Per-Coordinate State Quantization Freezes a Permutation-Recurrence Reservoir, but Never a Mixing One

A Preprint

Caitlyn Meeks

AuxiLab Tenerife — research@auxi.cafe

August 12, 2026

registered 2026-08-12 | internal run log #40, #41 | 2 seeds × 19 cells

headline confirmed, two anchors missed | promoted 2026-08-12

Abstract

Rounding an echo state network's state coordinate-wise onto a b-bit grid after every step degrades prediction smoothly but, in an earlier single-seed study on a sparse mixing reservoir, never produced any exact-state determinism at all, at any tested resolution down to a single bit. This study asks whether that negative depends on the *mixing*. Replacing the sparse recurrent matrix with a pure delay line — W = ρ·P for a single random N-cycle permutation P, which never combines information across coordinates — and holding every other parameter fixed, **the same quantization dial drives the reservoir to near-complete determinism over its recent input, while an input-matched mixing control shows none.** At ρ=0.8, the fraction of recurring 20-character contexts mapping to a single state reaches 1.0000 and 0.9944 at b=2 and 0.9831 and 0.9379 at b=4 across two independently drawn reservoirs, against a registered bar of 0.9; the matched mixing control reads exactly 0.0000 in every cell it was measured in, at both spectral radii and both seeds. One cell — seed 0 at b=2 — reaches determinism of exactly 1.0 and is a finite-state automaton over a 20-character window in the strict sense; the replication seed's best cell falls just short of that threshold while clearing the registered bar comfortably. Two pre-registered quantitative anchors missed decisively and resolved into an unanticipated structural result: exact-state determinism and linear decode depth respond to the identical dial with *opposite* spectral-radius sensitivities. Decode depth is clamped hard at ρ=0.8 (a 1.58–1.73 character range across the grid) and barely moves at ρ=0.95 (0.11–0.18 characters), while determinism is easier to trigger at the lower ρ. Exact-state collision and linear decodability are measurably different senses of "memory under a blurring dial," not two views of one quantity.

Keywords reservoir computing · echo state networks · state quantization · finite-state automata · permutation recurrence · delay lines · short-term memory · pre-registration · replication

---

§1 The question

An echo state network (ESN) [1] [2] keeps a high-dimensional state that is driven by an input stream and read out linearly. If that state is forced onto a finite grid after every update, the system becomes, in principle, a machine with finitely many reachable states — and the natural question is whether it behaves like one. An earlier study in this program applied per-coordinate rounding, x ←

`round(x·2^b)/2^b`, to the program's reference sparse-tanh reservoir and obtained a striking negative: bits per character (bpc) degraded smoothly as b fell, but the system never became deterministic over its recent input at any tested resolution — $b \in \{1, 2, 3, 4, 5, 6, 8, 10, 23\}$ and float32 — and its states never collided at all. The proposed explanation was re-injection — each step's rounding error is fed back through the recurrent matrix and re-mixed across all coordinates faster than the $\rho<1$ contraction can damp it.

That explanation makes a sharp prediction about a reservoir that *cannot* re-mix. In a pure delay line, $W = \rho \cdot P$ for a single random N-cycle permutation P, each coordinate's content is passed intact to exactly one successor coordinate; rounding error injected at step t sits in a rotating frame and decays as exactly ρ^k, never spread across the state. Quantization should then impose a hard memory horizon and produce a finite-state automaton, precisely where the mixing reservoir produced neither. That reservoir family — a permutation recurrence with orthogonal input codes — is already characterized in this program at full precision, so the comparison changes exactly one thing. Recovering finite-state descriptions from recurrent networks is a well-established goal with its own methods and literature [3]; the question here is different and narrowly empirical — not how to extract an automaton from a trained network, but which recurrence structure causes a quantized reservoir to collapse onto finitely many states in the first place. It is a question the earlier study explicitly left open.

---

§2 Method

The driver combines two existing systems without modifying either. The per-step rounding line is taken verbatim from the per-coordinate quantization study; the reservoir construction, the linear decode probe, and the bpc evaluation are taken verbatim from the permutation-recurrence study. N=1024, leak 1.0, bias scale 0.2 for tanh arms and 0.0 for linear arms, washout 200, trained and evaluated on text8 [4] with 2M training and 500k validation and test characters in every cell that reports bpc or decode depth. The four instrument-only cells described below skip that collection entirely and run a single 500k-character determinism pass instead.

```
# recurrence, permutation arm: a single random N-cycle
#   (a pure delay line -- never combines coordinates)
W = rho * P

# recurrence, mixing control: the reference sparse matrix
#   (fan-in 10, rescaled to the same spectral radius)
W = rho * W_sparse

# input coding, BOTH arms: 27 exactly orthogonal Hadamard
#   codes, scaled by 1/sqrt(3)
Win = hadamard_codes(27) / sqrt(3)

# the quantization dial, applied after every update
x = round(x * 2**b) / 2**b
```

19 cells in three tiers. The **primary battery** is the permutation arm at $\rho \in \{0.8, 0.95\} \times b \in \{2, 4, 6, 8, \text{float32}\}$, with the full instrument set. The **control-square completion** is the sparse mixing arm at $\rho=0.95$ across the same b grid, also with the full instrument set. Four **instrument-only spot cells** measure determinism and cophenetic crispness only, with no bpc and no decode depth: the mixing arm at $\rho=0.8$, $b \in \{2, 4\}$ — closing the possibility that Hadamard input coding rather than the permutation caused any freezing, at the one grid corner whose only prior reference used random input coding — and a genuinely *linear* permutation arm at $\rho=0.8$, $b \in \{4, 8\}$, for which the rotating-frame argument is exact by construction rather than only approximately so under saturation.

Three instruments matter here. **Suffix-determinism** samples 40,000 positions from a 500k-character dump, groups them by their preceding k characters, and reports the fraction of multi-sample suffix classes whose members all sit at the identical state; a value of 1.0 at some finite k means identical recent context always drives the system to identical state, which is what it means for the closed-loop dynamics to be a finite-state automaton over that window. It is reported at the probe's maximum depth, k=20, and its resolution should be read with it: at that depth the 40,000 sampled positions fall into 39,788 distinct 20-character contexts, of which only **177 recur at all** and therefore enter the denominator. That count is a property of the text and the sampling seed, so it is identical in every cell of every arm reported here — which makes cells directly comparable — but it also means a single group is worth 0.0057, and differences finer than a few hundredths are at the instrument's resolution rather than resolved by it. The headline contrast below is between 1.0000 and 0.0000, far outside that concern. **Linear decode depth** fits ridge probes that recover the input character at increasing lags and integrates the resulting accuracy curve into a single figure in characters. **Validation bpc** is reported for both a closed-form ridge readout and a trained logistic readout.

---

§3 What we registered in advance

A 100k-character pilot ran before the predictions were finalized and was disclosed with the registration: it showed the permutation arm near 0.99 determinism and the mixing arm at exactly 0.0 at b=4. The pilot's direction was disclosed but explicitly not used to move the thresholds, which came from the accepted proposal.

- **P1** (headline): suffix-determinism at depth 20 rises above 0.9 for $b \le 4$ at $\rho=0.8$ in the permutation arm, against the mixing control's 0.0 at every b. b=6 at $\rho=0.8$ was designated in advance as a *boundary cell* and excluded from this judgment, because theory puts its dead-lag estimate at or past the probe's own depth ceiling, so a result either way there is uninformative. Named loss: the permutation arm also stays at 0.0, which would make the earlier study's re-injection explanation wrong in a different way.
- **P2** (ordinal): decode depth is strictly increasing in b at fixed ρ, with depth at b=4 at least 4 characters shallower than at float32 at $\rho=0.95$. Named loss: depth is inert in b down to b=2.
- **P3** (band): the ordinary-least-squares slope of decode depth against b at $\rho=0.8$ lands in [2, 4.5] characters per bit (theory: $1/\log_2(1/0.8) \approx 3.1$). The qualifying points were fixed in advance as every b whose depth is *strictly less* than the float32 depth at the same ρ, decided from this run's own float32 cell; fewer than two qualifying points would have made P3 unjudgeable rather than guessed.

The replication run re-registered the headline unchanged and added **R-P2**, promoting the first run's unregistered discovery to a testable claim: the decode-depth range across the b grid is larger at ρ=0.8 than at ρ=0.95, with the named loss being ranges within a factor of two of each other, or reversed. Its seed was drawn from a hardware random number generator and recorded before the run. Promotion required the headline to hit with zero reversals on both halves.

Disclosed mid-run amendment. The determinism probe carries an unconditional self-test that recomputes a sampled state from a cold start and refuses to write if it disagrees bit-for-bit with the stored one. It fired on the first float32 cell. Diagnosis with a standalone script showed a chunk-width non-associativity in float32 matrix-vector products of order $3 \cdot 10^{-8}$ — present identically in the mixing arm at the same position, so a general property of unquantized execution in this harness rather than anything about the permutation. At every $b \le 8$ the rounding step masks it completely, since $2^{-8} \approx 0.0039$ is five orders of magnitude coarser. The fix, committed before the run resumed, skips the determinism probe at float32 only and records it as null; bpc and decode depth are computed there unchanged, and no registered prediction needs determinism data at float32.

---

§4 What happened

All 19 cells completed in both runs, with no logistic readout tripping the program's divergence tripwire. As a construction check, the permutation arm at ρ=0.95 and float32 reproduces the archived full-precision cell from the permutation-recurrence study exactly — ridge 3.2247, logistic 2.7538, depth 11.05, each to the last recorded digit — confirming that the quantization line is a true no-op at float32 and that the hybrid is wired as intended.

Table 1. Suffix-determinism at depth 20, both seeds (seed 0 / hardware-drawn seed 899931030). Values are the fraction of the 177 recurring 20-character contexts that map to a single state; 1.0 — reached in exactly one cell here — means the closed-loop dynamics are a finite-state automaton over a 20-character window. Dashes are cells the tier did not run.

recurrence	ρ	b=2	b=4	b=6	b=8
permutation, tanh	**0.8**	**1.0000 / 0.9944**	**0.9831 / 0.9379**	**0.8531 / 0.8418**	**0.6384 / 0.6158**
permutation, tanh	0.95	0.7401 / 0.8192	0.4463 / 0.4576	0.2599 / 0.2599	0.1638 / 0.1751
sparse mixing, tanh	**0.8**	**0.0000 / 0.0000**	**0.0000 / 0.0000**	**—**	**—**
sparse mixing, tanh	**0.95**	**0.0000 / 0.0000**	**0.0000 / 0.0000**	**0.0000 / 0.0000**	**0.0000 / 0.0000**
permutation, linear	0.8	—	0.0282 / 0.0395	—	0.0056 / 0.0113

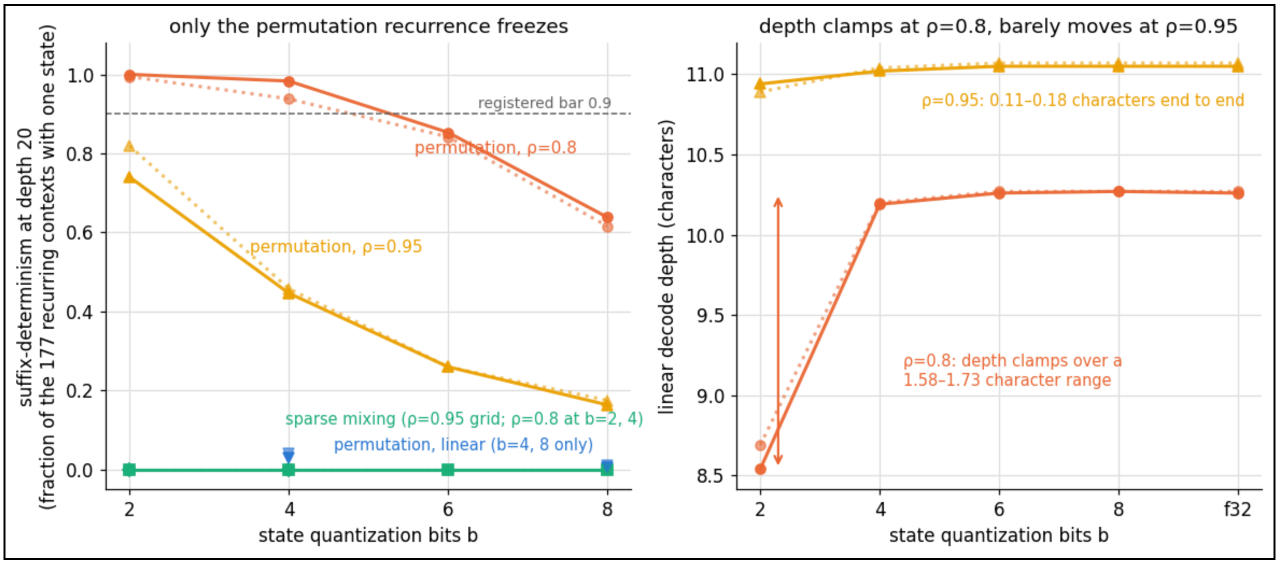

Figure 1. Left: only the permutation recurrence freezes. Solid lines are seed 0, dotted lines the hardware-drawn replication seed. The mixing control sits at exactly zero wherever it was measured, and the linear permutation arm — for which the rotating-frame argument is exact — barely leaves zero either. Markers are drawn without connecting lines where a tier ran only part of the b grid. Right: the same quantization dial clamps linear decode depth hard at ρ=0.8 and almost not at all at ρ=0.95, the opposite ordering to determinism.

P1 — confirmed, cleanly, on both seeds. At ρ=0.8 both b ≤ 4 cells clear the 0.9 bar decisively in both runs, and the mixing control reads exactly 0.0000 in all twelve cells that measure it — the four full-battery cells at ρ=0.95 and the two spot cells at ρ=0.8, in each seed. The "0.0 at every b" half of the prediction holds exactly, not approximately. The named loss did not fire. The registered boundary cell at b=6 landed at 0.8531 and 0.8418, just under the bar and consistent with its designation, and b=8 (0.6384, 0.6158) shows the trend continuing smoothly rather than abruptly. Because the mixing control uses the same Hadamard input coding as the permutation arm and still reads zero at every corner, the effect is a property of the recurrence family and not of the input code.

P2 — missed decisively on its quantitative anchor, and the miss is the finding. The registered claim was strict increase in b at *fixed ρ*, i.e. at both radii, and it does not hold strictly at either. At ρ=0.8 it essentially holds (8.54 → 10.19 → 10.26 → 10.27 → 10.26 across b = 2, 4, 6, 8, float32 at seed 0, a genuine 1.7-character range), with the sole reversal at the saturated tail, where b=8's 10.27 exceeds float32's 10.26 by 0.01 — one unit in the last reported digit. At ρ=0.95 depth ties outright at b=6, b=8 and float32 in both seeds (11.05/11.05/11.05 and 11.07/11.07/11.07), so the increase is weak rather than strict there too. But the registered anchor — depth at b=4 at least 4 characters shallower than at float32, *at ρ=0.95* — misses by two orders of magnitude: 11.02 against 11.05, a gap of 0.03 characters. Across the entire ρ=0.95 grid, depth varies by 0.11 characters at seed 0 and 0.18 at the replication seed. The structure this reveals was not predicted: depth clamping under quantization is *spectral-radius dependent* — a range of about 1.7 characters at ρ=0.8 against 0.11 at seed 0, and 1.58 against 0.18 at the replication seed, ratios of about 15.5× and 8.8× — while determinism runs the other way, being easier to trigger at the *lower* ρ (1.0000 at ρ=0.8, b=2 against 0.7401 at ρ=0.95). R-P2 registered this split as a claim before the second run and it hit, with the ranges far outside its factor-of-two named loss.

P3 — missed, decisively, against the registered band. Applying the pre-registered qualifying rule, exactly two points qualify at ρ=0.8 in each seed — b=2 and b=4 — since neither b=6 nor b=8 is strictly shallower than float32 in either seed (at seed 0, b=6 ties it and b=8 exceeds it; at the replication seed, b=6, b=8 and float32 are all 10.27). Two points is the registered minimum, so P3 is judgeable rather than unjudgeable. The slopes are 0.825 and 0.755 characters per bit, below even the named-loss threshold of 1 and roughly four times shallower than the [2, 4.5] band the ρ^k-decay theory implies. The depth-versus-b relationship is real, reproducible across seeds to within 9%, and far shallower than the mechanism that governs exact-state determinism — consistent with P2's finding that decodability saturates almost immediately.

Table 2. Linear decode depth in characters and validation bits per character, both seeds (seed 0 / seed 899931030). Lower bpc is better. The float32 rows are the unquantized reference for each arm.

recurrence	ρ	b	decode depth	ridge val bpc	logistic val bpc
permutation	0.8	2	8.54 / 8.69	3.1955 / 3.1946	2.7213 / 2.7261
permutation	0.8	4	10.19 / 10.20	3.1618 / 3.1650	2.7010 / 2.7061
permutation	0.8	6	10.26 / 10.27	3.1468 / 3.1494	2.6901 / 2.6938
permutation	0.8	8	10.27 / 10.27	3.1450 / 3.1477	2.6888 / 2.6924
permutation	0.8	f32	10.26 / 10.27	3.1448 / 3.1476	2.6887 / 2.6924
permutation	0.95	2	10.94 / 10.89	3.2869 / 3.2816	2.8204 / 2.8155
permutation	0.95	4	11.02 / 11.04	3.2328 / 3.2323	2.7626 / 2.7656
permutation	0.95	6	11.05 / 11.07	3.2258 / 3.2271	2.7550 / 2.7591
permutation	0.95	8	11.05 / 11.07	3.2247 / 3.2266	2.7538 / 2.7586
permutation	0.95	f32	11.05 / 11.07	3.2247 / 3.2265	2.7538 / 2.7585
sparse mixing	0.95	2	9.25 / 9.17	3.2918 / 3.2859	2.8276 / 2.8215
sparse mixing	0.95	4	9.95 / 9.93	3.1682 / 3.1649	2.6994 / 2.6970
sparse mixing	0.95	6	10.02 / 10.02	3.1535 / 3.1491	2.6862 / 2.6835
sparse mixing	0.95	8	10.03 / 10.02	3.1526 / 3.1480	2.6853 / 2.6824
sparse mixing	0.95	f32	10.03 / 10.02	3.1525 / 3.1479	2.6852 / 2.6824

An unregistered spot check that ran the wrong way. The linear permutation cells were included because the rotating-frame argument is exact when there is no saturating nonlinearity: the Jacobian is diagonal everywhere by construction, not merely approximately so. They should therefore have frozen at least as hard as the tanh cells. They barely froze at all — 0.0282 and 0.0395 at b=4, 0.0056 and 0.0113 at b=8, against the tanh arm's 0.9831 and 0.9379 at the identical ρ and b. One reading, offered as a lead and not as a result: tanh's boundedness may be doing mechanistic work the diagonal-Jacobian argument does not capture, since squashing states into [−1,1] makes a fixed *absolute* rounding resolution a coarse *relative* one and drives collisions, whereas an unbounded linear state's wide dynamic range leaves the same absolute grid comparatively fine. These cells carried no pass/fail bar of their own and are not part of the promoted claim; the reading needs its own registered test.

---

§5 What it means

The earlier negative was about the mixing, and the open question it left is now closed. Under per-coordinate rounding, freezing is not a universal property of discretized reservoirs and not an artifact of a particular resolution; it depends on the recurrence. The contrast in Table 1 is not a near-miss — 1.0000 and 0.9944 against 0.0000 and 0.0000, on two independently drawn reservoirs, with input coding held identical.

But "no mixing ⇒ freezing" is not the mechanism, and two results here already rule it out. The obvious reading of the rotating-frame argument — that a recurrence which never combines coordinates lets rounding error decay in place until the state collapses onto finitely many points — is too strong. The linear permutation cells in §4 never combine coordinates either, and make the argument *exact* rather than approximate, yet they barely froze. And a companion study in this program, published alongside this one, replaces the permutation with an equally block-diagonal operator whose rotation numbers are irrational and finds freezing that is real but markedly weaker at every matched cell. Taken together, the defensible statement is narrower than the mechanism the registration assumed: a recurrence that never combines coordinates is apparently *necessary* for this kind of freezing, since the mixing control never leaves zero, but it is demonstrably not sufficient. What supplies the rest — bounded saturation, rational phase, or both — is not settled by this study.

Two senses of "memory under blurring" come apart. The registration treated exact-state determinism and linear decode depth as two views of one quantity, and predicted they would move together under the same dial. They do not. Turning ρ down makes exact-state collision *easier* and makes decode depth *more* sensitive to the grid; turning it up does the reverse on both counts. A system can be nearly automaton-like and still lose almost no linearly decodable history, and it can be far from automaton-like while its decodable history is clamped. This is a further instance of a dissociation this program has now recorded on several unrelated substrates — between what a state stores and what a linear reader can use — and the first surfaced by pairing these two particular instruments.

What this does not establish. Two seeds is the program's promotion minimum, not a tight interval; the values above are reported as a min–max across seeds rather than as a bootstrap confidence interval, and the family's total exposure in the program's ledger is two runs. Whether boundedness is a precondition for freezing rather than a complication it survives is a lead from an unregistered spot check, not a result. The shallow decode-depth slope is reproducible but unexplained: the ρ^k-decay argument that correctly predicted *which* family freezes does not predict *how fast* decodable depth responds, and the gap between 0.8 characters per bit and the theory's 3.1 is the clearest open quantitative question this study leaves. Both threads are recorded in the program's intake queue, and this entry's independent audit is still pending.

---

§6 References

1. H. Jaeger, "The 'echo state' approach to analysing and training recurrent neural networks," GMD Report 148, German National Research Center for Information Technology, 2001.
2. M. Lukoševičius, H. Jaeger, "Reservoir computing approaches to recurrent neural network training," *Computer Science Review* 3(3), 2009.
3. G. Weiss, Y. Goldberg, E. Yahav, "Extracting automata from recurrent neural networks using queries and counterexamples," *ICML* 2018.
4. M. Mahoney, "Large text compression benchmark" (text8), mattmahoney.net/dc/textdata.

§7 Provenance

- **b4ff1bc2** — the first pass; 19 cells, seed 0 (internal run log #40). Registered before running, with the pilot result disclosed. P1 confirmed with the matched control at exactly zero; P2 and P3 missed against their registered anchors and resolved into the unregistered determinism-versus-depth split ⇒ provisional, intake record filed, replication automatically queued. The float32 probe amendment above was journaled and committed before the run resumed.
- **cbe086b8** — seed 899931030, drawn from a hardware random number generator with the raw draw recorded before running; identical driver and identical 19-cell grid (internal run log #41). The headline and the newly-registered depth-split claim both confirmed with zero reversals ⇒ promoted to the program's confirmed-findings record per the pre-stated rule.

All claims judged strictly against the registered wording, misses included above; 2 seeds; replicated before publication; the program's independent audit of this entry is still pending. The float32 comparison cell for the permutation arm is cited from the permutation-recurrence study's archived run and was reproduced rather than assumed. The per-coordinate quantization study it answers is a single-seed result in the program's measured-context record and has no whitepaper of its own; its codebook-based successor is published separately. Internal designation: freezing the card index.