied comparator floors.

To also check the complete-model, clock and rank proofs, install NumPy in a Python environment and run:

```
python3 -I -B study/verify_study.py --workers 2
python3 -I -B supplementary-data/check_tables.py
```

The first command checks all 32 method results, including zero-floor controls, and recovers the six construction successes and five physical-width successes. The clock checker applies the published padded constants without rederiving their original grid proof. The second command checks saved table arithmetic and cost sums. Neither command authenticates historical durations or reruns fitting, search, audio processing or hardware. The README records tested environments, observed replay durations and corruption checks. Normal Python is required; there is no interpreter-binary hash gate. These are same-host relocated checks, not independent external reproduction.

The supplied `study/generate/` entrypoint accepts a new rational binary probability table and produces a checked recurrence floor. Its uniform HMM upper and bounded cover search provide a minimal reusable implementation; it does not reproduce the study's fitted-upper/correction policy. The README gives input formats and positive and negative examples. The unchanged historical producer sources and policy configuration are included separately for inspection.

The small `reproducibility/verify.py` example remains available as one complete Q11 candidate check. The LaTeX source rebuilds with `tectonicpaper.tex`. All submitted figure exports are unchanged.

F.2 HISTORICAL SOURCE MAP

The immutable 110-page revision16 report retains the original proofs, prior tables and negative results. Its PDF SHA256 is `acb26196071ef9d02416d9c264a30eaf0bd9dc550f526d433db80ec83dfbb437`. The accompanying `archive-map.md` maps every original appendix to the present paper or its retained record.

The paths below identify the retained research repository. The full repository and historical archives are not included in the submission ZIP. Its LaTeX source and prospective-study mathematical results use the self-contained entrypoints above.

Table F.1: Scientific entrypoints. A table summary does not replace the bound proof, source-bound reader or complete method record. A checked lower bound and a fitted physical upper have different roles even when stored together.

Claim or evidence	Current source
Acoustic confirmation, fitting, transfer and device pilot	Frozen inputs named in Appendix C; original protocols and full outcomes in revision16, Appendices B–I
Whole-class recurrence certificate and integer witnesses	Appendix A; acoustic saved-proof readers in the integrated v8 evidence package

Claim or evidence	Current source
Mesh limit and realization counterexamples	Appendix B; revision16 Appendices L, N and X; exact mechanism study
Six-construction candidate study	`results/rank-recurrence-separation-v1/` `results/candidate-validation-v1/`: frozen preparation, target generation, 32 method runs, proof audits and final decision/cost readback
Complete candidate rows and arm summaries	`notes/iteration148-summary/`: exact-fraction CSVs and original summary
Historical policy comparisons and adverse outcomes	Appendix E and `archive-map.md`; original report and raw studies remain unchanged

F.3 Scope of saved-proof replay

Two retained local packages cover the evidence. The v8 archive includes the acoustic application and the earlier studies through revision16 Appendix Z. Its twelve-stage saved-proof replay completed in 4,639.667 seconds with reads of the original checkout and network access denied. The separate historical supplement adds the complete candidate study and the separately archived learned-model sector proof. Its first isolated replay completed in 2,783.543 seconds with reads of the entire original workspace, network access and extraction writes denied. It reconstructed all eight targets, checked all 32 method proofs, recovered all 400 candidate decisions and preserved the original cost records.

These proof-replay durations do not replace the experimental costs in Appendix D. Replay runs no fitting or primary search. Full success requires the top-level completion record; an inventory or subset check is insufficient.

The package entrypoints, after extraction to canonical absolute paths, are:

```
/absolute/python -I -B /absolute/V8/audit_all.py \
  --reports /absolute/new-parent/v8-reports
/absolute/python -I -B /absolute/SUPPLEMENT/audit_supplement.py \
  --reports /absolute/new-parent/supplement-reports
```

The report directories must be new, their parent must exist, and the runtime must satisfy the respective package requirements. Both require normal Python, not optimized mode. The supplement additionally pins the interpreter binary. The historical runtime and raw audio are unbundled. These full-workflow checks establish same-host relocation, not independent external installation or cross-platform reproduction. The submission ZIP now checks the prospective-study mathematical results without this runtime restriction; it does not port the historical execution ledger or the acoustic workflow. Public hosting and unified project-code licensing remain unfinished. AudioMNIST recording-specific redistribution rights are unresolved; raw recordings are excluded from the proof packages.

Archive hashes, runtime pins and completion records are indexed in `memory-release-artifact-v8/README.md` and `memory-supplement-artifact-v1/README.md`; the historical reader guide is `notes/iteration149-reproduction-guide.md`.