# The Echo-State Property Beyond ρ = 1: Empirical Forgetting Survives to ρ = 1.60 in a Saturating Reservoir

A Preprint

**Caitlyn Meeks**
AuxiLab Tenerife — research@auxi.cafe

August 12, 2026

registered 2026-08-12 | internal run log #34, #35 | 2 seeds × 9 ρ values, zero reversals | replicated before publication | promoted 2026-08-12



**Abstract**

This study originated in an open-ended design brief. The one spectral parameter every reservoir in this program treats as load-bearing is the recurrent matrix's spectral radius ρ, which classical echo-state theory [1] keeps below 1. With a saturating tanh the state is bounded regardless of ρ, so the behavior past that boundary is a checkable question: we swept the reference sparse-tanh reservoir, otherwise completely unchanged, across ρ ∈ 0.50–2.00 (nine points, two seeds, the second hardware-drawn), adding an empirical forgetting test — fork a washed-in reservoir against a randomly reset copy, feed both identical text, and measure whether they converge. Three replicated findings. (1) **Full forgetting survives to ρ=1.60** — disagreement decays to exactly zero, 60–80% past the classical boundary — **then breaks sharply** between ρ=1.60 and 1.80 (final-disagreement ratio 0.000 → 0.45–0.46), the same narrow window in both seeds. (2) Ridge validation bits per character (bpc) degrades smoothly and monotonically from ρ=0.80 onward with no kink at either the classical or the empirical boundary, while decode depth keeps *improving* to an interior peak at exactly ρ=1.60 in both seeds — two separate mechanisms, and a further instance of the decodability–usability dissociation recorded elsewhere in this program. (3) Unplanned but registered as a named alternative in advance: **ρ=0.50 and ρ=0.80 beat the program's own fixed-ρ=0.95 anchor** (3.1472) at both seeds, by 0.0676/0.0670 bits at ρ=0.50 — the margins agree across seeds to 0.0006 bits. The standard baseline convention is not this architecture's own best setting; what, if anything, to change across the program is explicitly left open.

***Keywords*** reservoir computing · echo state property · spectral radius · stability boundary · decodability–usability dissociation · pre-registration

---

## §1 Introduction

The echo state property — the reservoir's state becoming a function of input history alone, independent of initial conditions — is classically guaranteed by keeping the spectral radius of the recurrent matrix below 1, and that sufficient condition has hardened, in practice, into a boundary that is rarely crossed [1][2]. For a *linear* reservoir, crossing it genuinely diverges. For a saturating tanh reservoir the state is confined to $[-1,1]^N$ by construction, so the interesting questions survive the crossing: does the forgetting property actually fail at ρ=1? Where does it really fail? And what happens to prediction and memory on the way?

---

## §2 Method

The reference sparse-tanh reservoir (N=1024, fan-in-10 sparse random W, leak 1.0, bias 0.2, random input map — wiring identical to the archived reference reservoir of the addressed-memory ESN study) with one parameter unlocked: ρ ∈ {0.50, 0.80, 0.95, 1.00, 1.10, 1.30, 1.60, 1.80, 2.00}. Standard text8 splits [3], 2M training characters; ridge and logistic validation bits per character (bpc) and U3 decode depth (a linear decode-depth probe, in the program's internal instrument numbering) on the identical code path as every other reservoir. One new diagnostic, making the classical claim empirical:

```
# empirical echo-state test, per rho:
# wash 200 steps of real text into the reservoir, then FORK:
#   copy A keeps the washed-in state
#   copy B is hard-reset to uniform-random in [-1,1]^N
# feed both the IDENTICAL next 8000 characters; measure RMS
# per-unit disagreement at 0/10/50/200/1000/8000 steps.
conv_ratio = disagreement(8000) / disagreement(0)
# 0   => full forgetting (ESP holds in practice)
# >0  => a permanent memory of an arbitrary initial state
```

The replication seed (1726372129) was drawn from a hardware random-number generator with the raw draw recorded in the registration before running.



---

## §3 Pre-registered predictions

A pilot run over the full grid was disclosed first; predictions were finalized after it and labeled pilot-informed:

- **RC1** (headline): conv_ratio ≤ 0.05 for every ρ ≤ 1.60 and ≥ 0.3 at ρ ∈ {1.80, 2.00} — a sharp, localized transition strictly between 1.60 and 1.80. Named loss: the window sits anywhere else.
- **RC2**: ridge validation bpc monotonically non-decreasing across every adjacent pair from ρ=0.80 through 2.00 — no special kink at the classical ρ=1 line.
- **RC3** (a predicted dissociation): decode depth has a strictly interior maximum while bpc has none over the same range — depth and bpc moving in opposite directions.
- **RC4** (reproduction check): the sweep's own ρ=0.95 cell must reproduce the on-disk reference-reservoir anchor (3.1472) within 0.02 bits. Named alternative, rated a major surprise in advance: *any* cell beats the anchor outright.
- **RCQ1/RCQ2** (bands): conv_ratio at ρ=2.00 stays ≥ 0.5; the best cell lands within 0.10 bits of the anchor.

The replication re-registered all of the above (RRC1–RRC4, both bands) for the fresh seed, with promotion pre-committed to RRC1 and RRC3 being confirmed cleanly, plus replication of RRC4's finding that cells beat the reference-reservoir anchor.

---

## §4 Results

**Table 1.** Both seeds (internal run log entries #34 / #35; values seed 0 / seed 1726372129). Reference-reservoir anchor, cited: ridge 3.1472. All 18 cells pass the U3 lag-1 gate; no logistic readout was voided by the program's fit-invalidating divergence rule.

| ρ | ridge val bpc | U3 depth | conv_ratio |
|---|---|---|---|
| **0.50** | **3.0796 / 3.0802** | **7.18 / 7.18** | **0.000 / 0.000** |
| **0.80** | **3.1073 / 3.1155** | **9.16 / 9.25** | **0.000 / 0.000** |
| 0.95 | 3.1472 / 3.1566 | 10.11 / 10.17 | 0.000 / 0.000 |
| 1.00 | 3.1649 / 3.1742 | 10.30 / 10.36 | 0.000 / 0.000 |
| 1.10 | 3.2064 / 3.2155 | 10.63 / 10.72 | 0.000 / 0.000 |
| 1.30 | 3.3084 / 3.3180 | 11.33 / 11.50 | 0.000 / 0.000 |
| 1.60 | 3.4678 / 3.4766 | 12.10 / 12.35 | 0.000 / 0.000 |
| 1.80 | 3.5482 / 3.5554 | 12.03 / 12.21 | 0.462 / 0.447 |
| 2.00 | 3.6045 / 3.6103 | 11.63 / 11.72 | 0.724 / 0.634 |

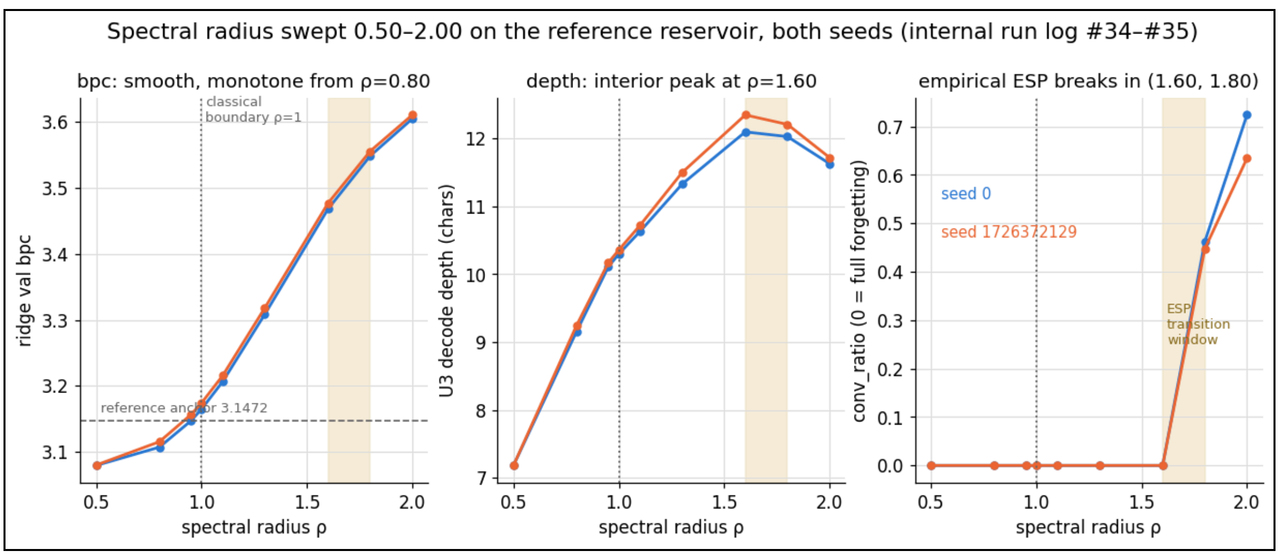


**Figure 1.** The sweep, both seeds (internal run log entries #34, #35). Dotted vertical line: the classical ρ=1 boundary. Shaded band: the empirically measured transition window (1.60, 1.80). Left panel's dashed line: the cited reference-reservoir anchor 3.1472 — the ρ=0.50 and ρ=0.80 cells sit below it in both seeds.

**The empirical forgetting boundary is sharp, and lies well beyond the classical line.** Disagreement between the forked copies decays to exactly zero at every ρ through 1.60, in both seeds — the reservoir state is a function of its input history alone even 60% past the classical boundary. At ρ=1.80 the two copies never converge (conv_ratio 0.462 / 0.447), and at ρ=2.00 the plateau deepens (0.724 / 0.634). The transition falls in the same narrow window in both seeds (RC1/RRC1 confirmed; RCQ1/RRCQ1 confirmed).

**Prediction and memory dissociate across the same sweep.** Ridge bpc worsens monotonically from ρ=0.80 through 2.00 — 7/7 adjacent pairs, both seeds, with no visible kink at either ρ=1 or the empirical transition — while decode depth keeps improving to an interior peak at exactly ρ=1.60 in both seeds, then turns down. The peak in depth lands almost exactly where forgetting starts to fail. Whatever bpc pays for is a smooth mechanism separate from the sharp one the forgetting test measures (RC2/RC3 and their replications confirmed; this reservoir *does* show the decodability–usability split that the threshold-fire-and-refractory reservoir did not).

**The reproduction check passed, and its named alternative was triggered.** At seed 0 the sweep's ρ=0.95 cell reproduces the on-disk reference-reservoir anchor exactly (3.1472, difference 0.0000; the archived anchor was itself computed at seed 0, so this is a like-for-like rebuild). But ρ=0.50 (3.0796 / 3.0802) and ρ=0.80 (3.1073 / 3.1155) beat the anchor outright at both seeds — margins 0.0676 / 0.0670 bits at ρ=0.50, a 0.0006-bit spread across independently drawn reservoirs, in a smooth four-point trend rather than an isolated cell. The registered 0.10-bit band (RCQ2) technically held; the band's own motivating question — was ρ=0.95 already near this architecture's optimum? — resolved clearly in the negative.



---

## §5 Discussion

**The classical line is not where this system breaks.** For a saturating tanh reservoir under real text drive, the empirical echo-state boundary sits far past the classical sufficient condition — and when it fails, it fails sharply, not gradually. Between the two boundaries lies a genuinely usable regime: at ρ=1.60 the reservoir decodes *deeper* than anywhere else on the grid while still fully forgetting its initial state — obtained at a real bpc cost, the same trade-off between decodability and usability that recurs throughout this program.

**The program's own baseline is not optimal, and this is now a confirmed finding.** The fixed-ρ=0.95 anchor (3.1472) cited as the standard baseline across dozens of registrations is a reproducible ~0.05–0.07 bit underperformance relative to lower ρ on this same architecture. What this paper does *not* say: whether the baseline convention should change across the program, whether any past conclusion needs re-reading (each was judged on its own registered terms), or whether ρ below 0.50 does better still — untested. Those questions are flagged in the program's records for a deliberate decision, not decided by this study. Uncertainty is reported as a min–max range across two seeds, and the program's independent audit of this entry 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. M. Mahoney, "Large text compression benchmark" (text8), mattmahoney.net/dc/textdata.


---

## §7 Provenance

- **907a2395** — the originating design brief; 9 cells, seed 0 (internal run log #34). Pilot run disclosed before registration; all four ordinal predictions and both bands confirmed, and RC4's named "would be a major surprise" alternative was triggered ⇒ provisional, intake record filed, replication automatically queued.
- **f6a50222** — seed 1726372129, hardware-drawn with the raw draw recorded before running (internal run log #35). All predictions confirmed, zero reversals; the margin over the reference-reservoir anchor replicated to 0.0006 bits ⇒ all three findings promoted to the program's confirmed-findings record per the pre-stated rule.

All claims judged strictly against the registered wording; 2 seeds; replicated before publication; the program's independent audit of this entry is still pending. Internal designation: rainbow connection.

---